

Sun Network QDR InfiniBand Gateway Switch

Administration Guide for Firmware Version 2.1



Part No.: E36260-01
March 2013, Revision A

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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

Copyright © 2013, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

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

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

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

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

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

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



Adobe PostScript

Contents

Using This Documentation xi

Troubleshooting the Gateway 1

Gateway Hardware Problems 1

InfiniBand Fabric Problems 3

Network Management Troubleshooting Guidelines 6

Understanding Routing Through the Gateway 7

Switch Chip Port to QSFP Connectors and Link LED Routes 8

QSFP Connectors and Link LEDs to Switch Chip Port Routes 8

Signal Route Through the Gateway 9

Switch GUIDs Overview 10

Understanding Administrative Commands 13

Configuration Overview 13

Oracle ILOM Command Overview 16

Hardware Command Overview 16

InfiniBand Command Overview 17

Administering the Chassis 19

Monitoring the Chassis 19

▼ Display Gateway General Health 20

▼ Display Power Supply Status 20

▼ Check Board-Level Voltages 21

- ▼ Display Internal Temperatures 22
- ▼ Display Fan Status 22
- ▼ Display Gateway Environmental and Operational Data 23
- ▼ Display Flash Drive Information 24
- ▼ Display Free Memory and File System Space 25
- ▼ Display Chassis FRU ID 26
- ▼ Display Power Supply FRU ID 26
- ▼ Display Gateway Firmware Versions 27
- Oracle Auto Service Request 28

Controlling the Chassis 28

- ▼ Restart the Management Controller 29
- ▼ Change the root User Password 29

Administering the I4 Switch Chip 31

Monitoring the I4 Switch Chip 31

- ▼ Display the Switch Chip Port to QSFP Connector Mapping 32
- ▼ Locate a InfiniBand Device and Connector From the GUID and Port 33
- ▼ Locate a Switch Chip or Connector From the GUID and Port 33
- ▼ Display Switch Chip Boot Status 34
- ▼ Display Link Status 35
- ▼ Display Connectors Configured With Autodisable Functionality 36
- ▼ Display Switch Chip Port Status 37
- ▼ Display Switch Chip Port Counters 38

Controlling the I4 Switch Chip 40

- ▼ Disable a Switch Chip Port 40
- ▼ Enable a Switch Chip Port 41

Autodisable Ports and Connectors 43

Autodisable Functionality 43

- ▼ Add Connectors to the Autodisable List 44
- ▼ Reenable Autodisabled Connectors 45
- ▼ Delete Connectors From the Autodisable List 46

Administering the Subnet Manager 49

Monitoring the Subnet Manager 49

- ▼ Display Subnet Manager Status 50
- ▼ Display the Subnet Manager's Location History 51
- ▼ Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm 51
- ▼ Display the Secret M_Key Configuration 52
- ▼ Verify M_Key Persistence 53
- ▼ Verify the Subnet Manager Configuration Integrity 54
- ▼ Display the Subnet Manager Log 55

Controlling the Subnet Manager 56

- ▼ Set the Subnet Manager Priority 56
- ▼ Set the Subnet Manager Prefix 57
- ▼ Enable Subnet Manager Controlled Handover 58
- ▼ Set the Subnet Manager Readable Management Key 59
- ▼ Clear the Subnet Manager Readable Management Key 60
- ▼ Set the Subnet Manager Routing Algorithm 61
- ▼ Enable the Subnet Manager 62
- ▼ Disable the Subnet Manager 63

Configuring Secure Fabric Management 64

Secret M_Key Overview 64

smsubnetprotection Command 65

- ▼ Prepare for Secret M_Key Functionality 66

Configuring Secret M_Keys 67

- ▼ Create the Secret M_Key Configuration 67

- ▼ Add Secret M_Keys to the Configuration 69
- ▼ Set the Secret M_Key to Use 72
- ▼ Delete Secret M_Keys From the Configuration 73
- ▼ Configure a Local Secret M_Key 75
- ▼ Set the Replication Password 76
- ▼ Enable or Disable Secret M_Key Persistence 76
- ▼ Enable or Disable Secret M_Key Functionality 77
- Considering Situations In Regard to Secret M_Key 78
 - ▼ Add a New Subnet Manager to the Secret M_Key Fabric 78
 - ▼ Merge Subnets Into One Secret M_Key Fabric 79
- Secret M_Key and Firmware Upgrade and Downgrade Guidelines 80

Administering the InfiniBand Fabric 83

Monitoring the InfiniBand Fabric 83

Displaying Information About Nodes 84

- ▼ Identify All Switches in the Fabric 84
- ▼ Identify All CAs in the Fabric 85
- ▼ Display the InfiniBand Fabric Topology (Simple) 86
- ▼ Display the InfiniBand Fabric Topology (Detailed) 86
- ▼ Display a Route Through the Fabric 87
- ▼ Display the Link Status of a Node 88
- ▼ Display Counters for a Node 89
- ▼ Display Low-Level Detailed Information About a Node 90
- ▼ Display Low-Level Detailed Information About a Port 91

Displaying Fabric Functionality Configuration Information 93

- ▼ Display the InfiniBand Fabric Partition Configuration 93
- ▼ Display the InfiniBand Fabric Configuration 94
- ▼ Display the smnodes List 95

- ▼ Display the Fabric Element Configuration 96
- ▼ Display the IPoIB Interface 96
- ▼ Display Configured Secret M_Keys 97

Controlling the InfiniBand Fabric 97

- Investigating Nodes 98
 - ▼ Perform Comprehensive Diagnostics for the Entire Fabric 98
 - ▼ Determine Changes to the InfiniBand Fabric Topology 99
 - ▼ Find 1x, SDR, or DDR Links in the Fabric 100
 - ▼ Determine Which Links Are Experiencing Significant Errors 100

Controlling Nodes 101

- ▼ Clear Data and Error Counters 102
- ▼ Reset a Port 102
- ▼ Set Port Speed 103
- ▼ Disable or Enable an External Port 104

Partitioning the InfiniBand Fabric 105

- `smpartition` Command 106
- `smpartition` Command Guidelines 107
- Partitions and P_Keys 108
 - ▼ Determine the Partitions and P_Keys 109
 - ▼ Create a User Partition 110
 - ▼ Add or Remove a Port From a Partition 112
 - ▼ Modify a Partition or Port 114
 - ▼ Delete a Partition 116
 - ▼ Remove User Partitions for Firmware Downgrade 117

Correlating Fabric Elements 119

- `createfabric` Command 120
 - ▼ Create a Fabric Element Configuration 120
 - ▼ Add or Delete an Element From the Configuration 122

▼ Modify an Element of the Configuration	123
Configuring the Fabric Director Node List	125
Fabric Director and Fabric Elements	125
fdconfig Command Overview	126
▼ Create a Fabric Configuration	127
▼ Add or Remove an Element From the Fabric Configuration	129
▼ Modify an Element of the Fabric Configuration	131
Performing Supportive Tasks	132
▼ Create the smnodes List	133
▼ Update the smnodes List	134
▼ Create the IPoIB Interface	135
▼ Delete the IPoIB Interface	136
Administering Gateway Resources	137
Installing Gateway Supportive Software (Linux)	137
▼ Acquire the BXOFED Software (Linux)	138
▼ Install the BXOFED Software (Linux)	139
▼ Acquire the ConnectX-2 Firmware	141
Creating VNICs Under Gateway Manual Mode (Linux)	142
Gateway Manual Mode Overview (Linux)	143
▼ Determine VNIC Configuration Parameters For Gateway Manual Mode (Linux)	144
▼ Determine VLAN Associations for Gateway Manual Mode (Linux)	145
▼ Configure and Create VNICs for Gateway Manual Mode (Linux)	146
VNIC Configuration on Linux Hosts in Gateway Manual Mode	147
Creating VNICs Under Host Manual Mode (Linux)	148
Host Manual Mode Overview (Linux)	149
Central Configuration File (Linux)	150
VNIC-Specific Configuration File (Linux)	151

- ▼ Determine VNIC Configuration Parameters for Host Manual Mode (Linux) 152
- ▼ Create the Central Configuration File (Linux) 154
- ▼ Create the VNIC-Specific Configuration Files (Linux) 156
- mlx4_vnic_confd Daemon (Linux) 157
- ▼ Configure and Create VNICs for Host Manual Mode (Linux) 158

Creating Virtual IO Adapters (Oracle Solaris) 160

- Ethernet Over InfiniBand Overview (Oracle Solaris) 161
- ▼ Verify That the Ethernet Over InfiniBand Driver Is Installed (Oracle Solaris) 162
- ▼ Install the Ethernet Over InfiniBand Driver (Oracle Solaris) 162
- ▼ Set Up VIOAs (Oracle Solaris) 163
- ▼ Configure IP Addresses on the Data Link (Oracle Solaris) 167
- ▼ Create Oracle Solaris VNICs With VLAN Tagging (Oracle Solaris) 169

Monitoring Gateway Resources 172

- ▼ Display LAG Information 173
- ▼ Display the VLANs 174
- ▼ Display the VNICs 176
- ▼ Display IO Adapters 177
- ▼ Display Gateway Ethernet Port Information 178
- ▼ Display Gateway Port Information 179
- ▼ Display the Gateway Instance Number 181
- ▼ Display Information About the Gateway 182

Controlling LAGs 183

- ▼ Enable LAG Mode 184
- ▼ Create LAGs 185
- ▼ Create a Global Gateway 186
- ▼ Add or Delete Connectors From a LAG 187
- ▼ Delete a LAG 188

▼ Disable LAG Mode	189
Controlling VLANs and VNICs	190
VNIC Allocation on the Gateway	191
▼ Create VLANs	192
▼ Delete VLANs	194
▼ Create VNICs	195
▼ Modify VNICs	197
▼ Enable VNICs	199
▼ Disable VNICs	200
▼ Delete VNICs	201
Controlling Gateway Ports and Parameters	202
▼ Enable a Gateway Ethernet Port	203
▼ Disable a Gateway Ethernet Port	204
▼ Set the Gateway Ethernet Port Parameters	205
▼ Set the Gateway Service Level	206
▼ Set the Gateway Instance Number	207
▼ Set the Gateway System Name	208
Index	209

Using This Documentation

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

- “Product Notes” on page xi
- “Related Documentation” on page xi
- “Feedback” on page xii
- “Access to Oracle Support” on page xii

Product Notes

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

http://docs.oracle.com/cd/E36256_01

Related Documentation

Documentation	Links
---------------	-------

Sun Network QDR InfiniBand Gateway Switch Firmware Version 2.1	http://docs.oracle.com/cd/E36256_01
---	---

Documentation	Links
Oracle Solaris 11 OS	http://www.oracle.com/goto/Solaris11/docs
Oracle Integrated Lights Out Manager (ILOM) 3.0	http://docs.oracle.com/cd/E19860-01
All Oracle products	http://docs.oracle.com

Feedback

Provide feedback on this documentation at:

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

Access to Oracle Support

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

Troubleshooting the Gateway

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

- [“Gateway Hardware Problems” on page 1](#)
- [“InfiniBand Fabric Problems” on page 3](#)
- [“Network Management Troubleshooting Guidelines” on page 6](#)
- [“Understanding Routing Through the Gateway” on page 7](#)
- [“Switch GUIDs Overview” on page 10](#)

Related Information

- [“Understanding Administrative Commands” on page 13](#)
- [“Administering the Chassis” on page 19](#)
- [“Administering the I4 Switch Chip” on page 31](#)
- [“Administering Gateway Resources” on page 137](#)
- [“Administering the InfiniBand Fabric” on page 83](#)
- [“Administering the Subnet Manager” on page 49](#)

Gateway Hardware Problems

This table lists situations that might occur with gateway 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.	<ol style="list-style-type: none">1. Check the power supply status. See “Display Power Supply Status” on page 20.2. Unplug the respective power cord, wait 15 minutes, then reattach the power cord.3. If the previous steps do not rectify the situation, replace the power supply. Refer to <i>Gateway Service</i>, servicing power supplies.

Situation	Corrective Steps
There is no network communication with the management controller.	<ol style="list-style-type: none"> 1. Verify that the management controller is powered without faults. Refer to <i>Gateway Service</i>, checking the chassis status LEDs. 2. Verify that there is a LinkUp condition at the NET MGT port. Refer to <i>Gateway Service</i>, checking the NET MGT port LEDs. 3. Verify that the DHCP server is providing the IP address you are using to access the management controller. See DHCP server documentation. 4. If you are able to access the management controller through the USB connector, restart the management controller. Refer to <i>Gateway Installation</i>, accessing the management controller and “Restart the Management Controller” on page 29. 5. If you are unable to access the management controller through the USB connector, power cycle the gateway. Refer to <i>Gateway Service</i>, powering off power supplies and <i>Gateway Service</i>, powering on power supplies. 6. If the previous steps do not rectify the situation, replace the gateway. Refer to <i>Gateway Service</i>, removing the gateway from the rack. Refer to <i>Gateway Installation</i>, installing the gateway into the rack.
The Attention LED on the gateway chassis is lit or the management controller seems dysfunctional.	<ol style="list-style-type: none"> 1. If you are unable to access the management controller, power cycle the gateway. Refer to <i>Gateway Service</i>, powering off the power supply and <i>Gateway Service</i>, 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 29. 3. Check overall gateway health. See “Display Gateway General Health” on page 20 4. Verify that the gateway is within operating temperatures and voltages. See “Display Gateway Environmental and Operational Data” on page 23. 5. If the previous steps do not rectify the situation, replace the gateway. Refer to <i>Gateway Service</i>, removing the gateway from the rack. Refer to <i>Gateway Installation</i>, installing the gateway into the rack.
The Attention LED on a fan is lit or the fan seems dysfunctional.	<ol style="list-style-type: none"> 1. Check the fan speed. See “Display Fan Status” on page 22. 2. If the previous step does not rectify the situation, replace the fan. Refer to <i>Gateway Service</i>, servicing fans. 3. If the previous step does not rectify the situation, install the fan into another available slot. Refer to <i>Gateway Service</i>, servicing fans. 4. If no other slots are available, replace the gateway. Refer to <i>Gateway Service</i>, removing the gateway from the rack. Refer to <i>Gateway Installation</i>, installing the gateway into the rack.

Related Information

- “InfiniBand Fabric Problems” on page 3
- “Network Management Troubleshooting Guidelines” on page 6
- “Understanding Routing Through the Gateway” on page 7
- “Switch GUIDs Overview” on page 10

InfiniBand Fabric Problems

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

Situation	Corrective Steps
After installation, no links are operational.	<ol style="list-style-type: none">1. Verify that there is at least one Subnet Manager active on the InfiniBand fabric. See “Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51.2. If no Subnet Manager is active, start the Subnet Manager within the gateway. Refer to <i>Gateway Installation</i>, starting the Subnet Manager.3. If the previous steps do not rectify the situation, restart the Subnet Manager. See “Disable the Subnet Manager” on page 63 and “Enable the Subnet Manager” on page 62.
After installation, not all links are operational.	<ol style="list-style-type: none">1. Determine which links are nonoperational. See “Display Link Status” on page 35.2. For links that are “Down”, disable and re-enable the respective ports. See “Disable a Switch Chip Port” on page 40 and “Enable a Switch Chip Port” on page 41.3. If the previous steps do not rectify the situation, disable the respective port. See “Disable a Switch Chip Port” on page 40.
There was a power outage during a firmware update.	<ol style="list-style-type: none">1. If you are able to access the management controller, restart the management controller. See “Restart the Management Controller” on page 29.2. If you are unable to access the management controller, power cycle the gateway. Refer to <i>Gateway Service</i>, removing the gateway from the rack. Refer to <i>Gateway Installation</i>, installing the gateway into the rack.3. Repperform the firmware upgrade. Refer to <i>Gateway Remote Management</i>, upgrading the gateway firmware.

Situation	Corrective Steps
Performance of the InfiniBand fabric seems diminished.	<ol style="list-style-type: none"> 1. Determine if there are errors or problems with the InfiniBand fabric. See: <ul style="list-style-type: none"> “Perform Comprehensive Diagnostics for the Entire Fabric” on page 98 “Find 1x, SDR, or DDR Links in the Fabric” on page 100 “Determine Which Links Are Experiencing Significant Errors” on page 100 2. Locate the affected nodes by the GUID provided in the output of the <code>ibdiagnet</code> command. See “Locate a Switch Chip or Connector From the GUID and Port” 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. Refer to <i>Gateway Service</i>, 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 or Enable an External Port” on page 104. <p>Temporary solution:</p> <ul style="list-style-type: none"> • If the problem still remains, disable the affected port. See “Disable or Enable an External Port” on page 104. <p>Permanent solution:</p> <ul style="list-style-type: none"> • If the problem still remains, replace the affected component or the gateway. Refer to <i>Gateway Service</i>, servicing data cables. Refer to remote port’s documentation for replacement procedures. Refer to <i>Gateway Service</i>, removing the gateway from the rack. Refer to <i>Gateway Installation</i>, installing the gateway into the rack.
An InfiniBand Link LED is blinking.	<ol style="list-style-type: none"> 1. Disconnect and properly reconnect both ends of the respective InfiniBand cable. Refer to <i>Gateway Service</i>, servicing the data cables. 2. If the LED is still blinking, determine the significance of the errors through use of the <code>ibdiagnet</code> command. See “Determine Which Links Are Experiencing Significant Errors” on page 100. 3. Determine which connectors map to the affected link by deconstructing the node’s GUID and port. See “Locate a Switch Chip or Connector From the GUID and Port” on page 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 or Enable an External Port” on page 104. 6. If the errors are still significant, swap the cable with a known good one or reconnect the cable to a known good remote port, and repeat from Step 2. 7. Depending upon what does or does not rectify the problem, replace that component. Refer to <i>Gateway Service</i>, servicing the data cables. Refer to remote port’s documentation for replacement procedures.

Situation	Corrective Steps
Some InfiniBand links are running at 1x or SDR.	<p>For a temporary solution:</p> <ol style="list-style-type: none"> Identify the suspect links using the <code>ibdiagnet</code> command. See “Find 1x, SDR, or DDR Links in the Fabric” on page 100. Look for text like this: <pre>-W- link with SPD=2.5 found at direct path "1,19" From: a Switch PortGUID=0x00066a00d80001dd Port=19 To: a Switch PortGUID=0x00066a00d80001dd Port=24</pre> Determine which connectors map to the affected link by deconstructing the node’s GUID and port. See “Locate a Switch Chip or Connector From the GUID and Port” on page 33. Verify the cable connection at both ends. Refer to <i>Gateway Service</i>, servicing the data cables. Disable and re-enable the respective ports. See “Disable or Enable an External Port” on page 104. If the previous steps do not rectify the problem, disable the port. See “Disable or Enable an External Port” on page 104. <p>For a permanent solution:</p> <ol style="list-style-type: none"> Perform the steps for a temporary solution, Steps 1 to Step 4. Swap the cable with a known good cable or reconnect the cable to a known good remote port, and repeat from Step 1. Depending upon what does or does not rectify the problem, replace that component or the gateway. Refer to <i>Gateway Service</i>, servicing the data cables. Refer to the remote port’s documentation for replacement procedures. Refer to <i>Gateway Service</i>, removing the gateway from the rack. Refer to <i>Gateway Installation</i>, installing the gateway into the rack.
There are errors on some InfiniBand links.	<ol style="list-style-type: none"> Clear the error counters. See “Clear Data and Error Counters” on page 102. Start a fabric stress test. Identify the suspect links using the <code>ibdiagnet</code> command. See “Determine Which Links Are Experiencing Significant Errors” on page 100. Look for text like this: <pre>-W- lid=0x0006 guid=0x0021283a8816c0a0 dev=48438 Port=34 Performance Monitor counter : Value link_recovery_error_counter : 0x1 symbol_error_counter : 0x25 (Increase by 3 during ibdiagnet)</pre> 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.	<ol style="list-style-type: none"> 1. You can find the location of a node in the gateway by deconstructing the node's GUID and port. See "Locate a Switch Chip or Connector From the GUID and Port" on page 33. 2. Use the <code>dcspport</code> 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

- ["Gateway Hardware Problems"](#) on page 1
- ["Network Management Troubleshooting Guidelines"](#) on page 6
- ["Understanding Routing Through the Gateway"](#) on page 7
- ["Switch GUIDs Overview"](#) on page 10

Network Management Troubleshooting Guidelines

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

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

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

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

Service	Command
HTTP	<code>set /SP/services/http secureredirect=disabled servicestate=enabled</code>
HTTPS	<code>set /SP/services/https servicestate=enabled</code>

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

Related Information

- [“Gateway Hardware Problems” on page 1](#)
- [“InfiniBand Fabric Problems” on page 3](#)
- [“Understanding Routing Through the Gateway” on page 7](#)
- [“Switch GUIDs Overview” on page 10](#)

Understanding Routing Through the Gateway

The tables in these topics describe the routing through the gateway. 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 8](#)
- [“QSFP Connectors and Link LEDs to Switch Chip Port Routes” on page 8](#)
- [“Signal Route Through the Gateway” on page 9](#)

Related Information

- [“Gateway Hardware Problems” on page 1](#)
- [“InfiniBand Fabric Problems” on page 3](#)
- [“Network Management Troubleshooting Guidelines” on page 6](#)
- [“Switch GUIDs Overview” on page 10](#)

Switch Chip Port to QSFP Connectors and Link LED Routes

Port	Connector	Port	Connector	Port	Connector	Port	Connector
1	1A-ETH-1(P1) 1A-ETH-2(P2)	10	13B	19	0B	28	4A
2	1A-ETH-3(P3) 1A-ETH-4(P4)	11	12A	20	0A	29	5B
3	0A-ETH-1(P1) 0A-ETH-2(P2)	12	12B	21	1B	30	5A
4	0A-ETH-3(P3) 0A-ETH-4(P4)	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 35](#)
- [“Display Switch Chip Port Status” on page 37](#)
- [“QSFP Connectors and Link LEDs to Switch Chip Port Routes” on page 8](#)
- [“Signal Route Through the Gateway” on page 9](#)

QSFP Connectors and Link LEDs to Switch Chip Port Routes

Connector Group	Connector A Port	Connector B Port	Connector Group	Connector A Port	Connector B Port
0	20	19	9	14	13
1	22	21	10	16	15

Connector Group	Connector A Port	Connector B Port	Connector Group	Connector A Port	Connector B Port
2	24	23	11	18	17
3	26	25	12	11	12
4	28	27	13	9	10
5	30	29	14	7	8
6	35	36	15	5	6
7	33	34	0	3 (ETH-1, ETH-2) 4 (ETH-3, ETH-4)	
8	31	32	1	1 (ETH-1, ETH-2) 2 (ETH-3, ETH-4)	

Related Information

- [“Display the Switch Chip Port to QSFP Connector Mapping”](#) on page 32
- [“Display Link Status”](#) on page 35
- [“Display Switch Chip Port Status”](#) on page 37
- [“Switch Chip Port to QSFP Connectors and Link LED Routes”](#) on page 8
- [“Signal Route Through the Gateway”](#) on page 9

Signal Route Through the Gateway

By combining the information from the tables in [“Understanding Routing Through the Gateway”](#) on page 7, it is possible to determine a route through the gateway. This topic describes a sample situation that might occur.

1. A route is initiated at connector 2A. The LED blinks.
2. Using [“QSFP Connectors and Link LEDs to Switch Chip Port Routes”](#) on page 8, it is determined that the link routes to the switch chip through port 24.
3. The Subnet Manager instructs the switch chip to use port 5 to forward the link.
4. Using [“Switch Chip Port to QSFP Connectors and Link LED Routes”](#) on page 8, it is determined that the link routes to connector 15A. The LED blinks.

Related Information

- [“Switch Chip Port to QSFP Connectors and Link LED Routes”](#) on page 8
- [“QSFP Connectors and Link LEDs to Switch Chip Port Routes”](#) on page 8

- “Display the Switch Chip Port to QSFP Connector Mapping” on page 32
- “Display a Route Through the Fabric” on page 87

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 gateway, GUIDs are modified to identify the node’s role and location. This table describes the GUID’s structure.

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

The five fields of the GUID are described as follows:

- **MAC** – Bits 63 through 16 are the Machine Allocation Code (MAC) address. A standard for network components, the MAC address is typically provided by manufacturers in a 6-byte, colon delimited string. For example, 00:11:22:33:44:55.
- **Type** – Bits 15 through 12 is the type of board on which the node resides.
- **Pos** – Bits 11 through 8 identify the position of the board within the gateway.
- **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) to 0x8 (Fabric card 8)	0xA (I4 chip 0) to 0xB (I4 chip 1)
Line card	0x1	0x0 (Line card 0) to 0x8 (Line card 8)	0xA (I4 chip 0) to 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) to 0xF (I4 chip 5)

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

```
Ca : 0x0021283bad45c000 ports 2 "SUN IB QDR GW switch gw-2 Bridge 0"
```

The GUID is 0x0021283bad45c000 or 0021283BAD45 C 0 0 0. Using the information provided in this topic:

- The MAC address is 0x0021283BAD45 or 00:21:28:3B:AD:45.
- The type is 0xC, or a gateway.
- The position is 0x0.
- The device is 0x0, or BridgeX chip 0.
- The number is 0.

Related Information

- *Gateway Reference*, `ibnetdiscover` command
- *Gateway Reference*, `ibnodes` command
- *Gateway Reference*, `ibswitches` command
- *Gateway Reference*, `ibhosts` command
- [“Identify All Switches in the Fabric” on page 84](#)
- [“Gateway Hardware Problems” on page 1](#)
- [“InfiniBand Fabric Problems” on page 3](#)
- [“Network Management Troubleshooting Guidelines” on page 6](#)
- [“Understanding Routing Through the Gateway” on page 7](#)

Understanding Administrative Commands

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

- [“Configuration Overview”](#) on page 13
- [“Oracle ILOM Command Overview”](#) on page 16
- [“Hardware Command Overview”](#) on page 16
- [“InfiniBand Command Overview”](#) on page 17

Related Information

- [“Troubleshooting the Gateway”](#) on page 1
- [“Administering the Chassis”](#) on page 19
- [“Administering the I4 Switch Chip”](#) on page 31
- [“Administering Gateway Resources”](#) on page 137
- [“Administering the InfiniBand Fabric”](#) on page 83
- [“Administering the Subnet Manager”](#) on page 49

Configuration Overview

After you have installed the gateway, you can configure its functionality. This table lists functions of the gateway 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 56
		Configure the prefix.	“Set the Subnet Manager Prefix” on page 57
		Enable controlled handover.	“Enable Subnet Manager Controlled Handover” on page 58
		Configure the M_Key.	“Set the Subnet Manager Readable Management Key” on page 59
		Set the routing algorithm, if not Fat Tree.	“Set the Subnet Manager Routing Algorithm” on page 61
2.	Configure the InfiniBand fabric.	Create a fabric configuration.	“Create a Fabric Configuration” on page 127
3.	Secure fabric management.	Create the smnodes list.	“Create the smnodes List” on page 133
		(optional) Configure IPOIB functionality.	“Create the IPoIB Interface” on page 135
		Create the fabric element configuration.	“Create a Fabric Element Configuration” on page 120
		Set the replication password.	“Set the Replication Password” on page 76
		Enable M_Key persistence.	“Enable or Disable Secret M_Key Persistence” on page 76
		Configure secret M_Keys.	“Configuring Secret M_Keys” on page 67
4.	Partition the InfiniBand fabric.	Enable secret M_Keys.	“Enable or Disable Secret M_Key Functionality” on page 77
		Collect GUIDs.	“Identify All Switches in the Fabric” on page 84 “Identify All CAs in the Fabric” on page 85 “Display Gateway Port Information” on page 179
		Update the smnodes file.	“Update the smnodes List” on page 134
		Create the partition information file.	“Determine the Partitions and P_Keys” on page 109
		Configure partitions and P_Keys.	“Create a User Partition” on page 110
5.	Configure LAGs.	Add GUIDs to the partitions.	“Add or Remove a Port From a Partition” on page 112
		Enable LAG mode.	“Enable LAG Mode” on page 184
		Create LAGs.	“Create LAGs” on page 185

Step	Functionality	Description	Links
6.	Enable host support of the gateway.	Acquire BXOFED software.	“Acquire the BXOFED Software (Linux)” on page 138
		Install BXOFED software.	“Install the BXOFED Software (Linux)” on page 139
7.	Determine the VNIC mode.	Use Gateway Manual Mode or Host Manual Mode.	“Gateway Manual Mode Overview (Linux)” on page 143 “Host Manual Mode Overview (Linux)” on page 149
		8. Configure VNICs.	
	Gateway Manual Mode	Create the MAC address file.	“Determine VNIC Configuration Parameters For Gateway Manual Mode (Linux)” on page 144
		Append the MAC address file.	“Determine VLAN Associations for Gateway Manual Mode (Linux)” on page 145
		Create VLANs.	“Create VLANs” on page 192
		Create VNICs.	“Create VNICs” on page 195
	Host Manual Mode	Determine the configuration file.	“Central Configuration File (Linux)” on page 150 “VNIC-Specific Configuration File (Linux)” on page 151
		Create the MAC address file.	“Determine VNIC Configuration Parameters for Host Manual Mode (Linux)” on page 152
		Create the configuration file.	“Create the Central Configuration File (Linux)” on page 154 “Create the VNIC-Specific Configuration Files (Linux)” on page 156
	Enable VLANs and VNICs.	“Configure and Create VNICs for Host Manual Mode (Linux)” on page 158	
9.	Save the configuration.	Back up the configuration for restoration at a later time, should it become necessary.	<i>Gateway Remote Management</i> , backing up the configuration.

Related Information

- [“Oracle ILOM Command Overview” on page 16](#)
- [“Hardware Command Overview” on page 16](#)
- [“InfiniBand Command Overview” on page 17](#)

Oracle ILOM Command Overview

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

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

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

Information about Oracle ILOM support of the gateway is available in the *Oracle Integrated Lights Out Manager (ILOM) 3.0 Supplement for the Sun Network QDR InfiniBand Gateway Switch*, available online at:

<http://www.oracle.com/pls/topic/lookup?ctx=E19671-01&id=homepage>

Overall information about Oracle ILOM 3.0 is available online at:

<http://www.oracle.com/pls/topic/lookup?ctx=E19860-01&id=homepage>

Related Information

- *Gateway Remote Management*, understanding Oracle ILOM commands
- “[Hardware Command Overview](#)” on page 16
- “[InfiniBand Command Overview](#)” on page 17

Hardware Command Overview

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

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

```
FabMan@gateway_name->command [arguments] [arguments] . . .
```

Related Information

- *Gateway Reference*, understanding hardware commands
- [“Oracle ILOM Command Overview” on page 16](#)
- [“InfiniBand Command Overview” on page 17](#)

InfiniBand Command Overview

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

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

```
FabMan@gateway_name->command [option] [option] ...
```

Related Information

- *Gateway Reference*, understanding InfiniBand commands
- [“Oracle ILOM Command Overview” on page 16](#)
- [“Hardware Command Overview” on page 16](#)

Administering the Chassis

These topics describe the administration of the gateway.

- [“Monitoring the Chassis” on page 19](#)
- [“Controlling the Chassis” on page 28](#)

Related Information

- [“Troubleshooting the Gateway” on page 1](#)
- [“Understanding Administrative Commands” on page 13](#)
- [“Administering the I4 Switch Chip” on page 31](#)
- [“Administering Gateway Resources” on page 137](#)
- [“Administering the InfiniBand Fabric” on page 83](#)
- [“Administering the Subnet Manager” on page 49](#)

Monitoring the Chassis

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

- [“Display Gateway General Health” on page 20](#)
- [“Display Power Supply Status” on page 20](#)
- [“Check Board-Level Voltages” on page 21](#)
- [“Display Internal Temperatures” on page 22](#)
- [“Display Fan Status” on page 22](#)
- [“Display Gateway Environmental and Operational Data” on page 23](#)
- [“Display Flash Drive Information” on page 24](#)
- [“Display Free Memory and File System Space” on page 25](#)
- [“Display Chassis FRU ID” on page 26](#)
- [“Display Power Supply FRU ID” on page 26](#)
- [“Display Gateway Firmware Versions” on page 27](#)

- “Oracle Auto Service Request” on page 28

Related Information

- “Monitoring the I4 Switch Chip” on page 31
- “Monitoring Gateway Resources” on page 172
- “Monitoring the InfiniBand Fabric” on page 83
- “Monitoring the Subnet Manager” on page 49
- “Controlling the Chassis” on page 28

▼ Display Gateway General Health

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

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

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

Related Information

- *Gateway Reference*, `showunhealthy` command
- “Display Gateway Environmental and Operational Data” on page 23

▼ Display Power Supply Status

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

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

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

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

- a. **Remove the power cord and affected power supply from the gateway chassis.**
Refer to *Gateway Service*, powering off the power supply and *Gateway Service*, removing the power supply.
- b. **Wait one minute.**
- c. **Reinstall the power supply.**
Refer to *Gateway Service*, installing the power supply.
- d. **Reattach the power cord.**
Refer to *Gateway Service*, powering on the power supply.
- e. **If the power supply's Attention LED lights or the checkpower command still reports Alert or Not Present for the power supply, replace the power supply.**
Refer to *Gateway Service*, servicing the power supplies.

Related Information

- [Gateway Reference, checkpower command](#)
- [“Check Board-Level Voltages” on page 21](#)
- [“Display Gateway Environmental and Operational Data” on page 23](#)

▼ Check Board-Level Voltages

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

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

```
FabMan@gateway_name->checkvoltages
Voltage ECB OK
Measured 3.3V Main = 3.30 V
Measured 3.3V Standby = 3.42 V
Measured 12V = 12.06 V
Measured 5V = 5.03 V
Measured VBAT = 3.17 V
Measured 1.0V = 1.01 V
Measured I4 1.2V = 1.22 V
Measured 2.5V = 2.51 V
Measured V1P2 DIG = 1.18 V
Measured V1P2 ANG = 1.18 V
Measured 1.2V BridgeX = 1.22 V
```

```
Measured 1.8V = 1.80 V
Measured 1.2V Standby = 1.20 V
All voltages OK
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, checkvoltages command
- [“Display Power Supply Status” on page 20](#)
- [“Display Gateway Environmental and Operational Data” on page 23](#)

▼ Display Internal Temperatures

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

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

```
FabMan@gateway_name->showtemps
Back temperature 29
Front temperature 29
SP temperature 47
Switch temperature 40, maxtemperature 42
Bridge-0 temperature 40, maxtemperature 41
Bridge-1 temperature 45, maxtemperature 46
All temperatures OK
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, showtemps command
- [“Display Fan Status” on page 22](#)
- [“Display Gateway Environmental and Operational Data” on page 23](#)

▼ Display Fan Status

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

- On the management controller, type.

```
FabMan@gateway_name->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
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `getfanspeed` command
- [“Display Gateway Environmental and Operational Data” on page 23](#)

▼ Display Gateway 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 gateway.

- On the management controller, type.

```
FabMan@gateway_name->env_test
Environment test started:
Starting Environment Daemon test:
Environment daemon running
Environment Daemon test returned OK
Starting Voltage test:
Voltage ECB OK
Measured 3.3V Main = 3.28 V
Measured 3.3V Standby = 3.37 V
Measured 12V = 12.06 V
Measured 5V = 5.03 V
Measured VBAT = 3.25 V
Measured 1.0V = 1.01 V
Measured I4 1.2V = 1.22 V
Measured 2.5V = 2.51 V
Measured V1P2 DIG = 1.18 V
Measured V1P2 ANG = 1.18 V
Measured 1.2V BridgeX = 1.22 V
Measured 1.8V = 1.80 V
Measured 1.2V Standby = 1.20 V
Voltage test returned OK
Starting PSU test:
```

```

PSU 0 present
PSU 1 present
PSU test returned OK
Starting Temperature test:
Back temperature 30.5
Front temperature 30.6
SP temperature 34.1
Switch temperature 48, maxtemperature 49
Bridge-0 temperature 40, maxtemperature 41
Bridge-1 temperature 45, maxtemperature 45
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
Bridge-0 OK
Bridge-1 OK
All Internal ibdevices OK
Onboard ibdevice test returned OK
Starting SSD test:
SSD test returned OK
Environment test PASSED
FabMan@gateway_name->

```

Related Information

- *Gateway Reference*, `env_test` command

▼ Display Flash Drive Information

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

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

```

FabMan@gateway_name->showdisk
Device Model      : 512MB NANDrive
Serial Number     : 0000000000A1PsB87FAa
Firmware Version: ND B642J

```

```
SMART supported : yes
SMART enabled  : yes

Average age (4 groups):
  group 0: 12817
  group 1:   533
  group 2:   715
  group 3:   954

Block erase count (4 groups):
  group 0: 12782181
  group 1:  532378
  group 2:  713173
  group 3:  952668
  Total  : 14980400

Total physical blocks: 4096

Estimated Life Usage      : 12.817 %

Estimated Remaining Life: 87.183 %

FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, showdisk command
- [“Display Gateway Environmental and Operational Data” on page 23](#)

▼ Display Free Memory and File System Space

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

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

```
FabMan@gateway_name->showfree -md
Memory:                305.6 MB   61.3%

Filesystems:
/dev/hda5 /             127.1 MB  32.2%
/dev/hda3 /var/log      10.1 MB  66.4%
/dev/hda2 /config       10.3 MB  67.6%
tmpfs     /dev/shm        249.2 MB 100.0%
tmpfs     /tmp           248.2 MB  99.6%
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, showfree command
- [“Display Gateway Environmental and Operational Data”](#) on page 23

▼ Display Chassis FRU ID

The showfruinfo command displays gateway chassis FRU ID information.

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

```
FabMan@gateway_name->showfruinfo
Sun_Man1R:
  UNIX_Timestamp32           : Fri Dec 14 01:24:28 2012
  Sun_Fru_Description        : ASSY,NM2-GW
  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            : 5111402
  Sun_Serial_Number          : 0110SJC-1010NG0040
  Serial_Number_Format       : 4V3F1-2Y2W2X4S
  Initial_HW_Dash_Level      : 03
  Initial_HW_Rev_Level       : 50
  Sun_Fru_Shortname          : NM2 gateway
  Sun_Hazard_Class_Code      : Y
  Sun_SpecPartNo             : 885-1655-01
Sun_FRU_LabelR:
  Sun_Serial_Number          : 1013AK208D
  FRU_Part_Dash_Number       : 541-4188-01
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, showfruinfo command
- [“Display Power Supply FRU ID”](#) on page 26

▼ Display Power Supply FRU ID

The showpsufpu command displays power supply FRU ID information.

- On the management controller, type.

```
FabMan@gateway_name->showpsufru slot
```

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

```
FabMan@gateway_name->showpsufru 0
Sun_SpecPartNo      : 885-1165-02
UNIX_Stamp32       : Sun Jan  3 15:35:39 2010
Sun_Part_Number    : 3002143
Sun_Serial_No      : BF15WA
Vendor_ID_Code     : 02a2
Initial_HW_Dash_Level : 02
PSU_Voltage 1     : 0x04b0 (1200)
PSU_Current 1    : 0x186a (6250)
PSU_Voltage 2     : 0x014a (330)
PSU_Current 2    : 0x012c (300)
Sun_Hazard_Class_Code : Y
IPMI_Board_Manufacturer : EMERSON
IPMI_Board_Product_Name : A237
IPMI_Board_Serial_Number: 1357ZHO-0952BF15WA
IPMI_Board_Part_Number : 300-2143-02
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, showpsufru command
- “Display Chassis FRU ID” on page 26

▼ Display Gateway Firmware Versions

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

- On the management controller, type.

```
FabMan@gateway_name->version
SUN DCS gw version: 2.1.2-1
Build time: Dec  7 2012 09:46:50
FPGA version: 0x34
SP board info:
Manufacturing Date: 2009.06.22
Serial Number: "NCD3R0391"
Hardware Revision: 0x0006
Firmware Revision: 0x0102
```

```
BIOS version: NOW1R112
BIOS date: 04/24/2009
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, version command
- *Gateway Remote Management*, version command

Oracle Auto Service Request

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

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

Related Information

- [“Troubleshooting the Gateway” on page 1](#)
- *Gateway Service*, detecting and managing faults

Controlling the Chassis

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

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

- [“Restart the Management Controller” on page 29](#)
- [“Change the root User Password” on page 29](#)

Related Information

- [“Controlling the I4 Switch Chip” on page 40](#)
- [“Controlling Gateway Ports and Parameters” on page 202](#)
- [“Controlling the InfiniBand Fabric” on page 97](#)

- [“Controlling the Subnet Manager” on page 56](#)
- [“Monitoring the Chassis” on page 19](#)

▼ Restart the Management Controller

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

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

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

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

The management controller power cycles and reinitializes.

Related Information

- [“Change the root User Password” on page 29](#)

▼ Change the root User Password

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

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

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

2. Type.

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

Related Information

- *Gateway Installation*, accessing the management controller
- [“Restart the Management Controller” on page 29](#)

Administering the I4 Switch Chip

These topics describe the administration of the I4 switch chip.

- [“Monitoring the I4 Switch Chip” on page 31](#)
- [“Controlling the I4 Switch Chip” on page 40](#)

Related Information

- [“Troubleshooting the Gateway” on page 1](#)
- [“Understanding Administrative Commands” on page 13](#)
- [“Administering the Chassis” on page 19](#)
- [“Administering Gateway Resources” on page 137](#)
- [“Administering the InfiniBand Fabric” on page 83](#)
- [“Administering the Subnet Manager” on page 49](#)

Monitoring the I4 Switch Chip

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

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

- [“Display the Switch Chip Port to QSFP Connector Mapping” on page 32](#)
- [“Locate a InfiniBand Device and Connector From the GUID and Port” on page 33](#)
- [“Locate a Switch Chip or Connector From the GUID and Port” on page 33](#)
- [“Display Switch Chip Boot Status” on page 34](#)
- [“Display Link Status” on page 35](#)
- [“Display Connectors Configured With Autodisable Functionality” on page 36](#)
- [“Display Switch Chip Port Status” on page 37](#)

- “Display Switch Chip Port Counters” on page 38

Related Information

- “Monitoring the Chassis” on page 19
- “Monitoring Gateway Resources” on page 172
- “Monitoring the InfiniBand Fabric” on page 83
- “Monitoring the Subnet Manager” on page 49
- “Controlling the I4 Switch Chip” on page 40

▼ Display the Switch Chip Port to QSFP Connector Mapping

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

- On the management controller, type.

```
FabMan@gateway_name->dcSPORT -port port
```

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

Alternatively, type.

```
FabMan@gateway_name->dcSPORT -connector connector
```

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

For example:

```
FabMan@gateway_name->dcSPORT -port 7  
DCS-GW Switch port 7 maps to connector 14A  
FabMan@gateway_name->dcSPORT -connector 14A  
Connector 14A maps to Switch port 7  
FabMan@gateway_name->
```

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

Related Information

- *Gateway Reference*, `dcSPORT` command
- [“Understanding Routing Through the Gateway”](#) on page 7

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

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

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

```
FabMan@gateway_name->dcSPORT -guid guid -port port
```

where:

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

For example:.

```
FabMan@gateway_name->dcSPORT -guid 0x00212856d102c0a0 -port 36
DCS-GW Switch port 36 maps to connector 6B
FabMan@gateway_name->
```

Related Information

- [“Switch GUIDs Overview”](#) on page 10
- *Gateway Reference*, `dcSPORT` command

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

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

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

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

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

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

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

For example:

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

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

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

From “Switch GUIDs Overview” on page 10, the `c0` means that the node is inside of a Sun Network QDR InfiniBand Gateway Switch. The `a` means that the node is located in I4 switch chip A, the only switch chip. The port is 21.

4. Use the tables in “Understanding Routing Through the Gateway” on page 7 to cross-reference the routing between ports and to connections.

For example, using “Switch Chip Port to QSFP Connectors and Link LED Routes” on page 8, I4 switch chip, port 21 routes to connector 1B.

Related Information

- “Switch GUIDs Overview” on page 10
- *Gateway Reference*, `dcSPORT` command

▼ Display Switch Chip Boot Status

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

- On the management controller, type.

```
FabMan@gateway_name->checkboot
Switch OK
Bridge-0 OK
Bridge-1 OK
All Internal ibdevices OK
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, checkboot command
- “Display Gateway Environmental and Operational Data” on page 23
- *Gateway Reference*, ibdevreset command

▼ Display Link Status

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

- On the management controller, type.

```
FabMan@gateway_name->listlinkup connected
Connector 2A Present <-> Switch Port 24 up (Enabled)
Connector 3A Present <-> Switch Port 26 up (Enabled)
Connector 6A Present <-> Switch Port 35 up (Enabled)
Connector 8A Present <-> Switch Port 31 up (Enabled)
Connector 11A Present <-> Switch Port 12 down (Enabled)
Connector 13A Present <-> Switch Port 9 down (Enabled)
Connector 14A Present <-> Switch Port 7 down (Enabled)
Connector 0A-ETH Present
  Bridge-0 Port 0A-ETH-1 (Bridge-0-2) up (Enabled)
  Bridge-0 Port 0A-ETH-2 (Bridge-0-2) up (Enabled)
  Bridge-0 Port 0A-ETH-3 (Bridge-0-1) up (Enabled)
  Bridge-0 Port 0A-ETH-4 (Bridge-0-1) up (Enabled)
Connector 1A-ETH Present
  Bridge-1 Port 1A-ETH-1 (Bridge-1-2) up (Enabled)
  Bridge-1 Port 1A-ETH-2 (Bridge-1-2) up (Enabled)
  Bridge-1 Port 1A-ETH-3 (Bridge-1-1) up (Enabled)
  Bridge-1 Port 1A-ETH-4 (Bridge-1-1) up (Enabled)
Connector 9B Present <-> Switch Port 13 up (Enabled)
FabMan@gateway_name->
```

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

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

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

Related Information

- *Gateway Reference*, `listlinkup` command
- [“Display the Link Status of a Node” on page 88](#)

▼ Display Connectors Configured With Autodisable Functionality

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

- On the management controller, type.

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

Related Information

- *Gateway Reference*, autodisable command
- [“Autodisabling Ports and Connectors” on page 43](#)

▼ Display Switch Chip Port Status

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

- On the management controller, type.

```
FabMan@gateway_name->getportstatus connector | ibdev port
```

where:

- *connector* is the name of the connector (0A to 15A, 0A-ETH, 1A-ETH, 0B to 15B).

Note – The subordinate names of the 10GbE connectors, 0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4, are also valid *connector* names.

- *ibdev* is the InfiniBand device name (Switch, Bridge-0-0, Bridge-0-1, Bridge-1-0, Bridge-1-1)
- *port* is the number of the port (1 to 36).

For example:

```
FabMan@gateway_name->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
FabMan@gateway_name->
```

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

Related Information

- *Gateway Reference*, `getportstatus` command
- [“Display the Link Status of a Node” on page 88](#)

▼ Display Switch Chip Port Counters

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

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

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

where:

- *port* is the number of the port (1 to 36).
- *connector* is the name of the connector (0A to 15A, 0A-ETH, 1A-ETH, 0B to 15B).

Note – The subordinate names of the 10GbE connectors, 0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4 are also valid *connector* names.

For example, for port 1:

```
FabMan@gateway_name->getportcounters 1
Port counters for internal Switch port 1
SymbolErrors.....0
LinkRecovers.....0
LinkDowned.....0
RcvErrors.....0
RcvRemotePhysErrors.....0
RcvSwRelayErrors.....0
XmtDiscards.....0
XmtConstraintErrors.....0
RcvConstraintErrors.....0
LinkIntegrityErrors.....0
ExcBufOverrunErrors.....0
VL15Dropped.....0
XmtData.....909556
RcvData.....737760
XmtPkts.....12830
RcvPkts.....10368
XmtWait.....0
FabMan@gateway_name->
```

For example, for connector 1A-ETH-1:

```
FabMan@gateway_name->getportcounters 1A-ETH-1
Port counters for connector 1A-ETH-1 Bridge-1 port Bridge-1-2
RX bytes.....8832
RX packets.....69
RX Jumbo packets.....0
RX unicast packets.....0
RX multicast packets.....69
RX broadcast packets.....0
RX no buffer.....0
RX CRC.....0
RX runt.....0
RX errors.....0
TX bytes.....8832
TX packets.....69
TX Jumbo packets.....0
TX unicast packets.....0
TX multicast packets.....69
```

```
TX broadcast packets.....0
TX errors.....0
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `getportcounters` command
- [“Display Counters for a Node”](#) on page 89

Controlling the I4 Switch Chip

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

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

- [“Disable a Switch Chip Port”](#) on page 40
- [“Enable a Switch Chip Port”](#) on page 41
- [“Autodisabling Ports and Connectors”](#) on page 43

Related Information

- [“Controlling the Chassis”](#) on page 28
- [“Controlling Gateway Ports and Parameters”](#) on page 202
- [“Controlling the InfiniBand Fabric”](#) on page 97
- [“Controlling the Subnet Manager”](#) on page 56
- [“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.

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

- On the management controller, type.

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

where:

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

Alternatively, type.

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

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

For example:

```
FabMan@gateway_name->disableswitchport Switch 14
Disable Switch port 14
Initial PortInfo:
# Port info: DR path slid 65535; dlid 65535; 0 port 14
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: DR path slid 65535; dlid 65535; 0 port 14
LinkState:.....Down
PhysLinkState:.....Disabled
FabMan@gateway_name->
```

Related Information

- [Gateway Reference](#), disableswitchport command
- [“Disable or Enable an External Port” on page 104](#)
- [“Enable a Switch Chip Port” on page 41](#)
- [“Autodisabling Ports and Connectors” on page 43](#)

▼ Enable a Switch Chip Port

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

- On the management controller, type.

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

where:

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

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

Alternatively, type.

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

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

For example:

```
FabMan@gateway_name->enableswitchport Switch 14
Enable Switch port 14
Initial PortInfo:
# Port info: DR path slid 65535; dlid 65535; 0 port 14
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: DR path slid 65535; dlid 65535; 0 port 14
LinkState:.....Down
PhysLinkState:.....PortConfigurationTraining
FabMan@gateway_name->
```

Related Information

- [Gateway Reference](#), enableswitchport command
- [“Disable or Enable an External Port”](#) on page 104
- [“Disable a Switch Chip Port”](#) on page 40
- [“Reenable Autodisabled Connectors”](#) on page 45

Autodisable Ports and Connectors

These topics instruct how to manage autodisable ports and connectors.

- [“Autodisable Functionality” on page 43](#)
- [“Add Connectors to the Autodisable List” on page 44](#)
- [“Reenable Autodisable Connectors” on page 45](#)
- [“Delete Connectors From the Autodisable List” on page 46](#)

Related Information

- [“Disable a Switch Chip Port” on page 40](#)
- [“Enable a Switch Chip Port” on page 41](#)

Autodisable Functionality

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

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

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

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

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

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

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

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

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

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

Related Information

- *Gateway Reference*, `autodisable` command
- *Gateway Reference*, `enableswitchport` command
- [“Add Connectors to the Autodisable List” on page 44](#)
- [“Reenable Autodisable Connectors” on page 45](#)
- [“Delete Connectors From the Autodisable List” on page 46](#)
- *Gateway Service*, servicing cables

▼ Add Connectors to the Autodisable List

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

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

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

where:

- *connector* is the name of the connector (0A to 15A and 0B to 15B).
- *cause* is the reason to disable the connector:
 - H – The error rate is higher than what the InfiniBand specification permits.
 - L – The link speed or width is less than optimal (less than 4x QDR).

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

```
FabMan@gateway_name->autodisable all
```

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

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

Related Information

- [Gateway Reference, autodisable command](#)
- [“Autodisable Functionality” on page 43](#)
- [“Reenable Autodisabled Connectors” on page 45](#)
- [“Delete Connectors From the Autodisable List” on page 46](#)

▼ **Reenable Autodisabled Connectors**

1. Investigate why the connector was autodisabled.

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

2. On the management controller, type.

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

where:

- *port* is the number of the port (5 to 36).

- *connector* is the number of the QSFP connector (0A to 15B).

For example:

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

3. If you did not replace the cable in [Step 1](#), you must retrain the links for the port.

a. Disable the port.

See [“Disable a Switch Chip Port”](#) on page 40.

b. Reenable the port.

Do not use the `--automatic` option when you reenable the port. See [“Enable a Switch Chip Port”](#) on page 41.

Related Information

- [Gateway Reference](#), `enableswitchport` command
- [“Autodisable Functionality”](#) on page 43
- [“Add Connectors to the Autodisable List”](#) on page 44
- [“Delete Connectors From the Autodisable List”](#) on page 46

▼ Delete Connectors From the Autodisable List

- On the management controller, type.

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

where:

- *connector* is the name of the connector (0A to 15A and 0B to 15B).
- *cause* is the reason to disable the connector:
 - H – The error rate is higher than what the InfiniBand specification permits.

- L – The link speed or width is less than optimal (less than 4x QDR).
Alternatively, to delete all connectors, type.

```
FabMan@gateway_name->autodisable none
```

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

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

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

Related Information

- *Gateway Reference*, `autodisable` command
- “Autodisable Functionality” on page 43
- “Add Connectors to the Autodisable List” on page 44
- “Reenable Autodisabled Connectors” on page 45

Administering the Subnet Manager

These topics describe the administration of the Subnet Manager.

- “Monitoring the Subnet Manager” on page 49
- “Controlling the Subnet Manager” on page 56
- “Configuring Secure Fabric Management” on page 64

Related Information

- “Troubleshooting the Gateway” on page 1
- “Understanding Administrative Commands” on page 13
- “Administering the Chassis” on page 19
- “Administering the I4 Switch Chip” on page 31
- “Administering Gateway Resources” on page 137
- “Administering the InfiniBand Fabric” on page 83

Monitoring the Subnet Manager

These topics describe how to monitor the Subnet Manager.

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

- “Display Subnet Manager Status” on page 50
- “Display the Subnet Manager’s Location History” on page 51
- “Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51
- “Display the Secret M_Key Configuration” on page 52
- “Verify M_Key Persistence” on page 53
- “Verify the Subnet Manager Configuration Integrity” on page 54

- “Display the Subnet Manager Log” on page 55

Related Information

- “Monitoring the Chassis” on page 19
- “Monitoring the I4 Switch Chip” on page 31
- “Monitoring Gateway Resources” on page 172
- “Monitoring the InfiniBand Fabric” on page 83
- “Controlling the Subnet Manager” on page 56
- “Configuring Secure Fabric Management” on page 64

▼ Display Subnet Manager Status

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

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

```
FabMan@gateway_name->getmaster
Local SM enabled and running, state MASTER
20120831 09:45:17 Master SubnetManager on sm lid 1 sm guid 0x3ba01234567a0a0 :
SUN DCS 36P QDR mnm2 10.172.144.66
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `getmaster` command
- “Display the Subnet Manager’s Location History” on page 51
- “Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51
- “Verify the Subnet Manager Configuration Integrity” on page 54
- “Display the Subnet Manager Log” on page 55

▼ Display the Subnet Manager’s Location History

- On the management controller, type.

```
FabMan@gateway_name->getmaster -l
Local SM enabled and running, state MASTER
SM monitoring history:
20130118 07:25:37 No Master SubnetManager seen in the system
20130118 07:25:55 Master SubnetManager on sm lid 1 sm guid 0x3baabba115c0a0 :
SUN IB QDR GW switch o4nm2-gw-3 10.172.144.71
20130118 07:27:13 Master SubnetManager on sm lid 10 sm guid 0x21284694e5a0a0 :
SUN DCS 36P QDR o4nm2-36p-1 10.172.144.65
.
.
.
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `getmaster` command
- [“Display Subnet Manager Status” on page 50](#)
- [“Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51](#)
- [“Verify the Subnet Manager Configuration Integrity” on page 54](#)
- [“Display the Subnet Manager Log” on page 55](#)

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

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

1. On the management controller, type.

```
FabMan@gateway_name->setsmpriority list
Current SM settings:
smpriority 8
controlled_handover TRUE
subnet_prefix 0xfe80000000000000
```

```
M_Key 0xababababe0001
Routing engine FatTree
FabMan@gateway_name->
```

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

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

See:

- “Set the Subnet Manager Priority” on page 56
- “Set the Subnet Manager Prefix” on page 57
- “Enable Subnet Manager Controlled Handover” on page 58
- “Set the Subnet Manager Readable Management Key” on page 59
- “Set the Subnet Manager Routing Algorithm” on page 61

Related Information

- *Gateway Reference*, `setsmpriority` command
- “Controlling the Subnet Manager” on page 56
- “Display Subnet Manager Status” on page 50
- “Display the Subnet Manager’s Location History” on page 51
- “Verify the Subnet Manager Configuration Integrity” on page 54
- “Display the Subnet Manager Log” on page 55

▼ Display the Secret M_Key Configuration

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

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

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

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

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

```
FabMan@gateway_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=o4nm2-gw-6
# time=15 Sep 04:03:00
# checksum=0
#! commit_number : 4
Mkey                Untrusted Mkey          Smkey                    Attribute
-----
0x00abcdefabcdef01  0x1aa45124fee612ae      0x15fc26aea300f831
0x00abcdefabcdef02  0x4ccd8230de6cd348      0x3fc7e6ad701a8a2a
0x00abcdefabcdef03  0x9baa1debcc74de5e      0x1b253003600d137b      C
FabMan@gateway_name->
```

In the output, the Attribute heading identifies:

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

Related Information

- *Gateway Reference*, smsubnetprotection command
- [“Configuring Secret M_Keys” on page 67](#)

▼ Verify M_Key Persistence

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

1. On the management controller, type.

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

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

See “Enable or Disable Secret M_Key Persistence” on page 76.

Related Information

- *Gateway Reference*, `localmkeypersistence` command
- “Enable or Disable Secret M_Key Persistence” on page 76

▼ Verify the Subnet Manager Configuration Integrity

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

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

● Verify the Subnet Manager configuration.

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

where:

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

For example, when the Subnet Manager has priority 6:

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

Related Information

- *Gateway Reference*, `smconfigtest` command
- “Display Subnet Manager Status” on page 50
- “Display the Subnet Manager’s Location History” on page 51
- “Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51
- “Display the Subnet Manager Log” on page 55

▼ Display the Subnet Manager Log

1. On the management controller, type.

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

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

2. Tap the space bar to display the next screen of the log.

3. Press the Q key to quit.

Related Information

- *Gateway Reference*, showsmlog command
- [“Display Subnet Manager Status” on page 50](#)
- [“Display the Subnet Manager’s Location History” on page 51](#)
- [“Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51](#)
- [“Verify the Subnet Manager Configuration Integrity” on page 54](#)

Controlling the Subnet Manager

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

These topics describe how to control the Subnet Manager:

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

- “Set the Subnet Manager Priority” on page 56
- “Set the Subnet Manager Prefix” on page 57
- “Enable Subnet Manager Controlled Handover” on page 58
- “Set the Subnet Manager Readable Management Key” on page 59
- “Clear the Subnet Manager Readable Management Key” on page 60
- “Set the Subnet Manager Routing Algorithm” on page 61
- “Enable the Subnet Manager” on page 62
- “Disable the Subnet Manager” on page 63

Related Information

- “Controlling the Chassis” on page 28
- “Controlling the I4 Switch Chip” on page 40
- “Controlling Gateway Ports and Parameters” on page 202
- “Controlling the InfiniBand Fabric” on page 97
- “Monitoring the Subnet Manager” on page 49
- “Controlling the Subnet Manager” on page 56

▼ Set the Subnet Manager Priority

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

1. **On the management controller, disable the Subnet Manager.**
See “Disable the Subnet Manager” on page 63.
2. **Set the Subnet Manager priority.**

```
FabMan@gateway_name->setsmpriority priority
```

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

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

3. **Enable the Subnet Manager.**
See “Enable the Subnet Manager” on page 62.

Related Information

- *Gateway Reference*, setsmpriority command
- “Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51
- “Set the Subnet Manager Prefix” on page 57
- “Enable Subnet Manager Controlled Handover” on page 58
- “Set the Subnet Manager Readable Management Key” on page 59
- “Clear the Subnet Manager Readable Management Key” on page 60
- “Set the Subnet Manager Routing Algorithm” on page 61
- “Enable the Subnet Manager” on page 62
- “Disable the Subnet Manager” on page 63

▼ Set the Subnet Manager Prefix

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

1. **On the management controller, disable the Subnet Manager.**
See “Disable the Subnet Manager” on page 63.

2. Set the Subnet Manager prefix.

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

3. Enable the Subnet Manager.

See “Enable the Subnet Manager” on page 62.

Related Information

- *Gateway Reference*, `setsubnetprefix` command
- “Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51
- “Set the Subnet Manager Priority” on page 56
- “Enable Subnet Manager Controlled Handover” on page 58
- “Set the Subnet Manager Readable Management Key” on page 59
- “Clear the Subnet Manager Readable Management Key” on page 60
- “Set the Subnet Manager Routing Algorithm” on page 61
- “Enable the Subnet Manager” on page 62
- “Disable the Subnet Manager” on page 63

▼ Enable Subnet Manager Controlled Handover

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

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

See “Disable the Subnet Manager” on page 63.

2. Enable controlled handover.

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


2. Set the readable management key.

```
FabMan@gateway_name->setsmmkey m_key
```

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

For example:

```
FabMan@gateway_name->setsmmkey 0xabababe0001
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xabababe
M_Key 0xabababe0001
Routing engine FatTree
FabMan@gateway_name->
```

3. Enable the Subnet Manager.

See “Enable the Subnet Manager” on page 62.

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

Related Information

- *Gateway Reference*, `setsmmkey` command
- “Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm” on page 51
- “Set the Subnet Manager Priority” on page 56
- “Set the Subnet Manager Prefix” on page 57
- “Enable Subnet Manager Controlled Handover” on page 58
- “Clear the Subnet Manager Readable Management Key” on page 60
- “Set the Subnet Manager Routing Algorithm” on page 61
- “Enable the Subnet Manager” on page 62
- “Disable the Subnet Manager” on page 63

▼ Clear the Subnet Manager Readable Management Key

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

2. Set the routing algorithm.

```
FabMan@gateway_name->setsmrouting fattree minhop
```

For example:

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

3. Enable the Subnet Manager.

See “Enable the Subnet Manager” on page 62.

Related Information

- *Gateway Reference*, setsmrouting command
- “Set the Subnet Manager Priority” on page 56
- “Set the Subnet Manager Prefix” on page 57
- “Enable Subnet Manager Controlled Handover” on page 58
- “Set the Subnet Manager Readable Management Key” on page 59
- “Clear the Subnet Manager Readable Management Key” on page 60
- “Enable the Subnet Manager” on page 62
- “Disable the Subnet Manager” on page 63

▼ Enable the Subnet Manager

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

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

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

Related Information

- *Gateway Reference*, `enablesm` command
- “Disable the Subnet Manager” on page 63
- “Set the Subnet Manager Priority” on page 56
- “Set the Subnet Manager Prefix” on page 57
- “Enable Subnet Manager Controlled Handover” on page 58
- “Set the Subnet Manager Readable Management Key” on page 59
- “Clear the Subnet Manager Readable Management Key” on page 60
- “Set the Subnet Manager Routing Algorithm” on page 61
- “Disable the Subnet Manager” on page 63

▼ Disable the Subnet Manager

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

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

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

Related Information

- *Gateway Reference*, `disablesm` command
- “Enable the Subnet Manager” on page 62
- “Set the Subnet Manager Priority” on page 56
- “Set the Subnet Manager Prefix” on page 57
- “Enable Subnet Manager Controlled Handover” on page 58
- “Set the Subnet Manager Readable Management Key” on page 59
- “Clear the Subnet Manager Readable Management Key” on page 60
- “Set the Subnet Manager Routing Algorithm” on page 61
- “Enable the Subnet Manager” on page 62

Configuring Secure Fabric Management

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

- “Secret M_Key Overview” on page 64
- “smsubnetprotection Command” on page 65
- “Prepare for Secret M_Key Functionality” on page 66
- “Configuring Secret M_Keys” on page 67
- “Set the Replication Password” on page 76
- “Enable or Disable Secret M_Key Persistence” on page 76
- “Enable or Disable Secret M_Key Functionality” on page 77
- “Considering Situations In Regard to Secret M_Key” on page 78

Related Information

- “Monitoring the Subnet Manager” on page 49
- “Controlling the Subnet Manager” on page 56

Secret M_Key Overview

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

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

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

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

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

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

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

Configuring secret M_Key for the first time requires:

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

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

Related Information

- [“smsubnetprotection Command” on page 65](#)
- [“Prepare for Secret M_Key Functionality” on page 66](#)
- [“Configuring Secret M_Keys” on page 67](#)
- [“Set the Replication Password” on page 76](#)
- [“Enable or Disable Secret M_Key Persistence” on page 76](#)
- [“Enable or Disable Secret M_Key Functionality” on page 77](#)
- [“Considering Situations In Regard to Secret M_Key” on page 78](#)

smsubnetprotection Command

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

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

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

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

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

Related Information

- [“Secret M_Key Overview” on page 64](#)
- [“Prepare for Secret M_Key Functionality” on page 66](#)
- [“Configuring Secret M_Keys” on page 67](#)
- [“Set the Replication Password” on page 76](#)
- [“Enable or Disable Secret M_Key Persistence” on page 76](#)
- [“Enable or Disable Secret M_Key Functionality” on page 77](#)
- [“Considering Situations In Regard to Secret M_Key” on page 78](#)

▼ Prepare for Secret M_Key Functionality

1. **Update the `smnodes` list for all `smnodes` in the InfiniBand fabric.**
See [“Update the `smnodes` List” on page 134](#).
2. **(Option) Configure IPoIB functionality.**
See [“Create the IPoIB Interface” on page 135](#).
3. **Create and propagate the fabric element configuration.**
See [“Create a Fabric Element Configuration” on page 120](#).
4. **Set the replication password.**
See [“Set the Replication Password” on page 76](#).
5. **Enable M_Key persistence.**
See [“Enable or Disable Secret M_Key Persistence” on page 76](#).
6. **Enable secret M_Key functionality.**
See [“Enable or Disable Secret M_Key Functionality” on page 77](#).

7. Configure the secret M_Keys.

See “Configuring Secret M_Keys” on page 67.

Related Information

- “Secret M_Key Overview” on page 64
- “smsubnetprotection Command” on page 65
- “Configuring Secret M_Keys” on page 67
- “Set the Replication Password” on page 76
- “Enable or Disable Secret M_Key Persistence” on page 76
- “Enable or Disable Secret M_Key Functionality” on page 77
- “Considering Situations In Regard to Secret M_Key” on page 78

Configuring Secret M_Keys

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

- “Create the Secret M_Key Configuration” on page 67
- “Add Secret M_Keys to the Configuration” on page 69
- “Set the Secret M_Key to Use” on page 72
- “Delete Secret M_Keys From the Configuration” on page 73
- “Configure a Local Secret M_Key” on page 75

Related Information

- “Secret M_Key Overview” on page 64
- “smsubnetprotection Command” on page 65
- “Prepare for Secret M_Key Functionality” on page 66
- “Set the Replication Password” on page 76
- “Enable or Disable Secret M_Key Persistence” on page 76
- “Enable or Disable Secret M_Key Functionality” on page 77
- “Considering Situations In Regard to Secret M_Key” on page 78

▼ Create the Secret M_Key Configuration

1. Prepare for secret M_Key functionality.

See “Prepare for Secret M_Key Functionality” on page 66.

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

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

3. Display the active secret M_Key configuration.

See “Display the Secret M_Key Configuration” on page 52.

4. Add an M_Key to the configuration.

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

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

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

For example:

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

5. Select the M_Key to use.

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

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

For example:

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

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

6. Display the pending M_Keys.

```
FabMan@gateway_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=o4nm2-gw-6
# time=15 Sep 03:38:33
# checksum=0
#! commit_number : 2
Mkey                               Untrusted Mkey                Smkey                           Attribute
-----                               -
0x00abcdefabcdef01                0x1aa45124fee612ae            0x15fc26aea300f831
0x00abcdefabcdef02                0x4ccd8230de6cd348            0x3fc7e6ad701a8a2a          S
FabMan@gateway_name->
```

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

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

```
FabMan@gateway_name->smsubnetprotection commit
FabMan@gateway_name->
```

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

Related Information

- *Gateway Reference*, smsubnetprotection command
- [“Add Secret M_Keys to the Configuration”](#) on page 69
- [“Set the Secret M_Key to Use”](#) on page 72
- [“Delete Secret M_Keys From the Configuration”](#) on page 73
- [“Configure a Local Secret M_Key”](#) on page 75

▼ Add Secret M_Keys to the Configuration

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

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

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

2. Display the active secret M_Key configuration.

See “Display the Secret M_Key Configuration” on page 52.

3. Add an M_Key to the configuration.

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

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

For example:

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

4. Repeat [Step 3](#) for any additional M_Keys you want to add to the configuration.
5. (Optional) If you will use one of the just added secret M_Keys, set the M_Key to use.

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

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

For example:

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

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

6. Display the pending M_Keys.

```
FabMan@gateway_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=o4nm2-gw-6
# time=15 Sep 03:47:12
# checksum=0
#! commit_number : 3
Mkey                               Untrusted Mkey                   Smkey                             Attribute
-----                               -
0x00abcdefabcdef01                0x1aa45124fee612ae               0x15fc26aea300f831
0x00abcdefabcdef02                0x4ccd8230de6cd348               0x3fc7e6ad701a8a2a           C
0x00abcdefabcdef03                0x9baa1debcc74de5e               0x1b253003600d137b           S
FabMan@gateway_name->
```

In the output, the Attribute heading identifies:

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

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

```
FabMan@gateway_name->smsubnetprotection commit
FabMan@gateway_name->
```

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

Related Information

- *Gateway Reference*, `smsubnetprotection` command
- [“Create the Secret M_Key Configuration” on page 67](#)
- [“Set the Secret M_Key to Use” on page 72](#)
- [“Delete Secret M_Keys From the Configuration” on page 73](#)
- [“Configure a Local Secret M_Key” on page 75](#)

▼ Set the Secret M_Key to Use

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

```
FabMan@gateway_name->smsubnetprotection start  
FabMan@gateway_name->
```

2. Display the active secret M_Key configuration.
See “Display the Secret M_Key Configuration” on page 52.
3. Determine which M_Key you will set as the new current.
4. Set the current M_Key.

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

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

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

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

5. Display the pending M_Keys.

```
FabMan@gateway_name->smsubnetprotection list modified  
# File_format_version_number 1  
# Sun DCS IB mkey config file  
# This file is generated, do not edit  
# secretmkey=enabled  
# nodeid=o4nm2-gw-6  
# time=15 Sep 03:38:33  
# checksum=0  
#! commit_number : 2  
Mkey                Untrusted Mkey                Smkey                Attribute  
-----  
0x00abcdefabcdef01  0x1aa45124fee612ae            0x15fc26aea300f831  
0x00abcdefabcdef02  0x4ccd8230de6cd348            0x3fc7e6ad701a8a2a    S  
0x00abcdefabcdef03  0x9baa1debcc74de5e            0x1b253003600d137b    C  
FabMan@gateway_name->
```

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

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

```
FabMan@gateway_name->smsubnetprotection commit
FabMan@gateway_name->
```

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

Related Information

- *Gateway Reference*, smsubnetprotection command
- [“Create the Secret M_Key Configuration” on page 67](#)
- [“Add Secret M_Keys to the Configuration” on page 69](#)
- [“Delete Secret M_Keys From the Configuration” on page 73](#)
- [“Configure a Local Secret M_Key” on page 75](#)

▼ Delete Secret M_Keys From the Configuration

Note – If you delete the secret M_Key also used by a node temporarily unavailable to the fabric, when that node becomes available, there is a secret M_Key mismatch, and the node is isolated. See [“Configure a Local Secret M_Key” on page 75](#).

1. (Optional) If you want to delete the current secret M_Key in use, set the new current M_Key to one you will not delete.
See [“Set the Secret M_Key to Use” on page 72](#).
2. Display the active secret M_Key configuration.
See [“Display the Secret M_Key Configuration” on page 52](#).
3. Determine which M_Key you will delete.
4. Initiate a secret M_Key configuration session on the management controller that hosts the master Subnet Manager.

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

5. Delete the M_Key from the configuration.

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

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

For example:

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

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

7. Display the pending M_Keys.

```
FabMan@gateway_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=o4nm2-gw-6
# time=15 Sep 04:18:20
# checksum=0
#! commit_number : 5
Mkey                Untrusted Mkey                Smkey                Attribute
-----
0x00abcdefabcdef01  0x1aa45124fee612ae            0x15fc26aea300f831
0x00abcdefabcdef02  0x4ccd8230de6cd348            0x3fc7e6ad701a8a2a      C
FabMan@gateway_name->
```

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

```
FabMan@gateway_name->smsubnetprotection commit
FabMan@gateway_name->
```

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

9. Display the active secret M_Key configuration.

See “Display the Secret M_Key Configuration” on page 52.

Related Information

- *Gateway Reference*, smsubnetprotection command

- “Create the Secret M_Key Configuration” on page 67
- “Add Secret M_Keys to the Configuration” on page 69
- “Set the Secret M_Key to Use” on page 72
- “Configure a Local Secret M_Key” on page 75

▼ Configure a Local Secret M_Key

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

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

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

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

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

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

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

See:

- “Enable or Disable Secret M_Key Persistence” on page 76
- “Restart the Management Controller” on page 29

Related Information

- *Gateway Reference*, smsubnetprotection command
- “Create the Secret M_Key Configuration” on page 67
- “Add Secret M_Keys to the Configuration” on page 69

- “Set the Secret M_Key to Use” on page 72
- “Delete Secret M_Keys From the Configuration” on page 73

▼ Set the Replication Password

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

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

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

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

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

Related Information

- *Gateway Reference*, `smsubnetprotection` command
- “Secret M_Key Overview” on page 64
- “`smsubnetprotection` Command” on page 65
- “Prepare for Secret M_Key Functionality” on page 66
- “Configuring Secret M_Keys” on page 67
- “Enable or Disable Secret M_Key Persistence” on page 76
- “Enable or Disable Secret M_Key Functionality” on page 77
- “Considering Situations In Regard to Secret M_Key” on page 78

▼ Enable or Disable Secret M_Key Persistence

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

- On the management controller, type.

```
FabMan@gateway_name->localmkeypersistence state
```

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

```
FabMan@gateway_name->localmkeypersistence enable
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, localmkeypersistence command
- “Verify M_Key Persistence” on page 53
- “Secret M_Key Overview” on page 64
- “smsubnetprotection Command” on page 65
- “Prepare for Secret M_Key Functionality” on page 66
- “Configuring Secret M_Keys” on page 67
- “Set the Replication Password” on page 76
- “Enable or Disable Secret M_Key Functionality” on page 77
- “Considering Situations In Regard to Secret M_Key” on page 78

▼ Enable or Disable Secret M_Key Functionality

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

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

```
FabMan@gateway_name->smsubnetprotection state
```

where *state* is either enablesecretmkey or disablesecretmkey. For example:

```
FabMan@gateway_name->smsubnetprotection enablesecretmkey
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, smsubnetprotection command
- “Secret M_Key Overview” on page 64
- “smsubnetprotection Command” on page 65
- “Prepare for Secret M_Key Functionality” on page 66

- “Configuring Secret M_Keys” on page 67
- “Set the Replication Password” on page 76
- “Enable or Disable Secret M_Key Persistence” on page 76
- “Considering Situations In Regard to Secret M_Key” on page 78

Considering Situations In Regard to Secret M_Key

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

- “Add a New Subnet Manager to the Secret M_Key Fabric” on page 78
- “Merge Subnets Into One Secret M_Key Fabric” on page 79
- “Secret M_Key and Firmware Upgrade and Downgrade Guidelines” on page 80

Related Information

- “Secret M_Key Overview” on page 64
- “smsubnetprotection Command” on page 65
- “Prepare for Secret M_Key Functionality” on page 66
- “Configuring Secret M_Keys” on page 67
- “Set the Replication Password” on page 76
- “Enable or Disable Secret M_Key Persistence” on page 76
- “Enable or Disable Secret M_Key Functionality” on page 77

▼ Add a New Subnet Manager to the Secret M_Key Fabric

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

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

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

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

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

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

3. **Verify that the Subnet Manager is disabled.**
See [“Disable the Subnet Manager”](#) on page 63.
4. **Set the Subnet Manager Priority to a low value, so that it is a standby Subnet Manager.**
See [“Set the Subnet Manager Priority”](#) on page 56.
5. **Enable the Subnet Manager.**
See [“Enable the Subnet Manager”](#) on page 62.
6. **If you installed a new switch or gateway, attach the signal and InfiniBand cables now.**
Refer to *Switch Installation* or *Gateway Installation*.
7. **Prepare the new Subnet Manager for M_Key functionality.**
See [“Prepare for Secret M_Key Functionality”](#) on page 66.
8. **Optional, change the priority of the new Subnet Manager.**
See [“Set the Subnet Manager Priority”](#) on page 56.

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

Related Information

- [“Merge Subnets Into One Secret M_Key Fabric”](#) on page 79
- [“Secret M_Key and Firmware Upgrade and Downgrade Guidelines”](#) on page 80

▼ Merge Subnets Into One Secret M_Key Fabric

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

1. **If the firmware on the switches and gateways are less than version 2.1, upgrade them.**
Refer to *Switch Remote Management* or *Gateway Remote Management*, upgrading the firmware.
2. **Configure identical secret M_Key lists on each subnet.**
See [“Prepare for Secret M_Key Functionality”](#) on page 66 and [“Create the Secret M_Key Configuration”](#) on page 67.

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

3. **Determine which Subnet Manager is to be the master Subnet Manager and which are to be the standby Subnet Managers.**

4. **Set the priority for each Subnet Manager.**

See “[Set the Subnet Manager Priority](#)” on page 56.

5. **Physically merge the subnets into one fabric.**

Related Information

- “[Add a New Subnet Manager to the Secret M_Key Fabric](#)” on page 78
- “[Secret M_Key and Firmware Upgrade and Downgrade Guidelines](#)” on page 80

Secret M_Key and Firmware Upgrade and Downgrade Guidelines

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

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

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

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

Related Information

- [“Add a New Subnet Manager to the Secret M_Key Fabric” on page 78](#)
- [“Merge Subnets Into One Secret M_Key Fabric” on page 79](#)

Administering the InfiniBand Fabric

These topics describe the administration of the InfiniBand fabric.

- [“Monitoring the InfiniBand Fabric”](#) on page 83
- [“Controlling the InfiniBand Fabric”](#) on page 97

Related Information

- [“Troubleshooting the Gateway”](#) on page 1
- [“Understanding Administrative Commands”](#) on page 13
- [“Administering the Chassis”](#) on page 19
- [“Administering the I4 Switch Chip”](#) on page 31
- [“Administering the Subnet Manager”](#) on page 49
- [“Administering Gateway Resources”](#) on page 137

Monitoring the InfiniBand Fabric

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

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

- [“Displaying Information About Nodes”](#) on page 84
- [“Displaying Fabric Functionality Configuration Information”](#) on page 93

Related Information

- [“Monitoring the Chassis”](#) on page 19
- [“Monitoring the I4 Switch Chip”](#) on page 31
- [“Monitoring Gateway Resources”](#) on page 172

- “Monitoring the Subnet Manager” on page 49
- “Monitoring Gateway Resources” on page 172
- “Controlling the InfiniBand Fabric” on page 97

Displaying Information About Nodes

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

- “Identify All Switches in the Fabric” on page 84
- “Identify All CAs in the Fabric” on page 85
- “Display the InfiniBand Fabric Topology (Simple)” on page 86
- “Display the InfiniBand Fabric Topology (Detailed)” on page 86
- “Display a Route Through the Fabric” on page 87
- “Display the Link Status of a Node” on page 88
- “Display Counters for a Node” on page 89
- “Display Low-Level Detailed Information About a Node” on page 90
- “Display Low-Level Detailed Information About a Port” on page 91

Related Information

- “Investigating Nodes” on page 98
- “Controlling Nodes” on page 101
- “Displaying Fabric Functionality Configuration Information” on page 93

▼ Identify All Switches in the Fabric

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

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

- On the management controller, type.

```
FabMan@gateway_name->ibswitches
Switch : 0x00212856cfe2c0a0 ports 36 "SUN IB QDR GW switch mnm34-98" enhanced
port 0 lid 6 lmc 0
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `ibswitches` command
- *Gateway Reference*, `ibhosts` command
- [“Identify All CAs in the Fabric” on page 85](#)

▼ Identify All CAs in the Fabric

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

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

1. On the management controller, type.

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

```
FabMan@gateway_name->ibhosts
Ca      : 0x00212856cd22c000 ports 2 "SUN IB QDR GW switch mnm34-97 Bridge 0"
Ca      : 0x00212856cd22c040 ports 2 "SUN IB QDR GW switch mnm34-97 Bridge 1"
Ca      : 0x0002c903000891aa ports 2 "mnm34-54 HCA-1"
Ca      : 0x00212800013ece9e ports 2 "mnm34-55 HCA-1"
Ca      : 0x0003ba000100e370 ports 2 "mnm34-60 HCA-1"
.
.
.
FabMan@gateway_name->
```

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

Related Information

- *Gateway Reference*, `ibhosts` command
- *Gateway Reference*, `ibswitches` command
- [“Identify All Switches in the Fabric” on page 84](#)

▼ Display the InfiniBand Fabric Topology (Simple)

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

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

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

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

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

Related Information

- *Gateway Reference*, `showtopology` command
- [“Determine Changes to the InfiniBand Fabric Topology” on page 99](#)
- [“Perform Comprehensive Diagnostics for the Entire Fabric” on page 98](#)

▼ Display the InfiniBand Fabric Topology (Detailed)

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

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

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

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

Related Information

- *Gateway Reference*, `ibnetdiscover` command
- [“Perform Comprehensive Diagnostics for the Entire Fabric” on page 98](#)

▼ Display a Route Through the Fabric

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

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

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

where:

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

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

For example:

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

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

For this example:

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

Related Information

- *Gateway Reference*, `ibtracert` command

▼ Display the Link Status of a Node

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

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

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

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

For example:

```
FabMan@gateway_name->ibportstate 15 23
PortInfo:
# Port info: Lid 15 port 23
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 15; dlid 65535; 0,23
LinkState:.....Active
PhysLinkState:.....LinkUp
LinkWidthSupported:.....1X or 4X
LinkWidthEnabled:.....1X or 4X
LinkWidthActive:.....4X
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.....10.0 Gbps
FabMan@gateway_name->
```

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

Related Information

- *Gateway Reference*, `ibportstate` command
- “Display Switch Chip Port Status” on page 37
- “Display Link Status” on page 35
- “Find 1x, SDR, or DDR Links in the Fabric” on page 100
- “Set Port Speed” on page 103

▼ Display Counters for a Node

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

- On the management controller, type.

```
FabMan@gateway_name->perfquery lid port
```

where:

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

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

For example:

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

```
FabMan@gateway_name->perfquery 15 23
# Port counters: Lid 15 port 23
PortSelect:.....23
CounterSelect:.....0x1b01
SymbolErrors:.....0
.
.
.
VL15Dropped:.....0
XmtData:.....20232
RcvData:.....20232
XmtPkts:.....281
RcvPkts:.....281
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `perfquery` command
- “Display Switch Chip Port Counters” on page 38
- “Clear Data and Error Counters” on page 102

▼ Display Low-Level Detailed Information About a Node

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

- On the management controller, type.

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

where *lid* is the LID of the node.

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

```
FabMan@gateway_name->smpquery switchinfo 15
# Switch info: Lid 15
LinearFdbCap:.....49152
RandomFdbCap:.....0
McastFdbCap:.....4096
LinearFdbTop:.....16
DefPort:.....0
DefMcastPrimPort:.....255
DefMcastNotPrimPort:.....255
LifeTime:.....18
StateChange:.....0
LidsPerPort:.....0
PartEnforceCap:.....32
InboundPartEnf:.....1
OutboundPartEnf:.....1
FilterRawInbound:.....1
FilterRawOutbound:.....1
EnhancedPort0:.....1
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `smpquery` command
- [“Display Low-Level Detailed Information About a Port” on page 91](#)

▼ Display Low-Level Detailed Information About a Port

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

- On the management controller, type.

```
FabMan@gateway_name->smpquery portinfo lid port
```

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

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

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

```
FabMan@gateway_name->smpquery portinfo 15 23
# Port info: Lid 15 port 23
Mkey:.....0x0000000000000000
GidPrefix:.....0x0000000000000000
Lid:.....0x0000
SMLid:.....0x0000
CapMask:.....0x0
DiagCode:.....0x0000
MkeyLeasePeriod:.....0
LocalPort:.....0
LinkWidthEnabled:.....1X or 4X
LinkWidthSupported:.....1X or 4X
LinkWidthActive:.....4X
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkState:.....Active
PhysLinkState:.....LinkUp
LinkDownDefState:.....Polling
ProtectBits:.....0
LMC:.....0
.
.
.
SubnetTimeout:.....0
RespTimeVal:.....0
LocalPhysErr:.....8
OverrunErr:.....8
MaxCreditHint:.....85
RoundTrip:.....16777215
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `smpquery` command
- [“Display Low-Level Detailed Information About a Node” on page 90](#)

Displaying Fabric Functionality Configuration Information

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

- “Display the InfiniBand Fabric Partition Configuration” on page 93
- “Display the InfiniBand Fabric Configuration” on page 94
- “Display the smnodes List” on page 95
- “Display the Fabric Element Configuration” on page 96
- “Display the IPoIB Interface” on page 96
- “Display Configured Secret M_Keys” on page 97

Related Information

- “Displaying Information About Nodes” on page 84
- “Partitioning the InfiniBand Fabric” on page 105
- “Correlating Fabric Elements” on page 119
- “Configuring the Fabric Director Node List” on page 125
- “Performing Supportive Tasks” on page 132

▼ Display the InfiniBand Fabric Partition Configuration

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

1. On the management controller, type.

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

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

For example:

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

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

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

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

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

Repeat as necessary.

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

See [“Modify a Partition or Port” on page 114](#).

Related Information

- [Gateway Reference](#), `smpartition` command
- [“Partitioning the InfiniBand Fabric” on page 105](#)
- [“Display the InfiniBand Fabric Topology \(Detailed\)” on page 86](#)

▼ Display the InfiniBand Fabric Configuration

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

1. On the management controller, type.

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

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

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

See:

- [“Add or Remove an Element From the Fabric Configuration” on page 129.](#)
- [“Modify an Element of the Fabric Configuration” on page 131](#)

Related Information

- *Gateway Reference*, `fdconfig` command
- [“Configuring the Fabric Director Node List” on page 125](#)

▼ Display the smnodes List

1. On the management controller, type.

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

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

See [“Update the smnodes List” on page 134.](#)

Related Information

- *Gateway Reference*, `smnodes` command
- [“Create the smnodes List” on page 133](#)
- [“Update the smnodes List” on page 134](#)

▼ Display the Fabric Element Configuration

1. On the management controller, type.

```
FabMan@gateway_name->createfabric list active
-----
Hostname                                ILOM IP Addr                Redundant IP Addr
-----
mnm12-gw-1                             123.45.67.89                -
mnm13-gw-1                             123.45.67.91                -
FabMan@gateway_name->
```

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

See [“Modify an Element of the Configuration”](#) on page 123.

Related Information

- [Gateway Reference](#), createfabric command
- [“Create a Fabric Element Configuration”](#) on page 120

▼ Display the IPoIB Interface

1. On the management controller, type.

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

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

See:

- [“Delete the IPoIB Interface”](#) on page 136
- [“Create the IPoIB Interface”](#) on page 135.

Related Information

- [Gateway Reference](#), create_ipoib command
- [“Delete the IPoIB Interface”](#) on page 136
- [“Create the IPoIB Interface”](#) on page 135

▼ Display Configured Secret M_Keys

1. On the management controller, type.

```
FabMan@gateway_name->smsubnetprotection list active
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=o4nm2-gw-6
# time=15 Sep 03:54:46
# checksum=378d9b09744e1d8b8ba6ae868c99d0c9
#! commit_number : 3
Mkey                               Untrusted Mkey                     Smkey                               Attribute
-----                               -
0x00abcdefabcdef01                0x1aa45124fee612ae                 0x15fc26aea300f831
0x00abcdefabcdef02                0x4ccd8230de6cd348                 0x3fc7e6ad701a8a2a
0x00abcdefabcdef03                0x9baa1debcc74de5e                 0x1b253003600d137b                C
FabMan@gateway_name->
```

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

See [“Configuring Secret M_Keys”](#) on page 67.

Related Information

- [Gateway Reference](#), smsubnetprotection command
- [“Set the Secret M_Key to Use”](#) on page 72
- [“Add Secret M_Keys to the Configuration”](#) on page 69
- [“Delete Secret M_Keys From the Configuration”](#) on page 73

Controlling the InfiniBand Fabric

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

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

- [“Investigating Nodes”](#) on page 98
- [“Controlling Nodes”](#) on page 101
- [“Partitioning the InfiniBand Fabric”](#) on page 105

- “Correlating Fabric Elements” on page 119
- “Configuring the Fabric Director Node List” on page 125
- “Performing Supportive Tasks” on page 132

Related Information

- “Controlling the Chassis” on page 28
- “Controlling the I4 Switch Chip” on page 40
- “Controlling Gateway Ports and Parameters” on page 202
- “Controlling the Subnet Manager” on page 56
- “Controlling Gateway Ports and Parameters” on page 202
- “Monitoring the InfiniBand Fabric” on page 83

Investigating Nodes

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

- “Perform Comprehensive Diagnostics for the Entire Fabric” on page 98
- “Determine Changes to the InfiniBand Fabric Topology” on page 99
- “Find 1x, SDR, or DDR Links in the Fabric” on page 100
- “Determine Which Links Are Experiencing Significant Errors” on page 100

Related Information

- “Controlling Nodes” on page 101
- “Partitioning the InfiniBand Fabric” on page 105
- “Correlating Fabric Elements” on page 119
- “Configuring the Fabric Director Node List” on page 125
- “Performing Supportive Tasks” on page 132

▼ Perform Comprehensive Diagnostics for the Entire Fabric

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

- On the management controller, type.

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

Related Information

- *Gateway Reference*, `ibdiagnet` command
- *Gateway Reference*, `ibdiagpath` command
- “Display the InfiniBand Fabric Topology (Detailed)” on page 86

▼ Determine Changes to the InfiniBand Fabric Topology

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

1. Take a snapshot of the fabric topology.

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

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

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

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

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

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

```
FabMan@gateway_name->matchtopology Sept8.topo  
Topology matching will take some time. Please wait!  
-I-----  
-I- Topology matching results  
-I-----
```

```
-I- The topology defined in ib_topology.topo perfectly matches the
discovered fabric.
```

```
-----
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, generatetopology command
- *Gateway Reference*, matchtopology command
- [“Display the InfiniBand Fabric Topology \(Simple\)” on page 86](#)

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

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

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

```
FabMan@gateway_name->ibdiagnet -lw 4x -ls 10 -pc -pm -skip all
```

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

Related Information

- *Gateway Reference*, `ibdiagnet` command
- [“Display the Link Status of a Node” on page 88](#)
- [“Determine Which Links Are Experiencing Significant Errors” on page 100](#)

▼ Determine Which Links Are Experiencing Significant Errors

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

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

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

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

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

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

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

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

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

Related Information

- *Gateway Reference*, `ibdiagnet` command
- [“Find 1x, SDR, or DDR Links in the Fabric” on page 100](#)
- [“Display Counters for a Node” on page 89](#)

Controlling Nodes

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

- [“Clear Data and Error Counters” on page 102](#)
- [“Reset a Port” on page 102](#)
- [“Set Port Speed” on page 103](#)
- [“Disable or Enable an External Port” on page 104](#)

Related Information

- [“Investigating Nodes” on page 98](#)
- [“Partitioning the InfiniBand Fabric” on page 105](#)
- [“Correlating Fabric Elements” on page 119](#)
- [“Configuring the Fabric Director Node List” on page 125](#)
- [“Performing Supportive Tasks” on page 132](#)

▼ Clear Data and Error Counters

When you are optimizing the InfiniBand fabric for performance, you might want to know how the throughput increases or decreases according to changes you are making. Alternatively, if you are troubleshooting a port, the `perfquery` command provides counters of errors occurring at that port. To determine the throughput as a function of time, or if the problem at the port has been resolved, the `ibdiagnet -pc` command enables you to reset the data counters and error counters to 0.

- On the management controller, type.

```
FabMan@gateway_name->ibdiagnet -pc
```

Related Information

- *Gateway Reference*, `ibdiagnet` command
- *Gateway Reference*, `perfquery` command
- “Display Counters for a Node” on page 89

▼ Reset a Port

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

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

- On the management controller, type.

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

where:

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

For example:

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

After PortInfo set:
# Port info: Lid 15 port 23
LinkState:.....Down
PhysLinkState:.....PortConfigurationTraining
FabMan@gateway_name->

```

Related Information

- [Gateway Reference](#), `ibportstate` command
- [“Disable or Enable an External Port” on page 104](#)

▼ Set Port Speed

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

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

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

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

where:

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

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

For example:

```
FabMan@gateway_name->ibportstate 15 23 speed 1
Initial PortInfo:
# Port info: Lid 15 port 23
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps

After PortInfo set:
# Port info: Lid 15 port 23
LinkSpeedEnabled:.....2.5 Gbps
FabMan@gateway_name->ibportstate 15 23 speed 7
Initial PortInfo:
# Port info: Lid 15 port 23
LinkSpeedEnabled:.....2.5 Gbps
After PortInfo set:
# Port info: Lid 15 port 23
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `ibportstate` command
- [“Disable the Link Status of a Node” on page 88](#)

▼ Disable or Enable an External Port

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

Note – If secret M_Key functionality is enabled, use the `-M m_key` option to the `ibportstate` command for external ports. Use the `disableswitchport` and `enableswitchport` commands for internal ports when secret M_Key functionality is enabled. See [“Disable a Switch Chip Port” on page 40](#) and [“Enable a Switch Chip Port” on page 41](#).

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

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

where:

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

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

For example:

```
FabMan@gateway_name->ibportstate 15 23 disable
Initial PortInfo:
# Port info: Lid 15 port 23
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 23
LinkState:.....Down
PhysLinkState:.....Disabled
FabMan@gateway_name->
```

Related Information

- [Gateway Reference, ibportstate command](#)
- [“Disable a Switch Chip Port” on page 40](#)
- [“Disable or Enable an External Port” on page 104](#)
- [“Reset a Port” on page 102](#)

Partitioning the InfiniBand Fabric

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

- [“smpartition Command” on page 106](#)
- [“smpartition Command Guidelines” on page 107](#)
- [“Partitions and P_Keys” on page 108](#)
- [“Determine the Partitions and P_Keys” on page 109](#)
- [“Create a User Partition” on page 110](#)
- [“Add or Remove a Port From a Partition” on page 112](#)
- [“Modify a Partition or Port” on page 114](#)

- “Delete a Partition” on page 116
- “Remove User Partitions for Firmware Downgrade” on page 117

Related Information

- “Investigating Nodes” on page 98
- “Controlling Nodes” on page 101
- “Correlating Fabric Elements” on page 119
- “Configuring the Fabric Director Node List” on page 125
- “Performing Supportive Tasks” on page 132

smpartition Command

You can partition your InfiniBand fabric using the `smpartition` command and its subcommands. The `smpartition` command is issued from the gateway or 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 that you can affect.

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

Refer to *Gateway 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

- *Gateway Reference*, `smpartition` command
- “smpartition Command Guidelines” on page 107
- “Partitions and P_Keys” on page 108
- “Determine the Partitions and P_Keys” on page 109
- “Create a User Partition” on page 110
- “Add or Remove a Port From a Partition” on page 112
- “Modify a Partition or Port” on page 114

- “Delete a Partition” on page 116
- “Remove User Partitions for Firmware Downgrade” on page 117

smpartition Command Guidelines

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

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

When you know the VNIC’s Ethernet connector, this table provides the respective BridgeX chip’s portname.

Ethernet Connector	BridgeX Portname
0A-ETH-1	Bridge-0-2
0A-ETH-2	Bridge-0-2
0A-ETH-3	Bridge-0-1
0A-ETH-4	Bridge-0-1
1A-ETH-1	Bridge-1-2

Ethernet Connector	BridgeX Portname
1A-ETH-2	Bridge-1-2
1A-ETH-3	Bridge-1-1
1A-ETH-4	Bridge-1-1

When you know the `Portname`, you can look in the output of the `showgwports` command for the respective `PortGUID`. Using the table and the `showgwports` command together, when you know the Ethernet connectors used for your VNICs, you can find the respective BridgeX chip's InfiniBand port GUIDs to add to the non-default partition.

Related Information

- [“smpartition Command” on page 106](#)
- [“Partitions and P_Keys” on page 108](#)
- [“Determine the Partitions and P_Keys” on page 109](#)
- [“Create a User Partition” on page 110](#)
- [“Add or Remove a Port From a Partition” on page 112](#)
- [“Modify a Partition or Port” on page 114](#)
- [“Delete a Partition” on page 116](#)
- [“Remove User Partitions for Firmware Downgrade” on page 117](#)

Partitions and P_Keys

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

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

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

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

Related Information

- “[smpartition Command](#)” on page 106
- “[smpartition Command Guidelines](#)” on page 107
- “[Determine the Partitions and P_Keys](#)” on page 109
- “[Create a User Partition](#)” on page 110
- “[Add or Remove a Port From a Partition](#)” on page 112
- “[Modify a Partition or Port](#)” on page 114
- “[Delete a Partition](#)” on page 116
- “[Remove User Partitions for Firmware Downgrade](#)” on page 117

▼ Determine the Partitions and P_Keys

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

See:

- “[Identify All Switches in the Fabric](#)” on page 84
- “[Identify All CAs in the Fabric](#)” on page 85
- “[Display Gateway Port Information](#)” on page 179

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 identifier to similar P_Keys in the text file.**

Related Information

- *Gateway Reference*, `ibswitches` command
- *Gateway Reference*, `ibhosts` command
- *Gateway Reference*, `showgwports` command
- “[smpartition Command](#)” on page 106
- “[smpartition Command Guidelines](#)” on page 107
- “[Partitions and P_Keys](#)” on page 108
- “[Create a User Partition](#)” on page 110
- “[Add or Remove a Port From a Partition](#)” on page 112
- “[Modify a Partition or Port](#)” on page 114
- “[Delete a Partition](#)” on page 116

- “Remove User Partitions for Firmware Downgrade” on page 117

▼ 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 114 for an example.

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

See “Create the smnodes List” on page 133.

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

See “Add or Delete an Element From the Configuration” on page 122.

3. Determine your partition P_Keys.

See “Determine the Partitions and P_Keys” on page 109.

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

```
FabMan@gateway_name->smpartition start
FabMan@gateway_name->
```

5. Create the user partition.

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

where:

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

<i>mtu</i> Number	1	2	3	4	5
MTU Value	256	512	1024	2048	4096

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

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

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

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

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

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

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

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

For example:

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

6. Verify the partition configuration.

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

7. Consider your next steps.

- If you want to add GUIDs, see [“Add or Remove a Port From a Partition”](#) on page 112.
- If you want to modify the configuration, see [“Modify a Partition or Port”](#) on page 114.
- If you want to delete the configuration altogether, see [“Delete a Partition”](#) on page 116.
- Otherwise, go to [Step 8](#).

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

```
FabMan@gateway_name->smpartition commit
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `smpartition` command
- [“Display the InfiniBand Fabric Partition Configuration”](#) on page 93
- [“smpartition Command”](#) on page 106
- [“smpartition Command Guidelines”](#) on page 107
- [“Partitions and P_Keys”](#) on page 108
- [“Determine the Partitions and P_Keys”](#) on page 109
- [“Add or Remove a Port From a Partition”](#) on page 112
- [“Modify a Partition or Port”](#) on page 114
- [“Delete a Partition”](#) on page 116
- [“Remove User Partitions for Firmware Downgrade”](#) on page 117

▼ Add or Remove a Port From a Partition

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

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

```
FabMan@gateway_name->smpartition start
FabMan@gateway_name->
```

2. Add ports to the partition.

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

where:

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

For example:

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

3. Or, remove ports from the partition.

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

See [Step 2](#) for variable definitions.

For example:

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

4. Verify the partition configuration.

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

```
0x00212800013e9313,  
0x00212800013e93f7;  
FabMan@gateway_name->
```

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

```
FabMan@gateway_name->smpartition commit  
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `smpartition` command
- “*smpartition Command*” on page 106
- “*smpartition Command Guidelines*” on page 107
- “*Partitions and P_Keys*” on page 108
- “*Determine the Partitions and P_Keys*” on page 109
- “*Create a User Partition*” on page 110
- “*Modify a Partition or Port*” on page 114
- “*Delete a Partition*” on page 116
- “*Remove User Partitions for Firmware Downgrade*” on page 117

▼ Modify a Partition or Port

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

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

```
FabMan@gateway_name->smpartition start  
FabMan@gateway_name->
```

2. Modify the partition or ports.

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

where:

- *partition_name* is an alphanumeric tag to the InfiniBand partition (optional).
- *p_key* is the partition key (1 to 7fff).

- *mtu* is the number that maps to the actual MTU (1 to 5).

<i>mtu</i> Number	1	2	3	4	5
MTU Value	256	512	1024	2048	4096

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

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

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

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

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

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

- *port* is the GUID of the port, or the special parameter, to modify:
 - ALL_CAS – All CAs in the InfiniBand fabric.
 - ALL_SWITCHES – All switches.
 - ALL_ROUTERS – All routers.
- *member* is the membership type (*full*, *limited*, or *both*) for the port.

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

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

3. Verify the partition configuration.

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

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

```
FabMan@gateway_name->smpartition commit
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `smpartition` command
- “[smpartition Command](#)” on page 106
- “[smpartition Command Guidelines](#)” on page 107
- “[Partitions and P_Keys](#)” on page 108
- “[Determine the Partitions and P_Keys](#)” on page 109
- “[Create a User Partition](#)” on page 110
- “[Add or Remove a Port From a Partition](#)” on page 112
- “[Delete a Partition](#)” on page 116
- “[Remove User Partitions for Firmware Downgrade](#)” on page 117

▼ Delete a Partition

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

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

```
FabMan@gateway_name->smpartition start
FabMan@gateway_name->
```

2. Delete the partition.

```
FabMan@gateway_name->smpartition delete -n partition_name | -pkey p_key
```

where:

- *partition_name* is an alphanumeric tag to the InfiniBand partition (optional).
- *p_key* is the partition key (2 to 7ffe).

Note – You cannot delete the pre-defined partitions with P_Keys 1 and 7fff.

For example:

```
FabMan@gateway_name->smpartition delete -n testpartition  
FabMan@gateway_name->
```

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

```
FabMan@gateway_name->smpartition commit  
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, *smpartition* command
- “*smpartition* Command” on page 106
- “*smpartition* Command Guidelines” on page 107
- “Partitions and P_Keys” on page 108
- “Determine the Partitions and P_Keys” on page 109
- “Create a User Partition” on page 110
- “Add or Remove a Port From a Partition” on page 112
- “Modify a Partition or Port” on page 114
- “Remove User Partitions for Firmware Downgrade” on page 117

▼ Remove User Partitions for Firmware Downgrade

You must perform these steps before downgrading the firmware to a version before 2.0.

Note – Only perform this procedure if you are downgrading the firmware.

Note – By removing user partitions, you might lose connectivity that was available only through those partitions. Additionally, you might gain undesired connectivity because all hosts become full members of the default partition.

1. On the management controller of the gateways where VNICs using user partitions were created, remove all VNICs that are using user partitions.

See “Delete VNICs” on page 201.

2. Remove all VLANs mapped to user partitions.

See “Delete VLANs” on page 194.

3. Remove all user-defined partitions.

See “Delete a Partition” on page 116.

4. Restore the default partition to default settings.

See “Modify a Partition or Port” on page 114, however use this command line to modify the partition.

```
FabMan@gateway_name->smpartition modify -pkey 0x7fff -port ALL_CAS  
-flag ipoib -m full
```

5. Consider your next step:

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

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

- If your InfiniBand fabric has multiple Subnet Managers, go to [Step 6](#).

6. Remove all entries from the Subnet Manager node list.

```
FabMan@gateway_name->smnodes delete IP_address [IP_address ...]
```

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

```
FabMan@gateway_name->smnodes delete 123.45.67.89 123.45.67.90  
FabMan@gateway_name->
```

7. Repeat [Step 6](#) on the management controllers of all Subnet Managers in the InfiniBand fabric.

8. Downgrade the firmware.

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

Related Information

- *Gateway Reference*, `smpartition` command
- *Gateway Reference*, `smnodes` command
- “`smpartition` Command” on page 106
- “`smpartition` Command Guidelines” on page 107
- “Partitions and P_Keys” on page 108
- “Determine the Partitions and P_Keys” on page 109
- “Create a User Partition” on page 110
- “Add or Remove a Port From a Partition” on page 112
- “Modify a Partition or Port” on page 114
- “Delete a Partition” on page 116
- “Delete VNICs” on page 201
- “Delete VLANs” on page 194

Correlating Fabric Elements

These topics help you to create a fabric element configuration.

- “`createfabric` Command” on page 120
- “Create a Fabric Element Configuration” on page 120
- “Add or Delete an Element From the Configuration” on page 122
- “Modify an Element of the Configuration” on page 123

Related Information

- “Investigating Nodes” on page 98
- “Controlling Nodes” on page 101
- “Partitioning the InfiniBand Fabric” on page 105
- “Configuring the Fabric Director Node List” on page 125
- “Performing Supportive Tasks” on page 132

createfabric Command

The fabric element configuration file enables both the partition daemon and the Fabric Director to communicate via Ethernet or IPoIB in a redundant fashion. This is accomplished by mapping a correlation of the management controller's host name with the IP address recognized by Oracle ILOM, and the IPoIB IP address of each element.

You can use the `createfabric` command and its subcommands to create or edit the fabric configuration file, add, modify, or delete fabric elements from the file, and complete the configuration file and commit it to use. You can also display the fabric element configuration file with the `createfabric` command.

The `createfabric` command is issued on each management controller that hosts a Subnet Manager. Like the Subnet Manager nodes list, the fabric element configuration file that is created on one management controller, must be duplicated on all other management controllers.

There are two fabric element configurations, the *active* configuration is the one currently in use, and the *in-progress* configuration is the one that you can affect with the `createfabric` command.

Note – The IP addresses of the fabric element configuration file must be in synchronization with IP addresses of the `smnodes` list. If you update the `smnodes` list, you must update the fabric element configuration file.

Related Information

- [Gateway Reference, createfabric command](#)
- [“Display the Fabric Element Configuration” on page 96](#)
- [“Create a Fabric Element Configuration” on page 120](#)
- [“Add or Delete an Element From the Configuration” on page 122](#)
- [“Modify an Element of the Configuration” on page 123](#)

▼ Create a Fabric Element Configuration

1. Verify the `smnodes` list.

See [“Display the smnodes List” on page 95](#).

2. Initiate a fabric element configuration session on the management controller.

```
FabMan@gateway_name->createfabric start empty
FabMan@gateway_name->
```

3. Add a fabric element to the configuration.

```
FabMan@gateway_name->createfabric add-element -name hostname -ilomIp IP_address  
[-redundantIP IP_address]
```

where:

- *hostname* is the host name of the element.
- *IP_address* is the IP address for the element recognized by Oracle ILOM or IPoIB.

For example:

```
FabMan@gateway_name->createfabric add-element -name mnm12-gw-1 -ilomIp  
123.45.67.89  
Element added  
FabMan@gateway_name->
```

4. Repeat [Step 3](#) for any additional fabric elements you want to add to the configuration.

5. Display the fabric element configuration.

```
FabMan@gateway_name->createfabric list modified  
Hostname                                ILOM IP Addr                Redundant IP Addr  
-----                                -  
mnm12-gw-1                              123.45.67.89                -  
FabMan@gateway_name->
```

6. End the configuration session and commit the in-progress configuration to the active fabric element configuration.

```
FabMan@gateway_name->createfabric complete  
FabMan@gateway_name->
```

7. Repeat [Step 2](#) through [Step 6](#) on all management controllers hosting Subnet Managers in the InfiniBand fabric.

Related Information

- *Gateway Reference*, `createfabric` command
- [“Display the Fabric Element Configuration” on page 96](#)
- [“createfabric Command” on page 120](#)
- [“Add or Delete an Element From the Configuration” on page 122](#)
- [“Modify an Element of the Configuration” on page 123](#)

▼ Add or Delete an Element From the Configuration

1. Verify the `smnodes` list.

See “Display the `smnodes` List” on page 95.

2. Initiate a fabric element configuration session on the management controller.

```
FabMan@gateway_name->createfabric start from-current
FabMan@gateway_name->
```

3. Add a fabric element to the configuration.

```
FabMan@gateway_name->createfabric add-element -name hostname -ilomIp IP_address
[-redundantIP IP_address]
```

where:

- `hostname` is the host name of the element.
- `IP_address` is the IP address for the element recognized by Oracle ILOM or IPoIB.

For example:

```
FabMan@gateway_name->createfabric add-element -name mnm13-gw-1 -ilomIp
123.45.67.90
Element added
FabMan@gateway_name->
```

4. Or, delete a fabric element from the configuration.

```
FabMan@gateway_name->createfabric delete-element -name hostname
```

where `hostname` is the host name of the element.

For example:

```
FabMan@gateway_name->createfabric delete-element -name mnm13-gw-1
Element deleted
FabMan@gateway_name->
```

5. Display the fabric element configuration.

```
FabMan@gateway_name->createfabric list modified
Hostname                                ILOM IP Addr                Redundant IP Addr
-----                                -
```

```

mm12-gw-1          123.45.67.89      -
mm13-gw-1          123.45.67.90      -
FabMan@gateway_name->

```

6. End the configuration session and commit the in-progress configuration to the active fabric element configuration.

```

FabMan@gateway_name->createfabric complete
FabMan@gateway_name->

```

7. Repeat Step 2 through Step 6 on all management controllers hosting Subnet Managers in the InfiniBand fabric.

Related Information

- *Gateway Reference*, createfabric command
- “Display the Fabric Element Configuration” on page 96
- “createfabric Command” on page 120
- “Create a Fabric Element Configuration” on page 120
- “Modify an Element of the Configuration” on page 123

▼ Modify an Element of the Configuration

1. Verify the smnodes list.

See “Display the smnodes List” on page 95.

2. Initiate a fabric element configuration session on the management controller.

```

FabMan@gateway_name->createfabric start from-current
FabMan@gateway_name->

```

3. Display the fabric element configuration.

```

FabMan@gateway_name->createfabric list modified
Hostname          ILOM IP Addr      Redundant IP Addr
-----
mm12-gw-1          123.45.67.89      -
mm13-gw-1          123.45.67.90      -
FabMan@gateway_name->

```

4. Modify a fabric element of the configuration.

```
FabMan@gateway_name->createfabric modify-element -name hostname [-ilomIp IP_address]
[-redundantIP IP_address]
```

where:

- *hostname* is the host name of the element as seen in the Hostname column.
- *IP_address* is the IP address for the element recognized by Oracle ILOM or IPoIB.

For example:

```
FabMan@gateway_name->createfabric modify-element -name mnm13-gw-1 -ilomIp
123.45.67.91
Element modified
FabMan@gateway_name->
```

5. Verify the fabric element configuration.

```
FabMan@gateway_name->createfabric list modified
```

Hostname	ILOM IP Addr	Redundant IP Addr
-----	-----	-----
mnm12-gw-1	123.45.67.89	-
mnm13-gw-1	123.45.67.91	-

```
FabMan@gateway_name->
```

6. End the configuration session and commit the in-progress configuration to the active fabric element configuration.

```
FabMan@gateway_name->createfabric complete
FabMan@gateway_name->
```

7. Repeat Step 2 through Step 6 on all management controllers hosting Subnet Managers in the InfiniBand fabric.

Related Information

- *Gateway Reference*, createfabric command
- “Display the Fabric Element Configuration” on page 96
- “createfabric Command” on page 120
- “Create a Fabric Element Configuration” on page 120
- “Add or Delete an Element From the Configuration” on page 122

Configuring the Fabric Director Node List

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

- [“Fabric Director and Fabric Elements” on page 125](#)
- [“fdconfig Command Overview” on page 126](#)
- [“Create a Fabric Configuration” on page 127](#)
- [“Add or Remove an Element From the Fabric Configuration” on page 129](#)
- [“Modify an Element of the Fabric Configuration” on page 131](#)

Related Information

- [“Investigating Nodes” on page 98](#)
- [“Controlling Nodes” on page 101](#)
- [“Partitioning the InfiniBand Fabric” on page 105](#)
- [“Correlating Fabric Elements” on page 119](#)
- [“Performing Supportive Tasks” on page 132](#)

Fabric Director and Fabric Elements

The Fabric Director monitors the InfiniBand fabric by polling the nodes or *elements* of the fabric. Each fabric element has a Fabric Director instance, and the two communicate with each other in a one-to-one relationship. The Fabric Director gathers information about the fabric element through the LDA and ENVD daemons running on the fabric element. From the information gathered, the Fabric Director creates a local SNMP `fabricMIB`.

The Fabric Director of the element hosting the master Subnet Manager is assigned the role of master. The master Fabric Director gathers information about all of the fabric elements by communicating out-of-band (management network) with the Fabric Director instances on those fabric elements. This master Fabric Director creates a model of the fabric that is also in the form of an SNMP MIB, as defined by the SUN-FABRIC-MIB specification.

By default, Fabric Directors run in nonmaster mode. Only when a fabric element’s Subnet Manager becomes the master Subnet Manager, does the fabric element’s corresponding Fabric Director become the master Fabric Director.

This table describes the `fabricMIB` tables and `fabricMIB` scalars generated by Fabric Directors.

Type	Generated fabricMIB Tables	Generated fabricMIB Scalars
All Fabric Directors	<ul style="list-style-type: none"> • fabricElemPortTable • fabricElemConnectorTable 	<ul style="list-style-type: none"> • fabricElemMgrIpAddress • fabricElemName • fabricElemNumPorts • fabricElemOperStatus • fabricElemNumConnectors
Master Fabric Directors	<ul style="list-style-type: none"> • fabricMgmtElemTable • fabricMgmtExtElemTable • fabricMgmtGatewayTable • fabricMgmtLinkTable 	<ul style="list-style-type: none"> • fabricMgmtFabricDescr • fabricMgmtFabricType • fabricMgmtFabricTopology • fabricMgmtFabricOperStatus • fabricMgmtFabricName

Related Information

- *Gateway Reference*, SUN-FABRIC-MIB mib
- *Gateway Reference*, fdconfig command
- [“Display the InfiniBand Fabric Configuration” on page 94](#)
- [“fdconfig Command Overview” on page 126](#)
- [“Create a Fabric Configuration” on page 127](#)
- [“Add or Remove an Element From the Fabric Configuration” on page 129](#)
- [“Modify an Element of the Fabric Configuration” on page 131](#)

fdconfig Command Overview

The fdconfig command is used to configure a list of fabric elements expected to be in the fabric, the fabric node list. The list is used by the Fabric Directors, and provides these attributes for each fabric element:

- Name – The host name of the fabric element as configured by Oracle ILOM.
- IP address – The out-of-band management network IP address of the management controller hosting the fabric element and Fabric Director.
- Fabric element type – The name of the type of fabric element.
 - sw36 – Sun Datacenter InfiniBand Switch 36
 - sw36gw – Sun Network QDR InfiniBand Gateway Switch
 - unknown – The element’s type is unknown.
- Role type – The purpose of the fabric element:
 - gateway – The element acts as a gateway.

- spineSwitch – The element acts as a spine switch.
- leafSwitch – The element acts as a leaf switch.
- unknown – The element’s role is unknown.
- Master – Whether the Fabric Director instance on the fabric element is the master.

Related Information

- [Gateway Reference](#), `fdconfig` command
- [“Display the InfiniBand Fabric Configuration”](#) on page 94
- [“Fabric Director and Fabric Elements”](#) on page 125
- [“Create a Fabric Configuration”](#) on page 127
- [“Add or Remove an Element From the Fabric Configuration”](#) on page 129
- [“Modify an Element of the Fabric Configuration”](#) on page 131

▼ Create a Fabric Configuration

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

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

```
FabMan@gateway_name->fdconfig start-fabric-config empty
FabMan@gateway_name->
```

2. Define the fabric name.

```
FabMan@gateway_name->fdconfig define-fabric-name name
```

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

```
FabMan@gateway_name->fdconfig define-fabric-name test
Fabric name updated
FabMan@gateway_name->
```

3. Define an element of the fabric.

```
FabMan@gateway_name->fdconfig define-element -name name -ip IP_address
-type type -role role1 [-role role2]
```

where:

- *name* is the host name of the element as configured by Oracle ILOM.
- *IP_address* is the IP address of an element.
- *type* is the identifier of the type of element:
 - *sw36* – Sun Datacenter InfiniBand Switch 36
 - *sw36gw* – Sun Network QDR InfiniBand Gateway Switch
 - *unknown* – The element type is unknown.
- *role* is the purpose of the element:
 - *gateway* – The element acts as a gateway.
 - *spineSwitch* – The element acts as a spine switch.
 - *leafSwitch* – The element acts as a leaf switch.
 - *unknown* – The element role is unknown.

For example:

```
FabMan@gateway_name->fdconfig define-element -name primary -ip 123.45.67.89 -type
sw36gw -role leafSwitch -role gateway
Element added
FabMan@gateway_name->
```

4. Verify the fabric configuration.

```
FabMan@gateway_name->fdconfig list-in-progress-fabric-config
Name          IP Addr          Type    Role(s)          isMaster  Fabricname
-----
primary       123.45.67.89    sw36gw  gateway          test
FabMan@gateway_name->
```

5. Consider your next step.

- If you want to add elements, perform [Step 3](#) to [Step 4](#) for each new element.
- If you want to abort the configuration altogether, type.

```
FabMan@gateway_name->fdconfig abort
In progress config aborted
FabMan@gateway_name->
```

- Otherwise, go to [Step 6](#).

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

```
FabMan@gateway_name->fdconfig complete-fabric-config
FabMan@gateway_name->
```

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

The configuration files and fabric MIBs are created.

Related Information

- *Gateway Reference*, `fdconfig` command
- “Display the InfiniBand Fabric Configuration” on page 94
- “Fabric Director and Fabric Elements” on page 125
- “`fdconfig` Command Overview” on page 126
- “Add or Remove an Element From the Fabric Configuration” on page 129
- “Modify an Element of the Fabric Configuration” on page 131

▼ Add or Remove an Element From the Fabric Configuration

1. Initiate a fabric configuration session.

```
FabMan@gateway_name->fdconfig start-fabric-config from-current
FabMan@gateway_name->
```

2. Add an element to the fabric configuration.

```
FabMan@gateway_name->fdconfig define-element -name name -ip IP_address
-type type -role role1 [-role role2]
```

where:

- *name* is the host name of the element as configured by Oracle ILOM.
- *IP_address* is the IP address of an element.
- *type* is the identifier of the type of element:
 - `sw36` – Sun Datacenter InfiniBand Switch 36
 - `sw36gw` – Sun Network QDR InfiniBand Gateway Switch
 - `unknown` – The element type is unknown.
- *role* is the purpose of the element:
 - `gateway` – The element acts as a gateway.
 - `spineSwitch` – The element acts as a spine switch.
 - `leafSwitch` – The element acts as a leaf switch.

- unknown – The element role is unknown.

For example:

```
FabMan@gateway_name->fdconfig define-element -name secondary -ip
123.45.67.90 -type sw36 -role leafSwitch
Element added
FabMan@gateway_name->
```

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

```
FabMan@gateway_name->fdconfig remove-element -name name
```

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

```
FabMan@gateway_name->fdconfig remove-element -name secondary
Element deleted
FabMan@gateway_name->
```

4. Verify the fabric configuration.

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

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

```
FabMan@gateway_name->fdconfig complete-fabric-config
FabMan@gateway_name->
```

6. Perform [Step 1](#) through [Step 5](#) on the management controllers of all elements of the fabric.

7. If you added a new element, create a fabric configuration on the management controller of that element using the same configuration information.

See [“Create a Fabric Configuration”](#) on page 127.

Related Information

- *Gateway Reference*, fdconfig command
- [“Display the InfiniBand Fabric Configuration”](#) on page 94
- [“Fabric Director and Fabric Elements”](#) on page 125

- “fdconfig Command Overview” on page 126
- “Create a Fabric Configuration” on page 127
- “Modify an Element of the Fabric Configuration” on page 131

▼ Modify an Element of the Fabric Configuration

You can modify the parameters of an element.

1. Initiate a fabric configuration session.

```
FabMan@gateway_name->fdconfig start-fabric-config from-current
FabMan@gateway_name->
```

2. Modify an element of the fabric configuration.

```
FabMan@gateway_name->fdconfig redefine-element -name name [-ip
IP_address|-type type|-role role1 [-role role2]]
```

where:

- *name* is the host name of the element as configured by Oracle ILOM.
- *IP_address* is the IP address of an element.
- *type* is the identifier of the type of element:
 - *sw36* – Sun Datacenter InfiniBand Switch 36
 - *sw36gw* – Sun Network QDR InfiniBand Gateway Switch
 - *unknown* – The element type is unknown.
- *role* is the purpose of the element:
 - *gateway* – The element acts as a gateway.
 - *spineSwitch* – The element acts as a spine switch.
 - *leafSwitch* – The element acts as a leaf switch.
 - *unknown* – The element role is unknown.

For example:

```
FabMan@gateway_name->fdconfig redefine-element -name secondary
-rrole spineSwitch
Element modified
FabMan@gateway_name->
```

3. Verify the fabric configuration.

```
FabMan@gateway_name->fdconfig list-in-progress-fabric-config
```

Name	IP Addr	Type	Role(s)	isMaster	Fabricname
primary	123.45.67.89	sw36gw	gateway		test
secondary	123.45.67.90	sw36	spineSwitch	no	test

```
FabMan@gateway_name->
```

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

```
FabMan@gateway_name->fdconfig complete-fabric-config  
FabMan@gateway_name->
```

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

Related Information

- [Gateway Reference](#), fdconfig command
- [“Display the InfiniBand Fabric Configuration”](#) on page 94
- [“Fabric Director and Fabric Elements”](#) on page 125
- [“fdconfig Command Overview”](#) on page 126
- [“Create a Fabric Configuration”](#) on page 127
- [“Add or Remove an Element From the Fabric Configuration”](#) on page 129

Performing Supportive Tasks

These tasks support the extended functionality.

- [“Create the smnodes List”](#) on page 133
- [“Update the smnodes List”](#) on page 134
- [“Create the IPoIB Interface”](#) on page 135
- [“Delete the IPoIB Interface”](#) on page 136

Related Information

- [“Investigating Nodes”](#) on page 98
- [“Controlling Nodes”](#) on page 101
- [“Partitioning the InfiniBand Fabric”](#) on page 105
- [“Correlating Fabric Elements”](#) on page 119

- “Configuring the Fabric Director Node List” on page 125

▼ Create the smnodes List

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

Note – If you are using DHCP for your management network, you must configure the DHCP server to assign a fixed and unique IP address to the Subnet Manager node, based upon the MAC address of that node. The algorithm to determine the IP address is user-determined. It is most important that the IP address is unique and consistent.

The list should have an entry for every Sun Datacenter InfiniBand Switch 36 and Sun Network QDR InfiniBand Gateway Switch that runs a Subnet Manager in your InfiniBand fabric.

Note – If the Subnet Manager nodes of your InfiniBand fabric ever change (disabled, added, and so on), you must update all copies of the Subnet Manager nodes file.

1. **Determine the IP addresses of all management controllers in your InfiniBand fabric.**
2. **On the management controller hosting the master Subnet Manager, propagate the Subnet Manager nodes file with the IP addresses of all Subnet Manager nodes.**

```
FabMan@gateway_name->smodes add IP_address IP_address ...
```

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

```
FabMan@gateway_name->smodes add 123.45.67.89 123.45.67.90
FabMan@gateway_name->
```

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

4. Determine how you will partition your InfiniBand fabric.

See “Determine the Partitions and P_Keys” on page 109.

Related Information

- *Gateway Reference*, `smodes` command
- “Display the smnodes List” on page 95
- “Update the smnodes List” on page 134
- “Create a User Partition” on page 110
- “Prepare for Secret M_Key Functionality” on page 66

▼ Update the smnodes List

1. On the management controller of the master Subnet Manager, display all the Subnet Manager nodes.

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

2. Add any missing Subnet Manager node IP addresses for the local fabric.

```
FabMan@gateway_name->smnodes add IP_address IP_address ...
```

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

```
FabMan@gateway_name->smnodes add 123.45.67.91 123.45.67.92
FabMan@gateway_name->
```

3. Or, delete any unnecessary Subnet Manager node IP addresses for the local fabric.

```
FabMan@gateway_name->smnodes delete IP_address IP_address ...
```

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

```
FabMan@gateway_name->smnodes delete 123.45.67.91 123.45.67.92
FabMan@gateway_name->
```

4. If you are combining fabrics, repeat [Step 2](#), adding all the Subnet Manager node IP addresses from the remote fabrics.
5. Verify the updated list.

```
FabMan@gateway_name->smnodes list
123.45.67.89
123.45.67.90
123.45.67.91
123.45.67.92
123.45.66.77
123.45.68.10
FabMan@gateway_name->
```

6. Repeat [Step 1](#) through [Step 5](#) for all management controllers hosting Subnet Manager nodes in the local fabric.
7. If you are combining fabrics, repeat [Step 1](#) through [Step 5](#) for all management controllers hosting Subnet Manager nodes in the remote fabrics.

Related Information

- [Gateway Reference](#), smodes command
- [“Display the smnodes List”](#) on page 95
- [“Create the smnodes List”](#) on page 133
- [“Create a User Partition”](#) on page 110
- [“Prepare for Secret M_Key Functionality”](#) on page 66

▼ Create the IPoIB Interface

This procedure creates an Internet protocol over InfiniBand interface.

Note – Only one IPoIB interface can be defined.

- On the management controller, type.

```
FabMan@gateway_name->create_ipoib -n if_name -m mask -t if_type p_key IP_address
```

where:

- *if_name* is the identifier of the IP interface.
- *mask* is the subnet mask.
- *if_type* is the type of interface (external, internal, or both).

- *p_key* is the partition key (1 to 7fff or default).
- *IP_address* is the IP address of the interface.

For example:

```
FabMan@gateway_name->create_ipoib -n newipoib -m 255.255.254.0 -t both 0x0001
123.45.67.89
Create interface:
Name          PKey    IP-address      Subnetmask      Interface  Interface-type
-----
newipoib      0x0001  123.45.67.89   255.255.254.0  ib0        both
FabMan@gateway_name->
```

Related Information

- [Gateway Reference](#), create_ipoib command
- [“Display the IPoIB Interface” on page 96](#)
- [“Delete the IPoIB Interface” on page 136](#)
- [“Prepare for Secret M_Key Functionality” on page 66](#)

▼ Delete the IPoIB Interface

- On the management controller, type.

```
FabMan@gateway_name->delete_ipoib if_name | p_key
```

where:

- *if_name* is the identifier of the IP interface.
- *p_key* is the partition key (1 to 7fff or default).

For example:

```
FabMan@gateway_name->delete_ipoib newipoib
Delete interface:
Name          PKey    IP-address      Subnetmask      Interface  Interface-type
-----
newipoib      0x0001  123.45.67.89   255.255.254.0  ib0        both
FabMan@gateway_name->
```

Related Information

- [Gateway Reference](#), delete_ipoib command
- [“Display the IPoIB Interface” on page 96](#)
- [“Create the IPoIB Interface” on page 135](#)

Administering Gateway Resources

These topics describe the host and BridgeX functionality of the gateway and how to configure, monitor, and control that functionality.

- [“Installing Gateway Supportive Software \(Linux\)” on page 137](#)
- [“Creating VNICs Under Gateway Manual Mode \(Linux\)” on page 142](#)
- [“Creating VNICs Under Host Manual Mode \(Linux\)” on page 148](#)
- [“Creating Virtual IO Adapters \(Oracle Solaris\)” on page 160](#)
- [“Monitoring Gateway Resources” on page 172](#)
- [“Controlling LAGs” on page 183](#)
- [“Controlling VLANs and VNICs” on page 190](#)
- [“Controlling Gateway Ports and Parameters” on page 202](#)

Related Information

- [“Troubleshooting the Gateway” on page 1](#)
- [“Understanding Administrative Commands” on page 13](#)
- [“Administering the Chassis” on page 19](#)
- [“Administering the I4 Switch Chip” on page 31](#)
- [“Administering the InfiniBand Fabric” on page 83](#)
- [“Administering the Subnet Manager” on page 49](#)

Installing Gateway Supportive Software (Linux)

Install the BXOFED software on each host of the InfiniBand fabric that utilizes the VNIC and VLAN features of the gateway. Within the BXOFED software package are drivers and applications you use to configure VNICs and VLANs on the hosts.

- [“Acquire the BXOFED Software \(Linux\)” on page 138](#)

- “Install the BXOFED Software (Linux)” on page 139
- “Acquire the ConnectX-2 Firmware” on page 141

Related Information

- “Creating VNICs Under Gateway Manual Mode (Linux)” on page 142
- “Creating VNICs Under Host Manual Mode (Linux)” on page 148
- “Creating Virtual IO Adapters (Oracle Solaris)” on page 160
- “Monitoring Gateway Resources” on page 172
- “Controlling LAGs” on page 183
- “Controlling VLANs and VNICs” on page 190
- “Controlling Gateway Ports and Parameters” on page 202

▼ Acquire the BXOFED Software (Linux)

1. **Open a web browser on a host that will receive the BXOFED software.**
2. **Go to this URL:**
`http://support.oracle.com`
Oracle’s My Oracle Support page is displayed.
3. **Sign in if you already have an account.**
The dashboard page is displayed.

Note – If you do not have an account, you must register.

4. **Click the Patches & Updates tab.**
The Patches & Updates page is displayed.
5. **In the Patch Search window, click the Product or Family (Advanced).**
The Patch Search window updates.
6. **In the Product Is field, type BridgeX.**
Possible products are suggested.
7. **Click the most appropriate link.**
The Release Is field might autopropagate with the most current version.

8. **In the Release drop-down menu, select the most current version of the BridgeX OFED software.**

For example, BridgeX OFED 1.5.1.

9. **Click outside of the drop-down menu.**

10. **Click Search.**

The Patch Search window expands with the search results.

11. **In the Patch Name column, click the patch number link respective to your platform.**

For example, 12621910. The Patch Search window reformats.

12. **Click Read Me to display the README file.**

13. **Click Download.**

The File Download window opens.

14. **Click the *filename.zip* link to initiate the download.**

For example, p12621910_151_Linux-x86-64.zip.

15. **Indicate where the file should be saved.**

The file is downloaded and saved.

16. **In your receiving directory, decompress the *filename.zip* file.**

The BXOFED software is in the BXOFED-1.5.1-*version_for* Oracle.tgz file. There are also README, release notes, installation guide, and user manual files in the *filename.zip* file.

17. **Read the README, release notes, and installation guide files for information on how to install the BXOFED software.**

Related Information

- *Gateway Remote Management*, acquiring the gateway firmware package
- [“Install the BXOFED Software \(Linux\)” on page 139](#)

▼ **Install the BXOFED Software (Linux)**

When you install the BXOFED software, any previous installations of OFED or BXOFED software are removed. Configuration files are not removed.

Note – If you are installing the BXOFED software on a cluster, install the software onto one of the cluster nodes, then install the `.rpm` files in the `OFED-1.5.1/RPMS` on all remaining cluster nodes using cluster-aware tools.

1. **Become superuser of the host that received the BXOFED software.**
2. **Change to the directory where you extracted the `.tgz` file.**
3. **Run the installation script.**

```
# ./BXOFED-1.5.1-1.6.3/install.pl
```

The script begins. Interactive menus direct you through the installation process. During the installation, two configuration files are created:

- `ofed.conf` – contains the names of the software modules installed and the configuration settings chosen during the installation.
- `ofed_net.conf` – contains the IPoIB configuration settings chosen during the installation.

The script finishes. This information is found in the respective locations:

- Man pages are installed in `/usr/share/man`.
- Documentation is installed under the `/usr/share/doc` directory.
- IPoIB configuration information is installed under the `/etc/sysconfig/network*` directory.
- The `openibd` daemon is installed under the `/etc/init.d` directory.
- BXOFED commands are located in the `/usr/bin` and `/usr/sbin` directories.
- BXOFED software installation information is displayed with the `/etc/infiniband/info` script.

4. **(Optional) If the `.tgz` file was extracted to a NFS shared directory for a cluster, then to install the BXOFED software onto any remaining nodes in that cluster.**
 - a. **Log in as superuser of a node to receive the BXOFED software.**
 - b. **Change to the directory where the `.tgz` file was extracted.**
 - c. **Install the BXOFED software automatically.**

```
# ./BXOFED-1.5.1-1.6.3/install.pl -c path/ofed.conf -n path/ofed_net.conf
```

where *path* is the directory path to the `ofed.conf` and `ofed_net.conf` files.

- d. **Repeat from [Step a](#) for all nodes to receive the BXOFED software.**

5. Reboot the Linux InfiniBand host(s).

Related Information

- *Gateway Remote Management*, upgrading the gateway firmware
- [“Acquire the BXOFED Software \(Linux\)” on page 138](#)

▼ Acquire the ConnectX-2 Firmware

For your host to properly interface with the gateway, the firmware of the ConnectX-2 chip in the HCA must be updated to version 2.7.000 or higher.

1. Open a web browser on the host that will receive the ConnectX-2 firmware.

2. Go to this URL:

<http://support.oracle.com>

Oracle’s My Oracle Support page is displayed.

3. Sign in if you already have an account.

The dashboard page is displayed.

Note – If you do not have an account, you must register.

4. Click the Patches & Updates tab.

The Patches & Updates page is displayed.

5. In the Patch Search window, click the Product or Family (Advanced).

The Patch Search window updates.

6. In the Product Is drop-down menu, select your HCA.

For example, Sun Dual Port 4x QDR InfiniBand (IB) HCA PCIe ExpressModule.

7. In the Release Is drop-down menu, select the latest firmware version.

For example, FW25408 v2.7.8130.

8. Click outside of the drop-down menu.

9. In the Platform Is drop-down menu, select the Oracle Solaris appropriate for your host.

For example, Oracle Solaris on x86-64 (64-bit).

10. Click outside of the drop-down menu.

11. Click Search.

The Patch Search window expands with the search results.

12. In the Patch Name column, click the respective patch number link.

For example, 12610332. The Patch Search window reformats.

13. Click Read Me to display the README file.

14. Click Download.

The File Download window opens.

15. Click the *filename.zip* link to initiate the download.

For example, p12610332__Solaris86-64.zip.

16. Indicate where the file should be saved.

The file is downloaded and saved.

17. In your receiving directory, decompress the *filename.zip* file.

The ConnectX-2 firmware is in the *fw-ConnectX2-rel-2_7_version.bin* file.

For example, *fw-ConnectX2-rel-2_7_8130-375-3697-01.bin*.

18. Refer to your HCA documentation for instructions on how to upgrade the ConnectX-2 firmware.

Creating VNICs Under Gateway Manual Mode (Linux)

These topics describe how to create VNICs under gateway manual mode:

- [“Gateway Manual Mode Overview \(Linux\)”](#) on page 143
- [“Determine VNIC Configuration Parameters For Gateway Manual Mode \(Linux\)”](#) on page 144
- [“Determine VLAN Associations for Gateway Manual Mode \(Linux\)”](#) on page 145
- [“Configure and Create VNICs for Gateway Manual Mode \(Linux\)”](#) on page 146
- [“VNIC Configuration on Linux Hosts in Gateway Manual Mode”](#) on page 147

Related Information

- [“Installing Gateway Supportive Software \(Linux\)”](#) on page 137
- [“Creating VNICs Under Host Manual Mode \(Linux\)”](#) on page 148

- [“Creating Virtual IO Adapters \(Oracle Solaris\)”](#) on page 160
- [“Monitoring Gateway Resources”](#) on page 172
- [“Controlling LAGs”](#) on page 183
- [“Controlling VLANs and VNICs”](#) on page 190
- [“Controlling Gateway Ports and Parameters”](#) on page 202

Gateway Manual Mode Overview (Linux)

The default and suggested means of creating and managing VNICs on the gateway is through use of BridgeX gateway administrative commands. To use the commands described in these topics, you must access them through the `/SYS/Gateway_Mgmt` or `/SYS/Fabric_Mgmt` Linux shell targets of the Oracle ILOM CLI. The VNICs you create through the commands are based on the InfiniBand hosts.

When creating VNICs, consider this:

- You must specify the MAC for the VNIC. If you do not, the default MAC for the VNIC is 00:00:00:00:00:00, which renders the VNIC unusable.
- You must ensure that each MAC is unique.
- If you do not have a list of unique global MACs, use locally administrated MACs.
- VNICs will not go to an up state without an active Subnet Manager present.
- Manually created VNICs are persistent, and survive reboots and power cycles.

When creating VNICs, consider that there are conditions from partitioning. See [“smpartition Command Guidelines”](#) on page 107.

Related Information

- [“Host Manual Mode Overview \(Linux\)”](#) on page 149
- [“Determine VNIC Configuration Parameters For Gateway Manual Mode \(Linux\)”](#) on page 144
- [“Determine VLAN Associations for Gateway Manual Mode \(Linux\)”](#) on page 145
- [“Configure and Create VNICs for Gateway Manual Mode \(Linux\)”](#) on page 146
- [“VNIC Configuration on Linux Hosts in Gateway Manual Mode”](#) on page 147

▼ Determine VNIC Configuration Parameters For Gateway Manual Mode (Linux)

MAC addresses, GUIDs, and connector names or LAGs are assigned to VNICs. You can create a repository of this information to assist in creating VNICs.

Note – This procedure creates a *gateway* MAC address list, used to create VNICs in gateway manual mode.

1. Determine if you will use unique global MAC addresses or locally administrated MAC addresses.
2. Create a list of MAC addresses, one per row.

Note – Each MAC address must be unique and not 00:00:00:00:00:00.

Note – Only even numbers are supported for the most significant byte of the MAC address (unicast).

3. To the left of each MAC address, provide the HCA port GUID to receive that MAC address.
4. To the left of each GUID and MAC address pair, provide the name of the connector or LAG where the host will be physically connected to the 10GbE network.

There are eight connector names, 0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4.

5. Use the list for when you create the VNICs.

An example entry in the gateway MAC address file might look like this:

```
0A-ETH-1 0003ba000100c70b 00:30:48:7d:de:e4
```

For this example:

- 0A-ETH-1 is the gateway connector.
- 0003ba000100c70b is the HCA port GUID.
- 00:30:48:7d:de:e4 is the MAC address.

Related Information

- “Determine VNIC Configuration Parameters for Host Manual Mode (Linux)” on page 152
- “Gateway Manual Mode Overview (Linux)” on page 143
- “Determine VLAN Associations for Gateway Manual Mode (Linux)” on page 145
- “Configure and Create VNICs for Gateway Manual Mode (Linux)” on page 146
- “VNIC Configuration on Linux Hosts in Gateway Manual Mode” on page 147

▼ Determine VLAN Associations for Gateway Manual Mode (Linux)

You must associate a VNIC with a VLAN, even if that VLAN identifier is `NO`. You also must partition the InfiniBand fabric and provide a partition key for each VLAN. If the VLAN identifier `NO` is to be associated with the `P_Key default`, you need not create the default partition, for it already exists.

Note – The association of VLAN identifier `NO` to partition key `default` survives firmware upgrades.

1. **If you have not done already, partition the InfiniBand fabric.**
See “Partitioning the InfiniBand Fabric” on page 105.
2. **Use the `smpartition` command on the management controller to list the GUIDs associated with the partition and partition key.**
See “Display the InfiniBand Fabric Partition Configuration” on page 93
3. **Return to the gateway MAC address list and determine which addresses will be associated with a VLAN and which will be associated with VLAN ID `NO`.**
4. **For each MAC address, append the VLAN identifier (`NO` or 2 to 4094) and partition key number to the right of the MAC address.**

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

Note – Given a *connector* and *vlan_id* combination (including `NO`), there is only one *p_key* permitted.

Note – Due to hardware limitations for MultiCast groups, there is a maximum of 1000 VLANs.

Even if the MAC address will not be associated with a VLAN (VLAN identifier NO), you must assign a partition key (`default` or `other`).

Note – When you associate a MAC with a VLAN, the GUID respective to the MAC must be a member of the partition associated with the partition key.

5. Use the list for when you create the VLANs and VNICs.

An example entry in the gateway MAC address list might look like this:

```
0A-ETH-1 0003ba000100c70b 00:30:48:7d:de:e4 3 default
```

For this example:

- `0A-ETH-1` is the gateway connector.
- `0003ba000100c70b` is the HCA port GUID.
- `00:30:48:7d:de:e4` is the MAC address.
- `3` is the VLAN identifier.
- `default` is the partition key.

Related Information

- [“Gateway Manual Mode Overview \(Linux\)” on page 143](#)
- [“Determine VNIC Configuration Parameters For Gateway Manual Mode \(Linux\)” on page 144](#)
- [“Configure and Create VNICs for Gateway Manual Mode \(Linux\)” on page 146](#)
- [“VNIC Configuration on Linux Hosts in Gateway Manual Mode” on page 147](#)

▼ Configure and Create VNICs for Gateway Manual Mode (Linux)

1. Determine the VNIC configuration parameters.

See [“Determine VNIC Configuration Parameters For Gateway Manual Mode \(Linux\)” on page 144](#).

2. Determine VLAN associations.

See [“Determine VLAN Associations for Gateway Manual Mode \(Linux\)” on page 145](#).

3. Create VLANs.

See “Create VLANs” on page 192.

4. Create VNICs.

See “Create VNICs” on page 195.

Related Information

- “Configure and Create VNICs for Host Manual Mode (Linux)” on page 158
- “Gateway Manual Mode Overview (Linux)” on page 143
- “Determine VNIC Configuration Parameters For Gateway Manual Mode (Linux)” on page 144
- “Determine VLAN Associations for Gateway Manual Mode (Linux)” on page 145
- “VNIC Configuration on Linux Hosts in Gateway Manual Mode” on page 147

VNIC Configuration on Linux Hosts in Gateway Manual Mode

If you are creating VNICs in Gateway Manual Mode, you must meet these conditions.

VNIC names appear frequently in the documentation. For example, the VNIC-specific configuration file contains the `DEVICE` parameter, which format is described as:

```
DEVICE=name
```

where *name* is suggested to be `ethXX`, and `XX` is a user-specified sequence number. VNIC names of `eth0`, `eth1`, ... to `eth7` have been problematic with Linux kernels.

A solution is to use this format for the VNIC *name*:

```
ethgateway_instance_VNIC_instance
```

where:

- *gateway_instance* is the gateway instance number as returned by the `showgwconfig` command.
- *VNIC_instance* is a sequential VNIC ID number assigned to the VNIC. You can view the VNIC ID numbers with the `showvnics` command.

An example VNIC name might be `eth03_03`.

Using this new format, the filename for the VNIC-specific configuration file becomes `ifcfg-ethgateway_instance_VNIC_instance`. For the previous example, the configuration filename is `ifcfg-eth03_03`.

For consistency, the VNIC-specific configuration file *must* contain these two parameters:

```
DEVICE=name
HWADDR=mac
```

where:

- *name* is `ethgateway_instance_VNIC_instance` as described previously.
- *mac* is the MAC assigned to the VNIC.

Related Information

- [“Gateway Manual Mode Overview \(Linux\)” on page 143](#)
- [“Determine VNIC Configuration Parameters For Gateway Manual Mode \(Linux\)” on page 144](#)
- [“Determine VLAN Associations for Gateway Manual Mode \(Linux\)” on page 145](#)
- [“Configure and Create VNICs for Gateway Manual Mode \(Linux\)” on page 146](#)

Creating VNICs Under Host Manual Mode (Linux)

These topics describe how to create VNICs under host manual mode:

- [“Host Manual Mode Overview \(Linux\)” on page 149](#)
- [“Central Configuration File \(Linux\)” on page 150](#)
- [“VNIC-Specific Configuration File \(Linux\)” on page 151](#)
- [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)” on page 152](#)
- [“Create the Central Configuration File \(Linux\)” on page 154](#)
- [“Create the VNIC-Specific Configuration Files \(Linux\)” on page 156](#)
- [“mlx4_vnic_conf daemon \(Linux\)” on page 157](#)
- [“Configure and Create VNICs for Host Manual Mode \(Linux\)” on page 158](#)

Related Information

- [“Installing Gateway Supportive Software \(Linux\)”](#) on page 137
- [“Creating VNICs Under Gateway Manual Mode \(Linux\)”](#) on page 142
- [“Creating Virtual IO Adapters \(Oracle Solaris\)”](#) on page 160
- [“Monitoring Gateway Resources”](#) on page 172
- [“Controlling LAGs”](#) on page 183
- [“Controlling VLANs and VNICs”](#) on page 190
- [“Controlling Gateway Ports and Parameters”](#) on page 202

Host Manual Mode Overview (Linux)

Before creating VNICs in host manual mode, you must install the BXOFED software onto the hosts. Afterwards, you create VNICs using static configuration files located on the hosts. These configuration files define the number of VNICs and the VNICs’ properties. The `mlx4_vnic_confd` daemon in the BXOFED software reads the configuration files and passes the relevant data to the `mlx4_vnic` module, which creates the VNICs.

The two types of configuration files for VNICs provide the same functionality:

- A central configuration file (`mlx4_vnic.conf`)
- VNIC-specific configuration files (`ifcfg-ethXX` where `XX` is the `eth` number of the VNIC)

If both forms of configuration files exist, the central configuration file has precedence and is the only file used.

When creating VNICs, consider that there are conditions from partitioning. See [“smpartition Command Guidelines”](#) on page 107.

Related Information

- [“Gateway Manual Mode Overview \(Linux\)”](#) on page 143
- [“Central Configuration File \(Linux\)”](#) on page 150
- [“VNIC-Specific Configuration File \(Linux\)”](#) on page 151
- [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)”](#) on page 152
- [“Create the Central Configuration File \(Linux\)”](#) on page 154
- [“Create the VNIC-Specific Configuration Files \(Linux\)”](#) on page 156
- [“mlx4_vnic_confd Daemon \(Linux\)”](#) on page 157
- [“Configure and Create VNICs for Host Manual Mode \(Linux\)”](#) on page 158

Central Configuration File (Linux)

The central configuration file is the `/etc/infiniband/mlx4_vnic.conf` file and consists of multiple single-line entries, each describing a VNIC. Each VNIC entry has this format and parameters:

```
name=name mac=mac ib_port=device:port [vid=vlan_ID] vnic_id=number bx=string eport=connector
```

where:

- *name* is the VNIC device name or eth number.
- *mac* is the MAC assigned to the VNIC.
- *device* is the device name retrieved from the `hca_id` field in the output of the `ibv_devinfo` command.
- *port* is the port number, either 1 or 2.
- *vlan_ID* is the VLAN identifier to assign to the VNIC (NO or 2 to 4094).

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

Note – The `vid` parameter is optional.

- *number* is a unique number to assign to the VNIC. The value is 1 to 32767.
- *string* is either the HCA port GUID or the system name.
- *connector* is the gateway connector assigned to the VNIC (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).

Related Information

- [“Host Manual Mode Overview \(Linux\)” on page 149](#)
- [“VNIC-Specific Configuration File \(Linux\)” on page 151](#)
- [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)” on page 152](#)
- [“Create the Central Configuration File \(Linux\)” on page 154](#)
- [“Create the VNIC-Specific Configuration Files \(Linux\)” on page 156](#)
- [“mlx4_vnic_confd Daemon \(Linux\)” on page 157](#)
- [“Configure and Create VNICs for Host Manual Mode \(Linux\)” on page 158](#)

VNIC-Specific Configuration File (Linux)

The `ifcfg-ethXX` file already exists as a means for the network service to derive information about a node. The `XX` represents the `eth` number of the VNIC. To configure VNICs, additional parameters are added to the base file format. For Red Hat Linux, the `ifcfg-ethXX` file has this format:

```
DEVICE=name
HWADDR=mac
BOOTPROTO=dhcp
ONBOOT=yes
BXADDR=string
BEXPORT=connector
VNICVLAN=vlan_ID
VNICIBPORT=device:port
```

where:

- *name* is the VNIC device name or `eth` number.

Note – The `DEVICE` parameter is optional. If the parameter is missing, the suffix of the configuration file name (`ethXX`) is used.

- *mac* is the MAC assigned to the VNIC.
- *string* is either the HCA port GUID or the system name.
- *connector* is the gateway connector assigned to the VNIC (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).
- *vlan_ID* is the VLAN identifier to assign to the VNIC (NO or 2 to 4094).

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

Note – The `VNICVLAN` parameter is optional.

- *device* is the device name retrieved from the `hca_id` field in the output of the `ibv_devinfo` command.
- *port* is the port number, either 1 or 2.

Note – Additional parameters used for regular `eth` interfaces can be appended to the `ifcfg-ethXX` file.

Related Information

- “Host Manual Mode Overview (Linux)” on page 149
- “Central Configuration File (Linux)” on page 150
- “Determine VNIC Configuration Parameters for Host Manual Mode (Linux)” on page 152
- “Create the Central Configuration File (Linux)” on page 154
- “Create the VNIC-Specific Configuration Files (Linux)” on page 156
- “mlx4_vnic_confd Daemon (Linux)” on page 157
- “Configure and Create VNICs for Host Manual Mode (Linux)” on page 158

▼ Determine VNIC Configuration Parameters for Host Manual Mode (Linux)

MAC addresses, GUIDs, and other parameters are assigned to VNICs. You can create a repository of this information to assist in creating VNICs.

Note – This procedure creates a *host* MAC address list, used to create VNICs in host manual mode.

You must associate a VNIC with a VLAN, even if that VLAN identifier is `NO`. You also must partition the InfiniBand fabric and provide a partition key for each VLAN. If the VLAN identifier `NO` is to be associated with the `P_Key default`, you need not create the default partition, for it already exists.

Note – The association of VLAN identifier `NO` to partition key `default` survives firmware upgrades.

1. **Determine if you will use unique global MAC addresses or locally administrated MAC addresses.**
2. **Create a numbered list, one number per row.**
3. **For each number, provide a MAC address to the right.**

Note – Each MAC address must be unique and not `00:00:00:00:00:00`.

Note – Only even numbers are supported for the most significant byte of the MAC address (unicast).

4. For each MAC address, open a terminal window and log in to the host to receive that MAC address.
5. Use the `ibv_devinfo` command to determine the `hca_id` value and HCA port GUID.

For example:

```
# ibv_devinfo
hca_id:mlx4_0
  fw_ver:          2.5.9266
  node_guid:       0003:ba00:0100:c708
  sys_image_guid: 0003:ba00:0100:c70b
.
.
.
#
```

In the example, the `hca_id` value is `mlx4_0`, and the HCA port GUID (`sys_image_guid`) is `003ba000100c70b`.

6. For each MAC address, append the host name, the respective `hca_id` value, the port (1 or 2), and the HCA port GUID to the right of the MAC address.
7. For each MAC address that will be associated with a VLAN, including VLAN ID NO, append the VLAN identifier (NO or 2 to 4094) to the right of the GUID.

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

Note – Due to hardware limitations for MultiCast groups, there is a maximum of 1000 VLANs.

8. For each MAC address, append the name of the gateway connector where the host will physically connect to the 10GbE network to the right of the GUID or VLAN identifier (if present).

Note – Given a `connector` and `vlan_id` combination, there is only one `p_key` permitted.

9. Use the list for when you create VLANs and VNICs.

An example entry in the host MAC address list might look like this:

```
1 00:30:48:7d:de:e4 cupcake mlx4_0 1 0003ba000100c70b 3 0A-ETH-1
```

For this example:

- 1 is the entry number. This will become the VNIC number and eth number.
- 00:30:48:7d:de:e4 is the MAC address.
- cupcake is the host name and is used to identify the location of the VNIC.
- mlx4_0 is the value of the hca_id field and is called the device name.
- 1 is the port of the HCA.
- 0003ba000100c70b is the HCA port GUID.
- 3 is the VLAN identifier.
- 0A-ETH-1 is the connector.

Related Information

- [“Determine VNIC Configuration Parameters For Gateway Manual Mode \(Linux\)” on page 144](#)
- [“Determine VLAN Associations for Gateway Manual Mode \(Linux\)” on page 145](#)
- [“Host Manual Mode Overview \(Linux\)” on page 149](#)
- [“Central Configuration File \(Linux\)” on page 150](#)
- [“VNIC-Specific Configuration File \(Linux\)” on page 151](#)
- [“Create the Central Configuration File \(Linux\)” on page 154](#)
- [“Create the VNIC-Specific Configuration Files \(Linux\)” on page 156](#)
- [“mlx4_vnic_conf Daemon \(Linux\)” on page 157](#)
- [“Configure and Create VNICs for Host Manual Mode \(Linux\)” on page 158](#)

▼ Create the Central Configuration File (Linux)

The `etc/infiniband/mlx4_vnic.conf` file is shared by all hosts in the InfiniBand fabric. You create the configuration file from the host MAC address file.

1. Use the host MAC address list that you created previously.

See [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)” on page 152](#).

2. Open a text editor to create the configuration file.

3. Create an entry in the configuration file with this format:

```
name=name mac=mac ib_port=device:port [vid=vlan_ID] vnic_id=number bx=string eport=connector
```

where:

- *name* is ethXX and XX is column one of the host MAC address file.
- *mac* is column two of the host MAC address list.
- *device* is column four of the host MAC address list.
- *port* is column five of the host MAC address list.
- *vlan_ID* is column seven of the host MAC address list.

Note – The *vid* parameter is optional.

- *number* is column one of the host MAC address list.
- *string* is column six of the host MAC address list.
- *connector* is column eight of the host MAC address list.

For example:

```
name=eth1 mac=00:30:48:7d:de:e4 ib_port=mlx4_0:1 vid=0 vnic_id=1 bx=0003ba000100c70b eport=0A-ETH-1
```

4. Repeat [Step 3](#) for all rows of the host MAC address list.

5. Save the file with the name `mlx4_vnic.conf`.

6. Copy the `mlx4_vnic.conf` file to the `/etc/infiniband` directory of all hosts in the InfiniBand fabric.

Related Information

- [“Create the VNIC-Specific Configuration Files \(Linux\)”](#) on page 156
- [“Host Manual Mode Overview \(Linux\)”](#) on page 149
- [“Central Configuration File \(Linux\)”](#) on page 150
- [“VNIC-Specific Configuration File \(Linux\)”](#) on page 151
- [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)”](#) on page 152
- [“mlx4_vnic_conf Daemon \(Linux\)”](#) on page 157
- [“Configure and Create VNics for Host Manual Mode \(Linux\)”](#) on page 158

▼ Create the VNIC-Specific Configuration Files (Linux)

Each host will have a VNIC-specific configuration file for the VNICs it will host. Each line of the MAC address file becomes a VNIC-specific configuration file. Repeat this procedure for each host in the InfiniBand fabric.

1. Use the host MAC address list that you created previously.

See “[Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)](#)” on page 152.

2. Look in column three of the host MAC address list to find entries for the respective host.

Note the row numbers in column one.

3. Open a text editor to create a configuration file.

4. Enter text from the entry appropriate for the respective host into the file with this format.

```
DEVICE=name
HWADDR=mac
BOOTPROTO=dhcp
ONBOOT=yes
BXADDR=string
BEXPORT=connector
VNICVLAN=vlan_ID
VNICIBPORT=device:port
```

where:

- *name* is ethXX and XX is column one of the host MAC address file.
- *mac* is column two of the host MAC address list.
- *string* is column six of the host MAC address list.
- *connector* is column eight of the host MAC address list.
- *vlan_ID* is column seven of the host MAC address list.

Note – The VNICVLAN parameter is optional.

- *device* is column four of the host MAC address list.

- *port* is column five of the host MAC address list.

For example:

```
DEVICE=eth1
HWADDR=00:30:48:7d:de:e4
BOOTPROTO=dhcp
ONBOOT=yes
BXADDR=0003ba000100c70b
BXEPORT=0A-ETH-1
VNICVLAN=0
VNICIBPORT=mlx4_0:1
```

5. Save the file with the name of `ifcfg-ethXX` where `XX` is the value of column one.

For example, `ifcfg-eth1`.

6. Repeat Step 3 to Step 5 for each entry appropriate for the host.
7. Repeat Step 2 to Step 6 for each host to have VNIC configuration files.

Related Information

- [“Create the Central Configuration File \(Linux\)”](#) on page 154
- [“Host Manual Mode Overview \(Linux\)”](#) on page 149
- [“Central Configuration File \(Linux\)”](#) on page 150
- [“VNIC-Specific Configuration File \(Linux\)”](#) on page 151
- [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)”](#) on page 152
- [“mlx4_vnic_confD Daemon \(Linux\)”](#) on page 157
- [“Configure and Create VNICs for Host Manual Mode \(Linux\)”](#) on page 158

mlx4_vnic_confD Daemon (Linux)

The `mlx4_vnic_confD` daemon is used to manage host administrated VNICs. The daemon is located in the `/etc/init.d` directory of each InfiniBand host. The syntax for the command line is:

```
mlx4_vnic_confD operation
```

where *operation* sets the daemon’s interaction with VNICs.

This table describes the *operations* supported by the `mlx4_vnic_confD` daemon.

Operation	Description
start	Starts, loads, or creates new host administrated VNICs.
stop	Stops all host administrated VNICs.
restart	Closes and then reopens all host administrated VNICs.
reload	Updates the system with the most current configuration files. This operation does not modify host administrated VNICs with unchanged configuration files.

Related Information

- [“Host Manual Mode Overview \(Linux\)” on page 149](#)
- [“Central Configuration File \(Linux\)” on page 150](#)
- [“VNIC-Specific Configuration File \(Linux\)” on page 151](#)
- [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)” on page 152](#)
- [“Create the Central Configuration File \(Linux\)” on page 154](#)
- [“Create the VNIC-Specific Configuration Files \(Linux\)” on page 156](#)
- [“Configure and Create VNICs for Host Manual Mode \(Linux\)” on page 158](#)

▼ Configure and Create VNICs for Host Manual Mode (Linux)

1. **On the management controller, set the maximum number of host created VNICs (Host Manual Mode) per host port per gateway connector.**

```
FabMan@gateway_name->sethostvniclimit connector -host hostname -port port -limit number
-vlanlist vlan_ID1 ... vlan_IDn
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).
- *hostname* is the host name of the management controller or system name recognized by the BridgeX manager.
- *port* is the port of the host.
- *number* is the maximum number of VNICs for the IOA.
- *vlan_ID* is the VLAN identifier (NO or 2 to 4094).

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

For example, to enable Host Manual Mode for the host `cupcake` and limit it to 4 VNICs:

```
FabMan@gateway_name->sethostvniclimit 0A-ETH-1 -host cupcake -port 1 -limit 4
-vlanlist 3
FabMan@gateway_name->
```

2. Repeat [Step 1](#) for all hosts in the fabric that need to be enabled and regulated.

3. Determine if you will use a central configuration file or VNIC-specific configuration files.

See:

- [“Central Configuration File \(Linux\)”](#) on page 150
- [“VNIC-Specific Configuration File \(Linux\)”](#) on page 151

4. Determine the VNIC configuration parameters.

See [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)”](#) on page 152.

5. Create the appropriate configuration files.

See:

- [“Create the Central Configuration File \(Linux\)”](#) on page 154
- [“Create the VNIC-Specific Configuration Files \(Linux\)”](#) on page 156

6. Log in as superuser of the host and start the `mlx4_vnic_conf` daemon.

```
# /etc/init.d/mlx4_vnic_conf start
```

The VNICs are created.

7. Repeat [Step 6](#) for all hosts of the InfiniBand fabric.

Related Information

- [“Configure and Create VNICs for Gateway Manual Mode \(Linux\)”](#) on page 146
- [“Host Manual Mode Overview \(Linux\)”](#) on page 149
- [“Central Configuration File \(Linux\)”](#) on page 150
- [“VNIC-Specific Configuration File \(Linux\)”](#) on page 151

- [“Determine VNIC Configuration Parameters for Host Manual Mode \(Linux\)”](#) on page 152
- [“Create the Central Configuration File \(Linux\)”](#) on page 154
- [“Create the VNIC-Specific Configuration Files \(Linux\)”](#) on page 156
- [“mlx4_vnic_conf Daemon \(Linux\)”](#) on page 157

Creating Virtual IO Adapters (Oracle Solaris)

These topics discuss steps to enable gateway functionality for Oracle Solaris 11 hosts.

Step	Description	Links
1.	Review background information.	“Ethernet Over InfiniBand Overview (Oracle Solaris)” on page 161
2.	Verify the driver is installed.	“Verify That the Ethernet Over InfiniBand Driver Is Installed (Oracle Solaris)” on page 162
3.	(Optional) Install the driver.	“Install the Ethernet Over InfiniBand Driver (Oracle Solaris)” on page 162
4.	Create the VIOAs.	“Set Up VIOAs (Oracle Solaris)” on page 163
5.	Create IP interfaces and assign IP addresses to the VIOAs.	“Configure IP Addresses on the Data Link (Oracle Solaris)” on page 167
6.	Create Oracle Solaris VNICs.	“Create Oracle Solaris VNICs With VLAN Tagging (Oracle Solaris)” on page 169

Related Information

- [“Installing Gateway Supportive Software \(Linux\)”](#) on page 137
- [“Creating VNICs Under Gateway Manual Mode \(Linux\)”](#) on page 142
- [“Creating VNICs Under Host Manual Mode \(Linux\)”](#) on page 148
- [“Monitoring Gateway Resources”](#) on page 172
- [“Controlling LAGs”](#) on page 183
- [“Controlling VLANs and VNICs”](#) on page 190
- [“Controlling Gateway Ports and Parameters”](#) on page 202

Ethernet Over InfiniBand Overview (Oracle Solaris)

The Ethernet over InfiniBand driver is available in current releases of the Oracle Solaris 11 image. This driver supports the Data Link Provider Interface over all InfiniBand ports of an Oracle Solaris 11 host connected to the gateway. The driver uses the IBA unreliable datagram mode to enable initialization, gateway handshake, heartbeat management, frame transmit and receive functions, multicast support, and statistical reporting. The driver is delivered through an IPS-based package with the file name of `ethernet-over-ib`.

Network interfaces in the Oracle Solaris 11 operating system correspond to instances of network hardware devices, and the interfaces are configured over data links. The Ethernet over InfiniBand driver supports virtual network hardware on HCAs and enables interfaces to be assigned to that virtual hardware.

Administrators create IP interfaces on top of data links. A data link represents a link object in the second layer of the OSI model. The respective physical link is directly associated with a device (physical or virtual) and device instance name. The device instance name is comprised of the device driver name augmented with the instance number, which has a value of 0 to n instances of network devices (physical or virtual) using that driver. For each virtual device and interface created on the gateway, there is a corresponding data link named `eoibX` created on the Oracle Solaris 11 host.

Associating the InfiniBand port of an HCA to a gateway Ethernet port, and assigning one or more MAC addresses to the pair creates a virtual IO adapter (VIOA). The `createvnic` command of the gateway is used to fulfill this objective. Oracle Solaris discovers the VIOA, binds an `eoibX` datalink instance to the VIOA, and then manages the VIOA as if it were a physical network interface card (NIC). Like a physical NIC, the VIOA represents the access path to an Ethernet port. Because more than one MAC address can be assigned to a VIOA, the VNIC management commands of the gateway are actually managing the MAC addresses within the gateway itself.

Related Information

- [“Verify That the Ethernet Over InfiniBand Driver Is Installed \(Oracle Solaris\)” on page 162](#)
- [“Install the Ethernet Over InfiniBand Driver \(Oracle Solaris\)” on page 162](#)
- [“Set Up VIOAs \(Oracle Solaris\)” on page 163](#)
- [“Configure IP Addresses on the Data Link \(Oracle Solaris\)” on page 167](#)
- [“Create Oracle Solaris VNICs With VLAN Tagging \(Oracle Solaris\)” on page 169](#)

▼ Verify That the Ethernet Over InfiniBand Driver Is Installed (Oracle Solaris)

1. Become superuser of the Oracle Solaris 11 host.
2. Type.

```
# pkg info ethernet-over-ib
Name: system/io/infiniband/ethernet-over-ib
Summary: Ethernet over InfiniBand (EoIB) Drivers
Description: InfiniBand device driver implementing Ethernet over InfiniBand
Category: System/Hardware
State: Installed
Publisher: solaris
Version: 0.5.11
Build Release: 5.11
Branch: 0.175.1.0.0.2.17991
Packaging Date: Mon Oct 24 10:50:29 2011
Size: 304.74 kB
FMRI:
pkg://solaris/system/io/infiniband/ethernet-over-ib@0.5.11,5.11-0.175.1.0.0.2.17991:20111024T105029Z
#
```

3. If there is no information returned, or an error message is displayed, install the Ethernet over InfiniBand driver.

See [“Install the Ethernet Over InfiniBand Driver \(Oracle Solaris\)”](#) on page 162.

Related Information

- [pkg command man page](#)
- [“Ethernet Over InfiniBand Overview \(Oracle Solaris\)”](#) on page 161
- [“Install the Ethernet Over InfiniBand Driver \(Oracle Solaris\)”](#) on page 162
- [“Set Up VIOAs \(Oracle Solaris\)”](#) on page 163
- [“Configure IP Addresses on the Data Link \(Oracle Solaris\)”](#) on page 167
- [“Create Oracle Solaris VNICs With VLAN Tagging \(Oracle Solaris\)”](#) on page 169

▼ Install the Ethernet Over InfiniBand Driver (Oracle Solaris)

1. Consider your first steps:

- If you know for certain that your Oracle Solaris 11 image contains the `ethernet-over-ib` package, go to [Step 4](#).
 - If you are not certain, go to [Step 2](#).
2. **Download the Oracle Solaris 11 operating system.**
See this URL for instructions.
<http://www.oracle.com/technetwork/server-storage/solaris11/downloads/index.html>
 3. **Install the Oracle Solaris 11 operating system according to the documentation included with the software.**
 4. **As superuser of the Oracle Solaris 11 host, install the Ethernet over InfiniBand package.**

```
# pkg install ethernet-over-ib
```

Follow the prompts.

Related Information

- [pkg command man page](#)
- [“Ethernet Over InfiniBand Overview \(Oracle Solaris\)” on page 161](#)
- [“Verify That the Ethernet Over InfiniBand Driver Is Installed \(Oracle Solaris\)” on page 162](#)
- [“Set Up VIOAs \(Oracle Solaris\)” on page 163](#)
- [“Configure IP Addresses on the Data Link \(Oracle Solaris\)” on page 167](#)
- [“Create Oracle Solaris VNICs With VLAN Tagging \(Oracle Solaris\)” on page 169](#)

▼ Set Up VIOAs (Oracle Solaris)

Note – If the VIOA is to be assigned to a VLAN, you cannot create the VIOA from the gateway. Instead, use the procedure in [“Create Oracle Solaris VNICs With VLAN Tagging \(Oracle Solaris\)” on page 169](#).

1. **Become superuser of the Oracle Solaris 11 host.**

For this procedure, the Oracle Solaris 11 host name is `solaris01`.

2. Display the HCA GUIDs and port GUIDs configured on the host.

```
# dladm show-ib
LINK          HCAGUID          PORTGUID          PORT STATE  PKEYS
ibp0          21280001A0A590  21280001A0A591  1   up     FFFF
ibp1          21280001A0A590  21280001A0A592  2   up     FFFF
#
```

3. Record the host name, HCA GUIDs, and port GUIDs.

In this example, the HCA GUID is 21280001A0A590 and the port GUIDs are 21280001A0A591 and 21280001A0A592.

4. From the management controller of the gateway, display the HCAs recognized by the gateway.

```
FabMan@gateway_name->ibhosts
Ca      : 0x0021280001A0A590 ports 2 "solaris01"
Ca      : 0x00212856cd22c040 ports 2 "SUN IB QDR GW switch mnm34-97 Bridge 1"
Ca      : 0x0002c903000891aa ports 2 "mnm34-54 HCA-1"
Ca      : 0x00212800013ece9e ports 2 "mnm34-55 HCA-1"
Ca      : 0x0003ba000100e370 ports 2 "mnm34-60 HCA-1"
.
.
.
FabMan@gateway_name->
```

5. Compare the output of the `ibhosts` command with the recorded host name, HCA GUIDs, and port GUIDs.

The port GUIDs are the HCA GUID +1 (port 1) and HCA GUID +2 (port 2), respectively.

6. Verify that the gateway recognizes the Oracle Solaris 11 host correctly.

If not, check the cabling between the gateway and the Oracle Solaris 11 host.

7. Determine if a VIOA has already been assigned to the Oracle Solaris 11 host.

```
FabMan@gateway_name->showvnics
ID STATE  FLG IOA_GUID          NODE          IID  MAC          VLN PKEY  GW
-----
4 UP      N 00:03:BA:00:01:00:E3:71 mnm34-60 0000 02:02:02:02:02:03 NO  ffff
0A-ETH-1
5 UP      N 00:03:BA:00:01:00:E3:71 mnm34-60 0002 02:02:02:02:02:04 NO  ffff
0A-ETH-1
```

```
3 DISABLED N 00:03:BA:00:01:00:E3:71 mnm34-60 0000 02:02:02:02:02:02 NO ffff
0A-ETH-1
6 UP N 00:03:BA:00:01:00:E3:72 mnm34-60 0000 02:02:02:02:02:05 NO ffff
0A-ETH-1
FabMan@gateway_name->
```

Note – Until the IP interface is created on the data link (using the `ipadm` command) at the Solaris host, the `STATE` of the VNIC is `WAIT-IOA`.

8. Examine the output by column:

- `STATE` – The state of the VIOA.
- `IOA_GUID` – The port GUID of the VIOA. The GUID might be listed several times.
- `NODE` – The host name of the node hosting the VIOA.
- `MAC` – The MAC address assigned to the VIOA. There might be several MAC addresses.
- `PKEY` – The partition key assigned.
- `GW` – The physical connector cabled to the node or host.

9. Determine your next steps.

- If no port GUID of the Oracle Solaris 11 host is listed, or there are no MAC addresses assigned to the port GUID, go to [Step 10](#).
- Otherwise, go to [“Configure IP Addresses on the Data Link \(Oracle Solaris\)”](#) on [page 167](#).

10. Determine your MAC addresses and whether they are to be global or locally administrated.

Note – Each MAC address must be unique and not `00:00:00:00:00:00`.

Note – Only even numbers are supported for the most significant byte of the MAC address (unicast).

11. Activate the VIOA by assigning a MAC address to it.

```
FabMan@gateway_name->createvnic connector -guid guid -mac mac_address
-pkey default
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).
- *guid* is the global unique identifier of the target port on the host associated with the VIOA.
- *mac_address* is the MAC address to be assigned to the VIOA.

Note – You can use the information from [Step 8](#) as a guide.

For example, to assign the MAC address of a0:a5:91:95:30:9a to the VIOA for port GUID 21280001A0A591 through connector 0A-ETH-2:

```
FabMan@gateway_name->createvnic 0A-ETH-2 -guid 0021280001A0A591 -mac
a0:a5:91:95:30:9a -pkey default
VNIC created
FabMan@gateway_name->
```

The VIOA is created and an Ethernet over InfiniBand data link is bound to the VIOA.

12. Verify that the VIOA is active.

```
FabMan@gateway_name->showvnics
```

ID	STATE	FLG	IOA_GUID	NODE	IID	MAC	VLN	PKEY	GW
7	WAIT-IOA	N	00:21:28:00:01:A0:A5:91	solaris01	0000	a0:a5:91:95:30:9a	NO	ffff	

4	UP	N	00:03:BA:00:01:00:E3:71	mnm34-60	0000	02:02:02:02:02:03	NO	ffff	
0A-ETH-1									
5	UP	N	00:03:BA:00:01:00:E3:71	mnm34-60	0002	02:02:02:02:02:04	NO	ffff	
0A-ETH-1									
3	DISABLED	N	00:03:BA:00:01:00:E3:71	mnm34-60	0000	02:02:02:02:02:02	NO	ffff	
0A-ETH-1									
6	UP	N	00:03:BA:00:01:00:E3:72	mnm34-60	0000	02:02:02:02:02:05	NO	ffff	
0A-ETH-1									

```
FabMan@gateway_name->
```

Note – Until the IP interface is created on the data link (using the `ipadm` command) at the Solaris host, the `STATE` of the VNIC is `WAIT-IOA`.

13. Configure an IP address on the data link.

See “[Configure IP Addresses on the Data Link \(Oracle Solaris\)](#)” on page 167.

Related Information

- `dladm` command man page
- *Gateway Reference*, `ibhosts` command
- *Gateway Reference*, `showvnics` command
- *Gateway Reference*, `createvnics` command
- “Ethernet Over InfiniBand Overview (Oracle Solaris)” on page 161
- “Verify That the Ethernet Over InfiniBand Driver Is Installed (Oracle Solaris)” on page 162
- “Install the Ethernet Over InfiniBand Driver (Oracle Solaris)” on page 162
- “Configure IP Addresses on the Data Link (Oracle Solaris)” on page 167
- “Create Oracle Solaris VNICs With VLAN Tagging (Oracle Solaris)” on page 169

▼ Configure IP Addresses on the Data Link (Oracle Solaris)

1. Become superuser of the Oracle Solaris 11 host.
2. Verify that the data link corresponding to the activated VIOA has been created.

For example:

```
# dladm show-phys
LINK          MEDIA          STATE    SPEED  DUPLEX    DEVICE
vnet0         Ethernet      up       0      unknown  vnet0
ibp0          Infiniband    up       32000  unknown  ibp0
ibp1          Infiniband    down     2000   unknown  ibp1
net5          Ethernet      up       10000  full     eoib0
net4          Ethernet      up       10000  full     eoib1
net6          Ethernet      unknown  10000  full     eoib2
#
```

The `dladm show-phys` command displays all data links corresponding to all hardware devices.

Similarly, the `dladm show-link` command displays the class, MTU, and state for all data links. For example:

```
# dladm show-link
LINK          CLASS    MTU    STATE    OVER
vnet0         phys    1500   up       --
ibp0          phys    65520  up       --
ibp1          phys    65520  down     --
```

```
net5          phys      1500    up      --
net4          phys      1500    up      --
net6          phys      1500    unknown --
#
```

3. Create an IP interface on the data link:

```
# ipadm create-ip link
```

where *link* is the name of the data link. For example:

```
# ipadm create-ip net6
#
```

Note – When you create the IP interface on the data link, the respective VNIC STATE (as seen with the `showvnics` command) goes from `WAIT-IOA` to `UP`.

4. Verify that the interface was created.

```
# ipadm show-if
IFNAME      CLASS      STATE      ACTIVE      OVER
net6        ip         down       no          --
#
```

Note – You must create the IP interface before an IP address is assigned to it.

5. Assign an IP address to the IP interface.

```
# ipadm create-addr -T static -a IP_address/24 IP_interface/v4
```

where:

- *IP_address* is the IP address to be assigned to the IP interface.
- *IP_interface* is the interface corresponding to the data link and VIOA.

For example, to assign the IP address of `192.168.17.24` to IP interface `net6`:

```
# ipadm create-addr -T static -a 192.168.17.24/24 net6/v4
#
```

6. Verify that the IP address was properly assigned.

```
# ipadm show-addr IP_interface/v4
```

where *IP_interface* is the interface corresponding to the data link and VIOA. For example:

```
# ipadm show-addr net6/v4
ADDROBJ    TYPE    STATE    ADDR
net6/v4    static  ok       192.168.17.24/24
#
```

Related Information

- [dladm command man page](#)
- [ipadm command man page](#)
- [“Ethernet Over InfiniBand Overview \(Oracle Solaris\)” on page 161](#)
- [“Verify That the Ethernet Over InfiniBand Driver Is Installed \(Oracle Solaris\)” on page 162](#)
- [“Install the Ethernet Over InfiniBand Driver \(Oracle Solaris\)” on page 162](#)
- [“Set Up VIOAs \(Oracle Solaris\)” on page 163](#)
- [“Create Oracle Solaris VNICs With VLAN Tagging \(Oracle Solaris\)” on page 169](#)

▼ Create Oracle Solaris VNICs With VLAN Tagging (Oracle Solaris)

Note – Follow this procedure to configure Ethernet over InfiniBand with VLAN tagging in Oracle Solaris. This procedure is independent of and unrelated to Host Manual Mode or Gateway Manual Mode. To create VIOAs on the gateway, see [“Set Up VIOAs \(Oracle Solaris\)” on page 163](#).

1. **Create the VIOAs to use as IP interfaces for the Solaris VNICs.**
See [“Set Up VIOAs \(Oracle Solaris\)” on page 163](#).
2. **On the management controller, create a VLAN to P_Key mapping for the VIOAs.**

```
FabMan@gateway_name->createvlan connector -vlan vlan_ID -pkey p_key
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH).
- *vlan_ID* is the VLAN identifier (NO or 2 to 4094).

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

- *p_key* is the partition identifier (1 to 7fff or default).

For example, to associate connector 1A-ETH-3 with VLAN 3 and default partition key:

```
FabMan@gateway_name->createvlan 1A-ETH-3 -vlan 3 -pkey default
FabMan@gateway_name->
```

3. If you are associating VLANs with the Solaris VNICs, perform these steps.

a. Set the maximum number of VNICs per host.

```
FabMan@gateway_name->sethostvniclimit connector -host hostname -port port -limit number
-vlanlist vlan_ID1 ... vlan_IDn
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).
- *hostname* is the host name of the management controller or system name recognized by the BridgeX manager.
- *port* is the port of the host.
- *number* is the maximum number of VNICs for the IOA.
- *vlan_ID* is the VLAN identifier (NO or 2 to 4094).

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

For example, to limit the host cupcake to 4 VNICs:

```
FabMan@gateway_name->sethostvniclimit 0A-ETH-1 -host cupcake -port 1 -limit 4
-vlanlist 3
FabMan@gateway_name->
```

b. Repeat [Step a](#) for all hosts in the fabric that need to be regulated.

4. On the Oracle Solaris 11 host, create a Oracle Solaris VNIC on the data link.

```
# dladm create-vnic -l IP_interface -m mac_address -v vlan_ID name
```

where:

- *IP_interface* is the interface corresponding to the data link and VIOA.
- *mac_address* is the MAC address to be assigned to the VIOA.
- *vlan_ID* is the VLAN identifier (NO or 2 to 4094).

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

- *name* is the name of the VNIC. For ease of identification, the VNIC name can be formatted as *eX_vnicY*, where:
 - *X* is the instance number of the Ethernet over InfiniBand IP interface.
 - *Y* is the instance of the VNIC, independent of the interface instance.

For example:

```
# dladm create-vnic -l net6 -m 2:8:20:7c:5d:d3 -v 3 e0_vnic0
#
```

5. Verify the creation of the VNIC.

```
# dladm show-vnic
LINK          OVER      SPEED    MACADDRESS      MACADDRTYPE  VID
e0_vnic0     net6      10000    2:8:20:7c:5d:d3  random       3
#
```

6. Display a list of all data links.

```
# dladm show-link
LINK          CLASS     MTU      STATE    OVER
net6          phys     1500    up       --
e0_vnic0     vnic     1500    up       net6
```

7. Configure the IP address of the data link `e0_vnic0`.

See [“Configure IP Addresses on the Data Link \(Oracle Solaris\)”](#) on page 167.

8. On the management controller, verify that the VNICs are properly associated and up.

```
FabMan@gateway_name->showvnic
ID STATE FLG IOA_GUID NODE IID MAC VLN PKEY GW
-----
0 UP N 0021280001CF37AC solaris01 00c0 02:42:42:42:12:03 NO 0xffff
0A-ETH-1
1 UP H 0021280001CF37AC solaris01 80c1 02:08:20:7C:5D:D3 3 0xffff
0A-ETH-1
FabMan@gateway_name->
```

Related Information

- [dladm command man page](#)
- [“Ethernet Over InfiniBand Overview \(Oracle Solaris\)” on page 161](#)
- [“Verify That the Ethernet Over InfiniBand Driver Is Installed \(Oracle Solaris\)” on page 162](#)
- [“Install the Ethernet Over InfiniBand Driver \(Oracle Solaris\)” on page 162](#)
- [“Set Up VIOAs \(Oracle Solaris\)” on page 163](#)
- [“Configure IP Addresses on the Data Link \(Oracle Solaris\)” on page 167](#)

Monitoring Gateway Resources

These topics describe how to monitor the gateway interfaces.

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

- [“Display LAG Information” on page 173](#)
- [“Display the VLANs” on page 174](#)
- [“Display the VNICs” on page 176](#)
- [“Display IO Adapters” on page 177](#)
- [“Display Gateway Ethernet Port Information” on page 178](#)
- [“Display Gateway Port Information” on page 179](#)
- [“Display the Gateway Instance Number” on page 181](#)

- “Display Information About the Gateway” on page 182

Related Information

- “Monitoring the Chassis” on page 19
- “Monitoring the I4 Switch Chip” on page 31
- “Monitoring the InfiniBand Fabric” on page 83
- “Monitoring the Subnet Manager” on page 49
- “Installing Gateway Supportive Software (Linux)” on page 137
- “Creating VNICs Under Gateway Manual Mode (Linux)” on page 142
- “Creating VNICs Under Host Manual Mode (Linux)” on page 148
- “Creating Virtual IO Adapters (Oracle Solaris)” on page 160
- “Controlling LAGs” on page 183
- “Controlling VLANs and VNICs” on page 190
- “Controlling Gateway Ports and Parameters” on page 202

▼ Display LAG Information

The `showlag` command displays general information about all LAGs or detailed information for a specific LAG.

1. On the management console, display general information about all LAGs.

```
FabMan@gateway_name->showlag
Name          Iport      Members Vnics Gw Port ID BANDWIDTH Distribution LACP  Mode
DiscoverPkey Type    SL
-----
-----
bxmtest0     Bridge-0-1 0        1000  2704      NA      NA      NA
0xffff      Global 1
FabMan@gateway_name->
```

2. Display detailed information about a specific LAG.

```
FabMan@gateway_name->showlag lagname
```

where *lagname* is the identifier of the LAG

For example

```
FabMan@gateway_name->showlag test
Name          enabled active RC
-----
0A-ETH-2     1         1    0x0
0A-ETH-3     1         1    0x0
FabMan@gateway_name->
```

3. If the information displayed is incorrect, modify the LAGs.

See “Add or Delete Connectors From a LAG” on page 187.

Related Information

- *Gateway Reference*, showlag command
- “Controlling LAGs” on page 183
- “Display the VLANs” on page 174
- “Display the VNICs” on page 176
- “Display IO Adapters” on page 177
- “Display Gateway Ethernet Port Information” on page 178
- “Display Gateway Port Information” on page 179
- “Display the Gateway Instance Number” on page 181
- “Display Information About the Gateway” on page 182

▼ Display the VLANs

You can use the showvlan command to identify which VLANs are associated with the gateway 10GbE connectors.

1. On the management controller, type.

```
FabMan@gateway_name->showvlan
Connector/LAG  VLN  PKEY
-----
0A-ETH-1      3    ffff
0A-ETH-1      0    ffff
```

0A-ETH-2	3	ffff
0A-ETH-2	0	ffff
0A-ETH-3	3	ffff
0A-ETH-3	0	ffff
0A-ETH-4	3	ffff
0A-ETH-4	0	ffff
1A-ETH-1	3	ffff
1A-ETH-1	0	ffff
1A-ETH-2	3	ffff
1A-ETH-2	0	ffff
1A-ETH-3	3	ffff
1A-ETH-3	0	ffff
1A-ETH-4	3	ffff
1A-ETH-4	0	ffff

FabMan@gateway_name->

Note – A VLAN identifier (VLN) of 0, -1, or NO indicates that there is no VLAN assigned.

Note – Your output will differ from that in the example.

The `showvlan` command lists the connectors on the left, and the associated VLANs and respective partition keys on the right.

2. If the information displayed is incorrect, modify the VLAN configuration.

See [“Controlling VLANs and VNICs”](#) on page 190.

Related Information

- [Gateway Reference](#), `showvlan` command
- [“Controlling VLANs and VNICs”](#) on page 190
- [“Display LAG Information”](#) on page 173
- [“Display the VNICs”](#) on page 176
- [“Display IO Adapters”](#) on page 177
- [“Display Gateway Ethernet Port Information”](#) on page 178
- [“Display Gateway Port Information”](#) on page 179
- [“Display the Gateway Instance Number”](#) on page 181
- [“Display Information About the Gateway”](#) on page 182

▼ Display the VNICs

The `showvnics` command displays information about and status of the VNICs.

1. On the management controller, type.

```
FabMan@gateway_name->showvnics
ID STATE   FLG IOA_GUID                NODE      IID  MAC                      VLN PKEY  GW
-----
-----
 16 UP      N 00:21:28:00:01:3E:CE:A0 mnm34-55  0000 02:02:02:02:02:04 3
ffff 1A-ETH-4
 13 UP      N 00:03:BA:00:01:00:E3:71 mnm34-60  0000 02:02:02:02:02:01 3
ffff 1A-ETH-3
 15 UP      N 00:21:28:00:01:3E:CE:9F mnm34-55  0000 02:02:02:02:02:03 3
ffff 1A-ETH-1
 14 UP      N 00:03:BA:00:01:00:E3:72 mnm34-60  0000 02:02:02:02:02:02 3
ffff 1A-ETH-2
FabMan@gateway_name->
```

Note – Your output will differ from that in the example.

For each VNIC, the `showvnics` command displays the VNIC ID, its state, any flags, its associated host GUID and node description, its MAC address, its associated VLAN ID, its partition key, and connector.

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

See “Modify VNICs” on page 197.

Related Information

- [Gateway Reference](#), `showvnics` command
- [“Controlling VLANs and VNICs”](#) on page 190
- [“Display LAG Information”](#) on page 173
- [“Display the VLANs”](#) on page 174
- [“Display IO Adapters”](#) on page 177
- [“Display Gateway Ethernet Port Information”](#) on page 178
- [“Display Gateway Port Information”](#) on page 179
- [“Display the Gateway Instance Number”](#) on page 181
- [“Display Information About the Gateway”](#) on page 182

▼ Display IO Adapters

The `showioadapters` command displays information about the IO adapters. IO adapters represent an HCA port GUID where an EoIB driver is active on the host side.

1. On the management controller, type.

```
FabMan@gateway_name->showioadapters
IOA_GUID                NODE          LID  FLAGS  GW
-----
0021280001CF205B        o4test56     1    HD     1A-ETH-2
0021280001CF3787        o4test57     2    HD     1A-ETH-2
0021280001CF205B        o4test56     1    HD     1A-ETH-1
.
.
.
0021280001CF3787        o4test57     2    HD     0A-ETH-3
0021280001CF205B        o4test56     1    HD     0A-ETH-4
0021280001CF3787        o4test57     2    HD     0A-ETH-4
FabMan@gateway_name->
```

Note – Your output will differ from that in the example.

The `FLAGS` indicate whether the VNIC was created by Host Manual Mode (HD), Gateway Manual Mode (ND), or other means (no flag).

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

See [“Modify VNICs”](#) on page 197 or [“Create Oracle Solaris VNICs With VLAN Tagging \(Oracle Solaris\)”](#) on page 169.

Related Information

- [Gateway Reference](#), `showioadapters` command
- [“Controlling VLANs and VNICs”](#) on page 190
- [“Display LAG Information”](#) on page 173
- [“Display the VLANs”](#) on page 174
- [“Display the VNICs”](#) on page 176
- [“Display Gateway Ethernet Port Information”](#) on page 178
- [“Display Gateway Port Information”](#) on page 179
- [“Display the Gateway Instance Number”](#) on page 181
- [“Display Information About the Gateway”](#) on page 182

▼ Display Gateway Ethernet Port Information

You can display Ethernet port information with the `setgwethport` command. The output is similar to that of the `getportstatus` command.

1. On the management controller, type.

```
FabMan@gateway_name->setgwethport connector
```

where *connector* is either 0A-ETH or 1A-ETH. For example:

Note – Your output will differ from that of the example.

```
FabMan@gateway_name->setgwethport 0A-ETH
Port status for connector 0A-ETH-1:
Adminstate.....Enabled
State.....Up
Link state.....Up
Protocol.....Ethernet
Link Mode.....XFI
Speed.....10Gb/s
MTU.....9600
Tx pause.....Global
Rx pause.....Global
Port status for connector 0A-ETH-2:
Adminstate.....Enabled
.
.
.
MTU.....9600
Tx pause.....Global
Rx pause.....Global
FabMan@gateway_name->
```

2. If the information displayed is incorrect, modify the gateway Ethernet port parameters.

See [“Set the Gateway Ethernet Port Parameters”](#) on page 205.

Related Information

- [Gateway Reference](#), `setgwethport` command
- [Gateway Reference](#), `getportstatus` command
- [“Set the Gateway Ethernet Port Parameters”](#) on page 205
- [“Display LAG Information”](#) on page 173

- “Display the VLANs” on page 174
- “Display the VNICs” on page 176
- “Display IO Adapters” on page 177
- “Display Gateway Port Information” on page 179
- “Display the Gateway Instance Number” on page 181
- “Display Information About the Gateway” on page 182

▼ Display Gateway Port Information

Use the `showgwports` command to display detailed information about the gateway ports.

● On the management controller, type.

```
FabMan@gateway_name->showgwports -v

BRIDGE DEVICES:
-----
NodeGUID                NodeDescription
-----
0x00212856d102c000 SUN IB QDR GW switch o4nm2-gw-6 10.172.144.77 Bridge 0
0x00212856d102c040 SUN IB QDR GW switch o4nm2-gw-6 10.172.144.77 Bridge 1

INTERNAL PORTS:
-----
Device  Port Portname  PeerPort  PortGUID                LID    IBState  GWState
Speed  VLs  MTU
-----
Bridge-0 1  Bridge-0-1  4    0x00212856d102c001 0x0002 Active  Up       40Gbs
2 4096
Bridge-0 2  Bridge-0-2  3    0x00212856d102c002 0x0003 Active  Up       40Gbs
2 4096
Bridge-1 1  Bridge-1-1  2    0x00212856d102c041 0x0004 Active  Up       40Gbs
2 4096
Bridge-1 2  Bridge-1-2  1    0x00212856d102c042 0x0005 Active  Up       40Gbs
2 4096

CONNECTOR 0A-ETH:
-----
Port      Bridge      Adminstate Link  State      MTU  TxPause  RxPause
Linkmode      Speed
-----
0A-ETH-1 Bridge-0-2 Enabled  Up  Up        9600 Global  Global    XFI
10Gb/s
0A-ETH-2 Bridge-0-2 Enabled  Up  Up        9600 Global  Global    XFI
10Gb/s
0A-ETH-3 Bridge-0-1 Enabled  Up  Up        9600 Global  Global    XFI
10Gb/s
0A-ETH-4 Bridge-0-1 Enabled  Up  Up        9600 Global  Global    XFI
10Gb/s

CONNECTOR 1A-ETH:
-----
Port      Bridge      Adminstate Link  State      MTU  TxPause  RxPause
Linkmode      Speed
-----
-----
```

1A-ETH-1	Bridge-1-2	Enabled	Up	Up	9600	Global	Global	XFI
10Gb/s								
1A-ETH-2	Bridge-1-2	Enabled	Up	Up	9600	Global	Global	XFI
10Gb/s								
1A-ETH-3	Bridge-1-1	Enabled	Up	Up	9600	Global	Global	XFI
10Gb/s								
1A-ETH-4	Bridge-1-1	Enabled	Up	Up	9600	Global	Global	XFI
10Gb/s								
FabMan@gateway_name->								

When the `-v` option is used, this information is provided:

- BridgeX devices – The node GUID and node description are provided.
- Internal ports – The device, port, port name, peer port, port GUID, LID, InfiniBand and gateway state, speed and MTU are displayed.
- Ethernet connectors – For each connector, the port, associated bridge, the administrative, link, and physical state, MTU, transmit and receive pause, link mode, and speed are provided.

Note – Your output will differ from that of the example.

Related Information

- *Gateway Reference*, `showgwports` command
- “Set the Gateway Ethernet Port Parameters” on page 205
- “Display LAG Information” on page 173
- “Display the VLANs” on page 174
- “Display the VNICs” on page 176
- “Display IO Adapters” on page 177
- “Display Gateway Ethernet Port Information” on page 178
- “Display the Gateway Instance Number” on page 181
- “Display Information About the Gateway” on page 182

▼ Display the Gateway Instance Number

Gateways use instance numbers to identify themselves to each other.

1. On the management controller, type.

```
FabMan@gateway_name->setgwinstance list
Current GW instance: 34
FabMan@gateway_name->
```

2. If the information displayed is incorrect, modify the gateway instance.

See “Set the Gateway Instance Number” on page 207.

Related Information

- *Gateway Reference*, setgwinstance command
- “Set the Gateway Instance Number” on page 207
- “Display LAG Information” on page 173
- “Display the VLANs” on page 174
- “Display the VNICs” on page 176
- “Display IO Adapters” on page 177
- “Display Gateway Ethernet Port Information” on page 178
- “Display Gateway Port Information” on page 179
- “Display Information About the Gateway” on page 182

▼ Display Information About the Gateway

You can display version, status, and operational information about the gateway with the showgwconfig command.

1. On the management controller, type.

```
FabMan@gateway_name->showgwconfig
BXN (pid 13724) is running
BXN versions: bxn_user 2.0.0898-0, BXN-API 1.6.0, bxn_libs 2.0.0898-0, bxn_main
1.31 mlx_bx_core 1.31
Parameter                Configured Value      Running Value
-----
GWInstance                42                    42
SystemName                None                  o4nm2-gw-4
EoIB Data SL              None                  1
EoIB Control SL           None                  2
Allow host VNIC config   None                  no
LAG mode                  yes                   yes
```

Default discover P_key	None	0xffff
System MAC	Not applicable	00:3b:aa:bb:a0:15
FabMan@gateway_name->		

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

See:

- [“Set the Gateway Instance Number” on page 207](#)
- [“Set the Gateway System Name” on page 208](#)
- [“Set the Gateway Service Level” on page 206](#)
- *Gateway Reference*, `allowhostconfig` command
- [“Enable LAG Mode” on page 184](#)
- [“Disable LAG Mode” on page 189](#)
- *Gateway Reference*, `setdefaultdiscpkey` command

Related Information

- *Gateway Reference*, `showgwconfig` command
- [“Display LAG Information” on page 173](#)
- [“Display the VLANs” on page 174](#)
- [“Display the VNICs” on page 176](#)
- [“Display IO Adapters” on page 177](#)
- [“Display Gateway Ethernet Port Information” on page 178](#)
- [“Display Gateway Port Information” on page 179](#)
- [“Display the Gateway Instance Number” on page 181](#)

Controlling LAGs

LAGs are a means of effectively increasing the bandwidth of an Ethernet network by grouping ports.

- [“Enable LAG Mode” on page 184](#)
- [“Create LAGs” on page 185](#)
- [“Create a Global Gateway” on page 186](#)
- [“Add or Delete Connectors From a LAG” on page 187](#)
- [“Delete a LAG” on page 188](#)
- [“Disable LAG Mode” on page 189](#)

Related Information

- “Display LAG Information” on page 173
- “Controlling VLANs and VNICs” on page 190
- “Controlling Gateway Ports and Parameters” on page 202
- “Controlling the Chassis” on page 28
- “Controlling the I4 Switch Chip” on page 40
- “Controlling the InfiniBand Fabric” on page 97
- “Controlling the Subnet Manager” on page 56
- “Installing Gateway Supportive Software (Linux)” on page 137
- “Creating VNICs Under Gateway Manual Mode (Linux)” on page 142
- “Creating VNICs Under Host Manual Mode (Linux)” on page 148
- “Creating Virtual IO Adapters (Oracle Solaris)” on page 160
- “Monitoring Gateway Resources” on page 172

▼ Enable LAG Mode

Note – LAG mode must be enabled on all InfiniBand gateways in the fabric before LAGs are created.

1. On the management controller, type.

```
FabMan@gateway_name->enablelagmode
VNICs for host nodes with EoIB driver version that do not support LAG will go
down if LAG mode is enabled.
Do you still want to enable LAG(y/n)?y
Stopping Bridge Manager..                [ OK ]
Starting Bridge Manager.                  [ OK ]
FabMan@gateway_name->
```

2. Repeat Step 1 for all InfiniBand gateways in the fabric.

Related Information

- *Gateway Reference*, enablelagmode command
- “Create LAGs” on page 185
- “Add or Delete Connectors From a LAG” on page 187
- “Delete a LAG” on page 188
- “Disable LAG Mode” on page 189

▼ Create LAGs

Once LAG mode is enabled, you can create LAGs.

1. Determine which connectors will be grouped into a LAG and what the LAG names will be.

Note – A connector is assigned to only one LAG at a time.

Note – If a VNIC has been assigned to a connector, that connector cannot be added to a LAG until the respective VNIC has been deleted.

2. On the management controller, create the first LAG.

```
FabMan@gateway_name->createlag lagname [-discoverykey p_key] [-datasl
sl] connector1 [connector2 [...connectorN]] [-n] [-lacpmode
active|passive|off] [-vniclimit 1|4]
```

where:

- *lagname* is the identifier of the LAG
- *connector1* is the name of the first connector to be added to the LAG (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).
- *connectorN* is the name of the last connector to be added to the LAG (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).
- *p_key* is the discovery partition key (1 to 7fff or default).
- *sl* is the data service level.

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

Note – The *-vniclimit* option sets the maximum number of VNICs to external port associations (1=1024 VNICs, 4=4096 VNICs).

For example:

```
FabMan@gateway_name->createlag 0A12 0A-ETH-1 0A-ETH-2
lag created
FabMan@gateway_name->
```

3. Repeat [Step 2](#) for any other LAGs to be created.

Related Information

- *Gateway Reference*, `createlag` command
- “Enable LAG Mode” on page 184
- “Add or Delete Connectors From a LAG” on page 187
- “Delete a LAG” on page 188
- “Disable LAG Mode” on page 189

▼ Create a Global Gateway

A Global Gateway is a way to use EoIB functionality internal to the InfiniBand subnet without connections to any specific external port of the gateway. Use the `createlag` command without specifying any connectors to create a Global Gateway.

Note – A Global Gateway can accommodate up to 8000 VNICs, and you can create a maximum of 16 Global Gateways per Sun Network QDR InfiniBand Gateway Switch. However, the Sun Network QDR InfiniBand Gateway Switch only supports a total of 8192 VNICs, regardless of how they are distributed to LAGs, Global Gateways, and external ports.

1. Enable LAG mode.

See “Enable LAG Mode” on page 184.

2. On the management controller, create the Global Gateway.

```
FabMan@gateway_name->createlag lagname [-discoverypkey p_key] [-datasl sl]
```

where:

- *lagname* is the identifier of the Global Gateway.
- *p_key* is the discovery partition key (1 to 7fff or default). Setting a discovery P_Key limits the Global Gateway’s visibility.
- *sl* is the data service level (0 to 15) and is used to discriminate between different classes of InfiniBand traffic.

Note – Use service level 1 (s1 1) only for low-latency, high-priority, small-message, low-bandwidth traffic. Use other service levels for regular, high-bandwidth traffic.

For example:

```
FabMan@gateway_name->createlag globalgate -datas1 2
lag created
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, createlag command
- “Enable LAG Mode” on page 184
- “Disable LAG Mode” on page 189

▼ Add or Delete Connectors From a LAG

You can add more connectors to a LAG to increase its bandwidth. Or, you can remove connectors from a LAG to free them up for other purposes.

Note – A connector is assigned to only one LAG at a time.

Note – If a VNIC has been assigned to a connector, that connector cannot be added to a LAG until the respective VNIC has been deleted.

Note – If the LAG is a Global Gateway, you cannot add any connectors to the LAG.

1. To add a connector to a LAG, on the management controller, type.

```
FabMan@gateway_name->addlagport lagname connector1 [connector2
[...connectorN]]
```

where:

- *lagname* is the identifier of the LAG
- *connector1* is the name of the first connector to be added to the LAG (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).

- *connectorN* is the name of the last connector to be added to the LAG (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).

For example:

```
FabMan@gateway_name->addlagport 0A12 0A-ETH-3
lag updated
FabMan@gateway_name->
```

2. Or, to delete a connector from a LAG, type.

```
FabMan@gateway_name->dellagport lagname connector1 [connector2
[...connectorN]]
```

where:

- *lagname* is the identifier of the LAG
- *connector1* is the name of the first connector to be removed from the LAG (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).
- *connectorN* is the name of the last connector to be removed from the LAG (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4).

For example:

```
FabMan@gateway_name->dellagport 0A12 0A-ETH-3
lag updated
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `dellagport` command
- [“Enable LAG Mode” on page 184](#)
- [“Create LAGs” on page 185](#)
- [“Delete a LAG” on page 188](#)
- [“Disable LAG Mode” on page 189](#)

▼ Delete a LAG

When a LAG is no longer needed, or its connectors must be re-purposed, you can delete the LAG.

Note – This same procedure deletes a Global Gateway.

- On the management controller, type.

```
FabMan@gateway_name->deletelag lagname
```

where *lagname* is the identifier of the LAG

For example:

```
FabMan@gateway_name->deletelag OA12  
LAG OA12 deleted  
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `deletelag` command
- “Enable LAG Mode” on page 184
- “Create LAGs” on page 185
- “Add or Delete Connectors From a LAG” on page 187
- “Disable LAG Mode” on page 189

▼ Disable LAG Mode

When you disable LAG mode, the LAG configurations remain. Upon re-enabling LAG mode, the LAGs become active again.

Note – LAG mode must be disabled for all InfiniBand gateways in the fabric.

1. On the management controller, type.

```
FabMan@gateway_name->disablelagmode  
Stopping Bridge Manager.. [ OK ]  
Starting Bridge Manager.. [ OK ]  
FabMan@gateway_name->
```

2. Repeat [Step 1](#) for all InfiniBand gateways in the fabric.

Related Information

- *Gateway Reference*, `disablelagmode` command
- “Enable LAG Mode” on page 184
- “Create LAGs” on page 185
- “Add or Delete Connectors From a LAG” on page 187

- “Delete a LAG” on page 188

Controlling VLANs and VNICs

These topics describe how you can control the gateway ports and parameters:

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

- “VNIC Allocation on the Gateway” on page 191
- “Create VLANs” on page 192
- “Delete VLANs” on page 194
- “Create VNICs” on page 195
- “Modify VNICs” on page 197
- “Enable VNICs” on page 199
- “Disable VNICs” on page 200
- “Delete VNICs” on page 201

Related Information

- “Controlling LAGs” on page 183
- “Controlling Gateway Ports and Parameters” on page 202
- “Controlling the Chassis” on page 28
- “Controlling the I4 Switch Chip” on page 40
- “Controlling the InfiniBand Fabric” on page 97
- “Controlling the Subnet Manager” on page 56
- “Installing Gateway Supportive Software (Linux)” on page 137
- “Creating VNICs Under Gateway Manual Mode (Linux)” on page 142
- “Creating VNICs Under Host Manual Mode (Linux)” on page 148
- “Creating Virtual IO Adapters (Oracle Solaris)” on page 160
- “Monitoring Gateway Resources” on page 172

VNIC Allocation on the Gateway

By default, VNICs are created and allocated by the management controller of the gateway using Gateway Manual Mode. The `createvnic` command creates the VNICs, and the resulting VNIC resources are controlled by the gateway.

Alternatively, individual hosts can create and allocate VNICs on the gateway. There are two ways to enable this functionality.

- **sethostvniclimit** – Running the `sethostvniclimit` command on the management controller of the gateway restricts the number of VNICs that can be created by specified host ports, for specified external ports 0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4. Using the `sethostvniclimit` command to allow host created VNICs is optimal, because the VNIC resources are controlled by the gateway.
- **allowhostconfig** – Running the `allowhostconfig` command on the management controller of the gateway allows all hosts connected to the gateway through the InfiniBand fabric to create unlimited VNICs on any external port (0A-ETH-1 to 0A-ETH-4 and 1A-ETH-1 to 1A-ETH-4). In this situation, VNIC resources are not controlled by the gateway. Use of the `allowhostconfig` command requires consideration of the impact on the gateway's performance.

Only one of these methods can be used at a time. You can display which method is available for use with the `showgwconfig` command. If in the output of the `showgwconfig` command the Running Value for the Allow host VNIC config parameter is `yes`, then unlimited host created VNICs functionality is enabled (`allowhostconfig`). If you want to restrict host created VNICs, use the `disallowhostconfig` command to change the Running Value to `no` (as seen with the `showgwconfig` command), and then use the `sethostvniclimit` command.

Related Information

- *Gateway Reference*, `sethostvniclimit` command
- *Gateway Reference*, `allowhostconfig` command
- *Gateway Reference*, `showgwconfig` command
- *Gateway Reference*, `disallowhostconfig` command
- “Create VNICs” on page 195
- “Enable VNICs” on page 199
- “Disable VNICs” on page 200
- “Delete VNICs” on page 201

▼ Create VLANs

When you create a VLAN, you create a mapping between a gateway connector, a VLAN identifier, and a partition key. You can create VLANs with the `createvlan` command. If a VNIC will be associated to a VLAN, you must create the VLAN before the VNIC.

Note – Due to hardware limitations for MultiCast groups, there is a maximum of 1000 VLANs.

Note – Regardless if you plan to use VLANs, you must configure a no VLAN to P_Key relationship. For example, the VLAN ID, NO, is assigned to the P_Key, default.

Note – The association of VLAN identifier NO to partition key default survives firmware upgrades.

1. If unknown, determine the connectors, VNICs, and partition keys to be associated with the VLANs.

See “Determine VLAN Associations for Gateway Manual Mode (Linux)” on page 145.

2. (Optional) Use the gateway MAC address list that you created earlier.

See “Determine VLAN Associations for Gateway Manual Mode (Linux)” on page 145.

3. Create the VLAN on the management controller.

```
FabMan@gateway_name->createvlan connector |lagname -vlan vlan_ID -pkey p_key
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH). This value is column one of the gateway MAC address list.
- *lagname* is the identifier of the LAG.
- *vlan_ID* is the VLAN identifier (NO or 2 to 4094). This value is column four of the gateway MAC address list.

Note – Do not use VLAN identifiers 0, 1, or 4095. According to the IEEE 802.1Q specification, VLAN ID 0 is used for priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

- *p_key* is the partition identifier (1 to 7fff or default). This value is column five of the gateway MAC address list.

For example, to associate connector 1A-ETH-3 with VLAN 3 and default partition key:

```
FabMan@gateway_name->createvlan 1A-ETH-3 -vlan 3 -pkey default
FabMan@gateway_name->
```

4. Repeat Step 3 for all VLANs to be associated with connectors.

Note – For the `createvlan` command, given a *connector* | *lagname* and *vlan_id* combination, there is only one *p_key* permitted.

5. Verify the VLAN configuration.

```
FabMan@gateway_name->showvlan
Connector/LAG  VLN  PKEY
-----
0A-ETH-1      0    ffff
0A-ETH-2      0    ffff
0A-ETH-3      0    ffff
0A-ETH-4      0    ffff
1A-ETH-1      3    ffff
1A-ETH-1      0    ffff
1A-ETH-2      3    ffff
1A-ETH-2      0    ffff
1A-ETH-3      3    ffff
1A-ETH-3      0    ffff
1A-ETH-4      3    ffff
1A-ETH-4      0    ffff
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `createvlan` command
- *Gateway Reference*, `showvlan` command
- “Delete VLANs” on page 194
- “Create VNICs” on page 195
- “Enable VNICs” on page 199

- “Disable VNICs” on page 200
- “Delete VNICs” on page 201
- “Display the VLANs” on page 174

▼ Delete VLANs

You can delete VLANs with the `deletevlan` command.

1. Delete all VNICs associated with the connector for the VLAN you are deleting.
See “Delete VNICs” on page 201.
2. On the management controller, type.

```
FabMan@gateway_name->deletevlan lagname |connector -vlan vlan_ID
```

where:

- *lagname* is the identifier of the LAG.
- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH).
- *vlan_ID* is the VLAN identifier (2 to 4094).

For example, to delete the association of connector 1A-ETH-1 to VLAN 3, type.

```
FabMan@gateway_name->deletevlan 1A-ETH-1 -vlan 3
FabMan@gateway_name->
```

3. Repeat from [Step 1](#) for all VLANs to be deleted.
4. Verify that the VLANs have been deleted.

```
FabMan@gateway_name->showvlan
Connector/LAG  VLN  PKEY
-----
0A-ETH-1      0    ffff
0A-ETH-2      0    ffff
0A-ETH-3      0    ffff
0A-ETH-4      0    ffff
1A-ETH-1      0    ffff
1A-ETH-2      0    ffff
1A-ETH-3      0    ffff
1A-ETH-4      0    ffff
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `deletevlan` command
- *Gateway Reference*, `showvlan` command
- “Create VLANs” on page 192
- “Create VNICs” on page 195
- “Enable VNICs” on page 199
- “Disable VNICs” on page 200
- “Delete VNICs” on page 201
- “Display the VLANs” on page 174

▼ Create VNICs

You must create VLANs before VNICs. All VNICs are assigned to a VLAN, even if the VLAN ID is NO. For consistency, you can assign the VLAN ID of NO to the P_Key of default.

1. **If unknown, determine the HCA port GUIDs, MAC addresses, the partition keys, and the connectors that will be assigned to the VNICs.**

See “Determine VNIC Configuration Parameters For Gateway Manual Mode (Linux)” on page 144.

2. **If a VLAN will be associated with the VNICs, create the VLANs before the VNICs.**

See “Create VLANs” on page 192.

3. **(Optional) Use the gateway MAC address list you created earlier.**

See “Determine VNIC Configuration Parameters For Gateway Manual Mode (Linux)” on page 144 or “Determine VLAN Associations for Gateway Manual Mode (Linux)” on page 145.

4. **Create the VNIC on the management controller.**

```
FabMan@gateway_name->createvnic connector|lagname <-guid guid>|<-host  
hostname -port port>|<-node string -port port> -mac mac -vlan vlan_ID -pkey  
p_key
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH). This value is column one of the gateway MAC address list.
- *lagname* is the identifier of the LAG.

- *guid* is the global unique identifier of the target port on the host associated with the VNIC. This value is column two of the gateway MAC address list.
- *hostname* is the host name of the management controller or system name recognized by the BridgeX manager.
- *port* is the port of the host.
- *string* is the character string in the node description field.
- *mac* is the machine address code. This value is column three of the gateway MAC address list.

Note – Only even numbers are supported for the most significant byte of the MAC address (unicast).

Note – If you do not assign a MAC address, the VNIC is not utilized. You must later assign a MAC address to the VNIC with the `modifyvnic` command if you want to use the VNIC.

- *vlan_ID* is the VLAN identifier. This value is column four of the gateway MAC address list.
- *p_key* is the partition identifier (default). This value is column five of the gateway MAC address list.

For example, to create a persistent VNIC for connector 1A-ETH-3 associated with VLAN 3 with partition key `default`, type.

```
FabMan@gateway_name->createvnic 1A-ETH-3 -guid 00:03:BA:00:01:00:E3:71 -mac
02:02:02:02:02:01 -vlan 3 -pkey default
FabMan@gateway_name->
```

Note – If the VNIC is to be non-persistent, append the `-n` option to the command line.

Note – For VNICs created by the `createvnic` command, the Network Interface names at the host which is set by the `DEVICE` parameter in the `ifcfg-ethX` (where *X* is identifier), the name should be different from `eth0` to `eth7`, since these names are reserved for devices created and controlled from the host. A suggested name scheme would be `ethgateway_instance_VNIC_ID`, but all names apart for `eth0` to `eth7` would be valid.

5. Repeat [Step 4](#) for all VNICs to be created.

6. Verify the status and configuration of the VNICs.

```
FabMan@gateway_name->showvnics
ID STATE   FLG IOA_GUID          NODE      IID  MAC              VLN PKEY  GW
-----
----  -
 16 UP      N 00:21:28:00:01:3E:CE:A0 mnm34-55  0000 02:02:02:02:02:04 3
ffff 1A-ETH-4
 13 UP      N 00:03:BA:00:01:00:E3:71 mnm34-60  0000 02:02:02:02:02:01 3
ffff 1A-ETH-3
 15 UP      N 00:21:28:00:01:3E:CE:9F mnm34-55  0000 02:02:02:02:02:03 3
ffff 1A-ETH-1
 14 UP      N 00:03:BA:00:01:00:E3:72 mnm34-60  0000 02:02:02:02:02:02 3
ffff 1A-ETH-2
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `createvnic` command
- *Gateway Reference*, `showvnics` command
- [“VNIC Allocation on the Gateway” on page 191](#)
- [“Create VLANs” on page 192](#)
- [“Delete VLANs” on page 194](#)
- [“Enable VNICs” on page 199](#)
- [“Disable VNICs” on page 200](#)
- [“Delete VNICs” on page 201](#)
- [“Display the VNICs” on page 176](#)

▼ Modify VNICs

You can use the `modifyvnic` command to change the VNIC’s MAC address or the VLAN with which the VNIC is associated.

1. **Identify the VNIC you want to modify.**
See [“Display the VNICs” on page 176](#).
2. **On the management controller, type:**

```
# modifyvnic connector vnic_ID -mac mac -vlan vlan_ID
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH).
- *vnic_ID* is the VNIC identifier.
- *mac* is the machine address code.
- *vlan_ID* is the VLAN identifier (NO or 2 to 4094).

For example, to modify VNIC 15 on connector 1A-ETH-1 so that it has a different MAC address, type:

```
# modifyvnic 1A-ETH-1 15 -mac 02:02:02:02:02:88
VNIC ID 15 modified
#
```

3. Repeat from [Step 1](#) for all VNICs to be modified.

4. Verify that the VNICs have been modified:

```
# showvnics
ID STATE FLG IOA_GUID NODE IID MAC VLN PKEY GW
-----
-----
16 UP N 00:21:28:00:01:3E:CE:A0 mnm34-55 0000 02:02:02:02:02:04 3
9234 1A-ETH-4
13 UP N 00:03:BA:00:01:00:E3:71 mnm34-60 0000 02:02:02:02:02:01 3
9234 1A-ETH-3
15 UP N 00:21:28:00:01:3E:CE:9F mnm34-55 0000 02:02:02:02:02:88 3
9234 1A-ETH-1
14 UP N 00:03:BA:00:01:00:E3:72 mnm34-60 0000 02:02:02:02:02:02 3
9234 1A-ETH-2
#
```

Related Information

- [Gateway Reference](#), `modifyvnic` command
- [Gateway Reference](#), `showvnics` command
- [“VNIC Allocation on the Gateway”](#) on page 191
- [“Create VLANs”](#) on page 192
- [“Delete VLANs”](#) on page 194
- [“Create VNICs”](#) on page 195
- [“Enable VNICs”](#) on page 199
- [“Disable VNICs”](#) on page 200
- [“Delete VNICs”](#) on page 201
- [“Display the VNICs”](#) on page 176

▼ Enable VNICs

VNICs are automatically enabled when they are created. You can use this procedure to enable a previously disabled VNIC.

1. Identify the VNIC you want to enable.

See “Display the VNICs” on page 176.

2. Enable the VNIC on the management controller.

```
FabMan@gateway_name->enablevnic connector vnic_ID
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH).
- *vnic_ID* is the VNIC identifier.

For example, to enable VNIC 15 on connector 1A-ETH-1, type.

```
FabMan@gateway_name->enablevnic 1A-ETH-1 15
FabMan@gateway_name->
```

3. Repeat from Step 1 for all VNICs to be enabled.

4. Verify that the VNICs have been enabled.

```
FabMan@gateway_name->showvnics
ID  STATE  FLG  IOA_GUID                NODE      IID  MAC                VLN  PKEY  GW
-----
16  UP      N    00:21:28:00:01:3E:CE:A0 mnm34-55  0000 02:02:02:02:02:04 3
ffff 1A-ETH-4
13  UP      N    00:03:BA:00:01:00:E3:71 mnm34-60  0000 02:02:02:02:02:01 3
ffff 1A-ETH-3
15  UP      N    00:21:28:00:01:3E:CE:9F mnm34-55  0000 02:02:02:02:02:03 3
ffff 1A-ETH-1
14  UP      N    00:03:BA:00:01:00:E3:72 mnm34-60  0000 02:02:02:02:02:02 3
ffff 1A-ETH-2
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, enablevnic command
- *Gateway Reference*, showvnics command
- “VNIC Allocation on the Gateway” on page 191

- “Create VLANs” on page 192
- “Delete VLANs” on page 194
- “Create VNICs” on page 195
- “Disable VNICs” on page 200
- “Delete VNICs” on page 201
- “Display the VNICs” on page 176

▼ Disable VNICs

If you believe that there is a problem with a VNIC, you can disable it.

1. Identify the VNIC you want to disable.

See “Display the VNICs” on page 176.

2. Disable the VNIC on the management controller.

```
FabMan@gateway_name->disablevnic connector vnic_ID
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH).
- *vnic_ID* is the VNIC identifier.

For example, to disable VNIC 15 on connector 1A-ETH-1, type.

```
FabMan@gateway_name->disablevnic 1A-ETH-1 15
FabMan@gateway_name->
```

3. Repeat from Step 1 for all VNICs to be disabled.

4. Verify that the VNICs have been disabled.

```
FabMan@gateway_name->showvnics
```

ID	STATE	FLG	IOA_GUID	NODE	IID	MAC	VLN	PKEY	GW
16	UP		N 00:21:28:00:01:3E:CE:A0	mnm34-55	0000	02:02:02:02:02:04	3		
ffff	1A-ETH-4								
13	UP		N 00:03:BA:00:01:00:E3:71	mnm34-60	0000	02:02:02:02:02:01	3		
ffff	1A-ETH-3								

```
15 DISABLED N 00:21:28:00:01:3E:CE:9F mnm34-55 0000 02:02:02:02:02:88 3
ffff 1A-ETH-1
14 UP N 00:03:BA:00:01:00:E3:72 mnm34-60 0000 02:02:02:02:02:02 3
ffff 1A-ETH-2
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, disablevnic command
- *Gateway Reference*, showvnics command
- “VNIC Allocation on the Gateway” on page 191
- “Create VLANs” on page 192
- “Delete VLANs” on page 194
- “Create VNICS” on page 195
- “Enable VNICS” on page 199
- “Delete VNICS” on page 201
- “Display the VNICS” on page 176

▼ Delete VNICS

Before you can delete a VLAN, you must delete the VNIC associated with that VLAN and connector.

1. Identify the VNIC you want to delete.

See “Display the VNICS” on page 176.

2. Delete the VNIC on the management controller.

```
FabMan@gateway_name->deletevnic connector vnic_ID
```

where:

- *connector* is the name of the connector (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH).
- *vnic_ID* is the VNIC identifier.

For example, to delete VNIC 15 on connector 1A-ETH-1, type.

```
FabMan@gateway_name->deletevnic 1A-ETH-1 15
FabMan@gateway_name->
```

3. Repeat from Step 1 for all VNICS to be deleted.

4. Verify that the VNICs have been deleted.

```
FabMan@gateway_name->showvnics
ID STATE   FLG IOA_GUID          NODE      IID  MAC              VLN PKEY  GW
-----
----  -
 16 UP      N 00:21:28:00:01:3E:CE:A0 mnm34-55  0000 02:02:02:02:02:04 3
ffff 1A-ETH-4
 13 UP      N 00:03:BA:00:01:00:E3:71 mnm34-60  0000 02:02:02:02:02:01 3
ffff 1A-ETH-3
 14 UP      N 00:03:BA:00:01:00:E3:72 mnm34-60  0000 02:02:02:02:02:02 3
ffff 1A-ETH-2
FabMan@gateway_name->
```

5. If you deleted the VNICs in order to delete VLANs, delete the VLANs.

See [“Delete VLANs” on page 194](#).

Related Information

- [Gateway Reference](#), deletevnic command
- [Gateway Reference](#), showvnics command
- [“VNIC Allocation on the Gateway” on page 191](#)
- [“Create VLANs” on page 192](#)
- [“Delete VLANs” on page 194](#)
- [“Create VNICs” on page 195](#)
- [“Enable VNICs” on page 199](#)
- [“Disable VNICs” on page 200](#)
- [“Display the VNICs” on page 176](#)

Controlling Gateway Ports and Parameters

These topics describe how you can control the ports and parameters of the Sun Network QDR InfiniBand Gateway Switch from Oracle:

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

- “Enable a Gateway Ethernet Port” on page 203
- “Disable a Gateway Ethernet Port” on page 204
- “Set the Gateway Ethernet Port Parameters” on page 205
- “Set the Gateway Service Level” on page 206
- “Set the Gateway Instance Number” on page 207
- “Set the Gateway System Name” on page 208

Related Information

- “Monitoring Gateway Resources” on page 172
- “Controlling LAGs” on page 183
- “Controlling VLANs and VNICs” on page 190
- “Controlling the Chassis” on page 28
- “Controlling the I4 Switch Chip” on page 40
- “Controlling the InfiniBand Fabric” on page 97
- “Controlling the Subnet Manager” on page 56
- “Installing Gateway Supportive Software (Linux)” on page 137
- “Creating VNICs Under Gateway Manual Mode (Linux)” on page 142
- “Creating VNICs Under Host Manual Mode (Linux)” on page 148
- “Creating Virtual IO Adapters (Oracle Solaris)” on page 160

▼ Enable a Gateway Ethernet Port

- On the management controller, type.

```
FabMan@gateway_name->enablegwport connector [-discoverpkey p_key]
```

where:

- *connector* is the name of the connector hosting the port (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH).
- *p_key* is the discovery partition key (1 to 7fff or default).

Note – The `-discoverypkey` option ensures that only host ports sharing the discovery partition are able to use this gateway port.

For example, to enable the port associated with connector 0A-ETH-3, type.

```
FabMan@gateway_name->enablegwport 0A-ETH-3
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `enablegwport` command
- “Disable a Gateway Ethernet Port” on page 204
- “Set the Gateway Ethernet Port Parameters” on page 205
- “Set the Gateway Service Level” on page 206
- “Set the Gateway Instance Number” on page 207
- “Set the Gateway System Name” on page 208

▼ Disable a Gateway Ethernet Port

You can disable an enabled gateway Ethernet port with the `disablegwport` command.

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

```
FabMan@gateway_name->disablegwport connector
```

where *connector* is the name of the connector hosting the port (0A-ETH-1 to 0A-ETH-4, 0A-ETH, 1A-ETH-1 to 1A-ETH-4, and 1A-ETH).

For example, to disable the port associated with connector 0A-ETH-3, type.

```
FabMan@gateway_name->disablegwport 0A-ETH-3
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `disablegwport` command
- “Enable a Gateway Ethernet Port” on page 203
- “Set the Gateway Ethernet Port Parameters” on page 205
- “Set the Gateway Service Level” on page 206
- “Set the Gateway Instance Number” on page 207

- “Set the Gateway System Name” on page 208

▼ Set the Gateway Ethernet Port Parameters

Use the `setgwethport` command to configure the gateway Ethernet ports.

- On the management controller, type.

```
FabMan@gateway_name->setgwethport connector [-mtu mtu] [-txpause pause] [-rxpause pause]
```

where:

- *connector* is the name of the connector (0A-ETH or 1A-ETH).
- *mtu* is the value of the MTU.
- *pause* is the priority flow control pause in quanta (0 to 255, Global, or None)

For example, to set the MTU of connector 1A-ETH to 9600, type.

```
FabMan@gateway_name->setgwethport 1A-ETH -mtu 9600
Port status for connector 1A-ETH-1:
Adminstate.....Enabled
State.....Up
Link state.....Up
Protocol.....Ethernet
Link Mode.....XFI
Speed.....10Gb/s
MTU.....9600
Tx pause.....Global
Rx pause.....Global
Port status for connector 1A-ETH-2:
.
.
.
Port status for connector 1A-ETH-4:
Adminstate.....Enabled
State.....Up
Link state.....Up
Protocol.....Ethernet
Link Mode.....XFI
Speed.....10Gb/s
MTU.....9600
Tx pause.....Global
Rx pause.....Global
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `setgwethport` command
- “Enable a Gateway Ethernet Port” on page 203
- “Disable a Gateway Ethernet Port” on page 204
- “Set the Gateway Service Level” on page 206
- “Set the Gateway Instance Number” on page 207
- “Set the Gateway System Name” on page 208
- “Display Gateway Ethernet Port Information” on page 178
- “Display Gateway Port Information” on page 179

▼ Set the Gateway Service Level

You can set the Ethernet data traffic and control traffic service levels with the `setgwsl` command.

Note – When enabling EoIB, consider that there are conditions from partitioning. See “[smpartition Command Guidelines](#)” on page 107.

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

```
FabMan@gateway_name->setgwsl eoib ctrl level
```

where *level* is the service level.

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

For example, to set the data traffic service level to 2, type.

```
FabMan@gateway_name->setgwsl eoib 2
Stopping Bridge Manager..          [ OK ]
Starting Bridge Manager.           [ OK ]
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, `setgwsl` command
- “Enable a Gateway Ethernet Port” on page 203
- “Disable a Gateway Ethernet Port” on page 204

- “Set the Gateway Ethernet Port Parameters” on page 205
- “Set the Gateway Instance Number” on page 207
- “Set the Gateway System Name” on page 208
- “Display Information About the Gateway” on page 182

▼ Set the Gateway Instance Number

By default, gateway instance numbers are set according to the 6 lowest bits of the gateway IP address, and must be unique for VNICs to be operational. When two or more gateways share the same InfiniBand fabric, they might have duplicate gateway instance numbers. In this situation, you must assign unique instance numbers to all of the gateways in the fabric.

1. On each gateway’s management controller, display the gateway’s instance number.

```
FabMan@gateway_name->setgwinstance list
```

2. Determine which gateway will receive a new instance number and what that number will be.

The gateway instance is a number from 0 to 63.

3. On that gateway’s management controller, type.

```
FabMan@gateway_name->setgwinstance instance
```

where *instance* is a number from 0 to 63. For example, to set the instance to 3:

```
FabMan@gateway_name->setgwinstance 3
Stopping Bridge Manager..                [ OK ]
Starting Bridge Manager..                 [ OK ]
FabMan@gateway_name->
```

Related Information

- *Gateway Reference*, setgwinstance command
- “Enable a Gateway Ethernet Port” on page 203
- “Disable a Gateway Ethernet Port” on page 204
- “Set the Gateway Ethernet Port Parameters” on page 205
- “Set the Gateway Service Level” on page 206
- “Set the Gateway System Name” on page 208

- [“Display the Gateway Instance Number” on page 181](#)

▼ Set the Gateway System Name

The BridgeX manager within the management controller recognizes a system name that is independent of the management controller’s host name. If no system name is configured, the host name is used. The system name is a maximum of 19 characters.

1. On the management controller, type.

```
FabMan@gateway_name->setgwssystemname systemname
```

where *systemname* is the BridgeX manager’s system name. For example:

```
FabMan@gateway_name->setgwssystemname brooklyn
Stopping Bridge Manager..                [ OK ]
Starting Bridge Manager.                  [ OK ]
FabMan@gateway_name->setgwssystemname --list
BXM system name set to --myname
FabMan@gateway_name->
```

2. Use the `--list` option of the `setgwssystemname` command to verify the name configuration.

```
FabMan@gateway_name->setgwssystemname --list
BXM system name set to brooklyn
FabMan@gateway_name->
```

Related Information

- [Gateway Reference, setgwssystemname command](#)
- [“Enable a Gateway Ethernet Port” on page 203](#)
- [“Disable a Gateway Ethernet Port” on page 204](#)
- [“Set the Gateway Ethernet Port Parameters” on page 205](#)
- [“Set the Gateway Service Level” on page 206](#)
- [“Set the Gateway Instance Number” on page 207](#)
- [“Display Information About the Gateway” on page 182](#)

Index

A

- acquiring
 - BXOFED software, 138
 - ConnectX-2 firmware, 141
- adding
 - autodisabled connectors, 44
 - connectors to LAGs, 187
 - element to configuration, 129
 - fabric elements to configuration, 122
 - ports to partition configuration, 112
 - secret M_Keys, 69
 - smnodes to list, 133
- addlagport command, 187
- administering
 - chassis, 19
 - gateway resources, 137
 - hosts, 137
 - I4 switch chip, 31
 - InfiniBand fabric, 83
 - Subnet Manager, 49
- administrative command overview, 13
- ASR, 28
- autodisable command, 36, 44, 46
- autodisabled connectors
 - adding, 44
 - deleting, 46
 - displaying configured, 36
 - overview, 43
 - reenabling, 45
- autodisabling ports and connectors, 43
- automatic service request, 28

B

- BXOFED
 - acquiring, 138
 - installing, 137, 139

C

- central configuration file, 150
- changing password
 - root user, 29
- chassis
 - administering, 19
 - controlling, 28
 - FRU ID, 26
 - monitoring, 19
- checkboot command, 34
- checking
 - board
 - temperature, 22
 - voltage, 21
 - chassis status, 23
 - environmental data, 23
 - fan
 - speed, 22
 - status, 22
 - flash drive status, 24
 - link status, 35
 - port counters, 38
 - port status, 37
 - power supply, 20
 - switch chip boot status, 34
- checkpower command, 20
- checkvoltages command, 21
- clear
 - data counters, 102
 - error counters, 102
 - management key, 60
- command
 - addlagport, 187
 - autodisable, 36, 44, 46
 - checkboot, 34
 - checkpower, 20
 - checkvoltages, 21

- create_ipoib, 96, 135
- createfabric, 120
 - add-element, 122
 - delete-element, 122
 - list active, 96
 - modify-element, 123
 - start, 120
- createlag, 185, 186
- createvlan, 192
- createvnic, 195
- dcSPORT, 32
- delete_ipoib, 136
- deletelag, 188
- deletevlan, 194
- deletevnic, 201
- dellagport, 187
- disablegwport, 204
- disablelagmode, 189
- disableSM, 63
- disableswitchport, 40
- disablevnic, 200
- dladm, 163, 167, 169
- enablegwport, 203
- enablelagmode, 184
- enableSM, 62
- enableswitchport, 41, 45
- enablevnic, 199
- env_test, 23
- fdconfig, 126
 - abort, 127
 - complete-fabric-config, 127
 - define-element, 129
 - define-fabric-name, 127
 - list-current-fabric-config, 94
 - redefine-element, 131
 - remove-element, 129
 - start-fabric-config empty, 127
- generatetopology, 99
- getfanspeed, 22
- getmaster, 51
- getportcounters, 38
- getportstatus, 37
- ibdiagnet, 98, 100, 102
- ibhosts, 85
- ibnetdiscover, 86
- ibportstate, 88, 102, 103, 104
- ibswitches, 84
- ibtracert, 87
- ipadm, 167
- listlinkup, 35
- localmkeypersistence, 53, 76
- matchtopology, 99
- modifyvnic, 197
- perfquery, 89
- reboot, 29
- setgwethport, 178, 205
- setgwinstance, 181, 207
- setgws1, 206
- setgwsystemname, 208
- setsmmkey, 59, 60
- setsmpriority, 51, 56, 58
- setsmrouting, 61
- setsubnetprefix, 57
- showdisk, 24
- showfree, 25
- showfruinfo, 26
- showgwconfig, 182
- showgwports, 179
- showioadapters, 177
- showlag, 173
- showpsufu, 26
- showsmlog, 55
- showtemps, 22
- showtopology, 86
- showunhealthy, 20
- showvlan, 174
- showvnics, 176
- smconfigtest, 54
- sminfo, 50
- smnodes
 - add, 133
 - delete, 134
 - list, 95
- smpartition, 106
 - add, 112
 - create, 110
 - delete, 116
 - list, 93
 - modify, 114
 - remove, 112
 - start, 110
- smpquery, 90, 91
- smsubnetprotection, 65, 67
 - add, 69
 - delete, 73
 - disablesecretmkey, 77
 - enablesecretmkey, 77
 - list active, 97

- list modified, 52
- set-current, 72
- setlocalsecretmkey, 75
- setreplicationpassword, 76
- version, 27
- configuration overview, 13
- configuring
 - Fabric Director node list, 125
 - IP address on data links, 167
 - IP interfaces on data links, 167
 - local secret M_Key, 75
 - secret M_Key, 67
- connector
 - autodisabled
 - adding, 44
 - deleting, 46
 - reenabling, 45
 - autodisabling, 43
 - displaying autodisable configured, 36
- controlling
 - chassis, 28
 - gateway parameters, 202
 - gateway ports, 202
 - I4 switch chip, 40
 - InfiniBand fabric, 97
 - LAGs, 183
 - nodes, 101
 - Subnet Manager, 56
 - VLANs, 190
 - VNICs, 190
- create_ipoib command, 96, 135
- createfabric
 - add-element command, 122
 - delete-element command, 122
 - list active command, 96
 - modify-element command, 123
 - start command, 120
- createfabric command, 120
- createlag command, 185, 186
- createvlan command, 192
- createvnic command, 195
- creating
 - central configuration file, 154
 - fabric configuration, 127
 - fabric element configuration, 120
 - Global Gateway, 186
 - IPoIB interface, 135

- LAGs, 185
- Oracle Solaris VNICs, 169
- partition configuration, 110
- secret M_Key configuration, 67
- VIOAs, 160, 163
- VLANs, 192
- VNICs, 142, 148, 195
- VNIC-specific configuration files, 156

D

- data counters
 - clearing, 102
- data links, 161
- dcsport command, 32
- delete_ipoib command, 136
- deletelag command, 188
- deletevlan command, 194
- deletevnic command, 201
- deleting
 - autodisabled connectors, 46
 - connectors from LAGs, 187
 - fabric elements from configuration, 122
 - IPoIB interface, 136
 - LAGs, 188
 - partition configuration, 116
 - secret M_Keys, 73
 - smnodes from list, 134
 - VLANs, 194
 - VNICs, 201
- dellagport command, 187
- disablegwport command, 204
- disablelagmode command, 189
- disablesm command, 63
- disableswitchport command, 40
- disablevnic command, 200
- disabling
 - gateway Ethernet port, 204
 - LAG mode, 189
 - port, 104
 - secret M_Key
 - functionality, 77
 - persistence, 76
 - Subnet Manager, 63
 - switch chip port, 40
 - VNICs, 200
- displaying
 - autodisable configured connectors, 36

- disabled links, 35
- fabric configuration, 94
- fabric configuration information, 93
- fabric elements, 96
- file system space, 25
- flash drive information, 24
- free memory, 25
- FRU ID
 - chassis, 26
 - power supply, 26
- gateway
 - Ethernet information, 178
 - firmware, 27
 - health, 20
 - information, 182
 - IO adapters, 177
 - port information, 179
 - service level, 179, 182
 - system name, 182
- instance number, 181
- IPoIB interface, 96
- LAGs, 173
- M_Key persistence, 53
- node
 - counters, 89
 - information, 84, 90
 - link status, 88
- partition configuration, 93
- port information, 91
- route, 87
- secret M_Key
 - configuration, 52
 - persistence, 53
- secret M_Keys, 97
- smnodes, 95
- Subnet Manager
 - configuration integrity, 54
 - controlled handover state, 51
 - location history, 51
 - log, 55
 - management key, 51
 - prefix, 51
 - priority, 51
 - routing algorithm, 51
 - status, 50
- topology
 - detailed, 86
 - simple, 86
- VLANs, 174

- VNICs, 176
- dladm command, 163, 167, 169

E

- element
 - adding, 129
 - modifying, 131
 - removing, 129
- enablegwport command, 203
- enablelagmode command, 184
- enablesm command, 62
- enableswitchport command, 41, 45
- enablevnic command, 199
- enabling
 - gateway Ethernet port, 203
 - LAG mode, 184
 - secret M_Key
 - functionality, 77
 - persistence, 76
 - Subnet Manager, 62
 - switch chip port, 41
 - VNICs, 199
- env_test command, 23
- EoIB, 161
 - installing driver, 162
 - verifying installation, 162
- error counters
 - clearing, 102
 - displaying, 89
- Ethernet
 - displaying information, 178
 - setting port, 205
- example route through gateway, 9

F

- fabric
 - diagnostics, 98
 - displaying configuration information, 93
 - element, 125
- Fabric Director, 125
- fabric elements, 119
 - adding elements to configuration, 122
 - configurations, 120
 - creating configuration, 120
 - deleting elements from configuration, 122
 - displaying, 96

- modifying elements of configuration, 123
- fan
 - checking
 - speed, 22
 - status, 22
- fdconfig
 - abort command, 127
 - complete-fabric-config command, 127
 - define-element command, 129
 - define-fabric-name command, 127
 - list-current-fabric-config command, 94
 - redefine-element command, 131
 - remove-element command, 129
 - start-fabric-config empty command, 127
- fdconfig command, 126
- firmware
 - downgrades and partitions, 117
 - gateway, 27
 - secret M_Key guidelines, 80
- FRU ID
 - chassis, 26
 - power supply, 26
- full membership, 109

G

- gateway
 - controlling
 - parameters, 202
 - ports, 202
 - displaying information, 179, 182
 - Ethernet port
 - disabling, 204
 - enabling, 203
 - manual mode, 142
 - configuration parameters, 144
 - overview, 143
 - VLAN associations, 145
 - VNIC conditions, 147
 - VNIC creation, 146
 - resources
 - administering, 137
 - monitoring, 172
 - troubleshooting, 1
- generatetopology command, 99
- getfanspeed command, 22

- getmaster command, 51
- getportcounters command, 38
- getportstatus command, 37
- Global Gateway, 186
- GUID
 - mapping
 - to CXP connector, 33
 - to switch chip, 33
 - nomenclature, 10
- guidelines
 - firmware
 - secret M_Key, 80
 - smpartition, 107

H

- hardware commands
 - overview, 16
- hardware problems, 1
- HCA identity, 85
- host
 - administering, 137
 - manual mode, 148
 - central configuration file, 150
 - configuration parameters, 152
 - overview, 149
 - VNIC creation, 158
 - VNIC-specific configuration file, 151

I

- I4 switch chip
 - administering, 31
 - controlling, 40
 - monitoring, 31
- ibdiagnet command, 98, 100, 102
- ibhosts command, 85
- ibnetdiscover command, 86
- ibportstate command, 88, 102, 103, 104
- ibswitches command, 84
- ibtracert command, 87
- identifying
 - HCAs, 85
 - switches, 84
- InfiniBand
 - commands, 17
 - creating configuration, 127
 - displaying configuration, 94

- fabric
 - administering, 83
 - controlling, 97
 - monitoring, 83
 - problems, 3
- installing
 - BXOFED software, 137, 139
 - EoIB driver, 162
- instance number
 - displaying, 181
 - setting, 207
- investigating nodes, 98
- ipadm command, 167
- IPoIB interface
 - creating, 135
 - deleting, 136
 - displaying, 96

L

- LAGs, 173
 - adding connectors, 187
 - controlling, 183
 - creating, 185
 - deleting, 188
 - deleting connectors, 187
 - disabling, 189
 - displaying, 173
 - enabling, 184
- limited membership, 109
- link
 - status, 35
 - status of node, 88
- listlinkup command, 35
- localmkeypersistence command, 53, 76

M

- management controller
 - restarting, 29
- mapping
 - GUID
 - to CXP connector, 33
 - to switch chip, 33
 - LED to switch chip port, 8
 - QSFP to switch chip port, 8, 32
 - switch chip port
 - to LED, 8
 - to QSFP, 8, 32

- matchtopology command, 99
- membership
 - full, 109
 - limited, 109
- mlx4_vnic_conf daemon, 157
- mode
 - gateway manual, 142
 - host manual, 148
- modifying
 - element, 131
 - fabric elements of configuration, 123
 - partition configuration, 114
 - port membership, 114
 - VNICs, 197
- modifyvnic command, 197
- monitoring
 - chassis, 19
 - gateway resources, 172
 - I4 switch chip, 31
 - InfiniBand fabric, 83
 - Subnet Manager, 49

N

- NET MGT port
 - troubleshooting, 6
- nodes
 - controlling, 101
 - counters, 89
 - displaying information, 84
 - information, 90
 - investigating, 98
 - link status, 88

O

- Oracle ILOM
 - command
 - overview, 16
- Oracle Solaris VNICs, 169
- overview
 - administrative commands, 13
 - autodisabled connectors, 43
 - Ethernet over InfiniBand, 161
 - gateway manual mode, 143
 - hardware commands, 16
 - host manual mode, 149
 - InfiniBand commands, 17
 - Oracle ILOM

- commands, 16
- secret M_Key, 64

P

- P_Keys, 108, 109
- partitioning, 108, 109
 - adding ports, 112
 - creating configuration, 110
 - deleting configuration, 116
 - displaying configuration, 93
 - firmware downgrades, 117
 - InfiniBand fabric, 105
 - modifying configuration, 114
 - removing ports, 112
 - starting configuration session, 110
- password, 29
- performing
 - supportive tasks, 132
- perfquery command, 89
- poor performing links, 100
- ports
 - autodisabling, 43
 - counters, 38
 - disabling, 104
 - information, 91
 - resetting, 102
 - speed, 103
 - status, 37
- power supply
 - checking
 - status, 20
 - FRU ID, 26
- preparing
 - secret M_Key, 66
- problem
 - hardware, 1
 - InfiniBand fabric, 3

Q

- QSFP connector
 - mapping
 - from GUID, 33
 - to switch chip port, 32

R

- reboot command, 29
- recovery errors, 100

- removing
 - element from configuration, 129
 - ports from partition configuration, 112
- resetting
 - port, 102
- restarting
 - management controller, 29
- route
 - QSFP to switch chip port, 8, 32
 - switch chip port to QSFP, 8, 32
 - through gateway, 7, 9
 - through the fabric, 87

S

- secret M_Key, 65
 - adding secret M_Keys, 69
 - conditions for situations, 78
 - configuring, 67
 - configuring local, 75
 - creating configuration, 67
 - deleting secret M_Keys, 73
 - disabling
 - functionality, 77
 - persistence, 76
 - displaying, 97
 - configuration, 52
 - persistence, 53
 - enabling
 - functionality, 77
 - persistence, 76
 - firmware guidelines, 80
 - merging subnets, 79
 - new Subnet Manager, 78
 - overview, 64
 - preparing, 66
 - setting replication password, 76
 - setting to use, 72
- secure fabric management, 64
- service level
 - setting, 206
- setgwethport command, 178, 205
- setgwinstance command, 181, 207
- setgws1 command, 206
- setgwsystemname command, 208
- setsmmkey command, 59, 60
- setsmpriority command, 51, 56, 58
- setsmrouting command, 61

- setsubnetprefix command, 57
- setting
 - Ethernet port, 205
 - instance number, 207
 - port speed, 103
 - replication password, 76
 - secret M_Key to use, 72
 - service level, 206
 - Subnet Manager
 - controlled handover, 58
 - management key, 59
 - prefix, 57
 - priority, 56
 - routing algorithm, 61
 - system name, 208
- showdisk command, 24
- showfree command, 25
- showfruinfo command, 26
- showgwconfig command, 182
- showgwports command, 179
- showioadapters command, 177
- showlag command, 173
- showpsufu command, 26
- showsmlog command, 55
- showtemps command, 22
- showtopology command, 86
- showunhealthy command, 20
- showvlan command, 174
- showvnics command, 176
- smconfigtest command, 54
- sminfo command, 50
- smnodes
 - adding to list, 133
 - deleting from list, 134
 - displaying, 95
- smnodes
 - add command, 133
 - delete command, 134
 - list command, 95
- smpartition
 - add command, 112
 - create command, 110
 - delete command, 116
 - list command, 93
 - modify command, 114
 - remove command, 112
 - start command, 110
- smpartition command, 106
 - guidelines, 107
- smpquery command, 90, 91
- smsubnetprotection
 - add command, 69
 - delete command, 73
 - disablesecretmkey command, 77
 - enablesecretmkey command, 77
 - list active command, 52, 97
 - set-current command, 72
 - setlocalsecretmkey command, 75
 - setreplicationpassword command, 76
 - start command, 67
- smsubnetprotection command, 65
- status
 - chassis, 23
 - fan, 22
 - flash drive, 24
 - power supply, 20
 - Subnet Manager, 50
- Subnet Manager
 - added to secret M_Key fabric, 78
 - clearing management key, 60
 - configuration integrity, 54
 - controlled handover state, 51
 - controlling, 56
 - disabling, 63
 - displaying
 - configuration integrity, 54
 - location history, 51
 - log, 55
 - management key, 51
 - prefix, 51
 - priority, 51
 - routing algorithm, 51
 - status, 50
 - enabling, 62
 - management key, 51, 59, 60
 - monitoring, 49
 - setting
 - controlled handover, 58
 - management key, 59
 - prefix, 57
 - priority, 56
 - routing algorithm, 61
- supportive tasks, 132
- switch

- GUIDs, 10
- identity, 84
- switch chip
 - checking boot status, 34
 - port
 - disabling, 40
 - enabling, 41
- symbol errors, 100
- system name
 - setting, 208

VNIC-specific configuration file, 151

T

- topology
 - change, 99
 - displaying
 - detailed, 86
 - simple, 86
- troubleshooting
 - gateway, 1
 - network management, 6

V

- verifying
 - EoIB driver installed, 162
- version command, 27
- VIOAs, 161
 - creating, 163
- VLANs
 - associations, 145
 - controlling, 190
 - creating, 192
 - deleting, 194
 - displaying, 174
- VNICs
 - allocation on gateway, 191
 - conditions for gateway manual mode, 147
 - configuration parameters, 144, 152
 - controlling, 190
 - creating, 142, 148, 195
 - gateway manual mode, 146
 - host manual mode, 158
 - deleting, 201
 - disabling, 200
 - displaying, 176
 - enabling, 199
 - modifying, 197

