Sun Network QDR InfiniBand Gateway Switch

Product Notes for Firmware Version 2.0

Part No.: E26705-05
February 2013, Revision A

Please
Recycle

Adobe PostScript

Contents

Using This Documentation

This document provide last-minute information regarding the installation, administration, and service of the Sun Network QDR InfiniBand Gateway Switch from Oracle.

Related Documentation

Documentation	Links
Sun Network QDR InfiniBand Gateway Switch Firmware Version 2.0	http://docs.oracle.com/cd/E26699_01
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 about 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.

Sun InfiniBand QDR Gateway Switch Product Notes

These product notes provide last-minute, late-breaking information regarding the gateway. These notes pertain to the release of the 2.0.8-1 firmware for the gateway.

Known Problems

The following table describes known problems with the gateway.

Bug	Description	Workaround
15560533	Setting an alert rule to `ipmipet` sometimes does not work.	There is no impact to the InfiniBand fabric. **Workaround**: After setting the alert rule to `ipmipet`, set the level to `disable`, and then set the level to the desired value.
15693303	BridgeX reset maximizes symbol error counter.	The `setgwinstance`, `setgwsl`, and `setgwsystemname` commands reset the BridgeX chip, causing symbol errors to maximize. **Workaround**: Follow the procedure described in "Symbol Errors on Internal Ports" on page 7.
15696509	IP addresses within the Subnet Manager nodes file does not match those assigned by DHCP.	By default, the DHCP server is not assured to assign a consistent IP address to a Subnet Manager node upon reboot. **Workaround**: Manually configure the DHCP server to assign a specific and unique IP address to each Subnet Manager node.
15703751	No way to set Fabric Monitor console timeout.	The Fabric Monitor will timeout after 15 minutes of idle time. **Workaround**: You must login to the web interface and launch the Fabric Monitor again.

Bug	Description	Workaround
15723028	BridgeX ports indicate random errors in Fabric Monitor.	Powering on the gateway randomly displays these errors at BridgeX ports BX_0 and BX_1: • Peer RcvErrors • Peer VL15Dropped • XmtDiscards There is no workaround at this time.
15747266	Fabric Monitor does not support Internet Explorer.	Though the Oracle ILOM web interface supports Internet Explorer, the Fabric Monitor does not. **Workaround**: Use Mozilla Firefox 3.0 or 3.6.
15747583	Oracle ILOM web interface not compatible with Mozilla Firefox 6.0.	Using the Mozilla Firefox 6.0 browser, only navigational control is possible with the Oracle ILOM web interface. No informational output is displayed. **Workaround**: Use Mozilla Firefox 3.0 or 3.6.
15747938	`smpartition start` command gives unclear message if the Subnet Manager is disabled.	If the Subnet Manager is disabled before starting a partition configuration session using the `smpartition start` command, confusing messages are displayed upon issuing the command. **Workaround**: Ensure that the Subnet Manager is enabled before starting a partition configuration session.
15747967	`createvnic` command should handle GUIDs with `0x` prefix correctly.	GUIDs can be displayed as `0021280001a1623a`, `00:21:28:00:01:a1:62:3a`, or `0x0021280001a1623a`. The latter version is not recognized by the `createvnic` command. **Workaround**: Truncate the `0x` prefix of the GUID for the `createvnic` command.
15750278	`createvlan` command does not support both hex and decimal P_Key value input.	The construct of the P_Key and its membership bit requires the value to be in hexadecimal. Providing and displaying the P_Key in decimal would complicate identifying the state of the membership bit. There is no workaround.
15750534	`--clear` option of `setgwethport` command does not restore default parameters.	From the command usage, the `--clear` option of the `setgwethport` command should restore to default parameters. This situation does not happen. **Workaround**: Use the `-clear` option.
15755116	`checkguidfilesftree` is displayed with the `help all` command, but is not supported.	Using the `help all` command within the Oracle ILOM `/SYS/Fabric_Mgmt` target displays the `checkguidfilesftree` command. This command has been deprecated and is not available.

Bug	Description	Workaround
15755133	`fdconfig` command is unavailable from the `/SYS/Fabric_Mgmt` target.	The `fdconfig` command is required for configuring the FabricMIB and should be available from the Oracle ILOM `/SYS/Fabric_Mgmt` target. The command is not available. **Workaround**: Log in as the `root` user, where the `fdconfig` command is available.
15755727	Need tool to display InfiniBand topology with physical InfiniBand entities.	At present, no command provides the relationship of InfiniBand fabric GUIDs, LIDs, and IB ports to each other in a simplified manner. **Workaround**: Use the `ibnetdiscover` command.
15756465	Usage of `deletevnic` command is wrong.	The usage output of the `deletevnic` command indicates that connectors 0A-ETH and 1A-ETH are supported. They are not.
15758329	Stopping or starting the `envd` daemon initiates a segmentation fault for the FabricMIB.	Stopping or starting the environment daemon on a gateway running as a fabric manager or monitored by a fabric manager changes the FabricMIB configuration due to the new status of that fabric element. Should an `snmpwalk` of the FabricMIB be issued at the same time the FabricMIB library reloads the updated configuration, internal file pointers might be mishandled, initiating a segmentation fault. There is no workaround at this time.
15758406	Fabric Monitor might fail to load.	During a firmware upgrade or downgrade, certain symbolic links might not be created or become lost. Consequently, the Fabric Monitor will fail to load when requested. **Workaround**: Reperform the upgrade or downgrade, using the `-force` option.
15761792	BridgeX manager dies with a `base port id collision detected` error message.	When two or more gateways within the InfiniBand fabric have adjacent gateway instance numbers (for example 16 and 17), and the gateway with the lower instance number (16) has more than one LAG enabled, there is a port ID collision between the two gateways. **Workaround**: Configure all gateways to have even instance numbers (2, 4, 6, ...) with the `setgwinstance` command.
15819653	Gateway hard hangs after reboot.	On rare occurrence when rebooting the gateway, the gateway might experience a hard hang. The last entry in the system log is: `kernel: mlx_bx_core 0000:05:00.0: go bit not cleared` **Workaround**: Power cycle the gateway by removing both power cords from the power supplies, and then reattaching them.

General Information and Issues

Reserved VLAN Identifiers

When creating VLANs, use VLAN identifiers 2 through 4094. According to the IEEE 802.1Q specification, VLAN ID 0 is used for the priority tag, and VLAN ID 1 is usually reserved for a switch or bridge management VLAN.

Jumbo Frames

Consider the following issues for Jumbo Frames:

- True Jumbo Frames are not supported.
- The maximum MTU is limited by the MTU of the InfiniBand fabric.
- The MTU of the HCA is 2048 by default, but can be increased up to 4096.
- The default MTU of the I4 switch chip and BridgeX chips is 4096.

The gateway supports Jumbo Frames up to the InfiniBand limit of 4096 bytes.

You must configure the entire InfiniBand fabric to use a 4096 MTU. This configuration includes configuring the Subnet Manager, InfiniBand devices, and the ConnectX HCAs. You can configure a ConnectX HCA by activating the `set_4k_mtu` parameter of the HCA's `mlx4_core` module.

Subnet Manager

Access to the VNICs through the gateway requires an active Subnet Manager for the InfiniBand fabric. The management controller within the gateway is already configured with a Subnet Manager. You can enable the Subnet Manager with the `enablesm` command.

Snapshot Dataset Information

The `normal`, `fruid`, and `full` datasets of the snapshot utility are currently equivalent and contain the same data in the snapshot.

Software Information and Issues

Main Board, Management Controller, and Chassis Serial Numbers

The gateway documentation describes how to retrieve the chassis serial number using the `showfruinfo` command or the `/SYS/MB` Oracle ILOM target. These methods actually display the serial number of the main board and the management controller respectively, and not the gateway chassis. The gateway chassis serial number is provided on the pull-out tab on the left side front of the gateway chassis, adjacent to power supply 0.

High Availability in Partitions

To allow communication fail-over between HCAs belonging to the same operating system instance, the HCA ports must be members of the same partition and have identical membership type (full or limited).

Having both full and limited port memberships within a partition for the same operating system instance creates a configuration instability that might cause subtle communication problems.

Email Alert Rules

You must specify the value of the `email_custom_sender` property of an email alert rule, because the alert does not use the `custom_sender` property of the `/SP/clients/smtp` target.

Time Zone Support

The following time zones are only supported in firmware versions 1.3.4, 2.0.7-2, and 2.0.8-1. If you upgrade or downgrade to a firmware version other than 1.3.4, 2.0.7-2, or 2.0.8-1 you must set an alternative time zone.

- America/Argentina/Salta
- America/Argentina/San_Luis

- America/Bahia_Banderas
- America/Kralendijk
- America/Lower_Princes
- America/Matamoros
- America/Metlakatla
- America/North_Dakota/Beulah
- America/Ojinaga
- America/Santa_Isabel
- America/Santarem
- America/Sitka
- Antarctica/Macquarie
- Asia/Ho_Chi_Minh
- Asia/Kathmandu
- Asia/Kolkata
- Asia/Novokuznetsk
- Pacific/Chuuk
- Pacific/Pohnpei

Firmware Update Considerations

If you are going to downgrade the firmware to a version earlier than 2.0, you must remove user partitions and depopulate the Subnet Manager nodes list. Refer to *Gateway Administration*, removing partitions for firmware downgrade.

If you are going to downgrade from firmware 2.0.*x* to 1.3.4_1 or earlier, you might see these type of messages in the `/var/log/message` file after the downgrade:

```
lda: Unknown config parameter: ErrLogCount=100; .

lda: Unknown config parameter: ErrLogTimeInterval=100; .
```

The `ErrLogCount` and `ErrLogTimeInterval` configuration parameters introduced in firmware 2.0.*x* are unknown to firmware version 1.3.4-1 and earlier LDAs. The LDA logs these messages and ignores them from there on. The messages appear once per LDA startup and are harmless.

If you later upgrade from firmware 1.3.4-1 to 2.0.*x* and the `ErrLogCount` and `ErrLogTimeInterval` configuration parameters are absent, the 2.0.*x* LDA uses the compiled default values. Consequently, no LDA messages regarding these configuration parameters are recorded.

Host Software and Firmware

To configure and use VNICs from an InfiniBand host, the host must have the following software and firmware installed:

- BXOFED software. See "Acquire the BXOFED Software" on page 16.
- ConnectX-2 firmware version 2.7.000 or higher. See "Acquire the ConnectX-2 Firmware" on page 18.

Both of these software packages are available from Oracle.

Symbol Errors on Internal Ports

When starting the gateway or when restarting the Bridge Manager (`bxm`), there might be 65535 symbol errors on the internal links between the I4 switch chip and the BridgeX chips. These errors are identified in hardware commands such as `getportcounters` or in InfiniBand commands such as `ibdiagnet`.

Workaround: Manually reset the error counters using the following commands:

```
# getportcounters 1 -R; getportcounters 2 -R; getportcounters 3 -R;
getportcounters 4 -R
```

Commands Available to Restricted Linux Shells

Typing `help all` within either the `/SYS/Switch_Diag`, `/SYS/Gateway_Mgmt`, or `/SYS/Fabric_Mgmt` restricted Linux shells lists the commands available to those shells.

In the output, the `checkguidfilesftree` command is listed, however, the command is not supported.

Conversely, the `fdconfig` command is not listed in the output, but is supported by the `/SYS/Fabric_Mgmt` restricted Linux shell.

Hardware Information and Issues

Ethernet Support

The gateway supports 10 Gb/sec Ethernet networks.

The following table lists supported hardware to construct one splitter cable for 10 GbE networks.

Description	Part Number	Quantity Needed
Optical splitter cable	X2127A	1
10 GbE QSFP transceiver	X2124A	1
10 GbE SFP+ transceiver	X2129A	4

Number of LAGs

The gateway supports a maximum of 16 LAGs.

Unusable Ports

The ports covered by the Do Not Remove plastic tab are unusable at the time of this document. Do not remove the tab.

Documentation Information and Issues

`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 conjuction with the `ALL_CAS` port grouping in the `smpartition add` command.
- To remove the `ipoib` flag from a partition, use the `smpartition modify -flag` command without any additional arguments. The command removes the `ipoib` flag successfully, but erroneously issues this warning:

  ```
  Missing port parameter

  Could not modify partition.
  ```

- If you use EoIB and VNIC functionality in a nondefault 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
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 nondefault partition.

VNIC Configuration on Linux Hosts in Gateway Manual Mode

If you are creating VNICs in Gateway Manual Mode, you must follow 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 `eth`*XX*, 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*:

`eth`*gateway_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-eth`*gateway_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 `eth`*gateway_instance_VNIC_instance* as described previously.
- *mac* is the MAC assigned to the VNIC.

Corrections to Creating Virtual I/O Adapters (Oracle Solaris)

In *Gateway Administration*, there are several errors in the procedures for creating VIOAs in Oracle Solaris. Those procedures are corrected here.

▼ 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 Gateway Administration, creating Oracle Solaris VNICs.

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 Oracle 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 14.

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 UP          N 00:21:28:00:01:A0:A5:91 solaris01 0000 a0:a5:91:95:30:9a NO
ffff 0A-ETH-2
```

```
 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 Oracle 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 14.

▼ 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
eoib2 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 `eoib2`:

```
# ipadm create-addr -T static -a 192.168.17.24/24 eoib2/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 eoib2/v4
ADDROBJ TYPE STATE ADDR
eoib2/v4 static ok 192.168.17.24/24
#
```

SNMP V3 Protocol Passwords

In sections of *Gateway Remote Management*, authentication and privacy passwords are described for SNMP service user accounts and SNMP V3 protocol commands. For the user accounts and commands, the MD5 and SHA authentication passwords are 8 to 12 characters in length, and the DES privacy password must be exactly 8 characters long.

Features and Functionality Documented

The features and functionality described in the gateway documentation has been updated to reflect the firmware starting with version 2.0. Upgrading your gateway firmware to the most current version helps increase gateway functionality. See "Upgrading the Gateway Firmware" on page 19.

▼ Acquire the BXOFED Software

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 autopropogate with the most current version.

8. **In the Release Is 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.**

▼ 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.**

Upgrading the Gateway Firmware

In *Gateway Remote Management*, firmware version numbers are provided as x.y, x.y.z, and x.y.z-w. Currently, these numbers are 2.0, 2.0.8, and 2.0.8-1 respectively. The following two procedures describe how to acquire and upgrade the firmware through the Oracle ILOM CLI.

▼ Acquire the Gateway Firmware Package (CLI)

1. **Open a web browser on a host that is on the same Ethernet network as the management controller to receive the firmware update.**

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 and Updates page is displayed.

5. **In the Patch Search window, click the Search tab.**

   The Patch Search window updates.

6. **Click the Product or Family (Advance) link.**

   The Patch Search window updates.

7. **In the Product Is drop-down menu, select Sun Network QDR Infiniband Gateway Switch.**

8. **In the Release Is drop-down menu, select Sun Network QDR Infiniband Gateway Switch** *x.y.z***.**

   Where *x.y.z* is the version number of the firmware package to be acquired. For example, 2.0.8.

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 respective patch number link.**

    For example, `16295011`. 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, `p16295011_208_Generic.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 firmware is in the `SUN_DCS_gw_`*x*`.`*y*`.z.tar.gz` file.

    The README file contains the latest information about the firmware release.

17. **Unpack the** `.gz` **file.**

```
$ gtar -zxvf SUN_DCS_gw_x.y.z.tar.gz
```

    The extracted files are displayed.

18. **Move the gateway firmware package (***filename*`.pkg`**) to a directory on a host that is accessible by Oracle ILOM.**

19. **Upgrade the gateway firmware.**

    See "Upgrade the Gateway Firmware (CLI)" on page 21.

▼ Upgrade the Gateway Firmware (CLI)

Note – If you are going to downgrade the firmware to a version earlier than 2.0, you must remove user partitions and depopulate the Subnet Manager nodes list. Refer to *Gateway Administration*, removing partitions for firmware downgrade.

---

1. **Open an SSH session as user `root` and connect to the management controller by specifying the controller's host name.**

   For example:

```
% ssh -l root gateway_name
root@gateway_name's password: password
#
```

   where *gateway_name* is the host name of the management controller. Initially, the *password* is `changeme`.

2. **If the Subnet Manager is running on the management controller, disable the Subnet Manager with the `disablesm` command.**

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

3. **Verify that there is at least 150 MB available in the `/tmp` directory.**

```
# df -h /tmp
Filesystem            Size  Used Avail Use% Mounted on
tmpfs                 250M  240K 249M 1% /tmp
#
```

   In this example, there are 249 MB available. If not enough space is available, you must delete files from the `/tmp` directory.

4. **Verify that there is at least 1 MB available in the `/config` directory.**

```
# df -h /config
Filesystem            Size  Used Avail Use% Mounted on
/dev/hda2              16M  3.6M   11M  25% /config
#
```

In this example, there are 11 MB available. If not enough space is available, you must delete files from the `/config` directory.

5. **Verify that there is at least 1 MB available in the `/var/log` directory.**

```
# df -h /var/log
Filesystem            Size  Used Avail Use% Mounted on
/dev/hda3              16M  3.6M   11M  25% /var/log
#
```

In this example, there are 11 MB available. If not enough space is available, you must delete files from the `/var/log` directory.

6. **Verify that there is at least 150 MB free memory available.**

```
# free -m
             total       used       free     shared    buffers     cached
Mem:           498        104        393          0         12         47
-/+ buffers/cache:         45        453
Swap:            0          0          0
#
```

In the `-/+ buffers/cache:` row of the `free` column, there should be at least 150 MB free memory. In this example, there are 453 MB available. If not enough memory is available, you must exit nonessential applications that are running.

7. **Start the Oracle ILOM shell.**

```
# spsh
Oracle(R) Integrated Lights Out Manager
Version ILOM 3.0 r47111
Copyright (c) 2010, Oracle and/or its affiliates. All rights reserved.
->
```

You are now in the Oracle ILOM shell.

You can use the `exit` command to return to the Linux shell.

8. **Begin the upgrade process.**

```
-> load -source URI/pkgname
```

where:

- *URI* is the uniform resource indicator for the host where the gateway firmware package is located. The FTP and HTTP protocols are supported.
- *pkgname* is the name of the firmware package in the transfer directory.

For example, using the FTP protocol:

```
-> load -source
ftp://root:changeme@123.45.67.99//tmp/sundcs_gw_repository_2.0.8_1.pkg
Downloading firmware image. This will take a few minutes.
```

Note – If you are upgrading from firmware version 2.0.X, you can use the `-force` option to disable version number checking, and force the upgrade.

The firmware is downloaded. The upgrade begins. A caution is displayed, and you are asked to commit to the upgrade.

```
NOTE: Firmware upgrade will upgrade firmware on SUN DCS gw Kontron module,
      I4 and BridgeX. Upgrade takes few minutes to complete.

      ILOM will enter a special mode to load new firmware. No other tasks
      should be performed in ILOM until the firmware upgrade is complete.

Are you sure you want to load the specified file (y/n)?
```

9. **Answer `y` to the prompt to commit to the upgrade.**

The upgrade begins.

```
Setting up environment for firmware upgrade. This will take few minutes.
Starting SUN DCS gw FW update

==========================
Performing operation: I4 A
==========================
I4 fw upgrade from 7.3.0(INI:1) to 7.4.1010(INI:1):
Upgrade started...
Upgrade completed.
INFO: I4 fw upgrade from 7.3.0(INI:1) to 7.4.1010(INI:1) succeeded

==========================
```

```
Performing operation: BX A
==========================
BX fw upgrade from 8.3.3166(INI:4) to 8.4.2812(INI:6):
Upgrade started...
Upgrade completed.
INFO: BX fw upgrade from 8.3.3166(INI:4) to 8.4.2812(INI:6) succeeded

==========================
Performing operation: BX B
==========================
BX fw upgrade from 8.3.3166(INI:4) to 8.4.2812(INI:6):
Upgrade started...
Upgrade completed.
INFO: BX fw upgrade from 8.3.3166(INI:4) to 8.4.2812(INI:6) succeeded

============================
Summary of Firmware update
============================
I4 status                 :  FW UPDATE - SUCCESS
I4 update succeeded on    :  A
I4 already up-to-date on  :  none
I4 update failed on       :  none
BX status                 :  FW UPDATE - SUCCESS
BX update succeeded on    :  A, B
BX already up-to-date on  :  none
BX update failed on       :  none

========================================
Performing operation: SUN DCS gw firmware update
========================================
SUN DCS gw Kontron module fw upgrade from 1.3.2-1 to 2.0.8-1:
Please reboot the system to enable firmware update of Kontron module. The
download of the Kontron firmware image happens during reboot.

After system reboot, Kontron FW update progress can be monitored in browser using
URL [http://GWsystem] OR at OS command line prompt by using command [telnet
GWsystem 1234] where GWsystem
is the hostname or IP address of SUN DCS GW.

Firmware update is complete.
->
```

10. Exit the Oracle ILOM CLI shell.

```
-> exit
exit
#
```

11. **Reboot the gateway to enable the new firmware.**

    Refer to *Gateway Administration*, restarting the entire gateway.

Note – The restart process takes between 4 to 5 minutes to complete.

You can monitor the update progress through:

- **Web browser** – `http://`*gateway_name*
- **CLI** – `telnet` *gateway_name* `1234`

where *gateway_name* is the host name or IP address of the management controller.

Note – The Oracle ILOM stack requires at least 2 minutes to become operational after a reboot.

The next time you log in to the gateway, this message is displayed:

```
FW upgrade completed successfully on Mon Oct 17 18:36:14 IST 2012.
Please run the "fwverify" CLI command to verify the new image.
This message will be cleared on next reboot.
```

12. **If the Subnet Manager was disabled in** Step 2**, log in to Oracle ILOM, access the restricted Linux shell, and enable the Subnet Manager.**

```
% ssh -l ilom-admin gateway_name
ilom-admin@gateway_name's password: password
-> show /SP/Fabric_Mgmt
NOTE: show on Fabric_Mgmt will launch a restricted Linux shell.
     User can execute switch diagnosis, SM Configuration and IB
     monitoring commands in the shell. To view the list of commands,
     use "help" at rsh prompt.

     Use exit command at rsh prompt to revert back to
     ILOM shell.
FabMan@gateway_name->enablesm
Starting IB Subnet Manager.                                [  OK  ]
Starting partitiond daemon.                                [  OK  ]
FabMan@gateway_name->
```

13. **Verify the firmware version.**

```
FabMan@gateway_name->version
SUN DCS gw version: 2.0.8-1
Build time: Feb  6 2013 09:47:52
FPGA version: 0x33
SP board info:
Manufacturing Date: 2009.06.24
Serial Number: "NCD3R0853"
Hardware Revision: 0x0006
Firmware Revision: 0x0103
BIOS version: NOW1R113
BIOS date: 07/20/2009
FabMan@gateway_name->
```

In the first line of the output is `SUN DCS gw version` *x.y.z-w*, where *x.y.z-w* is the version of the firmware upgraded (or downgraded). For example, 2.0.8-1.

14. **Verify the firmware integrity.**

```
FabMan@gateway_name->fwverify
Checking all present packages:
.......................................................................... OK
Checking if any packages are missing:
.......................................................................... OK
Verifying installed files:
.......................................................................... OK
FabMan@gateway_name->
```