

**Sun Network QDR InfiniBand Gateway
Switch Product Notes for Firmware
Version 2.2**



Part No: E76428-05
October 2019

Part No: E76428-05

Copyright © 2019, 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 documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then 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 that 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 about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

Access to Oracle Support

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

Référence: E76428-05

Copyright © 2019, 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 stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, 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 livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel 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 un risque de dommages corporels. Si vous utilisez ce logiciel ou ce 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 des applications dangereuses.

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 de 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle. 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

Using This Documentation	7
Product Documentation Library	7
Feedback	7
 Late-Breaking Information	9
Features and Functionality Documented	9
Known Problems	10
General Information and Issues	11
Jumbo Frames	11
Subnet Manager	12
Snapshot Dataset Information	12
Software Information and Issues	12
Partition Configuration Consistency Checking	12
Infiniband Switch Power Supplies With Part Number 7351491 Require Firmware 2.2.9-3 or Later	13
IPMI 2.0 is Reenabled After Upgrading From Firmware 2.1 to 2.2	13
New ip Command	13
New ifconfig Command	14
Management Ethernet Port Numbers Used in Firmware Version 2.2	14
Inactive FIPS Target	15
Main Board, Management Controller, and Chassis Serial Numbers	15
High Availability in Partitions	15
Email Alert Rules	15
Oracle Solaris EoIB Considerations After Upgrading to Firmware 2.2	15
Hardware Information and Issues	16
▼ Reset Undervoltage Condition	16
Ethernet Support	16
Unusable Ports	17

Documentation Information and Issues	17
Correct Description of the create_ipoib Command	17
Correct Description of the smsubnetprotection Command	17
Autodisable Command	17
Upgrade and Downgrade Paths Supported	19
Firmware Update Guidelines	19
▼ Determine Management Controller Type	20
disablegwport Command Connector Invalid	21
▼ Replace a Gateway	21
Multiple Subnet Managers in a Single Fabric	22
Upgrading the Gateway Firmware	23

Using This Documentation

- **Overview** – Provides last minute information about the gateway
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Advanced experience administering network hardware

Product Documentation Library

Documentation and resources for this product and related products are available at http://docs.oracle.com/cd/E76425_01/.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Late-Breaking Information

These product notes provide last-minute, late-breaking information regarding the gateway. These notes pertain to the 2.2.14-1 firmware for the gateway.

- [“Features and Functionality Documented” on page 9](#)
- [“Known Problems” on page 10](#)
- [“General Information and Issues” on page 11](#)
- [“Software Information and Issues” on page 12](#)
- [“Hardware Information and Issues” on page 16](#)
- [“Documentation Information and Issues” on page 17](#)

Features and Functionality Documented

Overall, Firmware Version 2.2 is similar to Firmware Version 2.1. Because of this, the Sun Network QDR InfiniBand Gateway Switch Firmware Version 2.2 Documentation library contains mostly documentation for Firmware Version 2.1. However, Firmware Version 2.2 has additional functionality:

- Toolbox CLI utility
- New Oracle ILOM targets and properties
- New Oracle ILOM web interface layout
- Bug fixes, updated procedures, and so on

To support these new features, these documents were either created or updated for Firmware Version 2.2.

- [Sun Network QDR InfiniBand Gateway Switch Documentation Supplement for Firmware Version 2.2](#) – Provides updates to installation, administration, Oracle ILOM, service, and command reference documents.
- [Sun Network QDR InfiniBand Gateway Switch Product Notes for Firmware Version 2.2](#) – This document.

- [Licensing Information User Manual Sun Network QDR InfiniBand Gateway Switch Firmware Version 2.2](#) – Provides new licensing information.

When seeking information about Firmware Version 2.2, read these documents first.

Known Problems

Bug	Description	Workaround
25371165	The IB firmware upgrade is not supported in the BUI.	When running firmware 2.2, the BUI in the firmware upgrade does not work. The BUI also does not work for upgrades and downgrades between 2.2 firmware versions or downgrades to the 2.1 firmware. However, the BUI in the firmware upgrade from 2.1 to 2.2 works. Workaround: Use the ILOM CLI upgrade and downgrade tool as a workaround.
25219512	The /conf/lag.conf file gets deleted when upgrading the firmware from 2.1.6 to 2.2.2.	When upgrading from firmware 2.1 to 2.2.2 or 2.2.4, LAGs are not preserved and need to be recreated. This issue was fixed in firmware version 2.2.5, so you will not experience this issue when you upgrade from 2.1 to 2.2.5 (or later). Workaround: Step 1. Prior to the upgrade from 2.1.x, perform the following: <ul style="list-style-type: none">■ Step a. Execute <code>showlag</code> to find which LAGs are created to find the parameters for the LAGs.■ Step b. Execute <code>showlag lag name</code> for each LAG in Step 1a. to find the connector members for the LAGs. The output of these commands have all the required information for recreating the LAGs after the firmware upgrade to 2.2.2 or 2.2.4. Step 2. After the upgrade to 2.2.2 or 2.2.4 from 2.1.x, perform the following: <ul style="list-style-type: none">■ Step a. To recreate the LAGs, use the command <code>createlag</code> for the LAGs in Step 1b, with the parameters from the output in Step 1a.<ul style="list-style-type: none">■ 0xffff is default for Discover PKey■ NA is default for data SL■ static (i.e. no LACP) is default for LACP■ 1k is default for VNIC limit and Default is default for Pkey If the <code>showlag</code> command in Step 1a shows the use of the default values for a LAG, only the LAG name and the connectors need to be provided as parameters for the <code>createlag</code> command. <ul style="list-style-type: none">■ Step b. Reboot the NM2 GW to enable the VLANs and VNICs on top of the LAGs.
24665004	The default location for the toolbox config files (<code>cli</code> and <code>fabric conf</code>) are not correct.	The default location for the toolbox config files are not in the config directory, so if using this default location, future firmware updates will overwrite the configuration and the user will need to re-enter the configuration.

Bug	Description	Workaround
18129290	<code>smpartition</code> command fails from firmware version 2.0 switch with firmware version 2.1 or 2.2 peer.	<p>Workaround: Do not use the default location for the toolbox config files, instead specify the location that will be used for future firmware versions, which is <code>/config/conf/</code> using the toolbox CLI.</p> <p>In a mixed firmware fabric where the master Subnet Manager is on a gateway with firmware version 2.0 and other gateways or switches in the fabric are firmware version 2.1 or 2.2, issuing the <code>smpartition commit</code> command fails.</p> <p>Workaround: When you use the <code>smpartition</code> commands to partition the mixed firmware fabric, issue the commands from a master Subnet Manager on a switch or gateway with firmware version 2.1 or 2.2.</p>
17159243	VNICs with MAC addresses having eight leading zeros are disabled after a reboot.	<p>VNICs with MAC addresses having four or more leading bytes which are zero (00:00:00:12:34) are not persistent after a reboot, including reboots that occur during a firmware downgrade or upgrade.</p> <p>Workaround: Do not create VNICs with MAC addresses having four or more leading bytes which are zero (eight or more leading zeros).</p>
15703751	No way to set Fabric Monitor console timeout.	<p>The Fabric Monitor will timeout after 15 minutes of idle time.</p> <p>Workaround: You must log in to the web interface and start the Fabric Monitor again.</p>

General Information and Issues

Jumbo Frames

Consider the following 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 `enableism` 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

Partition Configuration Consistency Checking

As of firmware 2.2.13, the consistency checking of the partition configuration for the switch with the Subnet Manager enabled has been improved. The contents of the current local partition configuration file is checked when a Subnet Manager becomes master, and during run-time whenever the contents is read due to that the master Subnet Manager needs to re-configure the IB subnet. If an inconsistency is detected during run-time, the Subnet Manager will restart. If an inconsistency is detected following the local Subnet Manager becoming master, then the Subnet Manager will refuse to configure the subnet until a new configuration version that passes the consistency check has been established. (typically by performing a new "smpartition commit" operation.) Note that there is no locally stored checksum for the partition configuration, hence the above checks do not include any checksum verification.

Any discovered inconsistencies are recorded in the `syslog` file, and an SNMP trap is generated (assuming an SNMP alert target has been defined in the switch ILOM).

In addition to the local consistency checks, there is also an automated periodic check that all switches defined in the `smnodes` list (normally all switches with enabled Subnet Manager) have a valid partition configuration and also that the checksum for the partition configuration is the same on each switch. The check is performed from the switch that at any point in time is hosting the master Subnet Manager and verifies that the checksum computed for the current partition configuration on the master is the same as the corresponding checksum computed for the local partition configuration on each switch with standby Subnet Manager. As with the local consistency checks, any detected problem is reported in the `syslog` file and generates an SNMP trap. (assuming an SNMP alert target has been defined in the switch ILOM.)

Consistent partition configuration for all smnodes can also be checked manually by invoking the `smpartition check` CLI command from the switch currently representing the master Subnet Manager. It is strongly recommended to run `smpartition check` commands after a new partition configuration has been committed (through the `smpartition commit` CLI commands). Also ensure that any detected inconsistencies during run-time are corrected immediately.

Infiniband Switch Power Supplies With Part Number 7351491 Require Firmware 2.2.9-3 or Later

Due to a newer FRU format, IB switches produced after March 2018 with power supplies with part number 7351491 must use firmware version 2.2.9 or later.

IPMI 2.0 is Reenabled After Upgrading From Firmware 2.1 to 2.2

When upgrading from firmware 2.1.x to 2.2.2-2.2.7, IPMI 2.0 is reenabled and must be disabled again.

IPMI is enabled and disabled in the switch's Oracle ILOM CLI or Web interface. The instructions are provided in the *Oracle Integrated Lights Out Manager (ILOM) 3.0 Supplement for the Sun Datacenter InfiniBand Switch 36 Firmware Version 2.1* guide, available at https://docs.oracle.com/cd/E76424_01/pdf/E36270.pdf:

New ip Command

Starting with firmware version 2.2.6, the `ip` command displays device attributes in the CLI.

Syntax

```
ip link show up device [eth0] [lo] [dummy0] [ib0]
```

where:

- `device` is the network device. Valid values include `eth0`, `lo`, `dummy0`, and `ib0`.

This command has a subcommand. This table describes the subcommand and provides the syntax.

Subcommand	Description
up	Displays interfaces that are running.

New ifconfig Command

Starting with firmware version 2.2.6, the `ifconfig` command shows information for an IP interface in the CLI.

Syntax

```
ifconfig [-a | interface]
```

where:

- *interface* is the name of the IP interface. Valid values include `eth0`, `lo`, `dummy0`, and `ib0`.

Options

Option	Description
-a	Displays the status of all interfaces, even those that are down.

Management Ethernet Port Numbers Used in Firmware Version 2.2

The following Ethernet port numbers are used by Daemons in Firmware version 2.2 and should be open across the Ethernet Management Network:

- Port Number 1150
- Port Number 1151
- Port Number 1152

These numbers are used for internal communication and management communication between Infiniband Switches in the Infiniband Subnet.

Inactive FIPS Target

Though the 2.2.14-1 firmware has an Oracle ILOM target of `/SP/services/fips`, this target is inactive.

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 for 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.

Oracle Solaris EoIB Considerations After Upgrading to Firmware 2.2

- Existing Oracle Solaris EoIB datalinks are disabled after upgrading to firmware version 2.2, and new EoIB datalinks are created.
- **If vanity naming is in use** – During an upgrade to firmware version 2.2, consider that you will need to know the names of the EoIB datalinks prior to the upgrade (for example

netX), perform the upgrade, and find the names of the equivalent EoIB datalinks after the upgrade (for example netY). Then disable Oracle Solaris VNICs, VLANs, and IP interfaces that existed over the old datalinks, rename the new datalinks with the old datalink names using the `dladm rename-link netY netX` command, and bring the VNICs, VLANs, and IP interfaces back up.

- **If vanity naming is not in use** – After the upgrade to firmware version 2.2, consider you will need to recreate the Oracle Solaris VNICs, VLANs, and IP interfaces on the new EoIB datalinks.

Hardware Information and Issues

▼ Reset Undervoltage Condition

If a power supply experiences a temporary brownout or undervoltage condition, the `checkpower` command might indicate an `Alert` status for the power supply. The `Alert` does not reset upon supply voltage returning to nominal values. If the `checkpower` command reports an `Alert`, yet supplied voltage is known to be sufficient, perform this task.

1. **Remove the power cord from the suspect power supply.**
2. **Remove the power supply from the switch chassis.**
3. **Wait for 1 minute.**
4. **Install the power supply into the switch chassis.**
5. **Attach the power cord to the power supply.**

Note - If the `checkpower` command still displays an `Alert` status for the power supply, refer to *Gateway Administration*, troubleshooting the switch, to troubleshoot the power supply.

Ethernet Support

The gateway supports 10 Gb/sec Ethernet networks.

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

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

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

Correct Description of the `create_ipoib` Command

In the *Sun Network QDR InfiniBand Gateway Switch Command Reference for Firmware Version 2.1*, the `p_key` parameter is incorrectly described as the partition key with valid values of 1 to 7fff or default. The correct description is that the `p_key` parameter is the 15-bits partition key plus the membership bit as the most significant bit, and when combined, makes a 16 bit parameter with a valid value of 1 to ffff or default.

Correct Description of the `smsubnetprotection` Command

In the *Sun Datacenter InfiniBand Switch 36 Command Reference for Firmware Version 2.1*, the `tid` (transactionID) is listed as an option for the `setreplicationpassword` command. This is not correct. `tid` is not a valid option for that command.

Autodisable Command

The information in this section replaces sections in the [Sun Datacenter InfiniBand Switch 36 Administration Guide for Firmware Version 2.1](#) and the [Sun Datacenter InfiniBand Switch 36 Command Reference for Firmware Version 2.1](#).

Each connector has an autodisable mode for each supported cause that is either on or off. If the mode is on for a specific connector and cause then an automatic disable action will persistently set the associated link state to disabled whenever the situation defined by cause occurs.

Setting the autodisable mode to off will prevent any subsequent autodisable operations for the relevant connector and cause, but will not impact the current link state.

Syntax

Subcommand Syntax	Descriptions
<code>autodisable add connector cause</code>	Turn auto-disabled mode on for the specified connector and cause.
<code>autodisable all</code>	Turn auto-disable mode on for all connectors, all causes.
<code>autodisable del connector cause</code>	Turn auto-disabled mode off for the specified connector and cause.
<code>autodisable list</code>	List connectors which have auto-disable mode turned on along with the specified cause(s).
<code>autodisable none</code>	Turn auto-disable mode off for all connectors, all causes.

where:

- Connector is the name of the connector (0A to 17A and 0B to 17B)
- Cause is the reason to disable the connector:
 - H – The error rate is higher than what the InfiniBand specification permits.
 - L – The link speed or width is less than optimal (Less than 4x QDR as long as the remote port supports 4x QDR).

Description

This hardware command creates an internal list of connectors that are to be automatically disabled, should the links through the connectors exhibit sub-standard performance. The autodisable list has two parts; one for connectors that experience high error rates on their links, the other for connectors that experience suboptimal link speed and width. A connector can be assigned to both parts of the list.

The H cause is configured for a high error rate such as the following:

- Symbol error count higher than allowed by the IB specification on the switch connector.
- Symbol error count higher than allowed by the IB specification on the remote HCA port (from firmware version 2.2.4).

Upgrade and Downgrade Paths Supported

This table identifies supported firmware version upgrade and downgrade paths.

From	To			
	2.1.2 - 2.1.6	2.1.7 - 2.1.9	2.2.2-2.2.7	2.2.9 and later
2.1.2 - 2.1.6	Upgrade or Downgrade	Upgrade	Upgrade	
2.1.7 - 2.1.9	Must downgrade to 2.1.6 first, and management controller must support.	Upgrade or Downgrade	Upgrade	
2.2.2 - 2.2.7	Must downgrade to 2.1.6 first, and management controller must support.	Must downgrade to 2.1.8 first.	Upgrade or Downgrade	Must first upgrade to 2.2.7-2 using the 2.2.7_2.pkg file. Then downgrade to required firmware version.
2.2.9 and later			Must first downgrade to 2.2.7-2 using the 2.2.7_2_signed.pkg file. Then downgrade to required firmware version.	Upgrade and Downgrade

For example, downgrading from firmware version 2.2.7 to 2.1.6 is supported, if the gateway is configured with an older management controller. However, downgrading directly from firmware version 2.2.7 to 2.1.2 is not supported. If the gateway is configured with an older management controller, you must first downgrade from firmware version 2.2.7 to 2.1.6 using the downgrade version of firmware 2.1.6. Then you can downgrade from 2.1.6 to 2.1.2.

See “[Firmware Update Guidelines](#)” on page 19 for instructions on upgrading and downgrading firmware.

Firmware Update Guidelines

- Now that version 2.2.9 and newer firmware is signed firmware (all earlier versions are unsigned), all upgrades and downgrades are performed using the 2.2.7-2 pkg files that are available from My Oracle Support (<https://support.oracle.com>) patch ID 26575828.
- You can downgrade from firmware 2.2.7 to 2.1.8 or 2.1.7. You downgrade to firmware 2.1.8 using a special downgrade version of firmware 2.1.8. After this is done, you can then downgrade from firmware 2.1.8 to 2.1.7. The special downgrade version of the 2.1.8 firmware is part of patch 22173637 available at My Oracle Support. Ensure that you read the README file included in the patch.

- You can downgrade from firmware 2.2.7 to 2.1.6 if your management controller is of the older type. See [“Determine Management Controller Type” on page 20](#). You downgrade to firmware 2.1.6 using a special downgrade version of firmware 2.1.6. After this is done, you can then downgrade from firmware 2.1.6 to any supported firmware. See [“Upgrade and Downgrade Paths Supported” on page 19](#). The special downgrade version of the 2.1.6 firmware is part of patch 20380280 available at My Oracle Support. Ensure that you read the README file included in the patch.
- Newer management controllers cannot be downgraded to 2.1.6 or lower. To determine if your management controller is the newer type, see [“Determine Management Controller Type” on page 20](#).
- For versions 2.1.x through 2.2.2-2.2.7, the initial firmware upgrade or downgrade might display an error at the beginning and then exit. This is normal. An installer in the firmware has set or reset prerequisite values. Repeat the upgrade or downgrade.
- For greater security, ensure that all Sun Network QDR InfiniBand Gateway Switches and Sun Datacenter InfiniBand Switch 36 switches are upgraded to the latest firmware version. A miss-match of firmware versions reduces or negates secret M_Key functionality, and compromises fabric security.

For example in a mixed 2.2.7 and 2.1.5 fabric, where SMs are running on both 2.2.7 and 2.1.5 switches, you must use the `-override-inconsistent-partition-configurations` option with the `smsubnetprotection` command. To do this is insecure.

▼ Determine Management Controller Type

1. **From the Oracle ILOM CLI, display the management controller part number.**

```
-> show /SYS/MB fru_part_number
/SYS/MB
  Properties:
    fru_part_number = 7014379
->
```

2. **In the output, record the 7-digit number of the `fru_part_number` property.**

For example, 7014379.

- If the part number is greater than 7057249, the management controller is the newer type, and you cannot downgrade to 2.1.6 or lower.
- If the part number is less than or equal to 7057249, the management controller is the older type, and you can downgrade to 2.1.6 or lower.

For this example, the part number is less than 7057249, and it is possible to downgrade to firmware 2.1.6 or lower.

disablegwport Command Connector Invalid

In the Sun Network QDR InfiniBand Gateway Switch documentation and in the help and usage provided by the `disablegwport` command itself, the `0A-ETH` and `1A-ETH` connector names are valid as command variables. This is incorrect. Only these connector names are valid as command variables:

- `0A-ETH-1`, `0A-ETH-2`, `0A-ETH-3`, `0A-ETH-4`
- `1A-ETH-1`, `1A-ETH-2`, `1A-ETH-3`, `1A-ETH-4`

Do not use `0A-ETH` or `1A-ETH` as command variables for the `disablegwport` command.

▼ Replace a Gateway

This procedure briefly outlines the steps to replace a gateway within an InfiniBand fabric. Refer to *Gateway Service*, *Gateway Administration*, and *Gateway Installation* for more information. This procedure assumes that a secret `M_Key` policy is in use and the fabric is partitioned.

1. **Set the priority of the gateway to be removed to the lowest.**
Wait for any handover to complete.
2. **Create a backup of the gateway configuration for the gateway to be removed.**
3. **Completely power off the gateway.**
4. **Remove the gateway from the fabric and management network.**
5. **Install and power on the replacement gateway, but do not connect it to the fabric or the fabric's management network.**
6. **Configure the replacement gateway with the same host name, IP address, as the gateway removed.**
7. **Restore the gateway configuration previously backed up.**
8. **Readout and store the SM priority from the restored gateway configuration using `setsmpriority list` and then set the SM priority to the lowest in the Infiniband Subnet.**
9. **Completely power off the replacement gateway.**

10. **Connect the replacement gateway to the fabric's management network and the fabric.**
11. **Power on the replacement gateway.**
Wait for any negotiation and propagation to complete.
12. **Setup the `smnodes` list and fabric configuration file for the replacement gateway, and ensure that the list and file are consistent with other gateway and switch lists and files.**
13. **Perform a `smpartition start` and `smpartition commit` from the master Subnet Manager.**
Wait for the partition configuration to propagate to the replacement gateway.
14. **Perform a `smsubnetprotection start` and `smsubnetprotection commit` from the master Subnet Manager.**
Wait for the secret `M_Key` policy to propagate to the replacement gateway.
15. **Set the priority of the replacement gateway to that originally of the removed gateway.**
Wait for any handover to complete.

Multiple Subnet Managers in a Single Fabric

When a fabric has multiple Subnet Managers, you must configure some parameters uniquely and some identically.

- **Subnet Manager Priority** – Subnet Managers can have different Priority values. The overall priority is determined from both the gateway's GUID and the Priority value. Configure the Subnet Manager(s) with the highest Priority value first, then configure any remaining Subnet Manager(s).
- **Subnet Manager Prefix** – All Subnet Manager must use the same prefix. Configure the standby Subnet Managers first, then configure the master Subnet Manager.
- **Subnet Manager Controlled Handover** – All Subnet Manager must use the same configuration for controlled handover. Configure the standby Subnet Managers first, then configure the master Subnet Manager.
- **Subnet Manager Routing Algorithm** – All Subnet Manager must use the same routing algorithm. Configure the standby Subnet Managers first, then configure the master Subnet Manager.

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 2.2, 2.2.14, and 2.2.14-1 respectively. The following two procedures describe how to acquire and upgrade the firmware through the 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. **From the More... drop-down menu, select Patches & Updates.**

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, 2.2.14.

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, 30303703. 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, p30303703_2214_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_2.2.14-1/SUN_DCS_GW/sundcs_gw_repository_2.2.14_1.pkg file.

The readme_SUN_DCS_GW_2.2.14-1.txt file contains the latest information about the firmware release.

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

18. **Upgrade the gateway firmware.**

See [“Upgrade the Gateway Firmware \(CLI\)” on page 24](#).

▼ Upgrade the Gateway Firmware (CLI)

Note - Before upgrading or downgrading the gateway firmware, see [“Firmware Update Guidelines” on page 19](#).

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

For example:

```
% ssh -l ilom-admin gateway_name
```



```
ilom-admin@gateway_name's password: password
->
```

where *gateway_name* is the host name of the management controller. Initially, the *password* is *ilom-admin*.

2. (Optional) If the Subnet Manager is running on the management controller, disable the Subnet Manager.

```
-> start /SYS/Fabric_Mgmt
Are you sure you want to start /SYS/Fabric_Mgmt (y/n)? y

NOTE: start /SYS/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->disable sm
Stopping partitiond-daemon.                [ OK ]
Stopping IB Subnet Manager..                [ OK ]
FabMan@gateway_name->exit
exit
->
```

3. 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, TFTP, and HTTP protocols are supported when upgrading from firmware 2.1. If upgrading from firmware 2.2, SFTP and SCP are also supported.
- *pkgname* is the name of the firmware package in the transfer directory.

For example, using the HTTP protocol:

```
-> load -source http://123.45.67.89/tmp/sundcs_gw_repository_2.2.14_1.pkg
Downloading firmware image. This will take a few minutes.
```

The firmware is downloaded. The upgrade begins. If you did not disable the Subnet Manager, a warning is displayed. 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.

WARNING: Subnet manager is running on the switch. Proceeding with installation will stop it.

4. Answer y to the prompt to commit to the upgrade.

The upgrade begins.

Are you sure you want to load the specified file (y/n)? **y**
Setting up environment for firmware upgrade. This will take a few minutes.
Starting SUN DCS gw FW update

```
=====
Performing operation: I4 A
=====
I4 A: I4 is already at the given version.
```

```
=====
Performing operation: BX A
=====
BX A: BX is already at the given version.
```

```
=====
Performing operation: BX B
=====
BX B: BX is already at the given version.
```

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

```
=====
Performing operation: SUN DCS gw firmware update
=====
SUN DCS gw Kontron module fw upgrade from 2.2.9-3 to 2.2.14-1:
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://system] OR at OS command line prompt by using command [telnet system 1234] where system is the hostname or IP address of SUN DCS 36P or GW.

Firmware update is complete.

5. Restart the gateway to enable the new firmware.

```
-> reset /SP
Are you sure you want to reset /SP (y/n)? y
Performing reset on /SP
Broadcast message from root@gateway_name
      (unknown) at 12:45 ...
The system is going down for reboot NOW!
-> Connection to gateway_name closed by remote host.
Connection to gateway_name closed.
-bash-3.2$
```



Caution - Do not power cycle or reset the Switch prior to completion of the reset process, as that will interrupt the firmware installation.

Note - The reset process takes about 4 to 5 minutes. However, it can take longer if management network or the server hosting the firmware has limited resources.

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 Wed Aug 15 09:06:10 CET 2019.
Please run the "fwverify" CLI command to verify the new image.
This message will be cleared on next reboot.
```

6. Access the restricted Linux shell, and verify the firmware version.

```
% ssh -l ilom-admin gateway_name
ilom-admin@gateway_name's password: password
-> start /SYS/Fabric_Mgmt
```

Are you sure you want to start /SYS/Fabric_Mgmt (y/n)? **y**

NOTE: start /SYS/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->version
SUN DCS gw version: 2.2.14-1
Build time: Aug 14 2019 09:07:27
FPGA version: 0x34
SP board info:
Manufacturing Date: 2012.01.10
Serial Number: "NCD8F0072"
Hardware Revision: 0x0007
Firmware Revision: 0x0000
BIOS version: SUN0R100
BIOS date: 06/22/2010
FabMan@gateway_name->
```

In the first line of the output for the version command 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, 2.2.14-1.

7. If you disabled the Subnet Manager in [Step 2](#), enable it now.

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

8. Verify the firmware integrity.

```
FabMan@gateway_name->fwverify -s
Checking all present packages:
.....
.....
..... OK

Checking if any packages are missing:
.....
.....
..... OK

Checking FW Coreswitch:
FW Version: 7.4.3002 OK
```

```
PSID: SUNX2826_I40_004 OK
Verifying image integrity OK
```

```
Checking FW Bridge-0:
  FW Version: 8.6.2010 OK
  PSID: SUNX2826_BX0_006 OK
  Verifying image integrity OK
```

```
Checking FW Bridge-1:
  FW Version: 8.6.2010 OK
  PSID: SUNX2826_BX1_006 OK
  Verifying image integrity OK
```

```
FabMan@gateway_name->
```

