

**Sun Network QDR InfiniBand Gateway
Switch Product Notes for Firmware
Version 2.1**

ORACLE®

Part No: E36258-07
March 2017

Part No: E36258-07

Copyright © 2013, 2017, 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.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

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: E36258-07

Copyright © 2013, 2017, 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 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, 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.

Accessibilité de la documentation

Pour plus d'informations sur l'engagement d'Oracle pour l'accessibilité à la documentation, visitez le site Web Oracle Accessibility Program, à l'adresse <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

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
 Sun Network QDR InfiniBand Gateway Switch Product Notes	 9
Known Problems	9
General Information and Issues	11
Jumbo Frames	11
Subnet Manager	11
Snapshot Dataset Information	11
Software Information and Issues	12
New Functionality for Commands	12
Main Board, Management Controller, and Chassis Serial Numbers	12
High Availability in Partitions	12
Email Alert Rules	12
Time Zone Support	13
Firmware Update Guidelines	14
Host Software and Firmware	15
Oracle Solaris EoIB Considerations After Upgrading to Firmware 2.1	15
Hardware Information and Issues	16
Ethernet Support	16
Number of LAGs	16
Unusable Ports	16
Documentation Information and Issues	16
Upgrade and Downgrade Paths Supported	16
Command Corrections	17
VNIC Limits for ETH Ports	27
VNIC Limits for the createLag Command	27

Command Output Description Incorrect	27
Understanding Switch Chip Port to Connector Routing	28
disablegwport Command Connector Invalid	29
Features and Functionality Documented	29
Correct Power Consumption	30
Node Description Format	30
Battery Service Sequence	30
▼ Replace a Gateway	31
Multiple Subnet Managers in a Single Fabric	32
Replication Password Information Incomplete	33
Temperature Sensor Thresholds Incorrect	33
README File Is Incomplete	33
Configuring Secure Fabric Management	34
▼ Acquire the BXOFED Software (Linux)	34
▼ Acquire the ConnectX-2 Firmware	35
Upgrading the Gateway Firmware	37

Using This Documentation

- **Overview** – Provide late-breaking information for the Sun Network QDR InfiniBand Gateway Switch from Oracle.
- **Audience** – Technicians, system administrators, and authorized service providers.
- **Required knowledge** – Advanced experience troubleshooting and replacing hardware.

Product Documentation Library

Documentation and resources for this product and related products are available at http://docs.oracle.com/cd/E36256_01.

Feedback

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

Sun Network QDR InfiniBand Gateway Switch

Product Notes

These product notes provide last-minute, late-breaking information regarding the gateway. These notes apply to the 2.1.9-1 firmware for the gateway.

Known Problems

Bug	Description	Workaround
24673478	A direct firmware upgrade from 1.3 to 2.1.7, 2.1.8, or 2.1.9 fails.	Upgrading a gateway switch directly from 1.3 to 2.1.7 or later fails. Workaround: Perform a firmware upgrade from 1.3 to 2.1.6, and then upgrade to 2.1.7, 2.1.8, or 2.1.9.
22195180	Firmware upgrade from 2.1.8 to 2.1.8 without -force option causes fwverify command to fail.	When reinstalling firmware 2.1.7, 2.1.8, or 2.1.9, neither as an upgrade or a downgrade, the process completes without errors. However, the fwverify command indicates that the installation has failed. Workaround: You must force a reinstallation from the Oracle ILOM CLI. Use the -force option to the Oracle ILOM load command. See “Upgrade the Gateway Firmware (CLI)” on page 39 .
20714591	All Subnet Managers segfault at 54 ip 080ba543 sp xxxxxxxx error 4 displayed.	An InfiniBand fabric that is not a pure Fat Tree topology might experience a segmentation fault across all Subnet Managers in the fabric. This can occur when cabling or decabling an active gateway or switch configured for Fat Tree topology, as the actual topology at that instant is not pure Fat Tree. Workaround: Either use the setsmrouting command to set the Subnet Managers to Min Hop routing, or ensure a pure Fat Tree topology and power off the gateway before cabling or decabling it. See “Replace a Gateway” on page 31 .
20649071	Fabric Director is filling up the gateway root file system.	If the hostname of a management controller exceeds 26 characters, the Fabric Director incessantly writes errors to the log file which can fill the root filesystem if left unchecked. Workaround: Do not use more than 26 characters for the hostname of the management controller.

Known Problems

Bug	Description	Workaround
		Note - The maximum number of characters for the node description field is 17.
18129290	smpartition command fails from firmware version 2.0 switch with firmware version 2.1 peer.	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, issuing the <code>smpartition commit</code> command fails. 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.
17408412	connector command not fully functional from Fabric_Mgmt restricted shell.	When the <code>info</code> or <code>dump</code> options of the <code>connector</code> command are issued in the <code>Fabric_Mgmt</code> , <code>Gateway_Mgmt</code> , or <code>Switch_Diag</code> restricted shells, the command fails. Workaround: Use the Fabric Monitor feature of the Oracle ILOM web interface to retrieve connector information.
17159243	VNICs with MAC addresses having eight leading zeros are disabled after a reboot.	VNICs with MAC addresses having four or more leading bytes which are zero (00:00:00:00:12:34) are not persistent after a reboot, including reboots that occur during a firmware downgrade or upgrade. Workaround: Do not create VNICs with MAC addresses having four or more leading bytes which are zero (eight or more leading zeros).
16847481	Setting time zone changes permissions of time zone file.	After setting the time zone, the permissions of the time zone file are changed and must be reset to their original values. Otherwise, the <code>fwverify</code> command returns an error. Workaround: Change the permissions on the time zone file. 1. After setting the time zone, become the <code>root</code> user. 2. Set the permissions for the time zone file. # chmod 644 /conf/localtime
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 <code>ibnetdiscover</code> command.
15703751	No way to set Fabric Monitor console timeout.	The Fabric Monitor will timeout after 15 minutes of idle time. Workaround: You must log in to the web interface and start the Fabric Monitor again.
15560533	Setting an alert rule to <code>ipmipet</code> sometimes does not work.	There is no impact to the InfiniBand fabric. Workaround: After setting the alert rule to <code>ipmipet</code> , set the level to <code>disable</code> , and then set the level to the desired value.

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 `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

New Functionality for Commands

Both the `smpartition` and `smsubnetprotection` command have new functionality. See [“smpartition Command” on page 18](#) and [“smsubnetprotection Command” on page 22](#).

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.

Time Zone Support

The following time zones are only supported in firmware versions 1.3.4 and newer 1.3.x releases, 2.0.7 and newer 2.0.x releases and all 2.1.x releases. If you upgrade or downgrade to a firmware version other than those identified, 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

In firmware version 2.1.2-2 through 2.1.9-1, these additional time zones are supported:

- Africa/Juba
- America/Creston
- Asia/Hebron

If you downgrade from firmware version 2.1.2-2 through 2.1.9-1, you must set an alternative time zone.

In firmware version 2.1.6-2, 2.1.7-2, and 2.1.9-1, these additional time zones are supported:

- Asia/Khandyga
- Asia/Ust-Nera
- Europe/Busingen

- Pacific/Bougainville

If you downgrade from firmware version 2.1.6-2, 2.1.7-2, or 2.1.9-1, you must set an alternative time zone.

Note - See [“Firmware Update Guidelines” on page 14](#) about downgrading from firmware version 2.1.7-2 or 2.1.9-1.

Firmware Update Guidelines

- When downgrading from firmware 2.1.7, 2.1.8, or 2.1.9, you must first downgrade to firmware 2.1.6 using a special 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 16](#). 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.

Note - The downgrade version of the 2.1.6 firmware requires initiating downgrade twice.

- For greater security, ensure that all Sun Network QDR InfiniBand Gateway Switches and Sun Datacenter InfiniBand Switch 36 switches are upgraded to firmware version 2.1.6 or later version. A miss-match of firmware versions reduces or negates secret M_Key functionality, and compromises fabric security.

For example, in a mixed 2.1.6 and 2.1.5 fabric, you must use the `-override-inconsistent-partition-configurations` option with the `smsubnetprotection` command. Doing this is not secure.
- If you are going to downgrade the firmware to a version earlier than 2.1.x, you must reconfigure secret M_Key functionality, if it was enabled. Refer to the [Sun Network QDR InfiniBand Gateway Switch Administration Guide for Firmware Version 2.1](#) for instructions on securing the fabric.
- 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 if these features were configured. Refer to the [Sun Network QDR InfiniBand Gateway Switch Administration Guide for Firmware Version 2.1](#) for instructions on removing partitions for firmware downgrade.
- If you are going to downgrade from firmware 2.0.x to 1.3.5 or earlier, you might see these types of messages in the `/var/log/message` file after the downgrade:

`lnda: Unknown config parameter: ErrLogCount=100; .`
`lnda: Unknown config parameter: ErrLogTimeInterval=100; .`

The `ErrLogCount` and `ErrLogTimeInterval` configuration parameters introduced in firmware 2.0.x are unknown to firmware version 1.3.5 and earlier LDAs. The LDA logs these messages and ignores them afterward. The messages appear once per LDA startup and are harmless.

If you later upgrade from firmware 1.3.5 or earlier 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 a Linux or Oracle Solaris InfiniBand host, the host must have the following software and firmware installed:

- **For Linux hosts** – BXOFED software. See [“Acquire the BXOFED Software \(Linux\)” on page 34](#).
- **For Oracle Solaris hosts** – ethernet-over-ib software. The software package is available in the Oracle Solaris 11 image.
- **For either host** – ConnectX-2 firmware version 2.7.8130 or newer. See [“Acquire the ConnectX-2 Firmware” on page 35](#).

All of these software packages are available from Oracle.

Oracle Solaris EoIB Considerations After Upgrading to Firmware 2.1

- Existing Oracle Solaris EoIB datalinks are disabled after upgrading to firmware version 2.1, and new EoIB datalinks are created.
- **If vanity naming is in use** – During an upgrade to firmware version 2.1, 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 2.1, consider you will need to recreate the Oracle Solaris VNICs, VLANs, and IP interfaces on the new EoIB datalinks.

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 10GbE 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

Upgrade and Downgrade Paths Supported

The README files of previous firmware releases have omitted supported upgrade and downgrade paths.

This table identifies supported firmware version upgrade and downgrade paths.

From	To				
	1.1.x	1.3.x	2.0.x	2.1.2 - 2.1.6	2.1.7 - 2.1.9
1.1.x	Upgrade or Downgrade	Upgrade or Downgrade			
1.3.x	Upgrade or Downgrade	Upgrade or Downgrade	Upgrade or Downgrade	Upgrade or Downgrade	<ul style="list-style-type: none"> ■ Upgrade ■ Downgrade to 2.1.6 first
2.0.x		Upgrade or Downgrade	Upgrade or Downgrade	Upgrade or Downgrade	<ul style="list-style-type: none"> ■ Upgrade ■ Downgrade to 2.1.6 first
2.1.2 - 2.1.6		Upgrade or Downgrade	Upgrade or Downgrade	Upgrade or Downgrade	<ul style="list-style-type: none"> ■ Upgrade ■ Downgrade to 2.1.6 first
2.1.7 - 2.1.9		<ul style="list-style-type: none"> ■ Upgrade ■ Downgrade to 2.1.6 first 	<ul style="list-style-type: none"> ■ Upgrade ■ Downgrade to 2.1.6 first 	<ul style="list-style-type: none"> ■ Upgrade ■ Downgrade to 2.1.6 first 	Upgrade or Downgrade

For example, upgrading from firmware version 1.1.2 to 1.3.5 is supported. However, upgrading from firmware version 1.1.2 to 2.1.5 is not. You must first upgrade from firmware version 1.1.2 to 1.3.x, and from firmware version 1.3.x to 2.1.5.

Similarly, downgrading from firmware version 2.1.6 to 2.0.7 is supported. However, downgrading from firmware version 2.1.9 to 2.0.7 is not. You must first downgrade from firmware version 2.1.9 to 2.1.6 using the downgrade version of firmware 2.1.6. Then you can downgrade from 2.1.6 to 2.0.7.

Note - See [“Firmware Update Guidelines” on page 14](#) for information on upgrading and downgrading firmware.

Command Corrections

Both the `smpartition` and `smsubnetprotection` commands have had updates or features that were mistakenly undocumented. The following two sections are updated and corrected versions of command reference information that supersede information currently provided in the [Sun Network QDR InfiniBand Gateway Switch Command Reference for Firmware Version 2.1](#) and the [Sun Network QDR InfiniBand Gateway Switch Administration Guide for Firmware Version 2.1](#).

smpartition Command

Manages the partition configuration.

Syntax

`smpartition subcommand [-h]`

This hardware command has subcommands that determine its functionality. This table describes the *subcommands* and provides their syntax.

Subcommand Syntax	Description
<code>peerversion</code>	Displays the firmware version of smnode peers of the master Subnet Manager.
<code>start [tid]</code>	Initiates a new configuration based upon a currently used configuration.
<code>create [tid tid] [-n partition_name] -pkey p_key [use_grh] [-m defmember] [-flag [ipoib [mtu mtu][rate rate][sl sl][scope scope]]]</code>	Creates a new partition. The -m option configures the default membership for the partition.
<code>delete [tid tid] -n partition_name -pkey p_key</code>	Deletes a partition.
<code>add [tid tid] -n partition_name -pkey p_key -port port ALL_CAS ALL_SWITCHES ALL_ROUTERS [-m member]</code>	Adds one or more ports to the partition. The -m option sets the membership for the port.
<code>remove [tid tid] -n partition_name -pkey p_key -port port ALL_CAS ALL_SWITCHES ALL_ROUTERS</code>	Removes one or more ports to the partition.
<code>modify [tid tid] -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]]</code>	Modifies a partition flag or port membership. The -m option sets the membership for the port.
<code>list active modified [no-page]</code>	Displays the active or modified configuration. By default, the output is displayed one page at a time, advanced by pressing the spacebar. The no-page option enables a continuous stream of output without page breaks.
<code>listcurrenttid</code>	Lists the current transaction ID.
<code>commit [tid tid]</code>	Commits the modified configuration to become the active configuration.
<code>abort [tid tid]</code>	Abruptly ends the configuration session. All modified configuration information is lost and the active configuration remains unchanged.

where:

- *tid* is the transaction ID (0 to 4294967295).
- *partition_name* is an alphanumeric tag to the InfiniBand partition (optional).

- *p_key* is the partition key (1 to 7fff or default).

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

- *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.

- *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 add, remove, or 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.

Description

This hardware command is used to manage the InfiniBand partitions and is available only on management controllers that are hosting the primary (or master) Subnet Manager. There are two configurations for the InfiniBand partition, the active configuration and the modified configuration. When configuring a partition, you must initiate the configuration session with the `smpartition start` command. During the session, you create a modified copy of the active configuration. To end the session, you must use the `smpartition commit` command to make the modified configuration the active configuration. Once committed, the active configuration is distributed to all Subnet Managers in the InfiniBand fabric where the management controller's IP addresses are listed in the Subnet Manager nodes file.

The Subnet Manager nodes file must exist in every management controller file system. The file contains a list of IP addresses of all active management controllers hosting a Subnet Manager in your fabric. The file should have an entry for every Sun Network QDR InfiniBand Gateway Switch and Sun Datacenter InfiniBand Switch 36 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 and the fabric element configuration file. Refer to the `smnodes` command and the `createfabric` command in the [Sun Network QDR InfiniBand Gateway Switch Command Reference for Firmware Version 2.1](#).

Options

This table describes the options to the `smpartition` command and their purposes.

Option	Purpose
<code>tid</code>	Specifies the transaction ID. The transaction ID adds an additional layer of security to the <code>smpartition</code> command. The identifier is a 32-bit unsigned integer, returned when the partition configuration session is started with the <code>smpartition start tid</code> command. This identifier is then required for all subsequent actions to the particular partition. Use of the transaction ID mediates changes to the partition by multiple users.
<code>-n</code>	Specifies the partition name.
<code>-pkey</code>	Specifies the partition key.
<code>use_grh</code>	If the <code>use_grh</code> option is used in the <code>smpartition create</code> command, a requirement of the partition is that Global Route Headers (GRH) are attached to InfiniBand messages and are used for path resolution requests made to the Subnet Manager. This option provides additional security for Engineered Systems.
<code>-m</code>	Specifies the membership type. If the <code>-m</code> option is used in the <code>smpartition create</code> command, the default membership type of the partition is specified. If the <code>-m</code> option is used with the <code>smpartition add</code> command or <code>smpartition modify</code> command, the membership type of the port is specified.

Option	Purpose
	If ports are added to the partition without specifying the membership type, the default membership type for the partition is applied to the port.
-port	<p>Specifies the port or ports to be acted upon:</p> <ul style="list-style-type: none"> ■ <i>port</i> – The GUID of the port to be acted upon. <p>Alternatively, one these special parameters is specified instead of a GUID.</p> <ul style="list-style-type: none"> ■ ALL_CAS – All CAs in the InfiniBand fabric. ■ ALL_SWITCHES – All switches. ■ ALL_ROUTERS – All routers.
-flag	<p>Specifies:</p> <ul style="list-style-type: none"> ■ <i>ipoib</i> – If present, IP over InfiniBand is to be supported. ■ <i>mtu</i> – Sets the MTU. ■ <i>rate</i> – Sets the throughput of a link (link width + link speed). ■ <i>sl</i> – Sets the service level. ■ <i>scope</i> – Sets the multicast address scope. <p>Note - The -flag option by itself disables IPoIB.</p> <p>If you use the -flag option in the <code>smpartition modify</code> command, you must restart the master Subnet Manager or perform a Subnet Manager handover after the <code>smpartition commit</code> command. Because this causes an interruption of service, if you want flag parameters different than the default, consider setting partition flags at the time of partition creation.</p>
-h	Provides help.

Example

This example shows how to display the active configuration of the InfiniBand partition with the `smpartition` command.

```
FabMan@gateway_name->smpartition list active
# Sun DCS IB partition config file
# This file is generated, do not edit
#! version_number : 16
Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=full;
SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
part1 = 0x9001,ipoib:
0x0002c90300089138=full,
0x0002c9030008923b=full,
0x0002c9030008923c=full,
0x0002c90300089103=limited,
0x0002c90300089104=full,
0x0002c90300089137=limited;
part2 = 0x9002,ipoib:
```

```

0x0003ba000100e389=full,
0x0002c903000890cb=limited,
0x0002c903000890cc=full,
0x0002c903000890c8=full,
0x0002c903000890c7=limited;
FabMan@gateway_name->

```

smsubnetprotection Command

Manages the secret M_Key.

Syntax

smsubnetprotection *subcommand* [-h]

This hardware command has subcommands that determine its functionality. This table describes the *subcommands* and provides their syntax.

Subcommand Syntax	Description
start [-force][-addononly -deleteonly][-override-inconsistent-partition-configurations][-override-unavailable-smnodes][tid]	Initiates a new configuration based upon a currently used configuration. Use the -force option or -override-unavailable-smnodes option to bypass the partition daemon check. See “Options” on page 26 .
list active modified	Displays a list of active secret M_Keys, the current secret M_Key, and the enabled status, or displays a list of pending M_Keys and the M_Key to be assigned to current status.
listlocalmkey	Displays the current local M_Key for an I4 switch chip without a corresponding Subnet Manager and its status.
listcurrenttid	Lists the current transaction ID.
setlocalsecretmkey <i>m_key</i>	Sets the secret M_Key locally for an I4 switch chip without a corresponding Subnet Manager.
clearlocalmkey	Clears the local secret M_Key.
add <i>m_key</i> [tid <i>tid</i>]	Adds an M_Key to the configuration.
delete <i>m_key</i> [tid <i>tid</i>]	Deletes an M_Key from the configuration.
undo [tid <i>tid</i>]	Reverts the previous add, delete, or set-current operation.
set-current <i>m_key</i> [tid <i>tid</i>]	Sets the current M_Key.
commit [-force][-override-inconsistent-partition-configurations][-override-unavailable-smnodes][tid <i>tid</i>]	Commits the modified configuration to become the active configuration. Use the -force option or

Subcommand Syntax	Description
<code>abort [tid <i>tid</i>]</code>	-override-unavailable-smnodes option to bypass the partition daemon check. See “Options” on page 26. Abruptly ends the configuration session. All modified configuration information is lost, and the active configuration remains unchanged.
<code>setreplicationpassword <i>password</i> [tid <i>tid</i>]</code>	Configures the replication (and encryption) password.
<code>enablesecretmkey [-force][-override-inconsistent-partition-configurations][-override-unavailable-smnodes]</code>	Enables secret M_Key functionality. Use the -force option or -override-unavailable-smnodes option to bypass the partition daemon check. See “Options” on page 26.
<code>disablesecretmkey [-force][-override-inconsistent-partition-configurations][-override-unavailable-smnodes]</code>	Disables secret M_Key functionality. Use the -force option or -override-unavailable-smnodes option to bypass the partition daemon check. See “Options” on page 26.

where:

- *m_key* is the management key (16 hexadecimal digits).
- *tid* is the transaction ID (0 to 4294967295).
- *password* is encryption string for M_Key replication (8 alphanumeric characters).

Description

This hardware command manages the secret M_Key and its implementation. The secret M_Key is a passphrase used by trusted Subnet Managers to securely perform activities (enabling ports, setting parameters, and so on) on the I4 switch chips as well as any end node in the InfiniBand fabric.

A readable M_Key is an M_Key operating in a mode, where the node that possesses the M_Key permits the value of the M_Key to be read through in-band operations on the InfiniBand fabric, without first specifying the current readable M_Key value. The secret M_Key is an M_Key that cannot be obtained in-band by way of the InfiniBand fabric without first knowing the current secret M_Key value.

Use the `smsubnetprotection` command and its subcommands to create and manage the list of secret M_Keys. When configuring a list of secret M_Keys, you first enable secret M_Key functionality with the `enablesecretmkey` subcommand. Then you initiate the configuration session on the master Subnet Manager with the `smsubnetprotection start` command. During the session, you add or delete secret M_Keys to the configuration, set the current secret M_Key, and list the M_Keys configured.

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

To end the session, you must use the `smsubnetprotection commit` command to make the configuration active. Once committed, the configuration is automatically distributed to all Subnet Managers in the InfiniBand fabric.

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

Should a local secret M_Key be created for an I4 switch chip without a corresponding Subnet Manager, that secret M_Key is only recognized by that I4 switch chip, and is unrecognized by the other I4 switch chips in the InfiniBand fabric.

Because of the complexity of the secret M_Key functionality, this table describes the impact of certain scenarios and actions you can take.

Scenario	Impact and Actions
Setting up secret M_Key in a mixed firmware fabric.	<p>If the master Subnet Manager has firmware 2.1, only other Subnet Managers with firmware 2.1 can administrate the fabric. For Subnet Managers with firmware 2.0 or lower, the fabric “disappears”.</p> <p>If the master Subnet Manager has firmware 2.0 or lower, you can only set up local secret M_Keys for the I4 switch chips on their respective Subnet Managers with firmware 2.1.</p> <p>Both situations are unsupported and not recommended.</p>
Downgrading firmware after secret M_Key has been enabled.	<p>If the master Subnet Manager is downgraded to firmware 2.0 or lower and there is a standby Subnet Manager with firmware 2.1, the secret M_Key is maintained through the standby Subnet Manager during the master Subnet Manager's reboot. After the reboot, the situation becomes as in the first scenario.</p> <p>If you downgrade any other Subnet Manager to firmware 2.0 or lower, the situation becomes as in the first scenario.</p> <p>Before you downgrade any firmware, disable secret M_Key.</p> <p>Note - Readable M_Key is not affected by a downgrade from firmware 2.1 to 2.0.</p>
Upgrading from a lower firmware version.	<p>Do not enable secret M_Key until all Subnet Managers in the fabric are at firmware version 2.1 or higher.</p>
Introducing a new Subnet Manager with firmware 2.1 or higher, yet no secret M_Key policy, into a secret M_Key fabric.	<p>Before introducing the new Subnet Manager to the fabric:</p> <ol style="list-style-type: none"> 1. Disable the new Subnet Manager. 2. Set the new Subnet Manager priority to the lowest. 3. Update the <code>smnodes</code> file with the <code>smnodes</code> command. 4. Enable the new Subnet Manager. <p>After introducing the new Subnet Manager to the fabric:</p> <ol style="list-style-type: none"> 1. Update the fabric configuration with the <code>fdconfig</code> command. 2. Update the fabric mapping with the <code>createfabric</code> command. 3. Perform a <code>smpartition start</code>, then <code>smpartition commit</code>, then <code>smsubnetprotection start</code>, and finally, <code>smsubnetprotection commit</code> from the master Subnet Manager. 4. Return the priority of the new Subnet Manager to its previous value.

Scenario	Impact and Actions
Secret M_Key values are mismatched.	<p>If you add a Subnet Manager with one set of secret M_Keys to a fabric with a different set of secret M_Keys, the added Subnet Manager is not recognized.</p> <p>Before introducing the new Subnet Manager to the fabric:</p> <ol style="list-style-type: none"> 1. Update the fabric's master Subnet Manager's list of known secret M_Keys to include the secret M_Keys already configured for the new Subnet Manager, with the <code>smsubnetprotection add</code> command. 2. Do not change the current secret M_Key. 3. Disable the new Subnet Manager. 4. Set the new Subnet Manager priority to the lowest. 5. Update the <code>smnodes</code> file with the <code>smnodes</code> command. 6. Enable the new Subnet Manager. <p>After introducing the new Subnet Manager to the fabric:</p> <ol style="list-style-type: none"> 1. Update the fabric configuration with the <code>fdconfig</code> command. 2. Update the fabric mapping with the <code>createfabric</code> command. 3. Perform a <code>smpartition start</code>, then <code>smpartition commit</code>, then <code>smsubnetprotection start</code>, and finally, <code>smsubnetprotection commit</code> from the master Subnet Manager. 4. Set the secret M_key policy as desired from the master Subnet Manager. 5. Return the priority of the new Subnet Manager to its previous value.
Merging two or more subnets into one fabric.	<p>If each subnet is configured with different secret M_Key policies, then the subnets will not “see” each other and will act independently.</p> <p>If one subnet is without a secret M_Key policy, then the subnet with a secret M_Key policy controls the subnet without.</p> <p>If each subnet is configured with identical secret M_Key policies, they merge into a single subnet.</p> <p>Before physically merging the subnets:</p> <ol style="list-style-type: none"> 1. Set the priority of one master Subnet Manager to lower than the other. 2. Configure the soon-to-be master Subnet Manager of the combined subnets with partition information from both subnets with the <code>smpartition</code> command. 3. Update the soon-to-be master Subnet Manager's list of known secret M_Keys to include the secret M_Keys already configured for the other subnet, with the <code>smsubnetprotection add</code> command. 4. Do not change the current secret M_Key. <p>After physically merging the subnets:</p> <ol style="list-style-type: none"> 1. Update the <code>smnode</code> files for all <code>smnodes</code> of both subnets with the <code>smnodes</code> command. 2. Configure both subnets with the new fabric configuration with the <code>fdconfig</code> command. 3. Correlate both subnets to the new fabric mapping with the <code>createfabric</code> command. 4. Perform a <code>smpartition start</code>, then <code>smpartition commit</code>, then <code>smsubnetprotection start</code>, and finally, <code>smsubnetprotection commit</code> from the now master Subnet Manager. 5. Set the secret M_key policy as desired from the master Subnet Manager.

This table describes each of the columns of the output of the `smsubnetprotection` command.

Column Heading	Description
Mkey	Secret M_Keys provided by the user for trusted devices.
Untrusted Mkey	Secret M_Keys generated from user input, for untrusted devices.
Smkey	SMKeys are used in communication between the Subnet Managers.
Attribute	The attribute of the M_Key. <ul style="list-style-type: none"> ■ C – The current secret M_Key. ■ S – The standby secret M_Key about to become current.

The `smsubnetprotection` command is available from the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI interface.

Options

This table describes the options to the `smsubnetprotection` command and their purposes.

Option	Purpose
-force	Specifies the action to bypass the partition daemon check and perform the operation even though not all smnodes are available or communicating with the management network. The -force option is synonymous with the -override-unavailable-smnodes option.
-addonly	Specifies that the session is only to add secret M_Keys to the configuration.
-deleteonly	Specifies that the session is only to delete secret M_Keys from the configuration.
-override-inconsistent-partition-configurations	Specifies that the check for partition consistency across smnodes is bypassed. Before updating the secret M_Key configuration, all smnodes to use that secret M_Key must have the same partition configuration. If not, the user is warned of such situation during the secret M_Key configuration update. This option overrides the check, and permits the secret M_Key configuration to be used, regardless of the consequences. Use of this option compromises the integrity of your fabric.
-override-unavailable-smnodes	Specifies the action to bypass the partition daemon check and perform the operation even though not all smnodes are available or communicating with the management network. The -override-unavailable-smnodes option is synonymous with the -force option.
tid	Specifies the transaction ID. The transaction ID adds an additional layer of security to the <code>smsubnetprotection</code> command. The identifier is a 32-bit unsigned integer, returned when the secret M_Key configuration is created (<code>smsubnetprotection start</code>) with the <code>tid</code> option. This identifier is then required for all subsequent actions to the secret M_Key configuration. Use of the transaction ID mediates changes to the secret M_Key configuration by multiple users.

Example

This example shows how to display the active secret M_Keys with the `smsubnetprotection` command.

```
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->
```

VNIC Limits for ETH Ports

The gateway documentation fails to state that there are a maximum of 1000 VNICS per each ETH port (0A-ETH-1, 0A-ETH-2, and so on).

VNIC Limits for the `createlag` Command

In the gateway documentation, the `-vniclimit` option for the `createlag` command is incompletely defined. What follows is the corrected description.

The `-vniclimit` option sets the maximum number of VNICS to external port associations per LAG (1 = 1000 VNICS, 4 = 4000 VNICS). If the limit is set at 4, then all four ports of the connector group are part of the LAG. The connector group (0A-ETH or 1A-ETH) is specified in the `createlag` command, and not the individual connectors.

Command Output Description Incorrect

In the [Sun Network QDR InfiniBand Gateway Switch Command Reference for Firmware Version 2.1](#), the PKEY column description is incorrect in the descriptions for the `showvlan`

and `showvnic` commands. The PKEY column lists the P_Key and the P_Key membership bit combined. The P_Key membership bit is the most significant bit of the two-byte value assigned to the respective P_Key.

To simplify, subtract 0x8000 from the value seen in the PKEY column to determine the actual P_Key value.

For example:

PKEY column value ffff - 0x8000 = P_Key 0x7fff

Understanding Switch Chip Port to Connector Routing

The [Sun Network QDR InfiniBand Gateway Switch Administration Guide for Firmware Version 2.1](#) has provided erroneous tables mapping switch chip ports to connectors, and vice-versa. These two tables provide corrected information.

Note - The `dcspport` command, the `listlinkup` command, and the Fabric Monitor correctly identify the switch chip port to connector routing.

Switch Chip Port to QSFP Connectors and Link LED Routes

Port	Connector	Port	Connector	Port	Connector	Port	Connector
1	1A-ETH-1(P1)	10	13B	19	0B	28	4A
	1A-ETH-2(P2)						
2	1A-ETH-3(P3)	11	12B	20	0A	29	5B
	1A-ETH-4(P4)						
3	0A-ETH-1(P1)	12	11A	21	1B	30	5A
	0A-ETH-2(P2)						
4	0A-ETH-3(P3)	13	9B	22	1A	31	8A
	0A-ETH-4(P4)						
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

Port	Connector	Port	Connector	Port	Connector	Port	Connector
9	13A	18	12A	27	4B	36	6B

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
2	24	23	11	12	17
3	26	25	12	18	11
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)	

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.

Features and Functionality Documented

The features and functionality described in the gateway documentation has been updated to reflect the 2.1.2-2 and later versions of the firmware. Upgrading your gateway firmware to the most current version helps increase gateway functionality. See [“Upgrading the Gateway Firmware” on page 37](#).

Correct Power Consumption

In the Electrical Specifications table of the [Sun Network QDR InfiniBand Gateway Switch Installation Guide for Firmware Version 2.1](#), the table incorrectly states the power requirement of 550 Watts. The correct value is 320 Watts maximum.

Node Description Format

The format of the node description in the [Sun Network QDR InfiniBand Gateway Switch Installation Guide for Firmware Version 2.1](#) is incorrect. The correct format for the switch chip and BridgeX chip are respectively:

```
SUN IB QDR GW switch hostname mc_IP
```

```
SUN IB QDR GW switch hostname mc_IP Bridge bridge_number
```

where:

- *hostname* is the hostname of the gateway and a maximum of 17 characters. Any additional characters are truncated.
- *mc_IP* is the IP address of the management controller in the gateway.
- *bridge_number* is the number of the BridgeX chip (0 or 1).

For example:

```
SUN IB QDR GW switch IBGateway-03 123.45.67.89
```

```
SUN IB QDR GW switch IBGateway-03 123.45.67.89 Bridge 0
```

Battery Service Sequence

In the [Sun Datacenter InfiniBand Switch 36 Service Manual for Firmware Version 2.1](#), the summary of the tasks to service the battery is in the incorrect sequence. This table corrects the sequence.

Step	Description	Sections
1.	Determine if the battery is faulty.	“Determine If the Battery is Faulty”.
2.	Power off both power supplies.	“Power Off a Power Supply”.
3.	Remove all data cables.	“Remove a Data Cable”.
4.	Remove the gateway from the rack.	“Remove the Gateway From the Rack”.

Step	Description	Sections
5.	Replace the battery.	“Replace the Battery” .
6.	Install the gateway in the rack.	“Installing the Gateway ” in the Sun Network QDR InfiniBand Gateway Switch Installation Guide for Firmware Version 2.1 . Note - Do not power on the gateway.
7.	Install all Data cables.	“Install a Data Cable”.
8.	Power on both power supplies.	“Power On a Power Supply”.

Note - You must completely power off the gateway before disconnecting the data cables. Similarly, you must attach all data cables before powering on the gateway.

▼ Replace a Gateway

This procedure briefly outlines the steps to replace a gateway within an InfiniBand fabric. For more information, refer to the following books:

- [Sun Network QDR InfiniBand Gateway Switch Service Manual for Firmware Version 2.1](#)
- [Sun Network QDR InfiniBand Gateway Switch Installation Guide for Firmware Version 2.1](#)
- [Sun Network QDR InfiniBand Gateway Switch Administration Guide for Firmware Version 2.1](#)

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 hostname, IP address, and lowest priority as the gateway removed.**
7. **Restore the gateway configuration previously backed up.**

8. **Completely power off the replacement gateway.**
9. **Connect the replacement gateway to the fabric's management network and the fabric.**
10. **Power on the replacement gateway.**
Wait for any negotiation and propagation to complete.
11. **Perform a `smpartition start` and `smpartition commit` from the master Subnet Manager.**
Wait for the partition configuration to propagate to the replacement gateway.
12. **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.
13. **Setup the `smnodes` file 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.**
14. **(Optional) Set the priority of the replacement gateway to the original priority 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 with the highest Priority value first, and then configure any remaining Subnet Managers.
- **Subnet Manager Prefix** – All Subnet Managers must use the same prefix. Configure the standby Subnet Managers first, and then configure the master Subnet Manager.
- **Subnet Manager Controlled Handover** – All Subnet Managers must use the same configuration for controlled handover. Configure the standby Subnet Managers first, and then configure the master Subnet Manager.
- **Subnet Manager Routing Algorithm** – All Subnet Managers must use the same routing algorithm. Configure the standby Subnet Managers first, and then configure the master Subnet Manager.

Replication Password Information Incomplete

The replication password is an eight alphanumeric character string used for encrypting communications between Subnet Managers nodes. All Subnet Managers must be configured with the same string, and you set the replication password using the `smsubnetprotection` command on the management controller of each Subnet Manager node. Because of the password's secure nature, is not readable. Therefore, you must remember the password for when adding Subnet Manager nodes in the future. Should you forget the replication password, you must reconfigure all Subnet Manager nodes with a new replication password.

Temperature Sensor Thresholds Incorrect

Some of the temperature sensor thresholds described in the documentation are incorrect. This table provides correct threshold values.

Sensor	Upper Critical Threshold	Upper Nonrecoverable Threshold
/SYS/MB/T_SP	80°C	85°C
/SYS/MB/T_BACK	60°C	65°C
/SYS/MB/T_FRONT	60°C	65°C
/SYS/MB/T_I4A	90°C	95°C
/SYS/MB/T_B0	90°C	95°C
/SYS/MB/T_B1	90°C	95°C

README File Is Incomplete

The README file accompanying the 2.1.2-2 firmware available for download from My Oracle Support is lacking this list of bugs fixed in the 2.1.2-2 firmware. The README file accompanying the 2.1.3-4 through 2.1.9-1 firmware is correct and inclusive of these bugs.

- 15787181 – Segmentation fault in `osm_mgrp_delete_port()`
- 15809823 – Subnet Manager should have a lease time of 60 seconds
- 15810464 – `getmaster` CLI command should show local SM state
- 15957702 – Switch running master Subnet Manager should trigger handover at reboot
- 15805689 – Default high error rate threshold should be increased to 3456 symbol errors per 24h
- 15815591 – BridgeX and I4 firmware versions should be checked by `fwverify`

- 15810720 – Do not allow to create partition without specifying P_Key

Configuring Secure Fabric Management

In the [Sun Network QDR InfiniBand Gateway Switch Administration Guide for Firmware Version 2.1](#) and the [Sun Network QDR InfiniBand Gateway Switch Command Reference for Firmware Version 2.1](#), text might instruct you to configure secret M_Key functionality before enabling the secret M_Keys. This is incorrect. You must first enable secret M_Key functionality before configuring the secret M_Keys.

▼ Acquire the BXOFED Software (Linux)

1. **Open a web browser on a host that will receive the BXOFED software.**
2. **Go to Oracle's My Oracle Support page.**
<http://support.oracle.com>
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 Product or Family (Advanced).**
The Patch Search window updates.
6. **In the Product Is field, type Sun Product Patches.**
7. **Click outside of the drop-down menu.**
8. **In the Release Is drop-down menu, scroll down to and select Mellanox 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 file, the Release Notes, and the Installation Guide for instructions 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 2.7.8130 or higher.

1. **Open a web browser on the host that will receive the ConnectX-2 firmware.**

2. **Go to Oracle's My Support page.**

<http://support.oracle.com>

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 Product or Family (Advanced).**
The Patch Search window updates.
6. **In the Product Is field, type the name of your HCA.**
Possible products are suggested as you type.
7. **In the Product Is drop-down menu, select your HCA.**
For example, Sun IB Dual Port 4x QDR PCIe ExpModule HCA.
8. **In the Release Is drop-down menu, select the appropriate firmware version.**
For example, FW25408 v2.7.8130.
9. **Click outside of the drop-down menu.**
10. **In the Platform Is drop-down menu, select the Oracle Solaris appropriate for your host.**
For example, Oracle Solaris on x86-64 (64-bit).
11. **Click outside of the drop-down menu.**
12. **Click Search.**
The Patch Search window expands with the search results.
13. **In the Patch Name column, click the respective patch number link.**
For example, 12610332. The Patch Search window reformats.
14. **Click Read Me to display the README file.**
15. **Click Download.**
The File Download window opens.
16. **Click the *filename.zip* link to initiate the download.**
For example, p12610332__Solaris86-64.zip.
17. **Indicate where the file should be saved.**
The file is downloaded and saved.
18. **In your receiving directory, decompress the *filename.zip* file.**

The ConnectX-2 firmware is in the `fw-ConnectX2-rel-version-part_number.bin` file. For example, `fw-ConnectX2-rel-2_7_8130-375-3697-01.bin`.

19. **Refer to your HCA documentation for instructions on how to upgrade the ConnectX-2 firmware.**

Upgrading the Gateway Firmware

In the [Oracle Integrated Lights Out Manager \(ILOM\) 3.0 Supplement for the Sun Network QDR InfiniBand Gateway Switch Firmware Version 2.1](#), firmware version numbers are provided as `x.y`, `x.y.z`, and `x.y.z-w`. Currently, these numbers are 2.1, 2.1.9, and 2.1.9-1 respectively. The following two procedures describe how to acquire and upgrade the firmware through the ILOM CLI.

Note - The gateway must have at least version 2.0 firmware installed before the following two procedures can be performed. Refer to the [Oracle Integrated Lights Out Manager \(ILOM\) 3.0 Supplement for the Sun Network QDR InfiniBand Gateway Switch Firmware Version 2.1](#) for instructions on installing the firmware.

▼ 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, 2.1.9.
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, 25467807. 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, p25467807_219_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_2.1.9-1/SUN_DCS_GW/sundcs_gw_repository_2.1.9_1.pkg file.
The readme_SUN_DCS_GW_2.1.9-1.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 39.](#)

▼ Upgrade the Gateway Firmware (CLI)

Note - Before upgrading or downgrading the gateway firmware, see [“Firmware Update Guidelines” on page 14](#).

Note - If you are going to downgrade the firmware to a version earlier than 2.1, you must disable secret M_Keys. Refer to the [Sun Network QDR InfiniBand Gateway Switch Administration Guide for Firmware Version 2.1](#) for instructions on disabling secret M_Key functionality.

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 the [Sun Network QDR InfiniBand Gateway Switch Administration Guide for Firmware Version 2.1](#) for instructions on removing partitions for firmware downgrade.

1. Consider your first step:

- If you are upgrading or downgrading from firmware version 2.0 or newer, go to [Step 2](#).
- If you are upgrading or downgrading from a firmware version earlier than 2.0, go to [Step 5](#).

2. 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`.

3. If the Subnet Manager is running on the management controller, disable the Subnet Manager.

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

4. Go to [Step 12](#).

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

```
% 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.

6. If the Subnet Manager is running on the management controller, disable the Subnet Manager.

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

7. 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.

8. 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.

9. 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, delete files from the /var/log directory.

10. 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.

11. 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.

12. 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. If you are upgrading from firmware 2.1.2-2 or newer, the TFTP protocol is 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_2.1.9_1.pkg
Downloading firmware image. This will take a few minutes.
```

Note - If you are experiencing version number contention, you can use the -force option to disable version number checking, and force the upgrade.

An error is displayed and the firmware upgrade process exits.

```
Error: Exiting Firmware update. Please see file /tmp/fw_upgrade_pre_check.err for details.
```

```
Firmware image update failed.  
load: Command Failed  
->
```

This is normal. An installer within the firmware package set prerequisite values.

Note - Do not restart or reboot the management controller or the gateway. Proceed to [Step 13](#).

13. Reload the firmware.

Retype the command as you did in [Step 12](#).

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.

Please make sure to upgrade in the order specified in the user documentation.
Are you sure you want to load the specified file (y/n)?

14. Answer y to the prompt to begin the upgrade.

Setting up environment for firmware upgrade. This will take 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 2.1.6-2 to 2.1.9-1:

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.

->

15. 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 (Thu Mar 16 20:10:49 2017):
The system is going down for reboot NOW!
-> Connection to gateway_name closed by remote host.
Connection to gateway_name closed.
```

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:

Please run the "fwverify" CLI command to verify the new image.
This message will be cleared on next reboot.

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

```
% ssh -l ilom-admin gateway_name
ilom-admin@gateway_name's password: password
-> show /SYS/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->version
SUN DCS gw version: 2.1.9-1
Build time: Jan 12 2017 09:30:46
FPGA version: 0x34
SP board info:
Manufacturing Date: 2012.01.10
Serial Number: "NCD8F0072"
Hardware Revision: 0x0007
Firmware Revision: 0x0102
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, 2.1.9-1.

17. Verify the firmware integrity.

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

```
FW Version: 7.4.3002 OK
PSID: SUNX2826_I40_002 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->
```

