Sun Blade 6000 Virtualized 40 GbE Network Express Module

User's Guide

Part No.: E21078-05
January 2015

Copyright © 2011, 2015 , Oracle and/or its affiliates. All rights reserved.

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

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

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

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

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications 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 on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2011, 2015 Oracle et/ou ses affiliés. Tous droits réservés.

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

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

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

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

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

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

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

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

Please
Recycle

Adobe PostScript

Contents

Using This Documentation xi

Understanding the NEM 1

Task Overview 1

Product Features 2

Understanding the Components 4

SAS Connections 4

Ethernet Connections 5

10GbE NIC Virtualization 8

Fabric Express Modules 8

Processor 9

Identifying Ports and LEDs 9

Connector Locations 10

External NEM Ports 11

Gigabit Ethernet Ports 11

SFP+ Ports 11

RJ-45 Ethernet Connector Port LEDs 13

NEM Front Panel and LEDs 14

Performing Hot-plug Insertion and Removal 17

Performing PCIe Hot-plug 17

Understanding PCIe Hot-plug 17

ATTN Button 18

Surprise Removal 19

Forcing NEM Graceful Removal 19

- Prepare the PCIe Interface for Hot-plug Removal 20

Performing SAS Interface Hot-plug 21

- Prepare SAS Interfaces for Hot-plug Removal 22

Restoring SAS Configurations After NEM Installation 22

Saving the Zoning Configuration to a Backup File 25

- Save the Zoning Configuration to a Backup File (Web) 25
- Save the Zoning Configuration to a Backup File (CLI) 26

Recovering Zoning Configurations 27

- Recover Zoning Configurations (Web) 27
- Recover Zoning Configurations (CLI) 28

Installing or Replacing a NEM 31

Before Installing a NEM 31

Verifying the CMM 32

- Determine the CMM Part Number Using ILOM 3.0x Web Interface 32
- Determine the CMM Part Number Using ILOM 3.1x Web Interface 32
- Determine the CMM Part Number Using `ipmitool` 33
- Find the CMM Part Number on the CMM Label 33

- Install a NEM 34

Verifying NEM Installation 36

- Verify Installation (CMM ILOM Web) 37
- Verify Installation (CMM ILOM CLI) 37

- Remove a NEM 40
- Replace a NEM 43

Installing and Removing SFP+ Optical Transceiver Modules 46

- Install an SFP+ Optical Transceiver Module 47
- Remove an Optical Transceiver SFP+ Module 48

SFP+ Connector Cabling Options 49

Enabling Jumbo Frames 51

Configuring Jumbo Frames (Solaris) 51

Jumbo Frames Overview 51

Jumbo Frames Configurations and Driver Statistics 52

- Enable Jumbo Frames (Solaris) 54
- Check Layer 2 Configuration 55
- Check Layer 3 Configuration 55

Configuring Jumbo Frames (Linux) 56

- Temporarily Configure Jumbo Frames (Linux) 56
- Permanently Enable Jumbo Frames (Linux) 57

- Enable Jumbo Frames (Windows) 59

Configuring Jumbo Frames (VMware ESX) 59

- Create a Jumbo Frames-Enabled vSwitch (VMWare) 60
- Create a Jumbo Frames-Enabled VMkernel Interface (VMWare 60

Configuring Static Link Aggregation 63

Example (Oracle Solaris 11 OS) 63

Example (Oracle Solaris 10 OS) 64

Example (Linux OS) 64

Booting an x86 Server Over the Network 67

Booting the NEM on an x86 Blade Server Over the Network 67

- Set Up the BIOS for Booting From the NEM 67
- Use the F12 Key to Initiate the Network Boot 68
- Use the F8 Key to Boot Off the PXE Server 69

- Boot a SPARC Blade Over the Network 69

Installing Drivers on Oracle Solaris OS SPARC or x86 Platform 75

- Download the `sxge` Driver 75
- Configure the Network Host Files 76
- Configure the `sxge` Device Driver Parameters 78

Installing Drivers on a Linux Platform 81

Downloading, Installing, and Removing Drivers 81

- Download and Install the Drivers 82
- Remove the Drivers 88

Configuring the Network Interface 90

Temporary Network Interface Configuration 90

Permanent Network Interface Configuration 91

- Bring the Interface Online Manually 91
- Configure the Network Interface File Automatically for Red Hat Linux 93
- Configure the Network Interface Automatically for SUSE Linux 95

Device Checking and Testing 97

- Configure the Syslog Parameter 100

Installing Drivers on a Windows Platform 103

Installing Drivers on a Windows Platform 103

Installing and Uninstalling the Network Controller 103

- Install the Sun Blade 6000 40 GbE Network Controller 104
- Uninstall the Sun Blade 6000 40 GbE Network Controller 105

Installing and Uninstalling the Sun Blade 6000 40 GbE VLAN Driver (Optional) 106

- Install the VLAN Driver 106
- Uninstall the VLAN Driver 107

Adding and Removing VLANs 108

- Add a VLAN 108
- Remove a VLAN 109

▼ Install the Enclosure Device on a Windows Server 2008 System 110

Dynamic and NIC Teaming 111

Installing Drivers on a VMware ESX Server Platform 113

Installing ESX Server Drivers 113

▼ Install a NEM Driver for a New ESX4.0 Installation 114

▼ Install a NEM Driver on an Existing ESX 4.0 Server 114

▼ Install a NEM Driver on an ESX/ESXi Server 115

▼ Configure the NEM Network Adapters 116

▼ Uninstalling the Drivers 116

Configuring Blade and Port Parameters and Managing the NEM 117

Configuring and Viewing Parameters 117

Defaults and Options 118

Setting Blade Parameters 119

Bandwidth 119

PCI Generation 120

MAC Address 121

Promiscuous 122

Receiver Side Scaling 122

SR-IOV 123

VLAN 124

Setting Port Parameters 125

Speed 125

Link Status 125

Saving User Configurations or Resetting to Defaults 126

Viewing Parameter Settings, Firmware Versions, and Configuration File 127

Showing Blade Parameters 128

Bandwidth 128

Gen 128
MAC Address 129
Promiscuous 130
RSS 130
SR-IOV 131
VLAN 131
Showing Port Parameters 132
Showing Firmware 133
Showing Configuration File 133
Configuring NEM for Oracle VM Para-Virtualization 134
▼ Configure Para-Virtualization Without Live Migration 134
▼ Configure Para-Virtualization With Live Migration 137
Managing the NEM Using CMM ILOM 139
ILOM Documentation and Updates 139
▼ Start NEM CLI from CMM ILOM CLI 140

Updating the NEM Firmware 143

Configuring Hot Plug Service 143
▼ Verify and Enable Hot Plug (Oracle Solaris OS) 144
▼ Verify and Enable Hot Plug (Linux OS) 144
Disabling Interfaces 145
Updating Firmware (Web) 145
▼ Obtain and Apply the Newest Firmware 146
▼ Reset the NEM (Web) 147
Updating the Firmware (CLI) 147
▼ Update NEM ILOM Firmware (CLI) 148
▼ Reset the NEM (CLI) 149

Glossary 153

Index 159

Using This Documentation

This document contains information about Oracle's Sun Blade 6000 Virtualized 40 GbE Network Express Module, including instructions for installing and replacing the product. Also, this document provides procedures for installing and configuring drivers on compatible operating systems.

- "Related Documentation" on page xi
- "Feedback" on page xii
- "Support and Accessibility" on page xii

Related Documentation

Documentation	Links
All Oracle products	http://www.oracle.com/documentation
Sun Blade 6000 Virtualized 40 GbE Network Express Module	http://docs.oracle.com/cd/E21077_01/
Sun Blade 6000 Modular System	http://docs.oracle.com/cd/E19938-01/
Oracle Integrated Lights Out Manager 3.0	http://docs.oracle.com/cd/E19860-01/

Documentation	Links
SAS1/SAS2 Configuration	http://docs.oracle.com/cd/E22513_01/
Oracle VM Server for x86	http://docs.oracle.com/cd/E15458_01/index.htm
Vendor VMware Publications	http://www.vmware.com/support/pubs/

Feedback

Provide feedback on this documentation at:

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

Support and Accessibility

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

Understanding the NEM

Oracle's Sun Blade 6000 Virtualized 40 GbE Network Express Module (NEM) is a multi-purpose connectivity module for the Sun Blade 6000 modular system. The NEM provides network and storage connectivity between the blades in an Oracle Sun Blade 6000 modular system chassis and external devices. The NEM supports connection to external devices through 10 GbE SFP+ ports and 10/100/1000 TPE ports. Also, the NEM supports Sun ASIC Dual 10GbE NIC virtualization.

- "Task Overview" on page 1
- "Product Features" on page 2
- "Understanding the Components" on page 4
- "Identifying Ports and LEDs" on page 9

Task Overview

Description	Link
Learn about the features of the NEM.	"Product Features" on page 2
Install or replace the NEM.	"Installing or Replacing a NEM" on page 31
Boot over the Network with the NEM.	"Booting an x86 Server Over the Network" on page 67
Install and configure Oracle Solaris OS drivers.	"Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
Install and configure Linux drivers.	"Installing Drivers on a Linux Platform" on page 81
Install and configure Windows drivers.	"Installing Drivers on a Windows Platform" on page 103
Install and configure ESX drivers.	"Installing Drivers on a VMware ESX Server Platform" on page 113

Description	Link
Understand and use the CMM CLI	"Configuring Blade and Port Parameters and Managing the NEM" on page 117
Use ILOM with the NEM.	"Connecting to NEM ILOM" on page 133

Product Features

Feature	Description
Sun ASIC	• 40GbE or 10GbE NIC virtualization • 32MB Flash memory • Internal ARM9S Processor • Sun RGMII interface through an Ethernet PHY and switch to CMM
LSI SAS2 x36 Expander	• Internal ARM9S Processor • 4MB SRAM, 8MB flash memory • Management in-band and via I^2C • Expander MII interface through an Ethernet PHY and switch to CMM
FPGA	• Altera EP3C • I^2C Fabric • Two I^2C to I^2C bridges • NEM control and status registers (NEM Regs) • PCA9553 compatible LED drivers • Virtual LED controller • PCA9555 compatible GPIO • HPL to blades • Hot-plug state machine (PCMM) • SCPGM engine for downloading FPGA configuration image from the CMM • Logic for power control, module reset, and miscellaneous functions
Management Interfaces	• 10/100BASE-T Ethernet port to CMM • I^2C to CMM • HPL to server module (blade)

Feature	Description
Ports	• Ten pass through 1GbE RJ-45 ports (supports 10/100/1000BASE-T) • Four SFP+ 10GbE ports configured to: 40GbE (all SFP+ active) 2x10GbE (SFP+ ports 0 and 1 active) 1x1GbE (SFP+ port 0 active) NOTE: At 1GbE, SFP+ to RJ-45 transceiver can be used in port 0 to support 1000Base-T (RJ-45).
Server Module Interfaces (per Blade)	• One x8 PCIe Gen2 interface to Sun ASIC • Two SAS2 interfaces to SAS2 Expanders • One 10/100/1000BASE-T Ethernet (RJ-45) port • One HPL to FPGA
Updates	All software and embedded firmware can be field upgraded
Indicators/Controls	• RJ-45 LEDs: Link/Activity and Speed • SFP+ LEDs: Link/Activity and Speed • SIS LEDs • LOC LED and push button • ATTN LED and button
Health	• Voltage monitoring • Temperature monitoring • Fault detection and reporting
Power Supplies	• 3.3V AUX from chassis midplane • 12V from chassis midplane • Other voltages generated on board

Feature	Description
Environmental	Cooling: top-to-back forced air Commercial Operating (C10 Chassis): • Humidity: 10 to 90% non-condensing • Temperature: 0 to 35°C (32 to 95°F) • Altitude: 0-10,000 ft (3048 meters) NEBS Operating (N6000 Chassis): • Humidity: 5 to 85% non-condensing (90% for short-term NEBS conditions) • Temperature normal operating: 0 to 45°C (32 to 113°F) • Temperature short-term NEBS emergency: 5 to 55°C (41 to 131°F • Altitude: 0 to 5,900 ft (1800 meters) at 45°C ambient (113F) and 0 to 13,000 (3960 meters) at 35°C ambient (95F) Non-operating: • Humidity: 5 to 95% non-condensing • Temperature: –40 to 70°C (-40 to 158°F) • Altitude: 39,370 ft (12000 meters)

Related Information

- "Understanding the Components" on page 4
- "Identifying Ports and LEDs" on page 9

Understanding the Components

- "SAS Connections" on page 4
- "Ethernet Connections" on page 5
- "10GbE NIC Virtualization" on page 8
- "Fabric Express Modules" on page 8
- "Processor" on page 9

SAS Connections

Each NEM provides 10 x2 LSI-SAS2 connections for server and storage blades. All NEM SAS interfaces support SAS (3 Gbps) and SAS2 (6 Gbps) data rates.

The Sun Blade Zone Manager running on the CMM provides a way to create and manage assignments. The zoning is programmed on the chassis CMM using the ILOM software.

Sun Blade Zone Manager only works with server and storage blades that are SAS2 enabled.

See the following documentation for more information on zoning SAS2 modules in the Sun Blade 6000 modular system:

- *Oracle Integrated Lights Out Manager (ILOM) CMM Administration Guide* for information on using Sun Blade Zone Manager.
- *Sun Blade Storage Module M2 Administration Guide* for information on zoning SAS2 disks in the chassis.
- *Sun Blade 6000 Network Express Module Product Notes* for information on supported server and storage modules.
- Server blade documentation for information on availability of SAS2 HBAs.

There are no external SAS connections on the NEM.

For the latest information about SAS1 and SAS2 compatibility, refer to the *SAS-1/SAS-2 Compatibility Upgrade Guide*, located at:

`http://docs.oracle.com/cd/E21077_01/`

Related Information

- "Ethernet Connections" on page 5
- "10GbE NIC Virtualization" on page 8
- "Fabric Express Modules" on page 8
- "Processor" on page 9
- "Product Features" on page 2
- "Identifying Ports and LEDs" on page 9

Ethernet Connections

The NEM provides the magnetics and RJ-45 connectors for the ten 10/100/1000 BASE-T Ethernet interfaces that come from Sun Blade 6000 server modules through the midplane.

Each server module has:

- One 10/100/1000BASE-T Ethernet pass through port.

- Two shared virtual ports. All ten blades are internally connected through each of the virtual ports.

The speed of the shared virtual ports can be configured as follows:

- Both 10Gb ports can connect externally to a switch and internally to all blades.
- One 40Gb port can be used for both external and internal connections, and the 10Gb port can be used for internal-connection only.

The NEM provides a 40GbE connection for each server module. Each server module appears to have its own 40GbE NIC through the virtualized NEM ASIC. The 40GbE NIC uses four physical SFP+ ports as one logical 40GbE port.

See "10GbE NIC Virtualization" on page 8 for further information on the NEM 10GbE virtualization.

The following illustration shows a schematic of the virtualized NEM Ethernet connections.

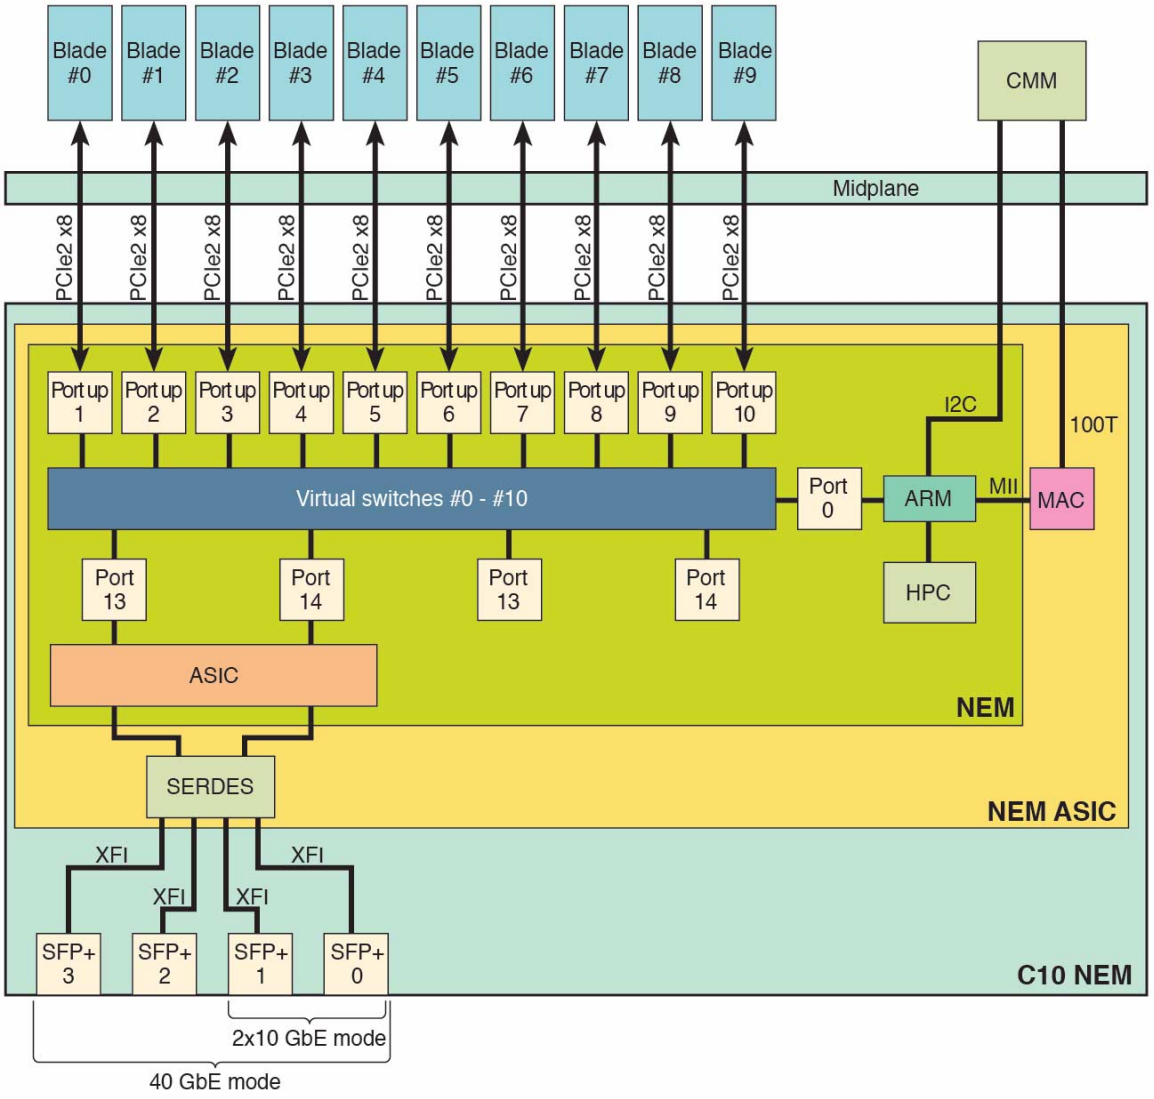

Related Information

- "SAS Connections" on page 4
- "10GbE NIC Virtualization" on page 8
- "Fabric Express Modules" on page 8
- "Processor" on page 9
- "Product Features" on page 2

- "Identifying Ports and LEDs" on page 9

10GbE NIC Virtualization

The virtualized ASIC allows up to 10 hosts to share each 10GbE network port, with a dedicated PCIe endpoint for each host, one per 10GbE network interface.

Shared I/O enables each server module to work as if a dedicated NIC were directly connecting the server module to the network. Each server module owns a virtual MAC that provides per-server-module statistics on Rx/Tx traffic. The MAC that interfaces to the 10 GbE network port is shared and hidden from the server modules.

Related Information

- "SAS Connections" on page 4
- "Ethernet Connections" on page 5
- "Fabric Express Modules" on page 8
- "Processor" on page 9
- "Product Features" on page 2
- "Identifying Ports and LEDs" on page 9

Fabric Express Modules

For server modules to access the 10 GbE functionality of the NEM, they must have a FEM installed in the server module. The *Sun Blade 6000 Virtualized 40 GbE Network Express Module Product Notes* lists which FEMs are required for each server module that is supported for the NEM.

Note – The Sun Blade Storage Module M2 does not have an Ethernet controller. Therefore the Ethernet ports on the NEM that correspond to slots containing a Sun Blade Storage Module M2 are not used.

Related Information

- "SAS Connections" on page 4
- "Ethernet Connections" on page 5
- "10GbE NIC Virtualization" on page 8
- "Processor" on page 9

- "Product Features" on page 2
- "Identifying Ports and LEDs" on page 9

Processor

The NEM has an internal ARM9S controller as its processor, which is responsible for control and management of the NEM ASICs. The NEM has a 10/100BASE-T Ethernet management port for connectivity to the Sun Blade 6000 CMM.

Refer to the "Connecting to NEM ILOM" on page 133 for more information on the ILOM server management application for the SP.

Note – The NEM does not have a chassis SP onboard that requires IP parameter configuration.

Related Information

- "SAS Connections" on page 4
- "Ethernet Connections" on page 5
- "10GbE NIC Virtualization" on page 8
- "Fabric Express Modules" on page 8
- "Product Features" on page 2
- "Identifying Ports and LEDs" on page 9

Identifying Ports and LEDs

- "Connector Locations" on page 10
- "External NEM Ports" on page 11
- "RJ-45 Ethernet Connector Port LEDs" on page 13
- "NEM Front Panel and LEDs" on page 14

Connector Locations

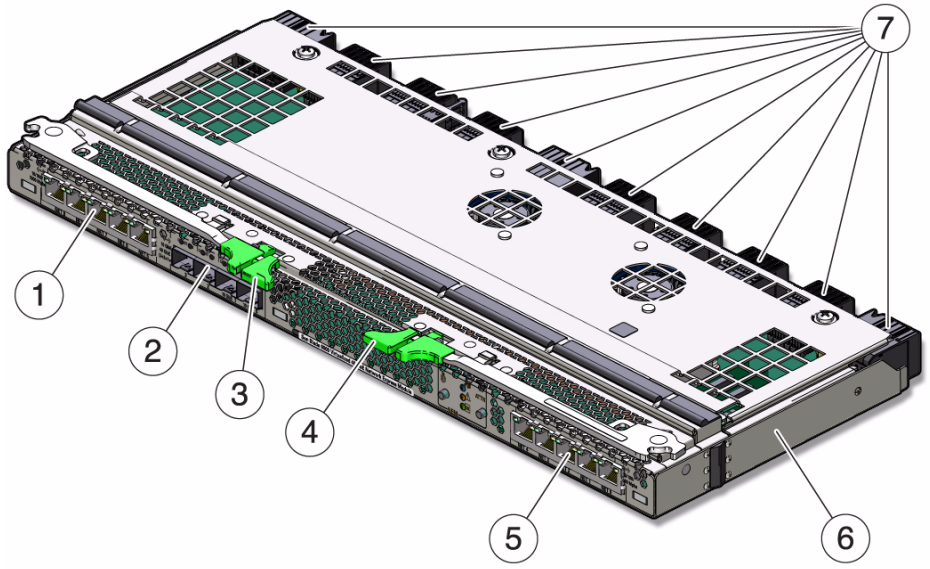

TABLE: NEM Components

Label	Description
1	10/100/1000Base-T Pass-Through Ethernet, RJ-45 connectors (5 total)
2	10GbE connectors (4 total) (Requires SFP+ module) See "SFP+ Ports" on page 11 and "SFP+ Connector Cabling Options" on page 49
3, 4	Ejector levers
5	10/100/1000Base-T Pass-Through Ethernet, RJ-45 connectors (5 total)
6	Chassis
7	Midplane connectors (10 total)

Related Information

- "External NEM Ports" on page 11
- "RJ-45 Ethernet Connector Port LEDs" on page 13
- "NEM Front Panel and LEDs" on page 14
- "SAS Connections" on page 4
- "Ethernet Connections" on page 5
- "10GbE NIC Virtualization" on page 8

- “Fabric Express Modules” on page 8
- “Processor” on page 9
- “Product Features” on page 2
- “Understanding the Components” on page 4

External NEM Ports

This topic describes the external ports on the NEM.

- “Gigabit Ethernet Ports” on page 11
- “SFP+ Ports” on page 11

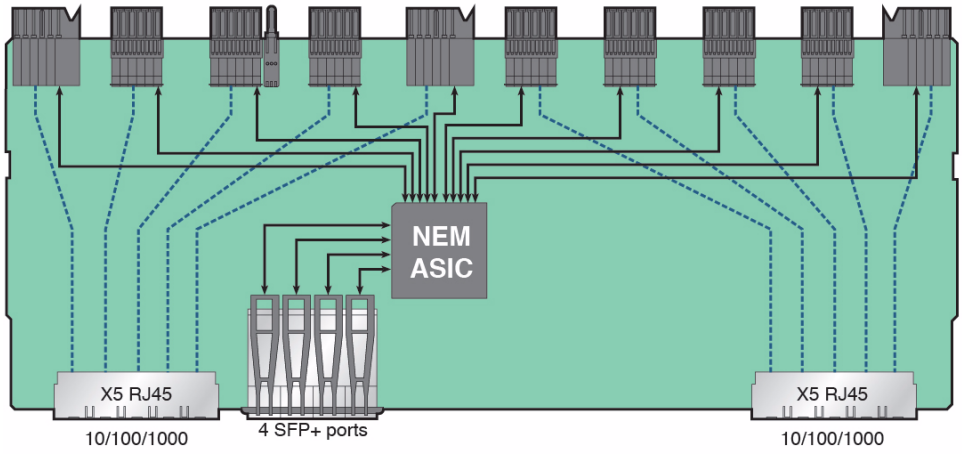

Gigabit Ethernet Ports

There are 10 pass-through RJ-45 Gigabit Ethernet ports on the NEM. The pass-through Gigabit Ethernet ports are strictly passive and isolated from other functional blocks, with no interaction.

SFP+ Ports

Four SFP+ ports provide virtualized 10GB connections to the server modules. SFP+ ports support optical module cables (short or long range) or twin-axial cables.

You can configure the four SFP+ ports as described in the following table.

Configuration	Ports and LED Behavior
40G	• All four SFP+ ports (P0 to P3) are active • Speed LED is located at upper right and the Link/Act LED is located at upper left • Speed LEDs turn on only when a transceiver is inserted into P0 • Link/Act LEDs turn on to indicate that the link is up when all four SFP+ ports have cables installed and their link partners are active. If any cables are removed, all four Link/Act LEDs will turn off.
10G	• P0 and P1 SFP+ ports are active • The Link/Act LED blinks off when there is activity on the link. • Link/Act LED at upper left turns on when a port has a cable installed and its link partner is active. • Speed LED is inactive (off)
1G	• P0 is active. Remaining SFP+ ports are inactive. • If using the SFP+ to RJ-45 transceiver adapter, ensure it is inserted in P0 • Link/Act LED at upper left turns on when the port has a cable installed and its link partner is active. • The Link/Act LED blinks off when there is activity on the link. • Speed LED is inactive (off)

Related Information

- "Connector Locations" on page 10
- "RJ-45 Ethernet Connector Port LEDs" on page 13
- "NEM Front Panel and LEDs" on page 14
- "SAS Connections" on page 4
- "Ethernet Connections" on page 5
- "10GbE NIC Virtualization" on page 8
- "Fabric Express Modules" on page 8
- "Processor" on page 9
- "Product Features" on page 2
- "Understanding the Components" on page 4
- "SFP+ Connector Cabling Options" on page 49
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46

RJ-45 Ethernet Connector Port LEDs

Each RJ-45 Ethernet port has two LEDs. The left LED is green and lights to show that a link has been established. It blinks on and off whenever there is network activity on that port.

On the RJ-45 connectors, the right LED is bi-color (amber and green) and indicates the speed of connection by the color it displays:

- 10 megabits per second, the right LED is off
- 100 megabits per second, the right LED displays one color
- 1000 megabits per second, the LED displays the other color

The green or amber color scheme varies from one server blade to another. The following table provides a chart for interpreting the link-speed relationships.

Sun Blade Server Model	10 MbE (Right LED)	100 MbE (Right LED)	1000 MbE (Right LED)
T3-1B	Off	Amber	Green
T4-1B	Off	Amber	Green
X6270 M2	Off	Amber	Green
X3-2B	Off	Amber	Green

When an Ethernet port is connected to an x86 server blade (server blades with a model number that starts with X) that has been put in WOL mode, the Link LED indicates when the system is in standby mode. It does this by blinking in a repeating, non-random pattern. It flashes ON for 0.1 second and OFF for 2.9 seconds. In standby mode, the system is functioning at a minimal level and is ready to resume full activity.

Note – SPARC-based server blades do not support the WOL mode. When an Ethernet port is connected to a SPARC server blade, the Link LED behaves as described in the table. SPARC-based blades are designated by a T in front of the server module number (for example, T6300).

Related Information

- "Connector Locations" on page 10
- "External NEM Ports" on page 11
- "NEM Front Panel and LEDs" on page 14
- "SAS Connections" on page 4

- “Ethernet Connections” on page 5
- “10GbE NIC Virtualization” on page 8
- “Fabric Express Modules” on page 8
- “Processor” on page 9
- “Product Features” on page 2
- “Understanding the Components” on page 4
- “SFP+ Connector Cabling Options” on page 49
- “Installing and Removing SFP+ Optical Transceiver Modules” on page 46

NEM Front Panel and LEDs

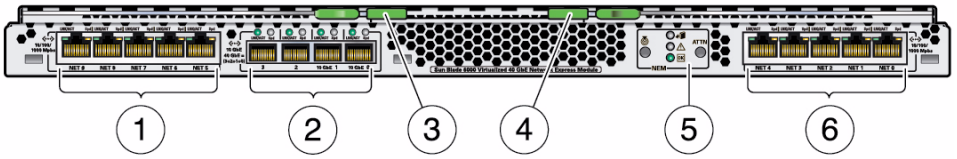

Label	Description
1	10/100/1000Base-T Pass-Through Ethernet, RJ-45 connectors (5 total)
2	10GbE connectors (4 total) (Requires SFP+ module) See “SFP+ Ports” on page 11 and “SFP+ Connector Cabling Options” on page 49.
3, 4	Ejector levers
5	SIS LEDs and ATTN and LOC LEDs and pushbuttons
6	10/100/1000Base-T Pass-Through Ethernet, RJ-45 connectors (5 total)

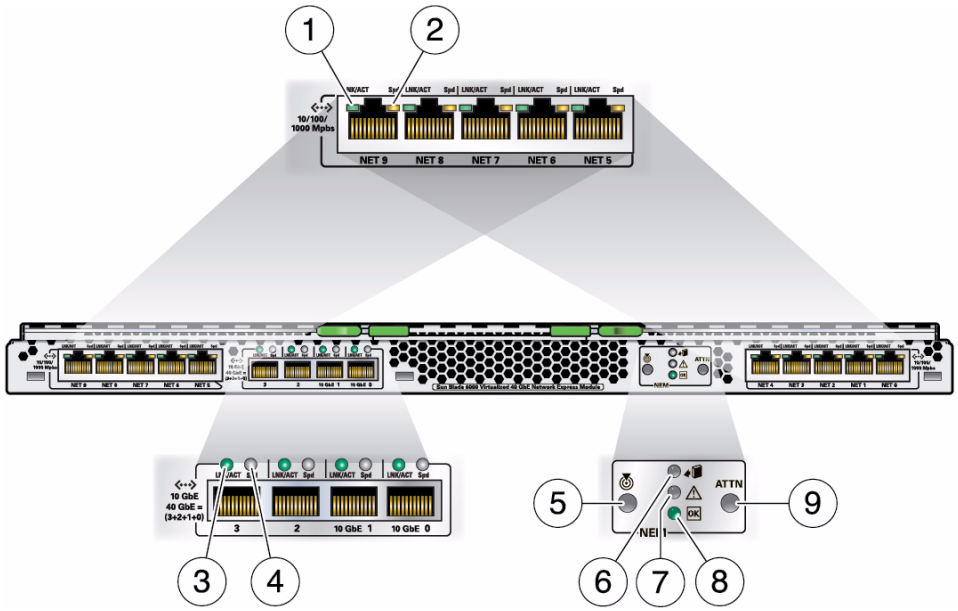

	LED/Button Name	Description
1	Activity/Link for 10/100/1000	See "RJ-45 Ethernet Connector Port LEDs" on page 13.
2	Speed LED for 10/100/1000	See "RJ-45 Ethernet Connector Port LEDs" on page 13.
3	SFP+ Activity/Link for 10GbE ports	The left green LED is a network activity indicator. It's on when the link is up and blinks when there is network activity. For variations in LED behavior, see table in "SFP+ Ports" on page 11.
4	SFP+ Link Speed for 40GbE ports	The right green LED is on for 40GbE and off for any other link configuration. For variations in LED behavior, see table in "SFP+ Ports" on page 11.
5	Locate button and LED (white)	Helps locate each NEM. • Push and release to make the Locate LED blink. • If the LED is blinking, push and release this button to stop the blinking. • Push and hold down the button for 5 seconds to initiate a "push-to-test" mode. This mode illuminates LEDs for 15 seconds. The LED can be enabled remotely.
6	Ready-to-Remove LED (blue)	Lights when NEM is ready to be removed. Press the Attention button and wait until the Ready-to-Remove LED lights before removing the NEM. Refer to "Performing Hot-plug Insertion and Removal" on page 17 for information on the behavior of the Ready-to-Remove LED during hot-plug insertion and removal.

	LED/Button Name	Description
7	Module Fault LED (amber)	Has two states: • Off: No fault. • On: An event has been acknowledged, and service action is required.
8	Module Activity (Power/OK) LED (green)	Has three states: • Off: Module is not configured or is offline. • On: Module is configured and online. • Blinking: Module is configuring or a hot-plug event is occurring. Refer to "Performing Hot-plug Insertion and Removal" on page 17 for information on the behavior of the Power/OK LED during hot-plug insertion and removal.
9	ATTN button	Prepares NEM and associated blades for NEM removal and alerts the blades when a new NEM has been inserted. Refer to "Performing Hot-plug Insertion and Removal" on page 17 for information on the behavior of the ATTN button during hot-plug insertion and removal.

Related Information

- "Connector Locations" on page 10
- "External NEM Ports" on page 11
- "RJ-45 Ethernet Connector Port LEDs" on page 13
- "SAS Connections" on page 4
- "Ethernet Connections" on page 5
- "10GbE NIC Virtualization" on page 8
- "Fabric Express Modules" on page 8
- "Processor" on page 9
- "Product Features" on page 2
- "Understanding the Components" on page 4
- "SFP+ Connector Cabling Options" on page 49
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46

Performing Hot-plug Insertion and Removal

These topics describe guidelines and procedures required to prepare PCIe and SAS interfaces for the NEM hot-plug actions. The NEM supports hot-plug insertion and removal. Hot-plug insertion and removal involves two separate interfaces to chassis components: PCIe interfaces and SAS interfaces.

These guidelines and procedures are not intended to be comprehensive removal and installation procedures. For step-by-step instructions on hot-plug installation or removal, see "Installing or Replacing a NEM" on page 31.

Note – Hot-plug operations are not supported for server modules with VMware ESX operating system installed.

- "Performing PCIe Hot-plug" on page 17
- "Performing SAS Interface Hot-plug" on page 21

Related Information

- "Installing or Replacing a NEM" on page 31
- "NEM Front Panel and LEDs" on page 14

Performing PCIe Hot-plug

- "Understanding PCIe Hot-plug" on page 17
- "Prepare the PCIe Interface for Hot-plug Removal" on page 20

Understanding PCIe Hot-plug

- "ATTN Button" on page 18

- “Surprise Removal” on page 19
- “Forcing NEM Graceful Removal” on page 19

ATTN Button

Press the ATTN button to initiate a hot-plug insertion or hot-plug removal. (See “NEM Front Panel and LEDs” on page 14 for the location of the ATTN button.) The green Power/OK LED slowly blinks while the action is taking place.

To abort the hot-remove or insert request, press the ATTN button a second time within five seconds after the green Power/OK LED begins blinking.

Caution – Hot-plug service has to be enabled on all blades in the chassis for hot-plug removal of a NEM.

An alternative to pushing the ATTN button is to use the Prepare to Remove and Return to Service commands in the ILOM.

The procedures for installing and removing the NEM describe the use of the ATTN button and ILOM hot-plug commands. See the following procedures:

- “Install a NEM” on page 34
- “Remove a NEM” on page 40
- “Replace a NEM” on page 43

Note – Successive pushes of the ATTN button (other than to cancel within five seconds) are not allowed. If a blade refuses to release the NEM and go offline after you press the ATTN button, force a graceful hot remove, as described in “Forcing NEM Graceful Removal” on page 19.

Related Information

- “Surprise Removal” on page 19
- “Forcing NEM Graceful Removal” on page 19
- “Prepare the PCIe Interface for Hot-plug Removal” on page 20
- “Prepare SAS Interfaces for Hot-plug Removal” on page 22
- “Install a NEM” on page 34

Surprise Removal

The NEM does not support surprise removal. Any attempt to remove the NEM without properly completing a graceful hot-plug removal, as described in "Remove a NEM" on page 40 and "Replace a NEM" on page 43 is considered a surprise removal. The blue Ready-to-Remove LED must be illuminated before you remove a NEM.

If you remove a host blade from the chassis, the operation of the NEM is not affected, and the NEM does not log the blade removal event. The remaining server blades in the chassis are unaffected by the removal of the blade.

Related Information

- "ATTN Button" on page 18
- "Forcing NEM Graceful Removal" on page 19
- "Prepare the PCIe Interface for Hot-plug Removal" on page 20
- "Prepare SAS Interfaces for Hot-plug Removal" on page 22
- "Install a NEM" on page 34
- "Remove a NEM" on page 40

Forcing NEM Graceful Removal

Note – Hot-plug service has to be enabled on all blades in the chassis for a hot-plug removal to properly initiate. If the Ready-to-Remove LED does not light, check the configuration of the blades, or power off all of them.

After you initiate a hot-plug removal, you must complete it while the blue Ready-to-Remove LED is illuminated. If the green Power/OK returns to solid green, this means that the NEM is no longer off-line.

If the NEM does not go offline, it is possible that some blades are ready for removal while others are not. Until the completion of the removal of all blades occurs, blades that were ready to be removed will no longer be detected by the NEM. It might be necessary to force blades that have denied immediate hot-plug removal to release their resources, so that you can proceed with the removal of the NEM.

Use the instructions in "Prepare the PCIe Interface for Hot-plug Removal" on page 20 to ensure that all blades have unplumbed their sxge interfaces, then continue with the ATTN button for hot removal.

Related Information

- "ATTN Button" on page 18
- "Surprise Removal" on page 19
- "Prepare the PCIe Interface for Hot-plug Removal" on page 20
- "Prepare SAS Interfaces for Hot-plug Removal" on page 22
- "Install a NEM" on page 34

▼ Prepare the PCIe Interface for Hot-plug Removal

- Perform this procedure before hot removing a NEM.
- For hosts in the chassis running Oracle Solaris 10 OS 09/10 or newer, the `hotplugd(1M)` daemon must be started before attempting to initiate a hot-plug action.
- For Linux OS, ensure the appropiate hotplug kernel module is already loaded.

1. **Ensure that all blades in the chassis are configured to allow hot-plug service.**
2. **Check to see if there is an Ethernet interface associated with the NEM on the server module. Perform the action associated with the appropriate operating system.**

   The following table lists the supported operating system types, the command or action required to view the Ethernet interfaces, and the name of the Ethernet device.

Operating System	Command or Action	Ethernet device name
Oracle Solaris 10 OS	**`ifconfig -a | grep sxge`**	`sxge` (associated with NEM MAC address)
Oracle Solaris 11 OS	**`dladm show-linkprop -p`** *mac-address*	Where *mac address* is the corresponding MAC address.
Linux, VMware	**`ifconfig -a | grep eth`**	`eth` (associated with NEM MAC address)
Windows	Click the Safely Remove Hardware icon on the lower right task bar.	Sun Blade 6000 10Gbe Networking Controller

   a. **If the output does not show any Ethernet interface associated with the NEM to be removed, proceed to one of the following:**
      - "Remove a NEM" on page 40
      - "Replace a NEM" on page 43

b. If the output shows an Ethernet interface associated with the NEM to be removed, go to Step 3.

3. Disable all interfaces between the blade(s) and the NEM.

Perform the action associated with the appropriate operating system.

The variable *x* in the commands in the table represents the Ethernet instance found in Step 2.

Operating System	Commands or Action
Oracle Solaris 10 OS	**ifconfig sxge***x* **down** **ifconfig sxge***x* **unplumb**
Oracle Solaris 11 OS	**ifconfig net***x* **down** **ifconfig net***x* **unplumb**
Linux	It is not necessary to bring the driver down to perform hot-plug. To bring the driver down: **ifconfig eth***x* **down**
Windows	1. Double-click the Safely Remove Hardware icon at the lower right of the Windows taskbar. The Safely Remove Hardware dialog appears. 2. Select the Sun Blade 6000 10Gbe Networking Controller and click Stop. 3. Click Close to finish. Now you can safely remove the NEM(s).

4. Repeat Step 2 to make sure that the Ethernet interface targeted is disabled.

5. Proceed with one of the following:

- "Remove a NEM" on page 40
- "Replace a NEM" on page 43

Performing SAS Interface Hot-plug

Removal and replacement activities can affect saved server-to-storage access configurations.

- "Prepare SAS Interfaces for Hot-plug Removal" on page 22
- "Restoring SAS Configurations After NEM Installation" on page 22
- "Saving the Zoning Configuration to a Backup File" on page 25
- "Recovering Zoning Configurations" on page 27

Note – The best way to ensure there is no interruption in service or accidental data loss when a NEM is removed is to have two NEMs installed in the chassis.

▼ Prepare SAS Interfaces for Hot-plug Removal

Note – The best way to ensure there is no interruption in service or accidental data loss when a NEM is removed is to have two NEMs installed in the chassis.

1. **Ensure that all blades in the chassis are configured to allow hot-plug service.**
2. **Back up the SAS zoning configurations. See** "Saving the Zoning Configuration to a Backup File" on page 25**.**
3. **For hosts in the chassis running Oracle Solaris 10 OS 09/10 or newer, start the** `hotplugd(1M)` **daemon.**
4. **Ensure that all I/O activity to the component has been stopped before proceeding with the hot-plug action.**

   If I/O activity is not stopped, interrupted service or data loss can occur.
5. **Ensure that the appropriate multi-path drivers are installed if your servers are not using a hardware RAID solution.**

Related Information

- "Restoring SAS Configurations After NEM Installation" on page 22
- "Saving the Zoning Configuration to a Backup File" on page 25
- "Recovering Zoning Configurations" on page 27

Restoring SAS Configurations After NEM Installation

After installing a NEM using hot-plug, you might need to perform some additional tasks to restore the chassis SAS configurations.

In the following table, find the action that corresponds with your hot-plug scenario and follow the user intervention instructions.

Action	Outcome	User Intervention
Insert a second NEM into an empty chassis slot.	The NEM is automatically registered with the CMM. The second NEM will obtain zoning configurations from the first NEM. Any future saved configurations remain unchanged if the NEM is removed and reinserted into the same slot. The external port configuration is not copied over from the existing NEM to the new NEM, but the NEM external SAS ports are not currently supported.	• None. When using a SAS2 REM configured for RAID, dual path support to storage will automatically be enabled. • When using a SAS2 REM not configured for RAID, dual path support might need to be configured at the OS level. Check your OS documentation for information on enabling SAS multipath support.
Remove and insert the same NEM into the same chassis slot.	Because the NEM is already registered with the CMM, the NEM zoning configurations will remain unchanged.	The host REM might identify the storage module virtual drives as a foreign configuration (MegaRAID), or inactive (Fusion-MPT IR). A foreign or inactive configuration can be recovered using the REM management software. • For MegaRAID, use the import feature described in the *LSI MegaRAID SAS Software User's Guide*. • For Fusion-MPT IR, use the activate feature described in the *LSI SAS2 Integrated RAID User's Guide*.
Remove and insert a NEM into a different slot in the same chassis.	Upon insertion into the chassis, the SAS zoning configurations on the NEM expanders are deleted. The NEM is automatically registered with the CMM. Any future saved configurations remain unchanged if the SAS2 NEM is removed and reinserted into the same slot.	• Restore zoning configurations from a backup of CMM configurations. See "Recovering Zoning Configurations" on page 27. —*Or*— • Use the CMM zoning utility to reassign storage module resources to SAS2 server blades.

Action	Outcome	User Intervention
Replace a failed NEM.	• If there are two NEMs, upon insertion in the chassis, the replacement NEM is automatically registered with the CMM. The replacement NEM obtains configurations from the first NEM, except for external port configurations. Any future saved configurations remain unchanged if the NEM is removed and inserted into the same slot. • If there is only one NEM, upon insertion into the chassis, the SAS zoning configurations of the replacement NEM expanders are deleted. The NEM is automatically registered with the CMM. Any future saved configurations remain unchanged if the NEM is removed and inserted into the same slot.	• If there are two NEMs, no further action is needed. • If there is only one NEM, restore zoning configurations from a backup of CMM configurations. See "Recovering Zoning Configurations" on page 27. *—Or—* Use the CMM zoning utility to reassign storage module resources to SAS2 server blades. • If there is only one NEM, the host REM might identify the storage module virtual drives as a foreign configuration (MegaRAID), or inactive (Fusion-MPT IR). A foreign or inactive configuration can be recovered using the REM management software. For MegaRAID, use the import feature described in the *LSI MegaRAID SAS Software User's Guide*. For Fusion-MPT IR, use the activate feature described in the *LSI SAS2 Integrated RAID User's Guide*.
Connect a SAS2 JBOD to the SAS-NEM external SAS ports.	Not currently supported.	Do not connect external SAS storage to the NEM external SAS ports.

Related Information

- "Prepare SAS Interfaces for Hot-plug Removal" on page 22
- "Saving the Zoning Configuration to a Backup File" on page 25
- "Recovering Zoning Configurations" on page 27

Note – The best way to ensure that there is no interruption in service or accidental data loss when a NEM is removed is to have two NEMs installed in the chassis.

Saving the Zoning Configuration to a Backup File

- "Save the Zoning Configuration to a Backup File (Web)" on page 25
- "Save the Zoning Configuration to a Backup File (CLI)" on page 26

▼ Save the Zoning Configuration to a Backup File (Web)

1. **Create and save the storage access configuration for the chassis using the instructions in the** *Oracle Integrated Lights Out Manager (ILOM) CMM Administration Guide***.**

2. **Make sure that you have permissions for one of the following roles: Admin, User Management, Console, Reset and Host Control, and Read Only (a,u,c,r,o).**

   These roles are required to perform a complete backup of the Oracle ILOM SP configuration.

3. **Open a web browser and log in to the CMM by typing:**

```
http://chassis-sp-ipaddress/
```

   where *chassis-sp-ipaddress* is the IP address of your chassis CMM.

   The ILOM login page appears.

4. **Log in to CMM ILOM.**

   The CMM ILOM main page is displayed.

5. **With CMM selected in the left pane, click the Maintenance tab for the CMM.**

   The CMM maintenance subtabs are displayed

6. **Click the Backup/Restore tab.**

   The Backup/Restore page is displayed.

7. **Select Backup from the Operating drop-down list.**

8. **Fill out the information on the page to create your backup file.**

   For complete instructions on using ILOM Backup/Restore, refer to the *Oracle Integrated Lights Out Manager (ILOM) 3.0 Web Interface Procedures Guide*.

Related Information

- "Prepare SAS Interfaces for Hot-plug Removal" on page 22
- "Restoring SAS Configurations After NEM Installation" on page 22
- "Recovering Zoning Configurations" on page 27

▼ Save the Zoning Configuration to a Backup File (CLI)

1. **Create and save the storage access configuration for the chassis using the instructions in the the** *Oracle Integrated Lights Out Manager (ILOM) CMM Administration Guide***.**

2. **Make sure that you have permissions for one of the following roles: Admin, User Management, Console, Reset and Host Control, and Read Only (a,u,c,r,o).**

   These roles are required to perform a complete backup of the Oracle ILOM SP configuration.

3. **Log in to the CMM ILOM CLI.**

4. **Change to the /SP/config directory.**

```
-> cd /SP/config
```

5. **If you want sensitive data, such as user passwords, SSH keys, certificates, and so forth, to be backed up, you must provide a passphrase.**

```
-> set passphrase=passphrase
```

6. **Initiate the Backup operation:**

```
-> set dump_uri=
transfer-method://username:password@ipaddress/directory-path/filename.config
```

   where:

   - *transfer-method* can be tftp, ftp, sftp, scp, http, or https
   - *username* is the name of the user account on the remote system. (Username is required for scp, sftp, and ftp. Username is not used for tftp, and it is optional for http and https.)
   - *password* is the password for the user account on the remote system. (Password is required for scp, sftp, and ftp. Password is not used for tftp, and it is optional for http and https.)
   - *ipaddress* is the IP address or the host name of the remote system.
   - *directory-path* is the storage location on the remote system.

- *filename* is the name assigned to the backup file.

For complete instructions on using ILOM Backup/Restore, refer to the *Oracle Integrated Lights Out Manager (ILOM) 3.0 CLI Procedures Guide*.

Recovering Zoning Configurations

This topic describes how to recover host-to-storage zoning configurations that might have been lost by accident or due to hardware replacement.

- "Recover Zoning Configurations (Web)" on page 27
- "Recover Zoning Configurations (CLI)" on page 28

▼ Recover Zoning Configurations (Web)

Note – This procedure is for advanced users and Oracle technicians. The CMM ILOM configuration backup file is an XML file. If you have multiple CMM ILOM configuration backup files and the latest version does not have the zoning configurations you need, you have the option of copying the storage assignments section of one file and pasting it into another. For this to work, your storage modules and server blades must be in the same physical slots for the zoning configurations you want to restore. For more information on performing this procedure, refer to the Oracle ILOM documentation.

You must have previously created a backup CMM ILOM configuration file that contains the zoning configurations you want to restore. See "Saving the Zoning Configuration to a Backup File" on page 25.

1. **Make sure that you have permissions for one of the following roles: Admin, User Management, Console, Reset and Host Control, and Read Only (a,u,c,r,o).**

   These roles are required to perform a complete backup of the ILOM SP configuration.

2. **Open a web browser and log in to the CMM by typing:**

```
http://chassis-sp-ipaddress/
```

   where *chassis-sp-ipaddress* is the IP address of your chassis processor.

   The ILOM login page is displayed.

3. **Log in to ILOM.**

4. **With CMM selected in the left pane, click the Maintenance tab for the CMM.**

   The CMM maintenance subtabs are displayed.

5. **Click the Backup/Restore tab.**

   The Backup/Restore page is displayed.

6. **Select Restore from the Operation drop-down list.**

   Fill out the information on the page to restore your backup file.

   For complete instructions on using ILOM Backup/Restore, refer to the *Oracle Integrated Lights Out Manager (ILOM) 3.0 Web Interface Procedures Guide*.

7. **To initiate the Restore operation, click Run.**

   The Restore operation executes.

Note – While the restore operation is executing, sessions on the ILOM SP will be temporarily suspended. The sessions will resume normal operation after the restore operation is complete. A restore operation typically takes two-to-three minutes to complete.

Related Information

- "Recover Zoning Configurations (CLI)" on page 28
- "Prepare SAS Interfaces for Hot-plug Removal" on page 22
- "Restoring SAS Configurations After NEM Installation" on page 22
- "Recovering Zoning Configurations" on page 27

▼ Recover Zoning Configurations (CLI)

Note – This procedure is for advanced users and Oracle technicians. The CMM ILOM configuration backup file is an XML file. If you have multiple CMM ILOM configuration backup files and the latest version does not have the zoning configurations you need, you have the option of copying the storage assignments section of one file and pasting it into another. For this to work, your storage modules and server blades must be in the same physical slots for the zoning configurations that you want to restore. For more information on performing this procedure, refer to the Oracle ILOM documentation.

You must have previously created a backup CMM ILOM configuration file that contains the zoning configurations you want to restore. See "Saving the Zoning Configuration to a Backup File" on page 25.

1. **Make sure that you have permissions for one of the following roles: Admin, User Management, Console, Reset and Host Control, and Read Only (a,u,c,r,o).**

   These roles are required to perform a complete backup of the ILOM SP configuration.

2. **Open a terminal window and establish an ssh connection to the CMM by typing:**

```
# ssh -l username cmm-ip
```

   where *username* is your ILOM username and *cmm-ip* is the IP address of the CMM.

   The login prompt is displayed.

3. **Log in as root and type the root password:**

```
/hostname/login: username password:xxxxxxxx
```

   where *xxxxxxxx* is your root password.

   After you have successfully logged in, the prompt is displayed.

4. **Change to the** `/SP/config` **directory.**

   -> **cd /SP/config**

5. **Initiate the restore operation by typing:**

```
-> set load_uri=
transfer-method://username:password@ipaddress-or-hostname/directory-path/filename.config
```

   where:

   - *transfer-method* can be tftp, ftp, sftp, scp, http, or https.
   - *username* is the name of the user account on the remote system. (Username is required for scp, sftp, and ftp. Username is not used for tftp, and it is optional for http and https.)
   - *password* is the password for the user account on the remote system. (Password is required for scp, sftp, and ftp. Password is not used for tftp, and it is optional for http and https.)
   - *ipaddress-or-hostname* is the IP address or the host name of the remote system.
   - *directory-path* is the storage location on the remote system.
   - *filename* is the name assigned to the backup file.

   The restore operation executes.

Related Information

- “Recover Zoning Configurations (Web)” on page 27
- “Prepare SAS Interfaces for Hot-plug Removal” on page 22
- “Saving the Zoning Configuration to a Backup File” on page 25
- “Restoring SAS Configurations After NEM Installation” on page 22
- “Prepare SAS Interfaces for Hot-plug Removal” on page 22
- “Restoring SAS Configurations After NEM Installation” on page 22

Installing or Replacing a NEM

These topics describe how to install or replace a NEM in a powered-on Sun Blade 6000 Series Chassis.

Caution – Damage to the NEM can occur as the result of careless handling or ESD. Always handle the NEM with care to avoid damage to electrostatic-sensitive components. To minimize the possibility of ESD-related damage, Oracle strongly recommends using both a workstation anti-static mat and an ESD wrist strap. You can get an ESD wrist strap from an electronics store or from Oracle as part number 250-1007.

- "Before Installing a NEM" on page 31
- "Install a NEM" on page 34
- "Verifying NEM Installation" on page 36
- "Remove a NEM" on page 40
- "Replace a NEM" on page 43
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46
- "SFP+ Connector Cabling Options" on page 49

Before Installing a NEM

If your Sun Blade 6000 modular system chassis has not yet been installed and configured, read the *Sun Blade 6000 Modular System Installation Guide* before following the installation instructions for the NEM.

The *Sun Blade 6000 Modular System Installation Guide* might be packaged with the Sun Blade 6000 modular system chassis. You can find this document at:

`http://docs.oracle.com/cd/E19938-01/`

To verify compatibility of NEM with chassis, blades, and optional components, refer to the *Sun Blade 6000 Virtualized 40 GbE Network Express Module Product Notes*.

Verifying the CMM

If you want to install a Sun Blade 6000 Virtualized 40 GbE Network Express Module on a Sun Blade 6000 A90-D chassis, you might have to replace the CMM.

If the NEM is installed in the chassis when you order it, the correct CMM will be installed in the chassis.

To determine whether or not you need to replace the CMM on your A90-D chassis, verify the CMM part number by using one of these procedures:

- "Determine the CMM Part Number Using ILOM 3.0x Web Interface" on page 32
- "Determine the CMM Part Number Using ILOM 3.1x Web Interface" on page 32
- "Determine the CMM Part Number Using ipmitool" on page 33
- "Find the CMM Part Number on the CMM Label" on page 33

▼ Determine the CMM Part Number Using ILOM 3.0x Web Interface

1. **Log into the CMM web interface.**
2. **Click Chassis on the left navigation bar.**

   The chassis inventory is displayed.
3. **Find the /CH/CMM field.**

   The CMM part number is displayed in this field.
4. **Do one of the following:**
   - If the product part number displayed is 541-4340-02, contact Oracle service to order a replacement CMM before you install the NEM.
   - If the product part number displayed is 7019286 or higher, the CMM is supported. Proceed with the NEM installation.

▼ Determine the CMM Part Number Using ILOM 3.1x Web Interface

1. **Log into the CMM web interface.**

   The first screen displayed is the Chassis Summary.
2. **Click Chassis View on the left navigation bar.**

   The chassis inventory is displayed.

3. **Find the /CH/CMM field.**

   The CMM part number is displayed in this field.

4. **Do one of the following:**

   - If the product part number displayed is 541-4340-02, contact Oracle service to order a replacement CMM before you install the NEM.
   - If the product part number displayed is 7019286 or higher, the CMM is supported. Proceed with the NEM installation.

▼ Determine the CMM Part Number Using `ipmitool`

This procedure assumes that `ipmitool` is installed and enabled on a remote system that is networked with the CMM.

Note – This method cannot be used to identify the CMM part number when the CMM has ILOM 3.1.x installed. You can use the web interface to identify the part number or locate the part number from the CMM part number label.

---

1. **Issue the following command:**

```
ipmitool -H cmm_ip -v -I lanplus -U username -P password fru print 3
```

   where *cmm_ip* is the IP address of the CMM and *username* and *password* are the Oracle ILOM CMM user name and password.

2. **Do one of the following:**

   - If the Product Part Number displayed is 541-4340-02, contact Oracle service to order a replacement CMM before you install the NEM.
   - If the Product Part Number displayed is 7019286 or higher, the CMM is supported. Proceed with the NEM installation.

▼ Find the CMM Part Number on the CMM Label

1. **To view the CMM part number label, remove the CMM partially from the chassis until you can view the part number label.**

   Refer to the *Sun Blade 6000 Modular System Service Manual* for detailed instructions on removing the CMM.

2. **Do one of the following:**
   - If the Product Part Number displayed is 541-4340-02 or less than 7019286, contact Oracle service to order a replacement CMM before you install the NEM.
   - If the Product Part Number displayed is 7019286 or higher, the CMM is supported. Proceed with the NEM installation.

Related Information

- "Install a NEM" on page 34
- "Verifying NEM Installation" on page 36
- "Remove a NEM" on page 40
- "Replace a NEM" on page 43
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46
- "SFP+ Connector Cabling Options" on page 49
- "Understanding the NEM" on page 1
- "Performing Hot-plug Insertion and Removal" on page 17

▼ Install a NEM

This procedure is for installing a NEM into an empty slot. If you are replacing a NEM, see "Replace a NEM" on page 43.

Note – Mixing different NEM types is not supported. NEM pairs must be the same.

You can insert either one or two NEMs in the Sun Blade 6000 chassis. If you insert only one, install it in the lower slot (NEM0).

1. **If you are installing a NEM into a powered-on chassis, read** "Performing Hot-plug Insertion and Removal" on page 17**.**
2. **Remove the NEM filler panel from the slot you plan to use.**
3. **Align the NEM with the vacant NEM slot.**

   Ensure that:

   - The RJ-45 port connectors of the NEM are facing you.
   - The NEM ejector levers are fully opened.
   - The NEM ejector levers are on the top of the module.

4. **Slide the NEM into the vacant NEM chassis slot until you feel it stop.**

   Support the weight of the NEM with one hand at the bottom of the NEM.

5. **Complete the installation by closing the ejector levers to secure the NEM in the chassis.**

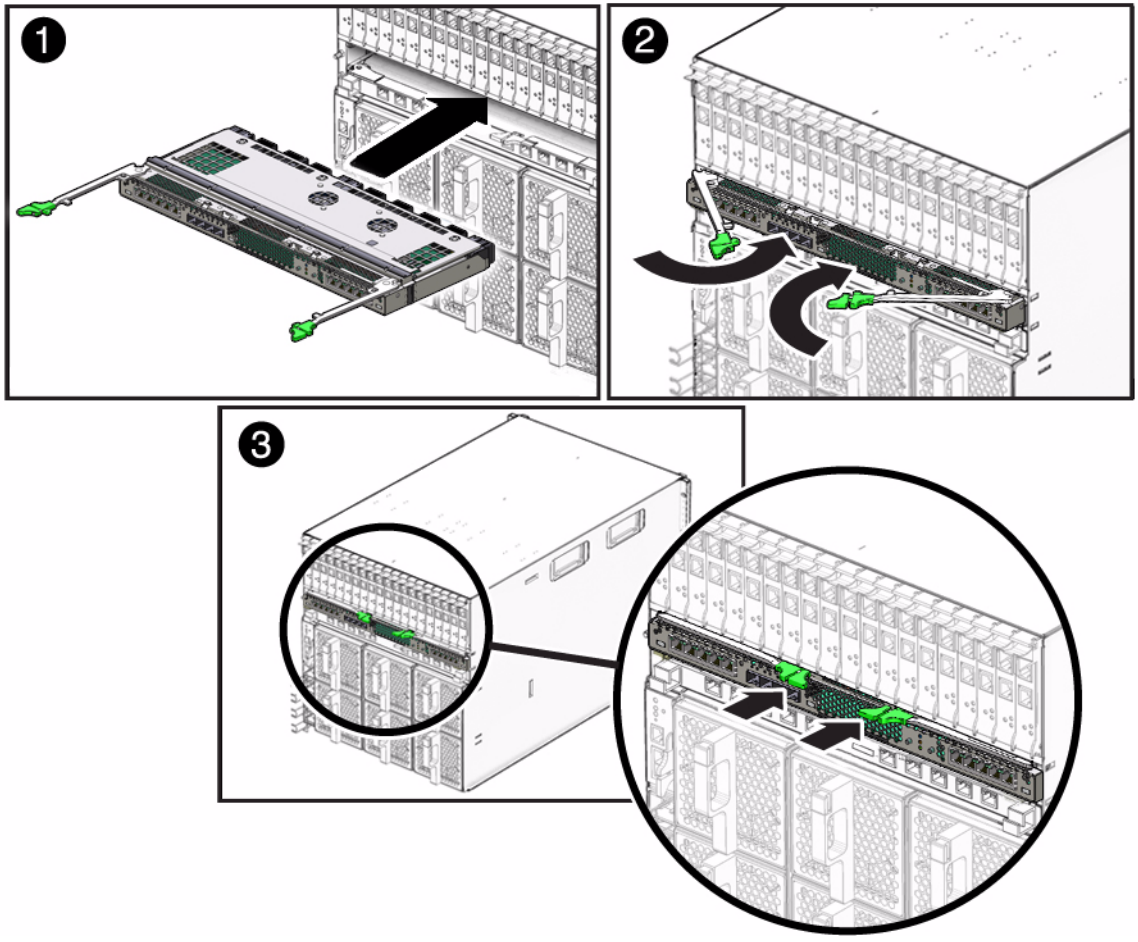

6. **Do one of the following:**

   - If you have installed a NEM in a powered-off chassis, go to Step 11.
   - If you are installing the NEM in a powered-on chassis, complete the rest of the procedure.

7. **After inserting the NEM into a chassis slot, wait until the Power/OK LED goes to standby blink.**

   Standby LED blinking indicates that the NEM is in standby mode.

8. **Press the ATTN button once.**

   The Power/OK LED blinks faster. Wait for the NEM to come online. It can take up to three minutes for the NEM to come online and become fully operational.

Note – If you push the ATTN button twice within five seconds, the green Power/OK LED remains in standby blink mode, and the request for bringing the NEM online is cancelled.

---

9. **View the behavior of the Power/OK LED or the ILOM interface.**

   The NEM is online and the PCIe interfaces are available to the blades when the following occurs:

   - The NEM green Power/OK LED illuminates solid.
   - In the ILOM web interface System Information —> Components screen: the Prepare-to-Remove option is available in the Actions drop-down list for the NEM.
   - In the ILOM CLI `show /CH/NEMx` output, the following status shows in the Properties section: `prepare_to_remove_status = NotReady`.

10. **Refer to** "Restoring SAS Configurations After NEM Installation" on page 22 **to see if any additional actions need to be taken for the SAS interfaces.**

11. **Connect cables, as needed. See** "SFP+ Connector Cabling Options" on page 49**.**

Related Information

- "Verifying NEM Installation" on page 36
- "Remove a NEM" on page 40
- "Replace a NEM" on page 43
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46
- "SFP+ Connector Cabling Options" on page 49
- "Understanding the NEM" on page 1
- "Performing Hot-plug Insertion and Removal" on page 17
- "Restoring SAS Configurations After NEM Installation" on page 22

Verifying NEM Installation

- "Verify Installation (CMM ILOM Web)" on page 37
- "Verify Installation (CMM ILOM CLI)" on page 37

▼ Verify Installation (CMM ILOM Web)

1. **In a web browser, type the IP address of the CMM in the location bar.**
2. **Log in to the CMM ILOM.**
3. **From the CMM ILOM left navigation bar, select CMM.**
4. **Select Components from the second row of tabs.**
5. **In the Component Name column, click on the NEM that you installed.**

   For example: /CH/NEM1.

   Information about the NEM is displayed.

   If the NEM information is not displayed in the Component Management screen, verify that the NEM has been properly seated in the chassis.
6. **Select the Storage tab.**
7. **Click the Detailed Setup button under Sun Blade Zone Manager.**

   The NEM SAS–2 ports are displayed at the bottom of the Manage Zoning window.

Related Information

- "Verify Installation (CMM ILOM CLI)" on page 37
- "Remove a NEM" on page 40
- "Replace a NEM" on page 43
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46
- "SFP+ Connector Cabling Options" on page 49
- "Understanding the NEM" on page 1

▼ Verify Installation (CMM ILOM CLI)

The NEM is automatically detected when the chassis is powered on.

You can connect to the ILOM CLI by either of two methods:

- Connect a terminal or PC running a terminal emulator directly to the CMM serial port on your chassis.

  or

- Connect to the chassis Ethernet network management port using a secure shell (ssh).

Instructions for setting up and using ILOM are in "Connecting to NEM ILOM" on page 133.

Note – The examples in this topic refer to NEM0. If you are installing NEM1, replace the “0” in these examples with “1.”

1. **Log in to the CMM ILOM.**
2. **Type the following command:**

   -> **show /CH/NEM***x*

   where x is 0 or 1, depending on the slot the NEM is installed in.

   The NEM and the NEM FRU information is displayed in the CLI.

```
  Targets:
         MB
         SAS
         SP
         PRSNT
         STATE
         ERR
         OK
         SERVICE
         OK2RM
         LOCATE
     Properties:
         type = Network Express Module
  ipmi_name = NEM0
         fru_name = SUN BLADE 6000 VIRTUALIZED MULTI-FABRIC 40GE
PCI-EX NEM
         fru_version = SAS 5.3.8.1
         fru_part_number = 541--4391-01
         fru_serial_number = 0328MSL-1113L20028
         fru_extra_1 = SAS 5.3.8.1
         fault_state = Ok
         load_uri = (none)
         clear_fault_action = (none)
         prepare_to_remove_status = NotReady
         prepare_to_remove_action = (none)
         return_to_service_action = (none)
     Commands:
         cd
         load
         reset
         set
         show
```

3. **Verify that the NEM has joined the SAS fabric by typing:**

```
-> show /STORAGE/sas_zoning
```

```
/STORAGE/sas_zoning
    Targets:
        BL0
        BL2
        BL3
        BL5
        BL7
        BL8
        BL9
        NEM0
        NEM1
    Properties:
        zone_management_state = enabled
        reset_password_action = (Cannot show property)
        reset_access_action = (Cannot show property)
    Commands:
        cd
        set
        show
```

4. **If the NEM does not appear in ILOM, verify that the NEM is properly seated in the chassis, then verify that you have the newest CMM ILOM firmware.**

   For more information about using CMM ILOM, see the ILOM documentation, at:

   http://docs.oracle.com/cd/E19860-01/

Related Information

- "Verify Installation (CMM ILOM Web)" on page 37
- "Remove a NEM" on page 40
- "Replace a NEM" on page 43
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46
- "SFP+ Connector Cabling Options" on page 49
- "Understanding the NEM" on page 1

▼ Remove a NEM

You can remove a NEM from a powered-on chassis using a hot-plug operation, or from a powered-off chassis. The server modules must be informed if you are replacing a non-redundant NEM.

This procedure is for removing a NEM if you do not plan to replace it. If you plan to replace a NEM, see "Replace a NEM" on page 43.

Caution – If the chassis is powered on and you are not replacing the NEM within 60 seconds, install a NEM filler panel to ensure proper system operation.

Note – Hot-plug service has to be enabled on all blades in the chassis for the hot-plug removal to properly initiate. If the Ready-to-Remove LED does not light, check the configuration of the blades, or power off all of them.

1. **If you are removing a NEM from a powered-on chassis, read the hot-plug guidelines in** "Performing Hot-plug Insertion and Removal" on page 17 **and follow the instructions in** "Prepare the PCIe Interface for Hot-plug Removal" on page 20 **and** "Prepare SAS Interfaces for Hot-plug Removal" on page 22**.**
2. **In the rear of the chassis, locate the NEM that you want to remove.**
3. **Do one of the following:**
   - If you are removing a NEM from a powered-off chassis, go to Step 7.
   - If you are removing a NEM from a powered-on chassis, continue with Step 4.
4. **Push the ATTN button once.**

Note – Pushing the ATTN button twice within five seconds will cause the Power/OK LED to return to solid green. This action will cancel the request for a NEM removal.

5. **View the behavior of the Power/OK LED and Ready-to-Remove LED or the ILOM interface to determine if the Ready-to-Remove action is successful.**

   If the action is successful:

   - The NEM green Power/OK begins to slowly blink.
   - The OK- to-Remove blue LED illuminates and the Power LED turns off.

- In the ILOM System Information —> Components screen: The Return-to-Service option is available in the Actions drop-down list for the NEM and the Ready-to-Remove status column for the NEM is Ready.
- In the ILOM CLI `show /CH/NEMx` command output, the following status is listed in the Properties section: `prepare_to_remove_status = Ready`
- /CH/NEM*x*/OK reads Off in either the ILOM web interface or CLI.

If the action is unsuccessful:

- The NEM green Power/OK LED begins to slowly blink, then returns to solid green, and the Ready-to-Remove LED does not go on.
- In the ILOM System Information —> Components screen, the Return to Service option does not become available in the Actions drop-down list and the Ready-to-Remove status column for the NEM still reads Not Ready.
- In the ILOM CLI `show /CH/NEMx` command output, the following status is listed in the Properties section: `prepare_to_remove_status = NotReady`
- /CH/NEM*x*/OK reads On in either the ILOM web interface or CLI.

6. **Do one of the following:**
   - If the Ready-to-Remove action is successful, proceed with Step 7.
   - If the Ready-to-Remove action is not successful, do not remove the NEM. Follow the instructions in "Forcing NEM Graceful Removal" on page 19 to remove the NEM.
7. **Label the cables so that you can install them in the same location, then remove all cables from the NEM.**

Note – Wait five minutes after CMM boot up or NEM insertion before removing the NEM.

---

8. **Press together and hold the buttons on both the right and left ejector levers.**
9. **To unlatch the NEM from the chassis, open the ejector levers by extending them outward.**
10. **Holding the open ejector levers, pull the NEM toward you until you can pull the rest of the module out by hand.**

    Support the weight of the NEM, with one hand at the bottom of the NEM.

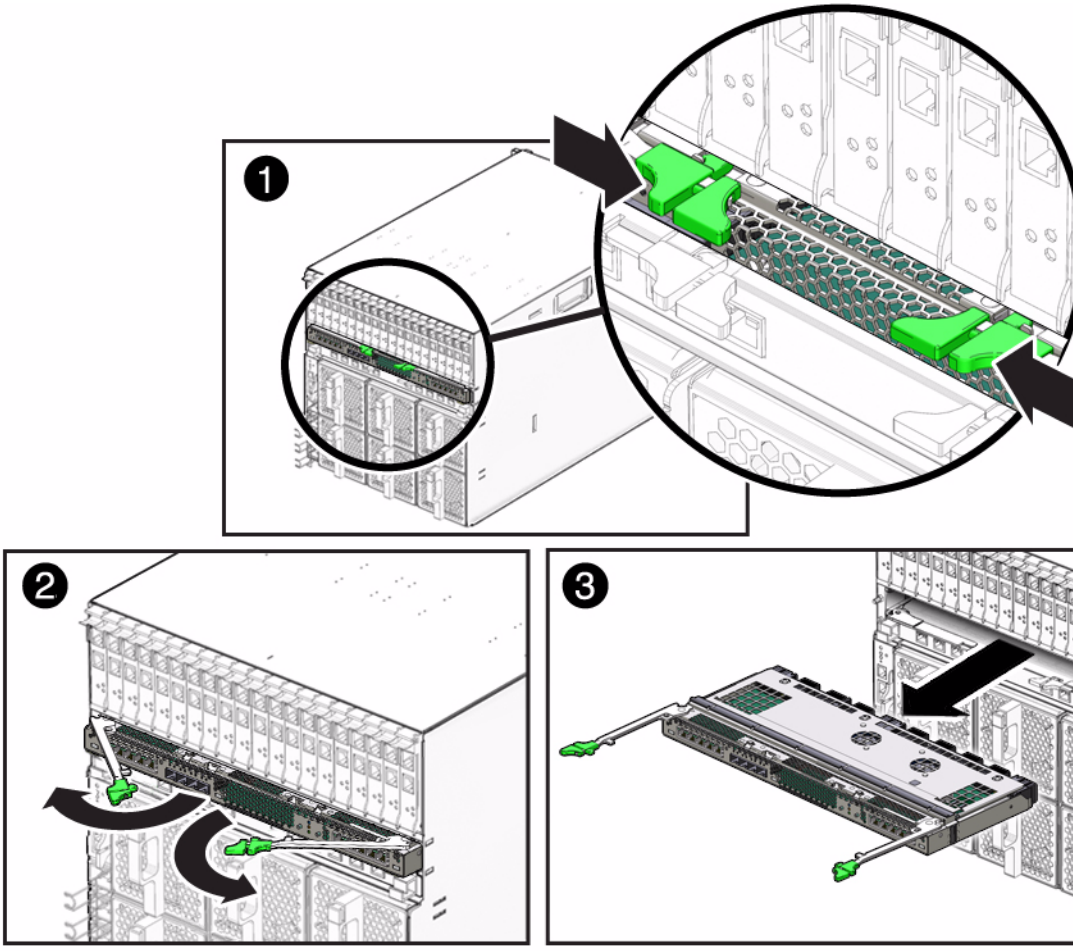

Related Information

- "Replace a NEM" on page 43
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46
- "SFP+ Connector Cabling Options" on page 49
- "Understanding the NEM" on page 1
- "Performing Hot-plug Insertion and Removal" on page 17

▼ Replace a NEM

The server modules must be informed if you are replacing a non-redundant NEM.

Caution – If the chassis is powered on and you are not replacing the NEM within 60 seconds, install an NEM filler panel to ensure proper system operation.

Caution – If you are replacing a functional SAS2 NEM (for example replacing a NEM with a NEM) and another SAS2 NEM is installed in the chassis, make sure that multipathing at the OS level is enabled on all active server blades in the chassis before you remove the existing NEM. This ensures that the primary path to disks is not lost, which can cause an OS panic. For servers running Windows 2003, which does not support multipathing, shut down I/O to all disks on the disk module that is not in hardware RAID volumes.

If you are removing a NEM from a powered on chassis, read the hot-plug guidelines "Performing Hot-plug Insertion and Removal" on page 17 and follow the instructions in "Prepare the PCIe Interface for Hot-plug Removal" on page 20.

Caution – Before replacing a NEM, make sure that you have backed up the SAS assignment configuration, as described in "Saving the Zoning Configuration to a Backup File" on page 25.

1. **In the rear of the chassis, locate the NEM that you want to remove.**
2. **Do one of the following:**
   - If you are replacing a NEM in a powered-off chassis, go to Step 6.
   - If you are replacing a NEM in a powered-on chassis, continue with Step 3.

Note – Hot-plug service has to be enabled on all blades in the chassis for a hot-plug removal to properly initiate. If the Ready-to-Remove LED does not light, check the configuration of the blades, or power off all of them.

3. **Push the ATTN button once.**

Note – Pushing the ATTN button twice within five seconds will cause the Power/OK LED to return to solid green. This action will cancel the request for a NEM removal.

4. **View the behavior of the Power/OK LED and Ready-to-Remove LED or the ILOM interface to determine if the Ready-to-Remove action is successful.**

   If the action is successful:

   - The NEM green Power/OK begins to slowly blink.
   - The OK- to-Remove blue LED illuminates and the Power LED turns off.
   - In the ILOM System Information —> Components screen: The Return-to-Service option is available in the Actions drop-down list for the NEM and the Ready-to-Remove status column for the NEM is Ready.
   - In the ILOM CLI `show /CH/NEMx` command output, the following status is listed in the Properties section: `prepare_to_remove_status = Ready`
   - /CH/NEM*x*/OK reads Off in either the ILOM web interface or CLI.

   If the action is unsuccessful:

   - The NEM green Power/OK LED begins to slowly blink, then returns to solid green, and the Ready-to-Remove LED does not go on.
   - In the ILOM System Information —> Components screen, the Return-to-Service option does not become available in the Actions drop-down list and the Ready-to-Remove status column for the NEM still reads Not Ready.
   - In the ILOM CLI `show /CH/NEMx` command output, the following status is listed in the Properties section: `prepare_to_remove_status = NotReady`
   - /CH/NEM*x*/OK reads On in either the ILOM web interface or CLI.

5. **Do one of the following:**

   - If the Ready-to-Remove action is successful, proceed with Step 6.
   - If the Ready-to-Remove action is not successful, do not remove the NEM. Follow the instructions in "Forcing NEM Graceful Removal" on page 19 to remove the NEM.

6. **Label the cables so that you can install them in the same location, then remove all cables from the NEM.**

7. **Press together and hold the buttons on both the right and left ejector levers.**

8. **To unlatch the NEM from the chassis, open the ejector levers by extending them outward.**

9. **Holding the open ejector levers, pull the NEM toward you until you can pull the rest of the module out by hand.**

   Support the weight of the NEM, with one hand at the bottom of the NEM.

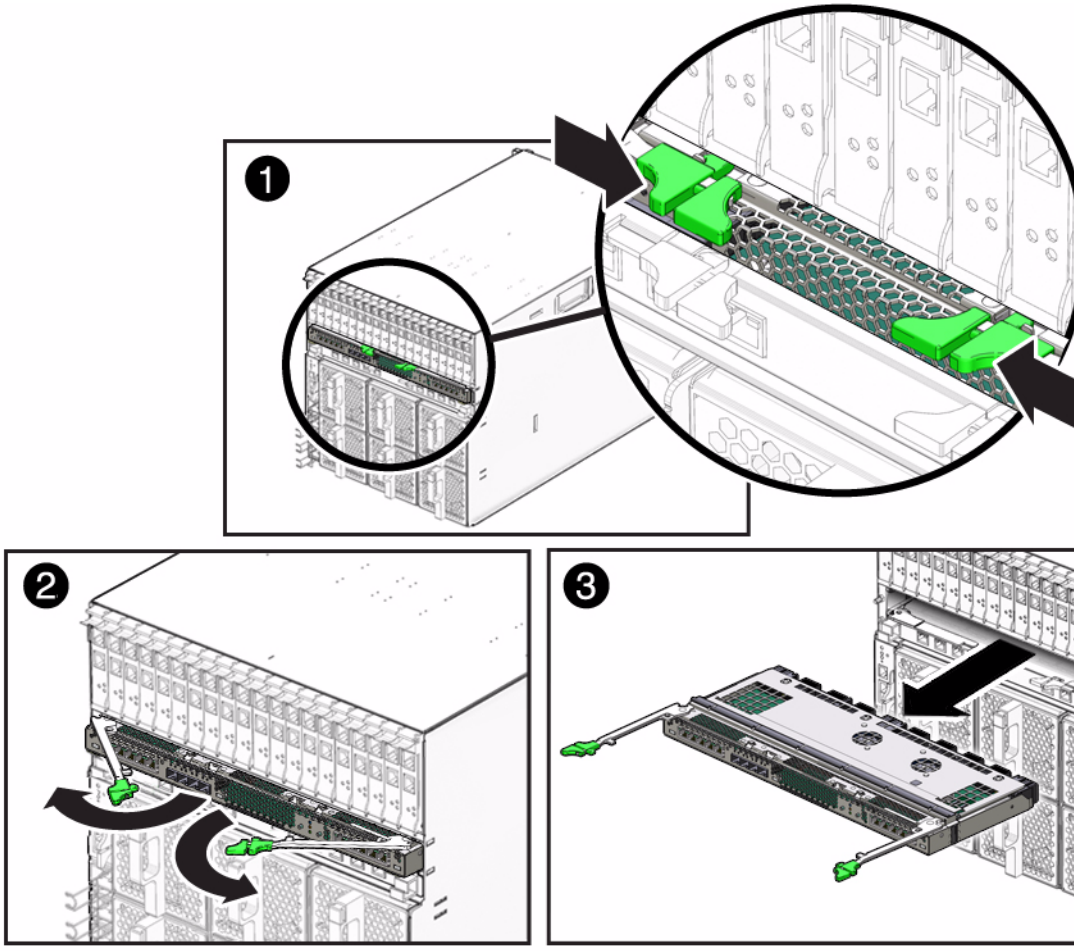

Note – Wait at least 10 seconds before installing the new NEM.

10. **Insert the new NEM into the chassis.**
11. **Do one of the following:**
    - If you have replaced a NEM in a powered-off chassis, go to Step 16.
    - If you have replaced a NEM in a powered-on chassis, continue with Step 12.
12. **After the NEM has been inserted into a chassis slot, wait until the Power/OK LED goes to standby blink.**

    Standby blink indicates that the NEM is in standby mode.

13. **Press the ATTN button once.**

    It can take up to three minutes for the NEM to come online.

Note – If you push the ATTN button twice within five seconds, the green Power/OK LED remains in standby blinking mode, and the request for bringing the NEM online is cancelled.

---

14. **View the behavior of the Power/OK LED or the ILOM interface.**

    The NEM is online and the PCIe interfaces are available to the blades when the following occurs:

    - The NEM green Power/OK LED illuminates solid.
    - In the ILOM web interface System Information —> Components screen, the Prepare-to-Remove option is available in the Actions drop-down list for the NEM.
    - In the ILOM CLI `show /CH/NEMx` output, the following status shows in the Properties section: `prepare_to_remove_status = NotReady`.

15. **See "Restoring SAS Configurations After NEM Installation" on page 22 to see if any additional actions need to be taken for the SAS interfaces.**

16. **Reconnect all the cables in their original locations.**

Related Information

- "Verifying NEM Installation" on page 36
- "Remove a NEM" on page 40
- "Installing and Removing SFP+ Optical Transceiver Modules" on page 46
- "SFP+ Connector Cabling Options" on page 49
- "Understanding the NEM" on page 1
- "Performing Hot-plug Insertion and Removal" on page 17
- "Restoring SAS Configurations After NEM Installation" on page 22

Installing and Removing SFP+ Optical Transceiver Modules

The NEM requires an SFP+ transceiver in at least one port to create a 10GbE connection. Before installing or removing an SFP+ module, refer to "10GbE NIC Virtualization" on page 8.

- “Install an SFP+ Optical Transceiver Module” on page 47
- “Remove an Optical Transceiver SFP+ Module” on page 48

▼ Install an SFP+ Optical Transceiver Module

1. **Remove the protective end cap for the optical transceiver.**
2. **Pull the locking handle into the full horizontal position until you feel the handle click into position.**
3. **Holding the optical transceiver (SFP+ module) by the edges, align the SFP+ module with the slot in the NEM and slide it into the opening.**
4. **Applying even pressure at both corners of the SFP+ module, push the module until it is firmly seated in the slot.**
5. **Push the handle up to the vertical position to lock the SFP+ module in place.**

Note – If you pull the locking handle down when the SFP+ module is installed, remove the SFP+ module entirely and reinstall it. The handle operates an internal lock. Pulling the handle down can disconnect the SFP+ module, even though it might appear to be connected.

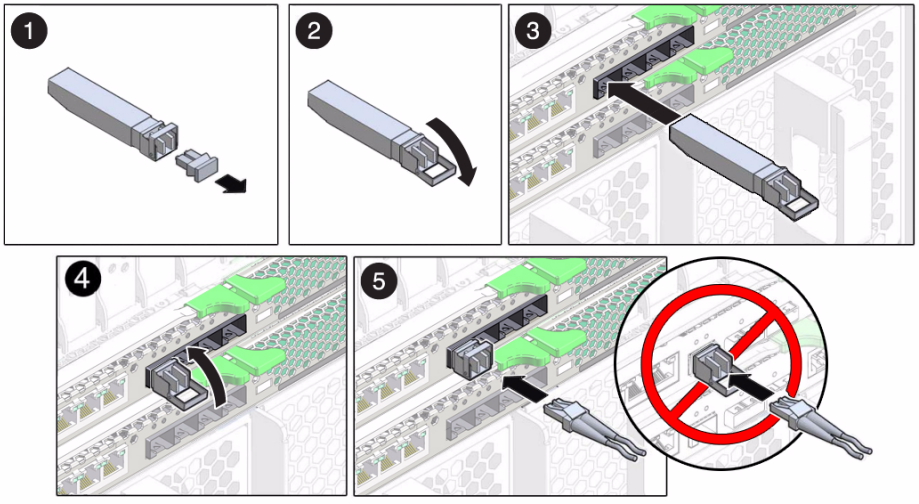

Related Information

- “Remove an Optical Transceiver SFP+ Module” on page 48
- “SFP+ Connector Cabling Options” on page 49

- "Identifying Ports and LEDs" on page 9
- "Remove a NEM" on page 40
- "Replace a NEM" on page 43
- "Understanding the NEM" on page 1

▼ Remove an Optical Transceiver SFP+ Module

1. **Press the optical cable connector latch down, and gently pull out the optical cable.**

Tip – Fully remove the cables from the transceivers. Optical equipment is sensitive. Do not leave cables partially in the transceivers.

---

2. **Pull down the SFP+ module latch into the open position.**
3. **Slide out the SFP+ module.**
4. **Replace the protective end caps if you plan to store the module.**

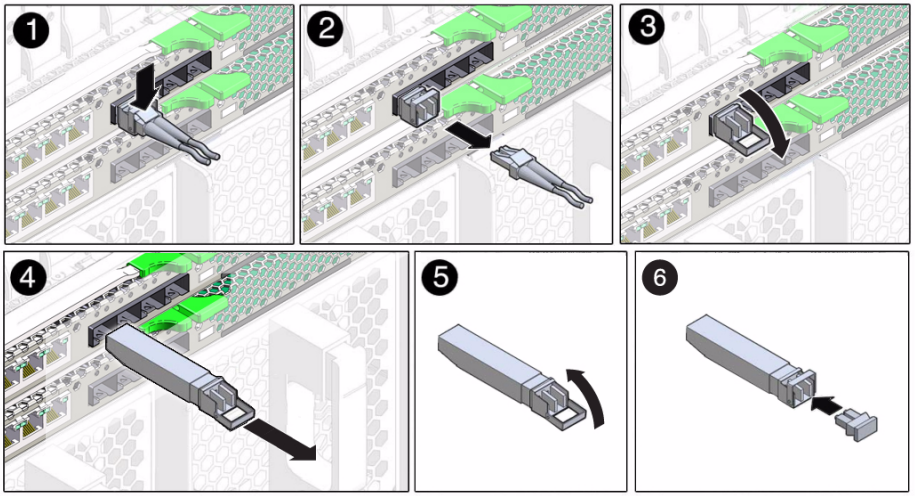

Related Information

- "Install an SFP+ Optical Transceiver Module" on page 47
- "SFP+ Connector Cabling Options" on page 49
- "Identifying Ports and LEDs" on page 9
- "SFP+ Connector Cabling Options" on page 49

- "Remove a NEM" on page 40
- "Replace a NEM" on page 43
- "Understanding the NEM" on page 1

SFP+ Connector Cabling Options

The following table shows the OM MMF and SMF cabling options available for the NEM SFP+ connectors.

SFP+ Connector Type	Cable type	Reach
Short range	OM 1 MMF	33 meters
Short range	OM 2 MMF	50 meters
Long range	OM 3 MMF	300 meters

Note – The optical transceiver for the NEM uses a LC. The connection on the other side of the cable might use a different connector type.

Note – For the ten 1G RJ-45 pass-through ports, maximum cable length is 90 meters, Cat5e or better.

Note – In 10GbE and 40GbE configurations, the 5-meter copper splitter (P/N X2125A-5M) is not supported on the SFP+ ports.

Tip – If LC cables are not fully inserted, link instability might occur. Ensure that a click is heard when inserting cables.

Caution – Avoid putting unnecessary stress on the connection. Do not bend or twist the cable near the connectors, and avoid sharp cable bends of more than 90 degrees.

Related Information

- "Identifying Ports and LEDs" on page 9
- "Remove a NEM" on page 40
- "Replace a NEM" on page 43

Enabling Jumbo Frames

These topics provide instructions for configuring Jumbo Frames, based on your OS:

- "Configuring Jumbo Frames (Solaris)" on page 51
- "Configuring Jumbo Frames (Linux)" on page 56
- "Enable Jumbo Frames (Windows)" on page 59
- "Configuring Jumbo Frames (VMware ESX)" on page 59

Configuring Jumbo Frames (Solaris)

- "Jumbo Frames Overview" on page 51
- "Jumbo Frames Configurations and Driver Statistics" on page 52
- "Enable Jumbo Frames (Solaris)" on page 54
- "Check Layer 2 Configuration" on page 55
- "Check Layer 3 Configuration" on page 55

Jumbo Frames Overview

By default, Jumbo Frames is disabled. Configuring Jumbo Frames enables Ethernet interfaces to send and receive packets larger than the standard 1500 bytes. However, the actual transfer size depends on the switch capability and the Ethernet NEM driver capability.

Note – Refer to the documentation that came with your switch for commands to configure Jumbo Frames support.

Related Information

- "Jumbo Frames Configurations and Driver Statistics" on page 52

- "Check Layer 2 Configuration" on page 55
- "Check Layer 3 Configuration" on page 55

Jumbo Frames Configurations and Driver Statistics

The Jumbo Frame feature is by default disabled. The Jumbo Frames configuration checking occurs at Layer 2 or Layer 3, depending on the configuration method.

The following examples show uses of the `kstat` command to display driver statistics.

- Display the receive packet counts on all four of the receive DMA channels on interface 1, for example:

```
# kstat -m sxge | grep rdc_pac
rdc_packets                     120834317
rdc_packets                     10653589436
rdc_packets                     3419908534
rdc_packets                     3251385018
```

  Using the `kstat sxge:1` command shows all the statistics the driver supports for that interface.

- Display driver statistics of a single DMA channel, for example:

```
# kstat -m sxge -n RDC_0
module: sxge                            instance: 0
name:   RDC_0                           class:    net
crtime                          134.619306423
ctrl_fifo_ecc_err               0
data_fifo_ecc_err               0
peu_resp_err                    0
rdc_bytes                       171500561208
rdc_errors                      0
rdc_packets                     120834318
rdc_rbr_empty                   0
rdc_rbrfull                     0
rdc_rbr_pre_empty               0
rdc_rbr_pre_par_err             0
rdc_rbr_tmout                   0
rdc_rcrfull                     0
rdc_rcr_shadow_full             0
rdc_rcr_sha_par_err             0
rdc_rcrthres                    908612
```

```
rdc_rcrto                       150701175
rdc_rcr_unknown_err             0
snaptime                        173567.49684462
```

- Display driver statistics of `sxge0` interface, for example:

```
# kstat -m sxge -n sxge0
module: sxge                            instance: 0
name:   sxge0                           class:    net
brdcstrcv                       0
brdcstxmt                       0
collisions                      0
crtime                          134.825726986
ierrors                         0
ifspeed                         10000000000
ipackets                        265847787
ipackets64                      17445716971
multircv                        0
multixmt                        0
norcvbuf                        0
noxmtbuf                        0
obytes                          1266555560
obytes64                        662691519144
oerrors                         0
opackets                        129680991
opackets64                      8719615583
rbytes                          673822498
rbytes64                        24761160283938
snaptime                        122991.23646771
unknowns                        0
```

- Display all driver statistics, for example:

```
# kstat -m sxge
```

Related Information

- "Jumbo Frames Overview" on page 51
- "Check Layer 2 Configuration" on page 55
- "Check Layer 3 Configuration" on page 55

▼ Enable Jumbo Frames (Solaris)

The following procedure describes how to enable Jumbo Frames in an Oracle Solaris OS 10 environment. By default, the feature is disabled.

1. **To the `/kernel/drv/sxge.conf` file, add the following to enable Jumbo Frames:**

```
accept-jumbo=1;
```

2. **Reboot the blade.**

```
% reboot -- -r
```

3. **Plumb the `sxge` interface to assign MTU size.**

   For example:

```
ifconfig sxgex plumb IP_address up mtu y
```

   where *x* can be *0*, *1*, *2*, or *3*, and *y* can be any number from 1501 to 9194.

4. **To verify that the MTU is assigned properly, use the following command:**

```
ifconfig -a
```

Related Information

- "Jumbo Frames Overview" on page 51
- "Jumbo Frames Configurations and Driver Statistics" on page 52
- "Check Layer 2 Configuration" on page 55
- "Check Layer 3 Configuration" on page 55

▼ Check Layer 2 Configuration

- **View the MTU configuration of an `sxge` instance at any time with the `ifconfig` command.**

```
# ifconfig -a
sxge0: flags=1000843 <UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu
9178 index 4
inet 192.168.1.29 netmask ffffff00 broadcast 192.168.1.255
ether 0:14:4f:62:1:3
```

Note that the MTU (9178) is 38 bytes fewer than the maximum Jumbo Frame size (9216). These 38 bytes include the 16 byte-hardware header, the Ethernet header, maximum payload, and the cyclic-redundancy-check checksum.

Related Information

- "Jumbo Frames Overview" on page 51
- "Jumbo Frames Configurations and Driver Statistics" on page 52
- "Check Layer 3 Configuration" on page 55

▼ Check Layer 3 Configuration

- **Check the Layer 3 configuration by using the `dladm` command with the `show-link` option.**

For example:

```
# dladm show-link
sxge0           type: non-vlan  mtu: 9194       device: sxge0
sxge1           type: non-vlan  mtu: 9194       device: sxge1
igb0            type: non-vlan  mtu: 1500       device: igb0
igb1            type: non-vlan  mtu: 1500       device: igb1
sxge2           type: non-vlan  mtu: 9194       device: sxge2
sxge3           type: non-vlan  mtu: 9194       device: sxge3
```

Related Information

- "Check Layer 2 Configuration" on page 55
- "Jumbo Frames Overview" on page 51
- "Enable Jumbo Frames (Solaris)" on page 54
- "Jumbo Frames Configurations and Driver Statistics" on page 52

Configuring Jumbo Frames (Linux)

By default, Linux configures Ethernet network interfaces to support only standard-size Ethernet frames (1500 bytes). The NEM hardware supports Ethernet Jumbo Frames up to 9216 bytes.

To enable network interface support of Ethernet Jumbo Frames, use the `ifconfig(8)` command to set the network interface MTU parameter to the desired frame size.

Note that there is no official or standard Jumbo Frame size specification. While the exact size chosen for your network's Jumbo Frame support is typically not important, it is important that you configure all communicating nodes on the network to have the same size (in case a packet-size error occurs and the packet gets discarded).

Note – You can use the commands shown in the following examples for both RHEL and SLES.

- "Temporarily Configure Jumbo Frames (Linux)" on page 56
- "Permanently Enable Jumbo Frames (Linux)" on page 57

▼ Temporarily Configure Jumbo Frames (Linux)

To temporarily enable jumbo frame support for a network interface, use the `ifconfig ethn mtu` *nnn* command. You can do this while the interface is up and running (and actively passing network traffic), but if you set the maximum frame size to a smaller value, you might disrupt incoming traffic from other nodes that are using an older (larger) value.

1. **Check the current frame size (MTU) value.**

```
host #> ifconfig eth2
eth2      Link encap:Ethernet  HWaddr 00:14:4F:29:00:1D
inet addr:10.1.10.156  Bcast:10.1.10.255  Mask:255.255.255.0
inet6 addr: fe80::214:4fff:fe29:1/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
RX packets:1 errors:0 dropped:0 overruns:0 frame:0
TX packets:30 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:150 (150.0 b)  TX bytes:7850 (7.6 KiB)
  Memory:fb000000-fc000000
```

Note that in this example, `eth2` (the NEM from previous examples) is currently running with the standard MTU of 1500 bytes.

2. **Set the desired new value.**

For a 9000 byte example:

```
host #> ifconfig eth2 mtu 9000
```

3. **Verify the setting.**

```
host #> ifconfig eth2
eth2      Link encap:Ethernet  HWaddr 00:14:4F:29:00:01
inet addr:10.1.10.156  Bcast:10.1.10.255  Mask:255.255.255.0
inet6 addr: fe80::214:4fff:fe29:1/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
RX packets:26 errors:0 dropped:0 overruns:0 frame:0
TX packets:40 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:3900 (3.8 KiB)  TX bytes:9352 (9.1 KiB)
Memory:fb000000-fc000000
```

Note that `ifconfig` reports the MTU size to be 9000 bytes. NFS 8KB pages will now flow (send or receive) as a single Ethernet packet.

Related Information

- "Permanently Enable Jumbo Frames (Linux)" on page 57

▼ Permanently Enable Jumbo Frames (Linux)

To automatically have Jumbo Frame support enabled (whenever the `sxge` driver is loaded), specify the MTU parameter in the `sxge` device's corresponding `ifcfg` file.

1. **Set the MTU parameter in the corresponding `ifcfg` file.**

   It is `ifcfg-eth2` for the examples in this document. For example:

```
host #> cat /etc/sysconfig/network-scripts/ifcfg-eth2
# Oracle NEM 40GbE
DEVICE=eth2
BOOTPROTO=static
HWADDR=00:14:4F:29:01:1D
IPADDR=10.1.10.156
NETMASK=255.255.255.0
MTU=9124
ONBOOT=no
```

2. **Restart the interface.**

   a. **If the driver is currently installed and running, take the interface down.**

```
host #> ifdown eth2
```

   b. **Bring the interface up.**

```
host #> ifup eth2
```

3. **Verify that MTU value for the sxge device is as specified.**

```
host #> ifconfig eth2
  eth2      Link encap:Ethernet  HWaddr 00:14:4F:29:00:1D
   inet addr:10.1.10.156  Bcast:10.1.11.255  Mask:255.255.255.0
   inet6 addr: fe80::214:4fff:fe29:1d/64 Scope:Link
   UP BROADCAST RUNNING MULTICAST  MTU:9124  Metric:1
   RX packets:0 errors:0 dropped:0 overruns:0 frame:0
   TX packets:34 errors:0 dropped:0 overruns:0 carrier:0
   collisions:0 txqueuelen:1000
   RX bytes:0 (0.0 b)  TX bytes:8532 (8.3 KiB)
   Memory:f9000000-fa000000
```

Related Information

- "Temporarily Configure Jumbo Frames (Linux)" on page 56

▼ Enable Jumbo Frames (Windows)

By default, the Jumbo Frames feature is disabled. When the Jumbo Frames feature is enabled, the miniport driver becomes capable of handling packet sizes of up to 9216 bytes. Bigger size packets are split into supported sizes by the driver before they are handled. When the feature is disabled, the driver handles packet sizes up to 1518 bytes.

1. **In the Device Manager, right-click on the Sun Blade 6000 40 GbE Networking Controller and select Properties.**
2. **Select the Advanced Tab, then select Jumbo Frames.**
3. **Select the enabled option, and click OK to enable the Jumbo Frames feature.**

Configuring Jumbo Frames (VMware ESX)

Use the following guidelines when configuring Jumbo Frames for VMware ESX servers:

- Any packet larger than 1500 MTU is a Jumbo Frame. ESX supports frames up to 9kB (9000 Bytes). Jumbo frames are limited to data networking only (virtual machines and the VMotion network) on ESX 3.5.
- Even though it is possible to configure Jumbo Frame for iSCSI Network, it is unsupported.
- Jumbo Frames must be enabled for each vSwitch or VMkernel interface through the command-line interface on your ESX Server 3 host.
- If you want to enable ESX Server to send larger frames out onto the physical network, the network must support Jumbo Frames end-to-end for Jumbo Frames to be effective.
- For more information on Jumbo Frames configuration for ESX, refer to: *ESX Configuration Guide for ESX 4.0* available at:

  `http://www.vmware.com/support/pubs/vs_pages/vsp_pubs_esx40_vc40.html`.

These topics describe how to configure Jumbo Frames:

- "Create a Jumbo Frames-Enabled vSwitch (VMWare)" on page 60

- "Create a Jumbo Frames-Enabled VMkernel Interface (VMWare" on page 60

▼ Create a Jumbo Frames-Enabled vSwitch (VMWare)

1. **Log in directly to your ESX Server host console.**
2. **Set the MTU size to the largest MTU size among all the virtual network adapters connected to the vSwitch.**

   Use the following command to set the MTU size for the vSwitch:

```
# esxcfg-vswitch -m MTU vSwitch
```

   where *MTU* is the MTU size and *vSwitch* is the name assigned to the vSwitch.

   This command sets the MTU for all uplinks on the vSwitch.

3. **Display a list of vSwitches on the host by typing:**

```
# esxcfg-vswitch -l
```

4. **Verify that the configuration of the vSwitch is correct.**

Related Information

- "Create a Jumbo Frames-Enabled VMkernel Interface (VMWare" on page 60

▼ Create a Jumbo Frames-Enabled VMkernel Interface (VMWare

1. **Log in directly to your ESX Server host console.**
2. **Create a VMkernel connection with Jumbo Frame support by typing:**

```
# esxcfg-vmknic -a -i ip address -n netmask -m MTU port-group-name
```

   where *ip address* is the ip address of the server, *netmask* is the netmask address, *MTU* is the MTU size, and *port-group-name* is the name assigned to the port group.

3. **Display a list of VMkernel interfaces by typing:**

```
# esxcfg-vmknic -l
```

4. **Verify that the configuration of the Jumbo Frame-enabled interface is correct.**

Note – ESX Server supports a maximum MTU size of 9000.

Related Information

- "Create a Jumbo Frames-Enabled vSwitch (VMWare)" on page 60

Configuring Static Link Aggregation

Refer to your OS documentation for commands and information about aggregating links.

Using static aggregation, you can aggregate the two NEM ports using OS-specific commands. Dynamic aggregation is not possible.

Link aggregation for the NEM ports is not configurable on either VMWare OS or Windows OS.

These topics provide examples of configuring static link aggregation:

- "Example (Oracle Solaris 11 OS)" on page 63
- "Example (Oracle Solaris 10 OS)" on page 64
- "Example (Linux OS)" on page 64

Example (Oracle Solaris 11 OS)

Before creating aggregation, if the sxge interfaces are plumbed, unplumb them using the `ipadm` command.

Refer to your OS documentation for commands and information about aggregating links.

The following example aggrgates two sxge ports into an aggregated device named aggr22 and assigns an IP address to it.

```
# dladm show-phys
LINK            MEDIA               STATE      SPEED  DUPLEX    DEVICE
net0             Ethernet            up         10000  full      sxge0
net1             Ethernet            up         10000  full      sxge1
net2             Ethernet            unknown    0      unknown   igb0
net3             Ethernet            up         1000   full      igb1
net4            Ethernet            up        10     full     usbecm0
```

```
# dladm create-aggr -l net0 -l net1 aggr22
# ipadm create-ip aggr22
# ipadm create-addr -T static -a local=192.168.1.138/24 aggr22
aggr22/v4
```

```
# ipadm show-addr
ADDROBJ           TYPE     STATE        ADDR
lo0/v4            static   ok           127.0.0.1/8
net3/v4           static   ok           10.134.167.138/24
net4/v4           static   ok           169.254.182.77/24
aggr22/v4         static   ok           192.168.1.138/24
lo0/v6            static   ok           ::1/128
net3/v6          addrconf ok          fe80::221:28ff:fea7:9895/10
```

Example (Oracle Solaris 10 OS)

Before creating aggregation, if the sxge interfaces are plumbed, unplumb them using the `ifconfig` command.

Refer to your OS documentation for commands and information about aggregating links.

The following example aggregates two sxge ports.

```
# dladm show-dev
usbecm0         link: up        speed: 10    Mbps       duplex: full
sxge0           link: unknown   speed: 10000 Mbps       duplex: full
sxge1           link: unknown   speed: 10000 Mbps       duplex: full
igb0            link: up        speed: 1000  Mbps       duplex: full
igb1          link: unknown   speed: 0     Mbps      duplex: unknown
```

```
# dladm create-aggr -d sxge0 -d sxge1 22
# ifconfig aggr22 plumb up
# ifconfig aggr22 192.168.1.172/24 up
```

Example (Linux OS)

Refer to your OS documentation for commands and information about aggregating links.

The following example creates a bond between sxge ports.

Use these commands to ensure that the bond module is loaded.

```
# lsmod | grep bonding
# modprobe bonding
```

Create a bond device comprising of two sxge ports (eth2 and eth3 in the example).

```
# ethtool -i eth2
driver: sxge
version: 0.12202011
firmware-version:
bus-info: 0000:a2:00.0
# ethtool -i eth3
driver: sxge
version: 0.12202011
firmware-version:
bus-info: 0000:a3:00.0
```

```
# ifconfig bond0 up
# ifenslave bond0 eth2 eth3
# ifconfig bond0 192.168.1.142/24 up
```

To check the slave interfaces of bond0:

```
# cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.6.0 (September 26, 2009)
Bonding Mode: load balancing (round-robin)
MII Status: up
MII Polling Interval (ms): 0
Up Delay (ms): 0
Down Delay (ms): 0
Slave Interface: eth2
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:21:28:c6:f2:ee
Slave queue ID: 0
Slave Interface: eth3
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:21:28:c6:f3:20
Slave queue ID: 0
```

Booting an x86 Server Over the Network

To perform a network boot via the NEM, make sure that the NEM and the boot server are on the same LAN. If you are loading an OS over the network, you have to configure the boot server so that it knows which OS to download to the blade and which IP address to use for the NEM's MAC.

- "Booting the NEM on an x86 Blade Server Over the Network" on page 67
- "Boot a SPARC Blade Over the Network" on page 69

Booting the NEM on an x86 Blade Server Over the Network

The NEM is designed to boot from the first detected PXE image. If there are two NEMs in the chassis and both are connected to a PXE server, you must use either the MAC address assignment booting method to boot from the correct NEM or the ESC key to jump to the next NEM PXE image.

If you want to use the F8 network booting method, follow the instructions in "Use the F8 Key to Boot Off the PXE Server" on page 69.

If you want to use the F12 network booting method, follow the instructions in both of the following procedures:

- "Set Up the BIOS for Booting From the NEM" on page 67
- "Use the F12 Key to Initiate the Network Boot" on page 68

▼ Set Up the BIOS for Booting From the NEM

1. **Power on the host blade server.**

2. **When the BIOS boot screen is displayed, press F2 (or Ctrl-E from a remote console) to enter the BIOS Setup Utility.**
3. **Navigate to the Boot tab.**
4. **Select Boot Device Priority and press Enter.**
5. **Select the first boot device position and presss Enter.**

   A drop-down list is displayed with the available boot devices listed.
6. **From the drop-down list, select** `PXE:ORACLE 40GNem v1.9 GPL etherboot.org` **as the first boot device.**
7. **Press F10 to Save and Exit the BIOS Setup Utility.**

   The host blade server reboots.

Note – The NEM is designed to boot from the first detected NEM port. If you have two NEMs in the chassis and both are connected to a PXE server, you must use either the MAC address assignment booting method to boot from the correct NEM or the ESC key to jump to the next NEM port.

---

8. **Connect the Ethernet cable to the NEM Ethernet port.**

   See the driver installation topic that corresponds with the operating system installed on the blade module.
9. **Boot the system from the PXE server to load an OS or install an OS.**

Related Information

- "Use the F12 Key to Initiate the Network Boot" on page 68
- "Use the F8 Key to Boot Off the PXE Server" on page 69

▼ Use the F12 Key to Initiate the Network Boot

1. **Power on the host server blade.**
2. **Follow the instructions in** "Set Up the BIOS for Booting From the NEM" on page 67 **to set up the NEM as the first network boot device.**
3. **Press F12 to boot from the network.**

Note – If the NEM does not detect a PXE server, the CPU host blade will attempt to boot from the on-board 1 Gbps NIC devices.

Related Information

- "Use the F8 Key to Boot Off the PXE Server" on page 69

▼ Use the F8 Key to Boot Off the PXE Server

1. **Power on the host server blade.**
2. **When prompted, press F8 to display the boot choices list.**
3. **Select the** `PXE:ORACLE 40GNem v1.12 GPL etherboot.org` **NIC from the boot list.**
4. **Press Enter to boot from the PXE server.**

Note – If the NEM does not detect a PXE server, the CPU host blade will attempt to boot from the on-board 1 Gbps NIC devices.

Related Information

- "Use the F12 Key to Initiate the Network Boot" on page 68

▼ Boot a SPARC Blade Over the Network

Use the following procedure to boot over the network using the NEM with a SPARC blade server.

1. **Add a MAC address-IP-address pair to the** `/etc/ethers` **and** `/etc/hosts` **files on the boot server.**

   This pair enables the boot server to answer RARP calls from the blade for resolving the IP address for the NEM's local MAC address.

2. **Add a client script on the boot server to specify which OS should be loaded to the SPARC blade.**

   Refer to the installation documentation for the operating system you want to install.

3. **After you have configured the boot server, do one of the following:**
   - If the SPARC blade is not running Oracle Solaris OS, log on to the console of the SPARC blade, power the blade on, and boot to the OpenBoot `ok` prompt.

- If the blade is already running Oracle Solaris OS, type the following in a terminal window to get to the `ok` prompt:

```
sync;sync;halt
```

4. **Determine all the network interfaces available for the blade to use.**

   Some of them will be the NEM's ports. For example:

```
{0} ok show-nets
a) /pci@400/pci@2/pci@0/pci@c/network@0
b) /pci@400/pci@2/pci@0/pci@2/network@0,1
c) /pci@400/pci@2/pci@0/pci@2/network@0
d) /pci@400/pci@1/pci@0/pci@c/pci@0/pci@3/network@0
e) /pci@400/pci@1/pci@0/pci@c/pci@0/pci@2/network@0
q) NO SELECTION
Enter Selection, q to quit:
```

5. **Change the directory to the applicable NEM port path.**

   Example:

```
{0} ok cd /pci@400/pci@1/pci@0/pci@c/pci@0/pci@2/network@0
```

6. **Check its properties:**

```
{0} ok .properties
```

   If it is a Sun Blade 6000 Virtualized 40 GbE Network Express Module, the output will resemble the following.

```
{0} ok .properties
assigned-addresses
                  82400010 00000000 10300000 00000000 00004000
                  82400018 00000000 10400000 00000000 00100000
                  82400030 00000000 10500000 00000000 00100000
local-mac-address  00 21 28 c6 fe 9e
reg               00400000 00000000 00000000 00000000 00000000
                  03400010 00000000 00000000 00000000 00004000
                  03400018 00000000 00000000 00000000 00100000
                  02400030 00000000 00000000 00000000 00100000
version           Sun Blade 6000 Virtualized 40GbE NEM FCode 0.92
                  11/05/03
board-model       511-1596-03
model             SUNW,pcie-sxge
compatible        pciex108e,2078.108e.6277.1
```

```
                    pciex108e,2078.108e.6277
                    pciex108e,2078.1
                    pciex108e,2078
                    pciexclass,020000
                    pciexclass,0200
address-bits        00000030
max-frame-size      00002400
network-interface-type ethernet
device_type         network
name                network
fcode-rom-offset    000010200
interrupts          00000001
cache-line-size     00000010
class-code          00020000
subsystem-id        00006277
subsystem-vendor-id 0000108e
revision-id         00000001
device-id           00002078
vendor-id           0000108e
```

7. **Type `device-end`, then perform the netboot.**

```
{0} ok device-end
```

```
{0} ok boot /pci@400/pci@0/pci@9/pci@0/pci@1/network@0
```

If DHCP will be used for resolving the IP address at the beginning of the boot process, add dhcp at the end of the boot command, as in the following example:

```
{0} ok  boot /pci@400/pci@0/pci@9/pci@0/pci@1/network@0:dhcp
```

If successful, output similar to the following is displayed on the blade's console.

```
Boot device: /pci@400/pci@0/pci@9/pci@0/pci@1/network@0:dhcp  File
and args:
SunOS Release 5.10 Version Generic_141444-09 64-bit
Copyright 1983-2009 Sun Microsystems, Inc.  All rights reserved.
Use is subject to license terms.
NOTICE: sxge1: link up, 10000 Mbps, full duplex
Configuring devices.
Using DHCP for network configuration information.
Reading ZFS config: done.
Setting up Java. Please wait...
Serial console, reverting to text install
Beginning system identification...
Searching for configuration file(s)...
Using sysid configuration file /sysidcfg
Search complete.
Discovering additional network configuration...
Completing system identification...
Starting remote procedure call (RPC) services: done.
NIS domain name is
System identification complete.
Starting Oracle Solaris OS installation program...
Executing JumpStart preinstall phase...
Searching for SolStart directory...
Checking rules.ok file...
Using begin script: install_begin
Using finish script: patch_finish
Executing SolStart preinstall phase...
Executing begin script "install_begin"...
Begin script install_begin execution compl
```

The net installation menu is displayed. The Oracle Solaris OS installation program is divided into a series of short sections where you'll be prompted to provide information for the installation. At the end of each section, you can change the

selections you've made before continuing.

Related Information

- "Set Up the BIOS for Booting From the NEM" on page 67

Installing Drivers on Oracle Solaris OS SPARC or x86 Platform

Oracle Solaris OS 10/09 and newer supports the NEM. These topics explain how to configure the `sxge` driver on a Oracle Solaris OS. The `sxge` GigabitEthernet driver is a multi-threaded, loadable, cloneable, GLD-based STREAMS driver.

- "Download the sxge Driver" on page 75
- "Configure the Network Host Files" on page 76
- "Configure the sxge Device Driver Parameters" on page 78

Related Information

- "Installing Drivers on a Linux Platform" on page 81
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

▼ Download the `sxge` Driver

The `sxge` device driver is not included with versions of Oracle Solaris OS 10, so it is necessary to download and install the driver. The driver is included in versions of Oracle Solaris 11 OS.

1. **If you are using a version of Oracle Solaris 10, download the newest patch that includes the driver from My Oracle Support. See** "Updating the NEM Firmware" on page 143**.**
2. **Install the sxge driver.**
   a. **Unzip the patch** `p12918641_100_Generic.zip` **or newer.**
   b. **Locate the firmware package, which is** `firmware/image-1.010.pkg` **or newer.**

3. **Configure the sxge driver.**

Related Information

- "Configure the Network Host Files" on page 76
- "Configure the sxge Device Driver Parameters" on page 78

▼ Configure the Network Host Files

This topic describes how to configure the network host files after you verify that the `sxge` driver is on your system.

1. **At the command line, use the `grep` command to search the `/etc/path_to_inst` file for `sxge` interfaces.**

   Example for Oracle Solaris OS SPARC:

```
# grep sxge /etc/path_to_inst
"/pci@400/pci@2/pci@0/pci@c/pci@0/pci@2/network@0" 0 "sxge"
"/pci@400/pci@2/pci@0/pci@c/pci@0/pci@3/network@0" 1 "sxge"
```

   Example for Oracle Solaris OS x86:

```
# grep sxge /etc/path_to_inst
"/pci@0,0/pci8086,340e@7/pci108e,2070@0/pci108e,2070@2/pci108e,62
77@0" 0 "sxge"
"/pci@0,0/pci8086,340e@7/pci108e,2070@0/pci108e,2070@3/pci108e,62
77@0" 1 "sxge"
```

   In these examples for SPARC and x86, respectively, the device instance is from a NEM installed in the chassis. The instance number is shown in *italics*.

2. **Set up the NEM's `sxge` interface.**

   Use the `ifconfig` command to assign an IP address to the network interface. Type the following at the command line, replacing *ip-address* and *netmask-address* with the NEM's IP address and the netmask address:

```
# ifconfig sxge0 plumb ip-address netmask netmask-address broadcast + up
```

   Refer to the `ifconfig(1M)` man page and the Oracle Solaris OS documentation for more information.

3. **(Optional) For a setup that remains the same after you reboot, create an** `/etc/hostname.sxge`*number* **file, where** *number* **is the instance number of the** `sxge` **interface you plan to use.**

   For example, to use the NEM's `sxge` interface in the Step 1's sample output, create an `/etc/hostname.sxge`*x* file, where *x* is the number of the sxge interface. If the instance number were `1`, the filename would be `/etc/hostname.sxge1`.

   Follow these guidelines for the host name:

   - The `/etc/hostname.sxge`*number* file must contain the host name for the appropriate `sxge` interface.
   - The host name must be different from the host name of any other interface. For example: `/etc/hostname.sxge0` and `/etc/hostname.sxge1` cannot share the same host name.
   - The host name must have an IP address listed in the `/etc/hosts` file.

     The following example shows the `/etc/hostname.sxge`*number* file required for a system named `zardoz-c10-bl1`.

```
# cat /etc/hostname.sxge0
zardoz-c10-bl1
```

4. **Create an appropriate entry in the** `/etc/hosts` **file for each active** `sxge` **interface.**

   For example:

```
# cat /etc/hosts
#
# Internet host table
#
127.0.0.1     localhost
129.168.1.29 zardoz-c10-bl1
```

Related Information

- "Configure the sxge Device Driver Parameters" on page 78
- "Installing Drivers on a Linux Platform" on page 81
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

▼ Configure the sxge Device Driver Parameters

The sxge device driver controls the NEM Ethernet interfaces. You can manually set the sxge driver parameters to customize each device in your system.

The sxge device driver configuration file is in the following directory:

```
/kernel/drv/sxge.conf
```

1. **Change a parameter by uncommenting the line in the `/kernel/drv/sxge.conf` file and providing a new value.**

   The following is an example of the content in the `/kernel/drv/sxge.conf` file. All parameters are listed and explained in this file. The default values are loaded when the `sxge` driver is started.

```
# cat /kernel/drv/sxge.conf
#
#
# driver.conf file for Sun 10/40Gb Ethernet Driver (sxge)
#
#
#------- Jumbo frame support ----------------------------------
# To enable jumbo support,
# accept-jumbo = 1;
#
# To disable jumbo support,
# accept-jumbo = 0;
#
# Default is 0.
#
#
#------- FMA Capabilities -----------------------------------
#
# Change FMA capabilities to non-default
#
# DDI_FM_NOT_CAPABLE 0x00000000
# DDI_FM_EREPORT_CAPABLE 0x00000001
# DDI_FM_ACCCHK_CAPABLE 0x00000002
# DDI_FM_DMACHK_CAPABLE 0x00000004
# DDI_FM_ERRCB_CAPABLE 0x00000008
#
# fm-capable = 0xF;
#
# default is DDI_FM_EREPORT_CAPABLE | DDI_FM_ERRCB_CAPABLE = 0x5
```

   In the following example, the FMA functionality is disabled.

```
fm-capable = 0x0;
```

2. **For the new parameters to be effective, reload the `sxge` driver or reboot the system.**

Related Information

- "Installing Drivers on a Linux Platform" on page 81
- "Installing Drivers on a Windows Platform" on page 103

- "Installing Drivers on a VMware ESX Server Platform" on page 113

Installing Drivers on a Linux Platform

- "Downloading, Installing, and Removing Drivers" on page 81
- "Configuring the Network Interface" on page 90
- "Device Checking and Testing" on page 97
- "Configure the Syslog Parameter" on page 100

Downloading, Installing, and Removing Drivers

This topic explains how to download, install, and remove the `sxge` and `sxgevf` Linux drivers.

The `sxge` 40Gbe/10Gbe driver is a multi-threaded and loadable driver supporting up to four transmit and four receive channels per function by default, in simultaneous (parallel) operation, taking advantage of up to eight CPUs/threads to spread the 40Gbe/10GbE network traffic and increase overall network throughput.

The `sxge` driver enables and disables the capability via API provided by SR-IOV core. If the hardware has SR-IOV capability, loading its PF driver would enable it and all VFs associated with the PF. Allocation of the VF can be dynamically controlled by the PF via registers encapsulated in the capability.

By default, this feature is not enabled, and the PF behaves as a traditional PCIe device. Once this feature is turned on, each VF's PCI configuration space can be accessed by its own bus, device, and function number (routing ID).

Each VF has PCI memory space, which is used to map its register set. The VF device driver operates on the register set, so that it can be functional and appear as a PCI device.

By default, blades are in traditional PCI-E device mode, and the `sxge` driver attaches to that device.

When a blade is in SR-IOV mode, the device is displayed as a PF device to the root domain (or `dom0`), and as a VF device attaching to the `sxgevf` driver in root and/or the guest domains.

The VFs are treated as hot-plugged PCI devices in the kernel, so that they are able to work in the same way as real PCI devices.

This topic covers the following procedures:

- "Download and Install the Drivers" on page 82
- "Remove the Drivers" on page 88

▼ Download and Install the Drivers

The following instructions apply to both drivers, `sxge` and `sxgevf`. When downloading and installing, specify which driver you want.

To use the `sxgevf` driver, you need to have the `sxge` driver already installed.

1. **Use the `ifconfig` command to obtain a list of the current Ethernet network interfaces.**

```
host #> ifconfig -a | grep eth
eth0      Link encap:Ethernet  HWaddr 00:14:4F:CA:15:68
eth1      Link encap:Ethernet  HWaddr 00:14:4F:CA:15:69
```

   After you have installed the NEM and installed and loaded a driver, a new `eth` device is displayed. This is the `eth` device for the NEM.

2. **Access the sxge drivers through one of the following methods:**

   - Locate and copy the `sxge` device driver `.zip` or `sxgevf` device driver `.zip` file from the Tools and Drivers CD for the server module operating system on which you want to install the driver.

     For Oracle Enterprise Linux, use the corresponding Red Hat Enterprise Linux driver.

     The drivers are located in the following directories:

     `/linux/drivers/rhel5.4`

     `/linux/drivers/sles10sp3`

     `/linux/drivers/sles11`

Note – Linux source files are available at `/linux/drivers/src`.

   - Download latest driver packages from the the following web site:

http://support.oracle.com

a. **At the My Oracle Support site, log in.**

   An account is required to download software. You must register if you do not have an account.

b. **Click the Patches & Updates tab.**

c. **Click "Product or Family" (Advance search).**

d. **In the Product field, type `Sun Blade 6000 Virtualized 40GbE NEM`.**

e. **In the Release field, select the newest release, then click the arrow next to the Sun Blade 6000 Virtualized 40GbE NEM folder.**

f. **Click the Search button.**

g. **Download the latest available firmware package, and extract it to an accessible folder on the network.**

3. **If you are downloading the `sxgevf` driver, make sure that SR-IOV mode is enabled before installing the driver. (See "Setting Blade Parameters" on page 119.)**

4. **Once the Linux drivers package is downloaded, uncompress and unpack the driver and other files.**

   In the following example, the download directory is `tge10`:

```
host #>mkdir tge10
host #>cp xxx/p12918641_100_Generic.zip
host #>cd tge10
host #>unzip p12918641_100_Generic.zip
host #>ls
readme.html
Sun_Blade_6000_Virtualized_40GbE_NEM_1_0_4.xml
Sun_Blade_6000_Virtualized_40GbE_NEM_1_0_4.zip
host #>unzip Sun_Blade_6000_Virtualized_40GbE_NEM_1_0_4.zip
host #>ls
host #>firmware fpga  license  linux  solaris  windows  release_note
```

5. **Select and install the appropriate OS driver package.**

   In the following example, RHEL6.0 Linux driver is selected for REHL version 6.0.

a. Navigate to the Linux directory that contains the driver package that is appropriate for your OS.

```
host #> cd linux/RHL6.0
host #>ls
kmod-sxge-1.0-0.x86_64.rpm sxge-debuginfo-1.0-0.x86_64.rpm
kmod-sxgevf-1.0-0.x86_64.rpm  sxgevf-debuginfo-1.0-0.x86_64.rpm
```

In the following example, the file `kmod-sxgevf-1.0-0.x86_64.rpm` is the package that will be installed.

If you are not sure which release your system is running, use the `lsb_release` command to show information about your host's operating system.

```
host #> lsb_release -a
LSB Version:
core-4.0-intel64:core-4.0-noarch:graphics-4.0-intel64:graphics-
4.0-noarch:printing-4.0-intel64:printing-4.0-noarch
Distributor ID: RedHatEnterpriseServer
Description: Red Hat Enterprise Linux Server release 6.0 (Santiago)
Release: 6.0
Codename: Santiago
host #> rpm -q kmod-sxgevf-1.0-0.x86_64
package kmod-sxgevf-1.0-0.x86_64 is not installed
```

b. Identify the OS-specific subdirectory, and verify that the driver you want to install (`sxge` or `sxgevf`) is not currently installed.

Note – If the driver is already installed, uninstall the driver to avoid complications. See "Remove the Drivers" on page 88 for instructions. The update command (`rpm -u`) is not supported for updating the drivers.

---

c. Install the appropriate package (`.rpm`) file.

```
host #> rpm -ivh --nodeps kmod-sxge-1.0-0.x86_64.rpm
Preparing...
########################################### [100%]
   1:kmod-sxge
########################################### [100%]
post Install Done
```

```
host #> rpm -ivh --nodeps kmod-sxgevf-1.0-0.x86_64.rpm
```

```
Preparing...
########################################### [100%]
   1:kmod-sxgevf ########################################### [100%]
post Install Done
```

Note – To load and use the sxgevf driver, you must have the sxge driver already installed.

---

6. Load the driver.

a. Verify that the NEM is available to the system (that is, it is actively on the PCIe I/O bus).

The following command sample has been edited for brevity. Lines marked [....] denote output that has been removed.

```
host #> lspci
[...]
13:00.0 PCI bridge: Oracle Corporation Unknown Device 2070
14:02.0 PCI bridge: Oracle Corporation Unknown Device 2070
14:03.0 PCI bridge: Oracle Corporation Unknown Device 2070
15:00.0 Ethernet controller: Sun Microsystems Computer Corp.
Unknown device 207a (rev 01)
15:00.2 Ethernet controller: Sun Microsystems Computer Corp.
Unknown device 207b (rev 01)
15:00.3 Ethernet controller: Sun Microsystems Computer Corp.
Unknown device 207b (rev 01)
16:00.0 Ethernet controller: Sun Microsystems Computer Corp.
Unknown device 207a (rev 01)
16:00.1 Ethernet controller: Sun Microsystems Compute Corp.
Unknown device 207b (rev 01)
16:00.2 Ethernet controller: Sun Microsystems Computer Corp.
Unknown device 207b (rev 01)
16:00.3 Ethernet controller: Sun Microsystems Computer Corp.
Unknown device 207b (rev 01)
[...]
```

```
host #> lspci
[...]
19:00.0 PCI bridge: Sun Microsystems Computer Corp.
Unknown device 2070 (rev 01)
1a:02.0 PCI bridge: Sun Microsystems Computer Corp.
Unknown device 2070 (rev 01)
1a:03.0 PCI bridge: Sun Microsystems Computer Corp.
Unknown device 2070 (rev 01)
1b:00.0 Ethernet controller: Sun Microsystems Computer Corp.
Unknown device 2078 (rev 01)
```

```
1c:00.0 Ethernet controller: Sun Microsystems Computer Corp.
Unknown device 2078 (rev 01)
[...]
```

- Device code 0x2078 is the NEM device when used as a regular PCI device.
- Device code 0x207a is the PF (using `sxge` driver) NEM device when SR-IOV is enabled.
- Device code 0x207b is the VF (using sxgevf driver) NEM device when SR-IOV is enabled.

  Presence of this line indicates the NEM is visible and available to the system.

b. Manually load the driver.

```
host #> modprobe sxge or sxgevf
```

c. Verify the driver is loaded.

```
host #> lsmod | grep sxge
sxge 71304 0
host#> modinfo sxge
filename:       /lib/modules/2.6.18-238.el5/extra/sxge/sxge.ko
version:        0.06032011
license:        GPL
description:    SXGE ethernet driver
author:
srcversion:     73779771D7EA5154342838B
alias:          pci:v0000108Ed0000207Asv*sd*bc*sc*i*
alias:          pci:v0000108Ed00002078sv*sd*bc*sc*i*
depends:
vermagic:       2.6.18-238.el5 SMP mod_unload gcc-4.1
parm:           debug:SXGE debug level (int)
[etc.]
```

```
host #> lsmod | grep sxgevf
sxge (or sxgevf) 71304 0
host#> modinfo sxgevf
filename:      /lib/modules/2.6.18-238.el5/extra/sxgevf/sxgevf.ko
version:        0.06032011
license:        GPL
description:    SXGEVF ethernet driver
author:
srcversion:     73779771D7EA5154342838B
alias:          pci:v0000108Ed0000207Asv*sd*bc*sc*i*
alias:          pci:v0000108Ed00002078sv*sd*bc*sc*i*
depends:
```

```
vermagic:       2.6.18-238.el5 SMP mod_unload gcc-4.1
parm:           debug:SXGEVF debug level (int)
[etc.]
```

- If you get the output shown in the previous examples, the drivers are loaded into memory and actively running.
- If the **modprobe** command fails, you will see the following output:

```
host #> modprobe sxge or sxgevf
FATAL: Module sxge not found.
FATAL: Module sxgevf not found.
```

This output indicates that you probably installed the wrong driver version. Uninstall the driver, then install the correct package for your Linux release.

If you are running a custom or patched kernel, you might have to build a custom driver to match your custom kernel.

7. Identify the NEM `eth` device.

After the NEM has been correctly installed, and the driver has been successfully installed and loaded, the new NEM `eth` device will be visible.

a. Execute the following command to view the available `eth` devices.

```
host #> ifconfig -a | grep eth
eth0      Link encap:Ethernet  HWaddr 00:14:4F:CA:15:68
eth1      Link encap:Ethernet  HWaddr 00:14:4F:CA:15:69
eth2      Link encap:Ethernet  HWaddr 00:14:4F:29:00:1D
```

In this example, `eth0` and `eth1` were previously present. The `eth2` device is for a new NEM. You can identify each `eth` device (NEM0 or NEM1) by matching the Ethernet MAC address with the one you recorded and saved when you physically installed the NEM into the chassis in "Install a NEM" on page 34.

b. Make sure that the `eth2` driver is the correct Ethernet driver for the NEM.

```
host #> ethtool -i eth2
driver: sxge
version: 0.06032011
firmware-version:
bus-info: 0000:1b:00.0
driver: sxgevf
version: 0.01202011
firmware-version:
bus-info: 0000:1c:00.0
```

c. For more detail on `eth3`, use the `ifconfig` command.

```
host #> ifconfig eth2
Link encap:Ethernet  HWaddr FA:CE:00:00:FC:0F
BROADCAST MULTICAST  MTU:1500  Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
Interrupt:194
```

The `eth2` device is active and available to the system, but has yet to be configured (assigned an IP address). See "Configuring the Network Interface" on page 90 for details on configuring the NEM for the Linux OS.

8. Load the VF driver to the guest OS.

Refer to the Oracle VM Server for x86 XEN documentation.

`http://docs.oracle.com/cd/E15458_01/index.htm`

Related Information

- "Remove the Drivers" on page 88
- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

▼ Remove the Drivers

Removing the Linux drivers is a two-step process.

1. **Use the `modprobe -r` command to unload the `sxge` and `sxgevf` drivers at any time, without actually uninstalling the drivers.**

```
host #> lsmod | grep sxge
sxge 168784 0
host #> modprobe -r sxge
#> lsmod | grep sxge
#>
```

```
host #> lsmod | grep sxgevf
sxgevf 168784 0
host #> modprobe -r sxgevf
#> lsmod | grep sxgevf
#>
```

   Once unloaded, you can manually load the drivers again, using the **`modprobe`** command. The drivers have not been uninstalled.

2. **Uninstall the drivers.**

   This command permanently removes the drivers and all related files from the system (you will have to reinstall the drivers again before you can use the NEM):

```
host #> rpm -q kmod-sxge-1.0-0.x86_64
kmod-sxge-1.0-0

host #> rpm -e kmod-sxge-1.0-0
```

```
host #> rpm -q kmod-sxgevf-1.0-0.x86_64
kmod-sxgevf-1.0-0

host #> rpm -e kmod-sxgevf-1.0-0
```

Note – Uninstalling the drivers does not unload the drivers. If you elected to skip Step 1 (you did not unload the drivers), loaded drivers will remain active in memory, and the NEM will remain useable until the system resets and reboots. This behavior can vary with different Linux installations.

Related Information

- "Download and Install the Drivers" on page 82
- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

Configuring the Network Interface

Note – The information in this topic is intended as a guideline for configuring network interfaces for the drivers. For more detailed information, refer to the administrative documentation for the Linux version installed on your system.

Before using the NEM network interface, you must first configure the network interface. Use the `ifconfig(8)` command to control the primary network interface options and values for any given network device (such as `eth2` for the NEM as demonstrated in the installation topic). At a minimum, you must assign a network TCP IP address and netmask to each network interface.

- "Temporary Network Interface Configuration" on page 90
- "Permanent Network Interface Configuration" on page 91
- "Bring the Interface Online Manually" on page 91
- "Configure the Network Interface File Automatically for Red Hat Linux" on page 93
- "Configure the Network Interface Automatically for SUSE Linux" on page 95

Temporary Network Interface Configuration

To temporarily configure the NEM Ethernet interface (for example, to test it), use the `ifconfig` command.

By assigning an IP network address and corresponding IP network address mask, you can manually bring the interface fully online. This temporary manual configuration is not preserved across a system reboot.

Related Information

- "Permanent Network Interface Configuration" on page 91
- "Bring the Interface Online Manually" on page 91
- "Configure the Network Interface File Automatically for Red Hat Linux" on page 93
- "Configure the Network Interface Automatically for SUSE Linux" on page 95
- "Remove the Drivers" on page 88
- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Windows Platform" on page 103

■ “Installing Drivers on a VMware ESX Server Platform” on page 113

Permanent Network Interface Configuration

For the NEM network interface to be automatically configured (that is, on every system boot), define the network interface information in the network device database.

Linux maintains a separate network interface configuration file for each network interface in the system. This configuration file is used to automatically configure each network interface when the system is first booted. These configuration files are plain-text files, which may be created and edited with your favorite text editor, as well as using the Linux system-specific System Administration GUI.

■ “Bring the Interface Online Manually” on page 91
■ “Configure the Network Interface File Automatically for Red Hat Linux” on page 93
■ “Configure the Network Interface Automatically for SUSE Linux” on page 95

▼ Bring the Interface Online Manually

1. Assign both an IP address and a netmask to bring the interface online.

```
host #> ifconfig eth2 10.1.10.156 netmask 255.255.255.0
```

The system switches the device online automatically when it has the required information.

2. **Verify that the device is online by using the `ifconfig` command.**

```
host #> ifconfig eth2
eth2  Link encap:Ethernet  HWaddr 00:14:4F:29:00:1D
inet addr:10.1.10.156  Bcast:10.1.10.255  Mask:255.255.255.0
inet6 addr: fe80::214:4fff:fe29:1/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
RX packets:2 errors:0 dropped:0 overruns:0 frame:0
TX packets:27 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:300 (300.0 b)  TX bytes:7854 (7.6 KiB)
Memory:fb000000-fc000000
```

This example shows configuring the newly-installed NEM interface `eth2` to be IP address `10.1.10.156`, which is declared in what used to be known as a Class C (8-bit/255-node) LAN.

Note that the `inet addr` shows the following:

- TCP IPv4 address `10.1.10.156` as specified in the `ifconfig` command
- IPv6 address has been automatically derived (in this example, Linux is configured to support IPv6 network communications)
- Reported state now shows up

3. **Verify that RX and TX packet counters are increasing, showing active traffic being routed through the newly configured NEM `eth2` network interface.**

   See the `ifconfig(8)` man page for more details and other options on using the `ifconfig` command to configure Ethernet interfaces.

4. **Use the `route(8)` command to show the current networks.**

```
host #> route
Kernel IP routing table
Destination     Gateway      Genmask Flags Metric Ref    Use Iface
10.1.10.0     *          255.255.255.0   U     0      0        0 eth2
10.8.154.0    *          255.255.255.0   U     0      0        0 eth1
default       ban25rtr0d0  0.0.0.0         UG    0      0        0 eth1
```

Note – In this example, `10.1.10` LAN traffic is being routed through the newly-configured NEM `eth2` network interface.

---

5. **If you want to temporarily switch the network device back to an offline or quiescent state, use the `ifconfig down` command.**

```
host #> ifconfig eth2 down
host #> ifconfig eth2
inet addr:10.1.10.156  Bcast:10.1.10.255  Mask:255.255.255.0
eth2  Link encap:Ethernet  HWaddr 00:14:4F:29:00:1D
BROADCAST MULTICAST  MTU:1500  Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
Memory:fb000000-fc000000

host #> route
Kernel IP routing table
Destination  Gateway     Genmask      Flags  Metric Ref  Use Iface
10.8.154.0   *            255.255.255.0   U     0      0      0 eth1
default      ban25rtr0d0  0.0.0.0         UG    0      0      0 eth1
```

Note the `10.1.10` local area network via `eth2` is no longer available, but that the `eth2` network interface itself is still present (packet counters zero).

Related Information

- "Configure the Network Interface File Automatically for Red Hat Linux" on page 93
- "Configure the Network Interface Automatically for SUSE Linux" on page 95
- "Download and Install the Drivers" on page 82
- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

▼ Configure the Network Interface File Automatically for Red Hat Linux

You can configure the network interface for the Oracle or Red Hat Enterprise Linux platform by using the GUI or manually editing the configuration file.

This procedure describes how to manually edit the configuration file. For instructions on configuring the network interface using the GUI, refer to the documentation for your OS version at:

- Oracle Enterprise Linux (OEL): `http://docs.oracle.com/cd/E37670_01/`

- RHEL: http://docs.redhat.com/docs/en-US/index.html

For OEL or RHEL systems, the interface configuration files are named `ifcfg-eth`*n* (for example., `ifcfg-eth2` for the `eth2` network device as shown in preceding examples). They reside in the `/etc/sysconfig/network-scripts` system directory.

1. **Create a configuration file, as shown in the following example:**

```
host #> ls -l /etc/sysconfig/network-scripts/
total 392-rw-r--r-- 3 root root   116 Oct 10 12:40 ifcfg-eth0
-rw-r--r-- 3 root root   187 Oct 10 12:40 ifcfg-eth1
-rw-r--r-- 3 root root   127 Oct 21 16:46 ifcfg-eth2
-rw-r--r-- 1 root root   254 Mar  3  2008 ifcfg-lo
[...]

host #>cat /etc/sysconfig/network-scripts/ifcfg-eth2
# Oracle NEM 40GbE
DEVICE=eth2
BOOTPROTO=static
HWADDR=00:14:4F:29:00:1D
IPADDR=10.1.10.156
NETMASK=255.255.255.0
ONBOOT=no
```

   This sample eth2 `ifcfg` file was created by using a text editor. The first line `# Oracle NEM 40GbE` is a comment that is useful in keeping track of different files. For this particular example, `ONBOOT=no` is specified, which means that the network interface is not brought online automatically when the system is booted. Specifying `ONBOOT=yes` would be the normal configuration.

2. **Use the `ifconfig` command or the shorthand `ifup` script to bring the network interface online (up) for usage once the system has booted (to at least run level 3).**

```
host #> ifconfig eth2 up
```

   or

```
host #> ifup eth2
```

Related Information

- "Bring the Interface Online Manually" on page 91
- "Configure the Network Interface Automatically for SUSE Linux" on page 95
- "Download and Install the Drivers" on page 82

- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

▼ Configure the Network Interface Automatically for SUSE Linux

You can configure the network interface for the SLES platform either by using the GUI or manually editing the configuration file.

For instructions on configuring the network interface using the GUI, refer to the documentation for your SLES version at:
`https://www.suse.com/documentation/`.

For Novell systems, the interface configuration files are named `ifcfg-eth-id` (for example, `ifcfg-eth-id-00:14:4F:29:00:1D` for the NEM network device as used in preceding examples), and reside in the `/etc/sysconfig/network` system directory.

1. **Create a configuration file, similar to the following example.**

```
host #> ls -l /etc/sysconfig/network
total 88[...]
-rw-r--r-- 1 root root   271 Oct 29 18:00
ifcfg-eth-id-00:14:4f:29:00:1D
-rw-r--r-- 1 root root   245 Oct 29 18:00
ifcfg-eth-id-00:14:4f:80:06:ef
-rw-r--r-- 1 root root   141 Apr 21  2008 ifcfg-lo
[...]

host #> cat /etc/sysconfig/network/ifcfg-eth-id-00:14:4f:29:00:1D
BOOTPROTO='static'BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='10.1.10.150'
NAME='Sun Microsystems Ethernet controller'
NETMASK='255.255.255.0'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
UNIQUE='DkES.he1wLcVzebD'
USERCONTROL='no'
_nm_name='bus-pci-0000:88:00.0'
```

This sample `ifcfg` file was created using the Network Setup Method GUI. Regardless of which method you employ to maintain the network device configuration database, after the appropriate `ifcfg` file has been properly created, it will automatically be applied whenever the system boots. All matched network interfaces are automatically configured.

2. **Use the `ifconfig` command or the shorthand `ifup` script to bring the network interface online for use after the system has booted (to at least run level 3).**

```
host #> ifconfig eth2 up
```

or

```
host #> ifup eth2
```

When you manually edit one of the `ifcfg` files, you might need to invoke an explicit (manual) `ifdown/ifup` sequence to apply the new configuration (for example, changing the IP address or netmask, changing the MTU, and so on).

Related Information

- "Bring the Interface Online Manually" on page 91

- "Configure the Network Interface File Automatically for Red Hat Linux" on page 93
- "Download and Install the Drivers" on page 82
- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

Device Checking and Testing

Once you have the NEM network interface device properly configured, online, and active, there are several ways you can verify the network interface operation:

- `ifconfig`: Use the `ifconfig` command to see if the RX/TX packet counts are increasing. The TX packet count indicates that the local system network services (or users) are queueing up packets to get sent over that interface; the RX packet count indicates that externally-generated packets have been received on that network interface.
- `route`: Use the `route` command to check that traffic for the network interface's network is being routed to that interface. If there are multiple network interfaces connected to a given LAN, traffic may be directed to one of the other interfaces, resulting in a zero packet count on the new interface.
- `ping`: If you know the IP address of another node on the network, use the `ping(8)` command to send a network packet to that node and get a response back.

```
host 39 #> ping tge30
PING tge30 (10.1.10.30) 56(84) bytes of data.
64 bytes from tge30 (10.1.10.30): icmp_seq=1 ttl=64 time=1.37 ms
64 bytes from tge30 (10.1.10.30): icmp_seq=2 ttl=64 time=0.148 ms
64 bytes from tge30 (10.1.10.30): icmp_seq=3 ttl=64 time=0.112 ms
64 bytes from tge30 (10.1.10.30): icmp_seq=4 ttl=64 time=0.074 ms
64 bytes from tge30 (10.1.10.30): icmp_seq=5 ttl=64 time=0.161 ms

--- tge30 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4001ms
rtt min/avg/max/mdev = 0.074/0.373/1.372/0.500 ms
```

By default, `ping` sends one ping packet out each second until it is stopped (by typing ^C). A slightly more thorough test would be a ping flood test. For example:

```
host #> ping -f -i 0 -s 1234 -c 1000 tge30
PING tge30 (10.1.10.30) 1234(1262) bytes of data.
--- tge30 ping statistics ---
1000 packets transmitted, 1000 received, 0% packet loss, time
1849ms
rtt min/avg/max/mdev = 0.048/0.200/0.263/0.030 ms, ipg/ewma
1.851/0.198 ms
```

This example sends out 1,000 ping packets (containing 1,234 bytes of data each or over a megabyte total) as fast as the other side responds. Note the 0% packet loss indicates a functional and sound network connection.

Check the network interface again, using `ifconfig`, to look for any apparent problems.

```
host #> ifconfig eth2
eth2 Link encap:Ethernet  HWaddr 00:14:4F:29:00:1D
inet addr:10.1.10.150  Bcast:10.1.10.255  Mask:255.255.255.0
inet6 addr: fe80::214:4fff:fe29:1/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
RX packets:2993 errors:0 dropped:0 overruns:0 frame:0
TX packets:2978 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:3286970 (3.1 MiB)  TX bytes:3287849 (3.1 MiB)
Memory:fb000000-fc000000
```

Note that no error, dropped, overrun, frame, carrier, or collision events are reported. Some network errors are expected even in normal operation, but usually are insignificant relative to the packet counts.

- `ethtool`: If `ifconfig` reports an accumulation of errors, then an extremely detailed report of NEM traffic (including error counts) is available by using the `ethtool(8)` command.

  The following is an excerpt of the `sxge` detailed statistics output.

```
host #> ethtool -S eth2
NIC statistics:
tx_frames: 7113048
tx_bytes: 4294967264
tx_frame_cnt_ovl: 0
tx_byte_cnt_ovl: 0
rx_frames: 5366
rx_bytes: 416808
rx_drops: 0
```

```
rx_drop_bytes: 0
rx_mcasts: 4501
rx_bcasts: 762
rx_frames_cnt_ovl: 0
rx_byte_cnt_ovl: 0
rx_drop_byte_ovl: 0
rx_drop_cnt_ovl: 0
rx_mcast_frame_cnt_ovl: 0
rx_bcast_frame_cnt_ovl: 0
rx_link_up: 1
rx_link_down: 0
rx_link_state: 1
rx_channel: 0
rx_bytes: 1416672
rx_dropped: 0
rx_errors: 0
rx_hw_pktcnt: 10136
rx_hw_pktdrop: 0
rx_rbr_empty: 0
rx_fifo_error: 0
rx_rcr_shadow_full: 0
rx_channel: 1
rx_packets: 303428
rx_bytes: 13957688
rx_dropped: 0
rx_errors: 0
rx_hw_pktcnt: 303428
rx_hw_pktdrop: 0
rx_rbr_empty: 0
rx_fifo_error: 0
rx_rcr_shadow_full: 0
rx_channel: 2
rx_packets: 35460186685
rx_bytes: 36560119545336
rx_dropped: 0
rx_errors: 0
rx_hw_pktcnt: 1100447649
rx_hw_pktdrop: 4989944
rx_rbr_empty: 9321
rx_fifo_error: 0
rx_rcr_shadow_full: 0
rx_channel: 3
rx_packets: 31126
rx_bytes: 4354871
rx_dropped: 0
rx_errors: 0
rx_hw_pktcnt: 31126
rx_hw_pktdrop: 0
rx_rbr_empty: 0
```

```
rx_fifo_error: 0
rx_rcr_shadow_full: 0
tx_channel: 0
tx_packets: 35574371264
tx_bytes: 37744382014305
tx_errors: 0
tx_channel: 1
tx_packets: 0
tx_bytes: 0
tx_errors: 0
tx_channel: 2
tx_packets: 0
tx_bytes: 0
tx_errors: 0
tx_channel: 3
tx_packets: 0
tx_bytes: 0
tx_errors: 0
```

Related Information

- "Configure the Syslog Parameter" on page 100
- "Downloading, Installing, and Removing Drivers" on page 81
- "Configuring the Network Interface" on page 90
- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

▼ Configure the Syslog Parameter

By default, most Linux systems are configured to ignore (discard without logging) debug-level syslog messages. To see a Linux driver's debug messages when they are enabled, configure the `syslog(2)` facility to capture and record debug-level messages.

The `syslog` configuration is usually stored in the `/etc/syslog.conf` file (see the `syslog.conf(5)` man page), and will usually contain an entry such as the following (excerpted from an RHEL5.3 `/etc/syslog.conf` file).

```
# Log anything (except mail) of level info or higher.
# Don't log private authentication messages!
*.info;mail.none;authpriv.none;cron.none/var/log/messages
```

1. **Change the last line of the entry to enable debug-level messages to be captured and logged.**

   For example, change `info` to `debug`.

```
*.debug;mail.none;authpriv.none;cron.none /var/log/messages
```

   Changes made to `/etc/syslog.conf` will not take effect until `syslogd` is restarted (for example, automatically when the system is first booted).

2. **If you want to make `syslogd` re-read the `/etc/syslog.conf` file without rebooting the system, use the command:**

```
kill -SIGHUP `cat /var/run/syslogd.pid`
```

   This command notifies the currently-running syslog daemon to re-read its configuration file. See the `syslogd(8)` man page for more information.

Related Information

- "Configuring the Network Interface" on page 90
- "Download and Install the Drivers" on page 82
- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Windows Platform" on page 103
- "Installing Drivers on a VMware ESX Server Platform" on page 113

Installing Drivers on a Windows Platform

- "Installing Drivers on a Windows Platform" on page 103
- "Dynamic and NIC Teaming" on page 111

Related Information

- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Linux Platform" on page 81
- "Installing Drivers on a VMware ESX Server Platform" on page 113

Installing Drivers on a Windows Platform

This topic describes the process for installing network, VLAN, and enclosure drivers on an x86 (Intel only; AMD is not supported) server module that supports Windows Server 2008 (32/64-bit).

- "Installing and Uninstalling the Network Controller" on page 103
- "Installing and Uninstalling the Sun Blade 6000 40 GbE VLAN Driver (Optional)" on page 106
- "Adding and Removing VLANs" on page 108
- "Install the Enclosure Device on a Windows Server 2008 System" on page 110

Installing and Uninstalling the Network Controller

- "Install the Sun Blade 6000 40 GbE Network Controller" on page 104

- "Uninstall the Sun Blade 6000 40 GbE Network Controller" on page 105

▼ Install the Sun Blade 6000 40 GbE Network Controller

1. **Unzip the drivers from the Tools and Drivers CD to a local file on your system or a remote location for a remote installation.**

   The drivers are located at: `\windows\w2k8\Sun Blade 6000 40GbE Network Express Module.msi`

Note – This file includes the networking drivers for both Windows Server 2008 32-bit and 64-bit.

---

2. **Navigate to the `Sun Blade 6000 40GbE Network Express Module.msi` file on your local or remote system and double-click it to start the installation.**

   The License Agreement page is displayed.

3. **Select "I Agree," then click Next to begin the installation.**

   The Confirm Installation page is displayed.

4. **Click Next.**

   The Select Installation Folder page is displayed.

5. **Select "Everyone" or "Just Me" and click Next.**

   The Installing Oracle Network Drivers page is displayed.

   When the installation finishes, the Installation Complete page is displayed.

6. **Click Close.**

   The Sun Blade 6000 40 GbE Networking Controller device is installed and visible in the Network Adapters section of the Windows Device Manager.

7. **Verify that the Sun Blade 6000 40 GbE Network Driver is installed.**

   a. **Click Start, then click Run.**

   b. **Type `devmgmt.msc` in the Open field, and click OK.**

      The Device Manager window is displayed.

   c. **Click Network adapters and verify that the Sun Blade 6000 40 GbE Networking Controller is in the Network adapters list.**

      Two controllers will be displayed if you have one NEM installed, and four will be displayed if you have two virtualized NEMs.

   d. **Exit the Windows Device Manager.**

Related Information

- "Uninstall the Sun Blade 6000 40 GbE Network Controller" on page 105
- "Installing and Uninstalling the Sun Blade 6000 40 GbE VLAN Driver (Optional)" on page 106
- "Adding and Removing VLANs" on page 108
- "Install the Enclosure Device on a Windows Server 2008 System" on page 110

▼ Uninstall the Sun Blade 6000 40 GbE Network Controller

1. **Before uninstalling the networking controller, remove any VLANs that are installed on the controller. See** "Remove a VLAN" on page 109**.**
2. **Open the Control Panel.**
3. **Double-click Programs and Features.**

   The Programs and Features dialog is displayed.
4. **Select Oracle Network Drivers.**
5. **Click Uninstall.**

   The networking controller is removed, unless you did not remove all VLANs installed on the Ethernet port. If VLANs remain on the Ethernet port, an error message is displayed.
6. **If you receive an error message:**

   a. **Click OK.**

   b. **Remove all VLANs installed on the Ethernet port. See** "Remove a VLAN" on page 109**.**

   c. **Return to** Step 2**.**

Related Information

- "Install the Sun Blade 6000 40 GbE Network Controller" on page 104
- "Installing and Uninstalling the Network Controller" on page 103
- "Installing and Uninstalling the Sun Blade 6000 40 GbE VLAN Driver (Optional)" on page 106
- "Adding and Removing VLANs" on page 108
- "Install the Enclosure Device on a Windows Server 2008 System" on page 110

Installing and Uninstalling the Sun Blade 6000 40 GbE VLAN Driver (Optional)

The Sun Blade 6000 40 GbE VLAN driver enables creation of VLANs on top of physical Ethernet ports. Install this driver only if you intend to create one or more VLANs. The VLAN driver is bundled with the Sun Blade 6000 40 GbE Networking Controller.

- "Install the VLAN Driver" on page 106
- "Uninstall the VLAN Driver" on page 107

▼ Install the VLAN Driver

1. **Install the Sun Blade 6000 40 GbE Networking Controller. See** "Install the Sun Blade 6000 40 GbE Network Controller" on page 104**.**
2. **Open the Managed Network Connections dialog box.**
   a. **Open the Control Panel.**
   b. **Double-click the Network Sharing Center to open it.**
   c. **Select Managed Network Connections on the left side of the panel under "Tasks."**
3. **Right-click one of the Sun Blade 6000 10 GbE Networking Controllers and select Properties.**

   The Local Area Connection Properties dialog is displayed.
4. **Click Install.**

   The Network Feature Type dialog is displayed.
5. **Select Protocol and click Add.**

   The Select Network Protocol dialog is displayed.
6. **Click Have Disk.**

   The Install From Disk dialog is displayed.

7. **Click Browse and choose one of the following paths, depending on the version of Windows Server that you have installed:**

```
C:\Program Files\Oracle America Inc\Oracle Network Drivers\
Win_2008\vlan\w2k8 {x86,x64}
```

   Use the x86 path if you have a 32-bit architecture or the x64 path if you have a 64-bit architecture.

   The Locate File dialog is displayed.

8. **Select either of the information setup files displayed, and click Open.**

   The Select Network Protocol dialog is displayed.

9. **Click OK.**

   The Sun Blade 6000 40 GbE VLAN driver installs.

   After the installation, the Sun Multiple VLAN Protocol Driver is viewable in the Local Area Connection Properties dialog for each Sun Blade 6000 40 GbE Network Controller installed on the system.

Related Information

- "Uninstall the VLAN Driver" on page 107
- "Installing and Uninstalling the Network Controller" on page 103
- "Installing and Uninstalling the Sun Blade 6000 40 GbE VLAN Driver (Optional)" on page 106
- "Adding and Removing VLANs" on page 108
- "Install the Enclosure Device on a Windows Server 2008 System" on page 110

▼ Uninstall the VLAN Driver

1. **Open the Managed Network Connections dialog.**

   a. **Open the Control Panel.**

   b. **Double-click the Network Sharing Center to open it.**

   c. **Select Managed Network Connections on the left side of the panel under "Tasks."**

2. **Right-click one of the Sun Blade 6000 40 GbE Networking Controllers and select Properties.**

   The Local Area Connection Properties dialog is displayed.

3. **Select Sun Multiple VLAN Protocol Driver and click Uninstall.**

   The VLAN driver is now uninstalled, and all VLANs that were installed on the controller are removed.

Related Information

- "Install the VLAN Driver" on page 106
- "Installing and Uninstalling the Network Controller" on page 103
- "Installing and Uninstalling the Sun Blade 6000 40 GbE VLAN Driver (Optional)" on page 106
- "Adding and Removing VLANs" on page 108
- "Install the Enclosure Device on a Windows Server 2008 System" on page 110

Adding and Removing VLANs

- "Add a VLAN" on page 108
- "Remove a VLAN" on page 109

▼ Add a VLAN

1. **Open the Managed Network Connections dialog.**

   a. **Open the Control Panel.**

   b. **Double-click the Network Sharing Center to open it.**

   c. **Select Managed Network Connections on the left side of the panel under "Tasks."**

2. **Right-click one of the Sun Blade 6000 40 GbE Networking Controllers and select Properties.**

Note – Note which controller port you use to create the VLAN. You will need to use the same controller port to remove the VLAN.

---

3. **Select Sun Multiple VLAN Protocol, and click Properties.**

   The Sun Multiple VLAN Protocol dialog is displayed.

4. **Select Add a VLAN, and type the VLAN number.**

5. **Click OK.**

   A new Local Area Connection icon called Sun Blade 6000 Virtualized 40 GbE VLAN Virtual Miniport Driver appears in the Network Connections window and in the Network Adapters section of Windows Device Manager. The physical LAN port used to create the virtual LAN loses its properties, because its been virtualized.

Related Information

- "Remove a VLAN" on page 109
- "Installing and Uninstalling the Network Controller" on page 103
- "Install the Enclosure Device on a Windows Server 2008 System" on page 110
- "Installing Drivers on a Windows Platform" on page 103

▼ Remove a VLAN

1. **Open the Managed Network Connections dialog.**

   a. **Open the Control Panel.**

   b. **Double-click the Network Sharing Center to open it.**

   c. **Select Managed Network Connections on the left side of the panel under "Tasks".**

2. **Right-click the Sun Blade 6000 40 GbE Networking Controller used to create the VLAN that you are removing.**

   For example, in the "Add a VLAN" on page 108 procedure, the network controller port used to create the VLAN is Local Area Connection 3 controller port.

3. **Select Properties.**

   The Local Area Connection Properties dialog is displayed.

4. **Select Sun Multiple VLAN Protocol, and click Properties.**

   The Sun Multiple VLAN Protocol dialog is displayed.

5. **Select Remove a VLAN, and choose the VLAN number from the drop-down list.**

6. **Click OK.**

   The Sun Blade 6000 Virtualized 40 Gbe VLAN Virtual Miniport Driver icon for the VLAN selected is no longer visible in the Managed Network Connections and Adapter section of the Device Manager.

Related Information

- "Add a VLAN" on page 108

- "Installing and Uninstalling the Network Controller" on page 103
- "Install the Enclosure Device on a Windows Server 2008 System" on page 110
- "Installing Drivers on a Windows Platform" on page 103

▼ Install the Enclosure Device on a Windows Server 2008 System

1. **Open the Windows Device manager.**

   a. **Click Start, then click Run.**

   b. **Type `devmgmt.msc` in the Open field and click OK.**

      The Windows Device Manager is displayed.

2. **Click on System Devices and find the Generic SCSI Enclosure Device.**

   One enclosure device is displayed if you have one virtualized NEM installed, and two enclosure devices are displayed if you have two virtualized NEMs installed.

3. **Right-click on the Generic SCSI Enclosure Device, and select Update driver.**

   The "How do you want to search for drive software?" page is displayed.

4. **Click "Browse my computer for driver software."**

   The Browse for driver software on your computer page is displayed.

5. **Click "Let me pick from a list of device drivers on my computer," and click Next.**

   The Select the device driver you want to install for this hardware page is displayed.

6. **Click Have Disk.**

   The Install From Disk page is displayed.

7. **Click Browse and navigate to the directory that contains the NEM information (`lsinodrv.inf`) file.**

   The path on the Tools and Drivers CD is as follows:

   DVDdrive:`\windows\w2k8\64-bit\lsinodrv.inf`

   The selected NEM driver is displayed.

8. **Click Next to install the NEM driver.**

   The Software Installing page is displayed as the device is installed. When software installation is completed, a message indicating that Windows has successfully updated your driver software is displayed.

9. **Click Close.**
10. **Verify that the enclosure device has been installed.**
    a. **Click Start, then click Run.**
    b. **Type `devmgmt.msc` in the Open field, and click OK.**
       The Windows Device Manager is displayed.
    c. **Click on System Devices and verify that the NEM is in the System Devices list.**
11. **If you want install a second enclosure device, repeat the procedure.**

Related Information

- "Installing and Uninstalling the Network Controller" on page 103
- "Installing and Uninstalling the Sun Blade 6000 40 GbE VLAN Driver (Optional)" on page 106
- "Installing Drivers on a Windows Platform" on page 103

Dynamic and NIC Teaming

Dynamic and NIC teaming are not possible with the NEM windows driver, because the driver does not support the LACP protocol.

Installing Drivers on a VMware ESX Server Platform

- "Installing ESX Server Drivers" on page 113
- "Configure the NEM Network Adapters" on page 116
- "Uninstalling the Drivers" on page 116

Related Information

- "Installing Drivers on Oracle Solaris OS SPARC or x86 Platform" on page 75
- "Installing Drivers on a Linux Platform" on page 81
- "Installing Drivers on a Windows Platform" on page 103

Installing ESX Server Drivers

Note – Not all available drivers and ESX versions are compatible with the NEM. Refer to the *Sun Blade 6000 Virtualized 40 GbE Network Express Module Product Notes* for compatibility.

For information about VMware ESX and available downloads, refer to:

`http://www.vmware.com`

The following topics describe how to install ESX drivers on either a new or existing ESX server installation. Choose the procedure that matches your installation and interface preference:

- "Install a NEM Driver for a New ESX4.0 Installation" on page 114
- "Install a NEM Driver on an Existing ESX 4.0 Server" on page 114
- "Install a NEM Driver on an ESX/ESXi Server" on page 115

▼ Install a NEM Driver for a New ESX4.0 Installation

1. **Obtain the sxge driver from the Tools and Drivers CD at the following location:**

   `/vmware/drivers/esx4.0/Vmware-esx-drivers-net-sxge-1.2.2.11-00000.i386.rpm`

2. **Place the ESX installation DVD in the DVD drive of the host.**
3. **Accept the terms of the license agreement.**
4. **Select a keyboard type.**
5. **When prompted for Custom Drivers, select Yes to install custom drivers.**
6. **Click Add to eject the ESX installation DVD.**
7. **Place the driver CD in the DVD drive of the ESX host.**
8. **Select the driver module to import drivers to the ESX host.**
9. **Click Next to continue.**

   A dialog box is displayed with the following message: Load the system drivers.

10. **Click Yes.**
11. **After the drivers are installed, replace the driver CD with the ESX installation DVD when prompted to continue the ESX installation.**

Related Information

- "Configure the NEM Network Adapters" on page 116

▼ Install a NEM Driver on an Existing ESX 4.0 Server

1. **Obtain the sxge driver from Tools and Drivers CD at the following location:**

   `/vmware/drivers/esx4.0/Vmware-esx-drivers-net-sxge-1.2.2.11-00000.i386.rpm`

2. **Power on the ESX or ESXi host, and log in as administrator.**
3. **Place the driver CD in the CD-ROM drive of the ESX host.**
4. **Mount the driver CD.**

5. **Navigate to** *cd-mountpoint*`/offline-bundle/` **and locate the** `sxge-vmware-driver-4-1-2-2-7-offline_bundle-193789.zip` **file.**

6. **Execute command** `esxupdate` **to update the driver:**

```
# esxupdate --bundle=
sxge-vmware-driver-4-1-2-2-7-offline_bundle-193789.zip
```

For more details on `esxupdate`, download the pdf document at:
`http://www.vmware.com/pdf/vsphere4/r40/vsp_40_esxupdate.pdf`

Related Information

- "Configure the NEM Network Adapters" on page 116

▼ Install a NEM Driver on an ESX/ESXi Server

1. **Obtain the sxge driver from the Tools and Drivers CD at the following location:**

   `/vmware/drivers/esx4.0/Vmware-esx-drivers-net-sxge-1.2.2.11-00000.i386.rpm`

2. **Power on the ESX or ESXi host.**

3. **Place the driver CD in the CD-ROM drive of the host where either the vSphere CLI package is installed or vMA is hosted.**

4. **Mount the driver CD.**

5. **Navigate to** *cd-mountpoint***/offline-bundle/ and locate the** `sxge-vmware-driver-4-1-2-2-7-offline_bundle-193789.zip` **file.**

6. **Execute command esxupdate to update the driver:**

```
# vihostupdate conn-options --install --bundle
sxge-vmware-driver-4-1-2-2-7-offline_bundle-193789.zip
```

For more details on `vihostupdate`, download the pdf document at:
`http://www.vmware.com/pdf/vsphere4/r40/vsp_40_vcli.pdf`

Related Information

- "Configure the NEM Network Adapters" on page 116

▼ Configure the NEM Network Adapters

1. **Log in to the ESX host using Virtual Infrastructure Client GUI.**
2. **Click the Configuration Tab, then select Network Adapters from the Hardware list on the left side of the GUI.**
3. **Select the Sun Blade 6000 10 GbE Network Controller driver that you want to configure.**

   Configure networking using the Sun Blade 6000 10 GbE Network Controller as you would any other network interface.

   For more information on network configuration for ESX, refer to *ESX Configuration Guide for ESX 4.0* available at:

   `http://www.vmware.com/support/pubs/vs_pages/vsp_pubs_esx40_vc40.html`

Related Information

- "Installing ESX Server Drivers" on page 113

▼ Uninstalling the Drivers

1. **Run the following command to obtain the VIB ID:**

```
esxupdate query -vib-view
```

2. **Execute the following command to remove a driver:**

```
esxupdate -b vibID remove
```

Configuring Blade and Port Parameters and Managing the NEM

These topics describe how to use CMM to configure blade and port parameters and to manage the NEM through the CMM ILOM. Also included is information about using Oracle Integrated Lights Out Manager (Oracle ILOM).

- "Configuring and Viewing Parameters" on page 117
- "Configuring NEM for Oracle VM Para-Virtualization" on page 134
- "Managing the NEM Using CMM ILOM" on page 139

Related Information

- "Installing or Replacing a NEM" on page 31
- "Saving the Zoning Configuration to a Backup File" on page 25
- "Updating the NEM Firmware" on page 143

Configuring and Viewing Parameters

- "Defaults and Options" on page 118
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Saving User Configurations or Resetting to Defaults" on page 126
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

Related Information

- "Configure Para-Virtualization Without Live Migration" on page 134
- "Managing the NEM Using CMM ILOM" on page 139

Defaults and Options

The following defaults and options are applicable.

Item	Description
Bandwidth	10% is the default for each of the 10 blade slots
Blade number	0 is default
Blade	If port number is not provided in set blade command, the default is to execute for both ports.
Blade parameters	• bandwidth • PCI generation • promiscuous • SR-IOV • VLAN
Configuration file	Empty by default
Function	0 is default
MAC address	Adds, clears, and deletes MAC addresses
Port	Sets default for both ports. The port1 is used only in 10G.
Port parameters	• speed • status

Related Information

- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Saving User Configurations or Resetting to Defaults" on page 126
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

Setting Blade Parameters

The following topics provide syntax and examples for configuring blade parameters using the CMM. See "Defaults and Options" on page 118 for a list of the defaults and options.

- "Bandwidth" on page 119
- "PCI Generation" on page 120
- "MAC Address" on page 121
- "Promiscuous" on page 122
- "Receiver Side Scaling" on page 122
- "SR-IOV" on page 123
- "VLAN" on page 124

Bandwidth

Sets a blade's port bandwidth. Bandwidth controls the percentage of available throughput that a blade will be allocated when the pipe is full. Under conditions where the pipe is not full, blades can exceed the bandwidth number.

- The total bandwidth of all available server blades directly attached to a NEM cannot exceed 100%.
- Increasing bandwidth of a blade to 100% requires lowering bandwidth of one or more of other blades.

Syntax: set -b | blade 0..9 -p 0 | 1 -bandw 1..100

For example, to set blade 3 port 0 bandwidth to 15%:

```
set -b 3 -p 0 -bw 15
```

If the new setting exceeds the available bandwidth, a message will be displayed. Adjust the bandwidth value to be within the available amount.

Related Information

- "Defaults and Options" on page 118
- "PCI Generation" on page 120
- "MAC Address" on page 121
- "Promiscuous" on page 122
- "Receiver Side Scaling" on page 122
- "SR-IOV" on page 123

- “VLAN” on page 124
- “Setting Port Parameters” on page 125
- “Viewing Parameter Settings, Firmware Versions, and Configuration File” on page 127
- “Saving User Configurations or Resetting to Defaults” on page 126

PCI Generation

Sets the PCIe bus switch to gen1 or gen2. The new setting will take effect only after a blade reboot.

Syntax: set -b | blade 0..9 -g | gen 1 | 2

For example, to set blade 3 to gen1:

```
set -b 3 -g 1
```

Blade 3 will be in PCIe gen1 after the blade is rebooted or the PCIe is retrained.

Related Information

- “Defaults and Options” on page 118
- “Bandwidth” on page 119
- “MAC Address” on page 121
- “Promiscuous” on page 122
- “Receiver Side Scaling” on page 122
- “SR-IOV” on page 123
- “VLAN” on page 124
- “Setting Port Parameters” on page 125
- “Viewing Parameter Settings, Firmware Versions, and Configuration File” on page 127
- “Saving User Configurations or Resetting to Defaults” on page 126

MAC Address

Sets the MAC address for a blade's port. Also, this command provides options to delete a MAC address or to clear all user-added MAC addresses for each blade.

Each blade can have up to 32 MAC addresses (16 per port), because each blade has two ports.

Syntax: set -b | blade 0..9 [-f | func 0..3] -p | port 0 | 1 -m xx:xx:xx:xx:xx:xx [-d | delete]

To add a new MAC address for port 0 of blade 2's function 0:

```
set blade 2 func 0 port 0 -m 0:1:12:23:45:56
```

To clear all user-added MAC addresses for port 0 of blade 3's function 1:

```
set -b 3 -f 1 -p 0 -m 0:0:0:0:0:0
```

To delete a specific MAC address for port 0 of blade three's function 2:

```
set -b 3 -f 2 -p 0 -m 0:ab:cd:1:2:03 -d
```

Related Information

- "Defaults and Options" on page 118
- "Bandwidth" on page 119
- "PCI Generation" on page 120
- "Promiscuous" on page 122
- "Receiver Side Scaling" on page 122
- "SR-IOV" on page 123
- "VLAN" on page 124
- "Setting Port Parameters" on page 125
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127
- "Saving User Configurations or Resetting to Defaults" on page 126

Promiscuous

Note – Use this option only for testing and performing diagnostics in non-production environments.

Sets a blade's port to promiscuous, which enables a blade to look at packets from all blades, rather than just its own packets. Optionally, you can provide a function number. You must provide a port number.

Syntax: set -b | blade 0..9 [-p | port 0 | 1] [-f | func 0..3] -r | promisc on | off

For example, to set port 0 and function 0 of blade 4 to promiscuous:

```
set -b 4 -p 0 -f 0 -r on
```

Related Information

- "Defaults and Options" on page 118
- "Bandwidth" on page 119
- "PCI Generation" on page 120
- "MAC Address" on page 121
- "Receiver Side Scaling" on page 122
- "SR-IOV" on page 123
- "VLAN" on page 124
- "Setting Port Parameters" on page 125
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127
- "Saving User Configurations or Resetting to Defaults" on page 126

Receiver Side Scaling

Sets blade to enable RSS mode. By default, RSS mode is on (enabled) for both ports. Each port can be configured independently, for example, one port with RSS on and the other off.

When RSS is off (disabled), packets are processed by only one CPU. When RSS is on, packets are processed by multiple CPUs, subject to the 4-tuple of the packet received.

Note – RSS uses the TCP 4-tuple to load balance (hash) the packets to a CPU.

RSS load balancing with in-order delivery leads to higher throughput because of parallel packet processing.

Syntax: set -b | blade 0..9 -p | port 0 | 1 rss on | off

For example, to set blade 3 and port 0 to RSS mode:

```
set -b 3 -p 0 rss on
```

Related Information

- "Defaults and Options" on page 118
- "Bandwidth" on page 119
- "PCI Generation" on page 120
- "MAC Address" on page 121
- "Promiscuous" on page 122
- "SR-IOV" on page 123
- "VLAN" on page 124
- "Setting Port Parameters" on page 125
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127
- "Saving User Configurations or Resetting to Defaults" on page 126

SR-IOV

Sets a blade to SR-IOV mode. The parameter of both ports will be changed.

Syntax: set -b | blade 0..9 -sr | sriov on | off

For example, to set blade 3 to SR-IOV:

```
set -b 3 -sr on
```

Related Information

- "Defaults and Options" on page 118
- "Bandwidth" on page 119
- "PCI Generation" on page 120
- "MAC Address" on page 121
- "Promiscuous" on page 122

- “Receiver Side Scaling” on page 122
- “VLAN” on page 124
- “Setting Port Parameters” on page 125
- “Viewing Parameter Settings, Firmware Versions, and Configuration File” on page 127
- “Saving User Configurations or Resetting to Defaults” on page 126

VLAN

Sets VLAN filtering for a blade’s port. You can specify the function setting and the port number. By default, function 0 will be applied and both ports of the blade will be set to the VLAN filtering.

Note – NEM does not support VLAN tagging.

Also, you can use the VLAN command to set multiple VLANs by using the `-vr` option, where you assign a group of VLANs to a blade.

Syntax: set -b | blade 0..9 [-f | func 0..3] [-p | port 0 | 1] -v | vlan on | off a,b,c... [-vr x:y]

Syntax: set -b | blade 0..9 [-f | func 0..3] [-p | port 0 | 1] -v | vlan on | off [a,b,c...] -vr x:y

For example, to set blade 5’s port 0 for VLANs 20, 22, 345, 654, and 40 to 50:

```
set -b 5 -p 0 -v on 20,22,345,654 -vr 40:50
```

Related Information

- “Defaults and Options” on page 118
- “Bandwidth” on page 119
- “PCI Generation” on page 120
- “MAC Address” on page 121
- “Promiscuous” on page 122
- “Receiver Side Scaling” on page 122
- “SR-IOV” on page 123
- “Setting Port Parameters” on page 125
- “Viewing Parameter Settings, Firmware Versions, and Configuration File” on page 127
- “Saving User Configurations or Resetting to Defaults” on page 126

Setting Port Parameters

The following topics provide syntax and examples for setting port parameters for speed. See "Defaults and Options" on page 118 for a list of the defaults and options.

- "Speed" on page 125
- "Link Status" on page 125

Speed

Sets speed for both ports. The speed of both ports will be changed regardless of the port number provided. The new setting will take effect only after reboot.

Port 1 is supported only for external traffic in 10G mode. It is in loopback for 1G and 40G modes, which allows blades to communicate through Port 1.

Syntax: set -p | port 0 | 1 -s | speed 1 | 10 | 40

For example, to set speed as 10G on port 0:

```
set port 0 -s 10
```

Related Information

- "Defaults and Options" on page 118
- "Link Status" on page 125
- "Setting Blade Parameters" on page 119
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127
- "Saving User Configurations or Resetting to Defaults" on page 126

Link Status

Enables you to manually force a port link up or down. The command provides an option to default the port link status to the actual physical link status.

Syntax: set -p | port 0 | 1 -l | link up | down | auto

For example, to force down the port 0 link:

```
set port 0 link down
```

For example, to default the link to the physical link status:

```
set -p 0 -l auto
```

Related Information

- "Defaults and Options" on page 118
- "Speed" on page 125
- "Setting Blade Parameters" on page 119
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127
- "Saving User Configurations or Resetting to Defaults" on page 126

Saving User Configurations or Resetting to Defaults

Saves user configurations (commands) to the configuration file or returns settings to factory defaults. Saving configuration settings maintains the same state after the next reboot.

To maintain consistency, use save on (to open) and save off (to save) your configuration file.

If required, you can return the file to factory default settings by using the save default options.

The factory default settings are as follows:

- all blades in Gen2
- SR-IOV disabled
- all VLANs enabled

Syntax: set -c | conf save on | off | default

To start recording `set` commands:

```
set -c save on
```

To stop recording `set` commands and to save the recorded commands into a configuration file:

```
set -c save off
```

To reset to factory settings and clear the configuration file:

Note – If you want to save your configurations for use later, back up the file before returning the file to factory settings, because your configurations will be deleted. See "Saving the Zoning Configuration to a Backup File" on page 25.

```
set -c save default
```

Related Information

- "Defaults and Options" on page 118
- "Setting Blade Parameters" on page 119
- "Setting Blade Parameters" on page 119
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127
- "Managing the NEM Using CMM ILOM" on page 139

Viewing Parameter Settings, Firmware Versions, and Configuration File

The following topics describe how to view the current blade and port parameter settings using the show commands.

- "Showing Blade Parameters" on page 128
- "Showing Port Parameters" on page 132
- "Showing Firmware" on page 133
- "Showing Configuration File" on page 133

Showing Blade Parameters

Shows the parameters for a blade. See "Defaults and Options" on page 118 for a list of the defaults and options.

- "Bandwidth" on page 128
- "Gen" on page 128
- "MAC Address" on page 129
- "Promiscuous" on page 130
- "SR-IOV" on page 131
- "VLAN" on page 131

Bandwidth

Shows the bandwidth for a blade.

Syntax: show -b | blade 0..9 -p | port 0 | 1 -bw | bandw

For example, to show bandwidth for blade 3, port 0:

```
show -b 3 -p 0 -bw
```

Related Information

- "Gen" on page 128
- "MAC Address" on page 129
- "Promiscuous" on page 130
- "RSS" on page 130
- "SR-IOV" on page 131
- "VLAN" on page 131
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Managing the NEM Using CMM ILOM" on page 139

Gen

Shows the PCIe bus switch as gen1 or gen2.

Syntax: show -b | blade 0..9 -g | gen

For example, to show blade 3's gen setting:

```
show -b 3 -g
```

Related Information

- "Bandwidth" on page 128
- "MAC Address" on page 129
- "Promiscuous" on page 130
- "RSS" on page 130
- "SR-IOV" on page 131
- "VLAN" on page 131
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Managing the NEM Using CMM ILOM" on page 139

MAC Address

Shows the factory default and any user-defined MAC address for a port of a blade. You can use different function numbers.

Syntax: show -b | blade 0..9 [-f | func 0..3] -p | port 0 | 1 -m | mac

Also, you can check all MAC addresses for all blades and functions.

Syntax: show -b all -p | port 0 | 1 -m | mac

For example, to show MAC address(es) of port 1 of blade 3's function 0:

```
show -b 3 -f 0 -p 1 -m
```

For example, to show MAC addresses of port 0 of all blades:

```
show -b all -p 0 -m
```

Related Information

- "Bandwidth" on page 128
- "Gen" on page 128
- "Promiscuous" on page 130
- "RSS" on page 130

- "SR-IOV" on page 131
- "VLAN" on page 131
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Managing the NEM Using CMM ILOM" on page 139

Promiscuous

Shows if a blade's port is set to promiscuous, which enables a blade to receive packets destined to any MAC address.

Syntax: show -b | blade 0..9 [-f | func 0..3] -p | port 0 | 1 -r | promisc

For example, to set port 0 and function 0 of blade 4 to promiscuous:

```
show -b 4 -f 0 -p 0 -r
```

Related Information

- "Bandwidth" on page 128
- "Gen" on page 128
- "MAC Address" on page 129
- "RSS" on page 130
- "SR-IOV" on page 131
- "VLAN" on page 131
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Managing the NEM Using CMM ILOM" on page 139

RSS

Shows if a blade's port is set to RSS mode. By default, RSS mode is on for both ports.

Syntax: show -b | blade 0..9 -p | port 0 | 1 rss

For example, to show blade 3 and port 0:

```
show -b 3 -p 0 rss
```

Related Information

- "Bandwidth" on page 128
- "Gen" on page 128
- "MAC Address" on page 129
- "Promiscuous" on page 130
- "VLAN" on page 131
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Managing the NEM Using CMM ILOM" on page 139

SR-IOV

Shows if a blade's port is set to SR-IOV.

Syntax: show -b | blade 0..9 -sr

For example, to show blade 3, port 0 SR-IOV:

```
show -b 3 -sr
```

Related Information

- "Bandwidth" on page 128
- "Gen" on page 128
- "MAC Address" on page 129
- "Promiscuous" on page 130
- "RSS" on page 130
- "VLAN" on page 131
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Managing the NEM Using CMM ILOM" on page 139

VLAN

Two different commands are available for showing VLANs. One shows if a VLAN is set for a a blade's port. Another shows if a specific VLAN is associated with one or more blades.

To show if a VLAN is set for a blade's port, specify the port number.

Syntax: show -b | blade 0..9 [-p | port 0 | 1] [-f | func 0..3] -v | vlan

For example, to show blade 1's function 0 VLANs:

```
show -b 1 -f 0 -v
```

To show if a specific VLAN is associated with blades, specify the VLAN number. You can specify a blade number to show that blade, or you can omit the blade number to check all blades for the VLAN.

Syntax: show -v | vlan 1..4094 -b | blade [0..9]

For example, to show if VLAN 5 is associated with blade 4:

```
show -v 5 -b 4
```

For example, to show all blades associated with VLAN 5:

```
show -v 5 -b
```

Related Information

- "Bandwidth" on page 128
- "Gen" on page 128
- "MAC Address" on page 129
- "Promiscuous" on page 130
- "RSS" on page 130
- "SR-IOV" on page 131
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Managing the NEM Using CMM ILOM" on page 139

Showing Port Parameters

Shows all port parameter settings. You can show properties for both ports or specify a port to see only its properties.

See "Defaults and Options" on page 118 for a list of the defaults and options.

Syntax: show -p | port [0 | 1]

For example, to show all properties for both ports:

```
show -p
```

Related Information

- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Defaults and Options" on page 118
- "Managing the NEM Using CMM ILOM" on page 139

Showing Firmware

Shows the current application firmware, boot loader, and FCODE/PXIE image versions.

Syntax: show -f | firmware

```
show firmware
```

Related Information

- "Updating the NEM Firmware" on page 143
- "Managing the NEM Using CMM ILOM" on page 139

Showing Configuration File

Shows the configuration file where commands are saved.

Syntax: show -c | conf

```
show conf
```

Related Information

- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Defaults and Options" on page 118
- "Managing the NEM Using CMM ILOM" on page 139

Configuring NEM for Oracle VM Para-Virtualization

If you are running Oracle VM with the NEM, use one of the following procedures to configure the NEM for para-virtualization (non SR-IOV mode).

- "Configure Para-Virtualization Without Live Migration" on page 134
- "Configure Para-Virtualization With Live Migration" on page 137

▼ Configure Para-Virtualization Without Live Migration

1. **Using VM, assign the MAC address to the bridge that the NEM port has been assigned.**

   This MAC address will be used in the guest's OS as the virtual port of the NEM.

   Alternatively, you can assign the MAC address by editing the `vm.cfg` file using XEN.

   Example: To edit `vm.cfg` file to set eth4 as the NEM port, you would enter the following in the file:

```
vif = ['type=netfront,mac=00:14:4F:85:BE:8f,bridge=eth4']
```

2. **Log in to NEM using the CMM CLI. See** "Start NEM CLI from CMM ILOM CLI" on page 140**.**

3. **Set each MAC address to the NEM port that will be used as the para-virtualized port, based on the firmware version:**

   - **For version 1.040 or older:**

```
set -b x -p y -f z -m xx:xx:xx:xx:xx:xx
```

where

x is the blade number (0-9) running Oracle VM guests

y is the NEM port number (0 or 1) assigned to the Oracle VM guest

z is always 0 for the para-virtualization environment

xx:xx:xx:xx:xx:xx MAC address (the same MAC address that you assigned to the guests using Oracle VM or by editing the `vm.cfg` file

For example, if NEM port 0, which is `eth1` in the DOM0, is assigned the MAC address of 00:21:28:40:38:FE for blade 3, you would enter the following command:

```
set -b 3 -p 0 -f 0 -m 00:21:28:40:38:FE
```

Note – Only 15 VM MAC addresses (including bonded MAC addresses) can be configured per NEM port.

- **For version 1.050 or newer:**

```
set -b x -p y -f z -m 00:11:22:33:xx:xx
```

where

x is the blade number (0-9) running Oracle VM guests

y is the NEM port number (0 or 1) assigned to the Oracle VM guest

z is always 0 for the para-virtualization environment

00:11:22:33:xx:xx is the MAC address (the same MAC address that you assigned to the guests using Oracle VM or by editing the `vm.cfg` file

For example, if NEM port 0, which is `eth1` in the DOM0, is assigned the MAC address of 00:21:28:40:38:FE for blade 3, you would enter one of the following commands:

```
set -b 3 -p 0 -f 0 -m 00:21:28:40:38:xx
```

Or

```
set -b 3 -p 0 -f 0 -m 00:21:28:40:xx:xx
```

Note – The last two bytes of a MAC address are used as mask bits that allow you to use up to 64K MAC addresses per NEM port per blade. For example, from 00:11:22:33:00:00 to 00:11:22:33:FF:FF.

---

4. **Disable promiscuous mode on the blade's port that is assigned to the guest. See** "Promiscuous" on page 122**.**
5. **Disable SR-IOV mode. See** "SR-IOV" on page 123**.**
6. **Repeat** Step 3 **to** Step 5 **for each port of each blade.**

Related Information

- "Configuring and Viewing Parameters" on page 117
- "Setting Port Parameters" on page 125
- "Setting Blade Parameters" on page 119
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

▼ Configure Para-Virtualization With Live Migration

1. Enable live migration on the NEMs.

From a network connected terminal, ssh to the Blade chassis CMM.

```
-> start /CH/NEM0/SP/cli
set -c save on
set -b 0 -p 0 -mg on
set -b 0 -p 1 -mg on
set -b 1 -p 0 -mg on
set -b 1 -p 1 -mg on
set -b 2 -p 0 -mg on
set -b 2 -p 1 -mg on
set -b 3 -p 0 -mg on
set -b 3 -p 1 -mg on
set -b 4 -p 0 -mg on
set -b 4 -p 1 -mg on
set -b 5 -p 0 -mg on
set -b 5 -p 1 -mg on
set -b 6 -p 0 -mg on
set -b 6 -p 1 -mg on
set -b 7 -p 0 -mg on
set -b 7 -p 1 -mg on
set -b 8 -p 0 -mg on
set -b 8 -p 1 -mg on
set -b 9 -p 0 -mg on
set -b 9 -p 1 -mg on
set -c save off
```

2. Run the following command before you start to configure Oracle VM and and any VMs using VM Manager:

```
% rmmod sxge
% insmod sxge.ko sxge_enable_live_migrate=1
```

Or you could pass a parameter to a kernel module using a `modprobe` configuration file. Create `/etc/modprobe.d/sxge.conf` that contains:

```
options sxge sxge_enable_live_migrate=1
```

Note – If you are not using a version that supports live migration, either upgrade your version or use the procedure for configuring without live migration. See "Configure Para-Virtualization Without Live Migration" on page 134.

3. **Log in to NEM using the CMM CLI. See** "Start NEM CLI from CMM ILOM CLI" on page 140**.**
4. **Disable promiscuous mode on the blade's port that is assigned to the guest. See** "Promiscuous" on page 122**.**
5. **Disable SR-IOV mode. See** "SR-IOV" on page 123**.**
6. **Repeat** Step 4 **to** Step 5 **for each port of each blade.**

Note – Only 15 VM MAC addresses (including bonded MAC addresses) can be configured per NEM port.

Related Information

- "Configuring and Viewing Parameters" on page 117
- "Setting Port Parameters" on page 125
- "Setting Blade Parameters" on page 119
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

Managing the NEM Using CMM ILOM

This topic contains information for using Oracle Integrated Lights Out Manager (Oracle ILOM) with the NEM.

The Oracle ILOM tool enables you to actively manage and monitor the NEM independently of the operating system state, providing you with a reliable LOM system. With Oracle ILOM, you can proactively learn about hardware errors and faults as they occur.

- "ILOM Documentation and Updates" on page 139
- "Start NEM CLI from CMM ILOM CLI" on page 140

Related Information

- "Defaults and Options" on page 118
- "Configuring and Viewing Parameters" on page 117
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127
- "Updating the NEM Firmware" on page 143

ILOM Documentation and Updates

Oracle ILOM documentation is divided into two categories:

- General Oracle ILOM information, located in the Oracle Integrated Lights Out Manager 3.0 documentation collection at:

  `http://docs.oracle.com/cd/E19860-01/`

- Information specific to the NEM, located in these online topics.

For product updates that you can download for the NEM, go to the following web site:

`http://docs.oracle.com/cd/E21077_01/`

This site contains updates for firmware and drivers, including CD-ROM .iso images.

Related Information

- "Related Documentation" on page xi
- "Saving the Zoning Configuration to a Backup File" on page 25
- "Recovering Zoning Configurations" on page 27

▼ Start NEM CLI from CMM ILOM CLI

Normally, you access Oracle ILOM through the Ethernet. When the NEM is installed in the chassis, its Oracle ILOM is automatically accessible on the same subnet as the CMM ILOM. Access to CMM ILOM is through the NEM CMM CLI.

1. **Log in to the CMM ILOM.**

   The default user is `root` and default password is `changeme`.

   Once you have successfully logged in, the CMM ILOM displays the ILOM default command prompt:

   ->

2. **Navigate to /**`CH/NEM`*n*`/SP/cli`**.**

```
cd /CH/NEMn/SP/cli
```

   where *n* is 0 or 1 for NEMs 0 or 1.

3. **Type the command `start`.**

   A prompt asking if you want to start is displayed.

4. **Type `y` to continue or `n` to cancel.**

   A welcome screen with version number is displayed.

5. **When you are done, type** `exit`**.**

   The NEM ILOM exits, and the CMM ILOM CLI prompt is displayed.

Related Information

- "Configuring and Viewing Parameters" on page 117
- "Setting Blade Parameters" on page 119
- "Setting Port Parameters" on page 125
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127
- "Updating the NEM Firmware" on page 143
- "Saving the Zoning Configuration to a Backup File" on page 25

- "Recovering Zoning Configurations" on page 27

Updating the NEM Firmware

These topics provide information and procedures for updating the Sun Blade 6000 Virtualized 40 GbE Network Express Module firmware. Two types of firmware require updating on the NEM. The SAS2 Expander firmware and the NEM SP ILOM firmware. You can update both through either the web interface or the CMM ILOM CLI.

The *SAS-1/SAS-2 Compatibility Upgrade Guide for the Sun Blade 6000 Modular System* is at: `http://docs.oracle.com/cd/E22513_01/html/E21704-01/`.

- "Configuring Hot Plug Service" on page 143
- "Disabling Interfaces" on page 145
- "Updating Firmware (Web)" on page 145
- "Updating the Firmware (CLI)" on page 147

Related Information

- "Start NEM CLI from CMM ILOM CLI" on page 140
- "ILOM Documentation and Updates" on page 139
- "Managing the NEM Using CMM ILOM" on page 139
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

Configuring Hot Plug Service

Before updating firmware, ensure that all blades in the chassis are configured with hot-plug service enabled (on). The hot-plug service must be enabled because the operator needs to take the NEM offline, and return it online for the updated firmware to take effect. This procedure is equivalent to a physical hot-plug operation.

The command to enable hot-plug service is dependent on the blade's OS. Refer to your OS documentation for the applicable command and detailed instructions.

For Oracle Solaris OS and Oracle Linux OS, see:

- "Verify and Enable Hot Plug (Oracle Solaris OS)" on page 144
- "Verify and Enable Hot Plug (Linux OS)" on page 144

▼ Verify and Enable Hot Plug (Oracle Solaris OS)

1. **Execute the `svcs hotplug` command on the blade console to determine if the hot-plug service is online (enabled).**
2. **If hot-plug is offline (disabled), execute the `svcadm enable hotplug` command as shown in the following example.**

```
# svcs hotplug
   STATE          STIME    FMRI
   disabled       14:18:57 svc:/system/hotplug:default
   # svcadm enable hotplug
   # svcs hotplug
   STATE          STIME    FMRI
   online         14:07:15 svc:/system/hotplug:default
```

▼ Verify and Enable Hot Plug (Linux OS)

1. **Open the `/etc/modprobe.conf` or `/etc/modprobe.conf.local` file in a text editor.**
2. **Add the following command:**

```
install acpiphp /bin/true
```

3. ***Either*, persistently load the pciehp module per OS as follows:**
   - In Oracle Linux or RHEL 5.x, enter the following commands at a terminal prompt:

```
# echo modprobe pciehp >> /etc/sysconfig/modules/hotplug.modules
# chmod +x /etc/sysconfig/modules/hotplug.modules
```

Note – This step is not required for Oracle Linux or RHEL 6.x because PCIe hot plug support is built-in.

- In SLES, open the `/etc/sysconfig/kernel` configuration file in a text editor and change:

  `MODULES_LOADED_ON_BOOT=""`

  to:

  `MODULES_LOADED_ON_BOOT="pciehp"`

4. ***Or*, execute the following command at a terminal prompt before performing a hot plug action:**

```
# modprobe pciehp
```

Disabling Interfaces

Disable all interfaces between blade(s) and the NEM that you are updating:

- Perform steps 2 through 4 of "Prepare the PCIe Interface for Hot-plug Removal" on page 20.
- "Performing SAS Interface Hot-plug" on page 21

Updating Firmware (Web)

- "Obtain and Apply the Newest Firmware" on page 146
- "Reset the NEM (Web)" on page 147

Related Information

- "Start NEM CLI from CMM ILOM CLI" on page 140
- "ILOM Documentation and Updates" on page 139
- "Managing the NEM Using CMM ILOM" on page 139
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

▼ Obtain and Apply the Newest Firmware

1. **Download the firmware that you need from the Oracle software download site:**

   http://support.oracle.com

Note – Firmware downgrade is not possible. If CMM ILOM detects firmware downgrade, the update will be blocked.

---

   a. **At the My Oracle Support site, log in.**

      An account is required to download software. You must register if you do not have an account.

   b. **Click the Patches & Updates tab.**

   c. **Click "Product or Family" (Advance search).**

   d. **In the Product field, type `Sun Blade 6000 Virtualized 40GbE NEM`.**

   e. **In the Release field, select the newest release, then click the arrow next to the Sun Blade 6000 Virtualized 40GbE NEM folder.**

   f. **Click the Search button.**

   g. **Download the latest available firmware package, and extract it to an accessible folder on the network.**

Note – If you are looking for the Oracle Solaris 10 OS sxge driver, see "Download the sxge Driver" on page 75.

---

2. **Log in to the CMM ILOM as a user with Administrator privileges.**

3. **From the left panel, select System Information -> Firmware.**

4. **Select the radio button for the NEM that you want to update.**

5. **Select Firmware Update from the Actions drop-down menu.**

6. **From the pop-up menu, select the Transfer Method that you want to use.**

   The following transfer methods are available: TFTP, FTP, SCP, HTTP, HTTPS, SFTP.

7. **Type the path to the new ILOM or SAS `.pkg` file in the Select File field, or click Browse to locate and select the `.pkg` file.**

8. **Click Update.**

   The firmware update process will take several minutes.

9. **If you are updating an additional firmware package, repeat the process to update the additional firmware package.**

10. **After updating all firmware, when you are ready to apply the firmware updates, power cycle the NEM. (Continue with the next topic** "Reset the NEM (Web)" on page 147**.)**

▼ Reset the NEM (Web)

To apply the firmware updates you downloaded, you must take the NEM offline, and bring it back online. The firmware update process does not automatically apply the changes or reset the NEM. A reset must be manually initiated.

1. **Log in to the CMM ILOM as a user with Administrator privileges.**
2. **From the left panel, select System Information -> I/O Modules.**
3. **Select the radio button adjacent to the NEM listed that you want to take offline.**
4. **From the Actions pull-down, choose Prepare-to-Remove.**
5. **Click OK in the pop-up menu and wait for the action to complete.**

   The Health status will display "Offline" when the Prepare-to-Remove action has completed.
6. **From the Actions pull-down, select Return-to-Service.**
7. **Click OK in the pop-up menu and wait (approximately three minutes) for the action to complete.**

   The Health status will display "OK" when the Return-to-Action has completed.

Related Information

- "Configuring and Viewing Parameters" on page 117
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

Updating the Firmware (CLI)

- "Update NEM ILOM Firmware (CLI)" on page 148
- "Reset the NEM (CLI)" on page 149

Related Information

- "Start NEM CLI from CMM ILOM CLI" on page 140
- "ILOM Documentation and Updates" on page 139
- "Managing the NEM Using CMM ILOM" on page 139
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

▼ Update NEM ILOM Firmware (CLI)

1. **Download the firmware that you need from the Oracle software download site:**

   `http://support.oracle.com`

   a. **At the Oracle Support site, log in.**

      An account is required to download software. You must register if you do not have an account.

   b. **Click the Patches & Updates tab.**

   c. **Click "Product or Family" (Advance search).**

   d. **In the Product field, type `Sun Blade 6000 Virtualized 40GbE NEM`.**

   e. **In the Release field, select the newest release, then click the arrow next to the Sun Blade 6000 Virtualized 40 GbE NEM folder.**

   f. **Click the Search button.**

   g. **Download the latest available firmware package, and extract it to an accessible folder on the network.**

2. **From a network connected terminal, log in to the CMM ILOM CLI as Administrator**

```
$ ssh username@cmm-ipaddress
```

   where *username* is the Administrator user name and *cmm-ipaddress* is the IP address of the ILOM CMM.

3. **Type the password.**

   The ILOM CLI prompt appears.

```
->
```

4. **Change directories to the blade slot containing the NEM to be upgraded:**

```
-> cd /CH/NEMx
```

where NEM*x* is the NEM 0 or 1 to be upgraded.

5. **Type the following command:**

```
-> load -source http://http_server_ipaddress/firmware-version.pkg
```

where *http-server-ipaddress* is the domain name or IP address of your HTTP server location where you copied the image file, and *firmware-version* is the name of the .pkg file.

The firmware update process can take several minutes.

6. **Repeat the process if you are updating another firmware package.**

7. **After updating all firmware, reset the NEM. (Continue with the next topic** "Reset the NEM (CLI)" on page 149**.)**

▼ Reset the NEM (CLI)

To apply the firmware updates you downloaded, you must take the NEM offline, and bring it back online. The firmware update process does not automatically apply the changes or reset the NEM. A reset must be manually initiated.

1. **Log in to the CMM ILOM as a user with Administrator privileges.**

2. **From the list, locate the NEM you want to take offline, then type the following commands at the CMM ILOM CLI:**

```
-> cd /CH/NEMx
```

where NEM*x* is the NEM 0 or 1.

```
-> set prepare_to_remove_action=true
```

Wait for the remove command to complete.

3. **Check the status of the NEM.**

```
-> show -d properties /CH/NEMx
```

where NEM*x* is the NEM 0 or 1.

4. **If the NEM is offline (Prepare-to-Remove-Status=Ready), then type the command to return it to service:**

```
-> set return_to_service_action=true
```

Wait for the command to complete.

5. **When the process completes, type the following command to ensure that the NEM is online and the newest SAS firmware version was installed:**

```
-> show -d properties /CH/NEMx
```

where NEM*x* is the NEM 0 or 1.

The following is an example of output that is displayed.

```
-> show -d properties /CH/NEM0
  /CH/NEM0
    Properties:
        type = Network Express Module
        ipmi_name = NEM0
fru_name = SUN BLADE 6000 VIRTUALIZED MULTI-FABRIC 40GE PCI-EX NEM
        fru_manufacturer = MiTAC International Corporation
        fru_version = SAS 5.3.8.1
        fru_part_number = 07052307
        fru_serial_number = 489089M+1123L2000F
        fru_extra_1 = SAS 5.3.8.1
        fault_state = OK
        load_uri = (none)
        clear_fault_action = (none)
        prepare_to_remove_status = NotReady
        prepare_to_remove_action = (none)
        return_to_service_action = (none)
```

6. **To display the system firmware version, type the commands shown in the following example, where NEM*x* is NEM0 or NEM1.**

```
-> start /CH/NEMx/SP/cli
     Are you sure you want to start /CH/NEM0/SP/cli (y/n)? y
          Welcome to Sun Blade 6000 Virtualized
```

```
                  Firmware version: 1.040
> show firmware
System firmware version:          1.040
Boot loader image version:        0.821
EEPROM image version:             1.400
OK
> exit
> Connection closed by foreign host

->
```

Related Information

- "Configuring and Viewing Parameters" on page 117
- "Viewing Parameter Settings, Firmware Versions, and Configuration File" on page 127

Glossary

A

ASIC application-specific integrated circuit

ATTN Attention push button and LED

C

CMM chassis management module (service processor in C10/C48 systems)

D

DDR Double Data Rate

DP D port of a PCIe switch

DRAM Dynamic Random Access Memory

E

ESD electrostatic discharge

F

FEM fabric express module

FPGA field programmable gate array

FRU field replaceable unit

G

GbE Gigabit Ethernet

H

HBA host bus adapter

HPL Hot-plug link

I

I²C inter-integrated circuit

ILOM Integrated Lights Out Management

L

LAN local area network

LC long connector

LOC Locate LED and push button

LOM Lights Out Management

M

MMF multimode fiber

MTU maximum transmission unit

N

NEM Network Express Module

NEM0, NEM1 Terms used by NEM management software to identify NEMs occupying slots in a chassis.

NIC Network Interface Card

NIU Network Interface Unit

O

OM optical module

OS operating system

P

PCIe Peripheral Component Interconnect Express

PEU PCIe unit (physical, data link, and switch transaction layers)

PF physical function device

R

RARP Reverse Address Resolution Protocol

RCR receiver

REM RAID Expansion Module

RSS Receive-Side Scaling

RSS enables packet receive-processing to scale with the available computer processors. Enabling RSS allows the network load from a network adapter to be balanced across CPUs while maintaining in-order delivery of data. This delivery is enabled with RSS by ensuring that the packets from a single TCP connection are processed on one CPU.

RX receive

S

SAS2 NEM A generic term that applies to any NEM that supports SAS2 connectivity

SEL system event log

SerDes Serializer/Deserializer

server module Any server module (blade) that will interoperate with a storage module (blade). For a server module to work with a NEM, the server module must have a SAS2 REM installed. The terms server module and server blade are used interchangeably.

SFP+ Small form-factor pluggable

SLES SUSE Linux Server

SMF single-mode fiber

SR-IOV Single Root I/O Virtualization is a PCI Express Extended capability that makes one physical device appear as multiple virtual devices.

storage module The Sun Blade 6000 Module M2. The terms storage module and storage blade are used interchangeably.

Sun ASIC Sun application-specific integrated circuit

T

TCP Transmission Control Protocol

TCP 4-tuple TCP destination port, TCP source port, Source IP address, and Destination IP address

TPE twisted-pair Ethernet

TX transmit

U

UNR upper-non-recoverable

V

VF virtual function device

VLAN Virtual Local Area Network

W

WOL Wake-on-LAN mode

Index

Numerics

10 GbE network interface, 8
10/100/1000 BASE-T Ethernet interfaces, 5
10/100BASE-T Ethernet management port, 9
40 GbE NIC, 6
5-meter copper splitter, 49

A

aggregation
  Linux OS, 65
  Oracle Solaris 10 OS, 64
  Oracle Solaris 11 OS, 63
AMD drivers, unsupported, 103
ARM9S controller, 9
ASIC, 8
assigning an IP address, 76
ATTN button, 16
ATTN button, hot-plug insertion or removal, 18

B

bandwidth
  set, 119
  show, 128
BIOS, setup for booting from NEM, 67
blade parameters, CMM CLI commands, 119
boot
  BIOS setup, 67
  ESC key, 67
  NEM from SPARC server over network, 69
  NEM from x86 server over network, 69
  server, configuring, 67
buttons, 14

C

cables
  length, 49
  optical modules, 11
  twin-axial, 11
CMM CLI commands
  bandwidth, 119
  configuration file, 133
  defaults and options, 118
  gen, 120
  link status, 125
  ports, 125
  promiscuous, 122
  save configuration, 126
  set commands, 119
  show, 127
  show firmware versions, 133
  show port parameters, 132
  speed, 125
  SR-IOV, 123
  VLAN, 124

D

Device code
  0x207, 86
  0x2078, 86
  0x207b, 86
disk hot plug actions, 21
documentation
  Sun Blade 6000 Modular system, xi
drivers
  Linux operating systems, 82
    configuring Jumbo Frames, 56
    configuring the network interface, 90
    removing drivers, 88
  Oracle Solaris operating systems
    configuring Jumbo Frames, 75
    configuring network host files, 76
    configuring sxge device driver, 78

VMware ESX server platforms, installing and configuring, 113
Windows operating systems, 103
installing network controller, 104
installing VLAN, 106
uninstall VLAN, 107
uninstalling network controller, 105

E

electrostatic-sensitive components, 31
environmental requirements, 4
ESC key, jumping to next NEM PXE image, 67
eth0 and eth1 devices, 87
Ethernet connections, 5
Ethernet interface, testing, 90
external SAS connections, unsupported, 5

F

F12 network booting method, 67
F8 network booting method, 67
fabric express modules (FEMs), 8
fault detection and reporting, 3
features, 1, 2
FEMs, 8
firmware versions, show, 133
firmware, updating, 143
FPGA, 2

G

gen parameter
set, 120
show, 128
Gigabit Ethernet ports, 11
graceful NEM removal, 19

H

health monitoring, 3
hosts, maximum, 8
hot-plug
actions and outcomes, 21
enabling on blades, 19
PCIe guidelines, 17
SAS guidelines, 21

I

I/O sharing, 8
ILOM
description, 143
indicators and controls, 3
installation
NEM, 34
preparation, 31
SFP+ module, 47
verifying, 37
Integrated Lights Out Manager, see ILOM, 139, 143
interfaces
10GbE network, 8
management, 2
server module, 3
internal processor, 9

J

Jumbo Frames
Linux, 56
Oracle Solaris OS, 52
VMware ESX servers, 59
Windows, 59

L

LC cables, 49
LEDs, 14
Activity/Link for 10/100/1000, 15
ATTN LED and push button, 16
LOC (Locate) and push button, 15
Module Activity, 16
RJ45 ports, 13
Service Action Required, 16
SFP+ Activity/Link for 40GbE ports, 15
Speed for 10/100/1000, 15
SRP+ Activity/Link for 10GbE ports, 15
link instability, 49
Linux operating systems
configuration file, 93
configuring Jumbo Frames, 56
configuring the network interface, 90
driver's debug messages, 100
drivers, 81
installing drivers, 82
removing drivers, 88
LOC (Locate) LED and push button, 15
LSI SAS2 x36 expander, 2

M

MAC address
  assignment booting method, 67
  show, 129
  virtual, 8
management interfaces, 2
managing the NEM, 117
modprobe command, 87
Module Activity LED, 16
Module Fault LED, 16
multipathing, 43

N

NEM ASIC, 6
NEM ASICs, control and management of, 9
NEM components
  10 GbE connections, 8
  Ethernet connections, 5
  fabric express modules, 8
  LSI-SAS2 connections, 4
  processor, 9
  virtualized multi-fabric, 8
NEM eth device, 87
NEM firmware, updating, 143
NEM hot-plug actions, 21
network adaptors, configuring, 116
network device database, 91
network host files, configuring, 76
network throughput, increase, 81

O

OK-to-Remove LED, 15
optical module cables, compatible, 11
Oracle ILOM
  connecting, 140
  description, 139
  documentation, 139
OS panic, 43
OS, loading over network, 67
overview, this guide, 1

P

parallel operation, configuring, 81
parameters, configuring, 117
pass-through Gigabit Ethernet ports, 11
PCIe
  endpoint, 8
  hot-plug guidelines, 17
  interface, NEM removal, 20
physical appearance, 10
port parameters
  set, 125
  show, 132
ports, 3, 11
  aggregating, 63
power supplies, 3
Power/OK LED, 16
Prepare to Remove command, 18
primary network interface, Linux OS, 90
primary path to disks, 43
promiscuous parameter
  set, 122
  show, 130
PXE
  image, booting the NEM, 67
  server, 67

R

Remove and Return to Service command, 18
removing a NEM, 40
removing an SFP+ module, 48
replacing a disk, 21
replacing a NEM, 43
replacing the storage module, 21
resetting the NEM
  from web after firmware update, 147, 149
RJ45 ports
  Ethernet port, 13
  Gigabit Ethernet ports, 11
roadmap, this guide, 1

S

SAS
  3 Gbps, 4
  configuring SAS2 modules, 5
  hot-plug guidelines, 21
  LSI-SAS2 connections, 4
  preparing for hot-plug removal, 22
  restoring configurations, 22
  SAS1 and SAS2 compatibility, 5
  SAS2 (6Gbps), 4

server module interfaces, 3
Service Action Required LED, 16
SFP+ module
  cabling, 49
  installation, 47
  installing, 47
  installing and replacing, 46
  ports, 11
  ports, configuring, 11
  removal, 48
show commands, CMM CLI, 127
SLES platform, configuring, 95
Solaris operating systems
  configuring network host files, 76
  configuring sxge device driver, 78
speed, set, 125
SR-IOV mode, 82
SR-IOV parameters
  set, 123
  show, 131
static aggregation, 63
statistics, per-server-module, 8
storage module hot-plug actions, 21
Sun ASIC, 2
Sun Blade 6000 CMM, 9
Sun Blade 6000 modular system chassis,
  installation, 31
Sun Blade Storage Module M2, 8
Sun Blade Zone Manager, 5
surprise NEM removal, 19
sxge and sxgevf drivers Linux
  configuring the network interface, 90
  downloading and installing, 82
  removing, 88
sxge driver Oracle Solaris OS
  configuring the network host files, 76
  downloading and installing, 75
  parameters, 78

T

temperature monitoring, 3
twin-axial cables, 11

V

verifying NEM installation
  with ILOM CLI, 37
  with ILOM web interface, 37
virtual MAC, 8
virtualized 10GB connections, via SFP+, 11
virtualized NEM Ethernet connections, diagram, 6
VLAN setting, show, 131
VMware ESX servers
  configuring Jumbo Frames, 59
  configuring network adaptors, 116
  installing drivers, 113
voltage monitoring, 3

W

wake-on-LAN mode, 13
Windows operating systems
  enabling Jumbo Frames, 59
  installing drivers, 103
  installing enclosure device, 110
  installing network controller drivers, 104
  installing VLAN drivers, 106

Z

zoning configurations
  programming, 5
  restoring, 27
  SAS2 modules, 5
  saving, 25, 26