

Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module User's Guide



Part No: 821-0875
September 2012, Rev A

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

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

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

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 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.

Contents

Using This Documentation	5
Related Books	5
About This Documentation (PDF and HTML)	6
We Welcome Your Comments	6
Change History	6
Overview of Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM User's Guide	7
Features of the Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM	9
Terminology	9
Features Overview	10
Components Overview	11
NEM Ports and LEDs	17
Performing Hot Plug Insertion and Removal	23
PCIe Hot Plug	23
SAS Interface Hot Plug	27
Installing or Replacing the Virtualized M2 NEM	39
Installing a NEM	39
Verifying NEM Installation	41
Removing a NEM	45
Replacing a NEM	48
Installing and Removing SFP+ Optical Transceiver Modules	51
Cabling the SPF+ Connectors	53
Bootting Over the Virtualized M2 NEM 10-Gigabit Ethernet Port	55
Bootting over the Network With an x86 Blade Server	55
Bootting over the Network With a SPARC Blade Server	61
Installing and Configuring the hxge Driver on a Solaris SPARC or x86 Platform	65
How to Configure the Network Host Files	65
Configuring the hxge Device Driver Parameters	66
Configuring the Jumbo Frames Feature	69

Installing and Configuring the hxge Driver on a Linux Platform	73
Installing and Removing the Driver on a Linux Platform	73
Configuring the Network Interface	78
Checking and Testing the hxge Device	82
Changing the hxge Driver Configuration	84
Troubleshooting the Driver	88
Configuring Jumbo Frames	91
Installing and Configuring Drivers on a Windows Platform	93
Installing Drivers on a Windows Platform	93
Enabling Jumbo Frames	113
Installing and Configuring Drivers on a VMware ESX Server Platform	115
Installing the ESX Server Drivers on an Existing ESX Server	115
Installing the ESX Server Drivers With a New ESX Installation	117
Configuring the Virtual NEM M2 Network Adapters	118
Configuring Jumbo Frames	119
ILOM Supplement	121
Accessing ILOM Documentation and Updates	121
Connecting to ILOM	122
Updating the NEM Firmware	127
Sun Blade Zone Manager	132
NEM Sensors	132
Enabling Private and Failover Mode	137
Using Hot Plug Commands	138
Fixing Problems with Oracle ILOM Using the Preboot Menu	141
Index	149

Using This Documentation

This section describes related documentation, submitting feedback to Oracle, and a document change history.

- “Related Books” on page 5
- “About This Documentation (PDF and HTML)” on page 6
- “We Welcome Your Comments” on page 6
- “Change History” on page 6

Related Books

The following is a list of documents related to Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module. These and additional support documents are available on the web at:

<http://download.oracle.com/docs/cd/E19530-01/>

Document	Description
<i>Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module Product Notes</i>	Important late-breaking information about the Virtualized M2 NEM.
<i>SFP+ Module Installation Guide for the Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module</i>	How to install the SFP+ modules for the Virtualized M2 NEM.
<i>Oracle Integrated Lights Out Manager (ILOM) 3.0 Documentation Collection</i>	Information on using ILOM to monitor the Virtualized M2 NEM.
<i>Sun Blade 6000 Modular System Safety and Compliance Guide</i>	Safety and compliance information about your the Sun Blade 6000 modular system.
<i>Sun Blade 6000 Modular System Product Notes</i>	Important late-breaking information about the Sun Blade 6000 modular system.
<i>Sun Blade Storage Module M2 Administration Guide</i>	How to perform administration tasks with the Sun Blade Storage Module M2.

Translated versions of some of these documents are available at the web site described previously in Simplified Chinese, Japanese, Korean, Spanish, and French. English documentation is revised more frequently and might be more up-to-date than the translated documentation.

About This Documentation (PDF and HTML)

This documentation set is available in both PDF and HTML. The information is presented in topic-based format (similar to online help) and therefore does not include chapters, appendices or section numbering.

We Welcome Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. To share your comments, go to <http://www.oracle.com/goto/docfeedback>.

Change History

The following lists the release history of this documentation set.

- July 2010. Initial publication.
- August 2010. Updated *Product Notes* with information for product software release 1.0.1.
- November 2010. Updated *Product Notes* with information for product software release 1.1.
- January 2011. Updated *Product Notes* for CR 7005233. Updated *User Guide* to add information about starting hotplugd (1M) service in Solaris 10 OS 09/10.
- April 2011. Updated *Product Notes* for CR 7020854. Updated *User Guide* to add information about updating SAS firmware.
- September 2011. Updated *Product Notes* information for the Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM Software release 1.2. Added new OS support. Added T3-1B SPARC blade support. Added CR 7075694. Updated *User Guide* to add information about configuring Linux for PCIe hot plug.
- August 2012. Updated *Product Notes* for CR 7183913 and 6932684.
- September 2012. Updated *User Guide* to add instructions for service processor firmware image recovery (CR 7102112). Updated *Product Notes* information for the Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM Software release 1.3.

Overview of Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM User's Guide

The following topics are covered in this book.

Description	Link
Learn about the features of the NEM.	“Features of the Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM” on page 9
Learn how to hotplug the NEM.	“Performing Hot Plug Insertion and Removal” on page 23
Learn how to install or replace the NEM.	“Installing or Replacing the Virtualized M2 NEM” on page 39
Find out how to boot over the Network with the NEM.	“Booting Over the Virtualized M2 NEM 10-Gigabit Ethernet Port” on page 55
Install and configure Solaris drivers.	“Installing and Configuring the hxge Driver on a Solaris SPARC or x86 Platform” on page 65
Install and configure Linux drivers.	“Installing and Configuring the hxge Driver on a Linux Platform” on page 73
Install and configure Windows drivers.	“Installing and Configuring Drivers on a Windows Platform” on page 93
Install and configure ESX drivers.	“Installing and Configuring Drivers on a VMware ESX Server Platform” on page 115
Use ILOM to monitor and manage the NEM.	“ILOM Supplement ” on page 121

Features of the Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM

The Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module (Virtualized M2 NEM) is a multi-purpose connectivity module for the Sun Blade 6000 modular system. The Virtualized M2 NEM supports connection to external devices through 10 GigabitEthernet (GbE) small form-factor pluggable (SFP)+ ports and 10/100/1000 twisted-pair Ethernet (TPE) ports. The NEM connects server modules (blades) in a Sun Blade 6000 modular system chassis with SAS-2 disk modules in the same chassis and also provides external SAS-2 connections.

This section covers the following topics:

- [“Terminology” on page 9](#)
- [“Features Overview” on page 10](#)
- [“Components Overview” on page 11](#)
- [“NEM Ports and LEDs” on page 17](#)

Note – If you are planning to use a Sun Blade Storage Module M2 for internal chassis host storage, you will also need to refer to its documentation additional installation and configuration details.

Terminology

The following terms are used in this document:

Term	Description
chassis	The Sun Blade 6000 Modular System blade enclosure.
storage module (or storage blade)	The Sun Blade Storage Module M2. The terms <i>storage module</i> and <i>storage blade</i> are used interchangeably.
server module (or server blade)	Any server module (blade) that will interoperate with a storage module (blade). For a server module to work with a Virtualized M2 NEM, the server module must have a SAS-2 REM installed. The terms <i>server module</i> and <i>server blade</i> are used interchangeably.

Term	Description
Virtualized M2 NEM	The Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module that plugs into a Sun Blade 6000 chassis (abbreviated Virtualized Multi-Fabric 10GbE M2 NEM).
Virtualized M2 NEM ASIC	A shortened reference to the ASICs embedded in the Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module that enable 10GbE virtualization.
Multi-Fabric NEM	A generic term that applies to any Network Express Module that provides a variety of interconnect options to server blades in a chassis. The Virtualized Multi-Fabric 10GbE M2 NEM is one example.
SAS-2 NEM	A generic term that applies to any Network Express Module that supports SAS-2 connectivity. The Virtualized Multi-Fabric 10GbE M2 NEM is one example of a SAS-2 NEM.
NEM 0, NEM 1	Terms used by NEM management software to identify Multi-Fabric NEMs occupying NEM slots in the chassis.
10 GbE	10 Gigabit Ethernet.

Features Overview

The key features of the Virtualized Multi-Fabric 10GbE M2 NEM are listed in the following table:

TABLE 1 Virtualized Multi-Fabric 10GbE M2 NEM Features Overview

Feature	Specifications
Service Processor (SP)	<ul style="list-style-type: none">Aspeed AST2000 (service processor)128MB DDR DRAM, 16MB flashSP Ethernet port through and Ethernet switch to the CMM
LSI SAS-2 x36 Expander	<ul style="list-style-type: none">Arm7-S Processor (SAS expander device)8MB SRAM (4MB fall back), 8MB flashManagement in-band and via I2C, expander Ethernet port through an Ethernet switch to CMM
Management Interfaces	<ul style="list-style-type: none">10/100BASE-T Ethernet port to the Chassis Monitoring Module (CMM)I2C to CMM

TABLE 1 Virtualized Multi-Fabric 10GbE M2 NEM Features Overview *(Continued)*

Feature	Specifications
Ports	<ul style="list-style-type: none"> ■ Four External x4 mini SAS-2 ports (external SAS connections are currently not supported) ■ Ten 10/100/1000BASE-T Ethernet (RJ-45) ports ■ Two SFP+ ports that support the SFP+ short range (SR) modules
Server Module Interfaces (per Blade)	<ul style="list-style-type: none"> ■ PCIe up to x8 ■ SAS-2 x2 ■ GbE x1
Updates	All software and embedded firmware can be field upgraded
Indicators/Controls	<ul style="list-style-type: none"> ■ Ethernet Link Status and Activity ■ SFP+ Link Status and Activity ■ SIS LEDs, Locate button, Attention button ■ Mini-SAS LED that is bi-functional for link and activity ■ Per Blade 10Gbps Link status in the NEM ILOM web interface
Health	<ul style="list-style-type: none"> ■ Voltage monitoring ■ Temperature monitoring ■ Fault detection
Power Supplies	<ul style="list-style-type: none"> ■ 3.3V_AUX from chassis midplane ■ 12V from chassis midplane ■ Other voltages generated on board
Environmental	Cooling: top-to-back forced air Operating: <ul style="list-style-type: none"> ■ Humidity: 10-90% non-condensing ■ Temperature: 0 to 35°C operating (32 to 95°F) ■ Altitude: 0-10,000 ft (3048 meters) Non operating <ul style="list-style-type: none"> ■ Humidity: 5-95% non-condensing ■ Temperature: -40 to 70°C operating (-40 to 158°F) ■ Altitude: 39,370 ft (12000 meters)

Components Overview

This section describes the major components of the Virtualized Multi-Fabric 10GbE M2 NEM.

The following topics are covered in this section:

- [“SAS Connections” on page 12](#)
- [“Ethernet Connections” on page 12](#)

- “10 GbE NIC Virtualization” on page 13
- “Fabric Express Modules” on page 17
- “Service Processor” on page 17

SAS Connections

Each Virtualized M2 NEM provides 10 x2 SAS-2 connections for server and storage blades .

One of the functions of the Virtualized Multi-Fabric 10GbE M2 NEM SAS-2 expander supports SAS zones, which enables chassis storage to be assigned to server blades. The Sun Blade Zone Manager running on the CMM (Chassis Monitoring Module) provides a way to create and manage the 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 SAS-2 enabled. Blade servers or storage blades installed in a chassis with a Virtualized M2 NEM that are not SAS-2 enabled will not be zoned using the Sun Blade Zone Manager.

See the following documentation for more information on zoning SAS-2 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 SAS-2 disks in the chassis.
- *Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module NEM Product Notes* for information on supported server and storage modules.
- Server blade documentation for information on availability of SAS-2 HBAs.

External SAS connections are not supported for the Virtualized M2 NEM.

Ethernet Connections

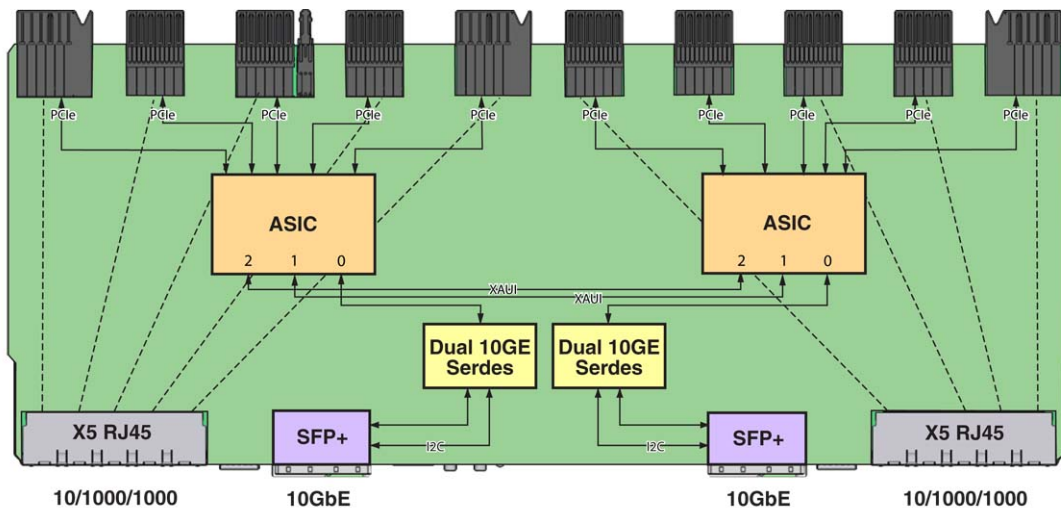
The Virtualized Multi-Fabric 10GbE M2 NEM provides the magnetics and RJ-45 connectors for the ten 10/100/1000 BASE-T Ethernet interfaces that come from Sun Blade 6000 server modules through the midplane. The Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM has one 10/100/1000BASE-T Ethernet port for each server module slot. There is no active circuitry on the Virtualized Multi-Fabric 10GbE M2 NEM for these GbE ports.

The Virtualized Multi-Fabric 10GbE M2 NEM also provides a 10 GbE connection for each server module. Each server module appears to have its own 10 GbE NIC through the Virtualized M2 NEM ASICs. All the server modules share two physical small form-factor pluggable (SFP+) 10 GbE ports via a dual channel 10 GbE Serializer/Deserializer (SerDes) per

ASIC. In a Sun Blade 6000 chassis, five server modules connect to a single Virtualized M2 NEM ASIC and share its 10 GbE port. Two Virtualized M2 NEM ASICs can be configured so that the ten server modules share one 10 GbE port for simplified cable aggregation.

See “[10 GbE NIC Virtualization](#)” on page 13 for further information on the NEM 10 GbE virtualization.

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



10 GbE NIC Virtualization

This section provides information on the 10 GbE NIC functionality for the Virtualized Multi-Fabric 10GbE M2 NEM.

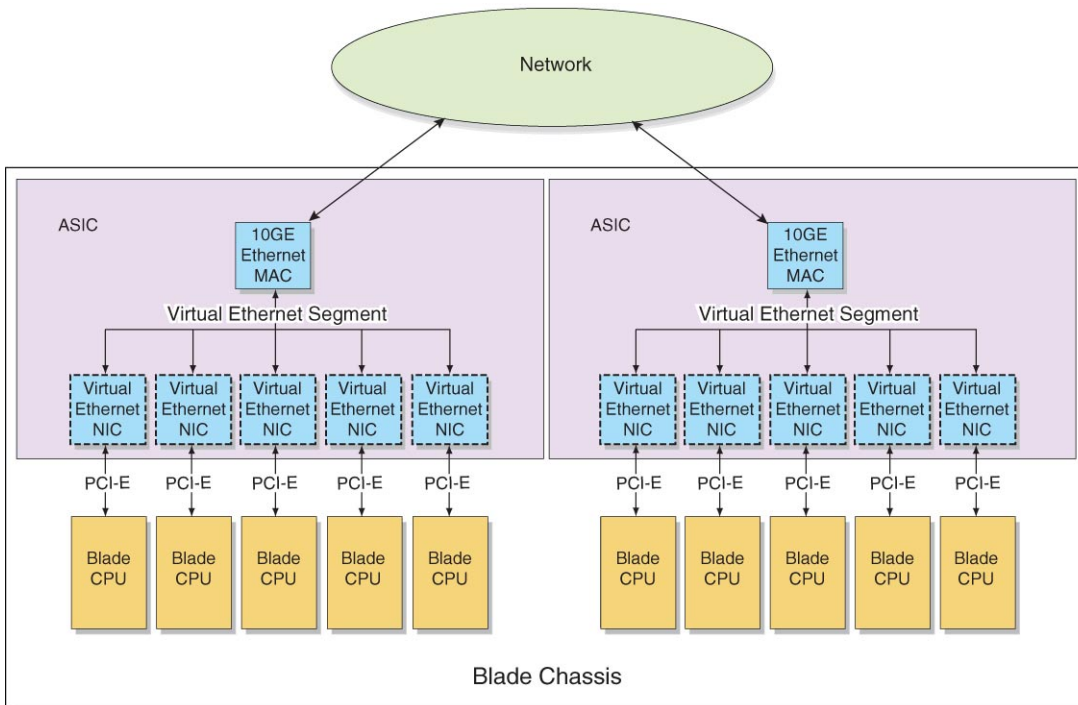
The Virtualized M2 NEM ASIC allows up to five hosts to share a single 10 GbE network port, with a dedicated PCIe endpoint for each host.

Shared I/O allows each server module to work as if a dedicated NIC connects 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 is hidden from the server modules. Only the service processor can access and configure this port.

A Virtualized Multi-Fabric 10GbE M2 NEM contains two Virtualized M2 NEM ASICs, which can operate in two different modes: bandwidth mode and connectivity mode.

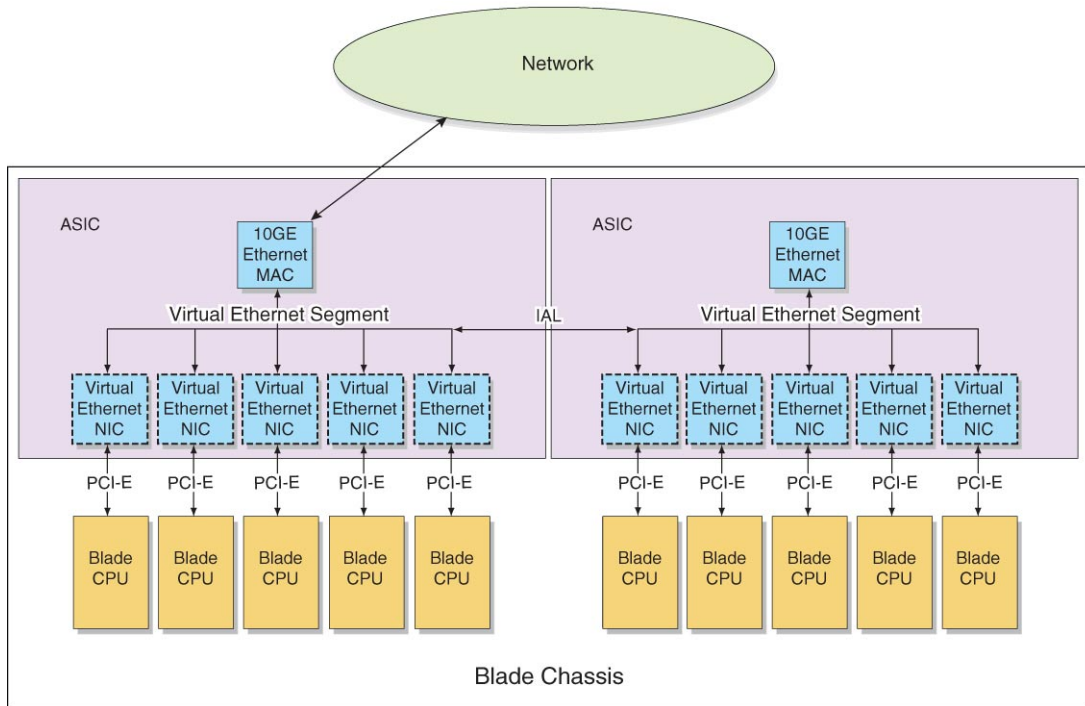
Bandwidth Mode

The Virtualized M2 NEM ASICs can operate in bandwidth mode where the Virtualized M2 NEM ASICs act without knowledge of each other. Each ASIC provides 10 GbE network access for the five hosts attached to it. Both SFP+ modules with optical cables should be installed, in order to provide 10 GbE connectivity for all ten server modules.



Connectivity Mode

Two Virtualized M2 NEM ASICs can also be interconnected so that a single 10 GbE port can act as the shared I/O for 10 server modules (connectivity mode). The inter-ASIC link (IAL) extends the virtual Ethernet segment to all the server modules. An example of this configuration is shown in the following illustration.



Private Mode

When private mode is enabled through the NEM ILOM interface, all 10 GbE traffic passes between the NEM ASICs via the IAL. If SFP+ modules are installed or removed, there is no effect on the IAL status.

If private mode is enabled and failover mode is also enabled, private mode takes precedence, and removal or insertion of SFP+ modules has no effect.

See [“Enabling Private and Failover Mode” on page 137](#) for instructions on enabling private mode.

Enabling or Disabling the IAL

The IAL is enabled or disabled due to a combination of three different factors:

- The number of SFP+ modules installed in the NEM
- Whether or not the NEM has failover mode enabled
- Whether or not the NEM has private mode enabled

The following table explains the IAL functionality:

TABLE 2 IAL Functionality

Mode	Original State of SFP+ Connections	Original State of IAL	Action taken on SFP+ Connections	Resulting State of SFP+ IAL
Static	One SFP+ module exists in the NEM prior to being inserted into the chassis.	IAL is disabled.	A second SFP+ module is added.	The insertion of the SFP+ module is ignored. The IAL is still disabled.
Static	Two SFP+ modules exist in the NEM prior to being inserted into the chassis.	IAL is disabled.	One SFP+ module is removed, disabled, or damaged.	The Virtualized M2 NEM ASIC stays in Bandwidth mode. One of the ASICs loses network connectivity. IAL is still disabled.
Failover	One SFP+ module exists in the NEM, and the NEM is alive in the chassis.	IAL is enabled.	A second SFP+ module is added.	The IAL dynamically disables itself and all traffic now flows through the two SFP+ connections.
Failover	Two SFP+ modules exist in the NEM, and the NEM is alive in the chassis.	IAL is enabled.	One SFP+ module is removed, disabled, or damaged.	IAL dynamically configures so traffic flows to the active SFP+ connection. IAL is enabled.
Private	Any number of SFP+ modules exist in the NEM.	IAL is enabled.	Any number of SFP+ modules are inserted or removed from the NEM.	IAL remains enabled — addition or removal of SFP+ modules does not affect IAL status.

Refer to [“Enabling Private and Failover Mode” on page 137](#) for instructions on how to enable failover mode through the NEM service processor.

Note – If IAL is in failover mode and an additional SFP+ module is installed, you must install a network—connected cable into the newly installed SFP+ connector to ensure 10GbE connectivity for all server modules.

For example, in the following scenario, IAL will be dynamically disabled:

- Ten server modules are installed in the Sun Blade 6000 chassis.
- One SFP+ module is installed in the Virtualized M2 NEM.
- Failover IAL mode is enabled.
- One additional SFP+ module is installed.

The five server modules attached to the ASIC that is connected to the new SFP+ module will lose 10GbE connectivity unless a fiber optic cable is connected to the new SFP+ connector.

Fabric Express Modules

For server modules to access the 10 GbE functionality of the Virtualized Multi-Fabric 10GbE NEM, they must have a fabric express module (FEM) installed in the server module. The [“Supported Server and Storage Modules” in Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module Product Notes](#) lists which FEMs are required for each server module that is supported for the Virtualized M2 NEM.

Note – The Sun Blade Storage Module M2 does not have an Ethernet controller. Therefore the Ethernet ports on the Virtualized Multi-Fabric 10GbE M2 NEM that correspond to slots containing a Sun Blade Storage Module M2 are not used.

Service Processor

The Virtualized Multi-Fabric 10GbE M2 NEM has an Aspeed AST2000 controller as its service processor (SP), which is responsible for the control and management of the Virtualized M2 NEM ASICs. The NEM also has a 10/100BASE-T Ethernet management port for connectivity to the Sun Blade 6000 CMM.

The following list describes the features of the service processor:

- The AST2000 controller has a 200MHz ARM9 CPU core and a rich set of features and interfaces. The BCM5241 receives a 25MHz reference clock.
- A Broadcom BCM5241 10/100BASE-T Ethernet PHY connected to the AST2000 controller’s 10/100M Fast Ethernet MAC provides the Ethernet management interface to the CMM.
- The SP can update Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM FPGA and firmware using GPIOs connected to the JTAG port of the FPGA.

Refer to [“ILOM Supplement” on page 121](#) for more information on the Integrated Lights Out Manager (ILOM) server management application for the SP.

NEM Ports and LEDs

This section covers the following topics:

- [“Physical Appearance of the Virtualized Multi-Fabric 10GbE M2 NEM” on page 18](#)
- [“External Virtualized M2 NEM Ports” on page 19](#)
- [“RJ45 Ethernet Connector Port LEDs” on page 19](#)
- [“NEM Front Panel and LEDs” on page 20](#)

Physical Appearance of the Virtualized Multi-Fabric 10GbE M2 NEM

The following illustration shows an overview of the Virtualized Multi-Fabric 10GbE M2 NEM.




TABLE 3 NEM Components

Label	Description
1	10/100/1000 TPE, RJ45 Connectors (5)
2	10GbE connector (1) (Requires SFP+ module)
3, 7	SAS-2 Connectors (2) (Not supported)
4	Serial management port
5, 6	Ejector Levers
8	10GbE connector (1) (Requires SFP+ module)

TABLE 3 NEM Components (Continued)

Label	Description
9	10/100/1000 TPE, RJ45 Connectors (5)
10	Virtualized Multi-Fabric 10GbE M2 NEM chassis
11	Midplane Connectors (10)

External Virtualized M2 NEM Ports

This section describes the external ports on the NEM.

External SAS x4 Ports

The external x4 wide mini-SAS-2 ports are not supported on the Virtualized M2 NEM.

Gigabit Ethernet Ports

There are 10 pass-through RJ-45 Gigabit Ethernet ports on the Virtualized Multi-Fabric 10GbE M2 NEM. The pass-through Gigabit Ethernet ports are strictly passive and isolated from the other functional blocks, with no interaction.

SFP+ Ports

There are two small form-factor, pluggable (SFP+) ports that provide virtualized 10 GB connections to the server modules.

RJ45 Ethernet Connector Port LEDs

Each RJ45 Ethernet port has two LEDs. The left LED is green and lights to show that a link has been established. It blinks off randomly 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. When the port is operating at 100 megabits per second, the right LED displays one color. When operating at 1000 megabits per second, it displays the other color. When operating at 10 megabits per second, the right LED is off. The green/amber color scheme varies from one server blade to another. The following table provides a chart for interpreting the link-speed relationships.

TABLE 4 Link-Speed LED Colors for Each Server Blade

Sun Blade Server Model	10 MbE (Right LED)	100 MbE (Right LED)	1000 MbE (Right LED)
T6320	Off	Green	Amber

TABLE 4 Link-Speed LED Colors for Each Server Blade (Continued)

Sun Blade Server Model	10 MbE (Right LED)	100 MbE (Right LED)	1000 MbE (Right LED)
T6340	Off	Amber	Green
X6270 M2	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 Wake-on-LAN (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 above. SPARC-based blades are designated by a T in front of the server module number (for example, T6300).

NEM Front Panel and LEDs

The following illustration shows the NEM front panel and LEDs, viewed from the back of the chassis. The following table describes the LED behavior.




TABLE 5 NEM Front Panel and LED Descriptions

Call-out	LED/Button Name	Description
1	Attention button	<p>Prepares NEM and associated blades for NEM removal and alerts the blades when a new NEM has been inserted.</p> <p>Refer to “Performing Hot Plug Insertion and Removal” on page 23 for information on the behavior of the Attention button during hot-plug insertion and removal.</p>
2	Locate button and LED (white)	<p>Helps locate each NEM.</p> <ul style="list-style-type: none"> 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 1 through 4 for 15 seconds. The LED can be enabled remotely.
3	Module Activity (Power/OK) LED (green)	<p>Has three states:</p> <ul style="list-style-type: none"> 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. <p>Refer to “Performing Hot Plug Insertion and Removal” on page 23 for information on the behavior of the Power/OK LED during hot-plug insertion and removal.</p>

TABLE 5 NEM Front Panel and LED Descriptions (Continued)

Call-out	LED/Button Name	Description
4	Module Fault LED (amber)	Has two states: <ul style="list-style-type: none">■ Off: No fault.■ On: An event has been acknowledged, and service action is required.
5	OK-to-Remove LED (blue)	Lights when it is safe to remove the NEM. Press the Attention button and wait until the OK-to-Remove LED lights before removing the NEM. Refer to “Performing Hot Plug Insertion and Removal” on page 23 for information on the behavior of the OK-to-Remove LED during hot-plug insertion and removal.
6	SAS Activity	Not supported.
7/ 8	Activity/Link for 10/100/1000 MbE	See “RJ45 Ethernet Connector Port LEDs” on page 19 .
9/10	Activity/Link for 10 GbE connection	The left green LED is for network activity status. It blinks on and off when there is network activity. The right green LED is for network link status. It is steady on when a 10GbE link is reached.
11	SFP+ Module	SFP+ modules are needed for 10GbE connectivity.

Performing Hot Plug Insertion and Removal

The Virtualized M2 NEM supports hot-plug insertion and removal. Hot plug insertion and removal involves two separate interfaces to chassis components: PCIe interfaces and SAS interfaces. This section describes guidelines and actions required to prepare PCIe and SAS interfaces for the NEM hot plug actions.

The guidelines and procedures in this section 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 the Virtualized M2 NEM” on page 39](#).

This section covers the following topics:

- [“PCIe Hot Plug” on page 23](#)
- [“SAS Interface Hot Plug” on page 27](#)

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

PCIe Hot Plug

This section covers the following topics:

- [“PCIe Hot Plug Guidelines” on page 23](#)
- [“How to Prepare the PCIe Interface for Hot Removal” on page 25](#)

PCIe Hot Plug Guidelines

This section covers the following guidelines for PCIe hot plug:

- [“Attention Button” on page 24](#)
- [“Surprise Removal” on page 24](#)
- [“Forcing NEM Graceful Removal” on page 24](#)

Attention Button

The Virtualized M2 NEM has an attention button. The attention button is pressed to initiate a hot-plug insertion or hot-plug removal. The green Power/OK LED slow blinks while the action is taking place. If the button is pressed a second time within 5 seconds after the green Power/OK LED begins blinking, the hot-remove or insert request is aborted.

An alternative to pushing the Attention 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:

- [“Installing a NEM” on page 39](#)
- [“Removing a NEM” on page 45](#)
- [“Replacing a NEM” on page 48](#)

Note – Successive pushes of the attention button (other than to cancel within 5 seconds) are not allowed. If one or more blades refuses to release the NEM and go offline after the first attention button press, force a graceful hot remove, as described in [“Forcing NEM Graceful Removal” on page 24](#).

Surprise Removal

The Virtualized M2 NEM does not support NEM surprise removal. Any attempt to remove the NEM without properly completing a graceful hot-plug removal, as described in [“Removing a NEM” on page 45](#) and [“Replacing a NEM” on page 48](#) is considered a surprise removal. The blue OK-to-Remove LED must be illuminated prior to removing the NEM.

If a host blade is removed from the chassis, the operation of the NEM is not affected and the NEM does not log the surprise removal event. The remaining server blades in the chassis are also not be affected by the removal of the blade.

Forcing NEM Graceful Removal

Once a hot plug removal has been initiated, it must be completed as indicated by the illumination of the blue OK-to-Remove LED. If the green Power/OK returns to solid green, this means that the NEM did not go off-line.

If the NEM did 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 Virtualized M2 NEM. It might be necessary to force blades that have denied immediate hot plug removal to release their resources, so that the rest of the system blades can proceed with the graceful removal of the NEM.

Use the instructions in [“How to Prepare the PCIe Interface for Hot Removal” on page 25](#) to ensure that all blades have unplumbed their hxge interfaces, and then continue with the ATTN button for hot removal.

▼ Configuring PCIe Device Hot Plug in Linux

Before attempting to perform a hot plug action on your the Virtualized M2 NEM, first configure any blade host running a Linux operating system specified in the workaround listed below to support PCIe hot plug

To configure the system for PCIe device hot plug:

1 Open the `/etc/modprobe.conf` or `/etc/modprobe.conf.local` file in a text editor.

2 Add the following command:

```
install acpihp /bin/true
```

3 *Either, persistently load the pciehp module per OS as follows:*

- In Oracle Linux or RHEL 5.x, enter the commands at a terminal prompt:

```
# echo modprobe pciehp >> /etc/rc.modules
```

```
# chmod +x /etc/rc.modules
```

Note – This step is not required for Oracle Linux or RHEL 6.x as PCIe hot plug support is built-in.

- In SLES, open the `/etc/sysconfig/kernel` configuration file in a text editor. 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
```

Next Steps [“How to Prepare the PCIe Interface for Hot Removal” on page 25](#)

▼ How to Prepare the PCIe Interface for Hot Removal

Before You Begin ■ Perform this procedure before hot removing a NEM.

- For hosts in the chassis running Oracle Solaris 10 OS 09/10 or later, the `hotplugd (1M)` daemon must be started before attempting to initiate a hot-plug action.

1 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
Solaris	<code>ifconfig -a grep hxge</code>	hxge (associated with Virtualized M2 NEM MAC address)
Linux	<code>ifconfig -a grep eth</code>	eth (associated with Virtualized M2 NEM MAC address)
Windows	Click the “Safely Remove Hardware” icon on the lower right task bar.	Sun Blade 6000 10Gbe Networking Controller

If the output does not show any Ethernet interface associated with the NEM that you want to remove, there is no existing interface. Proceed to one of the following procedures:

- [“Removing a NEM” on page 45](#)
- [“Replacing a NEM” on page 48](#)

If the output shows an Ethernet interface associated with the NEM to be removed, go to Step 2.

2 Disable all interfaces between the blade 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 number found in Step 1.

Operating System	Commands or Action
Solaris	<code>ifconfig hxgex down</code> <code>ifconfig hxgex unplumb</code>
Linux	<code>ifconfig ethx down</code>
Windows	<ol style="list-style-type: none">Double-click the Safely Remove Hardware icon at the lower right of the Windows taskbar. The Safely Remove Hardware dialog appears.Select the Sun Blade 6000 10 Gbe Networking Controller and click Stop.Click Close to finish. The NEM(s) can be safely removed.

- 3 Repeat Step 1 to make sure that the Ethernet interface targeted in Step 2 is disabled.
- 4 Proceed with one of the following procedures:
 - “Removing a NEM” on page 45
 - “Replacing a NEM” on page 48

SAS Interface Hot Plug

Removal and replacement activities can affect saved server-to-storage access configurations. This section covers procedures the following topics:

- “How to Prepare SAS Interfaces for Hot Plug Removal” on page 27
- “Restoring SAS Configurations After NEM Installation” on page 27
- “Saving the Zoning Configuration to a Backup File Using the CMM ILOM” on page 30
- “Recovering Zoning Configurations” on page 33

Note – The best way to ensure there is no interruption in service or accidental data loss when a Virtualized M2 NEM is removed is to have two Virtualized M2 NEMs installed in the chassis.

▼ How to Prepare SAS Interfaces for Hot Plug Removal

Before You Begin For hosts in the chassis running Oracle Solaris 10 OS 09/10 or later, the `hotplugd(1M)` daemon must be started before attempting to initiate a hot-plug action.

- 1 Back up the SAS zoning configurations. See “Saving the Zoning Configuration to a Backup File Using the CMM ILOM” on page 30.
- 2 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.
- 3 Ensure that the appropriate multi-path drivers are installed if your servers are not using a hardware RAID solution.

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 Virtualized M2 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 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.	<ul style="list-style-type: none"> ■ None. When using a SAS-2 REM configured for RAID, dual path support to storage will automatically be enabled. ■ When using a SAS-2 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 Virtualized M2 NEM into the same chassis slot.	Since 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 <i>LSI MegaRAID SAS Software User's Guide</i> . For Fusion-MPT IR, use the activate feature described in the <i>LSI SAS2 Integrated RAID User's Guide</i> .
Remove and insert a Virtualized M2 NEM into a different slot in the same chassis..	<p>Upon insertion into the chassis, the SAS zoning configurations on the NEM expanders are deleted.</p> <p>The NEM is automatically registered with the CMM. Any future saved configurations remain unchanged if the SAS-2 NEM is removed and reinserted in the same slot.</p>	<ul style="list-style-type: none"> ■ Restore zoning configurations from a backup of CMM configurations. See “Recovering Zoning Configurations” on page 33 —Or— ■ Use the CMM zoning utility to re-assign storage module resources to SAS-2 server blades.

Action	Outcome	User Intervention
Replace a failed Virtualized M2 NEM.	<ul style="list-style-type: none"> ■ If there are two Virtualized M2 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 in the same slot. ■ If there was only one Virtualized M2 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. 	<ul style="list-style-type: none"> ■ If there were two Virtualized M2 NEM, no further action is needed.. ■ If there was only one Virtualized M2 NEM, restore zoning configurations from a backup of CMM configurations. See “Recovering Zoning Configurations” on page 33 —Or— ■ If there was only one Virtualized M2 NEM use the CMM zoning utility to re-assign storage module resources to SAS-2 server blades. ■ If there was only one Virtualized M2 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 <i>LSI MegaRAID SAS Software User's Guide</i>. For Fusion-MPT IR, use the activate feature described in the <i>LSI SAS2 Integrated RAID User's Guide</i>.
Connect a SAS-2 JBOD to the SAS-NEM external SAS ports.	Not currently supported.	Do not connect external SAS storage to the NEM external SAS ports.

Saving the Zoning Configuration to a Backup File Using the CMM ILOM

This section contains the following procedures:

- “How to Save the Zoning Configuration to a Backup File Using the CMM CLI” on page 30
- “How to Save the Zoning Configuration to a Backup File Using the CMM Web Interface” on page 31

▼ How to Save the Zoning Configuration to a Backup File Using the CMM CLI

Before You Begin

- Create and save the storage access configuration for the chassis using the instructions in the Sun Blade Zone Manager section of the *Oracle Integrated Lights Out Manager (ILOM) CMM Administration Guide*.
- 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 in order to perform a complete backup of the ILOM SP configuration.

1 Log into the CMM ILOM CLI.

2 Change to the /SP/config directory.

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

3 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
```

4 To initiate the Backup operation, enter the command:

```
-> set  
dump_uri=transfer_method://username:password@ipaddress/directorypath/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.
- *directorypath* 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*.

▼ How to Save the Zoning Configuration to a Backup File Using the CMM Web Interface

- Before You Begin**
- Create and save the storage access configuration for the chassis using the instructions in the Sun Blade Zone Manager section of the *Oracle Integrated Lights Out Manager (ILOM) CMM Administration Guide*.
 - 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 in order to perform a complete backup of the ILOM SP configuration.

1 Open a web browser and log in to the CMM by entering the following URL:

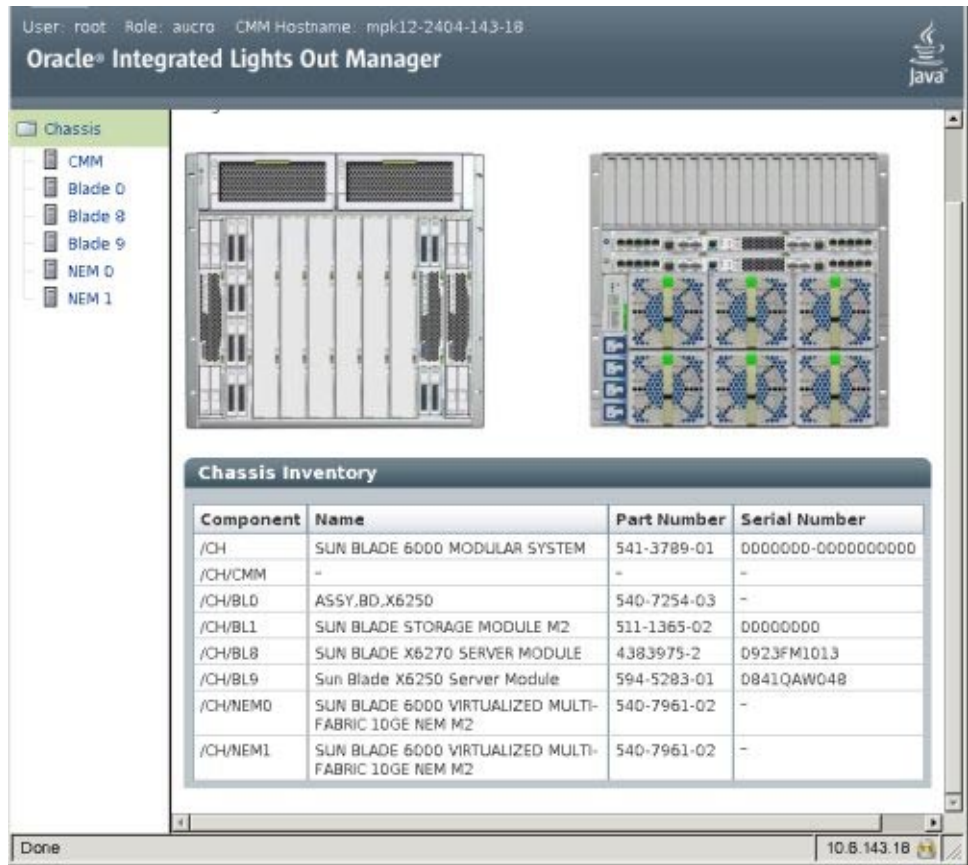
`http://chassis_sp_ipaddress/`

Where *chassis_sp_ipaddress* is the IP address of your chassis CMM.

The ILOM login page appears.

2 Log in to ILOM.

The CMM ILOM main page is displayed.

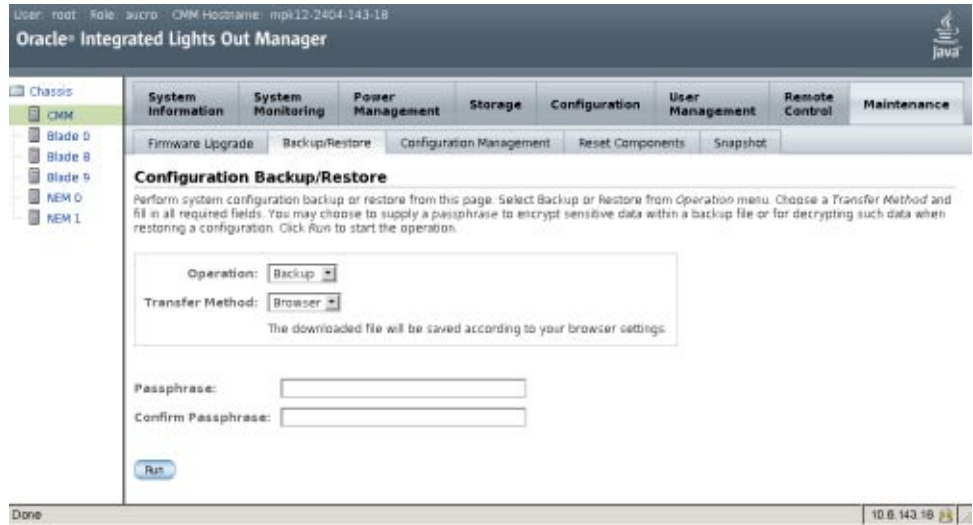


3 With CMM selected in the left pane, click the Maintenance tab for the CMM.

The CMM maintenance sub-tabs are displayed

4 Click the Backup/Restore tab.

The Backup/Restore page is displayed.



5 Select Backup from the Operating drop-down list.

6 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*.

Recovering Zoning Configurations

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

This section covers the following procedures

- [“How to Recover Zoning Configurations Using the Web Interface” on page 34](#)
- [“How to Recover Zoning Configurations Using the CLI” on page 37](#)

▼ How to Recover Zoning Configurations Using the Web Interface

Note – For advanced users or 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.

Before You Begin

- 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).
- 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 Using the CMM ILOM”](#) on page 30.

1 Open a web browser and log in to the CMM by entering the following URL:

`http://chassis_sp_ipaddress/`

Where *chassis_sp_ipaddress* is the IP address of your chassis service processor.

The ILOM login page appears.

2 Log in to ILOM.

The CMM ILOM main page is displayed.

User: root Role: auro CMM Hostname: mpk12-2404-143-18

Oracle® Integrated Lights Out Manager

Chassis

- CMM
- Blade 0
- Blade 8
- Blade 9
- NEM 0
- NEM 1

Chassis Inventory

Component	Name	Part Number	Serial Number
/CH	SUN BLADE 6000 MODULAR SYSTEM	541-3789-01	0000000-0000000000
/CH/CMM	-	-	-
/CH/BL0	ASSY,BD,X6250	540-7254-03	-
/CH/BL1	SUN BLADE STORAGE MODULE M2	511-1365-02	00000000
/CH/BL8	SUN BLADE X6270 SERVER MODULE	4383975-2	0923FM1013
/CH/BL9	Sun Blade X6250 Server Module	594-5283-01	0841QAW048
/CH/NEM0	SUN BLADE 6000 VIRTUALIZED MULTI-FABRIC 10GE NEM M2	540-7961-02	-
/CH/NEM1	SUN BLADE 6000 VIRTUALIZED MULTI-FABRIC 10GE NEM M2	540-7961-02	-

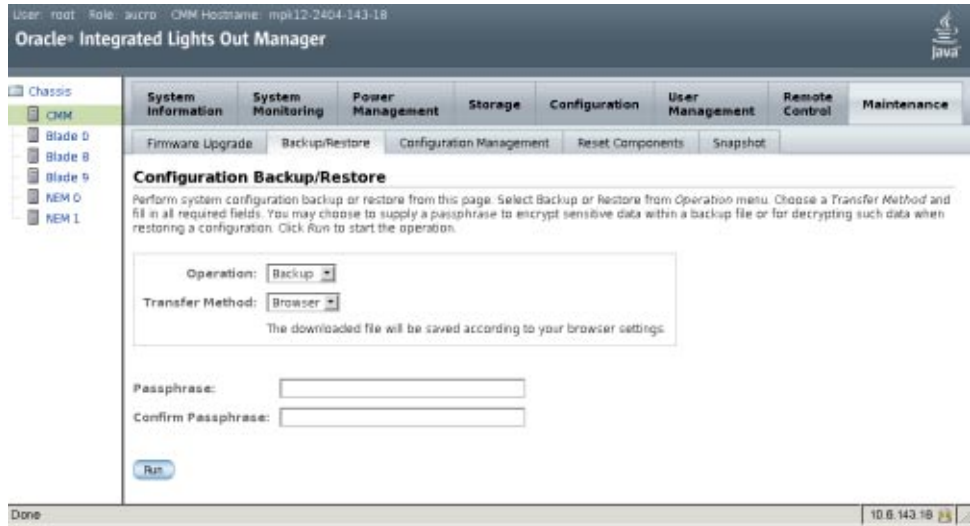
Done 10.6.143.18

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

The CMM maintenance sub-tabs are displayed.

4 Click the Backup/Restore tab.

The Backup/Restore page is displayed.



5 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*.

6 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 once the Restore operation is complete. A Restore operation typically takes two to three minutes to complete.

▼ How to Recover Zoning Configurations Using the CLI

Note – For advanced users or 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.

Before You Begin

- 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).
- 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 Using the CMM ILOM”](#) on page 30.

- 1 **Open a terminal window and establish an ssh connection to the CMM by entering the following command:**

```
# ssh -l username cmm_ip
```

Where *username* is your ILOM user name and *cmm_ip* is the IP address of the CMM.

The login prompt is displayed.

- 2 **Log in as root and enter the root password:**

```
/hostname/login: username
```

```
password: xxxxxxxx
```

After you have successfully logged in, the prompt will be displayed:

```
->
```

- 3 **Change to the /SP/config directory.**

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

- 4 **To initiate the Restore operation, enter the command:**

```
-> set load_uri=
```

```
transfer_method://username:password@ipaddress_or_hostname/directorypath/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.
- *directorypath* is the storage location on the remote system.
- *filename* is the name assigned to the backup file.

The Restore operation executes.

Installing or Replacing the Virtualized M2 NEM

This section describes how to replace a Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module (Virtualized M2 NEM) in a powered-on Sun Blade 6000 Series Chassis.



Caution – Damage to the NEM can occur as the result of careless handling or electrostatic discharge (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 antistatic mat and an ESD wrist strap. You can get an ESD wrist strap from any reputable electronics store or from Oracle as part number 250-1007.

This chapter contains the following sections:

- [“Installing a NEM” on page 39](#)
- [“Verifying NEM Installation” on page 41](#)
- [“Removing a NEM” on page 45](#)
- [“Replacing a NEM” on page 48](#)
- [“Installing and Removing SFP+ Optical Transceiver Modules” on page 51](#)
- [“Cabling the SPF+ Connectors” on page 53](#)

Installing a NEM

You can insert either one or two Virtualized M2 NEMs in the Sun Blade 6000 chassis. If you insert only one, install it in the lower slot (NEM 0).

Before installing your NEM, remove the NEM filler panel in the slot you plan to use.

Note – The Virtualized M2 NEM can be paired with a pass-through Gigabit Ethernet NEM or another Virtualized M2 NEM only. Mixing of Multi-Fabric NEM M2 types is not supported.

▼ How to Install a NEM

This procedure is for installing a NEM into an empty slot. If you are replacing a NEM, see [“Replacing a NEM” on page 48](#).

Before You Begin If you are installing a NEM into a powered on chassis, read [“Performing Hot Plug Insertion and Removal” on page 23](#) before installing the NEM.

1 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.

2 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.

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

The following illustration shows how to install the NEM.



4 Do one of the following:

- If you have installed a NEM in a powered off chassis, go to [Step 9](#).
- If you are installing the NEM in a powered on chassis, complete the rest of the procedure.

- 5 **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.

- 6 **Press the ATTN button once or use the ILOM Return to Service command**

See [“Using Hot Plug Commands” on page 138](#) for information on how to use the ILOM Return to Service command.

It can take up to 3 minutes for the NEM to come online and become fully operational.

Note – If you push the ATTN button twice within 5 seconds, the green Power/OK LED remains in standby blink mode, and the request for bringing the NEM online is cancelled.

- 7 **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.

- 8 **Refer to [“Restoring SAS Configurations After NEM Installation” on page 27](#) to see if any additional actions need to be taken for the SAS interfaces.**

- 9 **Connect cables, as needed. See [“Cabling the SPF+ Connectors” on page 53](#).**

Verifying NEM Installation

The following topics are covered in this section:

- [“How to Verify Installation With the CMM ILOM Web Interface” on page 41](#)
- [“How to Verify Installation With the CMM ILOM CLI” on page 43](#)

▼ How to Verify Installation With the CMM ILOM Web Interface

- 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.

User: root Role: auroa CMM Hostname: mpk12-2404-143-18

Oracle® Integrated Lights Out Manager

Chassis

- CMM
- Blade 0
- Blade 8
- Blade 9
- NEM 0
- NEM 1

Component Management

View component information, prepare to install or remove a component, update firmware, or clear fault status from this page. To modify a component, select the radio button next to that component, then choose an option from the Action drop down list. Components without radio buttons cannot be modified. Choosing the Prepare to Remove action shuts down the selected component and lights its blue Ready to Remove LED. To view further details, click on a Component Name.

Component Status

Actions

Filter: All Components

	Component Name	Type	Fault Status	Ready to Remove Status
-	/CH	Chassis	OK	-
-	/CH/CMM	Chassis Monitoring Module	-	-
-	/CH/CMM/SP	Service Processor	-	-
-	/CH/CMM/NET0	Network Interface	-	-
-	/CH/MIDPLANE	Power Distribution Board	-	-
<input type="radio"/>	/CH/BL0	Blade FRU	OK	-
<input type="radio"/>	/CH/BL1	Blade FRU	Faulted	-
-	/CH/BL1/MB	Motherboard	-	-
<input type="radio"/>	/CH/BL8	Blade FRU	OK	-
<input type="radio"/>	/CH/BL9	Blade FRU	OK	-
<input type="radio"/>	/CH/NEM0	Network Express Module	Faulted	Not Ready
-	/CH/NEM0/MB	Motherboard	-	-
<input type="radio"/>	/CH/NEM1	Network Express Module	OK	Not Ready

42

Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module User's Guide • September 2012, Rev A

Information about the NEM displays, as shown in the following example:



If the NEM does not appear 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.

▼ How to Verify Installation With the 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 documented in [“ILOM Supplement” on page 121](#).

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

1 Log in to the CMM ILOM.

2 Enter the following command:

```
-> show /CH/NEMx
```

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

The NEM and the NEM Field Replaceable Unit (FRU) information appears in the CLI.

Targets:

```
MB
SAS
SP
PRSNT
STATE
ERR
OK
SERVICE
OK2RM
LOCATE
```

Properties:

```
type = Network Express Module
ipmi_name = NEM1
system_identifier = mpk12-2404-143-17
fru_name = SUN BLADE 6000 VIRTUALIZED MULTI-FABR IC 10GE NEM M2
fru_version = FW 3.0.10.16, SAS 5.3.2.0
fru_part_number = 540-7961-02
fault_state = Faulted
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 To verify whether the NEM has joined the SAS fabric, type the following command:

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

```
/STORAGE/sas_zoning
```

Targets:

```
BL0
```

```
BL6
BL7
BL8
BL9
NEM0
NEM1
```

Properties:

```
zone_management_state = disabled
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 and verify that you have installed the latest CMM ILOM firmware.

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

<http://www.oracle.com/pls/topic/lookup?ctx=ilom30&id=homepage>

Removing a NEM

This procedure is for removing a NEM if you do not plan to replace it. If you plan to replace a NEM, see “Replacing a NEM” on page 48.

▼ How to Remove a NEM

You can remove a NEM from a powered-on chassis using a hot-plug operation, as well as from a powered-off chassis. 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.

Before You Begin

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 23 and follow the instructions in “How to Prepare the PCIe Interface for Hot Removal” on page 25 and “How to Prepare SAS Interfaces for Hot Plug Removal” on page 27.

- 1 In the rear of the chassis, locate the NEM that you want to remove.
- 2 Do one of the following:
 - If you are removing a NEM from a powered-off chassis, go to [Step 6](#).

- If you are removing a NEM from a powered-on chassis, follow the rest of the procedure.

3 Push the ATTN button once or use the ILOM Prepare to Remove command.

See “Using Hot Plug Commands” on page 138 for information on how to use the ILOM Prepare to Remove command.

Note – Pushing the ATTN button twice within 5 seconds will cause the Power/OK LED to return to solid green. This will cancel the request for a NEM removal.

4 View the behavior of the Power/OK LED and OK-to-Remove LED or the ILOM interface.

- **The Ready to Remove action is successful if the following occurs within a 65 minute time-frame:**

- The NEM green Power/OK begins to slow blink.

Note – Slow blink is 1 Hz.

- 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
- If /CH/NEMx/OK reads Off in either the ILOM web interface or CLI.

- **The Ready to Remove action is not successful if the following occurs:**

- The NEM green Power/OK LED begins to slow blink, then returns to solid green, and the OK-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
- If /CH/NEMx/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 24](#) to remove the NEM.


6 Remove all cables from the NEM.

Note – If applicable, wait 5 minutes after CMM boot up or NEM insertion before removing 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.

The following illustration shows how to remove a NEM.



Replacing a NEM

If a NEM fails, you need to replace it. 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 SAS-2 NEM (for example replacing a Multi-Fabric NEM with a 10GbE Virtualized M2 NEM) and another SAS-2 NEM is also 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 IO to all disks on the disk module that are not in hardware RAID volumes.



How to Replace a NEM

Before You Begin

If you are removing a NEM from a powered on chassis, read the hot-plug guidelines [“Performing Hot Plug Insertion and Removal” on page 23](#) and follow the instructions in [“How to Prepare the PCIe Interface for Hot Removal” on page 25](#)



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 Using the CMM ILOM” on page 30](#).

- 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, follow the rest of the procedure.
- 3 Push the ATTN button once or use the ILOM Prepare to Remove command.

See [“Using Hot Plug Commands” on page 138](#) for information on how to use the ILOM Prepare to Remove command.

Note – Pushing the ATTN button twice within 5 seconds will cause the Power/OK LED to return to solid green. This will cancel the request for a NEM removal.

4 View the behavior of the Power/OK LED and OK-to-Remove LED or the ILOM interface.

- **The Ready to Remove action is successful if the following occurs within a 65 minute time-frame:**
 - The NEM green Power/OK begins to slow¹ blink.
 - The OK-to-Remove blue LED illuminates and the Power OK button 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
 - If /CH/NEMx/OK reads Off in either the ILOM web interface or CLI.
- **The Ready to Remove action is not successful if the following occurs:**
 - The NEM green Power/OK LED begins to slow blink, then returns to solid green, and the OK-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
 - If /CH/NEMx/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 24](#) 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.

¹ Slow blink is 1 Hz

- 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.

The following illustration shows how to remove a 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, follow the rest of the procedure.

- 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 or use the ILOM Return to Service command.**

See [“Using Hot Plug Commands” on page 138](#) for information on how to use the ILOM Return to Service command.

It can take up to 3 minutes for the NEM to come online.

Note – If you push the ATTN button twice within 5 seconds, the green Power/OK LED remains in standby blink 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 27](#) to see if any additional actions need to be taken for the SAS interfaces.

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

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 [“10 GbE NIC Virtualization” on page 13](#).

This section covers the following topics:


- [“How to Install an SFP+ Optical Transceiver Module” on page 51](#)
- [“How to Remove an Optical Transceiver SFP+ Module” on page 52](#)

▼ How to Install an SFP+ Optical Transceiver Module

- 1 **Pull the locking handle into the full horizontal position until you feel the handle click into position.**

- 2 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.
- 3 Applying even pressure at both corners of the SFP+ module, push the module until it is firmly seated in the slot.
- 4 Push the handle closed to lock the SFP+ module in place.


Note – If you pull the locking handle down when the SFP+ module is installed, you should 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.



▼ How to Remove an Optical Transceiver SFP+ Module

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



Cabling the SPF+ Connectors

The following table shows the optical module (OM) multimode fiber (MMF) and single-mode fiber (SMF) cabling options available for the Virtualized Multi-Fabric 10GbE M2 NEM SPF+ connectors.

TABLE 6 Cabling Guidelines for SPF+ Connectors

SPF+ Connector Type	Cable type	Reach
Short range (SR)	OM 1 MMF	33 meters
	OM 2 MMF	50 meters
	OM 3 MMF	300 meters

Note – The optical transceiver for the Virtualized Multi-Fabric 10GbE M2 NEM uses a long connector (LC). The connection on the other side of the cable might use a different connector type.



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.

Booting Over the Virtualized M2 NEM 10-Gigabit Ethernet Port

In order to perform a network boot via the Virtualized M2 NEM, make sure that the Virtualized M2 NEM and the boot server are on the same LAN. If you are loading an OS using PXE, you also need configure the boot server so that it knows which OS should be downloaded to the blade and what IP address should be used for the Virtualized M2 NEM's MAC.

The following topics are covered in this section:

- [“Booting over the Network With an x86 Blade Server” on page 55](#)
- [“Booting over the Network With a SPARC Blade Server” on page 61](#)

Booting over the Network With an x86 Blade Server

The x86 server blade system BIOS only recognizes that at least one Virtualized NEM M2 is available as a network boot device. If more than one Virtualized M2 NEMs is installed, only one Virtualized M2 NEM device is displayed in the BIOS boot list

The Virtualized M2 NEM is designed to boot from the first detected PXE server. If there are two NEMs in the chassis, and both are connected to a PXE server, you must use MAC address assignment booting to boot from the correct NEM.

If you plan to use the F12 network booting method, follow the instructions in these two procedures:

- [“How to Set Up the BIOS for Booting From the Virtualized M2 NEM” on page 55](#)
- [“How to Use the F12 Key to Initiate the Network Boot” on page 59](#)

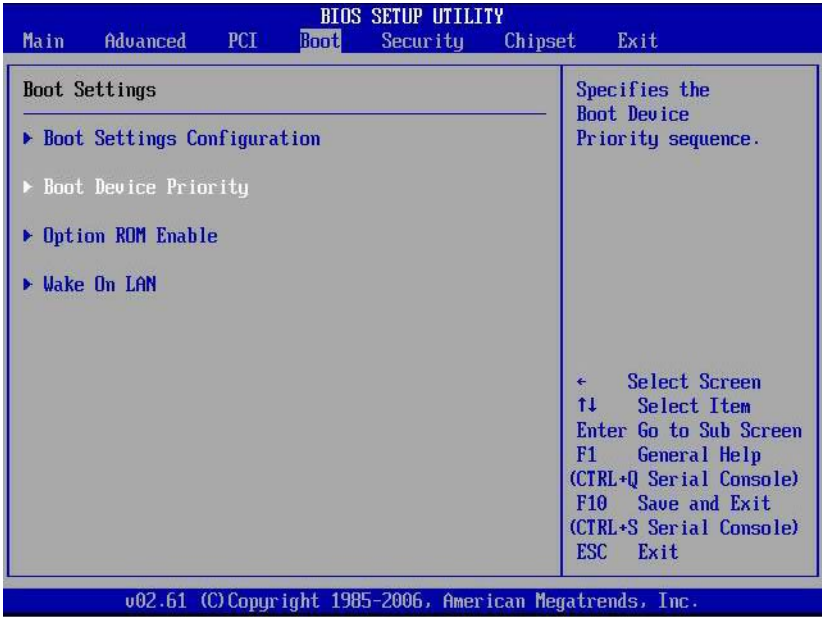
If you plan to use the F8 network booting method, follow the instructions in [“How to Use the F8 Key to Boot Off the PXE Server” on page 60](#).

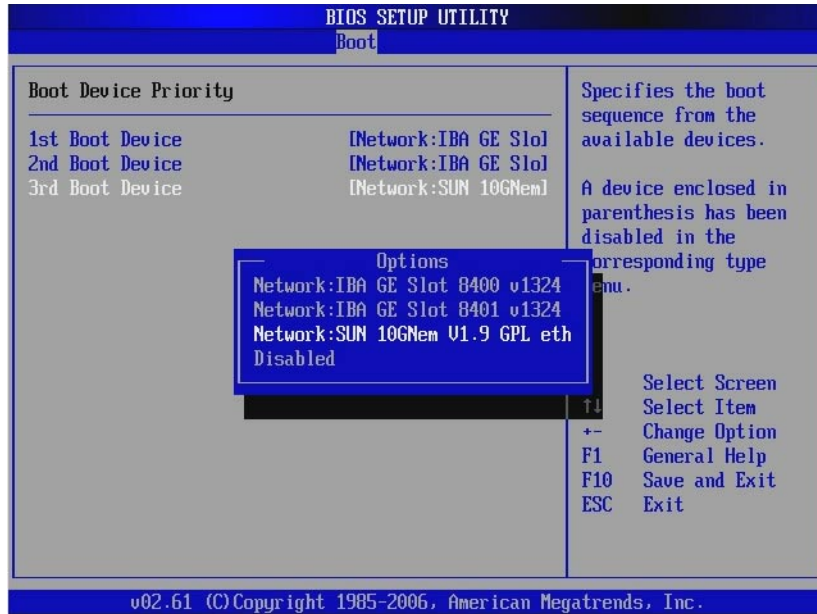
▼ How to Set Up the BIOS for Booting From the Virtualized M2 NEM

- 1 Power on the host blade server.

- 2 When the BIOS boot screen appears, 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.

The following two screens show examples of Intel and AMD-based BIOS Setup Utility screens. The actual screens might vary by platform.

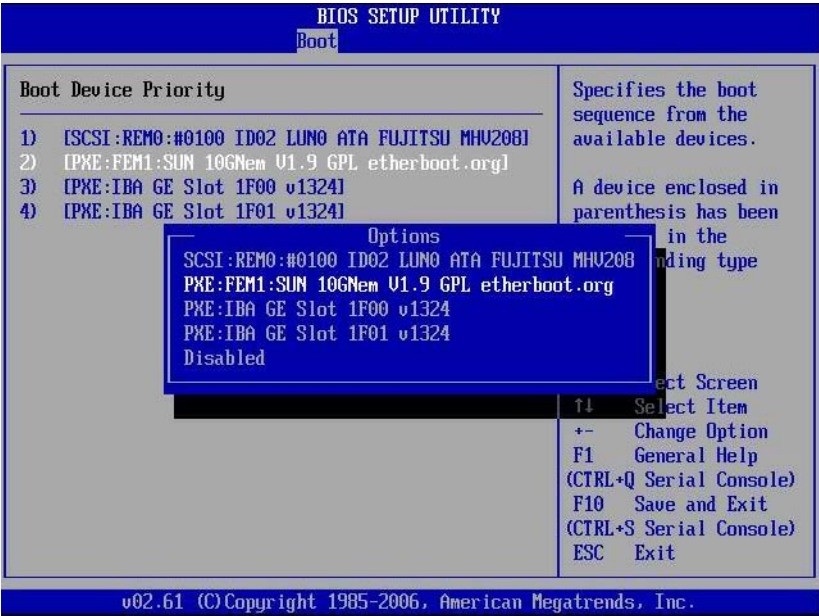





- 5 Select the first boot device position and press Enter.

A drop-down list appears with the available boot devices listed.

- 6 From the drop-down list, select **SUN 10GNem V1.9 GPL etherboot.org** as the first boot device.
- The following two screens show examples of Intel and AMD-based BIOS Setup Utility screens. The actual screens might vary by platform.





7 Press F10 to Save and Exit the BIOS Setup Utility.

The host blade server reboots.

Note – The Virtualized M2 NEM is designed to boot from the first detected PXE server. If there are two NEMs in the chassis, and both are connected to a PXE server, you must use MAC address assignment booting to boot from the correct NEM.

8 Connect the Ethernet cable to the NEM Ethernet port.

See the driver installation section in this chapter 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.

▼ How to Use the F12 Key to Initiate the Network Boot

1 Power on the host server blade.

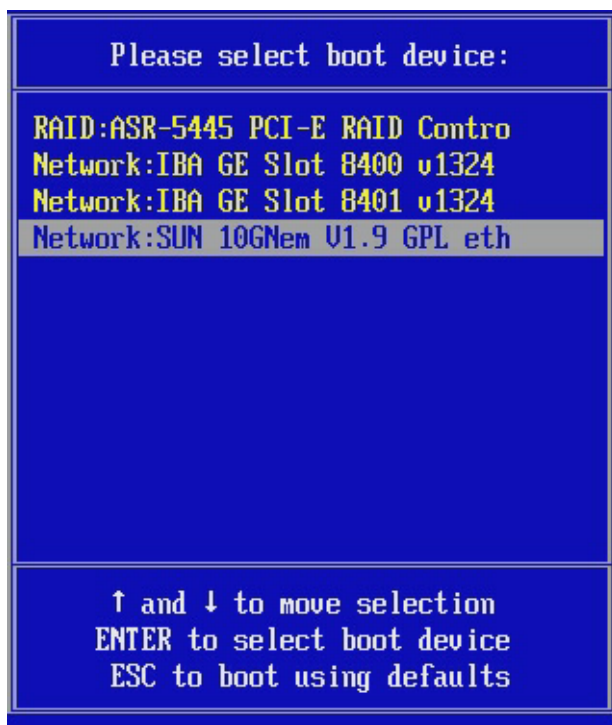
2 Follow the instructions in “[How to Set Up the BIOS for Booting From the Virtualized M2 NEM](#)” on [page 55](#) to set up the Virtualized M2 NEM as the first network boot device.

3 Press F12 to boot from the network.

Note – If the Virtualized M2 NEM doesn't detect a PXE Server, the CPU Host Blade will attempt to boot from the On-Board 1Gbps NIC devices.

▼ How to 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 SUN 10GNem V1.9 GPL etherboot.org NIC from the boot list.
- 4 Press Enter to boot from the PXE server.

Note – If the Virtualized M2 NEM doesn't detect a PXE Server, the CPU Host Blade will attempt to boot from the On-Board 1Gbps NIC devices.

Booting over the Network With a SPARC Blade Server

Use the following procedure to boot over the network using the Virtualized M2 NEM with a SPARC blade server.

▼ How to Boot Over the Network Using the OpenBoot on a SPARC System

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

This enables the boot server to answer RARP calls from the blade for resolving the IP address for the Virtualized M2 NEM 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 for details.

- 3 After the boot server has been configured, do one of the following:

- If the SPARC blade is not currently running Solaris: Logon to the console of the SPARC blade, power the blade on, and boot to the OpenBoot ok prompt.
- If the blade is already running Solaris, type the following in a terminal window to go to the ok prompt:
`sync;sync;halt`

- 4 Type `show-nets` to see all the network interfaces available for the blade to use. One of them should be the Virtualized M2 NEM. For example:

```
{0} ok show-nets
a) /pci@7c0/pci@0/pci@8/network@0
b) /pci@780/pci@0/pci@1/network@0,1
c) /pci@780/pci@0/pci@1/network@0
q) NO SELECTION
Enter Selection, q to quit:
```

The path with a single network node should be the Virtualized M2 NEM.

- 5 Change directory to that path.

```
cd /pci@7c0/pci@0/pci@8/network@0
```

- 6 Type command `.properties` to check its properties. If it is a Virtualized M2 NEM, you should see:

```
{0} ok .properties
assigned-addresses      82520010 00000000 11000000 00000000 01000000
                        82520018 00000000 10400000 00000000 00008000
                        82520020 00000000 00000000 00000000 00008000
                        82520030 00000000 10600000 00000000 00200000
local-mac-address       00 14 4f 63 00 09
```

```

phy-type      xgf
reg           00520000 00000000 00000000 00000000 00000000
              03520010 00000000 00000000 00000000 01000000
              03520018 00000000 00000000 00000000 00008000
              03520020 00000000 00000000 00000000 00008000
              02520030 00000000 00000000 00000000 00100000
version       Sun Blade 6000 Virtualized Multi-Fabric 10GbE NEM FCode
1.2 09/01/14
board-model   501-7995-04
model         SUNW,pcie-hydra
compatible    pciex108e,aaaa.108e.aaaa.1
              pciex108e,aaaa.108e.aaaa
              pciex108e,aaaa.1
              pciex108e,aaaa
              pciexclass,020000
              pciexclass,0200
address-bits   00000030
max-frame-size 00002400
network-interface-type ethernet
device_type    network
name           network
fcode-rom-offset 00006200
interrupts     00000001
cache-line-size 00000010
class-code     00020000
subsystem-id   0000aaaa
subsystem-vendor-id 0000108e
revision-id    00000001
device-id      0000aaaa
vendor-id      0000108e

```

7 Type device-end then do the netboot.

```

{0} ok device-end
{0} ok /pci@400/pci@0/pci@9/pci@0/pci@1/network@0:dhcp

```

```

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

```

If successful, you should see something similar to the following 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: hxge1: 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 Solaris 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 appears. The Solaris 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.

Installing and Configuring the hxge Driver on a Solaris SPARC or x86 Platform

The Virtualized NEM supports Solaris 10/09 and later. The hxge device Solaris is included with this version of Solaris, so it is not necessary to install the driver.

This section explains how to configure the hxge driver on a Solaris system. The hxge GigabitEthernet driver (hxge(7D)) is a multi-threaded, loadable, clonable, GLD-based STREAMS driver.

This section covers the following topics:

- “How to Configure the Network Host Files” on page 65
- “Configuring the hxge Device Driver Parameters” on page 66
- “Configuring the Jumbo Frames Feature” on page 69

▼ How to Configure the Network Host Files

This section describes how to configure the network host files after you install the hxge driver on your system.

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

```
# grep hxge /etc/path_to_inst
"/pci@7c,0/pci10de,5d@e/pci108e,aaaa@0" 0 "hxge"
```

In this example, the device instance is from a Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module installed in the chassis. The instance number is shown in *italics* in this example.

- 2 Set up the NEM's hxge interface.**

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

```
# ifconfig hxge0 plumb ip-address netmask netmask-address broadcast + up
```

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

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

To use the NEM's hxge interface in the Step 1 example, create an `/etc/hostname.hxge x` file, where x is the number of the hxge interface. If the instance number were 1, the filename would be `/etc/hostname.hxge1`.

Follow these guidelines for the host name:

- **The `/etc/hostname.hxge $number$` file must contain the host name for the appropriate hxge interface.**
- **The host name must be different from the host name of any other interface. For example: `/etc/hostname.hxge0` and `/etc/hostname.hxge1` 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.hxge $number$` file required for a system named `zardoz-c10-bl1`.

```
# cat /etc/hostname.hxge0
zardoz-c10-bl1
```

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

For example:

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

Configuring the hxge Device Driver Parameters

The hxge device driver controls the Virtualized M2 NEM Ethernet interfaces. You can manually set the hxge driver parameters to customize each device in your system.

The following procedures describe the two ways to set the hxge device driver parameters:

- [“How to Set Parameters Using the `hxge.conf` File” on page 67](#)
- [“How to Specify Parameter Values Using the `ndd` Utility” on page 68](#)

Note – If you use the `ndd` utility, the parameters are valid only until you reboot the system. This method is good for testing parameter settings.

▼ How to Set Parameters Using the hxge.conf File

The hxge device driver configuration file is in the following directory

/kernel/drv/hxge.conf

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

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

```
# cat /kernel/drv/hxge.conf
#
#
# driver.conf file for Sun 10Gb Ethernet Driver (hxge)
#
#----- Jumbo frame support -----
# To enable jumbo support,
# accept-jumbo = 1;
#
# To disable jumbo support,
# accept-jumbo = 0;
#
# Default is 0.
#
#----- Receive DMA Configuration -----
#
# rxdma-intr-time
#     Interrupts after this number of NIU hardware ticks have
#     elapsed since the last packet was received.
#     A value of zero means no time blanking (Default = 8).
#
# rxdma-intr-pkts
#     Interrupt after this number of packets have arrived since
#     the last packet was serviced. A value of zero indicates
#     no packet blanking (Default = 0x20).
#
# Default Interrupt Blanking parameters.
#
# rxdma-intr-time = 0x8;
# rxdma-intr-pkts = 0x20;
#
#----- Classification and Load Distribution Configuration -----
#
# class-opt.****.***
#     These variables define how each IP class is configured.
#     Configuration options includes whether TCAM lookup
#     is enabled and whether to discard packets of this class
#
#     supported classes:
#     class-opt-ipv4-tcp class-opt-ipv4-udp class-opt-ipv4-sctp
#     class-opt-ipv4-ah class-opt-ipv6-tcp class-opt-ipv6-udp
#     class-opt-ipv6-sctp class-opt-ipv6-ah
```

```
#          Configuration bits (The following bits will be decoded
#          by the driver as hex format).
#
#
#          0x10000:          TCAM lookup for this IP class
#          0x20000:          Discard packets of this IP class
#
# class-opt-ipv4-tcp = 0x10000;
# class-opt-ipv4-udp = 0x10000;
# class-opt-ipv4-sctp = 0x10000;
# class-opt-ipv4-ah = 0x10000;
# class-opt-ipv6-tcp = 0x10000;
# class-opt-ipv6-udp = 0x10000;
# class-opt-ipv6-sctp = 0x10000;
# class-opt-ipv6-ah = 0x10000;
#
#
#----- FMA Capabilities -----
#
# Change FMA capabilities to non-default
#
# DDI_FM_NOT_CAPABLE 0x00000000
# DDI_FM_EREPOR_T_CAPABLE 0x00000001
# DDI_FM_ACCCHK_CAPABLE 0x00000002
# DDI_FM_DMACHK_CAPABLE 0x00000004
# DDI_FM_ERRCB_CAPABLE 0x00000008
#
# fm-capable = 0xF;
#
# default is DDI_FM_EREPOR_T_CAPABLE | DDI_FM_ERRCB_CAPABLE = 0x5
```

In the following example, the NEM discards the TCP traffic for this blade system. In other words, the hxge driver does not receive any TCP traffic.

```
class-opt-ipv4-tcp = 0x20000;
```

In the following example, the FMA functionality will be disabled.

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

- 2 For the new parameters to be effective, reload the hxge driver or reboot the system.

▼ How to Specify Parameter Values Using the ndd Utility

This section describes how to modify and display parameter values using the ndd utility.

Before you use the ndd utility to get or set a parameter for a hxge device, you must specify the device instance for the utility.

- 1 Use the **ifconfig** command to identify the instance associated with hxge device.

```
# ifconfig -a
hxge0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
inet 192.168.1.29 netmask ffffffff broadcast 192.168.1.255
ether 0:14:4f:62:1:3
```

2 List all of the parameters supported by the hxge driver.

```
# ndd -get /dev/hxge0 ?
?                               (read only)instance          (read only)
rxdma_intr_time                (read and write)
rxdma_intr_pkts                (read and write)
class_opt_ipv4_tcp              (read and write)
class_opt_ipv4_udp              (read and write)
class_opt_ipv4_ah               (read and write)
class_opt_ipv4_sctp             (read and write)
class_opt_ipv6_tcp              (read and write)
class_opt_ipv6_udp              (read and write)
class_opt_ipv6_ah               (read and write)
class_opt_ipv6_sctp             (read and write)
```

Note that read only parameters cannot be changed.

3 Display the value of a parameter.

The following is an example for the `rxdma_intr_time` parameter:

```
# ndd -get /dev/hxge0 rxdma_intr_time
8
```

4 Modify the value of a parameter.

The following is an example for the `rxdma_intr_time` parameter. It changes the `rxdma_intr_time` parameter from `0x8` to `0x10`:

```
# ndd -set /dev/hxge0 rxdma_intr_time 0x10
# ndd -get /dev/hxge0 rxdma_intr_time
10
```

Configuring the Jumbo Frames Feature

This section describes how to enable the Jumbo Frames feature. It contains the following sections:

- [“Jumbo Frames Overview” on page 69](#)
- [“Checking the Jumbo Frames Configurations” on page 70](#)
- [“How to Enable Jumbo Frames in a Solaris Environment Using `hxge.conf`” on page 71](#)
- [“How to Check Layer 2 Configuration” on page 72](#)
- [“How to Check Layer 3 Configuration” on page 72](#)

Jumbo Frames Overview

Configuring Jumbo Frames enables the 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 Virtualized Multi-Fabric 10GbE M2 NEM driver capability.

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

Checking the Jumbo Frames Configurations

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 of the four receive DMA channels on interface 1, for example:

```
# kstat -m hxge | grep rdc_pac
rdc_packets          120834317rdc_packets          10653589436
rdc_packets          3419908534
rdc_packets          3251385018
# kstat -m hxge | grep rdc_jumbo      rdc_jumbo_pkts          0
rdc_jumbo_pkts      0
rdc_jumbo_pkts      0
rdc_jumbo_pkts      0
```

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

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

```
# kstat -m hxge -n RDC_0
module: hxge          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_jumbo_pkts        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 `hxge0` interface, for example:

```
# kstat -m hxge -n hxge0
module: hxge                               instance: 0
name:   hxge0                             class:   net
brdcstrcv          0
brdcstxmt          0
collisions         0
crtime             134.825726986
ierrors            0
ifspeed            100000000000
ipackets           265847787
ipackets64         17445716971
multircv           0
multixmt           0
norcbuf            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 hxge
```

▼ How to Enable Jumbo Frames in a Solaris Environment Using `hxge.conf`

1 Enable jumbo frames using the `hxge.conf` file .

For example:

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

Note that the maximum (default) jumbo frame size is 9216 bytes including a 16 byte hardware header. It is not recommended to change this size. However, it may be changed by including the following line in the `/etc/system` file.

```
set hxge_jumbo_frame_size = value
```

The value should be between 1500 and 9216.

2 Reboot the system.

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

▼ How to Check Layer 2 Configuration

- View the maximum transmission unit (MTU) configuration of an `hxge` instance at any time with the `ifconfig` command.

```
# ifconfig -a
hxge0: flags=1000843 <UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 9178 index 4
inet 192.168.1.29 netmask ffffffff 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 cyclic redundancy check (crc) checksum.

▼ How to Check Layer 3 Configuration

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

For example:

```
# dladm show-link
nge0          type: non-vlan  mtu: 1500      device: nge0
nge1          type: non-vlan  mtu: 1500      device: nge1
nxge0         type: non-vlan  mtu: 1500      device: nxge0
nxge1         type: non-vlan  mtu: 1500      device: nxge1
hxge0         type: non-vlan  mtu: 9178      device: hxge0
```


Installing and Configuring the hxge Driver on a Linux Platform

This section covers the following topics:

- [“Installing and Removing the Driver on a Linux Platform” on page 73](#)
- [“Configuring the Network Interface” on page 78](#)
- [“Checking and Testing the hxge Device” on page 82](#)
- [“Changing the hxge Driver Configuration” on page 84](#)
- [“Troubleshooting the Driver” on page 88](#)
- [“Configuring Jumbo Frames” on page 91](#)

Installing and Removing the Driver on a Linux Platform

This section explains how to download, install, and remove the Linux hxge driver. The hxge 10 Gigabit Ethernet driver (`hxge(1)`) is a parallel multi-threaded and loadable driver supporting up to four transmit channels and four receive channels in simultaneous operation, taking advantage of up to eight CPUs to spread the 10 GbE network traffic and increase overall network throughput.

This section covers the following topics:

- [“How to Download and Install the Driver on a Linux Platform” on page 73](#)
- [“How to Remove the Driver From a Linux Platform” on page 77](#)

▼ How to Download and Install the Driver on a Linux Platform

- 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 Virtualized Multi-Fabric 10GbE M2 NEM hardware and installed and loaded the driver, a new eth device appears. This will be the eth device for the NEM.

2 Access the hxge drivers through one of the following methods:

- **Locate and copy the hxge 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 (OEL) use the corresponding Red Hat Enterprise Linux (RHEL) driver.

The drivers are located in one of the following directories:

/linux/drivers/rhel5.4

/linux/drivers/sles10/sp3

/linux/drivers/sles11

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

- **Download latest drivers packages from the Drivers and Firmware link on the following web site: <http://oracle.com/goto/10gbenemm2>.**

3 Once the Linux drivers file is downloaded, uncompress and unpack the drivers. In this example, the download directory is tge10:

```
host #>mkdir tge10
host #>cp xxx/SB6000_Virt_MF_10GbE_NEM_Tools_And_Drivers_Linux_48230ai.bz2 tge10
host #>cd tge10
host #>tar -xjf SB6000_Virt_MF_10GbE_NEM_Tools_And_Drivers_Linux_48230ai.bz2
```

4 Select and install the appropriate OS driver package.

In this step, RHEL5.4 Linux driver is selected for OEL version 5.4.

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

```
host #> cd linux/drivers/rhel5.4
host #> ls -l
total 897
-rw-r--r-- 1 root root 895938 Mar 8 2010 hxge-1.2.7_rhel54-1.x86_64.rpm
-rw-r--r-- 1 root root 1485 Mar 8 2010 readme.html
```

The file `hxge-1.2.7_rhel54-1.x86_64.rpm` is the package which is to be installed.

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

```
host #> lsb_release -a
LSB Version: :core-3.1-amd64:core-3.1-ia32:core-3.1-noarch:graphics-3.1-amd64:
graphics-3.1-ia32:graphics-3.1-noarch
Distributor ID: EnterpriseEnterpriseServer
Description: Enterprise Linux Enterprise Linux Server release 5.4 (Carthage)
```

```

Release:      5.4
Codename:     Carthage

host #> rpm -q hxge
package hxge is not installed

```

b. Identify the OS-specific subdirectory, and verify that no hxge driver is currently installed.

Note – If an hxge driver is already installed, uninstall the driver, in order to avoid complications. See [“How to Remove the Driver From a Linux Platform” on page 77](#) for instructions on removing the driver. The update command (rpm -u) is not supported for updating the hxge driver.

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

```

host #> rpm -ivh hxge-1.2.7_rhel54-1.x86_64.rpm
Preparing... ##### [100%]
 1:hxge ##### [100%]
post Install Done

```

Once you have installed the hxge driver, you can immediately load the driver. If the NEM is physically and electrically installed, the driver automatically attaches to it and makes it available to the system. Alternatively, on the next system reset and reboot, the hxge driver automatically loads if there are any NEM devices present and detected.

5 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 command sample below has been edited for brevity. Sections marked [...] denote output that has been removed.

```

host #> lspci
[...]
00:1f.0 ISA bridge: Intel Corporation 82801JIR (ICH10R) LPC Interface Controller
00:1f.2 SATA controller: Intel Corporation 82801JI (ICH10 Family) SATA AHCI Controller
00:1f.3 SMBus: Intel Corporation 82801JI (ICH10 Family) SMBus Controller
13:00.0 Ethernet controller: Sun Microsystems Computer Corp. Unknown device aaaa (rev 01)
19:00.0 Ethernet controller: Sun Microsystems Computer Corp. Unknown device aaaa (rev 01)
1f:00.0 Ethernet controller: Intel Corporation 82575EB Gigabit Network Connection (rev 02)
1f:00.1 Ethernet controller: Intel Corporation 82575EB Gigabit Network Connection (rev 02)
[...]

```

Device code 0xAAAA (Unknown device aaaa (rev 01) output) is the Virtualized M2 NEM Device; presence of this line indicates the NEM is visible and available to the system.

b. Manually load the hxge driver.

```

host #> modprobe hxge

```

c. Verify the driver is loaded.

```

host #> lsmod | grep hxge
hxge                175440  0

```

```
host#> modinfo hxge
filename:      /lib/modules/2.6.18-164.el5/kernel/drivers/net/hxge.ko
version:      1.2.7
license:      GPL
description:   Sun Microsystems(R) 10 Gigabit Network Driver
author:       Sun Microsystems, james.puthukattukaran@sun.com
srcversion:   270F053A5DE6A454D1D224D
alias:        pci:v0000108Ed0000AAAAsv*sd*bc*sc*i*
depends:
vermagic:     2.6.18-164.el5 SMP mod_unload gcc-4.1
parm:         enable_jumbo:enable jumbo packets (int)
[etc.]
```

- If you get the output above, the driver is loaded into memory and actively running.
- If the `modprobe` command fails, you will receive the following output:

```
host #> modprobe hxge
FATAL: Module hxge not found.
```

This indicates that you probably installed the wrong driver version. Uninstall the `hxge` driver and 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.

6 Identify the NEM eth device.

Once the NEM has been correctly installed, and the `hxge` software 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; `eth2`, is for a new NEM Ethernet network interface device. 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 [“Installing a NEM” on page 39](#).

b. Make sure that the `eth2` driver is the correct Ethernet driver for the Virtualized M2 NEM.

```
host #> ethtool -i eth2
driver: hxge
version: 1.2.7
firmware-version: N/A
bus-info: 0000:19:00.0
```

c. For more detail on `eth2`, use the `ifconfig` command.

```
host #> ifconfig eth2

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:f9000000-fa000000
```

The eth2 device is active and available to the system, but has not yet been configured (assigned an IP address). See the next section for details on configuring the NEM for the Linux OS.

▼ How to Remove the Driver From a Linux Platform

Removing the Linux hxge driver is a two-step process.

- 1 **Use the `modprobe -r` command to unload the hxge driver at any time, without actually uninstalling the driver.**

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

Once unloaded, you can manually load the hxge driver again, using the `modprobe` command; the driver has not been uninstalled.

- 2 **Uninstall the hxge driver.**

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

```
host #> rpm -q hxge
hxge-1.2.7_rhel54-1
```

```
host #> rpm -e hxge
Uninstall Done.
```

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

Configuring the Network Interface

Note – The information in this section is intended as a guideline for configuring network interfaces for the `hxge` driver. For more detailed information, refer to the administrative documentation for the Linux version that you have installed.

Before using the NEM Hydra 10GbE 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 Virtualized M2 NEM as demonstrated in the installation section). At a minimum, you must assign a network (TCP) IP address and netmask to each network interface.

This section covers the following topics:

- [“About Temporary `hxge` Network Interface Configuration” on page 78](#)
- [“About Permanent `hxge` Network Interface Configuration” on page 78](#)
- [“How to Bring the Interface Online Manually” on page 79](#)
- [“How to Configure the Network Interface File Automatically for Red Hat Linux” on page 80](#)
- [“How to Configure the Network Interface Automatically for SUSE Linux” on page 81](#)

About Temporary `hxge` Network Interface Configuration

To temporarily configure the Virtualized M2 NEM Ethernet interface (for example, to test it out), use the `ifconfig` command.

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

About Permanent `hxge` Network Interface Configuration

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

Linux maintains a separate network interface configuration file for each network interface possible 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.

▼ How to Bring the Interface Online Manually

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

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

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

2 Verify 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 Ethernet 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) local area network (or 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 also support IPv6 network communications)
- Reported state now shows up

3 RX (receive) and TX (transmit) packet counters are increasing, showing active traffic being routed through the newly configured Virtualized M2 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 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 (but no longer in up state, packet counters now zero again).

▼ How to 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 Oracle Enterprise Linux (OEL) version at

<http://www.oracle.com/technology/tech/linux/install/index.html> or RHEL version at:
<http://docs.redhat.com/docs/en-US/index.html>

For OEL or RHEL systems, the interface configuration files are named `ifcfg-ethn` (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
# Sun NEM/Hydra 10GbE
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 hand using a text editor. The first line `# Sun NEM Hydra 10GbE` 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 (up) 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 runlevel 3).

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

or

```
host #> ifup eth2
```

▼ How to Configure the Network Interface Automatically for SUSE Linux

You can configure the network interface for the SUSE Linux Server (SLES) platform either by using the GUI or manually editing the configuration file.

This chapter describes how to manually edit the configuration file. For instructions on configuring the network interface using the GUI, refer to the documentation for your SLES version at: <http://www.novell.com/documentation/suse.html>

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 proceeding examples), and reside in the `/etc/sysconfig/network` system directory. For example:

- 1 Create a configuration file, as shown in 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, once 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 (up) for use once the system has booted (to at least runlevel 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).

Checking and Testing the hxge Device

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

Options for verifying the network interface operation include:

- `ifconfig`: Use the `ifconfig` command to see if the RX/TX (receive/transmit) 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 network (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 name (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 (for example, 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 indicating 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:255.255.0
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 errors, dropped, overruns, frame, carrier, or collision events are reported. Some network errors are expected even in normal operation, but should be insignificant relative to the packet counts.

- `ethtool`: If `ifconfig` reports an accumulation of errors, then an extremely detailed breakdown of NEM traffic details (including error counts of all sorts) may be obtained using the `ethtool(8)` command.

The following is an excerpt of the total hxge detailed statistics output.

```

host #> ethtool -S eth2
NIC statistics:
Rx Channel #: 0
Rx Packets: 3008
Rx Bytes: 3289580
Rx Errors: 0
Jumbo Packets: 0
ECC Errors: 0
RBR Completion Timeout: 0
PEU Response Error: 0
RCR Shadow Parity: 0
RCR Prefetch Parity: 0
RCR Shadow Full: 0
RCR Full: 0
RBR Empty: 0
RBR Full: 0
RCR Timeouts: 3008

```

```

RCR Thresholds: 0
Packet Too Long Errors: 0
No RBR available: 0
RVM Errors: 0
Frame Errors: 0
RAM Errors: 0
CRC Errors: 0
[...]

```

Changing the hxge Driver Configuration

The hxge network interface configuration generally concerns parameters that are external to the driver/interface parameters such as the IP network address. There is also a set of configuration parameters that are internal to the hxge driver. These parameters are not normally changed, and improper setting of these parameters can easily result in a dysfunctional network interface.



Caution – Improper setting of the hxge driver configuration parameters can adversely affect performance. Only experienced administrators should attempt to change an hxge driver configuration parameter.

To see a list of the available hxge driver configuration parameters, use the `modinfo(8)` command.

```

host #> modinfo hxge
filename:      /lib/modules/2.6.18-164.el5/kernel/drivers/net/hxge.ko
version:      1.2.7
license:      GPL
description:   Sun Microsystems(R) 10 Gigabit Network Driver
author:       Sun Microsystems, john.doe@oracle.com
srcversion:   270F053A5DE6A454D1D224D
alias:        pci:v0000108Ed0000AAAAsv*sd*bc*sc*i*
depends:
vermagic:     2.6.18-164.el5 SMP mod_unload gcc-4.1
parm:         enable_jumbo:enable jumbo packets (int)
parm:         intr_type:Interrupt type (INTx=0, MSI=1, MSIx=2, Polling=3) (int)
[...]
parm:         tcam_ipsec_ipv6:IPsec over IPv6 class (int)
parm:         tcam_stcp_ipv6:STCP over IPv6 class (int)
parm:         debug:Debug level (0=none,...,16=all) (int)

```

Each `parm:` line identifies an hxge driver configuration parameter that the system administrator can override when loading the hxge driver.

There are two ways to configure the driver parameters:

- [“How to Temporarily Configure the hxge Driver Parameters” on page 85](#)
- [“How to Permanently Configure the hxge Driver Parameters” on page 85](#)

The driver configuration parameters are described here: [“hxge Driver Configuration Parameters” on page 86](#).

▼ How to Temporarily Configure the hxge Driver Parameters

To temporarily change the hxge driver configuration, use the `modprobe(8)` command to specify a parameter value when loading the driver. An hxge driver parameter can only be specified (that is, changed to a non-standard value) on initially loading the driver. If the hxge driver is already loaded, you must first unload it using `modprobe -r hxge` before you load it with a different parameter specification.

1 Check to see if the hxge driver is already loaded.

```
host #> lsmod | grep hxge
hxge                175440
0
```

In this example, the driver is loaded.

2 Unload the currently active driver.

```
host #> modprobe -r hxge
```

You can also use the `rmmod(8)` command.

3 Manually load the hxge driver, specifying desired hxge parameters and values. For example, to enable detailed driver activity logs (and fill up the root partition).

```
host #> modprobe hxge debug=0x2001
```

▼ How to Permanently Configure the hxge Driver Parameters

For the hxge driver to be automatically configured every time the driver is loaded:

- **Add the hxge driver configuration to the `modprobe.conf(5)` file at `/etc/modprobe.conf`, using the `options` command.**

For example, to automatically (always) disable DMA channel spreading whenever the hxge driver is loaded, add the following line to the `/etc/modprobe.conf` file:

```
options hxge tcam=0
```

Here is a sample `modprobe.conf` file that disables the receive DMA channel spreading:

```
host #> cat /etc/modprobe.conf
alias eth0 igb
alias eth1 igb
alias scsi_hostadapter ahci
options hxge tcam=0
```

hxge Driver Configuration Parameters

The actual list of `hxge` driver parameters is subject to change from release to release. The following table lists the driver configuration parameters for the version 0.0.9 `hxge` driver. The table also lists the accepted values and defaults for the parameter, where applicable.

TABLE 7 Driver Configuration Parameters

Parameter	Description	Values	Default
<code>enable_jumbo</code>	Controls <code>hxge</code> driver runtime support of jumbo frames. <code>hxge</code> jumbo frame support is automatically enabled as needed (depending on the network interface-specified MTU value).	0 = No 1 = Yes	Automatic
<code>intr_type</code>	Controls what sort of interrupt mechanism (if any) is selected by the <code>hxge</code> driver. The <code>hxge</code> driver will automatically select the best (highest potential performance) interrupt support mechanism when it is initially loaded and started.	0 = INTx 1 = MSI 2 = MSIx 3 = Polling	Determined by system hardware support (MSIx is "best").
<code>rbr_entries</code>	Specifies how many 4KB receive buffers the <code>hxge</code> driver will allocate per receive channel (the NEM supports four parallel independent receive channels).		4096
<code>rcr_entries</code>	Specifies how many receive pointers (effectively, packets; jumbo packets can require up to 3 RCR entries per single jumbo packet) the <code>hxge</code> driver will allocate per receive channel.		8192
<code>rcr_timeout</code>	Magic internal unitless number. Do not change this number unless told to by a qualified Sun agent.		
<code>rcr_threshold</code>	Magic internal unitless number. Do not change this number unless told to by a qualified Sun agent.		
<code>rx_dma_channels</code>	Specifies how many receive DMA channels the <code>hxge</code> driver should attempt to activate when the driver is initialized and brought online. Each DMA channel represents an independent receive processing stream (interrupt and CPU with separate dedicated buffer pool, system resources permitting) capability.	1 = minimum 4 = maximum	4

TABLE 7 Driver Configuration Parameters (Continued)

Parameter	Description	Values	Default
tx_dma_channels	Specifies how many transmit DMA channels the hxge driver should attempt to activate when the driver is initialized and brought online.	1 = minimum 4 = maximum	4
num_tx_descs	Specifies how many transmit descriptors the hxge driver should allocate per transmit channel. Each transmit packet requires a transmit descriptor.		1024
tx_buffer_size	Specifies the small transmit buffer size. For transmit packets smaller than this value, the hxge driver will coalesce all the packet fragments together into a single pre-allocated tx_buffer_size hxge buffer; for transmit packets larger than this size, the hxge driver will construct a scatter/gather pointer list for the hardware to decipher.		256
tx_mark_ints	Magic internal unitless number. Do not change this number unless told to by a qualified Sun agent.		
max_rx_pkts	Specifies the maximum number of receive packets (queued by the NEM network engine) will be processed on any one receive interrupt, before the hxge driver (interrupt service routine) will dismiss the interrupt, releasing the interrupted CPU to perform other actions.		64
vlan_id	Specifies the implicit VLAN ID that the hxge driver will assign to non-VLAN-tagged packets.		4094
debug	Controls hxge printout verbosity of hxge driver progress, actions, and events. Normally, only significant or serious (error) information is printed out. Note: Read “Troubleshooting the Driver” on page 88 before changing this parameter.	0x2002 = normal operation (don't print DBG messages) 0x2001 = debug operation (print debug messages)	2002
strip_crc	Controls whether the hxge driver or the NEM network engine strips off the packet CRC.	0 = disable 1 = enable	0
enable_vmac_ints	Controls whether or not the hxge driver enables VMAC interrupts.	0 = disable 1 = enable	0
promiscuous	Controls whether or not the hxge driver enables the NEM engine to run in promiscuous mode.	0 = disable 1 = enable	0

TABLE 7 Driver Configuration Parameters (Continued)

Parameter	Description	Values	Default
chksum	Controls whether or not the hxge driver enables the NEM engine hardware checksumming capability.	0 = no hardware checksumming, 1 = hardware receive packet checksumming 2 = hardware transmit checksumming 3 = both	3
tcam	Controls whether or not the hxge driver enables the Virtualized M2 NEM ASIC hardware engine, spreading the receive traffic across multiple (up to 4) parallel independent receive streams (interrupts, CPUs). This might also be referred to as DMA channel spreading. See also rx_dma_channels in this table.	0 = disable 1 = enable	1
tcam_seed	Magic internal unitless number. Do not change this number unless told to by a qualified Oracle agent.		
tcam_tcp_ipv4	Controls whether or not the hxge driver enables DMA channel spreading for IPv4 UDP traffic.	0 = disable 1 = enable	1

Troubleshooting the Driver

This section describes the hxge driver debug messaging parameter that can be used to troubleshoot issues with the driver. The following topics are covered in this section:

- [“Debug Messaging Parameter Overview” on page 88](#)
- [“How to Set the Debug Drive Parameter Dynamically” on page 89](#)
- [“How to Configure the Syslog Parameter” on page 90](#)

Debug Messaging Parameter Overview

The Linux hxge driver has a built-in message and event logging facility, controlled by a message level parameter and logged through the system’s syslog(2) facility, usually to the file /var/log/messages).

The debug messaging parameter has two possible modes:

- **0x2002 to disable debug messaging, but still list error messages:** The hxge driver is configured by default (when it is initially loaded into the kernel memory) to list top-level startup messages and error events and messages.

For example, every time the hxge driver is loaded and started on an hxge network device, it will print a copyright statement, such as:

```
kernel: Sun Microsystems(R)10 Gigabit Network Driver-version 1.2.7
```

```
kernel: Copyright (c) 2009 Sun Microsystems.
```

and list the Ethernet MAC address(es) configured for that hxge device:

```
kernel: hxge: ...Initializing static MAC address 00:14:4f:62:00:1d
```

- **0x2001 to enable debug messaging:** Debug messaging provides large amounts of internal packet flow and event tracing, including specific details on each Ethernet packet sent or received by the hxge driver.

As a 10GbE network is easily capable of flowing over a million packets a second (at a 10GbE rate, a one KB packet is about a microsecond of wire time), this represents a potentially overwhelming load on the kernel's syslog facility's ability to buffer and write system messages to disk.



Caution – Do not enable the hxge driver's debug message logging facility unless necessary to diagnose a problem with the driver. The debug message logging has the potential to adversely affect system operation.

The default value for the message level parameter is 0x2002.

▼ How to Set the Debug Drive Parameter Dynamically

The Linux hxge driver's message logging can be statically specified via the debug driver configuration parameter. See [Table 7](#). This will set the messaging level when the driver is initially loaded into memory and initializes itself. This message level remains in effect until the driver is unloaded or dynamically overridden. The debug configuration parameter can only be specified when the driver is first loaded into kernel memory.

In addition to the static debug driver configuration parameter, the currently-running Linux hxge driver's message logging can be dynamically controlled via the `ethtool(8)` utility, using the `-s` switch.

- 1 To set the debug driver configuration parameter dynamically, use the following command:

```
ethtool -s ethn msglvl parametervalue
```

For example, still using eth2 as in previous examples, to dynamically turn on debug messaging in the currently-running hxge driver, use the command:

```
ethtool eth2 -s msglvl 0x2001
```

and to dynamically turn debug messaging back off again, leaving only error messages to be logged, use the command:

```
ethtool eth2 -s msglvl 2002
```

- 2 **Configure the syslog parameter.** See [“How to Configure the Syslog Parameter” on page 90](#).

▼ How to Configure the Syslog Parameter

By default, most Linux systems are configured to ignore (discard without logging) debug-level syslog messages. To see the Linux hxge driver's debug messages when they are enabled, the syslog(2) facility must also be configured 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 **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 notifies the currently-running syslog daemon to re-read its configuration file (see the syslogd(8) man page for more information).

Configuring Jumbo Frames

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

To enable `hxge` network interface support of Ethernet jumbo frames, use the `ifconfig(8)` command to set the network interface maximum transmission unit (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 – The commands shown in the following examples can be used for both RHEL and SLES.

This section covers the following topics:

- [“How to Temporarily Configure Jumbo Frames Support” on page 91](#)
- [“How to Permanently Enable Jumbo Frame Support” on page 92](#)

▼ How to Temporarily Configure Jumbo Frames Support

To temporarily enable (or change) jumbo frame support for an `hxge` network interface, use an `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` now reports the MTU size to be 9000 bytes. NFS 8KB pages will now flow (send or receive) as a single Ethernet packet.

▼ How to Permanently Enable Jumbo Frame Support

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

1 Set the MTU parameter in the corresponding `ifcfg` file (`ifcfg-eth2` for the examples in this document). For example:

```
host #> cat /etc/sysconfig/network-scripts/ifcfg-eth2
# Sun NEM/Hydra 10GbE
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 `hxge` 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
```

Installing and Configuring Drivers on a Windows Platform

The following topics are covered in this section:

- “Installing Drivers on a Windows Platform” on page 93
- “Enabling Jumbo Frames” on page 113

Installing Drivers on a Windows Platform

This section describes the process for installing network, VLAN, and enclosure drivers on an x86 (Intel or AMD) server module that supports Windows Server 2008 (32/64-bit).

The following sections describe the procedures for installing Windows drivers:

- “Installing and Uninstalling the Sun Blade 6000 10GbE Network Controller” on page 93
- “Installing and Uninstalling the Sun Blade 6000 10GbE VLAN Driver (Optional)” on page 99
- “Adding and Removing VLANs” on page 105
- “How to Install the Enclosure Device on a Windows Server 2008 System” on page 109

Installing and Uninstalling the Sun Blade 6000 10GbE Network Controller

This section covers the following procedure:

- “How to Install the Sun Blade 6000 10GbE Network Controller” on page 93
- “How to Uninstall the Sun Blade 6000 10GbE Networking Controller” on page 97

▼ How to Install the Sun Blade 6000 10GbE 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_10Gbe_Networking_Controller.msi

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

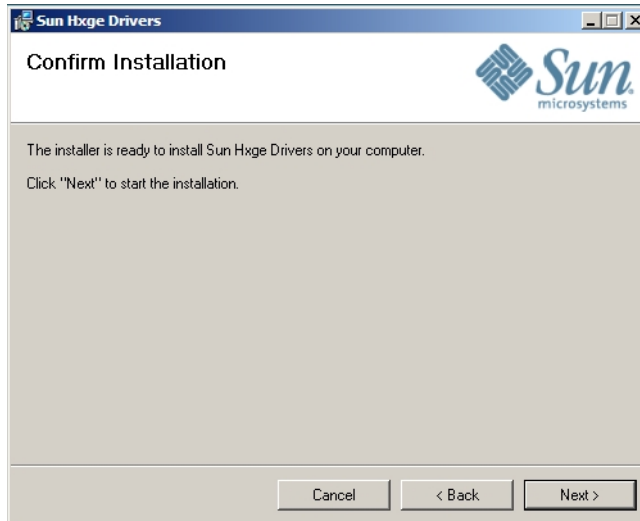
- 2 **Navigate to the `Sun_Blade_6000_10GbE_Networking_Controller.msi` file on your local or remote system and double-click it to start the installation.**

The License Agreement page displays.



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

The Confirm Installation page displays.



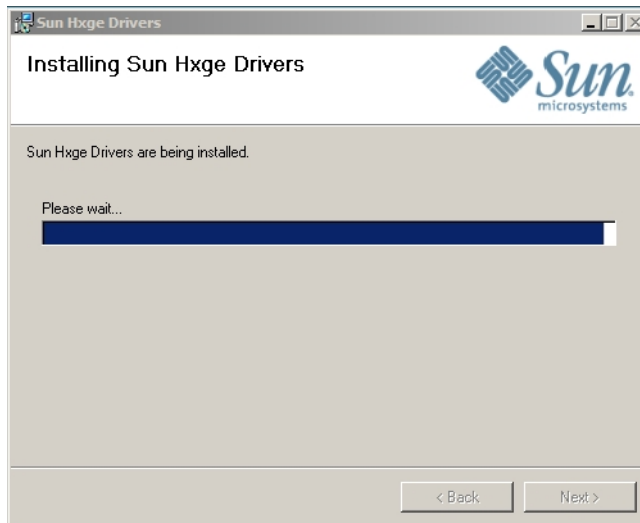
4 Click Next.

The Select Installation Folder page displays.



5 Select “Everyone” or “Just Me” and click Next.

The Installing Sun Hxge Drivers page displays.



When the installation finishes, the Installation Complete page displays.



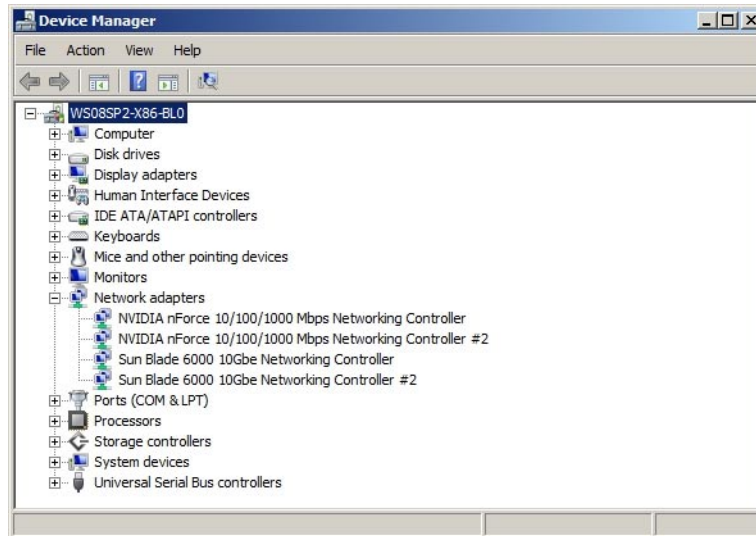
6 Click Close.

The Sun Blade 6000 10GbE Networking Controller device is installed and visible in the Network Adapters section of the Windows Device Manager.

7 Verify that the Sun Blade 6000 10GbE 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 displays.



- c. Click **Network adapters** and verify that the **Sun Blade 6000 10GbE Networking Controller** is in the **Network adapters** list.

One controller will be displayed if you have one Virtualized M2 NEM installed, and two will be displayed for two Virtualized M2 NEMs.

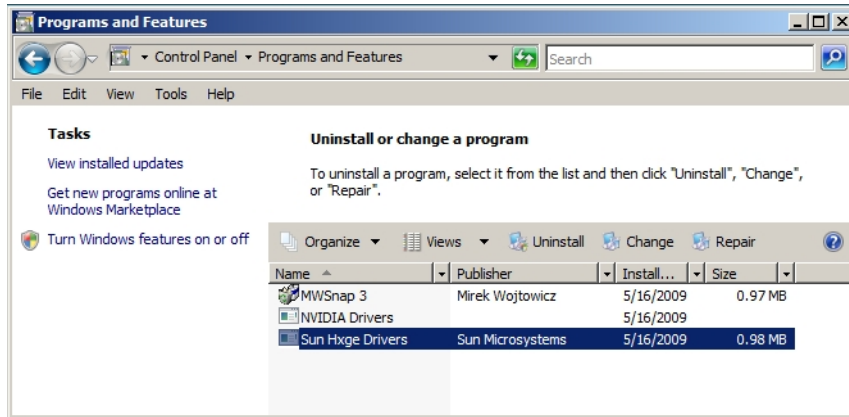
- d. Exit the **Windows Device Manager**.

▼ How to Uninstall the Sun Blade 6000 10GbE Networking Controller

- 1 Before uninstalling the networking controller, remove any VLANs that were installed on the controller. See [“How to Remove a VLAN” on page 107](#).
- 2 Open the **Control Panel**.

3 Double-click Programs and Features.


The Programs and Features dialog is displayed.



4 Select Sun Hxge 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, the following error message is displayed.



6 If you get the error message:

- a. Click OK.
- b. Remove all VLANs installed on the Ethernet port (see [“How to Remove a VLAN”](#) on page 107).
- c. Return to Step 2.

Installing and Uninstalling the Sun Blade 6000 10GbE VLAN Driver (Optional)

The Sun Blade 6000 10GbE VLAN driver enables creation of Virtual Local Area Networks (VLANs) on top of physical Ethernet ports. You only need to install this driver if you intend to create one or more VLANs. The VLAN driver is bundled with the Sun Blade 6000 10GbE Networking Controller.

This section covers the following topics:

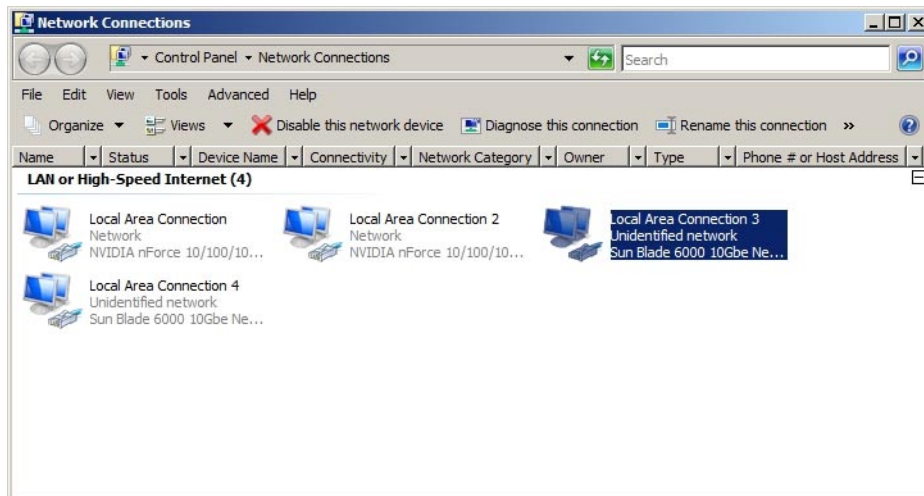
- [“How to Install the VLAN Driver” on page 99](#)
- [“How to Uninstall the VLAN Driver” on page 103](#)

▼ How to Install the VLAN Driver

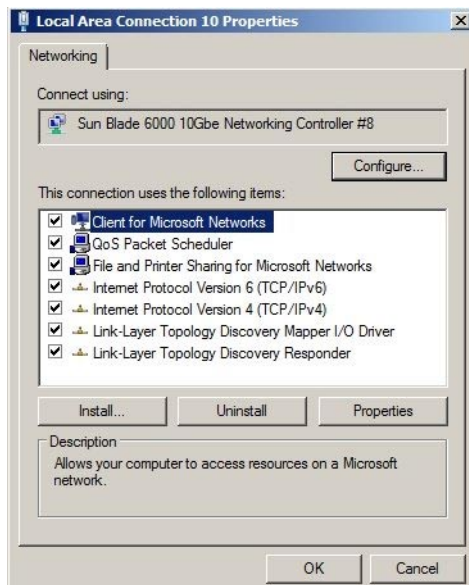
- 1 Install the Sun Blade 6000 10GbE Networking Controller. See [“How to Install the Sun Blade 6000 10GbE Network Controller” on page 93](#).
- 2 Open the Managed Network Connections dialog.
 - a. Open the Control Panel.
 - b. Double-click the Network Sharing Center to open it.
 - c. Select Managed Network Connections on the left side of the panel under “Tasks”.

3 Right-click one of the Sun Blade 6000 10 GbE Networking Controllers and select Properties.

In the following example, Local Area Connection 3 is selected.

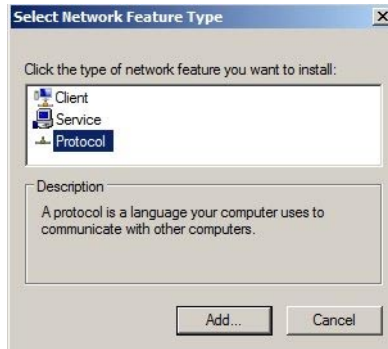


4 The Local Area Connection Properties dialog displays.

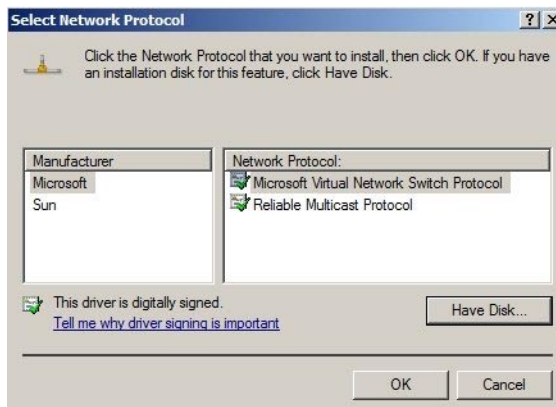


5 Click Install.

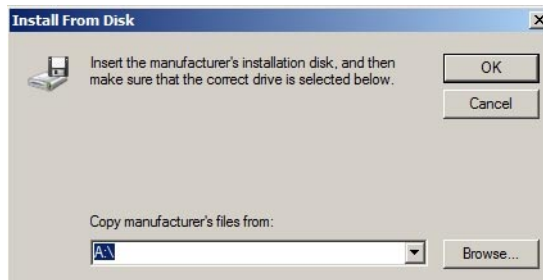
The Network Feature Type dialog displays.

**6 Select Protocol and click Add.**

The Select Network Protocol dialog.

**7 Click Have Disk.**

The Install From Disk dialog displays.

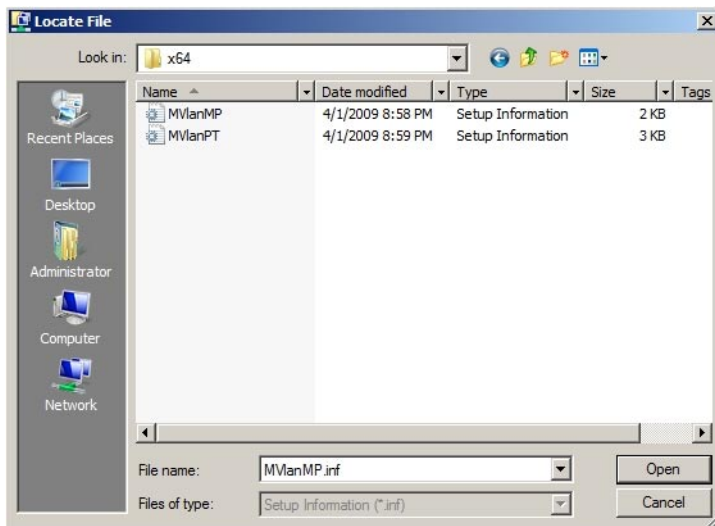


- 8 Click Browse and choose one of the following paths, depending on the version of Windows Server that you have installed:**

C:\Program Files(x86)\Sun Microsystems\Sun Hxge Drivers\NDIS6\vlan\w2k8\{x86,x64}

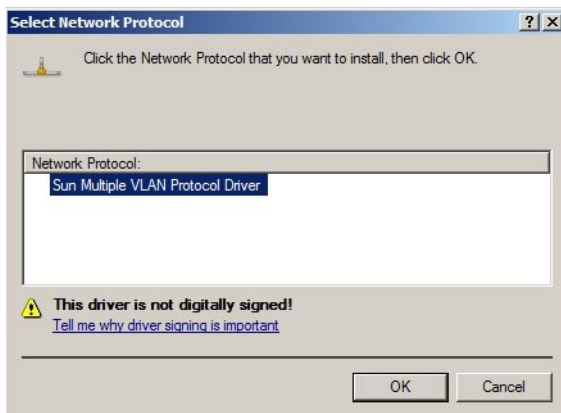
Use the x86 path if you have a 32-bit architecture and the x64 path if you have a 64-bit architecture.

The Locate File dialog displays.



- 9 Select either of the information setup files displayed and click Open.**

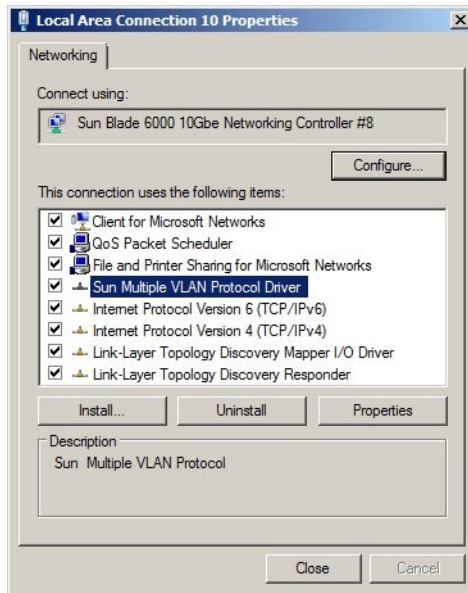
The Select Network Protocol dialog displays.



- 10 Click OK.**

The Sun Blade 6000 10GbE 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 10GbE Network Controller installed on the system.

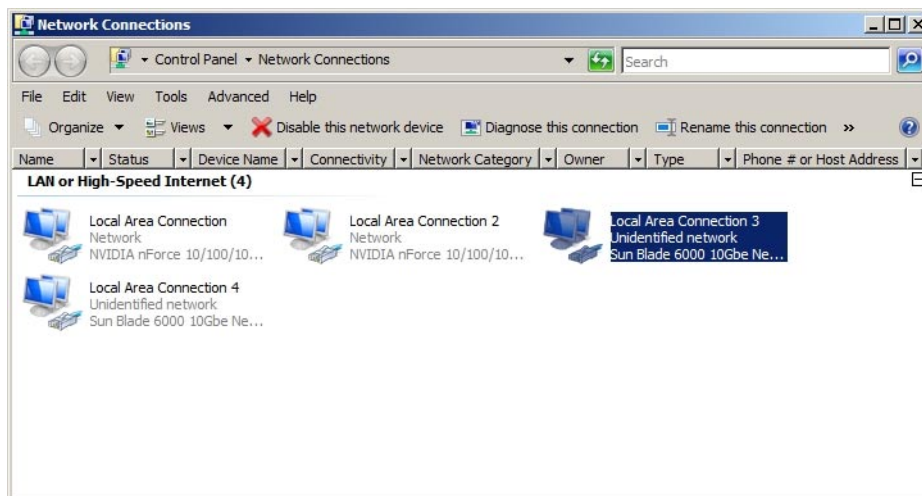


▼ How to Uninstall the VLAN Driver

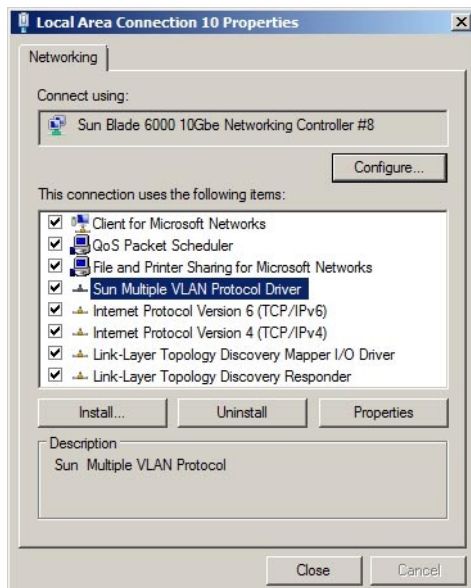
- 1 Open the Managed Network Connections dialog.
 - a. Open the Control Panel.
 - b. Double-click the Network Sharing Center to open it.
 - c. Select Managed Network Connections on the left side of the panel under "Tasks".

2 Right-click one of the Sun Blade 6000 10GbE Networking Controllers and select Properties.

In the following example, Local Area Connection 3 is selected.



3 The Local Area Connection Properties dialog displays.



4 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.

Adding and Removing VLANs

This section covers the following topics:

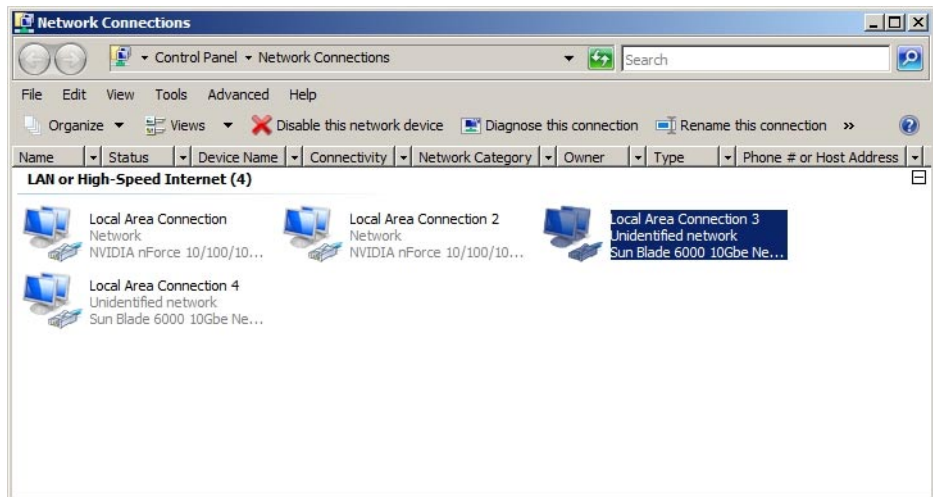
- “How to Add a VLAN” on page 105
- “How to Remove a VLAN” on page 107

▼ How to 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 10GbE Networking Controllers and select Properties.

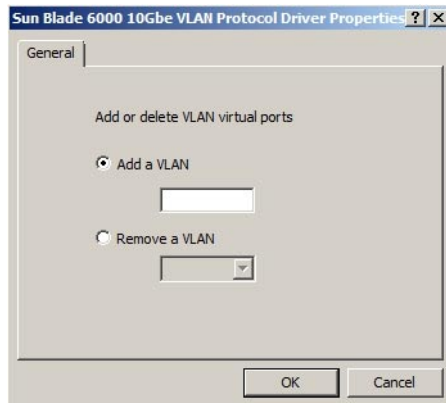
Note – Note which controller port you use to create the VLAN. You will need to use the same controller port to remove the VLAN.

In the following example, the VLAN will be added to the Local Area Connection 3 controller port.



3 Select Sun Multiple VLAN Protocol and click Properties.

The Sun Multiple VLAN Protocol dialog displays.

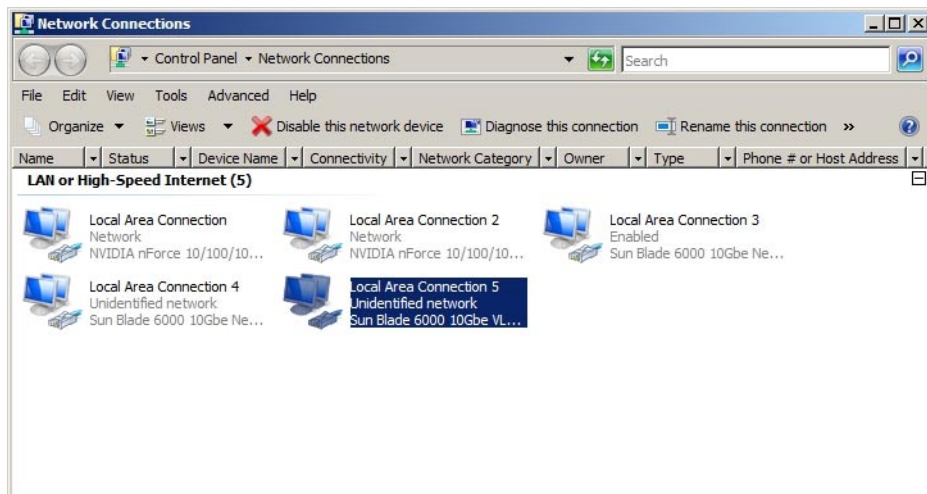


4 Select Add a VLAN and enter the VLAN number.

5 Click OK.

A new Local Area Connection icon called Sun Blade 6000 10GbE VLAN Virtual Miniport 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.

In the example below, the virtualized LAN is shown as Local Area Connection 5.

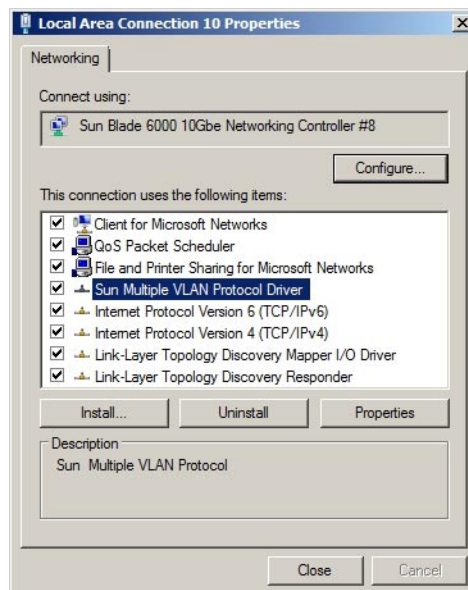


▼ How to Remove a VLAN

- 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 10GbE Networking Controller used to create the VLAN that you are removing.
 For example, in the “How to Add a VLAN” on page 105 procedure, the network controller port used to create the VLAN is Local Area Connection 3 controller port.

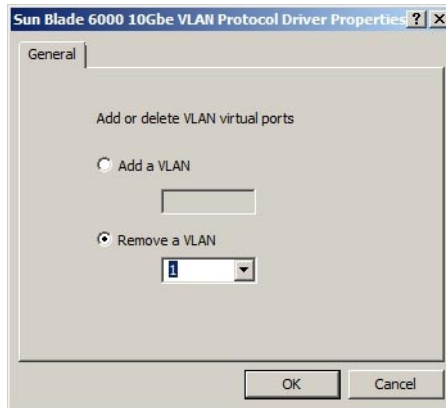
3 Select Properties.

The Local Area Connection Properties dialog displays.



4 Select Sun Multiple VLAN Protocol and click Properties.

The Sun Multiple VLAN Protocol dialog displays.



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

6 Click OK.

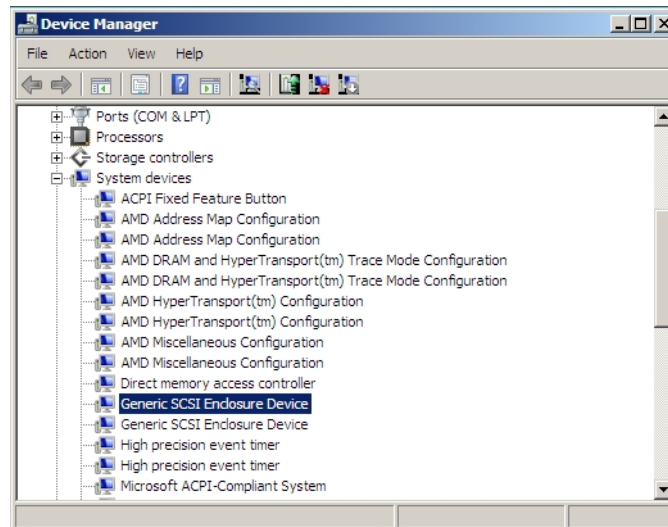
The Sun Blade 6000 VLAN Virtual Miniport icon for the VLAN selected is no longer visible in the Managed Network Connections and Adapter section of the Device Manager.

▼ How to 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 displays.

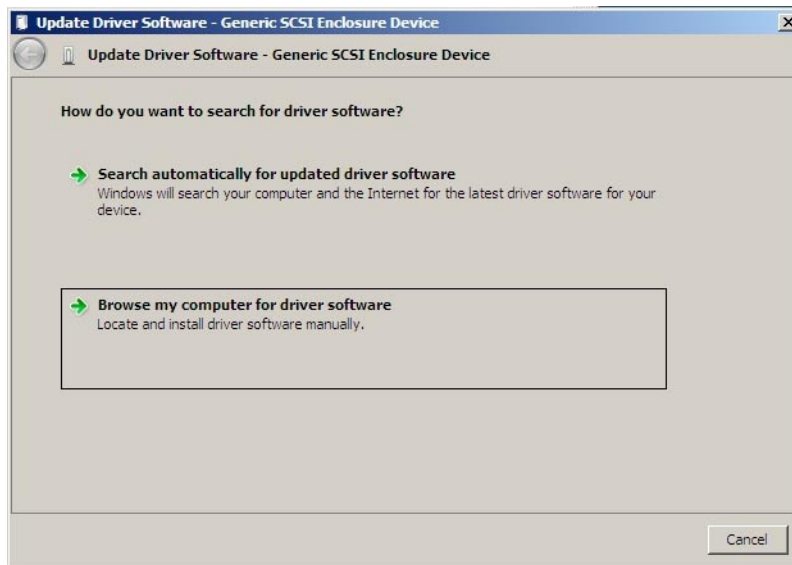


2 Click on System Devices and find the Generic SCSI Enclosure Device.

One enclosure device is displayed if you have one Virtualized M2 NEM installed and two enclosure devices are displayed if you have two Virtualized M2 NEMs installed.

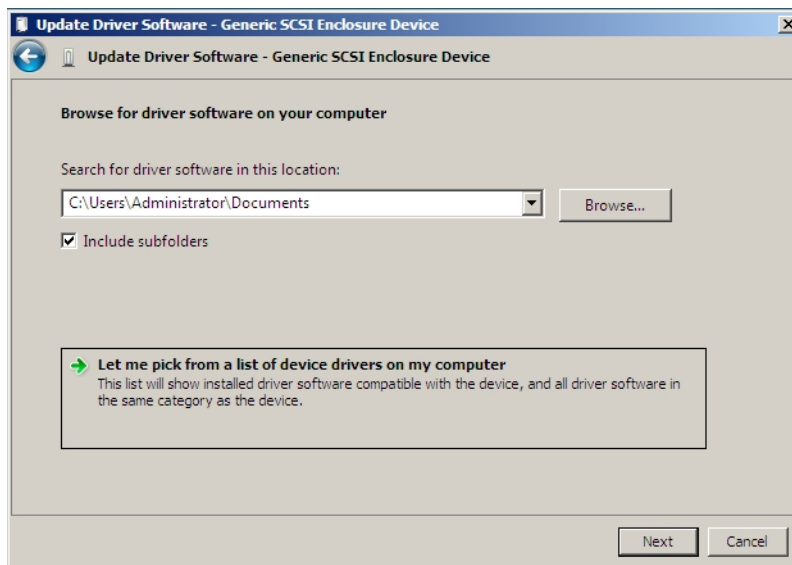
3 Right-click on the Generic SCSI Enclosure Device, and select Update driver.

The How do you want to search for drive software? page displays.



4 Click "Browse my computer for driver software".

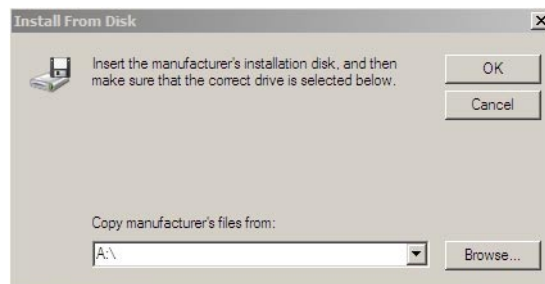
The Browse for driver software on your computer page displays.



- 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 displays.



- 6 Click Have Disk.
The Install From Disk page displays.

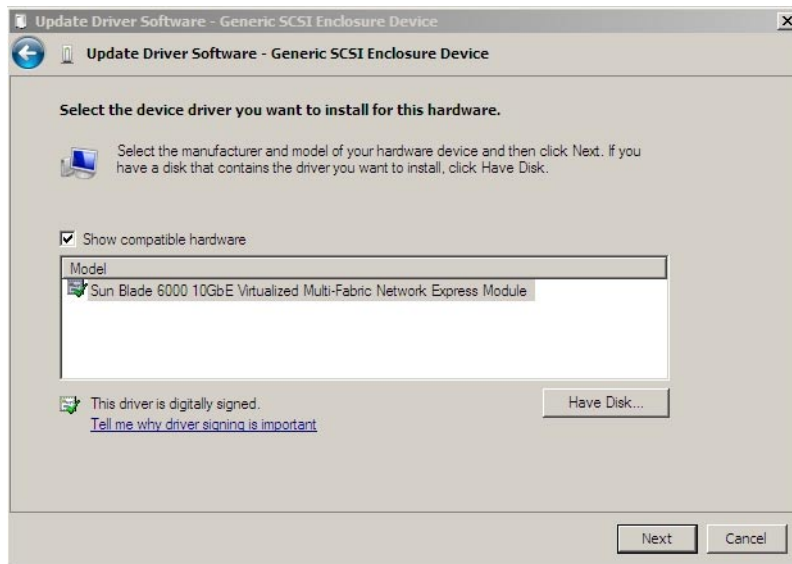


- 7 Click Browse and navigate to the directory that contains the Sun Blade 6000 10GbE Virtualized Multi-Fabric Network Express Module information (lsinodrv.inf) file.

The path on the Tools and Drivers CD is:

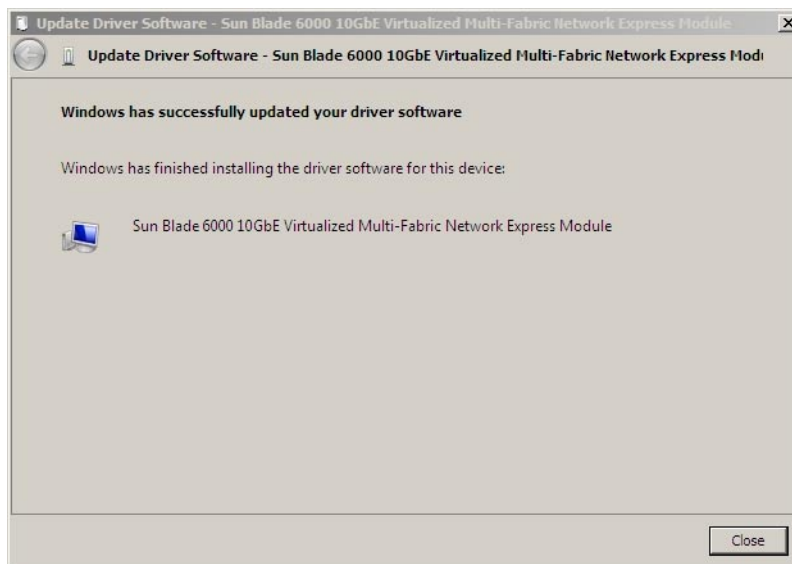
DVDdrive:\windows\w2k8\64-bit\lsinodrv.inf

The selected Sun Blade 6000 10GbE Virtualized Multi-Fabric- Network Express Module driver displays.



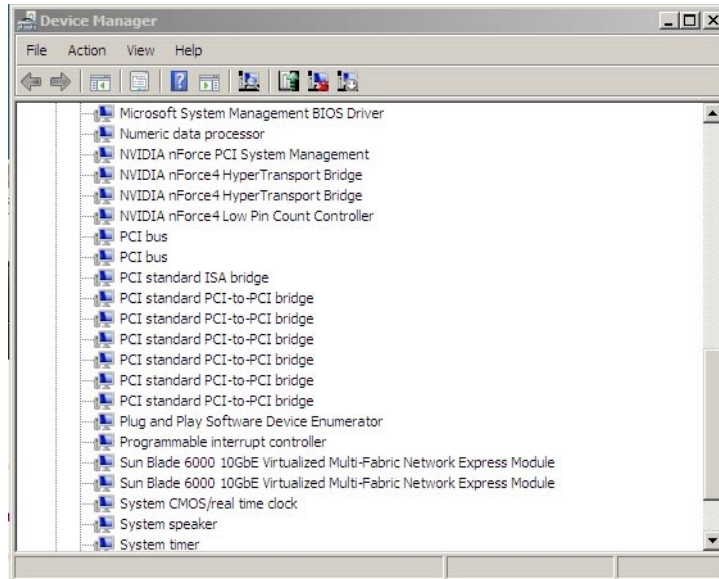
- 8 Click Next to install the Sun Blade 6000 10GbE Virtualized Multi-Fabric Network Express Module driver.

The Software Installing page displays as the device is installed, then the Windows has successfully updated your driver software page displays.



- 9 Click Close to complete the installation.
- 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 displays.



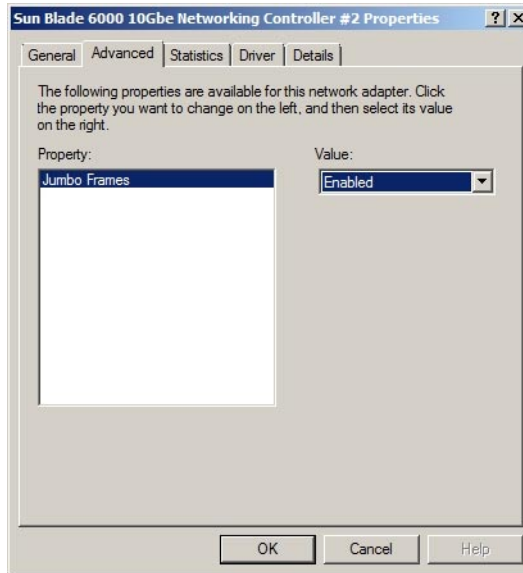
- c. Click on System Devices and verify that the Sun Blade 6000 10GbE Virtualized Multi-Fabric Network Express Module is in the System Devices list.
- d. Repeat this procedure to install a second enclosure device.

Enabling Jumbo Frames

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 of up to 1518 bytes.

▼ How to Enable Jumbo Frames

- 1 In the Device Manager, right-click on the Sun Blade 6000 10Gbe Networking Controller and select Properties.
- 2 Select the Advanced Tab, then select Jumbo Frames as shown below.



- 3 Select the enabled option and click OK to enable the Jumbo Frames feature.

Installing and Configuring Drivers on a VMware ESX Server Platform

This section describes the process for installing and configuring hxge drivers on a Sun Blade 6000 server module with VMware ESX Server installed.

The following topics are covered in this section:

- “Installing the ESX Server Drivers on an Existing ESX Server” on page 115
- “Installing the ESX Server Drivers With a New ESX Installation” on page 117
- “Configuring the Virtual NEM M2 Network Adapters” on page 118
- “Configuring Jumbo Frames” on page 119

Installing the ESX Server Drivers on an Existing ESX Server

The following topics are covered in this section:

- “How to Install a Virtualized M2 NEM Driver on an Existing ESX/ESXi 4.0 Server Using vihostupdate” on page 115
- “How to Install a Virtualized M2 NEM Driver on an Existing ESX 4.0 Server Using esxupdate” on page 116

▼ How to Install a Virtualized M2 NEM Driver on an Existing ESX/ESXi 4.0 Server Using vihostupdate

- 1 Obtain the hxge driver from one of the following:
 - The Virtualized M2 NEM Tools and Drivers CD
`/vmware/drivers/esx4.0/Vmware-esx-drivers-net-hxge-1.2.2.11- 00000.i386.rpm`
 - Download the hxge driver from the following site: http://www.vmware.com/support/vsphere4/doc/drivercd/esx-hxge_400.4.1.2.2.7.html and create a driver CD.
- 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 **hxge-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  
hxge-vmware-driver-4-1-2-2-7-offline_bundle-193789.zip
```

For more details on vihostupdate, see http://www.vmware.com/pdf/vsphere4/r40/vsp_40_vcli.pdf

▼ How to Install a Virtualized M2 NEM Driver on an Existing ESX 4.0 Server Using esxupdate

- 1 Obtain the hxge driver from one of the following:
 - The Virtualized M2 NEM Tools and Drivers CD
/vmware/drivers/esx4.0/Vmware-esx-drivers-net-hxge-1.2.2.11- 00000.i386.rpm
 - Download the hxge driver from the following site: http://www.vmware.com/support/vsphere4/doc/drivercd/esx-hxge_400.4.1.2.2.7.html and create a driver CD.
- 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 **hxge-vmware-driver-4-1-2-2-7-offline_bundle-193789.zip** file.
- 6 Execute command **esxupdate** to update the driver:

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

For more details on esxupdate, see http://www.vmware.com/pdf/vsphere4/r40/vsp_40_esxupdate.pdf

Installing the ESX Server Drivers With a New ESX Installation

This section covers instructions for installing the drivers when ESX is not already installed.

▼ How to Install a Virtualized M2 NEM Driver With a New ESX4.0 Installation

- 1 Obtain the hxge driver from one of the following:
 - The Virtualized M2 NEM Tools and Drivers CD:
`/vmware/drivers/esx4.0/Vmware-esx-drivers-net-hxge-1.2.2.11-00000.i386.rpm`
 - Download the hxge driver from the following site: http://www.vmware.com/support/vsphere4/doc/drivercd/esx-hxge_400.4.1.2.2.7.html and create a driver CD.
- 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 displays 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.

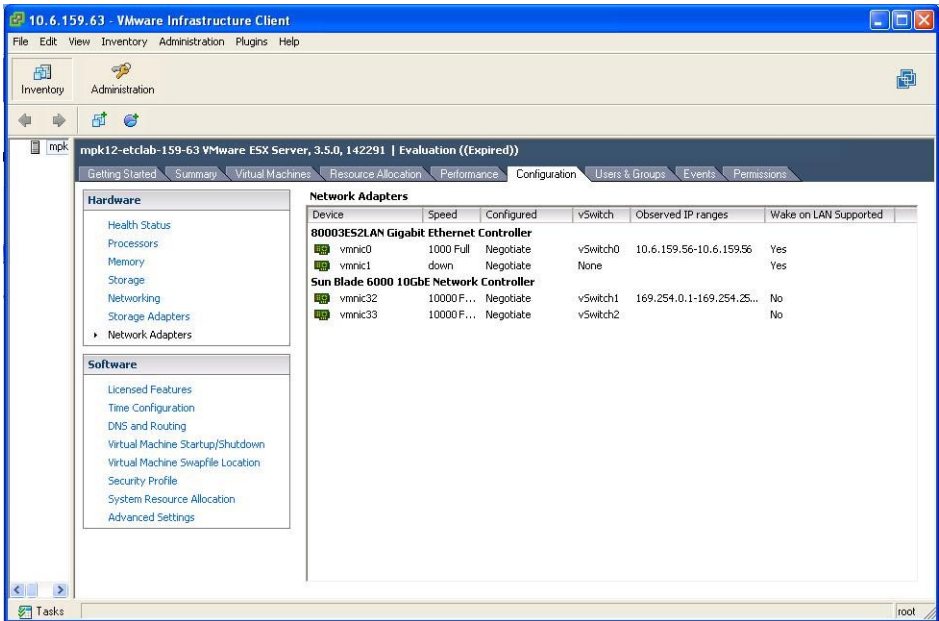
Configuring the Virtual NEM M2 Network Adapters

The following topics are covered in this section:

- “How to Configure the Virtual NEM M2 Network Adapters” on page 118

▼ How to Configure the Virtual NEM M2 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 10GbE Network Controller driver that you want to configure.

Configure Networking using the Sun Blade 6000 10GbE 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

Configuring Jumbo Frames

Note the following guidelines for configuring Jumbo Frames for VMware ESX server.

- 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.
- It is possible to configure Jumbo Frame for iSCSI Network, but it is not supported at this time.
- Jumbo Frames must be enabled for each vSwitch or VMkernel interface through the command-line interface on your ESX Server 3 host.
- To allow 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.

The following topics are covered in this section:

- “How to Create a Jumbo Frames-enabled vSwitch” on page 119
- “How to Create a Jumbo Frames-enabled VMkernel Interface” on page 120

▼ How to Create a Jumbo Frames-enabled vSwitch

- 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. To set the MTU size for the vSwitch, run the following command:

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

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

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

- 3 To display a list of vSwitches on the host, run the following command:
esxcfg-vswitch -l
- 4 Check that the configuration of the vSwitch is correct.

▼ How to Create a Jumbo Frames-enabled VMkernel Interface

- 1 Log in directly to your ESX Server host console.
- 2 To create a VMkernel connection with Jumbo Frame support, run the following command:
`esxcfg-vmknic -a -i ip address -n netmask -m MTU portgroup name`
where *ip address* is the ip address of the server, *netmask* is the netmask address, *MTU* is the MTU size and *portgroup name* is the name that has been assigned to the port group.
- 3 To display a list of VMkernel interfaces, run the following command:
`esxcfg-vmknic -l`
- 4 Check that the configuration of the Jumbo Frame-enabled interface is correct.

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

ILOM Supplement

This supplement contains information for using Oracle Integrated Lights Out Manager (ILOM) with the Virtualized M2 NEM.

The ILOM tool enables you to actively manage and monitor the NEM independently of the operating system state, providing you with a reliable Lights Out Management (LOM) system. With ILOM, you can proactively learn about hardware errors and faults as they occur. You can:

- Remotely control the power state of the NEM
- View the graphical and non-graphical consoles for the host
- View the current status of sensors and indicators on the system
- Determine the hardware configuration of your system
- Receive generated alerts about system events in advance via IPMI PETs, SNMP traps, or email alerts.
- Configure zoning of SAS devices using the Sun Blade Zone Manager
- Update NEM firmware from the CMM ILOM.

The following topics are covered in this chapter:

- “Accessing ILOM Documentation and Updates” on page 121
- “Connecting to ILOM” on page 122
- “Updating the NEM Firmware” on page 127
- “Sun Blade Zone Manager” on page 132
- “NEM Sensors” on page 132
- “Enabling Private and Failover Mode” on page 137
- “Using Hot Plug Commands” on page 138

Accessing ILOM Documentation and Updates

ILOM documentation is divided into two categories:

- Generalized ILOM information, located in the Sun Integrated Lights Out Manager 3.0 Documentation Collection at: <http://www.oracle.com/pls/topic/lookup?ctx=ilom30&id=homepage>
- Information specific to the Virtualized M2 NEM, located in this supplement

For product updates that you can download for the Virtualized M2 NEM, see the following Web site:

<http://oracle.com/goto/10gbenemm2>

This site contains updates for firmware and drivers, as well as CD-ROM .iso images.

Connecting to ILOM

This section describes how to connect to the ILOM located in the Virtualized M2 NEM.

The following topics are covered in this section:

- “ILOM Connection Options” on page 122
- “How to Find the ILOM’s IP Address” on page 123
- “Connecting to ILOM Through Ethernet” on page 124
- “Connecting to ILOM Through a Serial Connector” on page 125

ILOM Connection Options

Normally, you connect to the ILOM through the Ethernet. When the NEM is installed in the chassis, its ILOM is automatically accessible on the same subnet as the CMM ILOM. You can also connect to the ILOM through the CMM or NEM serial connector.

Note – The CMM ILOM in the chassis has an Ethernet switch that supports connections to the Virtualized M2 NEM and its ILOMs. To use this connection, you must be connected to the same subnet as the ILOM, and you must know the Virtualized M2 NEM ILOM’s Ethernet address.

You can connect to the ILOM on your NEM using one of the methods shown in the following table:

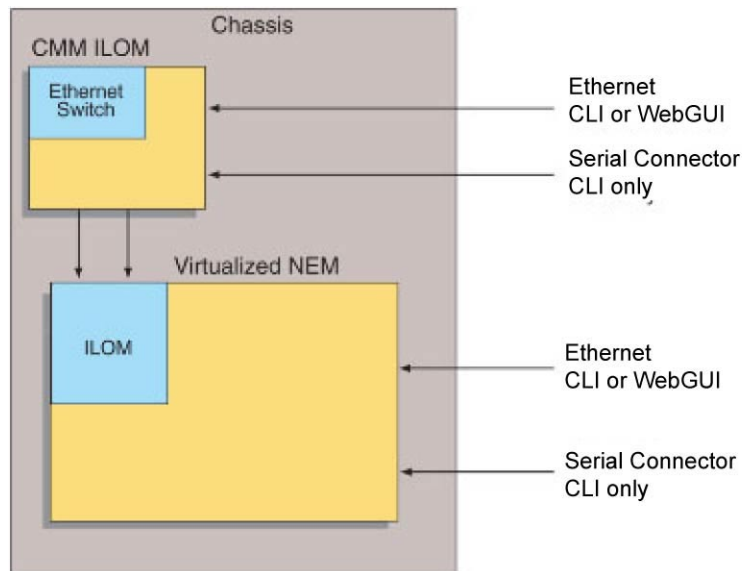
TABLE 8 ILOM Connection Options

Option Number	Connection Method	User Interface	More Information
1	Ethernet	Web interface or CLI	“How to Connect to the CLI Through Ethernet” on page 124 “How to Connect to the Web Interface Through Ethernet” on page 125
2	Serial - CMM	CLI only	“How to Connect to ILOM Through the Chassis (CMM) Serial Connector” on page 125

TABLE 8 ILOM Connection Options (Continued)

Option Number	Connection Method	User Interface	More Information
3	Serial - NEM	CLI only	“How to Connect to the ILOM Using the NEM Serial Connector” on page 127

The following illustration shows a graphical representation of the ILOM connection options.



▼ How to Find the ILOM's IP Address

1 Log in to the CMM.

See the *Oracle Integrated Lights Out Manager (ILOM) CMM Administration Guide*, 820-0052 for details.

2 Type the command:

```
show /CH/NEMn/SP/network
```

where *n* is the NEM number.

3 The CLI displays information about the NEM, including its IP address.

For example:

```
-> show /CH/NEM0/SP/network

/CH/NEM0/SP/network
Targets:

Properties:
  type = Network Configuration
  commitpending = (Cannot show property)
  ipaddress = IPaddress
  ipdiscovery = dhcp
  ipgateway = IPgateway
  ipnetmask = 255.255.255.0
  macaddress = Macaddress
  pendingipaddress = IPaddress
  pendingipdiscovery = dhcp
  pendingipgateway = IPgateway
  pendingipnetmask = 255.255.255.0

Commands:
  cd
  set
  show

->
```

Connecting to ILOM Through Ethernet

The following topics are covered in this section:

- [“How to Connect to the CLI Through Ethernet” on page 124](#)
- [“How to Connect to the Web Interface Through Ethernet” on page 125](#)

▼ How to Connect to the CLI Through Ethernet

Before You Begin If you do not know the ILOM's IP address, find it as described in [“How to Find the ILOM's IP Address” on page 123](#).

1 Start your SSH client.

2 To log in to the ILOM, type:

```
$ ssh root@ipaddress
```

where *ipaddress* is the address of the ILOM.

3 Type your password when prompted.

The default is changeme.

The CLI command prompt appears.

▼ **How to Connect to the Web Interface Through Ethernet**

- 1 **Type the IP address of the ILOM into your web browser.**

The login screen appears.

- 2 **Type your user name and password.**

When you first try to access the WebGUI, it prompts you to type the default user name and password. The default user name and password are:

- **Default user name: root**
- **Default password: changeme**

The default user name and password are in lowercase characters.

- 3 **Click Log In.**

The WebGUI appears.

Connecting to ILOM Through a Serial Connector

This section covers the following topics:

- [“How to Connect to ILOM Through the Chassis \(CMM\) Serial Connector” on page 125](#)
- [“How to Connect to the ILOM Using the NEM Serial Connector” on page 127](#)

▼ **How to Connect to ILOM Through the Chassis (CMM) Serial Connector**

- 1 **Connect a serial cable from the serial port on the chassis to a terminal device.**

The terminal device can be an actual terminal, a laptop running a terminal emulator, or a terminal server. It must be set to the following:

- **8N1: eight data bits, no parity, one stop bit**
- **9600 baud (default, can be set to any standard rate up to 57600)**

■ **Disable software flow control (XON/XOFF)**

The cable requires the following pin assignments.

Pin	Signal Description
1	Request To Send (RTS)
2	Data Terminal Ready (DTR)
3	Transmit Data (TXD)
4	Ground
5	Ground
6	Receive Data (RXD)
7	Data Carrier Detect (DCD)
8	Clear To Send (CTS)

2 Press Enter on the terminal device.

This establishes the connection between the terminal device and the CMM ILOM.

The CMM ILOM login prompt appears.

SUNCMMnnnnnnnnnnnnnnnn login:

The first string in the prompt is the default host name. It consists of the prefix SUNCMM and the CMM ILOM's SP address.

3 Log in to the CMM ILOM.

The default user is root and the default password is changeme.

Once you have successfully logged in, the CMM ILOM displays the ILOM default command prompt:

->

4 Navigate to /CH/NEM_{*n*}/SP/c_{*l*}i.

Where *n* is 0 through 1 for NEMs 0 through 1 respectively.

5 Enter the command start.

A prompt appears.

6 Enter y to continue or n to cancel.

If you entered y, the NEM ILOM prompts for its password.

Note – The CMM ILOM logs on to the NEM ILOM using the username in /CH/NEM n /SP/cli/user (where n is the NEM number). The default is root.

7 Enter the ILOM password.

The NEM ILOM's CLI interface appears. You are now connected to the NEM ILOM.

8 When you are done, type exit.

The NEM ILOM exits and the CMM ILOM CLI prompt appears.

▼ How to Connect to the ILOM Using the NEM Serial Connector

1 Connect a terminal or terminal emulator to the serial management port on the Virtualized M2 NEM.

The ILOM login prompt appears.

2 Enter the user name and password when prompted.

The default user is root and the default password is changeme.

The server module ILOM prompt appears.

->

3 When you are done, exit the ILOM by typing:

-> exit

Updating the NEM Firmware

There are two types of firmware to update on the NEM. SAS-2 Expander firmware and NEM SP ILOM firmware. You can update both the ILOM and SAS expander firmware through the Sun Blade 6000 chassis CMM ILOM.

The following procedures are covered in this section:

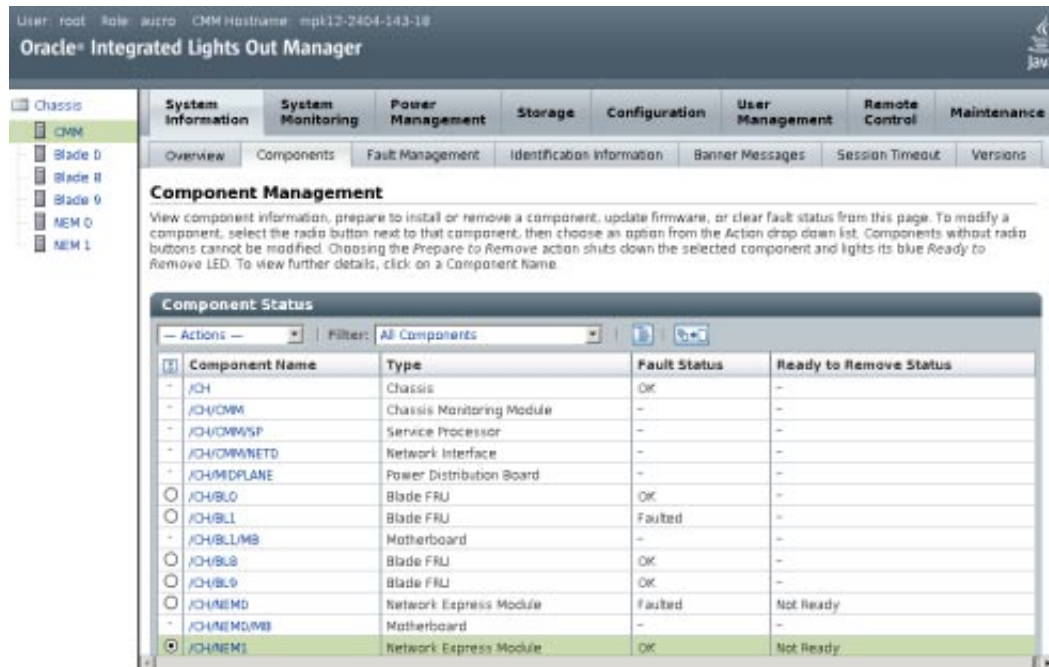
- [“How to Update NEM Firmware Using the Web Interface” on page 128](#)
- [“How to Update NEM ILOM Firmware Using the Command-Line Interface” on page 130](#)

▼ How to Update NEM Firmware Using the Web Interface

Note – NEM ILOM and SAS firmware updates are performed from the CMM. Attempting to perform a NEM SAS firmware update from the NEM SP either through the network or serial port will cause the update to fail.

- 1 **Download the firmware that you need from the Oracle software download site:**
<http://support.oracle.com>
 - a. **At the Oracle Support site, login.**
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 10GbE VMF NEM M2.**
 - e. **In the Release field, click “Select up to 10” and then click the arrow next to the Sun Blade 6000 10GbE VMF NEM M2 folder.**
 - f. **Select (check) the latest two software releases.**
 - g. **Click the Search button.**
 - h. **Download the latest available firmware package and extract it to an accessible folder on the network.**
- 2 **Log into the CMM ILOM as any user with Administrator privileges.**
- 3 **Click on the CMM in the left panel.**
- 4 **Select System Information -> Components**
- 5 **Select the NEM for which you want to upgrade the firmware.**
For example: /CH/NEM1.

6 Select Firmware Update from the Actions drop-down menu.



7 Select the transfer method that you want to use.

The following transfer methods are available: TFTP, FTP, SCP, HTTP, HTTPS, SFTP.



8 Enter the path to the new ILOM or SAS .pkg file in the Select Image File to Upload field or click Browse to locate and select the .pkg file.

9 Click Update.

The firmware update process can take several minutes.

- 10 If you are updating an additional firmware package, repeat the process to update the additional firmware package.**

▼ **How to Update NEM ILOM Firmware Using the Command-Line Interface**

Note – NEM ILOM and SAS firmware updates are performed from the CMM. Attempting to perform a NEM SAS firmware update from the NEM SP either through the network or serial port will cause the update to fail.

1 Download the firmware that you need from the Oracle software download site:

<http://support.oracle.com>

a. At the Oracle Support site, login.

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 10GbE VMF NEM M2.

e. In the Release field, click “Select up to 10” and then click the arrow next to the “Sun Blade 6000 10GbE VMF NEM M2” folder.

f. Select (check) the latest two software releases.

g. Click the Search button.

h. 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 Enter the password.

The ILOM CLI prompt appears.

->

4 Change directories to the blade slot containing the NEM to be upgraded:

-> **cd /CH/NEM x**

where NEM x is the NEM to be upgraded.

5 Enter the following command:

- **For ILOM firmware**

-> **load -source http://http_server_ipaddress/firmware-version.pkg**

- **For SAS expander firmware:**

-> **load -source http://http_server_ipaddress/firmware-version.pkg**

Where *http_server_ipaddress* is the domain name or IP address of your HTTP server 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 When the process completes, ensure that the proper firmware version was installed. Enter the following command:**

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

The following is an example of output that is displayed.

```
/CH/NEM1
Properties:
  type = Network Express Module
  ipmi_name = NEM1
  system_identifier = SUNSP-0000000000
  fru_name = SUN BLADE 6000 VIRTUALIZED MULTI-FABRIC 10GE NEM M2
  fru_version = FW 3.0.10.16, SAS 5.3.4.0
  fru_part_number = 540-7961-02
  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)
```

Sun Blade Zone Manager

Sun Blade Zone Manager is available through the Sun Blade 6000 CMM. This utility allows you to assign storage devices located on Sun Blade Storage Modules with SAS-2 enabled server blades installed in the Sun Blade 6000 chassis.

A Virtualized M2 NEM (or another SAS-2 NEM) must be installed in the chassis to allow a server blade to be zoned with storage devices on a storage blade. At this time, zoning of SAS-2 external ports on the NEM is not supported.

For information on setting up storage zoning, refer to the *Oracle Integrated Lights Out Manager (ILOM) CMM Administration Guide*.

NEM Sensors

The NEM includes a number of sensors that generate entries in the system event log (SEL) when the sensor crosses a threshold. Many of these readings are used to adjust the fan speeds and perform other actions, such as illuminating LEDs and powering off the NEM.

These sensors can also be configured to generate IPMI PET traps as, described in the *Sun Integrated Lights Out Manager 2.0 User's Guide*.

This section describes the sensors and provides details about their operation.



Caution – Do not use any interface other than the ILOM CLI or WebGUI to alter the state or configuration of any sensor or LED. Doing so could void your warranty.

[“NEM Sensor List” on page 132](#) lists the sensors. [“Sensor Details” on page 133](#) provides details about each sensor.

NEM Sensor List

The path to the following sensors is in the following format `/NEM/sensor`

For example: `/NEM/CMM/PRST`

NEM presence

- `NEM/CMM/PRSNT`

Temperature

- `MB/T-AMB`

- MB/T_LINK0
- MB/T_LINK1
- MB/T_SASE

Voltage

- MB/V_+1V
- MB/V_+1V2
- MB/V_+1V5
- MB/V_+1V8
- MB/V_+3V3
- MB/V_+12V
- MB/V_+1V2STBY
- MB/V_+1V8STBY
- MB/V_+3V3STBY
- MB/V_+5VSTBY
- MB/V_+12VSTBY

OEM

- SLOTID

Power unit

- VPS

Sensor Details

The following sections provide detailed information about the sensors.

- [“Slot and Presence Sensors” on page 133](#)
- [“Mainboard Temperature Sensors” on page 134](#)

Slot and Presence Sensors

The following sensors indicate the location and presence of the NEM.

slotid

This is a sensor indicating which slot the NEM is installed in. Values are from 0 to 1.

cmm/prsnt

This is a sensor indicating whether CMM ILOM is present.

TABLE 9 cmm.prsnt

Reading	State	Event	Description	Action
0x0001	Device Absent	Yes	CMM ILOM is absent.	none
0x0002	Device Present	Yes	CMM ILOM is present.	none

Mainboard Temperature Sensors

These temperature sensors are monitored, but are not used as inputs to the fan control algorithm and are not used to turn off system power when they are non-recoverable. No event will be generated for these sensor readings.

mb.t_amb

This sensor monitors the ambient temperature from the internal temperature sensor on the mainboard.

t_link0

This sensor monitors the internal temperature from the chip located on Link0 side of NEM mainboard.

t_link1

This sensor monitors the internal temperature from the chip located on Link1 side of NEM mainboard.

t_sase

This sensor monitors the internal temperature from SAS-2 expander's temperature located on NEM mainboard.

NEM Voltage Sensors

All mainboard voltage sensors are configured to generate the same events, and faults are handled in the same way.

mb.v_+1v

This sensor monitors the 1V main input that is active when the power is on.

mb.v_+1v2

This sensor monitors the 1.2V core input that is active when the power is on. It is not monitored when the host is powered off.

mb.v_+1v5

This sensor monitors the 1.5V core input that is active when the power is on. It is not monitored when the host is powered off.

mb.v_+1v8

This sensor monitors the 1.8V core input that is active when the power is on. It is not monitored when the host is powered off.

mb.v_+3v3

This sensor monitors the 3.3V main input that is active when the power is on.

mb.v_+12v

This sensor monitors the 12V main input that is active when the power is on.

mb.v_+1v2stby

This sensor monitors the 1.2V core input voltage that is active when the standby power is on. It is not monitored when the standby power is off.

mb.v_+1v8stby

This sensor monitors the 1.8V core input voltage that is active when the standby power is on. It is not monitored when the standby power is off.

mb.v_+3v3stby

This sensor monitors the 3.3V core input voltage that is active when the standby power is on. It is not monitored when the standby power is off.

mb.v_+5vstby

This sensor monitors the 5V core input voltage that is active when the standby power is on. It is not monitored when the standby power is off.

mb.v_+12vstby

This sensor monitors the 12V core input voltage that is active when the standby power is on. It is not monitored when the standby power is off.

vps

This sensor reports the NEM's power consumption in Watts.

ILOM Temperature Events

The ILOM SP logs temperature events of the NEM. It also shuts down the NEM if an Upper-Non-Recoverable (UNR) temperature event is reached or exceeded. The event log shows this event as being crossed before shutdown.

The Virtualized NEM is shut down when the NEM reaches the overtemp (UNR) limit. The amber Service LED is lit when any temperature threshold in the following table is crossed.

Sensor	Upper Non-critical	Upper critical limit	Unrecoverable critical limit
MB/T_AMB	60° C	70° C	80° C
MB/T_LINK0	95° C	110° C	120° C
MB/T_LINK0	95° C	110° C	120° C
MB/T_SASE	95° C	110° C	120° C

The following message an example of a NEM shut down due to a SAS Expander overtemp condition:

```
9 Tue May 5 21:45:55 1970 IPMI Log critical ID = 7 : pre-init timestamp :
Temperature : MB/T_SASE : Upper Non-recoverable going high : reading 121 >=
threshold 120 degrees C *
```

The following message an example of a NEM shut down due to an ambient temperature overtemp condition:

```
9 Tue May 5 21:45:55 1970 IPMI Log critical ID = 7 : pre-init timestamp :
Temperature : MB/T_AMB0 : Upper Non-recoverable going high : reading 121 >=
threshold 120 degrees C*
```

The following message an example of a NEM shut down due to a Link overtemp condition:

```
9 Tue May 5 21:45:55 1970 IPMI Log critical ID = 7 : pre-init timestamp :
Temperature : MB/T_LINK0 : Upper Non-recoverable going high : reading 121 >=
threshold 120 degrees C*
```

You might also see the events as shown in the following example:

```
Temperature : MB/T_AMB : Upper Critical going high

Temperature : MB/T_AMB : Upper Non-critical going high
```

These events are only for warning/logging purposes. They do not cause a shutdown. The amber service LED will be lit when these events are experienced.

Enabling Private and Failover Mode

You can enable or disable privacy mode and failover mode either through the web interface or the CLI. See [“10 GbE NIC Virtualization” on page 13](#) for more information on when to enable private or failover mode.

Note – Both private mode and failover mode can be in effect at the same time. However, each state must be enabled one at a time. Enabling them together can cause an error. Private mode will take precedence when they are both enabled.

This section covers the following procedures:

- [“How to Enable Failover or Private Mode Using the Web Interface” on page 137](#)
- [“How to Enable Failover or Private Mode Using the CLI” on page 138](#)

▼ How to Enable Failover or Private Mode Using the Web Interface

- 1 Log into the NEM ILOM web interface.
- 2 Select the Switch Management tab.



- 3 Do one of the following:
 - To enable port failover, check the Enable box for NEM port failover.

- To enable private mode, check the Enable box for NEM private mode.

Do not enable both features in the same screen.

- 4 Click Save.
- 5 If necessary, enable the second feature and click Save.

▼ How to Enable Failover or Private Mode Using the CLI

- 1 Display the current network mode.

```
-> show /SWITCH/network/  
/SWITCH/network  
Targets:  
  
Properties:  
    failover = disabled  
    private_mode = disabled  
  
Commands:  
    cd  
    set  
    show
```

- 2 Set the value in the appropriate mode.

- To set failover mode:

```
-> set /SWITCH/network failover=enabled|disable
```

- To set private mode:

```
-> set /SWITCH/network private_mode=enabled|disable
```

Do not enable both features at the same time.

- 3 If you want to enable the second feature, wait at least 60 seconds.

Using Hot Plug Commands

The ILOM has the following commands available for hot removing and inserting the NEM:

- Prepare to Remove (CLI: `prepare_to_remove_action`): This command prepares the PCIe interfaces of the NEM for hot-removal.
- Return to Service (CLI: `return_to_service_action`): This command restores the PCIe interface connections of a NEM that has just been installed.

This section contains the following procedures:

- “How to Prepare the NEM for Hot Plug Using the Web Interface” on page 139
- “How to Prepare the NEM for Hot Plug Using the Command-Line Interface” on page 140

▼ How to Prepare the NEM for Hot Plug Using the Web Interface

This procedure describes how to prepare the NEM for hot-plug insertion or removal using the web interface. If you are returning a NEM to service after a hot-insertion, you will need to use the CLI command `return_to_service_action = enable` rather than using the web interface. See “How to Prepare the NEM for Hot Plug Using the Command-Line Interface” on page 140.

Before You Begin Read the instructions in “Installing or Replacing the Virtualized M2 NEM” on page 39 before using these commands to perform a hot-plug insertion or removal of the NEM.

- 1 Log into the CMM ILOM web interface.
- 2 Click on CMM, then click on the System Information —> Components tab.
- 3 Click on the radio button next to the NEM that you want to remove or return to service.

For example, `/CH/NEM0`.

The screenshot shows the Oracle Integrated Lights Out Manager (ILOM) web interface. The top navigation bar includes tabs for System Information, System Monitoring, Power Management, Storage, Configuration, User Management, Remote Control, and Maintenance. The left sidebar shows a tree view with Chassis, CMM, Blade 0, Blade 8, Blade 9, NEM 0, and NEM 1. The main content area is titled 'Component Management' and contains a table of components.

Component Status

Component Name	Type	Fault Status	Ready to Remove Status
/CH	Chassis	OK	--
/CH/CMM	Chassis Monitoring Module	--	--
/CH/CMM/SP	Service Processor	--	--
/CH/CMM/NET0	Network Interface	--	--
/CH/MIDPLANE	Power Distribution Board	--	--
/CH/BL0	Blade FRU	OK	--
/CH/BL1	Blade FRU	Faulted	--
/CH/BL1/MB	Motherboard	--	--
/CH/BL8	Blade FRU	OK	--
/CH/BL9	Blade FRU	OK	--
/CH/NEM0	Network Express Module	Faulted	Not Ready

4 Choose one of the following from the Actions menu:

- **Prepare to Remove:** Prepares the NEM for hot removal.

Note – Wait for the OK-to-remove LED to turn on before removing the NEM.

- **Return to Service:** Returns the NEM to normal operation after installing it.

▼ How to Prepare the NEM for Hot Plug Using the Command-Line Interface

Note – Currently, there is a known issue when attempting to issue `prepare_to_remove` and `return_to_service` commands in the NEM ILOM CLI. Therefore, Oracle recommends that these CLI commands be performed in the CMM ILOM CLI instead of the NEM ILOM CLI. Refer to “[Prepare-to-Remove and Return-to-Service Does Not Work on NEM ILOM \(6948103\)](#)” in *Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 Network Express Module Product Notes* for updated information about this issue.

1 Log into the CMM ILOM CLI.

2 Change directories to the NEM to be removed or the NEM that was just inserted:

-> `cd /CH/NEMx`

where NEMx is the NEM to be removed or inserted.

3 Execute one of the following commands:

- `prepare_to_remove_action =enable|disable`

Note – If enabling `prepare_to_remove_action`, wait for the OK-to-remove LED to turn on before removing the NEM.

- `return_to_service_action = enable|disable`

Fixing Problems with Oracle ILOM Using the Preboot Menu

Description	Link
Accessing the preboot menu	“Accessing the Preboot Menu” on page 141
Preboot menu command summary	“Preboot Menu Command Summary” on page 143
Using the <code>edit</code> command and configuring the preboot menu for remote access	“Using the <code>edit</code> Command and Configuring the Preboot Menu for Remote Access” on page 144
Recovering the SP firmware image	“How to Recover the SP Firmware Image” on page 147

The Oracle ILOM preboot menu is a utility that can be used to fix problems with Oracle ILOM that cannot be fixed while it is running. It allows you to interrupt the Oracle ILOM boot process, configure settings, then continue booting Oracle ILOM. Among other things, it allows you to reset the Oracle ILOM root password to factory defaults, restore Oracle ILOM access to the serial port, and recover the SP firmware image.

Accessing the Preboot Menu

To access the preboot menu, you must boot the SP and interrupt the boot process.

There are two ways to interrupt the Oracle ILOM boot process: manually using the Locate button, or by typing **xyzy** during a pause in the bootstrap process.

The first method requires you to have physical access to the NEM.

The second method can be done remotely.

However, note the following:

- You must use a terminal or a terminal emulator. You cannot use an SSH, or an RKVMS session.
- For a remote connection, some preboot menu settings must be configured first. Until they are, you must use the Locate button the first time you access the preboot menu. Then you can configure settings for later remote access. For more about using a remote connection, see [“Using the `edit` Command and Configuring the Preboot Menu for Remote Access” on page 144](#).

See Also

- [“Accessing the Preboot Menu” on page 141](#)
- [“Preboot Menu Command Summary” on page 143](#)

▼ How to Access the Preboot Menu

- 1 **Connect a terminal or a computer running terminal emulation software to the NEM's serial management port using an RJ45 serial cable.**

For details on how to log into Oracle ILOM using a serial connection, refer to [“Connecting to ILOM” on page 122](#).

- 2 **Once logged in, reboot Oracle ILOM by entering the command:**

-> **reset /SP**

- 3 **Interrupt the Oracle ILOM boot process using one of these methods:**

- After resetting the SP, press and hold the Locate button on the NEM front panel until the preboot menu appears.
- Type in **xyzy** when you see the message:

```
Booting linux in
  n seconds...
```

Note – You cannot interrupt the Oracle ILOM boot process by typing **xyzy** until you have configured the settings as described in [“Using the edit Command and Configuring the Preboot Menu for Remote Access” on page 144](#). One of these settings sets the value of *n*, which is the amount of time in seconds that the system waits for your input.

The Oracle ILOM preboot menu appears as shown here.

```
Booting linux in 10 seconds...
```

```
          ILOM Pre-boot Menu
```

```
-----
Type "h" and [Enter] for a list of commands, or "?" [Enter] for
command-line key bindings. Type "h cmd" for summary of 'cmd' command.
```

```
Warning: SP will warm-reset after 300 seconds of idle time.
Set 'bootretry' to -1 to disable the time-out.
```

```
Preboot>
```

- 4 **Use the preboot menu commands to configure settings, reset the password, and so on.**
- 5 **When you are done, enter the boot command to exit the preboot menu and start Oracle ILOM.**

See Also

- [“Preboot Menu Command Summary” on page 143](#)
- [“Using the edit Command and Configuring the Preboot Menu for Remote Access” on page 144](#)

Preboot Menu Command Summary

The preboot menu includes the following commands.

Command	Description
boot	Boots the Oracle ILOM. The preboot menu exits, and Oracle ILOM boots. Note – This command executes a modified boot sequence that does not offer the choice to select the diagnostic level, or to interrupt the boot sequence and return to the preboot menu. To execute the normal boot sequence, use the <code>reset warm</code> command instead.
vers	Displays version information including the hardware type, board rev, Oracle ILOM rev, revisions of PBSW and recovery U-Boot. Shows the checksum integrity of the images, and the preference between redundant images.
help	Displays a list of commands and parameters.
show	Displays a list of SP settings.
edit	Starts an interactive dialog that prompts and changes settings one by one. See “Using the edit Command and Configuring the Preboot Menu for Remote Access” on page 144 for details.
diag	Runs the U-boot diagnostic tests in manual mode. For more on U-boot diagnostic tests, refer to the diagnostics guide for your system.
net	<p>{ config dhcp ping flash }</p> <ul style="list-style-type: none"> ■ config - Starts a dialog that allows you to change Oracle ILOM's network settings. ■ dhcp - Changes the network addressing from static to DHCP. Note – You must set <code>ipdiscovery=dhcp</code> using the <code>net config</code> command first. ■ ping - Sends a ping. ■ flash - Downloads an Oracle ILOM firmware image. See “How to Recover the SP Firmware Image” on page 147. Type help net command for more details on these commands.
reset	<p>{[warm] cold }. Resets the SP and the host.</p> <ul style="list-style-type: none"> ■ warm - Resets the SP without affecting a running host. ■ cold - Resets the SP and the host. It has the effect of powering off the server.

Command	Description
unconfig	<p>{ ilom_conf most all }</p> <p>Causes Oracle ILOM to erase any configuration information and return the values to defaults the next time it boots.</p> <ul style="list-style-type: none">■ <code>ilom_conf</code> - Resets configuration settings but preserves SP network and baudrate, preferred, and <code>check_physical_presence</code> settings.■ <code>most</code> - Resets the SP data storage, but preserves network settings and baudrate, preferred, and <code>check_physical_presence</code> settings.■ <code>all</code> - Resets all SP data storage and settings. Booting Oracle ILOM restores other defaults. <p>Note – None of these options erases the dynamic FRU PROMs.</p>

See Also

- [“Using the edit Command and Configuring the Preboot Menu for Remote Access” on page 144](#)

Using the edit Command and Configuring the Preboot Menu for Remote Access

This section shows how to use the `edit` command to change preboot menu settings. As an example, it also shows how to set the `bootdelay` and `check_physical_presence` settings so that you can interrupt the Oracle ILOM boot process using the **xyzy** command.

Until the `bootdelay` and `check_physical_presence` settings are set to the values shown in this procedure, the only way to interrupt the Oracle ILOM boot process is to hold the Locate button down while Oracle ILOM is booting.

- `bootdelay` is optional, but setting it to a larger value gives you more time to enter the required command.
- `check_physical_presence` must be set to **no**.

Next Steps

- [“How to Use the edit Command and to Configure the Preboot Menu for Remote Access” on page 145](#)

▼ How to Use the `edit` Command and to Configure the Preboot Menu for Remote Access

1 Access the preboot menu.

For more information, see [“Accessing the Preboot Menu” on page 141](#).

2 At the preboot prompt, enter the command:

```
Preboot> edit
```

The preboot menu enters edit mode. In edit mode, the preboot menu displays its selections one by one, offering you a chance to change each one.

- To change a setting, type the new value, and then press Enter.
- To skip to the next setting, press Enter.

3 Press Enter to move through the settings until the `bootdelay` setting appears.

4 To change the `bootdelay` setting, type `3`, `10`, or `30`, and press Enter.

This specifies the number of seconds the SP boot process waits for your input.

The preboot menu redisplay the `bootdelay` setting with the new value.

5 Press Enter.

The next setting appears.

6 Press Enter to move through the settings until the `check_physical_presence` setting appears.

To change the `check_physical_presence` setting, type `no`, and then press Enter.

The preboot menu redisplay the `check_physical_presence` setting with the new value.

7 Press Enter.

The preboot menu asks you to confirm your changes:

```
Enter 'y[es]' to commit changes: [no]
```

8 Enter `y` to exit the edit session and save your changes.

If you want to exit without saving your changes, enter `n`.

The following display shows an edit session where the `bootdelay` and `check_physical_presence` settings are changed. See the table that follows for a description of edit command settings:

```
Preboot> edit
```

Press Enter by itself to reach the next question.

Press control-C to discard changes and quit.

```
Values for baudrate are {[ 9600 ]| 19200 | 38400 | 57600 | 115200 }.
Set baudrate? [9600]
Values for serial_is_host are {[ 0 ]| 1 }.
Set serial_is_host? [0]
Values for bootdelay are { -1 | 3 | 10 | 30 }.
Set bootdelay? [30] 10
Set bootdelay? [10]
Values for bootretry are { -1 | 30 | 300 | 3000 }.
Set bootretry? [<not set>]
Values for preferred are {[ 0 ]| 1 }.
Set preferred? [<not set>]
Values for preserve_conf are {[ yes ]| no }.
Set preserve_conf? [yes]
Values for preserve_users are {[ yes ]| no }.
Set preserve_users? [no]
Values for preserve_password are {[ yes ]| no }.
Set preserve_password? [yes]
Values for check_physical_presence are {[ yes ]| no }.
Set check_physical_presence? [no] no
Set check_physical_presence? [no]
Enter 'y[es]' to commit changes: [no] y
Summary: Changed 2 settings.
Preboot>
```

Setting	Description
baudrate	Sets the baud rate of the serial port. Selections include 9600,19200, 38400, 57600, and 115200.
serial_is_host	If this is set to 0, the serial port connects to Oracle ILOM. If this is set to 1, the serial port connects to the host.
bootdelay	The number of seconds the bootstrap process waits for the user to enter xyzy before booting the SP.
bootretry	The number of seconds the preboot menu waits for user input before timing out and starting the SP. Set to -1 to disable the timeout.
preferred	Unused.
preserve_conf	Setting this to no duplicates the function of the <code>unconfig ilom_conf</code> command, which resets many Oracle ILOM configuration settings, but preserves SP network, baudrate, and <code>check_physical_presence</code> settings the next time the SP is booted.
preserve_users	Setting this to no duplicates the function of the <code>unconfig users</code> command, which resets user information to the default value next time the SP is booted.

Setting	Description
<code>preserve_password</code>	Setting this to <code>no</code> duplicates the function of the <code>unconfig password</code> command, which resets the root password to the default next time the SP is booted.
<code>check_physical_presence</code>	If this is set to <code>Yes</code> , you must press and hold the Locate button to interrupt the SP boot process. If it is set to <code>No</code> , the boot process prompts you to interrupt it. For details, see “Using the <code>edit</code> Command and Configuring the Preboot Menu for Remote Access” on page 144.

▼ How to Recover the SP Firmware Image

The preboot menu provides the ability to recover the Oracle ILOM firmware image by updating (flashing) the SP firmware.

Normally, if the host is running, you can update the SP using the Oracle ILOM CLI or the web interface.

If the host is powered off and the SP firmware image becomes corrupted (making the server SP inaccessible using Oracle ILOM), you can use the following procedure to update it using the preboot menu.

Before You Begin You must have a valid `.flash` SP firmware image file on a TFTP server. In addition, the TFTP server must be accessible over a network connection to your NEM's SP. This file is located on the Tools and Drivers DVD ISO image (available with the Virtualized M2 NEM software release 1.0.1) from the patches and updates section on the Oracle support site: <http://support.oracle.com>

Note – Updating the SP firmware using the preboot menu requires a `.flash` file instead of the `.pkg` file used to update the SP from Oracle ILOM.

- 1 Access the preboot menu.**
For more information, refer to “Accessing the Preboot Menu” on page 141.
- 2 At the preboot prompt, enter the command:**

```
Preboot> net dhcp
```

This configures a DHCP network. You need to be network-connected to access your tftp server.
- 3 Enter the command:**

```
Preboot> net ping tftpIPAddress
```

where *tftpIPAddress* is the IP address of a TFTP server.

This checks to see if the TFTP server is accessible over the network.

4 Enter the command:

```
Preboot> net flash tftpIPAddress path/ILOM-version-Sun_Fire_X4338.flash
```

where:

- *tftpIPAddress* is the IP address of a tftp server
- *path* is the path to the file relative to `/tftpboot`
- *version* is the version of SP firmware

For example:

```
Preboot> net flash 10.8.173.25 images/ILOM-3_0_x_x_rxxx-Sun_Server_X4338.flash
```

This downloads and flashes the firmware image. After a series of messages, the preboot prompt appears.

5 Restart the SP by entering the command:

```
Preboot> reset
```

The preboot menu exits, and the service processor reboots.

Index

A

- accessing, preboot menu, 141
- ASIC, 13
- assigning an IP address, 65
- attention button, 21

B

- booting ILOM, 143
- buttons, 20

C

- cabling SFP+ connectors, 53
- configuring the preboot menu for remote access, 145–147

D

- diagnostics, U-boot, 143
- disk hot plug actions, 27
- drivers
 - Linux operating systems
 - configuring Jumbo Frames, 91
 - configuring the network interface, 78
 - installing, 73
 - removing, 77
 - Solaris operating system
 - configuring hxge device driver, 66
 - configuring Jumbo Frames, 69

- drivers, Solaris operating system (*Continued*)
 - configuring network host files, 65
- Windows operating system
 - installing, 93, 115

E

- edit command, preboot menu, 144
- Ethernet connections, 12

F

- fabric express modules (FEMs), 17
- failover mode, enabling, 137
- features, 9
- features overview, 10

G

- Gigabit Ethernet ports, 19
- graceful NEM removal, 24

H

- hot plug actions and outcomes, 27

I**ILOM**

- booting, 143
 - changing settings, 143
 - connecting, 122
 - description, 121
 - documentation, 121
 - fixing problems, 141
 - private mode, enabling, 137
 - sensors, 132
 - showing settings, 143
 - using the preboot menu, 141
 - version, 143
- ILOM firmware image, recovering, 147–148
- installing an SFP+ module, 51
- Integrated Lights Out Manager, *see* ILOM, 121
- inter-ASIC links
- description, 14
 - functionality, 15

J**Jumbo Frames**

- checking configurations, 70
- enabling in a Solaris SPARC environment, 71

L**LEDs, 20**

- Locate, 21
- Module Activity, 21
- Service Action Required, 22

Linux operating systems

- configuring Jumbo Frames, 91
- configuring the network interface, 78
- installing drivers, 73
- removing drivers, 77

Locate LED and button, 21**M****Module Activity LED, 21****Module Fault LED, 22****N****NEM components**

- 10 GbE connections, 13
 - Ethernet connections, 12
 - virtualized NEM ASIC, 13
- NEM hot plug actions, 27
- network settings, ILOM, 143

O**OK-to-Remove LED, 22****P**

- parameter values, modifying and displaying, 68
- physical appearance, 18
- ports, 19
- Power/OK LED, 21
- preboot menu
- accessing, 141
 - edit command, 144

R

- recovering SP firmware image, 147–148
- remote access, configuring the preboot menu, 145–147
- removing a NEM, 48
- removing an SFP+ module, 52
- replacing a disk, 27
- replacing a NEM, 27
- replacing the storage module, 27

S

- Service Action Required LED, 22
- service processor, 17

- SFP+ module
 - cabling, 53
 - installation, 51
 - removal, 52
- SFP+ ports, 19
- Solaris operating system
 - configuring hxge device driver, 66
 - configuring Jumbo Frames, 69
 - configuring network host files, 65
- SP firmware, recovering, 147–148
- storage module hot plug actions, 27
- surprise NEM removal, 24

T

- terms used, 9

U

- U-boot diagnostics, 143

V

- verifying NEM installation with ILOM CLI, 43
- verifying NEM installation with ILOM web interface, 42

W

- Windows operating system
 - enabling Jumbo Frames, 113
 - installing drivers, 93, 115
 - installing enclosure device, 109
 - installing network controller drivers, 93
 - installing VLAN drivers, 99

