



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

Sun Microsystems, Inc.
www.sun.com

Part No. 820-6561-13
January 2010, Revision A

Submit comments about this document by clicking the Feedback[+] link at: <http://docs.sun.com>

Copyright © Copyright © 2010 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Solaris and Sun Blade Virtualized MultiFabric GbE NEM are trademarks or registered trademarks of Sun Microsystems, Inc., or its subsidiaries, in the U.S. and other countries.

LSI™. LSI Logic™.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

Use of any spare or replacement CPUs is limited to repair or one-for-one replacement of CPUs in products exported in compliance with U.S. export laws. Use of CPUs as product upgrades unless authorized by the U.S. Government is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2010 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuels relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains listés à l'adresse <http://www.sun.com/patents> et un ou les brevets supplémentaires ou les applications de brevet en attente aux Etats - Unis et dans les autres pays.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems Sun Microsystems, le logo Sun, Solaris et Sun Blade Virtualized MultiFabric GbE NEM sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., ou ses filiales, aux Etats-Unis et dans d'autres pays.

LSI™. LSI Logic™.

Ce produit est soumis à la législation américaine en matière de contrôle des exportations et peut être soumis à la réglementation en vigueur dans d'autres pays dans le domaine des exportations et importations. Les utilisations, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers les pays sous embargo américain, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exhaustive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites. L'utilisation de pièces détachées ou d'unités centrales de remplacement est limitée aux réparations ou à l'échange standard d'unités centrales pour les produits exportés, conformément à la législation américaine en matière d'exportation. Sauf autorisation par les autorités des Etats-Unis, l'utilisation d'unités centrales pour procéder à des mises à jour de produits est rigoureusement interdite.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.



Please
Recycle



Adobe PostScript

Contents

Preface ix

1. Features of the Virtualized Multi-Fabric 10GbE NEM 1

Terminology 2

Features Overview 3

Overview of the Virtualized Multi-Fabric 10GbE NEM Components 4

SAS Connections 4

Ethernet Connections 5

10 GbE NIC Virtualization 6

Virtualized NEM ASIC 7

Enabling or Disabling the IAL 8

Fabric Express Modules 10

Service Processor 10

Physical Appearance of the Virtualized Multi-Fabric 10GbE NEM 11

External Virtualized NEM Ports 12

External SAS x4 Ports 12

Gigabit Ethernet Ports 12

SFP+ Ports 12

NEM Front Panel and LEDs 12

RJ45 Ethernet Connector Port LEDs 14

2. Installing or Replacing the Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module 17

Installing a NEM 18

- ▼ To Install a NEM 18

Verifying NEM Installation 20

- ▼ To Verify Installation With the CMM ILOM Web Interface 20
- ▼ To Verify Installation With the CMM ILOM CLI 22

Removing a NEM 23

- ▼ To Remove a NEM 23

Replacing a NEM 25

- ▼ To Replace a NEM 25

Installing SFP+ Optical Transceiver Modules 26

- ▼ To Install an SFP+ Optical Transceiver Module 27
- ▼ To Remove an Optical Transceiver SFP+ Module 28

Cabling the SFP+ Connectors 29

3. Installing and Configuring the Virtualized Multi-Fabric 10GbE NEM Drivers 31

Booting Over the NEM 10-Gigabit Ethernet Port 32

Booting over the Network With an x86 Blade Server 32

Important Considerations for Booting Over the Network With an x86 Blade Server Over the Virtualized NEM 32

- ▼ To Set Up the BIOS for Booting From the Virtualized NEM 33
- ▼ To Use the F12 Key to Initiate the Network Boot 36
- ▼ To Use the F8 Key to Boot Off the PXE Server 36

Booting over the Network With a SPARC Blade Server 37

- ▼ To Boot Over the Network Using the OpenBoot on a SPARC System 38

Installing and Configuring the hxge Driver on a Solaris SPARC or x86 Platform 40

Installing and Removing the Driver on a Solaris Platform 41

▼ To Download and Install the Driver on a Solaris Platform	41
▼ To Remove the Driver From a Solaris Platform	42
Configuring the Network Host Files	43
▼ To Configure the Network Host Files	43
Configuring the <code>hxge</code> Device Driver Parameters	44
▼ To Set Parameters Using the <code>hxge.conf</code> File	44
▼ To Specify Parameter Values Using the <code>ndd</code> Utility	47
Configuring the Jumbo Frames Feature	48
Jumbo Frames Overview	48
Checking Jumbo Frames Configurations	48
Enabling Jumbo Frames in a Solaris Environment	50
Installing and Configuring the <code>hxge</code> Driver on a Linux Platform	52
Installing and Removing the Driver on a Linux Platform	52
▼ To Download and Install the Driver on a Linux Platform	53
▼ To Remove the Driver From a Linux Platform	58
Configuring the Network Interface	59
Temporary <code>hxge</code> Network Interface Configuration	59
Permanent <code>hxge</code> Network Interface Configuration	61
Checking and Testing the <code>hxge</code> Device	65
Changing the <code>hxge</code> Driver Configuration	67
Troubleshooting the Driver	72
Debug Messaging Parameter Overview	72
Setting the Debug Messaging Parameter	73
Configuring Jumbo Frames	75
▼ To Temporarily Configure Jumbo Frames Support	75
▼ To Permanently Enable Jumbo Frame Support	76
Installing and Configuring Drivers on a Windows Platform	77
Installing Drivers on a Windows Platform	78

Installing and Uninstalling the Sun Blade 6000 10GbE Network Controller	78
Installing and Uninstalling the Sun Blade 6000 10GbE VLAN Driver (Optional)	84
Adding and Removing VLANs	90
Installing the Virtualized Multi-Fabric 10GbE NEM Enclosure Device	93
Enabling Jumbo Frames	104
▼ To Enable Jumbo Frames	105
Installing and Configuring Drivers on a VMware ESX Server Platform	106
Installing the ESX Server Drivers on an Existing ESX Server	106
▼ To Install a Virtualized NEM Driver on an Existing ESX 3.5 Server	106
▼ To Install a Virtualized NEM Driver on an Existing ESX/ESXi 4.0 Server Using vihostupdate	107
▼ To Install a Virtualized NEM Driver on an Existing ESX 4.0 Server Using esxupdate	108
Installing the ESX Server Drivers on With a New ESX Installation	109
▼ To Install a Virtualized NEM Driver With a New ESX3.5 Installation	109
▼ To Install a Virtualized NEM Driver With a New ESX4.0 Installation	110
Configuring the Virtual NEM Network Adapters	111
▼ To Configure the Virtual NEM Network Adapters	111
Configuring Jumbo Frames	112
Jumbo Frames Guidelines	112
▼ To Create a Jumbo Frames-enabled vSwitch	113
▼ To Create a Jumbo Frames-enabled VMkernel Interface	113
4. The Common Array Manager	115
About Enclosure Management	115
CAM	116

	CAM Agent	116
	Problem Resolution	117
	Obtaining CAM Software	117
	Using CAM With Disk Blades and Virtualized Multi-Fabric 10GbE NEM	118
	Monitoring Component Health	118
	Upgrading Expander Firmware	119
	▼ To Upgrade Expander Firmware	119
A.	Using the <code>lsiutil</code> Utility to Save and Restore HBA Persistence Mappings	127
	Where to Obtain the <code>lsiutil</code> Software	127
	Installing <code>lsiutil</code>	128
	Using the Interactive <code>lsiutil</code> Menus to Clear Your LSI Host Bus Adapter Non-Present Persistence Mapping	128
	▼ To Clear Your Non-Present Persistence Mappings	128
	Saving a Snapshot of Your Host Bus Adapter Persistence Mappings	131
	▼ To Save a Snapshot of Your Host Bus Adapter Persistence Mappings	131
B.	ILOM Supplement	135
	Documentation and Updates	135
	Product Updates	136
	Connecting to ILOM	136
	Option 1 - Ethernet	137
	▼ To Find the ILOM's IP Address	137
	▼ To Connect to the CLI	138
	▼ To Connect to the WebGUI	139
	Option 2 - Chassis Serial Connector	139
	▼ To Connect to ILOM Through the Chassis Serial Connector	139
	Option 3 - NEM Serial Connector	141

▼ To Connect to the ILOM Using the NEM Serial Connector	141
Temperature, Voltage, and Fan Sensors	142
Board Sensor List	142
Sensor Details	143
slotid	143
nem/cmm/prsnt	143
nem/ok	143
nem/ok2rm	144
nem/service	144
nem/locate	144
Mainboard Temperature Sensors	145
mb.t_amb[0..2]	145
NEM Voltage Sensors	145
mb.v_+12v	145
mb.v_+5v	145
mb.v_+3v3	145
mb.v_+1v2	146
ILOM Proxy	147
Starting the ILOM Proxy CLI Program	147
▼ To Start the ILOM Proxy CLI Program	147
Enable or Disable IAL Dynamic Mode	148
▼ To Enable or Disable IAL Dynamic Mode Using the Web Interface	148
▼ To Enable or Disable Dynamic Mode Using the CLI	149
Index	151

Preface

This document describes how to install and configure the Sun Blade™ 6000 Virtualized Multi-Fabric 10GbE Network Express Module (NEM) in a powered-on Sun Blade 6000 Modular System.

These instructions are designed for enterprise system administrators with experience installing network hardware and software.

How This Document Is Organized

[Chapter 1](#) provides an overview of the features of the Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module (NEM).

[Chapter 2](#) describes how to install or replace the NEM and verify that it has been installed correctly.

[Chapter 3](#) provides information on installing and configuring the NEM drivers.

[Chapter 4](#) describes the Common Array Manager (CAM) application.

[Appendix A](#) provides instructions for using the `lsiutil` application.

[Appendix B](#) provides information on using the Integrated Lights Out Manager (ILOM) with the NEM.

Related Documentation

The documents listed as online are available at:

(<http://docs.sun.com>)

Document	Part Number	Available
<i>Sun Blade 6000 Modular System Installation Guide</i>	820-0050	Printed and online
<i>Sun Blade 6000 Modular System Service Manual</i>	820-0051	Online
<i>Sun Blade 6000 Modular System Safety and Compliance Guide</i>	820-0053	Online
<i>Sun Blade 6000 Modular System Product Notes</i>	820-0055	Online
<i>Sun Blade 6000 Disk Module Administrator's Guide</i>	820-1702	Online
<i>Sun Blade 6000 Disk Module Configuration Guide</i>	820-6547	Online
<i>Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module Product Notes</i>	820-6564	Online
<i>SFP+ Module Installation Guide for the Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module</i>	820-6752	Printed and online
<i>Integrated Lights Out Manager (ILOM) 2.0 User Guide</i>	820-1188	Online

Documentation, Support, and Training

Sun Function	URL
Documentation	(http://docs.sun.com/)
Support	(http://www.sun.com/support/)
Training	(http://www.sun.com/training/)

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

(<http://www.sun.com/hwdocs/feedback>)

Please include the title and part number of your document with your feedback:

Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module User's Guide,
Sun part number: 820-6561-13.

Features of the Virtualized Multi-Fabric 10GbE NEM

The Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module (NEM) is a multi-purpose connectivity module for the Sun Blade 6000 modular system. The 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 and connects server modules (blades) in a Sun Blade 6000 modular system chassis with disk modules in the same chassis.

This chapter includes the following sections:

- [“Features Overview” on page 3](#)
- [“Overview of the Virtualized Multi-Fabric 10GbE NEM Components” on page 4](#)
- [“Physical Appearance of the Virtualized Multi-Fabric 10GbE NEM” on page 11](#)

Note – If you have not already read the *Sun Blade 6000 Disk Module Configuration Guide*, you should do so before proceeding with this manual.

Terminology

The following terms are used in this document:

chassis	The Sun Blade 6000 Modular System blade enclosure.
disk module (or disk blade)	The Sun Blade 6000 Disk Module. The terms <i>disk module</i> and <i>disk blade</i> are used interchangeably.
server module (or server blade)	Any server module (blade) that will interoperate with a disk module (blade). Examples are the Sun Blade X6220, X6240, X6250, X6440, X6450, T6300, and T6320 server modules. The terms <i>server module</i> and <i>server blade</i> are used interchangeably.
Virtualized NEM	The Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module that plugs into a Sun Blade 6000 chassis (abbreviated Virtualized Multi-Fabric 10GbE NEM).
Virtualized NEM ASIC	A shortened reference to the ASICs embedded in the Sun Blade 6000 Virtualized Multi-Fabric 10GbE 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 NEM is one example.
SAS NEM	A generic term that applies to any Network Express Module that supports SAS connectivity. The Virtualized Multi-Fabric 10GbE NEM is one example of a SAS 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 NEM are listed in [TABLE 1-1](#).

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

Feature	Specifications
Management	<ul style="list-style-type: none">• Arm7-S Processor (SAS expander device)• 2MB SRAM, 2MB FLASH, 8KB Serial EEPROM• Aspeed AST2000 (service processor)• 128MB DDR DRAM, 16MB FLASH
Management Interfaces	<ul style="list-style-type: none">• 10/100BASE-T Ethernet port to the chassis management module (CMM)• I2C to CMM
Ports	<ul style="list-style-type: none">• Four External x4 mini SAS ports• Ten 10/100/1000BASE-T Ethernet (RJ-45) ports• Two SFP+ ports that support the SFP+ short range (SR), long range (LR) and long range multi-mode (LR-M) modules
Server Module Interfaces (per Blade)	<ul style="list-style-type: none">• PCIe up to x8• SAS x2• GbE x1
Updates	Most software and 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
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	<ul style="list-style-type: none">• Cooling: top-to-back forced air• Humidity: 10-90% non-condensing• Temperature: 5-35°C operating (-40-70°C storage)• Altitude: 0-10000 ft (3048 meters)

Overview of the Virtualized Multi-Fabric 10GbE NEM Components

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

The following topics are covered in this section:

- [“SAS Connections” on page 4](#)
- [“Ethernet Connections” on page 5](#)
- [“10 GbE NIC Virtualization” on page 6](#)
- [“Service Processor” on page 10](#)

SAS Connections

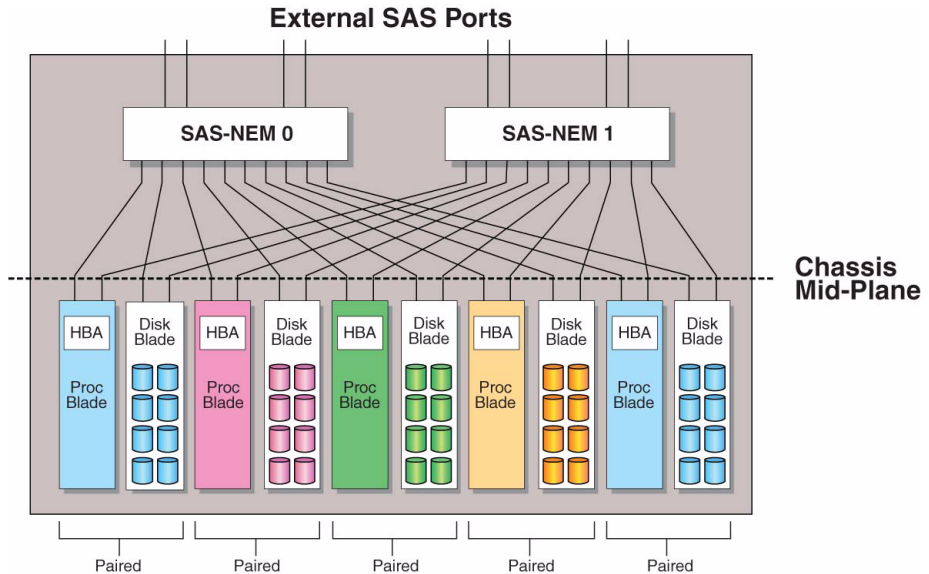
Each Virtualized Multi-Fabric 10GbE NEM has 10 SAS x2 connections to the modules in the chassis through the chassis mid-plane. There is one SAS x2 connection for each module slot. External SAS connections are not currently supported.

One of the functions of the Virtualized Multi-Fabric 10GbE NEM SAS expander is to provide zoning for the disks that are visible to the server modules in the chassis. This includes disks on the server modules, disks on disk modules, and external disks. Firmware on the Virtualized Multi-Fabric 10GbE NEM SAS expander controls zoning (that is, which disks can be seen by each individual server module).

The default zoning is such that the server modules and disk modules must be placed in the chassis in pairs. The pairs can be in slots 0+1, 2+3, 4+5, 6+7, or 8+9. Other pairings, such as 1+2 or 5+6, do not couple the server and storage modules.

Thus, in the recommended use of default zoning, a server module in an even-numbered slot, n , sees all the disks in slot $n+1$, but no other disks. [FIGURE 1-1](#) shows the default zoning configuration for server modules and disk modules.

FIGURE 1-1 Schematic of Chassis With Server Modules (Processor Blades) and Disk Modules (Disk Blades) in the Default Paired Zoning



The disk module that connects to a server module in slot n (where $n = 0, 2, 4, 6,$ or 8) must be in slot $n+1$.

Note – A server module that is not paired with a disk module works in any slot.

For more information on the Sun Blade 6000 disk module, refer to the documentation site at:

(<http://docs.sun.com/app/docs/prod/blade.6000disk#hic>)

Ethernet Connections

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

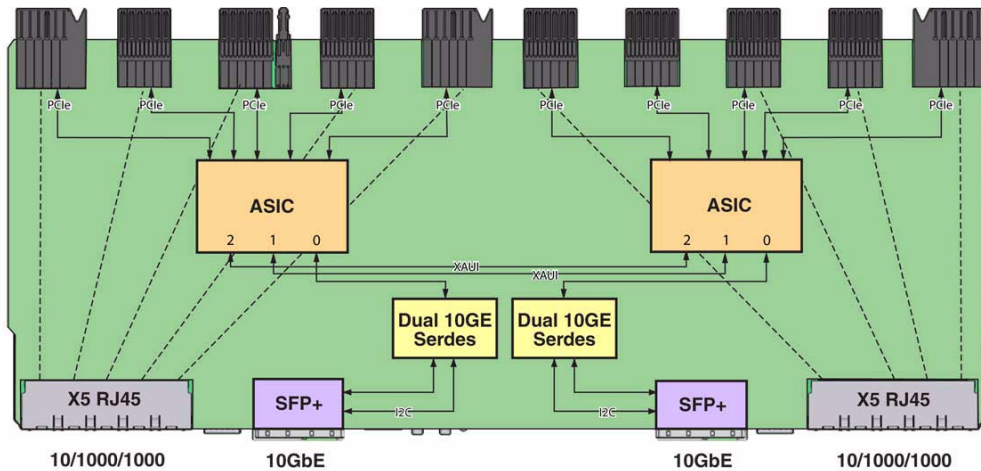
The Virtualized Multi-Fabric 10GbE NEM also provides a 10 GbE connection for each server module. Each server module appears to have its own 10 GbE NIC through the Virtualized NEM ASICs. All the server modules share two physical small form-factor pluggable (SFP+) 10 GbE ports via a dual channel 10 GbE

Serializer/Deserializer (SerDes) per ASIC. In a Sun Blade 6000 chassis, five server modules connect to a single Virtualized NEM ASIC and share its 10 GbE port. Two Virtualized 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 6 for further information on the NEM 10 GbE virtualization.

FIGURE 1-2 shows a schematic of the NEM Ethernet connections.

FIGURE 1-2 Schematic of the Virtualized NEM



10 GbE NIC Virtualization

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

This section covers the following topics:

- “[Virtualized NEM ASIC](#)” on page 7
- “[Enabling or Disabling the IAL](#)” on page 8
- “[Fabric Express Modules](#)” on page 10

Virtualized NEM ASIC

The Virtualized 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 NEM contains two Virtualized NEM ASICs, which can operate in two different modes:

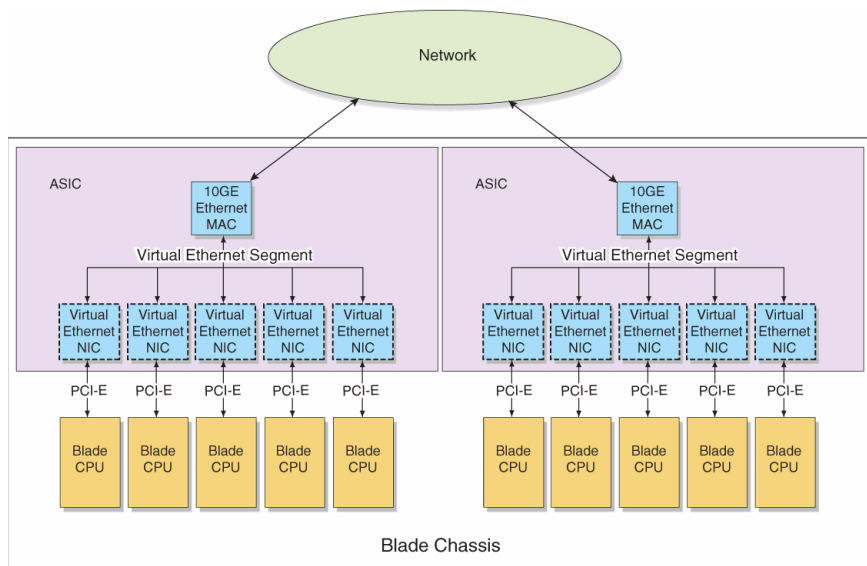
- “Bandwidth Mode” on page 7
- “Connectivity Mode” on page 8

Bandwidth Mode

The Virtualized NEM ASICs can operate in bandwidth mode where the Virtualized NEM ASICs act without knowledge of each other or connectivity (IAL) mode where all server modules share a single 10GbE uplink. Each ASIC provides 10 GbE network access for the five hosts attached to it.

FIGURE 1-3 shows how the ASICs operate in bandwidth mode.

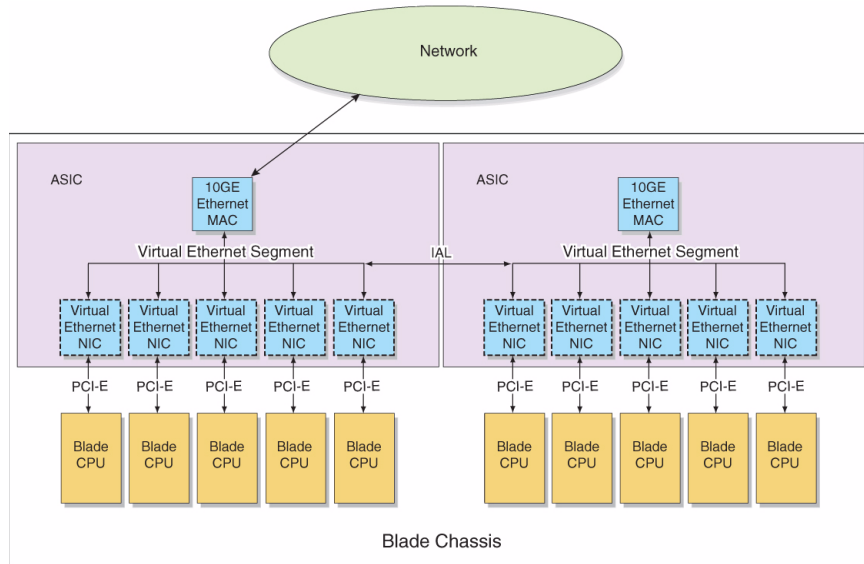
FIGURE 1-3 Blade Network Topology With 2 x 10 GbE



Connectivity Mode

Two Virtualized 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 illustration of this configuration is shown in [FIGURE 1-4](#).

FIGURE 1-4 Blade Network Topology With 1 x 10 GbE



Because of the symmetry of the IAL, either Virtualized NEM ASIC can be connected to the external network. This configuration can be used for higher aggregation per 10 GbE port (1x10 GbE for 10 server modules).

Enabling or Disabling the IAL

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

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

Refer to [TABLE 1-2](#) for more information on IAL.

TABLE 1-2 explains the IAL functionality:

TABLE 1-2 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 NEM ASIC stays in Bandwidth mode. One of the ASICs loses network connectivity. IAL is still disabled.
Dynamic	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.
Dynamic	Two SFP+ modules exist in the NEM, and the NEM is alive in the chassis.	IAL is disabled.	One SFP+ module is removed, disabled, or damaged.	IAL dynamically configures so traffic flows to the active SFP+ connection. IAL is enabled.

Refer to [“Enable or Disable IAL Dynamic Mode” on page 148](#) for instructions on how to enable dynamic mode through the NEM service processor.

Note – If IAL is in dynamic mode and an additional SFP+ module is installed, you must install a fiber optic 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 NEM.
- Dynamic 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 some 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. “Required FEMs” on page 10 shows which FEMs are required for each server module. Refer to the Virtualized NEM product web pages on (<http://sun.com>) for updates to the list of FEMs required for the server modules.

TABLE 1-3 Required FEMs

Server Blade Module	FEM Part Number
T6300	FEM not needed
T6320	X4835A
T6340	X4835A
X6220	FEM not needed
X6240	X4263A
X6250	X4681A
X6440	X4263A
X6450	X4681A

Note – The Sun Blade 6000 Disk Module does not have an Ethernet controller. Therefore the Ethernet ports on the Virtualized Multi-Fabric 10GbE NEM that correspond to slots containing Sun Blade 6000 Disk Modules are not used.

Service Processor

The Virtualized Multi-Fabric 10GbE NEM has an Aspeed AST2000 controller as its service processor (SP), which is responsible for the control and management of the Virtualized 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 NEM FPGA and firmware using GPIOs connected to the JTAG port of the FPGA.

Refer to [Appendix B](#) for more information on the Integrated Lights Out Manager (ILOM) server management application for the SP.

Physical Appearance of the Virtualized Multi-Fabric 10GbE NEM

FIGURE 1-5 shows an overview of the Virtualized Multi-Fabric 10GbE NEM.

FIGURE 1-5 NEM Component Overview




TABLE 1-4 NEM Components

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

TABLE 1-4 NEM Components (*Continued*)

Label	Description
5, 6	Ejector Levers
7	SAS Connectors (2) (Not currently supported)
8	10GbE connector (1) (Requires SFP+ module)
9	10/100/1000 TPE, RJ45 Connectors (5)
10	Virtualized Multi-Fabric 10GbE NEM chassis
11	Midplane Connectors (10)

External Virtualized NEM Ports

This section describes the external ports on the NEM.

External SAS x4 Ports

External SAS is not currently supported for the Virtualized NEM.

Gigabit Ethernet Ports

There are 10 pass-through RJ-45 Gigabit Ethernet ports on the Virtualized Multi-Fabric 10GbE 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.

NEM Front Panel and LEDs

[FIGURE 1-6](#) shows the NEM front panel and LEDs, viewed from the back of the chassis. [TABLE 1-5](#) describes the LED behavior.

FIGURE 1-6 NEM Front LEDs and Buttons




TABLE 1-5 NEM Front Panel and LED Descriptions

	LED/Button Name	Description
1	Attention button	Not currently supported.
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.
4	Module Fault LED (amber)	<p>Has two states:</p> <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)	Not currently supported.
6	SAS Activity	Not currently supported.

TABLE 1-5 NEM Front Panel and LED Descriptions (Continued)

	LED/Button Name	Description
7/ 8	Activity/Link for 10/100/1000 MbE	See TABLE 1-6 .
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.

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. [TABLE 1-6](#) provides a chart for interpreting the link-speed relationships.

TABLE 1-6 Link-Speed LED Colors for Each Server Blade

Sun Blade Server Model	10 MbE (Right LED)	100 MbE (Right LED)	1000 MbE (Right LED)
X6220	Off	Amber	Green
X6250	Off	Amber	Green
X6450	Off	Amber	Green
T6300	Off	Green	Amber
T6320	Off	Green	Amber
T6340	Off	Amber	Green

When an Ethernet port is connected to an x64 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 [TABLE 1-6](#). SPARC-based blades are designated by a T in front of the server module number (for example, T6300).

Installing or Replacing the Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module

This chapter describes how to replace a Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module (Virtualized NEM) in a powered-on Sun Blade 6000 Series Chassis. This chapter also includes instructions to verify that the replacement Virtualized NEM has been installed correctly.

This chapter contains the following sections:

- “Installing a NEM” on page 18
- “Verifying NEM Installation” on page 20
- “Removing a NEM” on page 23
- “Replacing a NEM” on page 25
- “Installing SFP+ Optical Transceiver Modules” on page 26
- “Cabling the SPF+ Connectors” on page 29



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, Sun 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 Sun as part number 250-1007.

Installing a NEM

You can insert either one or two Virtualized Multi-Fabric 10GbE NEMs in the Sun Blade 6000 chassis. If you insert only one, it should go 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 NEM can be paired with a pass-through Gigabit Ethernet NEM or another Virtualized NEM only. Mixing of Multi-Fabric NEM types is not supported.

▼ 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 25](#).

1. Power off the chassis.

Refer to the chassis documentation for information.

2. 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.


3. 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.

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

[FIGURE 2-1](#) shows how to install the NEM.

FIGURE 2-1 Installing a Virtualized Multi-Fabric 10GbE NEM



5. Power on the chassis.

Refer to the chassis documentation for information on powering on the chassis.

6. Execute the procedure [“Saving a Snapshot of Your Host Bus Adapter Persistence Mappings”](#) on page 131 on all server blades in the chassis that meet all three of these requirements:

- The server blade is running CAM or a CAM agent.
- The server blade uses an LSI host bus adapter.
- The server blade is running the Solaris OS.

7. Connect cables, as needed. See [“Cabling the SPF+ Connectors”](#) on page 29.

Verifying NEM Installation

The following topics are covered in this section:

- [“To Verify Installation With the CMM ILOM Web Interface” on page 20](#)
- [To Verify Installation With the CMM ILOM CLI](#)

▼ To Verify Installation With the CMM ILOM Web Interface

1. From the CMM ILOM left navigation bar, select CMM.
2. In a web browser, type the IP address of the CMM in the location bar.
3. Select Components from the second row of tabs ([FIGURE 2-2](#)).
4. Select /CH/NEM0 ([FIGURE 2-3](#)).

If the NEM does not appear in ILOM, verify that the NEM is properly seated in the chassis.

FIGURE 2-2 CMM ILOM Screen for Viewing NEM Information

Sun™ Integrated Lights Out Manager

Java™

Chassis

CMM

Blade 0

Blade 1

Blade 2

Blade 3

Blade 4

Blade 5

Blade 6

Blade 7

Blade 8

Blade 9

System Information

System Monitoring

Configuration

User Management

Remote Control

Versions

Session Time-Out

Components

Identification Information

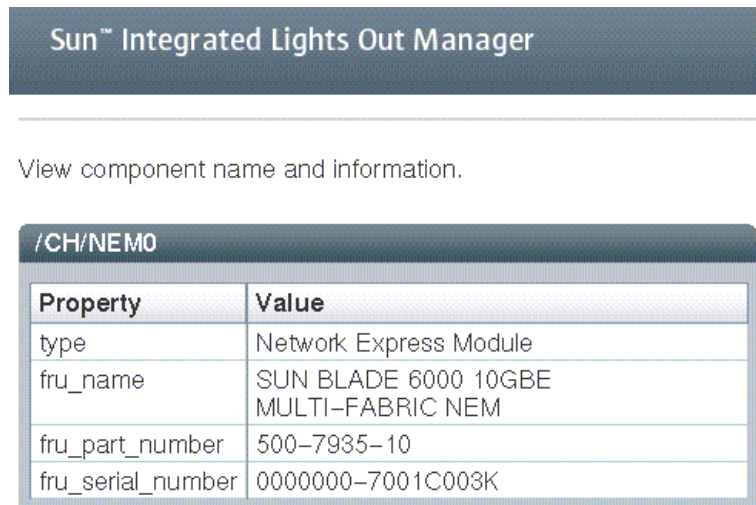
Component Management

View component information from this page. To view further details, click on a Component Name.

Component Management Status

Component Name	Type
/CH	Chassis
/CH/CMM	Chassis Monitoring Module
/CH/CMM/NET0	Network Interface
/CH/CMM/SP	Service Processor
/CH/CMM/MB	Motherboard
/CH/BL0	Blade
/CH/BL0/SP	Service Processor
/CH/BL0/SEEPROM	PROM
/CH/NEMO	Network Express Module

FIGURE 2-3 FRU Information for NEM0 in the Web Browser



The screenshot shows the Sun Integrated Lights Out Manager interface. At the top, there is a header bar with the text "Sun™ Integrated Lights Out Manager". Below this, a line of text reads "View component name and information." Underneath, there is a table titled "/CH/NEM0". The table has two columns: "Property" and "Value". The rows in the table are: "type" with value "Network Express Module", "fru_name" with value "SUN BLADE 6000 10GBE MULTI-FABRIC NEM", "fru_part_number" with value "500-7935-10", and "fru_serial_number" with value "0000000-7001C003K".

/CH/NEM0	
Property	Value
type	Network Express Module
fru_name	SUN BLADE 6000 10GBE MULTI-FABRIC NEM
fru_part_number	500-7935-10
fru_serial_number	0000000-7001C003K

▼ 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 [Appendix B](#).

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

1. Access the CMM ILOM.

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

```
-> show /CH/NEM0
Targets:
SEEPROM
SP

Properties:
type = Network Express Module
fru_name = SUN BLADE 6000 10GBE Hydra NEM
fru_part_number = 501-7935-02
fru_serial_number = 0000000-7001C003K

Commands:
cd
show
```

2. 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://docs.sun.com/app/docs/prod/blade.6000mod#hic>)

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 25.

▼ To Remove a NEM

1. Power off the chassis.

Refer to the chassis documentation for information.

2. In the rear of the chassis, locate the NEM that you want to remove.

3. Remove all cables from the NEM.


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

5. To unlatch the NEM from the chassis, open the ejector levers by extending them outward.
6. 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.

FIGURE 2-4 shows how to remove a NEM.

FIGURE 2-4 Removing a NEM



7. Power on the chassis.

Refer to the chassis documentation for information.

Replacing a NEM

If a NEM fails, you need to replace it. You must power off the chassis before you remove the NEM. The server modules must be informed if you are replacing a non-redundant NEM.



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



Caution – If you are replacing a functional SAS-NEM (for example replacing a Multi-Fabric NEM with a 10GbE Virtualized Multi-Fabric NEM), 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.

▼ To Replace a NEM

1. Power off the chassis.

Refer to the chassis documentation for information.

2. In the rear of the chassis, locate the NEM that you want to remove.

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

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

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

6. **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.

[FIGURE 2-4](#) shows how to remove a NEM.

7. **Determine if the following are present in the chassis:**

- A server blade that is running CAM or a CAM agent.
- A server blade that uses an LSI host bus adapter.

If *both* of the above conditions are true, continue with [Step 8](#). If *either* one of the above conditions is not true, continue to [Step 11](#).

8. **Power on the chassis.**

Refer to the chassis documentation for information.

9. **Follow the procedure “[Using the Interactive lsiutil Menus to Clear Your LSI Host Bus Adapter Non-Present Persistence Mapping](#)” on page 128 on all server blades in the chassis that meet both of these requirements listed in [Step 7](#).**

10. **Power off the chassis.**

11. **Insert the new NEM into the chassis.**

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

13. **Power on the chassis.**

14. **Execute the procedure “[Saving a Snapshot of Your Host Bus Adapter Persistence Mappings](#)” on page 131 on all server blades in the chassis that meet *all three* of these requirements:**

- The server blade is running CAM or a CAM agent.
- The server blade uses an LSI host bus adapter.
- The server blade is running the Solaris OS.

Installing 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 “[Connectivity Mode](#)” on page 8.

▼ 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.


FIGURE 2-5 Installing an Optical Transceiver SFP+ Module



▼ To Remove an Optical Transceiver SFP+ Module

1. Press the optical cable connector latch down and gently pull out the optical cable (FIGURE 2-6).
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.

FIGURE 2-6 Removing an Optical Module from the NEM



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 NEM SPF+ connectors.

TABLE 2-1 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 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.

Installing and Configuring the Virtualized Multi-Fabric 10GbE NEM Drivers

The drivers for the Virtualized Multi-Fabric 10GbE NEM are available on the NEM Tools and Drivers CD. You can access the Tools and Drivers CD at:

(<http://www.sun.com/servers/blades/downloads.jsp>)

Install the Virtualized Multi-Fabric 10GbE NEM drivers that are appropriate for the server module operating system. The locations of the drivers for each operating system are described in the corresponding operating system section.

This chapter also contains information on booting over the 10GbE network.

This chapter contains the following sections:

- “Booting Over the NEM 10-Gigabit Ethernet Port” on page 32
- “Installing and Configuring the hxge Driver on a Solaris SPARC or x86 Platform” on page 40
- “Installing and Configuring the hxge Driver on a Linux Platform” on page 52
- “Installing and Configuring Drivers on a Windows Platform” on page 77
- “Installing and Configuring Drivers on a VMware ESX Server Platform” on page 106

Booting Over the NEM 10-Gigabit Ethernet Port

In order to perform a network boot via the Virtualized NEM, make sure that the Virtualized 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 NEM's MAC.

The following procedures describe how to boot over the network using the system BIOS or OpenBoot.

- [“Booting over the Network With an x86 Blade Server” on page 32](#)
- [“Booting over the Network With a SPARC Blade Server” on page 37](#)

Booting over the Network With an x86 Blade Server

This section covers the following topics:

- [“Important Considerations for Booting Over the Network With an x86 Blade Server Over the Virtualized NEM” on page 32](#)
- [“To Set Up the BIOS for Booting From the Virtualized NEM” on page 33](#)
- [“To Use the F12 Key to Initiate the Network Boot” on page 36](#)
- [“To Use the F8 Key to Boot Off the PXE Server” on page 36](#)

Important Considerations for Booting Over the Network With an x86 Blade Server Over the Virtualized NEM

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

The Virtualized 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:

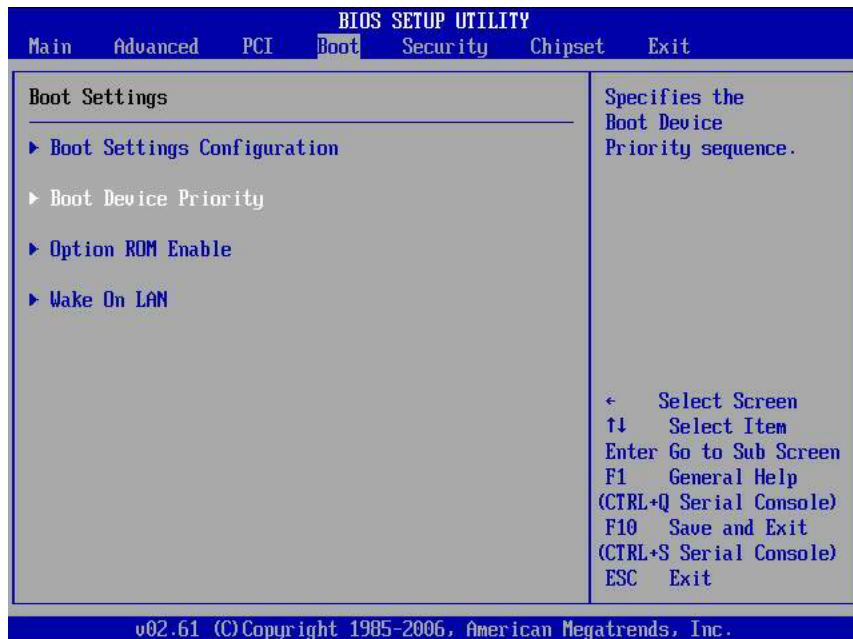
- “To Set Up the BIOS for Booting From the Virtualized NEM” on page 33
- “To Use the F12 Key to Initiate the Network Boot” on page 36

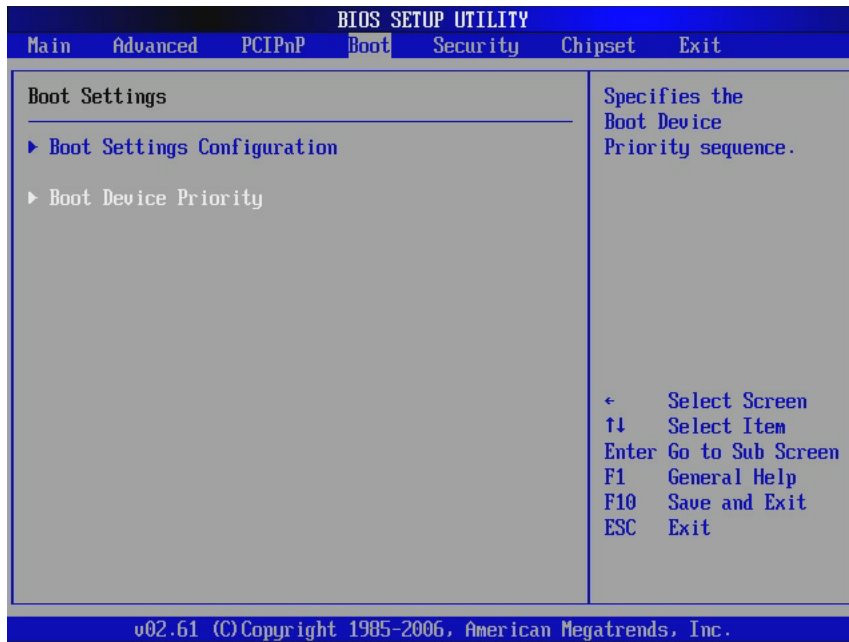
If you plan to use the F8 network booting method, follow the instructions in “To Use the F8 Key to Boot Off the PXE Server” on page 36.

▼ To Set Up the BIOS for Booting From the Virtualized 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.



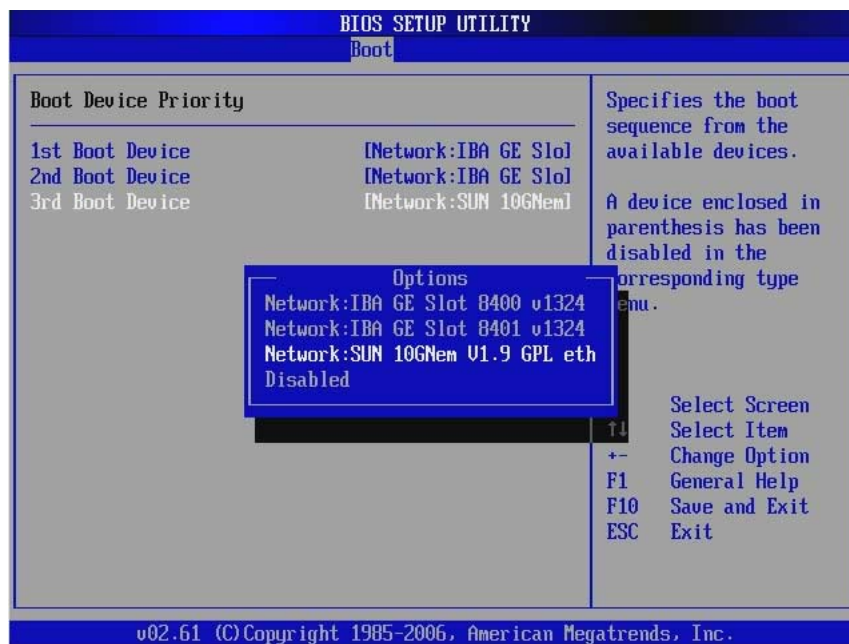
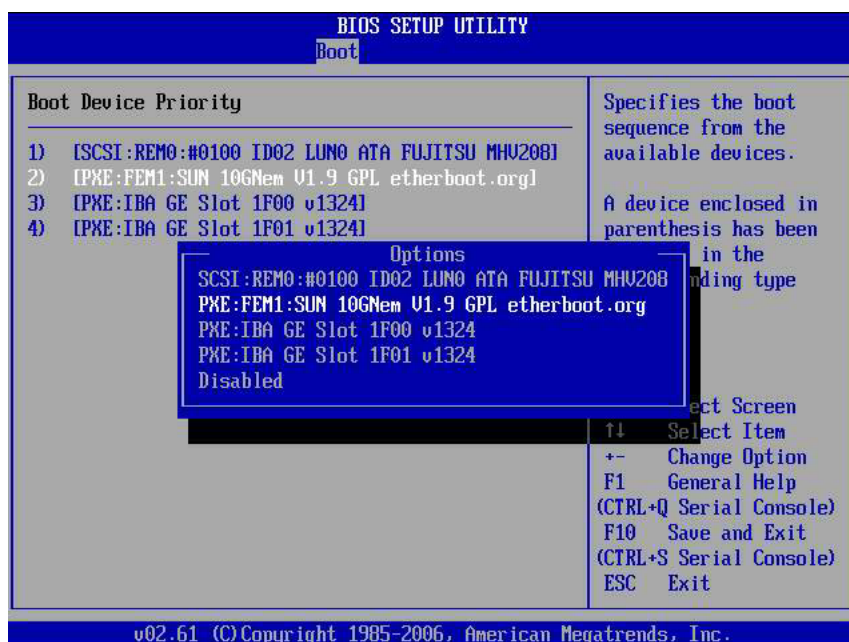


5. Select the first boot device position and presss Enter.

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 plaform.



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

The host blade server reboots.

Note – The Virtualized 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 server from the PXE server to load an OS or install an OS.

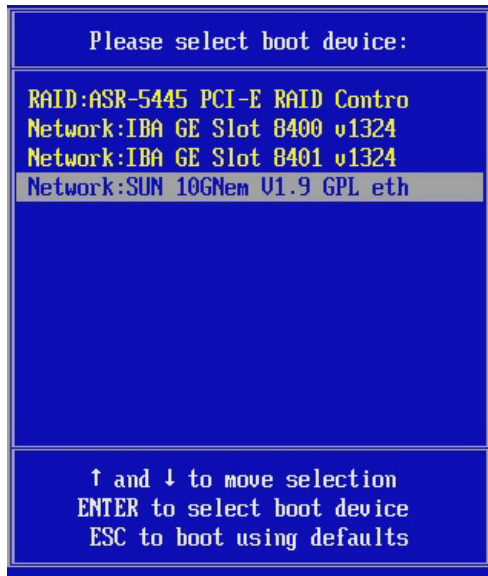
▼ To Use the F12 Key to Initiate the Network Boot

1. Power on the host server blade.
2. Follow the instructions in [“To Set Up the BIOS for Booting From the Virtualized NEM” on page 33](#) to set up the Virtualized NEM as the first network boot device.
3. Press F12 to boot from the network.

Note – If the Virtualized NEM doesn’t detect a PXE Server, the CPU Host Blade will attempt to boot from the on-board 1 Gbps NIC devices.

▼ 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 NEM doesn't detect a PXE Server, the CPU Host Blade will attempt to boot from the on-board 1 Gbps NIC devices.

5. Install and configure the hxge driver. See the driver installation section in this chapter that corresponds with the operating system installed on the blade module.
6. Install and configure the hxge driver.
See the driver installation section in this chapter that corresponds with the operating system installed on the blade module.

Booting over the Network With a SPARC Blade Server

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

▼ 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 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 NEM. For example:

```
{0} ok show-netsa) /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 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 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 netboot.

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

If successful, we should see something as follows 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 5.11 Version snv_116 64-bit  
Copyright 1983-2009 Sun Microsystems, Inc. All rights reserved.  
Use is subject to license terms.  
Configuring /dev  
Using DHCP for network configuration information.  
USB keyboard  
Setting up Java. Please wait...  
Reading ZFS config: done.
```

The net installation menu appears which allows you to which allows you to choose the installation configuration.

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

If your system uses the Solaris SPARC or x86 operating system, you need to download and install the hxge device driver for Solaris platforms.

This section explains how to download, install, and 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:

- [“Installing and Removing the Driver on a Solaris Platform” on page 41](#)
- [“Configuring the Network Host Files” on page 43](#)
- [“Configuring the hxge Device Driver Parameters” on page 44](#)
- [Configuring the Jumbo Frames Feature](#)

Installing and Removing the Driver on a Solaris Platform

▼ To Download and Install the Driver on a Solaris Platform

1. **Locate and copy the `hxge` device driver software from the Tools and Drivers CD for the server module on which you want to install the driver.**

The Solaris drivers are located in one of two directories:

```
/solaris/x64/  
/solaris/sparc/
```

2. **Uncompress the gzipped tar file. For example:**

```
# gunzip hxge.tar.gz
```

3. **Unpack the tar file. For example:**

```
# tar -xvf hxge.tar
```

4. **Install the software packages by typing the following at the command line:**

```
#!/usr/sbin/pkgadd -d .
```

A menu similar to the following displays:

```
The following packages are available:  
  
  1 SUNWhxge SUN 10Gb hxge NIC Driver  
    (i386) 11.10.0,REV=2008.04.24.11.05  
  
Select package(s) you wish to process (or 'all' to process  
all packages). (default: all) [?,??,q]:
```

5. **Select the packages to install.**

- Press Return or type `all` to accept the default and install all packages.
- Type the specific numbers, separated by a space, if you prefer not to install any optional packages.

The following is an example of the displayed output:

```
SUN 10Gb hxge NIC Driver(i386) 11.10.0,REV=2008.04.24.11.05
Copyright 2008 Sun Microsystems, Inc.All rights reserved.
Use is subject to license terms.
Using </> as the package base directory.
## Processing package information.
## Processing system information.
3 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.
This package contains scripts which will be executed with
super-user permission during the process of installing this package
Do you want to continue with the installation of<SUNWhxge>[y,n,?]
```

6. Type **y** to continue the installation.

The following is an example of the displayed output for a successfully installed driver:

```
Installing SUN 10Gb hxge NIC Driver as <SUNWhxge>
## Installing part 1 of 1.
/kernel/drv/amd64/hxge
/kernel/drv/hxge
[ verifying class <none> ]
[ verifying class <renamenew> ]
## Executing postinstall script.
System configuration files modified but hxge driver not loaded or
attached.
Installation of <SUNWhxge> was successful.
```

▼ To Remove the Driver From a Solaris Platform

1. Determine which the driver packages were installed.

```
# pkginfo | grep SUNWhxge
system      SUNWhxge SUN 10Gb hxge NIC Driver
```

2. Remove the driver packages with the **pkgrm** command, followed by the names of an installed packages.

```
# pkgrm SUNWhxge any additional package names
```

Configuring the Network Host Files

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

▼ To Configure the Network Host Files

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.hxgenumber` 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.hxgex` 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.hxgenumber` 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.hxgnumber` file required for a system named `zardoz-c10-b11`.

```
# cat /etc/hostname.hxge0
zardoz-c10-b11
```

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-b11
```

Configuring the `hxge` Device Driver Parameters

The `hxge` device driver controls the Virtualized 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:

- [“To Set Parameters Using the `hxge.conf` File” on page 44](#)
- [“To Specify Parameter Values Using the `ndd` Utility” on page 47](#)

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.

▼ 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.

▼ 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 ffffffff00 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 48](#)
- [“Checking Jumbo Frames Configurations” on page 48](#)
- [“Enabling Jumbo Frames in a Solaris Environment” on page 50](#)

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 NEM driver capability.

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

Checking 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          120834317
rdc_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
crttime                  134.825726986
ierrors                  0
ifspeed                  10000000000
ipackets                  265847787
ipackets64               17445716971
multircv                 0
multixmt                 0
norcvbuf                 0
noxmtbuf                 0
obytes                   1266555560
obytes64                 662691519144
oerrors                  0
opackets                  129680991
opackets64               8719615583
rbytes                   673822498
rbytes64                 24761160283938
snaptime                 122991.23646771
unknowns                  0
```

- Display all driver statistics, for example:

```
# kstat -m hxge
```

Enabling Jumbo Frames in a Solaris Environment

This section describes how to enable Jumbo Frames in a Solaris environment.

▼ 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
```

▼ 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.

▼ 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 52](#)
- [“Configuring the Network Interface” on page 59](#)
- [“Configuring Jumbo Frames” on page 75](#)

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.

▼ 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:E2:BA:34
eth1      Link encap:Ethernet  HWaddr 00:14:4F:E2:BA:35
```

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

2. 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.

The drivers are located in one of the following directories:

```
/linux/drivers/rhel5.1
/linux/drivers/rhel5.2
/linux/drivers/rhel4.7
/linux/drivers/sles10sp2
```

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

3. Uncompress and unpack the driver download file, if necessary.

The following example shows all the files extracted into a `tge10` subdirectory (The command sample below has been edited for brevity. Sections marked {....} denote output that has been removed for this example):

```
host #> mkdir tge10
host #> cd tge10Archive: sun_10_Gigabit_Ethernet_driver_update_10.zip
host #> unzip sun_10_Gigabit_Ethernet_driver_update_10.zip
creating: Linux/
creating: Linux/RHEL4.7/
creating: Linux/RHEL4.7/2.6.9-78.ELlargesmp/
[...]
creating: Linux/RHEL4.7/2.6.9-78.ELsmp/
[...]
creating: Linux/src/
extracting: Linux/src/hxge_src.zip
creating: Linux/RHEL5/
[...]
creating: Linux/SUSE10-SP2/
```

```
[...]
creating: Docs/
inflating: Docs/x8_Express_Dual_10GBE_Low_Profile_Release_Notes.pdf
inflating: Docs/x8_Express_Quad_Gigabit_Ethernet_Express_Module_Release_Notes.p
df
inflating: Docs/x8_Express_Quad_Gigabit_Ethernet_Low_Profile_Release_Notes.pdf
inflating: Docs/x8_Express_Dual_10GBE_Express_Module_Release_Notes.pdf
creating: Firmware/
[...]
creating: Windows/
[...]
host #>
```

4. Select and install the appropriate OS driver package.

- a. Navigate to the Linux directory that contains the driver-specific files. For example:

```
host #> ls -l
total 47908
drwxr-xr-x  2 root root    4096 Sep 25 16:31 Docs
drwxr-xr-x  2 root root    4096 Sep 25 18:46 Firmware
drwxr-xr-x 12 root root    4096 Sep 25 16:31 Linux
-rw-r--r--  1 root root 48984046 Oct  7 12:13
sun_10_Gigabit_Ethernet_driver_update_10.zip
drwxr-xr-x  2 root root    4096 Sep 25 16:31 Windows

host #> cd Linux
host #> ls -ldrwxr-xr-x 4 root root 4096 Sep 25 16:30 RHEL4.7
total 40
drwxr-xr-x 2 root root 4096 Sep 25 16:30 RHEL5.2
drwxr-xr-x 2 root root 4096 Sep 25 16:30 src
drwxr-xr-x 3 root root 4096 Sep 25 16:31 SUSE10-SP2
```

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
Distributor ID: RedHatEnterpriseServer
Description:    Red Hat Enterprise Linux Server release 5.2 (Tikanga)
Release:        5.2
Codename:       Tikanga
```

- b. Identify the OS-specific subdirectory and verify that no `hxge` driver is currently installed; for example: RHEL5.2 for RedHat Enterprise Linux 5 Update 2, or SUSE10-SP2 for Novell's SuSE Linux Enterprise Server 10 Service Pack 2

The following example uses RHEL5.2:

```
host #> cd RHEL5.2
host #> ls -l-rw-r--r-- 1 root root 2752340 Oct  7 12:35
hxge-1.1.1_rhel52-1.x86_64.rpm
total 2692
host #> rpm -q hxge
package hxge is not installed
```

Note – If an `hxge` driver is already installed, uninstall the driver, in order to avoid complications. See [“To Remove the Driver From a Linux Platform” on page 58](#) 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-1.1.1_rhel52-1.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 make 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
[...]
07:00.0 Ethernet controller: Sun Microsystems Computer Corp. Unknown device abcd
(rev 01)
[...]
84:00.0 Ethernet controller: Intel Corporation 82575EB Gigabit Network
Connection (rev 02)
[...]
85:00.0 Ethernet controller: Sun Microsystems Computer Corp. Unknown device aaaa
(rev 01)
```

Device code `0xAAAA` (Unknown device `aaaa (rev 01)` output) is the Virtualized 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                168784  0
host #>modinfo hxgefilename:
/lib/modules/2.6.18-92.el5/kernel/drivers/net/hxge.ko
version:            1.1.1
license:            GPL
description:        Sun Microsystems(R) 10 Gigabit Network Driver
author:             Sun Microsystems, <james.puthukattukaran@sun.com>
srcversion:         B61926D0661E6A268265A9C
alias:              pci:v0000108Ed0000AAAsv*sd*bc*sc*i*
depends:
[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:E2:BA:34
eth1      Link encap:Ethernet  HWaddr 00:14:4F:E2:BA:35
eth2      Link encap:Ethernet  HWaddr 00:14:4F:62:01:08
```

In this example, `eth0` and `eth1` were previously present; `eth2` is new, this is the NEM Ethernet network interface device. If you had installed two NEMs, you would also see an `eth3` device representing the other NEM. You can

identify each eth device (NEM0 or NEM1) by matching the Ethernet MAC address with the one you recorded and saved when you physically installed the NEM into the chassis in [Chapter 2](#).

- b. Make sure that the eth2 driver is the correct Ethernet driver for the Virtualized NEM.**

```
host #> ethtool -i eth2
driver: hxge
version: 1.1.1
firmware-version: N/A
bus-info: 0000:85:00.0
```

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

```
host #>ifconfig eth2
eth2      Link encap:Ethernet  HWaddr 00:14:4F:62:01:08
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:b7000000-b8000000
```

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.

▼ 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
host #> lsmod | grep hxge
host #>
```

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-1.1.1_rhel52-1
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 being able to use 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 NEM Hydra as demonstrated in the installation section). At a minimum, you must assign a network (TCP) IP address and netmask to each network interface.

Temporary `hxge` Network Interface Configuration

To temporarily configure the NEM Hydra 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.

▼ To Bring the Interface Online Manually

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

```
host #>ifconfig eth2 10.1.10.150 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:29:00:01
       inet addr:10.1.10.150  Bcast:10.1.10.255  Mask:255.255.255.0
       inet6 addr: fe80::214:4fff:fe29:1/64  Scope:Link
       UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
       RX packets: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 `10.1.10.150`, 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 `10.1.10.150` 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 NEM Hydra `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:10.1.10.150  Bcast:10.1.10.255  Mask:255.255.255.0
eth2  Link encap:Ethernet  HWaddr 00:14:4F:29:00:01
      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).

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.

Permanently Configure the Network Interface for the Red Hat Linux Platform

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

- For instructions on configuring the network interface using the GUI, refer to the documentation for your RHEL version at: (<https://www.redhat.com/docs/>)
- For instructions on configuring the network interface manually, see “To Configure the Network Interface File Automatically” on page 62.

▼ To Configure the Network Interface File Automatically

For RedHat 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:29:00:00
IPADDR=10.1.10.150
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
```

Permanently Configure the Network Interface for a SUSE Platform

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

- For instructions on configuring the network interface using the GUI, refer to the documentation for your SLES version at:
(<http://www.novell.com/documentation/suse.html>)

- For instructions on configuring the network interface manually, see [“To Configure the Network Interface Automatically”](#) on page 64.

▼ To Configure the Network Interface Automatically

For Novell systems, the interface configuration files are named `ifcfg-eth-id` (for example, `ifcfg-eth-id-00:14:4F:29:00:01` 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:29:00:01
-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:29:00:01
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.helwLcVzebD'
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.

- `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:29:00:01
      inet addr:10.1.10.150 Bcast:10.1.10.255 Mask:255.255.255.0
      inet6 addr: fe80::214:4fff:fe29:1/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:2993 errors:0 dropped:0 overruns:0 frame:0
      TX packets:2978 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:3286970 (3.1 MiB)  TX bytes:3287849 (3.1 MiB)
      Memory:fb000000-fc000000
```

Note that no 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/2.6.18-92.el5/kernel/drivers/net/hxge.ko
version:      1.1.1
license:      GPL
description:   Sun Microsystems(R) 10 Gigabit Network Driver
author:       Sun Microsystems, <james.puthukattukaran@sun.com>
srcversion:   B61926D0661E6A268265A9C
alias:        pci:v0000108Ed0000AAAAsv*sd*bc*sc*i*
depends:
```

```

vermagic:      2.6.18-92.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:          rbr_entries:No. of RBR Entries (int)
parm:          rcr_entries:No. of RCR Entries (int)
[...]
parm:          tcam_udp_ipv6:UDP over IPv6 class (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:

- [“To Temporarily Configure the `hxge` Driver Parameters” on page 68](#)
- [“To Permanently Configure the `hxge` Driver Parameters” on page 69](#)

▼ 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 #>lsmmod | grep hxge
hxge          148824  0

```

In this example, the driver is loaded.

2. Unload the currently active driver.

```

host #>modprobe -r hxge

```

You can also use the `rmmmod(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
```

▼ 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 e1000e
alias eth1 e1000e
alias scsi_hostadapter ata_piix
options hxge tcam=0
```

hxge Driver Configuration Parameters

The actual list of hxge driver parameters is subject to change from release to release. [TABLE 3-1](#) 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 3-1 Driver Configuration Parameters

Parameter	Description	Values	Default
enable_jumbo	Controls hxge driver runtime support of jumbo frames. hxge jumbo frame support is automatically enabled as needed (depending on the network interface-specified MTU value).	0 = No 1 = Yes	Automatic
intr_type	Controls what sort of interrupt mechanism (if any) is selected by the hxge driver. The hxge 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").
rbr_entries	Specifies how many 4KB receive buffers the hxge driver will allocate per receive channel (the NEM supports four parallel independent receive channels).		4096
rcr_entries	Specifies how many receive pointers (effectively, packets; jumbo packets can require up to 3 RCR entries per single jumbo packet) the hxge driver will allocate per receive channel.		8192
rcr_timeout	Magic internal unitless number. Do not change this number unless told to by a qualified Sun agent.		
rcr_threshold	Magic internal unitless number. Do not change this number unless told to by a qualified Sun agent.		
rx_dma_channels	Specifies how many receive DMA channels the hxge 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
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

TABLE 3-1 Driver Configuration Parameters (*Continued*)

Parameter	Description	Values	Default
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 72 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
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

TABLE 3-1 Driver Configuration Parameters (*Continued*)

Parameter	Description	Values	Default
<code>tcam</code>	Controls whether or not the <code>hxge</code> driver enables the Virtualized 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 <code>rx_dma_channels</code> in this table.	0 = disable 1 = enable	1
<code>tcam_seed</code>	Magic internal unitless number. Do not change this number unless told to by a qualified Sun agent.		
<code>tcam_tcp_ipv4</code>	Controls whether or not the <code>hxge</code> 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 72](#)
- [“Setting the Debug Messaging Parameter” on page 73](#)
- [“To Configure the Syslog Parameter” on page 74](#)

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 1.1.1
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.

Setting the Debug Messaging Parameter

The Linux hxge driver's message logging can be statically specified via the debug driver configuration parameter. See [“Changing the hxge Driver Configuration” on page 67](#). 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.

▼ To Set the Debug Drive Parameter Dynamically

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

```
ethtool -s ethn msglvl parameter value
```

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 [“To Configure the Syslog Parameter” on page 74](#).

▼ 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 (excerpted from an RHEL5.3 /etc/syslog.conf file).

```
# Don't log private authentication messages!

# Log anything (except mail) of level info or higher.
*.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 transition 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.

▼ 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:29:00:01
          inet addr:10.1.10.150  Bcast:10.1.10.255  Mask:255.255.255.0
          inet6 addr: fe80::214:4fff:fe29:1/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets: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:10.1.10.150  Bcast:10.1.10.255  Mask:255.255.255.0
          inet6 addr: fe80::214:4fff:fe29:1/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU: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.

▼ 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:29:00:00
IPADDR=10.1.10.150
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 `hgxe` device is as specified.

```
host #>ifconfig eth2
eth2      Link encap:Ethernet  HWaddr 00:14:4F:29:00:01
          inet6 addr: fe80::214:4fff:fe29:1/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:9124  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:12 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 b)  TX bytes:2670 (2.6 KiB)
          Memory:fb000000-fc000000
```

Installing and Configuring Drivers on a Windows Platform

The following topics are covered in this section:

- [“Installing Drivers on a Windows Platform” on page 78](#)

- [“Enabling Jumbo Frames” on page 104](#)

Installing Drivers on a Windows Platform

This section describes the process for installing network, VLAN, and enclosure drivers on an x64 (Intel or AMD) server module that supports Windows Server 2003 (32/64-bit) or Windows Server 2008 (32/64-bit).

Note – The installation process shown in this section applies to all Windows Server 2003 and Windows Server 2008 server modules. The section on installing the enclosure device is separated into two procedures, one for each operating system.

The following sections describe the procedures for installing Windows drivers:

- [“Installing and Uninstalling the Sun Blade 6000 10GbE Network Controller” on page 78](#)
- [“Installing and Uninstalling the Sun Blade 6000 10GbE VLAN Driver \(Optional\)” on page 84](#)
- [“Installing the Virtualized Multi-Fabric 10GbE NEM Enclosure Device” on page 93](#)

Installing and Uninstalling the Sun Blade 6000 10GbE Network Controller

▼ 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/w2k3/Sun_Blade_6000_10Gbe_Networking_Controller.msi`

Note – This installer installs the networking drivers for both Windows Server 2003 32-bit/64-bit and Windows Server 2008 32-bit/64-bit.

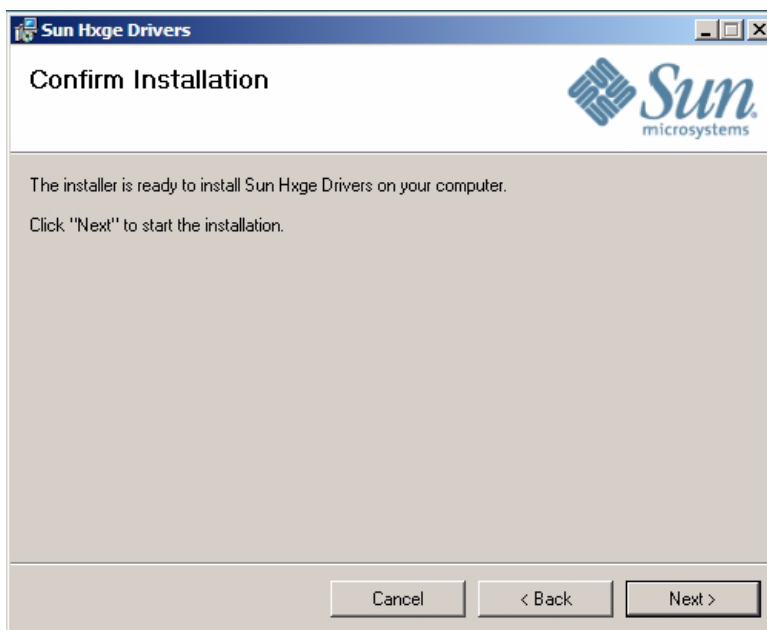
2. **Navigate to the `Sun_Blade_6000_10Gbe_Networking_Controller.msi` file on your local or remote system and double-click it to start the installation.**

The License Agreement page displays.



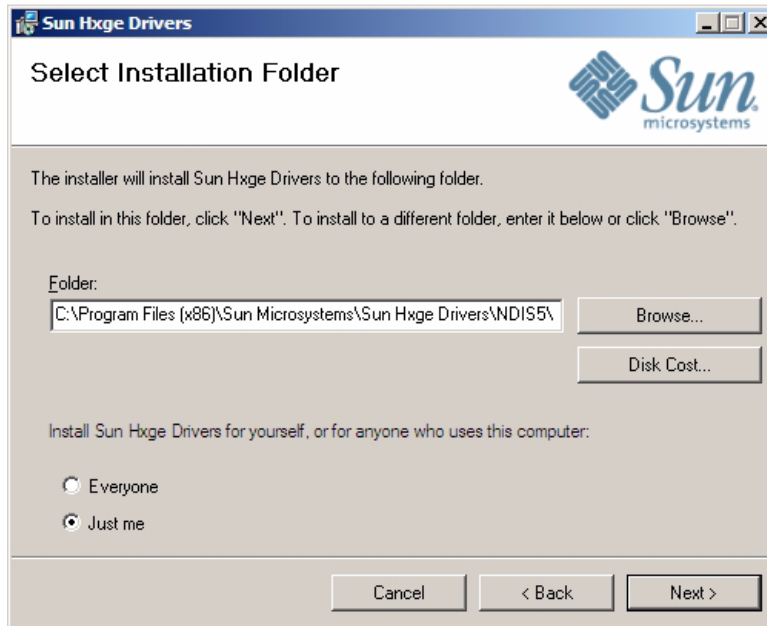
3. Select "I Agree", then click Next to begin the installation.

The Confirm Installation page displays.



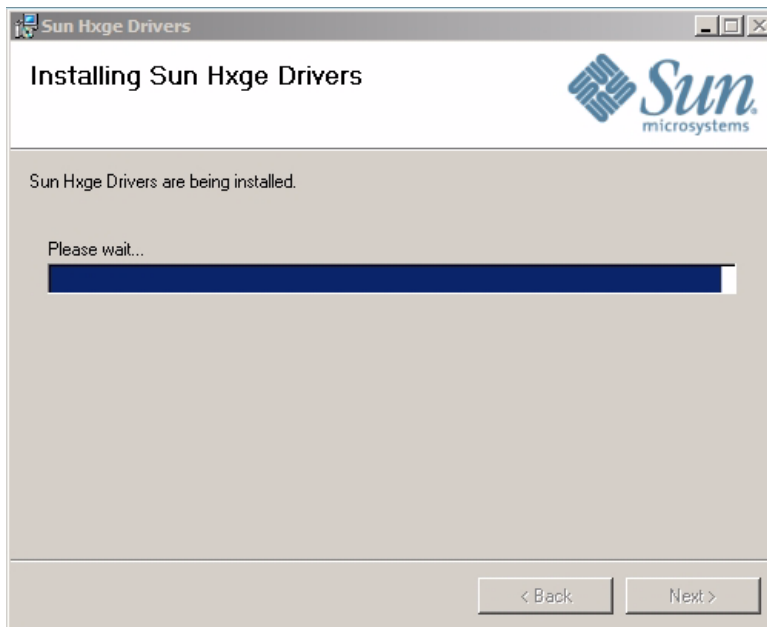
4. Click Next.

The Select Installation Folder page displays.

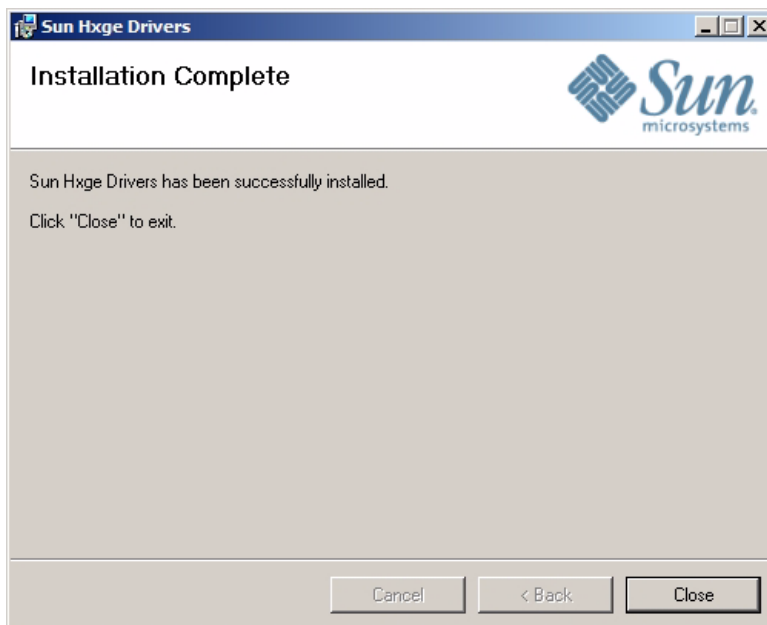


5. Select "Everyone" or "Just Me" and click Next.

The Installing Sun Hxge Drivers page displays.



When the installation finishes, the Installation Complete page displays.



6. Click Close.

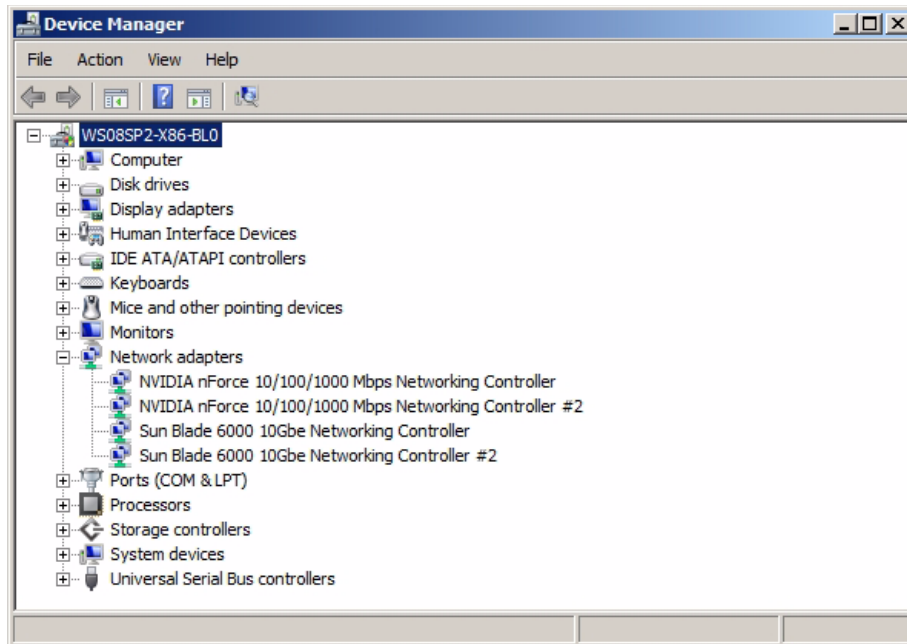
The Sun Blade 6000 10Gbe Networking Controller device is installed and visible in the Network Adapters section of the Windows Device Manager.

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 NEM installed, and two will be displayed for two Virtualized NEMs.

d. Exit the Windows Device Manager.

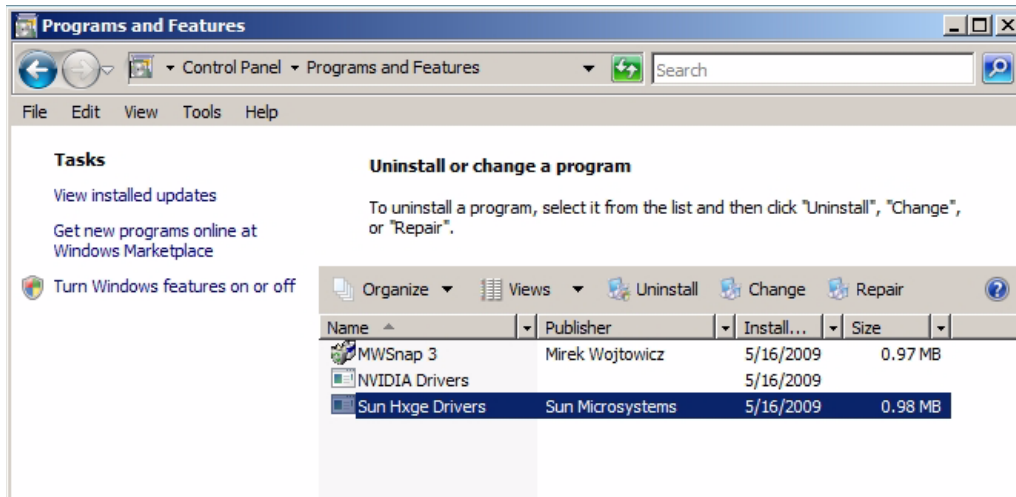
▼ To Uninstall the Sun Blade 6000 10Gbe Networking Controller

1. Before uninstalling the networking controller, remove any VLANs that were installed on the controller. See ["To Remove a VLAN"](#) on page 92.

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 [“To Remove a VLAN”](#) on page 92).
- c. Return to [Step 2](#).

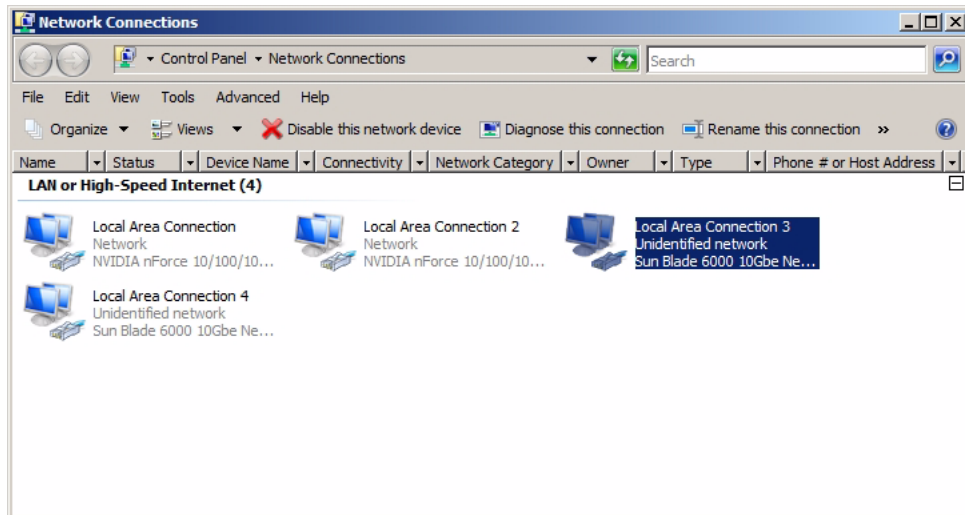
Installing and Uninstalling the Sun Blade 6000 10GbE VLAN Driver (Optional)

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.

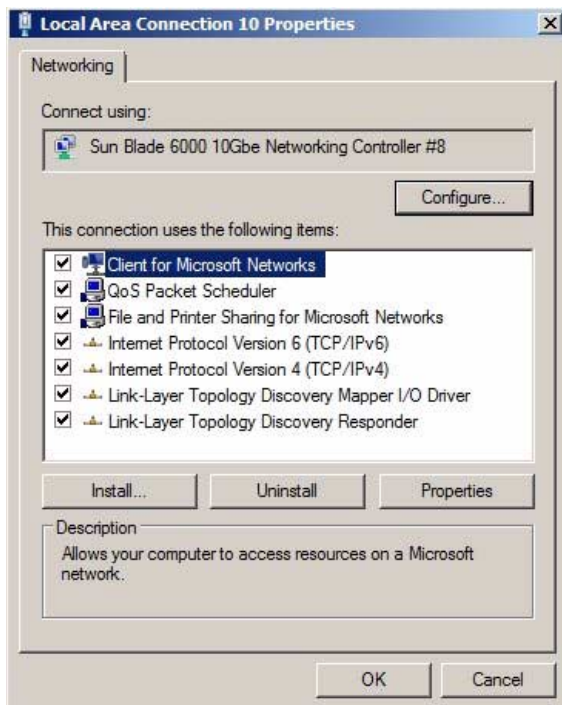
▼ To Install the VLAN Driver

1. **Install the Sun Blade 6000 10GbE Networking Controller.** See [“Installing and Uninstalling the Sun Blade 6000 10GbE Network Controller”](#) on page 78.
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.

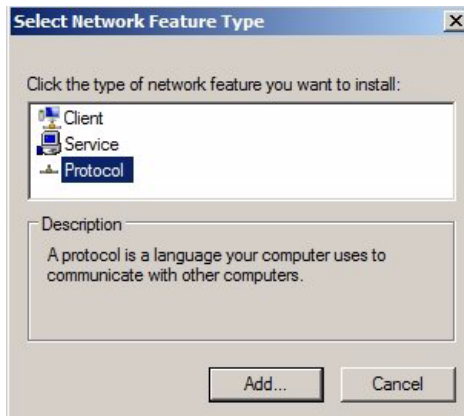


4. **The Local Area Connection Properties dialog displays.**



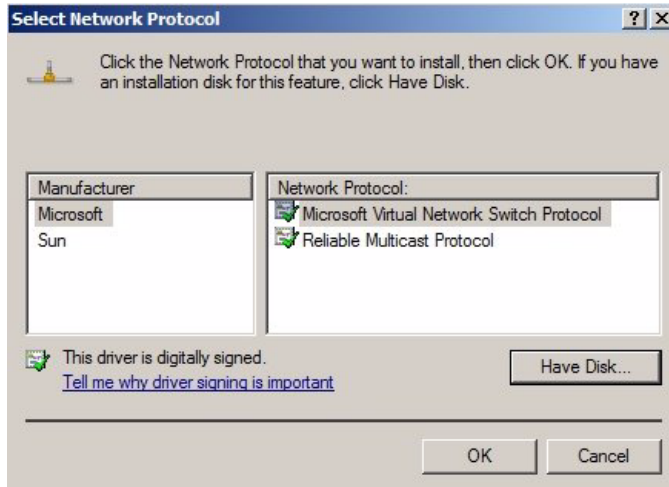
5. Click Install.

The Network Feature Type dialog displays.



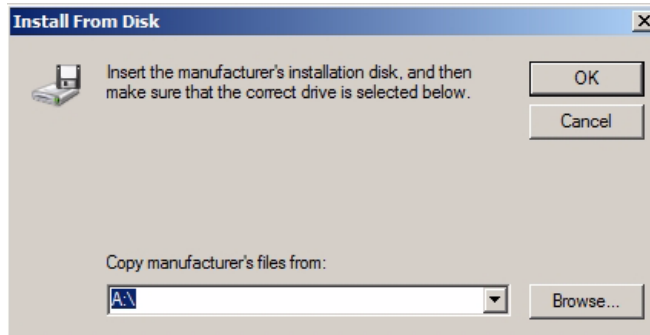
6. Select Protocol and click Add.

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:

■ **Windows Server 2003:**

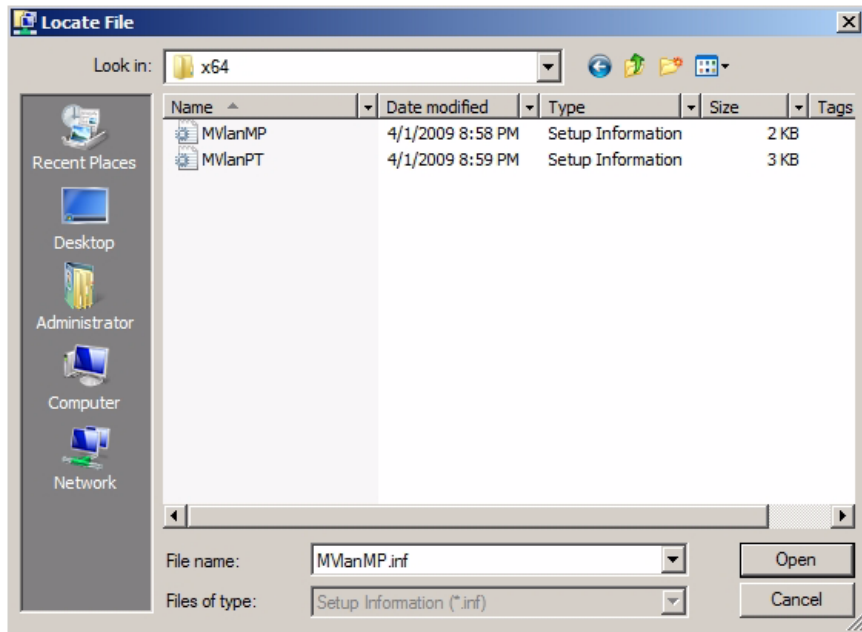
`C:\Program Files(x86)\Sun Microsystems\Sun Hxge Drivers\NDIS5\vlan\w2k3\{x86,x64}`

■ **Windows Server 2008:**

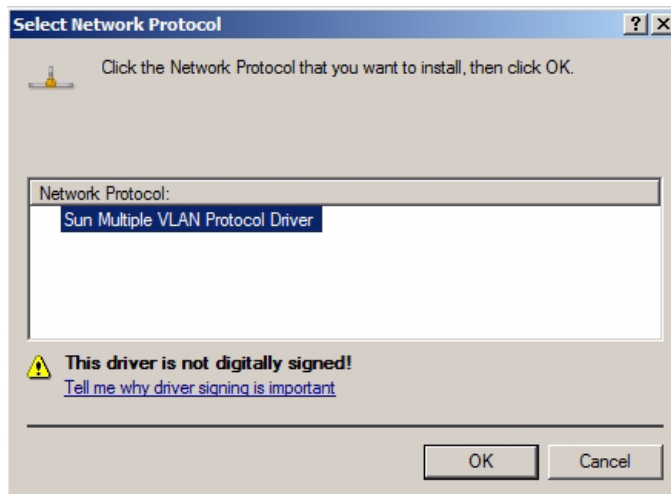
`C:\Program Files(x86)\Sun Microsystems\Sun Hxge Drivers\NDIS6\vlan\w2k8\{x86,x64}`

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.



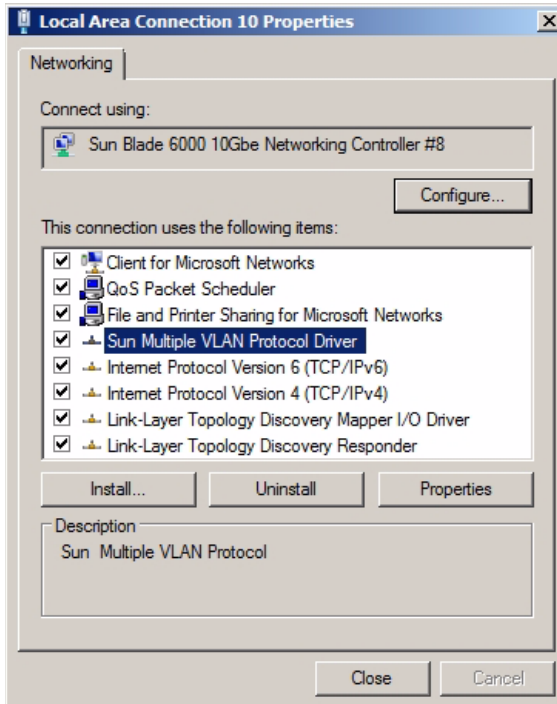
9. Select either of the information setup files displayed and click **Open**.
The Select Network Protocol dialog displays.



10. Click OK.

The Sun Blade 6000 10GbE VLAN driver installs.

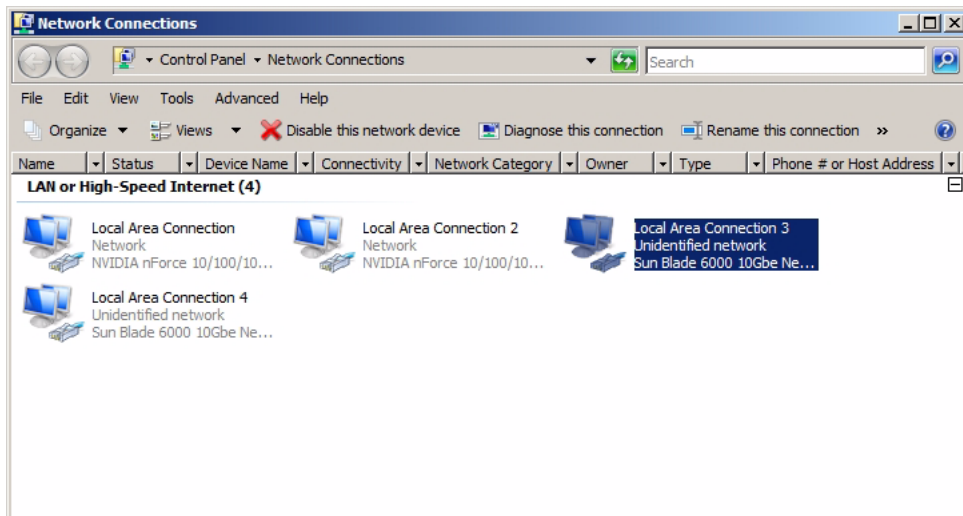
After the installation, the Sun Multiple VLAN Protocol Driver is viewable in the Local Area Connection Properties dialog for each Sun Blade 6000 10GbE Network Controller installed on the system.



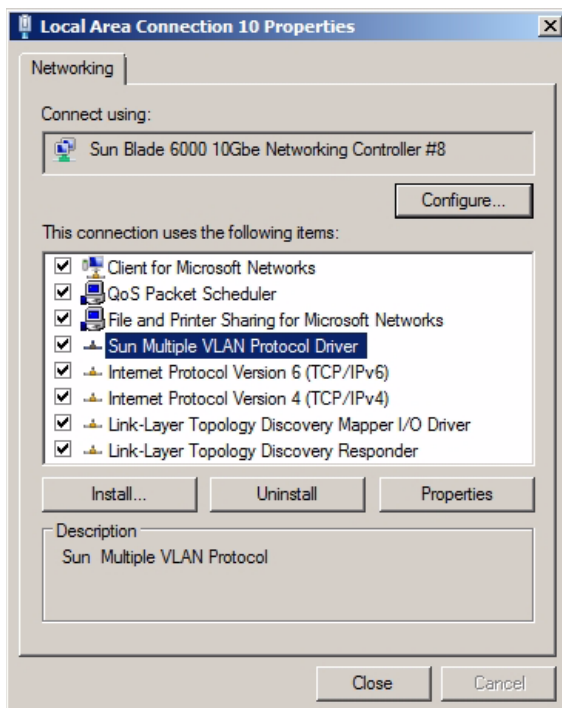
▼ **To Uninstall the VLAN Driver**

- 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.

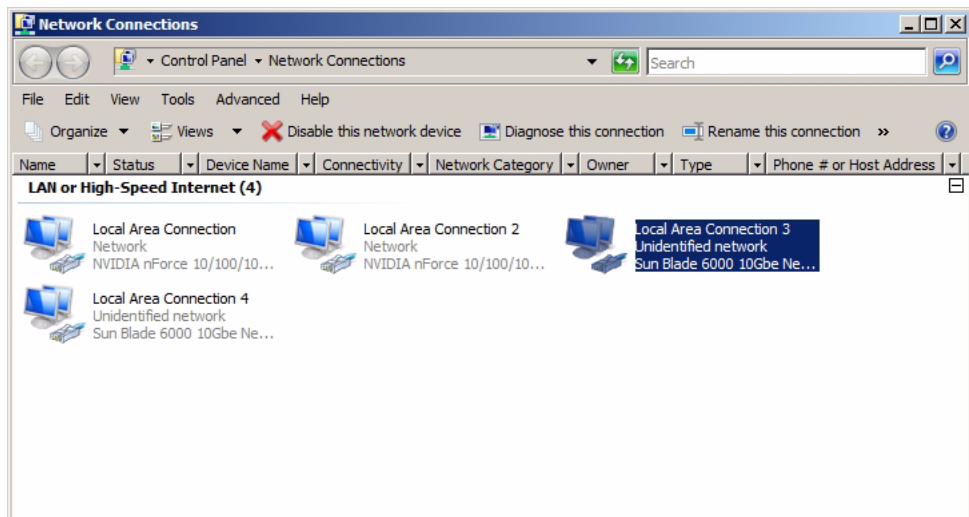
Adding and Removing VLANs

▼ 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.

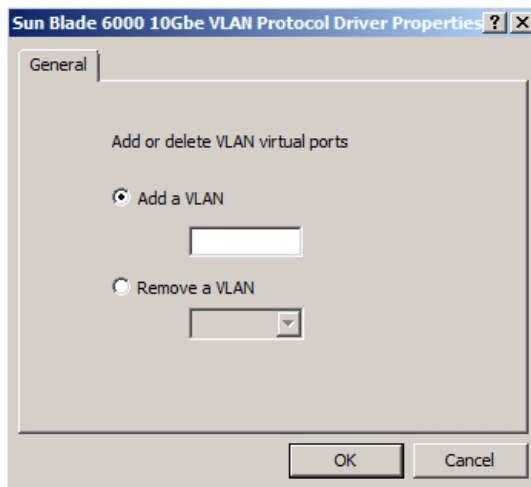
Note – Note which controller port you use to create the VLAN. You will need to use the same controller port to remove the VLAN.

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.

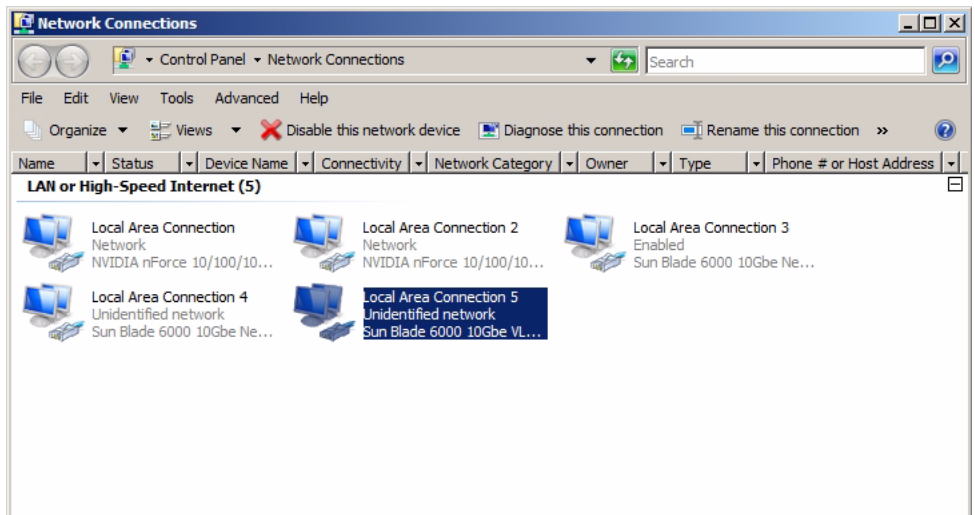


4. Select **Add a VLAN** and enter the VLAN number.

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.



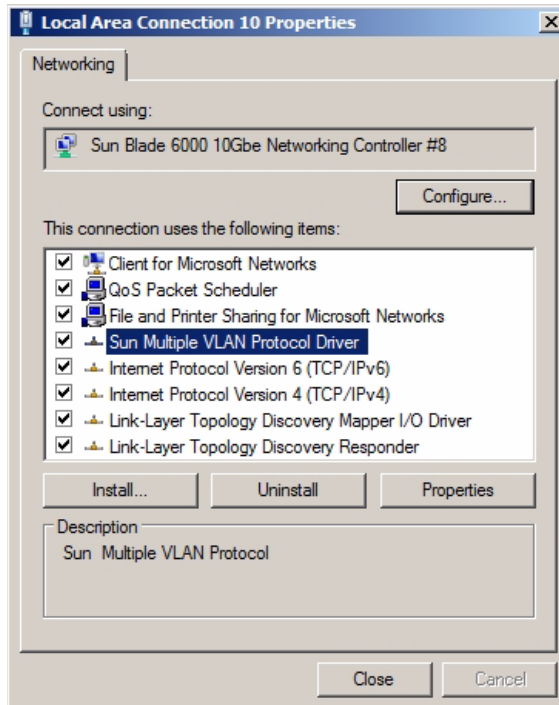
▼ To Remove a VLAN

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 [Step 2](#) of the "To Add a VLAN" procedure, the network controller port used to create the VLAN is Local Area Connection 3 controller port.

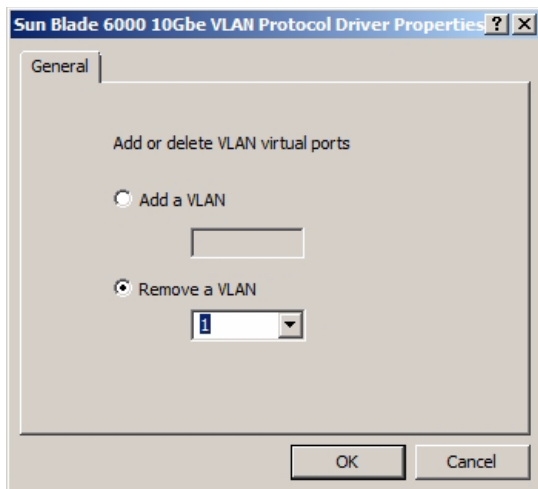
3. Select Properties.

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.

Installing the Virtualized Multi-Fabric 10GbE NEM Enclosure Device

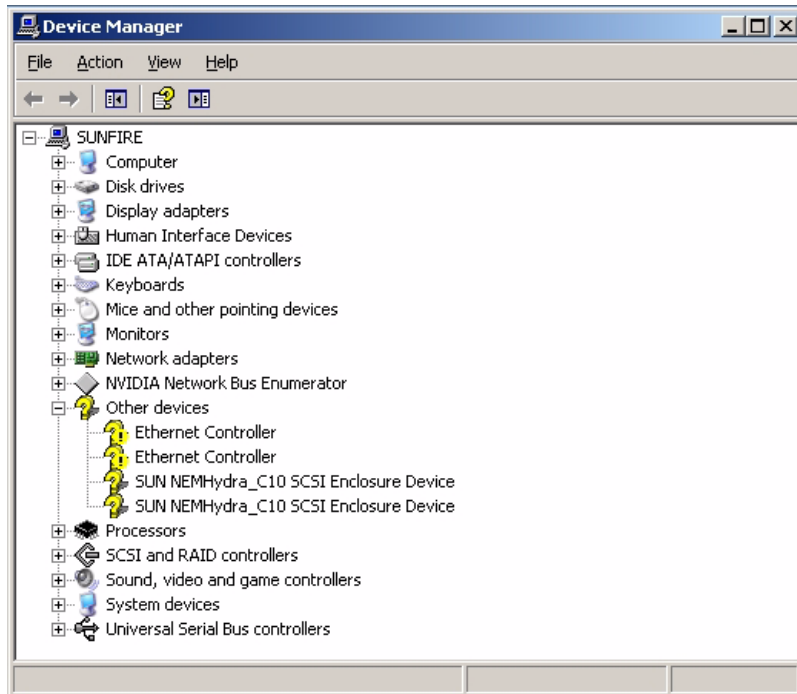
Use one of the following procedures to install the NEM enclosure device.

- [“To Install the Enclosure Device on a Windows Server 2003 System” on page 93](#)
- [“To Install the Enclosure Device on a Windows Server 2008 System” on page 99](#)

▼ To Install the Enclosure Device on a Windows Server 2003 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 Other devices and find the Sun NEMHydra C10 SCSI Enclosure Device.

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

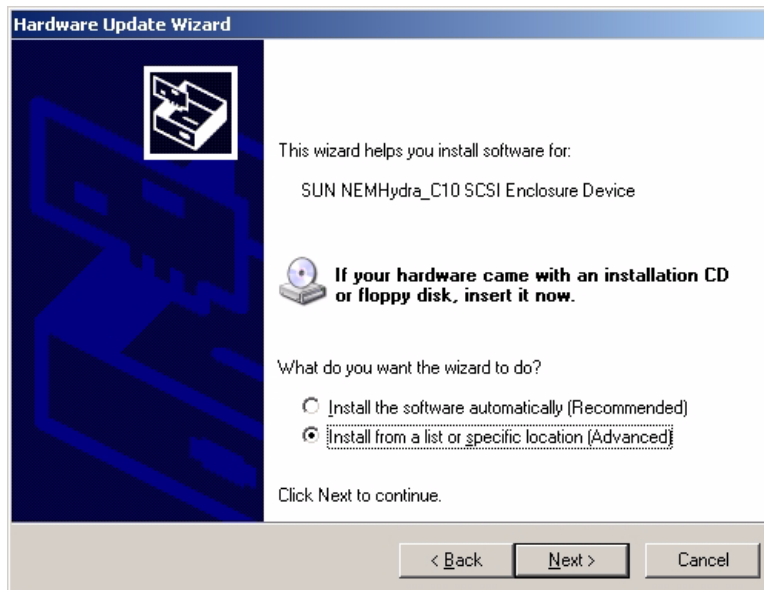
3. Right-click on the Sun NEMHydra_C10 SCSI Enclosure Device, and select Update driver.

The Hardware Update Wizard displays.



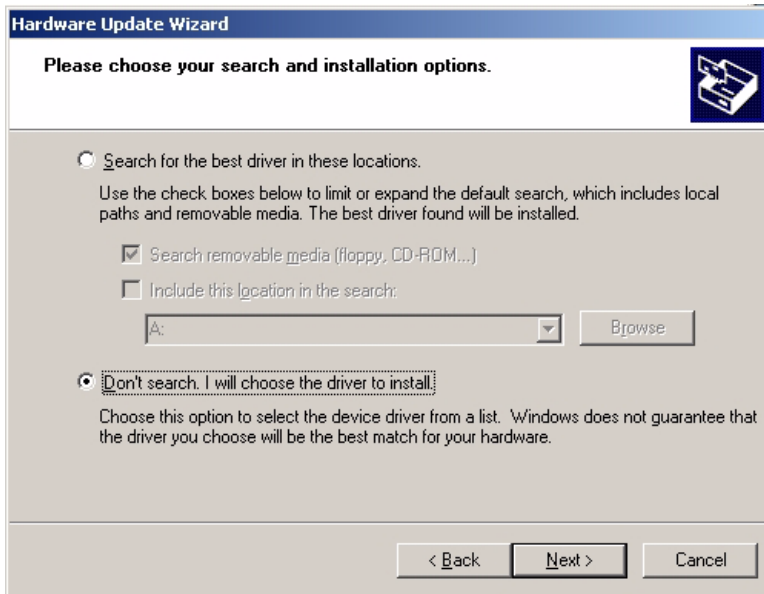
4. Select “No, not this time,” then click Next.

The Installation choices page displays.

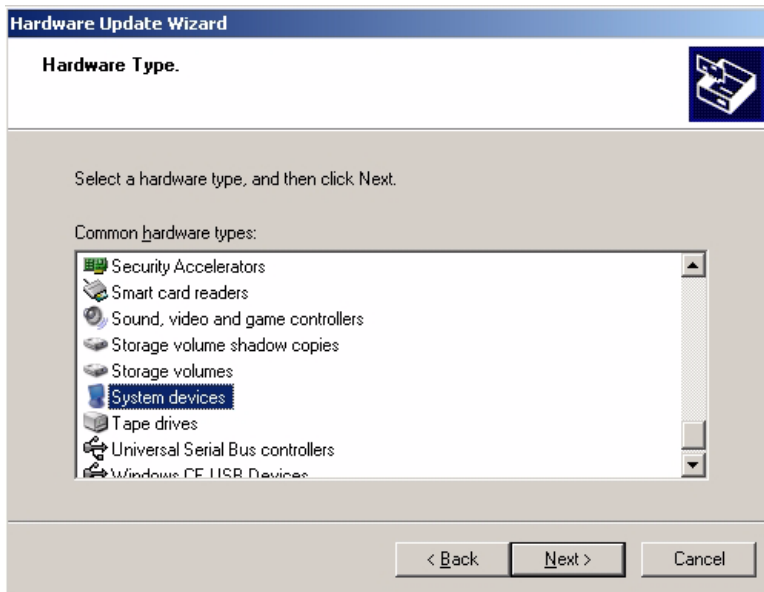


5. Select “Install from a list or specific location (Advanced),” then click Next.

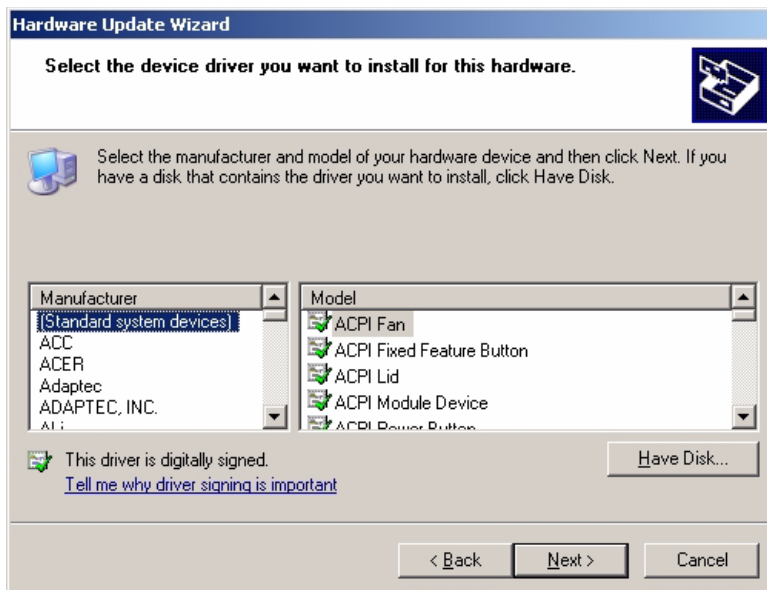
The search and installation options page displays.



6. Select "Don't search. I will choose the driver to install," then click Next.
The Hardware Type page displays.

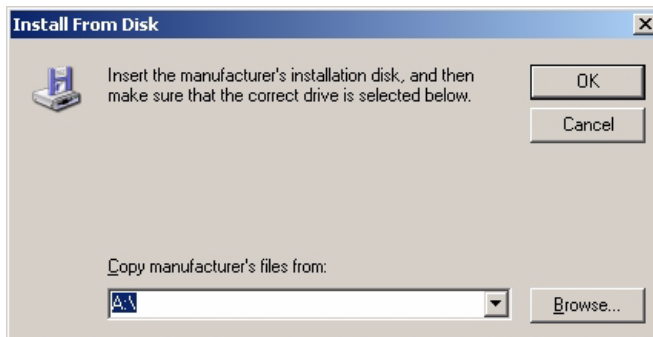


7. Select System devices, then click Next.
The Select the device driver page displays.



8. Click Have Disk.

The Install From Disk page displays.

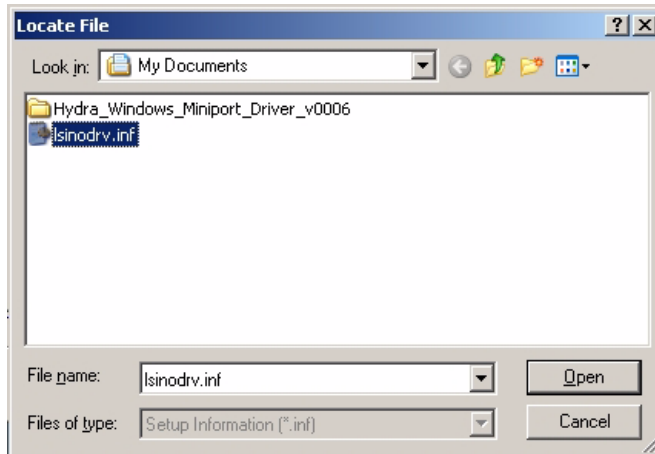


9. Click Browse and navigate to the directory that contains the Virtualized Multi-Fabric 10GbE NEM Enclosure Device information (lsinodrv.inf) file.

The path on the Tools and Drivers CD is:


DVDdrive:\windows\w2k3\{32-bit, 64-bit}\lsinodrv.inf

The Locate File dialog box displays.



10. Select the `lisnodrv.inf` file, then click **Open.**

The Software Installing page displays as the device is installed, then the Completing the Hardware Update Wizard page displays.



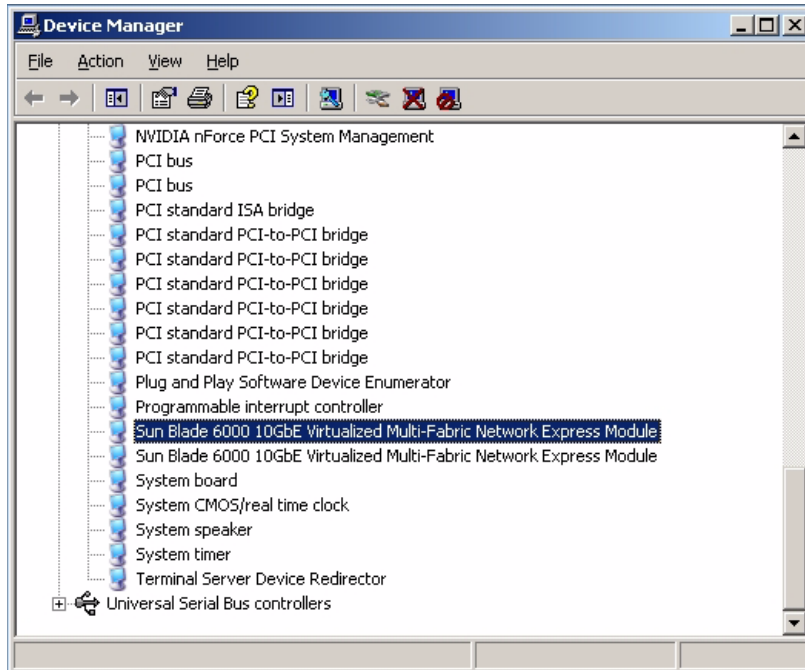
11. Click **Finish to complete the installation.**

12. 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 an additional enclosure device.

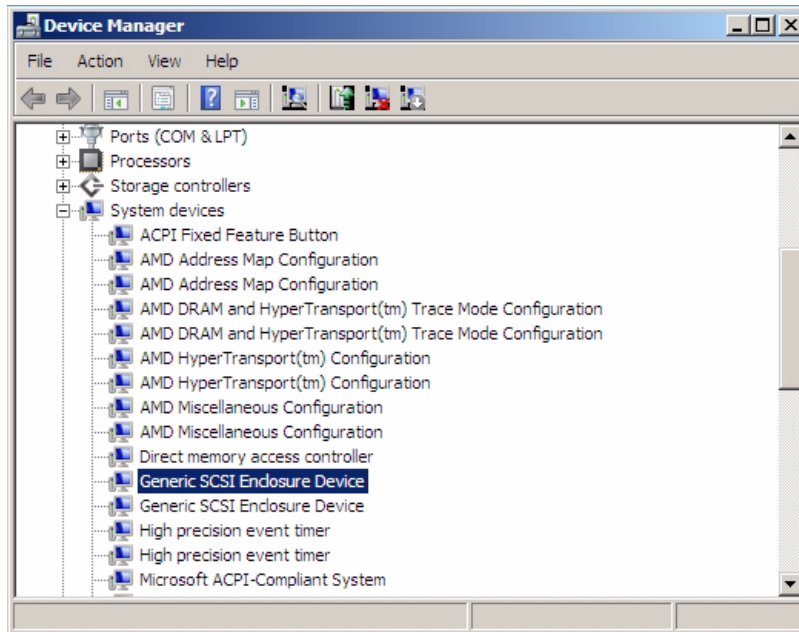
▼ To Install the Enclosure Device on a Windows Server 2008 System

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.

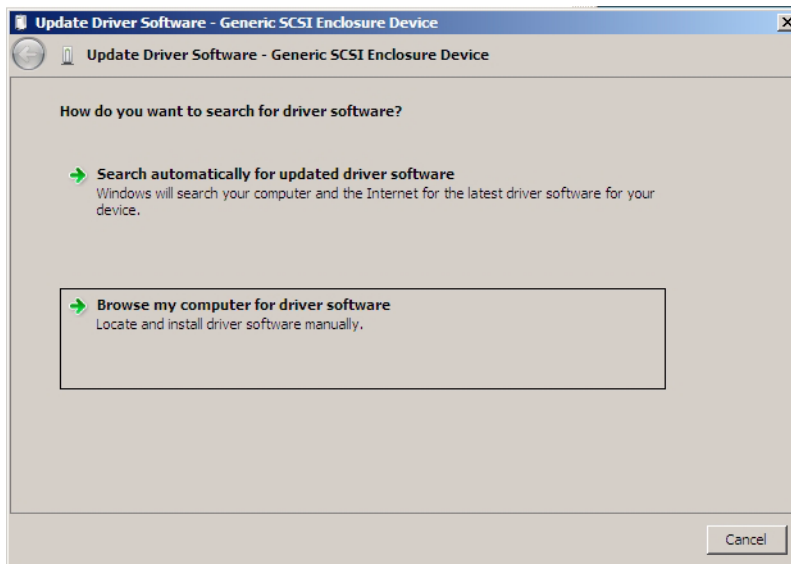


2. Click on System Devices and find the Generic SCSI Enclosure Device.

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

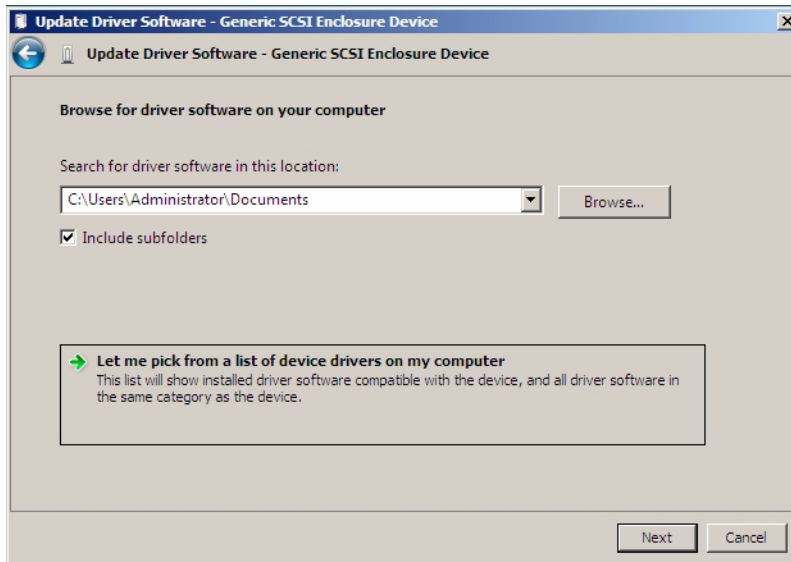
3. Right-click on the Generic SCSI Enclosure Device, and select Update driver.

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



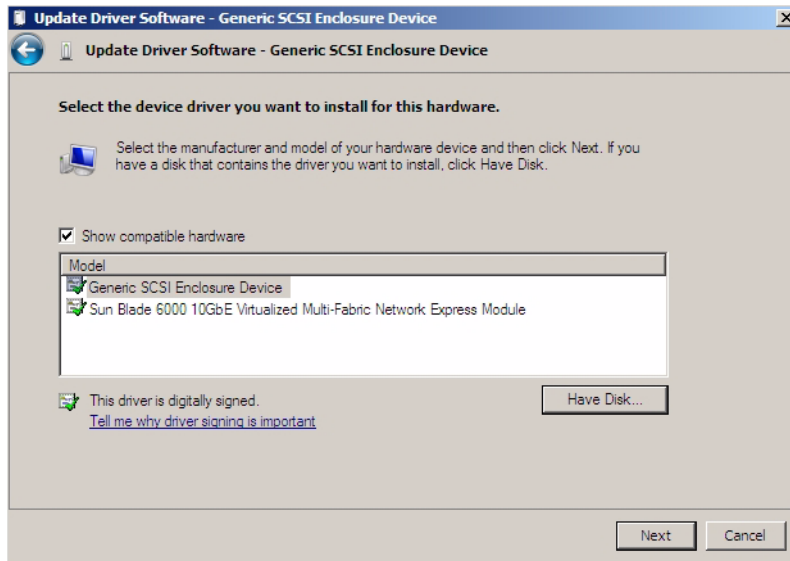
4. Click **“Browse my computer for driver software”**.

The Browse for driver software on your computer page displays.



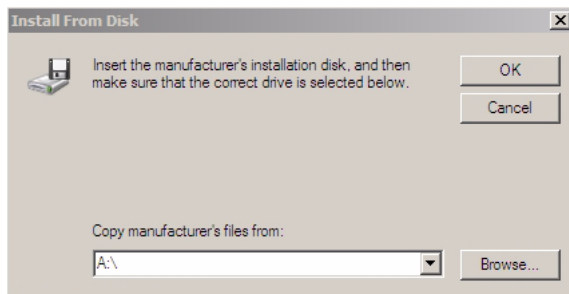
5. Click **“Let me pick from a list of device drivers on my computer”**, and click **Next**.

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



6. Click Have Disk.

The Install From Disk page displays.

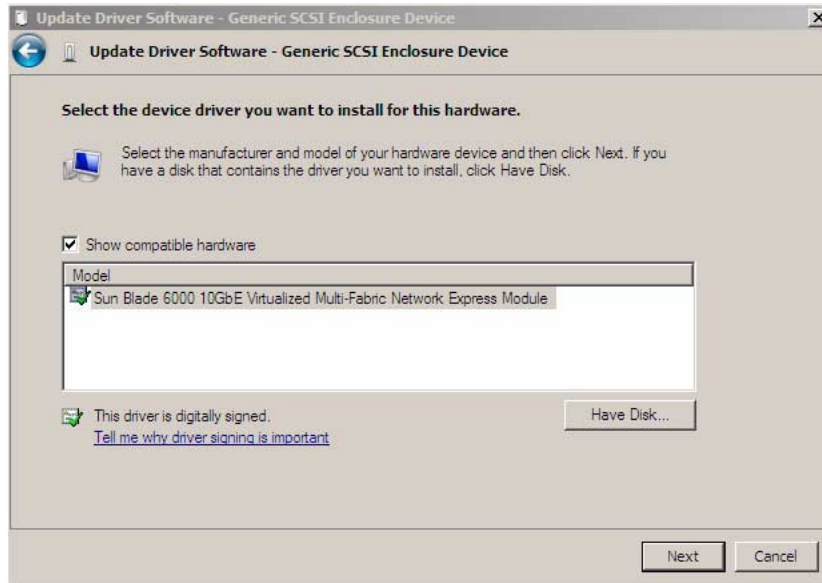


7. Click Browse and navigate to the directory that contains the Sun Blade 6000 10GbE Virtualized Multi-Fabric Network Express Module information (lsinodrv.inf) file.

The path on the Tools and Drivers CD is:

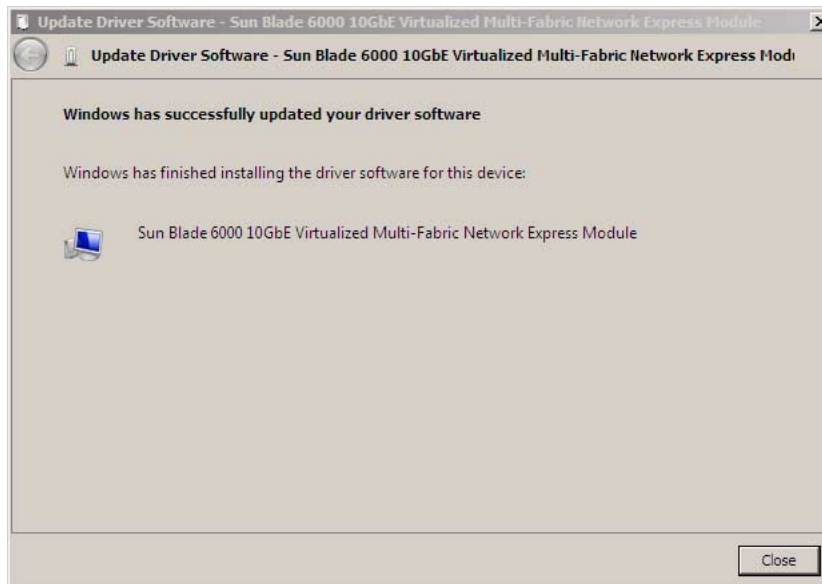
DVDdrive:\windows\w2k8\{32-bit,64-bit}\lsinodrv.inf

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.



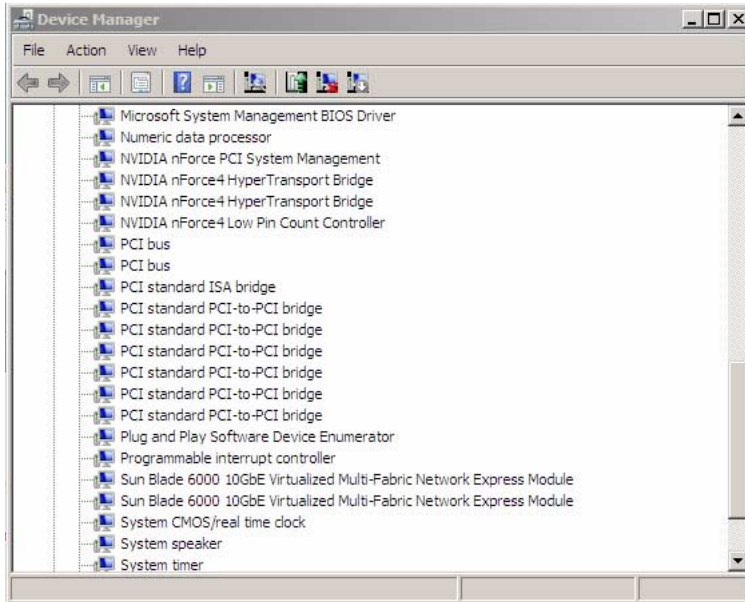
9. Click Close to complete the installation.

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.

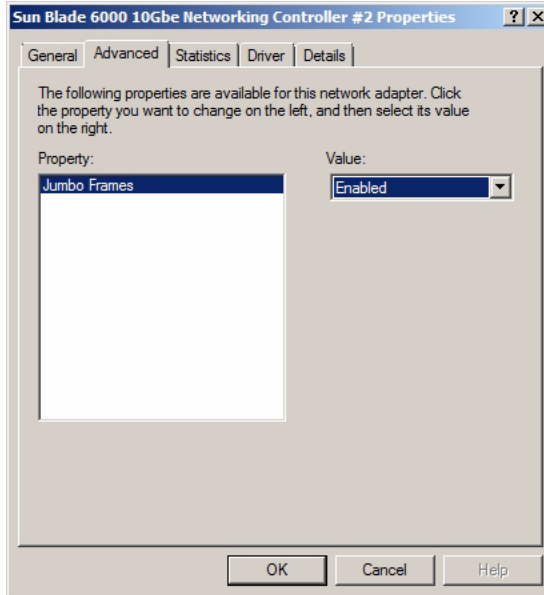
Enabling Jumbo Frames

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.

▼ 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 106
- “Configuring the Virtual NEM Network Adapters” on page 111
- “Configuring Jumbo Frames” on page 112

Installing the ESX Server Drivers on an Existing ESX Server

This section describes the procedures for installing the VMware ESX Server drivers on a server module with ESX Server already installed. This section contains the following procedures:

- “To Install a Virtualized NEM Driver on an Existing ESX 3.5 Server” on page 106
- “To Install a Virtualized NEM Driver on an Existing ESX/ESXi 4.0 Server Using vihostupdate” on page 107
- “To Install a Virtualized NEM Driver on an Existing ESX 4.0 Server Using esxupdate” on page 108

▼ To Install a Virtualized NEM Driver on an Existing ESX 3.5 Server

1. Obtain the hxge driver from one of the following:

- The Virtualized NEM Tools and Drivers CD:
`/vmware/drivers/esx3.5u4/Vmware-esx-drivers-net-hxge-1.2.2.11-00000.i386.rpm`
- Download the hxge driver from the following site:
(<http://download3.vmware.com/software/esx/esx350-hxge-1.2.2.11-201041.iso>) and create a driver CD.

2. Start the ESX Server machine and log in to service console as root.

3. Place the driver CD in the CD-ROM drive of the server.

4. Mount the driver CD.

5. Do one of the following:

- If you created a driver CD from the VMware download site, **change the directory to the VMupdates/RPMS/**
- If you are using the Tools and Drivers CD, change to
`/vmware/drivers/esx3.5u4/`

6. Do one of the following:

- If the NEM is already installed in the modular system chassis, run the following command.
`esxupdate update`
The server reboots after the driver update.
- If the NEM is not already installed in the modular system, run the `esxupdate` command with the `---noreboot` option, shut down the host machine manually, install the NEM, and then boot the host machine.
This allows you to install in a single boot.

Note – Remove the driver CD from the CD-ROM drive before the server reboots.

7. After the server reboots, log in to ESX Server.

8. Do the following to verify that the driver is installed successfully:

- a. Run the `esxupdate query` command and verify that the information about the driver is mentioned in the resulting message.
- b. View the PCI ID XML file in the `/etc/vmware/pciid/` directory.
The driver information is available in the file.
- c. Check for the latest version of the driver module in
`/usr/lib/vmware/vmkmod/`
- d. To verify that the driver is loaded and functioning, run the `vmkload_mod -l` command.
The driver is listed in the displayed list.

▼ To Install a Virtualized NEM Driver on an Existing ESX/ESXi 4.0 Server Using `vihostupdate`

1. Obtain the `hxge` driver from one of the following:

- The Virtualized 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 the 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)

▼ To Install a Virtualized NEM Driver on an Existing ESX 4.0 Server Using esxupdate

1. Obtain the hxge driver from one of the following:
 - The Virtualized 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 the 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 on With a New ESX Installation

This section describes the procedures for installing the VMware ESX Server drivers during an ESX server installation. This section contains the following procedures:

- “To Install a Virtualized NEM Driver With a New ESX3.5 Installation” on page 109
- “To Install a Virtualized NEM Driver With a New ESX4.0 Installation” on page 110

▼ To Install a Virtualized NEM Driver With a New ESX3.5 Installation

1. Obtain the `hxge` driver from one of the following:

- The Virtualized NEM Tools and Drivers CD:
`/vmware/drivers/esx3.5u4/Vmware-esx-drivers-net-hxge-1.2.2.1-00000.i386.rpm`
- Download the `hxge` driver from the following site:
(<http://download3.vmware.com/software/esx/esx350-hxge-1.2.2.11-201041.iso>) and create a driver CD.

2. Place the driver CD in the CD-ROM drive of the host machine.

3. Start the host machine.

4. When prompted for an upgrade or installation, press Enter for graphical mode.

5. Choose the language you prefer.

6. Select a keyboard type.

7. After you are prompted to swap the driver CD with the ESX Server installation CD, insert the ESX Server 3.5 Update 3 installation CD and continue with ESX Server installation.

See

(http://www.vmware.com/pdf/vi3_301_201_installation_guide.pdf) for detailed ESX installation instructions.

8. After ESX Server is installed and the system reboots, log in to ESX Server.

9. Verify that the driver is installed successfully:

- a. Run the `esxupdate query` command.

A message containing the information about the driver appears.

- b. View the PCI ID XML file in the `/etc/vmware/pciid/` directory.

The driver information is available in the file.

- c. Check for the latest version of the driver module in the following directory:

`/usr/lib/vmware/vmkmod/`

- d. To verify that the driver is loaded and functioning, enter the following command:

```
vmkload_mod -l
```

The driver is listed in the displayed list.

▼ To Install a Virtualized NEM Driver With a New ESX4.0 Installation

1. Obtain the hxge driver from one of the following:

- The Virtualized 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.htm) 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.

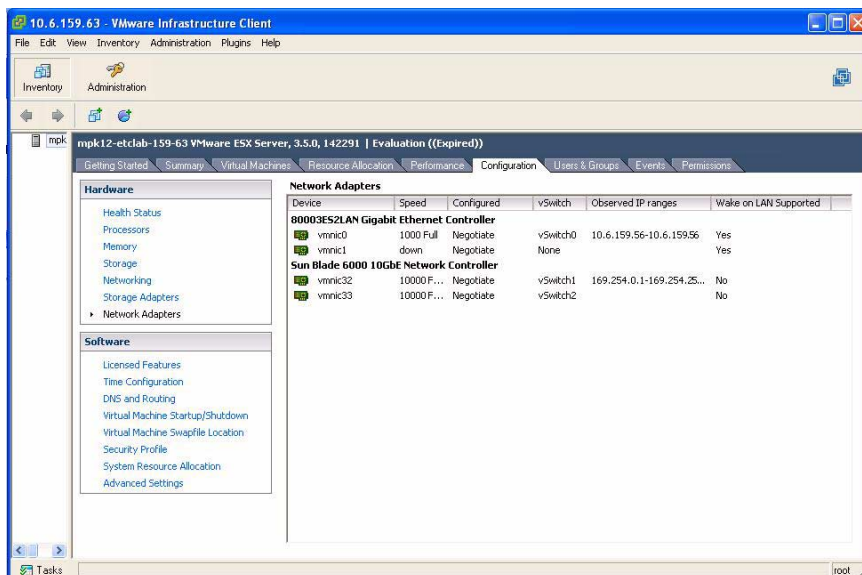
6. Click Add to eject the ESX installation DVD.

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 Network Adapters

▼ To Configure the Virtual NEM Network Adapters

1. Log in to the ESX host using Virtual Infrastructure Client GUI.
2. Click the Configuration Tab, then select Network Adapters from the Hardware list on the left side of the GUI.



3. Select the Sun Blade 6000 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

ESX Server 3 Configuration Guide available at:

(http://www.vmware.com/support/pubs/vi_pages/vi_pubs_35u2.html) or

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

The following topics are covered in this section:

- “Jumbo Frames Guidelines” on page 112
- “To Create a Jumbo Frames-enabled vSwitch” on page 113
- “To Create a Jumbo Frames-enabled VMkernel Interface” on page 113

Jumbo Frames Guidelines

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 3i/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 Server 3 Configuration Guide available at:

(http://www.vmware.com/support/pubs/vi_pages/vi_pubs_35u2.html) or

ESX Configuration Guide for ESX 4.0 available at:

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

▼ 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.

▼ 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 addresss 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.

The Common Array Manager

This chapter contains the following sections:

- [“About Enclosure Management” on page 115](#)
- [“CAM” on page 116](#)
- [“Obtaining CAM Software” on page 117](#)
- [“Using CAM With Disk Blades and Virtualized Multi-Fabric 10GbE NEM” on page 118](#)
- [“Upgrading Expander Firmware” on page 119](#)

About Enclosure Management

Your Sun Blade 6000 Virtualized Multi-Fabric 10GbE NEM supports a powerful set of enclosure management features accessible from a management client that supports SES-2 (SCSI Enclosure Services). The SAS expander on the Virtualized Multi-Fabric 10GbE NEM is SES-2 compliant. These enclosure management features are available through the management software called the Sun Common Array Manager (CAM), which provides a system administrator with the following capabilities:

- Event and fault monitoring
- E-mail alert notification
- FRU identification and status
- Enclosure reset
- Enclosure firmware upgrade
- Fault isolation
- Service Advisor wizard for problem resolution

- Sun Auto Service Request (ASR) uses fault telemetry 24/7 to automatically initiate a service request and begin the problem resolution process as soon as a problem occurs

CAM

CAM is a software application that allows you to manage SAS-NEMs and disk blades in a chassis from one central point. It is written in Java so you can run it on any platform.

Note – CAM operates primarily through disk blades. SAS-NEMs, including the Virtualized Multi-Fabric 10GbE NEM, are managed as if they were subcomponents of the disk blades. To use CAM in the Sun Blade 6000 modular system, you must register the disk blades in the chassis. You cannot register a SAS-NEM. However, if at least one disk blade is registered, then the SAS-NEMs will be visible to CAM and CAM can monitor them and update their expander firmware. This chapter is written around managing a disk blade because that is the way that CAM works. Management of your NEMs occurs indirectly through disk blade management.

In a Sun Blade 6000 chassis, one server blade should be chosen to host the CAM software.

CAM Agent

A second piece of software, called the CAM Agent, must be installed on every server blade that connects to a disk blade. The agent software is OS dependent. There are versions for Linux, Solaris, and Windows.

When a server blade connects to a disk blade, an additional CAM Agent plug-in is required for CAM to recognize the disk blade.

Note – The CAM installation program handles the installation of all three pieces of software.

The CAM Agent communicates with the disk module through the server blade's SAS host bus adapter. Both LSI and Adaptec controllers are supported.

At the initial release of the Sun Blade 6000 disk module, the most important functionality of CAM is firmware management. The CAM Agent can report the current versions of expander firmware and can update the expander firmware on both the disk module and the SAS-NEM.

The CAM Agent can monitor temperature and voltage on both the disk blade and the SAS-NEM.

It can also present the topology of your system and report FRU IDs.

Problem Resolution

CAM provides two helpful tools for resolving problems that might occur with the system.

- **Auto Service Request (ASR)** monitors the system's health and performance and automatically notifies the Sun Technical Support Center when critical events occur. Critical alarms generate an Auto Service Request case. These notifications enable Sun Service to respond faster and more accurately to critical on-site issues.

To use ASR, you must provide Sun online account information to enroll your CAM software to participate in the ASR service. After you enroll CAM with ASR, you can choose which systems you want to be monitored and enable them individually.

- **Service Advisor** is a troubleshooting wizard that provides information and procedures for replacing system components.

Obtaining CAM Software

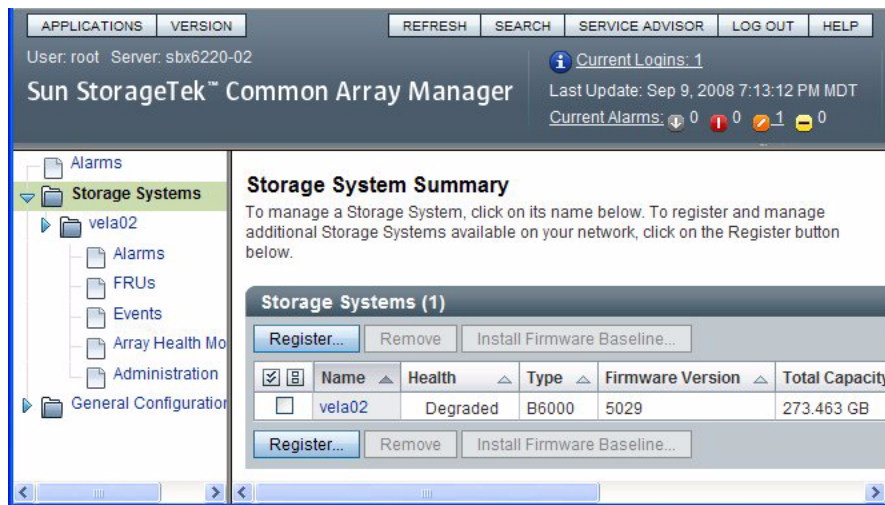
Support for the NEM is available beginning with CAM version 6.1.2. To obtain the latest CAM software and expander firmware, go to the Sun software download site:

<http://www.sun.com/download/index.jsp>

Click on the tab **View by Category**, and then choose **Systems Administration -> Storage Management**. Look for the StorageTek Common Array Manager (CAM) Software release 6.1.2 or later.

Using CAM With Disk Blades and Virtualized Multi-Fabric 10GbE NEM

CAM provides both a browser and command-line interface. When using the browser interface, you set up user accounts. When an authorized user logs in, information is laid out in pages. A handy navigation tree lists available information. See the following figure.



You use the navigation tree on the left to move among pages within an application. You can click a link to get details about a selected item (such as alarms, FRUs, events, or array health). You can also sort and filter information displayed on a page. When you place your pointer over a button, tree object, link, icon, or column, a tool tip provides a brief description of the object.

Monitoring Component Health

CAM can monitor voltage and temperature on installed disk blades and Virtualized Multi-Fabric 10GbE NEM and can raise alarms (including notification) when thresholds are exceeded. Choose the Array Health Monitoring item from the navigation tree to learn about this capability.

The monitoring capabilities of CAM are fully documented elsewhere. Refer to the *Sun StorageTek Common Array Manager User Guide for the J4000 Array Family* (820-3765) found at:

(<http://docs.sun.com/app/docs/prod/stor.armgr#hic>)

Note – CAM also has extensive online documentation that you can access using the Help button in the upper right of the browser interface.

Upgrading Expander Firmware

Both the Sun Blade 6000 disk modules and the Sun Blade 6000 Virtualized Multi-Fabric 10GbE NEM enclosures contain SAS expanders with upgradable firmware. You should keep this firmware at the latest released versions. CAM includes firmware upgrade capabilities for these components.

Note – The Sun Blade 6000 disk modules and Sun Blade 6000 Virtualized Multi-Fabric 10GbE NEMs must always be upgraded to the same firmware revision level.

▼ To Upgrade Expander Firmware

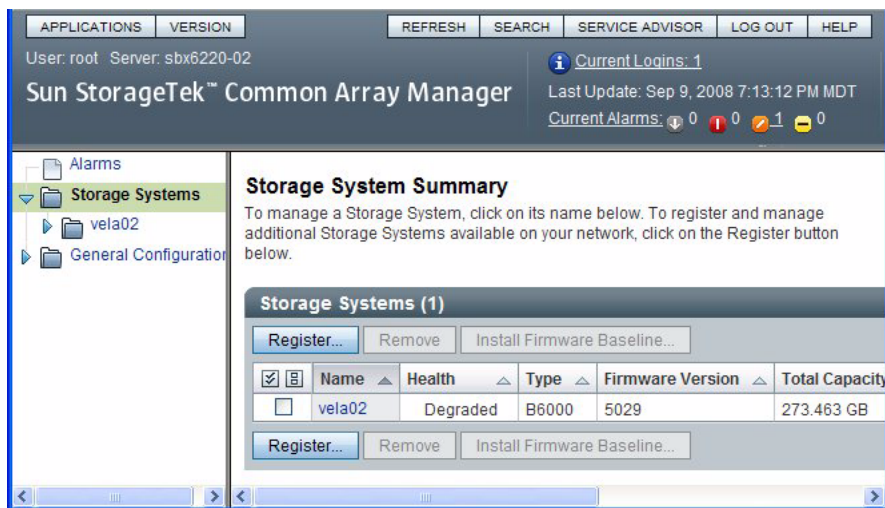
This procedure assumes that you have already registered your disk modules using the CAM interface.

Note – Registering the disk blades automatically makes CAM aware of the Virtualized Multi-Fabric 10GbE NEM.

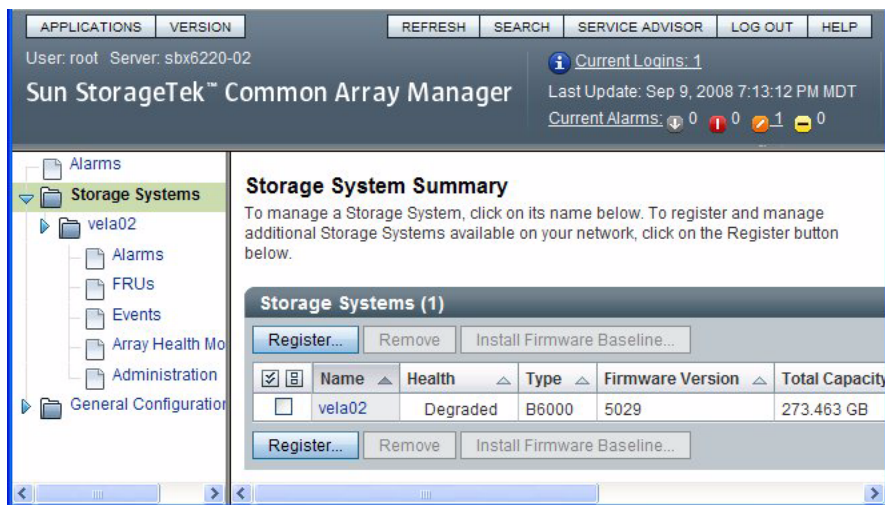
1. Open the CAM browser interface on the server that hosts the CAM software.

You see the registered storage systems listed in the opening summary page. In this example, there is only one registered storage system, a disk blade registered with the name “vela02.”

You can see that vela02’s health is degraded and also that there is a major (orange) alarm in the Current Alarms list.

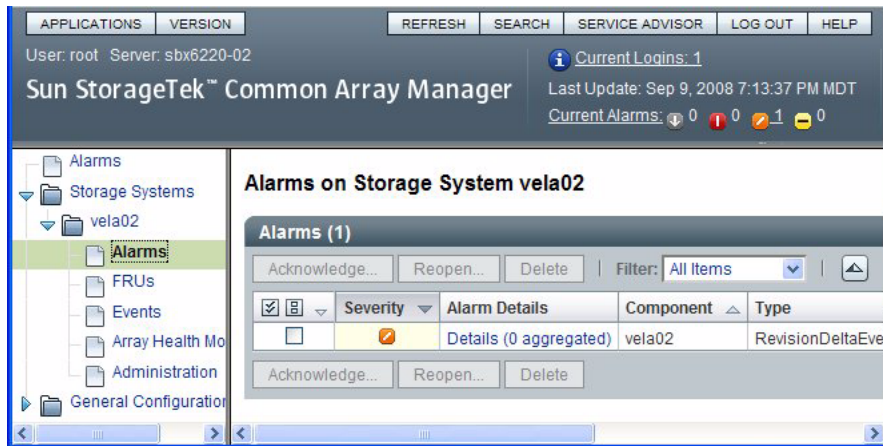


2. Click the arrow to expand vela02 in the navigation tree.



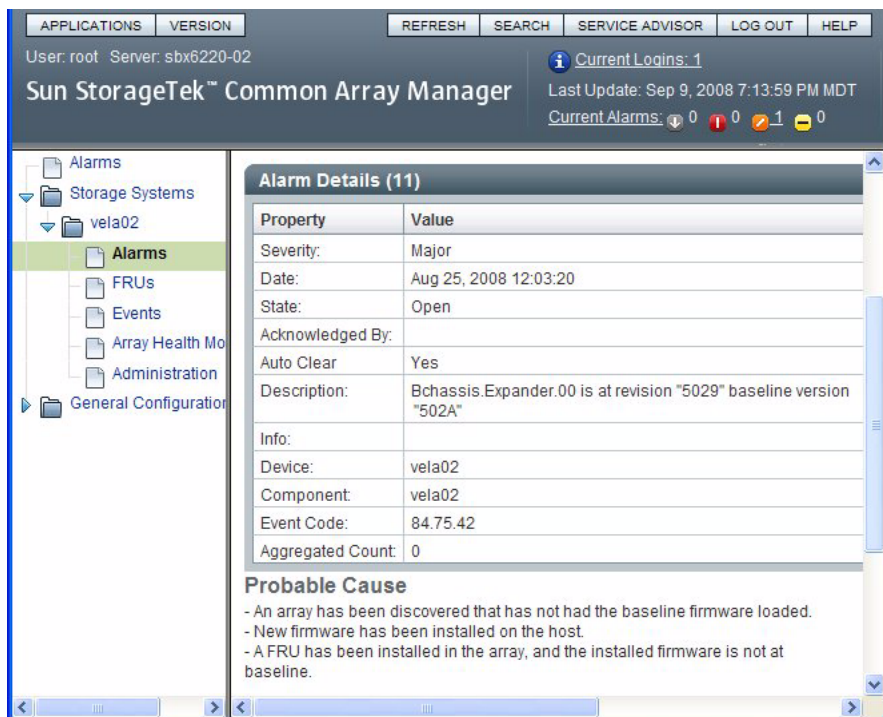
3. Select the Alarms item in the vela02 tree.

The Alarms page for vela02 opens. You can see one major severity (orange) alarm of type RevisionDeltaEvent.



4. Click the link under Alarm Details.

The Alarm Details page opens.



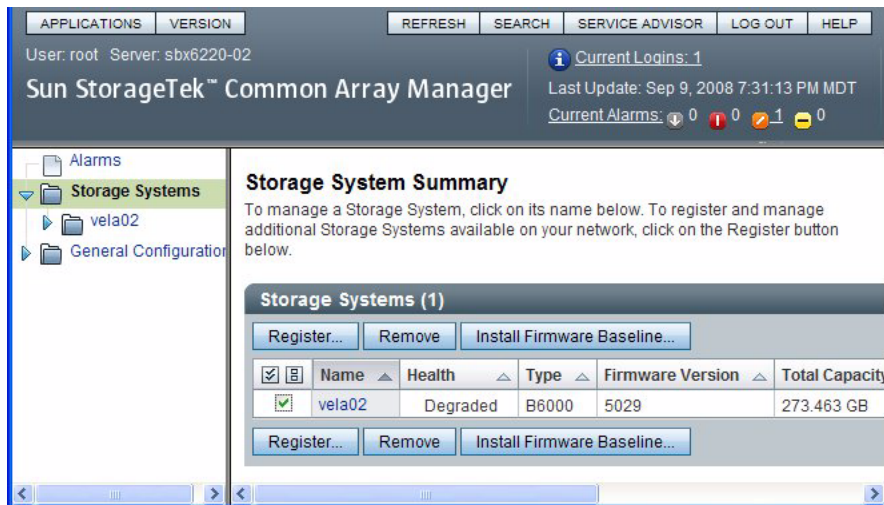
You can see from the description that the installed version of expander firmware (5029) is not up to date.

Note – You will get this same alarm if the firmware on *any* of the expanders (two on the disk blade and one on each of the NEMs) is not current for any reason.

Looking at the Probable Cause, you see various possible ways that the degraded condition that raised the alarm might have been produced.

5. Return to the Storage System Summary page by choosing Storage Systems in the navigation tree and check the check box to the left of the vela02 name.

This selects the vela02 storage system, then enables the firmware update button.



Note – In this example, there is only one registered storage system. If there were more, the buttons would be enabled only for the checked system.

6. Click the Install Firmware Baseline button to update the expander firmware.

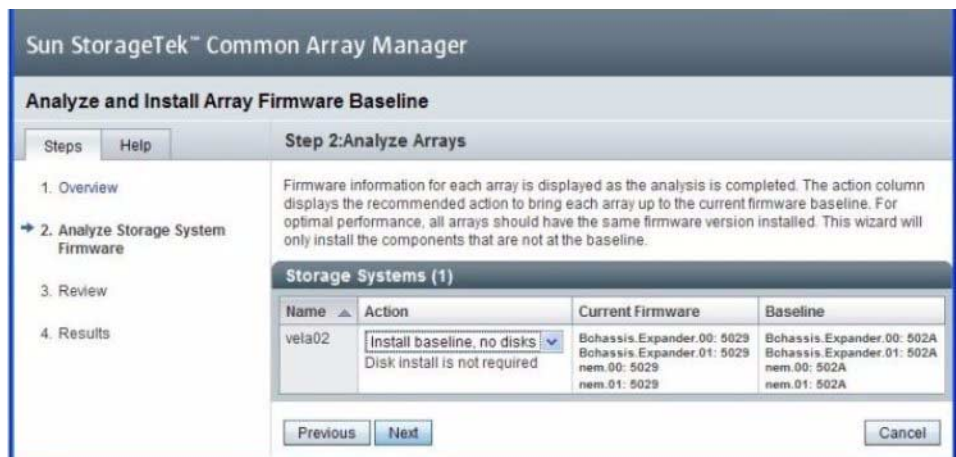
The Analyze and Install Array Firmware Baseline wizard opens.

Note – CAM updates the firmware on *any* expander that is not at the current level.



7. Click Next.

The next screen shows the current firmware versions and the baseline (correct) versions for each expander.



8. Accept the default Action and click Next.

The next screen shows a review of what you have selected.

Sun StorageTek™ Common Array Manager

Analyze and Install Array Firmware Baseline

Steps Help Step 3: Review

1. Overview
2. Analyze Storage System Firmware
→ 3. Review
4. Results

Review the current selections and then click Finish to perform any specified firmware install.

⚠ All management operations for Storage System should stop prior to firmware install. Disk I/O must be quiesced if disk firmware is being installed.

Storage Systems (1)

Name	Action	Current Firmware	Baseline
vela02	Install baseline, no disks Disk install is not required	Bchassis.Expander.00: 5029 Bchassis.Expander.01: 5029 nem.00: 5029 nem.01: 5029	Bchassis.Expander.00: 502A Bchassis.Expander.01: 502A nem.00: 502A nem.01: 502A

Previous Finish Cancel

9. If you need to change the Action, click Previous. Otherwise, click Finish.

A series of screens appear, with the Status field showing the current expander being updated. When the process is completed, the Firmware Install completed message appears in the Status Field.

Sun StorageTek™ Common Array Manager

Analyze and Install Array Firmware Baseline

Steps Help Step 4: Results

1. Overview
2. Analyze Storage System Firmware
3. Review
→ 4. Results

Monitor the status of the array firmware install process displayed in the Status and Progress fields. When the array firmware install process is complete, click Close.

All operations have completed.
Status for each Storage System is displayed in the table below.

Storage Systems (1)

Name	Status	Progress
vela02	Firmware install completed	<div></div>

Close

10. Click Close.

The Storage System Summary page reappears. You can see that the disk blade's Health has changed to OK and the alarm is turned off.

APPLICATIONS

VERSION

REFRESH

SEARCH

SERVICE ADVISOR

LOG OUT

HELP

User: root Server: bl189

Current Logins: 1

Last Update: Sep 11, 2008 10:01:38 AM PT

Current Alarms: 0 0 0 0

Alarms

Storage Systems

vela02

General Configuration

Storage System Summary

To manage a Storage System, click on its name below. To register and manage additional Storage Systems available on your network, click on the Register button below.

Storage Systems (1)

Register...

Remove

Install Firmware Baseline...

<input checked="" type="checkbox"/>	Name	Health	Type	Firmware Version	Total Capacity
<input type="checkbox"/>	vela02	OK	B6000	502A	410.195 GB

Register...

Remove

Install Firmware Baseline...

Using the `lsiutil` Utility to Save and Restore HBA Persistence Mappings

The procedures in this manual for replacing a Virtualized Multi-Fabric 10GbE NEM requires the use of a program called `lsiutil`. This appendix provides information how to use the application.



Caution – The `lsiutil` application has many features and capabilities. Executing certain combinations of commands can cause catastrophic damage. When you use this application, be sure you execute the procedures exactly as they are documented here, command by command. Do not skip commands or steps and do not add commands or steps that are not in the documented procedure.

The following topics are covered in this chapter:

- [“Where to Obtain the `lsiutil` Software” on page 127](#)
- [“Installing `lsiutil`” on page 128](#)
- [“Using the Interactive `lsiutil` Menus to Clear Your LSI Host Bus Adapter Non-Present Persistence Mapping” on page 128](#)
- [“Saving a Snapshot of Your Host Bus Adapter Persistence Mappings” on page 131](#)

Where to Obtain the `lsiutil` Software

To use `lsiutil` for the procedure in this appendix, you must have version 1.60 at minimum.

To obtain the latest `lsiutil` software, go to the disk blade download site:

Installing `lsiutil`

Download `LSIUtiKit1.60.zip` and unzip it in your root directory. It will create a main directory (`LSIUtiKit1.60`) with subdirectories for DOS, Linux, Solaris, Solaris x86, and Windows. If you are running `lsiutil` on a SPARC blade, use the Solaris subdirectory. Otherwise, use the Solaris x86 subdirectory.

There is an OS-specific version of `lsiutil` in each subdirectory.

Using the Interactive `lsiutil` Menus to Clear Your LSI Host Bus Adapter Non-Present Persistence Mapping

The procedure for clearing your LSI host bus adapter's non-present persistence mappings requires using the interactive menus of `lsiutil`.

▼ To Clear Your Non-Present Persistence Mappings

1. **Log in to your server as root.**
2. **Change to the directory that contains the appropriate version of `lsiutil` for your application. For example:**

```
# cd LSIUtiKit1.60/Linux
```
3. **Enter the application name in the appropriate context. For example:**

```
# ./lsiutil
```

The opening screen appears.


```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
[root@bl189 ~]# cd LSIUtilKit1.60/Linux
[root@bl189 Linux]# ./lsiutil

LSI Logic MPT Configuration Utility, Version 1.60, July 11, 2008

1 MPT Port found

      Port Name          Chip Vendor/Type/Rev    MPT Rev  Firmware Rev  IOC
1.  /proc/mpt/ioc0      LSI Logic SAS1068E B2    105      01185b00      0

Select a device: [1-1 or 0 to quit]
```

4. Type 1 and press Enter.

The main menu appears.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Select a device: [1-1 or 0 to quit] 1

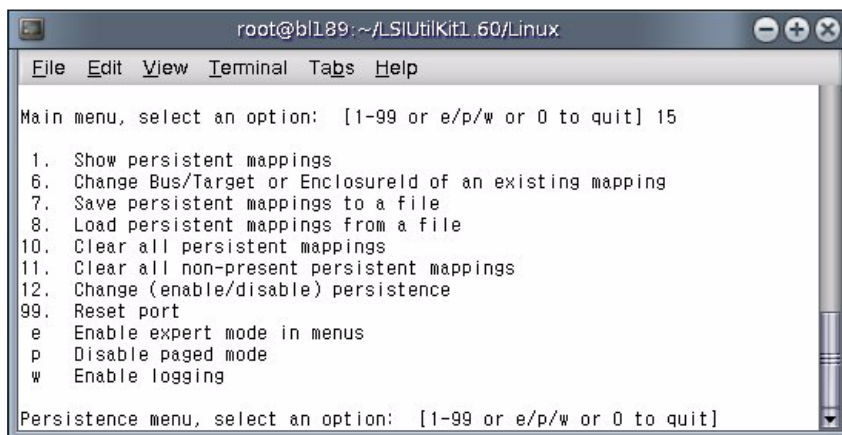
1.  Identify firmware, BIOS, and/or FCode
2.  Download firmware (update the FLASH)
4.  Download/erase BIOS and/or FCode (update the FLASH)
8.  Scan for devices
10. Change IOC settings (interrupt coalescing)
13. Change SAS IO Unit settings
16. Display attached devices
20. Diagnostics
21. RAID actions
22. Reset bus
23. Reset target
42. Display operating system names for devices
45. Concatenate SAS firmware and NVDATA files
59. Dump PCI config space
60. Show non-default settings
61. Restore default settings
66. Show SAS discovery errors
69. Show board manufacturing information
97. Reset SAS link, HARD RESET
98. Reset SAS link
99. Reset port
   e Enable expert mode in menus
   p Enable paged mode
   w Enable logging

Main menu, select an option: [1-99 or e/p/w or 0 to quit]
```

You cannot see all the commands in the menu because they are too numerous, but you can enter any option number if you know what it is. In our case we want to open the Persistence menu. This is done with option 15.

5. Type 15 and press Enter.

The Persistence menu opens.

A screenshot of a terminal window titled 'root@bl189:~/LSIUtilKit1.60/Linux'. The window has a menu bar with 'File', 'Edit', 'View', 'Terminal', 'Tabs', and 'Help'. The main text area shows a 'Main menu, select an option: [1-99 or e/p/w or 0 to quit] 15'. Below this is a list of options: 1. Show persistent mappings, 6. Change Bus/Target or Enclosureid of an existing mapping, 7. Save persistent mappings to a file, 8. Load persistent mappings from a file, 10. Clear all persistent mappings, 11. Clear all non-present persistent mappings, 12. Change (enable/disable) persistence, 99. Reset port, e Enable expert mode in menus, p Disable paged mode, w Enable logging. At the bottom, it says 'Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]'.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help

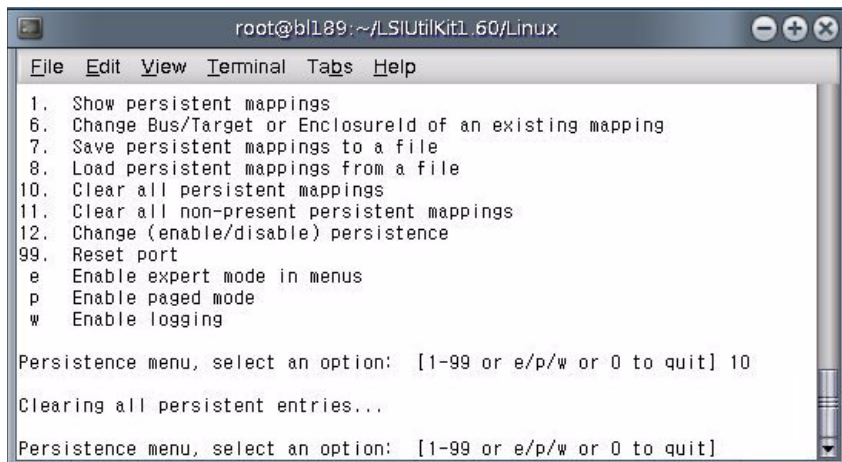
Main menu, select an option: [1-99 or e/p/w or 0 to quit] 15

1. Show persistent mappings
6. Change Bus/Target or Enclosureid of an existing mapping
7. Save persistent mappings to a file
8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
e Enable expert mode in menus
p Disable paged mode
w Enable logging

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]
```

6. Type 11 (Clear all non-present persistence mappings) and press Enter.

This clears any entries in the host bus adapter's persistence map that refer to target objects that are no longer present.

A screenshot of the same terminal window as before, but now showing the 'Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 10'. The list of options is the same, but option 10 is highlighted. Below the list, it says 'Clearing all persistent entries...'. At the bottom, it says 'Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]'.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help

1. Show persistent mappings
6. Change Bus/Target or Enclosureid of an existing mapping
7. Save persistent mappings to a file
8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
e Enable expert mode in menus
p Enable paged mode
w Enable logging

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 10

Clearing all persistent entries...

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]
```

7. Type and enter 0 three times to exit the lsiutil application.

Saving a Snapshot of Your Host Bus Adapter Persistence Mappings

▼ To Save a Snapshot of Your Host Bus Adapter Persistence Mappings

1. Log in to your server as root.

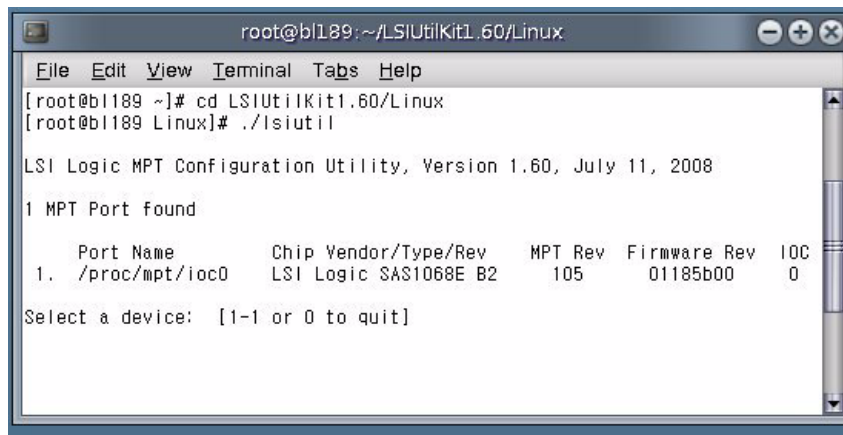
2. Change to the directory that contains the appropriate version of lsiutil for your OS. For example:

```
# cd LSIUtilKit1.60/Linux
```

3. Enter the application name in the appropriate context. For example:

```
# ./lsiutil
```

The opening screen appears.



The screenshot shows a terminal window titled "root@bl189:~/LSIUtilKit1.60/Linux". The terminal output is as follows:

```
[root@bl189 ~]# cd LSIUtilKit1.60/Linux
[root@bl189 Linux]# ./lsiutil

LSI Logic MPT Configuration Utility, Version 1.60, July 11, 2008

1 MPT Port found

      Port Name          Chip Vendor/Type/Rev    MPT Rev  Firmware Rev  IOC
1.  /proc/mpt/ioc0      LSI Logic SAS1068E B2    105      01185b00      0

Select a device: [1-1 or 0 to quit]
```

4. Type **1** and press **Enter**.

The main menu appears.

A screenshot of a Linux terminal window. The title bar shows the user is 'root' at 'bl189' in the directory '~/LSIUtilKit1.60/Linux'. The terminal has a menu bar with 'File', 'Edit', 'View', 'Terminal', 'Tabs', and 'Help'. The main text displays a menu titled 'Select a device: [1-1 or 0 to quit] 1'. The menu lists 23 numbered options, followed by three options labeled 'e', 'p', and 'w'. The options are: 1. Identify firmware, BIOS, and/or FCode; 2. Download firmware (update the FLASH); 4. Download/erase BIOS and/or FCode (update the FLASH); 8. Scan for devices; 10. Change IOC settings (interrupt coalescing); 13. Change SAS IO Unit settings; 16. Display attached devices; 20. Diagnostics; 21. RAID actions; 22. Reset bus; 23. Reset target; 42. Display operating system names for devices; 45. Concatenate SAS firmware and NVDATA files; 59. Dump PCI config space; 60. Show non-default settings; 61. Restore default settings; 66. Show SAS discovery errors; 69. Show board manufacturing information; 97. Reset SAS link, HARD RESET; 98. Reset SAS link; 99. Reset port; e. Enable expert mode in menus; p. Enable paged mode; w. Enable logging. At the bottom, it says 'Main menu, select an option: [1-99 or e/p/w or 0 to quit]'.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Select a device: [1-1 or 0 to quit] 1

 1. Identify firmware, BIOS, and/or FCode
 2. Download firmware (update the FLASH)
 4. Download/erase BIOS and/or FCode (update the FLASH)
 8. Scan for devices
10. Change IOC settings (interrupt coalescing)
13. Change SAS IO Unit settings
16. Display attached devices
20. Diagnostics
21. RAID actions
22. Reset bus
23. Reset target
42. Display operating system names for devices
45. Concatenate SAS firmware and NVDATA files
59. Dump PCI config space
60. Show non-default settings
61. Restore default settings
66. Show SAS discovery errors
69. Show board manufacturing information
97. Reset SAS link, HARD RESET
98. Reset SAS link
99. Reset port
  e Enable expert mode in menus
  p Enable paged mode
  w Enable logging

Main menu, select an option: [1-99 or e/p/w or 0 to quit]
```

You cannot see all the commands in the menu, but you can enter any command number if you know what it is. In our case we want to open the Persistence menu. This is done with command 15.

5. Type 15 and press Enter.

The Persistence menu opens.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Main menu, select an option: [1-99 or e/p/w or 0 to quit] 15
1. Show persistent mappings
6. Change Bus/Target or EnclosureId of an existing mapping
7. Save persistent mappings to a file
8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
e Enable expert mode in menus
p Disable paged mode
w Enable logging
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]
```

6. Type 7 and press Enter.

You are prompted for the name of the file where you want to store the host bus adapter's persistent mappings snapshot.

7. Enter a file name. For example, `HBA33_persistence_file`.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Main menu, select an option: [1-99 or e/p/w or 0 to quit] 15
1. Show persistent mappings
6. Change Bus/Target or EnclosureId of an existing mapping
7. Save persistent mappings to a file
8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
e Enable expert mode in menus
p Disable paged mode
w Enable logging
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 7
Enter persistent mapping filename: HBA33_persistence_file
```

The file is saved to the current directory. You get confirmation of this with a statement of the number of persistent entries saved.

8. Type and Enter 0 three times to exit the `lsiutil` application.



Caution – You *must* save the persistence mappings snapshot to external media as you will not be able to reinstall it from a local disk.

ILOM Supplement

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

- 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.

This appendix provides information about the following topics:

- [“Documentation and Updates” on page 135](#)
- [“Connecting to ILOM” on page 136](#)
- [“Temperature, Voltage, and Fan Sensors” on page 142](#)
- [“ILOM Proxy” on page 147](#)
- [“Enable or Disable IAL Dynamic Mode” on page 148](#)

Documentation and Updates

ILOM documentation is divided into two categories:

- Generalized ILOM information, located in the *Sun Integrated Lights Out Manager 2.0 User's Guide*
- Information specific to the Virtualized NEM, located in this supplement

Product Updates

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

(<http://www.sun.com/downloads>)

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 NEM.

Note – The Sun Blade 6000 Modular System chassis has its own ILOM, called a CMM ILOM, which is a separate entity from the NEM ILOM. It is described in the *Sun Integrated Lights Out Manager 2.0 Administration Guide*

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

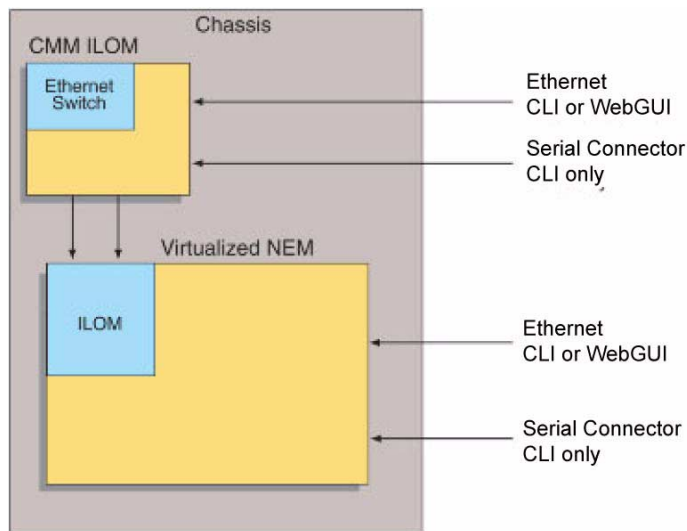
TABLE B-1 ILOM Connection Options

Option Number	Connection Method	User Interface	More Information
1	Ethernet	Web interface or CLI	Option 1 - Ethernet
2	Serial - CMM	CLI only	Option 2 - Chassis Serial Connector
3	Serial - NEM	CLI only	Option 3 - NEM Serial Connector

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

[FIGURE B-1](#) shows a graphical representation of the ILOM connection options.

FIGURE B-1 ILOM Connection Options



Option 1 - Ethernet

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 connect to the ILOM using SSH to connect to the command line interface (CLI), or using a browser to connect to the web GUI.

- To connect to the CLI, see [“To Connect to the CLI” on page 138](#).
- To connect to the web GUI, see [“To Connect to the WebGUI” on page 139](#).

If you do not know the ILOM's IP address, find it as described in the following procedure.

▼ To Find the ILOM's IP Address

1. Log in to the CMM.

See the Integrated *Lights Out Manager (ILOM) Administration Guide For the Sun Blade 6000 Modular System*, 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/NEMN/SP/network
Targets:

Properties:
  type = Network Configuration
  commitpending = (Cannot show property)
  ipaddress = IPaddress
  ipdiscovery = dhcp
  ipgateway = IPgateway
  ipnetmask = 255.255.252.0
  macaddress = Macaddress
  pendingipaddress = Ipaddress
  pendingipdiscovery = dhcp
  pendingipgateway = IPgateway
  pendingipnetmask = 255.255.252.0

Commands:
  cd
  set
  show

->
```

▼ To Connect to the CLI

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.

▼ To Connect to the WebGUI

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.

Option 2 - Chassis Serial Connector

The chassis serial connector connects to the chassis CMM ILOM. The CMM ILOM provides a command to connect to the NEM ILOM(s).

▼ To Connect to ILOM Through the Chassis 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

Pin	Signal Description
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 SUNSP and the CMM ILOM's MAC 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/NEMn/SP/cli`.

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/BLn/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.

The following display shows an example:

```
-> cd /CH/NEM0/SP/cli
/CH/NEM0/SP/cli
```

```
-> start
Are you sure you want to start /CH/NEM0/SP/cli (y/n)? y
Password:          Type the password to the CMM ILOM.

Sun(TM) Integrated Lights Out Manager

Version 2.0.3.12

Copyright 2006 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.

Warning: password is set to factory default.

-> cd /NEM          This command is only available on the NEM ILOM.
/NEM

-> exit             Type this command to exit the NEM ILOM and return to the CMM
ILOM.
Connection to 10.6.153.33 closed.
```

Option 3 - NEM Serial Connector

▼ 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 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. **Type `start /NEMx` to start the NEM.**

Note – Start `/NEMx` command will not start the NEM until the ATTN button is pressed manually.

4. **When you are done, exit the ILOM by typing:**

`-> exit`

Temperature, Voltage, and Fan 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.

[“Board Sensor List” on page 142](#) lists the sensors. [“Sensor Details” on page 143](#) provides details about each sensor.

Board Sensor List

Board Sensors

- slotid
- nem/cmm/prsnt
- nem/ok
- nem/ok2rm
- nem/service
- nem/locate
- nem[0..1].prsnt
- mb/t_amb[0..2]
- mb/v_+1v2
- mb/v_+3v3
- mb/v_+5v
- mb/v_+12v

Sensor Details

The following sections provide detailed information about the sensors.

slotid

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

nem/cmm/prsnt

This is a sensor indicating whether CMM ILOM is present.

TABLE B-2 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

nem/ok

This is an LED indicator sensor shows the power state of the NEM. It does not generate any events.

TABLE B-3

Reading	State	Event	Description
0x0001	State Deasserted	No	This indicates that nem/ok LED is OFF.
0x0002	State Asserted	No	This indicates that nem/ok LED is ON.

nem/ok2rm

This is an LED indicator sensor that shows the immediate state of the nem/ok2rm LED on the NEM. It does not generate any events.

TABLE B-4 sys.ok2rm Sensors

Reading	State	Event	Description
0x0001	State Deasserted	No	This state indicates that the nem/ok2rm LED is OFF.
0x0002	State Asserted	No	This state indicates that the nem/ok2rm LED is ON.

nem/service

This is an LED indicator sensor that shows the immediate state of the nem/service LED on the NEM. It does not generate any events.

TABLE B-5 sys.alert Sensors

Reading	State	Event	Description
0x0001	State Deasserted	No	This state indicates that the Service LED is OFF.
0x0002	State Asserted	No	This state indicates that the Service LED is ON or BLINKING

nem/locate

This is an LED indicator sensor that shows the immediate state of the Locate LED on the NEM. It does not generate any events.

TABLE B-6 sys.locate Sensors

Reading	State	Event	Description
0x0001	State Deasserted	No	This state indicates that the nem/led LED is OFF.
0x0002	State Asserted	No	This state indicates that the nem/led LED is FAST BLINKING.

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[0..2]`

There are three temperature sensors on the motherboard.

These sensors monitor the ambient temperature from the internal temperature sensor on the mainboard.

- `MB/T_AMB[0]` represents SAS ambient temperature.
- `MB/T_AMB[1]` represents Hydra 0 ASIC ambient temperature.
- `MB/T_AMB[2]` represents Hydra 1 ASIC ambient temperature.

NEM Voltage Sensors

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

`mb.v_+12v`

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

`mb.v_+5v`

This sensor monitors the 5V main input that is active when the power is on.

`mb.v_+3v3`

This sensor monitors the 3.3V 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.

TABLE B-7 mb.v_+1v2

Threshold	Direction	Event	Description	Action
Lower Non-Critical	Assert	Yes	Voltage has decreased below lower non-critical threshold.	System Alert LED is SLOW
Lower Non-Critical	Deassert	Yes	Voltage has returned to normal from lower non-critical.	System Alert LED is OFF
Lower Critical	Assert	Yes	Voltage has decreased below lower critical threshold.	System Alert LED is SLOW
Lower Critical	Deassert	Yes	Voltage has returned to lower non-critical from lower critical.	System Alert LED is SLOW
Lower Non-Recoverable	Assert	Yes	Voltage has decreased below lower non-recoverable threshold.	System Alert LED is SLOW
Lower Non-Recoverable	Deassert	Yes	Reading has returned to lower critical from lower non-recoverable.	System Alert LED is SLOW
Upper Non-Critical	Assert	Yes	Voltage has increased above upper non-critical threshold.	System Alert LED is SLOW
Upper Non-Critical	Deassert	Yes	Voltage has returned to normal from upper non-critical.	System Alert LED is OFF
Upper Critical	Assert	Yes	Voltage has increased above upper critical threshold.	System Alert LED is SLOW
Upper Critical	Deassert	Yes	Voltage has returned to upper non-critical from upper critical.	System Alert LED is SLOW
Upper Non-Recoverable	Assert	Yes	Voltage has increased above upper non-recoverable threshold.	System Alert LED is SLOW
Upper Non-Recoverable	Deassert	Yes	Voltage has returned to upper critical from upper non-recoverable.	System Alert LED is SLOW

ILOM Proxy

In addition, a *proxy* ILOM CLI program is provided on the CMM to augment the core ILOM program's ability to manage Virtualized NEMs and Sun Blade disk modules.

To take advantage of the module management capabilities offered by the proxy ILOM CLI, verify that you have the latest version of the proxy program installed. Refer to the *Sun Blade 6000 Disk Module Configuration Guide* for details about Multi-Fabric NEM software and firmware requirements.

Starting the ILOM Proxy CLI Program

Instructions for setting up and using ILOM are documented in the *Sun Integrated Lights Out Manager 2.0 User's Guide*.

Prerequisite: To start the proxy CLI, a Virtualized Multi-Fabric 10GbE NEM module must be installed in the chassis in NEM slot 0.

▼ To Start the ILOM Proxy CLI Program

1. **Start the proxy program through the ILOM CLI. 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.
 - Connect to the Ethernet network management port using a secure shell (SSH).
2. **After you are connected to the ILOM CLI, start the proxy CLI by typing the following command:**

```
-> start /CH/NEM0/SP/cli
```

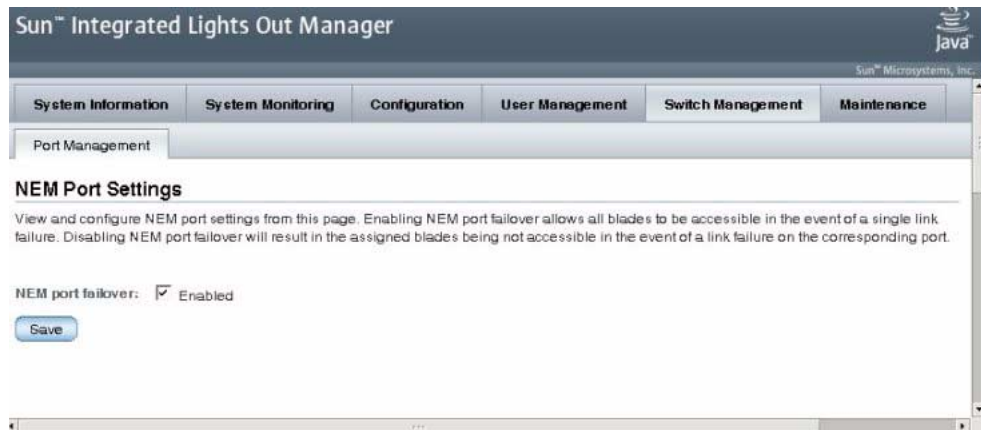
Note – The proxy CLI commands are case sensitive.

Enable or Disable IAL Dynamic Mode

You can enable or disable IAL dynamic mode either through the web interface or the CLI. See [“Enabling or Disabling the IAL” on page 8](#) for more information on when to enable IAL dynamic mode.

▼ To Enable or Disable IAL Dynamic Mode Using the Web Interface

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



3. At the NEM port failover field, do one of the following:
 - Check the Enable box to enable dynamic mode.
 - Uncheck the Enable box to disable dynamic mode.

▼ To Enable or Disable Dynamic Mode Using the CLI

1. Display the current IAL mode.

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

2. Set the value in the appropriate mode.

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


Index

Numerics

10 GbE connections, 6

A

ASIC, 7

assigning an IP address, 43

attention button, 13

B

buttons, 12

C

cabling SFP+ connectors, 29

CAM

about enclosure management, 115

agent, 116

monitoring component health, 118

obtaining CAM software, 117

overview, 116

problem resolution, 117

upgrading enclosure firmware, 119

using with disk blades and Multi-Fabric
NEMs, 118

D

default zoning requirements, 4

drivers

ESX server operating system

configuring Jumbo Frames, 112

installing drivers, 106

Linux operating systems

configuring Jumbo Frames, 75

configuring the network interface, 59

installing, 53

removing, 58

Solaris operating system

configuring hxge device driver, 44

configuring Jumbo Frames, 48

configuring network host files, 43

installing, 41

removing, 42

Windows operating system

installing, 78

dynamic mode, enabling, 148

E

ESX server operating system

configuring Jumbo Frames, 112

configuring NEM network adapters, 110

installing drivers

existing ESX 3.5 server, 106

existing ESX 4.0 server with esxupdate, 108

existing ESX/ESXi 4.0 server with
vihostupdate, 107

new ESX 3.5 server, 109

new ESX 4.0 server, 110

Ethernet connections, 5

Ethernet LEDs, 14

expander firmware

upgrading, 119

F

fabric express modules (FEMs), 10

features, 1

features overview, 3
firmware
 upgrading expander firmware, 119

G

Gigabit Ethernet ports, 12

I

ILOM

- connecting, 136
- description, 135
- documentation, 135
- dynamic mode, 148
- proxy program, 147
- sensors, 142

installing an SFP+ module, 27

Integrated Lights Out Manager, see ILOM, 135

inter-ASIC links

- description, 8
- functionality, 9

J

Jumbo Frames

- checking configurations, 48
- enabling in a Solaris SPARC environment, 50

L

Layer 2 configuration, checking, 51

Layer 3 configuration, checking, 52

LEDs, 12

- Ethernet, 14
- Locate, 13
- Module Activity, 13
- Service Action Required, 13

Linux operating systems

- configuring Jumbo Frames, 75
- configuring the network interface, 59
- installing drivers, 53
- removing drivers, 58

Locate LED and button, 13

lsiutil

- installing, 128
- obtaining, 127
- using to clear LSI persistence mapping, 128
- using to save snapshot of HBA persistence mappings, 131

M

Module Activity LED, 13

Module Fault LED, 13

N

NEM components

- 10 GbE connections, 6
- Ethernet connections, 5
- virtualized NEM ASIC, 7

network booting with Solaris blade server, 37

network booting with x86 blade server, 32

O

OK-to-Remove LED, 13

P

parameter values

- modifying and displaying, 47

physical appearance, 11

ports, 12

Power/OK LED, 13

PXE booting with Solaris blade server, 37

PXE booting with x86 blade server, 32

R

removing a NEM, 25

removing an SFP+ module, 28

S

SAS expanders

- upgrading firmware, 119

Service Action Required LED, 13

service processor, 10

SFP+ module

- cabling, 29
- installation, 27
- removal, 28

SFP+ ports, 12

Solaris operating system

- configuring hxge device driver, 44
- configuring Jumbo Frames, 48
- configuring network host files, 43
- installing drivers, 41
- removing drivers, 42

T

terms used, 2

U

upgrading enclosure firmware, 119

upgrading expander firmware, 119

V

verifying NEM installation with ILOM CLI, 20, 22

W

Windows operating system

- enabling Jumbo Frames, 104

- installing drivers, 78

- installing enclosure device, 93, 99

- installing network controller drivers, 78

- installing VLAN drivers, 84

