



SGI® 10-Gigabit Ethernet Network
Adapter User's Guide

007-4669-002

CONTRIBUTORS

Written by Lori Johnson

Illustrated by Dan Young and Chrystie Danzer

Engineering contributions Arthur Kepner, Edward Mascarenhas, Steve Modica, Alain Renaud, James Yarbrough

COPYRIGHT

© 2004, 2005 Silicon Graphics, Inc. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the US government or any contractor thereto, it is acquired as "commercial computer software" subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is Silicon Graphics, Inc., 1600 Amphitheatre Pkwy 2E, Mountain View, CA 94043-1351.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, Altix, IRIX, Onyx, Origin, and Silicon Graphics Tezro are registered trademarks and Silicon Graphics Prism is a trademark of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

Intel is a registered trademark of Intel Corporation. Linux is a registered trademark of Linus Torvalds, used with permission by Silicon Graphics, Inc. UNIX is a registered trademark of The Open Group in the United States and other countries.

FCC WARNING

This equipment has been tested and found compliant with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

ATTENTION

This product requires the use of external shielded cables in order to maintain compliance pursuant to Part 15 of the FCC Rules.

VDE 0871/6.78

This equipment has been tested to and is in compliance with the Level A limits per VDE 0871.

EUROPEAN UNION STATEMENT

This device complies with the European Directives listed on the "Declaration of Conformity" which is included with each product. The CE mark insignia displayed on the device is an indication of conformity to the aforementioned European requirements.



International Special Committee on Radio Interference (CISPR)

This equipment has been tested to and is in compliance with the Class A limits per CISPR publication 22.

Canadian Department of Communications Statement

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Attention

Cet appareil numérique n'émet pas de perturbations radioélectriques dépassant les normes applicables aux appareils numériques de Classe A prescrites dans le Règlement sur les interférences radioélectriques établi par le Ministère des Communications du Canada.

Japanese Compliance Statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Compliance Statement in Chinese

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

New Features

This guide contains the new Chapter 4, "Performance Tuning" on page 15.

Record of Revision

Version	Description
001	May 2004 Initial publication
002	August 2005 Revision

Contents

About this Guide	xi
Important Information	xi
Scope of this Guide	xii
Related Publications	xii
Obtaining Publications	xiii
Conventions	xiii
Product Support	xiv
Reader Comments	xiv
1. Features and Capabilities of the SGI 10-Gbit Ethernet Network Adapter	1
Features	2
Cabling	2
Configuration Limits	3
Tools	3
2. Connecting the Adapter to a Network	5
Installing the Adapter	5
Connecting to the Network	5
3. Operating the Adapter	7
Verifying Functionality	7
Using LEDs to Determine Functionality	7
Verifying Adapter Recognition	8
IRIX Adapter Recognition Verification	8
Linux Adapter Recognition Verification	9
007-4669-002	ix

Enabling the Adapter	10
Verifying that the Adapter is Properly Configured and Enabled	10
IRIX Adapter Configuration Verification	10
Linux Adapter Configuration Verification	11
Resetting the Adapter	11
Configuration	12
IRIX Configuration	12
Linux Configuration	12
MTU Size Settings	12
4. Performance Tuning	15
IRIX Performance Tuning	15
Jumbo Frames and IRIX	15
Read/Write Size and IRIX	16
Socket Buffer Size and IRIX	16
Multibuffer Mode and IRIX	16
SGI ProPack 3 Performance Tuning	17
Socket Read and Write Buffer Sizes for SGI ProPack 3	17
Jumbo Frames and SGI ProPack 3	17
Read/Write Size and SGI ProPack 3	17
TCP/IP Socket Buffer Size and SGI ProPack 3	18
Glossary	19
Index	21

About this Guide

This guide describes the SGI 10-Gigabit (Gbit) Ethernet network adapter. This guide shows you how to connect the adapter to an Ethernet network and explains how to operate the adapter.

The SGI 10-Gbit Ethernet network adapter operates under the versions of the IRIX and SGI ProPack for Linux operating systems, as specified in the release notes.

You can use the SGI 10-Gbit Ethernet network adapter in addition to your current adapter.

This guide assumes that you have general knowledge of Ethernet networks and the system in which the adapter is installed.

Important Information



Warning: Never look into the end of a fiber optic cable to confirm that light is being emitted (or for any other reason).

Do not use any type of magnifying device, such as a microscope, eye loupe, or magnifying glass. Such activity causes a permanent burn on the retina of the eye. Optical signal cannot be determined by looking into the fiber end.

Most fiber optic laser wavelengths (1300 nm and 1550 nm) are invisible to the eye and cause permanent eye damage. Shorter wavelength lasers (for example, 780 nm) are visible and can cause significant eye damage.

Use only an optical power meter to verify light output.

Scope of this Guide

This guide is written to facilitate installation of the adapter and does not cover detailed points of network configuration. It contains the following chapters:

- Chapter 1, "Features and Capabilities of the SGI 10-Gbit Ethernet Network Adapter", summarizes features, cabling, configuration limits, and tools.
- Chapter 2, "Connecting the Adapter to a Network", shows you how to connect the adapter to your network.
- Chapter 3, "Operating the Adapter", explains how to verify installation of the adapter and software, how to reset the adapter, how to set parameters to improve performance, and how to set configuration parameters.

Related Publications

This guide is part of a document set that fully supports the installation, operation, and service of the adapter. For more information about installing and servicing the adapter, see the user's guide for the system in which the adapter is installed.

Also see the following:

- *IRIX Admin: Networking and Mail*
- *Linux Configuration and Operations Guide*
- *The Network Administrators' Guide*
- The Intel Corporation *10 Gigabit Ethernet Technology Overview* white paper:

http://www.intel.com/network/connectivity/resources/doc_library/white_papers/pro10gbe_lr_sa_wp.pdf

- IRIX man pages that are part of the 10-Gbit Ethernet product:
 - `xg(7M)`
 - `xgcmd(1M)`
- Standard Linux man pages that are useful for any Ethernet device:
 - `ethtool(8)`
 - `ifconfig(8)`

– ip(8)

Obtaining Publications

You can obtain SGI documentation as follows:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- If it is installed on your SGI system, you can use InfoSearch, an online tool that provides a more limited set of online books, release notes, and man pages. With an IRIX system, enter `infosearch` at a command line or select **Help > InfoSearch** from the Toolchest.
- On IRIX systems, you can view release notes by entering either `grelnotes` or `relnotes` at a command line.
- On Linux systems, you can view release notes on your system by accessing the `README.txt` file for the product. This is usually located in the `/usr/share/doc/productname` directory, although file locations may vary.
- You can view man pages by typing `man title` at a command line.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)

[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

Product Support

SGI provides a comprehensive product support and maintenance program for its products:

- If you are in North America, contact the Technical Assistance Center at +1 800 800 4SGI or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this publication, contact SGI. Be sure to include the title and document number of the publication with your comments. (Online, the document number is located in the front matter of the publication. In printed publications, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Use the Feedback option on the Technical Publications Library Web page:
<http://docs.sgi.com>
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.

- Send mail to the following address:

Technical Publications
SGI
1500 Crittenden Lane, M/S 535
Mountain View, California 94043-1351

SGI values your comments and will respond to them promptly.

Features and Capabilities of the SGI 10-Gbit Ethernet Network Adapter

The 10-Gbit Ethernet technology is an extension of Gigabit Ethernet (1000-Base-T) technology that allows over-the-wire speeds up to 10 Gbits per second (Gbps), which is theoretically ten times the rate of existing technology. Given the limitations of the bus and the CPUs, realistic performance with the SGI 10-Gbit Ethernet network adapter will not exceed 4.5–6 Gbps.

The 10-Gbit Ethernet technology is targeted at backbone networks and interserver connectivity. It provides an upgrade path for high-end workstations that require more bandwidth than Gigabit Ethernet can provide.

The adapter is supported in the following systems:

- Silicon Graphics Prism
- SGI Altix 3000
- SGI Altix 350
- SGI Origin 350
- Silicon Graphics Onyx 4
- Silicon Graphics Onyx 350
- SGI Origin 3000 with IX brick or PX brick
- Silicon Graphics Onyx 3000 with IX brick or PX brick
- Silicon Graphics Tezro

This chapter includes the following sections:

- "Features" on page 2
- "Cabling" on page 2
- "Configuration Limits" on page 3
- "Tools" on page 3

Features

The adapter includes the following features:

- Full-duplex Gigabit Ethernet interface as defined in the IEEE P802.3ae approved standard
- Jumbo frame support (up to 9600 bytes)
- 16 DMA channels (8 for reading, 8 for writing)
- Single LC fiber connection
- Support for PCI 64-bit 66-MHz and PCIX 66/100/133-MHz
- Universal dual-voltage signaling (3.3 V and 5 V)
- Compliance with PCI Local Bus Specification (Revision 2.3) and PCI-X Addendum to the PCI Local Bus Specification (Revision 1.0a) standards
- Optical cabling 850-nm (10GBASE-SR)

Cabling

The adapter is an LC connector and uses a 10GBASE-SR transceiver at 850-nm. It is connected to the network using a multimode fiber (MMF) cable of 50-micron. The cable (which is not included in the shipment) must have a quality in the range from 400-MHz * km to 2000-MHz * km, depending on its length, as shown in Table 1-1.

Table 1-1 10-Gbit Cable Standards

Diameter (Microns)	Modal Bandwidth (MHz * km)	Range (Meters)
50	400	2 to 66
50	500	2 to 82
50	2000	2 to 300

Configuration Limits

The number of the 10-Gbit Ethernet network adapters supported varies by system. Consult with your SGI representative to determine the currently supported maximum for your configuration.

Tools

The SGI 10-Gbit Ethernet network adapter product contains the following IRIX tools:

- `xg(7F)` driver
- `xgcmd(1M)` diagnostic tool to interrogate and control both the driver itself and the interfaces under its control

The following standard Linux commands are useful with any Ethernet device:

- `ethtool(8)`
- `ifconfig(8)`
- `ip(8)`

For more information, see the man page associated with each tool.

Connecting the Adapter to a Network

This chapter discusses the following:

- "Installing the Adapter"
- "Connecting to the Network"

For information about configuring the interfaces, see *IRIX Admin: Networking and Mail* and *The Network Administrators' Guide*.

Installing the Adapter

The installation instructions for the SGI 10-Gbit Ethernet network adapter are different for different systems. Refer to the following sources for installation instructions:

- Altix 3000, Onyx 3000, or Origin 3000: The adapter must be installed by an SGI certified service provider.
- Origin 350, Onyx 350, and Altix 350: See the instructions for installing a PCI card in the user's or owner's guide that came with the system.
- Silicon Graphics Tezro: See the instructions for installing a PCI card in the user's or owner's guide that came with the workstation.

Connecting to the Network

To connect the adapter to a network, do the following:

1. Remove the protective end caps and **save them**.



Caution: 10-Gbit optics are very sensitive. If you plan on leaving them disconnected for any length of time, you must replace the end caps. The optics on the SGI 10-Gbit Ethernet network adapter cannot be cleaned.

2. Insert the LC connector on one end of the fiber-optic cable into the adapter, as shown in Figure 2-1. Ensure that the connector is inserted completely into the jack.

Note: If the network connects to an Ethernet switch, consult the operating manual for the switch to ensure that the switch port is enabled and configured correctly and as the correct adapter type (10GBASE-SR).

3. Insert the connector on the other end of the fiber-optic cable into the connector on the Ethernet switch, or another computer system (as appropriate).

Figure 2-1 shows the connector.




Figure 2-1 Fibre Optic Connections

Operating the Adapter

This chapter describes various issues that may occur when using the SGI 10-Gbit Ethernet network adapter in a 10-Gbit Ethernet network. It includes the following sections:

- "Verifying Functionality"
- "Resetting the Adapter" on page 11
- "Configuration" on page 12
- "MTU Size Settings" on page 12

Verifying Functionality

This section explains the following:

- "Using LEDs to Determine Functionality"
- "Verifying Adapter Recognition" on page 8
- "Enabling the Adapter" on page 10
- "Verifying that the Adapter is Properly Configured and Enabled" on page 10

Using LEDs to Determine Functionality

The SGI 10-Gbit Ethernet network adapter has light-emitting diodes (LEDs) that indicate whether the adapter is configured correctly and connected to an active Ethernet. Table 3-1 describes the functions of the LEDs.

Table 3-1 LEDs on the SGI 10-Gbit Ethernet Network Adapter

LED	State	Purpose
ACT	Blinking	Data detected
	Off	No data detected
LINK	On	Good link
	Off	No link: faulty cable, faulty connector, or communication mismatch

During normal operation, the link LED is on; the data LED blinks whenever the adapter is receiving traffic.

Note: At the time of release, a bug in the driver prevents the link light from behaving correctly. The link light may be on or off due to this bug. This problem will be corrected in a future release.

Verifying Adapter Recognition

IRIX Adapter Recognition Verification

Under IRIX, the network interface name for the SGI 10-Gbit Ethernet network adapter is `xgN`, where `N` is the number of the adapter: 1 for the first adapter, 2 for the second adapter (if installed), and so on. Use the commands in the examples below to display the network interface names.

Note: IRIX uses `/etc/ioconfig.conf` to keep adapters and unit numbers consistent between reboots and new hardware installs. As a result, if a adapter is removed, it's interface name will be reserved until it is removed from `/etc/ioconfig.conf`. As a result, if you move adapters around the system or remove adapters, you may see some interface numbers get skipped.

To verify that the operating system has located the adapter, enter the following:

```
irix% /bin/hinv
```

A line similar to the following should appear, where `N` is the number of the adapter:

```
10 Gigabit Ethernet: xgN, module 001c02, 66 MHz PCI bus 2 slot 1
```


For example, the first adapter would appear as `xg1`:

```
10 Gigabit Ethernet: xg1, module 001c02, 66 MHz PCI bus 2 slot 1
```

Linux Adapter Recognition Verification

Unlike IRIX, under Linux all Ethernet adapters all share the prefix `ethN`.

To verify that the adapter has been recognized, do the following:

1. Use the `lspci` command to ensure that the device has been recognized:

```
[root@linux root]# /sbin/lspci
.....
02:02.0 Ethernet controller: S2io Inc.: Unknown device 5831 (rev 03)
```

Note: At this point (before the module is loaded), the `/proc/net/dev` file will not display the device.

On SGI ProPack 3 systems, you can also use the `hinv` command to display device information.

2. Load the `s2io` module:

```
[root@linux root]# insmod s2io
```

The device will now appear in the `/proc/net/dev` file:

```
[root@linux root]# cat /proc/net/dev | grep eth1
eth1:      0          0  0  0  0  0          0          0  ....
[root@linux root]#
```

The device also appears in `ifconfig` output as well.

To make this module load at system boot time, enter the following line into the `modules.conf` file:

```
alias eth1 s2io
```

3. Use the `ifconfig` command to show the new `ethN` interface. For example:

```
[root@linux root]# ifconfig eth1
eth1      Link encap:Ethernet  HWaddr 08:00:69:14:72:AD
          BROADCAST MULTICAST  MTU:1500  Metric:1
```

```
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
Interrupt:62
```

Enabling the Adapter

To enable the adapter, enter the following:

```
# ifconfig IPaddress broadcast broadcastaddress netmask netmaskvalue
```

For other systems to see the new address, you must enter the new hosts addresses in DNS or in host files or NIS as required for your system.

For details, see the operating-system specific documentation about networking.

Verifying that the Adapter is Properly Configured and Enabled

IRIX Adapter Configuration Verification

To verify that the network interface is configured properly and is enabled on, enter the following on an IRIX system:

```
irix% /usr/etc/netstat -ia
```

Columns with the following headings should appear:

```
Name Mtu Network Address
```

In the Name column, the xg number should appear. If it is followed by an asterisk (*), the interface is disabled.

In the Mtu column, the size of the current maximum transmission unit (MTU) should appear. The MTU size is set via the -mtu switch of the ifconfig command. If no size is specified by the -mtu switch, the adapter defaults to an MTU size of 1500.

In the Network column, the IP network address should appear.

In the Address column, the canonical MAC address of the adapter should appear, which looks similar to the following:

```
08:00:69:0b:e0:41
```

In this address, the organizationally unique identifier (OUI) of the adapter vendor is represented by the first three sets of numbers (for example, 08:00:69). The last three sets vary, depending on the system.

See the `netstat` man page for more details.

Linux Adapter Configuration Verification

To verify that the network interface is configured properly and is enabled on, enter the following on a Linux system:

```
ifconfig -a
```

For example, the following shows that the MAC address is 00:0C:FC:00:01:24, the IP address is 10.2.2.2, the card is up, and the MTU is 9600:

```
[root@linux root]# ifconfig -a
eth1      Link encap:Ethernet  HWaddr 00:0C:FC:00:01:24
          ^^^^^^^^^^^^^^^^^^^^^
          MAC address
          inet addr:10.2.2.2  Bcast:10.2.2.255  Mask:255.255.255.0
          ^^^^^^^
          IP address
          UP BROADCAST RUNNING MULTICAST  MTU:9600  Metric:1
          ^^                               ^^^
          the card is up                    MTU
          RX packets:5 errors:0 dropped:0 overruns:0 frame:0
          TX packets:5 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:414 (414.0 b)  TX bytes:322 (322.0 b)
          Interrupt:59
```

Resetting the Adapter

In the unlikely event that you need to reset the adapter, enter the following, where *N* is the adapter number:

- IRIX:

```
# ifconfig xgN down
# ifconfig xgN up
```

- Linux:

```
# ifconfig ethN down  
# ifconfig ethN up
```

Configuration

IRIX and Linux use different configuration tools.

IRIX Configuration

The configuration of the adapter is controlled by a number of parameters, which are discussed in the IRIX `/etc/config/xgcmd.options` file.

To change the configuration of your adapter, use the `xgcmd` command. For more information, see the `xgcmd` man page.

Linux Configuration

To change the configuration of your adapter, use the `ethtool(8)` command. For more information, see the `ethtool(8)` man page.

MTU Size Settings

The maximum transmission unit (MTU) size is controlled by the `mtu mtu_size` switch of the `ifconfig` command. The most common MTU sizes are 1500 bytes (standard-size Ethernet frames) and 9000 (jumbo Ethernet frames). The adapter supports an MTU size of up to 9600 bytes. Configuring the adapter to use jumbo frames will increase network bandwidth and reduce CPU load, but only if the network supports jumbo frames.

To configure the MTU size, follow these steps:

1. At the command prompt, enter the following command:

```
netstat -i
```

The output will display information about the network adapters currently installed in the system, as follows:

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
xg1	1500	10.50.1	system.name	0	0	1	0	0

2. To bring down the adapter, enter the following command, where *N* is the number of the adapter:

```
ifconfig xgN down
```

3. To change the MTU size of the 10-Gbit Ethernet adapter, enter the following command, where *N* is the number of the adapter:

```
ifconfig xgN mtu 9600
```

4. To restart the adapter, enter the following:

```
ifconfig xgN up
```

5. Enter the following command to verify that the MTU size has been changed:

```
netstat -i
```

The output should look similar to the following:

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
xg1	9600	10.50.1	system.name	0	0	3	0	0

Performance Tuning

This chapter discusses performance tuning topics for IRIX and SGI ProPack 3 for Linux.

For good TCP performance, the socket buffers used by applications must be at least as large as the bandwidth-round trip delay product between the two endpoints. Normally, larger socket buffers are called for with 10-Gbit Ethernet than when lower-bandwidth network interface cards are used.

IRIX Performance Tuning

This section discusses the following:

- "Jumbo Frames and IRIX" on page 15
- "Read/Write Size and IRIX" on page 16
- "Socket Buffer Size and IRIX" on page 16
- "Multibuffer Mode and IRIX" on page 16

There is no one tuning configuration that is best for all environments. This section gives guidelines for the following cases:

- A *point-to-point configuration*, in which IRIX is transmitting to or receiving from IRIX
- A *multiclient configuration*, in which an IRIX system is connected via 10-Gbit Ethernet to a switch that fans out to multiple clients via 1-Gbit Ethernet

Jumbo Frames and IRIX

In general, an MTU of 9000 bytes (known as *jumbo frames*) gives the best TCP throughput performance and scaling results. Avoid an MTU of 1500 bytes if possible.

Read/Write Size and IRIX

In a point-to-point configuration, use a buffer length of 513920 bytes.

In a multiclient configuration, use a buffer length of 64240 bytes.

Socket Buffer Size and IRIX

Socket buffer size is set either:

- By the application
- Via the `tcp_sendspace` and `tcp_recvspace` tunable parameters

In a point-to-point configuration, use one of the following socket buffer sizes shown in Table 4-1.

Table 4-1 Socket Buffer Size in a Point-to-Point Configuration

MTU	Socket Buffer Size
1500 bytes	2048 KB
9000 bytes	4096 KB

In a multiclient configuration, use a socket buffer size of 2048 KB regardless of the MTU size.

`tcp_delwake_count` is a tunable parameter used to delay the wakeup of the receiving process for TCP input segments until a certain amount of data has been placed on the socket queue. This count is the number of bytes to be accumulated in the socket receive buffer before the receiving process is awakened. Change this value according to the configuration:

- In a point-to-point configuration, use 5840
- In a multiclient configuration, use 11680

Multibuffer Mode and IRIX

The IRIX multibuffer mode is dynamically enabled when the MTU is greater than 1500 bytes.

SGI ProPack 3 Performance Tuning

This section discusses the following:

- "Socket Read and Write Buffer Sizes for SGI ProPack 3" on page 17
- "Jumbo Frames and SGI ProPack 3" on page 17
- "Read/Write Size and SGI ProPack 3" on page 17
- "TCP/IP Socket Buffer Size and SGI ProPack 3" on page 18

Socket Read and Write Buffer Sizes for SGI ProPack 3

The largest-allowed socket read and write buffer sizes are controlled by the following files:

- Read: `/proc/sys/net/core/rmem_max`
- Write: `/proc/sys/net/core/wmem_max`

Jumbo Frames and SGI ProPack 3

Using a large maximum transmission unit (MTU) is necessary for the best 10-Gbit Ethernet performance. Generally, the bigger the MTU, the better. The driver supports MTUs as large as 9600 bytes.

Using a standard 1500-byte MTU will result in 10-Gbit Ethernet performance that is, at best, only about 50% higher than Gigabit Ethernet.

Read/Write Size and SGI ProPack 3

Applications should read large buffers from and write large buffers to the network for the best throughput and to reduce CPU utilization.

For example, an application that uses `recv(2)` calls with 32-KB buffers will generally have better throughput than if the application were to use twice as many `recv` calls with 16-KB buffers.

TCP/IP Socket Buffer Size and SGI ProPack 3

In SGI ProPack 3, `/proc/sys/net/core/rmem_max` and `/proc/sys/net/core/wmem_max` are both set to 524288 bytes, which is usually large enough to provide good performance.

If you to adjust the socket buffers, use `sysctl(8)` command.

Glossary

CD-ROM (CD)

A flat metallic disk that contains information that you can view and copy onto your own hard disk; you cannot change or add to the disk. CD-ROM is an abbreviation for compact disc read-only memory.

Ethernet

A communication network used to connect computers.

gigabit (Gbit)

A communication rate of 2^{30} bits per second.

host

Any system connected to the network.

hostname

The name that uniquely identifies each host (system) on the network.

IP address

A number that uniquely identifies each host (system) on a TCP/IP network.

LED

Light-emitting diode, a light on a piece of hardware that indicates status or error conditions.

MAC

Medium access control, also called the physical layer.

MAC address

The physical address of the SGI 10-Gbit Ethernet Network adapter, which is distinct from the IP address.

MTU

Maximum Transmission Unit is a configuration parameter that controls the size of the Ethernet frames that the SGI 10-Gigabit Ethernet network adapter can transmit and receive.

man (manual) page

An online document that describes how to use a particular IRIX command. Also called reference page.

multiclient configuration

A TCP/IP configuration in which the IRIX system is connected via 10-Gbit Ethernet to a switch that fans out to multiple clients via 1-Gbit Ethernet.

NIS

Network Information Service, a distributed database mechanism for user accounts, host names, mail aliases, and so on.

PCI

Peripheral Component Interconnect, a bus specification. The PCI bus is a high-performance local bus used to connect peripherals to memory and a microprocessor. Many vendors offer devices that plug into the PCI bus.

point-to-point configuration

A TCP/IP configuration in which IRIX is transmitting to or receiving from IRIX.

reference page

See *man (manual)* page.

TCP/IP

A standard networking protocol that is included in the IRIX software.

Index

1000-Base-SX, 2

A

adapter enabling, 10
adapter recognition, 8
Altix systems, 1

C

cabling, 2
capabilities, 1
configuration limits, 3
connector, 2, 6

D

diagnostics, 3
DMA channels, 2
driver, 3

E

enabling the adapter, 10, 11
/etc/ioconfig.conf, 8

F

features, 1
fiber connection, 2
fibre optic connections, 6
fibre type, 2

007-4669-002

frame sizes, 2
frames, 12
functionality verification, 7

H

hinv, 8

I

IEEE standards, 2
ifconfig, 10, 12
installation, 5
interface names, 8
IP address, 10
IRIX performance tuning, 15

J

jumbo frames
 IRIX performance tuning, 15
 MTU sizes supported, 12
 SGI ProPack 3 performance tuning, 17
 sizes supported, 2

L

LC connector, 2, 6
LC fiber connection, 2
LEDs, 7

M

- MAC address, 10
- MMR, 2
- MTU sizes, 12, 15
- multibuffer mode and tuning, 16
- multiclient configuration, 15

N

- netstat, 10–12
- network connection, 5
- network interface names, 8

O

- Onyx systems, 1
- operating, 7
- Origin systems, 1

P

- PCI-X bus interface, 2
- performance tuning
 - IRIX, 15
 - SGI ProPack 3, 17
- point-to-point configuration, 15
- ports, 2
- /proc/sys/net/core/rmem_max, 19
- /proc/sys/net/core/wmem_max, 19

R

- read/write size
 - IRIX performance tuning, 16
 - SGI ProPack 3 performance tuning, 17
- resetting the adapter, 11
- rmem_max, 19

S

- SGI ProPack 3 performance tuning, 17
- signaling, 2
- socket buffer size
 - IRIX performance tuning, 16
 - SGI ProPack 3 performance tuning, 17
- standard-size Ethernet frames, 12
- supported systems, 1
- sysctl, 18

T

- TCP/IP and tuning, 18
- Tezro systems, 1
- tools, 3
- troubleshooting with LEDs, 7

U

- universal dual-voltage signaling, 2

V

- verifying adapter recognition, 8
- verifying functionality, 7

W

- wmem_max, 19

X

- xg driver, 3
- xgcmd diagnostic tool, 3