

SGI ProPack™ for Linux® Start Here

007-4558-004

CONTRIBUTORS

Written by Julie Boney

Edited by Susan Wilkening

Graphics by Chrystie Danzer

Production by Glen Traefald

COPYRIGHT

© 2003 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 USA 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, IRIX, Origin, Onyx, Onyx2, and XFS are registered trademarks and Altix, NUMAflex, NUMAlink, OpenMP, Performance Co-Pilot, SGI Advanced Linux, SGI ProPack, and SGIconsole are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

SGI Advanced Linux Environment 2.1 is based on Red Hat Linux Advanced Server 2.1 for the Itanium Processor, but is not sponsored by or endorsed by Red Hat, Inc. in any way.

Cray is a registered trademark of Cray, Inc. FLEXlm is a registered trademark of Macrovision Corporation. Java is a registered trademark of Sun Microsystems, Inc. in the United States and/or other countries. KAP/Pro Toolset and VTune are trademarks and Intel, Itanium, and Pentium are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Linux is a registered trademark of Linus Torvalds, used with permission by Silicon Graphics, Inc. MIPS is a registered trademark of MIPS Technology, Inc. PostScript is a registered trademark of Adobe Systems, Inc. QLogic is a registered trademark of QLogic Corporation. Red Hat is a registered trademark and Red Hat Linux Advanced Server 2.1 and RPM are trademarks of Red Hat, Inc. Windows is a registered trademark of Microsoft Corporation in the United States and/or other countries.

Cover Design By Sarah Bolles, Sarah Bolles Design, and Dany Galgani, SGI Technical Publications.

Record of Revision

Version	Description
001	February 2003 Original publication
002	February 2003 Updated for SGI ProPack v2.1.1 for Linux release
003	April 2003 Updated for SGI ProPack v2.1.2 for Linux release
004	June 2003 Updated for SGI ProPack v2.2 for Linux release

New Features in This Guide

This update of *SGI ProPack for Linux Start Here* supports the 2.2 release of SGI ProPack for Linux.

Major Documentation Changes

Changes in this guide include the following:

- A change has been made in the Linux base distribution to a distribution based on Linux Red Hat Advanced Server 2.1. This new base distribution is the SGI Advanced Linux Environment 2.1.
- The kernel is now installed to `/boot/efi/EFI/sgi/`. This change is part of the upgrade to elilo v3.3. For more details about this change, see “Extensible Firmware Interface (EFI)” in Chapter 3.
- The Updates and Boot CDs have been eliminated from the list of CDs provided, and their functionality folded into the Installation CDs.
- A new section, “Upgrading from Previous SGI ProPack Releases to SGI ProPack 2.2” in Chapter 2 has been added.

Contents

Record of Revision	iii
Figures	xi
Tables	xiii
Related Publications	xvi
Intel Compiler Documentation	xvi
Other Intel Documentation	xvi
SGI Documentation	xvi
Obtaining SGI Publications	xviii
Conventions	xviii
Reader Comments	xviii
1. Release Features	1
Software Introduction	1
CD Contents	2
Installation Overview	4
2. Software Planning and Installation	5
Software Planning	5
Table of Disk Partitions	6
Linux Kernel Crash Dump	6
Installing the SGI Advanced Linux Environment 2.1 Base OS	6
Upgrading from Previous SGI ProPack Releases to SGI ProPack 2.2	10
Installing SGI ProPack for Linux	11
Installing the System Controller Software 1.1 CD	12
Upgrading Your Software	13
Recovering a Damaged Root Filesystem	13
Manually Mounting the Root Filesystem	15

3.	Product Support	.17
	Operating System Enhancements.	.21
	CpuMemSets Support.	.21
	Cpuset Support	.22
	Comprehensive System Accounting (CSA).	.22
	Partitioning.	.23
	I/O Subsystems.	.23
	Persistent PCI-X Bus Numbering	.23
	Persistent Naming of Ethernet Devices	.25
	XSCSI Subsystem	.27
	XSCSI-SCSI Subsystem	.27
	XFS Filesystem	.27
	XVM Volume Manager	.28
	HPC Application Tools and Support.	.28
	Message Passing Toolkit	.28
	Performance Co-Pilot	.29
	System Management	.29
	PROM Chips	.29
	Extensible Firmware Interface (EFI)	.29
	FLEXlm	.32
	SGIconsole	.32
	System Controller Firmware	.33
	NUMA Tools	.33
	dlook Command	.34
	dplace Command	.34
4.	Performance Tuning	.35
	Determining System Configuration	.35
	Single Processor Code Tuning	.36
	Getting the Correct Results	.37
	Using ddd	.38
	Managing Heap Corruption Problems	.39
	Using Tuned Code.	.40

Determining Tuning Needs	41
Using Compiler Options Where Possible	41
Tuning the Cache Performance	42
Managing Memory	43
Using <code>chattr</code> to Change Stack and Data Segments	43
Multiprocessor Code Tuning	44
Parallelizing Your Code	44
Solving Bottlenecks in Code	46
Fixing False Sharing	46
Using <code>dplace</code> and <code>runon</code>	47
<code>dplace</code> for MPI Codes	47
<code>dplace</code> for OpenMP Codes	48
Environment Variables for Performance Tuning	48
Profilers and Performance Tools	49
Profiling with <code>pfmon</code>	49
The <code>profile.pl</code> script	50
<code>profile.pl</code> with MPI programs	50
Using VTune for Remote Sampling	51
Using GuideView	51
Other Performance Tools	52
Index	55

Figures

Figure 1-1	SGI ProPack for Linux Release CD Contents	3
-------------------	---	---

Tables

Table 2-1	devfs Disk Partitions	6
Table 3-1	SGI ProPack v2.2 for Linux Products	18
Table 3-2	EFI Commands	30

About This Guide

This guide provides information about the SGI ProPack for Linux release. It is divided into the following chapters:

- Chapter 1, "Release Features," describes the major features of this release, the CD contents, and documentation information.
- Chapter 2, "Software Planning and Installation," describes the flowcharts of disk partitions and file configurations and layouts that you need to get started, and provides instructions for installing SGI ProPack for Linux.
- Chapter 3, "Product Support," documents the product components that are supported on SGI Altix 3000 systems.
- Chapter 4, "Performance Tuning," documents tuning issues for single processor and multiprocessor programs.

The information in this guide, other SGI ProPack for Linux documentation, and all other documentation included in the RPMs on the distribution CDs can be found on the CD titled "SGI ProPack v2.2 for Linux - Documentation CD." To access the information on the documentation CD, open the `index.html` file with a web browser. Because this online file can be updated later in the release cycle than this document, you should check it for the latest information.

Note: The release notes, which contain the latest information about software and documentation in this release, are on the SGI ProPack for Linux Documentation CD in the `root` directory, in a file named `README.TXT`.

Related Publications

Documents listed in this section contain additional information that might be helpful.

Intel Compiler Documentation

Documentation for the Intel compilers is located on your system in the `/docs` directory of the directory tree where your compilers are installed. If you have installed the Intel compilers, the following documentation is available:

- *Intel C++ Compiler User's Guide* (`c_ug_lnx.pdf`).
- *Intel Fortran Compiler User's Guide* (`for_ug_lnx.pdf`).
- *Intel Fortran Programmer's Reference* (`for_prgr.pdf`).
- *Intel Fortran Libraries Reference* (`for_lib.pdf`).

Other Intel Documentation

The following references describe the Itanium (previously called "IA-64") architecture and other topics of interest:

- *Intel Itanium 2 Processor Reference Manual for Software Development and Optimization*, available online at <http://developer.intel.com/design/itanium/manuals>.
- *Intel Itanium Architecture Software Developer's Manual*, available online at <http://developer.intel.com/design/itanium/manuals>.
- *Introduction to Itanium Architecture*, available online at <http://shale.intel.com/softwarecollege/CourseDetails.asp?courseID=13> (secure channel required).

SGI Documentation

The following SGI documentation is available:

- *Linux Device Driver Programmer's Guide*
Provides information on programming, integrating, and controlling drivers.

- *Message Passing Toolkit: MPI Programmer's Manual*
Describes industry-standard message passing protocol optimized for SGI computers.
- *Origin 2000 and Onyx2 Performance Tuning and Optimization Guide*
Contains information specific to MIPS/IRIX systems, but the general guidelines in the document are hardware and operating system independent.
- *Performance Co-Pilot for IA-64 Linux User's and Administrator's Guide*
Describes the Performance Co-Pilot (PCP) software package of advanced performance tools for SGI systems running the Linux operating system.
- *Linux Resource Administration Guide*
Provides a reference for people who manage the operation of SGI computer systems running the Linux operating system.
- *SGI Altix 3000 User's Guide*
Provides an overview of the architecture and describes the major components of the SGI Altix 3000 family of servers and superclusters. It also describes the standard procedures for powering up and powering down the system, provides basic troubleshooting information, and includes important safety and regulatory specifications.
- *SGI ProPack v2.2 for Linux Release Notes*
Provide the latest information about software and documentation in this release. The release notes are on the SGI ProPack for Linux Documentation CD in the `root` directory, in a file named `README.TXT`.
- *SGIconsole 2.0 Start Here*
Provides an introduction to SGIconsole.
- *SGI L1 and L2 Controller Software User's Guide*
Describes how to use the L1 and L2 controller commands at your system console to monitor and manage the SGI Altix 3000 family of servers and superclusters.
- *XFS for Linux Administration*
Describes XFS, an open-source, fast recovery, journaling filesystem that provides direct I/O support, space preallocation, access control lists, quotas, and other commercial file system features.

Obtaining SGI Publications

You can obtain SGI documentation in the following ways:

- 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.
- You can view man pages by typing `man <title>` on a command line.

Conventions

The following conventions are used throughout this publication:

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.)
<code>manpage(x)</code>	Man page section identifiers appear in parentheses after man page names.
GUI element	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, 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
1600 Amphitheatre Parkwy, M/S 535
Mountain View, CA 94043-1351
- Send a fax to the attention of “Technical Publications” at +1 650 932 0801.

SGI values your comments and will respond to them promptly.

Release Features

This chapter provides an introduction to SGI ProPack for Linux and describes the CD contents, provides documentation information, and provides a brief overview of the installation process.

Software Introduction

The open-source and collaborative development environment of the Linux community provides a new and different model for computer manufacturers to deliver an operating system for a large server or supercomputer. Using Linux as the operating system for very large computer systems has the benefit of providing better software protection and better integration between a system manufacturer's operating system and high-performance computing (HPC) applications and codes for both independent software vendors and end users. Also, the collaboration on OS improvements in the Linux community as well as among the various computer manufacturers has enabled Linux to evolve and improve in multiple dimensions faster than would otherwise be possible for a single company working on its own operating system.

The SGI ProPack for Linux product includes capabilities and performance improvements ideal for enabling technical and creative users to solve their big compute and data problems by using the Linux operating system and Itanium processors. This product adds to or enhances features in base Linux distributions based on Red Hat Linux Advanced Server 2.1. SGI ProPack for Linux is designed to run on any SGI Altix 3000 series system. SGI hardware platforms and OS configuration settings supported by SGI in this release are documented at the following URL:

<http://support.sgi.com/linux>

CD Contents

The CDs that you receive with SGI ProPack for Linux are as follows (Figure 1-1 lists the contents of each CD):

- SGI Advanced Linux Environment 2.1
 - SGI Advanced Linux Environment 2.1 CD1 - Installation CD 1 of 2
 - SGI Advanced Linux Environment 2.1 CD2 - Installation CD 2 of 2
 - SGI Advanced Linux Environment 2.1 CD3 - Source Code CD 1 of 2
 - SGI Advanced Linux Environment 2.1 CD4 - Source Code CD 2 of 2
- SGI ProPack v2.2 for Linux
 - SGI ProPack v2.2 for Linux CD1 - Open/Free Source Software
 - SGI ProPack v2.2 for Linux CD2 - Proprietary Software

Note: This CD contains software for use on SGI systems only. Redistribution is not permitted. Please see the shrink-wrap license agreement.

- SGI ProPack v2.2 for Linux CD3 - Documentation CD
- System Controller Software 1.1 CD (for IRIX and Linux)

Note: This software is licensed for use only on Origin and Onyx 3000 series systems (SGI systems based on MIPS processors) and on SGI Altix 3000 series systems (SGI systems based on Intel Itanium processors). This CD contains software provided under restricted or SGI proprietary licensing terms. The licensing terms for items on this CD allow users to install this software only on SGI systems.

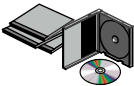
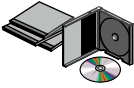
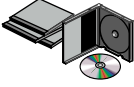
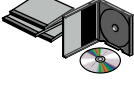
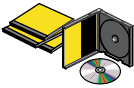
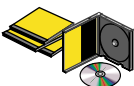

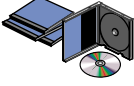
SGI Advanced Linux Environment 2.1		
	SGI Advanced Linux Environment 2.1 Installation CD 1 of 2 CD1	Interactive installation procedures
	SGI Advanced Linux Environment 2.1 Installation CD 2 of 2 CD2	
	SGI Advanced Linux Environment 2.1 Source Code CD 1 of 2 CD3	Source code for base OS and common open source applications
	SGI Advanced Linux Environment 2.1 Source Code CD 2 of 2 CD4	
SGI ProPack v2.2 for Linux		
	SGI ProPack v2.2 for Linux Open/Free Source Software CD1	GPL/LGPL and other open source licensed software, latest SGI platform and NUMA support, various kernel performance improvements, CpuMemSets and performance measuring tools, XSCSI infrastructure, QL-SCSI, XFS, CSA, LKCD, kdb
	SGI ProPack v2.2 for Linux Proprietary Software CD2	System PROM, XVM, cpuset and NUMAtools (dlook, dplace, and so on), Array Services, FLEXIm, system partitioning software, MPT
	SGI ProPack v2.2 for Linux Documentation CD CD3	Collection of documentation, FAQs, HOWTOs, and man pages
System Controller Software 1.1 CD		
	System Controller Software 1.1 CD	System Controller software/firmware (for IRIX and Linux)

Figure 1-1 SGI ProPack for Linux Release CD Contents

For a complete list of the RPMs included on the SGI Advanced Linux Environment 2.1 and SGI ProPack v2.2 for Linux CDs listed in Figure 1-1 on page 3, see the release notes, which contain the latest information about software and documentation in this release. The release notes are on the SGI ProPack for Linux Documentation CD in the `root` directory, in a file named `README.TXT`.

SGI maintains the following website for open source information that describes projects related to its open source efforts:

<http://oss.sgi.com>

You can also access open source information (such as LKCD) from the following website:

<http://sourceforge.net>

Installation Overview

SGI ProPack for Linux will most likely come preinstalled on your SGI platform. If you should need to install it, be aware that you must first install the SGI Advanced Linux Environment 2.1 CDs, then the SGI ProPack v2.2 Open/Free Source Software (CD1), and finally, the SGI ProPack v2.2 Proprietary Software (CD2). Installation is described in Chapter 2, “Software Planning and Installation”.

Note: There is no Updates CD for the v2.2 release. All updates are being folded back into the SGI Advanced Linux Environment 2.1 CDs.

Software Planning and Installation

This chapter describes the planning that you need to do to get started and provides instructions for installing SGI ProPack for Linux.

Your SGI Altix 3000 system comes with a base Linux distribution (SGI Advanced Linux Environment 2.1) and the SGI ProPack software preinstalled. This chapter describes how to install the software from the CD if it should ever become necessary to reinstall it.

For security reasons, Linux requires a root password for login. The default password for your preinstalled software is `sgisgi`. After you have logged in, change this root password to a string of your own choice.

The SGI ProPack software works only with the SGI AdvancedLinux Environment 2.1. Earlier versions of these distributions, or any other distributions, are not compatible with SGI ProPack software.

Before you install or configure your system, please read Chapter 1, “Release Features” so that you understand the features of SGI ProPack for Linux software and how to configure them. You might also want to familiarize yourself with some or all of the documentation listed in “Related Publications” on page xvi.

Software Planning

This section provides disk and file information you need to have about the software that has been shipped to you.

Table of Disk Partitions

SGI ships systems from the factory with the partition layout shown in Table 2-1.

Table 2-1 devfs Disk Partitions

Device	Mount Point	Size
/dev/xscsi/pci01.03.0-1/target1/lun0/part1	/boot/efi	500 MB
/dev/xscsi/pci01.03.0-1/target1/lun0/part2	swap	9 GB
/dev/xscsi/pci01.03.0-1/target1/lun0/part3	/	25 GB

For more information on device naming, see “Persistent PCI-X Bus Numbering” on page 23 and “XSCSI Subsystem” on page 27.

Linux Kernel Crash Dump

Linux Kernel Crash Dump (LKCD) creates files in `/var/dump`. The `/var` directory size will grow as Performance Co-Pilot (PCP) and Comprehensive System Accounting (CSA) log to their own directories in `/var`. These files do not usually exceed 100 MB. For default LKCD configuration details, see the `lkcd_config(1)` man page.

Note: If there is an `lkcd` dump, LKCD could easily exceed multiple gigabytes in `/var`. This is why the default root filesystem is larger. For this reason, some customers may wish to make a separate `/var/dump` filesystem or change the configuration of `lkcd`.

Installing the SGI Advanced Linux Environment 2.1 Base OS

This section provides steps for installing the SGI Advanced Linux Environment 2.1 CD1 and CD2 set. For console setup options and instructions, see the section titled “Monitoring Your Server” in the *SGI Altix 3000 User Guide*.

Note: The installation procedure has buttons that allow you to go back to the previous screen or to quit the installation. To use these buttons, press the Tab key to highlight the one you want to use and press Enter.

1. Insert SGI Advanced Linux Environment 2.1 CD1 into the system's CD-ROM drive and restart the system.
2. While the system is powering up, check the device mapping table. It should look similar to the following:

Device mapping table

```
fs0 : Pci(2|1)/Ata(Primary, Master)/CDROM(Entry1) This is the CD-ROM
fs1 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part1, Sig00000000)
fs2 : Pci(1|1)/Scsi(Pun0/Lun2)/HD(Part1, Siggl)
blk0 : Pci(2|1)/Ata(Primary, Master)
blk1 : Pci(2|1)/Ata(Primary, Master)/CDROM(Entry1)
blk2 : Pci(1|1)/Scsi(Pun0/Lun1)
blk3 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part1, Sig00000000)
blk4 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part2, Sig00000000)
blk5 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part3, Sig00000000)
blk6 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part4, Sig00000000)
blk7 : Pci(1|1)/Scsi(Pun0/Lun2)
blk8 : Pci(1|1)/Scsi(Pun0/Lun2)/HD(Part1, Siggl0)
```

3. At the Shell> prompt, type the CD-ROM device name, as follows:

```
Shell> fs0:      Type this to change to the CD-ROM device
fs0:\>
```

4. Boot the CD by typing **elilo** and pressing Enter. When the ELILO boot : prompt appears, press Enter again.

```
fs0:\> elilo      This is the boot loader
LoadPe: using PE image entry point
LoadPe: using PE image entry point
ELILO boot:      Press ENTER to boot the Linux kernel.
```

5. The CD will then boot and start the installation process.
6. Select the language you would prefer to use for the installation and as the system default. To continue, select **OK**.

Note: Console support for Asian languages is not available at this time. SGI recommends that customers in Asia set the default system language to English.

7. The next screen lets you choose your installation method. Currently, **Local CDROM** is the only choice. To continue, select **OK**.
8. The next screen is the SGI ProPack information screen. To continue, select **OK**.
9. The next screen asks you if you have a mouse installed. For an SGI Altix 3000 system, select **None**. To continue, select **OK**.
10. You will now see a window titled **Installation type**. If you wish to upgrade, select **Upgrade Existing System** and skip now to “Upgrading from Previous SGI ProPack Releases to SGI ProPack 2.2” on page 10 for further instructions. Otherwise, to completely install the system, wiping out any previous data on the disk, choose **Custom** and continue to the next step in this section.
11. In the partition screen, select the **Autopartition** option and select **OK**. When warning messages appear, select **IGNORE**.
12. When asked which partitioning scheme to use, select the **Remove all partitions on this system** option. If you have two or more SCSI disks installed in your SGI Altix 3000 system, the installer asks you which drives you wish to use for installation. From the factory, SGI Altix 3000 systems come with the entire operating system installed on one system disk. If you would like the entire operating system installed on one disk, be sure that only one of the disks is selected for installation. Both devices are selected by default. To continue, select **OK**.
13. When the warning message appears about losing all data on the devices, select **Yes** to confirm the partition layout.
14. The next screen allows you to modify the partition layout as needed for your requirements. If you change the partition layout, ensure that the `/boot/efi` filesystem partition type remains set to `vfat`, the swap partition is set to `swap`, and all other file system types are set to `xfs`. To complete the partition layout, select **OK**.
15. If you use Bootp or DHCP to configure your network interfaces, select **OK**. If you use static IP addresses, fill in the required information and then to continue, select **OK**.

Note: If you plan to have an Ethernet board in bus 1 slot 2 of the IX-brick, directly to the right of the IO9, see “Persistent Naming of Ethernet Devices” on page 25. (This does not affect most users.)

16. If you chose to use a static IP address on the previous screen, you are prompted for a hostname for the system. Enter the hostname and to continue, select **OK**.
17. On the next screen you will be asked if you wish to configure a firewall on this system. Select the appropriate option based on your system's need and to continue, select **OK**.
18. The next screen allows you to select additional languages to support on the system.

Note: If you select additional languages, you must select a default language. The default language will be used on your system once installation is complete. If you choose to install other languages, you can change your default language after installation. Console support for Asian languages is not available at this time. SGI recommends that customers in Asia set the default system language to English.

19. The next screen allows you to select which time zone this system is in. Select the appropriate information and to continue, select **OK**.
20. The next screen asks you for the root password for this system. Enter the password you would like to use for root access and to continue, select **OK**.
21. The next screen allows you to create a user for the system. Creating a user allows you to log in to the system as a user other than `root`. If you choose to create a user, fill in the required boxes and to continue, select **OK**. Leaving the boxes blank and selecting **OK** continues without creating another user.
22. If you created a user in the previous screen, the next screen allows you to add additional users. Add other users as necessary. To continue, select **OK**.
23. The next screen lets you choose your authentication configuration. To continue, select **OK**.
24. The **Package Group Selection** screen appears. All packages should be selected for installation. Select **Everything** (located at the bottom of the list). To continue, press the Tab key to highlight **OK** and press Enter.
25. The **Installation to begin** screen is informational. To continue, select **OK**.
26. The installation process will now begin. During the process, you will be prompted to insert SGI Advanced Linux Environment 2.1 CD2. Do this when instructed and to continue, select **OK**.
27. When the installation has finished, select **OK**. The system will reboot. You can remove the CD from the drive.

You are not done yet! Only the base environment has been updated. After the system restarts, to upgrade SGI ProPack, proceed to “Installing SGI ProPack for Linux” on page 11.

Upgrading from Previous SGI ProPack Releases to SGI ProPack 2.2

This section explains how to upgrade from previous SGI ProPack releases, including SGI ProPack 2.1, 2.1.1, or 2.1.2, to this SGI ProPack 2.2 release. For console setup options and instructions, see the section titled “Monitoring Your Server” in the *SGI Altix 3000 User Guide*.

Upgrading to version 2.2 is different from previous upgrades because the base environment has changed. For this reason, it is necessary to boot from SGI Advanced Linux Environment 2.1 CD1.

After you have followed steps 1 through 9 in “Installing the SGI Advanced Linux Environment 2.1 Base OS” on page 6, the **Installation Type** window appears. To upgrade, perform the following steps:

1. Select **Upgrade Existing System**, then select **OK** and press **Enter**.
2. Ignore the warning message by selecting **Ignore** and pressing **Enter**. If you have more than one SCSI disk in your system, this warning might be repeated for each SCSI disk.
3. If your system has more than one root filesystem, select the root partition you wish to upgrade then select **OK** and press **Enter**. If you have only one existing root filesystem, you will see a window stating that it will be upgraded. Select **OK** and press **Enter**.
4. A **Customize Packages to Upgrade** window appears. Select **No** and press **Enter**.
5. You will see a window stating that packages to be upgraded are being found. It will take a few minutes to get past this point.
6. An **Upgrade to begin** window appears. Ensure that **OK** is selected and press **Enter**.
7. During the system upgrade process, you might be prompted to insert SGI Advanced Linux Environment 2.1 CD2. When instructed, put that CD in the drive and press **Enter**.

8. When the installation has finished, select **OK**. The system will reboot. You can remove the CD from the drive.

You are not done yet! Only the base environment has been updated. After the system restarts, to upgrade SGI ProPack, proceed to “Installing SGI ProPack for Linux” on page 11.

Installing SGI ProPack for Linux

This section provides steps for installing the SGI ProPack v2.2 for Linux CD1 (Open/Free Source Software) and the SGI ProPack v2.2 for Linux CD 2 (Proprietary Software). For console setup options and instructions, see the section titled “Monitoring Your Server” in the *SGI Altix 3000 User Guide*.

1. Log in to the system as `root`.
2. Insert SGI ProPack v2.2 for Linux CD1 into the system and enter the following command to mount the CD:

```
mount /dev/cdrom /mnt/cdrom
```
3. Enter `/mnt/cdrom/INSTALL`.
4. The SGI Welcome screen for SGI ProPack will appear. To continue, select **OK**.
5. The next screen lets you choose your installation type. Select **Yes** and select **OK**.
6. The **Package Group Selection** screen appears. This screen allows you to select the type of package group you want to install. You select a package group by using the up and down arrow keys and pressing the space bar to select the one you want. When you select a package group, RPMs for that package group will be installed after you press the Tab key to highlight **OK** and press Enter. All packages are selected by default. If you change the selection, ensure that **SGI Proprietary** remains selected. If you will be using SGI licensed software, such as XVM, be sure to select **SGI Licensed**. Otherwise, leave it unselected. To continue, select **OK**.
7. The **Installation to Begin** screen appears. It tells you that a log of the installation will be placed in `/tmp/sgi-install.log`. Press the Tab key to highlight **OK** and press Enter.
8. The installation begins. You will see the **Package Installation** screen, which tells you which packages are being installed and logs the time it takes to install them.

Note: If you started the SGI ProPack installation from the system console, you may see a message stating "INIT: *version number* reloading" as the packages are being installed. This message is normal and is not an error.

9. During the installation process, you will be asked to insert SGI ProPack v2.2 for Linux CD2. Insert the disk when prompted and press Enter.
10. After the installation is complete, the **Complete** screen appears. Press Enter. You are returned to the `root` prompt.
11. After the software installation is complete, reboot your system to begin using the newly installed SGI ProPack for Linux software environment. You can reboot by typing `reboot` and pressing Enter. You can then remove the CD from the drive.

Installing the System Controller Software 1.1 CD

This section provides steps for installing the System Controller Software 1.1 (for IRIX and SGI Linux). This software is licensed for SGI Origin and Onyx 3000 series systems based on MIPS processors and for SGI Altix 3000 systems.

1. Log in to the system as `root`, using the root password you used in the previous installation process.
2. Insert the System Controller Software 1.1 CD into the system and enter the following command to mount the CD:

```
mount /dev/cdrom /mnt/cdrom
```
3. Change directories to `/mnt/cdrom/RPMS/ia64` and enter the following command:

```
./install
```
4. The installation begins. You will see which packages are being installed. After the installation is complete, you are returned to the `root` prompt.
5. After the software installation is complete, remove the System Controller Software 1.1 CD.

Upgrading Your Software

When you are ready to upgrade your SGI ProPack software, follow the instructions outlined in the section titled “Installing SGI ProPack for Linux” on page 11, using the newer SGI ProPack CD set.

To reinstall all of the software on your system, follow the entire set of installation instructions listed previously, using the latest SGI ProPack CD set.

Recovering a Damaged Root Filesystem

If your root filesystem should become damaged, use the following recovery steps:

1. Insert SGI Advanced Linux Environment 2.1 CD1 into the system’s CD-ROM drive and restart the system.
2. While the system is powering up, check the device mapping table. It should look similar to the following:

Device mapping table

```
fs0  : Pci(2|1)/Ata(Primary, Master)/CDROM(Entry1) This is the CD-ROM
fs1  : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part1, Sig00000000)
fs2  : Pci(1|1)/Scsi(Pun0/Lun2)/HD(Part1, Sigg1)
blk0 : Pci(2|1)/Ata(Primary, Master)
blk1 : Pci(2|1)/Ata(Primary, Master)/CDROM(Entry1)
blk2 : Pci(1|1)/Scsi(Pun0/Lun1)
blk3 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part1, Sig00000000)
blk4 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part2, Sig00000000)
blk5 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part3, Sig00000000)
blk6 : Pci(1|1)/Scsi(Pun0/Lun1)/HD(Part4, Sig00000000)
blk7 : Pci(1|1)/Scsi(Pun0/Lun2)
blk8 : Pci(1|1)/Scsi(Pun0/Lun2)/HD(Part1, Sigg10)
```

3. At the Shell> prompt, type the CD-ROM device name, as follows:

```
Shell> fs0:           Type this to change to the CD-ROM device
fs0:\>
```

4. Boot the CD by typing **elilo linux rescue**, as follows:

```
fs0:\> elilo linux rescue           This starts the boot loader
```

5. The CD will then boot and start the rescue process.

6. Select the language you prefer to use for the rescue process. To continue, select **OK**.

Note: Console support for Asian languages is not available at this time. SGI recommends that customers in Asia set the language to English.

7. The next screen lets you choose the media containing the rescue image. Currently, **Local CDROM** is the only choice. To continue, select **OK**.
8. On the next screen, you have the option of mounting your root drive. If your root filesystem has been damaged, it is important to attempt to mount it before running the `xfstools` `repair` command. This ensures that valuable data from the journal is recovered.

If you select **Continue** and the installer hangs or crashes, you will have to restart this process and select **Skip**. Then you can attempt to mount the filesystem manually. See the section titled “Manually Mounting the Root Filesystem”.

9. If all goes well, the next screen informs you that your system can be found under `/mnt/sysimage`. To enter the rescue shell, select **OK**.
10. At this point, administrators can inspect and attempt to repair or recover any damaged or missing software. What follows is an example of how you might repair a corrupted root filesystem.

While it is important to mount your filesystem to recover any data contained in the journal, you must then unmount the filesystem before running the `xfstools` `repair` command. The `mount` command issued without arguments lists currently mounted filesystems, as follows:

```
sh-2.05# mount
```

```
rootfs on / type rootfs (rw)
devfs on /dev type devfs (rw)
/dev/root.old on / type ext2 (rw)
none on /dev type devfs (rw)
/proc on /proc type proc (rw)
/dev/pts on /dev/pts type devpts (rw)
/tmp/cdrom on /mnt/source type iso9660 (ro)
/tmp/loop0 on /mnt/runtime type cramfs (ro)
/dev/xscsi/pci01.03.0-1/target1/lun0/part3 on /mnt/sysimage type xfs (rw)
/dev/xscsi/pci01.03.0-1/target1/lun0/part1 on /mnt/sysimage/boot/efi type vfat (rw)
none on /mnt/sysimage/dev/pts type devpts (rw)
none on /mnt/sysimage/proc type proc (rw)
```

The line ending with `/mnt/sysimage type xfs (rw)` also contains the device corresponding to your root disk. In this case, `/dev/xscsi/pci01.03.0-1/target1/lun0/part3`.

11. Enter the following command to unmount the filesystem:

```
umount /mnt/sysimage
```

12. Run `xfs_repair` on the filesystem, as follows:

```
xfs_repair /dev/xscsi/pci01.03.0-1/target1/lun0/part3
```

13. After `xfs_repair` completes, the filesystem has been repaired and you can reset your system.

Manually Mounting the Root Filesystem

If the installer could not locate your root filesystem, you can locate and mount it manually. This section describes that process. For details, see *XFS for Linux Administration*. In rare cases, mounting a corrupted filesystem could cause the mount command to hang or have trouble. If that happens, see *XFS for Linux Administration* for instructions on how to run `xfs_repair` without mounting first.

1. The following example assumes that your root filesystem is on partition 3 and that it resides on a disk in the first disk bay of the IX-brick. To view the partitions available on the system, issue the following command:

```
sh-2.05# cat /proc/partitions
```

```
major minor #blocks name rio rmerge rsect ruse wio wmerge wsect wuse running use aveq
 5      0   125470 xscsi/pci01.01.0/target0/lun0/disc 218 4524 18968 10952 0 0 0 0 0 10952
10952
 4      0  35843686 xscsi/pci01.03.0-1/target1/lun0/disc 32 101 692 69 2 14 512 10 0 80 80
 4      1    513008 xscsi/pci01.03.0-1/target1/lun0/part1 0 0 0 0 0 0 0 0 0 0 0
 4      2    9438208 xscsi/pci01.03.0-1/target1/lun0/part2 0 0 0 0 0 0 0 0 0 0 0
 4      3  25891840 xscsi/pci01.03.0-1/target1/lun0/part3 28 29 584 50 2 14 512 10 0 61 61
 4     16  35843686 xscsi/pci01.03.0-1/target2/lun0/disc 4 72 108 19 0 0 0 0 0 19 19
 4     17    513008 xscsi/pci01.03.0-1/target2/lun0/part1 0 0 0 0 0 0 0 0 0 0 0
 4     18    9438208 xscsi/pci01.03.0-1/target2/lun0/part2 0 0 0 0 0 0 0 0 0 0 0
 4     19  25891840 xscsi/pci01.03.0-1/target2/lun0/part3 0 0 0 0 0 0 0 0 0 0 0
```

2. When the system is running normally, the XSCSI pathname in the following example represents the system disk in bay 1 of the IX-brick. This example assumes that you put the root filesystem on partition 3.

```
/dev/xscsi/pci01.03.0-1/target1/lun0/part3
```

3. To manually mount the filesystem, enter the following commands:

```
sh-2.06# mkdir /mnt/sysimage
```

```
sh-2.07# mount -t xfs /dev/xscsi/pci01.03.0-1/target1/lun0/part3 /mnt/sysimage
```

Product Support

This chapter documents the product components that are supported on SGI Altix 3000 systems. (For a list of the products, see Table 3-1 on page 18.)

Descriptions of the product components are grouped in this chapter as follows:

- Operating system enhancements
- I/O subsystems
- HPC application tools and support
- System management tools
- NUMA tools

Software provided by SGI for the SGI ProPack for Linux release consists of a kernel RPM for the SGI Altix 3000 product, SGI Advanced Linux Environment 2.1 RPMs, and value-added software developed by SGI to run specifically on SGI systems. Table 3-1 provides a description of the SGI ProPack for Linux products.

Table 3-1 SGI ProPack v2.2 for Linux Products

Product	Description
Application performance measuring tools	<p>The following tools perform program optimization:</p> <p>VTune - This tool, developed and supported by Intel, uses the performance measurement facilities of the Itanium processor to take profiles based on elapsed time or other architected events within the processor. These profiles can be used to measure, tune, and improve application performance.</p> <p>pfmon - This tool, available as open source and licensed under the GPL, provides a command line interface to control the performance measurement facilities of the Itanium processor. Data generated by this tool can be post-processed to produce a variety of reports describing application performance. These reports can be used to measure, tune, and improve application performance. The pfmon package also includes a library interface, <code>libpfmon.a</code>, that can be used to create customized performance measurement tools.</p>
Array Services	Provides a set of tools with kernel support that simplify the management of systems and parallel applications for clusters of SGI systems.
CpuMemSets	Provides the kernel support and infrastructure for implementing processor and memory placement.
cpuset	Provides for naming a set of CPUs.
CSA	Provides jobs-based accounting of per-task resources and disk usage for specific login accounts on Linux systems.
FLEXlm	Provides a floating license, run-time environment. Includes daemons suitable for serving floating licenses.

Table 3-1 SGI ProPack v2.2 for Linux Products **(continued)**

Product	Description
IOC4 driver	Driver that supports the Internal IDE CD-ROM, NVRAM, and Real-Time Clock.
kdb	Supports kernel debugging of the running system, either directly from the keyboard or over a serial console.
Kernel partitioning support	Provides the software infrastructure necessary to support a partitioned system, including cross-partition communication support.
Kernel performance improvements	Includes community-based patches such as the O(1) scheduler back-ported from the 2.5.x kernel and included in the SGI kernel, several patches that reduce lock contention on the Big Kernel Lock (BKL), and FRlocks (Fast-Reader Locks), which reduce contention on the <code>xtime_lock</code> (used by <code>gettimeofday</code>).
L1 and L2 System Controller firmware	Provides support for managing and monitoring the power, cooling, and testing functions for a brick and system compute rack.
LKCD	Provides system crash dump analysis tools including <code>lcrash</code> and all user level scripts required for saving and configuring system crash dumps.
MPT	Provides industry-standard message passing libraries optimized for SGI computers.
NUMA tools	Provides a collection of NUMA related tools (<code>dlook</code> , <code>dplace</code> , and so on).
Performance Co-Pilot collector infrastructure	Provides performance monitoring and performance management services targeted at large, complex systems.
PROM	Allows the CPU to boot and allows you to perform system administration and software installations.
runon	Enables running a command on a particular CPU or set of CPUs.

Table 3-1 SGI ProPack v2.2 for Linux Products **(continued)**

Product	Description
XFS	Provides a high-performance filesystem for Linux.
XSCSI infrastructure	Provides disk and ATAPI CD-ROM support for the QLogic QLA12160 SCSI Host Bus Adapter. In addition, some QLogic fibre channel host bus adapter cards are supported. For details on which cards are supported, see the <i>SGI Altix 3000 User's Guide</i> . The infrastructure includes support for robust error handling, failover, and SAN. It can be configured with or without the Linux SCSI layers. If configured without the Linux SCSI layer, XSCSI will make the familiar disk device names (<code>/dev/sda</code> , etc).
XVM	Provides software volume manager functionality such as disk striping and mirroring.

SGI does not support the following:

- Software obtained from other places (that is, not released by SGI).
- Other releases, updates, or patches from Red Hat.
- Software patches, drivers, or other changes obtained from the Linux community or other vendors.
- Kernels recompiled or reconfigured to run with parameter settings or other modules as not specified by SGI. In particular, the following kernel components are not supported due to quality, functionality, or performance issues: ext3, ReiserFS, LVM, and md. You should use XFS and XVM instead.
- Unsupported hardware configurations and devices.

Operating System Enhancements

Building on the Linux operating system's rapid expansion and improvements for general commercial and enterprise environments, SGI has focused on improving Linux capabilities and performance specifically for high performance computing's (HPC's) big compute and big data environments. Thus, SGI has leveraged its experience with NUMAflex and HPC from its IRIX operating systems and MIPS processor-based systems and concentrated on the Linux kernel improvements specifically important to HPC environments.

CpuMemSets Support

CpuMemSets provides a Linux kernel facility that enables system services and applications to specify on which CPUs they can be scheduled, and from which nodes they can allocate memory. The SGI ProPack kernel and library installation automatically include support for CpuMemSets.

The default configuration makes all CPUs and all system memory available to all applications. You can use the CpuMemSets facility to restrict any process, process family, or process virtual memory region to a specified subset of the system CPUs and memory.

The `runon` command relies on CpuMemSets to enable a user to run a specified command on a specified list of CPUs. Both a C shared library and Python language module are provided to access the CpuMemSets system interface.

In future releases of SGI ProPack for Linux, SGI anticipates providing additional facilities that provide other system services that provide convenient access to the CpuMemSets facility.

For further documentation and details on CpuMemSets support, see the chapter titled "CPU Memory Sets and Scheduling" in the *Resource Administration Guide for Linux* or browse the files that are installed as part of the CpuMemSets RPM, as listed by the following command:

```
rpm -ql CpuMemSets
```

In particular, see the man pages for `cpumemsets` and `runon`.

CpuMemSets is an SGI open source project, also available from the following location:

```
http://oss.sgi.com/projects/cpumemsets
```

Cpuset Support

The cpuset system is primarily a workload manager tool that permits a system administrator to restrict the number of processors that a process or set of processes can use.

A cpuset is a named set of CPUs, which can be defined to be restricted or open. A restricted cpuset allows only processes that are members of the cpuset to run on the set of CPUs. An open cpuset allows any process to run on its CPUs, but a process that is a member of the cpuset can run only on the CPUs belonging to the cpuset. A cpuset is defined by a cpuset configuration file and a name.

A system administrator can use cpusets to create a division of CPUs within a larger system. Such a divided system allows a set of processes to be contained to specific CPUs, reducing the amount of interaction and contention those processes have with other work on the system. In the case of a restricted cpuset, the processes that are attached to that cpuset will not be affected by other work on the system; only those processes attached to the cpuset can be scheduled to run on the CPUs assigned to the cpuset. An open cpuset can be used to restrict processes to a set of CPUs so that the effect these processes have on the rest of the system is minimized.

Comprehensive System Accounting (CSA)

The port of Comprehensive System Accounting (CSA) software packages from IRIX to Linux is the result of an open source collaboration between SGI and Los Alamos National Laboratory (LANL) to provide jobs-based accounting of per-task resources and disk usage for specific login accounts on Linux systems.

Providing extensive system accounting capabilities is often important for very large systems, especially when the system will be shared or made available for other organizations to use. CSA uses a Job Containers feature, which provides on Linux the notion of a *job*. A job is an inescapable container and a collection of processes that enables CSA to track resources for any point of entry to a machine (for example, interactive login, cron job, remote login, batched workload, and so on).

CSA on Linux is an SGI open source project, also available from the following location:

<http://oss.sgi.com/projects/csa>

For further documentation and details on CSA support, see the chapter titled “Comprehensive System Accounting” in the *Resource Administration Guide for Linux*.

Partitioning

SGI provides the ability to divide a single SGI Altix 3000 system into a collection of smaller system partitions. Each partition runs its own copy of the operating system kernel and has its own system console, root filesystem, IP network address, and physical memory. All partitions in the system are connected via the SGI high-performance NUMAlink interconnect, just as they are when the system is not partitioned. Thus, a partitioned system can also be viewed as a cluster of nodes connected via NUMAlink.

Benefits of partitioning include fault containment and the ability to use the NUMAlink interconnect and global shared memory features of the Altix 3000 to provide high-performance clusters.

For further documentation and details on partitioning, see the *SGI ProPack v2.2 for Linux Release Notes*.

I/O Subsystems

Although some HPC workloads might be mostly CPU bound, others involve processing large amounts of data and require an I/O subsystem capable of moving data between memory and storage quickly, as well as having the ability to manage large storage farms effectively. The SCSI subsystem, XFS filesystem, XVM volume manager, and data migration facilities were leveraged from IRIX and ported to provide a robust, high-performance, and stable storage I/O subsystem on Linux.

The following sections describe persistent PCI-X bus numbering, persistent naming of Ethernet devices, the XSCSI subsystem, the XSCSI-SCSI subsystem, the XFS filesystem, and the XVM Volume Manager.

Persistent PCI-X Bus Numbering

Persistent PCI-X bus numbering ensures that bus numbers can remain the same across reboots in case of faulty hardware or reconfiguration. During platform initialization, as buses are discovered, they are assigned a logical bus number. Each logical bus number is

unique, systemwide. The default number of buses supported by SGI Altix 3000 systems is 256 (numbered 0 to 255).

By default, bus numbers are allocated starting from the lowest C-brick module ID to which the I/O brick is connected. An I/O brick is either an IX-brick or a PX-brick on SGI Altix 3000 systems. Each I/O brick is allocated 0x10 buses, although the current I/O bricks support only six buses each. Therefore, bus numbers are not contiguous across the system. Bus numbers are sparse and have holes in them.

An I/O brick has six physical buses. These buses are numbered 0x1 through 0x6, left to right, looking at the back of the I/O brick. If there is more than one I/O brick on the system, the buses on the next I/O brick are numbered 0x11 through 0x16. The rationale for this numbering is that the bus numbers of each I/O brick are stamped on the back of the brick and they start from 1. Therefore, the rightmost digit of a bus number corresponds to the actual stamped number on the I/O brick.

If you have only one I/O brick, you do not need persistent bus numbering. However, if you have more than one I/O brick, persistent bus numbering is strongly recommended, so that if an I/O brick fails to boot, your bus numbers are still the same.

To make use of persistent bus numbering, you must supply the `ioconfig=` parameter to the kernel.

The `ioconfig=` parameter takes a comma-separated list of I/O brick numbers as an argument. The following manual boot example uses `elilo`. It tells the kernel that the I/O brick represented by `101.01` is assigned bus numbers 0x1 through 0x6 and the I/O brick represented by `101.02` is assigned bus numbers 0x11 through 0x16.

```
Shell> elilo vmlinux ioconfig="101.01,101.02"root=/dev/xscsi/pci01.03.0-1/target1/lun0/part3
```

You can find the numbers to use with `ioconfig` by looking at the components of your system from your L2 controller, as in the following example:

```
l2-pumpkin-001-L2>pwr
001c11:
power appears on
001c27:
power appears on
101i01:
power appears on
101p02:
power appears on
```

In this example, 101i01 is an IX-brick and represents the 101.01 in the previous `elilo` boot example. The 101p02 notation is a PX-brick and represents 101.02.

Most customers will have their system set up to boot automatically. This means that you should update your `elilo.conf` file with the `ioconfig` parameter. The `elilo.conf` file is available in the `/boot/efi/EFI/sgi` directory on your Linux system. The following is an example `elilo.conf` file:

```
prompt
timeout=50
relocatable
default=sgilinux
append="ioconfig=101.01,101.02"

image=vmlinuz-2.4.19-sgi2lr4
    label=sgilinux
    read-only
    root=/dev/xscsi/pci01.03.0-1/target1/lun0/part3
```

In the previous example, pay special attention to the `append=` line. That is, notice where the `ioconfig` information goes when using `elilo.conf`.

Persistent Naming of Ethernet Devices

Persistent naming of Ethernet devices is an SGI proprietary mechanism and is supported on SGI Altix 3000 systems. Persistent naming refers to the mechanism that ensures that the Gigabit Ethernet card on the IO9 interface of an SGI Altix 3000 system is set up to always be `eth0`. It guarantees that the base Ethernet device number is assigned to the correct MAC address on an SGI Altix 3000 system even when multiple Ethernet devices are present in the system.

The `/etc/sysconfig/networking/eth0_persist` file contains the mapping of Ethernet device numbers to MAC addresses. If the file does not exist, it is created by the `/etc/rc.d/init.d/eth_persist` script, which is run at boot time. To ensure that `eth0` is indeed assigned to the MAC address of the IO9 Ethernet card, it might be necessary to edit the file after the first time an SGI Altix 3000 system has been brought up after a clean install.

Besides ensuring that the mapping of Ethernet device numbers to MAC addresses persists as cards are added to a system, persistent naming also allows system administrators to control the way in which Ethernet devices are numbered. For example,

if the Ethernet card with device number ethX is lost and the system administrator tries to recover by using the Ethernet card with device number ethY, it is possible to force the latter card to take on Ethernet device number ethX by editing the `/etc/sysconfig/networking/eth0_persist` file accordingly.

Following is a sample `/etc/sysconfig/networking/eth0_persist` file:

```
eth0 08:00:69:13:dc:ec
eth1 08:00:69:13:72:e8
```

The content of this file results in the following configuration:

```
[root]# ifconfig -a
eth0 Link encap:Ethernet HWaddr 08:00:69:13:DC:EC
      inet addr:128.162.246.125 Bcast:128.162.246.255 Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:843 errors:0 dropped:0 overruns:0 frame:0
      TX packets:1245 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:100
      RX bytes:386044 (376.9 Kb) TX bytes:126741 (123.7 Kb)
      Interrupt:59

eth1 Link encap:Ethernet HWaddr 08:00:69:13:72:E8
      BROADCAST MULTICAST MTU:1500 Metric:1
      RX packets:136 errors:0 dropped:0 overruns:0 frame:0
      TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:100
      RX bytes:8850 (8.6 Kb) TX bytes:0 (0.0 b)
      Interrupt:63
```

If the slot directly to the right of the IO9 is populated by an Ethernet board, that is, bus 1, slot 2, and if you re-install Linux after it has been placed there or if you remove the `eth0_persist` persistent naming file, the IO9 could become eth1 instead of eth0. This is not recommended because it could affect your product licenses. If the Ethernet board is in any other IX-brick slot, you will not encounter this problem.

If it is necessary to have an Ethernet board in IO9 bus 1 slot 2, install the OS, including SGI ProPack, with the Ethernet board removed. After SGI ProPack is installed, you can replace the Ethernet board in bus 1 slot 2.

XSCSI Subsystem

The SGI XSCSI subsystem on Linux leverages from IRIX functionality to provide more robust error handling, failover, and storage area network (SAN) infrastructure support as well as long-term large system performance tuning. XSCSI takes advantage of specific features of SGI architecture that standard open source drivers cannot without rewriting.

The naming convention of XSCSI device names is shown in the following example:

```
/dev/xscsi/pci01.03.0-1/target1/lun0/part1
```

Components of the XSCSI device name are as follows:

pci01	System bus number 0x1
03	Device number 3 on that bus
0-1	Logical unit 0 port 1

Notice that the device number (slot number), logical unit, and port number are fixed. These will never change. However, the system bus number could change because of a hardware problem (such as the I/O brick not booting) or a reconfiguration.

Persistent PCI-X bus numbering (see “Persistent PCI-X Bus Numbering” on page 23), if enabled, provides persistent naming to prevent bus number changes even when the hardware fails or is reconfigured. If you use XSCSI names for mounting or locating devices and you also use persistent bus numbering, your XSCSI device names will always be persistent across reboots.

XSCSI-SCSI Subsystem

The XSCSI-SCSI subsystem provides SCSI device emulation of XSCSI devices. It allows programs written for the SCSI disk driver, the SCSI tape driver, and the SCSI generic driver to use devices controlled by the XSCSI drivers.

XFS Filesystem

The SGI XFS filesystem provides a high-performance filesystem for Linux. XFS is an open-source, fast recovery, journaling filesystem that provides direct I/O support, space preallocation, access control lists, quotas, and other commercial file system features. Although other filesystems are available on Linux, performance tuning and

improvements leveraged from IRIX make XFS particularly well suited for large data and I/O workloads commonly found in HPC environments.

For more information on the XFS filesystem, see *XFS for Linux Administration*.

XVM Volume Manager

The SGI XVM Volume Manager provides a logical organization to disk storage that enables an administrator to combine underlying physical disk storage into a single logical unit, known as a logical volume. Logical volumes behave like standard disk partitions and can be used as arguments anywhere a partition can be specified.

A logical volume allows a filesystem or raw device to be larger than the size of a physical disk. Using logical volumes can also increase disk I/O performance because a volume can be striped across more than one disk. Logical volumes can also be used to mirror data on different disks.

HPC Application Tools and Support

SGI has ported HPC libraries, tools, and software packages from IRIX to Linux to provide a powerful, standards-based system using Linux and Itanium 2-based solutions for HPC environments. The following sections describe some of these tools, libraries, and software.

Message Passing Toolkit

The SGI Message Passing Toolkit (MPT) provides industry-standard message passing libraries optimized for SGI computers. On Linux, MPT contains MPI and SHMEM APIs, which transparently utilize and exploit the low-level capabilities within SGI hardware, such as its block transfer engine (BTE) for fast memory-to-memory transfers and the hardware memory controller's fetch operation (fetchop) support. Fetchops enable direct communication and synchronization among multiple MPI processes while eliminating the overhead associated with system calls to the operating system.

Parallel workloads, such as MPI jobs, can be launched, monitored, and controlled across a cluster or partitioned system using the SGI Array Services software. Array Services provides the notion of an array session, which is a set of processes that can be running on

different cluster nodes or system partitions. Array Services is implemented using Process Aggregates (PAGGs), which is a kernel module that provides process containers. PAGGs has been open-sourced by SGI for Linux.

For more information on the Message Passing Toolkit, see the *Message Passing Toolkit: MPI Programmer's Manual*.

Performance Co-Pilot

The SGI Performance Co-Pilot software was ported from IRIX to Linux to provide a collection of performance monitoring and performance management services targeted at large, complex systems. Integrated with the low-level performance hardware counters and with MPT, Performance Co-Pilot provides such services as CPU, I/O, and networking statistics; visualization tools; and monitoring tools.

System Management

This section describes system management tools. These include the hardware and software environment required to boot the system, license management, system console, and system controllers.

PROM Chips

Programmable read-only memory (PROM) chips are placed in your computer at the factory with software programmed into them that allows the CPU to boot and allows you to perform system administration and software installations. The PROM chips are not part of your disk or your operating system; they are the lowest level of access available for your system. You cannot erase them or bypass them. For more information on PROM, see the `prom(1)` man page.

Extensible Firmware Interface (EFI)

SGI Altix 3000 systems provide the Extensible Firmware Interface (EFI), a supporting platform to provide input to the CPU and to handle its output. In addition, the EFI controls the server's boot configuration, maintaining the boot menu in durable, nonvolatile memory.

SGI ProPack v2.2 uses the `elilo-3.3` package, which is fully compliant with EFI specification 1.10 with regard to where the bootloader (`elilo.efi`) should be located in the EFI system partition. According to that specification, the bootloader must be located in a dedicated vendor directory. On SGI Altix systems, that directory is `/boot/efi/EFI/sgi/`. As of version 3.3, `elilo` looks for its configuration file **only** in the directory from which it is loaded. Because this also applies to the kernel, the kernel image, which was previously installed in `/boot/efi/`, is now installed in `/boot/efi/EFI/sgi/`. For further documentation and details about this change, see the `elilo` documentation in `/usr/share/docs/`.

From the EFI prompt, you can perform basic file-management tasks (including text editing), make configuration changes, or write scripts that execute at boot time. For a summary of EFI commands, see Table 3-2.

Table 3-2 EFI Commands

EFI Command	Description
<code>alias [-bdv] [sname] [value]</code>	Sets or gets alias settings
<code>attrib [-b] [+/- rhs] [file]</code>	Views or sets file attributes
<code>bcfg</code>	Configures boot driver and load options
<code>cd [path]</code>	Updates the current directory
<code>cls [background color]</code>	Clears screen
<code>comp file1 file2</code>	Compares two files
<code>cp file [file] ... [dest]</code>	Copies files or directories
<code>date [mm/dd/yyyy]</code>	Gets or sets date
<code>dblk device [Lba] [blocks]</code>	Performs hex dump of block I/O devices
<code>dh [-b] [-p prot_id] [handle]</code>	Dumps handle information
<code>dmpstore</code>	Dumps variable store
<code>echo [-on -off] [text]</code>	Echoes text to stdout or toggles script echo
<code>edit [file name]</code>	Edits a file
<code>endfor</code>	Script only: Delimits loop construct

Table 3-2 EFI Commands (continued)

EFI Command	Description
endif	Script-only: Delimits IF THEN construct
err [<i>level</i>]	Sets or displays error level
exit	Exits
flash <i>filename</i>	Flashes PROM on C-brick
for var in <i>set</i>	Script-only: Indicates loop construct
getmtc	Gets next monotonic count
goto <i>label</i>	Script-only: Jumps to label location in script
guid [-b] [<i>sname</i>]	Dumps known guid IDs
help [-b] [<i>internal command</i>]	Displays this help
if [not] <i>condition</i> then	Script-only: Indicates IF THEN construct
load <i>driver_name</i>	Loads a driver
ls [-b] [<i>dir</i>] [<i>dir</i>] ...	Obtains directory listing
map [-bdvr] [<i>sname</i> [:]] [<i>handle</i>]	Maps shortname to device path
mem [<i>address</i>] [<i>size</i>] [;MMIO]	Dumps memory or memory mapped I/O
memmap [-b]	Dumps memory map
mkdir <i>dir</i> [<i>dir</i>] ...	Makes directory
mm <i>address</i> [<i>width</i>] [; <i>type</i>]	Modifies memory: Mem, MMIO, IO, PCI
mode [<i>col row</i>]	Sets or gets current text mode
mount <i>BlkDevice</i> [<i>sname</i> [:]]	Mounts a filesystem on a block device
mv <i>sfile dfile</i>	Moves files
pause	Script-only: Prompts to quit or continue
pci [<i>bus dev</i>] [<i>func</i>]	Displays PCI device(s) info
reset [cold/warm] [<i>reset string</i>]	Indicates cold or warm reset

Table 3-2 EFI Commands (**continued**)

EFI Command	Description
<code>rm <i>file/dir</i> [<i>file/dir</i>]</code>	Removes file or directories
<code>set [-bdv] [<i>sname</i>] [<i>value</i>]</code>	Sets or gets environment variable
<code>setsize <i>newsize fname</i></code>	Sets the files size
<code>stall <i>microseconds</i></code>	Delays for <i>x</i> microseconds
<code>time [<i>hh:mm:ss</i>]</code>	Gets or sets time
<code>touch [<i>filename</i>]</code>	Views or sets file attributes
<code>type [-a] [-u] [-b] <i>file</i></code>	Types file
<code>ver</code>	Displays version information
<code>vol <i>fs</i> [<i>volume label</i>]</code>	Sets or displays volume label

FLEXlm

FLEXlm is a flexible license management system from Macrovision that lets independent software vendors (ISVs) license their products and helps system administrators install and manage licenses with minimal overhead. It supports a wide range of licensing options, including simple node-locked licenses and floating licenses with redundant servers.

To build licensed software, ISVs must purchase a set of keys from Macrovision. System administrators can install license servers anywhere. Products purchased from SGI are typically licensed using FLEXlm.

For more information, visit

<http://www.macrovision.com/solutions/esd/flexlm/flexlm.shtml>

SGIconsole

SGIconsole is a combination of hardware and software that provides console management and allows monitoring of multiple SGI servers running the IRIX operating

system and SGI ProPack for Linux. These servers include SGI partitioned systems and large, single-system-image servers, including SGI Altix 3000 servers.

SGIconsole consists of an 1U rackmountable SGI server based on the Intel Pentium processor, a serial multiplexer or Ethernet hub, and a software suite that includes the Console Manager package and Performance Co-Pilot, which provides access to common remote management tools for hardware and software.

Console Manager is a graphical user interface for the SGIconsole management and monitoring tool used to control multiple SGI servers. SGIconsole also has a command line interface. For more information on SGIconsole, see the *SGIconsole Start Here*.

System Controller Firmware

The L1 and L2 controllers are system controller firmware used in SGI systems.

The L1 controller is embedded in each brick in SGI Origin and Onyx 3000 series systems and in SGI Altix 3000 systems. It provides power and control sequencing, along with temperature and power monitoring for each brick.

The L2 controller is a rack-level controller that monitors and controls the bricks in its rack. All L2 controllers in a system are networked together and they consolidate the control and monitoring information from each brick to provide system-level control and monitoring.

For more information on the L1 and L2 system controller firmware, see the *SGI L1 and L2 Controller Software User's Guide*.

NUMA Tools

This section describes the commands that are currently provided with the collection of NUMA related tools.

dlook Command

The `dlook` command displays the memory map and CPU use for a specified process. The following information is printed for each page in the virtual address space of the process:

- The object that owns the page (file, SYSV shared memory, device driver, and so on)
- Type of page (RAM, FETCHOP, IOSPACE, and so on)
- If RAM memory, the following information is supplied:
 - Memory attributes (SHARED, DIRTY, and so on)
 - Node on which that the page is located
 - Physical address of page (optional)

Optionally, the amount of elapsed CPU time that the process has executed on each physical CPU in the system is also printed.

dp1ace Command

The `dp1ace` command binds a related set of processes to specific CPUs or nodes to prevent process migrations. In some cases, this tool improves performance because of the occurrence of a higher percentage of memory accesses to the local node.

Performance Tuning

This chapter describes tuning issues for single processor and multiprocessor programs. It contains the following sections:

- See “Single Processor Code Tuning” on page 36 for details about performance tuning of single processor programs.
- See “Multiprocessor Code Tuning” on page 44 for details about performance tuning of multiprocessor programs.
- See “Profilers and Performance Tools” on page 49 for details about tools that can be used to tune performance of either type of code.

For a list of available references, see “Related Publications” on page xvi.

Determining System Configuration

To determine the details of the system you are running, you can browse files from the `/proc` pseudo-filesystem (see the `proc(5)` man page for details). Following is some of the information you can obtain:

- `/proc/cpuinfo`: displays processor information, one entry per processor. Use this to determine clock speed and processor stepping.
- `/proc/meminfo`: provides a global view of system memory usage, such as total memory, free memory, swap space, and so on.
- `/proc/discontig`: shows memory usage (in pages).
- `/proc/pal/cpu0/cache_info`: provides detailed information about L1, L2, and L3 cache structure, such as size, latency, associativity, line size, and so on. Other files in `/proc/pal/cpu0` provide information about the Translation Lookaside Buffer (TLB) structure, clock ratios, and other details.
- `/proc/version`: provides information about the installed kernel.

- `/proc/perfmon`: if this file does not exist in `/proc` (that is, if it has not been exported), performance counters have not been started by the kernel and none of the performance tools that use the counters will work.
- `/proc/mounts`: provides details about the filesystems that are currently mounted.
- `/proc/modules`: contains details about currently installed kernel modules.

You can also use the `uname` command, which returns the kernel version and other machine information. In addition, the `topology` command is a Bourne shell script that uses information in `/dev/hw` to display system configuration information. See the `topology(1)` man page for details.

The currently recommended compilers for SGI Altix 3000 systems are `efc` and `ecc` (the Intel Fortran compiler and Intel C/C++ compiler). Other compilers (such as the GNU set of compilers) can be used; however, `efc` and `ecc` provide the best performance on Linux systems. Examples in this chapter, unless noted otherwise, use the `efc` or `ecc` compiler.

Single Processor Code Tuning

Several basic steps are used to tune performance of single-processor code:

- Get the expected answers and then tune performance. For details, see “Getting the Correct Results” on page 37.
- Use existing tuned code, such as that found in math libraries and scientific library packages. For details, see “Using Tuned Code” on page 40.
- Determine what needs tuning. For details, see “Determining Tuning Needs” on page 41.
- Use the compiler to do the work. For details, see “Using Compiler Options Where Possible” on page 41.
- Consider tuning cache performance. For details, see “Tuning the Cache Performance” on page 42.
- Set environment variables to enable higher-performance memory management mode. For details, see “Managing Memory” on page 43.
- Change stack and data segments. For details, see “Using `chatr` to Change Stack and Data Segments” on page 43.

Getting the Correct Results

One of the first steps in performance tuning is to verify that the correct answers are being obtained. Once the correct answers are obtained, tuning can be done. You can verify answers by initially disabling specific optimizations and limiting default optimizations. This can be accomplished by using specific compiler options and by using debugging tools.

The following compiler options emphasize tracing and porting over performance:

- `-O`: the `-O0` option disables all optimization. The default is `-O2`.
- `-g`: the `-g` option preserves symbols for debugging.
- `-mp`: the `-mp` option limits floating-point optimizations and maintains declared precision.
- `-IPF_fltacc`: the `-IPF_fltacc` option disables optimizations that affect floating-point accuracy.
- `-r`; `-i`: the `-r8` and `-i8` options set default real, integer, and logical sizes to 8 bytes, which are useful for porting Cray codes. **This explicitly declares intrinsic and external library functions.**

Some debugging tools can also be used to verify that correct answers are being obtained. The following debuggers can be used:

- `gdb`: the GNU project debugger. This is useful for debugging programs written in C, C++, and Fortran 95. When compiling with C and C++, include the `-g` option on the compiler command line to produce the `dwarf2` symbols database used by `gdb`.

When using `gdb` for Fortran debugging, include the `-g` and `-O0` options. Do not use `gdb` for Fortran debugging when compiling with `-O1` or higher.

The debugger to be used for Fortran 95 codes can be downloaded from http://sourceforge.net/project/showfiles.php?group_id=56720. (Note that the standard `gdb` compiler does not support Fortran 95 codes.) To verify that you have the correct version of `gdb` installed, use the `gdb -v` command. The output should appear similar to the following:

```
GNU gdb 5.1.1 FORTRAN95-20020628 (RC1)
Copyright 2002 Free Software Foundation, Inc.
```

For a complete list of `gdb` commands, see the `gdb` user guide online at http://sources.redhat.com/gdb/onlinedocs/gdb_toc.html or use the `help` option. Note that current instances of `gdb` do not report `ar.ec` registers correctly. If you are debugging rotating, register-based, software-pipelined loops at the assembly code level, try using `idb` instead.

- `idb`: the Intel debugger. This is a fully symbolic debugger for the Linux platform. The debugger provides extensive support for debugging programs written in C, C++, FORTRAN 77, and Fortran 90. At this time, `idb` cannot be used to debug multithreaded or multiprocessor programs on Itanium systems.

Running `idb` with the `-gdb` option on the shell command line provides `gdb`-like user commands and debugger output.

- `ddd`: a GUI to a command line debugger. It supports `gdb` and `idb`. For details about usage, see “Using `ddd`” on page 38.

Using `ddd`

The `ddd` tool is a GUI to an arbitrary command line debugger. To instantiate `ddd`, use the `--debugger` option to specify the debugger used (for example, `--debugger "idb"`). The default debugger used is `gdb`.

The resulting screen is divided into panes that show:

- Array inspection
- Source code
- Disassembled code
- A command line window to the debugger engine

These panes can be switched on and off from the **View** menu.

Some commonly used commands can be found on the menus. In addition, the following actions can be useful:

- Select an address in the assembly view, click the right mouse button, and select `lookup`. The `gdb` command is executed in the command pane and it shows the corresponding source line.
- Select a variable in the source pane and click the right mouse button. The current value is displayed. Arrays are displayed in the array inspection window. You can print these arrays to PostScript by using the **Menu > Print Graph** option.

- You can view the contents of the register file, including general, floating-point, NaT, predicate, and application registers by selecting **Registers** from the **Status** menu. The **Status** menu also allows you to view stack traces or to switch threads.

Managing Heap Corruption Problems

Two methods can be used to check for heap corruption problems in programs that use `glibc malloc/free` dynamic memory management routines: environment variables and Electric Fence.

Set the `MALLOC_CHECK_` environment variable to 1 to print diagnostic messages or to 2 to abort immediately when heap corruption is detected.

Electric Fence is a `malloc` debugger for Linux and Unix. It stops your program on the exact instruction that overruns (or underruns) a `malloc()` memory buffer and displays the source-code line that causes the bug. It aligns either the start or end of an allocated block with an invalid page, causing segmentation faults to occur on buffer overruns or underruns at the point of error. It can also detect accesses to previously freed regions of memory.

Underruns and overruns cannot be simultaneously detected. The default behavior is to place inaccessible pages immediately after allocated memory, but the complementary case can be enabled by setting the `EF_PROTECT_BELOW` environment variable. To use Electric Fence, link with the `libefence` library, as shown in the following example:

```
% cat foo.c
#include <stdio.h>
#include <stdlib.h>
int main (void)
{
    int i;
    int * a;
    float *b;

    a = (int *)malloc(1000*sizeof (int));
    b = (float *)malloc(1000*sizeof (float));

    a[0]=1;
    for (i=1 ; i<1001;i++)
        {
            a[i]=a[i-1]+1;
        }
    for (i=1 ; i<1000;i++)
```

```
        {
          b[i]=a[i-1]*3.14;
        }

    printf("answer is %d %f \n"a[999],b[999]);
}
```

Compile and run the program as follows (note the error when it is compiled with the library call):

```
% ecc foo.c
% ./a.out
answer is 1000 3136.860107
% ecc foo.c -lefence
% ./a.out
```

```
Electric Fence 2.2.0 Copyright (C) 1987-1999 Bruce Perens
Segmentation fault
%
```

To avoid potentially large core files, the recommended method of using Electric Fence is from within a debugger. See the `efence` man page for additional details.

Using Tuned Code

Where possible, use code which has already been tuned for optimum hardware performance.

The following mathematical functions should be used where possible to help obtain best results:

- **MKL:** Intel's Math Kernel Library. This library includes BLAS, LAPACK, and FFT routines.
- **VML:** the Vector Math Library, available as part of the MKL package (`libmkl_vml_itp.so`).
- **Standard Math library:** standard math library functions are provided with the Intel compiler's `libimf.a` file. If the `-lm` option is specified, `glibc libm` routines are linked in first.

Documentation is available for MKL and VML at the following website:

http://intel.com/software/products/perflib/index.htm?iid=ipp_home+software_libraries&

Determining Tuning Needs

Determine what points in your code might benefit from tuning. Use the following tools:

<code>time</code>	Use this command to obtain an overview of user, system, and elapsed time.
<code>gprof</code>	Use this tool to obtain an execution profile of your program (a <code>pcsamp</code> profile). Use the <code>-p</code> compiler option to enable <code>gprof</code> use.
VTune	This Intel performance monitoring tool is a Linux-server, Windows-client application. It supports remote sampling on all Itanium and Linux systems.
<code>pfmon</code>	This performance monitoring tool is designed for Itanium and Linux. It uses the Itanium Performance Monitoring Unit (PMU) to do counting and sampling on unmodified binaries.

For information about VTune and `pfmon`, see “Profilers and Performance Tools” on page 49.

Using Compiler Options Where Possible

Several compiler options can be used to optimize performance. For a short summary of `efc` or `ecc` options, use the `-help` option on the compiler command line. Use the `-dryrun` option to show the driver tool commands that `efc` or `ecc` generate. This option does not execute tools.

Use the following options to help tune performance:

<code>-ftz</code>	Flushes underflow to zero to avoid kernel traps. Enabled by default at <code>-O3</code> optimization.
<code>-fno-alias</code>	Assumes no pointer aliasing. Other aliasing options include <code>-ansi_alias</code> and <code>-fno_fnalias</code> . Note that alias assertions may generate incorrect code.
<code>-ip</code>	Generates single file, interprocedural optimization.
<code>-ipo</code>	Generates multifile, interprocedural optimization.

- O3 Enables -O2 optimizations plus more aggressive optimizations, including loop transformation and prefetching. Level 3 optimization may not improve performance for all programs.
- opt_report Generates an optimization report and places it in the file specified in `-opt_report_file`.
- override_limits This is an undocumented option that sometimes allows the compiler to continue optimizing when it has hit an internal limit.
- prof_gen and -prof_use Generates and uses profiling information. These options require a three-step compilation process:
 1. Compile with proper instrumentation using `-prof_gen`.
 2. Run the program on one or more training datasets.
 3. Compile with `-prof_use`, which uses the profile information from the training run.
- S Compiles and generates assembly listing in `.s` files and does not link. The assembly listing can be used in conjunction with the output generated by the `-opt_report` option to try to determine how well the compiler is optimizing loops.

Tuning the Cache Performance

There are several actions you can take to help tune cache performance:

- Avoid large power-of-2 strides and dimensions that cause cache thrashing.
- Use strides of 1 wherever possible.
- Cache bank conflicts can occur if there are two accesses to the same 16-byte-wide bank at the same time. Try different padding of arrays if the output from the `pfmon -e L2_OZQ_CANCEL_S1_BANK_CONF` command and the output from the `pfmon -e CPU_CYCLES` command shows a high number of bank conflicts relative to total CPU cycles. These can be combined into one command:

```
% pfmon -e CPU_CYCLES,L2_OZQ_CANCEL_S1_BANK_CONF a.out
```

A maximum of four performance monitoring events can be counted simultaneously.

- Group together data that is used at the same time and do not use vectors in your code, if possible.
- Try to avoid the use of temporary arrays and minimize data copies.

Managing Memory

Codes that frequently allocate and deallocate memory through `glibc malloc/free` calls may accrue significant system time due to memory management overhead. By default, `glibc` strives for system-wide memory efficiency at the expense of performance. A detailed discussion of this issue can be found at <http://www.linuxshowcase.org/ezolt.html>.

To enable the higher-performance memory management mode, set the following environment variables:

```
% setenv MALLOC_TRIM_THRESHOLD_ -1
% setenv MALLOC_MMAP_MAX_ 0
```

Because allocations in `efc` using the `malloc` intrinsic use the `glibc malloc` command internally, these environment variables are also applicable in Fortran codes using, for example, Cray pointers with `malloc/free`. But they do not work for Fortran 90 allocatable arrays, which are managed directly through Fortran library calls.

Using `chatr` to Change Stack and Data Segments

The `chatr` command changes the ELF header of a program so that when the program is run, instructions in the stack and data segments are or are not executable.

The default is to allow the stack and data segments to be executable. For certain workloads, a performance improvement can be obtained by turning off the default behavior and making the stack and code segments unexecutable. This typically benefits workloads that do a large number of `fork()` calls. Workloads without a large number of `fork()` calls will probably not see a performance improvement.

You can change this behavior system-wide by using the `sysctl` command. The following command disallows executable stacks and data segments:

```
% sysctl vm.executable_stacks=0
```

The following command allows executable stacks and data segments:

```
% sysctl vm.executable_stacks=1
```

Note that some programs may now fail with a `SEGV` error even though they worked correctly before. Usually, these are programs that store instructions in the stack or data segment and then branch to those instructions (for example, older versions of the X server that load graphics drivers into the data segment, or Java JIT compilers). You can use the `chattr` command to make these programs work correctly, regardless of the value of `vm.executable_stacks`.

For details about usage, see the `chattr(1)` man page.

Multiprocessor Code Tuning

Before beginning any multiprocessor tuning, first perform single processor tuning. This can often obtain good results in multiprocessor codes also. For details, see “Single Processor Code Tuning” on page 36.

Multiprocessor tuning consists of the following major steps:

- Choose the parallelization methodology for your code. For details, see “Parallelizing Your Code” on page 44.
- Analyze your code to make sure it is parallelizing properly. For details, see “Solving Bottlenecks in Code” on page 46.
- Check to determine if false sharing exists. For details, see “Fixing False Sharing” on page 46.
- Tune for data placement. For details, see “Using `dplace` and `runon`” on page 47.
- Use environment variables to assist with tuning. For details, see “Environment Variables for Performance Tuning” on page 48.

Parallelizing Your Code

The first step in multiprocessor performance tuning is to choose the parallelization methodology that you want to use for tuning. You can use the following options:

- Use the Message Passing Interface (MPI) from the SGI Message Passing Toolkit (MPT). MPI is optimized and more scalable for SGI Altix 3000 series systems than

generic MPI libraries. It takes advantage of the SGI Altix 3000 architecture and SGI Linux NUMA features.

Use the `-lmpi` compiler option to use MPI. For a list of environment variables that are supported, see the `mpi(1)` man page. Those variables that are valid for IRIX systems only are so noted on the man page.

`MPIO_DIRECT_READ` and `MPIO_DIRECT_WRITE` are supported under Linux for local XFS filesystems in SGI MPT version 1.6.1 and beyond.

- Use OpenMP. When using C, C++, or Fortran code with OpenMP directives, you can use the following compiler options:

`efc -openmp` or `ecc -openmp`

These options use the Intel OpenMP front-end that is built into the Intel compilers. The resulting executable file makes calls to `libguide.so`, which is the Intel OpenMP run-time library.

`guide` An alternate command to invoke the Intel compilers to use OpenMP code. Use `guidec` (in place of `ecc`), `guideefc` (in place of `efc`), or `guidec++` to translate code with OpenMP directives into code with calls to `libguide`. See “Other Performance Tools” on page 52 for details.

The `-openmp` option to `efc` is the Intel long-term OpenMP compiler for Linux. However, if you have performance problems with this option, using `guide` might provide improved performance.

- Use the compiler to invoke automatic parallelization. Use the `-parallel` and `-par_report` option to the `efc` or `ecc` compiler. These options show which loops were parallelized and the reasons why some loops were not parallelized. If a source file contains many loops, it might be necessary to add the `-override_limits` flag to enable automatic parallelization. The code generated by `-parallel` is based on the OpenMP API, and the standard OpenMP environment variables and Intel extensions apply.

Determine the amount of code that is parallelized. Use the following formula to calculate the amount of code that is parallelized:

$$p = N(T(1) - T(N)) / T(1)(N - 1)$$

In this equation, $T(1)$ is the time the code runs on a single CPU and $T(N)$ is the time it runs on N CPUs. Speedup is defined as $T(1)/T(N)$.

If $speedup/N$ is less than 50% (that is, $N > (2-p) / (1-p)$), stop using more CPUs and tune for better scalability.

CPU activity can be displayed with the `top` or `vmstat` commands or accessed by using the PCP tools (for example, `pmval kernel.percpu.cpu.user`) or by using Performance Co-Pilot visualization tools on a remote IRIX workstation. Hardware configuration information can be displayed using the `/proc/cpuinfo` or `/proc/meminfo` script. Detailed cache (and other) information can be displayed using the `/proc/pal/cpuX` script, where *X* is a CPU number.

Solving Bottlenecks in Code

For details about the different profiling tools that can pinpoint the bottlenecks in your code, see “Profilers and Performance Tools” on page 49.

Fixing False Sharing

If the parallel version of your program is slower than the serial version, false sharing might be occurring. False sharing occurs when two or more data items that appear not to be accessed by different threads in a shared memory application correspond to the same cache line in the processor data caches. If two threads executing on different CPUs modify the same cache line, the cache line cannot remain resident and correct in both CPUs, and the hardware must move the cache line through the memory subsystem to retain coherency. This causes performance degradation and reduction in the scalability of the application. If the data items are only read, not written, the cache line remains in a shared state on all of the CPUs concerned. False sharing can occur when different threads modify adjacent elements in a shared array.

You can use the following methods to verify that false sharing is happening:

- Use the performance monitor to look at output from `pfmon` and the `BUS_MEM_READ_BRILL_SELF` and `BUS_RD_INVALID_ALL_HITM` events.
- Use `pfmon` to check `DEAR` events to track common cache lines.
- Use the PCP `pmshub` utility to monitor cache traffic and CPU utilization.

Note: The `pmshub` utility may not be available for this release. Consult your release notes for information about its availability.

If false sharing is a problem, try the following solutions:

- Use the HW counter to run a profile that monitors storage to shared cache lines. This will show the location of the problem.
- Revise data structures or algorithms.
- Check shared data, static variables, common blocks, and private and public variables in shared objects.
- Use critical regions to identify the part of the code that has the problem.

Using `dplace` and `runon`

The `dplace` command binds processes to specified CPUs in a round-robin fashion. Once bound to a process, they do not migrate. This is similar to `_DSM_MUSTRUN` on IRIX systems. `dplace` numbering is done in the context of the current CPU memory set.

The `runon` command restricts execution to the listed set of CPUs; however, processes are still free to move among listed CPUs.

For more details on these commands, see the following sections.

`dplace` for MPI Codes

For best placement and for CPU-data affinity, use one of the following methods:

- Use `MPI_DSM_MUSTRUN` in a `cpuset`.
- Set the `MPI_DSM_CPULIST` environment variable.
- Use the `dplace -s1` command.

When using SGI MPI-compiled codes, use the `-s1` option and provide the full path to `dplace`, as in the following example:

```
% mpirun -np 32 /usr/bin/dplace -c16-47 -s1 ./a.out
```

Using the `-s1` option directs `dplace` to skip placement of the lightweight MPI “shepherd” process.

dplace for OpenMP Codes

When using OpenMP codes from Intel, use the `-x6` option, as in the following example:

```
% setenv OMP_NUM_THREADS 4  
% dplace -x6 -c5-8 ./a.out
```

The `-x6` option directs `dplace` to skip placement of the two lightweight OpenMP shepherd processes.

Notice the distinction between `dplace` and `runon`:

```
% setenv OMP_NUM_THREADS 4  
% dplace -x6 -c5-8 ./a.out  
% runon 5-8 ./a.out
```

The `dplace` command allows you to skip placement of the shepherd processes, while `runon` does not. `dplace` also locks each thread to its corresponding CPU; `runon` restricts the threads to CPUs 5 through 8, but they are free to migrate among those CPUs.

For information about `dplace` and its usage, see the *Linux Resource Administration Guide*.

Environment Variables for Performance Tuning

You can use several different environment variables to assist in performance tuning. For details about environment variables used to control the behavior of MPI, see the `mpi(1)` man page.

Several OpenMP environment variables can affect the actions of the OpenMP library. For example, some environment variables control the behavior of threads in the application when they have no work to perform or are waiting for other threads to arrive at a synchronization semantic; other variables can specify how the OpenMP library schedules iterations of a loop across threads. The following environment variables are part of the OpenMP standard:

- `OMP_NUM_THREADS` (The default is the number of CPUs in the system.)
- `OMP_SCHEDULE` (The default is `static`.)

- `OMP_DYNAMIC` (The default is `false`.)
- `OMP_NESTED` (The default is `false`.)

In addition to the preceding environment variables, Intel provides several OpenMP extensions, two of which are through the use of the `KMP_LIBRARY` variable.

The `KMP_LIBRARY` variable sets the run-time execution mode, as follows:

- If set to `serial`, single-processor execution is used.
- If set to `throughput`, CPUs yield to other processes when waiting for work. This is the default and is intended to provide good overall system performance in a multiuser environment. This is analogous to the IRIX `_DSM_WAIT=YIELD` variable.
- If set to `turnaround`, worker threads do not yield while waiting for work. This is analogous to the IRIX `_DSM_WAIT=SPIN` variable. Setting `KMP_LIBRARY` to `turnaround` may improve the performance of benchmarks run on dedicated systems, where multiple users are not contending for CPU resources.

If your program gets a segmentation fault immediately upon execution, you may need to increase `KMP_STACKSIZE`. This is the private stack size for threads. The default is 4 MB. You may also need to increase your shell stack size limit.

Profilers and Performance Tools

Several tools are available to help determine areas for performance improvements. The following sections describe some of those tools.

Profiling with `pfmon`

The `pfmon` tool is a performance monitoring tool designed for Linux. It uses the Itanium Performance Monitoring Unit (PMU) to count and sample on unmodified binaries. In addition, it can be used for the following tasks:

- To monitor unmodified binaries in its per-process mode.
- To run system-wide monitoring sessions. Such sessions are active across all processes executing on a given CPU.
- Launch a system-wide session on a dedicated CPU or a set of CPUs in parallel.

- Monitor activities happening at the user level or at the kernel level.
- Collect basic event counts.
- Sample program or system execution, monitoring up to four events at a time.

To see a list of available options, use the `pfmon -help` command.

The `profile.pl` script

The `profile.pl` script handles the entire user program profiling process, including using `dplace`. Typical usage is as follows:

```
% profile.pl -c0-3 -x6 command args
```

This script designates processors 0 through 3. The `-x6` option is necessary only for OpenMP codes.

The result is a profile taken on the `CPU_CYCLES` PMU event and placed into `profile.out`. This script also supports profiling on other events such as `IA64_INST_RETIRED`, `L3_MISSES`, and so on; enter `pfmon -l` for a complete list of PMU events. The script handles running the command under the performance monitor, creating a map file of symbol names and addresses from the executable and any associated dynamic libraries, and running the profile analyzer.

See the `profile.pl(1)`, `analyze.pl(1)`, and `makemap.pl(1)` man pages for details.

`profile.pl` with MPI programs

For MPI programs, use the `profile.pl` command with the `-s1` option, as in the following example:

```
% mpirun -np 4 profile.pl -s1 -c0-3 test_prog </dev/null
```

The use of `/dev/null` ensures that MPI programs run in the background without asking for TTY input.

Using VTune for Remote Sampling

The Intel VTune performance analyzer does remote sampling experiments. The VTune data collector runs on the Linux system and an accompanying GUI runs on an IA-32 Windows machine, which is used for analyzing the results.

For details about using VTune, see the following URL:

```
http://developer.intel.com/software/products/vtune/vtune61/index.htm.
```

Note: VTune may not be available for this release. Consult your release notes for details about its availability.

Using GuideView

GuideView is a graphical tool that presents a window into the performance details of a program's parallel execution. GuideView is part of the KAP/Pro Toolset, which also includes the Guide OpenMP compiler and the Assure Thread Analyzer. GuideView is not a part of the default software installation on your system.

GuideView uses an intuitive, color-coded display of parallel performance bottlenecks which helps pinpoint performance anomalies. It graphically illustrates each processor's activity at various levels of detail by using a hierarchical summary.

Statistical data is collapsed into relevant summaries that indicate where attention should be focused (for example, regions of the code where improvements in local performance will have the greatest impact on overall performance).

To gather programming statistics, use the `-O3`, `-openmp`, and `-openmp_profile` compiler options. This causes the linker to use `libguide_stats.a` instead of the default `libguide.a`. The following example demonstrates the compiler command line used to produce a file named `swim`:

```
% efc -O3 -openmp -openmp_profile -o swim swim.f
```

To obtain profiling data, run the program, as in this example:

```
% export OMP_NUM_THREADS=8  
% ./swim < swim.in
```

When the program finishes, the `swim.gvs` file is produced and it can be used with GuideView. To invoke GuideView with that file, use the following command:

```
% guideview -jpath=your_path_to_Java -mhz=998 ./swim.gvs.
```

The graphical portions of GuideView require the use of Java. Java 1.1.6-8 and Java 1.2.2 are supported and later versions appear to work correctly. Without Java, the functionality is severely limited, but text output is still available and is useful, as the following portion of the text file that is produced demonstrates:

```
Program execution time (in seconds):
cpu           :           0.07 sec
elapsed      :          69.48 sec
  serial     :           0.96 sec
  parallel   :          68.52 sec
cpu percent   :           0.10 %
end

Summary over all regions (has 4 threads):
# Thread      #0      #1      #2      #3
Sum Parallel  :  68.304  68.230  68.240  68.185
Sum Imbalance :   1.020   0.592   0.892   0.838
Sum Critical Section:  0.011   0.022   0.021   0.024
Sum Sequential :   0.011  4.4e-03  4.6e-03  1.6e-03
Min Parallel  :  -5.1e-04 -5.1e-04  4.2e-04 -5.2e-04
Max Parallel  :   0.090   0.090   0.090   0.090
Max Imbalance :   0.036   0.087   0.087   0.087
Max Critical Section:  4.6e-05  9.8e-04  6.0e-05  9.8e-04
Max Sequential :   9.8e-04  9.8e-04  9.8e-04  9.8e-04
end
```

Other Performance Tools

The following performance tools also can be of benefit when you are trying to optimize your code:

- Guide OpenMP Compiler is an OpenMP implementation for C, C++, and Fortran from Intel.
- Assure Thread Analyzer from Intel locates programming errors in threaded applications with no recoding required.

For details about these products, see:

<http://developer.intel.com/software/products/threading>

Note: These products have not been thoroughly tested on SGI systems. SGI takes no responsibility for the correct operation of third party products described or their suitability for any particular purpose.

Index

A

Array Services, description, 18
Autopartition option, 8

B

Booting the CD, 7, 13
Bus numbering, 23
Buttons for returning, 7

C

Cache performance, 42
CD contents, 2
Changing the ELF header, 43
Compilers, 36
Comprehensive System Accounting (CSA) support, 22
CpuMemSet support, 21
Cpuset system, 22
CSA support, 22

D

Disk partition table, 6
dlook command, 34
dplace command, 34

for MPI, 47
for OpenMP, 48

E

EFI commands, 30
EFI support, 29
Electric Fence debugger, 39
Environment variables, 48
Ethernet device naming, 25
Extensible Firmware Interface (EFI) support, 29

F

False sharing, 46
Filesystem recovery, 13
Firewall configuration, 9
FLEXlm support, 32

G

gdb tool, 37
GNU debugger, 37
GuideView tool, 51

H

- hardware platforms, 1
- HPC support
 - libraries and tools, 28
 - Message Passing Toolkit, 28
 - Performance Co-Pilot (PCP), 29

I

- idb tool, 38
- Installation
 - additional users, 9
 - CDs, 11
 - log, 11
 - method selection, 8, 14
 - open/free source software, 11
 - overview, 4
 - package group selection, 9, 11
 - password for access, 9
 - proprietary software, 11
 - reboot, 12
 - SGI Linux Environment 7.2 base OS, 6
 - SGI ProPack for Linux, 11
 - system controller software, 12
 - time zone, 9
 - upgrading, 13
 - user creation, 9
- I/O subsystems
 - for HPC systems, 23
 - XFS filesystem support, 27
 - XSCSI support, 27
 - XSCSI-SCSI support, 27
 - XVM support, 28
- IP address selection, 8

L

- L1 and L2 controllers, 33
- Language selection, 7, 9, 14
- Linux Kernel Crash Dump (LKCD), 6
- LKCD for file configuration, 6

M

- Message Passing Toolkit
 - for parallelization, 44
 - support, 28
 - using dplace, 47
 - using profile.pl, 50

N

- NUMA tools, 33
 - dllook command, 34
 - dplace command, 34

O

- Open source website, 4
- OpenMP
 - environment variables, 48
 - Guide OpenMP Compiler, 52
 - with dplace, 48
- OS configuration settings, 1
- OS enhancements
 - Comprehensive System Accounting (CSA), 22
 - CpuMemSets, 21
 - cpuset support, 22
 - for HPC environments, 21
 - partitioning, 23

P

Package group selection, 9, 11
 Parallelization
 automatic, 45
 using OpenMP, 45
 using MPI, 44
 Partition selection, 8
 Password requirement, 5
 Performance
 Assure Thread Analyzer, 52
 Guide OpenMP Compiler, 52
 GuideView, 51
 measuring, 18
 VTune, 51
 Performance Co-Pilot support, 29
 Persistent bus numbering, 23
 Persistent naming
 bus numbers, 23
 Ethernet devices, 25
 XSCSI devices, 27
 pfmon tool, 49
 Product list, 17
 Product support, 17, 35
 profile.pl script, 50
 Profiling
 pfmon, 49
 profile.pl, 50
 PROM support, 29

R

Release features, 1
 Remove partitions, 8
 Rescue mode, 13
 Root access password, 9

S

Screen
 Complete, 12
 Installation to Begin, 11
 Package Group Selection, 9, 11
 Package Installation, 11
 Welcome, 11
 SGIconsole support, 32
 Software
 installation, 5
 introduction, 1
 planning, 5
 System controller firmware support, 33
 System Controller Software
 CD, 12
 System management
 EFI, 29
 FLEXlm, 32
 PROM chips, 29
 SGIconsole, 32
 system controller firmware, 33
 tools, 29

T

Tools
 Assure Thread Analyzer, 52
 Guide OpenMP Compiler, 52
 GuideView, 51
 pfmon, 49
 profile.pl, 50
 VTune, 51
 Tuning
 cache performance, 42
 ddd tool, 38
 debugging tools
 Electric Fence, 39
 gdb, 37

- idb, 38
- determining needs, 41
- dplace, 47
- Electric Fence, 39
- environment variables, 48
- false sharing, 46
- heap corruption, 39
- managing memory, 43
- multiprocessor code, 44
- parallelization, 44
- profiling
 - GuideView, 51
 - mpirun command, 50
 - pfmon, 49
 - profile.pl script, 50
 - VTune analyzer, 51
- recommended compilers, 36
- single processor code, 36
- solving bottlenecks, 46
- system configuration, 35
- using chatr, 43
- using compiler options, 41
- using dplace, 47
- using math functions, 40
- using runon, 47
- verifying correct results, 37

U

- Unsupported elements, 20
- Upgrading from previous releases, 10
- Upgrading software, 13

V

- VTune performance analyzer, 51

X

- XFS filesystem support, 27
- XSCSI support, 27
- XSCSI-SCSI support, 27
- XVM support, 28