

SGI® ProPack™ for Linux® Start Here

007-4558-012

CONTRIBUTORS

Written by Julie Boney and Terry Schultz

Graphics by Chrystie Danzer

Production by Karen Jacobson

COPYRIGHT

© 2003, 2004, 2005, Silicon Graphics, Inc. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, Altix, IRIX, OpenGL, Origin, Onyx, Onyx2, and XFS are registered trademarks and NUMAflex, NUMAlink, OpenGL Multipipe, OpenGL Performer, OpenGL Vizserver, OpenGL Volumizer, OpenMP, Performance Co-Pilot, SGI Advanced Linux, SGI ProPack, SGIconsole, SHMEM, and Silicon Graphics Prism are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

SGI Advanced Linux Environment 3 is based on Red Hat Enterprise Linux AS 3, but is not sponsored by or endorsed by Red Hat, Inc. in any way.

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 and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries. Windows is a registered trademark of Microsoft Corporation in the United States and/or other countries. All other trademarks mentioned herein are the property of their respective owners.

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
005	August 2003 Updated for SGI ProPack v2.2.1 for Linux release
006	October 2003 Updated for SGI ProPack v2.3 for Linux release
007	February 2004 Updated for SGI ProPack v2.4 for Linux release
008	May 2004 Updated for SGI ProPack 3 for Linux release
009	August 2004 Updated for the SGI ProPack 3 for Linux Service Pack 1 release
010	November 2004 Updated for the SGI ProPack 3 for Linux Service Pack 2 release

- 011 January 2005
 Updated for the SGI ProPack 3 for Linux Service Pack 3 release
- 012 February 2005
 Updated for the SGI ProPack 3 for Linux Service Pack 4 release

New Features in This Guide

This update of *SGI ProPack for Linux Start Here* supports the SGI ProPack 3 for Linux Service Pack 4 release.

Major Documentation Changes

Changes in this guide for the SGI ProPack 3 for Linux SP4 release include the following:

- Updated information about software CDs in “CD Contents” on page 2.
- Updated information in “Installation Overview” on page 5.
- Updated information in “Documentation Updated for This Release” on page 6.

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	xx
Conventions	xx
Reader Comments	xx
1. Release Features	1
Software Introduction	2
CD Contents	2
Installation Overview	5
Documentation Updated for This Release	6
2. Software Planning and Installation	7
Software Planning	8
Table of Disk Partitions	8
Linux Kernel Crash Dump	8
Latest PROM Firmware	9
Installing SGI Linux Software Over the Network - Recommended	9
Set up the NFS Server	9
Set up the TFTP Server	11
Set up the DHCP Server	12
Install the SGI Advanced Linux Environment 3	13
Install SGI ProPack 3 SP4	14

Installing the SGI Advanced Linux Environment 3 Base OS	14
License Manager Updates	18
Upgrading from Previous SGI ProPack Releases to SGI ProPack 3 for Linux Service Pack 4	18
Installing SGI ProPack 3 for Linux Service Pack 4	19
Alternate SGI ProPack Installation Methods	20
CD Media to Create ISO Images	21
Install from ISO Images on Local Disk	21
Install from Remote CDs	22
Installing the System Controller Software 1.9 CD	23
Obtaining the Latest SGI Altix System Firmware	24
Upgrading Your Software	26
Recovering a Damaged Root Filesystem	26
Manually Mounting the Root Filesystem	26
3. Product Support	27
SGI ProPack for Linux Products	27
Operating System Enhancements	33
CpuMemSets Support	33
Cpuset Support	34
Comprehensive System Accounting (CSA)	34
Partitioning	35
I/O Subsystems	35
Persistent PCI-X Bus Numbering	36
Persistent Naming of Ethernet Devices	37
XSCSI Subsystem	39
XSCSI-SCSI Subsystem	39
Using Standard Linux SCSI Drivers on XSCI Device	40
XFS Filesystem	40
XVM Volume Manager	40
Tape Stream Driver	40
HPC Application Tools and Support	41
Message Passing Toolkit	41
Performance Co-Pilot	41

System Management	42
PROM Chips	42
Extensible Firmware Interface (EFI)	42
FLEXlm	45
SGIconsole	45
System Controller Firmware.	46
L1 and L2 System Controllers	46
L3 System Controller	47
NUMA Data Placement Tools	47
dlook Command	47
dplace Command	48
Graphics APIs and Software	48
OpenGL Multipipe	48
OpenGL Multipipe Software Development Kit (SDK)	49
OpenGL Performer	49
OpenGL Vizserver	50
OpenGL Volumizer	50
4. Performance Tuning	51
Index	53

Figures

Figure 1-1	SGI ProPack 3 for Linux Service Pack 3 Release CD Contents	. . . 4
-------------------	--	---------

Tables

Table 2-1	devfs Disk Partitions	8
Table 3-1	SGI ProPack 3 for Linux Service Pack 4 Products	28
Table 3-2	EFI Commands	43

About This Guide

This guide provides information about the SGI ProPack 3 for Linux Service Pack 4 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 systems.

Note: Performance tuning information for single processor and multiprocessor programs now resides in the *Linux Configuration and Operations Guide*. Detailed tuning information can be found in the *Linux Application Tuning Guide* available with this release.

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 3 for Linux Service Pack 4 - 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 3 for Linux Service Pack 4 Documentation CD (CD4) 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_prg.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/itanium2/manuals/251110.htm>.
- *Intel Itanium Architecture Software Developer's Manual*, available online at <http://developer.intel.com/design/itanium/manuals/iiasdmanual.htm>.
- *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- Porting to SGI Altix Systems*
Provides information on programming, integrating, and controlling drivers.

- *Porting IRIX Applications to SGI Altix Platforms: SGI ProPack for Linux*
Provides information about porting an application to the SGI Altix platform.
- *Message Passing Toolkit (MPT) User's Guide*
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 Configuration and Operations Guide*
Provides information on how to perform system configuration and operations for SGI ProPack servers.
- *Linux Resource Administration Guide*
Provides a reference for people who manage the operation of SGI ProPack servers and contains information needed in the administration of various system resource management features such as Comprehensive System Accounting (CSA), Array Services, CPU memory sets and scheduling, and the Cpuset System.
- *SGI Altix 350 System User's Guide*
Provides an overview of the Altix 350 system components, and it describes how to set up and operate this system.
- *SGI Altix 350 Quick Start Guide*
Guides a knowledgeable user through the installation, setup, and simple configuration of most SGI Altix 350 systems.
- *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 Altix 3700 Bx2 User's Guide*
This guide provides an overview of the architecture and descriptions of the major components that compose the SGI Altix 3700 Bx2 family of servers. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.

- *Silicon Graphics Prism Visualization System User's Guide*
Provides an overview of the Silicon Graphics Prism Visualization System components, and it describes how to set up and operate this system.
- *SGI ProPack 3 for Linux Service Pack 4 Release Notes*
Provides 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 2.1 Start Here*
Provides an introduction to SGIconsole and information about setting up and configuring SGIconsole hardware and software.
- *Console Manager for SGIconsole Administrator's Guide*
Provides information about the Console Manager software graphical interface allows you to control multiple SGI servers, SGI partitioned systems, and large single-system image servers.
- *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 and SGI Altix 350 systems.
- *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.
- *XVM Volume Manager Administrator's Guide*
Describes the configuration and administration of XVM logical volumes using the XVM Volume Manager.
- *Event Manager User Guide*
Provides information about the Event Manger application that collects event information from other applications. This document describes the Event Manager application, the application programming interface that you can use to access it, the procedures that you can use to communicate with it from another application, and the commands that you can use to control it.
- *Embedded Support Partner User Guide*
Provides information about using the Embedded Support Partner (ESP) software suite to monitor events, set up proactive notification, and generate reports. This revision of the document describes ESP version 3.0, which is the first version of ESP that supports the Linux operating system.

- *Linux Application Tuning Guide*
Provides information about tuning application programs on SGI Altix systems. Application programs include Fortran and C programs written on SGI Linux systems with the compilers provided by Intel.
- *SCSL User's Guide*
Provides information about the scientific libraries on SGI Altix systems and SGI IRIX systems. Topics include discussions about BLAS, LAPACK, and FFT routines.
- *SGI OpenGL Multipipe User's Guide*
Describes the OpenGL Multipipe product, which allows you to run single-pipe applications in a multipipe environment without modification.
- *OpenGL Performer Getting Started Guide*
Introduces the most important concepts and classes in the OpenGL Performer library.
- *OpenGL Performer Programmer's Guide*
Provides information about OpenGL Performer that is a programming interface (with ANSI C and C++ bindings) for creating real-time graphics applications and offers high-performance, multiprocessed rendering in an easy-to-use 3D graphics toolkit.
- *SGI OpenGL Vizserver Administrator's Guide*
This document is intended for system administrators and gives information about installing, configuring, tuning, and troubleshooting OpenGL Vizserver.
- *SGI OpenGL Vizserver User's Guide*
Provides information about OpenGL Vizserver, you can interact with images and three-dimensional models utilizing the capabilities of remote SGI servers.
- *SGI OpenGL Volumizer 2 Programmer's Guide*
Documents OpenGL Volumizer 2, a C++ volume rendering toolkit optimized for SGI scalable servers. It provides the developer with the tool set needed to solve the problems inherent in high-quality, interactive volume rendering of large datasets. This guide gives an introduction to the OpenGL Volumizer 2 application programming interface (API) and examples of its use.
- *OpenGL Volumizer 2 Reference Pages*
Provides reference page information for OpenGL Volumizer 2.

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
1500 Crittenden Lane, M/S 535
Mountain View, CA 94043-1351

SGI values your comments and will respond to them promptly.

Release Features

This chapter provides an introduction to SGI ProPack 3 for Linux Service Pack 4 and describes the CD contents, provides documentation information, and provides a brief overview of the installation process. It covers the following topics:

- “Software Introduction” on page 2
- “CD Contents” on page 2
- “Installation Overview” on page 5
- “Documentation Updated for This Release” on page 6

Software Introduction

The SGI ProPack 3 for Linux SP4 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 SGI Advanced Linux Environment 3, the base Linux distribution based on Red Hat Enterprise Linux AS 3. SGI ProPack 3 for Linux SP4 is required by and designed to run on any SGI Altix system or Silicon Graphics Prism system.

SGI hardware platforms and OS configuration settings supported by SGI in this release are documented at the following URL:

<http://www.sgi.com/servers/altix/configs.html>

CD Contents

The CDs that you receive with SGI ProPack 3 for Linux Service Pack 4 are as follows (Figure 1-1 lists the contents of each CD):

- SGI Advanced Linux Environment 3
 - SGI Advanced Linux Environment 3 CD1 - Installation CD 1 of 4
 - SGI Advanced Linux Environment 3 CD2 - Installation CD 2 of 4
 - SGI Advanced Linux Environment 3 CD3 - Installation CD 3 of 4
 - SGI Advanced Linux Environment 3 CD4 - Installation CD 4 of 4
 - SGI Advanced Linux Environment 3 CD5 - Source Code CD 1 of 3
 - SGI Advanced Linux Environment 3 CD6 - Source Code CD 2 of 3
 - SGI Advanced Linux Environment 3 CD7 - Source Code CD 3 of 3
- SGI ProPack 3 for Linux
 - SGI ProPack 3 for Linux Service Pack 4 CD1 Open/Free Source Software CD 1 of 2
 - SGI ProPack 3 for Linux Service Pack 4 CD2 Open/Free Source Software CD 2 of 2
 - SGI ProPack 3 for Linux Service Pack 4 CD3 Proprietary Software CD

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

- SGI ProPack 3 for Linux Service Pack 4 CD4 Open/Free Source Code CD
- SGI ProPack 3 for Linux Service Pack 4 CD5 Documentation CD
- System Controller Software 1.9 CD (for IRIX and Linux)

Note: The System Controller Software 1.9 CD is **not** shipped by default. You can download the System Controller Software 1.9 package from <http://support.sgi.com/> or order it on CD-ROM (P/N 812-1138-007). Contact your SGI account representative if you need assistance obtaining the software package.

Note: The SGI System Controller software is licensed for use only on Origin and Onyx 3000 series systems (SGI systems based on MIPS processors) and on SGI Altix 350 systems, SGI Altix 3000 series systems (SGI systems based on Intel Itanium processors), and Silicon Graphics Prism systems. 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.

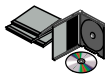

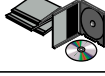




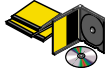




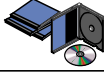
* SGI Advanced Linux Environment 3		
	SGI Advanced Linux Environment 3 Installation CD 1 of 4 CD1	Interactive installation procedures
	SGI Advanced Linux Environment 3 Installation CD 2 of 4 CD2	
	SGI Advanced Linux Environment 3 Installation CD 3 of 4 CD3	
	SGI Advanced Linux Environment 3 Installation CD 4 of 4 CD4	
	SGI Advanced Linux Environment 3 Source Code CD 1 of 3 CD5	Source code for base OS and common open source applications
	SGI Advanced Linux Environment 3 Source Code CD 2 of 3 CD6	
	SGI Advanced Linux Environment 3 Source Code CD 3 of 3 CD7	
* SGI ProPack 3 for Linux Service Pack 4		
	SGI ProPack 3 for Linux Service Pack 4 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 3 for Linux Service Pack 4 Open/Free Source Software CD2	
	SGI ProPack 3 for Linux Service Pack 4 Proprietary Software CD CD3	System PROM, XVM, cpuset and NUMAtools (dlook, dplace, and so on), Array Services, FLEXIm, system partitioning software, MPT
	SGI ProPack 3 for Linux Service Pack 4 Open/Free Source Code CD CD4	Source RPMs for the GPL/LGPL and other open source licensed software that ships on the SGI ProPack CD1
	SGI ProPack 3 for Linux Service Pack 4 Documentation CD CD5	Collection of documentation, FAQs, HOWTOs, and man pages
* System Controller Software 1.9 CD		
	System Controller Software 1.9 CD	System Controller software/firmware (for IRIX and Linux). This CD is only included with the ProPack software if you have requested it. You can always download the latest Software Controller software via Supportfolio.

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

For a complete list of the RPMs included on the SGI Advanced Linux Environment 3 and SGI ProPack 3 for Linux Service Pack 3 CDs listed in Figure 1-1 on page 4, see a file called `RPMs.TXT` that is available on the SGI ProPack 3 for Linux Service Pack 3 Doc CD in the root directory.

For the latest information about software and documentation in this release, see the release notes that are on the SGI ProPack 3 for Linux Service Pack 3 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 3 for Linux SP3 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 3 CDs, then SGI ProPack 3 Service Pack 3 Open/Free Source Software CDs 1 of 2 (CD1) 2 of 2 (CD2). The SGI ProPack 3 Service Pack 3 Open/Free Source Software CDs will then require the SGI ProPack 3 Service Pack 3 Proprietary Software CD (CD3). The SGI ProPack 3 Service Pack 3 Open/Free Source Code CD (CD4) and SGI ProPack 3 for Linux Service Pack 3 Documentation CD (CD5) are optional. Installation is described in Chapter 2, “Software Planning and Installation”.

This release also supports installing your Linux software over the network as described in “Installing SGI Linux Software Over the Network - Recommended” on page 9.

Note: There is no Updates CD for this release. All updates to SGI Advanced Linux Environment 3 have been folded into the SGI ProPack 3 SP3 CDs. To upgrade a system from SGI ProPack 3 to SGI ProPack 3 SP3, simply follow the steps outlined in the section entitled “Installing SGI ProPack 3 for Linux Service Pack 4” in Chapter 2.

Documentation Updated for This Release

The following manuals have been updated for this release:

None for this release.

For a description of these manuals, see “SGI Documentation” on page xvi.

Software Planning and Installation

This chapter describes the planning that you need to do to get started and provides instructions for installing SGI ProPack 3 for Linux Service Pack 4.

Your SGI Altix system or Silicon Graphics Prism system comes with a base Linux distribution (SGI Advanced Linux Environment 3) and the SGI ProPack 3 SP4 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 3 for Linux SP4 software works only with the SGI Advanced Linux Environment 3. Earlier versions of these distributions, or any other distributions, are not compatible with SGI ProPack 3 for Linux SP4 software.

Before you install or configure your system, please read Chapter 1, “Release Features” so that you understand the features of SGI ProPack 3 for Linux SP4 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.

This chapter covers the following topics:

- “Software Planning” on page 8
- “Installing SGI Linux Software Over the Network - Recommended” on page 9
- “Installing the SGI Advanced Linux Environment 3 Base OS” on page 14
- “License Manager Updates” on page 18
- “Upgrading from Previous SGI ProPack Releases to SGI ProPack 3 for Linux Service Pack 4” on page 18
- “Installing SGI ProPack 3 for Linux Service Pack 4” on page 19
- “Alternate SGI ProPack Installation Methods” on page 20

- “Installing the System Controller Software 1.9 CD” on page 23
- “Upgrading Your Software” on page 26
- “Recovering a Damaged Root Filesystem” on page 26
- “Manually Mounting the Root Filesystem” on page 26

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 36 and “XSCSI Subsystem” on page 39.

Linux Kernel Crash Dump

Linux Kernel Crash Dump (LKCD) creates files in `/var/log/dump` directory. To save crash dumps to a different location, change the `DUMPDIR` value in `/etc/sysconfig/dump` file.

The `/var/log` directory size will grow as Performance Co-Pilot and Comprehensive System Accounting (CSA) log to their own directories in `/var/log`. 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 the `/var/log/dump` directory. This is why the default root filesystem is larger. For this reason, some customers may wish to make a separate `/var/log/dump` filesystem or change the configuration of `lkcd`.

Latest PROM Firmware

Unless you have the most current PROM, the kernel may fail to boot during the installation procedure. The prom RPM is on ProPack 3 SP4 CD3 in SGI/RPMS directory. For information on obtaining the latest PROM, see “Obtaining the Latest SGI Altix System Firmware” in Chapter 2.

Installing SGI Linux Software Over the Network - Recommended

You can install the SGI Advanced Linux Environment 3 Base OS and SGI ProPack 3 for Linux SP4 over the network. Because of the large number of CD swaps required when installing from physical CD media, SGI recommends installing over the network. You will need DHCP, TFTP, and NFS servers for this task. They can be on the same system or on different systems.

Set up the NFS Server

Perform the following steps to set up the NFS server to install SGI ProPack 3 SP4 over the network:

1. To set up the NFS server for network booting, you need to create separate `sgiadvlinux-3` and `propack-3SPX` (where X is 1, 2, and so on) directories, as follows:

```
mkdir -p /some_path/install/sgiadvlinux-3
mkdir -p /some_path/install/propack-3SP4
```
2. Copy the SGI Advanced Linux Environment 3 ISO images into the `/some_path/install/sgiadvlinux-3` directory. You can hard link the images if they are in the same filesystem.

The `sgiadvlinux` directory must **not** contain any other images, they will confuse the installer. These are copied as raw ISO images, the NFS installer can look inside an ISO to get the files it needs.

3. Copy the contents of the first three SGI ProPack 3 SP4 CDs to directories in `/some_path/install/propack-3SP4`.
4. If you are using physical CD media, perform the following:
 - a. Insert the SGI ProPack 3 for Linux Service Pack 4 CD1 into the system's CD-ROM drive, and perform the following:

```
mount /dev/cdrom /mnt/cdrom
mkdir /some_path/install/propack-3SP4/cd1
/bin/cp -a /mnt/cdrom/. /some_path/install/propack-3SP4/cd1
umount /mnt/cdrom
```
 - b. Insert the SGI ProPack 3 for Linux Service Pack 4 CD2 into the system's CD-ROM drive, and perform the following:

```
mount /dev/cdrom /mnt/cdrom
mkdir /some_path/install/propack-3SP4/cd2
/bin/cp -a /mnt/cdrom/. /some_path/install/propack-3SP4/cd2
umount /mnt/cdrom
```
 - c. Insert the SGI ProPack 3 for Linux Service Pack 4 CD3 into the system's CD-ROM drive, and perform the following:

```
mount /dev/cdrom /mnt/cdrom
mkdir /some_path/install/propack-3SP4/cd3
/bin/cp -a /mnt/cdrom/. /some_path/install/propack-3SP4/cd3
umount /mnt/cdrom
```
5. If using ISO images of the SGI ProPack 3 for Linux Service Pack 4 CDs, perform the following:

Note: For information on creating ISO images from CD media, see “CD Media to Create ISO Images” on page 21.

a. cd /directory/containing/iso/files

- b. For ProPack 3 Service Pack 4 CD1, perform the following:

```
mount -o ro,loop propack-3SP4-cd1-open.iso /mnt/cdrom
mkdir /some_path/install/propack-3SP4/cd1
/bin/cp -a /mnt/cdrom/. /some_path/install/propack-3SP4/cd1
umount /mnt/cdrom
```


- c. For ProPack 3 Service Pack 4 CD2, perform the following:

```
mount -o ro,loop propack-3SP4-cd2-open.iso /mnt/cdrom
mkdir /some_path/install/propack-3SP4/cd2
/bin/cp -a /mnt/cdrom/. /some_path/install/propack-3SP4/cd2
umount /mnt/cdrom
```

- d. For ProPack 3 Service Pack 4 CD3, perform the following:

```
mount -o ro,loop propack-3SP4-cd3-prop.iso /mnt/cdrom
mkdir /some_path/install/propack-3SP4/cd3
/bin/cp -a /mnt/cdrom/. /some_path/install/propack-3SP4/cd3
umount /mnt/cdrom
```

6. Add entries to export the directories via NFS to the target hosts in the `/etc/exports` file. Then run command `exportfs -a` to implement the changes (see the `exports(5)` and `exportfs(8)` man pages for more information.)

Make sure that `/some_path/install/sgiadvlinux-3` and `/some_path/install/propack-3SP4` are exported over NFS.

Set up the TFTP Server

Perform the following steps to set up the TFTP server to install SGI ProPack 3 SP4 over the network:

1. Decide where the tftp server is going to fetch its images from. For Linux, this is specified in `/etc/xinetd.d/tftp` under `server_args`, option `-s`. For example,

```
server                = /usr/sbin/in.tftpd
server_args           = -s /some_path/tftpboot
disable               = no
```

2. Create the `sgiadvlinux-3` directory for tftp. There are no ProPack 3 SP4 files on the tftp server.

```
mkdir -p /some_path/tftpboot/sgiadvlinux-3
```

3. On the NFS server, mount `/some_path/install/sgiadvlinux-3/sgiadvlinux-3-cd1-inst.iso` over loopback. Within `sgiadvlinux-3-cd1-inst.iso`, mount `images/boot.img` over a second loopback. Copy the contents of `boot.img` (contains directory `efi`) to the `tftpboot` directory on the tftp server, as follows:

```
mount -o ro,loop
/some_path/install/sgiadvlinux-3/sgiadvlinux-3-cd1-inst.iso
mkdir /mnt/tmp1
mkdir /mnt/tmp2
/mnt/tmp1
mount -o ro,loop /mnt/tmp1/images/boot.img /mnt/tmp2
/bin/cp -a /mnt/tmp2/efi /some_path/tftpboot/sgiadvlinux-3
umount /mnt/tmp2 /mnt/tmp1
```

4. Edit `/some_path/tftpboot/sgiadvlinux-3/efi/boot/elilo.conf` and add this line to the "install" entry, as follows:

```
append="console=ttyS0 askmethod"
```

Set up the DHCP Server

Perform the following steps to set up the DHCP server to install SGI ProPack 3 SP4 over the network: (The syntax for this will vary depending on which DHCP server you are using.)

1. Find the MAC from the output of command `ifconfig eth0` on the system you want to install (labeled `HWaddr`). This can also be found from the EFI boot messages or in the `efi` boot menu next to the network boot option.

An example `ifconfig eth0` output is, as follows:

```
eth0      Link encap:Ethernet  HWaddr 08:00:69:13:F0:B4
```

An example EFI message is, as follows:

```
Initializing supported network interfaces
```

```
tg0: ASIC|Revision|MAC Address      : 5701|B5|08:00:69:13:f0:b4
tg0: Link Negotiation|Advertisement : On|<H10 F10 H100 F100 H1000
F1000>
tg0: Link|Speed|Duplex|Rx/Tx FlowCtrl: Up|1000|Full|Off/Off
```

Example EFI boot menu option is, as follows:

```
[Pci(4|0)/Mac(08006913F0B4)/NicName(tg0) ]
```

2. Create client entry in `/etc/dhcpd.conf` for the system you want to install using the MAC address from Step 1. Include a fixed-address entry with the `ipaddress` of the client. Include a `next-server` entry with the `tftp` server hostname. And include a `filename` entry with the path to the boot file on the `tftp` server (Note: Do not include the `/some_patch/tftpboot` prefix from the `tftp` configuration file, that is automatically added by the `tftp` server.

An example `dhcpd.conf` entry for a subnet containing the client with MAC address in the example above:

```
subnet 192.12.1.0 netmask 255.255.255.0 {
}

# Begin group options
group {
    option subnet-mask 255.255.255.0;
    option broadcast-address 192.12.1.255;
    option routers 192.12.1.1;
    option domain-name "your.co.com";
    option domain-name-servers 192.12.1.10, 192.12.1.11;

    # Device information
    host eva00 {
        hardware Ethernet 08:00:69:13:F0:B4;
        fixed-address 192.12.1.25;
        next-server tftphost.your.co.com;
        filename "/efi/boot/bootia64.efi";
    }
}
```

- Restart `dhcpd` server to pick up the new configuration. (Use `/etc/init.d/dhcpd restart` in `sgiadvlinux`).

Install the SGI Advanced Linux Environment 3

To install the SGI Advanced Linux Environment 3 over the network, perform the following:

- Boot the Altix system to be installed, select network boot.
- EFI issues DHCP requests, the DHCP server responds with the IP address and the next server and filename.

Note: Like all DHCP software, the DHCP server must be on the same network as the Altix system or you must have a DHCP relay on the same network as the Altix system to forward the requests to the final DHCP server.

When it asks for the installation source, select **NFS**.

3. Enter the IP address, netmask, gateway, DNS servers of the Altix system. Selecting DHCP may or may not work, there are timing problems with EFI DHCP and the spanning tree algorithm on some Cisco routers.
4. Enter the IP address of the NFS server and the path to the merged install directory, for example, `/some_path/install/sgiadvlinux-3` (on a test machine, it looks as follows: `/big/linux_distributions/sgi/install/sgiadvlinux-3`)

Everything else is the same as a normal install, except it is completely hands off. No need to swap the SGI Advanced Linux Environment 3 CDs.

Install SGI ProPack 3 SP4

To install SGI ProPack 3 SP4 over the network, perform the following:

1. Reboot the system after installing SGI Advanced Linux Environment 3.
2. Log into the system and mount the Propack 3 SP4 directory on the NFS server as `/mnt/source`, as follows:

```
mkdir -p /mnt/source
mount nfs.server:/some_path/install/propack-3SP4 /mnt/source
```

3. Run the installer, as follows:

```
/mnt/source/INSTALL
```

In a few minutes, your system will be ready.

Installing the SGI Advanced Linux Environment 3 Base OS

This section provides steps for installing the SGI Advanced Linux Environment 3 CD1, CD2, CD3, and CD4 set.

Note: Installation of the SGI Advanced Linux Environment 3 Base OS using physical CD media may require up to 15 CD swaps due to RPM dependencies.

For console setup options and instructions, see the section titled “Monitoring Your Server” in the *SGI Altix 3000 User Guide*, the section titled “Operating Your System” in

the *SGI Altix 350 System User's Guide*, or chapter 1, "Getting Started with the Silicon Graphics Prism" in the *Silicon Graphics Prism Visualization System User's Guide*.

When a Silicon Graphics Prism Visualization System comes up after the installation of SGI ProPack 3 SP4, you **cannot** login from the console. You need to boot the system with a serial console connected and then perform start operations in that console.

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 3 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,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. Type `cd efi\boot` to the boot image directory, as follows:

```
fs0:\> cd efi\boot
fs0:\efi\boot>
```

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

```
fs0:\efi\boot> elilo  
ELILO boot:
```

Press **ENTER** to boot the Linux kernel.

Note: If the kernel fails to boot here with a message about the PROM being too old, see “Obtaining the Latest SGI Altix System Firmware” in Chapter 2.

5. The CD will then boot and start the first "CD Found" screen. Skip the media test and start the installation. You may follow the screen instruction to check your CDs. To begin the installation process insert CD1 into the CDROM drive and press **Continue**.
6. Select the language you would prefer to use for the installation and as the system default. To continue, select **OK**.

Note: SGI recommends that English be used during installation. Other languages may not display properly on non-UTF-8 capable terminals.

7. The next screen is the SGI ProPack information screen. To continue, select **OK**.
8. You will now see a window titled **Installation type**. Currently, **Custom** is the only choice. To continue, select **OK**.
9. In the partition screen, select the **Autopartition** option and press Enter on your keyboard. When warning messages appear, select **IGNORE**.
10. 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**.
11. When the warning message appears about losing all data on the devices, select **Yes** to confirm the partition layout.
12. 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**.

13. 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 37. (This does not affect most users.)

14. If you chose to use a static IP address on the previous screen, you are prompted for Gateway and DNS for the system, fill in the required information and then to continue, select **OK**.
15. The next screen asks for the hostname of the system. If you had chosen the DHCP (in step 13), you may select the system hostname automatically via DHCP.
16. 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**.
17. 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.

18. The next screen allows you to select which time zone this system is in. Select the appropriate information and to continue, select **OK**.
19. 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**.
20. 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.
21. The **Installation to begin** screen is informational. To continue, select **OK**.
22. The installation process will now begin. During the process, you will be prompted to insert other SGI Advanced Linux Environment 3 CDs. Do this when instructed and to continue, select **OK**.
23. 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 installed. After the system restarts, proceed to “Installing SGI ProPack 3 for Linux Service Pack 4” on page 19.

License Manager Updates

The license manager was upgraded in the SGI ProPack v2.3 release to FLEXlm 9.0. FLEXlm licenses from any release earlier than SGI ProPack v2.3 may not work on SGI ProPack 3 SP4. You may need replacement FLEXlm licenses for SGI applications in certain situations.

Applications running on SGI ProPack 3 SP4 (and the prior releases SGI ProPack v2.3 and v2.4) use a different license identifier on SGI Altix series systems than applications running on SGI ProPack releases prior to SGI ProPack v2.2.1.

If you have any licensed SGI applications running on your Altix system, please contact SGI to request a new license key or contact your local support center and request a new key at this location:

<http://www.sgi.com/support/supportcenters.html>

Otherwise, these applications may not allow access after you install SGI ProPack 3. New licenses generated for FLEXlm 9.0 will work on SGI ProPack releases running SGI Advanced Linux Environment 3. SGI strongly recommends that you replace the licenses for your SGI applications before installing SGI ProPack 3 to avoid any possible interruption in service. Third party applications are unaffected by this change.

Upgrading from Previous SGI ProPack Releases to SGI ProPack 3 for Linux Service Pack 4

Upgrades from prior SGI ProPack releases to the SGI ProPack 3 for Linux Service Pack 4 release are only supported from the SGI ProPack 3 for Linux release. To upgrade a system currently running SGI ProPack 3, simply follow the steps in “Installing SGI ProPack 3 for Linux Service Pack 4” on page 19.

To install SGI ProPack 3 SP4 on a system running SGI ProPack v2.3 or v2.4, you must do a fresh install. For information on how to manage a fresh install while retaining such things as modified configuration files, see the release notes.

Note: The release notes, which contain the latest information about software and documentation in this release, are on the SGI ProPack 3 for Linux Service Pack 4 Documentation CD in the root directory, in a file named README.TXT.

Installing SGI ProPack 3 for Linux Service Pack 4

This section provides steps for installing the SGI ProPack 3 for Linux SP4 CD1 (Open/Free Source Software CD 1 of 2), SGI ProPack 3 for Linux SP4 CD2 (Open/Free Source Software CD 2 of 2), and the SGI ProPack 3 for Linux SP4 (CD3 (Proprietary Software)).

For information on alternate installation methods, see “Alternate SGI ProPack Installation Methods” on page 20.

For console setup options and instructions, see the section titled “Monitoring Your Server” in the *SGI Altix 3000 User Guide* or chapter 1, “Installation and Operation” in the *SGI Altix 350 System User’s Guide*, or “Connecting an L1 System Console” in chapter 1 of the *Silicon Graphics Prism User’s Guide*.

1. Log in to the system as `root`.
2. Insert SGI ProPack 3 for Linux Service Pack 4 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. If you want to install software from CD1, CD2, and CD3 select **ALL**, or if you only want to install software from CD1 and CD2 select **Open Source Only**.
6. The next screen lets you to customize the selection of software installed. Check the **Customize software selection** box if you want to do so.
7. The **Package Group Selection** screen appears if you have selected to customize software selection in previous screen. 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. If you have selected Open Source Only in the installation type screen (item 5), you will only have three groups - "Base updates", "Kernel Development", and "Software Development", and they are all selected by default. If you have selected **ALL** in the installation type screen (item 5), you will have three more groups - "SGI Proprietary", "SGI Licensed" and "SGI Cluster Infrastructure". The last two are not selected by default. 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**.

8. The **Installation to Begin** screen appears. It tells you that a log of the installation will be placed in `/tmp/sgi-install.log`. Press Enter. (**OK** is highlighted by default).
9. 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.

10. During the installation process, you will be asked to insert SGI ProPack 3 for Linux Service Pack 4 CD2 and CD3. Insert each disk when prompted and press Enter.
11. After the installation is complete, the **Complete** screen appears. Press Enter. You are returned to the `root` prompt.
12. After the software installation is complete, reboot your system to begin using the newly installed SGI ProPack 3 for Linux Service Pack 4 software environment. You can reboot by typing `reboot` and pressing Enter. You can then remove the CD from the drive.

Alternate SGI ProPack Installation Methods

In addition to using CD media and direct physical access to the CD drive on an Altix system, there are alternate methods you can use to install SGI ProPack, including:

- Install from ISO Images on Local Disk
- Install from Remote CDs

Note: SGI Advanced Linux Environment 3 or SGI ProPack 3 must already be installed on a system to use this installation method.

Note: ISO images for SGI ProPack 3 for Linux Service Pack 4 are available for download from Supportfolio for customers with support contracts. You will need the ISO images for SGI ProPack CD1, CD2, and CD3. In the examples below, the ISO images have been downloaded to a directory called `/tmp/propack_iso/`.

CD Media to Create ISO Images

If you have CDs, but want to ISO images, you can use the CD media to create ISO images, as follows:

Insert the CD into the system and mount it with the following command:

```
mount -t iso9660 -o ro /dev/cdrom /mnt/cdrom
```

Then enter the following command to create the ISO image:

```
dd if=/dev/cdrom of=/tmp/propack_iso/propack_cd1.iso
```

You may see I/O errors at the end of the copy, but these can be safely ignored.

Install from ISO Images on Local Disk

SGI ProPack can be installed by copying the ISO images to the local drive on the system. This installation method allows you to install SGI ProPack without physical access to the system and without having to swap CD media. It also shields the installation from any network hiccups.

1. After copying the ISO images to your system, login to the system as root and execute the following commands to mount the first ISO image to a new directory:

```
mkdir -p /mnt/source
```

```
mount -o loop /tmp/propack_iso/propack-3SP4-cd1-open.iso  
/mnt/source
```

2. Next, execute the following command to kick off the installation process:

```
/mnt/source/INSTALL
```

3. Follow steps 4-12 in the section entitled “Installing SGI ProPack 3 for Linux Service Pack 4” on page 19. When the installation asks for CD2, open another shell window and execute the following commands to unmount SGI ProPack CD1 and then mount CD2:

```
umount /mnt/source
```

```
mount -o loop /tmp/propack_iso/propack-3SP4-cd2-open.iso  
/mnt/source
```

Repeat the unmount and mount process again in the second shell window when asked for SGI ProPack CD3.

Note: The "not mounted" error messages from the unmount command can be safely ignored.

4. After the installation is complete, you are returned to the root prompt. Review the `elilo.conf` file (`/boot/efi/efi/sgi/elilo.conf`) to confirm that the new kernel is listed as the default kernel to boot. Unmount the last ISO image and reboot by typing `reboot` and pressing **Enter**.

Install from Remote CDs

SGI ProPack can also be installed using CDs exported from another client machine. The following steps assume a machine to which you have easy physical access such as a desktop PC running Linux and running Linux and a remote Altix machine. In the steps below, these machines are referred to as `desktop_pc` and `altix_server`, respectively.

Note: NFS differences on the client and server machines as well as network hiccups might prevent this method from working.

1. Insert the SGI ProPack CD1 media into the cd drive on `desktop_pc` and mount the CD:

```
mount /dev/cdrom /mnt/cdrom
```

2. Export `/mnt/cdrom` on the `desktop_pc` to the Altix machine by adding an entry in the `/etc/exports` file:

```
/mnt/cdrom    altix_server(*,sync)
```

After adding the entry to `/etc/exports`, use the following commands to stop and then start the NFS daemon:

```
/sbin/server nfs stop
```

```
/sbin/server nfs start
```

See the `exports(5)` manpage for more information on exporting NFS filesystems.

3. Use the `telnet` command to connect into `altix_server` and mount the remote CD exported from `desktop_pc` to a local directory like `/mnt/source` with the following command:

```
mount -o ro desktop_pc:/mnt/cdrom /mnt/source
```

4. From the same telnet window, start the ProPack installation:

```
/mnt/source/INSTALL
```

Note: The installation copies quite a bit of data over to `/tmp` on the local machine, so this could take several minutes depending on the speed of your network.

5. When the installer asks for SGI ProPack CD2 and CD3, open another shell to `altix_server` and unmount `/mnt/source`, swap the CDs in `desktop_pc` and then mount `desktop_pc:/mnt/cdrom` again to `/mnt/source`.
6. After the installation is complete, you are returned to the root prompt. Review the `elilo.conf` file (`/boot/efi/efi/sgi/elilo.conf`) to confirm that the new kernel is listed as the default kernel to boot. Unmount the last ISO image and reboot by typing `reboot` and pressing **Enter**.

Installing the System Controller Software 1.9 CD

This section provides steps for installing the System Controller Software 1.9 (for IRIX and SGI ProPack for Linux). This software is licensed for SGI Origin and Onyx 3000 series systems based on MIPS processors and for SGI Altix 350 systems and the SGI Altix 3000 family of servers and superclusters.

Note: The System Controller Software 1.9 CD is **not** shipped by default. You can download the System Controller Software 1.9 package from <http://support.sgi.com/> or order it on CD-ROM (P/N 812-1138-007). Contact your SGI account representative if you need assistance obtaining the software package.

1. Log in to the PC workstation or laptop computer where the System Controller Software 1.x" package resides, as `root`, using the root password you used in the previous installation process.
2. Insert the System Controller Software 1.9 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 1.9 CD.

Obtaining the Latest SGI Altix System Firmware

There are three levels of firmware to consider on your SGI Altix system, as follows:

- L2 controller software (only configurations of SGI Altix systems with routers have L2 controllers)
- L1 controllers on routers, IX-bricks, PX-bricks, C-bricks, and so on.
- The system PROM (each node has a system PROM and each C-brick has two nodes).

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.

You can always download the latest L1/L2 system controller software and PROM firmware via Supportfolio at <http://support.sgi.com/linux/>

You can also find system controller software by using the search term “system controller” at: http://support.sgi.com/search_request/downloads/index.

PROM is released as patches and the latest firmware can also be found at: http://support.sgi.com/browse_request/linux_patches_by_os

Note: Newer PROM firmware is often incompatible with previous versions of the PROM firmware. If you decide to manually flash PROMs, you must upgrade the PROM on all nodes at once or some nodes may hang during boot. For more information on this topic, please see the SGI ProPack 3 for Linux release notes.

To update or “flash” the system PROM, you first need to download the prom RPM from Supportfolio as described above.

Note: The PROM image on your hard drive is automatically updated when you upgrade your system to the latest SGI ProPack for Linux release. The installation of the `snprom` RPM (step 1) happens automatically as part of the ProPack installation. Flashing it (steps 2-4) must be performed.

Once you have the latest PROM RPM, flash the PROM, as follows:

1. Install the update PROM image with the following command:

```
# rpm -Uvh RPM_NAME...
```

2. Reboot the system and enter the EFI shell. Access `/boot/efi` (typically `fs0`) from the EFI Shell> prompt:

```
Shell> fs0:
```

3. Issue the `ls` command and locate the PROM image binary:

```
fs0:\> ls
```

```
snprom.bin
```

Verify the PROM version with the following command:

```
fs0:\> flash -v snprom.bin
```

```
SGI PROM Flashing Utility
```

```
Version of prom image in file 3.32
```

Verify that the current PROM on the system is older with the following command:

```
fs0:\> flash -R  
Running PROM version: 3.25
```

4. Flash all the nodes with the latest PROM image:

```
fs0:\> flash -a snprom.bin
```

Caution: Always allow the flash process to complete before attempting any other system activities. Flashing the PROMs may take several minutes (depending on system size). Interrupting a flash process can result in disabled processors and nodes.

Instructions for updating (flashing) the L1/L2 system controller firmware are also included in the *SGI Altix 3000 User's Guide* and the *SGI L1 and L2 Controller Software User's Guide*.

Upgrading Your Software

When you are ready to upgrade your SGI ProPack software, follow the instructions outlined in the section titled “Installing SGI ProPack 3 for Linux Service Pack 4” on page 19, 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

Information about recovering a damaged root filesystem has been moved to Chapter 2, “System Operation” in the *Linux Configuration and Operations Guide*.

Manually Mounting the Root Filesystem

Information about manually mounting the root filesystem has been moved to Chapter 2, “System Operation” in the *Linux Configuration and Operations Guide*.

Product Support

This chapter documents the product components that are supported on the SGI Altix 350 system, SGI Altix 3000 systems, and the Silicon Graphics Prism Visualization System. (For a list of the products, see Table 3-1 on page 28.)

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

- “SGI ProPack for Linux Products” on page 27
- “Operating System Enhancements” on page 33
- “I/O Subsystems” on page 35
- “HPC Application Tools and Support” on page 41
- “System Management” on page 42
- “NUMA Data Placement Tools” on page 47
- “Graphics APIs and Software” on page 48

SGI ProPack for Linux Products

Software provided by SGI for the SGI ProPack 3 for Linux Service Pack 4 release consists of a kernel RPM for the SGI Altix 350, SGI Altix 3000, and Silicon Graphics Prism products, SGI Advanced Linux Environment 3.0 RPMs, and value-added software developed by SGI to run specifically on SGI systems. Table 3-1 provides a description of the SGI ProPack 3 for Linux SP4 products.

Table 3-1 SGI ProPack 3 for Linux Service Pack 4 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. For more information on application tuning, see the <i>Linux Application Tuning Guide</i>.</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. For more information, see the Array Services chapter in the <i>Linux Resource Administration Guide</i> .
CpuMemSets	Provides the kernel support and infrastructure for implementing processor and memory placement. For more information, see the CPU Memory Sets and Scheduling chapter in the <i>Linux Resource Administration Guide</i> .
Cpuset System	The Cpuset System is primarily a workload manager tool permitting a system administrator to restrict the number of processors that a process or set of processes may use. A system administrator can use cpusets to create a division of CPUs within a larger system. For more information, see the Cpuset System chapter in the <i>Linux Resource Administration Guide</i> .

Table 3-1 SGI ProPack 3 for Linux Service Pack 4 Products **(continued)**

Product	Description
CSA	Provides jobs-based accounting of per-task resources and disk usage for specific login accounts on Linux systems. Linux CSA application interface library allows software applications to manipulate and obtain status about Linux CSA accounting methods. For more information, see the CSA chapter in the <i>Linux Resource Administration Guide</i> .
FLEXlm	Provides a floating license, run-time environment. Includes daemons suitable for serving floating licenses.
IOC4 serial driver	Driver that supports the Internal IDE CD-ROM, NVRAM, and Real-Time Clock. Serial ports are supported on the IOC4 base I/O chipset and the following device nodes are created: /dev/ttyIOC4/0 /dev/ttyIOC4/1 /dev/ttyIOC4/2 /dev/ttyIOC4/3
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. For more information on system partitioning, see the <i>Linux Configuration and Operations Guide</i> .
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(2)</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. For more information on system controller commands, see the <i>SGI L1 and L2 Controller Software User's Guide</i> .

Table 3-1 SGI ProPack 3 for Linux Service Pack 4 Products **(continued)**

Product	Description
LKCD	Provides system crash dump analysis tools including <code>lcrash(1)</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. For more information on MPT, see the <i>Message Passing Toolkit: MPI Programmer's Manual</i> .
NUMA tools	Provides a collection of NUMA related tools (<code>dlook(1)</code> , <code>dplace(1)</code> , and so on). For more information on NUMA tools, see the <i>Linux Configuration and Operations Guide</i> .
Performance Co-Pilot collector infrastructure	Provides performance monitoring and performance management services targeted at large, complex systems. For more information on Performance Co-Pilot, see <i>Performance Co-Pilot for IA-64 Linux User's and Administrator's Guide</i> .
PROM	Allows the CPU to boot and allows you to perform system administration and software installations. For more information on the levels of PROM in your system and how to flash PROM firmware, see "Obtaining the Latest SGI Altix System Firmware" on page 24.
runon	Enables running a command on a particular CPU or set of CPUs. For more information on the runon command, see the <code>runon(8)</code> man page and the <i>Linux Configuration and Operations Guide</i> and the <i>Linux Application Tuning Guide</i> .
XFS	Provides a high-performance filesystem for Linux. For more information on XFS, see the <i>XFS for Linux Administration</i> .

Table 3-1 SGI ProPack 3 for Linux Service Pack 4 Products **(continued)**

Product	Description
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> or the <i>SGI Altix 350 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 (/dev/sda, etc).
XVM	Provides software volume manager functionality such as disk striping and mirroring. For more information on XVM, see the <i>XVM Volume Manager Administrator's Guide</i> .
Tape Stream Driver	The SGI tape stream (TS) tape driver is low-level software that supports all of tape backup applications on SGI Altix systems. These applications include the Data Migration Facility (DMF), the Tape Management Facility (TMF), OpenVault, and xfsdump(8)/xfsrestore(8), the XFS filesystem incremental dump/restore utilities. For more information on the TS driver, see the <i>Linux Configuration and Operations Guide</i> .
OpenGL Multipipe	OpenGL Multipipe enables owners of SGI high performance visualization solutions to achieve higher levels of performance with existing single pipe applications by bringing the power of scalable visualization to applications. For more information, see the <i>SGI OpenGL Multipipe User's Guide</i> .
OpenGL Multipipe SDK	OpenGL Multipipe Software Development Kit (SDK) is an application programming interface (API) layer for OpenGL that provides a simplified approach to managing graphics applications across multiple graphics subsystems or pipes.
OpenGL Performer	OpenGL Performer is a powerful and comprehensive programming interface for developers creating real-time visual simulation and other professional performance-oriented 3D graphics applications. For more information, see the <i>OpenGL Performer Getting Started Guide</i> .

Table 3-1 SGI ProPack 3 for Linux Service Pack 4 Products **(continued)**

Product	Description
OpenGL Vizserver	OpenGL Vizserver is a technical and creative computing solution designed to deliver superior visualization, collaboration functionality, and advanced performance to any existing client desktop. For more information, see the <i>SGI OpenGL Vizserver User's Guide</i> and <i>SGI OpenGL Vizserver Administrator's Guide</i> .
OpenGL Volumizer	OpenGL Volumizer offers a cross platform, high-level volume rendering API designed for interactive, high quality, scalable visualization of large volumetric data sets. For more information, see the <i>SGI OpenGL Volumizer 2 Programmer's Guide</i> .

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 and LVM. SGI supports md for root and swap mirroring only; other md configurations are not supported. 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 *Linux Resource Administration Guide* 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(2)` and `runon(8)`.

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. For further documentation and details on cpusets, see chapter 5, “Cpuset System” in the *Linux Resource Administration Guide*.

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).

The Linux CSA application interface library allows software applications to manipulate and obtain status about Linux CSA accounting methods.

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 *Linux Resource Administration Guide*.

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 *Linux Configuration and Operations Guide*.

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-sgi21r4
    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 36), 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.

Using Standard Linux SCSI Drivers on XSCSI Device

For information on using standard Linux SCSI device drivers on XSCSI devices, see the *Linux Configuration and Operations Guide*.

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.

This release adds a new XVM multi-host failover feature. For more information on this new feature and XVM Volume Manager in general, see *XVM Volume Manager Administrator's Guide*.

Tape Stream Driver

The SGI tape stream (TS) tape driver is low-level software that supports all of tape backup applications on SGI Altix systems. These applications include the Data Migration Facility (DMF), the Tape Management Facility (TMF), OpenVault, and

`xfsdump(8)/xfsrestore(8)`, the XFS filesystem incremental dump/restore utilities. For more information on the TS driver, see the *Linux Configuration and Operations Guide*.

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 (MPT) User's Guide*.

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.

For more information on Performance Co-Pilot, see the *Performance Co-Pilot for IA-64 Linux User's and Administrator's Guide*.

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 and the L1 controller, see the *SGI L1 and L2 Controller Software User's Guide*.

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 3 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
<code>endif</code>	Script-only: Delimits IF THEN construct
<code>err [level]</code>	Sets or displays error level
<code>exit</code>	Exits
<code>flash filename</code>	Flashes PROM on C-brick
<code>for var in set</code>	Script-only: Indicates loop construct
<code>getmtc</code>	Gets next monotonic count

Table 3-2 EFI Commands (continued)

EFI Command	Description
<code>goto label</code>	Script-only: Jumps to label location in script
<code>guid [-b] [sname]</code>	Dumps known guid IDs
<code>help [-b] [internal command]</code>	Displays this help
<code>if [not] condition then</code>	Script-only: Indicates IF THEN construct
<code>load driver_name</code>	Loads a driver
<code>ls [-b] [dir] [dir] ...</code>	Obtains directory listing
<code>map [-bdvr] [sname[:]] [handle]</code>	Maps shortname to device path
<code>mem [address] [size] [;MMIO]</code>	Dumps memory or memory mapped I/O
<code>memmap [-b]</code>	Dumps memory map
<code>mkdir dir [dir] ...</code>	Makes directory
<code>mm address [width] [;type]</code>	Modifies memory: Mem, MMIO, IO, PCI
<code>mode [col row]</code>	Sets or gets current text mode
<code>mount BlkDevice [sname[:]]</code>	Mounts a filesystem on a block device
<code>mv sfile dfile</code>	Moves files
<code>pause</code>	Script-only: Prompts to quit or continue
<code>pci [bus dev] [func]</code>	Displays PCI device(s) info
<code>reset [cold/warm] [reset string]</code>	Indicates cold or warm reset
<code>rm file/dir [file/dir]</code>	Removes file or directories
<code>set [-bdv] [sname] [value]</code>	Sets or gets environment variable
<code>setsize newsize fname</code>	Sets the files size
<code>stall microseconds</code>	Delays for <i>x</i> microseconds
<code>time [hh:mm:ss]</code>	Gets or sets time
<code>touch [filename]</code>	Views or sets file attributes

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

EFI Command	Description
<code>type [-a] [-u] [-b] file</code>	Types file
<code>ver</code>	Displays version information
<code>vol fs [volume label]</code>	Sets or displays volume label

FLEXIm

FLEXIm 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 FLEXIm.

For more information, visit

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

SGIconsole

SGIconsole is a suite of software that allows you to manage multiple SGI servers running the IRIX or Linux operating system. These servers include SGI servers, SGI partitioned systems and large, single-system-image servers, including legacy SGI servers and Silicon Graphics graphical systems. SGIconsole can also provide console services for non SGI servers.

SGIconsole consists of software suite including the Console Manager package and Performance Co-Pilot, which provides access to common remote management tools for hardware and software. SGIconsole can manage SGI Altix and SGI Origin servers better than any other console server product. SGIconsole can communicate with L1, L2, module System Controller (MSC), and rackmount system multimodule System Controller (MMSC) system controllers and provide reset/nonmaskable interrupt (NMI)/power up/power down functionality for SGI servers.

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

This section describes the system controller firmware used in SGI Altix systems.

L1 and L2 System Controllers

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.

The following references describe how to install and connect to an L1 or L2 controller:

- Procedure 3-10, “Determining the Hardware Address of the L2 System Controller” in chapter 3 of the *SGIconsole 2.0 Start Here* describes how the L2 system controllers come online after a system boots and how to set the eth1 port on the SGI 1100 server to match when a DHCP server is running on a subnet.
- Procedure 3-2, “Determining the L2 Address of a New L2 Controller” in chapter 3 of the *Console Manager for SGIconsole Administrator’s Guide* describes how to find the L2 address of a new L2 controller.
- Procedure 4-1, “Connecting to a System Console” in chapter 4 of the *Console Manager for SGIconsole Administrator’s Guide* describes how to connect to a system console.
- Procedure 2-2, “Connecting to the L1 Controller” in chapter 2 of the *Linux Configuration and Operations Guide* describes how to connect to an L1 controller and from there an L2 controller with or without SGIconsole installed.

- Procedure 2-3, “Connecting to the L2 Controller” in chapter 2 of the *Linux Configuration and Operations Guide* describes how to connect to the L2 controller with or without SGIconsole installed.

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

L3 System Controller

The L3 controller is a Linux software package that resides on a PC workstation or laptop computer. It provides remote support, system controller consoles, and system-level maintenance functions. The L3 controller software is obtained through the "System Controller Software 1.x" release.

NUMA Data Placement Tools

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

d1ook Command

The `d1ook(1)` 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(1)` 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.

For more information on NUMA tools, see chapter 4, “NUMA Tools” in the *Linux Configuration and Operations Guide* and chapter 5, “Data Placement Tools” in the *Linux Application Tuning Guide*.

Graphics APIs and Software

This section describes graphics APIs and software supported on SGI ProPack 3 SP3 and covers the following topics:

- “OpenGL Multipipe” on page 48
- “OpenGL Multipipe Software Development Kit (SDK)” on page 49
- “OpenGL Performer” on page 49
- “OpenGL Vizserver” on page 50
- “OpenGL Volumizer” on page 50

OpenGL Multipipe

OpenGL Multipipe enables owners of high performance visualization solutions from Silicon Graphics to create high-resolution SGI Reality Center environments and achieve higher levels of performance while continuing to use existing single pipe applications. It brings the power of scalable visualization to new applications and enables organizations to create better products, generate new ideas, and increase operational efficiency.

OpenGL Multipipe enables high levels of screen space scaling and low-to-moderate levels of performance scaling without application modification. OpenGL Multipipe uses a flexible architecture that enables a logical screen to be built up from individual physical screens that can be located anywhere within the logical screen space. This enables the creation of large walls with 2D arrays of displays, wide panoramic screens with overlapping displays for edge blending and the use of a scalable graphics

compositor to drive a single screen with increased performance. For more information about OpenGL Multipipe, visit <http://www.sgi.com/products/software/multipipe/>.

OpenGL Multipipe Software Development Kit (SDK)

OpenGL Multipipe SDK is an API layer for OpenGL that provides a simplified approach to managing graphics applications across multiple graphics subsystems or pipes and focused on delivering run-time scalability on SGI scalable graphics servers.

Applications written using the OpenGL Multipipe SDK API can run seamlessly from desktop single-processor, single-pipe systems up to large multiprocessor, multipipe scalable graphics systems. A simple configuration file defines the desired configuration of the application, such that the same application can support normal office environments, SGI Reality Center rooms, walls, immersive desk displays, and domes. In addition, the same application can take advantage of scalable graphics systems by combining the power of multiple graphics pipes into a single display output, which in turn can be scaled to drive multiple display devices. OpenGL Multipipe SDK provides a straightforward solution for graphics applications that were designed for single-processor, single-pipe applications to scale seamlessly when run on graphics systems with multiple processors and multiple graphics pipes. For more information about OpenGL Multipipe SDK, visit <http://www.sgi.com/products/software/multipipe/sdk/>.

OpenGL Performer

OpenGL Performer is a powerful and comprehensive programming interface for developers creating real-time visual simulation and other professional performance-oriented 3D graphics applications. The toolkit simplifies development of applications used for visual simulation, manufacturing, simulation-based design, virtual reality, scientific visualization, interactive entertainment, broadcast video, architectural walk-through, and computer aided design.

The latest release is built atop the industry standard OpenGL graphics library, interoperates with OpenGL Volumizer, OpenGL Multipipe SDK, and OpenGL Vizserver, includes both ANSI C and C++ bindings, and is available for the IRIX operating system, and the Linux, Windows XP, and Windows 2000 operating systems. For more information about OpenGL Performer, visit <http://www.sgi.com/software/performer/>.

OpenGL Vizserver

OpenGL Vizserver is a technical and creative computing solution designed to deliver superior visualization, collaboration functionality, and advanced performance to any client, whether on a desktop workstation or wireless tablet. OpenGL Vizserver allows users to remotely view and interact with large data sets from any other system at any location in an organization and to collaborate with multiple colleagues using these same applications and data.

OpenGL Vizserver enables a single SGI advanced visualization system to act as a multiuser graphics server which distributes remote and collaborative visualization sessions to client systems running a mix of desktop operating systems. With OpenGL Vizserver, graphics processing is handled entirely on the SGI advanced visualization system using full 3D hardware acceleration; the client machine simply decompresses and displays the image. This leverages existing internal infrastructure and allows customers to leverage the power of the world-class SGI advanced visualization systems on common desktop workstations such as Silicon Graphics visual workstations, Sun Solaris systems or Intel architecture PCs running the Linux, Microsoft Windows NT, Windows 2000, or Windows XP operating systems. For more information about OpenGL Vizserver, visit <http://www.sgi.com/products/software/vizserver/>.

OpenGL Volumizer

OpenGL Volumizer is a cross platform, high-level volume rendering API for the energy, manufacturing, medical, and sciences markets. OpenGL Volumizer is a revolutionary graphics API designed for interactive, high quality, scalable visualization of large volumetric data sets. OpenGL Volumizer provides a high-level interface to OpenGL hardware to allow application writers and researchers to visualize multiple gigabytes of volumetric data. The API uses OpenGL for volume rendering and hence allows standard graphics applications to treat volumetric and surface data in a similar fashion.

OpenGL Volumizer is a library of C++ classes that facilitates the manipulation and display of volumetric data sets common in geo-science, medical, scientific and engineering applications. It provides a layer of functionality that sits on top of OpenGL and integrates seamlessly into higher-level toolkits and applications. For more information on OpenGL Volumizer, visit <http://www.sgi.com/products/software/volumizer/>.

Performance Tuning

Note: The performance tuning information previously in chapter 4 is now available in Chapter 3, “Performance Tuning” in the new *Linux Configuration and Operations Guide* available with the SGI ProPack 3 for Linux Service Pack 3 release. Detailed application performance tuning information is available in the *Linux Application Tuning Guide*.

Index

A

Array Services, description, 28
Autopartition option, 16

B

Booting the CD, 15
Bus numbering, 36
Buttons for returning, 15

C

CD contents, 2
Comprehensive System Accounting (CSA) support,
34
CpuMemSet support, 33
Cpuset system, 34
CSA support, 34

D

Disk partition table, 8
dlook command, 47
dplace command, 48

E

EFI commands, 43
EFI support, 42
Ethernet device naming, 37
Extensible Firmware Interface (EFI) support, 42

F

FLEXlm support, 45

G

Graphics APIs
OpenGL Multipipe, 48
OpenGL Multipipe SDK, 49
OpenGL Performer, 49
OpenGL Vizserver, 50
OpenGL Volumizer, 50

H

hardware platforms, 2
HPC support
libraries and tools, 41
Message Passing Toolkit, 41
Performance Co-Pilot (PCP), 41

I

Installation

- CDs, 19
- log, 20
- open/free source software, 19
- overview, 5
- package group selection, 19
- password for access, 17
- proprietary software, 19
- reboot, 20
- SGI Linux Environment 7.2 base OS, 14
- SGI ProPack for Linux, 19
- system controller software, 23
- time zone, 17
- upgrading, 26

I/O subsystems

- for HPC systems, 35
- XFS filesystem support, 40
- XSCSI support, 39
- XSCSI-SCSI support, 39
- XVM support, 40

IP address selection, 17

L

- L1 and L2 controllers, 46
- Language selection, 16
- Linux Kernel Crash Dump (LKCD), 8
- LKCD for file configuration, 8

M

- Message Passing Toolkit support, 41

N

- NUMA tools, 47
 - dllook command, 47
 - dplace command, 48

O

- Open source website, 5
- OS configuration settings, 2
- OS enhancements
 - Comprehensive System Accounting (CSA), 34
 - CpuMemSets, 33
 - cpuset support, 34
 - for HPC environments, 33
 - partitioning, 35

P

- Package group selection, 19
- Partition selection, 16
- Password requirement, 7
- Performance
 - measuring, 28
- Performance Co-Pilot support, 41
- Persistent bus numbering, 36
- Persistent naming
 - bus numbers, 36
 - Ethernet devices, 37
 - XSCSI devices, 39
- Product list, 27
- Product support, 27
- PROM support, 42

R

Release features, 1
Remove partitions, 16
Root access password, 17

S

Screen
 Complete, 20
 Installation to Begin, 20
 Package Group Selection, 19
 Package Installation, 20
 Welcome, 19
SGIconsole support, 45
Software
 installation, 7
 planning, 8
System controller firmware support, 46
System Controller Software
 CD, 23
System management
 EFI, 42
 FLEXlm, 45
 PROM chips, 42
 SGIconsole, 45
 system controller firmware, 46
 tools, 42

U

Unsupported elements, 32
Upgrading from previous releases, 18
Upgrading software, 26

X

XFS filesystem support, 40
XSCSI support, 39
XSCSI-SCSI support, 39
XVM support, 40

