

SGI ProPack™ 4 for Linux® Start Here

007-4727-001

CONTRIBUTORS

Written by Terry Schultz

Production by Karen Jacobson

COPYRIGHT

© 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 ProPack, SGIconsole, SHMEM, and Silicon Graphics Prism are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

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. SUSE LINUX and the SUSE logo are registered trademarks of Novell, Inc. 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 2005 Original publication

Contents

Record of Revision	iii
Tables	ix
Related Publications	xii
SUSE/Novell Documentation	xii
Intel Compiler Documentation	xiii
Other Intel Documentation	xiii
SGI Documentation	xiii
Obtaining SGI Publications	xvii
Conventions	xvii
Reader Comments	xvii
1. Release Features	1
Software Introduction	1
CD Contents	2
System Controller Software 1.9 CD	3
SGI Open Source Information	4
Installation Overview	4
2. Software Planning and Installation	5
Software Planning	6
Table of Disk Partitions	6
Systems with Multiple Ethernet Configurations	6
SLES9 SP1 Kernels	6
Linux Kernel Crash Dump	8
Latest PROM Firmware	8
License Manager Updates	9
Read Me First - Installing SLES9 with Service Pack 1 on Altix	9

How To Install SGI ProPack 4 for Linux with YaST 12
Instructing YaST to Look at Your Install Media 13
Installing from CD 13
Installing from ISO Images 13
Other Installation Sources 14
Installing SGI ProPack with YaST 15
Installing SGI ProPack with YaST - Using the Text (non-graphical) Interface 15
Installing SGI ProPack with YaST - Using the Graphical Interface 16
Installing the SN2 Kernel 17
Notes on Installing SGI ProPack with YaST 18
SLES and SGI ProPack Network Installation for Experts 18
Setting up the Network Install Server 18
Using AutoYaST 19
Installing the System Controller Software 1.9 CD 21
Obtaining and Installing the Latest SGI Altix System Firmware 22
Reinstalling Your Software 24
3. Product Support 25
SGI ProPack for Linux Products 25
Operating System Enhancements. 29
Cpuset Support 29
Comprehensive System Accounting (CSA). 29
Partitioning. 30
I/O Subsystems. 30
Persistent PCI Bus Numbering 31
Persistent IP Addressing of Ethernet Interfaces 32
XFS Filesystem 33
XVM Volume Manager 33
HPC Application Tools and Support. 34
Message Passing Toolkit 34
Performance Co-Pilot 35

System Management	35
PROM Chips	35
Extensible Firmware Interface (EFI)	35
FLEXlm	38
SGIconsole	38
System Controller Firmware.	39
L1 and L2 System Controllers	39
L3 System Controller	40
NUMA Tools	40
dlook Command	40
dplace Command	41
Index	43

Tables

Table 2-1	Disk Partitions	6
Table 2-2	SLES9 SP1 Kernels	7
Table 3-1	SGI ProPack 4 for Linux Products	26
Table 3-2	EFI Commands	36

About This Guide

This guide provides information about the SGI ProPack 4 for Linux release. It is divided into the following chapters:

- Chapter 1, “Release Features,” describes the major features of this release, the CD contents, SGI Open Source software, SGI system controller software, 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.

Online versions of this document, the SGI ProPack 4 release notes, which contain the latest information about software and documentation in this release, the list of RPMs distributed with SGI ProPack 4, and a useful migration guide, which contains helpful hints and advice for customers moving from earlier versions of SGI ProPack to SGI ProPack 4, can be found in the `/docs` directory on the SGI ProPack 4 Open/Free Source CD.

Related Publications

Documents listed in this section contain additional information that might be helpful, as follows:

- “SUSE/Novell Documentation” on page xii
- “Intel Compiler Documentation” on page xiii
- “Other Intel Documentation” on page xiii
- “SGI Documentation” on page xiii

SUSE/Novell Documentation

Documentation for SUSE LINUX Enterprise Server 9 is located on your system in the `/usr/share/doc` directory. There are also several helpful documents provided on SLES9 CD1 in the `/docu/en` directory.

The following documentation is available:

- SUSE LINUX Enterprise Server 9 Installation and Administration (manual.pdf), available online at:
http://www.novell.com/documentation/sles9/pdfdoc/sles_9_admin_guide/sles_9_admin_guide.pdf
- SUSE LINUX Enterprise Server 9 Start-Up Guide (startup.pdf), available online at:
http://www.novell.com/documentation/sles9/pdfdoc/sles_9_startup_guide/sles_9_startup_guide.pdf
- SUSE LINUX Enterprise Server 9 Package Descriptions available online at
<http://www.novell.com/products/linuxenterpriseserver/packages.html>
- Migrating from Red Hat to SUSE LINUX Enterprise Server 9: A Novell Migration Study available online at
http://www.novell.com/resourcecenter/ext_item.jsp?itemId=13485
- Novell YES CERTIFIED Bulletins for SGI available online at
<http://developer.novell.com/yesssearch/Search.jsp> (select "SGI" as the Company and "SUSE LINUX Enterprise Server 9 for Intel Itanium Process" as the Novell Product and click Search)

Novell makes SUSE LINUX Enterprise Server documentation available online at:
<http://www.novell.com/documentation/sles9/index.html>

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:

- *SGI ProPack 3 to SGI ProPack 4 Migration Guide*
This guide contains information about migrating from SGI ProPack 3 for Linux operating system to SGI ProPack 4 for Linux operating system.
- *SGI ProPack 4 for Linux 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`.

- *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 3700Bx2 User's Guide*
This guide provides an overview of the architecture and descriptions of the major components that compose the SGI Altix 3700Bx2 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 User's Guide*
Provides an overview of the Silicon Graphics Prism components, and it describes how to set up and operate this system.
- *SIconsole 2.1 Start Here*
Provides an introduction to SIconsole and information about setting up and configuring SIconsole hardware and software.
- *Console Manager for SIconsole 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 4 for Linux 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 1
- “CD Contents” on page 2
- “System Controller Software 1.9 CD” on page 3
- “SGI Open Source Information” on page 4
- “Installation Overview” on page 4

Software Introduction

The SGI ProPack 4 for Linux product includes capabilities and performance improvements ideal for enabling technical and creative users to solve their big compute and data problems by using the Linux operating system and Itanium processors. This product adds to or enhances features in SUSE LINUX Enterprise Server 9 (SLES9) Service Pack 1 (SP1), the base Linux distribution based on SLES 9 SP1. SGI ProPack 4 for Linux is designed to run on any SGI Altix system. The Silicon Graphics Prism Visualization System will be supported in a future release of SGI ProPack 4 for Linux.

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

This section describes the CDs that you receive with SGI ProPack 4 for Linux, as follows:

- **Open Source CD**

The Open Source CD contains open source software licensed under GPL/LGPL or other open source licenses, and has files necessary for installation of the following major products:

- Array Session Module
- Linux Comprehensive System Accounting (CSA)
- Linux jobs (Commands, init scripts, libraries and PAM module)
- LSI Fusion-MPT host adapter management utility
- NUMA tools to manage CPU placement of processes (*bctfree*, *dplace*, *dlook* utilities)
- Performance Co-Pilot, system-level performance monitoring and performance management tool
- Open Source SGI Profile Tools
- SGI ProPack configuration recommendations

For a complete list of the RPMs included on the Open Source CD, see a file called `RPMS.TXT` that is available on the SGI ProPack 4 for Linux Proprietary CD in the root directory.

- **Proprietary CD**

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

The Proprietary CD contains proprietary software for your SGI Altix system and has files necessary for installation of the following major products:

- Flexible File I/O (FFIO) system
- Cpuset Facility workload manager tool
- Embedded Support Partner (ESP) software suite to monitor events, set up proactive notification, and generate reports
- Performance Co-Pilot visualization tools

- Intel Compiler Runtime Libraries
- Array Services
- Message-Passing Toolkit (MPT)
- Scientific Computing Software Library (SCSL)
- XVM Command line tools and graphical user interface (GUI)
- SGI Itanium 2 PROM (System controller software)
- SGI ProPack documentation

For a complete list of the RPMs included on and SGI ProPack 4 for Linux see a file called `RPMS.TXT` that is available on the SGI ProPack 4 for Linux Proprietary 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 4 for Linux Documentation CD in the root directory, in a file named `README.TXT`.

The SGI ProPack 4 for Linux release notes get installed to the following location on a system running SGI ProPack 4:

```
/usr/share/doc/sgi-propack-4/README.txt
```

System Controller Software 1.9 CD

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 330 systems, SGI Altix 350 systems, SGI Altix 1330 and Altix 1350 cluster 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 Open Source Information

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 4 for Linux will most likely come preinstalled on your SGI platform. Installation is described in Chapter 2, “Software Planning and Installation”.

This release also supports installing your Linux software over the network as described in “SLES and SGI ProPack Network Installation for Experts” on page 18.

Software Planning and Installation

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

Your SGI Altix system comes with a base Linux distribution (SUSE LINUX Enterprise Server 9 (SLES9) Service Pack 1 (SP1)) and the SGI ProPack 4 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 4 for Linux software works only with the SLES9 SP1 environment. Earlier versions of this distribution, or any other distribution, are not compatible with SGI ProPack 4 for Linux software.

Before you install or configure your system, please read Chapter 1, “Release Features” so that you understand the features of SGI ProPack 4 for Linux software. You might also want to familiarize yourself with some or all of the documentation listed in “Related Publications” on page xii.

This chapter covers the following topics:

- “Software Planning” on page 6
- “License Manager Updates” on page 9
- “Read Me First - Installing SLES9 with Service Pack 1 on Altix” on page 9
- “How To Install SGI ProPack 4 for Linux with YaST” on page 12
- “SLES and SGI ProPack Network Installation for Experts” on page 18
- “Installing the System Controller Software 1.9 CD” on page 21

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 Disk Partitions

Device	Mount Point	Size
/dev/sda1	/boot/efi	500 MB
/dev/sda2	swap	9 GB
/dev/sda3	/	25 GB

Systems with Multiple Ethernet Configurations

On SGI ProPack 4 with systems with multiple Ethernet connections, there is no guarantee that Ethernet MAC address of the base I/O IO9 or IO10 card becomes `eth0`. Putting a dual-port Gigabit Ethernet card in the PCI slot right next to the IO9 card in an Altix 350 system renders the IO9 tg3 port `eth2`. The dual-port GigaBit Ethernet card gets `eth0` and `eth1`. This can make SUSE network installs very confusing. See the SGI ProPack 4 for Linux release notes for a work around to this problem.

SLES9 SP1 Kernels

Unlike earlier versions of SGI ProPack, SGI ProPack 4 for Linux does **not** contain any kernel packages. SLES9 SP1 contains four (4) kernel variants, all based on the same source, but built with slightly different configuration options, that are outlined in Table 2-2. SGI ProPack 4 provides kernel modules for each of the variants for all SGI software contained in the release.

Table 2-2 SLES9 SP1 Kernels

Kernel Variant	Configuration Differences
kernel-default	ia64 generic kernel maximum 128 CPU support 16KB page size no KDB support
kernel-sn2	sn2 specific kernel maximum 512 CPU support 16KB page size KDB support (must be configured on)
kernel-64k-pagesize	ia64 generic kernel maximum 128 CPU support 64KB page size no KDB support does not work with XFS filesystems of 512 byte blocks, but does work if XFS filesystem is created with larger blocks
kernel-debug	ia64 generic kernel maximum 128 CPU support 16KB page size KDB support (enabled by default) kprobes support dprobes support slab/spinlock debugging enabled

During installation, if you do not explicitly select to install any other kernels, your SGI Altix system will boot with the kernel-default variant after installation. You must explicitly select the other kernels during the SLES9 SP1 installation if you want to use any of the other kernel variants.

Both the kernel-default ("generic") and kernel-sn2 ("sn2") kernels are certified and supported on SGI Altix systems; however, SGI recommends using the sn2 kernel in most cases since this kernel is configured to support larger single system images (maximum

512 processors) and to include the kernel debugger (that is, KDB) to aid customer support should that become necessary.

For more information on installing the sn2 kernel, see “Installing the SN2 Kernel” on page 17.

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 `DUMPPDIR` value in `/etc/sysconfig/dump` file.

Note: Linux Kernel Crash Dump (LKCD) does not currently produce dumps reliably. LKCD is not enabled by default and SGI does not recommend enabling it unless explicitly requested to do so by SGI support personnel. KDB is required to gather information on system panics and hangs at this time.

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 the SGI ProPack 4 Proprietary CD in SGI/RPMS directory.

Note: Your SGI system must have PROM version 4.04 flashed on the system before installing SLES9 SP1. Without PROM version 4.04, the system will fail to boot after installation. SGI Altix software support contract customers can download PROM version 4.04 or later from the following area on Supportfolio:
https://support.sgi.com/browse_request/suse_patch_browse_results

For information on obtaining the latest PROM, see “Obtaining and Installing the Latest SGI Altix System Firmware” in Chapter 2.

License Manager Updates

Please backup your license key files prior to installing SLES9 SP1 and/or SGI ProPack 4. A license key file is **not** included in these RPMs. Your original license key files will need to be restored after you complete your installation.

In terms of FLEXlm software and licensing support, the Redhat and SUSE distributors do not ship any code. ISVs are responsible for obtaining and including FLEXlm components into their products. This includes FLEXlm license server components. SGI, as part of SGI ProPack, includes FLEXlm license server components for SGI products.

There are no differences in the FLEXlm license server components between SGI Propack 3 and SGI Propack 4.

SGI recommends FLEXlm version 9.2.3 (or greater) for the "it64_re3" platform. This information is also available to all customers at this location:

http://www.sgi.com/support/licensing/linux_faq.html

If you encounter any license related problems, please contact your local support center at this location:

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

Read Me First - Installing SLES9 with Service Pack 1 on Altix

This section provides some information about installing SLES9 Service Pack 1 on SGI Altix systems. You need to have SLES9 Service Pack 1 installed before you can install SGI ProPack 4 for Linux.

In general, the documentation that comes with SUSE applies to installing software on an Altix system. This section provides some notes to get you started, as follows:

Note: Please read all the notes below before attempting installation.

- You can find installation documentation on the first SLES9 CD in the `docu` directory. Start by reading `preparation.pdf` file. This is the "SUSE LINUX Enterprise Server Architecture-Specific Information" document.

- It is important to read the SLES9 Service Pack 1 release notes as well. You will find them in a text file named `Notes` in the root of the first SLES9 Service Pack 1 CD. Make sure you read the "**Using the bootable SP CD**" section. This is the suggested method of installing on an Altix.
- The YaST installer does not allow you to remove older partitions and create a clean root by default. To do so, you must run the YaST Partitioner in "Expert Mode" (see page 22 of the *SLES9 Installation and Administration* manual). The default filesystem type for partitions is `reiserfs`. To create partitions with the XFS filesystem type, manually "Edit" the partition and then select XFS. Boot partitions on SGI Altix systems should always use the FAT filesystem.
- There are several different methods to install SLES9 SP1. To use the graphical YaST installation, SGI recommends using the `vnc` or `ssh` methods, assuming the SGI Altix system is connected to a network.

Boot the system from the EFI prompt with the following options:

```
bootia64 console=ttyS0,38400n8 insmod=tg3 vnc=1
```

Provide the IP address that the server should use and answer any other network related questions. When the network setup is complete, you will be able to connect to the server using one of the steps below:

1. From another system on the network, use `vncviewer` or another `vnc` program to connect to the SGI Altix server:

```
vncviewer ip_address_of_server:1
```

2. From another system on the network, point a Java-enabled browser at the server:

```
http://ip_address_of_server:5801
```

To use the graphical YaST installer with `ssh`, boot the system from the EFI prompt with the following options:

```
bootia64 console=ttyS0,38400n8 insmod=tg3 usessh=1
```

After answering questions related to the network setup, you should eventually see the messages similar to the following:

```
Starting SSH daemon ...
```

```
/sbin/ifconfig eth0
```

```
eth0      Link encap:Ethernet  HWaddr 08:00:69:13:DC:F6
```

```
inet addr:###.###.###.### Bcast:###.###.###.255
Mask:255.255.255.0
```

```
*** sshd has been started ***
```

```
*** login using 'ssh -X root@###.###.###.###' ***
```

```
*** run 'yast' to start the installation ***
```

Once the SSH daemon is running, use the `ssh` command to connect to the system and then enter `yast` to start the installation, as follows:

```
ssh -X root@ip_address_of_server
```

```
yast
```

If you choose to install over the serial console or an L2 controller with `xterm`, note that the SGI Altix serial console is not 100% clean at this time and you may run into "screen corruption". Often this can be cleared up by entering `Ctrl-L` from the command line to re-draw the screen. When using `xterm`, remember to set the encoding to UTF-8. The following command line options have yielded fairly good results:

```
xterm -sb -en UTF-8 -geometry 130x40 -title "SGI console" -name "SGI
console" -e "cu -s 38400 -l /dev/ttyS0"
```

- After your system is running SLES9 SP1, you may wish to change the kernel you use to one of the several variants that SLES9 SP1 provides. After a default installation, your system will run the kernel-default variant.

Note: SGI recommends that you install the `sn2` kernel because it supports larger CPU counts (maximum 512) and it is also configured to support use of the kernel debugger (`kdb`) commands, which are used to diagnose system panics and hangs. For more information on installing the `sn2` kernel, see "Installing the SN2 Kernel" on page 17.

- If you plan to set up a system to do network installations, you may find the "Network Install Server Setup" in this document helpful. This section describes how to integrate SGI ProPack 4 with SUSE operating system repositories.
- For information on AutoYaST, see "Using AutoYaST" on page 19.

How To Install SGI ProPack 4 for Linux with YaST

These instructions assume that SUSE SLES9 with SLES9 Service Pack 1 has already been installed on your machine. If you do not yet have SUSE SLES9 with Service Pack 1 installed, please see “Read Me First - Installing SLES9 with Service Pack 1 on Altix” on page 9.

Note: SUSE provides extensive documentation on using YaST. See the `manual.pdf` in the `docu` directory of the first SLES9 CD.

If you want to set up a network install repository that allows you to install SUSE SLES9, SLES 9 Service Pack 1, and SGI ProPack 4 all in one session, see “SLES and SGI ProPack Network Installation for Experts” on page 18.

The two primary steps to installing SGI ProPack 4 using the SUSE YaST install tool are, as follows:

- “Instructing YaST to Look at Your Install Media” on page 13
- Do the actual installation or upgrade using YaST described in “Installing SGI ProPack with YaST” on page 15.

Below, you will find detailed instructions for telling YaST about your source media for various types of media (physical CD, CD ISO images, NFS, and so on). Choose the section that matches the source media you are using. After that, you will move on to the section titled “Installing SGI ProPack with YaST” on page 15. These instructions assume you are logged in as root. If you wish to use the graphical interface of YaST, please be sure your `DISPLAY` environment variable is set properly. YaST will provide the text interface when it is unable to use graphics. If you are using the text based installer from a GNOME terminal window, you might find setting your `TERM` environment variable to `xterm` may help.

Depending on if you are using the graphical interface or the text interface to YaST, you will either click on menu choices with your mouse, or select them using your arrow keys, the `TAB` key, and the `ENTER` key.

Instructing YaST to Look at Your Install Media

This section describes how to instruct YaST to find your install media. The instructions vary depending on your source. It covers these topics:

- “Installing from CD” on page 13
- “Installing from ISO Images” on page 13
- “Other Installation Sources” on page 14

Installing from CD

To install SGI ProPack 4 for Linux from CD, perform the following steps:

1. Insert the first SGI ProPack CD in to the CDROM drive.
2. Start up `yast2`
3. From the **YaST Control Center** menu, select or click on **Software**
4. Select or click on **Change Source of Installation**
5. Select or click on **Add**
6. Select **CD...** from the window
YaST will now attempt to mount the CD and add it to the list of media sources
7. Select or click on **Finish**.
8. Proceed to “Installing SGI ProPack with YaST” on page 15.

Note: You do not need to tell YaST about the 2nd SGI ProPack CD for this install method. YaST will prompt for it automatically when it needs it.

Installing from ISO Images

These instructions show you how to install SGI ProPack 4 for Linux using ISO images. There are many ways to do this. This section just provides one example. The example assumes the images are named `propack1.iso` and `propack2.iso`.

To install SGI ProPack 4 from ISO images, perform the following steps:

1. Copy the ISO images to the `/tmp` directory on your machine.
2. Make mount points for your ISO images. An example is, as follows:

```
mkdir /mnt/propack1
mkdir /mnt/propack2
```
3. Mount the ISO images, an example is, as follows:

```
mount -o ro,loop /tmp/propack1.iso /mnt/propack1
mount -o ro,loop /tmp/propack2.iso /mnt/propack2
```
4. Start `yast2`
5. From the **YaST Control Center** menu, select or click on **Software**.
6. Select or click on **Change Source of Installation**.
7. Select or click on **Add**.
8. Select **Local Directory...**
9. Type in the path to your first CD. In this example, `/mnt/propack1`. Then, select or click on **OK**.
10. Type in the path to your second CD. In this example, `/mnt/propack2`. Then, select or click on **OK**.
11. Select or click on **Finish**.
12. Proceed to “Installing SGI ProPack with YaST” on page 15.

Other Installation Sources

If you are an experienced YaST user, you may also find some of the other install methods useful such as NFS, HTTP, or FTP.

For NFS, you tell YaST to change the source of installation and choose **NFS...** as the source. Then you provide the NFS server name and path to the files.

If you have setup a HTTP or FTP server that provides the packages, you may use those too.

If you wish to install SUSE and SLES over the network all in one setting and, see the section titled “SLES and SGI ProPack Network Installation for Experts” on page 18.

After you have your source set up, you should proceed to either the “How To Install SGI ProPack 4 for Linux with YaST” on page 12.

Installing SGI ProPack with YaST

Read this section if you are installing SGI ProPack for the first time.

After you have told YaST how to find SGI ProPack in “Instructing YaST to Look at Your Install Media” on page 13, you are ready to perform the actual installation as described in this section. These instructions are separated in to two sections, as follows:

- “Installing SGI ProPack with YaST - Using the Text (non-graphical) Interface” on page 15
- “Installing SGI ProPack with YaST - Using the Graphical Interface” on page 16
- “Installing the SN2 Kernel” on page 17
- “Notes on Installing SGI ProPack with YaST” on page 18

Installing SGI ProPack with YaST - Using the Text (non-graphical) Interface

To install SGI ProPack 4 for Linux with YaST using the text interface, perform the following steps:

1. Make sure you are in the YaST install program at the **YaST Control Center** menu.
2. Select **Software**.
3. Select **Install and Remove Software**.
4. Select **Filter** (it is on the upper left part of the screen).
5. Select **Selections**.
6. Now you are in a **Package Categories** window. Scroll down to the bottom of the categories list and you will see package categories starting with **SGI ProPack**.
7. Select all of the **SGI ProPack** categories. You can do this, for example, by highlighting the category and hitting the **ENTER** key.
8. Verify that each **SGI ProPack** category now has a “+” in front of it. This means that YaST will install packages from those categories.
9. Select **OK**.
10. Select **Accept** to start installing the packages.

11. Although no kernels are installed in SGI ProPack at this time, the install process will make some adjustments to which kernel modules get loaded by default when booting. In addition, certain system tuning adjustments are made. It is best to reboot after SGI ProPack installation so these changes take effect.
12. See “Notes on Installing SGI ProPack with YaST” on page 18, if installation does not proceed.

Installing SGI ProPack with YaST - Using the Graphical Interface

To install SGI ProPack 4 for Linux with YaST using the graphical interface, perform the following steps.

1. Make sure you are in the YaST install program at the **YaST Control Center** window.
2. Click on **Software** (it may already be highlighted).
3. Click on **Install and Remove Software**.
4. Click on the pull-down menu to the right of **Filter** on the upper left part of the screen.
5. Click on **Selections**.
6. Make sure you are working with the left hand window with the **Selection** title and not the individual package window on the right.
7. Notice that there is a list of product selections. Each selection has a box to the left of it. Scroll down so you can see the **SGI ProPack** selections at the end of the list.
8. Select all of the **SGI ProPack** categories by clicking on the box to the left of the name.
9. Verify that each **SGI ProPack** has a check mark in the box. This means that YaST will install packages from those categories.
10. Now click **Accept** on the lower right to start installing the packages.
11. Although no kernels are installed in SGI ProPack at this time, the install process will make some adjustments to which kernel modules get loaded by default when booting. In addition, certain system tuning adjustments are made. It is best to reboot after SGI ProPack installation so these changes take effect.

Installing the SN2 Kernel

After your system is running SLES9 SP1, you may wish to change the kernel you use. SLES9 SP1 has several kernels available to choose from. By default, you get the "default" kernel. You may wish to install the sn2 kernel instead.

SGI Altix systems pre-installed with SLES9 SP1 are setup to run the sn2 kernel by default, so you do not need to install the sn2 kernel on top of a pre-installed SGI Altix.

To install the sn2 version of the kernel, insert the first SLES9 SP1 CD in to your CD-ROM. It should automatically mount to `/media/cdrom`. If not, you can mount it there by hand. Then run the following command to install the kernel-sn2 rpm: (replace the X's with the version that actually resides there)

```
% rpm -Uvh /media/cdrom/ia64/update/SUSE-CORE/9/rpm/ia64/kernel-sn2-X.X.X-X.XXX.ia64.rpm
```

You can also use the YaST tool to install it. If filtering by RPM groups, select `kernel-sn2` and have YaST install it for you.

If you manually install the sn2 version of the kernel, you will need to go to one CD to get the kernel and another CD to get the modules:

1. Insert the SLES9 SP1 CD1 into the CD-ROM drive and CD to the correct directory (`/suse/ia64/`).
2. Install sn2 kernel rpm, as follows:

```
rpm -Uvh kernel-sn2-*.rpm
```
3. Insert the SGI ProPack 4 Open/Free Source Software CD into the CD-ROM drive and CD to correct directory (`/sgi/ia64`)
4. Install SGI kernel modules for that kernel (on the SGI ProPack 4 CDs), as follows:

```
rpm -Uvh km-sn2-*.rpm
```

SGI kernel modules for proprietary software, such as XVM, is located in the `/sgi/ia64` directory on the SGI ProPack Proprietary Software CD.

If you use the sn2 kernel, which has KDB support compiled into the kernel, you may want to add the following modules to the `MODULES_LOADED_ON_BOOT` variable in `/etc/sysconfig/kernel` file, as follows:

```
kdbm_pg
kdbm_task
```

```
kdbm_vm  
xfsidbg
```

These kernel modules add additional KDB commands to aid in debugging.

Notes on Installing SGI ProPack with YaST

In some cases, you could be presented with some conflicts. These are a result of unresolved dependencies from previous installation or installation attempts. If you have unresolved dependencies, you will have to resolve them before installation of SGI ProPack 4 will proceed. If the only thing YaST is complaining about is that packages will be added, it is usually safe to let it take the default action. Please note that if packages are added due to these dependencies, YaST will need to install them from the operating system CDs or other media - so you may have to have the media available if you run in to this situation. Most users should not experience this problem.

SLES and SGI ProPack Network Installation for Experts

This section describes how to set up a network install server and provides some information on using AutoYaST. It covers these topics:

- “Setting up the Network Install Server” on page 18
- “Using AutoYaST” on page 19

Setting up the Network Install Server

This section provides you some information on installing SUSE SLES9, SLES9 Service Pack 1, and SGI ProPack 4 for Linux all in one session over the network. This is helpful for systems administrators who have several machines to install and keep up to date.

To do this, you will need to set up an NFS repository that contains the installation files.

Look at the SLES 9 SP 1 release notes for how to setup an NFS repository. You can find the release notes in the text file called `Notes` in the root of the first SLES9 Service Pack 1 CD. The section in the SLES9 SP1 release notes dealing with setting up an installation server is titled "Setting up an installation server."

Once you have an install server set up to handle SLES9 with SP1, you can then add SGI ProPack 4, as well. Simply follow the instructions in the SLES 9 SP1 release notes section titled "Integrating the Service Pack into an installation server". Instead of doing it for a Service Pack, do it for SGI ProPack 4, instead. The process is the same in both cases.

If you are doing manual network installs, you will now be able to tell YaST to configure the software to be installed and you will see the SGI ProPack entries when you filter by selections. Note that SGI ProPack 4 is **not** installed by default. You have to tell YaST you want to install it.

Important: Clients making use of your install server repository should have individual sources listed as their installation source (configured using the **Change Source of Installation** in YaST). See the following example.

For example, do **not** set up an install source as `/install/sles9+sp1+pp4/`. Rather, list each out such as the following:

```
/install/sles9+sp1+pp4/SUSE-CORE-Version-9/CD1
/install/sles9+sp1+pp4/SUSE-SLES-Version-9/CD1
/install/sles9+sp1+pp4/SUSE-SLES-9-Service-Pack-Version-1/CD1
/install/sles9+sp1+pp4/SGI-ProPack-4-for-Linux/CD1
```

Using AutoYaST

AutoYaST is a method to install and configure a system based control file settings. With AutoYaST, you could install and configure a system over the network automatically. It provides similar functionality to kickstart provided by Red Hat.

You can find information on how to use AutoYaST at the following locations:
<http://www.suse.de/~nashif/autoinstall/>

You can find information about AutoYaST in chapter 3 of *SUSE LINUX Enterprise Server Installation and Administration* manual. This resides in `manual.pdf` file in the `docu` directory of the first SLES9 CD. Documentation is also available on an installed SLES9 system in `/usr/share/doc/packages/autoyast2/html`.

This information is meant to augment the instructions already available, as follows:

- If you are using AutoYaST for the first time, it is best to start by selecting **Create Reference Profile** in the **Tools** menu of YaST's **Autoinstallation** setup program.

- Always be sure your installation sources point at individual components. These are configured using the **Change Source of Installation** facility within YaST. If your install source points at the top of an NFS repository (`"/install/sles9+sp1+pp4/"` in the example below), package selections will not work properly as products that are not already installed and will not be listed. Here is an example list of valid sources, as follows:

```
/install/sles9+sp1+pp4/SUSE-CORE-Version-9/CD1  
/install/sles9+sp1+pp4/SUSE-SLES-Version-9/CD1  
/install/sles9+sp1+pp4/SUSE-SLES-9-Service-Pack-Version-1/CD1  
/install/sles9+sp1+pp4/SGI-ProPack-4-for-Linux/CD1
```

- By default, YaST will not set up the EFI Boot Manager entry during automatic installation. To change this, use YaST's **Autoinstallation** setup program and enable this feature in **Boot Loader Configuration**.
- In SLES9 SP1, it can be very confusing to get the proper `root=` entry in the target system's `elilo` configuration. The `root=` in the `elilo` configuration is not fixed up when you tell **Autoinstallation** how to partition target disks. Instead, you need to go to the **System -> Boot Loader Configuration** menu, and configure the settings. The only way to change the `root=` entry is to use the **Edit Configuration Files** option. This allows you to modify the actual target `elilo.conf` file. By default, the root is set to the root of the system from which you are running YaST. Sometimes, not what you want.
- There is a bug in SLES9 SP1 where the `KERNEL_MODULES_LOADED_ON_BOOT` variable will be zeroed out when you install SLES9 + SLES9 Service Pack 1 + SGI ProPack in one session. See the SGI ProPack 4 for Linux release notes for details and for a workaround.
- If you want to configure networking, you will find the **Autoinstallation** tool offers a way to configure your network card. However, the network card SGI uses is not available in the list to select from. It may be easier to just let the hardware get detected on the first boot.
- For Altix, the instructions provided for booting over the network do not exactly match. You will need a `dhcp` and `tftp` server and you can copy the installer kernel, `initrd`, and the `elilo.conf` file from the first SLES 9 SP1 CD. After some modification, you should be able to boot this over the network with the same result as booting from the first SLES CD.
- If you are booting a system over the network for autoinstallation, you might want something similar to the following in your append line of the `elilo` configuration on your `tftp` server. There are of course many other ways to do this.

```
append="ramdisk_size=131072 console=ttyS0,115200n8 escdelay=250  
autoyast=nfs://128.162.243.5/install/ install=nfs://128.162.243.5/install/sles9-sp1-pp4
```

The major configuration actions are, as follows;

- `install=` points it at the SLES NFS repository to get the YaST components to launch the installer.
- `autoyast=` points to the directory that contains your autoyast profile. Note trailing `"/`. The file name it retrieves is the IP address converted to hex as described in the autoyast web documentation.
- These instructions will change significantly depending on which features of YaST. For example, you could have a rules file or a default file.
- Note that you could also choose to use SLP. SLP is documented by SUSE and is a method for clients to discover the location of the install server.

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 330 systems, SGI Altix 350 systems, SGI Altix 1330 and Altix 1350 cluster 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 and Installing 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: Your SGI system must have PROM version 4.04 flashed on the system before installing SLES9 SP1. Without PROM version 4.04, the system will fail to boot after installation. SGI Altix software support contract customers can download PROM version 4.04 or later from the following area on Supportfolio:
https://support.sgi.com/browse_request/suse_patch_browse_results

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 SGI 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.

The **all the nodes** comment in step 4 includes both partitioned and unpartitioned systems. For partitioned systems, this means that it is important to upgrade all nodes (including all nodes in all partitions) before resetting any partitions. In other words, do not reset any partitions until all nodes in all partitions have been upgraded.

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*.

Reinstalling Your Software

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.

Product Support

This chapter documents the product components that are supported on the SGI Altix 330 systems, SGI Altix 350 systems, SGI Altix 1330 and Altix 1350 cluster systems, and SGI Altix 3000 series systems. (For a list of the products, see Table 3-1 on page 26.)

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

- “SGI ProPack for Linux Products” on page 25
- “Operating System Enhancements” on page 29
- “I/O Subsystems” on page 30
- “HPC Application Tools and Support” on page 34
- “System Management” on page 35
- “NUMA Tools” on page 40

SGI ProPack for Linux Products

Software provided by SGI for the SGI ProPack 4 for Linux release consists of kernel modules for SGI software built against the kernels in SUSE LINUX Enterprise Server 9 Service Pack 1 (SLES9 SP1) and value-add software developed by SGI specifically to run on SGI Altix systems.

Table 3-1 SGI ProPack 4 for Linux Products

Product	Description
Application performance measuring tools	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. For more information on VTune, go to the following web location: http://developer.intel.com/software/products/vtune/
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> .
Cpuset System	The Cpuset System is primarily a workload manager tool permitting a system administrator to restrict the number of processors and memory resources that a process or set of processes may use. A system administrator can use cpusets to create a division of CPUs and memory resources within a larger system. For more information, see the "Cpusets on SGI ProPack 4 for Linux" chapter in the <i>Linux Resource Administration Guide</i> .
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.

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

Product	Description
IOC4 serial driver	<p>Driver that supports the Internal IDE CD-ROM, NVRAM, and Real-Time Clock.</p> <p>Serial ports are supported on the IOC4 base I/O chipset and the following device nodes are created:</p> <pre>/dev/ttyIOC4/0 /dev/ttyIOC4/1 /dev/ttyIOC4/2 /dev/ttyIOC4/3</pre>
Kernel partitioning support	<p>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>.</p>
L1 and L2 System Controller firmware	<p>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></p>
MPT	<p>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>.</p>
NUMA tools	<p>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 Application Tuning Guide</i>.</p>
Performance Co-Pilot collector infrastructure	<p>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>.</p>
PROM	<p>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 and Installing the Latest SGI Altix System Firmware" on page 22.</p>
XVM	<p>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>.</p>

SGI does not support the following:

- Base Linux software not released by SUSE/Novell for SLES9 SP1 or other software not released by SGI.
- Other releases, updates, or patches not released by SUSE/Novell for SLES9 SP1 or by SGI for SGI ProPack software.
- Software patches, drivers, or other changes obtained from the Linux community or vendors other than SUSE/Novell and SGI.
- Kernels recompiled or reconfigured to run with parameter settings or other modules as not specified by SUSE/Novell and SGI.
- 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.

Cpuset Support

The cpuset facility is primarily a workload manager tool permitting a system administrator to restrict the number of processors and memory resources that a process or set of processes may use. A cpuset defines a list of CPUs and memory nodes. A process contained in a cpuset may only execute on the CPUs in that cpuset and may only allocate memory on the memory nodes in that cpuset. Essentially, cpusets provide you with a CPU and memory containers or "soft partitions" within which you can run sets of related tasks. Using cpusets on an SGI Altix system improves cache locality and memory access times and can substantially improve an applications performance and runtime repeatability. Restraining all other jobs from using any of the CPUs or memory resources assigned to a critical job minimizes interference from other jobs on the system. For example, Message Passing Interface (MPI) jobs frequently consist of a number of threads that communicate using message passing interfaces. All threads need to be executing at the same time. If a single thread loses a CPU, all threads stop making forward progress and spin at a barrier. Cpusets can eliminate the need for a gang scheduler.

Cpusets are represented in a hierarchical virtual file system. Cpusets can be nested and they have file-like permissions. For detailed information on cpusets, see Chapter 6, "Cpusets on SGI ProPack 4 for Linux" 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 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. This section covers the following topics:

- “Persistent PCI Bus Numbering” on page 31
- “Persistent IP Addressing of Ethernet Interfaces” on page 32
- “XFS Filesystem” on page 33
- “XVM Volume Manager” on page 33

Persistent PCI Bus Numbering

Persistent PCI 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 and for IX-bricks and PX-bricks only, bus numbers are allocated starting from the lowest C-brick module ID to which the I/O brick is connected. For the NUMALink connected PA-bricks it is the lowest module ID of the PA-brick. Since you can have both IX-bricks and PX-bricks on a system along with PA-bricks, it is the lowest of all these combinations.

An I/O brick can be any one of the following:

- IX-brick
The IX-brick is PX-brick with a BaseIO card (IO9 or IO10) in bus1 slot1 of the brick.
- PX-brick
The PX-brick is a XIO connected PCI/PCI-X I/O brick. There are 6 pci/pci-x buses, each with 2 pci/pcix slots.
- PA-brick
The PA-brick, also called the PA expansion module, is a NUMALink connected PCI/PCI-X I/O Brick. There are 4 pci/pci-x buses, two with 2 slots, and 2 with a single slot. Note that the PA-brick is supported only in NUMALink-4 Altix configurations.

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.

On an SGI Altix 3000 series systems, the IX-bricks and PX-brick I/O bricks have six physical buses and a PA-brick I/O brick has four 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 right most digit of a bus number corresponds to the actual stamped number on the I/O brick. An SGI Altix 350 system has 2 buses per module.

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 set the `PciBusList1`, `PciBusList2`, and `PciBusList3` PROM environment variables. Each variable is a comma separated list of the `moduleid` of I/O Bricks representing the order you want the persistent bus numbering.

There are three `PciBusList#` variables to support persistent bus numbering for 16 I/O bricks. (This is because PROM variables are limited to 48 characters. That gives us six `moduleid` per variable to support up to 16 I/O bricks). Some examples are, as follows:

```
POD> setallenv PciBusList1 "001i01,001p04,101p10"
```

or

```
POD> setallenv PciBusList1 "001i01,001p04,101p10,010i10,011p12,012p11"
```

```
POD> setallenv PciBusList2 "102i01"
```

The PROM clearance `PciBusList1` variable is used to clear the persistent bus numbering.

Persistent IP Addressing of Ethernet Interfaces

An Ethernet interface can be given a persistent internet addresses by associating its permanent MAC address, such as 08:00:69:13:f1:aa, with an internet protocol (IP)

address, for example 192.168.20.1. An interface with a persistent IP address will be given the same IP address each time the system is booted.

SGI ProPack 4 is built on top of SLES9. SLES9 has its own implementation of Ethernet persistent naming, that is, the Media Access Control (MAC) address is embedded in the name of the configuration file of the associated Ethernet interface.

For example, if you want to ensure that the Ethernet interface whose MAC address is 08:00:69:13:f1:aa always uses the IP address 192.168.20.1, edit the `/etc/sysconfig/network/ifcfg-eth-id-08:00:69:13:f1:aa` file so that it contains the following:

```
BOOTPROTO="static"  
STARTMODE="onboot"  
IPADDR="192.168.20.1"  
NETMASK="255.255.255.0"  
NETWORK="192.168.20.0"  
BROADCAST="192.168.20.255"
```

For more information on Ethernet device configuration files and instructions for editing them, see `/usr/share/doc/packages/sysconfig/README`.

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*.

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 memory mapping within and between partitions for fast memory-to-memory transfers and the hardware memory controller's fetch operation (fetchop) support. Fetchops and other shared memory techniques enable ultra fast communication and synchronization between MPI processes in a parallel application.

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)

The Extensible Firmware Interface (EFI), a supporting platform to provide input to the CPU and to handle its output, is provided by SLES9 SP1, the base Linux operating system for SGI Altix systems running SGI ProPack 4. EFI also controls the server's boot configuration, maintaining the boot menu in durable, nonvolatile memory.

SLES9 SP1 uses the `elilo-3.4-217.2` package, which places the bootloader (`elilo.efi`) and configuration file (`elilo.conf`) in the `/boot/efi/efi/SuSE/` directory on SGI Altix systems.

Note: When booting from SLES9 CDs, use the `bootia64` command instead of `elilo`. Once the system is running SLES9 SP1, use `elilo` to boot from EFI.

Note: If you have installed multiple kernel images and want to boot with one that is not currently the system default (`vmlinuz` in `/boot/efi/efi/SuSE`), simply copy the `vmlinuz` and `initrd` files for the kernel you wish to use from `/boot` to `/boot/efi/efi/SuSE`.

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

Table 3-2 EFI Commands (continued)

EFI Command	Description
<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
<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

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

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

FLEXlm

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

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

For more information, visit

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

SGIconsole

SGIconsole is a combination of hardware and software that provides console management and allows monitoring of multiple SGI servers running the IRIX operating system and SGI ProPack for Linux. These servers include SGI partitioned systems and

large, single-system-image servers, including SGI Altix 350 systems and the SGI Altix 3000 family of servers and superclusters.

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

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

System Controller Firmware

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 Tools

This section describes the commands that are currently provided with the collection of NUMA related tools that can help you with tuning applications on your system.

Note: Performance tuning information for single processor and multiprocessor programs resides in *Linux Application Tuning Guide*.

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.

dplace Command

The `dplace(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 5, “Data Placement Tools” in the *Linux Application Tuning Guide*.

Index

A

Array Services, description, 26

B

Bus numbering, 31

C

CD contents, 2

Comprehensive System Accounting (CSA) support,
29

Cpuset system, 29

CSA support, 29

D

Disk partition table, 6

dlook command, 40

dplace command, 41

E

EFI commands, 36

EFI support, 35

Extensible Firmware Interface (EFI) support, 35

F

FLEXlm support, 38

H

hardware platforms, 1

HPC support

libraries and tools, 34

Message Passing Toolkit, 34

Performance Co-Pilot (PCP), 35

I

Installation

CDs, 21

overview, 4

system controller software, 21

upgrading, 24

I/O subsystems

for HPC systems, 30

XFS filesystem support, 33

XVM support, 33

L

L1 and L2 controllers, 39

Linux Kernel Crash Dump (LKCD), 8

LKCD for file configuration, 8

M

Message Passing Toolkit
support, 34

N

NUMA tools, 40
dllook command, 40
dplace command, 41

O

Open source website, 4
OS configuration settings, 1
OS enhancements
Comprehensive System Accounting (CSA), 29
cpuset support, 29
for HPC environments, 29
partitioning, 30

P

Password requirement, 5
Performance
measuring, 26
Performance Co-Pilot support, 35
Persistent bus numbering, 31
Persistent naming
bus numbers, 31
Ethernet devices, 32
XSCSI devices, 33
Product list, 25
Product support, 25
PROM support, 35

R

Release features, 1

S

SGIconsole support, 38
Software
installation, 5
introduction, 1
planning, 6
System controller firmware support, 39
System Controller Software
CD, 21
System management
EFI, 35
FLEXlm, 38
PROM chips, 35
SGIconsole, 38
system controller firmware, 39
tools, 35

U

Unsupported elements, 28
Upgrading from previous releases, 21

X

XFS filesystem support, 33
XVM support, 33