

SGI ProPack™ 5 for Linux® Service Pack 2
Start Here

007-4837-003

COPYRIGHT

© 2006, 2007, SGI. 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 SGI.

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

SGI, the SGI logo, Altix, IRIX, OpenGL, Origin, Onyx, Onyx2, and XFS are registered trademarks and NUMAflex, NUMALink, Performance Co-Pilot, RASC, SGI ProPack, SGIconsole, SHMEM, Silicon Graphics Prism, and Supportfolio are trademarks of SGI, 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, Pentium, and Xeon are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Linux is a registered trademark of Linus Torvalds in several countries. MIPS is a registered trademark of MIPS Technology, Inc. Novell is a registered trademark, and SUSE is a trademark of Novell, Inc., in the United States and other countries. QLogic is a registered trademark of QLogic Corporation. PostScript is a registered trademark of Adobe Systems, Inc. Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries. Scali Manage is a trademark of Scali 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	July 2006 Original publication.
002	April 2007 Updated to support the SGI ProPack 5 for Linux Service Pack 1 release.
003	July 2007 Updated to support the SGI ProPack 5 for Linux Service Pack 2 release.

New Features in This Guide

This update of *SGI ProPack 5 for Linux Service Pack 2 Start Here* supports the SGI ProPack 5 for Linux Service Pack 2 release.

Major Documentation Changes

Major changes in this guide for the SGI ProPack 5 for Linux SP2 release include the following:

- Updated information in “System Controller Software 1.17 CD” on page 4.
- Updated information in “Documentation Updated for This Release” on page 6.
- Updated information in “Latest PROM Firmware” on page 14.
- Updated information in “License Manager Updates” on page 15.
- Added information about MVAPICH in “MVAPICH” on page 58.

Contents

Record of Revision	iii
Figures	xi
Tables	xiii
Helpful Online Resources	xvi
Related Publications	xviii
Novell Documentation	xviii
Red Hat Documentation	xviii
Intel Compiler Documentation	xix
Other Intel Documentation	xix
SGI Documentation	xx
Obtaining SGI Publications	xxiv
Conventions	xxiv
Reader Comments	xxiv
1. Release Features	1
Software Introduction	1
CD Contents	2
Software Package License Information	3
SGI ProPack 5 for Linux CDs	3
Release Information, Documentation, and Release Notes	3
System Controller Software 1.17 CD	4
Getting Updates	4
SLES10 Updates	5
RHEL5 Updates	5
SGI Open Source Information	6
Installation Overview	6
Documentation Updated for This Release	6

2.	Software Planning and Installation	9
	Software Planning10
	Table of Disk Partitions11
12
	Systems with Multiple Ethernet Configurations12
	XVM Volumes Access in Rescue Mode12
	Adding the <code>xfs_quota</code> Variable to the <code>initrd</code> Boot Image12
	<code>irqbalance</code> Utility on SGI ProPack 5 SP2 Systems Running RHEL513
	SLES10 SP1 or RHEL5 Kernels13
	Linux Kernel Crash Dump13
	Latest PROM Firmware14
	Getting the Latest Firmware/Software for Your SGI Altix XE System15
	License Manager Updates15
	Read Me First - Installing SLES10 SP1 or RHEL5 on Your SGI System16

Installing SLES10 SP1 and SGI ProPack 5 Service Pack 2	19
Installation Overview	19
Installation Procedure	20
EFI Boot Manager	20
Language Selection	25
License Agreement	25
Installation	26
Installation of Add-on Products (especially SGI ProPack 5 SP2)	28
Clock and Time Zone	29
Partitioning Set Up	29
Software Selection and System Tasks	31
Setting Default Runlevel	32
Performing the Installation	32
Setting Up Your Hostname and Domain Name	33
Setting Up Your root Password	33
Network Configuration	33
Customer Center	35
Online Update	36
Service	36
Users	36
Release Notes	36
Hardware	37
SLES and SGI ProPack Network Installation for Experts	37
Setting up the Network Install Server	37
Using AutoYaST	38
Upgrading From an Earlier SGI ProPack 4 Release to SGI ProPack 5 SP2	39
Upgrading for SGI ProPack 5 SP1 to SGI Propack 5 SP2	39
Installing RHEL 5 and SGI ProPack 5 SP2	42
Installing the System Controller Software 1.17 CD	43
Obtaining and Installing the Latest SGI Altix System Firmware	44
Reinstalling Your Software	47
3. Product Support	49
SGI ProPack 5 for Linux Service Pack 2 Products	49

Operating System Enhancements.54
Cpuset Support54
Comprehensive System Accounting (CSA).55
Partitioning.55
I/O Subsystems.56
Persistent IP Addressing of Ethernet Interfaces56
PCI Domain Support for SGI Altix 450 and 4700 Systems56
XSCSI Naming Systems on SGI ProPack Systems56
XFS Filesystem57
XVM Volume Manager57
HPC Application Tools and Support.57
Message Passing Toolkit58
DAPL for Altix.58
MVAPICH58
Performance Co-Pilot60
System Management60
Embedded Support Partner61
PROM Chips61
Reconfigurable Application-Specific Computing61
Extensible Firmware Interface (EFI)62
SGIconsole65
System Controller Firmware66
L1 and L2 System Controllers66
L3 System Controller67
NUMA Data Placement Tools.67
dlook Command68
dplace Command68
Index.69

Figures

Figure 2-1	EFI Boot Manager Menu	21
Figure 2-2	Boot Screen	22
Figure 2-3	Network Configuration Screen	23
Figure 2-4	Starting the VNC Server	24
Figure 2-5	VNC Server via a Web Browser	25
Figure 2-6	Installation Mode Window	27
Figure 2-7	Import Untrusted GnuPG Key Screen	28
Figure 2-8	Server Connect Error Message	34
Figure 2-9	Server and Directory Window	35

Tables

Table 2-1	Disk Partitions for SGI ProPack 5 SP2 with SLES10 SP1	11
Table 2-2	Disk Partitions for SGI ProPack 5 SP2 with RHEL5	11
Table 2-3	Disk Partitions for x86_x64 Systems	11
Table 3-1	SGI ProPack 5 for Linux Service Pack 2 Products	50
Table 3-2	EFI Commands	63

About This Guide

This guide provides information about the SGI ProPack 5 for Linux Service Pack 2 (SP2) 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 5 SP2 release notes, which contain the latest information about software and documentation in this release, the list of RPMs distributed with SGI ProPack 5 SP2, and a useful migration guide, which contains helpful hints and advice for customers moving from earlier versions of SGI ProPack to SGI ProPack 5 SP2, can be found in the `/docs` directory on the on the SGI ProPack 5 for Linux Service Pack 2 CD.

This Preface covers the following topics:

- “Helpful Online Resources” on page xvi
- “Novell Documentation” on page xviii
- “Red Hat Documentation” on page xviii
- “Intel Compiler Documentation” on page xix
- “Other Intel Documentation” on page xix
- “SGI Documentation” on page xx
- “Obtaining SGI Publications” on page xxiv

Helpful Online Resources

This section describes the electronic support services available for your SGI system running SGI ProPack 5 for Linux SP2, as follows:

- Supportfolio

Supportfolio, the SGI support web site, including the SGI Knowledgebase, SGI YaST Update Online (YOU) server, SUSE YaST Online Server, links for software supports and updates, as described in this listing.

<https://support.sgi.com/login>

- Supportfolio SGI Knowledgebase

The SGI Knowledgebase available from Supportfolio has answers to common questions, problem resolutions, how-to documents, and many other useful articles for system administrators.

<https://support.sgi.com/kb/publisher.asp>

- SGI YaST Online Update Server

The SGI YaST Update Online (YOU) server is available for your convenience. Get updates to SGI ProPack 5 for Linux and SUSE Linux Enterprise Server 10 (SLES10) using the YaST install tool right from your system. The *YaST Online Update - Patch Media Instructions* describes how to update your system using patch media. Use your Supportfolio account name and password, and point YaST to this location:

<https://update.sgi.com/>

- Supportfolio Software Updates

Check for new updates to your purchased SGI software, and download or order CDs all in one convenient place. Supportfolio Software Updates provides one-stop shopping to check for and get new updates for your SGI software.

<https://support.sgi.com/software?cmd=order>

- SGI Customer Support Center

You can open calls to the SGI Customer Support Center on-line, keep track of all your cases, and receive automatic update information when your case is updated via the Supportfolio Problem Resolution Tools at:

<https://support.sgi.com/caselist>

- SGI Software Licensing

If you need to get your license for SGI software, you can order your software licenses conveniently via e-mail or on-line.

You can send a blank email to license@sgi.com, and you will receive a simple template to fill out. You will quickly receive your new license back with complete instructions for installing it back via E-Mail.

Or, you can go on-line to the SGI Key-O-Matic at:

<http://www.sgi.com/support/licensing/>

Related Publications

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

- “Novell Documentation” on page xviii
- “Red Hat Documentation” on page xviii
- “Intel Compiler Documentation” on page xix
- “Other Intel Documentation” on page xix
- “SGI Documentation” on page xx

Novell Documentation

Documentation for SUSE Linux Enterprise Server 10 is located on your system in the `/usr/share/doc` directory. There are also several helpful documents provided on SLES10 CD1 in the `/docu/en` directory.

The following documentation is available:

Note: Novell makes SUSE Linux Enterprise Server documentation available online at: <http://www.novell.com/documentation/>

- SUSE Linux Enterprise Server 10 Installation and Administration (manual.pdf)
- SUSE Linux Enterprise Server 10 Start-Up Guide (startup.pdf)
- SUSE Linux Enterprise Server 10 Package Descriptions available online at <http://www.novell.com/products/linuxenterpriseserver/packages.html>
- Novell YES CERTIFIED Bulletins for SGI available online at <http://developer.novell.com/yesearch/Search.jsp> (select "SGI" as the Company and "SUSE Linux Enterprise Server 10" as the Novell Product and click Search)

Red Hat Documentation

Documentation for RHEL5 is located in the `/usr/share/doc` directory. It contains package-specific documentation. If you install the `zisosfs` module, additional

documentation in the form of READMEs can be found in
`/usr/share/doc/zisofs-tools-1.x.x`.

On the RHEL5 DVD, `README-{lang}` and `RELEASE-NOTES-{lang}` files are found in the root directory.

Note: Red Hat makes Red Hat Enterprise Linux 5 (RHEL5) documentation available online at:

<http://www.redhat.com/docs/manuals/enterprise/RHEL-5-manual/index.html>

The following documentation is available:

- *Release Notes for x86_64 architecture*
- *Release Notes ia64 architecture*
- *Red Hat Enterprise Linux Deployment Guide*
- *Red Hat Enterprise Linux Installation Guide*
- *Red Hat Virtualization*

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).
- Intel Xeon processor documentation is available at <http://www.intel.com/design/Xeon/documentation.htm>

SGI Documentation

The following SGI documentation is available:

- *SGI ProPack 4 to SGI ProPack 5 Migration Guide*
This guide contains information about migrating from SGI ProPack 4 for Linux operating system to SGI ProPack 5 for Linux operating system..
- *SGI ProPack 5 for Linux Service Service Pack 1 Release Notes*
Provides the latest information about software and documentation in this release. The release notes are on the SGI ProPack for Linux CD in the `root` directory, in a file named `README.TXT`.
- *SGI ProPack 4 for Linux Device Driver Porting Guide*
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.
- *REACT Real-Time for Linux Programmer's Guide*
Describes the facilities of REACT real-time for Linux that supports real-time programs.
- *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 Altix servers.
- *Linux Resource Administration Guide*
Provides a reference for people who manage the operation of SGI Altix 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.
- *Guide to Administration, Programming Environments, and Tools Available on SGI Altix XE Systems*
Explains how to perform general system configuration and operations and describes programming environments and tools available for SGI Altix XE systems.
- *SGI Altix ICE System Administrator's Guide*
Describes how to perform general system installation, configuration, and operations on the SGI Altix ICE 8000 series systems.
- *SGI Altix ICE 8000 System User's Guide*
Describes the features of the SGI Altix ICE 8000 series system, as well as, troubleshooting, upgrading, and repairing.
- *SGI Altix XE210 System User's Guide*
Describes the features of the Altix XE210 server system as well as troubleshooting, upgrading, and repairing.
- *SGI Altix XE240 System User's Guide*
Describes the features of the Altix XE240 server system as well as troubleshooting, upgrading, and repairing.
- *SGI Altix 330 System User's Guide*
Provides an overview of the Altix 330 system components, and it describes how to set up and operate this system
- *Important Note for SGI Altix XE1200 Cluster Systems*
Describes the Silicon Graphics (SGI) and third party customer manuals that may be used with your Altix XE1200 cluster 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.
- *SGI Altix 450 System User's Guide*
Provides an overview of the Altix 450 system components, and it describes how to set up and operate this system.
- *SGI Altix 4700 User's Guide*
This guide provides an overview of the architecture and descriptions of the major components that compose the SGI Altix 4700 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.
- *Reconfigurable Application-Specific Computing User's Guide*
Provides information about the SGI reconfigurable application-specific software computing (RASC) program that delivers scalable, configurable computing elements for the SGI Altix family of servers and superclusters.
- *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 SGI Altix 350 systems, the SGI Altix 3000 and SGI Altix 4000 family of servers and superclusters.
- *XFS for Linux Administration*

Describes XFS, an open-source, fast recovery, journaling filesystem that provides direct I/O support, space preallocation, access control lists, quotas, and other commercial file system features.

- *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.

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
- 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
1140 East Arques Avenue
Sunnyvale, CA 94085-4602

SGI values your comments and will respond to them promptly.

Release Features

This chapter provides an introduction to SGI ProPack 5 for Linux Service Pack 2 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 1.17 CD” on page 4
- “Getting Updates” on page 4
- “SGI Open Source Information” on page 6
- “Installation Overview” on page 6
- “Documentation Updated for This Release” on page 6

Software Introduction

The SGI ProPack 5 for Linux Service Pack 2 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 Intel processors. This product adds to or enhances features in the base SUSE Linux Enterprise Server 10 (SLES10) Service Pack 1 (SP1) Linux distribution or Red Hat Enterprise Linux 5 (RHEL5) distribution. SGI ProPack 5 SP2 is designed to run on any SGI Altix, SGI Altix XE, or SGI Altix ICE 8200 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/>

CD Contents

This section describes the contents of the SGI ProPack 5 for Linux Service Pack 2 CD, as follows:

- 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 (*bcfree*, *dplace*, *dlook* utilities)
- Performance Co-Pilot, system-level performance monitoring and performance management tool
- Open Source SGI Profile Tools
- SGI ProPack configuration recommendations
- Flexible File I/O (FFIO) system
- Cpuset Facility workload manager tool
- Embedded Support Partner (ESP) software suite to monitor events, set up pro-active notification, and generate reports
- Performance Co-Pilot visualization tools
- Intel Compiler Runtime Libraries
- Array Services
- Message-Passing Toolkit (MPT)
- XVM Command line tools and graphical user interface (GUI)
- SGI Itanium 2 PROM (System controller software)
- SGI ProPack documentation
- REACT Real-time for Linux Libraries and Documentation
- libarena: Custom Memory Allocator Interface

Note: Not all SGI ProPack 5 SP2 features are available on x86_64 systems. For the list of supported products on the x86_64 architecture, see “SGI ProPack 5 for Linux Service Pack 2 Products” on page 49.

The Scientific Computing Software Library (SCSL) software is no longer delivered with the SGI ProPack release. SCSL continues to be available via web release in courtesy mode. Login to Supportfolio Online at <http://support.sgi.com>. Click on the Downloads link under Supportfolio areas on the left. Enter **scs1** in the **Enter a Product type or keywords** search box and then select **Linux OS** in the **Select the download type** list.

Software Package License Information

SGI ProPack contains software packages released under a variety of licenses. The licensing terms for all proprietary software packages allow users to install the software only on SGI systems. See the `PACKAGE_LICENSES.txt` file that is available in the `/docs` directory of the SGI ProPack 5 for Linux CD for more details.

SGI ProPack 5 for Linux CDs

The following CDs ship with each SGI Propack 5 for Linux Service Pack 2 release:

- SGI ProPack 5 for Linux Service Pack 2 for Intel Itanium CD
- SGI ProPack 5 for Linux Service Pack 2 for AMD64/EM64T CD

Release Information, Documentation, and Release Notes

For a complete list of RPMs included in the SGI ProPack 5 SP2 release, see the file called `RPMs.txt` that is available in the `/docs` directory on the SGI ProPack 5 for Linux Service Pack 2 CD.

For the latest information about software and documentation in this release, see the release notes that are in a file named `README.TXT` that is available in `/docs` directory on the SGI ProPack 5 for Linux Service Pack 2 CD.

The SGI ProPack 5 for Linux release notes get installed to the following location on a system running SGI ProPack 5 SP2:

```
/usr/share/doc/packages/sgi-propack-5/README.txt
```

System Controller Software 1.17 CD

The System Controller software provides a system control platform for field service and system maintenance applications.

Note: The System Controller Software 1.17 CD is **not** shipped by default. You can download the System Controller Software 1.17 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.

For installation instructions, see “Installing the System Controller Software 1.17 CD” on page 43.

Note: The SGI System Controller software is licensed for use on the SGI Altix 3000 and SGI Altix 4000 series systems. For license information, see “Software Package License Information” on page 3. The System Controller software does not apply to SGI Altix XE systems.

Getting Updates

This section describes how to get software updates for SGI ProPack 5 for Linux SP2 system.

Detailed instructions for downloading and installing SGI ProPack 5 for Linux SP2 updates are available from the SGI ProPack 5 SP2 product pages on Supportfolio: https://support.sgi.com/content_request/961100/index.html.

Note: You must have a Supportfolio account to access this page and the SGI update server.

SLES10 Updates

SGI SLES10 SP1 customers are encouraged to use the YaST Online Update tool to keep the software on their systems updated with the latest maintenance fixes.

Software updates for SLES10 SP1 are provided by Novell. In order to configure your system to enable online updates, you must register with Novell. For more help with this, you can use the **Novell Customer Center Configuration** tool available in the Software section of the YaST Control Center.

Customers that received SGI ProPack 5 SP2 pre-installed on their new SGI Altix system may run the following command to add the proper update source for SGI ProPack 5 SP2:

```
$ /usr/sbin/sgi-propack-update-source
```

This script will prompt you for your Supportfolio login and then add the proper SGI update source for the given architecture.

For additional information on how to configure YaST Online Update to download updates for SGI ProPack 5 SP2, login to Supportfolio and then point your browser to http://support.sgi.com/content_request/961100/index.html.

RHEL5 Updates

Software updates for RHEL5 are provided by Red Hat. As a Red Hat Enterprise Linux customer, you receive access to the Red Hat Network. This allows you to download updates for your system. For more information on the Red Hat Network and getting updates, see *Red Hat Network Overview* and the **Update** link at <https://www.redhat.com/rhn/>.

You can also use the `yum` software package manager to get updates, as follows:

```
% yum update
```

The `yum` software package manager checks for updates and then updates your system appropriately. You must be registered with RHN, for `yum` to get RHN RHEL5 updates.

You can also use the Package Update (`pup`), which is a graphical tool for locating and installing updates, as follows:

```
% pup
```

You can configure the `updatesd(8)` notifier daemon to notify you of updates and potentially install the updates for you. Configuration is done via the `yum-updatesd.conf(5)` file.

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 5 for Linux SP2 will most likely come pre-installed 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 37.

For information on installing RHEL5 and SGI ProPack 5 SP2, see “Installing RHEL 5 and SGI ProPack 5 SP2” on page 42.

Documentation Updated for This Release

The following manuals are new or have been updated for this release:

- *Linux Application Tuning Guide*, 007-4639-007
 - Added information about using `dplace` command with `fork`, `exec`, and `pthread_create` calls in “`dplace` Command” in Chapter 5.
 - Added information about “Linux Shared Memory Accounting” in Chapter 7.
- *REACT Real-Time for Linux Programmer’s Guide*, 007-4746-006

Contains information about the following:

- The ability to configure a group that has permission to perform REACT work without having root access. See "Granting Users Access to Select REACT Features" in Chapter 9.
- "Linux Capabilities and the SGI ProPack Captools Package" in Chapter 9

The following now apply to Altix XE systems:

- "Shielding a CPU from Timer Interrupts" in Chapter 4
- Chapter 5, "Using the Frame Scheduler"
- Chapter 8, "User-Level Interrupts"
- Appendix A, "Example Application"

The following changes are documented in Chapter 10, "SGI Linux Trace":

- The ability to start and stop the `tracedaemon` program from inside the `racevisualizer` program by selecting the stoplight toolbar icon
- The ability to filter events based on CPU
- A new `slt-cpu.all` file that combines from all of the individual `slt-cpu.N` files when `tracedaemon` is run from inside `tracevisualizer`
- The ability to record and display frame scheduler events
- The event graph now displays the end time and the format of the time ruler in either microseconds (us) or nanoseconds (ns)
- *Message Passing Toolkit (MPT) User's Guide, 007-3773-004*
 - Added a new chapter about MPI performance profiling in Chapter 7, "MPI Performance Profiling".
- *SGI Altix ICE System Administrator's Guide*
 - New for this release, describes how to use SGI Tempo systems management software to install, provision, configure, and operate your system.

Software Planning and Installation

This chapter describes the planning that you need to do to get started and provides instructions for installing SGI ProPack 5 for Linux Service Pack 2.

Your SGI Altix system comes with a base Linux distribution of SUSE Linux Enterprise Server 10 (SLES10) SP1 or Red Hat Enterprise Linux 5 (RHEL5) and the SGI ProPack 5 for Linux software preinstalled. This chapter describes how to install the software from the CD if it should ever become necessary to reinstall it.

Note: For information on how to install, configure, and operate software on an SGI Altix ICE 8200 system, see the *SGI Altix ICE System Administrator's Guide*.

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 5 for Linux SP2 software works only with the SLES10 SP1 environment or RHEL5 environment. Earlier versions of these distributions, or any other distributions, are not compatible with SGI ProPack 5 for Linux software. Note that if you are currently running SGI ProPack 4 for Linux with SLES9 SP1 (or later), you cannot update only the SLES9 base operating system to SLES10 SP1. You must also update the SGI ProPack 4 software with SGI ProPack 5 SP2 in order to continue using the SGI ProPack features.

Before you install or configure your system, please read Chapter 1, “Release Features” so that you understand the features of SGI ProPack 5 for Linux SP2 software. You might also want to familiarize yourself with some or all of the documentation listed in “Related Publications” on page xviii.

This chapter covers the following topics:

- “Software Planning” on page 10
- “License Manager Updates” on page 15

- “Read Me First - Installing SLES10 SP1 or RHEL5 on Your SGI System” on page 16
- “Installing SLES10 SP1 and SGI ProPack 5 Service Pack 2” on page 19
- “SLES and SGI ProPack Network Installation for Experts” on page 37
- “Upgrading From an Earlier SGI ProPack 4 Release to SGI ProPack 5 SP2” on page 39
- “Installing the System Controller Software 1.17 CD” on page 43

Software Planning

This section provides disk and file information you need to have about the software that has been shipped to you and covers the following topics:

- “Table of Disk Partitions” on page 11
- “Systems with Multiple Ethernet Configurations” on page 12
- “XVM Volumes Access in Rescue Mode” on page 12
- “Adding the xfs_quota Variable to the initrd Boot Image” on page 12
- “irqbalance Utility on SGI ProPack 5 SP2 Systems Running RHEL5” on page 13
- “SLES10 SP1 or RHEL5 Kernels” on page 13
- “Linux Kernel Crash Dump” on page 13
- “Latest PROM Firmware” on page 14
- “Getting the Latest Firmware/Software for Your SGI Altix XE System” on page 15

Table of Disk Partitions

SGI ships Altix systems with Itanium processors from the factory running SGI Propack 5 SP2 with SLES 10 SP1 with the partition layout as shown in Table 2-1

Table 2-1 Disk Partitions for SGI ProPack 5 SP2 with SLES10 SP1

Partition	Type	Mount Point	Size
part1	VFAT	/boot/efi	500 MB
part2	SWAP	/swap	10 GB
part3	XFS	/	25 GB
part4	XFS	/data1	remainder of disk

SGI ships Altix systems with Itanium processors or Altix XE systems with Xeon processors from the factory running SGI Propack 5 SP2 with RHEL5 with the partition layout as shown in Table 2-2.

Table 2-2 Disk Partitions for SGI ProPack 5 SP2 with RHEL5

Partition	Type	Mount Point	Size
/dev/sda1	ext3	/boot	100 MB
/dev/sda2	LVM	LVM	35 GB
/dev/sda3	ext3	/data1	remainder of disk

SGI ships Altix XE systems with Xeon processors from the factory with the partition layout shown in Table 2-3.

Table 2-3 Disk Partitions for x86_x64 Systems

Partition	Type	Mount Point	Size
/dev/sda1	SWAP	swap	8 GB
/dev/sda2	XFS	/	25 GB
/dev/sda3	XFS	/data1	remainder of disk

Systems with Multiple Ethernet Configurations

On SGI Altix systems running SGI ProPack 5 SP2 with multiple Ethernet connections, there is no guarantee that the 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 5 for Linux SP2 release notes for a work around to this problem.

Note: This section does not apply to SGI Altix XE systems.

XVM Volumes Access in Rescue Mode

You cannot access XVM Volumes in Rescue Mode on SGI ProPack 5 SP2 systems. It is possible to boot into rescue mode from the SLES10 SP1 DVD, but that DVD does not contain the `xvm` kernel module. This makes it impossible to perform maintenance on XVM volumes from rescue mode at this time. SGI recommends that you do not migrate any system directories to XVM since there is no way to access them from rescue mode.

Note: This section does not apply to systems running SGI ProPack 5 SP2 with RHEL5. For this release, RHEL5 does not support XFS or XVM.

Adding the `xfs_quota` Variable to the `initrd` Boot Image

You need to add `xfs_quota` to the `INITRD_MODULES` variable in the `/etc/sysconfig/kernel` file when `rootflags=quota` is set or the `initrd` boot loader will panic your system at boot time.

Note: This section does not apply to systems running SGI ProPack 5 SP2 with RHEL5. For this release, RHEL5 does not support XFS.

`irqbalance` Utility on SGI ProPack 5 SP2 Systems Running RHEL5

On SGI ProPack 5 SP2 systems running RHEL5, the `irqbalance` utility is configured off. This utility distributes device interrupts across all the CPUs on a multiprocessor system with the purpose of spreading the load. On a distributed shared memory (DSM) multiprocessor computer system, such as an SGI Altix system, the distance (that is, the latency) to a given CPU from a given device, varies quite substantially as you add router hops. Because of this, the PROM on an SGI Altix system, does not pick a random CPU to process the interrupts of a device; it picks the optimum CPU to handle the interrupts. This `irqbalance` utility breaks this optimum CPU assignment. Therefore, it is configured off.

SLES10 SP1 or RHEL5 Kernels

For ia64 systems, SLES10 SP1 contains two kernel variants, `kernel-default` and `kernel-debug`. During installation, if you do not explicitly select to install `kernel-debug`, your SGI Altix system will boot with the `kernel-default` variant after installation. For `x86_64` systems, SLES10 SP1 contains `kernel-default` (1CPU), `kernel-debug`, `kernel-smp` (multi-CPU) and `kernel-xen`. SGI recommends and pre-installs `kernel-smp` for `x86_64` systems.

RHEL5 has two kernel images, `kernel` and `kernel-xen` for ia64 and `x86_64` systems. The `kernel-xen` variant is only provided as a technology preview for ia64 systems, but is a fully supported kernel for `x86_64` systems.

Linux Kernel Crash Dump

In the event of a system crash on systems running SGI ProPack 5 SP2 running SLES10 SP1, the Linux Kernel Crash Dump (LKCD) facility will save a crash dump that will appear in the `/var/log/dump` directory after your system reboots. SGI support personnel can analyze this file to determine the cause of the system crash.

If you wish to save crash dumps to a different location, edit the `DUMPDIR` variable in `/etc/sysconfig/dump` and then run the following command: `lkcd config`.

Note: Depending on which kernel you have installed, the system may enter the KDB kernel debugger after a crash. If a `kdb>` prompt appears on the console, type `sr d` to initiate a crash dump.

The KDB kernel debugger is for ia64 systems only. For x86_64 systems, use kexec based kdump.

In the event of a system crash on systems running SGI ProPack 5 SP2 running RHEL5, the system uses the kexec/kdump crash dump utility. The kexec/kdump utility creates a Linux kernel crash dump at `/var/crash/timestamp/vmcore` if kexec/kdump is installed, enabled, and configured. After SGI ProPack 5 SP2 is installed, your system is set up to take dumps when needed. In order to analyze a crash dump, you need the `debuginfo` packages installed.

Latest PROM Firmware

Note: This section only applies to ia64 systems. For x86_64 systems, see “Getting the Latest Firmware/Software for Your SGI Altix XE System” on page 15.

Unless you have the most current PROM, the kernel may fail to boot during the installation procedure. The prom version is printed when the system resets in the prom banner.

The prom RPM is on the SGI ProPack 5 for Linux Service Pack 2 CD in the `/sgi/ia64` directory.

Note: SGI Altix 450 or SGI Altix 4700 systems running SGI ProPack 5 for Linux Service Pack 2 should use the PROM located in `shub2snprom.bin`. All other SGI systems should use the PROM located in `shub1snprom.bin`. This note does **not** apply to SGI Altix XE systems (instead see “Getting the Latest Firmware/Software for Your SGI Altix XE System” on page 15). Software support contract customers can download the most recent PROM versions from SGI Supportfolio: <https://support.sgi.com/>

You can determine the PROM level installed on your machine, as follows:


```
% cat /proc/sgi_prominfo/node*/version
4.74
SGI SAL Version 4.74 rel070313 IP41 built 10:14:15 AM Mar 13, 2007
```

On an SGI Altix 4700 system, the output from the version file looks similar to the following:

```
% cat /proc/sgi_prominfo/node*/version
SGI SAL Version 1.21 rel070201 IP65 built 08:38:57 AM Feb 1, 2007
```

The PROM RPM on the SGI ProPack 5 for Linux Service Pack 2 CD contains all of the appropriate flash images for SGI systems running SGI ProPack 5. The flash utility chooses the appropriate image for your machine from this RPM.

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

Getting the Latest Firmware/Software for Your SGI Altix XE System

For information on the basic input/output system (BIOS), Baseboard Management Controller (BMC) firmware, Intelligent Platform Management Interface (IPMI) control utility, Field Replaceable Unit/Sensor Data Record (FRUSDR) software, Hierarchical Storage Controller (HSC) software, LSI RAID firmware, Scali Manage system management software, Infiniband drivers, INTEL MPI Runtime Environment Kit for Linux, and Java Runtime Environment (JRE), see the **SGI Altix XE Firmware/Software** page on SGI Supportfolio at:
https://support.sgi.com/content_request/691346/index.html

License Manager Updates

Please backup your license key files prior to installing SLES10 SP1 and/or SGI ProPack 5 SP2. 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 software licensing support, the Redhat and SUSE distributors do not ship any code. ISVs are responsible for obtaining and including software licensing components into their products.

With SGI ProPack 5 for Linux Service Pack 2, SGI products have migrated to a new software licensing mechanism called LK. LK was developed by SGI for SGI products only. For more information, see the specific product release notes or the FAQ, as follows:

<http://www.sgi.com/support/licensing/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 SLES10 SP1 or RHEL5 on Your SGI System

Note: For information on how to install, configure, and operate software on an SGI Altix ICE 8200 system, see the *SGI Altix ICE System Administrator's Guide*.

This section provides some information about installing SLES10 SP1 on SGI Altix systems. You need to have SLES10 SP1 or RHEL5 installed before you can install SGI ProPack 5 for Linux Service Pack 2.

Note: If you are installing Red Hat Enterprise Linux 5 (RHEL5), you can proceed to "Installing RHEL 5 and SGI ProPack 5 SP2" on page 42.

For installation instructions, see "Installing SLES10 SP1 and SGI ProPack 5 Service Pack 2" on page 19.

Note: If you are currently running SLES10 GA, SGI recommends that you do an online update of SLES10 to get it to the latest level and then install SGI ProPack 5 SP2. For installation instructions, see "Installing SLES10 SP1 and SGI ProPack 5 Service Pack 2" on page 19

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

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

- **Find installation documentation.**

You can find installation documentation on the SLES10 SP1 DVD in the `docu/en` directory. It contains these `.pdf` documentation files:

```
apparmor-admin.pdf  preparation.pdf  sles-admin.pdf  startup.pdf
```

- **Find and read the SLES10 SP1 release notes.**

It is important to read the SLES10 SP1 release notes as well. You will find them in a text file named `RELEASE-NOTES.en.txt` in the `root` directory of the SLES10 SP1 DVD. The YaST installation tool can also get the latest version of the release notes directly from Novell. On the installed system, the latest version of the release notes are at the following location:

```
/usr/share/doc/release-notes/SUSE_Linux_Enterprise_Server_10/RELEASE-NOTES.en.html
```

- **How to erase everything and start with a clean root drive.**

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 the *SLES10 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.

- **Find the best installation method for SLES10 SP1 for your site.**

There are several different methods to install SLES10 SP1. To use the graphical YaST installation, SGI recommends using the `vnc` or `ssh` methods, assuming the SGI system is connected to a network.

Insert the media in the DVD drive and reset the machine. The BIOS on `x86_64` systems should already be configured to attempt booting from CD first. The SLES10 SP1 media should boot when the system comes back up.

To use this mode on `ia64` systems, SGI suggests the **EFI Shell [Built-in]** from the **Boot Maintenance** menu. A list of `fs#` devices will be presented before you get to the `Shell>` prompt. One from the list, will have the word `CDROM` in the display. If the `fs#` associated with the `CDROM` is 1, enter the following commands:

```
Shell> fs1:  
fs1:\> cd efi\boot
```

You are now ready to launch the installer. Issue the following command:

```
fs1:\efi\boot> bootia64
```

You now will see a list of choices including **Installation**, **Installation - Safe Settings**, and so on. Keep the default **Installation** choice highlighted. In the **boot:** prompt, your selection depends on the type of system you are booting.

If you prefer to install using `vnc` or `ssh`, enter **vnc=1** or **usessh=1** in the **Boot:** line.

If you are installing software on an Altix system you need to specify information about the serial console on the boot line. If you wish to use `vnc` or `ssh`, you can supply those options as well. An example boot line specifying the serial console for installation and the use of `vnc` is similar to the following:

For ia64:

```
boot: console=ttySG0,38400n8 vnc=1
```

For x86_64:

If you are using a VGA monitor/keyboard to install, you do not need to supply any options.

If you are using a serial console, the setting depends on how the system is configured. An example is, as follows:

```
boot: console=ttyS1,38400n8 vnc=1
```

Note: If `ttyS1` does not work, try `ttyS0`. The serial cable might be attached to the other serial port

For ia64 systems, if you choose to install over the serial console or an L2 controller with `xterm`, note that the serial console may not be 100% clean and that you may run into "screen corruption". Often this can be cleared up by entering `Ctrl-L` from the command line, which re-draws 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 -e "cu -s 38400 -l /dev/ttySG0"
```

- **Find out how to perform a network installation.**

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 5 with SUSE operating system repositories.

- **Find information on AutoYaST.**

For information on AutoYaST, see “Using AutoYaST” on page 38.

Installing SLES10 SP1 and SGI ProPack 5 Service Pack 2

To install SUSE Linux Enterprise Server 10 Service Pack 1 (SLES10 SP1) from the DVD media, perform the procedures described in this section.

Note: Before starting a fresh install of SLES 10 SP1 and SGI ProPack 5 SP2 over the existing operating system, it is a good idea to remove the production root drive from your system. Perform the fresh install over the operating system on your backup root drive. This provides a fall back route to your previous operating system, in case of application compatibility issues.

Installation Overview

A general overview of the procedure steps are, as follows:

- Power down your system and replace the production root drive with the backup drive.
- Insert the SLES 10 SP1 DVD in your DVD-ROM, and boot your system.
- Start the installation, using **vnc=1** option to start a VNC server and be sure to set the **console=** setting appropriately to use the serial console (see “Read Me First - Installing SLES10 SP1 or RHEL5 on Your SGI System” on page 16).
- On an SGI Altix XE system, make sure that BIOS is configured to boot from the DVD before the hard drive.
- Using a VNC client or compatible web browser that can run the VNC client java applet, connect to YaST remotely.
- Configure the partition layout of the drive, by changing the partition type of the root partition from **Reiser** to **XFS**.
- Modify the default software selections, by adding the **Apache Server** to the installation.

Installation Procedure

Note: Read Part 1 “Deployment” in *SUSE Linux Enterprise Server Installation and Administration* manual that resides in `sles-admin.pdf` file in the `docu` directory of the SLES10 SP1 DVD before beginning this installation procedure.

To install SUSE Linux Enterprise Server 10 (SLES10) SP1 from the DVD media, perform the following steps:

Note: These steps correspond to Section 3.2.2 “Selecting the Source of the Installation Data” of *SUSE Linux Enterprise Server Installation and Administration* manual. In this example, the source of installation will be “**Installing from SUSE Linux Enterprise Media**”. (DVD)

1. Insert the SLES 10 SP1 Installation DVD in the DVD drive and reboot your system.

EFI Boot Manager

2. When the system returns to the **EFI Boot Manager Menu**, boot the system on the SLES10 SP1 Install DVD via the Boot CDROM/DVD boot option as shown in Figure 2-1.

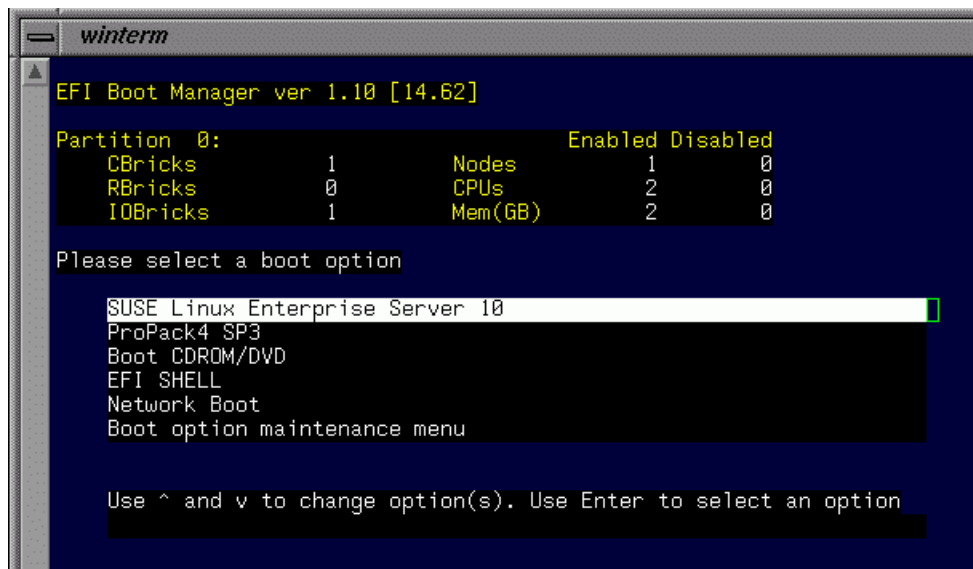


Figure 2-1 EFI Boot Manager Menu

3. If there is not a **Boot CDR0M/DVD** option, it can be added to the boot menu as follows:
 - a. Select **Boot option maintenance menu**
 - b. Select **Add a boot option**
 - c. Highlight the **Removable media boot** line
 - d. Enter **Boot CDR0M/DVD** for the name
 - e. For **New Boot Option Data**, enter **n**
 - f. Save to nvram
 - g. Choose **Exit** on the next two screens to return to the main menu
 - h. Select your new **Boot CDR0M/DVD** option to boot from the CD
4. From the boot screen, select **Installation** and add the appropriate **Boot Options** to start the installation, as shown in Figure 2-2.

Note: This section corresponds to the Section 3.3, “The Boot Screen” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

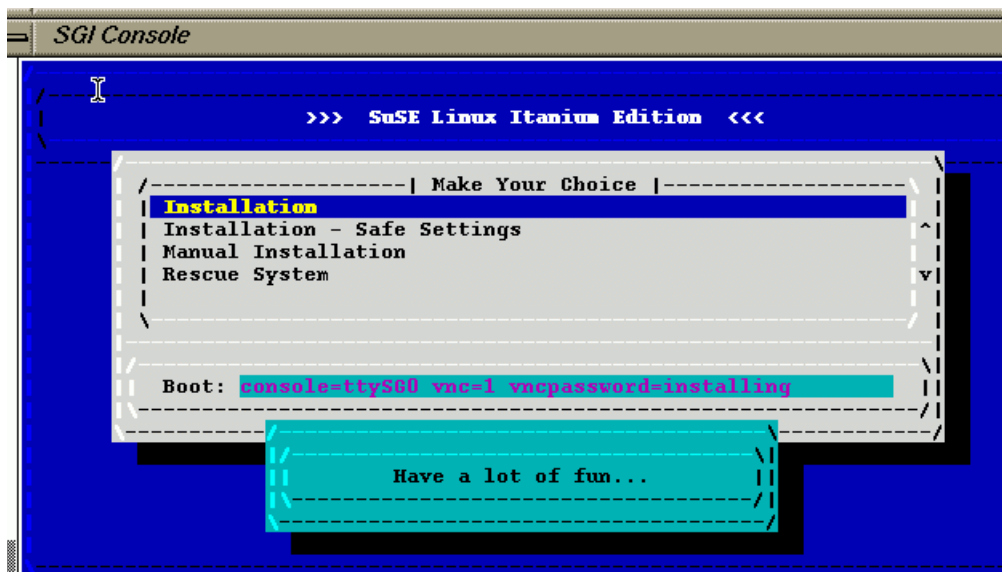


Figure 2-2 Boot Screen

5. These Boot Options will start the installation via the serial console and finish by using YaST on a remote system using VNC:
 - **console=ttySG0,38400** - Select the serial console as the output source for ia64 systems. Sets the baud rate to 38400.
 - **console=ttyS1,38400** - Select the serial console as the output source for x86_64 systems. Sets the baud rate to 38400. Set ttyS1 or ttyS0 depending on the port to which the serial cable is connected.
 - **vnc=1** - **Start a VNC Server**, to allow using a web browser, or VNC client program.
 - **vncpassword=your_password** - The password the VNC server will use to authenticate the user. Must be at least 6 characters.

There is no **console=** selection for VGA screen installs.

Note: For more examples of Boot Options, see Section 4.4.3, “Using Custom Boot Options” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

Once the proper Boot Options are added, return starts the install program, and brings up the **Network Configuration Screen** as shown in Figure 2-3.

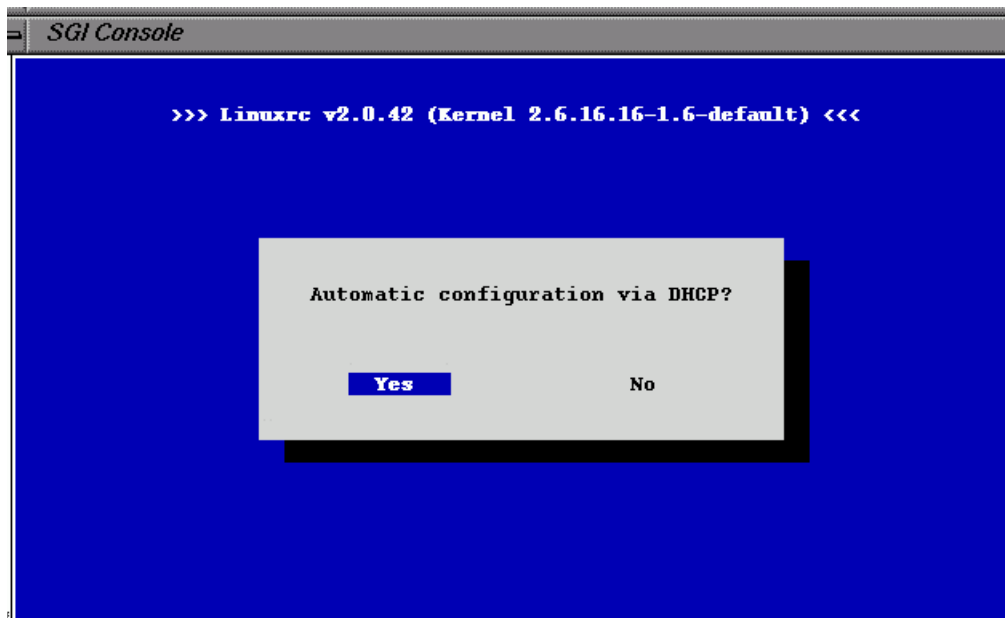


Figure 2-3 Network Configuration Screen

6. Choose DHCP by clicking on **Yes** or enter your information manually.

Note: If you manually enter the information, you will need system IP address, netmask, broadcast, and nameserver (if applicable).

There may be a delay here, while the system displays **Loading Data Into Ramdisk**. The install program starts the VNC Server, once the network has been configured as shown in Figure 2-4.

```
SGI Console
Loading Installation System (100960 kB) - 100%
starting hald... ok
starting syslogd (logging to /dev/tty4)... ok
starting klogd... ok
Starting ia32el done
starting slpd to announce VNC...
starting yast...
RProbing connected terminal...

Initializing virtual console...

Found a Linux console terminal on /dev/console (80 columns x 24 lines).

starting VNC server...
A log file will be written to: /tmp/vncserver.log ...

***
*** You can connect to 192.168.1.100, display :1 now with vncviewer
*** Or use a Java capable browser on http://192.168.1.100:5801/
***

(When YaST2 is finished, close your VNC viewer and return to this window.)

*** Starting YaST2 ***
Xlib: extension "XInputExtension" missing on display ":0.0".
□
```

Figure 2-4 Starting the VNC Server

Note: Control now changes to a web browser on a remote system for the duration of the install. Use the URL specified above if you will be using a web browser, otherwise specify the <IP address:number> shown above when connecting using a VNC client.

Figure 2-5 shows a connection to the VNC Server via a web browser running remotely:

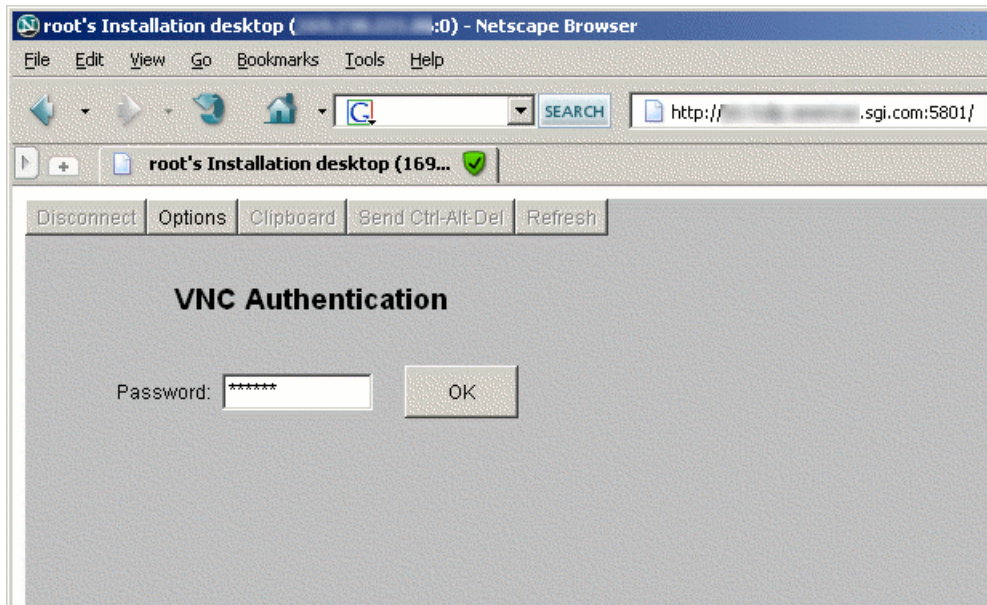


Figure 2-5 VNC Server via a Web Browser

The browser connects to your SGI Altix system, and the VNC Server prompts for the password set earlier in the Boot Options (see step 6).

Language Selection

Note: See Section 3.4, “Language Selection” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

7. Select the appropriate language from the **Language Selection** window and then click on the **Next** button to continue.

License Agreement

Note: See Section 3.6 “License Agreements” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

8. The **License Agreement** window appears. Click on the **Next** button to continue.

Installation

Note: See Section 3.7, “System Analysis” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

9. From the **Installation Mode** window as shown in Figure 2-6, click on the **New Installation button** and the **Include Add-On Products from Separate Media** button, since you will also be adding SGI ProPack 5 SP2 software and then click on the **Next** button to continue.

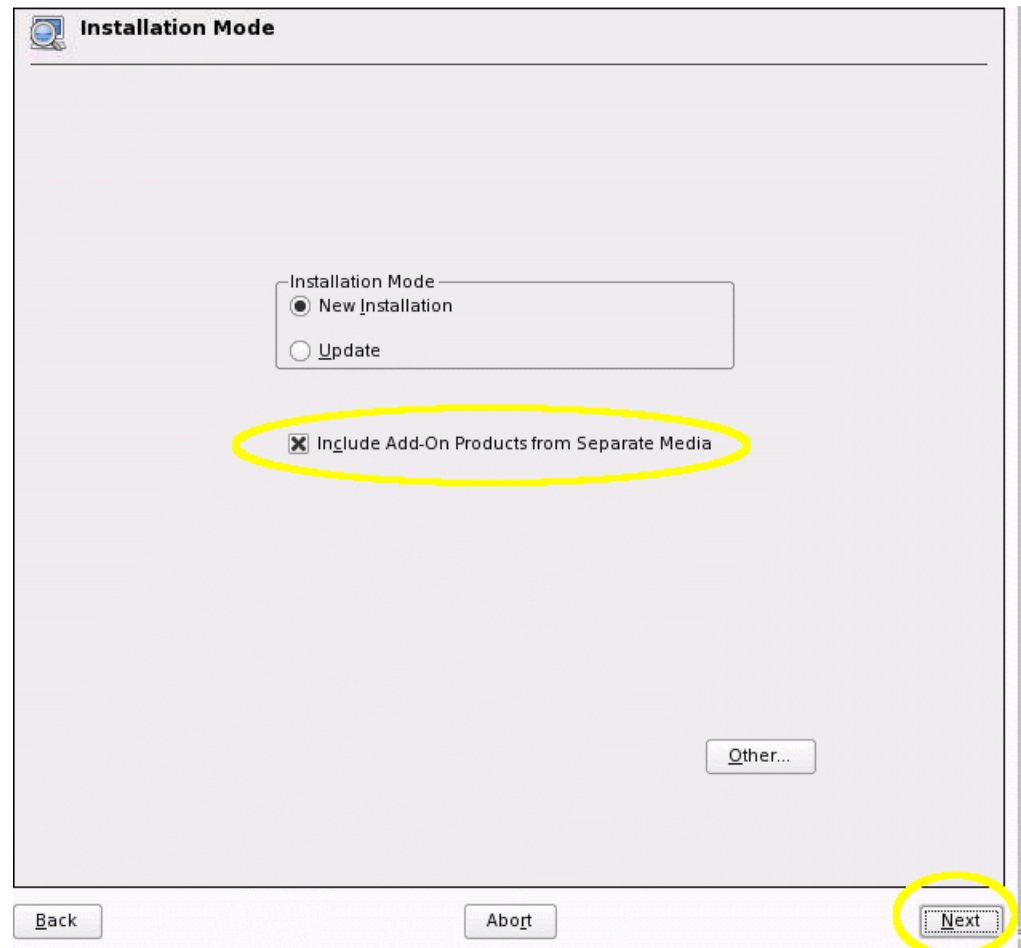


Figure 2-6 Installation Mode Window

10. From the **Add-On Product Installation** window, click on the **Add** button and then click the **Next** button to continue.
11. From the **Add-On Product Media** window, click on the **CD** button and then the **Next** button to continue.

Installation of Add-on Products (especially SGI ProPack 5 SP2)

12. Remove and replace the SLES10 SP1 DVD with the SGI ProPack 5 SP2 CD when prompted and then from the **Insert the Add-on Product DVD** screen, click on the **Continue** button to continue.
13. From the **Import Untrusted GnuPG Key** screen as shown in Figure 2-7, click on the **Import** button to import the SGI Key into YaST.

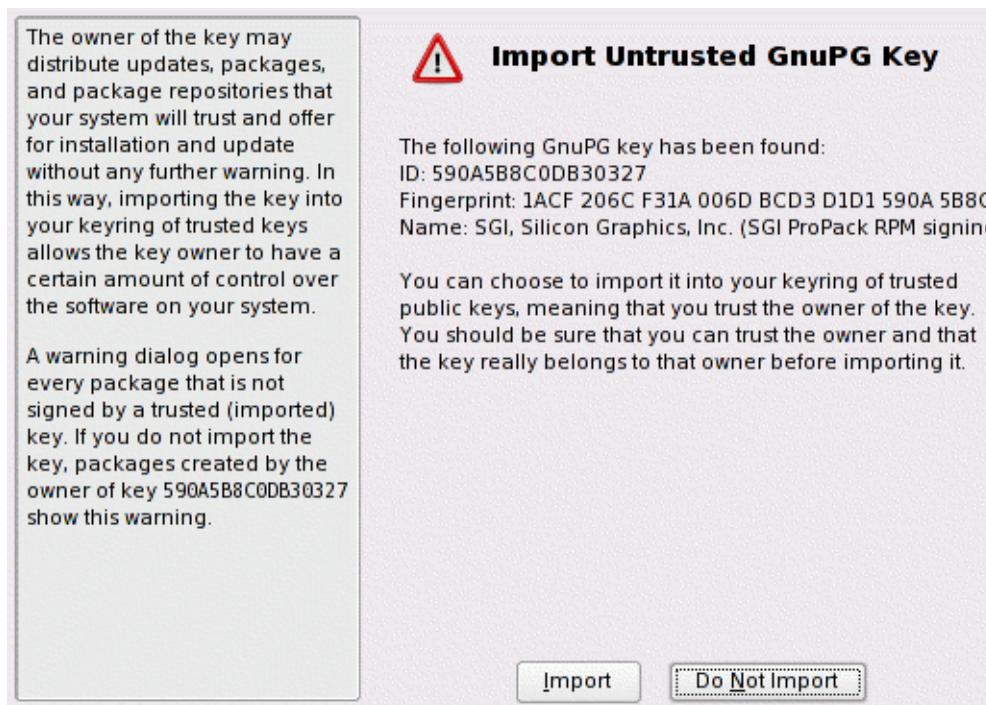


Figure 2-7 Import Untrusted GnuPG Key Screen

14. Choose **SGI ProPack** pattern to install SGI ProPack. Optionally, choose **SGI REACT** if you are an SGI Real-Time customer.
15. Verify that SGI ProPack 5 is now selected for installation as an **Add-On Product** (and **SGI REACT** if needed) and then click on the **Next** button to continue.

Tip: Additional Add-On Products can be added at this point.

Clock and Time Zone

Note: See Section 3.8, “Time Zone” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

16. From the **Clock and Time Zone** window, choose the appropriate **Region** and **Time Zone** and then click on the **Next** button to continue.

Partitioning Set Up

Note: See Section 3.9.1, “Partitioning” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

17. From the **Installation Settings** window, click on the **Change->Partitioning** pull down menu, to bring up the **Expert Partitioner Tool**.

Note: The first time, the tool asks how to initialize the configuration.

The following steps show partitioning on a very basic level. Read the partitioning section in the *SUSE Linux Enterprise Server Installation and Administration* manual. Additional information can be found in Section 7.5.8 "Partitioner". The YaST partitioning utility is very powerful and flexible, but must be used with care, by knowledgeable personnel. Select **Create Custom Partition Setup** if you need to tailor the size or number of partitions beyond what YaST has chosen or you wish to delete existing partitions and create new ones. Follow these guidelines in all cases:

- The EFI boot partition file system format type for ia64 systems should be **FAT**.
- If creating a new partition structure, always create the EFI boot partition for ia64 systems first or the system may not boot.
- Pay attention to whether or not you wish the partition to be formatted. The default is to format the partition.
- SGI suggests using **XFS** for all root filesystems.
- Always review the actions shown in the **Installation Settings** Window.

18. From the **Suggested Partitioning** window, click on the **Base Partition Setup on This Proposal** button to reuse the existing partitions. YaST defaults to the **Reiser** filesystem type, so you will need to change the root filesystem type to **XFS** and format both it and the boot partition before starting the install and then click on the **Next** button to continue.
19. From the **Expert Partitioner** window, select the root partition as shown and click on the **Edit** button and change the filesystem type to **XFS**.
20. From the **Edit Existing Partition** window, perform the following:
 - Verify that the **Format** button is selected.
 - Set the filesystem type to **XFS** using the pull-down menu
 - Verify the **Mount Point** is set to **"/"**

Note: Changes made at this point are not yet being written to the drive. No action will be taken until **Accept** is clicked in the **Installation Settings** Window to accept the final configuration.

Click the **OK** button which will close the **Edit** window and then go back to the **Expert Partitioner** window.

21. Click on the **Edit** button on the **Expert Partitioner** window.
22. For ia64 systems only, from the **Edit Existing Partition** window, perform the following:
 - Verify that the **Format** button is selected
 - Set the filesystem type to **FAT** using the drop-down menu
 - Verify the **Mount Point** is set to **"/boot/efi"**
 - Clicking the **OK** button, will close the **Edit** window, and go back to the **Expert Partitioner**
23. Verify the following settings in the **Expert Partitioner** window:
 - The boot and root partitions now have an **F** in the **F** column, indicating they will be formatted
 - The filesystem type for the root partition is now **XFS**.
 - Click on the **Finish** button to return to the **Installation Settings** screen.

24. Verify the partitioning information in the **Partitioning** section of the **Installation Settings** screen.

Caution: Pay special attention to the actions scheduled to be performed by the **Partitioning** section. There may be some default actions picked by the **Expert Partitioner** that you did not notice, if you have existing partitions that you want to keep. In the example above, the settings are as specified.

Software Selection and System Tasks

Note: See Section 3.9.2, “Software” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

25. Click the **Change->Software** menu to bring up the **Software Selection and System Tasks** screen.
26. The left panel of the **Software Selection and System Tasks** screen shows packages grouped by **Filter: Patterns**. Check marks show which packages are presently selected for installation. To install SGI ProPack 5 as an Add-On product, perform the following:
 - Choose **SGI ProPack** pattern to install SGI ProPack.
 - Optionally, choose **SGI REACT** if you are an SGI Real-Time customer.

Note: To install SGI ProPack 5 kernel modules for other kernel variants (for example, debug), click the **SGI ProPack Kernel Modules** pattern for that kernel variant.

27. YaST reports all the package changes which your selection has triggered. Click the **Continue** button to continue.

The **Installation Settings** screen now reflects the desired partitioning and formatting of the disk and the addition of **SGI ProPack**.

28. If the language needs to be changed, select **Language**. See Section 3.9.3, “Language” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

Setting Default Runlevel

Note: See Section 3.9.4, “Default Runlevel” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

29. Click on the **Expert** button on the **Installation Settings** screen.
30. Select **Change-->Default Runlevel** from the **Installation Settings** screen.
31. From the **Set Default Runlevel** Window, select **3: Full multiuser with network** and then click the **OK** button to continue.

The **Installation Settings** screen now shows **Default Runlevel** as **3: Full multiuser with network**.

Performing the Installation

Note: See Section 3.9.9, “Performing the Installation” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

32. Examine all settings on the **Installation Settings** screen and click **Accept** if satisfied they are all correct.

Tip: The **Expert** tab has additional configuration options.

33. A **Confirm Installation** window appears. This is the last opportunity to make any configuration changes. Any disk formatting will happen now if the **Install** button is clicked. Use the **Back** button to double check if needed.
34. Replace the SGI ProPack 5 SP2 CD with the SLES10 SP1 and then click the **OK** button to continue.

The **Package Installation** window appears. After partitioning is done, the selected software packages are installed

35. When prompted for the SGI ProPack 5 for Linux CD, insert it in the drive and click the **OK** button.

Once the install of all packages is done, the system will finish the basic installation and prepare the system for initial boot. When the reboot dialog box appears, click the **OK** button and the system will reboot.

Once the system reboots, you should find a SUSE Linux Enterprise Server 10 SP1 boot option in your boot manager.

Without manual intervention, the system will boot automatically into SLES10 SP1. If you are using VNC, the VNC server will be restarted.

36. Reconnect to the VNC server using a VNC client or VNC capable web browser (see step 5).

Setting Up Your Hostname and Domain Name

Note: See Section 3.10.1, “Hostname” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

37. Enter your **Hostname** and **Domain Name** and then click the **Next** button to continue.

Setting Up Your root Password

Note: See Section 3.10.2, “root Password” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

38. Enter the password for the root account and then click the **Next** button to continue.

Network Configuration

Note: See Section 3.10.3, “Network Configuration” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

39. Review the information provided in the **Network Configuration** window. If everything is correct, click **Use Following Configuration** button and then click the **Next** button to continue. If the network information is not correct, click the **Change** button and use the options in the pull-down window to make any necessary changes for your specific network configuration.

Note: If your network requires a proxy, make sure it is set correctly before proceeding to the **Test Internet Connection** window.

40. From the **Test Internet Connection** window, verify that the device you will use to connect to the Internet is listed on the screen and click the **Next** button to continue. If the device is not correct, make any modifications that are necessary using the **+** button.

If the system is able to access the Internet, it will try to download the latest SLES10 SP1 release notes and any available product updates. If you are installing SGI ProPack 5 SP2 with SLES10 SP1, the error dialog shown in Figure 2-8 appears when the installer goes to check for SGI ProPack 5 updates.




Figure 2-8 Server Connect Error Message

41. Existing SGI customers with Supportfolio accounts should click on the **Yes** button to display the window shown in Figure 2-9.
42. Uncheck the **Anonymous** box and enter your Supportfolio user name and password and then click the **OK** button. This step lets your system know where the SGI ProPack online update source is located. If you are unable to configure the online update sources, you can configure it at a later time via the Software section in the YaST Control Center.

Server and Directory

Protocol

FTP HTTP HTTPS SMB/CIFS

Server Name

update.sgi.com

Directory on Server

SGI-ProPack/5/ia64

Authentication

Anonymous

User Name

supportfolio_username

Password

OK Cancel

Figure 2-9 Server and Directory Window

Customer Center

Note: See Section 3.10.4, "Customer Center" in the *SUSE Linux Enterprise Server Installation and Administration* manual.

43. If your Internet connection is successful, the **Novell Customer Center Configuration** window appears. SGI recommends that you check the **Config Now** button and register your system with Novell. If you have your registration number that came with your SLES10 SP1 purchase, you may provide that information when prompted to register your subscription with Novell. If you do not have your registration number at this time, the installer will use a temporary registration code so that you can complete this step. You can re-register using your official registration code at a later time.

Online Update

Note: See Section 3.10.5, "Online Update" in the *SUSE Linux Enterprise Server Installation and Administration* manual.

44. If you were able to complete the previous step and register your product with Novell, the **Online Update** window appears. If you wish to download available updates for SLES10 SP1 and SGI ProPack 5 SP2, click the **Run Update** button and click the **Next** button. You will be given the opportunity to select any updates you wish to install. When ready, click the **Accept** button to download and install the updates. After the updates have been installed, click the **Next** button on the **Patch Download and Installation** window. If you choose not to install updates at this time, click the **Skip Update** button and then the **Next** button.

Service

Note: See Section 3.10.6, "Service" in the *SUSE Linux Enterprise Server Installation and Administration* manual.

45. From the **Installation Settings** window, click the Use Following Configuration button and then click the **Next** button to continue.

Users

Note: See Section 3.10.7, "Users" in the *SUSE Linux Enterprise Server Installation and Administration* manual.

46. From the **User Authentication Method** window, click the **Local (/etc/passwd)** button then click the **Next** button to continue.
47. From the **New Local User** window, fill in the appropriate information and then click the **Next** button to continue.

Release Notes

Note: See Section 3.10.7, "Release Notes" in the *SUSE Linux Enterprise Server Installation and Administration* manual.

48. The **Release Notes** window appears and then click the **Next** button to continue.

Hardware

Note: See Section 3.10.10, “Hardware” in the *SUSE Linux Enterprise Server Installation and Administration* manual.

49. From the **Hardware Configuration** window verify the information. If correct, click the **Use Following Configuration** button and click the **Next** button to continue.
50. From the **Installation Completed** window, click the **Finish** button.

The system will now reboot.

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 37
- “Using AutoYaST” on page 38

Setting up the Network Install Server

This section provides you with some information on installing SUSE SLES10 SP1 and SGI ProPack 5 for Linux Service Pack 2 all in one session over the network. This is helpful for systems administrators who have several machines to install and keep up to date.

Note: For information on setting up a network install server, see the “Installing Add-On Products” section in Chapter 7, “System Configuration with Yast” in the *SUSE Linux Enterprise Server Installation and Administration* manual that resides in the `sles-admin.pdf` file in the `docu` directory of the SLES10 SP1 DVD and see “Chapter 4. Configuration and Installation Options” in *AutoYaST* available at http://www.suse.de/~ug/autoyast_doc/index.html.

In SLES10 SP1, you can create a special `add_on_products` file that directs YaST to look at additional install sources. You may list a URL to the SGI ProPack 5 install source in this file.

To create an SGI ProPack 5 SP2 install source, you simply mount the SGI ProPack 5 SP2 CD and copy all of the contents in to a repository. An example of how to do this is, as follows:

```
mkdir /path/to/new-propack-repository
mount /dev/cdrom /mnt
cp -a /mnt/* /path/to/new-propack-repository/
umount /mnt
```

In the `add_on_products` file that exists in your SLES10 SP1 network repository, you just add a URL pointer to your new target. The URL can be `nfs`, `http`, `ftp`, and so. An example entry for the `add_on_products` file is similar to the following:

```
nfs://server/path/to/new-propack-repository
```

Using AutoYaST

AutoYaST is a tool that allows you to install and configure a system based on 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 in *AutoYaST* available at http://www.suse.de/~ug/autoyast_doc/index.html.

You can find information about AutoYaST in Chapter 5, “Automated Installation”, of the *SUSE Linux Enterprise Server Installation and Administration* manual. This resides in the `sles-admin.pdf` file in the `/docu/en` directory of the SLES10 SP1 DVD.

Documentation is also available in `/usr/share/doc/packages/autoyast2/html/index.html` on systems running SLES10 SP1.

Caution: With the initial SLES10 SP1 release, add-on-products handling with AutoYaST is extremely problematic and SGI recommends you do not use it.

Upgrading From an Earlier SGI ProPack 4 Release to SGI ProPack 5 SP2

This section describes how to upgrade from a prior release.

Note: The best way to upgrade a system that is running SLES9 and SGI ProPack 4 is to install both SLES10 SP1 and SGI ProPack 5 SP2 at the same time (see “Installing SLES10 SP1 and SGI ProPack 5 Service Pack 2” on page 19).

Updates from SLES 9 to SLES 10 are supported starting from one of the following bases:

- SLES9 GA
- SLES9 SP3
- SLES9 SP3 plus the latest patches from the maintenance Web

Note: Update a system by starting the SLES10 GA installation system and choosing **Update** instead of **New Installation**. If your system is not already running SLES9 SP3 and SGI ProPack 4 SP3, use the YaST Online Update tool to update your system **BEFORE** attempting to install SLES10 GA and SGI ProPack 5 SP1.

Upgrading for SGI ProPack 5 SP1 to SGI Propack 5 SP2

This section describes how to upgrade your SLES 10 GA system running the SGI ProPack 5 SP1 release to SGI ProPack 5 Service Pack 2.

During this upgrade, your kernel and several other packages may be upgraded to a newer SLES10 SP1 Online Update version. If you have been keeping up-to-date with SLES10 SP1 online updates, you may already have the newer updated packages installed. The upgrade will work fine in either case.

Caution: Do not use **System Update** to update SGI ProPack. The SLES10 SP1 version of **System Update** is not currently supported by SUSE and will not work correctly.

You need to pay careful attention to the upgrade instructions. If you simply were to mark a pattern for upgrade, you will find only a few packages would be installed. You need to work around a problem in YaST to get the install to proceed properly.

SGI suggests that existing customers upgrade their system from SGI ProPack 5 SP1 to SGI ProPack 5 SP2 using media because this ensures you get new packages for SGI ProPack 5 SP2. A simple online-update will not automatically offer new packages.

To upgrade from SGI ProPack 5 SP1 to SGI ProPack 5 SP2, perform the following steps:

1. Remove any old SGI ProPack 5 SP1 install source that may be present. There are several ways to do this. The following example uses the `yast` command.
 - a. Start the graphical or text version of `yast`.
 - b. Select **Software -> Installation Source**.
 - c. Check for an install source named **SGI ProPack 5 for Linux**
 - d. If the SGI ProPack 5 install source exists, select it and delete it.
 - e. Select `Finish` to exit out of the tool.
2. Add ProPack 5 SP2 add-on product, as follows:
 - a. Start the graphical or text version of `yast`.
 - b. Select **Software -> Add-on Product**.

Note: At this point, you need to select and define your media. This example shows the use of a physical CD.

- c. Select **CD...**

You will be prompted to insert the media.
- d. YaST will now add the install source for you.

Note: You will see an error message that states: "An error occurred while connecting to the server. Details: Cannot create the installation source Try again?"

This does **NOT** mean the add-on product failed to add. Instead, this is giving you a chance to enter your authentication information for getting updates/patches from SGI.

e. Answer **Yes** to the re-try question.

You are now given a chance to set up authentication details

f. Unselect the **Anonymous** box.

g. Enter your SupportFolio login name for **User name**.

h. Enter your SupportFolio password for **Password**.

YaST will now add the update source in addition to the ProPack 5 SP2 install source itself.

3. You are now software manager. Perform the following steps:

a. Select **Filter** and then **Patterns**.

Caution: At this point, do not mark the patterns themselves for upgrade. That will cause the upgrade to not complete properly. See the introductory comments to this section (“Upgrading for SGI ProPack 5 SP1 to SGI Propack 5 SP2” on page 39).

b. In Text mode, perform the following:

Highlight the **SGI ProPack** pattern and select **OK**.

Highlight one of the packages that are part of this pattern (any one will do).

Tab to **Actions** and hit **Enter** .

Select **All Listed Packages...**

Select **Update All**

Now all the packages in this pattern should have '>' next to them.

c. In Graphical mode, perform the following:

Highlight the **SGI ProPack** pattern and right-click on any package in the list.

Select **All In List** and then **Update if newer version available**

4. If you have additional SGI ProPack or REACT patterns to update, go to the Caution in step 3a. and repeat the steps for each additional pattern.

5. Select **Accept** to install.

Your system will now be updated.

Installing RHEL 5 and SGI ProPack 5 SP2

This section describes how to install Red Hat Enterprise Linux 5 (RHEL5) and SGI ProPack 5 for Linux Service Pack 2. There are several ways to install SGI ProPack 5 SP2 for RHEL5. The following example shows how to install it in graphical mode, which is the easiest method.

Note: It is also possible to install SGI ProPack 5 SP2 for RHEL5 using yum in text mode. See the `README-yum-install` file in the `docs` directory of the SGI ProPack 5 SP2 CD.

1. To install RHEL 5, see Red Hat Enterprise Linux 5 *Installation Guide for x86, Itanium, AMD64, and Intel Extended Memory 64 Technology (Intel EM64T)*
2. If you are not using a VGA display, this example assumes that you have a workstation running UNIX and can display X11 graphics. You should connect from your workstation to the target server with `ssh` using the `-X` option to pass along graphics to your desktop like this:

```
% ssh -X root@name-of-host
```
3. Make sure you have registered with RHN, that is, Red Hat Network. If you are not registered, perform this command, as follows:

```
% /usr/bin/rhn_register
```
4. Insert the SGI ProPack 5 for Linux Service Pack 2 into your CD-ROM and enter the following command to mount the CD:

```
% mount /dev/cdrom/mnt
```
5. Change directory to the `/mnt` directory, as follows:

```
% cd /mnt
```
6. Run the install script, as follows:

```
% ./INSTALL
```

7. Check the box just to the left of SGI ProPack and then click the **FORWARD** button to continue.

A screen may appear (if you have not pre-installed `sysstat`), as follows:

```
Dependencies Added
```

```
...
```

```
Will continue in 20 seconds.
```

You can wait the 20 seconds or click the **Continue** button.

8. When installation completes, a screen appears, as follows:
The software was successfully installed.
9. Click on the **CLOSE** button to close this screen.

Installing the System Controller Software 1.17 CD

This section provides steps for installing the System Controller Software 1.17 (for IRIX and SGI ProPack for Linux). This software is licensed for the SGI Altix 3000 and 4000 series of servers. The System Controller software provides a system control platform for field service and system maintenance applications.

Note: The System Controller Software 1.17 CD is **not** shipped by default. You can download the System Controller Software 1.17 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 1.17 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 1.17 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 systems with routers have L2 controllers)
- L1 controllers on routers, IX-bricks, PX-bricks, C-bricks, Individual Rack Units (IRU) and so on.
- The system PROM (each node has a system PROM and each C-brick has two nodes).

Note: This section only applies to SGI Altix systems with Intel Itanium processors.

On an SGI Altix 450 or SGI Altix 4700 system, the L1 firmware consists of three parts: the boot image, A image, and B image. For more information, see the “Upgrading L1 Firmware” section in either the *SGI Altix 450 User’s Guide* or *SGI Altix 4700 User’s Guide*.

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.

On an SGI Altix 450 or SGI Altix 4700 system, the L2 system controller is an application that provides control over multiple L1s and communication to other L2s. For more information, see Chapter 3, “System Control” in the *SGI Altix 450 User’s Guide* or Chapter 2, “System Control” in the *SGI Altix 4700 User’s Guide*, respectively.

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: SGI Altix 450 or SGI Altix 4700 systems running SGI ProPack 5 for Linux Service Pack 2 should use PROM located in `shub2snprom.bin`. All other SGI systems should use the PROM located in `shub1snprom.bin`. This note does **not** apply to SGI Altix XE systems (instead see “Getting the Latest Firmware/Software for Your SGI Altix XE System” on page 15). Software support contract customers can download the most recent PROM versions from SGI Supportfolio: <https://support.sgi.com/>

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, instructions on how to flash the PROM are, as follows:

Note: Instructions for flashing the PROMs are applicable to ia64 systems only.

1. If you have already installed SGI ProPack 5, you already have a version of the PROM firmware RPM in place. You can find and install the latest version of the PROM rpm by using the **online_update** mechanism, which will retrieve the latest prom RPM if a newer one is available.
2. If you do not have SGI ProPack 5 installed, or wish to upgrade the PROM prior to installation of SGI ProPack 5, retrieve the appropriate PROM RPM as described earlier in this section.
3. To install the RPM manually, use the following command:

```
rpm -Uvh snprom-XX.X-XXXX.ia64.rpm
```

4. Because there is no way to flash the PROM from Linux, you need to reboot your system and flash from EFI. As the system resets, you will see various messages from the system PROM. Eventually, you will be presented with a boot menu. It will look similar to the following:

```
EFI Boot Manager ver 1.10 [14.62]
```

```
Partition 1:
  CBricks      1      Nodes      1      0
  RBricks      0      CPUs       2      0
  IOBricks     1      Mem(GB)    2      0
```

```
Please select a boot option
```

```
SUSE Enterprise Linux Server 10
EFI Shell [Built-in]
Boot option maintenance menu
```

```
Arrow down to EFI Shell and hit enter.
```

5. Arrow down to EFI Shell [Built-in] and hit the ENTER key to get to the Shell> prompt.
6. You will now be at the Shell> prompt. If you just have one operating system installed and you do not have a CD or DVD in your DVD drive, you can follow this example. If you have a disk in the DVD drive or multiple operating systems, you may need to change to a different EFI filesystem to find the appropriate PROM. This example assumes fs0. Change to fs0, as follows:

```
Shell> fs0:
```

7. If you have an SGI Altix 4000 series system, such as SGI Altix 4700, perform the following command (otherwise, skip this step):

```
cd efi\sgi
```

8. At this point, SGI Altix 4000 series customers are in the efi\sgi directory and Altix 3000 series customers are in \. You can verify the version of the PROM you are currently running, as follows:

```
fs0:\> flash -R
Running PROM version: X.XX
```

9. You can check the version of the PROM binary by running the appropriate command, as follows:

- For SGI Altix 4000 series systems:

```
fs0:\efi\sgi> flash -V ip65prom.bin
SGI PROM Flashing Utility
Version of prom image in file X.XX
```
- For all other Altix systems:

```
fs0:\> flash -V snprom.bin
SGI PROM Flashing Utility
Version of prom image in file X.XX
```

10. You are now ready to flash the PROM on all the nodes, as follows:

- For SGI Altix 4000 series systems:

```
fs0:\efi\sgi> flash -a ip65prom.bin
```
- For all other Altix systems:

```
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 10 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.

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 3000 series, SGI Altix 4000 series, SGI Altix ICE 8000 series, and SGI Altix XE systems. (For a list of the products, see Table 3-1 on page 50.)

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

- “SGI ProPack 5 for Linux Service Pack 2 Products” on page 49
- “Operating System Enhancements” on page 54
- “I/O Subsystems” on page 56
- “HPC Application Tools and Support” on page 57
- “System Management” on page 60
- “NUMA Data Placement Tools” on page 67

SGI ProPack 5 for Linux Service Pack 2 Products

Software provided by SGI for the SGI ProPack 5 for Linux Service Pack 2 release consists of kernel modules for SGI software built against the kernels in SUSE Linux Enterprise Server 10 Service Pack 1 (SLES10 SP1) or Red Hat Enterprise Linux 5 (RHEL5) and value-add software developed by SGI specifically to run on SGI Altix or SGI Altix XE systems and some additional third-party software (see VTune in Table 3-1 on page 50).

Table 3-1 SGI ProPack 5 for Linux Service Pack 2 Products

Product	Architecture Supported	Description
Application performance measuring tools	ia64	<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 an application's performance. For more information on VTune, go to the following web location:</p> <p>http://developer.intel.com/software/products/vtune/</p>
Array Services	ia64	<p>Array Services includes administrator commands, libraries, daemons, and kernel extensions that support the execution of parallel applications across a number of hosts in a cluster, or array. The Message Passing Interface (MPI) of SGI ProPack uses Array Services to launch parallel applications. For information on MPI, see the Message Passing Toolkit (MPT) User's Guide</p> <p>The secure version of Array Services is built to make use of secure sockets layer (SSL) and secure shell (SSH).</p> <p>For more information on standard Array Services or Secure Array Services (SAS), see the Array Services chapter in the <i>Linux Resource Administration Guide</i>.</p>
Cpuset System	ia64 and x86_64	<p>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 5 for Linux" chapter in the <i>Linux Resource Administration Guide</i>.</p>
CSA	ia64	<p>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>.</p>

Table 3-1 SGI ProPack 5 for Linux Service Pack 2 Products **(continued)**

Product	Architecture Supported	Description
DAPL for Altix	ia64	A Direct Access Programming Library (DAPL) is provided for high performance communication between processes on a partitioned or single host Altix system. This can be used with Intel MPI via the remote direct memory access (RDMA) MPI device. See <code>dapl_xpmem(3)</code> for more information.
Embedded Support Partner	ia64 and x86_64	ESP is a support infrastructure which is integrated into the operating system. ESP components capture and store system level events, system configuration information, system availability information, provide an electronic logbook of activities performed on the system, log calls to the ESP database at SGI and other support related activities.
IOC4 serial driver	ia64	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: <code>/dev/ttyIOC4/0</code> <code>/dev/ttyIOC4/1</code> <code>/dev/ttyIOC4/2</code> <code>/dev/ttyIOC4/3</code>
Kernel partitioning support	ia64	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> .
libarena	x86_64	<code>libarena</code> is a custom memory allocator interface and implementation. Four allocators are provided: flat LIFO arena allocator, object pool allocator and two <code>malloc(3)</code> wrappers: one which returns the pointers unadulterated and one which obeys the requested, arbitrary alignment. These can be used directly, or through their exported prototype interfaces.

Table 3-1 SGI ProPack 5 for Linux Service Pack 2 Products **(continued)**

Product	Architecture Supported	Description
L1 and L2 System Controller firmware	ia64	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>
MPT	ia64	Provides industry-standard message passing libraries optimized for SGI computers. For more information on MPT, see the <i>Message Passing Toolkit (MPT) User's Guide</i> .
MVAPICH	x86_64	MVAPICH is open source software developed largely by the Network-Based Computing Laboratory (NBCL) at Ohio State University. MVAPICH develops the Message Passing Interface (MPI) style of process-to-process communications for computing systems employing Infiniband and other Remote Direct Memory Access (RDMA) interconnects. For more information, see "MVAPICH" on page 58.
NUMA tools	ia64 and x86_64	Provides a collection of NUMA related tools (dlook(1), dplace(1), and so on). For more information on NUMA tools, see the <i>Linux Application Tuning Guide</i> .
Performance Co-Pilot collector infrastructure	ia64 and x86_64	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	ia64 and x86_64	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 44.
Reconfigurable Application-Specific Computing	ia64	Reconfigurable application-specific software computing (RASC) program delivers scalable, configurable computing elements for the SGI Altix family of servers and superclusters and Silicon Graphics Prism visualization systems. For more information, see the <i>Reconfigurable Application-Specific Computing User's Guide</i> .

Table 3-1 SGI ProPack 5 for Linux Service Pack 2 Products **(continued)**

Product	Architecture Supported	Description
REACT real-time for Linux	ia64	Support for real-time programs. For more information, see the <i>REACT Real-Time for Linux Programmer's Guide</i> .
SGI Histx	ia64	SGI Histx is a performance analysis tool designed to complement <code>pfmon</code> , which is included with the SGI ProPack software for the Altix family of servers. The software is designed to run on Altix machines only and is only supported on systems running SGI ProPack 5 SP2 with SLES10. Used internally by SGI developers and benchmarkers, as well as SGI customers, this product is offered as a service to customers with a no-fee end user proprietary license via the Supportfolio website. For more information on SGI Histx, see http://www.sgi.com/products/software/histx.html
Utilities	ia64 and x86_64	<code>udev_xscsi</code> , is a <code>udev</code> helper for doing XSCSI device names. <code>sgtools</code> , a set of tools for SCSI disks using the Linux SG driver and <code>lsiutil</code> , the LSI Fusion-MPT host adapter management utility.
XVM	ia64 and x86_64	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> .

SGI does not support the following:

- Base Linux software not released by Novell for SLES10 SP1, Red Hat for RHEL5, or other software not released by SGI.
- Other releases, updates, or patches not released by Novell for SLES10, Red Hat for RHEL5, or by SGI for SGI ProPack software.
- Software patches, drivers, or other changes obtained from the Linux community or vendors other than Novell, Red Hat, and SGI.
- Kernels recompiled or reconfigured to run with parameter settings or other modules as not specified by Novell, Red Hat, 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.

In addition to their traditional use to control the placement of jobs on the CPUs and memory nodes of a system, cpusets also provide a convenient mechanism to control the use of Hyper-Threading Technology.

For detailed information on cpusets, see Chapter 6, "Cpusets on SGI ProPack 5 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 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 SGI Altix 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 IP Addressing of Ethernet Interfaces” on page 56
- “PCI Domain Support for SGI Altix 450 and 4700 Systems” on page 56
- “XSCSI Naming Systems on SGI ProPack Systems” on page 56
- “XFS Filesystem” on page 57
- “XVM Volume Manager” on page 57

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. For more information, see “Persistent Network Interface Names” in the *Linux Configuration and Operations Guide*.

PCI Domain Support for SGI Altix 450 and 4700 Systems

On an SGI Altix 450 and 4700 system, a *PCI domain* is a functional entity that includes a root bridge, subordinate buses under the root bridge, and the peripheral devices it controls. Separation, management, and protection of PCI domains is implemented and controlled by system software. For more information, see “PCI Domain Support for SGI Altix 450 and 4700 Systems” in the *Linux Configuration and Operations Guide*.

XSCSI Naming Systems on SGI ProPack Systems

The XSCSI subsystem on SGI ProPack 3 systems was an I/O infrastructure that leveraged technology from the IRIX operating system to provide more robust error handling, failover, and storage area network (SAN) infrastructure support, as well as long-term,

large system performance tuning. This subsystem is not necessary for SGI ProPack 5 systems. However, XSCSI naming is available on SGI ProPack 5 systems. For more information, see “XSCSI Naming Systems on SGI ProPack Systems” in 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 the *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*.

DAPL for Altix

A Direct Access Programming Library (DAPL) is provided for high performance communication between processes on a partitioned or single host Altix system. This can be used with Intel MPI via the remote direct memory access (RDMA) MPI device. See `dapl_xpmem(3)` for more information and the "Use XPMEM DAPL Library with MPI" section in Chapter 6, "Performance Tuning" in the *Linux Application Tuning Guide*.

MVAPICH

Note: For the SGI ProPack 5 for Linux SP2 release, MVAPICH is **only** supported on SGI Altix ICE 8200 systems.

MVAPICH is open source software developed largely by the Network-Based Computing Laboratory (NBCL) at Ohio State University. MVAPICH develops the Message Passing Interface (MPI) style of process-to-process communications for computing systems employing Infiniband and other Remote Direct Memory Access (RDMA) interconnects.

MVAPICH software is typically used across the network of a cluster computer system for improved performance and scalability of applications.

MVAPICH applications use the Infiniband network of SGI Altix ICE 8200 systems for interprocess RDMA communications. SGI Altix ICE 8200 systems are configured with two Infiniband fabrics, designated as `ib0` and `ib1`. In order to maximize performance, SGI advises that the `ib0` fabric be used for all MPI traffic, including MVAPICH MPI. The `ib1` fabric is reserved for storage related traffic. The default configuration for MVAPICH MPI is to use only the `ib0` fabric. For more information on fabric management on an SGI Altix ICE 8200 system, see the *SGI Altix ICE System Administrator's Guide*.

For more descriptions including the *MVAPICH User Guide* and other MVAPICH publications, see <http://mvapich.cse.ohio-state.edu>.

MVAPICH is installed as part of the Open Fabrics Enterprise Distribution (OFED) on SGI Altix ICE 8200 systems.

A simple test to insure installation of MVAPICH is to enter the command:

```
% rpm -qa | grep mvapich
mvapich_gcc
```

MVAPICH executables, libraries, and man pages can be found in the the following directory:

```
/usr/mpi/gcc/mvapich-0.9.9
```

Build options for MVAPICH are those assembled in the OFED1.2 distribution and include single-rail `ch_gen2` (ibverbs) with ROMIO.

Use of MVAPICH is a multiple step effort in that you must first build your MPI-based application using MVAPICH libraries and then execute your application with the MVAPICH-supplied launch tool, `mpirun_rsh`. A large number of tuning options and launch tool options are available. See the MVAPICH man pages, the *MVAPICH User Guide*, and the MVAPICH web site for information on these options.

It is strongly recommended that MVAPICH applications be built with an MVAPICH-supplied compiler script: `mpicc`, `mpicc`, `mpicxx`, `mpif77`, or `mpif90`. For example, to build an MVAPICH application perform the following:

```
% /usr/mpi/gcc/mvapich-0.9.9/bin/mpicc -o myapp mmpi.c
```

The compiler scripts provide links to all necessary libraries and include files. Applications can be launched using the `mpirun_rsh` command, which is available in the same MVAPICH executables directory. There are a number of options available for launching distributed applications across a cluster. For more information on the use of this tool, enter the the following command:

```
% /usr/mpi/gcc/mvapich-0.9.9/bin/mpirun_rsh --help
```

A very simplified example of its use is, as follows:

```
% /usr/mpi/gcc/mvapich-0.9.9/bin/mpirun_rsh -rsh -np 4 sv0 sv1 sv2 sv3 myapp
```

where you have requested the creation of an MPI process named `myapp` on each of the four hosts `sv0`, `sv1`, `sv2`, and `sv3` using `rsh`.

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.

Embedded Support Partner

Embedded Support Partner (ESP) is a support infrastructure which is integrated into the operating system. ESP components capture and store system level events, system configuration information, system availability information, provide an electronic logbook of activities performed on the system, log calls to the ESP database at SGI and other support related activities. For more information, see *Embedded Support Partner User Guide*.

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*.

Reconfigurable Application-Specific Computing

The RASC program leverages more than 20 years of SGI experience accelerating algorithms in hardware. Rather than using relatively fixed implementations, such as graphics processing units (GPUs), RASC uses field programmable gate array (FPGA) technology to develop a full-featured reconfigurable computer. The RASC program also addresses the ease of use and performance issues present in typical reconfigurable computer environments.

To address performance issues, RASC connects FPGAs into the NUMALink fabric making them a peer to the microprocessor and providing both high bandwidth and low latency. By attaching the FPGA devices to the NUMALink interconnect, RASC places the FPGA resources inside the coherency domain of the computer system. This placement allows the FPGAs extremely high bandwidth (up to 6.4GB/s/FPGA), low latency, and hardware barriers. These features enable both extreme performance and scalability. The RASC product also provides system infrastructure to manage and reprogram the contents of the FPGA quickly for reuse of resources.

With the RASC 2.1 release both block DMA and stream DMA are supported. For more information, see Chapter 5, “Direct I/O” in the *Reconfigurable Application-Specific Computing User’s Guide*.

RASC hardware implementation for SGI Altix 4700 systems is based on blade packaging. The new RASC blade has high-performance FPGAs with 200K logic cells and increased memory resources with 10 synchronous static RAM dual in-line memory modules (SSRAM DIMMs). For legacy systems, optional brick packaging is available for the latest blade-based RASC hardware.

RASC defines a set of APIs through the RASC Abstraction Layer (RASCAL). The abstraction layer can abstract the hardware to provide deep and wide scaling or direct and specific control over each hardware element in the system. In addition RASC provides a FPGA-aware version of GNU Debugger (GDB) that is based on the standard Linux version of GDB with extensions for managing multiple FPGAs. The RASC debug environment does not require learning new tool sets to quickly debug an accelerated application.

RASC supports the common hardware description languages (HDLs) for generating algorithms. RASC provides templates for Verilog- and VHDL-based algorithms. Several 3rd-party high-level language tool vendors are developing RASC interfaces and templates to use in place of traditional hardware design languages.

The “Getting Started with RASC Programming” section in chapter 1 or the *Reconfigurable Application-Specific Computing User’s Guide* provides an overview of FPGA programming.

The RASC tutorial (see “Tutorial Overview” in Chapter 7) steps you through the entire RASC design flow: integrating the algorithm with core service; simulating behavior on the algorithm interfaces; synthesizing the algorithm code; generating a bitstream; transferring that bitstream and metadata to the Altix platform; executing an application; and using GDB to debug an application on the Altix system and FPGA simultaneously.

For detailed information on the SGI RASC program, see the *Reconfigurable Application-Specific Computing 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 SLES10, the base Linux operating system

for SGI Altix systems running SGI ProPack 5. EFI also controls the server's boot configuration, maintaining the boot menu in durable, non-volatile memory.

SLES10 uses `elilo` 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 SLES10 CDs, use the `bootia64` command instead of `elilo`. Once the system is running SLES10, 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

Table 3-2 EFI Commands (continued)

EFI Command	Description
dmpstore	Dumps variable store
echo [-on -off] [text]	Echoes text to stdout or toggles script echo
edit [file name]	Edits a file
endfor	Script only: Delimits loop construct
endif	Script-only: Delimits IF THEN construct
err [level]	Sets or displays error level
exit	Exits
flash filename	Flashes PROM on C-brick
for var in set	Script-only: Indicates loop construct
getmtc	Gets next monotonic count
goto label	Script-only: Jumps to label location in script
guid [-b] [sname]	Dumps known guid IDs
help [-b] [internal command]	Displays this help
if [not] condition then	Script-only: Indicates IF THEN construct
load driver_name	Loads a driver
ls [-b] [dir] [dir] ...	Obtains directory listing
map [-bdvr] [sname[:]] [handle]	Maps shortname to device path
mem [address] [size] [;MMIO]	Dumps memory or memory mapped I/O
memmap [-b]	Dumps memory map
mkdir dir [dir] ...	Makes directory
mm address [width] [;type]	Modifies memory: Mem, MMIO, IO, PCI
mode [col row]	Sets or gets current text mode
mount BlkDevice [sname[:]]	Mounts a filesystem on a block device

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

EFI Command	Description
<code>mv <i>sfile dfile</i></code>	Moves files
<code>pause</code>	Script-only: Prompts to quit or continue
<code>pci [<i>bus dev</i>] [<i>func</i>]</code>	Displays PCI device(s) info
<code>reset [<i>cold/warm</i>] [<i>reset string</i>]</code>	Indicates cold or warm reset
<code>rm <i>file/dir</i> [<i>file/dir</i>]</code>	Removes file or directories
<code>set [-bdv] [<i>sname</i>] [<i>value</i>]</code>	Sets or gets environment variable
<code>setsize <i>newsiz</i> <i>fname</i></code>	Sets the files size
<code>stall <i>microseconds</i></code>	Delays for <i>x</i> microseconds
<code>time [<i>hh:mm:ss</i>]</code>	Gets or sets time
<code>touch [<i>filename</i>]</code>	Views or sets file attributes
<code>type [-a] [-u] [-b] <i>file</i></code>	Types file
<code>ver</code>	Displays version information
<code>vol <i>fs</i> [<i>volume label</i>]</code>	Sets or displays volume label

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 and 450 systems and the SGI Altix 3000 and 4000 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 Level 1 (L1) and Level 2 (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. An SGI Altix 4700 system can have multiple L1 host connections depending on the size of the system. The L1 controller 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.

System control on SGI Altix 450 and 4700 systems has changed from previous SGI Altix systems. All SGI Altix 450 and 4700 system enclosures contain an embedded microprocessor board and display assembly called the system controller. This microprocessor runs an embedded version of the Linux operating system. The system controller runs off standby power and is running as long as the enclosure is connected to an active power source. On an SGI Altix 450 and 4700 system, the L2, or Level 2 system controller, is an application that provides control over multiple L1s and communication to other L2s.

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

- The “Operating the L2 Controller on Altix 4700 Systems” section in the *SGI L1 and L2 Controller Software User’s Guide*.
- 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 level 3 (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 data placement 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.

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 5, “Data Placement Tools” in the *Linux Application Tuning Guide*.

Index

A

Array Services, description, 50

C

CD contents, 2

Comprehensive System Accounting (CSA) support, 55

Cpuset system, 54

CSA support, 55

D

Disk partition table, 11

dlook command, 68

dplace command, 68

E

EFI commands, 63

EFI support, 62

Extensible Firmware Interface (EFI) support, 62

H

hardware platforms, 1

HPC support

libraries and tools, 57

Message Passing Toolkit, 58

Performance Co-Pilot, 60

I

Installation

CDs, 43

installing SLES9 SP2, 16

overview, 6

read me first, 16

setting up a network install server, 37

system controller software, 43

upgrading, 47

I/O subsystems

for HPC systems, 56

XFS filesystem support, 57

XVM support, 57

L

L1 and L2 controllers, 66

Launchword

/usr/bin/nr

http

[//www.intel.com/design/Xeon/documentation.htm](http://www.intel.com/design/Xeon/documentation.htm), xx

License manager updates, 15

Linux Kernel Crash Dump (LKCD), 13

LKCD for file configuration, 13

M

Message Passing Toolkit
support, 58
Multiple Ethernet connection, 12
MVAPICH, 59

N

NUMA tools, 67
dllook command, 68
dplace command, 68

O

Online Support Services
Supportfolio, xvi
Open source website, 6
OS configuration settings, 1
OS enhancements
Comprehensive System Accounting (CSA), 55
cpuset support, 54
for HPC environments, 54
partitioning, 55

P

Password requirement, 9
PCI domain, 56
Performance
measuring, 50
Performance Co-Pilot support, 60
Persistent naming
Ethernet devices, 56
XSCSI devices, 57
Product list, 49

Product support, 49
PROM support, 61

R

Reconfigurable Application-Specific Computing, 61
Release features, 1

S

SGIconsole support, 65
Software
installation, 9
introduction, 1
planning, 10
Supportfolio, xvi
System controller firmware support, 66
System Controller Software
CD, 43
System management
EFI, 62
PROM chips, 61
SGIconsole, 65
system controller firmware, 66
tools, 60

U

Unsupported elements, 53
Upgrading a system, 39
Upgrading from previous releases, 43

X

XFS filesystem support, 57

XVM support, 57

XVM Volumes in Rescue Mode, 12

