



SGI ProPack™ 3 to SGI ProPack™ 4 for Linux® Migration Guide

007-4719-001

COPYRIGHT

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

LIMITED RIGHTS LEGEND

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

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo and Altix are registered trademarks and SGI ProPack for Linux is a trademark of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

Intel is a registered trademark of Intel Corporation, in the United States and other countries. Linux is a registered trademark of Linus Torvalds, used with permission by Silicon Graphics, Inc. SUSE LINUX and the SUSE logo are registered trademarks of Novell, Inc. All other trademarks mentioned herein are the property of their respective owners.

Record of Revision

Version	Description
001	February 2005 Original publication. Supports the SGI ProPack 4 for Linux release.

Contents

1. Migrating from SGI ProPack 3 to SGI ProPack 4 for Linux	1
About This Guide	1
Audience	2
SUSE LINUX Migration Documentation	2
Migration Overview	2
Differences between SGI ProPack 3 and SGI ProPack 4	3
Latest PROM Firmware	4
Topology	4
Booting a Booting a SLES9 SP1 System	4
Linux Threads Behavior	5
Kernel Placement	5
User IDs/Group IDs	6
CpuMemSets	6
Cpusets	6
Installing and Configuring Array Services	6
Comprehensive System Accounting	6
Partitioning Software	7
Persistent IP Addressing of Ethernet Interfaces	7
Gigabit Ethernet Driver	7
Bonding Driver	7
runon Command	8
KDB Kernel Debugger sn2kdb command	8
Hardware Information	8
xfsdump/xfrestore Commands	8
007-4719-001	v

Contents

YaST Installer	8
Per-Node Memory Available Statistics	9
Kernels	10
kernel-sn2-2.6.5-7.*	10
kernel-64k-pagesize-2.6.5-7.*	11
kernel-default-2.6.5-7.*	11
kernel-debug-2.6.5-7.*	11
Major Documentation Changes for ProPack 4	12
Obtaining Publications	12
Reader Comments	13

Migrating from SGI ProPack 3 to SGI ProPack 4 for Linux

This guide contains information about migrating from SGI ProPack 3 for Linux operating system to SGI ProPack 4 for Linux operating system.

Note: The SGI ProPack 4 for Linux main download page, caveats page, and information about which configuration files you should save off before migrating to the SGI ProPack 4 release are available at the following location on Supportfolio:http://support.sgi.com/content_request/329516/index.html

SGI ProPack contains numerous libraries, utilities, and other extensions designed to supplement a base Linux distribution with additional capabilities specifically for HPC, storage, and visualization environments.

Starting with ProPack 4, SGI ProPack will be available to run on top of SUSE SLES 9 SP1 (a standard, off-the-shelf Linux distribution) rather than SGI Advanced Linux Environment (an SGI built Linux distribution).

ProPack 4 will be certified and available with the 2.6 based Linux distribution of SLES 9 SP1 (or later) only going forward. ProPack 4 (and later) is not supported with other Linux distributions or SGI Advanced Linux Environment.

About This Guide

This guide provides information about specific differences between SGI ProPack 3 and SGI ProPack 4. For general information about migrating to SUSE LINUX Enterprise Server 9 (SLES9), see "SUSE LINUX Migration Documentation" on page 2.

This guide contains the following sections:

- "Audience" on page 2
- "SUSE LINUX Migration Documentation" on page 2
- "Migration Overview" on page 2
- "Differences between SGI ProPack 3 and SGI ProPack 4" on page 3

- "Major Documentation Changes for ProPack 4" on page 12

Audience

This guide is written for system administrators who are responsible for installing, configuring, and administering the SGI ProPack for Linux operating system running on the SGI Altix family of servers and superclusters.

SUSE LINUX Migration Documentation

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

SUSE provides information about migrating to SLES9 at the following location:
http://www.novell.com/resourcecenter/ext_item.jsp?itemId=13485.

SUSE also provides information about migrating applications to SUSE Linux at the following location:
http://www.novell.com/coolsolutions/cooldev/features/a_red_hat_migration_cdev.html.

Migration Overview

This section describes general differences between SGI ProPack 3 for Linux and SGI ProPack 4 for Linux.

SGI ProPack 3 was tied to the SGI Advanced Linux Environment 3, an SGI operating system over which SGI had control. SGI ProPack 4 runs layered over SLES 9 SP1, with some additions by SGI. SGI does **not** provide its own version of the kernel software.

Kernel fixes and updates are provided by SUSE.

Disk devices are quite different in ProPack 4 versus ProPack 3. The SUSE YaST installation tool and other SUSE tools do not recognize XSCSI device names. Device names are of the `/dev/sdX` variety. For most SCSI and Fibre channel disks, persistent naming can be achieved by looking through `/dev/disk/`. Although an XSCSI

compatibility mechanism is set up under `udev`, XSCSI paths should not be used in general except for working with Fibre channel device names.

The SLES9 kernel is far more modular than ProPack 3 kernel. In ProPack 4 with SLES9 SP1 as the base system, a typical system can have upwards of 50 kernel modules loaded.

Differences between SGI ProPack 3 and SGI ProPack 4

This section describes some important differences between SGI ProPack 3 and SGI ProPack 4 and covers these topics:

- "Latest PROM Firmware" on page 4
- "Topology" on page 4
- "Bootting a Bootting a SLES9 SP1 System" on page 4
- "Linux Threads Behavior" on page 5
- "Kernel Placement" on page 5
- "User IDs/Group IDs" on page 6
- "CpuMemSets" on page 6
- "Cpusets" on page 6
- "Installing and Configuring Array Services" on page 6
- "Comprehensive System Accounting" on page 6
- "Partitioning Software" on page 7
- "Persistent IP Addressing of Ethernet Interfaces" on page 7
- "Gigabit Ethernet Driver" on page 7
- "Bonding Driver" on page 7
- "`runon` Command" on page 8
- "KDB Kernel Debugger `sn2kdb` command" on page 8
- "Hardware Information" on page 8

- "xfsdump/xfrestore Commands" on page 8
- "YaST Installer" on page 8
- "Per-Node Memory Available Statistics" on page 9
- "Kernels" on page 10

Latest PROM Firmware

Unless you have the most current PROM, the kernel may fail to boot during the installation procedure. The prom RPM is on ProPack 4 Proprietary CD in SGI/RPMS directory.

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

For information on obtaining the latest PROM, see the "Obtaining the Latest SGI Altix System Firmware" section in Chapter 2, "Software Planning and Installation", in the *SGI ProPack 4 for Linux Start Here*.

Topology

Topology information is available in `/proc/sgi_sn/sn_topology`. `/dev/hw` is not available on ProPack 4 systems. You may access this information in several different formats using `/usr/bin/topology` or `/usr/bin/gtopology`.

Booting a Booting a SLES9 SP1 System

Use `efi/boot/bootia64` instead of `elilo` to boot from the SLES9 or SLES9 SP1 CDs (no boot option should be applied and any other boot option like `console=ttyS0` should be placed in the input field of the boot menu that appears after `bootia64` has been started). Once the system is running, SLES9 (and/or ProPack 4), use `elilo` to boot from `efi`.

If you have installed multiple kernel images and you 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 from the `/boot` directory to `/boot/efi/efi/SuSE` directory.

You should never change the `/boot/efi/efi/suse/elilo.conf` file directly. Instead, you can change the `/etc/elilo.conf` file and run the `elilo` command. This helps prevent the installer software from clobbering your changes.

Linux Threads Behavior

With SGI ProPack 3, to use the old Linux threads `pthread` model, you needed to set the `LD_ASSUME_KERNEL` variable, as follows

```
export LD_ASSUME_KERNEL=2.4.19
```

With SGI ProPack 4 on SLES9 SP1, to get the same thread behavior, set the `LD_ASSUME_KERNEL` variable, as follows:

```
export LD_ASSUME_KERNEL=2.4.21
```

This is documented in the SLES9 release notes. Applications compiled on SGI ProPack 3 (and some compiled on SGI ProPack v2.4) will run with this environment variable set. For more information about applications built on SGI ProPack for Linux using the Native Posix Thread Library, see the "Profiling with `profile.pl`" in the *Linux Application Tuning Guide*.

Kernel Placement

In SGI ProPack 3, when you installed a new kernel with the `rpm -ivh` command, the new kernel would appear in the `/boot/efi/efi/sgi` directory and but would not overwrite old kernels already there. With SLES9, if you install a kernel with `rpm -ivh`, the old kernel is wiped out in `/boot/efi/efi/suse` directory. There is only one kernel and only one `initrd(4)` boot loader block device stored in the `/boot/efi/efi/suse` directory. The `rpm` listing will still show multiple kernels and they will still be available in the `/boot` directory. However, `elilo` will only be able to find the "latest kernel you installed" at this time. This is also how YaST/YaST2 installer handles kernel installation. For more information about booting different installed kernels, see "Bootting a Booting a SLES9 SP1 System" on page 4.

User IDs/Group IDs

SGI ProPack 3 user/group IDs start at 500 in the `/etc/login.defs` file. SUSE user/group IDs start at 1000 in the `/etc/login.defs` file. If you are running both SGI ProPack 3 and SLES9 SP1 and SGI ProPack 4 on the same system, set the UIDs to 1000 in SGI ProPack 3 to avoid problems when switching between SGI ProPack 3 and SGI ProPack 4.

CpuMemSets

CpuMemSets functionality in SGI ProPack 3 for Linux is contained within the new `cpuset` implementation on SGI ProPack 4. For more information, see "Cpusets".

Cpusets

The `cpuset` implementation is new on SGI ProPack 4. For detailed information, see Chapter 6, "Cpusets on SGI ProPack 4 for Linux" in the *Linux Resource Administration Guide*. You can also find information at the following location:

`/usr/share/doc/packages/cpuset/cpuset.html`

Installing and Configuring Array Services

An account must exist on all hosts in the array for the purposes of running certain Array Services commands. This is controlled by the `/usr/lib/array/arrayd.conf` configuration file. The default is to use the user account `arraysvcs` on SGI ProPack 4 systems. The `arraysvcs` user account should be created by default during installation. The user account `guest` is used on SGI ProPack 3 systems. The account name can be changed in `arrayd.conf`. For more information, see the `arrayd.conf(5)` man page and Chapter 3, "Array Services in the *Linux Resource Administration Guide*."

Comprehensive System Accounting

In Comprehensive System Accounting (CSA), user exits allow you to tailor the `csarun` or `csaperiod` procedures to the specific needs of your site by creating scripts to perform additional site-specific processing during daily accounting. You need to create user exit files owned by user `csaacct` (this was user `adm` for SGI ProPack 3) with execute permission if your site uses the accounting user exits. The

`csaacct` user account is created by default during installation. For more information, see Chapter 2, “Comprehensive System Accounting” in the *Linux Resource Administration Guide*.

Partitioning Software

On SGI ProPack 4 systems, if you intend to use the cross-partition functionality of XPMEM, you will need to add `xpc` to the line in the `/etc/sysconfig/kernel` file that begins with `MODULES_LOADED_ON_BOOT`. Once that is added, you may either reboot the system or issue a `modprobe xpc` command to get the cross-partition functionality to start working. For more information on using `modprobe`, see the `modprobe(8)` man page. For more information on partitioning software, see the “System Partitioning” section in the *Linux Configuration and Operations Guide*.

Persistent IP Addressing of Ethernet Interfaces

The mechanism for associating persistent IP addresses with Ethernet interfaces is different in SLES9 SP1 than it was in SGI ProPack 3. In particular, the `/etc/sysconfig/networking/eth0_persist` file is no longer used. For information on Persistent IP Addressing of Ethernet Interfaces in SGI ProPack 4, see Chapter 3 of the *SGI ProPack 4 for Linux Start Here*.

Gigabit Ethernet Driver

The default Gigabit Ethernet driver for SGI ProPack 4 is called `tg3-new`. On SGI ProPack 3 systems it is `tg3`. The new driver performs better and has important bug fixes and is therefore recommended by SGI. The new driver will be used by default and you can verify that it is in use by checking the driver name returned by `ethtool -i interface_name`.

Bonding Driver

The network configuration files in SLES9 (and therefore in SGI ProPack 4) do not support the “master/slave” notation which is required to configure bonding devices automatically during network initialization. It is necessary to configure the bonding interfaces manually. For details, see `/usr/src/linux/Documentation/networking/bonding.txt` on your system.

runon Command

The `runon(1)` command is a deprecated package on SGI ProPack 4. You can use the `taskset(1)` command to retrieve or set the CPU affinity of a process. For more information, see the `taskset(1)` man page.

KDB Kernel Debugger `sn2kdb` command

The `sn2kdb` command has been replaced with the `archkdb` command.

Hardware Information

The Altix hardware inventory command, `hinv(1)` is not available on SLES9 SP1 with ProPack 4. Instead, please use the SUSE `hwinfo` command (`hwinfo --short`, `hwinfo --disk`, and `sitar --format=html`).

Note: The `hwinfo` command can take a very long time to complete when run as root. Run as a non-root user for quicker results.

For more information, see the `hwinfo(8)` man page and information in `/usr/share/doc/packages/hwinfo`.

`xfsdump/xfsrestore` Commands

`xfsdump(8)` and `xfsrestore(8)` are not loaded by default. Install the `xfsdump` and `xfsprogs-devel` rpms if you need these commands.

YaST Installer

The YaST installer does not allow you to remove older partitions and creating a clean root by default. To do so, use **Expert** mode. The default filesystem type is `reiserfs`. To select XFS, you will need to select **Edit** to edit the partition. Boot partitions should use the FAT filesystem.

With the exception of `fdisk`, `xvm`, and some `sysctl` options, it is best to let YAST handle options instead of trying to modify files by hand.

Per-Node Memory Available Statistics

The per-node memory available statistics that were available in `/proc/discontig` are now part of `sysfs`. For example:

```
% cat /sys/devices/system/node/node0/{meminfo,numastat}
```

```
Node 0 MemTotal:      1972384 kB
Node 0 MemFree:      1294912 kB
Node 0 MemUsed:      677472 kB
Node 0 HighTotal:    0 kB
Node 0 HighFree:    0 kB
Node 0 LowTotal:    1972384 kB
Node 0 LowFree:    1294912 kB
HugePages_Total:    0
HugePages_Free:    0
Hugepagesize:    262144 kB
numa_hit 3070703
numa_miss 0
numa_foreign 0
interleave_hit 1
local_node 3069479
other_node 1224
```

The following is text for the new `numa.mem` metrics:

```
numa.mem.util.total [per-node total memory]
numa.mem.util.free [per-node free memory]
numa.mem.util.used [per-node used memory]
numa.mem.util.hightotal [per-node highmem total]
numa.mem.util.highfree [per-node highmem free]
numa.mem.util.lowtotal [per-node lowmem total]
numa.mem.util.lowfree [per-node lowmem free]
numa.mem.util.hugepages_total [per-node hugepages total]
numa.mem.util.hugepages_free [per-node hugepages free]
numa.mem.util.hugepagessize [per-node hugepage size]
numa.mem.alloc.numa_hit [per-node count of number of times a process
  wanted to alloc memory on the local node and succeeded]
numa.mem.alloc.numa_miss [per-node count of times a process wanted to
  alloc on the local node but got memory on another node]
numa.mem.alloc.numa_foreign [count of times a process on another node
```

```
wanted to alloc on that node, but got this node]
numa.mem.alloc.interleave_hit [count of the number of time interleaving
wanted to allocate on this node and succeeded]
numa.mem.alloc.local_node [count of times a process ran on this node
and got memory on this node]
numa.mem.alloc.other_node [count of times a process ran on this node
and got memory from another node]
```

Kernels

The ProPack CDs no longer contain any kernel packages. SLES9 SP1 contains four different kernels and they all ship as part of the SLES9 SP1 set of ISO images. This section gives a brief explanation of the differences between each kernel. It covers these kernels:

- "kernel-sn2-2.6.5-7.*" on page 10
- "kernel-64k-pagesize-2.6.5-7.*" on page 11
- "kernel-default-2.6.5-7.*" on page 11
- "kernel-debug-2.6.5-7.*" on page 11

kernel-sn2-2.6.5-7.*

You must explicitly select the sn2 kernel for installation during the install of SLES9/SLES9 SP1.

Note: SGI recommends that you install the sn2 kernel because the kdb debugger used to diagnose system panics and hangs is only available with the sn2 kernel.

The kernel-sn2-2.6.5-7.* kernel has these characteristics:

- sn2 specific kernel
- Maximum 512 CPU support
- 16KB page size
- KDB support (must be configured on)

kernel-64k-pagesize-2.6.5-7.*

The `kernel-64k-pagesize-2.6.5-7.*` kernel has these characteristics:

- ia64 generic kernel
- Maximum 128 CPU support
- 64KB page size
- No KDB support
- Does not work with XFS filesystems of 512 byte blocks, but does work if XFS filesystem is created with larger blocks

kernel-default-2.6.5-7.*

The `kernel-default-2.6.5-7.*` kernel has these characteristics:

- ia64 generic kernel
- Maximum 128 CPU support
- 16KB page size
- no KDB support

kernel-debug-2.6.5-7.*

The `kernel-debug-2.6.5-7.*` kernel has these characteristics:

- ia64 generic kernel
- Maximum 128 CPU support
- 16KB page size
- KDB support (enabled by default)
- kprobes support
- dprobes support
- slab/spinlock debugging enabled

Major Documentation Changes for ProPack 4

The following is a list of SGI ProPack for Linux manuals that were updated for the SGI ProPack 4 for Linux release and a description of those updates.

- *SGI ProPack 4 for Linux Start Here*

This guide provides information about the SGI ProPack 4 for Linux release including release features, software planning and installation information and a description of product components.

It includes important information about how I/O subsystems work on ProPack 4, in several sections of chapter 3, as follows:

- “Persistent PCI Bus Numbering” that ensures that bus numbers can remain the same across reboots in case of faulty hardware or reconfiguration.
- “Persistent IP Addressing of Ethernet Interfaces” so that 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

- *Linux Configuration and Operations Guide*

Explains how to perform system configuration and operations under the Linux operating system specific to the SGI Altix family servers and superclusters. It contains new information specific to Altix systems running the SGI ProPack 4 for Linux operating system. Changes were made throughout the manual for SGI ProPack 4, in particular “System Partitioning” section in Chapter 1.

- *Linux Resource Administration Guide*

This manual contains information needed in the administration of various system resource management features. It contains several new chapters specific to SGI ProPack 4 for Linux, as follows:

- Chapter 6, “Cpusets on SGI ProPack 4 for Linux”
- Appendix B, “SGI ProPack 4 Cpuset Library Functions”

Obtaining Publications

To obtain SGI documentation, go to the SGI Technical Publications Library at <http://docs.sgi.com>.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. 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 can be found on the back cover.)

You can contact us in any of the following ways:

- Send e-mail to the following address:

`techpubs@sgi.com`

- Use the Feedback option on the Technical Publications Library World Wide Web page:

`http://docs.sgi.com`

- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

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

We value your comments and will respond to them promptly.