



SGI® Performance Suite 1.5 Start Here

007-5680-006

COPYRIGHT

© 2010, 2011, 2012, 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

Accelerate, Altix, ICE, IRIX, NUMAflex, NUMAlink, OpenGL, Performance Co-Pilot, REACT, SGI, the SGI logo, SHMEM, Silicon Graphics, Supportfolio, UV, and XFS are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and other countries.

FLEXlm is a registered trademark of Macrovision Corporation. Intel 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 the U.S. and other countries. Novell and SUSE are registered trademarks of Novell, Inc., in the United States and other countries. QLogic is a registered trademark of QLogic Corporation. Platform Manager is a trademark of Platform Computing, Inc. 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. 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	October 2010 Original publication. Supports the SGI Performance Suite 1.0 release.
002	February 2011 Updated to support the SGI Performance Suite 1.1 release.
003	April 2011 Updated to support the SGI Performance Suite 1.2 release.
004	November 2011 Updated to support the SGI Performance Suite 1.3 release.
005	May 2012 Updated to support the SGI Performance Suite 1.4 release.
006	November 2012 Updated to support the SGI Performance Suite 1.5 release.

New Features in This Guide

This update of *SIG Performance Suite 1.5 Start Here* supports the SIG Performance Suite 1.5 release.

Major Documentation Changes

Changes in this guide for the SIG Performance Suite 1.5 release include the following:

- Removed references to the Embedded Support Partner.

Contents

Record of Revision	iii
Tables	ix
Helpful Online Resources	xi
SUSE Documentation	xii
Red Hat Documentation	xiii
Intel Compiler Documentation	xiv
Other Intel Documentation	xiv
SGI Documentation	xiv
SGI Performance Suite Software Manuals	xiv
Additional SGI Software Manuals	xvi
SGI Hardware Manuals	xviii
Obtaining SGI Publications	xix
Conventions	xx
Reader Comments	xx
1. Release Features	1
Software Introduction	1
SGI Performance Suite Software Bundles	2
SGI® MPI	2
SGI® REACT™	3
SGI® UPC	4
SGI® Accelerate™	5
Software Package License Information	6
SGI Performance Suite CDs	6
Release Information and Release Notes	7
Finding SGI Performance Suite Manuals	7

Getting Updates	8
SLES 11 SP2 Updates	8
RHEL 6.3 Updates	8
SGI Performance Suite Updates	9
SGI Foundation Software 2.7 Updates	10
SGI Open Source Information.	11
Documentation Updated for This Release	11
2. SGI Performance Suite Software Installation	13
Installing or Upgrading SGI Performance Suite	13
Systems with Multiple Ethernet Configurations	13
XVM Volumes Access in Rescue Mode	14
Adding the <code>xfst_quota</code> Variable to the <code>initrd</code> Boot Image.	14
Superpages	14
SLES 11 SP2 or RHEL 6.3 Kernels.	16
Getting the Latest Firmware for Your SGI System	16
License Information	17
3. SGI Performance Suite Overview	19
SGI Performance Suite 1.5 Products	19
Performance Suite Products	22
Cpuset Support	22
Partitioning.	23
HPC Application Tools and Support.	23
Message Passing Toolkit	23
Performance Co-Pilot	24
NUMA Data Placement Tools.	24
<code>dlook</code> Command	24
<code>dplace</code> Command	25
<code>taskset</code> Command	25
Index.	27

Tables

Table 3-1	SGI Performance Suite Products	20
------------------	--	----

About This Guide

This guide provides information about the SGI Performance Suite 1.5 release.

The SGI Performance Suite CD includes the following in the `/docs` directory:

- A PDF version of this document, the *SGI Performance Suite Start Here*.
- A text file named `product-release-readme.txt`, which contains the release notes. The release notes contain the latest information about the software and documentation in this release.
- A text file named `product-release-rpms.txt`, which lists the RPMs distributed with SGI Performance Suite bundle.

Helpful Online Resources

This section describes the electronic support services available for your SGI system running SGI Performance Suite software. These are as follows:

- Supportfolio
Supportfolio is the SGI support website. It hosts the SGI Knowledgebase, and it hosts links for software support and updates.
<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 Update Server

The SGI update server is available for your convenience. Get updates to the SGI Performance Suite software releases using update tools (for example, YaST, zypper and yum) available in your operating system environment. Refer to the product release notes for more information about specific product details on the SGI update server.

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

- Supportfolio Software Updates

You can check for updates to your purchased SGI software, and you can download or order CDs all in one convenient place. Supportfolio Software Updates provide 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 online, keep track of all your cases, and receive automatic update information when your case is updated via the Supportfolio Problem Resolution Tools.

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

- SGI Software Licensing

You can order SGI software licenses conveniently via e-mail or online. If you send a blank e-mail to license@sgi.com, SGI responds via e-mail by sending you a template to fill out. SGI sends you your license quickly with complete instructions for license installation.

Alternatively, go online to the SGI Key-O-Matic.

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

SUSE Documentation

Documentation for SUSE Linux Enterprise Server 11 (SLES11) is located on your system in the `/docu/en` directory.

For SUSE Linux Enterprise Server 11 package descriptions, see the following website:

<http://www.suse.com/products/server/technical-information/#Package>

For SUSE documentation, see the following website:

<https://www.suse.com/documentation>

The following SLES11 documentation is available:

- *SUSE Linux Enterprise Server 11 Installation Quick Start*
art_sles_installquickstart.pdf
- *SUSE Linux Enterprise Server 11 Deployment Guide*
book_sle_deployment.pdf
- *SUSE Linux Enterprise Server 11 Administration Guide*
book_sle_admin.pdf
- *SUSE Linux Enterprise Server SLES11: Storage Administration Guide*
stor_admin.pdf

Red Hat Documentation

Documentation for Red Hat Enterprise Linux 6.3 (RHEL 6.3) or RHEL 6.2 is located in the `/usr/share/doc` directory. It contains package-specific documentation. If you install the `zifofs` module, additional documentation in the form of READMEs can be found in `/usr/share/doc/zifofs-tools-1.x.x`.

On the RHEL6 DVD, `README-{ lang }` and `RELEASE-NOTES-{ lang }` files are found in the root directory.

Note: Red Hat Enterprise Linux 6 (RHEL 6) documentation is available at the following website:

<http://www.redhat.com/docs/manuals/enterprise/RHEL-6-manual/index.html>

The following documentation is available:

- *Release Notes for x86_64 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 that resides within each compiler's installation directory. If you 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

Intel server processor information is available online at:

<http://www.intel.com/support/processors/sb/CS-029807.htm>

SGI Documentation

Online and postscript versions of SGI documentation are available at the SGI Technical Publications Library at the following website:

<http://docs.sgi.com>

The following topics describe the SGI documentation:

- “SGI Performance Suite Software Manuals” on page xiv
- “Additional SGI Software Manuals” on page xvi
- “SGI Hardware Manuals” on page xviii

SGI Performance Suite Software Manuals

The following documentation describes the SGI Performance Suite software:

- *SGI Performance Suite Product Bundles Release Notes*

Provides the latest information about software and documentation in this release. Release notes are in a file in the `/docs` directory on the product CD for each SGI Performance Suite software bundle.

- *Message Passing Toolkit (MPT) User Guide*
Describes the industry-standard message passing protocol optimized for SGI computers.
- *MPInside Reference Guide*
Documents the SGI MPInside MPI profiling tool.
- *REACT Real-Time for Linux Programmer's Guide*
Describes the facilities of REACT real-time for Linux that support real-time programs.
- *Performance Co-Pilot 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.
- *Unified Parallel C (UPC) User Guide*
Documents the SGI implementation of the Unified Parallel C (UPC) parallel extension to the C programming language standard.
- *SGI UV GRU Development Kit Programmer Guide*
Documents the SGI UV global reference unit (GRU) development kit. It describes the application program interface (API) that allows direct access to GRU functionality.
- *Linux Resource Administration Guide*
Provides a reference for people who manage the operation of SGI servers and contains information needed in the administration of various system resource management features such as Array Services, CPU memory sets and scheduling, and the Cpuset System.
- *Linux Application Tuning Guide for SGI X86-64 Based Systems*
Provides information about tuning application programs on SGI systems. Application programs include Fortran and C programs written on SGI Linux systems with the compilers provided by Intel.
- *SGI UV Systems Linux Configuration and Operations Guide*
Provides information on how to perform system configuration and operations under the Linux operating system for SGI UV 100, SGI UV 1000, and SGI UV 2000 services systems.

Additional SGI Software Manuals

The following documentation describes current SGI software products:

- *SGI Foundation Software X.X Start Here*
Provides information about the SGI Foundation Software X.X release that supports SGI differentiated server solutions.
- *SGI UV System Management Node Administrator Guide*
Describes the system management node (SMN) installation, configuration and operation for SGI UV 2000, SGI UV 1000 and SGI UV 100 series systems.
- *SGI Management Center for SGI ICE*
Describes how to perform general system installation, configuration, and operations on the SGI ICE 8400 and SGI 8200 systems. For other SGI hardware platforms, see the *SGI Management Center System Administrator Guide* and *SGI Management Center Installation and Configuration*.
- *SGI Management Center for SGI ICE X*
Describes how to perform general system installation, configuration, and operations on the SGI ICE X series systems. For other SGI hardware platforms, see the *SGI Management Center System Administrator Guide* and the *SGI Management Center (SMC) Installation and Configuration* guide.
- *SGI Management Center (SMC) Installation and Configuration*
This guide is intended for system administrators. It describes how to install and configure the SGI Management Center. A companion manual, the *SGI Management Center (SMC) System Administrator Guide*, describes general cluster administration. For SGI ICE platforms, use one of the following manuals: the *SGI Management Center for SGI ICE* (for SGI ICE 8400 or SGI ICE 8200 systems) or the *SGI Management Center for SGI ICE X*.
- *SGI Management Center (SMC) System Administrator Guide*
This guide describes how to monitor and control a cluster using the SGI Management Center. A companion manual, the *SGI Management Center (SMC) Installation and Configuration Guide*, describes how to install and configure the SGI Management Center.
For SGI ICE platforms, use one of the following manuals: the *SGI Management Center for SGI ICE* (for SGI ICE 8400 or SGI ICE 8200 systems) or the *SGI Management Center for SGI ICE X*.
- *SGI Management Center Quick Start Guide*
The guide explains how to deploy your SGI Management Center software.
- *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 Guide*
Describes the configuration and administration of XVM logical volumes using the XVM Volume Manager.
- *CXFS 6 Administration Guide for SGI InfiniteStorage*
Documents CXFS running on a storage area network (SAN). CXFS is clustered XFS, a parallel-access, shared, clustered filesystem for high-performance computing environments.
- *CXFS 6 Client-Only Guide for SGI InfiniteStorage*
Provides an overview of the installation and configuration procedures for CXFS client-only nodes running SGI CXFS clustered filesystems.
- *Guaranteed-Rate I/O Version 2 for Linux Guide*
Provides information about GRIO version 2, the second-generation guaranteed-rate I/O product from SGI.
- *High Availability Extension and SGI InfiniteStorage*
Provides information about creating resources for the high-availability (HA) SGI resource agents that SGI provides for use with the SUSE Linux Enterprise High Availability Extension (HAE) product.
- *DMF 5 Administrator Guide for SGI InfiniteStorage*
Describes how to administer the Data Migration Facility (DMF). DMF is a hierarchical storage management system for SGI environments. Its primary purpose is to preserve the economic value of storage media and stored data. This manual describes DMF components, features, and architecture; it focuses on the tasks of configuration and ongoing administrative duties.
- *SGI InfiniteStorage Appliance Manager User Guide*
Describes the operation of SGI InfiniteStorage Appliance Manager, which is a web-based interface that lets you configure, manage, and monitor a storage server.
- *SGI UV CMC Software User Guide*
Documents the chassis manager controller (CMC) command line interface (CLI) for the SGI UV 2000, SGI UV 1000, and SGI UV 100 systems.
- *Windows on SGI Altix UV Software Install Guide*
Describes how to install Windows Server 2008 R2 Datacenter software on an SGI UV 2000, SGI UV 1000, or SGI UV 100 system. This manual does not apply to SGI UV 20 or SGI UV 10 systems.

SGI Hardware Manuals

The following documentation describes current SGI hardware products:

- *SGI Altix UV 10 System User's Guide*
This guide provides an overview of the SGI UV 10 system architecture and describes the major components. 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 UV 100 System User's Guide*
This guide provides an overview of the SGI UV 100 system architecture and describes the major components. 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 UV 1000 System User's Guide*
This guide provides an overview of the SGI UV 1000 system architecture and describes the major components. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.
- *SGI UV 2000 System User Guide*
This guide provides an overview of the SGI UV 2000 system architecture and describes the major components. 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 ICE 8200 Series System Hardware User's Guide*
This guide describes the features of the SGI ICE 8200 systems and provides operating instructions and general troubleshooting information.
- *SGI Altix ICE 8400 Series System Hardware User's Guide*
This guide describes the features of the SGI ICE 8400 systems and provides operating instructions and general troubleshooting information.
- *SGI ICE X System Hardware User Guide*
This guide describes the features of the SGI ICE X systems and provides operating instructions and general troubleshooting information.
- *SGI Rackable C1104-2TY9 System User's Guide*
This guide describes the features of the SGI Rackable C1104-2TY9 server system as well as troubleshooting, upgrading, and repairing.
- *SGI Rackable C2108-TY11 System User's Guide*

This guide describes the features of the SGI Rackable C2108-TY11 server system as well as troubleshooting, upgrading, and repairing.

- *SGI Rackable C3108-TY11 System User's Guide*
This guide describes the features of the SGI Rackable C3108-TY11 server system as well as troubleshooting, upgrading, and repairing.
- *SGI C1103-TY12 System User's Guide*
This guide provides an overview of the SGI C1103-TY12 system architecture and describes the major components. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.
- *SGI CloudRack X2 System User's Guide*
This guide provides an overview of the SGI CloudRack X2 workgroup cluster server's architecture and describes the major components. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.
- *SGI Origin 400 System User's Guide*
This manual is written for system technicians who are responsible for troubleshooting, upgrading, and repairing SGI Origin 400 systems. This document provides reference information, feature information, and step-by-step instructions for adding and replacing components in the SGI Origin 400 system.

Obtaining SGI Publications

You can obtain SGI documentation in the following ways:

- Visit the SGI Technical Publications Library at the following website:

<http://docs.sgi.com>

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.

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
46600 Landing Parkway
Fremont, CA 94538

SGI values your comments and responds to them promptly.

Release Features

This chapter introduces SGI Performance Suite 1.5 release, describes its contents, and provides documentation information. It covers the following topics:

- “Software Introduction” on page 1
- “SGI Performance Suite Software Bundles” on page 2
- “Finding SGI Performance Suite Manuals” on page 7
- “Getting Updates” on page 8
- “SGI Open Source Information” on page 11
- “Documentation Updated for This Release” on page 11

Software Introduction

The SGI Performance Suite accelerates Linux application performance. It consists of the components described in the following sections:

- “SGI® MPI” on page 2
- “SGI® REACT™” on page 3
- “SGI® UPC” on page 4
- “SGI® Accelerate™” on page 5

SGI supports the SGI Performance Suite 1.5 release on the SUSE Linux Enterprise Server 11 Service Pack 2 (SLES 11 SP2) base distribution, the Red Hat Enterprise Linux Server 6.3 (RHEL 6.3) base distribution, and the RHEL 6.2 base distribution. SGI Performance Suite software (with the exception of SGI REACT) also runs on the Community ENTERprise Operating System 6 (CentOS 6) base distribution. The SGI support team answers questions relating to any of these products, but they do not provide technical support for the CentOS operating system.

For information about SGI support for SGI hardware platforms and operating system configuration settings, see the following URL:

<http://www.sgi.com/servers/>

SGI Performance Suite Software Bundles

This section describes the main contents of the SGI Performance Suite release. Manuals associated with the SGI Performance Suite software can be found at the SGI Technical Publications Library located at the following website:

<http://docs.sgi.com>

SGI® MPI

SGI® MPI is a scalable, high-performance message passing interface (MPI) environment including the following software packages:

- Message Passage Toolkit (MPT) software stack (`sgi-mpt`)
- PerfBoost

SGI PerfBoost uses a wrapper library to run applications compiled against other MPI implementations under the SGI Message Passing Toolkit (MPT) product on SGI platforms.
- `memacct`

Memory accounting utility (see the `memacct(3)` man page).
- `sgi-mpt-shmem`

Contains the SGI SHMEM™ library functions. The SHMEM library functions enable low-latency, high-bandwidth communication in highly parallelized, scalable programs.
- `xpmem`

SGI's NUMALink technology, the XPC kernel module, and the XPMEM kernel module allow users to create a very low latency, very large, shared-memory cluster for optimized use of Message Passing Interface (MPI) software and the SHMEM programming model. The globally addressable, cache coherent, shared memory is exploited by MPI and SHMEM to deliver high performance.

- MPInside

MPInside is a MPI profiling tool which provides valuable information to help MPI application developers to optimize their application.

For more information, see these related manuals:

- *Message Passing Toolkit (MPT) User Guide.*
- *MPInside 3.1 Reference Manual.* This manual ships with the MPInside software module only.
- *Linux Application Tuning Guide for SGI X86-64 Based Systems.*
- *SGI UV Systems Linux Configuration and Operations Guide.* This guide contains more information about SHMEM.

SGI® REACT™

SGI® REACT™ is a real-time performance solution for Linux and includes the following software packages:

- Linux REACT stack
- External interrupt driver
- REACT/Pro Frame Scheduler

The REACT/Pro Frame Scheduler is an execution manager that schedules `pthread`s or processes on one or more CPUs in a predefined cyclic order (see the `frs(3)` man page).

- Kernel barrier (`kbar`) facility

The kernel barrier (`kbar`) facility provides for the fast wakeup of many blocked user threads. When the barrier is signaled, the operating system uses a configurable number of CPUs to quickly wake all the threads blocked on it (see the `kbar(3)` man page).

- `sgi-shield` kernel module

Enables the use of the `cpu_shield` routine that controls timer interrupts on select CPUs.

- User level interrupt facility (see the `uli(3)` man page).

For more information, see the *REACT Real-Time Programmer Guide*.

SGI® UPC

SGI® UPC is an implementation of the Unified Parallel C (UPC) parallel extension to the C programming language standard and includes the following software packages:

- `sgi-upc`
- `sgi-upc-devel`
- `xpmem`

Cross process memory mapping. See `xpmem` descriptions in “SGI® MPI” on page 2.

- `sgi-array-services`

SGI 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.

- `sgi-mpt`

SGI MPT runtime environment. See “SGI® MPI” on page 2.

- `sgi-mpt-shmem`

Contains the SGI SHMEM™ library functions. The SHMEM library functions enable low-latency, high-bandwidth communication in highly parallelized, scalable programs. Also see `xpmem` descriptions in “SGI® MPI” on page 2.

- `libgru`

SGI UV global reference unit (GRU) libraries.

- `superpages`

Note: The `superpages` module is only supported on SGI® UV™ 2000, SGI UV 1000, and SGI UV 100 series systems. For more information on superpages, see “Superpages” on page 14.

- `numatools`

SGI NUMAlink memory placement tools.

- `libcpuset`

The `cpuset` library provides the mechanisms needed to create, destroy and manage cpusets, and to control the placement of processes and threads within cpusets.

- `libbitmask`

This bitmask library supports multi-word bitmask operations for applications programmed in the C programming language. It works in conjunction with Linux kernel support for processor and memory placement on multiprocessor SMP and NUMA systems. The cpuset library depends on this bitmask library.

For more information, see these related manuals:

- *Unified Parallel C (UPC) User Guide*
- *Message Passing Toolkit (MPT) User Guide*
- *Linux Resource Administration Guide*
- *SGI UV GRU Development Kit Programmer Guide*
- *Linux Application Tuning Guide for SGI X86-64 Based Systems*

SGI® Accelerate™

SGI® Accelerate™ includes optimized software libraries and runtime performance tools that increase application performance. It includes the following software packages:

- `cpusets`

The cpuset facility is primarily a workload manager tool permitting a system administrator to restrict the number of processor and memory resources that a process or set of processes may use. A cpuset defines a list of CPUs and memory nodes.

- `numatools`

See the numatools description in “SGI® UPC” on page 4.

- `FFIO`

Cpusets and NUMA tools provide controls for dedicated resource use which in turn helps optimize performance on a system shared by many users. For I/O intensive applications, FFIO helps improve performance without additional coding or re-tooling of the software logic.

- `sgi-pcp`

SGI Performance Co-Pilot UV hub performance counters.

- `libgru development kit`

SGI UV Global Reference Unit (GRU) is part of the SGI UV Hub application-specific integrated circuit (ASIC) Library Development Kit.

- **SGISOLVE**

The PSLDLT libraries under the SGISOLVE umbrella contain two direct solvers, PSLDLT and PSLDU, and one iterative solver, DIterative (deprecated), for solving symmetric, general unsymmetric, double precision, double complex, 32-bit and 64-bit integer, sparse linear systems of equations. These solvers are optimized and parallelized for both IN-CORE and OUT-OF-CORE execution on the SGI SMP platforms.

For more information, see these related manuals:

- *Linux Resource Administration Guide*
- *SGI UV GRU Development Kit Programmer's Guide*
- *Performance Co-Pilot Linux User's and Administrator's Guide*
- *Linux Application Tuning Guide for SGI X86-64 Based Systems*

Software Package License Information

The SGI Performance Suite contains software packages released under a variety of licenses. The licensing terms for all proprietary software packages allow you to install the software on SGI systems only. For licensing information, see the `PACKAGE_LICENSES.txt` file that resides in the `/docs` directory on the product CD.

SGI Performance Suite CDs

The following CDs ship with the SGI Performance Suite 1.5 release:

- SGI® MPI CD
- SGI® REACT® CD
- SGI® UPC CD
- SGI® Accelerate™ CD

You only receive the CDs associated with the software bundles you ordered. If you order the entire SGI Performance Suite, you receive all four CDs.

Release Information and Release Notes

For a complete list of RPMs included in the SGI Performance Suite 1.5 release, see the file named *product-release-rpms.txt* in the */docs* directory on the product CD.

For the latest information about software and documentation in this release, see the release notes file named *product-release-readme.txt* in the */docs* directory on the product CD.

Finding SGI Performance Suite Manuals

The SGI Performance Suite manuals are installed in the following location on your SGI system:

```
/usr/share/doc/packages/sgi-perfsuite-1/
```

The directory for each manual contains a PDF and an HTML copy of the manual. For example:

```
sys:/usr/share/doc/packages/sgi-perfsuite-1 # ls
figures      index.html      LX_UVConfigOps_AG  MPIinside_UG
PS_Start_Here GRUDEV_PG      LX_86_AppTune      LX_Resource_AG  MPT_UG
UPC_UG
```

The following example shows the content of the *PS_Start_Here* directory:

```
/usr/share/doc/packages/sgi-perfsuite-1 # ls
pdf sgi_html
```

You can also access the following local URL on your SGI system using the Mozilla Firefox browser, as follows:

```
# firefox -no-remote file:///usr/share/doc/packages/sgi-perfsuite-1/index.html
```

Note: The SGI documentation is available at SGI Technical Publications Library at <http://docs.sgi.com>.

Getting Updates

The following topics explain how to get software updates for your system:

- “SLES 11 SP2 Updates” on page 8
- “RHEL 6.3 Updates” on page 8
- “SGI Performance Suite Updates” on page 9
- “SGI Foundation Software 2.7 Updates” on page 10

Detailed instructions for downloading and installing updates are available from the SGI product pages on Supportfolio at the following website:

<https://support.sgi.com>

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

SLES 11 SP2 Updates

SLES 11 SP2 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 SLES 11 SP2 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. The **Novell Customer Center Configuration** module has been moved from the Software to the Support menu in YaST.

RHEL 6.3 Updates

Software updates for RHEL 6.3 are provided by Red Hat, Inc. 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 the following location:

<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 RHEL 6 updates.

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

SGI Performance Suite Updates

Customers that received SGI Performance Suite software pre-installed on their new SGI system can log in as the root user and run one of the following commands to add the proper update source. Each product has a unique command.

For example, the following update commands apply to RHEL 6 platforms:

- For SGI MPI run:

```
% /usr/sbin/sgi-mpi-update-source
```
- For SGI REACT run:

```
% /usr/sbin/sgi-react-update-source
```
- For SGI UPC, run:

```
% /usr/sbin/sgi-upc-update-source
```
- For SGI Accelerate run:

```
% /usr/sbin/sgi-accelerate-update-source
```

Each script prompts you for your Supportfolio login and adds the proper SGI update source for the given architecture. The preceding commands apply only to systems running RHEL 6 as the base operating system.

For information about software updates, see the release notes file named `product-release-readme.txt` in the `/docs` directory on the product CD.

For information about how to download and install SGI Performance Suite 1.5 updates, log into Supportfolio and navigate to the SGI Performance Suite product page.

The Supportfolio website is as follows:

<https://support.sgi.com>

The SGI Performance Suite product website is as follows:

https://support.sgi.com/content_request/786853/index.html

SGI Foundation Software 2.7 Updates

If you received SGI Foundation Software 2.7 pre-installed on your new SGI system, you can log in as the root user and run the following command to add the proper update source for SGI Foundation Software 2.7:

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

This script prompts you for your Supportfolio login and adds the proper SGI update source for the given architecture.

For information about how to download and install the SGI Foundation Software 2.7 updates, log into Supportfolio and navigate to the SGI Foundation Software product pages.

The Supportfolio website is as follows:

<https://support.sgi.com>

The SGI Foundation Software product website is as follows:

https://support.sgi.com/content_request/194480/index.html

For more information about the SGI Foundation Software 2.7 release, see the *SGI Foundation Software 2.7 Start Here*.

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 Open SpeedShop) from the following website:

<http://sourceforge.net>

Documentation Updated for This Release

The following manuals have been updated for the SGI Performance Suite 1.5 release:

- *REACT Real-Time for Linux Programmer Guide*, 007-4746-018
- *Message Passing Toolkit (MPT) User Guide*, 007-3773-021.

SGI Performance Suite Software Installation

Your SGI system comes with a base Linux distribution SUSE Linux Enterprise Server 11 Service Pack 2 (SLES 11 SP2) or Red Hat Enterprise Linux Server 6.3 (RHEL 6.3) and SGI Performance Suite software preinstalled.

Installing or Upgrading SGI Performance Suite

Installation instructions for the SGI Performance Suite software reside in the product release notes on the product CD. On the CD, the release notes reside in a file named *product-release-readme.txt* in the */docs* directory.

For example, for the SGI Accelerate 1.5 product, this file is *SGI-Accelerate-1.5-readme.txt*.

Upgrade instructions also reside in the release notes.

Systems with Multiple Ethernet Configurations

When doing a network installation on a system with multiple Ethernet interfaces, there is no guarantee about initial interface names. The first interface discovered is named `eth0`, and so on. This can make network installations confusing. If this is a problem, a workaround is to remove the Ethernet cards that are **not** being used to install software over the network. After the installation is complete, you can reinstall the Ethernet cards. A persistent naming mechanism ensures that an interface, once named, retains its name across system reboots.

XVM Volumes Access in Rescue Mode

You cannot access XVM Volumes in Rescue Mode. It is possible to boot into rescue mode from the operating system 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 because there is no way to access them from rescue mode.

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. If you do not, the `initrd` boot loader panics your system at boot time. This technique is SLES specific.

Superpages

The superpages mechanism on SGI UV systems allows you to map large areas of reserved memory into your application. `libsuperpages` is a library for allocating and mapping SGI UV superpages (see the `libsuperpages(3)` man page).

Superpages are defined in an EFI firmware variable, as a kernel boot option that can only be adjusted by the superuser. When enabled, at boot time, the SGI UV reserves the amount of specified memory to be used via superpages, and it is not available to the operating system for its regular activities. The UV BIOS passes superpages addresses to the superpages driver as an Advanced Configuration and Power Interface (ACPI) device.

The `superpages_define -w` command creates the EFI variable to define superpages. You must be a superuser (**root user**) to execute this command. After execution, to enable superpages, you need to reboot the system. The system BIOS then reserves the system memory for superpages rather than providing it to the operating system. The `superpages_define` command is described in more detail in the examples that follow.

The `superpages` command allows you to test the functionality of an SGI UV system's superpages. For a usage statement, enter `superpages -h`.

The `superpages_define` command is used to configure superpages. For a usage statement, enter `superpages_define -h`.

The `superpages_define -L` option displays valid superpage sizes, as follows:

```
uv-sys:~ # superpages_define -L
Valid superpage sizes:
[1] 0x200000 2M (2097152)
[2] 0x400000 4M (4194304)
[3] 0x1000000 16M (16777216)
[4] 0x2000000 32M (33554432)
[5] 0x8000000 128M (134217728)
[6] 0x20000000 512M (536870912)
[7] 0x40000000 1G (1073741824)
[8] 0x100000000 4G (4294967296)
[9] 0x400000000 16G (17179869184)
[10] 0x1000000000 64G (68719476736)
[11] 0x4000000000 256G (274877906944)
[12] 0x10000000000 1T (1099511627776)
```

The `superpages_define -c` option displays the available memory ranges for superpage allocation, as follows:

```
uv-sys:~ # superpages_define -c
available memory ranges for superpage allocation:
0x100000000-0x87fffffff 0x780000000 30G
0x1000000000-0x17feffffff      0x7ff000000 31G
available memory: 61G (63472M)
```

The following command shows the 4G areas of memory that can be used as superpages:

```
uv-sys:~ # superpages_define -p 4G

in /etc/superpages.conf format:
# for pagesize: 0x100000000 4G
size_4G 4G #node0 0x100000000:3 #node1 0x400000000:4 #node2
0x1000000000:4 #node                3 0x1400000000:3
# 14 4G superpages
```

The following command shows the 4G areas of memory that can be used as superpages, but it limits such areas to two per node:

```
uv-sys:~ # superpages_define -p 4G -n 2

in /etc/superpages.conf format:
# for pagesize: 0x100000000 4G
```

```
size_4G 4G #node0 0x100000000:2 #node1 0x400000000:2 #node2  
0x1000000000:2 #node3 0x1400000000:2  
# 8 4G superpages
```

After you are satisfied with your memory selection, type commands similar to the following to define superpages:

```
uv-sys:~ # superpages_define -p 4G -n 2 > /etc/superpages.conf  
uv-sys:~ # superpages_define -w
```

When you reboot your SGI UV system, those chosen areas are superpages.

You can use huge pages to optimize the performance of your MPI application. For more information, see the *Message Passing Toolkit (MPT) User Guide*.

SLES 11 SP2 or RHEL 6.3 Kernels

For `x86_64` systems, SLES 11 SP2 contains `kernel-default` and `kernel-trace`. The `kernel-default` variant provides multiprocessor support. The SGI Performance Suite provides kernel module packages for both the `default` and `trace` kernel variants.

For `x86_64` systems, RHEL 6.3 has one kernel image, `kernel`, and `kvm` is the only supported virtualization.

Getting the Latest Firmware for Your SGI System

For information on the latest firmware for your SGI system, see the *SGI Foundation Software 2.7 Start Here* available at the following website:

<http://docs.sgi.com>

License Information

Important: Remember to back up your license key files before you install SLES11 SP2, RHEL 6.3, and/or SGI software products. A license key file is not included in these RPMs.

You need to restore your original license key files after you complete your installation.

In terms of software licensing support, the operating system distributors do not ship any licensing code. Independent software vendors (ISVs) are responsible for obtaining and including software licensing components into their products. SGI products ship with SGI licensing software.

For more information, see the specific product release notes or the FAQ at the following website:

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

If you encounter any license-related problems, contact your local support center at the following website:

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

SGI Performance Suite Overview

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

- “SGI Performance Suite 1.5 Products” on page 19
- “Performance Suite Products” on page 22
- “HPC Application Tools and Support” on page 23
- “NUMA Data Placement Tools” on page 24

SGI Performance Suite 1.5 Products

The SGI Performance Suite software release consists of kernel modules for SGI software built against the kernels in SUSE Linux Enterprise Server 11 SP2 or Red Hat Enterprise Linux Server 6.3 and value-add software developed by SGI.

For information about how these products are bundled, see “SGI Performance Suite Software Bundles” on page 2.

Table 3-1 SGI Performance Suite Products

Product	Description
Array Services	<p>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) uses Array Services to launch parallel applications.</p> <p>The secure version of Array Services is uses the secure sockets layer (SSL) and secure shell (SSH).</p> <p>For information about MPI, see the <i>Message Passing Toolkit (MPT) User Guide</i>.</p> <p>For more information about standard Array Services or Secure Array Services (SAS), see the <i>Linux Resource Administration Guide</i>.</p>
Cpuset System	<p>The Cpuset System is primarily a workload manager tool that enables a system administrator to restrict the number of processors and memory resources that a process or set of processes can use. A system administrator can use cpusets to create a division of CPUs and memory resources within a larger system.</p> <p>For more information, see the <i>Linux Resource Administration Guide</i>.</p>
IOC4 serial driver	<p>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:</p> <pre>/dev/ttyIOC4/0 /dev/ttyIOC4/1 /dev/ttyIOC4/2 /dev/ttyIOC4/3</pre>
Kernel partitioning support	<p>Provides the software infrastructure necessary to support a partitioned system, including cross-partition communication support. Partitioning is supported on SGI UV 2000, SGI UV 1000, and SGI UV 100 systems only.</p> <p>For more information about system partitioning, see the <i>SGI UV Linux Configuration and Operations Guide</i>.</p>

Table 3-1 SGI Performance Suite Products (**continued**)

Product	Description
MPT	Provides industry-standard message passing libraries optimized for SGI computers. For more information, see the <i>Message Passing Toolkit (MPT) User Guide</i> .
NUMA tools	Provides a collection of NUMA related tools that includes <code>dlook(1)</code> , <code>dplace(1)</code> , and others. For more information, see the <i>Linux Application Tuning Guide for SGI X86-64 Based Systems</i> .
Performance Co-Pilot collector infrastructure	Provides performance monitoring and performance management services targeted at large, complex systems.
REACT real-time for Linux	Support for real-time programs. For more information, see the <i>REACT Real-Time for Linux Programmer Guide</i> .
Utilities	<code>udev_xsci</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.

SGI does not support the following:

- Base Linux software not released by Novell for SLES 11 SP2 or by Red Hat for RHEL 6.3 or other software not released by SGI.
- SGI does not provide technical support for the CentOS 6 operating system.
- Other releases, updates, or patches not released by Novell for SLES 11 SP2, Red Hat for RHEL 6.3 or by SGI for SGI Performance Suite 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 or Red Hat and SGI.
- Unsupported hardware configurations and devices.

Performance Suite Products

SGI Performance Suite provides application acceleration components for software developers and end-users. SGI Accelerate, SGI MPI, SGI REACT, and SGI UPC contain libraries and tools that enable software developers to develop, profile, and tune applications for faster performance. End users benefit from running their applications with the runtime acceleration tools supplied in SGI Accelerate and SGI MPI. The following sections describe some key components.

Cpuset Support

The cpuset facility is primarily a workload manager tool that permits 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 can execute only on the CPUs in that cpuset and can 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 system improves cache locality and memory access times and can substantially improve an application's 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 more information about cpusets, see the *Linux Resource Administration Guide* and the *Linux Application Tuning Guide for SGI X86-64 Based Systems*.

Partitioning

SGI provides the ability to divide a single SGI UV 100, SGI UV 1000, or SGI UV 2000 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 systems to provide high-performance clusters.

For further documentation and details on partitioning, see the *SGI UV Systems Linux Configuration and Operations Guide*.

HPC Application Tools and Support

SGI has ported HPC libraries, tools, and software packages from IRIX to Linux to provide a powerful, standards-based system using Linux and Xeon-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 that are 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 about the Message Passing Toolkit, see the *Message Passing Toolkit (MPT) User Guide*.

Performance Co-Pilot

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

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 tune applications on your system.

For information about performance tuning information for single processor and multiprocessor programs, see the *Linux Application Tuning Guide for SGI X86-64 Based Systems*.

`dlook` Command

The `dlook(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 run on each physical CPU in the system is also printed.

dplace Command

The `dplace(1)` command binds a related set of processes to specific CPUs or nodes to prevent process migrations. In some cases, this tool improves performance because of the occurrence of a higher percentage of memory accesses to the local node.

taskset Command

The `taskset(1)` command sets or retrieves the CPU affinity of a running process (given its PID) or launches a new command with a given CPU affinity. *CPU affinity* is a scheduler property that bonds a process to a given set of CPUs on the system. The Linux scheduler honors the given CPU affinity, and the process does not run on any other CPUs. Note that the Linux scheduler also supports natural CPU affinity; the scheduler attempts to keep processes on the same CPU as long as practical for performance reasons. Therefore, forcing a specific CPU affinity is useful only in certain applications.

For more information about NUMA tools, see the *Linux Application Tuning Guide for SGI X86-64 Based Systems*.

Index

A

Array Services, description, 20

C

CD contents, 2

Cpuset system, 22

D

dlook command, 24

dplace command, 25

G

Getting software updates, 8

H

hardware platforms, 2

HPC support

libraries and tools, 23

Message Passing Toolkit, 23

Performance Co-Pilot, 24

L

License manager updates, 17

M

Message Passing Toolkit
support, 23

N

NUMA tools, 24

dlook command, 24

dplace command, 25

O

Online Support Services
Supportfolio, xi

Open source website, 11

OS configuration settings, 2

OS enhancements

cpuset support, 22

for HPC environments, 22

partitioning, 23

P

Performance Co-Pilot support, 24

Product list, 19

Product support, 19

R

Release features, 1

RHEL updates, 8

S

SGI Foundation 2 updates, 10

SLES11 updates, 8

Software

 installation, 13

 introduction, 1

Software license information, 6

Supportfolio, xi

T

taskset command, 25

U

Unsupported elements, 21

X

XVM Volumes in Rescue Mode, 14