

SGI® Performance Suite

Extending High Performance Technical Computing for Large Scale Servers and Clusters

Key Features

SGI® Accelerate

SGI® MPI

SGI® REACT™



Overview

SGI Performance Suite takes Linux® performance software to the next level. While hardware and processor technology continue to scale, managing software performance has become increasingly complex. SGI® continues to extend technical computing performance for large scale servers and clusters.

SGI Performance Suite is designed to accelerate the performance of SGI® UV™, SGI® ICE™, and SGI® Rackable® servers. Some examples of the numerous benefits of the tools in SGI Performance Suite include:

- The SGI Message Passing Toolkit MPI library helped SGI ICE achieve world-record performance in the Standard Performance Evaluation Corporation's (SPEC) MPI2007 benchmark tests for high performance technical computing. The benchmarks measure performance on real-world applications in a variety of fields including computational fluid dynamics, electromagnetics, geophysics, ray tracing, and hydrodynamics.
- Multi-rail InfiniBand™ support which enables MPI and SHMEM traffic to run on multiple InfiniBand fabrics.
- NUMA Tools data placement tool enabled SGI UV to achieve:
 - World-record SPEC OMP (SPECCompG_base2012)
 - World-record SPEC CPU rate (SPECint_rate_base2006 and SPECfp_rate_base2006) performance
 - Top x86-64 STREAM result

Key Features and Benefits

SGI Performance Suite provides application acceleration components for software developers and end-users. SGI® Accelerate, SGI® MPI and SGI® REACT™ 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.

SGI®
Accelerate

Accelerate applications with optimized software libraries and tools

SGI® MPI

SGI's scalable, high-performance MPI environment

SGI®
REACT™

Hard real-time performance solutions for standard distribution Linux



SGI Accelerate – Optimizes Application Performance

SGI Accelerate helps accelerate application performance through tools that tune applications at runtime without recompiling, and through libraries which optimize performance with specialized algorithms.

- **Cpusets** – Enables a system administrator to restrict the number of processor and memory resources that a process, or set of processes can use.
- **NUMA Tools** – Enables data placement to specific memory locations to limit communications overhead.
- **Linkless Flexible File I/O** – Improves performance for I/O intensive applications without recompiling or retooling the software logic.
- **SGI UV GRU Development Kit** – allows developers to take advantage of the hardware acceleration features of the Global Reference Unit (GRU) within the SGI UV Hub.
- **SGISOLVE** – A collection of solvers optimized for in-core and out-of-core execution on SGI systems.

SGI MPI – High Performance MPI Environment

SGI MPI contains complementary tools enabling application acceleration, with SGI Message Passing Toolkit as the core MPI performance engine.

SGI MPI includes:

- **SGI Message Passing Toolkit** – Provides support for the MPI-3.1 and OpenSHMEM-1.3 standards. It runs on all SGI platforms.
- **SGI MPI PerfBoost** – A tool which enables applications compiled with MPI libraries other than SGI MPI to utilize the SGI MPI engine without relinking or recompiling. SGI has seen up to 70% performance boosts for certain applications.
- **Perfcatcher** – A profiling tool for SGI MPI and OpenSHMEM.
- **MPInside** – A profiling and performance analysis tool that captures communication analytics, and pinpoints bottlenecks and load imbalances during the application run. MPInside works with multiple MPI libraries include SGI MPI, OpenMPI, Intel MPI, and IBM Platform MPI.
- **MPIplace** – A runtime tool which utilizes input from MPInside and the job scheduler in order to optimize process and data placement onto a set of nodes and CPUs for an application’s particular communication pattern based on an earlier run.
- **Checkpoint/Restart** – Provides job management control to SGI MPI applications, by allowing the user to stop a job and restart it from the checkpoint later.

- **Array Services** – Accelerates SGI MPI job launches, enabling launches across thousands of nodes in seconds.
- **Workload Scheduler Integration** – Supported by Altair PBS Professional, Adaptive Computing Moab, SLURM, IBM Platform LSF, and Univa Grid Engine.
- **Check MPT** – A Nagios plugin enabling administrators to see status update messages including error messages from different SGI MPI jobs. Check MPT works with the SGI Management Center 3 Nagios interface.

SGI REACT – Hard Real-Time Performance

SGI REACT is the only hard real-time performance solution for standard Linux distributions, where no special custom kernels are required. SGI has successfully achieved a real-time latency guarantee of 30 µsec for 2-256 core systems.

SGI REACT includes:

- **Linux Trace Tool Next Generation** – Provides profiling information at the Linux kernel level in order provide trace information for analyzing the impact of kernel operations on application performance.
- **REACT** – Allows for easy generation and configuration of a real-time system.
- **React Frame Rate Scheduler** – An execution manager that schedules processes on 1 or more CPUs.
- **SGI Linux Trace Debugger** – A debugger for real-time applications.

SGI Performance Suite – System Requirements

Systems	SGI UV, SGI ICE, SGI Rackable
Operating Systems	Red Hat Enterprise Linux 6, 7 CentOS 6, 7 SUSE Linux Enterprise Server 11, 12
Interconnects	InfiniBand™ Ethernet Intel® Omni-Path Architecture

About SGI

SGI is a global leader in high performance solutions for compute, data analytics and data management that enable customers to accelerate time to discovery, innovation, and profitability.

For More Information

Please contact an SGI sales representative at 1-800-800-7441 or visit www.sgi.com.

Global Sales and Support: sgi.com

