Datasheet

Linux® Software for the SGI® Altix™ 3000 Family

Industry-Standard 64-Bit Linux with SGI ProPack™ Extensions: Breakthrough Scalability and Performance for Servers and Superclusters

Features

- Industry-leading scalability: Powers world's first 64p node and first shared-memory supercluster
- Designed for HPC: Solve compute- and data-intensive problems quickly and effectively
- Industry-standard 64-bit Linux distribution: Binary compatible with existing Linux software

Linux offerings from SGI are specifically designed for demanding high-performance computing problems. These industry-standard, fully supported software solutions leverage SGI's extensive experience in scaling, data-handling, and large-system management. With SGI, you get the flexibility of Linux backed by 20 years of high-performance computing research and development, all focused on solving your toughest problems quickly and effectively.

The SGI Altix 3000 family of servers and superclusters scales to single-OS nodes of up to 64 processors and shared-memory superclusters of thousands of processors. The key to scaling Linux effectively to these levels is focusing on achieving good performance on real-world applications. SGI engages in an exhaustive iterative process of testing, performance measurement, and tuning of demanding application codes, using many of the same tools and technologies SGI applied to scaling IRIX systems over the years.

SGI has long been a supporter of the open-source community and is committed to developing the performance functionality, and reliability of Linux. With its focus on high-productivity computing environments, SGI has contributed to Linux scalability, scheduling, memory usage, I/O, and other efforts critical to high-demand application performance. These investments have resulted in many enhancements and features in today's standard Linux distributions, along with additional tools to maximize computing productivity.

Enhanced Operating System for HPC

The Linux operating environment from SGI gives you all the tools you need to solve compute- and data-intensive problems easily and cost effectively, while maintaining compatibility with currently available Linux software.

Industry standards: The Linux community offers a broad selection of applications and tools to solve all kinds of business, technical, and systems administration problems. Linux software for SGI® systems is designed and configured to ensure compliance with industry standards, so that you can take advantage of the wealth of software available.

Binary compatible with Red Hat® Linux®, SGI systems will run commercial Linux applications, standards-compatible open-source software, and your own cluster codes.

High-performance NUMA support: SGI has 20 years of experience in designing, scaling, and tuning high-performance systems specifically for scientific and technical applications. The SGI Altix 3000 family represents the third-generation SGI® NUMA architecture, and SGI has leveraged its expertise to ensure that Linux can take full advantage of its flexibility and performance.

Global memory access: To handle large compute problems effectively, intelligent data handling is at least as important as fast processors; in fact, in many applications, processors can spend more time waiting for data than computing. Linux software from SGI supports 64-bit addressing and large global shared memory, allowing users to run very large complex solver models entirely in memory. The SGI Altix 3000 family also supports powerful clustering capabilities, including the unique ability to address global shared memory across cluster nodes without performance penalties. Common functions (e.g., sharing a data array between two nodes) that require multiple I/O operations on a traditional cluster are performed with a single memory operation on the SGI Altix 3000 supercluster.

XFS™: XFS is a journaled 64-bit filesystem designed to support very large files and filesystems without performance degradation. First introduced by SGI on the IRIX® operating system in 1993, XFS was ported to Linux and released into the open-source community in 2000.

Powerful Middleware for High-Productivity Computing

To supplement the standard Linux distribution, SGI provides a robust set of tools and features specifically designed to maximize performance on demanding scientific and technical applications.

Resource management: Maximizing performance on compute and data-intensive applications requires more than just fast hardware. Resource management tools can help deliver excellent, repeatable performance on real-world applications by allowing for the control and tuning of processors and memory. By increasing the system's productivity, you get to solutions faster at the lowest possible cost.

- Message Passing Toolkit (MPT): Effectively utilizing resources in a large multiprocessor system can be a complex undertaking. SGI's Message Passing Toolkit is an optimized set of the MPI and SHMEM parallel programming libraries, tuned to give your application access to the full power of the SGI Altix 3000 family architecture.
- CPU sets and memory placement: These features enable system services and applications to specify on which CPUs they may be scheduled, and from which nodes they may allocate memory. This gives users maximum flexibility in resource allocation and can help deliver fast repeatable run times on mission-critical work.
- Cluster software (Array Services): The Array Services software package contains a library, a system daemon, and a set of commands that enable developers to define and administer cluster configurations and manage the set of jobs running on the cluster.
- Performance Co-Pilot™: Performance Co-Pilot software tools track performance at the system-resource level to help identify potential areas for efficiency optimization. Invaluable to SGI engineers during the Linux scaling process, this tool can provide enormous benefit to system administrators and programmers who need optimal performance.

Data management: As processor performance increases, system bottlenecks tend to move to other system resources. To deliver good performance, real-world, data-intensive applications require intelligent management of data and I/O. These features have been designed to maximize I/O performance, robustness, and flexibility in a cost-effective way.

- XSCSI: XSCSI is the SCSI midlayer that enables sustained I/O in excess of 7GB per second on the SGI Altix 3000 systems, solving a common Linux bottleneck in support of demanding HPC workloads.
- CXFS™: CXFS is a shared filesystem for high-performance computing environments. It provides data sharing over a storage area network (SAN), allowing multiple computers simultaneous direct access to a common shared filesystem with local filesystem performance. Multiple systems share a single data file, and a single copy of the data is maintained. This saves disk space, eliminates the need for lengthy network-based file transfers, and reduces version-control problems. CXFS clients are available for systems running the Linux, IRIX, Windows NT®, Windows® 2000 and Solaris™ operating systems.
- XVM: XVM is an enhanced logical volume manager designed for production I/O workloads, easily handling issues related to very large files and rapidly changing requirements. It supports performance and redundancy features such as striping, mirroring, and concatenation, and allows you to adjust volume configurations on the fly. Volumes configured with redundant paths can automatically reroute disk requests in the event of component failure, keeping your mission-critical application up and running.

- Hierarchical storage management tools: The volume of data available to a typical supercomputing installation is increasing exponentially, and managing the data and the associated storage resources is becoming an increasingly complex problem. The SGI® Data Migration Facility (DMF) and Tape Management Facility (TMF) automatically migrate infrequently used data to less-expensive, near-line storage systems according to user-configurable management rules. Files are retrieved transparently, and always appear as local regardless of media location. Currently deployed on IRIX systems in hundreds of large data centers, DMF and TMF are now available to Linux users who require high-performance, efficient use of storage resources.

  System management: These flexible system management tools help you use resources efficiently and without interruption, maximizing productivity and availability and minimizing total cost of ownership.

- Partitioning: Using partitions, system managers can effectively divide and isolate resources to maximize resilience and eliminate single points of failure. The ability to access large shared memory across partitions is a unique and powerful capability of the SGI Altix 3000 family architecture.

- Linux FailSafe™: A powerful high-availability solution for mission-critical applications, FailSafe provides a simple way to achieve highly available applications without recoding, recompiling, or buying expensive extra hardware. With FailSafe, you get a robust clustering environment with resilience from any single point of failure.

- Comprehensive System Accounting (CSA): Large systems in multiuser environments can have complicated management and funding requirements. CSA provides methods for collecting per-process resource usage data, monitoring disk usage, and charging fees to specific login accounts according to configurable parameters.

Development Tools

Powerful development tools from SGI, the open source community, Intel, and other third parties help you get optimal performance and flexibility from the SGI Altix 3000 family, without a complex or cumbersome development process. Intel® Itanium® Processor Family compilers for C, C++, and Fortran 95 are available from SGI and Intel, and GNU C and Fortran 77 alternatives are also available. SGI, third parties, and the open source community offer a wide variety of libraries, debuggers, and performance analysis tools to help you build the best possible solutions.

Support and Services

With SGI, you get a complete, fully supported solution designed for high-productivity computing environments. Unlike most vendors, SGI supports all of the software shipped with its Linux systems, and works closely with third parties and with the open source community to ensure the availability of a rich and robust software offering. SGI also offers appropriate services to implement and integrate Linux applications in your environment. For more information on available services, please see *www.sgi.com/support*. For a useful list of Linux software, please see *www.sgi.com/go/linux/dir*.

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America (52) 5267-1387
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 6771.0290

3434 (06/30/2003)

J14327