SGI ICE X

World Record Cluster Performance

Highlights

The world's fastest supercomputer just got faster

The world-renowned SGI quality and performance you love

Flexible to fit your workload

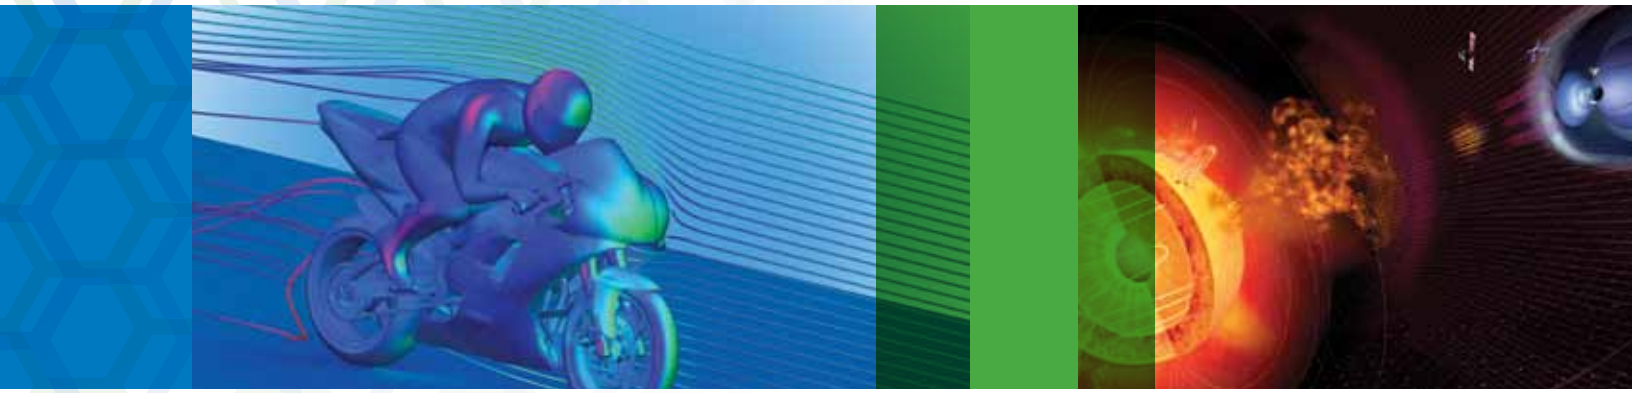

Petascale-Class, With Clear Roadmap to Exascale Performance

SGI ICE X is the next generation of the world's fastest distributed memory platform for over three years running. This performance leadership is proven time and again, not just in the lab, but at customer sites including the largest and fastest pure compute InfiniBand cluster in the world.

The system combines the powerful Intel® Xeon® processor E5-2600 family platform with a unique board and interconnect design. Running on standard Linux®, SGI ICE X delivers up to 2,304 processor cores — over 53 teraflops of compute power — per rack, easily scaling from 36 to tens of thousands of nodes to address the world's most challenging compute problems.

With its innovative design, SGI ICE X provides nearly a 2.5x increase in density, saving precious space in your data center. This results in a nearly 5x improvement in performance density over the previous generation, making your applications fly.

SGI ICE X is designed to minimize system overhead and communication bottlenecks that can rob efficiency and scalability. For example, the system offers the highest performance and most scalable system for CFD, and is now optimized for SGI OpenFOAM®.

Installs Production-Ready in Hours or Days, Not Weeks or Months

Power up and go, at scale, to solve even the most data-intensive problems.

Built entirely on industry-standard hardware and software components, SGI ICE X enables access to the full spectrum of the Linux ecosystem.

SGI Performance Suite optimizes the performance of Linux applications. And reliability, availability and serviceability are ensured with SGI Management Suite, enabling ease of administration, including power management, which is all-important in today's green environment.

Ultimate Flexibility and Seamless Scalability

SGI ICE X is the only system in its class offering expandability within and across technology generations while maintaining uninterrupted production workflow, and is the only platform capable of seamless scalability from tens of teraflops to tens of petaflops. User choice is available in topology and interconnect, power, cooling, CPUs and memory. The integrated bladed design offers rack-level redundant power and cooling via air, warm water or cold water, for enhanced reliability and availability. The result is a system with unmatched efficiency, performance and overall value.

Also available with storage and visualization options, SGI ICE X sets a new standard for flexibility, simplicity and ease-of-use in scale-out computing. Along with industry-leading professional services and support, SGI enables customers to efficiently overcome the challenges of complex data intensive workflows, leading to accelerated results.

Configuration Specifications

sgi.com/icex

Compute Blades	IP-113	IP-115
Processors	• Intel® Xeon® Processor E5-2600 Family	• Intel® Xeon® Processor E5-2600 Family
Memory	• 16 DDR3 DIMM slots (8 per CPU socket) • 4, 8, 16 and 32GB 1600 MT/s ECC Registered DIMMs	• 16 DDR3 DIMM slots per twin blade (4 per CPU socket) • 4, 8, 16 and 32GB 1600 MT/s ECC Registered DIMMs
Storage	• Two 2.5" SATA drives (HDD or SSD) per blade	• Two 2.5" SATA drives (HDD or SSD) per twin blade; one per logical node
FDR IB Mezzanine Card	• Single Port, Dual Port or Dual Single Port	• Dual Single Port
Cooling	• Traditional Heat Sinks	• SGI Cold Sinks
Topology Options	• Single or dual plane all-to-all, fat tree, hypercube and enhanced hypercube	• Single plane all-to-all, fat tree, hypercube and enhanced hypercube

Blade Enclosures	Standard	Premium
Integrated Switch	• Single 36 port FDR IB ASIC with 18 ports external	• Dual 36 port FDR IB ASIC with 48 ports external
Administrative Network	• Dedicated GigE network (redundancy optional), chassis management controller	

Racks	D-Rack (For IP-113 Compute Blade)	M-Rack (For IP-115 Compute Blades)
Dimensions	• 24"W x 40"D (42U standard) • Can optionally be extended to 48U	28"W x 40"D (42U standard)
Blade Enclosure Support	• Up to two blade enclosure pairs (72 total blade slots)	• Up to two blade enclosure pairs (72 total blade slots)
Power	• Up to 5+1 redundant 3000W power supplies per blade enclosure pair	• Up to 11+1 redundant 3000W power supplies per blade enclosure pair
Cooling	• Open-looped airflow or water (optional)	• Closed-loop airflow with warm water support capability

Storage		
InfiniteStorage InfiniBand Solutions	• High performance shared file systems • IP over InfiniBand	• Native InfiniBand block level access • Native InfiniBand SAN supported with CXFS

Hierarchical System Management	Tier 1: System Administration Controller	Tier 2: Rack Leader Controller (RLC)	Tier 3: Chassis Management Controller	Tier 4: Baseboard Management Controller
	• One per SGI ICE system • Provisions out software to RLC • Pulls aggregated cluster management data from RLC • Utilizes C1104-RP7 "6017" system	• One per two-blade enclosure pair • Holds blade boot images • Runs fabric management software • Aggregates cluster management data for rack • Utilizes C1104-RP7 "6017" system	• Two or four per blade enclosure pair • Controls master power to all compute nodes • Monitors power and blade enclosure environment	• One per compute node • Controls board-level hardware • Monitors compute node environment

Service Node Options	• Login Node • Gateway Node • Batch Node • Storage Node	• OSS Node • MDS Node	Service nodes can be optionally configured with: • GPUs such as NVIDIA® Quadro® FX, NVIDIA Quadro and NVIDIA Tesla® • Hard Disk Drives (SAS and/or SATA) • I/O cards (various)

Software Development	
Programming Languages and Debuggers	• C & C++: Intel® C++ Compiler, GNU GCC • Debuggers: Intel® Debugger included with Intel® compilers, GNU GDB, Rogue Wave Software® TotalView® Team, Allinea DDT, Intel® Inspector XEC • Fortran: Intel® Fortran Compilers (Fortran 95), GNU GCC (Fortran77) • Performance Analysis: Intel® VTune Amplifier XE, Intel® Trace Analyzer & Collector
Libraries	• SGI MPI • OpenMP included with Intel® compilers • Intel® Math Kernel Library • Intel® Parallel Building Blocks • Intel® Integrated Performance Primitives • Intel® MPI Library

System Software	
Operating Systems	• SUSE® Linux Enterprise Server 11 • Red Hat® Enterprise Linux 6
Cluster Solution Stack	• SGI Foundation Software 2: Optimized drivers and system monitoring • SGI Management Suite: Cluster management software • SGI Performance Suite optimized application performance package consisting of SGI Accelerate, SGI MPI, SGI REACT and SGI UPC • Altair® PBS Professional™: Job scheduling and workload management

Global Sales and Support: sgi.com/global