SGI® ICE™ X

One of the World's Fastest Commercial Supercomputers

Highlights

191 Pure Compute Teraflops per Rack

Scale from 36 to Thousands of Nodes

Streamlined, Flexible Architecture

Highly Space and Energy Efficient

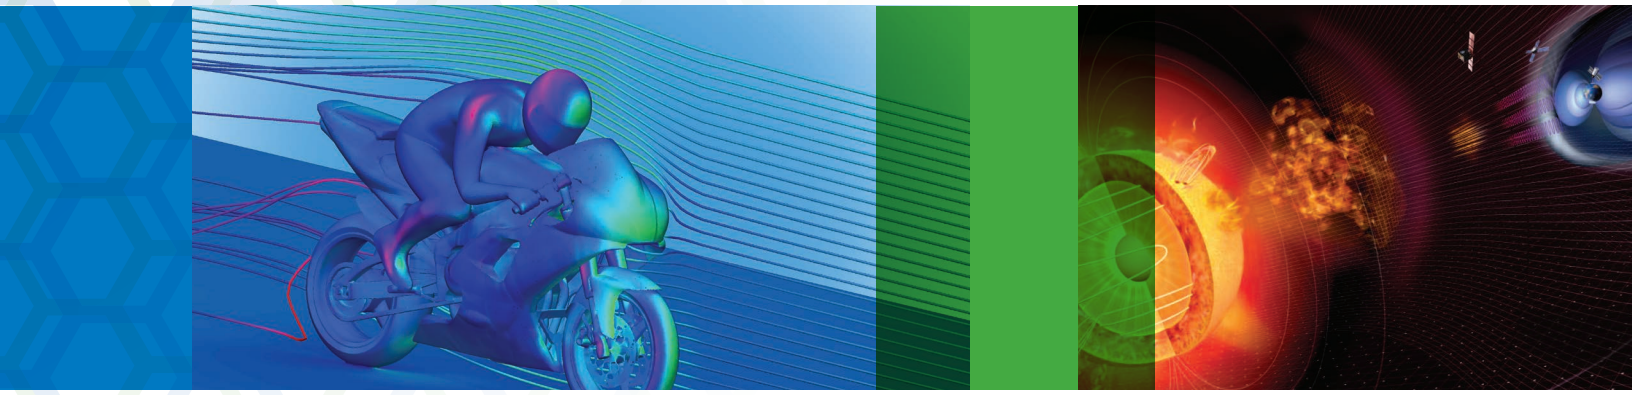

Petascale Performance Leadership

SGI's ICE X system is the latest generation of the record holding ICE family. This new addition continues the ICE legacy by making it the world's fastest distributed memory supercomputer for over 5 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 can be configured with compute nodes comprising of Intel® Xeon® processor E5-2600 v3 series exclusively or with compute nodes comprising of both Intel® Xeon® processor E5-2600 v3 series and Intel® Xeon Phi™ coprocessors. Running on standard Linux®, the SGI ICE X can deliver over 191 teraflops per rack and scale from 36 to tens of thousands of nodes to solve the world's most challenging compute problems.

SGI's ICE X is designed to minimize system overhead and communication bottlenecks that can rob efficiency and scalability. The system, for example, offers the highest performance and scalability for computer fluid dynamics (CFD), demonstrating up to 3,072 cores processing in parallel.

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

Power up and go, at scale, in less time and with less effort.

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

SGI Performance Suite optimizes the performance of Linux applications while SGI Foundation Suite optimizes drivers and system monitoring. SGI Management Suite helps ease administration, including all-important power management, and maintain high reliability, availability and serviceability.

Ultimate Flexibility and Seamless Scalability

The SGI ICE X is the only system in its class offering expandability within and across technology generations while maintaining uninterrupted production workflow. It is the only platform capable of seamless scalability from tens of teraflops to tens of petaflops. Clusters can be architected in a variety of topologies with choice of switch and single or dual plane InfiniBand blade connection.The integrated switch blade 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, the SGI ICE X system 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 and accelerate results.

Configuration Specifications

sgi.com/icex

Compute Blades	IP-131	IP-133	IP-119
Processors/Coprocessors	• Intel® Xeon® Processor E5-2600 v3 Series	• Intel® Xeon® Processor E5-2600 v3 Series	• Intel® Xeon® Processor E5-2600v2 Series • Intel® Xeon Phi™ Coprocessor
Memory	• 16 DDR4 DIMM slots (8 per CPU socket) • 8, 16 and 32GB 2133 MT/s ECC Registered DIMMs	• 16 DDR4 DIMM slots per twin blade (4 per CPU socket) • 8,16 and 32GB 2133 MT/s ECC Registered DIMMs	• 8 DDR3 DIMM slots (4 per CPU socket) • 4, 8, 16 and 32GB 1866 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	• Two 2.5" SATA drives (HDD or SSD) per blade
FDR IB Mezzanine Card	• Single port, dual port or dual single port	• Dual single port	• Single port, dual port, or dual single port
Cooling	• Traditional heat sinks	• SGI cold 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	• Single or dual 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-131 Compute Blade)	M-Rack (For IP-133 and 119 Compute Blades)
Dimensions	• 24"W x 49.5"D, air cooled • 24"W x 54.75"D, water cooled (42U Standard) • Can optionally be extended to 48U	• 32"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 17+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

*Use of turbo mode for extended periods of time may result in power not being N+1 depending on the configutarion, application mix and data center environment.

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. Visit sgi.com for more information.

Global Sales and Support: sgi.com/global