

SGI® ICE™ X

One of the World's Fastest
Commercial Supercomputers


Highlights

223 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 architecture is the basis for the world's fastest distributed memory supercomputers 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 v4 series exclusively or with compute nodes comprising of both Intel® Xeon® processor E5-2600 v4 series and Intel® Xeon Phi™ coprocessors. Running on standard Linux®, the SGI ICE X can deliver over 223 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 computational 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 Software optimizes drivers and system monitoring. SGI Management Suite helps ease administration, including all-important power management, to maintain high reliability, availability, and serviceability.

Ultimate Flexibility and Seamless Scalability

The SGI ICE platforms are the only systems in their class offering expandability within and across technology generations while maintaining an 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 set 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.



sgi

Configuration Specifications

Compute Blades		IP-131		IP-133	
Processors/Coprocessors		• Intel® Xeon® Processor E5-2600 v4 Series		• Intel® Xeon® Processor E5-2600 v4 Series	
Memory		• 16 DDR4 DIMM slots (8 per CPU socket) • 8, 16, 32, and 64GB 2400 MT/s ECC Registered DIMMs		• 16 DDR4 DIMM slots per twin blade (4 per CPU socket) • 8, 16, 32, and 64GB 2400 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-131 Compute Blade)		M-Rack (For IP-133 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		• High performance shared file systems		• Native InfiniBand block level access	
InfiniBand Solutions		• IP over InfiniBand		• Native InfiniBand SAN supported with CXFS™	
Hierarchical System Management		Tier 1: System Administration Controller • One per SGI ICE system • Provisions out software to RLC • Pulls aggregated cluster management data from RLC • Utilizes C1104-RP7 "6017" system		Tier 2: Rack Leader Controller (RLC) • 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	
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, PGI® PGC:++® GNU GCC • Debuggers: Intel® Debugger included with Intel® compilers, PGI® PGDBG®, Rogue Wave Software® TotalView® Team, Allinea DDT, Allinea Forge, GNU GDB • Fortran: Intel® Fortran Compilers, PGI® PGFORTRAN® GNU GCC • Performance Analysis: Intel® VTune Amplifier XE, Intel® Trace Analyzer & Collector, PGI® PGPROF®			
Libraries		• SGI MPI • OpenMP included with Intel® compilers • Intel® Math Kernel Library • Intel® Parallel Building Blocks • Intel® Integrated Performance Primitives • Intel® MPI Library			

*Use of turbo mode for extended periods of time may result in power not being N+1 depending on the configuration, 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

©2014-2016 Silicon Graphics International Corp. All rights reserved. SGI, the SGI logo, ICE, CXFS and OpenFOAM are registered trademarks or trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries. Intel, Xeon and the Intel Xeon logo are registered trademarks of Intel Corporation. All other trademarks are property of their respective holders. Specifications subject to change without notice. 10092013 4330 16032016

