

# SGI® ICE™ XA

## The World's Most Powerful Distributed-Memory Supercomputer

### Key Features

223 Pure Compute Teraflops per Rack

Scale from Hundreds to Tens  
of Thousands of Nodes

Industry-leading Power & Cooling Efficiency

Optimum Blade, Infiniband and Intel®  
Omni-Path Topology, and Linux® OS Flexibility



In today's world, High Performance Computing (HPC) is an essential technology for government institutions, universities, and commercial enterprises to solve complex problems in areas ranging from life, earth, and space sciences, to engineering and manufacturing, to national security.

SGI is a global leader in HPC and has been equipping customers with advanced, innovative solutions for over 20 years. SGI ICE is SGI's flagship scale-out platform combining speed, scale, and efficiency, to provide the most powerful distributed-memory supercomputer in the world.

### Key Features

To address the increasing size and volume of HPC workloads, SGI has developed a 6th generation supercomputer, SGI ICE XA. Extending SGI's technology leadership, ICE XA enables scientists, researchers, and engineers to achieve and accelerate computational breakthroughs at petascale and with a high return on investment.

- Fully leverage leading commercial applications as well as in-house developed codes using industry-standard, state-of-the-art x86 architectures, Linux®, and integrated InfiniBand and Intel Omni-Path technologies.
- Enjoy optimum flexibility in processor and accelerator mix, blade and storage configuration, and network topology to meet specific workload requirements.
- Reduce operational costs through superior power and cooling efficiency coupled with advanced SGI software for system health and power management.

- Power up and be production-ready in hours to days, not weeks to months.
- Grow existing ICE clusters by adding next generation nodes and without user interruption.
- Capitalize on SGI's path to Exascale and continuing innovation to solve Terascale and Petascale problems.

All ICE XA supercomputers are fully integrated and tested prior to leaving SGI's U.S.-based factory. And for solution designs which include Lustre-based storage and intelligent data management, installation, and system support with 24x7 remote monitoring, SGI Professional and Support Service teams are truly the best in the industry.

### Speed

SGI ICE XA, based on compute nodes featuring the latest Intel® Xeon® E5-2600 v4 and Intel Xeon Phi™ processors, is designed to run complex HPC workloads at petaflop speed. Successive SPEC MPI2007 world records reflect SGI's expertise in maximizing system performance of Xeon based cluster platforms. SGI Application Engineers have extensive expertise tuning systems to specific workload environments. What matters most is that SGI's performance leadership extends from the laboratory to customer sites, enabling some of the most powerful commercial supercomputing solutions in the world.



### Advanced HPC Software Accelerates Workloads

Behind the blazing speed of SGI ICE XA lies innovative hardware coupled with advanced HPC software.

- **SGI Performance Suite** provides a variety of application acceleration components including specialized libraries, plus a high performance MPI environment.
- **SGI Development Suite** provides an advanced software environment for developing, debugging, and analyzing performance of technical computing applications.
- **SGI Management Suite** is a comprehensive suite of tools for high speed system provisioning, proactive health management, and power resource management at a node level.

### Scale

#### Grow Seamlessly with Extreme Density

SGI ICE XA can provide up to 223 teraflops of pure compute per rack and grow to tens of thousands of nodes with minimal increase in system overhead.

#### Blade, InfiniBand Topology, and OS Flexibility

The highly differentiated architecture of SGI ICE XA provides optimum flexibility in system configuration.

- **Compute Blade:** Choose from four blade types. The IP-125CS is a quad node Intel Xeon E5 v4 compute blade with 2-sockets per node (total 8 sockets), delivering maximum Intel Xeon processor capabilities. The IP-139CS is a dual node Intel Xeon E5 v4 compute blade with 2-sockets per node (total 4 sockets), coupled with HDD/SSD drives or PCIe slots to jointly increase processing and I/O capabilities. The IP-139CS-SXM2 blade supports up to 4 NVIDIA® Tesla® P100 (NVIDIA Pascal™) GPU accelerators interconnected via NVIDIA NVLink™. The IP-143CS is a quad node Intel Xeon Phi processor compute blade with one socket per node (total 4 sockets). All blades utilize an innovative cold sink technology to provide high node density at proper heat levels.
- **Blade Enclosure:** 10.5U blade enclosures provide power, cooling, system control, and network fabric for up to 9 compute blades via an integrated midplane. ICE XA can be expanded by simply adding enclosures, with up to four blade enclosures in a single 42U high rack. The enclosure is also designed to support future blade technologies.
- **Topology:** SGI ICE XA supports both Infiniband and Intel Omni-Path interconnect technologies with complete flexibility in topology. Choose from All-to-All, Fat Tree, Hypercube, or Enhanced Hypercube fabrics with single or dual plane to best meet performance, system size, budget, and application needs.

- **Operating system:** ICE XA runs standard SUSE® Linux Enterprise Server, Red Hat® Enterprise Linux®, or CentOS. Different Linux operating systems can be provisioned on different nodes, thereby allowing a broad range of Linux-based HPC applications to run simultaneously. SGI is one of the largest contributors to the Linux community and brings deep expertise deploying and optimizing Linux systems.

### Scale Out Live with Next-gen Technologies

SGI's ICE platform features "Live Integration" whereby systems can be expanded using next generation processors, accelerators, and networking, without interrupting users. This highly valuable capability has been demonstrated best at NASA's Ames Research Center, whereby twenty live platform upgrades, spanning five generations of ICE innovation, has saved them tens of millions of user computational hours of productivity by avoiding downtime that is typically associated with generational upgrades.

### Efficiency

#### Cutting-Edge Power and Cooling

Energy efficiency and green computing is an increasingly important requirement for supercomputers. SGI has been a principal innovator and champion in this area for many years delivering industry leading Megaflop per Watt technology for x86 platforms.

ICE XA extends SGI's leadership in power and cooling efficiency with:

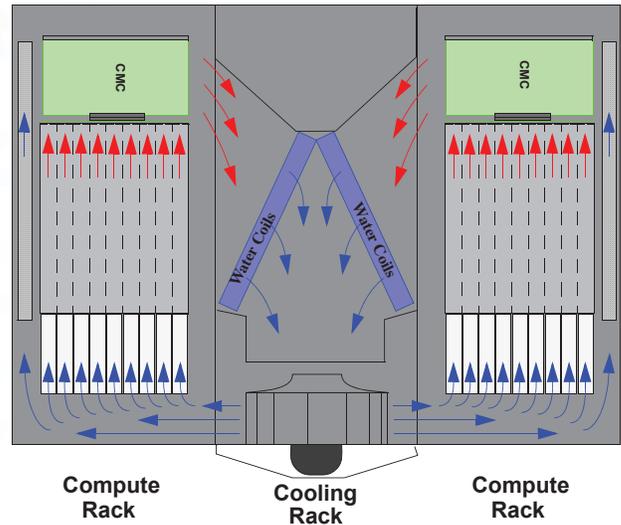
- **New E-Cell design:** Featuring 2nd generation SGI Cell technology, ICE XA systems are deployed as E-Cells. Two compute E-racks and a unified cooling rack are placed in a sealed E-Cell, with multiple cells connected to form large systems. A cooling rack within the cell draws hot air via an air-to-water heat exchanger and recirculates it to cool the compute racks. This "Closed-Loop Airflow" ensures no air from within the cell is mixed with data center air. In addition, the cell is always water-cooled.
- **Warm water cooling:** The E-Cell utilizes facility-supplied water for cooling and will not add any heat to the data center if the water temperature is within 45-90 degrees Fahrenheit (7-32C). This high 90F "Room Neutral" water temperature is often 30% or more efficient than industry competitors and can save millions in cooling costs. An air-to-water heat exchanger is provided with all E-Cells, and a water-to-water heat exchanger is deployed when cold sinks are utilized in the compute and switch blades.

- **Power supply efficiency:** Each blade enclosure in an E-rack is configured with up to 9 power supplies and supports N+1 power redundancy. The number of loaded power supplies can also be controlled so as to not draw unnecessary power and waste energy. The highly efficient design enables ICE XA to achieve a 94% power distribution at 50% load using 80 PLUS Platinum certified power supplies.

### Path to Exascale

Regarded as the next frontier in high performance computing, SGI is on a leading path to deliver supercomputers capable of an exaflop. Managing power consumption at a job level, utilizing higher facility water temperatures and full liquid immersion, and bringing high-speed SSD storage close to processors to accelerate I/O and reduce power consumption are among many areas of continuing innovation. While these ICE platforms will be built for only the largest of HPC environments, SGI's Exascale technology will extend broadly to solving compute problems at Tera- and Petascale.

### SGI ICE XA E-Cell



*Closed Loop airflow insulates E-Cell from data center environment*

Compute Blades	IP-125CS	IP-139CS	IP-139CS-SXM2	IP-143CS
Processors	Intel® Xeon® Processor E5-2600 v4 Series	Intel® Xeon® Processor E5-2600 v4 Series	Intel® Xeon® Processor E5-2600 v4 Series	Intel® Xeon Phi™ Processor
Compute Nodes/ Blade	Four 2-socket CPU nodes	Two 2-socket CPU nodes	One 2-socket CPU node	Four 1-socket CPU nodes
Memory/ Node	64-512GB/ node; 8 DDR4 DIMM slots (4 per CPU socket)	64-512GB /node; 8 DDR4 DIMM slots (4 per CPU socket)	64-512GB /node; 8 DDR4 DIMM slots (4 per CPU socket)	48-384GB/ node; 6 DDR4 DIMM slots (6 per CPU socket)
Memory Capacities	8, 16, 32 and 64GB 2400 MT/s ECC Registered DIMMs	8, 16, 32 and 64GB 2400 MT/s ECC Registered DIMMs	8, 16, 32 and 64GB 2400 MT/s ECC Registered DIMMs	8, 16, 32 and 64GB 2400 MT/s ECC Registered DIMMs
Coprocessor/ Accelerators (optional)	Not Applicable	• NVIDIA® Tesla® K40 GPU Accelerator	• Up to 4 NVIDIA® Tesla® P100 (NVIDIA Pascal™) GPU Accelerators	Not Applicable
I/O (optional)	Up to 8 Low Profile IO Cards (Infiniband, Intel® True Scale, Fibre Channel, Ethernet (1,10 & 40Gb)	Up to 4 Low Profile IO Cards (Infiniband, Intel® True Scale, Fibre Channel, Ethernet (1,10 & 40Gb)	Not Applicable	Not Applicable
Storage (optional)	Up to eight 2.5" SATA drives (HDD or SSD) per blade (2 per logical node)	Up to eight 2.5" SATA drives (HDD or SSD) per blade (4 per logical node)	Up to four 2.5" SATA drives (HDD or SSD) per blade	Up to eight 2.5" SATA drives (HDD or SSD) per blade (2per logical node)
Mezzanine Card	Single or Dual Port FDR, EDR Switch-IB™ 2, Intel® Omni-Path	Single or Dual Port FDR, EDR Switch-IB™ 2, Intel® Omni-Path	Single or Dual Port FDR, EDR Switch-IB™ 2, Intel® Omni-Path	Single or Dual Port FDR, EDR Switch-IB™ 2, Intel® Omni-Path
Cooling	SGI ColdSinks	SGI ColdSinks	SGI ColdSinks	SGI ColdSinks
Topologies Options	Single or dual plane all to all, fat tree, hypercube and enhanced hypercube	Single or dual plane all to all, fat tree, hypercube and enhanced hypercube	Single or dual plane all to all, fat tree, hypercube and enhanced hypercube	Single or dual plane all to all, fat tree, hypercube and enhanced hypercube

<b>Blade Enclosure</b>					
<b>Integrated Infiniband Switch</b>	<b>Standard</b>			<b>Premium</b>	
	Single 36 port FDR or EDR IB ASIC with 18 ports external			Dual 36 port FDR or EDR IB ASIC with 36 ports external	
<b>Integrated Omni-Path Switch</b>	Single 48 port Omni-Path ASIC with 30 ports external				
<b>Administrative Network</b>	Dedicated GigE network (redundancy optional), chassis management controller				
<b>Storage</b>					
<b>InfiniteStorage InfiniBand Solutions</b>	<ul style="list-style-type: none"> <li>• High performance shared file systems</li> <li>• IP over IB</li> <li>• Native IB block level access</li> <li>• Native IB SAN supported with CXFS</li> </ul>				
<b>Racks</b>					
	<b>D-Rack (For I/O &amp; Support Nodes)</b>			<b>E-Rack (For Compute Blades)</b>	
<b>Specifications</b>	83.2"H x 24.0"W x 49.5"D for air cooled rack. Water cooled rack is 54.8"D. 42U w/ 19" standard EIA mounting rails. Can optionally be extended to 48U		89.6"H x 24.0"W x 48.0"D. Room neutral up to 32°C datacenter supply water		
<b>Blade Enclosure Support</b>	n/a- used for I/O & support nodes		Up to four blade enclosures (36 logical nodes each)		
<b>Power</b>	Single and three-phase PDUs		Up to 8+1 redundant 3000W power supplies per blade enclosure		
<b>Cooling</b>	Open looped airflow or water (optional)		closed loop airflow/ water		
<b>System Management</b>					
<b>HSM</b>	Tier 1: System Administration Controller <ul style="list-style-type: none"> <li>• One per ICE system</li> <li>• Provisions out software to RLC</li> <li>• Pulls aggregated cluster management data from RLC</li> </ul>	Tier 2: Rack Leader Controller (RLC) <ul style="list-style-type: none"> <li>• One per eight blade enclosures (two E-Racks)</li> <li>• Holds blade boot images</li> <li>• Runs fabric management software</li> <li>• Aggregates cluster management data for rack</li> </ul>	Tier 3: Chassis Management Controller <ul style="list-style-type: none"> <li>• One per enclosure</li> <li>• Controls master power to all compute nodes</li> <li>• Monitors power and blade enclosure environment</li> </ul>	Tier 4: Baseboard Management Controller <ul style="list-style-type: none"> <li>• One per compute node</li> <li>• Controls board-level hardware</li> <li>• Monitors compute node environment</li> </ul>	
<b>Service Node Options</b>	<ul style="list-style-type: none"> <li>• Login Node</li> <li>• Gateway Node</li> <li>• Batch Node</li> </ul>	<ul style="list-style-type: none"> <li>• Storage Node</li> <li>• OSS Node</li> <li>• MDS Node"</li> </ul>	Service nodes can be optionally configured with: <ul style="list-style-type: none"> <li>• GPUs such as NVIDIA® Tesla™</li> <li>• I/O connectivity (various)</li> <li>• Hard Disk Drives (SAS and/or SATA)</li> <li>• OSS Node</li> </ul>		
<b>System Software</b>					
<b>Operating Systems</b>	<ul style="list-style-type: none"> <li>• SUSE® Linux Enterprise Server 11, 12</li> <li>• Red Hat® Enterprise Linux 6, 7</li> <li>• CentOS 6, 7</li> </ul>				
<b>Cluster Solution Stack</b>	<ul style="list-style-type: none"> <li>• SGI Management Suite: System Management software consisting of SGI Foundation Software and SGI Management Center</li> <li>• SGI Performance Suite optimized application performance package consisting of SGI Accelerate, SGI MPI, SGI REACT</li> </ul>				
<b>HPC Workload Managers</b>	<ul style="list-style-type: none"> <li>• Altair® PBS Professional™: Workload management tool with support for SGI Power Management</li> <li>• Moab® HPC Suite</li> <li>• Slurm Workload Manager</li> </ul>				
<b>Software Development</b>					
<b>Programming Languages and Debuggers</b>	<ul style="list-style-type: none"> <li>• C &amp; C++: Intel® C++ Compiler, PGI C/C++, GNU GCC</li> <li>• Debuggers: Intel® Debugger included with Intel® compilers, PGI® PGDBG®, Rogue Wave Software® TotalView® Team, Allinea DDT, GNU GDB</li> <li>• Fortran: Intel® Fortran Compilers, PGI Fortran, GNU GCC</li> <li>• Performance Analysis: Intel® VTune Amplifier XE, Intel® Trace Analyzer &amp; Collector, PGI® PGPROF®</li> <li>• Debugger and Profiler Allinea® Forge</li> </ul>				
<b>Libraries</b>	<ul style="list-style-type: none"> <li>• SGI MPI</li> <li>• OpenMP included with Intel® compilers</li> <li>• Intel® Math Kernel Library</li> <li>• Intel® Parallel Building Blocks</li> <li>• Intel® Integrated Performance Primitives</li> <li>• Intel® MPI Library</li> </ul>				

### 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

For more information about how SGI ICE XA can benefit your organization, visit [www.sgi.com](http://www.sgi.com) or call 1-800-800-7441.

Global Sales and Support: [sgi.com](http://sgi.com)

©2016 Silicon Graphics International Corp. All rights reserved. SGI, the SGI logo and ICE XA 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. Lustre is a registered trademark of Xyratex Technology, Ltd. in the United States. Other names may be claimed as property of others. All other trademarks are property of their respective holders. 05112014 4529 01122016

