SGI
High Performance Computing

- *Accelerate time to discovery, innovation, and profitability*

Typical Use Cases for SGI HPC Products

Large scale-out, distributed memory models:
Best in Class speed, scale and efficiency

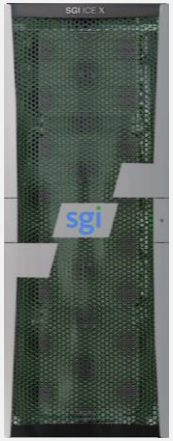

SGI ICE

Fat-node cluster, large models.
Ease of use for small depts & businesses

SGI UV

Customers that value factory integrated solutions for Technical Computing and Big Data.

SGI Rackable

SGI InfiniteData Cluster

For user files, fast scratch, active archiving

SGI InfiniteStorage

SGI HPC customer examples

- NASA ‘Pleiades’ 200,000+ ICE cores, 5.3PF, 160 racks
- Total ‘Pangea’ worlds largest commercial HPC, 110,000+ ICE cores, 2.3PF
- AWE 80,000+ ICE cores, 1.8PF machine and another at 1PF
- DOD 73,000+ Rackable cores
- JAMSTEC 64TB shared memory UV
- Pawsey Centre 100PB storage system, MWA and ASKAP data
- Idaho National Labs 2262 MFlops/Watt
- IT4 2.4PF, largest Phi installation in Europe
- USPS real time fraud detection scanning billions of pieces of mail and saving USPS millions of dollars

HPC Solutions – SGI ICE XA

Achieve and Accelerate Computational Breakthroughs

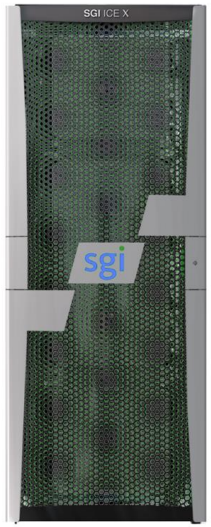

The World's Fastest Supercomputer

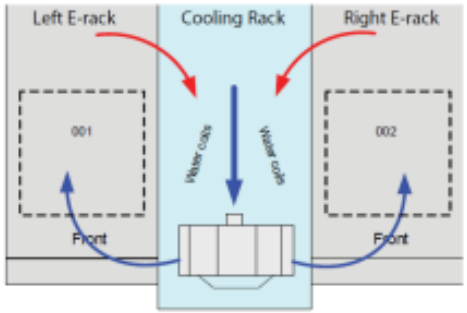

- Supercomputer with the world record performance crown – 5 years running!
- Seamlessly expand existing systems with new technology and without interruptions
- 2nd generation E-cell warm water cooling technology saves $M in energy costs
- Flexibility in blades and network topology to best meet customer needs

Notable Features of a “Cell”

“Closed-Loop Airflow” Environment

- Room Neutral – No air exchange with data center
- Reduces load on CRAC units in data center
- Always water-cooled

Warm Water Cooling Support

- Broad range of acceptable temperatures for additional cost savings (up to 32.2° C/ 90° F)/ energy reuse
- Contains air-to-water heat exchanger
- Contains a water-to-water heat exchanger for cold sinks

Large, “Unified” Cooling Racks for Efficiency

- Compute racks do not have their own cooling at the rack level
- Decreases power costs associated with cooling
- All cooling elements utilize one water source

Acoustically Treated

- Leave your ear plugs out and…

Have a conversation in your data center again!

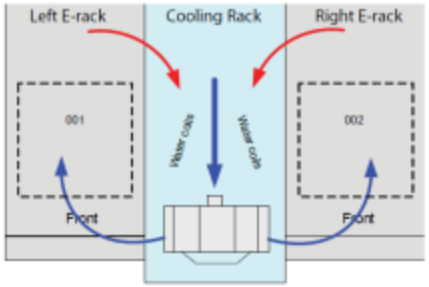

Air Circulation E-Cell


CDU


Blower
(One of three)


Cooling Rack
(Blowers removed)


Water Coil (One of two)

Best in Class HPC

Market Leading Energy Efficiency - Hardware Innovation

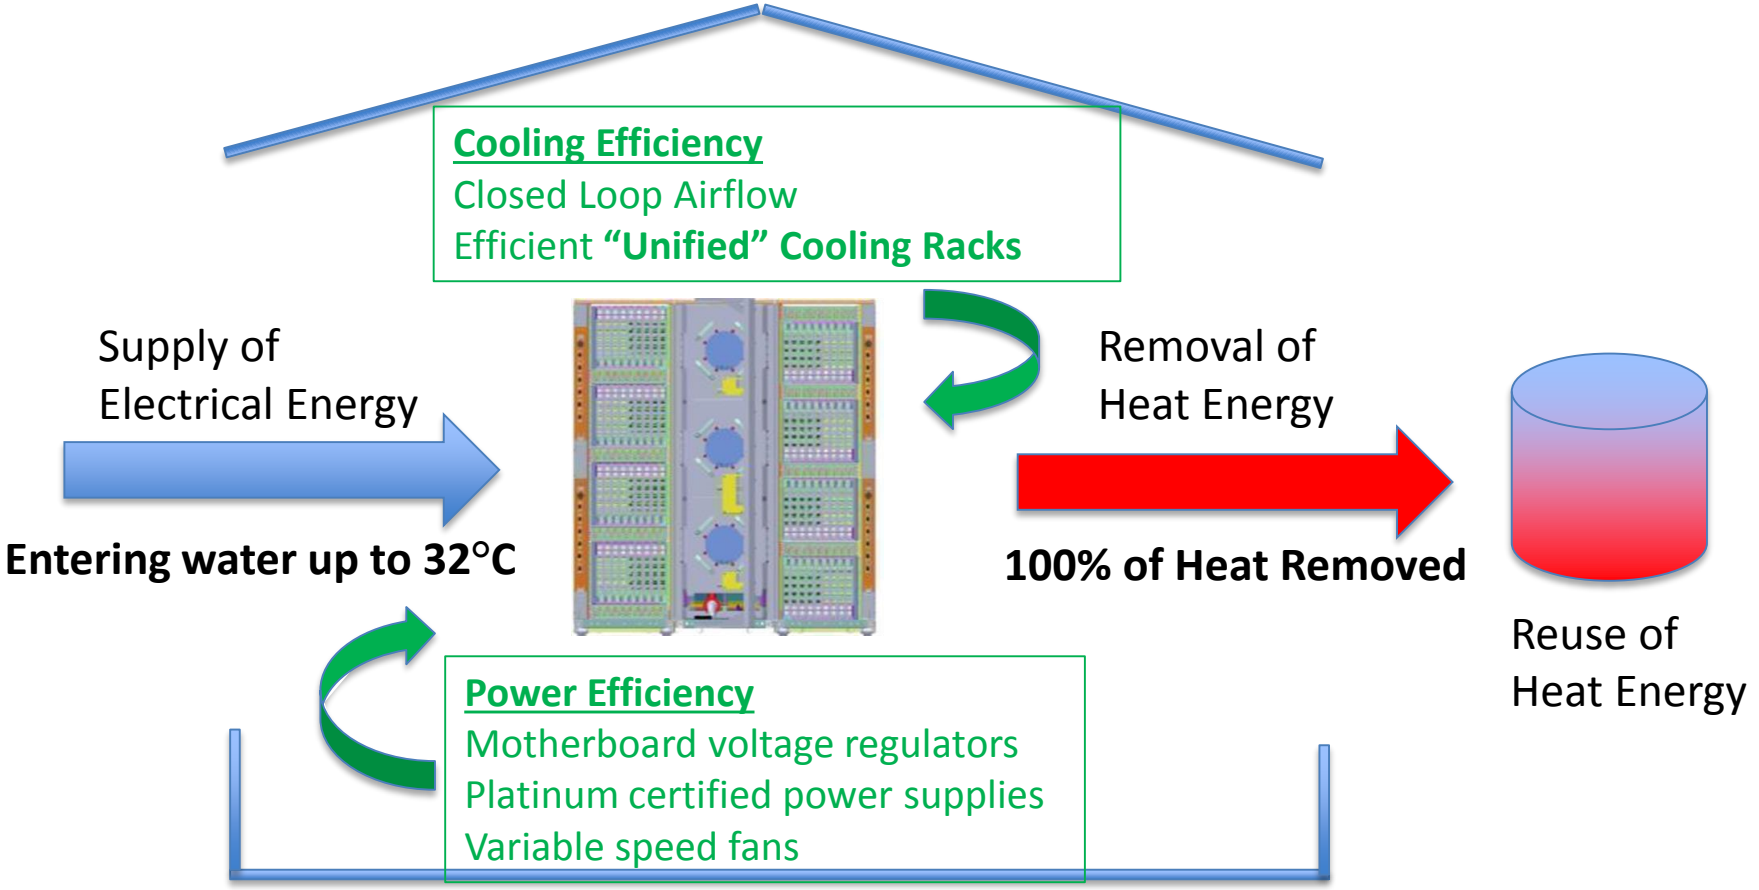

Industry Leading Efficiency 2262 Mflops/Watt – Idaho National Labs (Falcon)

HPC Solutions – SGI UV

Accelerate the Pace of Innovation

The Big Brain Supercomputer

- World's most powerful scale up systems for in-memory computing – up to 48TB!
- Leverage 7th gen technology,100+ patents, and 20+ years of in-memory expertise
- Choice of systems for compute- and data-intensive workloads
- Single system simplicity for ease of management and reducing IT overhead

Long History of SGI Shared Memory Systems: 1996-2015

Origin 2000, 1996
MIPS R10K

Origin 3000, 1999
MIPS R12K

Altix 3000, 2003
Intel Itanium

Altix 4700, 2006
Intel Itanium

UV1, 2010
Intel E7

UV 2000, 2012
Intel E5

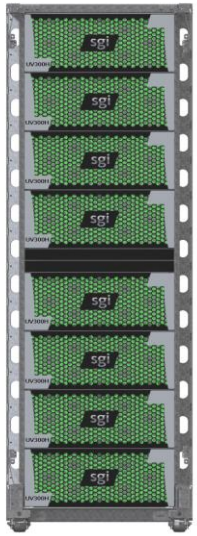

UV 300, 2015
Intel E7

UV 300

UV 300 Multi-Chassis Configurations

- Fabric: **NUMAlink 7**
  - Cache coherent hardware protocol
  - Optimised for small (cache line sized) transfers
  - Very low latency remote memory access
  - Each HARP2 ASIC has 16 links
  - Each link 14GT/s (7.47GB/sec)
- Topology
  - **All-to-all topology** for up to 8 chassis (one rack, 32 sockets) with maximum of 1 hop between nodes
  - Adaptive Routing
- Single Rack System
  - 32 CPU sockets
  - 64 Memory Risers
  - 12 DIMMs / Memory Riser
  - 768 DIMMs -> **24TB memory** (32GB DIMMs)
  - Up to 92 PCI-E slots
- Multi Rack, Architectural Max-Config System
  - Up to 2000+ CPUs
  - Up to **48TB Memory**
  - Up to 188 PCI-E slots

UV 300 NUMAlink Fabric & Topology

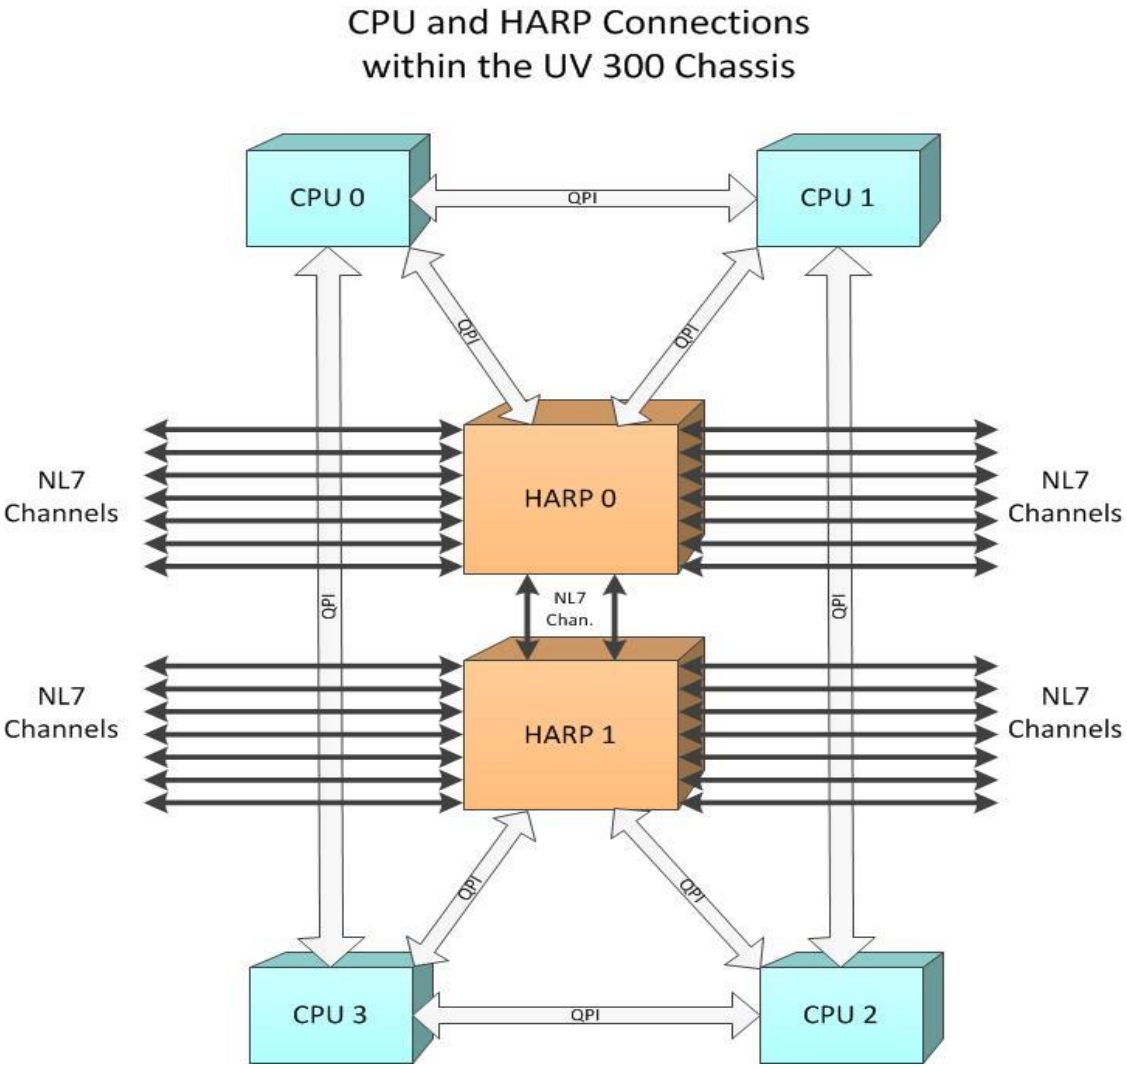

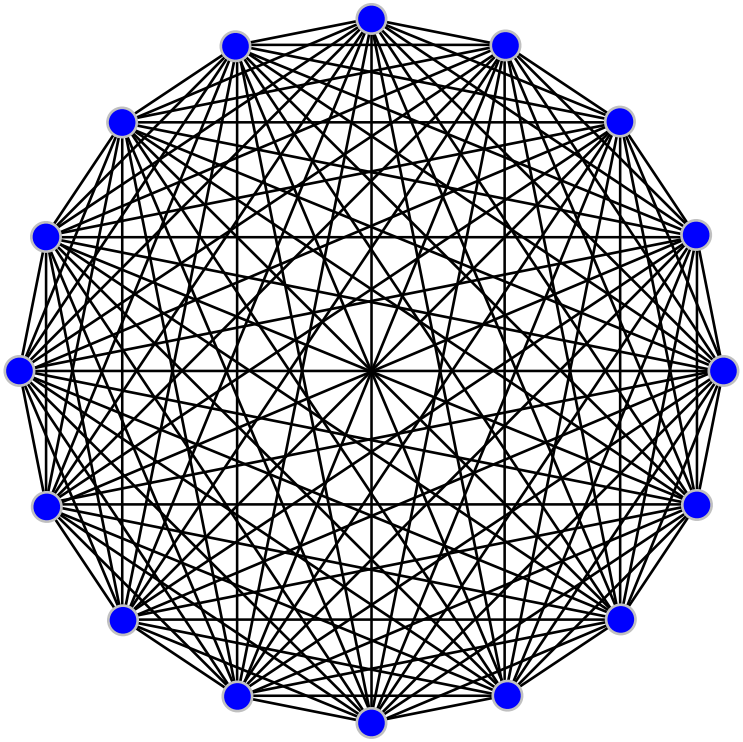

- All-to-all up to 32 sockets
- **Max hop count = 1**

Memory Bandwidth

*a[j] = b[j]+scalar*c[j];*

Sub. Date	Machine ID	number of cpus	TRIAD (10^6 Bytes/sec)
2012.08.14	SGI_Altix_UV_2000	2048	7,139,690
2011.04.05	SGI_Altix_UV_1000	2048	5,859,367
2006.07.10	SGI_Altix_4700	1024	4,350,166
2013.03.26	Fujitsu_SPARC_M10-4S	1024	4,002,703
2011.06.06	ScaleMP_Xeon_X6560_64B	768	2,259,709
2004.12.22	SGI_Altix_3700_Bx2	512	1,119,913
2003.11.13	SGI_Altix_3000	512	1,007,828
2003.10.02	NEC_SX-7	32	872,259
2008.04.07	IBM_Power_595	64	805,804
2013.09.12	Oracle_SPARC_T5-8	128	642,884
1999.12.07	NEC_SX-5-16A	16	583,069
2009.08.10	ScaleMP_XeonX5570_vSMP_16B	128	445,869
1997.06.10	NEC_SX-4	32	436,954
2004.08.11	HP_AlphaServer_GS1280-1300	64	431,450
1996.11.21	Cray_T932_321024-3E	32	359,270
2014.04.24	Oracle_Sun_Server_X4-4	60	245,068
2007.04.17	Fujitsu/Sun_Enterprise_M9000	128	227,059
2002.10.16	NEC_SX-6	8	213,024
2006.07.23	IBM_System_p5_595	64	206,243
2013.09.17	Intel_XeonPhi_SE10P	61	174,663

Triad is designed to stress local memory bandwidth since the vectors may be allocated in an aligned manner such that no communication is required to perform the computation.

STREAM source code is freely available from http://www.cs.virginia.edu/stream/.

HPC Solutions – Rackable

Delivering Top Value and Performance

Powerful, Flexible, and Efficient

- HPC clusters featuring value-driven density, performance, and capacity
- Unparalleled configuration flexibility in chassis and components
- Incorporates several system features designed to optimise power utilization and system cooling
- Scales from Teraflops to Petaflops
- Fully integrated and factory tested

Best in Class HPC

Resiliency and High Availability

80% of correctable memory errors turn into uncorrectable memory errors – the primary cause of server failures

Log → Analysis → Proactive Action → Resolution

Resulting in → Minimized downtime, improved overall solution robustness, and maximizing user productivity

AFRL DSRC
32 Racks & 73,000+ Cores
100% Monthly Availability

NOAA
12 Racks & 27,000+ Cores
100% Monthly Availability

HPC Solutions – SGI Software

Optimizing HPC Investments

SGI Management Suite

- Auto-provision at massive scale
- Health management to provide industry leading resiliency
- Best in class power management

SGI Performance Suite

- Full SGI MPI 3.0 compliant library and performance analysis tools
- Accelerate applications with SGI Performance Boost
- Hard Real-time library for Linux

Multi stage, hybrid, real time, data stream processing

or Real time analytics or SKA straw man

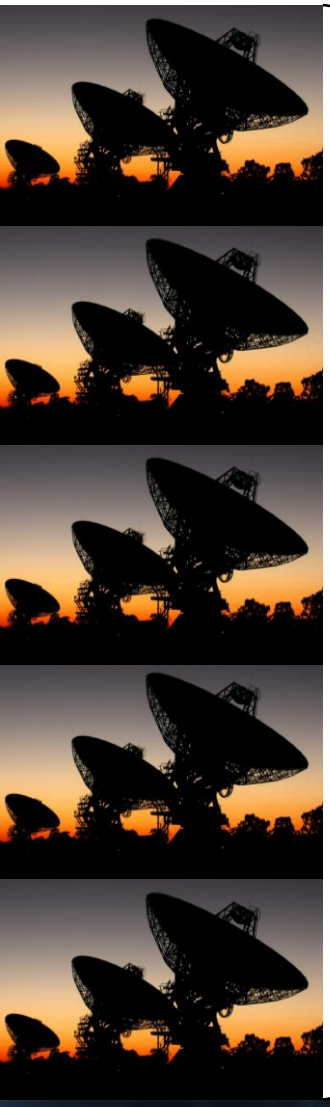

8TB/s

Embarrassingly Parallel
10x Data reduction

800GB/s

PowerMX with FPGA

Embarrassingly Parallel
10x Data reduction

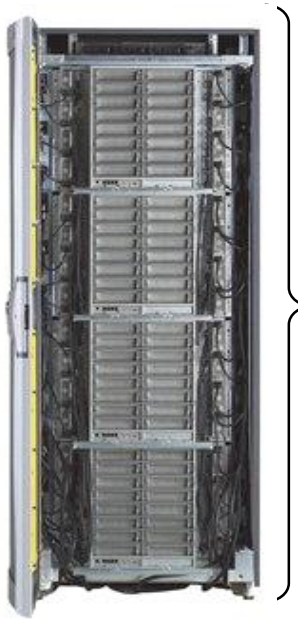

80GB/s

ICE XA

Thread intensive, in memory computing

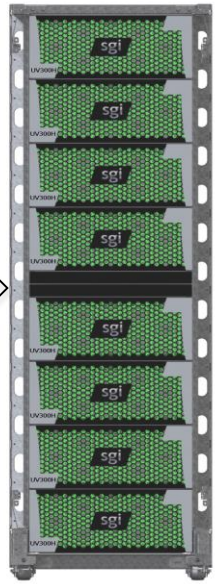

SGI UV 3000

- 64TB memory
- 12TB/s memory bandwidth
- 30 second buffer

Best in Class HPC

World Class Services

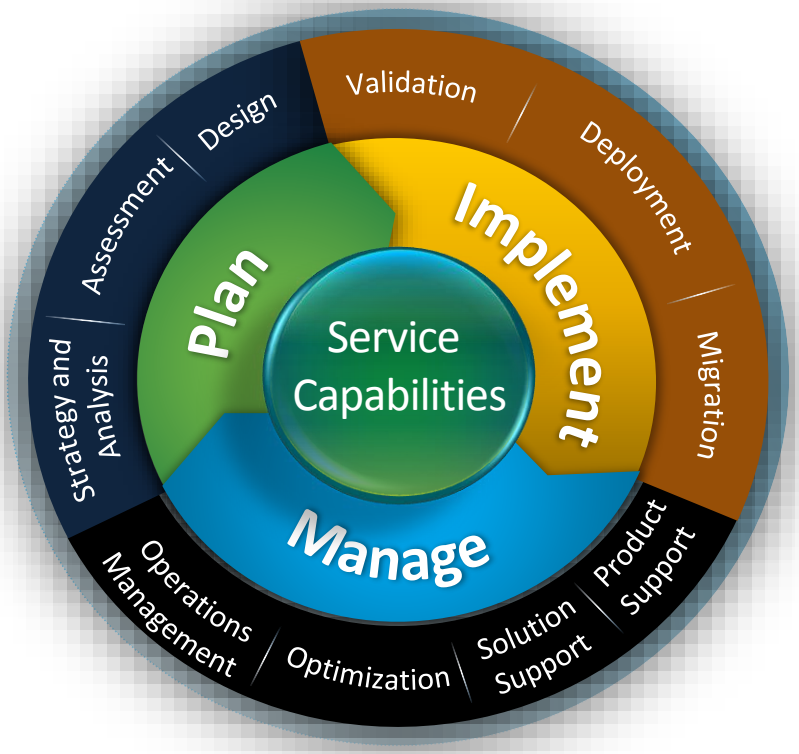

Streamlined Deployment

Faster time to productivity

Reduce stress on existing IT resources

Seamless Integration

Tera to Petaflop class solutions

Run full production while adding performance

Customer Intimacy

Customer excellence index averages 90%

High level of subject matter expertise

"Thanks to SGI Services team and their dedication and knowledge of our environment, we were able to stay focused on what we needed to deliver and maximize the return of our investment"
-Ryan Quick, Principal Architect