

# SGI<sup>®</sup> Speed and Scale

*Jesús Martínez Chavolla*  
*General Manager – México*



# Markets

## HPC

---

Commercial  
Scientific  
Modeling & Simulation

## Big Data

---

Hadoop  
In-memory  
Analytics  
Archive

## Cloud

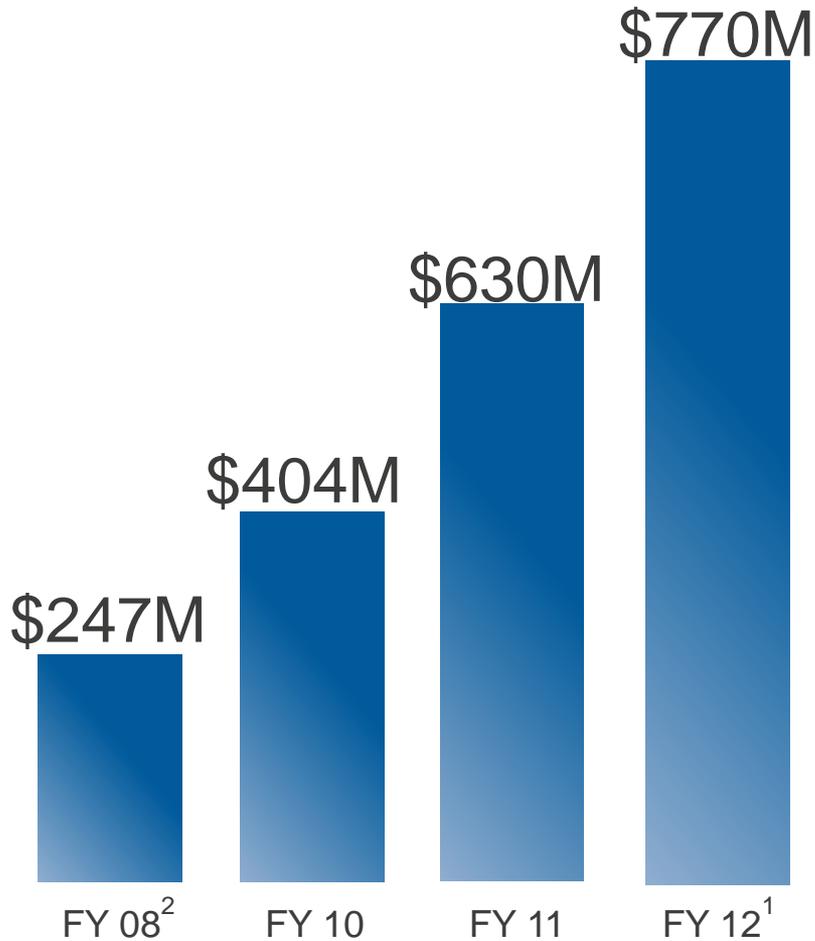
---

Public  
Private  
Government

sgi

Providing Customers with **Trusted** Technical & Commercial  
Computing Solutions

# Building Business Momentum



- World Leader Provider of Technical Computing
- Distribution in 50 countries
- 1,500+ employees
- Net Assets: \$538M<sup>3</sup>
- Cash & Investments: \$143M<sup>3</sup>
- Debt Free<sup>3</sup>

1 Revenue published FY2012

2 Rackable acquired substantially all the assets of Silicon Graphics, Inc., changed name to SGI in May 2009, filed a 6-month stub year-end on 6-26-09

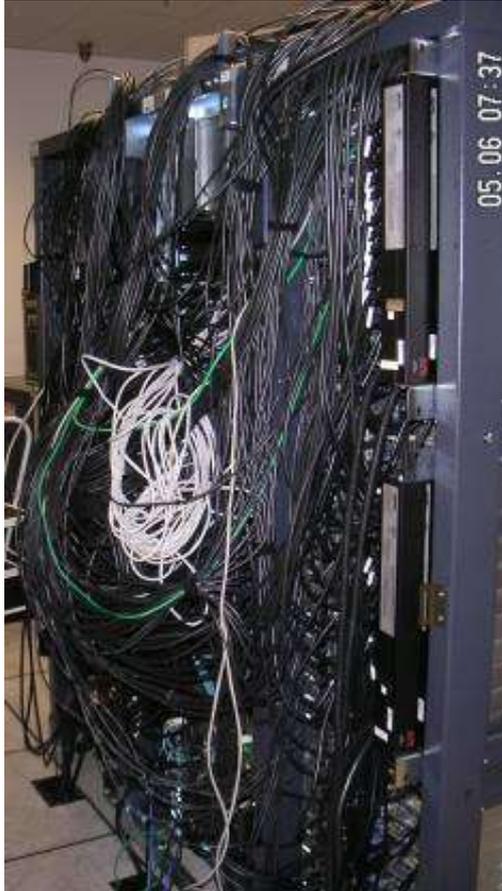
3 As of June 2012

**SGI<sup>®</sup> ICE<sup>™</sup> X**



# Breakthrough Reliability

## Cable-free, Redundant Components



**Traditional racked cluster**

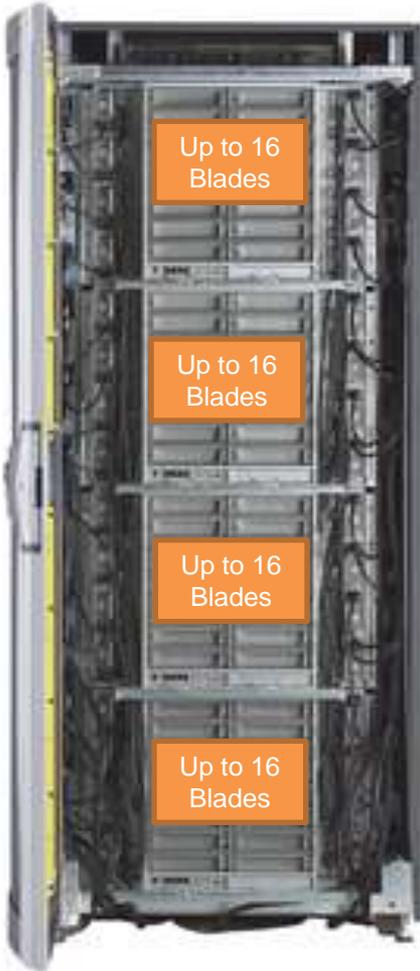
- Altix ICE racks and no cables in back
- Redundant, hot swap power and cooling
- Fully buffered DIMMs to reduce transient errors
- Blade design provides rapid serviceability
- InfiniBand FDR backplane for high signal reliability
- First \*over 1PF peak\* InfiniBand pure compute connected CPU cluster
- World's fastest distributed memory system
- First and only vendor capable of live, large-scale compute capacity integration



**Altix ICE cluster**

# Altix ICE & ICE X

**Performance Density: Up to 2304 Cores and 53.4 TFlops per Rack**



**SGI® ICE X Compute DUAL XEON E5 Blade Up to 16-Core, 144GB, 2-IB**



**SGI® Altix® ICE Compute DUAL XEON E7 Blade Up to 24-Core, 96GB, 2-IB**



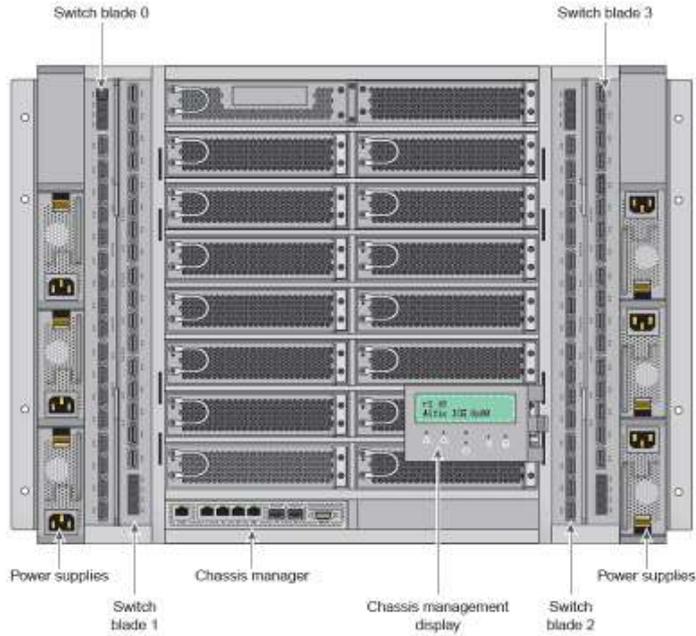
**SGI® Altix® ICE Compute DUAL Interlagos Blade Up to 32-Core, 128GB, 2-IB**



**M CELL RACK**



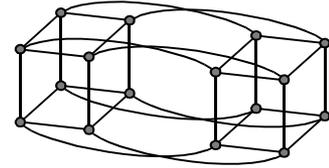
# Flexibility in Networking Topologies



**Robust integrated switch blade design enables industry-leading bisectional bandwidth at ultra-low latency!**

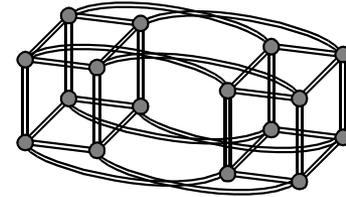
## ■ Hypercube Topology:

- Lowest network infrastructure cost
- Well suited for "nearest neighbor" type MPI communication patterns



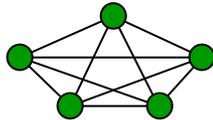
## ■ Enhanced Hypercube Topology:

- Increased bisectional bandwidth per node at only a small increase in cost
- Well suited for larger node count MPI jobs



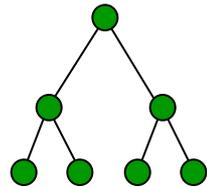
## ■ All-to-All Topology:

- Maximum bandwidth at lowest latency for up to 128 nodes
- Well suited for "all-to-all" MPI communication patterns.



## ■ Fat Tree Topology:

- Highest network infrastructure cost. Requires external switches.
- Well suited for "all-to-all" type MPI communication patterns

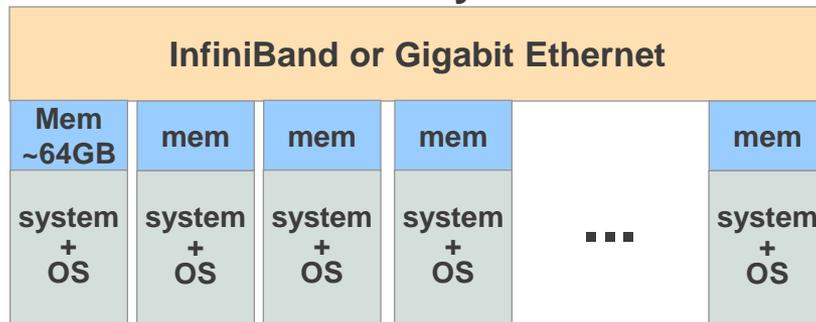


# SGI® UV2



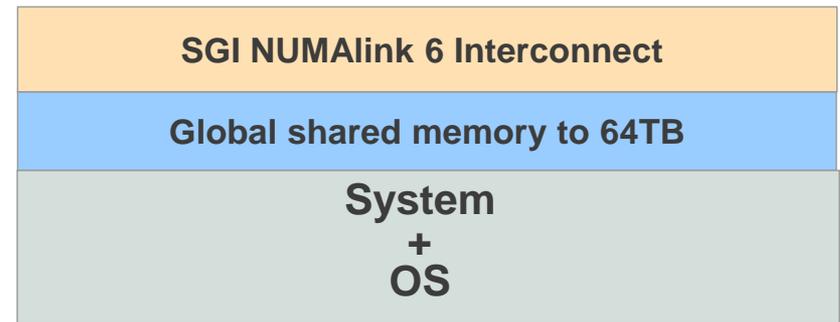
# SGI UV Shared Memory Architecture

## Commodity Clusters



- Each system has own memory and OS
- Nodes communicate over commodity interconnect
- Inefficient cross-node communication creates bottlenecks
- Coding required for parallel code execution

## SGI UV Platform



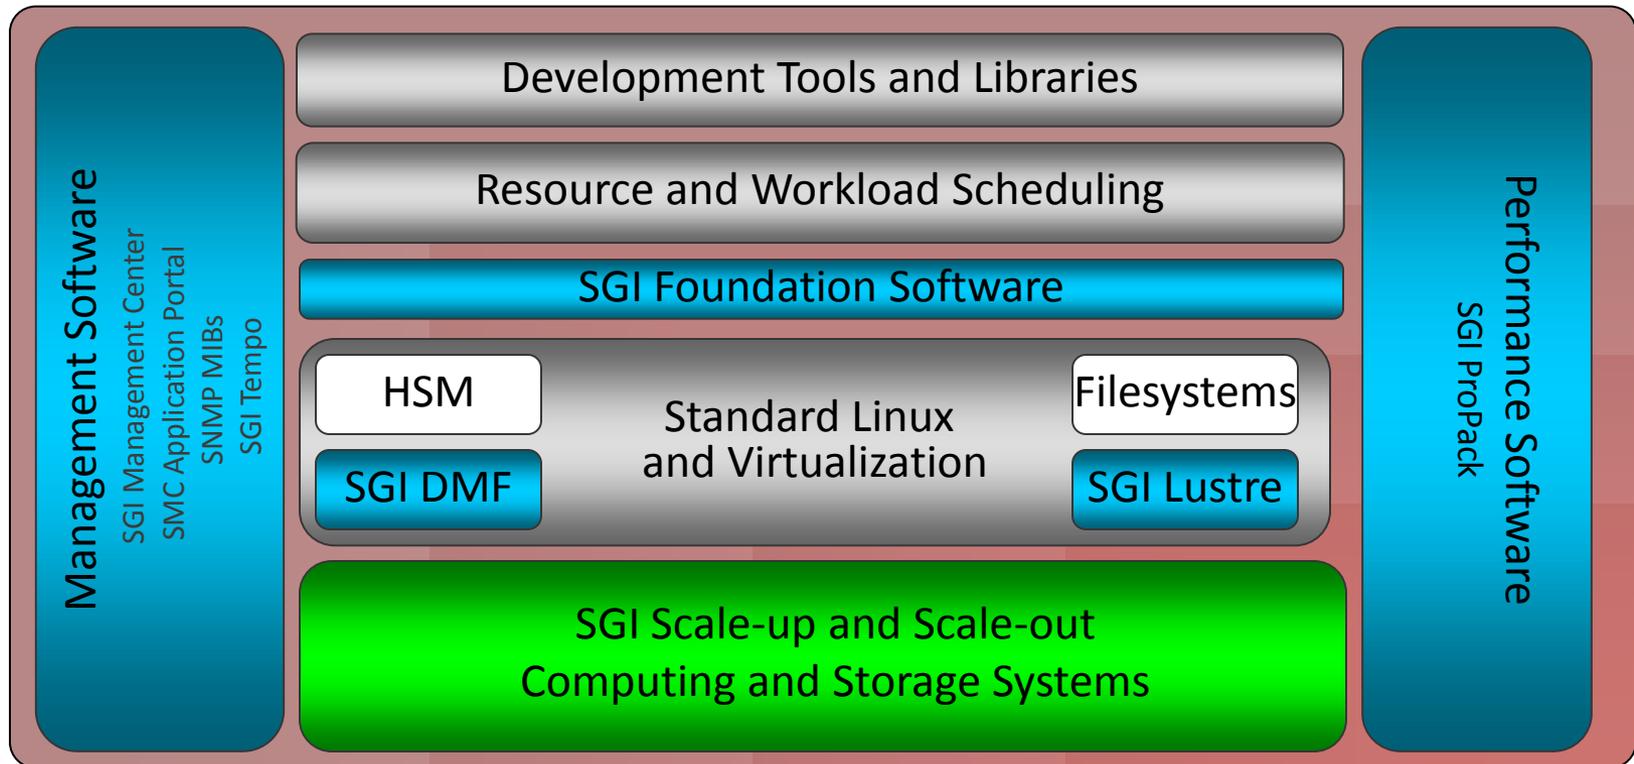
- All nodes operate on one large shared memory space
- Eliminates data passing between nodes
- Big data sets fit entirely in memory
- Less memory per node required
- Simpler to program
- High performance, low cost, easy to deploy
- Latency 80 ns, interconnect 37GB/sec

# SGI UV vs. the Competition

SGI UV vs HP, IBM, and Oracle/SUN	SGI UV1000	SGI UV2000	HP Superdome 2	IBM P795	Oracle/SUN
Processor Type	Intel E7	Intel E5-4600	Itanium 9300	Power 7	Sparc T4
# of Sockets	4 - 256	6 - 512	4 - 32	4 - 32	4 - 64
# Cores	16 - 2560	32 - 4096	16 - 128	256	256
Max Memory	16TB	64TB	4TB	8TB	4TB
# of Partitions	2 - 32	2 - 64	2 - 16	2 - 254	2 - 24



# SGI Software Environment



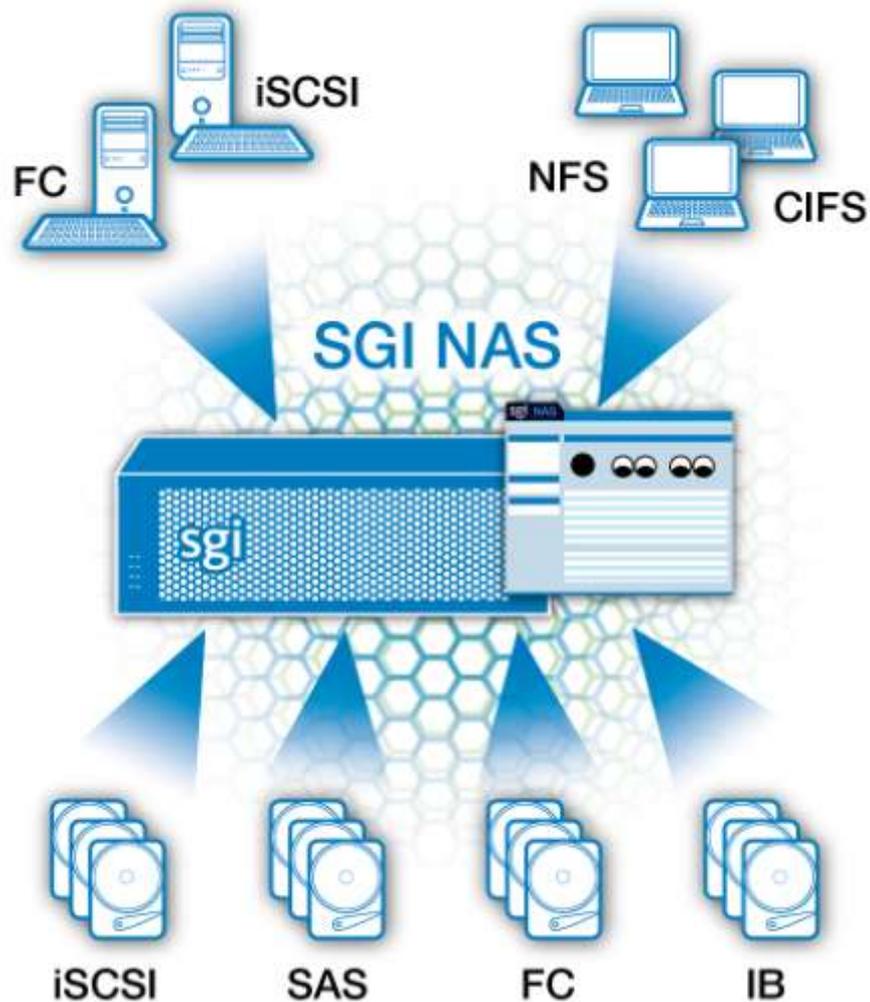
  SGI products

  Third party product (available from and/or integrated by SGI)

# SGI® NAS



# SGI NAS = Flexibility



- What is it?

- Enterprise-class open storage NAS/SAN software
- Integrated with SGI storage and server hardware

## The Result:

### A no-compromise Enterprise-class open storage solution:

- Hybrid Storage Pool for high performance
- Compression & Deduplication
- Unlimited file size & snapshots
- Synchronous and Asynchronous replication
- End-to-end 128bit data integrity



# SGI NAS

- Scalable from 20TB to 2PB with the same system
- Unparalleled density, scalability and flexibility



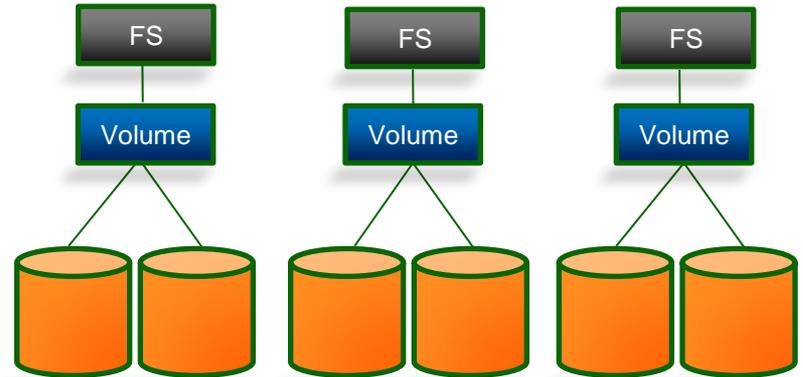
# SGI NAS Key Features

- Supports NAS and SAN
  - NFS/CIFS and iSCSI/FC
- Highly scalable capacity and performance
  - Benchmarked at over 2,500MB sec
  - Takes advantage of current technology eg SSD for performance
- Unlimited snapshots and clones
  - With fully integrated search
- Full active/active HA configuration
  - No single point of failure
- Heterogeneous block and file replication
  - Allows easy DR configurations
  - Synchronous or Asynchronous
- Internal RAID-Z and spares
  - Raid 5 & 6
- VM integration
  - Vmware, XEN and Hyper-V
- Integrated Windows back up & restore
- Cloud storage capabilities
- Compression, Dedupe and in box virus scan
  - All in band – not post process
- End to end data integrity
  - No silent data corruption
- Thin provisioning
- Non-disruptive volume growth
- DTrace based analytics

# Storage Pool

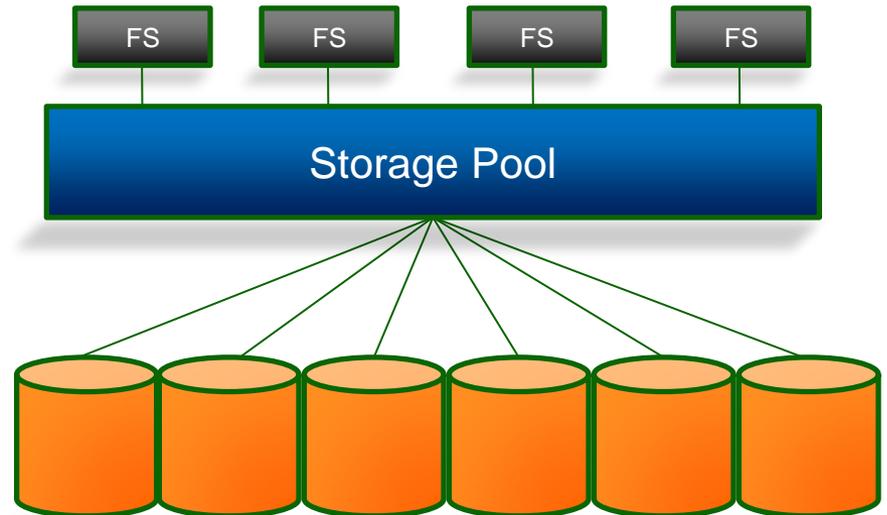
## Traditional file system:

- *Per Volume one file system*
- *Hard to switch drives to another volume*
- *Extensive planning necessary*
- *Every file system has a limited bandwidth*



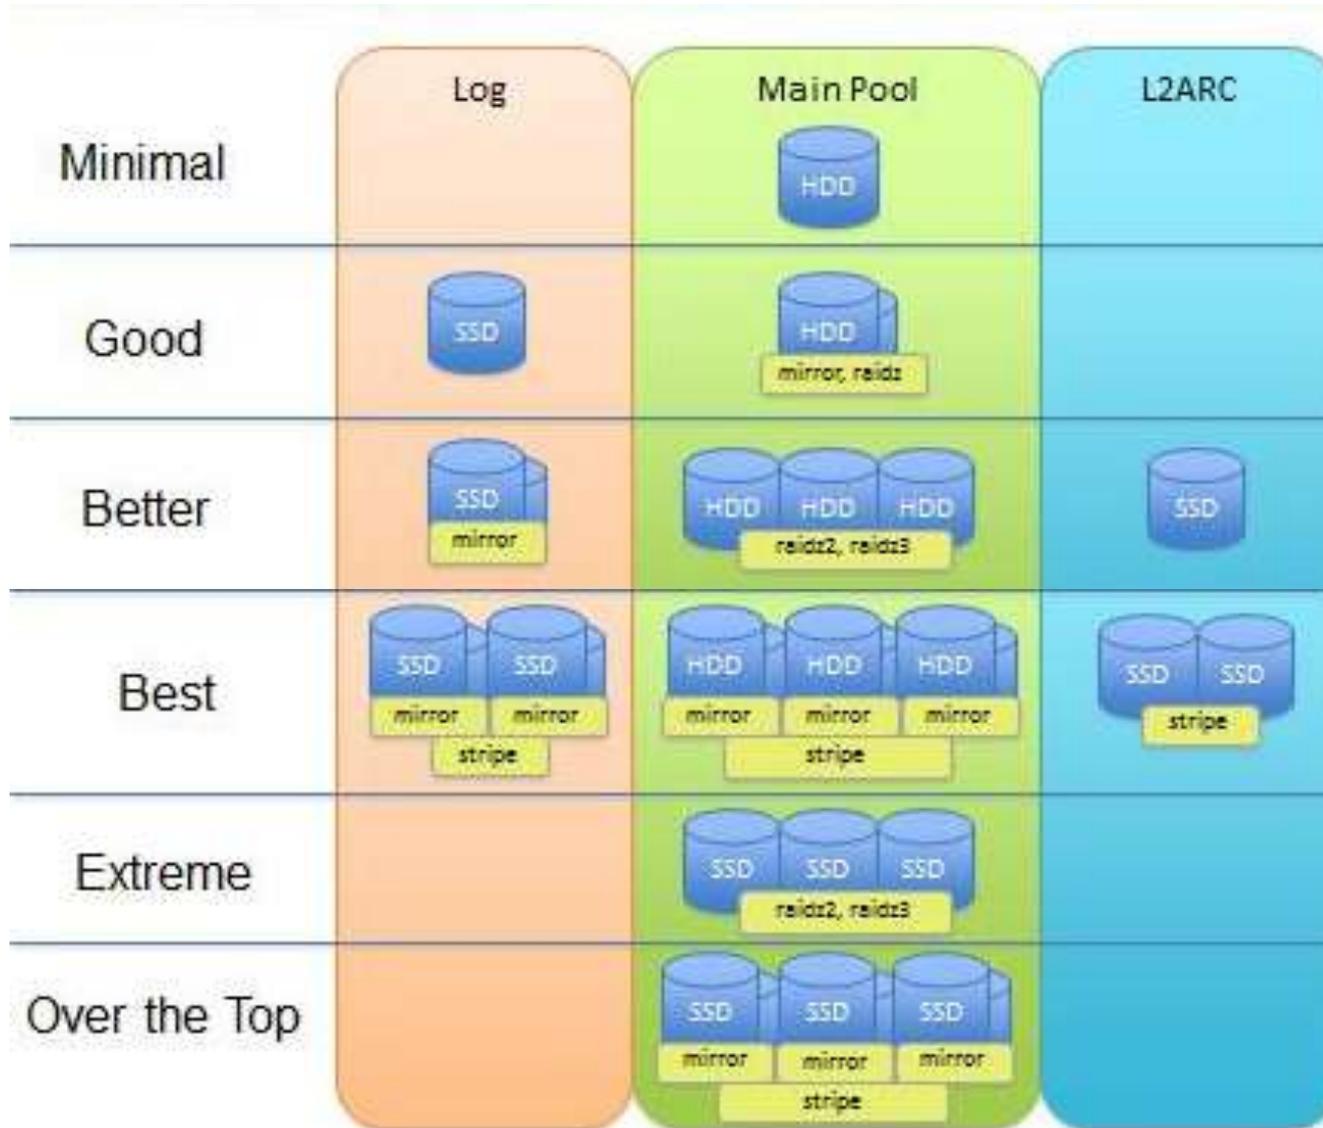
## SGI NAS ZFS:

- *No partitions/ volumes*
- *Extension automatically (thin provisioning)*
- *Uses bandwidth and I/O all pool members*
- *File systems serve themselves from the pool*
- *File system independent of volumes/ drives*



# Hybrid storage pool configurations

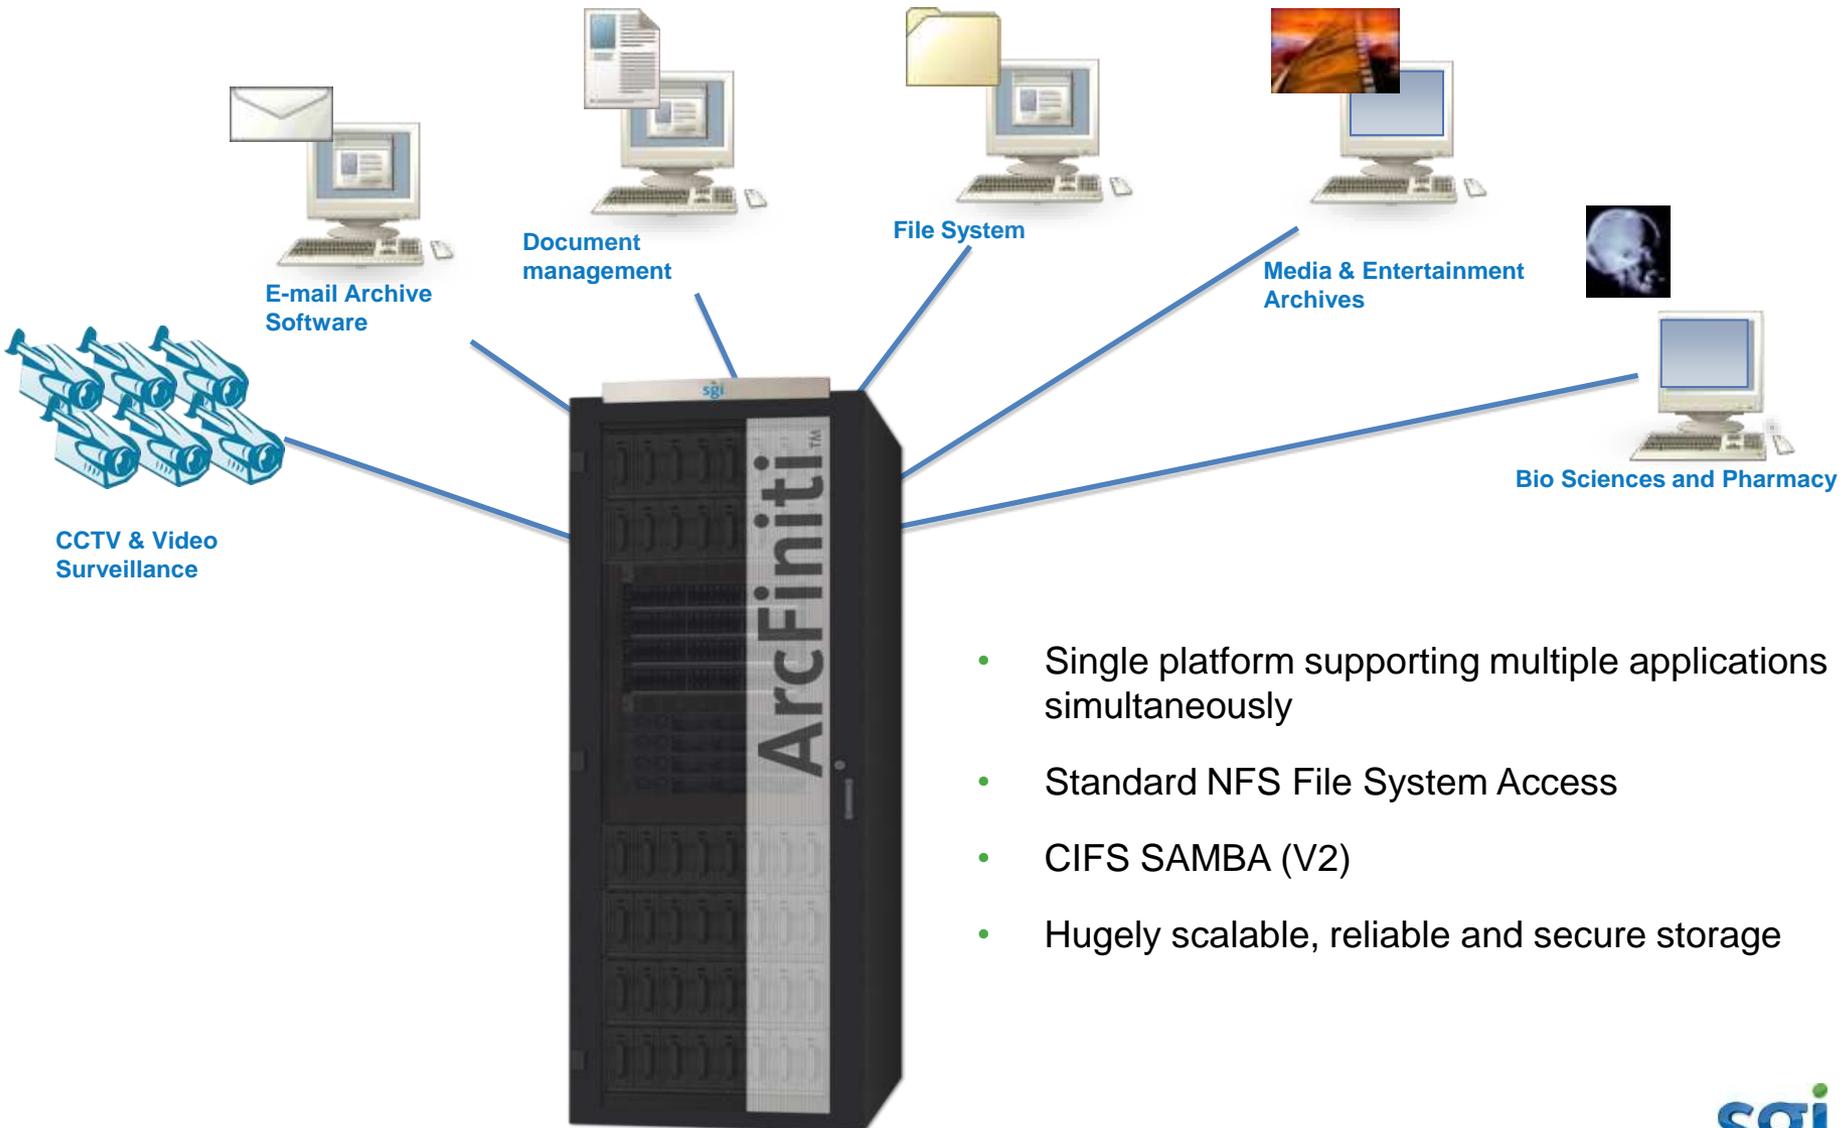
*Good, Better, Best, and Then Some*



# SGI® ArcFiniti™

*Fully Integrated, File Based Archive Solution*

# Objective: Providing a true archive for all persistent file data



- Single platform supporting multiple applications simultaneously
- Standard NFS File System Access
- CIFS SAMBA (V2)
- Hugely scalable, reliable and secure storage

# ArcFiniti™: Virtualized Tiers

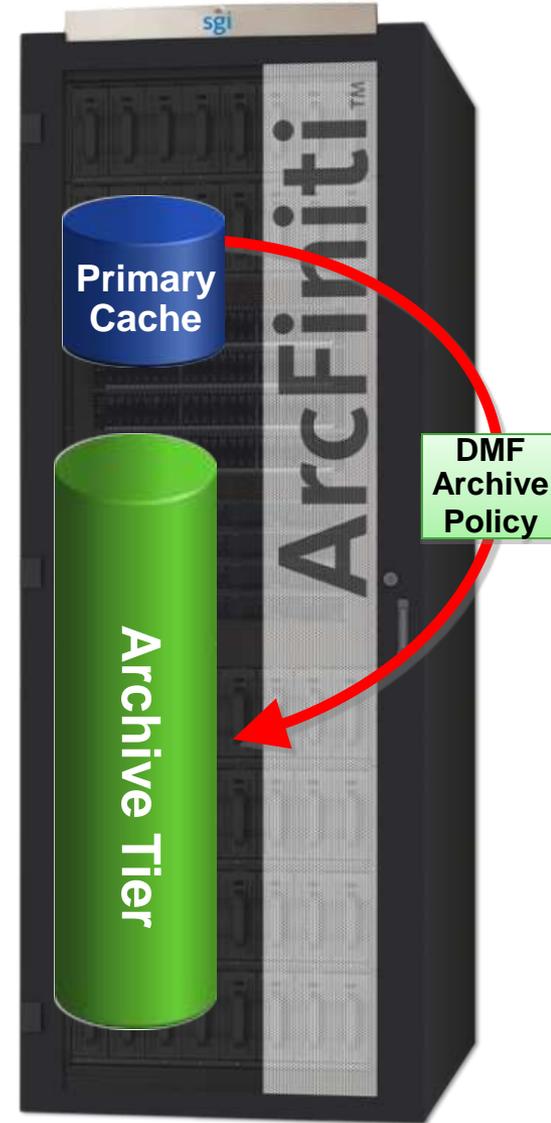
## Enhanced MAID: SGI's Differentiation

### POWER MANAGED RAID® Software

- Maximum of 25% OR 50% of drives powered on
- RAID protection (RAID5 3+1)
- Increases drive life by more than 6x
- Contributes to >3.8M hour disk drive MTBF

### DISK AEROBICS® Software

- Assures drive health and data integrity
- Predictive drive maintenance
- Avoids lengthy RAID rebuilds with proactive replacement
- Disk scrubbing for continuous data integrity



# Ultra Scalability – Useable capacity

## Five ArcFiniti Configurations

Model Number	A156	A234	A468	A936	A1404
Max Usable Capacity	178 TB	256 TB	490 TB	970 TB	1.44 PB
Primary Disk Cache	10TB* 22TB - Opt	10TB* 22TB - Opt	10TB* 22TB - Opt	22TB* 34TB - Opt	34TB* 46TB - Opt
Power Managed Archive Tier (Usable)	156 TB	234 TB	468 TB	936 TB	1.4 PB
GigE Ports	2	2	2	4	6
10GigE Ports	2	2	2	4	6

\* Standard configuration. Optional 12TB of additional cache capacity available

# ICE Cube Air

## Modular Data Center



# ICE Cube Air is about Integration



The process has already started...with the IT and the

# Cube® Air [\(Video\)](#)

- A modular data center that is built to order
- Eco-logical™ Air Cooling, Evaporative, Direct Expansion (DX) Cooling, PUE < 1.06
- Customizable 3rd Party Hardware Support
- Advanced Monitoring and Controls
- Installation and maintenance by SGI Worldwide Support
- Runs with outside air and evaporative cooling in most climates
- Supplemental direct expansion (DX) or chilled water cooling available
- ICE Cube Air uses 1 percent of the water
  - 'garden hose' –potable water, 2 GPM
  - No 'blow-back' (harmful chemicals, etc.)

# Efficiency Improvement Quantification

<b>EFFICIENCY IMPROVEMENT ANNUAL COST SAVINGS FOR VARIOUS LOADS</b>					
<b>LOAD (KW)</b>	<b>1%</b>	<b>2%</b>	<b>3%</b>	<b>4%</b>	<b>5%</b>
250	\$2,190	\$4,380	\$6,570	\$8,760	\$10,950
500	\$4,380	\$8,760	\$13,140	\$17,520	\$21,900
750	\$6,570	\$13,140	\$19,710	\$26,280	\$32,850
1000	\$8,760	\$17,520	\$26,280	\$35,040	\$43,800
1250	\$10,950	\$21,900	\$32,850	\$43,800	\$54,750
1500	\$13,140	\$26,280	\$39,420	\$52,560	\$65,700

**ASSUMES \$.10/KW-HR**

Source: Tier1 Research

# Monitoring Panel (Sample)

192.168.10.10
DISCONNECT
Tuesday, December 07, 2010 7:49:39 PM
CLOSE APP.

REAL TIME DATA
DAMPERS
DOORS
PSYCHROMETRIC CHART
SET POINTS
CONFIGURATION

### TEMPERATURES

	DB °C	RH %	OA STATUS		
OA01	20.80	30.50	<span style="background-color: yellow; padding: 2px;">B</span>		
	DB °C	RH %	AVERAGE DB °C	AVERAGE RH %	DELTA T °C
RI11	20.80	31.30			
RI12	21.10	30.40	21.00	31.80	10.10
R011	30.20				
R012	31.90		31.00		

### FAN DATA/PRESSURES

Fan output: 15.00 %

	Pa	m3/s	rpm	W	
F11	<span style="background-color: green; color: white; padding: 2px;">OK</span>	23	0	327	21
F12	<span style="background-color: green; color: white; padding: 2px;">OK</span>	26	0	327	19
F13	<span style="background-color: green; color: white; padding: 2px;">OK</span>	30	0	327	18
F14	<span style="background-color: green; color: white; padding: 2px;">OK</span>	27	0	327	17

Total Air Flow: 1.10 m3/s

Rack Pressure Drop: 0 Pa

### POWER MONITOR

	main	pdu	
Voltage	<span style="border: 1px solid black; padding: 2px;">208</span>	<span style="border: 1px solid black; padding: 2px;">208</span>	Volt
Current	<span style="border: 1px solid black; padding: 2px;">313.60</span>	<span style="border: 1px solid black; padding: 2px;">295.85</span>	A
CosPhi	<span style="border: 1px solid black; padding: 2px;">0.00</span>	<span style="border: 1px solid black; padding: 2px;">0.00</span>	
Power	<span style="border: 1px solid black; padding: 2px;">65.23</span>	<span style="border: 1px solid black; padding: 2px;">61.54</span>	kW

Total power: 65.23 kW

Total IT power: 61.54 kW

PUE: 1.059

### STATUS

Doors	<span style="background-color: red; color: white; padding: 2px;">OPEN</span>	
Light switch	<span style="background-color: red; color: white; padding: 2px;">ON</span>	VESDA PREALARM
Lights	<span style="background-color: yellow; padding: 2px;">ON</span>	<span style="background-color: green; color: white; padding: 2px;">NORMAL</span>
Evap step 1	<span style="background-color: orange; padding: 2px;">OFF</span>	VESDA FIREALARM
Evap step 2	<span style="background-color: orange; padding: 2px;">OFF</span>	<span style="background-color: green; color: white; padding: 2px;">NORMAL</span>
Filter 1	<span style="background-color: green; color: white; padding: 2px;">NORMAL</span>	PLC BATTERY MODE
LEAK	<span style="background-color: green; color: white; padding: 2px;">NORMAL</span>	<span style="background-color: green; color: white; padding: 2px;">NORMAL</span>

# Customers Choose SGI

Optimized  
Configurations



Rapid Large Scale  
Deployment



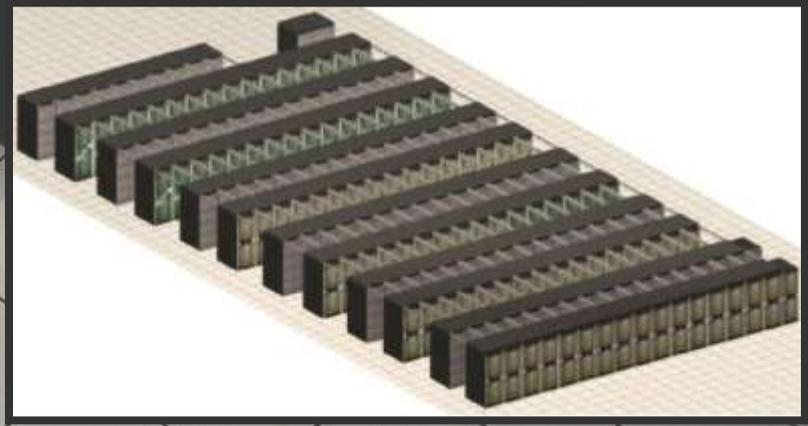
Integrated  
Solutions





**182 Racks**

- 1.3PF
- 4.1MW
- Largest IB Cluster



# Customers that are Breaking Barriers with ICE

"Our new SGI Altix ICE system was up and running the afternoon it arrived, and we began running benchmarks within 48 hours"  
Matthew Bate, Professor of Theoretical Astrophysics at the University of Exeter.

## New Mexico Computing Applications Center (NMCAC)



- Scientific research, broad range of codes
- 14,336 processor cores, 28TB memory, 172TB SGI® Infinite Storage 4500
- System operational in 48 hrs, will drive groundbreaking research for the state

## UFSC SINMEC



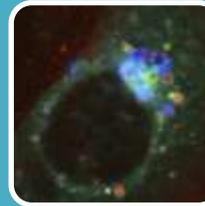
- Engineering analysis for oil & gas
- 512 processor cores, 512GB memory
- Ultra-deep water research

## Idaho National Laboratory



- Broad range of science codes , including computational chemistry, CFD
- 2,048 processor cores, 4TB memory
- System operational in 1-1/2 days

## IFREMER



- Scientific applications include WaveWatchIII, MARS, OPA, HYCOM, CORIOLIS
- 256 processor cores, augments Altix 450 system
- Oceanographic research

## National Oceanographic Centre Southampton, NOCS



- Oceanographic applications
- 768 processor-core, 1GB memory / core
- Runs multiple instances of the NEMO model , enabling NOCS to focus on their key scientific goals.

## GENCI (CINES)



- Research includes climatology, aeronauts research, energy, life & materials science.
- 23,640 processor cores, 4GB memory / core
- Resource for scientists & engineers, also connected to the RENATER French high-speed network and to EC infrastructures.

# Customers that are Breaking Barriers with ICE

"It quickly became clear to us that for the same number of processors, the performance of the SGI Altix ICE system was in a league of its own."

Henrik Diamant, Head of CFD, Honda Racing F1 Team

## HLRN



- North German advanced scientific research in diverse fields
- 25,000 processor cores, 90 TB memory and 2.3 PB of SGI InfiniteStorage
- Two phases, maximal performance of 312 Teraflops per second upon completion

## University of Exeter



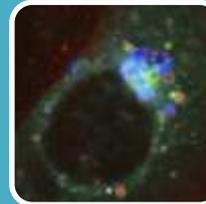
- Astrophysics research, custom in-house applications
- 1408 processor cores, 16TB memory
- Gas cloud simulation processing time dropped from 1 year to 6 weeks

## TOTAL Exploration



- World's Largest Commercial HPC System
- 2.3PF/17,920 processor cores, 41TB memory, 48 Altix XE Lustre file servers
- The world's largest hybrid system w/64 NVIDIA® Tesla™ S1070's (256 GPUs)!

## TU Dresden



- Bioinformatics, Chemistry/Pharma
- 512 processor cores, 2GB memory per core
- Drive research across multi-purpose scientific computing center

## Honda F-1 Racing Team



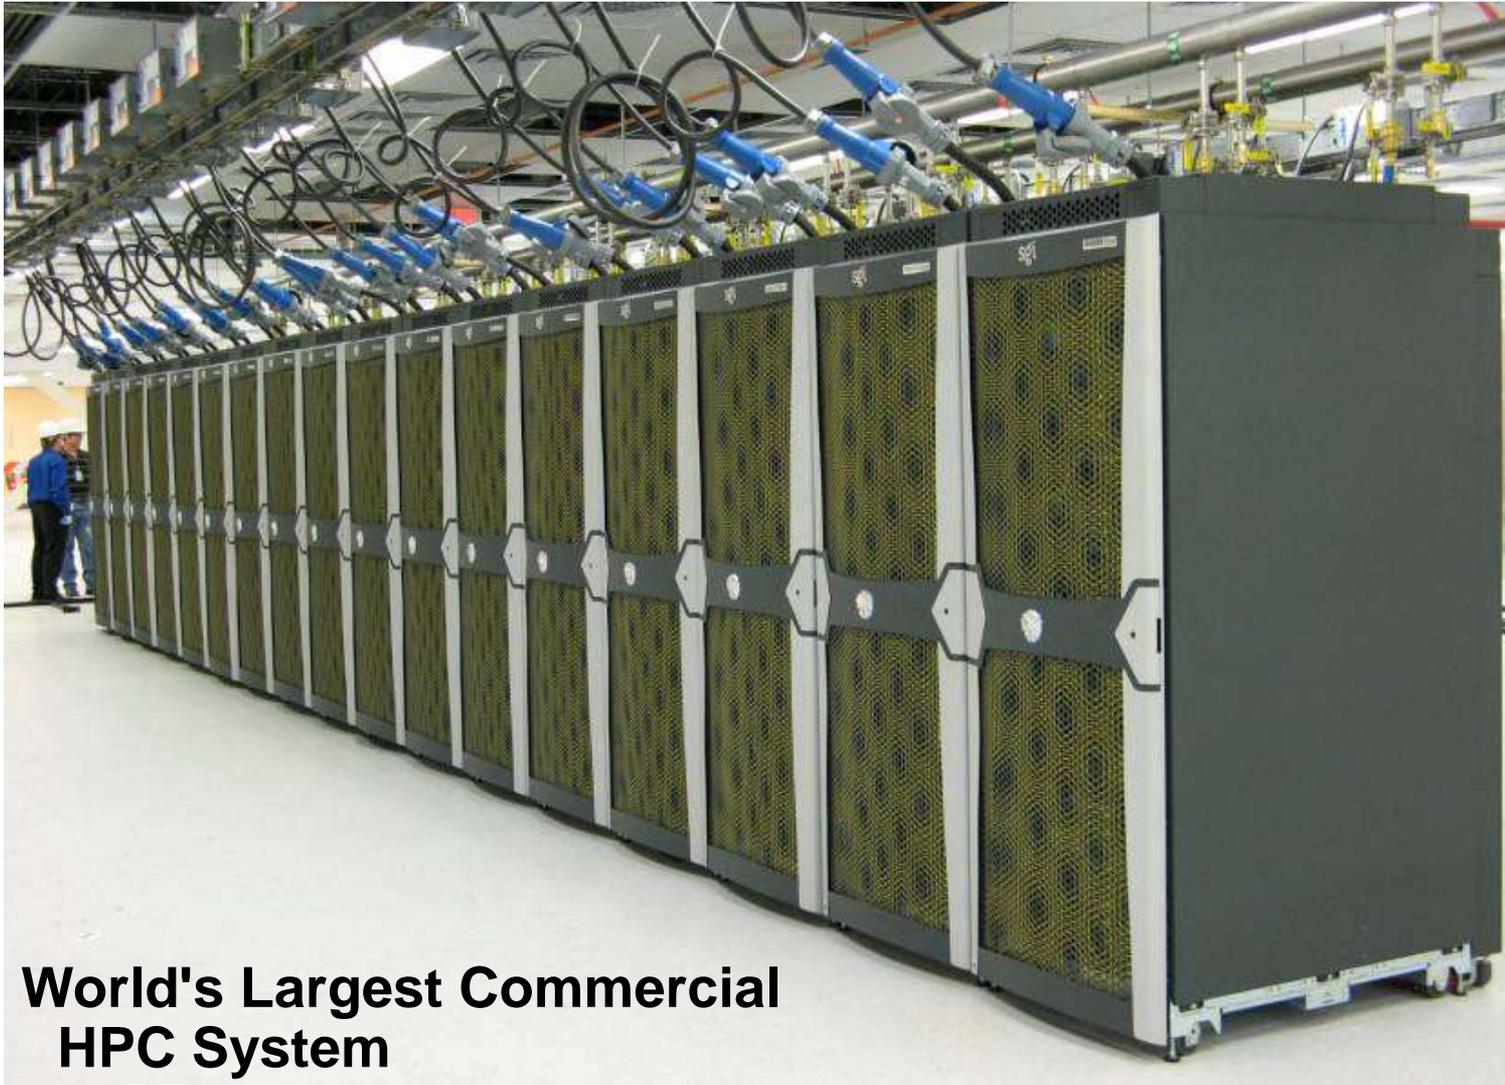
- CFD Whole F1 Car modeling
- 1024 processor cores, 24TB memory
- CFD calculations ran 18% faster on the SGI Altix ICE

## University of Arizona



- Wide spectrum of scientific and engineering research projects
- 1,520-cores, 2TB memory
- Hybrid computing solution, integrated with SGI® Altix® 4700 systems

# SGI ICE - Total



**World's Largest Commercial  
HPC System**

- 2.3PF
- 2.5MW

# Customers That Are... Breaking Barriers with ICE

## NASA-Ames



- The world's largest InfiniBand connected cluster system!
- 111,104 processor cores, 185.3TB memory and 40 PB SGI Infinite Storage
- Greater than 1PF sustained Linpack performance

## U.S. Army - TACOM



- CFD, impact analysis, and custom-code
- 1024 processor cores, 16TB memory
- Original "classic" and recent EX server with 64-blades, each with 32GB memory, and an IS4500 IO subsystem with 40TB scratch space

## NASA-Langley



- CFD and other sciences applications
- 3,072 processor cores, 2GB memory per core, 64TB storage
- Central compute resource for Langley facility

## University of Queensland



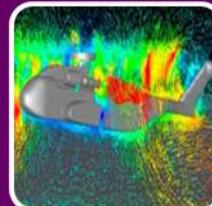
- Compute resource to drive earth studies
- Computational geodynamics, geo mechanics, earthquake physics and HPC Software.
- 512 cores, 2TB memory - Nexis 9000 (14.4TB)

## U. S. Navy – Naval Research Lab



- HPC applications – weather, CFD codes
- 1,500 processor cores, 6TB memory
- Variety of projects, including research in high speed I/O, Networks and Communications

## Sikorsky Aircraft



- CFD applications – LS-DYNA and Fluent
- 4 blades, 16 processor cores, augments an Altix XE cluster
- Sikorsky Aero Engineering, aerodynamics problems for a mix of helicopter programs



**GRACIAS...**