SGI UV: Bringing the Big Brain Computer to Galaxy

James Reaney
Senior Director, Research Markets
reaney@sgi.com

Agenda

- Scientific Computing Challenges
- SGI UV: The Big Brain for No-Limit Computing
- What does this mean for Galaxy?

Scientific Computing Challenges

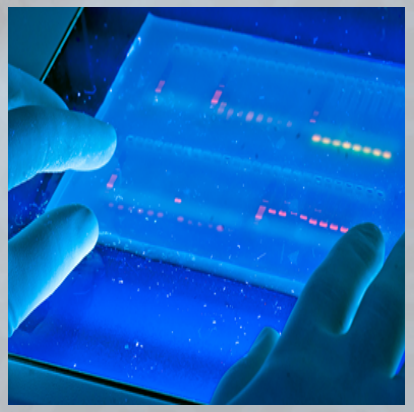

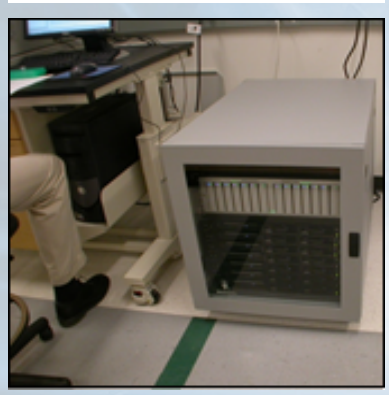

sgi

The Road to Exascale

An Insatiable Need For Computing

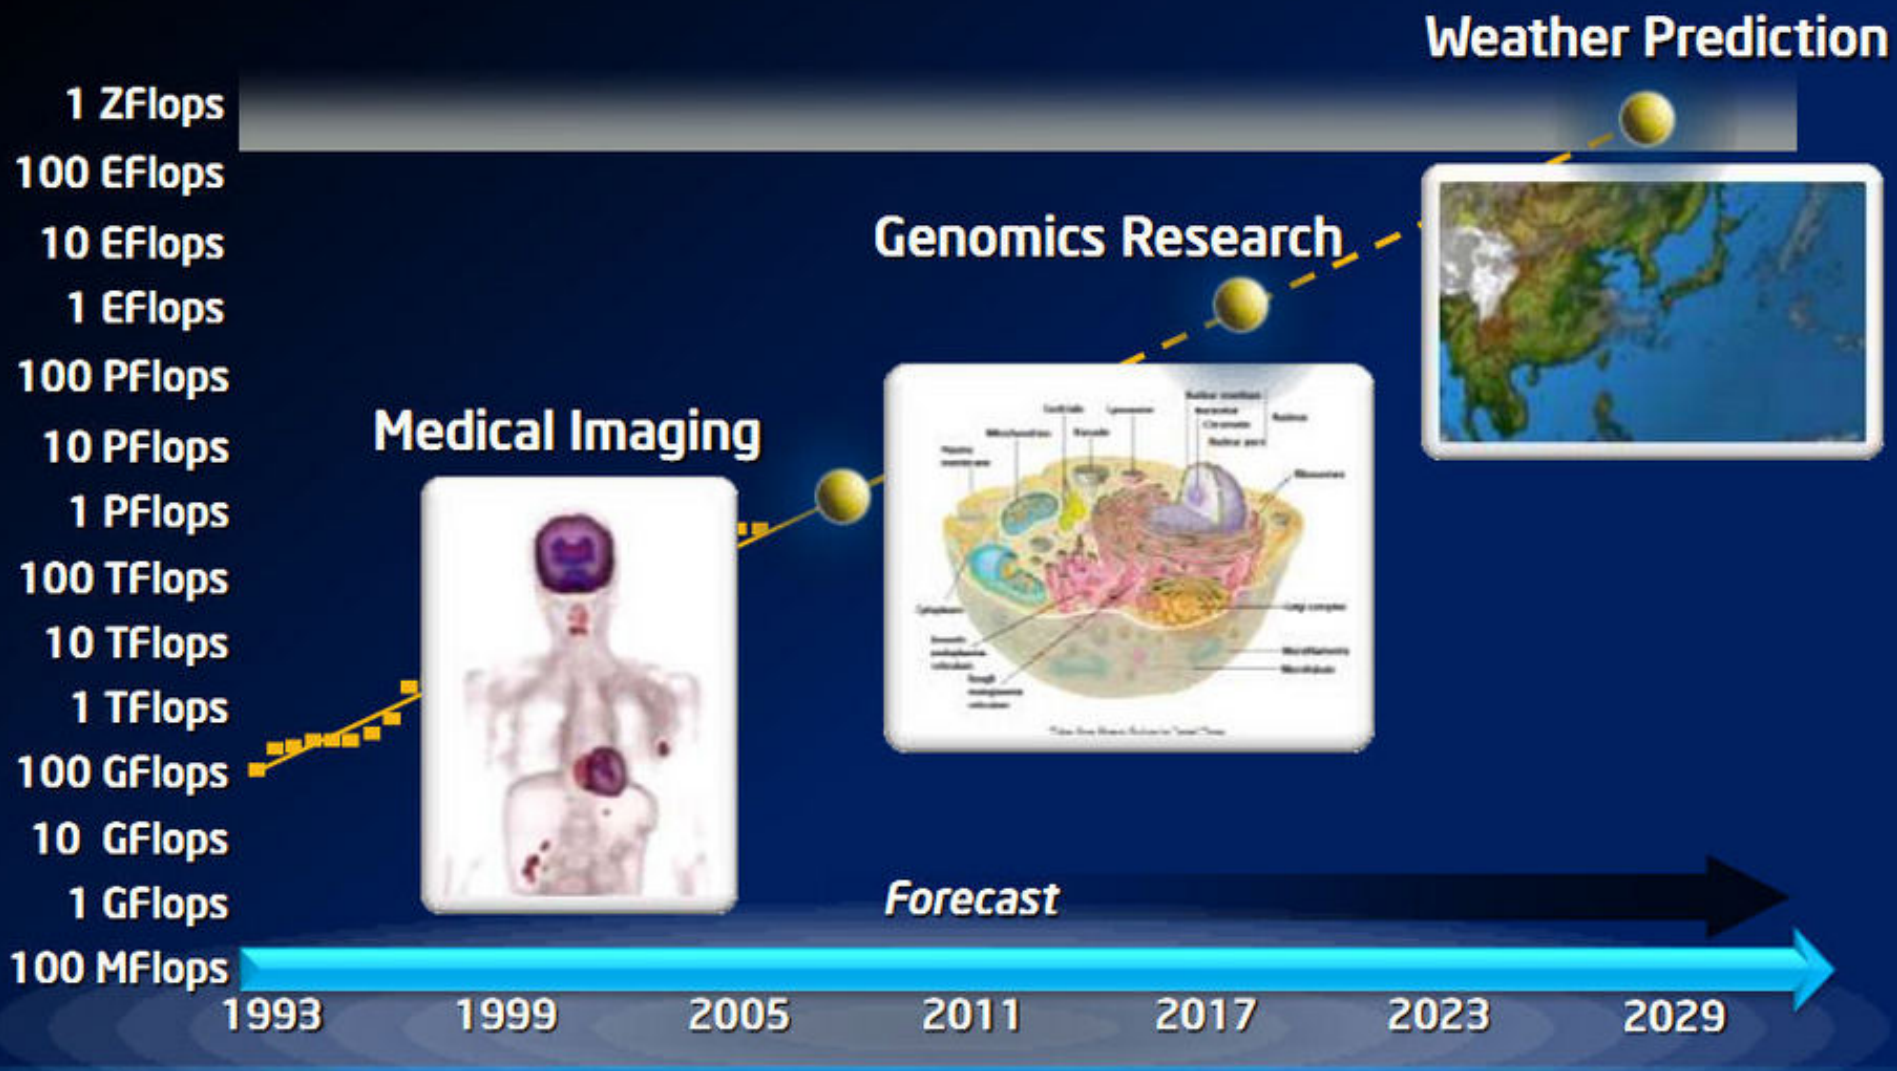

Exascale Problems Cannot Be Solved Using the Computing Power Available Today

Source: www.top500.org

intel

What does Exascale look like?

Systems	2009	2011	2015	2018
System Peak Flops/s	2 Peta	20 Peta	100-200 Peta	1 Exa
System Memory	0.3 PB	1 PB	5 PB	10 PB
Node Performance	125 GF	200 GF	400 GF	1-10 TF
Node Memory BW	25 GB/s	40 GB/s	100 GB/s	200-400 GB/s
Node Concurrency	12	32	O(100)	O(1000)
Interconnect BW	1.5 GB/s	10 GB/s	25 GB/s	50 GB/s
System Size (Nodes)	18,700	100,000	500,000	O(Million)
Total Concurrency	225,000	3 Million	50 Million	O(Billion)
Storage	15 PB	30 PB	150 PB	300 PB
I/O	0.2 TB/s	2 TB/s	10 TB/s	20 TB/s
MTTI	Days	Days	Days	O(1Day)
Power	6 MW	~10 MW	~10 MW	~20 MW

Exascale Requirements

Petascale Machine of 2010: TFLOP of Compute

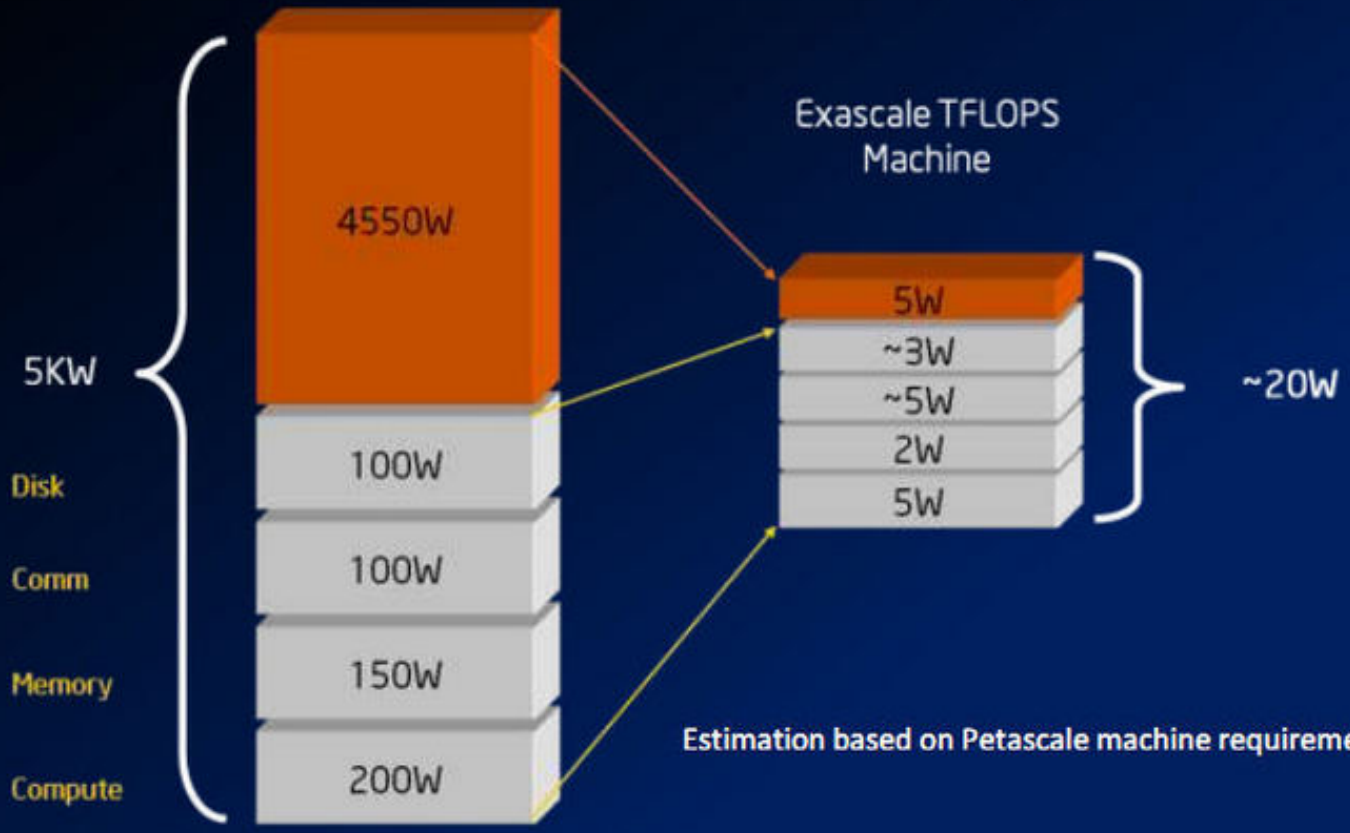

Compute 40x
Memory 75X
Comms 20x
Disk/Storage 33x
Other 900x

Estimation based on Petascale machine requirements circa 2010.

Visceral Focus on System Power Efficiency Improvement

Actual quotes from ILMN execs

- we are entering… "*the supersonic age of genomics*"
- the "demand for factory-scale sequencing of the human genome is *about to explode*"

  – Jay Flatley, CEO Illumina

- "Tens of thousands of samples are required…. You are trying to find needles in haystacks, and *you have to look at lots and lots of needles* to fundamentally understand the genetic basis of human disease"
- "Scientists are clearly finding clinical utility in the genome…. This creates a feedback loop where more discovery uncovers more clinical utility in the genome, which leads to an *increasing number of clinical researchers adopting these technologies*."

  – Christian Henry, SVP and CCO Illumina

“Needle in a Haystack” Analogy

Scientific Computing Needs

The right infrastructure can address a lot of these scientific computing challenges


SGI® UV™:
The Big Brain for
No-Limit Computing
sgi

SGI Technical Computing Servers

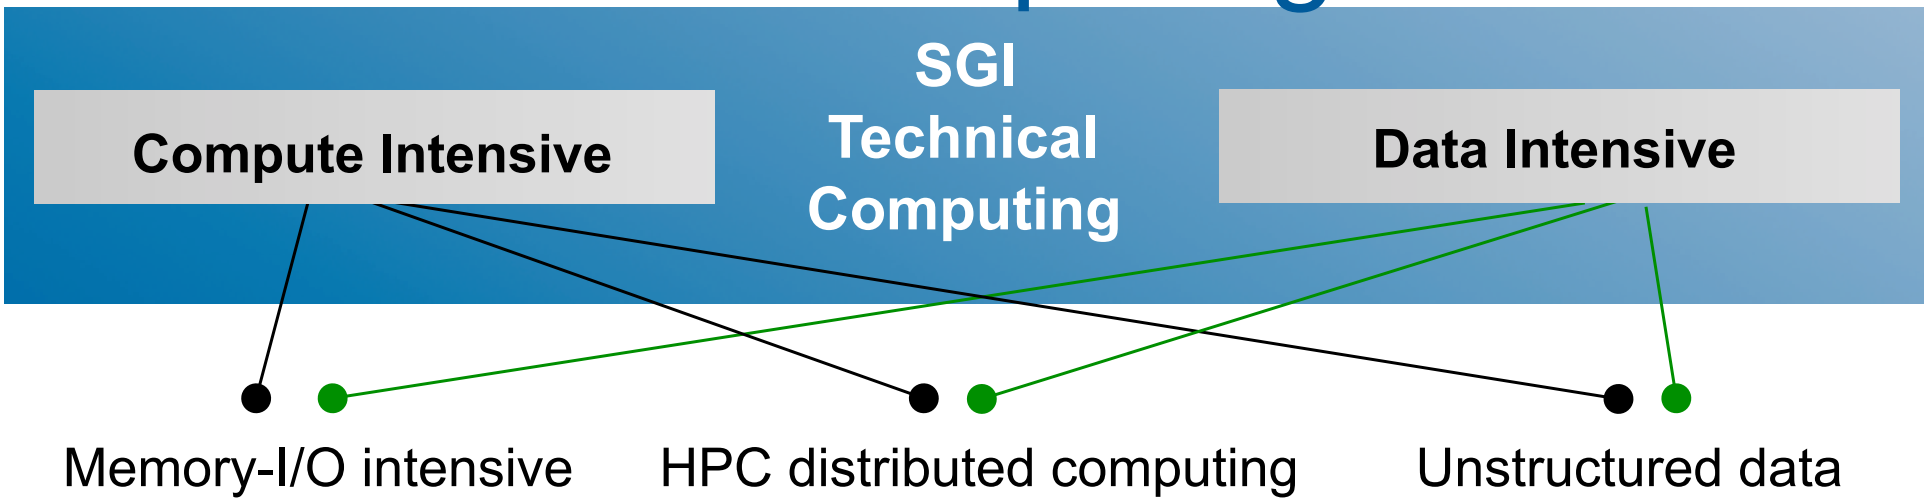

SGI® UV™
Large Shared-Memory, ccNUMA
Single System Image

SGI® ICE™ X
Very high-density
bladed cluster

SGI® Rackable™
Optimized cluster design

SGI UV 2: The Big Brain Computer

World-leading Capability for Data Intensive Work
Intel® Xeon® Processor E5-4600 Family

- World's Largest Shared-Memory System for Data-Intensive Applications
- Focus on Research, Not IT Problems
- No-Limit Computing, Built on Industry Standards

SGI UV 2: The Big Brain Computer

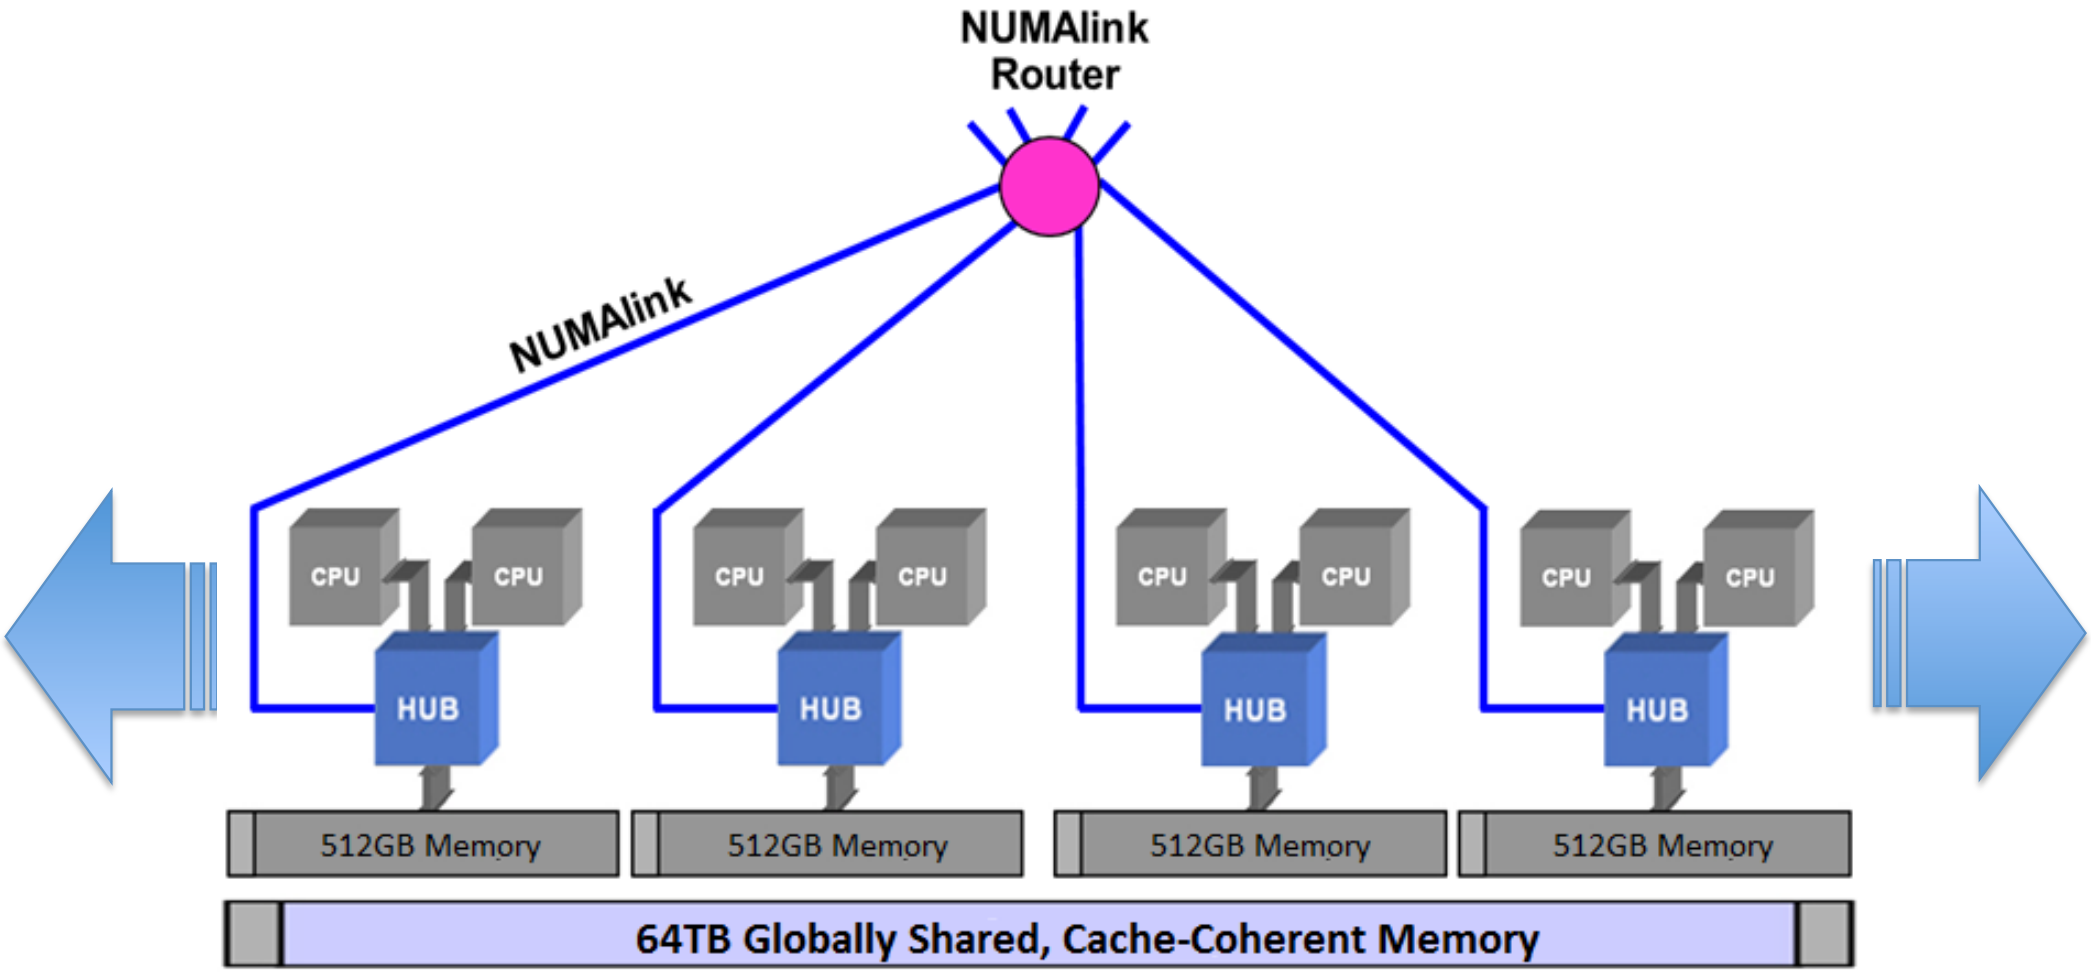

UV Comparison

Feature	Standard Scale-out Servers	SGI UV
Architecture Reference Terms	Scale Out, Clustered Distributed Memory	Scale Up, Single System Image Shared Memory
System Limit	16 cores, 512GB memory	4096 cores, 64TB memory
CPU	x86	x86 Intel® Xeon®
GPU or Intel Xeon Phi®	1 or 2, maximum	16 or 32
Memory, Storage, Networking, Operating Systems	Industry standard	Industry standard
Interconnect Fabric	Ethernet or InfiniBand	SGI NUMAlink™
External I/O Fabric	Ethernet, InfiniBand, or Fibre-Channel	Ethernet, InfiniBand, or Fibre-Channel

SGI UV2 IRU

Notes

- IRU: 10U 19” wide by 27” deep
- 8 blades / IRU: 16 sockets, 128 DIMM slots
- 1 or 2 CMCs in rear of IRU
- Three 12VDC 3037W, 200-240VAC or 277VAC N+1 Power Supplies
- Nine hot-pluggable, 119mm 12VDC axial cooling fans
- Two NUMAlink® backplane connections

UV2000 Compute Blade: 2 CPU, 16 DIMM IO expansion building blocks

PCIe / Base IO Bulkhead

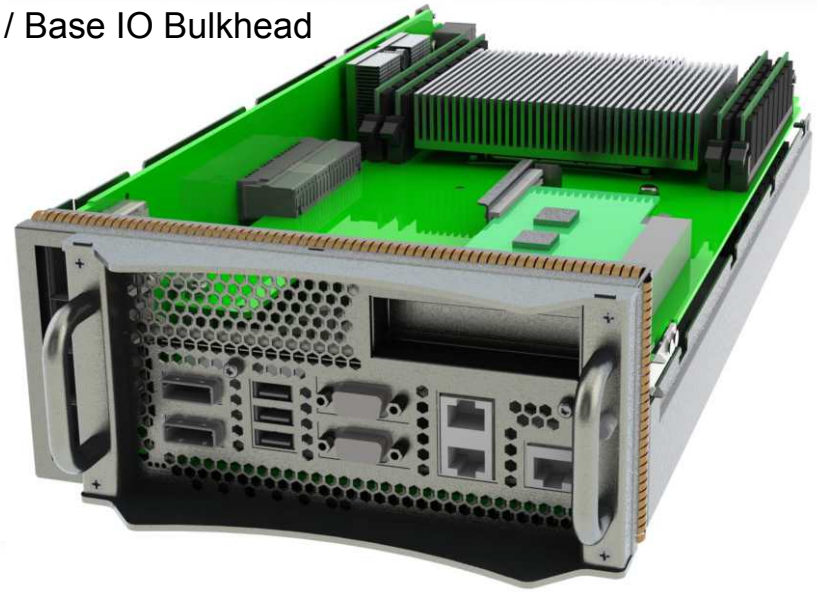

2 * PCIe
Low Profile

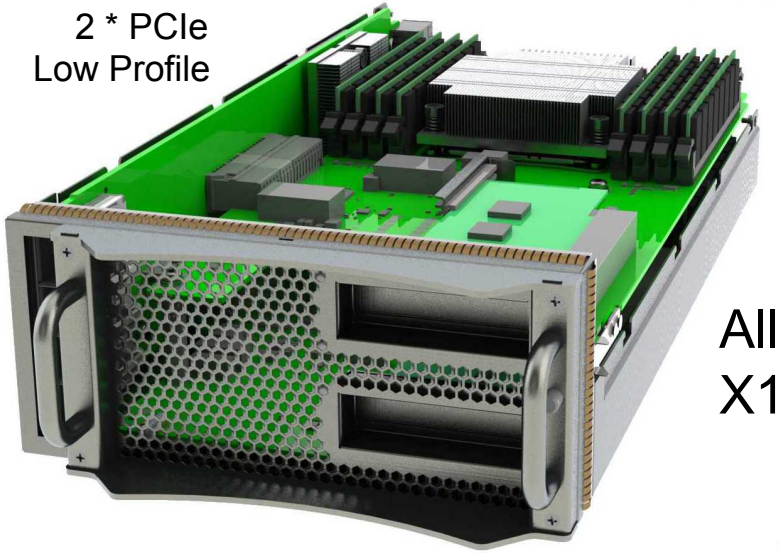

All PCIe Slots
X16, Gen3

2* 2.5” HDD/SSD

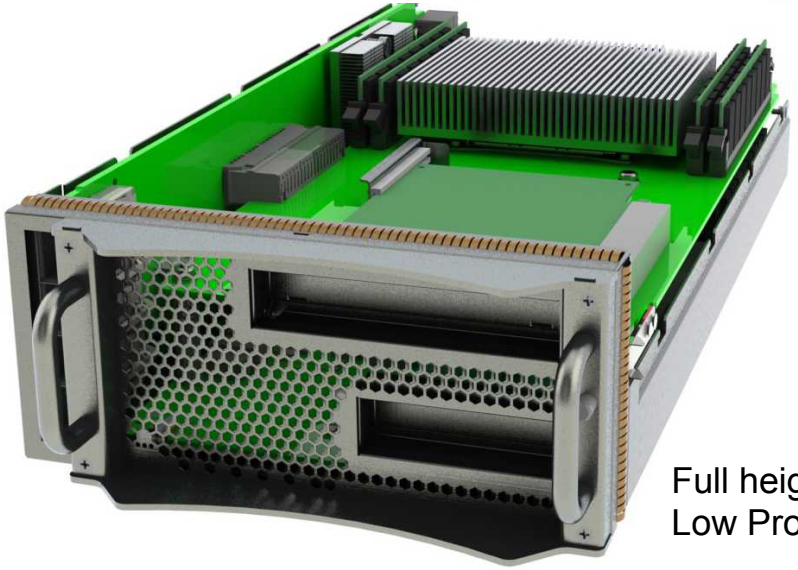

Full height half depth PCIe
Low Profile PCIe

SGI UV 2

Start small and grow … or start big.

Modular Design, Configuration Flexibility Supports GPU, Intel® Xeon® Phi™ and NVIDIA® Tesla® K20

UV 20

16-32 core
32GB-1.5TB

16-128 core
32GB-4TB

UV 2000

64-512 core
256GB-16TB

256- 4096 core
Up to 64TB

UV: World's Biggest Compute Node

Scale-Up In Memory Compute Platform

"Big Brain" Computing

Cache-Coherent, Global Shared Memory System

Single Operating System: Unmodified Standard Linux

4096 Xeon Cores, 64TB Shared Memory

128k cores, 8 PB Addressable Memory

NUMAlink® 6 Interconnect Topology

Designed for Research Computing

Enables Science not possible with commodity servers

Open System Design: Runs standard Linux & Apps

sgi

SGI UV 2

Open platform – Intel® Xeon® processor E5-4600 product family and off-the-shelf Linux®.

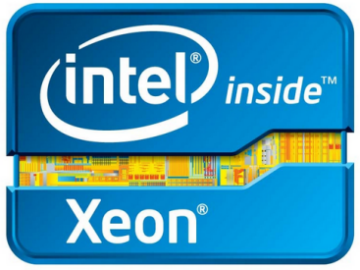

- Intel® Xeon® Processor E5-4600 product family
- SGI is major Linux contributor
  - Red Hat® Enterprise Linux 6
  - SUSE® Linux® Enterprise Server 11
- Standard Management, Storage Interfaces
  - SGI Infinite Storage Solutions
  - Interface with common management software schemes

SGI Performance Suite

SGI Accelerate

- Accelerate applications with optimized software libraries and tools
- Tune applications without recompiling
- Optimize performance with specialized algorithms

SGI MPI

- SGI's scalable, high performance MPI environment
- More than just an MPI library
- Includes runtime MPI acceleration, profiling, checkpoint/restart and more

SGI REACT

- Hard real-time performance for Linux
- Only hard real-time solution for standard distribution Linux
- No custom Linux kernel needed

SGI UPC

- SGI's optimized Unified Parallel C compiler environment
- Scales across SGI Numalink and InfiniBand

Top500 performance for standard distribution Linux

Workload Management Solutions

	Key Features
Altair PBS Professional 	Commercial-grade, highly scalable HPC workload management and job scheduling solution supporting green computing, security certification, and policy driven job scheduling
Adaptive Computing Moab 	Multi-dimensional, policy-based workload management system that integrates scheduling, managing, monitoring, and reporting of cluster workloads
Simple Linux Utility for Resource Management 	Open source resource management tool developed by Lawrence Livermore Labs

What does this mean for Galaxy?

sgi

Sounds great. How do I get one?

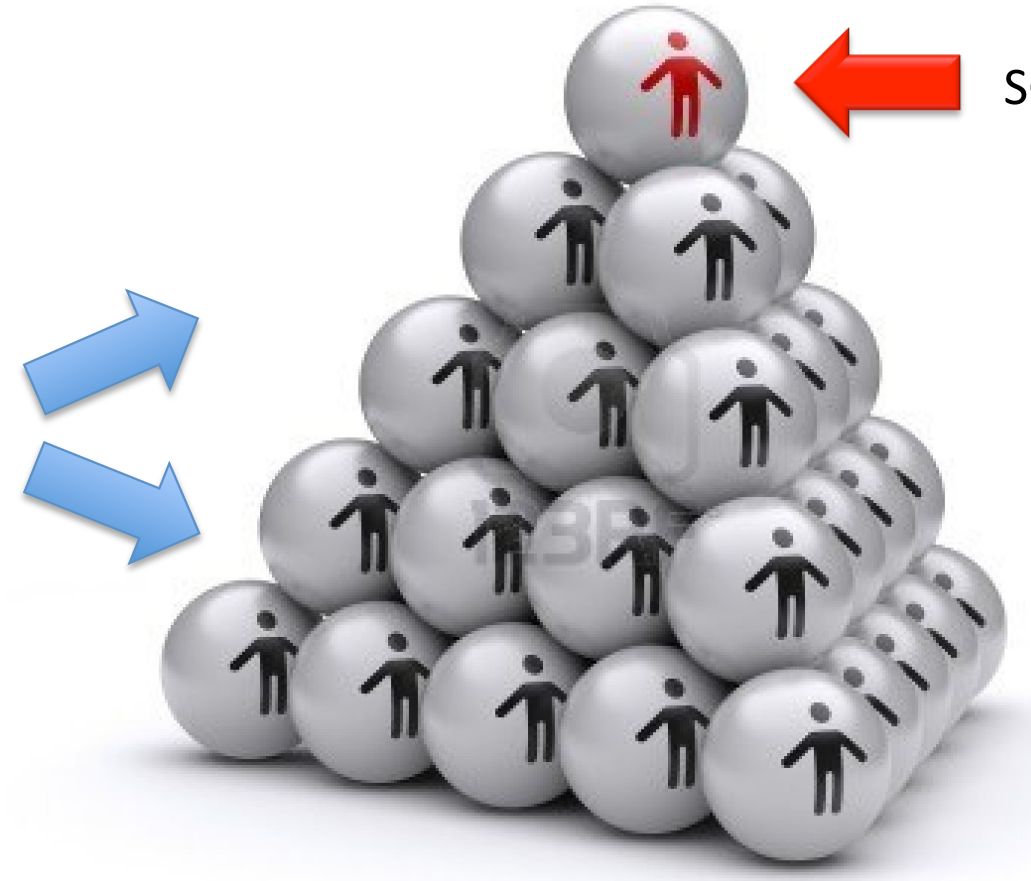

SGI Research Customers

Everyone Else!

Go work at one of these places

TECHNICAL UNIVERSITY OF DENMARK DTU

sgi

Contact SGI sales representative

Slipstream Appliance: Galaxy Edition

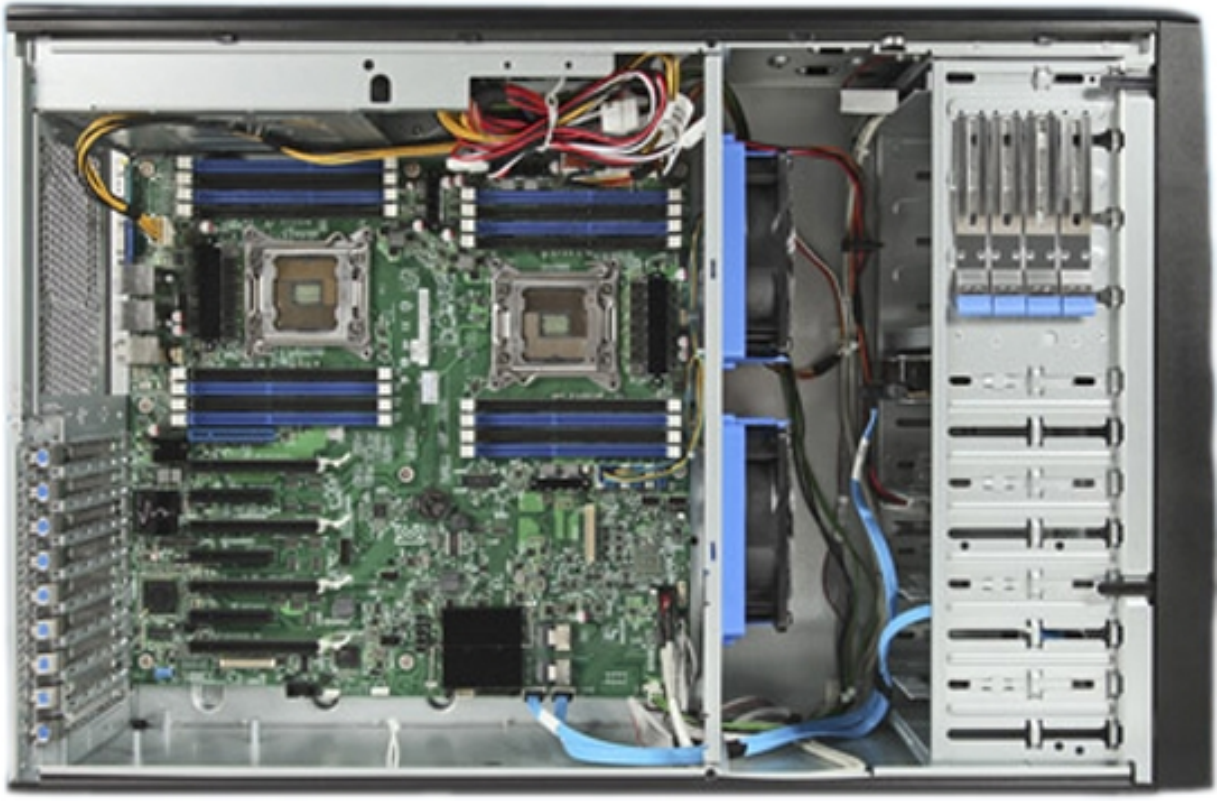

- Intel® Xeon® Ivy Bridge E5-2600v2
- 512GB to 1TB ECC RAM
- 2x120GB Data Center SSD
- 8x4TB SAS 6Gbps Enterprise HDD
- Dual-port 10GbE standard

- SGE, Galaxy, underlying tools are all pre-installed and tuned for optimal performance
- Pre-loaded datasets for rapid deployments and ease-of-use
- Auto-update support

- Installation & setup included
- 3YR Warranty Standard
- Premium Galaxy support and custom services available

Galaxy

BIOTEAM
Enabling Science

intel®

sgi

sgi

Slipstream Appliance: Galaxy Edition

Galaxy

BIOTEAM
Enabling Science

intel

sgi

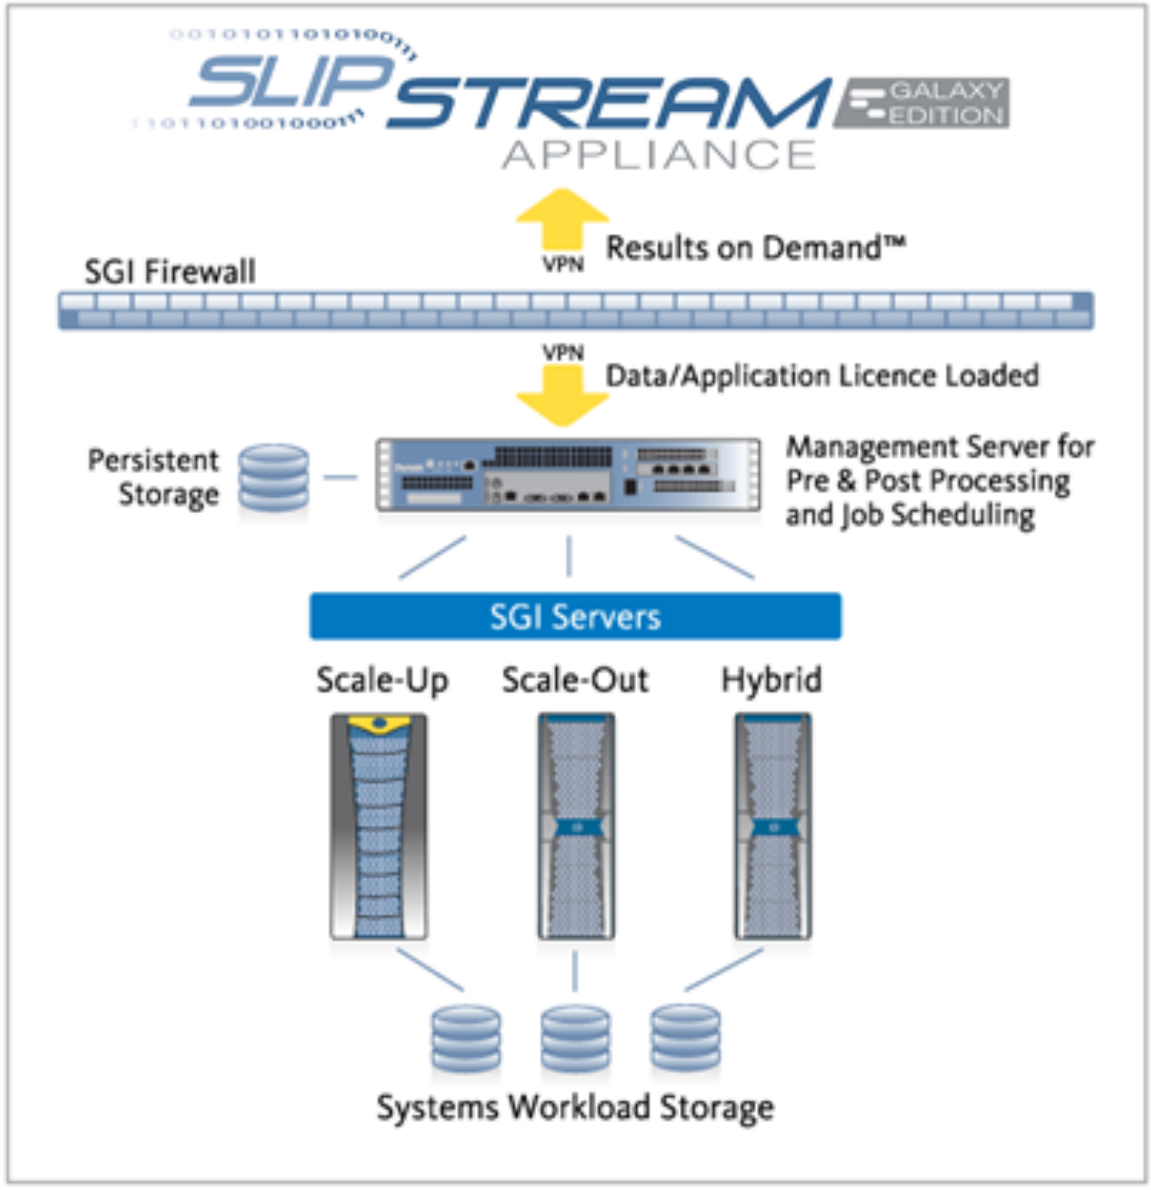

sgi

MORE
HEADROOM

sgi