



HIGH PERFORMANCE COMPUTING FROM SUN

Update for IDC HPC User Forum,
Norfolk, VA – April 2008

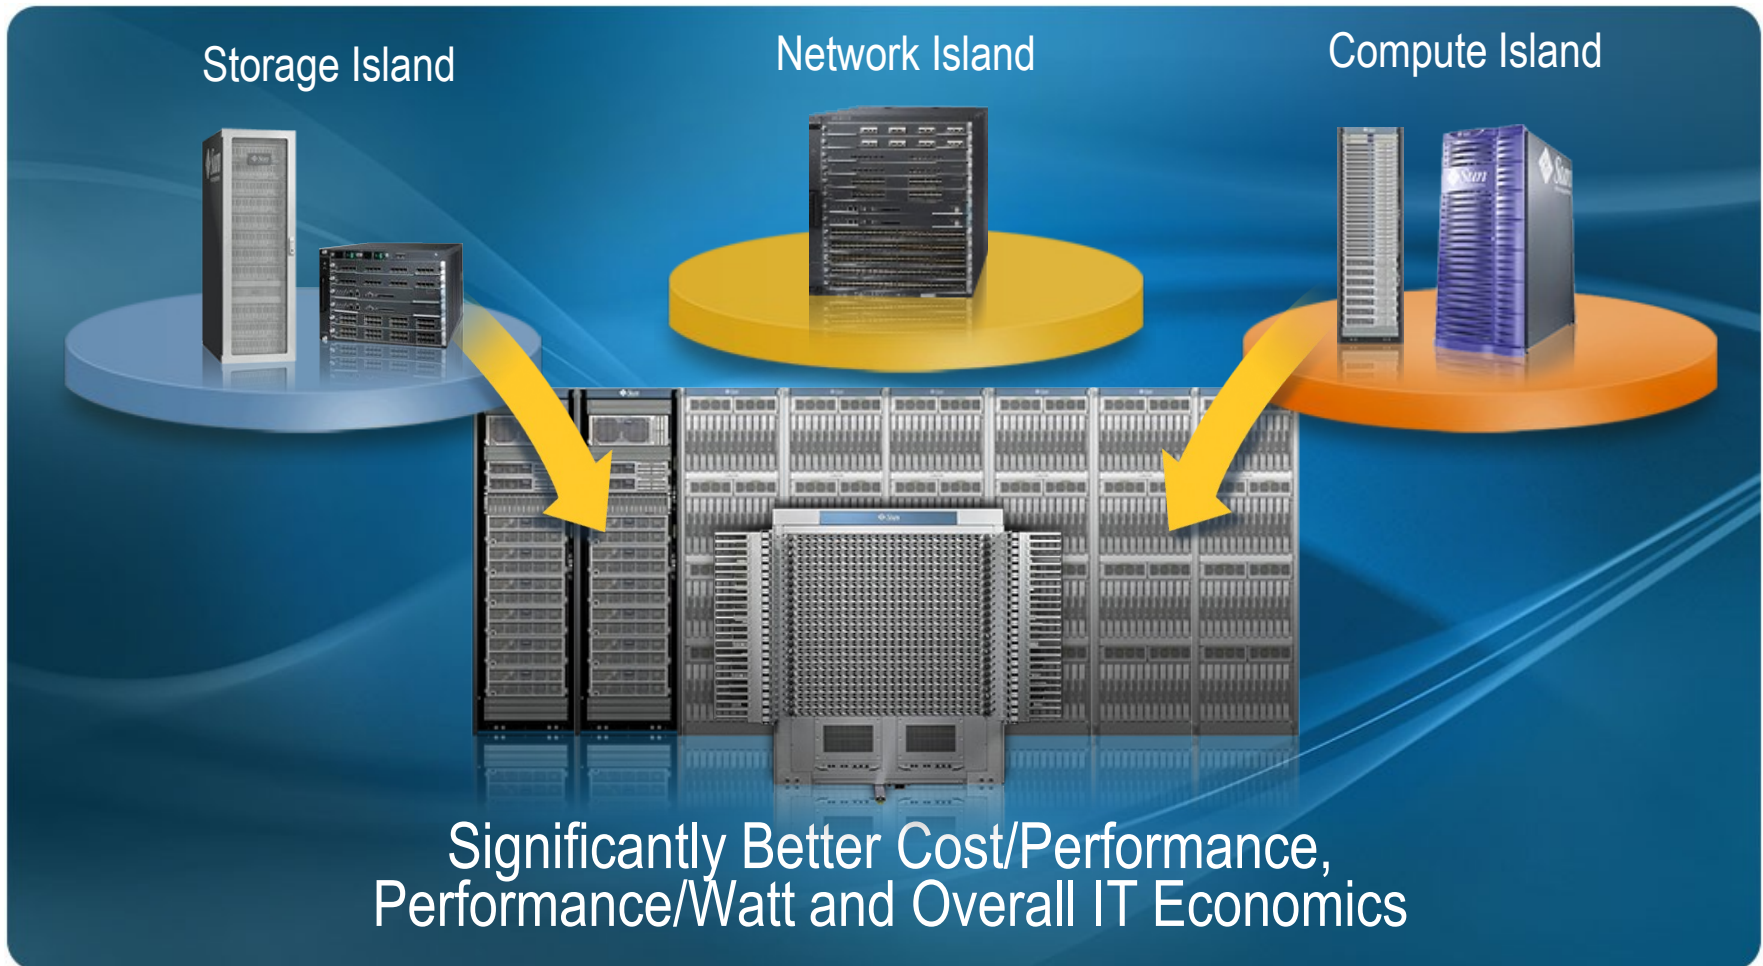
Bjorn Andersson

Director, HPC and Integrated Systems
Sun Microsystems



Sun Constellation System

Integrating the Islands



TACC: World's Top Supercomputer for Open Science



Open for Business!
 (see portal.teragrid.org)

Sun Fire X4500

- 1.7 Petabytes
- 72 GB/sec total bandwidth



Sun Fire X4600

- 25 Systems
- 800 Cores



Sun Blade 6048

- 3,936 blades
- 15,744k CPUs
- 62,976 Cores
- 125 TB / RAM

Sun Datacenter Switch 3456

- Dual Redundant
- 110 Tb/sec bisectional bandwidth





TeraGrid™ User Portal

[Login](#)

Welcome, Guest User

[Home](#) [Resources](#) [Documentation](#) [Training](#) [Consulting](#) [Allocations](#)
[Systems Monitor](#) [HPC Queue Prediction](#) [Science Gateways](#) [Data Collections](#) [User Responsibilities](#)

TeraGrid Systems Monitor

[Refresh](#)


High Performance Computing Systems

Name	Institution	System	CPUs	Peak TFlops	Memory TBytes	Disk TBytes	Load	Jobs*†		
								R	Q	O
Ranger	TACC	Sun Constellation	62976	504.00	123.00	1730.00		226	94	111
Abe	NCSA	Dell Intel 64 Linux Cluster	9600	89.47	9.38	100.00		44	17	50
Lonestar	TACC	Dell PowerEdge Linux Cluster	5840	62.16	11.60	106.50		83	182	1
Queen Bee	LONI	Dell Intel 64 Linux Cluster	5440	50.70	5.31	100.00		41	1	0
Big Red	IU	IBM e1350	3072	30.60	6.00	266.00		237	0	718
BigBen	PSC	Cray XT3	4136	21.50	4.04	100.00		24	201	26
Blue Gene	SDSC	IBM Blue Gene	6144	17.10	1.50	19.50		2	0	12
Tungsten	NCSA	Dell Xeon IA-32 Linux Cluster	2560	16.38	3.75	109.00		113	578	69
DataStar p655	SDSC	IBM Power4+ p655	2176	14.30	5.75	115.00		46	421	81
TeraGrid Cluster	NCSA	IBM Itanium2 Cluster	1744	10.23	4.47	60.00		86	26	0
Lear	Purdue	Dell EM64T Linux Cluster	1024	6.60	2.00	28.00		64	9	1
Cobalt	NCSA	SGI Altix	1024	6.55	3.00	100.00		45	600	0
Frost	NCAR	IBM BlueGene/L	2048	5.73	0.51	6.00		8	1	0
TeraGrid Cluster	SDSC	IBM Itanium2 Cluster	524	3.10	1.02	48.80		34	126	16
DataStar p690	SDSC	IBM Power4+ p690	192	1.30	0.88	115.00		4	37	12
TeraGrid Cluster	UC/ANL	IBM Itanium2 Cluster	128	0.61	0.24	4.00		1	0	0
NSTG	ORNL	IBM IA-32 Cluster	56	0.34	0.07	2.14		0	0	1
Rachel	PSC	HP Alpha SMP	128	0.31	0.50	6.00		8	3	0
Total:			108812	840.98	183.02	3015.94		1066	2296	1098

The Sun Constellation System

The World's Most Scalable General Purpose Computer

- New levels of performance and scalability
 - > Up to 2 PetaFLOPs
 - > Up to 1.8 PetaBytes RAM
 - > Up to 0.7 ExaBytes disk
- Open industry standards
 - > Linux, Solaris, OpenMPI, open InfiniBand interfaces and management
 - > X64 computing architecture
 - > InfiniBand DDR interconnect
- New levels of efficiency
 - > Provides a 6:1 reduction in physical ports and cables
 - > 20% smaller footprint than competition
 - > Eliminates 100s of discrete switching elements
- New levels of reliability
 - > Dramatically reduced complexity: 6x fewer cables, up to 50% fewer racks




Scale with Switches

Examples with Sun Datacenter Switch 3456


1 Core Switch

Servers: 3,456
PFLOPS: 0.5




2 Core Switches

Servers: 6,912
PFLOPS: 1.0




3 Core Switches

Servers: 10,368
PFLOPS: 1.5



4 Core Switches

Servers: 13,824
PFLOPS: 2.0



Previewed at SC07:
1RU 3x24 Switch



Scale with Blades

Examples with Sun Blade 6000 Family
(Available with AMD, Intel or SPARC CPUs)



Single
Blade

10 x Scaling

Sun
Blade 6000

48 x Scaling

Single Rack
768 Cores
6.6 TFLOPS
Peak
(with QC)

~300 x Scaling


4 Core
Switches
13,824 Servers
2 PFLOPS
Peak

Scale with Storage

Example with Sun Fire X4500 and Lustre file system



- Sun Fire X4500 - Industry's first data server
- Best server data throughput and storage density
- Runs Lustre parallel file system
- Transparent connection to Archive
- Standard platform and common systems management capabilities



The Sun HPC Cluster Platform

Open, Seamless and Comprehensive



Sun Services

Insight	Visualization, Workstations, Thin Clients, Remote Access
Developer	Compilers, Debuggers, Optimization Tools, Libraries
Management	Workload, Systems and Cluster Management
OS	Linux or Solaris
Interconnect	InfiniBand or Ethernet
Storage/Archive	Cluster Storage, Backup, Archive, File Systems, HSM
Compute	Racks or Blades Variety of CPU architectures



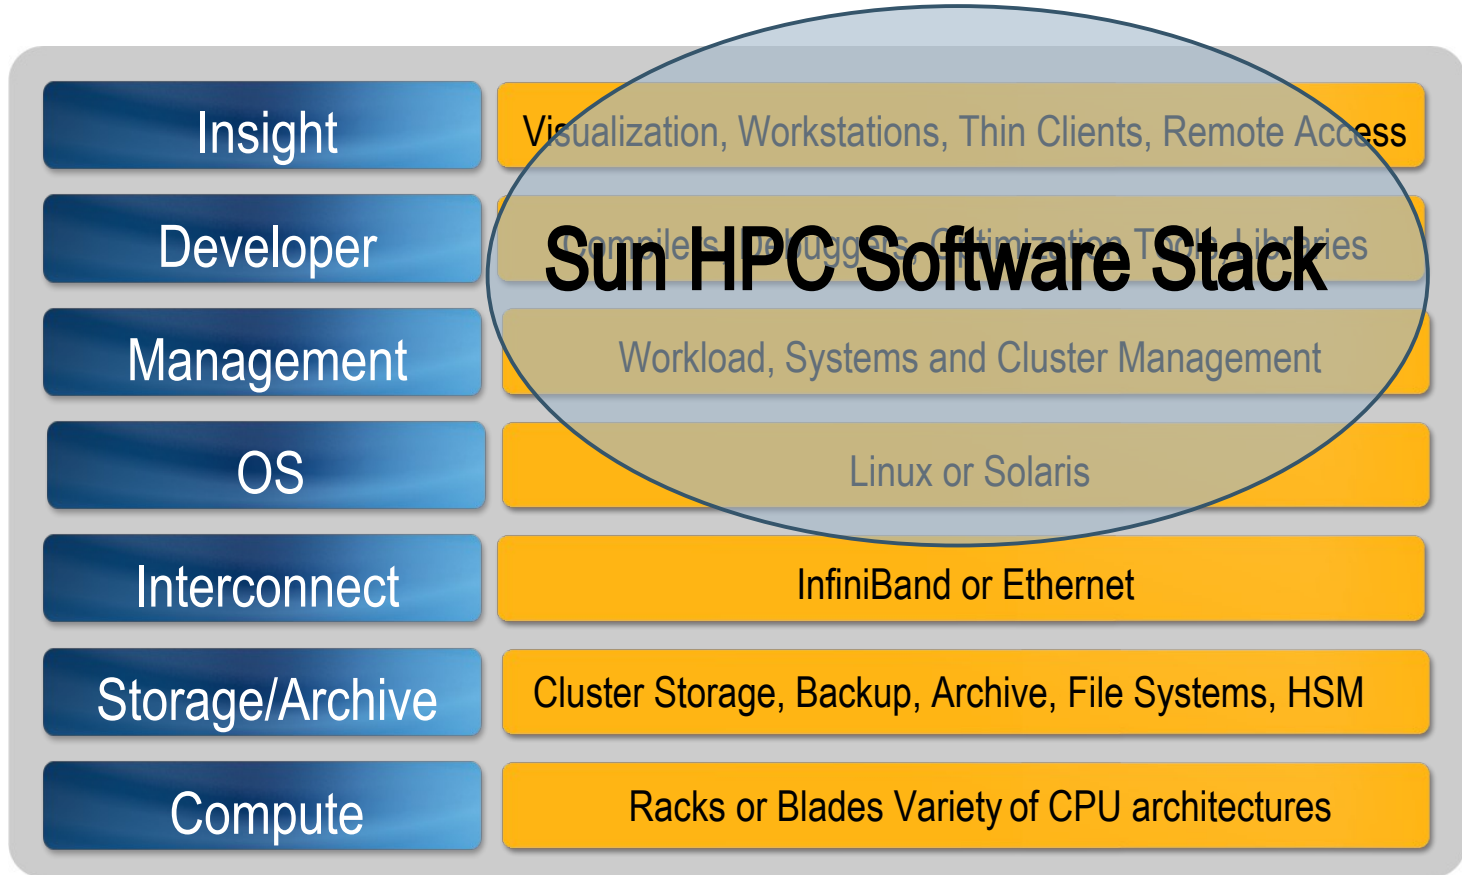
Sun Customer Ready

The Sun HPC Cluster Platform

Open, Seamless and Comprehensive



Sun Services



Sun Customer Ready

Lustre™ File System

World's Largest Network-Neutral Data Storage and Retrieval System


- The world's most scalable parallel filesystem
- 10,000's of clients
- over 50% of Top30 run Lustre
- 15% of Top500 run Lustre
- Open source, multi platform

Lustre User Group 2008

Sonoma, CA, April 28-30, 2008

Registration: <https://www.regonline.com/builder/site/Default.aspx?eventid=181696>

Questions: LUG2008@sun.com



Project Blackbox

Sun Modular Datacenter S20



High Capacity

- 820 CPU's, 3,280 Cores
 - > Sun Blade 6000 systems and Sun Fire X2200 servers
- 2,240 cores and 17,920 compute threads!
 - > UltraSPARC T2 servers
- 3 petabytes of disk
 - > Sun Fire X4500 servers

High Performance


- 31 Teraflops Peak
 - > 210 Sun Blade 6000 nodes
- ~170,000 web ops/sec
 - > 78 x Sun Fire T2000 servers



Stanford Linear Accelerator Center (SLAC)

Sun Constellation System

Reference Designs



Storage

Compute

Networking



HIGH PERFORMANCE COMPUTING FROM SUN.

Bjorn Andersson

sun.com/hpc – Main HPC landing page
hpc.sun.com – HPC Community Portal
blogs.sun.com/hpc – Blog with HPC news
Radio HPC – Sun HPC podcasts on iTunes

