

Jaguar: Petascale Science Delivered

Presented by

Buddy Bland

Leadership Computing Facility Project Director
National Center for Computational Sciences



Jaguar: World's most productive scientific computer

Designed for science from the ground up



Peak performance	2.595 PF
System memory	362 TB
Disk space	10.7 PB
Disk bandwidth	240+ GB/s
Interconnect bandwidth	532 TB/s


What does the system look like?

Jaguar combines the existing 263 TF Cray XT4 system at ORNL's NCCS with an upgraded 2,332 TF Cray XT5 to create a 2.6 PF system


System attribute	XT5	XT4
AMD Opteron processors	37,544	7,832
Node architecture	Dual socket SMP with 6-core processors	Single socket with 4-core processors
Memory per core/node (GB)	2/16	2/8
Total system memory (TB)	300	62
Disk capacity (TB)	10,000	750
Disk bandwidth (GB/s)	240	44
Interconnect	SeaStar2+ 3D torus	SeaStar2+ 3D torus

Jaguar's Cray XT5 nodes

- 12-core SMP
- Two 2.6 GHz AMD Opteron "Istanbul" six-core processors
- 125 Gflops per node
- Four 4 GB DDR2-800 DIMMS with four empty DIMM slots per node
- OpenMP support



Jaguar's Cray XT5 blades



- 4 nodes per blade
- 500 GF
- 64 GB
- 4 network connections

24 Cray XT5 blades make a Jaguar cabinet

One Jaguar XT5 cabinet:

- 11.9 TF
- 192 Opteron processors
- 1,152 Opteron cores
- 1,536 GB memory
- 96 nodes
- 1•4•24 section of 3D torus

200 cabinets complete the XT5 side of the system



Compute partition

- Nodes: 18,688
- Opteron processors: 37,376
- Processor cores: 224,256
- Peak TF: 2.3
- Memory: 300 TB
- Interconnect: 25•32•24 3D torus

Service and I/O partition

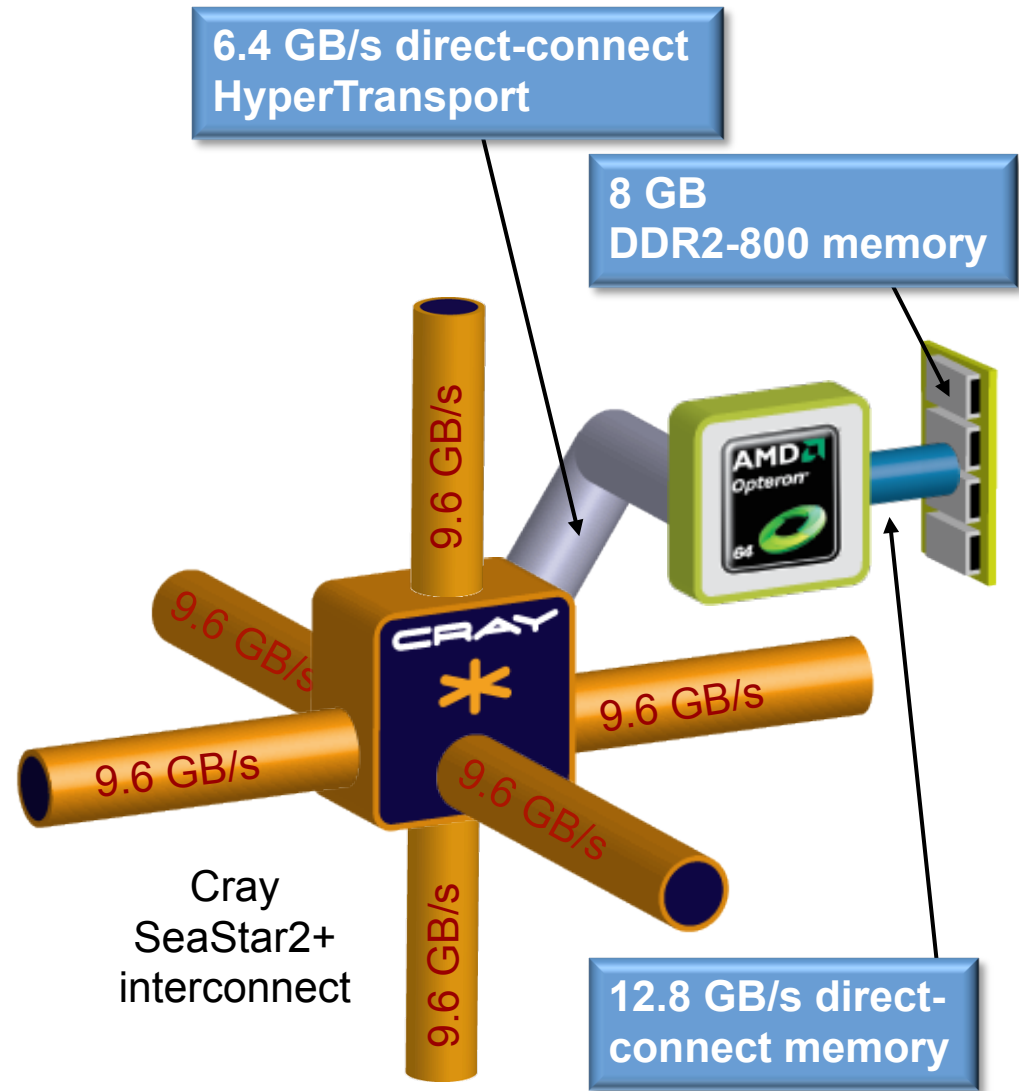
- Nodes: 256
- InfiniBand cards: 192 DDR
- Memory per SIO node: 8 GB
- 10 Gigabit Ethernet: 10

Power and cooling


- 480 V power supplies
- R134a refrigerant
- Single large fan per cabinet

How is the XT4 side different?

- Basic node design is the same as XT5, except it has one AMD Opteron quad-core processor per node instead of two
- Same 2 GB of DDR2-800 memory per core
- Same Seastar2+ interconnect



The XT4 board has four nodes



- 4 nodes per blade
- 134.4 GF
- 32 GB
- 4 network connections

Memory DIMMS
4 × 2 GB per socket

Voltage regulators

And 24 Cray XT4 blades make a Jaguar cabinet




One Jaguar XT4 cabinet:

- 3.2 TF
- 96 Opteron processors
- 384 Opteron cores
- 768 GB memory
- 96 nodes
- 1•4•24 section of 3D torus

ECOphlex liquid cooling

- Liquid-cooled design exhausts heat to R134a before it leaves the cabinet. **Replaces 100 CRAC units!**
- Saves about 900 kW of power in air movement alone



- Phase change of liquid to gas removes heat much more efficiently




480 V power saves money

- **Power supply efficiency**
 - Worked with Cray to select highest-efficiency power supplies
- **480 V power to the cabinet saves the inefficiency of converting from 480 V to 208 V to 48 V**
- **Keeping the voltage high saved \$1M in installation and \$500K in electrical costs**

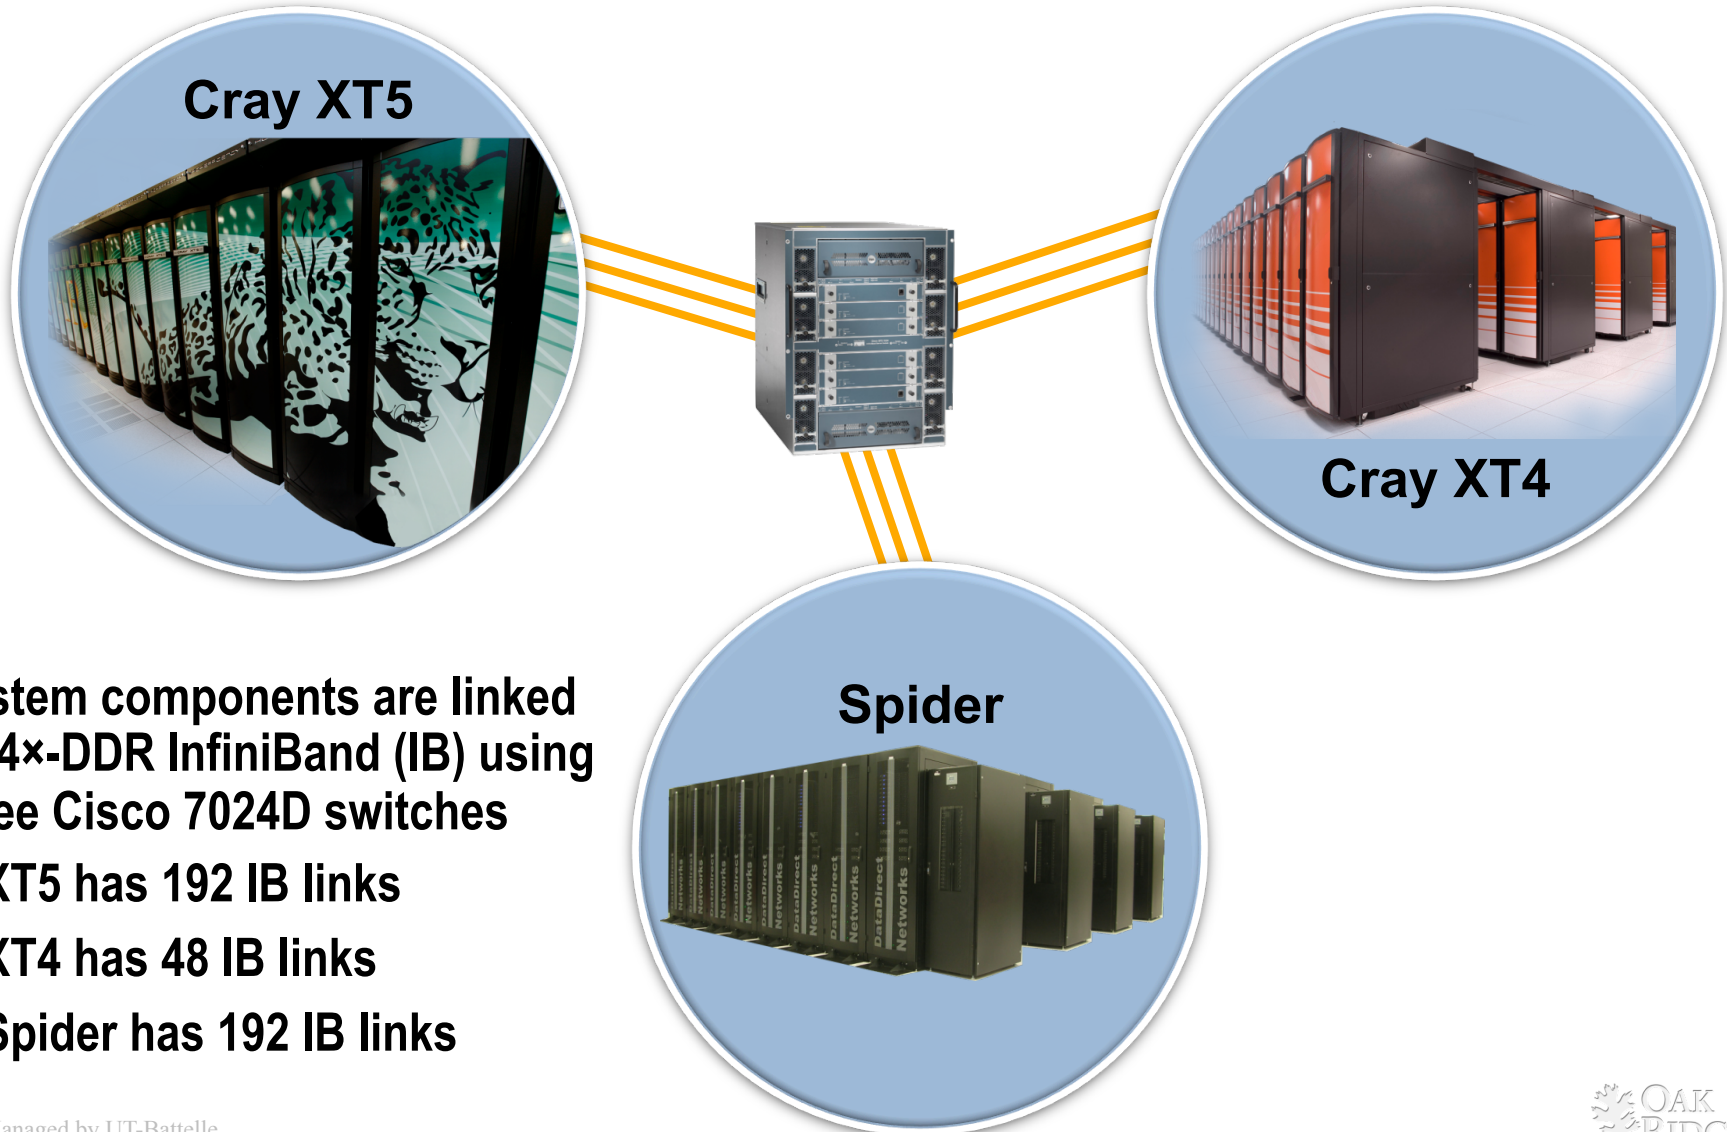


Centerwide file system



- “Spider” is being installed to provide a shared, parallel file system for all systems
 - Based on Lustre file system
- Bandwidth of more than 240 GB/s
- More than 10 PB of RAID6 capacity
 - 13,440 1 TB SATA drives
- 192 storage servers
 - 3 TB of memory
 - 14 TF
- Available from all systems via our high-performance scalable I/O network
 - More than 3,000 InfiniBand ports
 - More than 3 miles of cables
 - Scales as storage grows
- Engineered for high availability

Jaguar combines the 2.3 PF Cray XT5 with the existing 263 TF Cray XT4



Cray XT5

Cray XT4

Spider

System components are linked by 4x-DDR InfiniBand (IB) using three Cisco 7024D switches

- XT5 has 192 IB links
- XT4 has 48 IB links
- Spider has 192 IB links

Contact

Arthur S. (Buddy) Bland

**Leadership Computing Facility Project Director
National Center for Computational Sciences**

(865) 576-6727

blandas@ornl.gov

