

# Leadership Computing at ORNL Current Status and Plans

**LEADERSHIP  
COMPUTING FACILITY**  
NATIONAL CENTER FOR COMPUTATIONAL SCIENCES



*presented by*  
Buddy Bland  
Project Director

SOS 11, Key West  
June 12-14, 2007

Oak Ridge National Laboratory  
U.S. Department of Energy

# Since the last SOS Meeting

## Hardware


- ✓ Upgraded Cray XT3 from single to dual-core processors and doubled the memory
- ✓ Installed a 68 cabinet XT4 system
- ✓ Combined the XT3 and XT4 systems into a single 119TF system
- Upgrading HPSS with new Sun/STK 8500 tape libraries

## Software

- ✓ Worked with Cray on Compute Node Linux (CNL) to scale to 23,000+ cores on Jaguar
  - ✓ **Running several applications with performance comparable to Catamount**
- ✓ Working with SNL to convert Catamount to run on multi-core processors
- ✓ Working on external Lustre file system
- ✓ Upgraded HPSS to version 1.6

# Jaguar – 119 TF

- Upgraded to Cray XT4
  - Installed 68 XT4 cabinets in November
  - Moved users to XT4 system in February
  - Moved XT3 to 2<sup>nd</sup> floor and combined systems in March
- Systems have been combined and are in acceptance testing
- System has 11,508 dual-core processors and 46 TB of memory



# Review of 2006 – Phoenix Operations

- Phoenix – Cray X1E
  - **Largest Cray vector system in the world**
  - **Usage of 6.3 million hours**
  - **System availability: 96.5%**
  - **Scheduled availability: 99%**



# NCCS Roadmap for Leadership Computing

Mission: Deploy and operate the computational resources needed to tackle global challenges

- Future Energy
- Understanding the universe
- Nanoscale materials
- Climate Change
- Computational Biology

Vision: Maximize scientific productivity and progress on the largest scale computational problems

- Providing world class computational resources and specialized services
- Providing a stable hardware/software path of increasing scale to maximize productive applications development
- Work with users to scale applications to take advantage of systems







# LCF System Specifications

	119 TF XT4	250+ TF	1000 TF
<b>Compute Processors</b>	11,508 Dual-core 2.6 GHz Optrons	Replaces dual-core with quad-core Optrons	~23,000 quad-core Optrons
<b>SIO Processors</b>	198 I/O nodes	Total 198 Optron	~ 500 quad-core Optron
<b>Memory per Socket/Total</b>	4 GB / 45 TB total sys	8 GB / 69 TB total sys	8 GB / ~175 TB
<b>Interconnect Bandwidth per Socket</b>	Mixed Seastar 1 and Seastar 2 1.8 or 4.0 GB/s	Seastar 2 4.0 GB/s	Gemini <i>Faster</i>
<b>Disk Space</b>	900 TB	900 TB	5 - 15 PB
<b>Disk Bandwidth</b>	55 GB/s	55 GB/s	240 GB/s
<b>Date</b>	1/2007	4Q 2007	4Q 2008

# LCF Users and Usage

- Users come from universities, laboratories, and industry
- Usage across virtually all science domains



# INCITE: 2006 and 2007



## 2006

- Expanded to include SC high end computing resources at PNNL, ORNL and ANL in addition to LBNL and multiple year requests.
- Received 43 proposals requesting over 95 million processor hours.
  - **60% from Universities**
  - **40% had funding from other federal research agencies**
- 15 awards for over 18.2 million processor hours



## 2007

- Expanded in 2007 to include 80% of resources at ORNL Leadership Computing Facilities in addition to 10% of NERSC and 5% of PNNL
- Call issued July 27, 2006
- 88 new proposals received requesting over 180 Million processor hours
- 20 renewal proposals received requesting over 75 Million processor hours
- The proposals represented the following scientific disciplines: accelerator physics, astrophysics, chemical sciences, climate research, computer science, engineering physics, environmental science, fusion energy, life sciences, materials science, nuclear physics and nuclear engineering.
- Eleven of the new proposals were from industry
- 45 INCITE awardees receive a total of over 95 Million processor hours for 2007


2007 INCITE Allocations: 45 projects, 95 million hrs

For NCCS: 28 projects, 75 million hrs




# 250 TF Upgrade

- Jaguar will be upgraded to 250 TF in Fall 2007
  - **Replace dual-core Opteron with quad-core processors**
  - **Maintain 2 GB of memory per core**
  - **Each core support 4 FLOP/s per clock**
- Operating system changes to “Compute Node Linux”
  - **4-core, 8GB shared memory nodes**
  - **Adds OpenMP programming model option**




# 1000 TF System

- Cray “Baker” system is the first in Cray’s “Cascade” line designed for DARPA HPCS program
- FY 2009 install
- ~100,000 cores
- 2 GB per core
- Gemini interconnect
- Compute Node Linux OS



# Infrastructure – Data Analysis & Visualization



“Fat nodes” support legacy visualization and analysis (IDL, AVS, etc.)



Distributed clusters support the largest analysis requirements (VisIt, EnSight, Paraview, etc.)

A large rectangular area composed of a grid of small squares. The left portion is orange and the right portion is red.

Tera/Petascale  
compute system



Moving toward a model of using a portion of the large computational resource for analysis

# Questions?

