ORNL Cray XT3/4 Overview

Jeffrey S. Vetter

And a cast of dozens …

Cray, NCCS, Pat Worley, Sadaf Alam, Weikuan Yu

http://ft.ornl.gov, http://nccs.gov

Software Development Tools for Petascale Computing Workshop on Aug 1-2

Current NCCS Resources

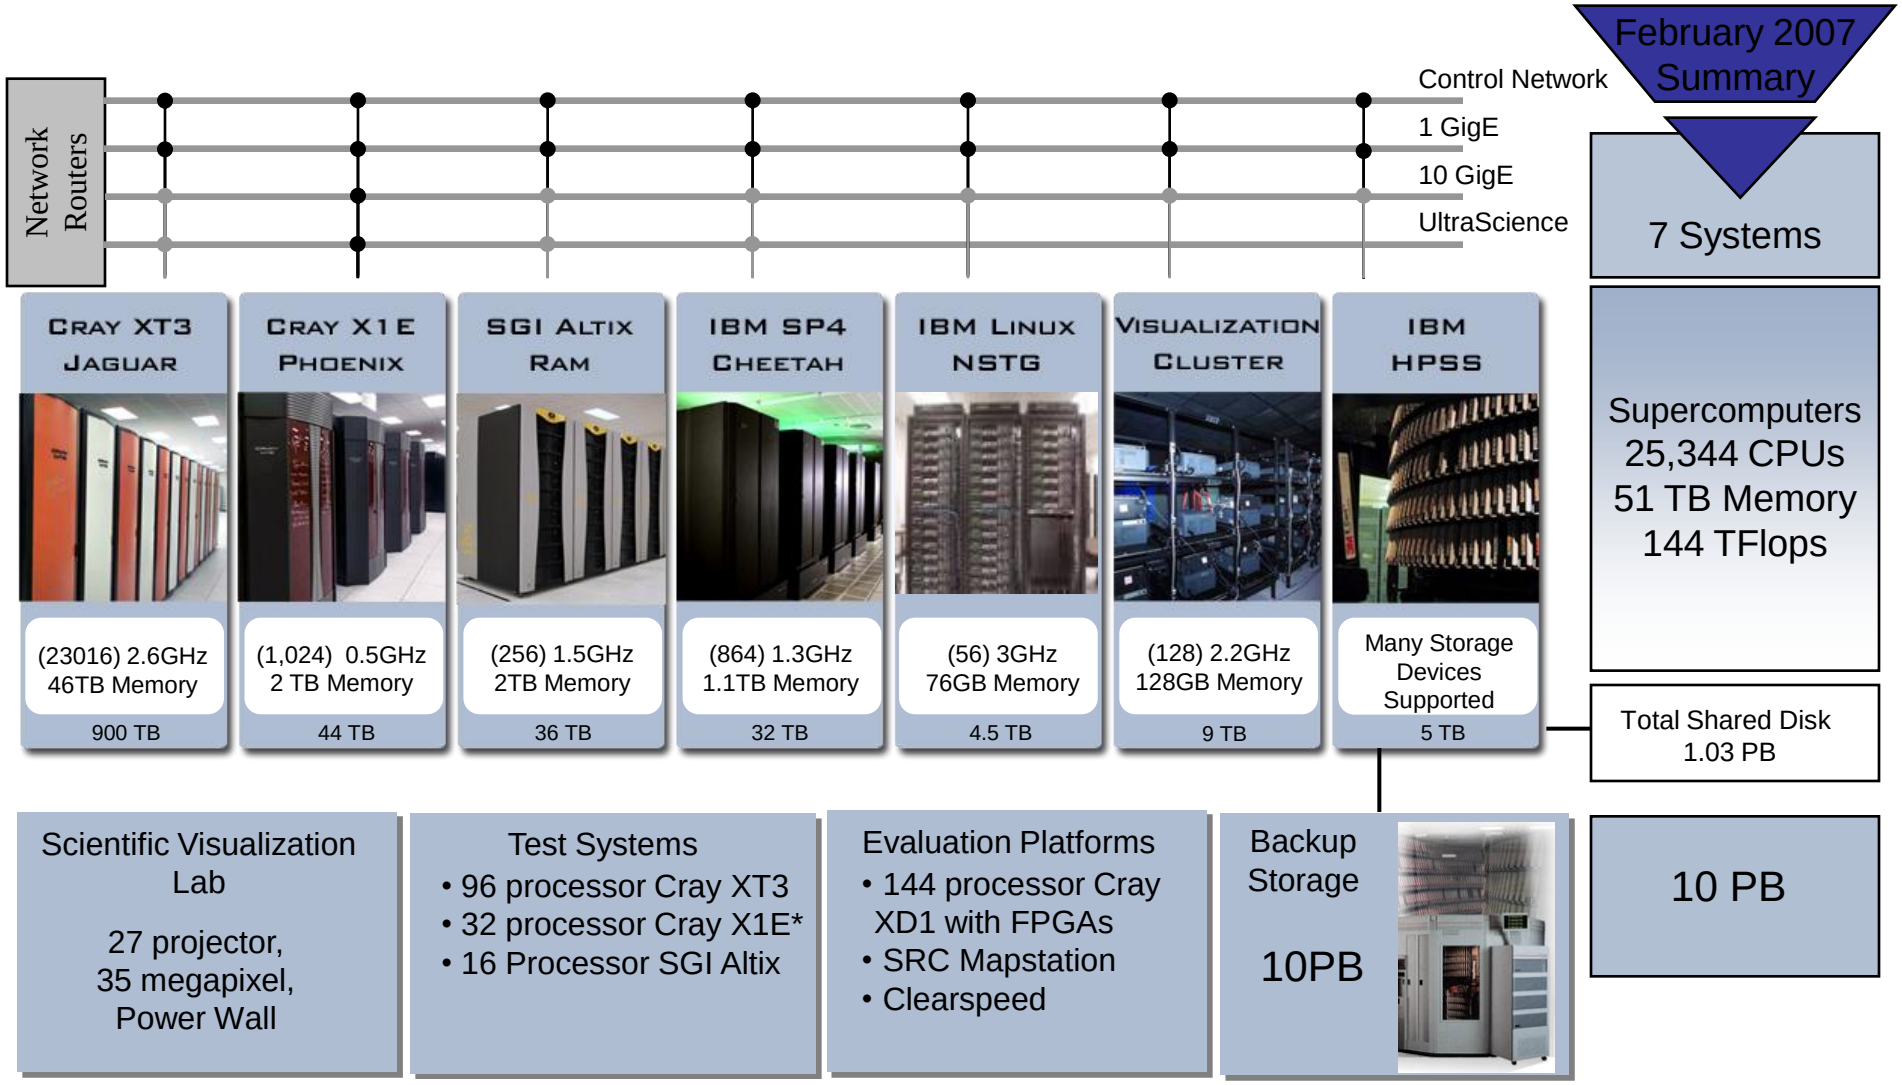

 CSM Experimental Computing Lab houses IBM CELL, GPUs, Clearspeed, FPGAs, etc

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

Cray XT4: 119 TF
Cray X1E: 18 TF

FY2007

Cray XT4: 250 TF
Cray X1E: 18 TF

FY2008

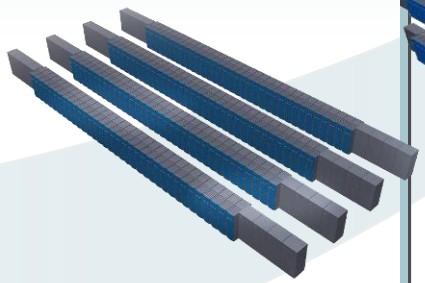

Cray HPCS-0: 1 PF leadership class system for science

FY2009

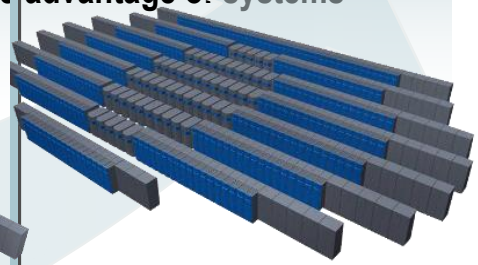

Future leadership class sustained PF system for science

FY2011

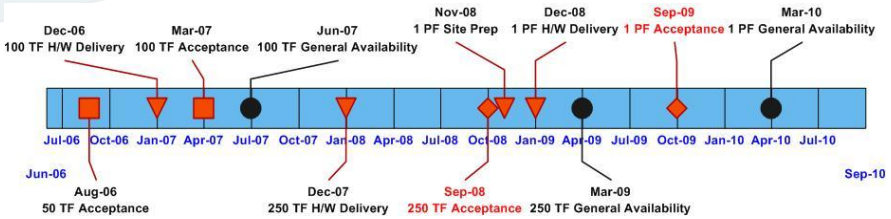

Jaguar – Cray XT4 with 11,706 Dual-Core AMD Opteron Processors

- 119 TF peak performance
- 46 TB main memory
- 2.5 MW of power (20 KW per rack)
- 124 cabinets – 96 Opterons per cabinet
- Air cooled bottom-to-top
- Single 3000 CFM variable speed fan per rack
- 2.3 miles of interconnect cables
- Upgrade to Quad-core processors in Fall, 2007

1000 TF Cray "Baker" system in 2008

System configuration

- 1 PF peak
- ~24,000 quad-core processors
- ~50 KW per cabinet
- ~7 MW power
- Over 3,000 watts/ft^2
- 40+ heat exchange units (chilled water to R-134a)

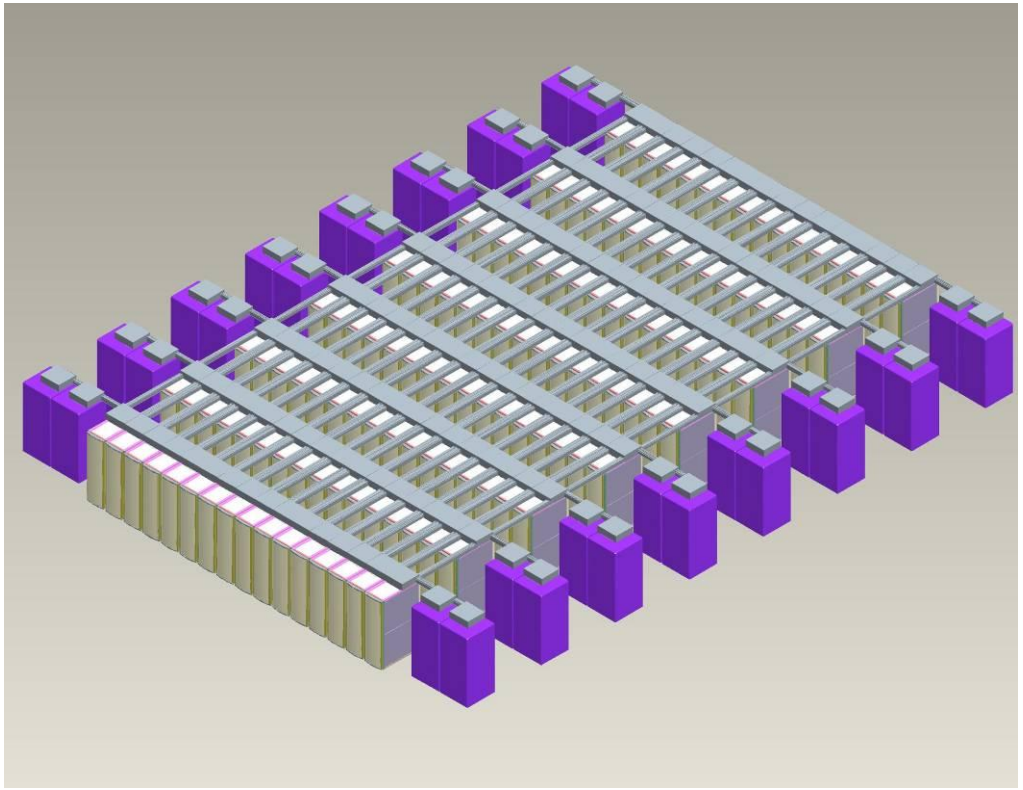

1 PF Cray system in 2008

Used by permission: Cray, Inc.

Cray XT3 Processing Element: Measured Performance

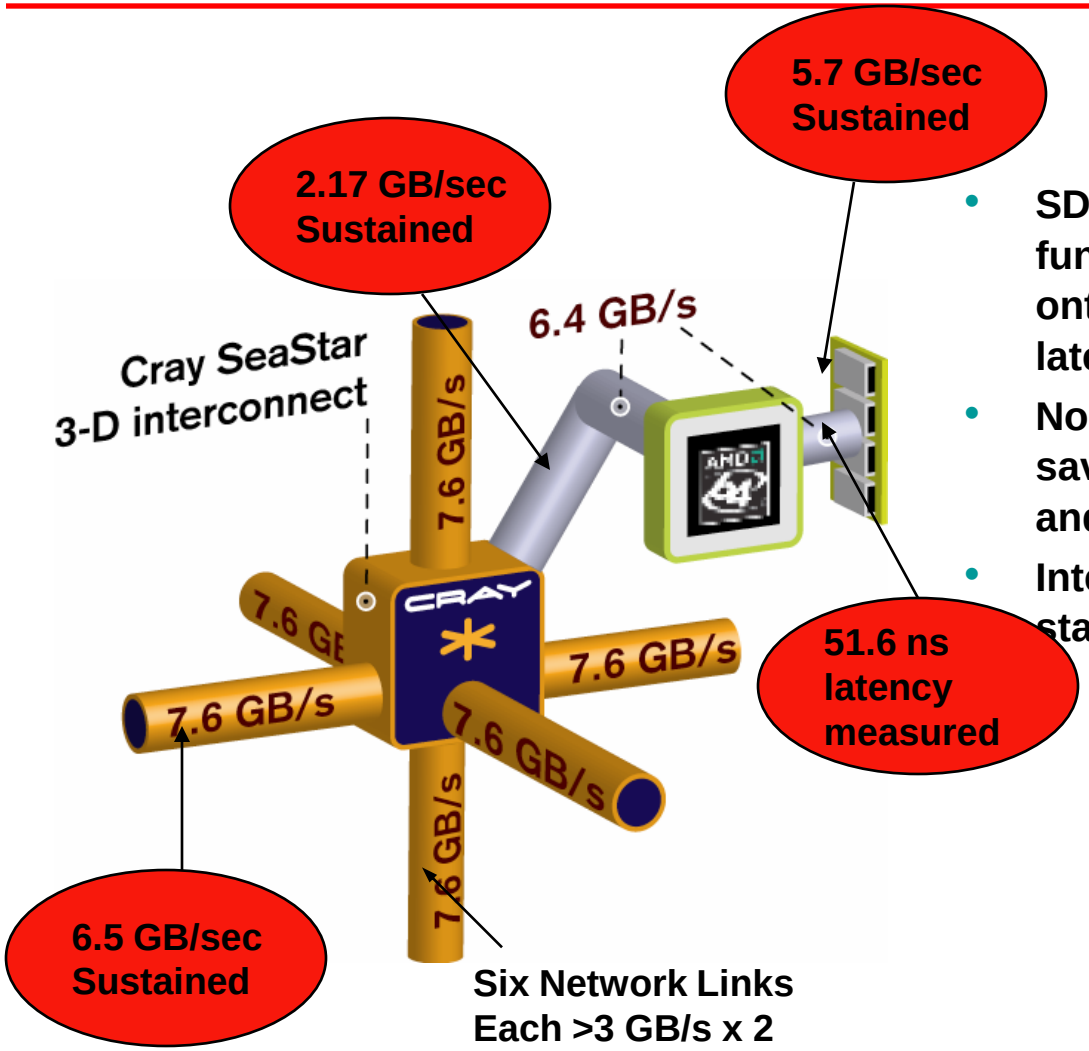

- SDRAM memory controller and function of Northbridge is pulled onto the Opteron die. Memory latency reduced to <60 ns
- No Northbridge chip results in savings in heat, power, complexity and an increase in performance
- Interface off the chip is an open standard (HyperTransport)

The Cray XT4 Processing Element: Providing a bandwidth-rich environment

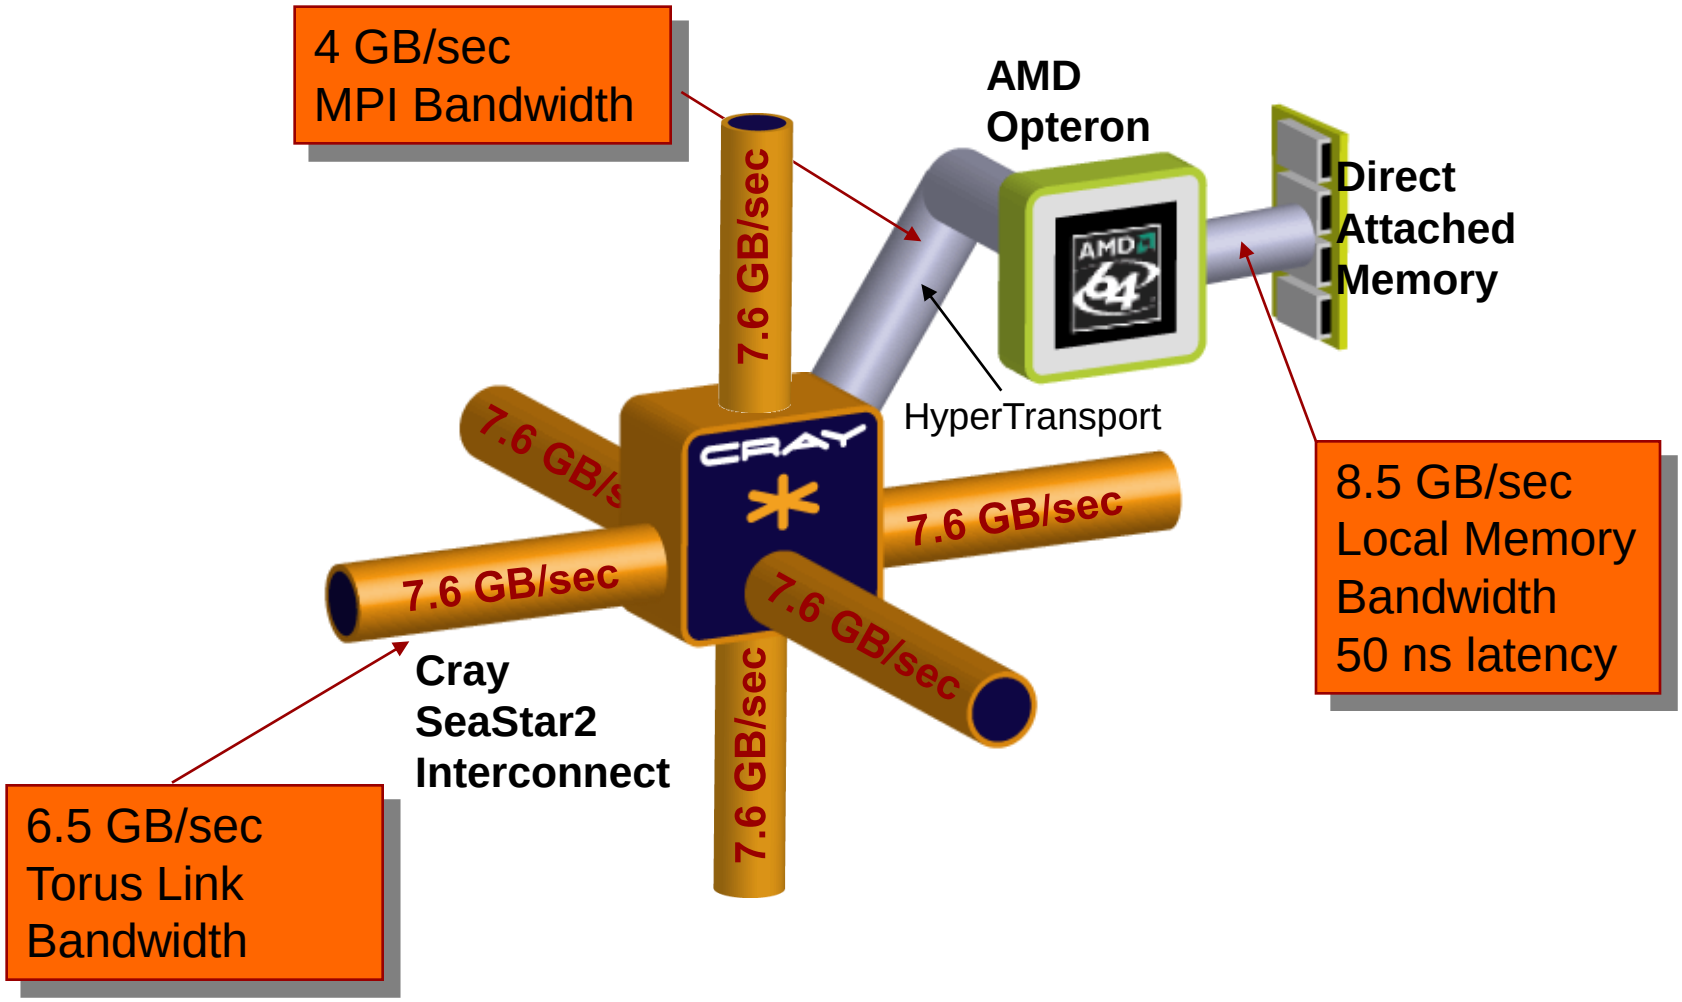

I/O

I/O Configuration

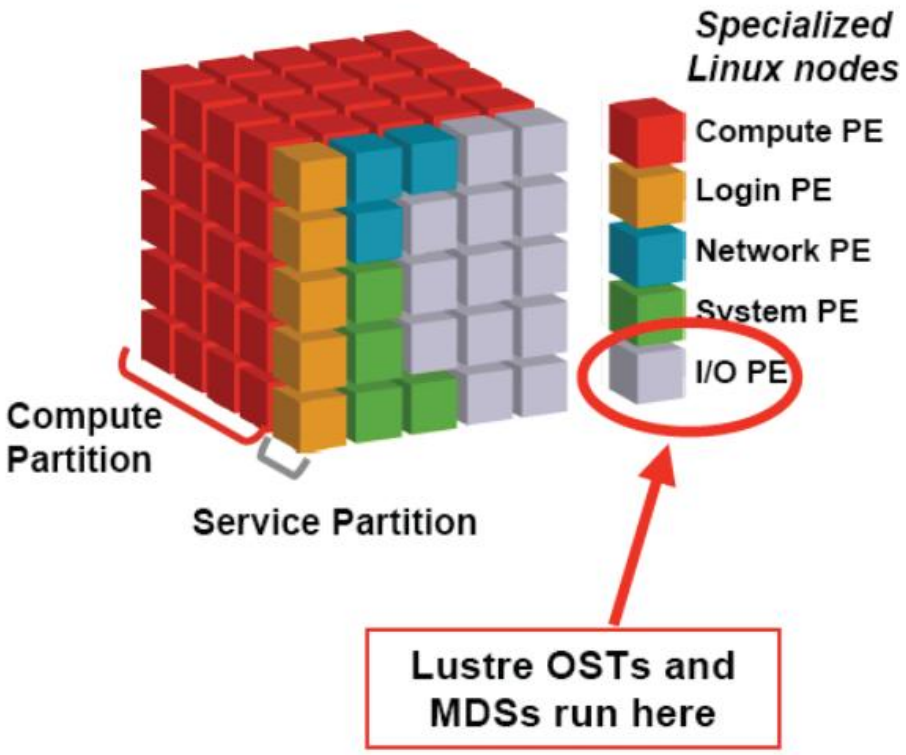

- Lustre filesystems
  - Serviced by 80 I/O nodes
  - /lustre/scr144
    - 144 OSTs
    - Peak 72 GB/s
    - Target ~48 GB/s
    - Early results
      - Read 45 GB/s
      - Write 25 GB/s
  - /lustre/scr72[a,b]
    - 72 OSTs each
    - Default scratch

User view of I/O

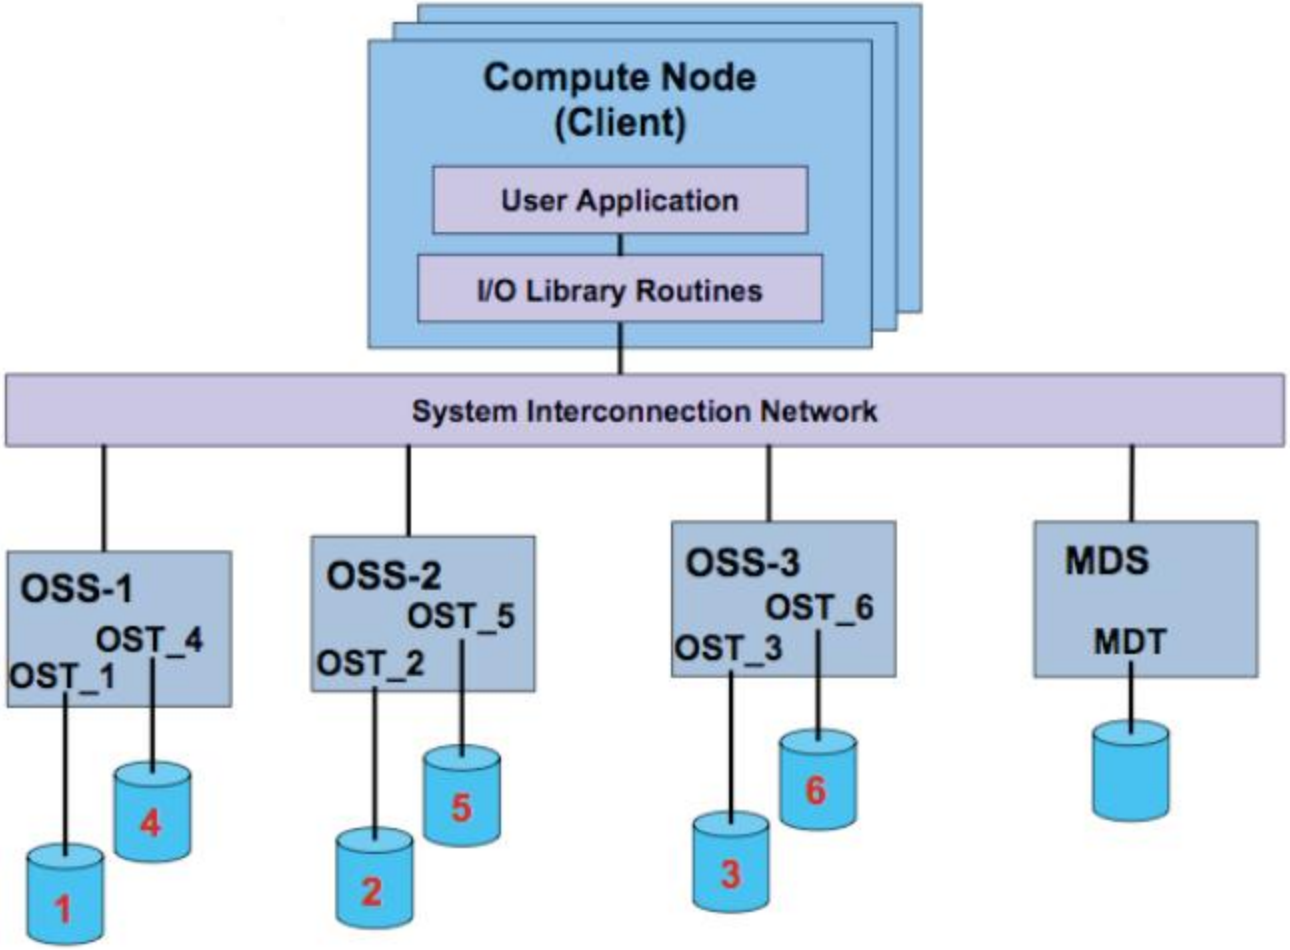

Performance

Performance > Matrix Multiply

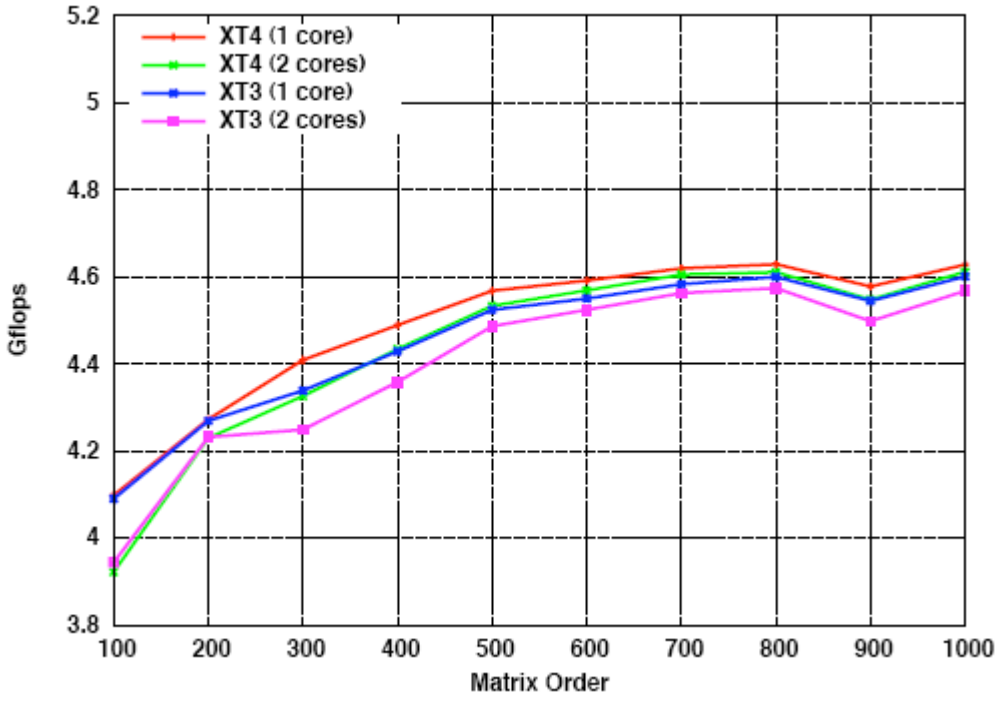

Performance > POP

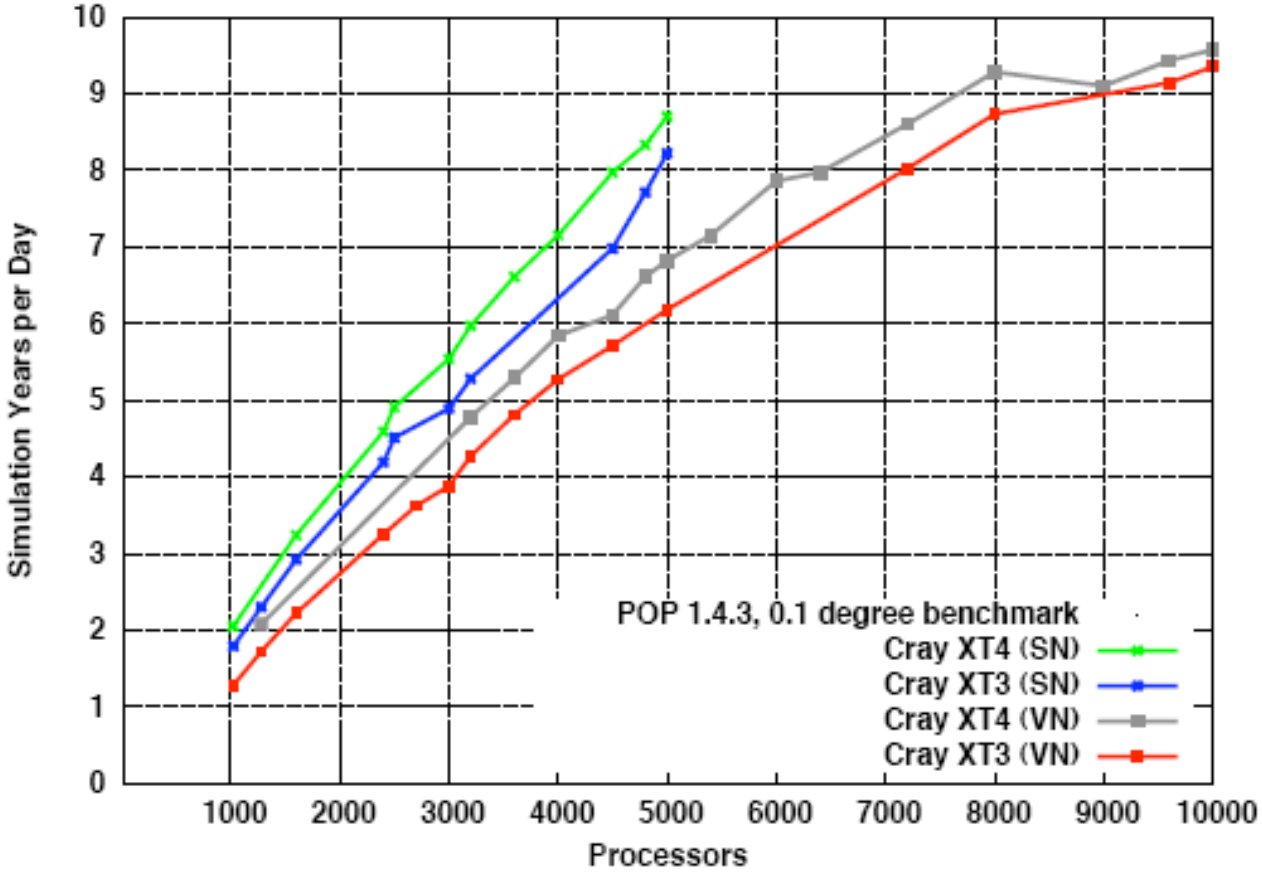

Performance > POP > Phases

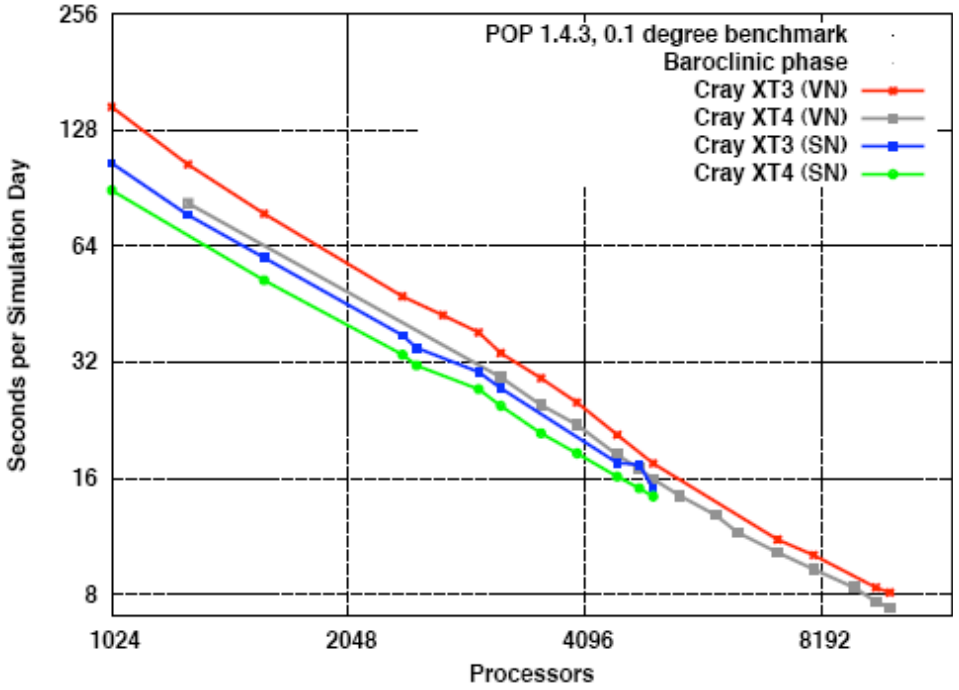

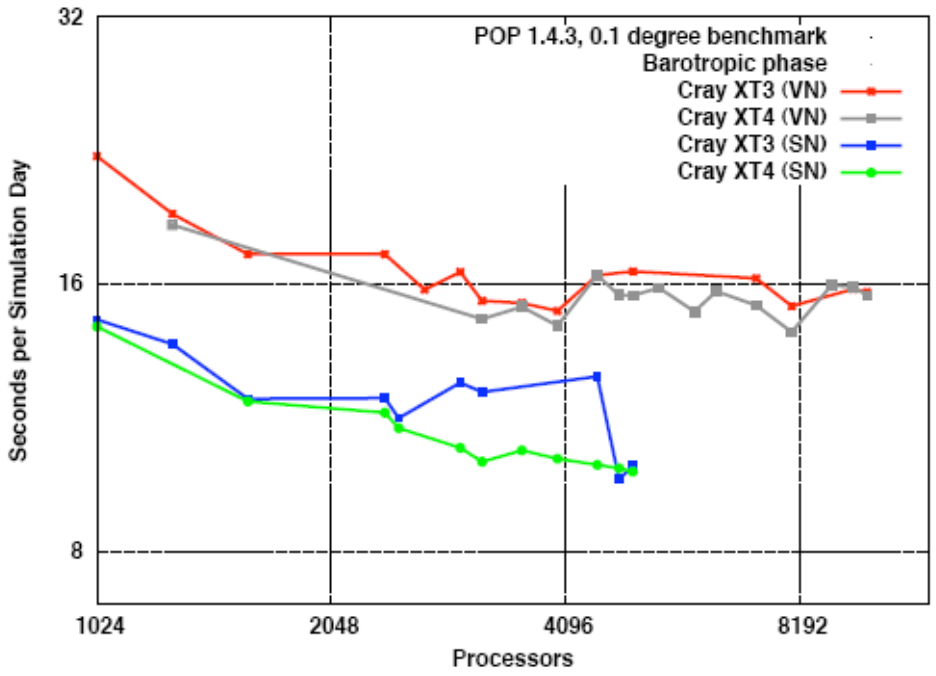

Performance Considerations

- Compiler options
- Page Size
  - 4KB v. 2MB (def)
- SN v. VN
- Process Mapping
  - Manage logical to physical placement of tasks
  - Cray
    - MPICH_RANK_REORDER_METHOD
    - Wrap, smp-style
- Collectives
  - MPI_COLL_OPT_ON

Software

Cray XT3/4 Software

- Operating system
  - Catamount
    - Lightweight OS
  - Compute Node Linux (in testing)
    - Derived from Linux
    - Targeting quad core release
    - More functionality
      - User threads
- Filesystem
  - Lustre
- Tools
  - Performance, debugging
    - Cray PAT and Apprentice
    - Tau
    - PAPI
    - MPIP
    - Totalview
- Compilers
  - PGI, Pathscale
- PBS/Moab batch scheduler

Cray Performance Analysis Infrastructure

- **CrayPat**
  - pat_build: Utility for automatic application instrumentation
    - No source code modification required
  - run-time library for measurements
    - transparent to the user
  - pat_report:
    - Performance reports
    - Generation of file for performance visualization
  - pat_help: Runtime help utility

- **Cray Apprentice^2**
  - Graphical performance analysis and visualization tool
    - Can be used off-line on Linux system

Performance Analysis with CrayPat & Cray Apprentice^2

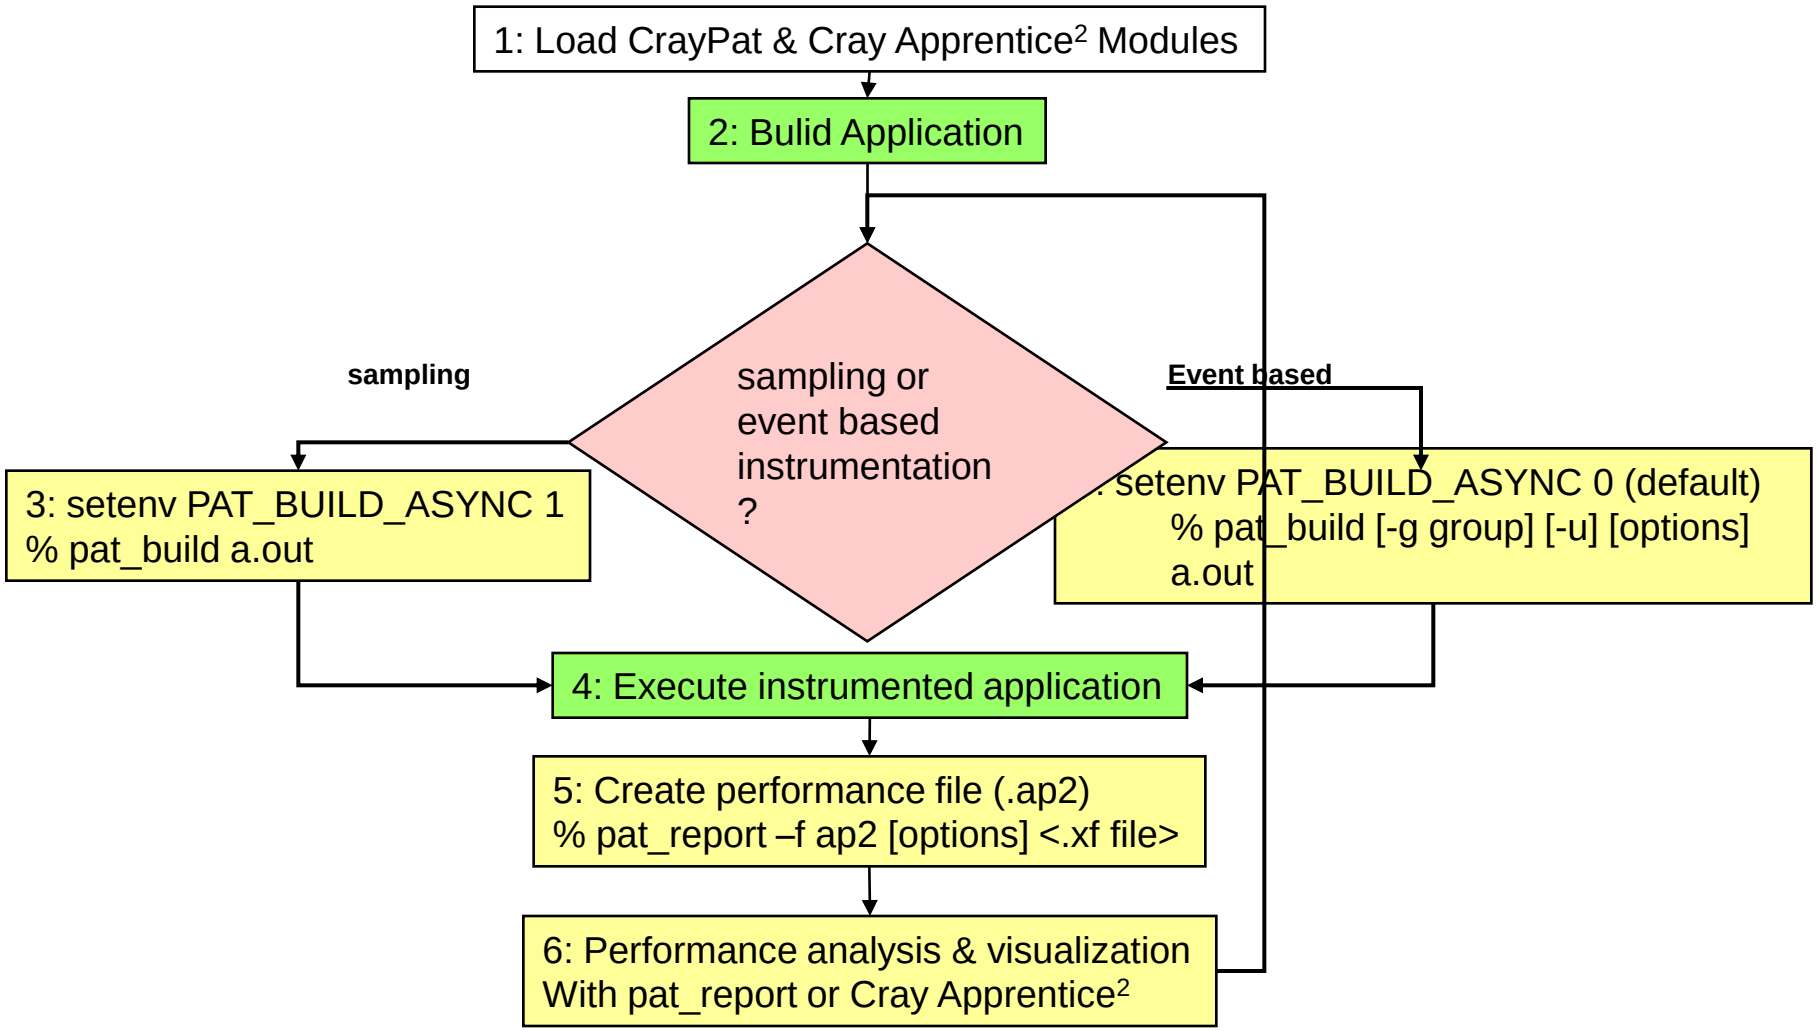

Performance Metrics Available in pat_report

- Profile by groups
  - Threshold
  - Load imbalance information
- Function Profile
  - Flat profile
  - Call Tree view
  - Callers view
  - Hardware counters information
- MPI Profiler
  - MPI Load balance
  - MPI Stats by bin
- I/O Statistics
  - Read and Write Statistics
- Heap Statistics
  - High water mark
  - Memory leaks

CrayPat API

- CrayPat performs automatic instrumentation at function level
- The CrayPat API can be used for fine grain instrumentation
  - Fortran
    - call PAT_region_begin(id, “label”, ierr)
    - DO Work
    - call PAT_region_end(id, ierr)
  - C
    - include <pat_api.h>
    - …
    - ierr = PAT_region_begin(id, “label”);
    - DO_Work();
    - ierr = PAT_region_end(id);

Cray Apprentice²

Call Graph Profile

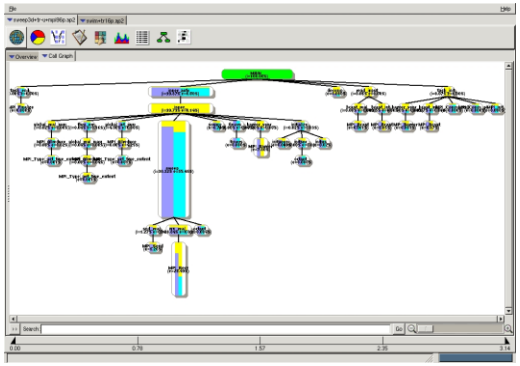

Function Profile

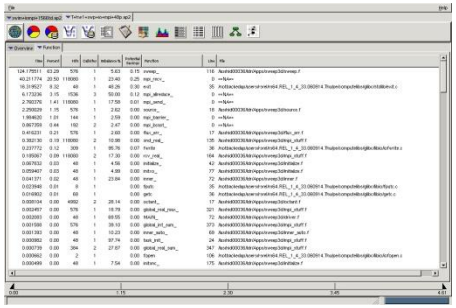

Function Overview

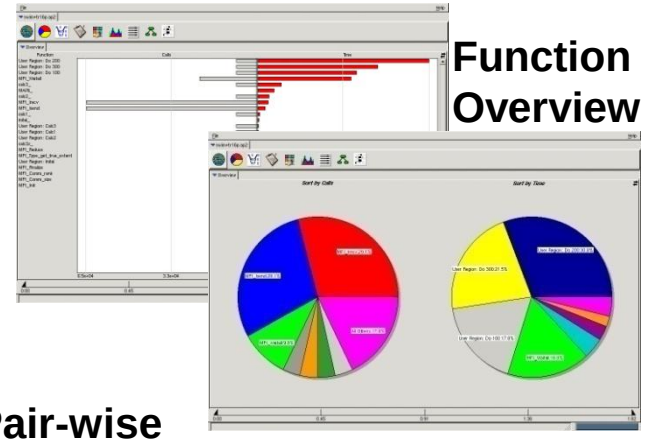

Load balance views

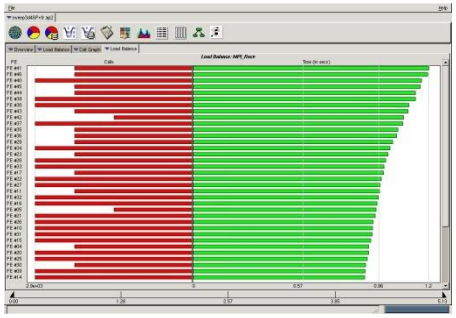

Pair-wise Communication View

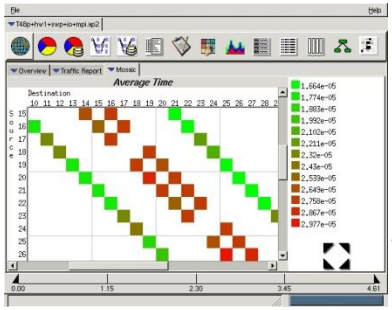

Source code mapping

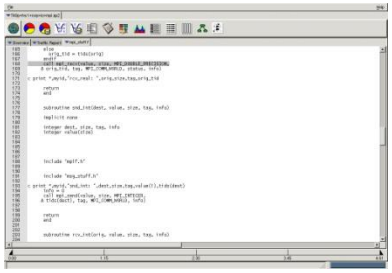

Time Line & I/O Views

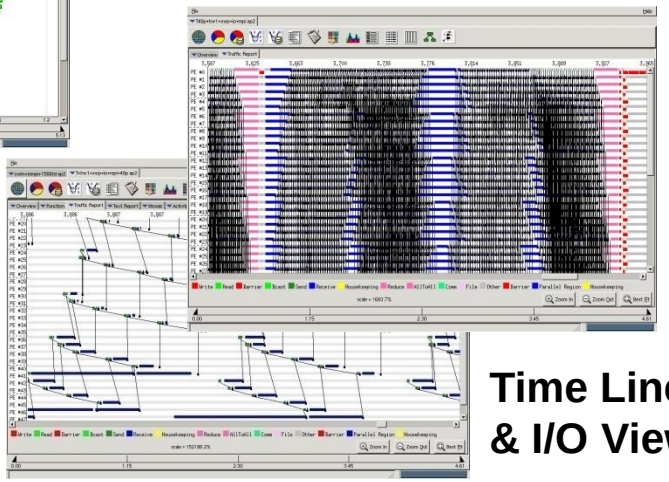

Communication & I/O Activity View

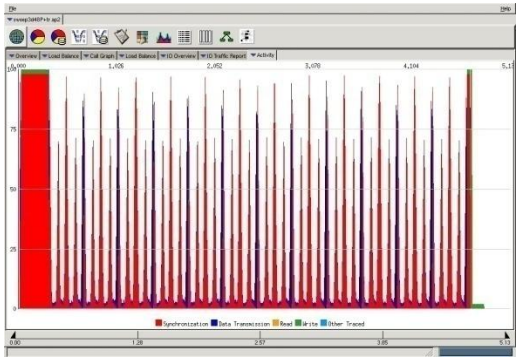

Usage

NCCS Demographics in 2006

2006 Projects			
Accelerator physics	1	Engineering	1
Astrophysics	3	Fusion	4
Chemistry	1	High energy physics	1
Climate change	3	Biology	2
Combustion	1	Materials science	2
Computer science	2	Nuclear physics	1

LCF Users by Affiliation

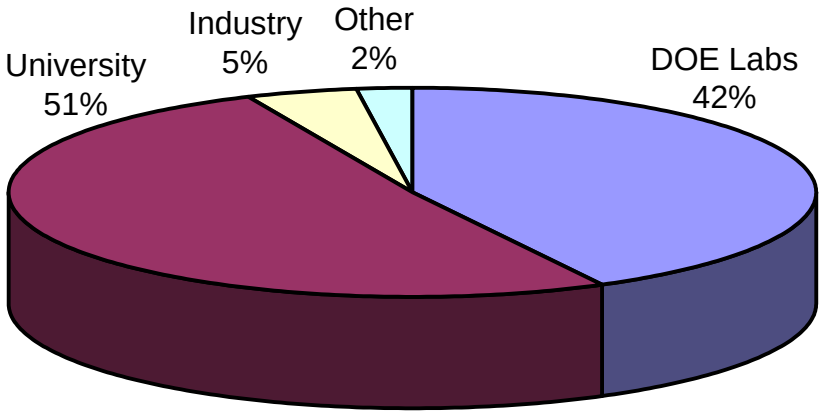

FY 2006 Phoenix Usage by Affiliation

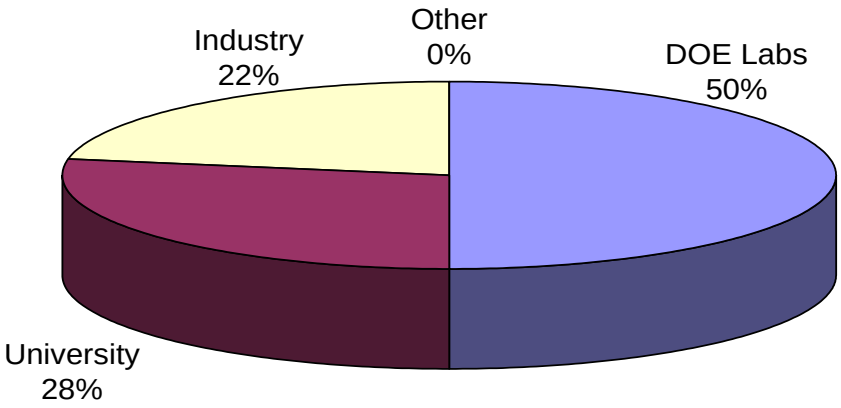

FY 2006 Jaguar Usage by Affiliation

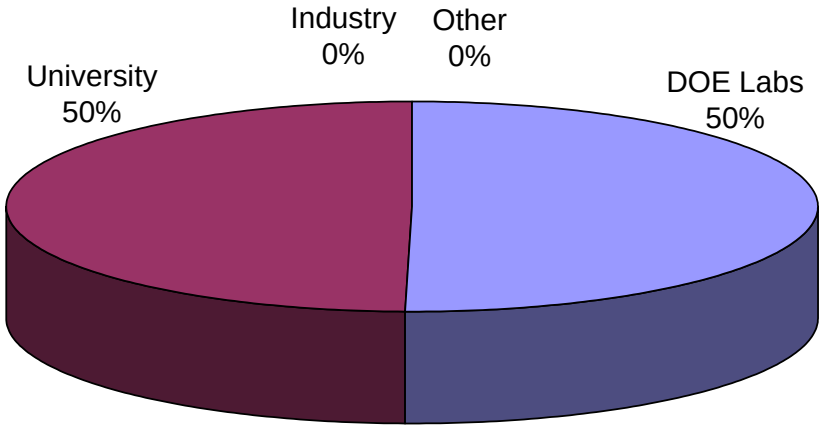

Simulation Job Size Distribution for Science Applications on the ORNL Cray XT3 in 2006

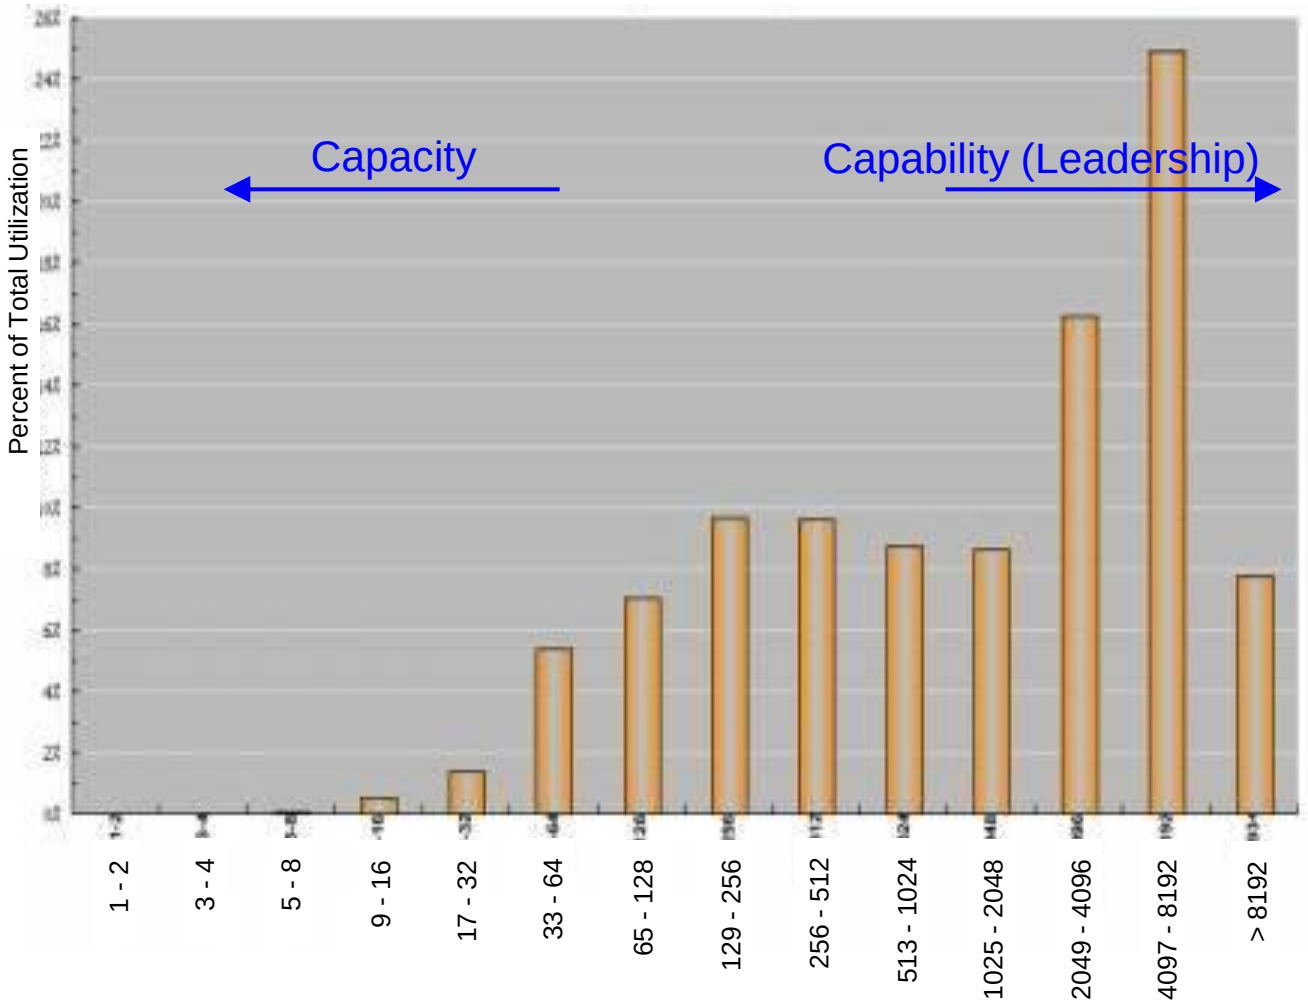

2007 INCITE Allocations at the ORNL LCF

Breakdown by Discipline

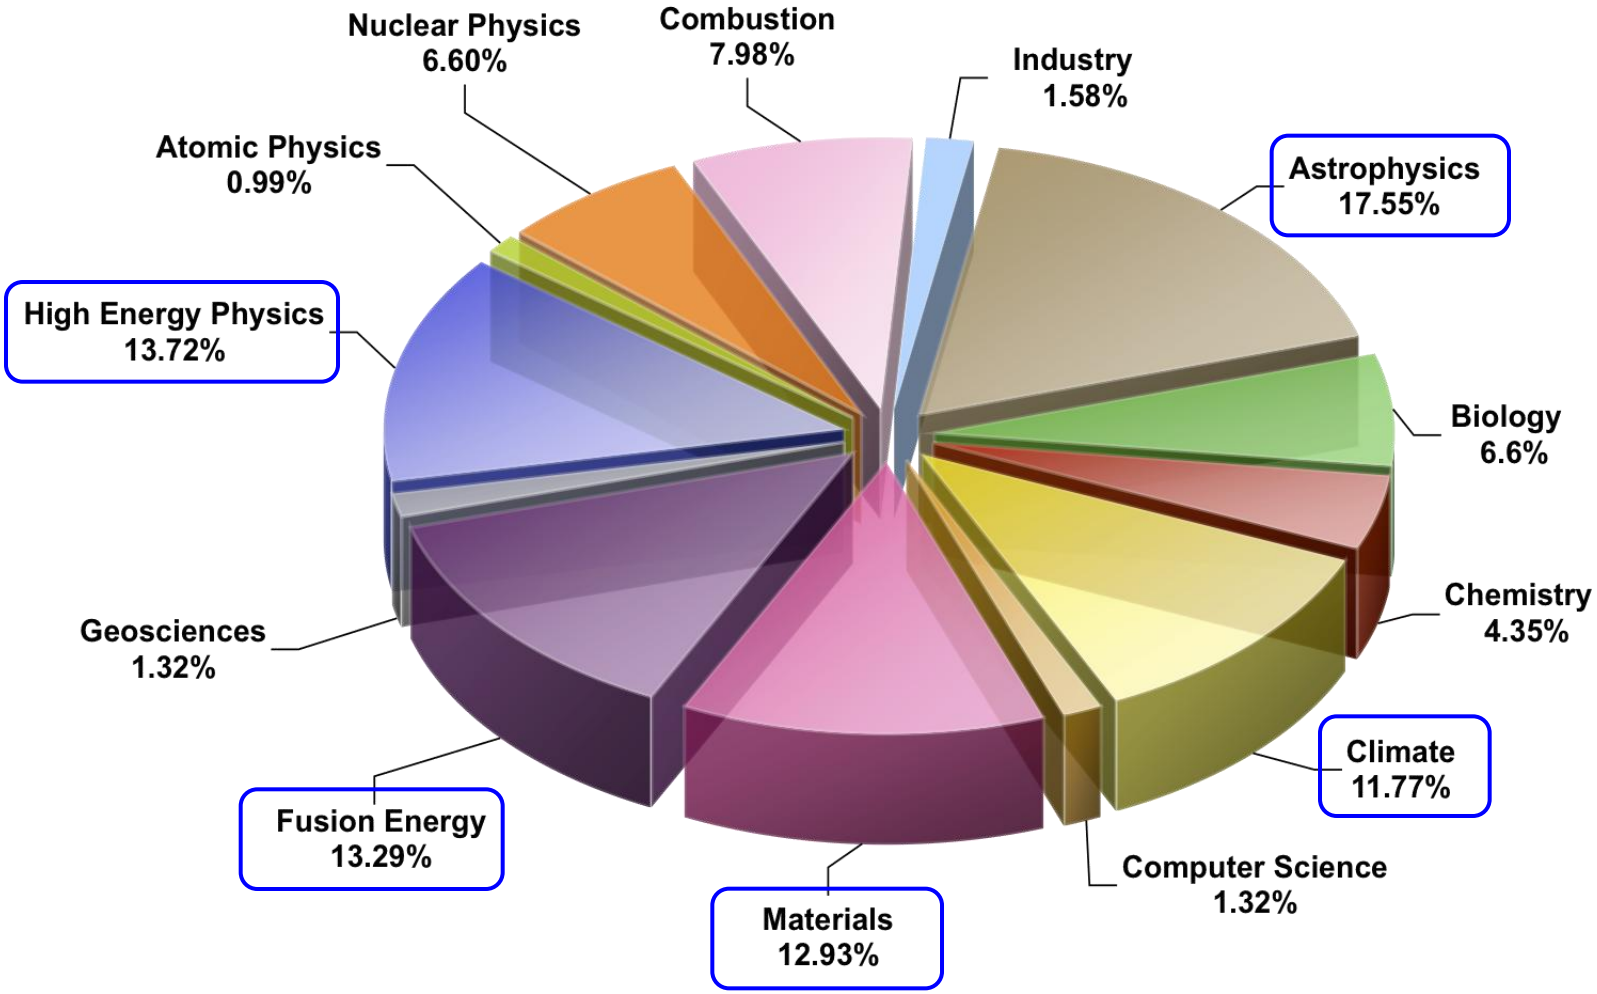

Circled domains make up 70% of total allocation

NERSC Cray XT4

- Very similar to ORNL platform
- Main differences
  - XT4 nodes (homogeneous system)
  - GPFS filesystem