



**Argonne**  
NATIONAL  
LABORATORY

*... for a brighter future*

***ALCF***

*Argonne Leadership  
Computing Facility*



U.S. Department  
of Energy

UChicago ▶  
Argonne<sub>LLC</sub>



A U.S. Department of Energy laboratory  
managed by UChicago Argonne, LLC

# ***ALCF Data Analytics and Visualization Resources***

*William (Bill) Allcock  
Argonne National Laboratory  
Leadership Computing Facility*

# Argonne Leadership Computing Facility

*Established 2006. Dedicated to breakthrough science and engineering.*

## ■ Blue Gene/L Capabilities

- Being decommissioned
- No jobs run after 7/31/08

## ■ Currently in Production

- 111 TF Blue Gene/P system
- Fast 0.8 PB file system
- Initial 16 PB tape archive
- Supports 20 INCITE projects

## ■ For early 2009 production

- 445 TF Blue Gene/P upgrade
- 8 PB next generation file system
- Supports even more challenging INCITE science projects

In 2004 DOE selected the ORNL, ANL and PNNL team based on a competitive peer review

- ORNL to deploy a series of Cray X-series systems
- ANL to deploy a series of IBM Blue Gene systems
- PNNL to contribute software technology



111 TF Endeavour BG/P System

# *DOE INCITE Program*

## *Innovative and Novel Computational Impact on Theory and Experiment*

- **Solicits large computationally intensive research projects**
  - To enable high-impact scientific advances
- **Open to all scientific researchers and organizations**
  - Scientific Discipline Peer Review
  - Computational Readiness Review
- **Provides large computer time & data storage allocations**
  - To a small number of projects for 1-3 years
  - Academic, Federal Lab and Industry, with DOE or other support
- **Primary vehicle for selecting principal science projects for the Leadership Computing Facilities**

# So...

- We have a cool computational resource
  - Equal to a stack of laptops 4000 feet high
- We have a smokin' storage resource
  - Equal to a 1 billion song IPOD, that can write over 1000 DVDs per minute
- How do you build a virtuoso visualization resource to match?

## ***EUREKA!***



# Base Building Block

## ■ Base building block is 3U:

- (2) SuperMicro 1U 6015-RU servers
  - (2) XEON E5405 2.00 GHz quad core processors
  - 32 GB RAM: (8) 4 rank, 4GB DIMMS
  - (1) Myricom 10G CX4 NIC
  - (2) small system disks; No local scratch space
  - 32 GFlops per server
- (1) NVidia 1U S4 External GPU
  - (4) Quadro FX5600
  - 2072 single precision GFlops per S4
    - 128 cores \* 1.35 GHz \* 3 Flops per clock \* 4 cards
- Servers connect to the S4 via a 16x PCIe V2.0 card
- Logically, we have (2) 1U servers, each with (2) Nvidia Quadro FX5600 graphics cards in it.

# System Totals

## ■ 4 racks

- 13 “building blocks” per rack, 52 total
- (104) 1U 2.0 GHz Xeon, 8 core, 32 GB RAM servers
- (52) 1U NVidia S4 external GPUs
  - *Largest S4 installation in the world at this time*
  - *(208) Quadro FX5600 high end graphics cards*
- Over 111 TF peak FLOP rating
  - *Includes the GPUs*
- Over 3.2 TB of RAM (5% of intrepid RAM)
- No local scratch space
  - *Data access is all via the central parallel file system*
  - *Myricom 10G (10 Gbs) NIC*

# *Challenges / Unique Aspects*


## ■ No local scratch space

- If you could have been a fly on the wall during that discussion...
- We put a lot of time, money, and effort into a 70+ GB/s parallel file system
- We believe we can sustain 500 MB/s...
- If we only get half of that, we are still faster than your average internal RAID array.
- 100+ RAID arrays is a huge admin burden.

# Challenges / Unique Aspects

## ■ Cooling

- (39) 1U boxes / 20KW in a rack
- Using passive thermal management racks
  - *Solid back door*
  - *“smoke stack” rejects heat directly to the return plenum*
- Spacing between the racks
  - *The picture on the right is to scale*
  - *20KW needs 2200 – 2500 CFM of air*
  - *Our tiles provide approx. 2000 CFM at our plenum pressure*
  - *Spacing allows the rack to “steal” air from adjacent tiles*
  - *Because of racks, no hot aisle / cold aisle mixing*





# Software Resources

- Software stack will be driven by our users
  - So if we don't have what you need, please speak up
- To date, here are the applications we intend to install and support:
  - VisIT
  - ParaView
  - VMD
  - vI3
  - VisTrails
  - And good old gnuplot 😊
- Note that we don't list any commercial apps
  - No user driven demand to date

# *People Resources*

- Your first line of defense for all help at ALCF
  - [support@alcf.anl.gov](mailto:support@alcf.anl.gov)
- The above will get you access to:
  - Your friendly neighborhood Catalyst
  - Application Performance Engineering and Data Analytics (APEDA) group
  - ALCF operations team
  - Advance Integration Group
- And of course we are partnered with the Math and Computer Science Division at ANL



**Argonne**  
NATIONAL  
LABORATORY

*... for a brighter future*

**ALCF**

*Argonne Leadership  
Computing Facility*



U.S. Department  
of Energy

UChicago ►  
Argonne<sub>LLC</sub>



**Office of  
Science**

U.S. DEPARTMENT OF ENERGY

A U.S. Department of Energy laboratory  
managed by UChicago Argonne, LLC

**Questions?**