# Effectively Targeting the Oak Ridge Leadership Computing Facility

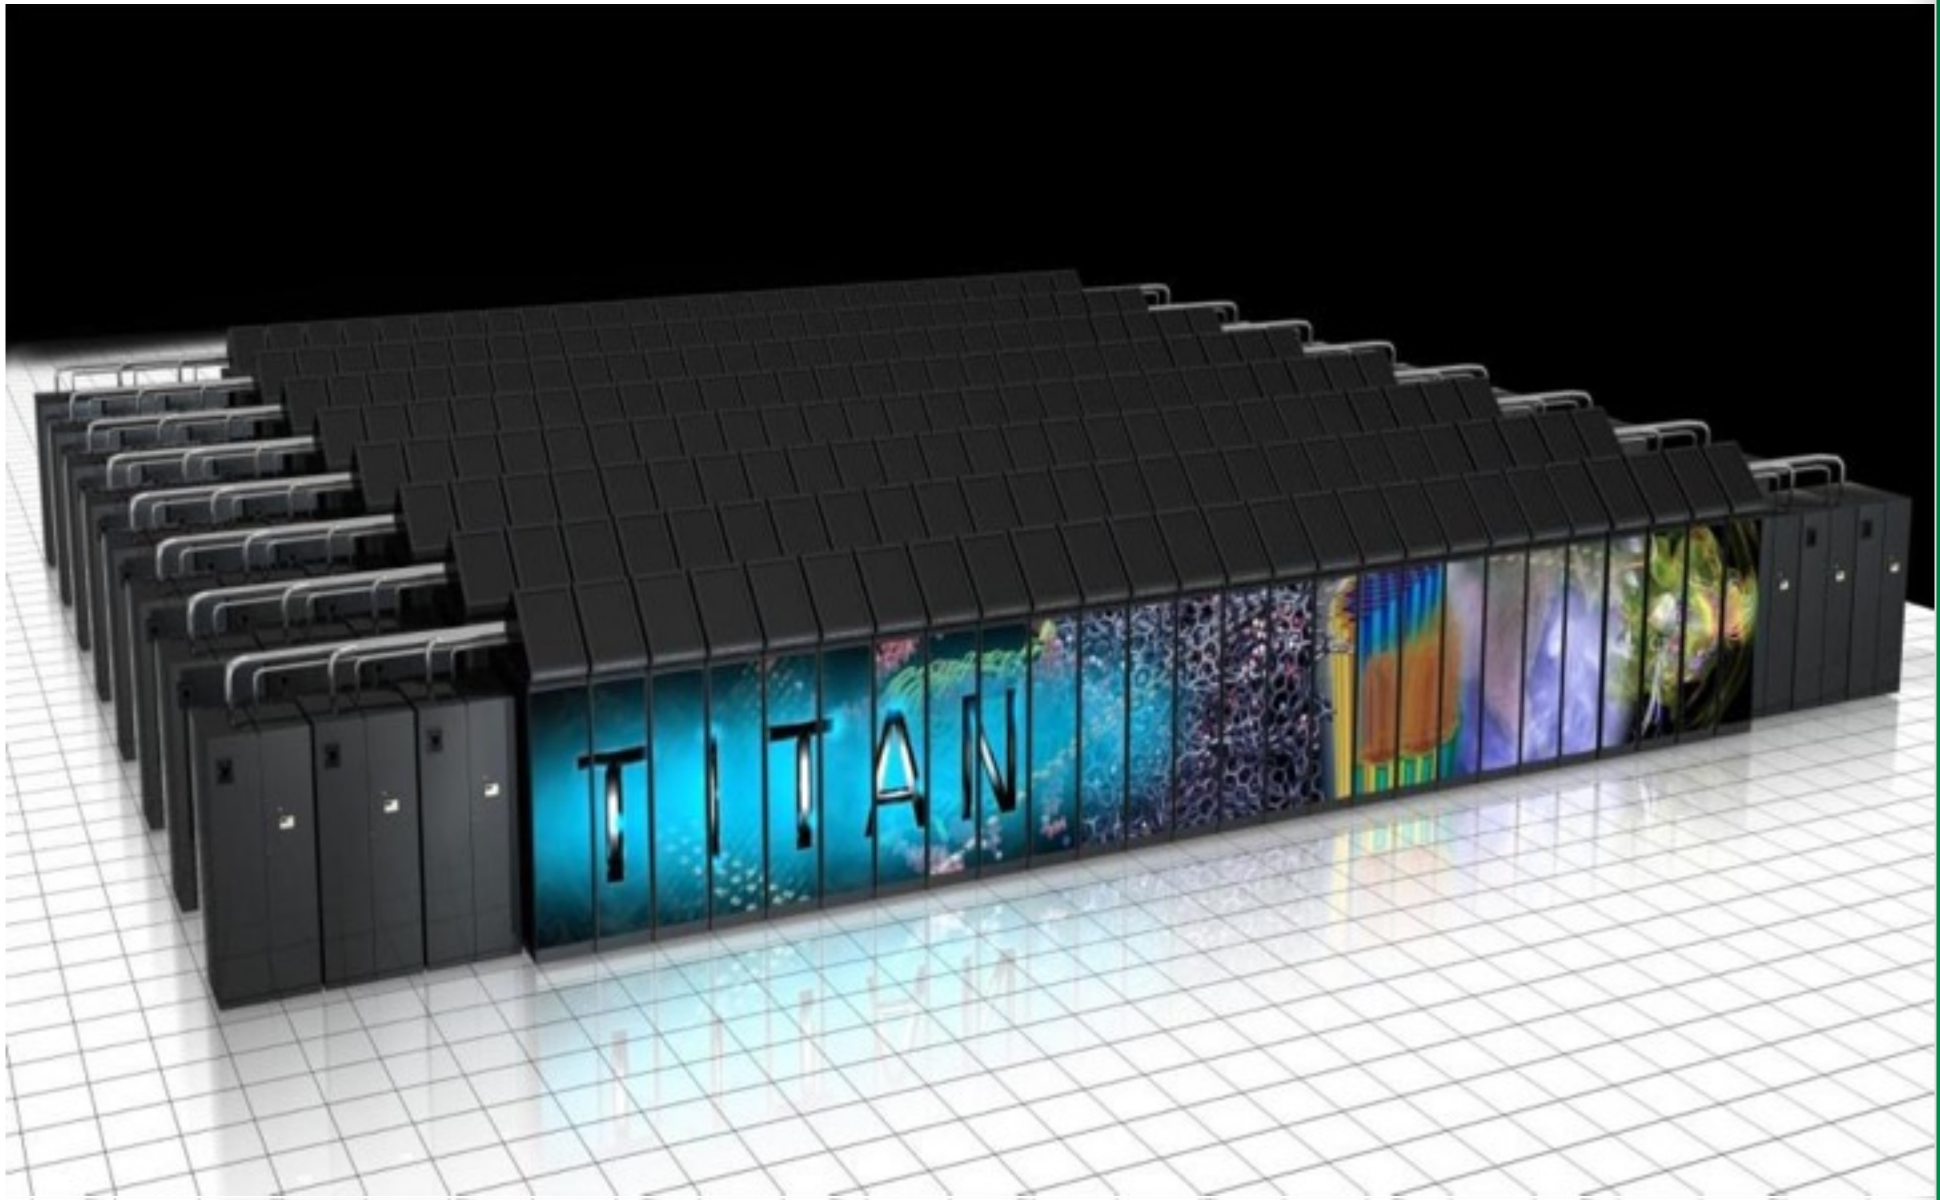



Bronson Messer

Scientific Computing (OLCF) & Theoretical Physics Groups
Oak Ridge National Laboratory



ORNL is managed by UT-Battelle
for the US Department of Energy


OAK RIDGE National Laboratory

# Outline

- Mission of the Oak Ridge Leadership Computing Facility (OLCF)
- Jaguar to Titan – Why GPUs?
- What's Next? - Summit (and Cori & Aurora)
  - A little about portability
- Some nuts & bolts about how to really compute at OLCF

# A Little About ORNL…

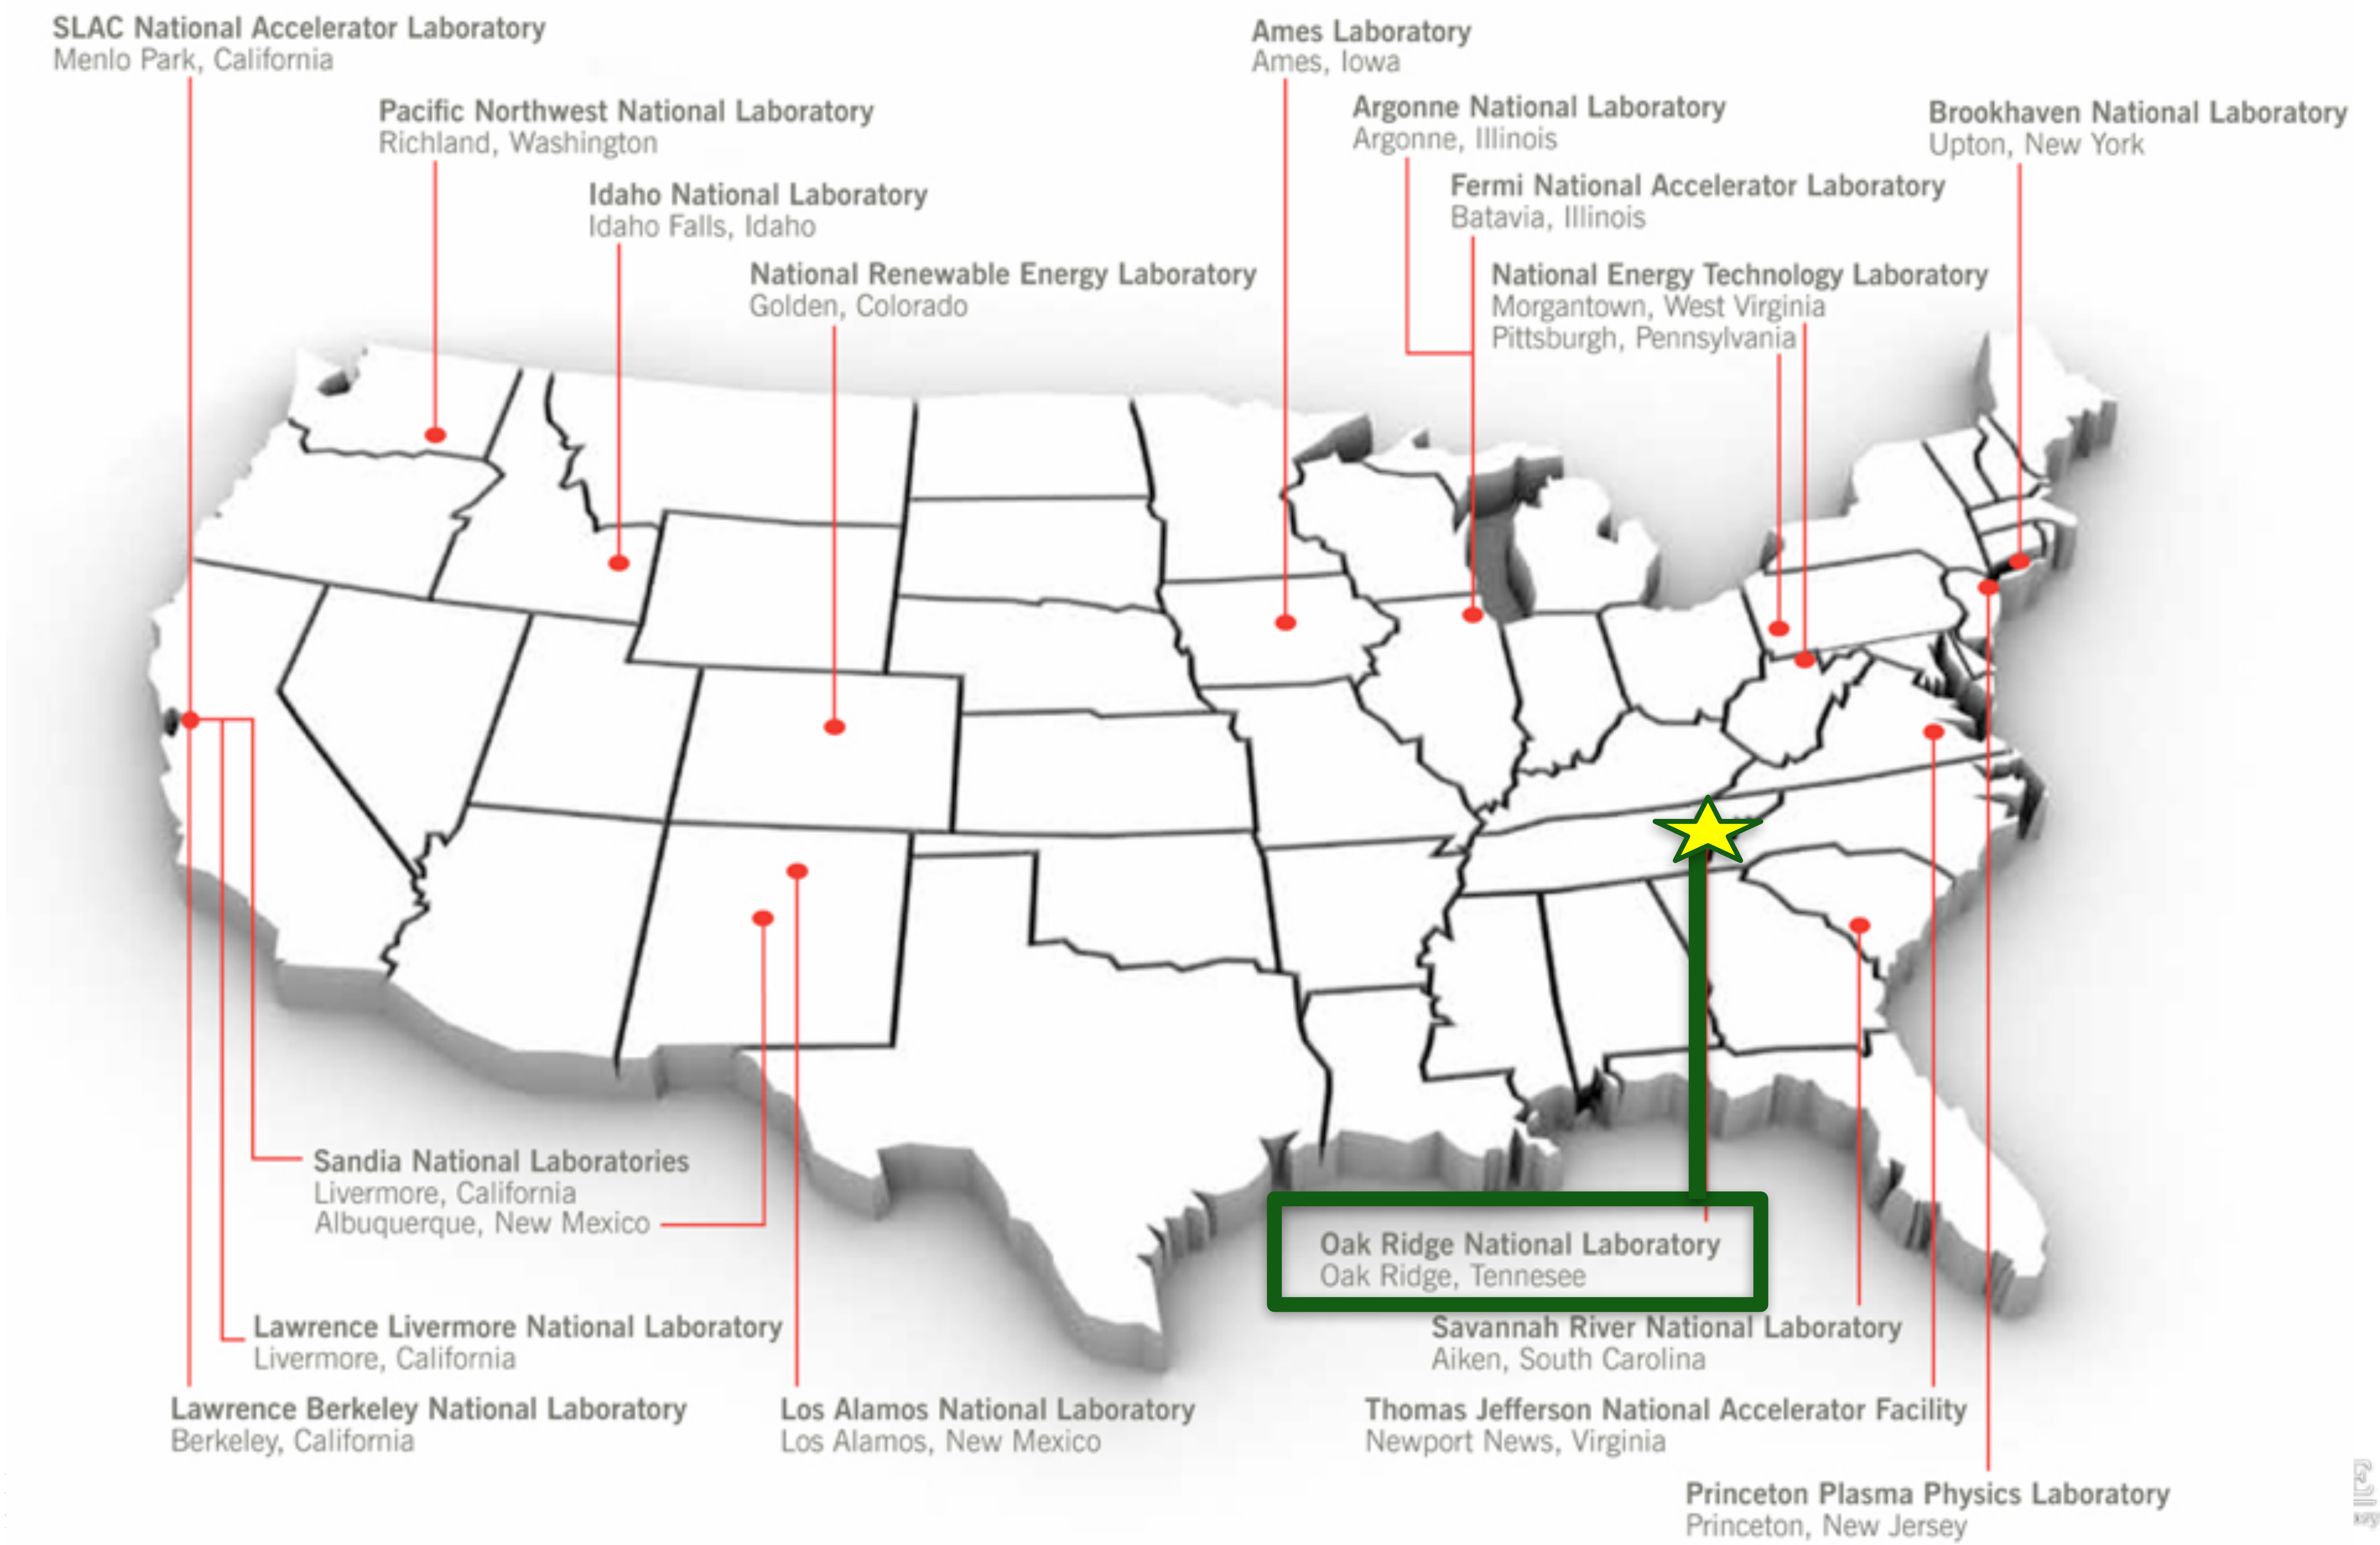

# ORNL is DOE's largest science and energy laboratory

- $1.5B budget
- 4,500 employees
- $750M modernization investment
- Nation's largest materials research portfolio
- 3,200 research guests annually
- Nation's most diverse energy portfolio
- Forefront scientific computing facilities
- World's most intense neutron source
- World-class research reactor
- 1,888 journal articles published in CY14
- Managing billion-dollar US ITER project
- 198 invention disclosures in FY15
- 66 patents issued in FY15

# What is the Leadership Computing Facility (LCF)?

- Collaborative DOE Office of Science user-facility program at ORNL and ANL
- Mission: Provide the computational and data resources required to solve the most challenging problems.
- 2-centers/2-architectures to address diverse and growing computational needs of the scientific community
- Highly competitive user allocation programs (INCITE, ALCC).
- Projects receive 10x to 100x more resource than at other generally available centers.
- LCF centers partner with users to enable science & engineering breakthroughs (Liaisons, Catalysts).



Oak Ridge National Laboratory
Argonne National Laboratory

# ORNL increased system performance by 1,000x (2004-2010)

Hardware scaled from single-core through dual-core to quad-core and dual-socket , 12-core SMP nodes

Scaling applications and system software was the biggest challenge

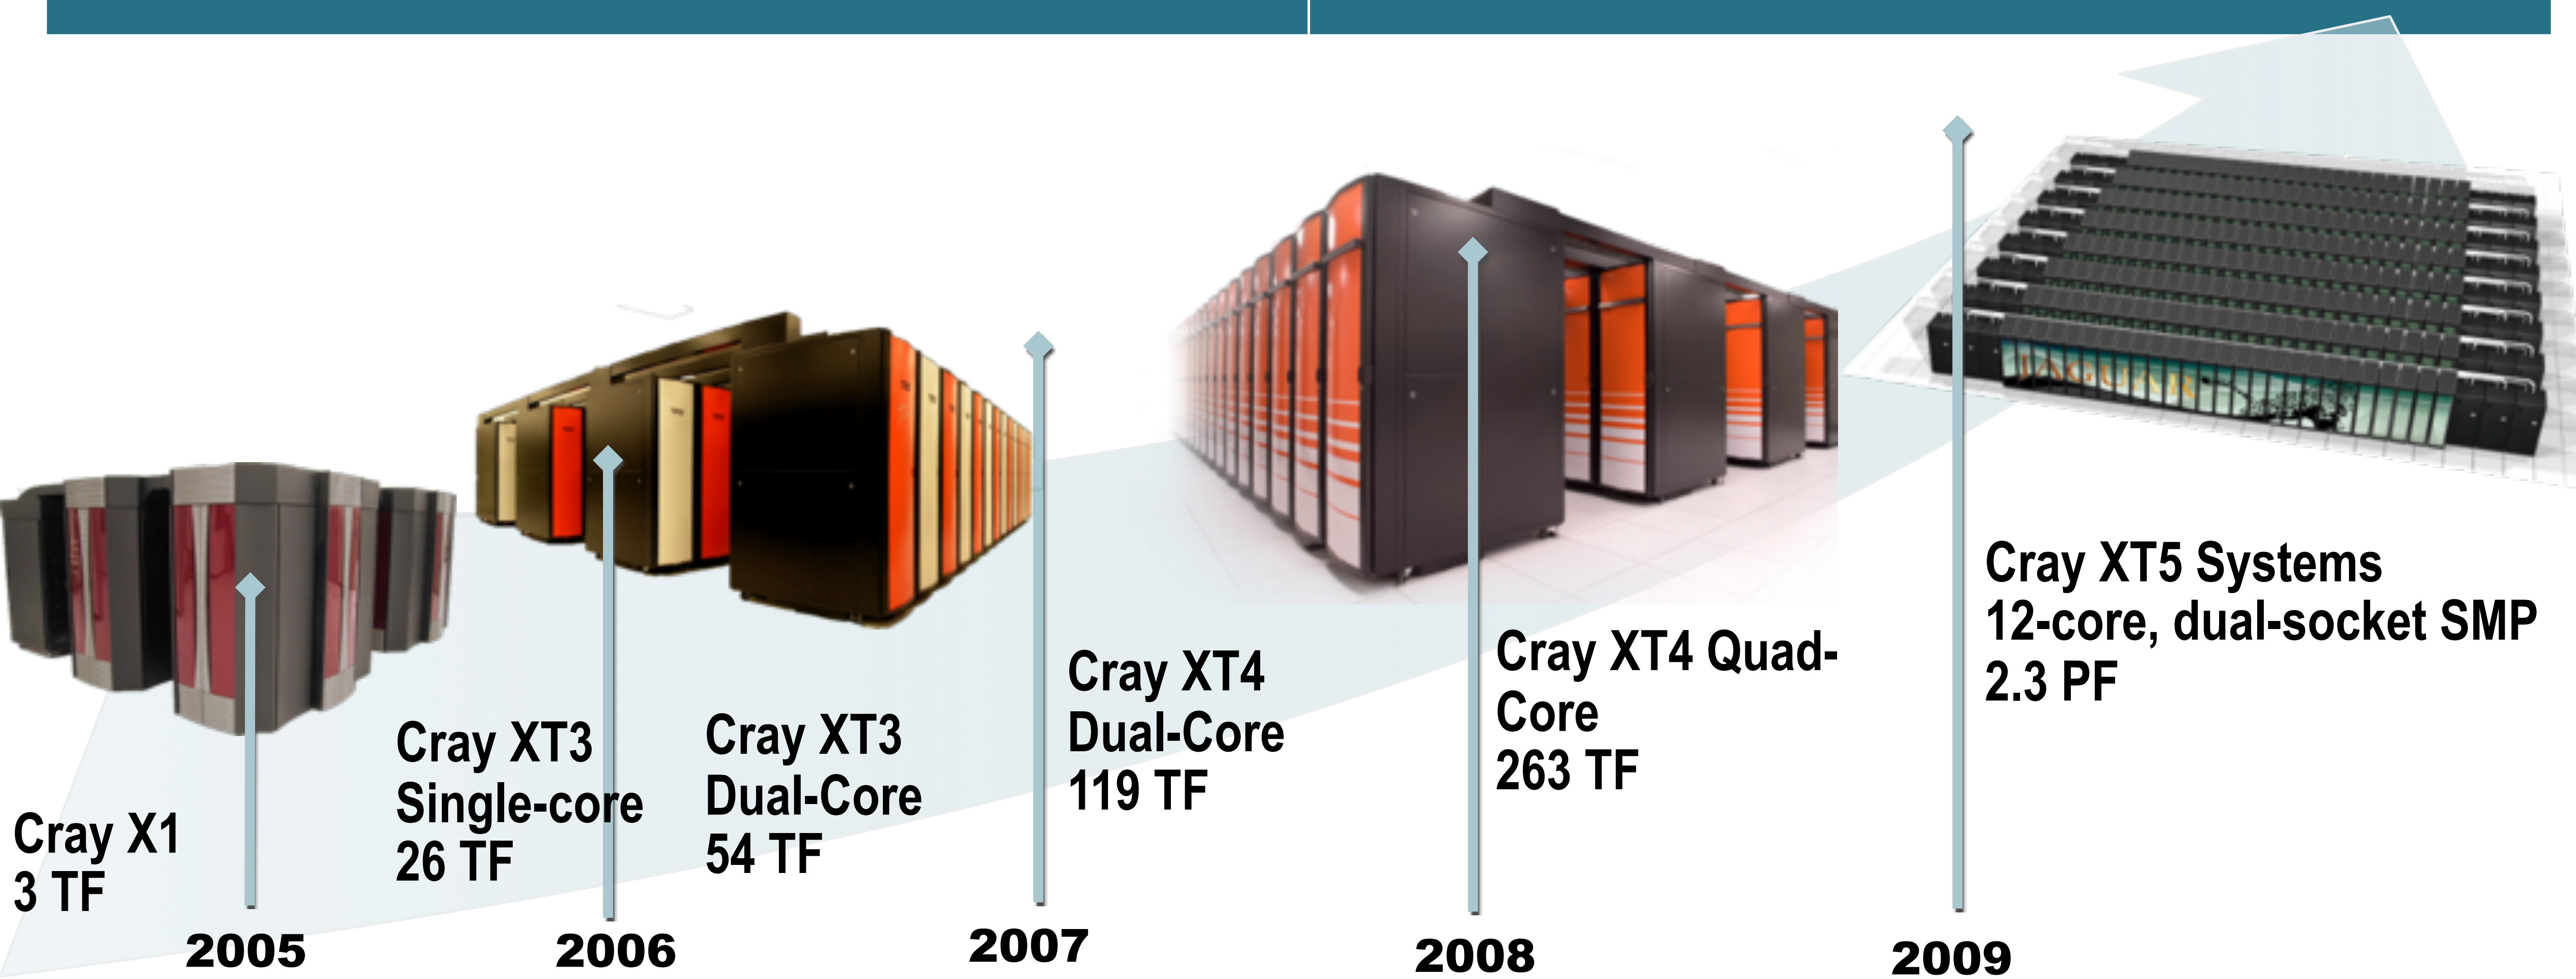

# Architectural Trends – No more free lunch

- CPU clock rates quit increasing in 2003
- $P = CV^2 f$
  Power consumed is proportional to the frequency and to the square of the voltage
- Voltage can't go any lower, so frequency can't go higher without increasing power
- Power is capped by heat dissipation and $$$
- Performance increases have been coming through increased parallelism

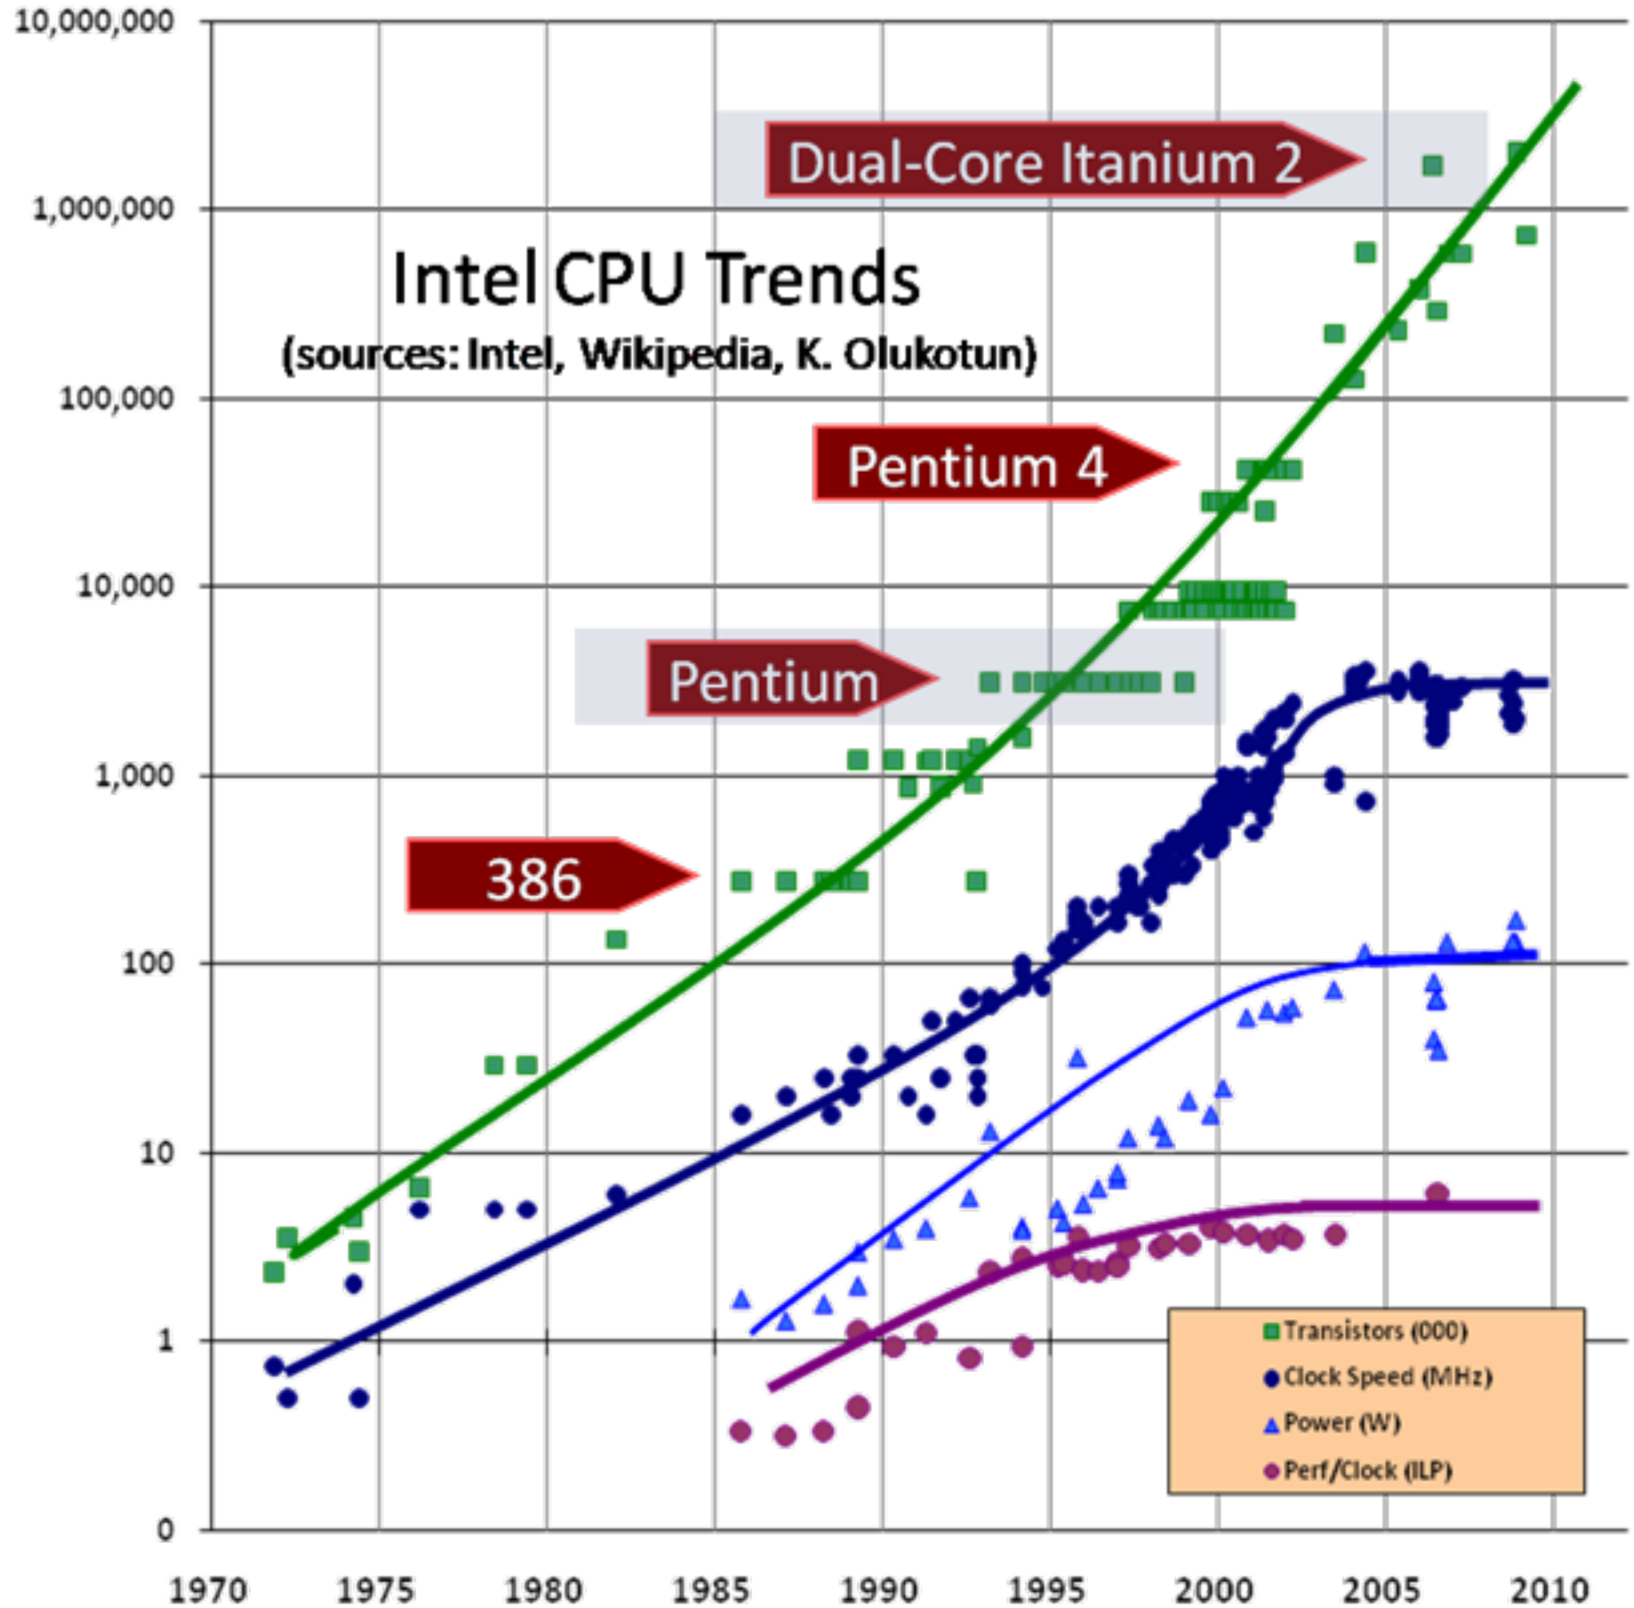


Herb Sutter: Dr. Dobb's Journal:
http://www.gotw.ca/publications/concurrency-ddj.htm

# The Effects of Moore's Law and Slacking [1] on Large Computations

Chris Gottbrath, Jeremy Bailin, Casey Meakin, Todd Thompson,
J.J. Charfman
Steward Observatory, University of Arizona

[1] This paper took 2 days to write

**Abstract**

We show that, in the context of Moore's Law, overall productivity can be increased for large enough computations by 'slacking' or waiting for some period of time before purchasing a computer and beginning the calculation.

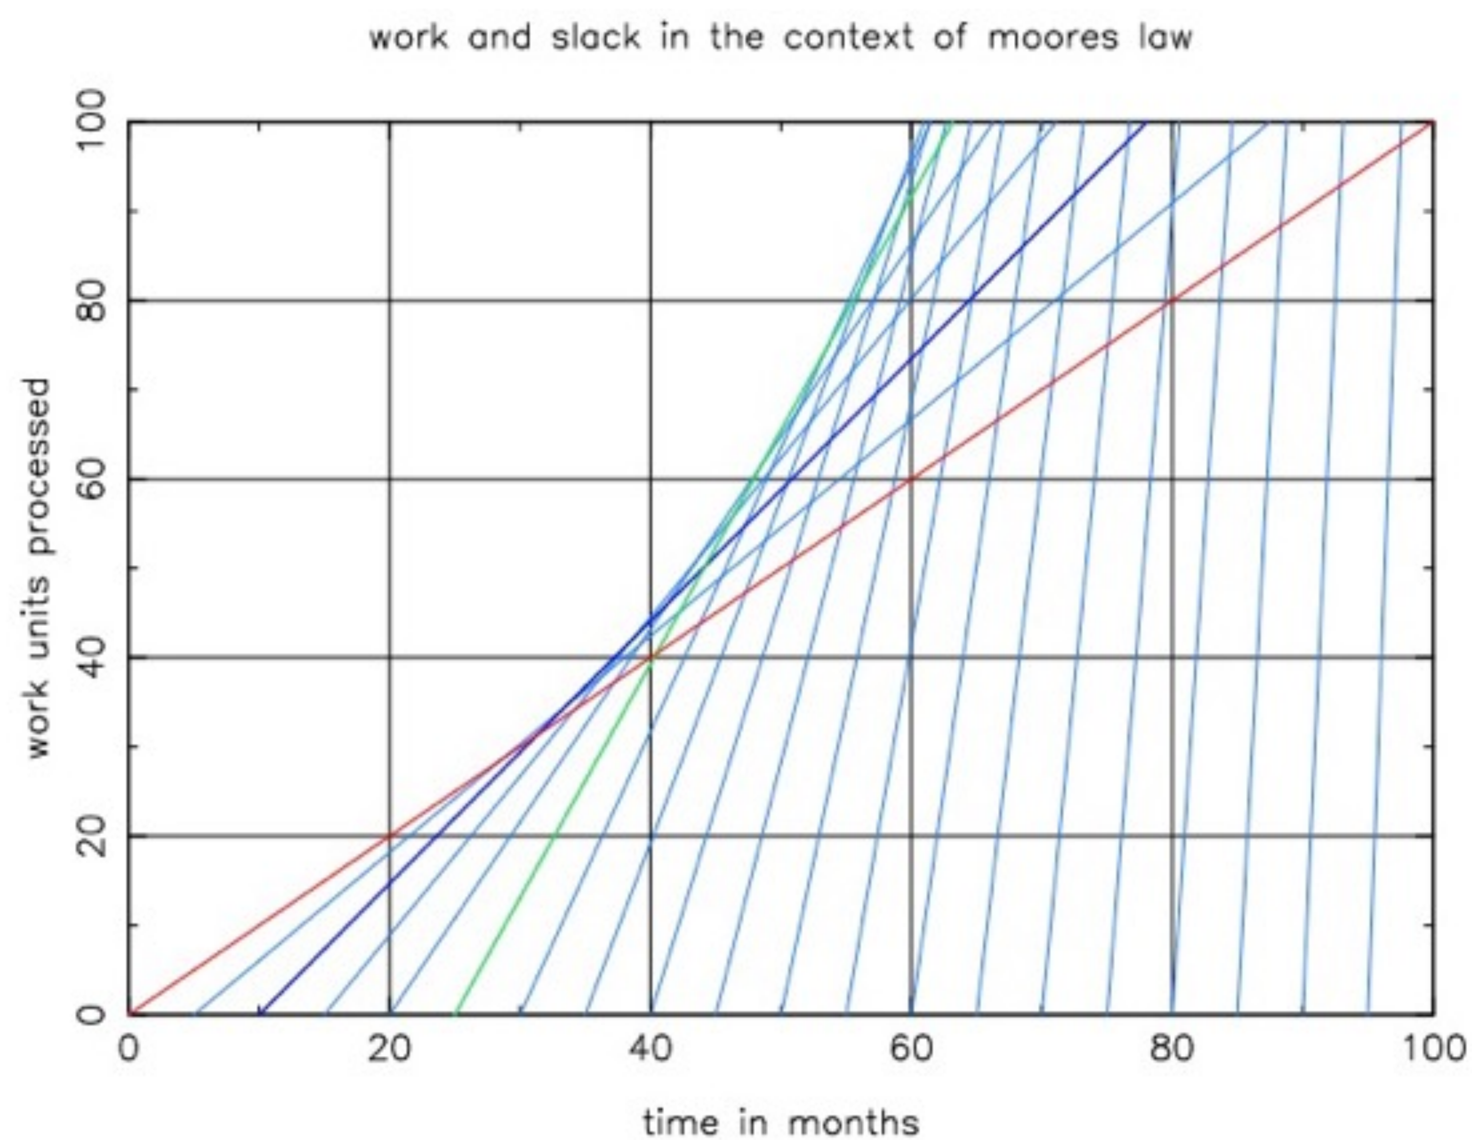


**astro-ph/9912202**

# GPUs provided a path forward using hierarchical parallelism

- **Expose more parallelism through code refactoring and source code directives**
  - Doubles performance (relative to CPUs) of many codes
- **Use right type of processor for each task**
- **Data locality: Keep data near processing**
  - GPU has high bandwidth to local memory for rapid access
  - GPU has large internal cache
- **Explicit data management: Explicitly manage data movement between CPU and GPU memories**

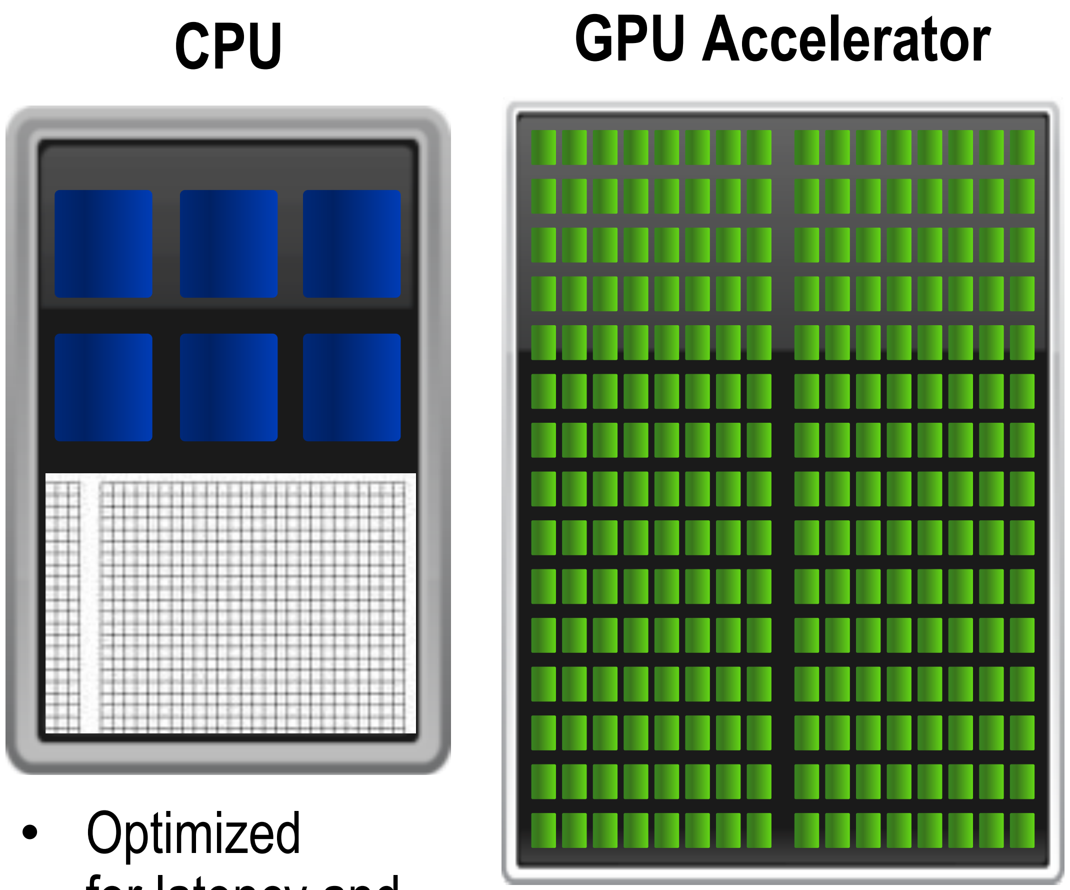


**CPU**
- Optimized for latency and sequential multitasking

**GPU Accelerator**
- Optimized for throughput and many simultaneous tasks
- 10× performance per socket
- 5× more energy-efficient systems

# ORNL's "Titan" hybrid system: Cray XK7 with AMD Opteron and NVIDIA Tesla

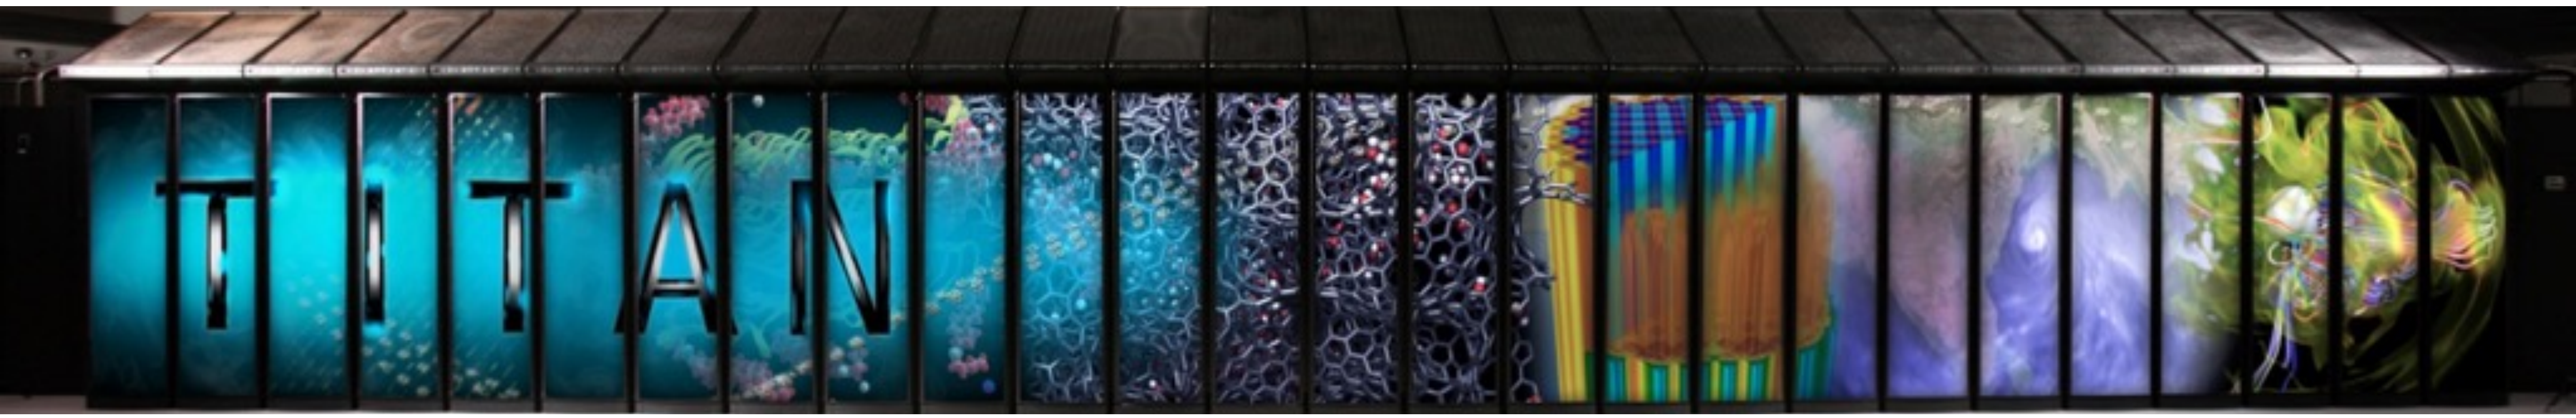


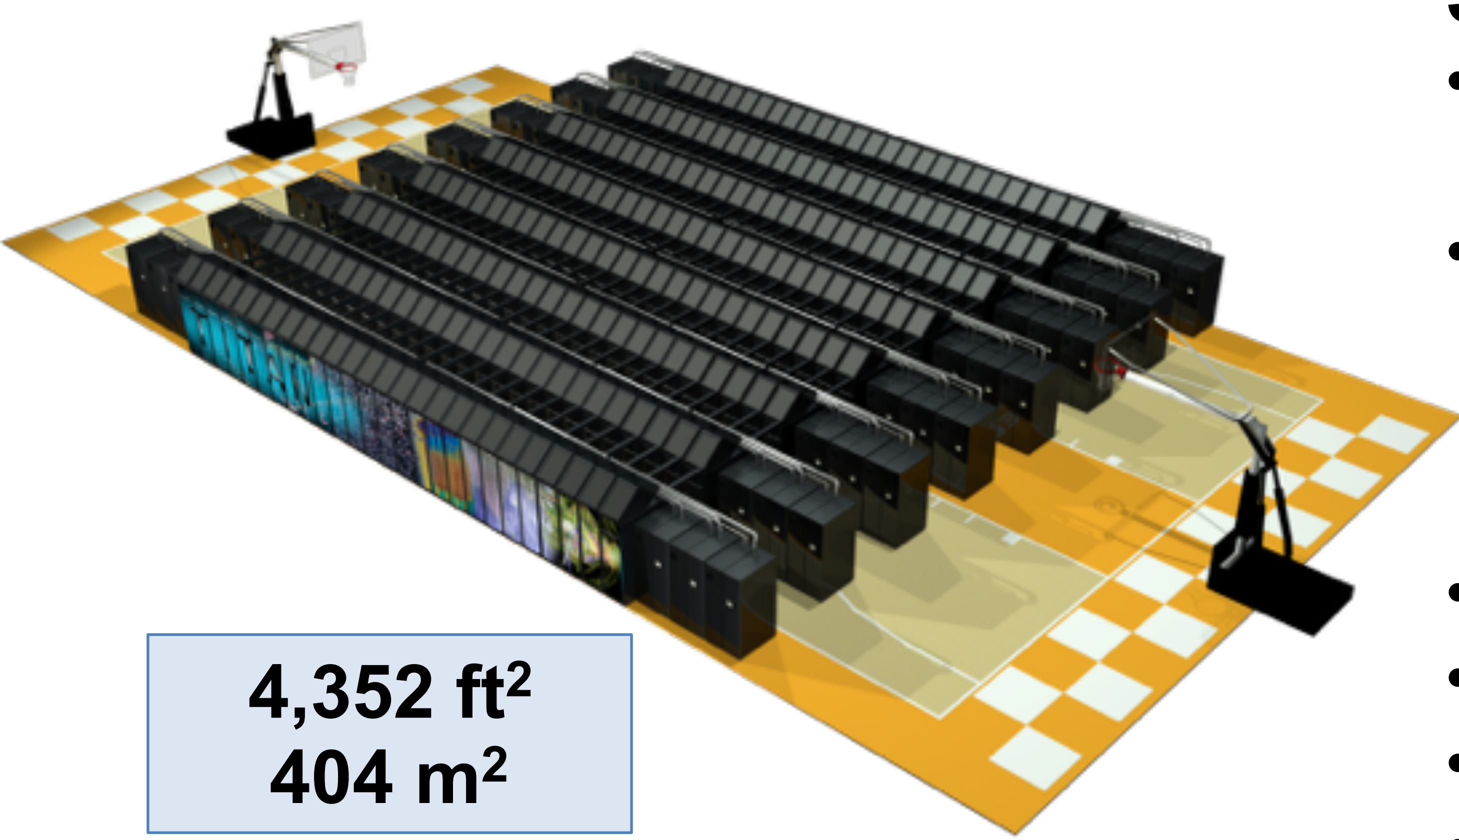

**4,352 ft$^2$**
**404 m$^2$**

**SYSTEM SPECIFICATIONS:**

- Peak performance of 27.1 PF
  - 24.5 GPU + 2.6 CPU
- 18,688 Compute Nodes each with:
  - 16-Core AMD Opteron CPU
  - NVIDIA Tesla "K20x" GPU
  - 32 + 6 GB memory
- 512 Service and I/O nodes
- 200 Cabinets
- 710 TB total system memory
- Cray Gemini 3D Torus Interconnect
- 8.8 MW peak power

# Impact on applications

- We began planning for Titan in 2009
- At the time there were no large-scale GPU systems deployed anywhere
- Furthermore, OLCF had little previous institutional knowledge of GPUs other than scattered individuals
- However, the consensus was that codes will require restructuring for memory locality, threading and heterogeneity to get to exascale—we decided to do it now


- Additionally, we didn't want a machine delivered that had no functioning application software
- Therefore, we selected a small set of applications for early porting, to spearhead an effort to move codes to Titan
- We went through a process to selected a diverse set of codes to give broad coverage to represent use cases for our users
- At the same time we wanted to capture institutional knowledge as lessons learned for going forward

# Center for Accelerated Application Readiness (CAAR)

**WL-LSMS**
Illuminating the role of material disorder, statistics, and fluctuations in nanoscale materials and systems.

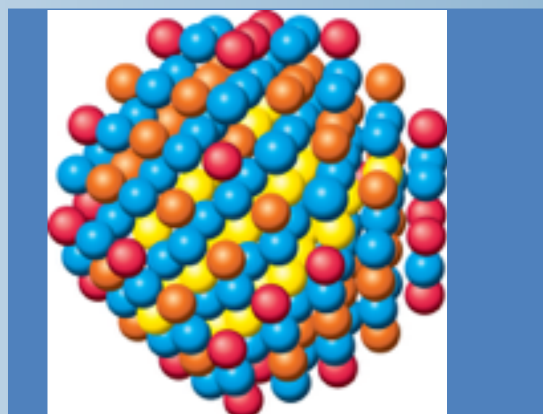

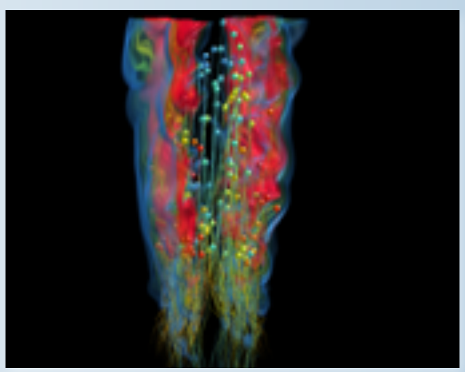

**S3D**
Understanding turbulent combustion through direct numerical simulation with complex chemistry.
.

**NRDF**
Radiation transport – important in astrophysics, laser fusion, combustion, atmospheric dynamics, and medical imaging – computed on AMR grids.

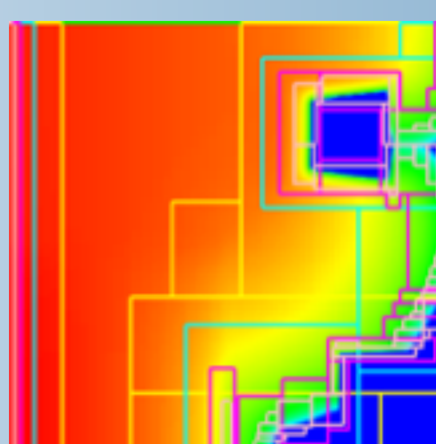

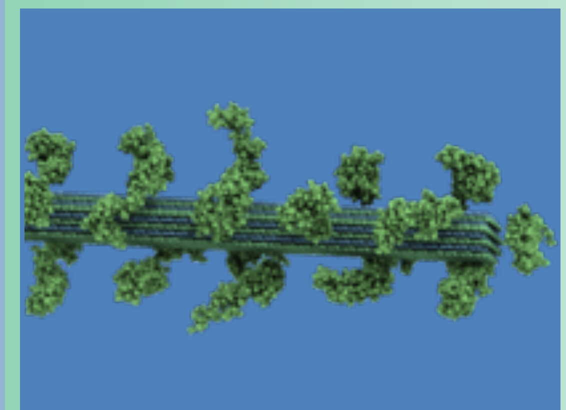

**LAMMPS**
A molecular description of membrane fusion, one of the most common ways for molecules to enter or exit living cells.

**CAM-SE**
Answering questions about specific climate change adaptation and mitigation scenarios; realistically represent features like precipitation patterns / statistics and tropical storms.

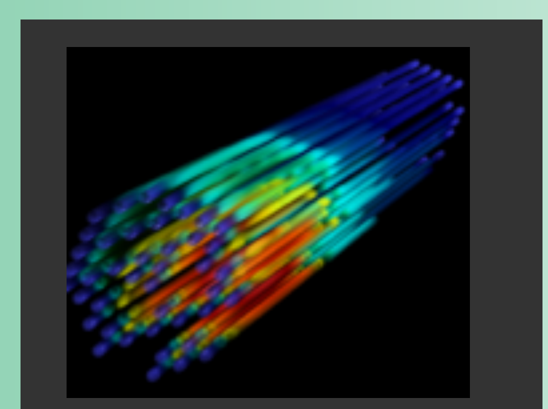

**Denovo**
Discrete ordinates radiation transport calculations that can be used in a variety of nuclear energy and technology applications.

# Performance results

| | XK6 (w/ GPU) vs. XK6 (w/o GPU) | XK6 (w/ GPU) vs. XE6 | Cray XK6: Fermi GPU plus Interlagos CPU<br>Cray XE6: Dual Interlagos and no GPU |
|---|---|---|---|
| **Application** | **Performance Ratio** | **Performance Ratio** | **Comment** |
| **S3D** | 1.5 | 1.4 | • Turbulent combustion<br>• **6% of Jaguar workload** |
| **Denovo** | 3.5 | 3.3 | • 3D neutron transport for nuclear reactors<br>• **2% of Jaguar workload** |
| **LAMMPS** | 6.5 | 3.2 | • High-performance molecular dynamics<br>• **1% of Jaguar workload** |
| **WL-LSMS** | 3.1 | 1.6 | • Statistical mechanics of magnetic materials<br>• **2% of Jaguar workload**<br>• 2009 Gordon Bell Winner |
| **CAM-SE** | 2.6 | 1.5 | • Community atmosphere model<br>• **1% of Jaguar workload** |

# Some lessons learned

- Up to 1-3 person-years required to port each code
  - Likely shorter today, given better tools, etc.
  - Also pays off for other systems—the ported codes often run significantly faster CPU-only (Denovo 2X, CAM-SE >1.7X)
- We estimate 70-80% of developer time was spent in code restructuring, regardless of whether using CUDA / OpenCL / OpenACC / …
- Each code team must make its own choice of using CUDA vs. OpenCL vs. OpenACC, based on the specific case—may be different conclusion for each code

# More lessons learned

- ***Code changes that have global impact on the code are difficult to manage, e.g., data structure changes. An abstraction layer may help, e.g., C++ objects/ templates***
- Tools (compilers, debuggers, profilers) were lacking early on in the project but are becoming more available and are improving in quality
- Debugging and profiling tools were useful in some cases (Allinea DT, CrayPat, Vampir, CUDA profiler)
- Science codes are under active development—porting to GPU can be pursuing a “moving target,” challenging to manage

# 2017 OLCF leadership system

## Hybrid CPU/GPU architecture

**Vendor: IBM (Prime) / NVIDIA™ / Mellanox Technologies®**

**5X-10X Titan's Application Performance**

**Approximately 4,500 nodes, each with:**

- Multiple IBM POWER9 CPUs and multiple NVIDIA Tesla® GPUs using the NVIDIA Volta architecture
- CPUs and GPUs connected via high speed NVLink
- Large coherent memory: over 512 GB/node (HBM + DDR4)
  - all directly addressable from the CPUs and GPUs
- An additional 800 GB of NVRAM, which can be configured as either a burst buffer or as extended memory
- Over 40 TF peak performance

**Dual-rail Mellanox® EDR-IB full, non-blocking fat-tree interconnect**

**IBM Elastic Storage (GPFS™) - 1TB/s I/O and 120 PB disk capacity.**

# CAAR II


- New application readiness activity for Summit
- Changes in management, etc. from CAAR I
  - Call for proposals
    - better coordination with code teams
    - code teams have skin in the game
  - CSEEN postdocs work with OLCF SciComp members
    - training in code development
    - science projects using new code
- Increased emphasis on platform portability

# CAAR (II) Applications

| Domain | Application | Methods | PI | Institution |
|---|---|---|---|---|
| *Astrophysics* | FLASH | Grid, AMR | Bronson Messer | ORNL |
| *Chemistry* | DIRAC | Particle, LA | Lucas Visscher | VUA |
| *Climate Science* | ACME (N) | Unstr Mesh | David Bader | LLNL |
| *Engineering* | RAPTOR | Kokkos | Joseph Oefelein | SNL |
| *Materials Science* | QMCPACK | MC | Paul Kent | ORNL |
| *Nuclear Physics* | NUCCOR | Particle | Gaute Hagen | ORNL |
| *Plasma Physics* | XGC (N) | PIC, PETSc | CS Chang | PPPL |
| *Seismic Science* | SPECFEM | Unstr Mesh | Jeroen Tromp | Princeton |
| *Astrophysics* | HACC(N,A) | Grid | Salman Habib | ANL |
| *Biophysics* | NAMD (N) | Particle | Klaus Schulten | UIUC |
| *Chemistry* | NWCHEM (N) | Particle, LA | Karol Kowalski | PNNL |
| *Chemistry* | LSDALTON | Particle, LA | Poul Jørgensen | Aarhus |
| *Plasma Physics* | GTC (N) | PIC | Zhihong Lin | UCI |

Selected via a call for proposals in Feb. 2015

# Early-access system being installed at OLCF now

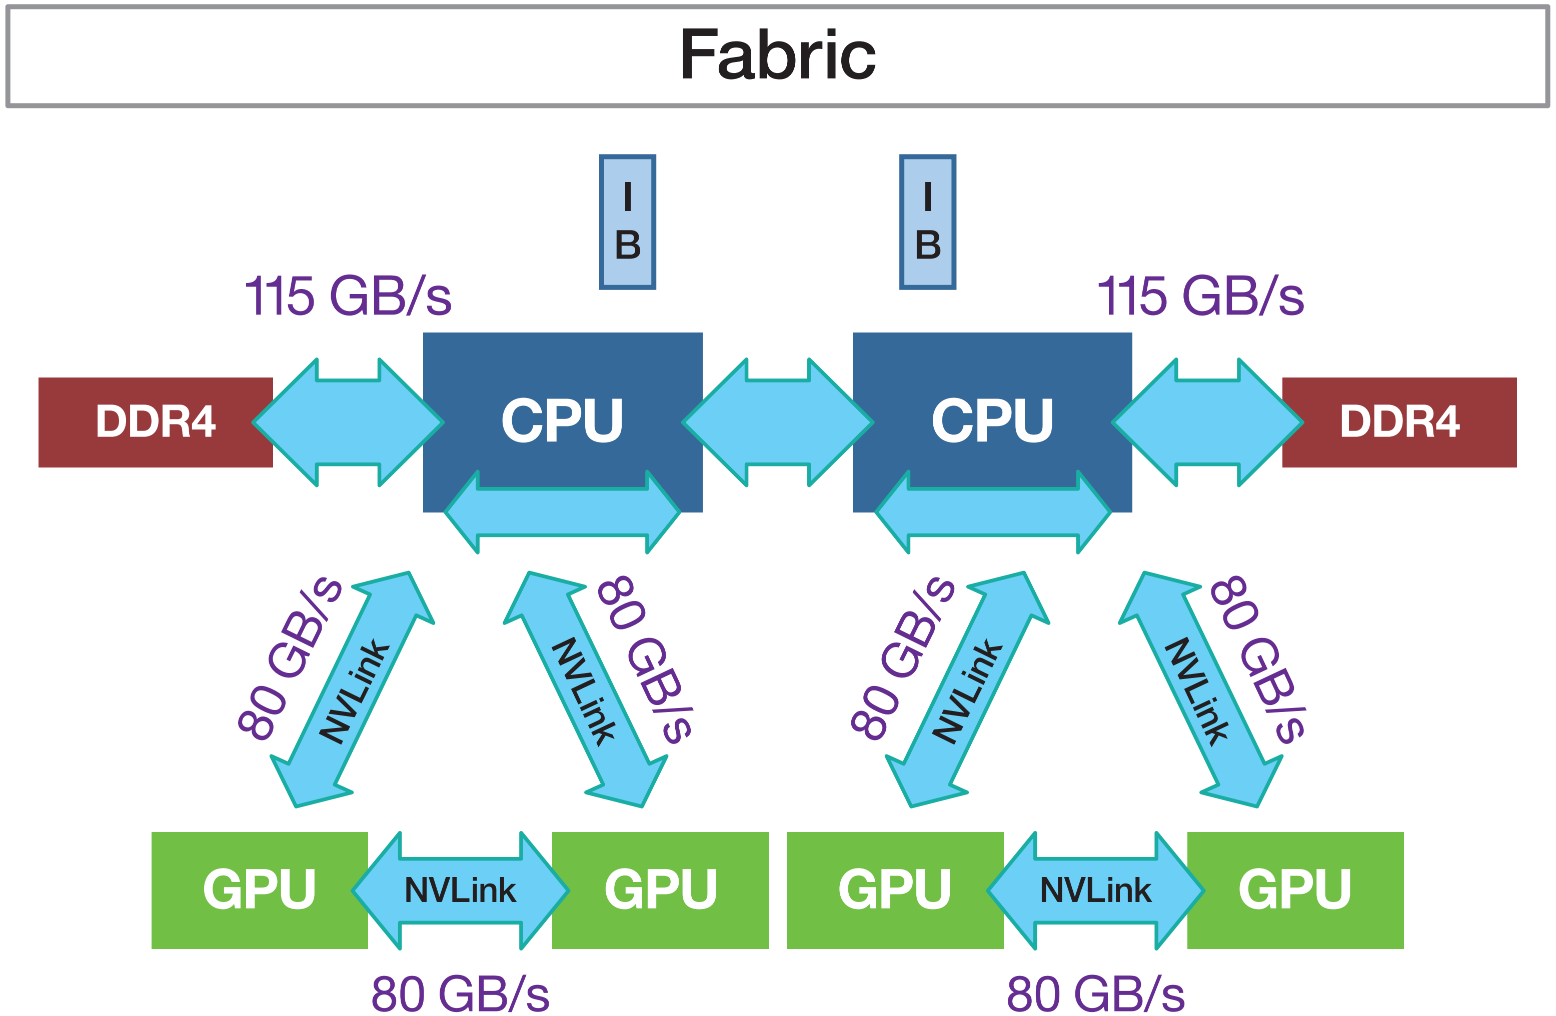


## IBM Power System S822LC for High Performance Computing

| IBM Power System S822LC for High Performance Computing at a glance | |
|---|---|
| **System configurations (8335-GTB)** | |
| Microprocessors | Two 8-core 3.25 GHz POWER8 processor cards or<br>two 10-core 2.86 GHz POWER8 processor cards |
| Level 2 (L2) cache | 512 KB L2 cache per core |
| Level 3 (L3) cache | 8 MB L3 cache per core |
| Level 4 (L4) cache | Up to 64 MB per socket |
| Memory Min/Max | 4 GB, 8 GB, 16 GB, 32 GB DDR4 modules, 128 GB to 1 TB total memory |
| Processor to Memory Bandwidth | 115 GB/sec per socket, 230 GB/sec per system (Max sustained memory bandwidth to L4 cache from SCM)<br>170 GB/sec per socket, 340 GB/sec per system (Max peak memory bandwidth to DIMMs from L4 cache) |

# What's next? - Two architecture paths for today and future Leadership Systems

**Power concerns for large supercomputers are driving the largest systems to either Hybrid or Many-core architectures**

**Hybrid Multi-Core (like Titan/Summit)**

- CPU / GPU hybrid systems
- Multiple CPUs and GPUs per node
- Smaller number of very powerful nodes
- Data movement issues to be much easier than previous systems – NVLINK coherent shared memory within a node
- Multiple levels of memory – on package, DDR, and persistent

**Many Core (Mira/Aurora)**

- 10's of thousands of nodes with millions of cores
- Homogeneous cores
- Multiple levels of memory – on package, DDR, and non-volatile
- Intel OmniPath interconnect

# Next-gen systems: Cori, Aurora, Summit


| Attributes | 2016 Cori (NERSC) | 2018 Aurora (Argonne) | 2018 Summit (Oak Ridge) |
|---|---|---|---|
| Peak PF | >30 | 180 | 200 |
| Power MW | <3.7 | 13 | 13.5 |
| Processors | Intel Xeon Phi (KNL) and Haswell | Intel Xeon Phi (KNH) | IBM Power9 + NVIDIA Volta |
| Sys. Mem. | ~1PB DDR4 + HBM + 1.5PB persistent | >7 PB HBM + Local + persistent | > 1.7 PB HBM+Local + 2.8 PB persistent |
| Nodes | 9,300 compute + 1,900 data | > 50,000 | ~4,500 |
| File System | 28 PB @ 744 GB/s Lustre | 150 PB @ 1 TB/s Lustre | 250PB @ 2.5 TB/s GPFS |

# Architecture and performance portability

## Portability is important for leadership-class computational science teams

- PIs often have access to multiple architectures at any given time
- Code development teams often target users on multiple platforms
- Applications have much longer lifespans than computer architectures
  - Portability can assist applications in smoother transition to future architectures
- Porting to modern parallel platforms is labor intensive
  - Exposing parallelism through refactoring is the primary task

## Approach: ASCR facilities coordinate application readiness activities

- Points of contact at ALCF, NERSC and OLCF for applications
- Formulation of joint performance portability plan
- Collaboration on porting of common applications to reduce duplication of effort
- Common resources for application readiness programs

# API standardization: Past experience

1990s shared memory

1990s distributed memory

2010s manycore

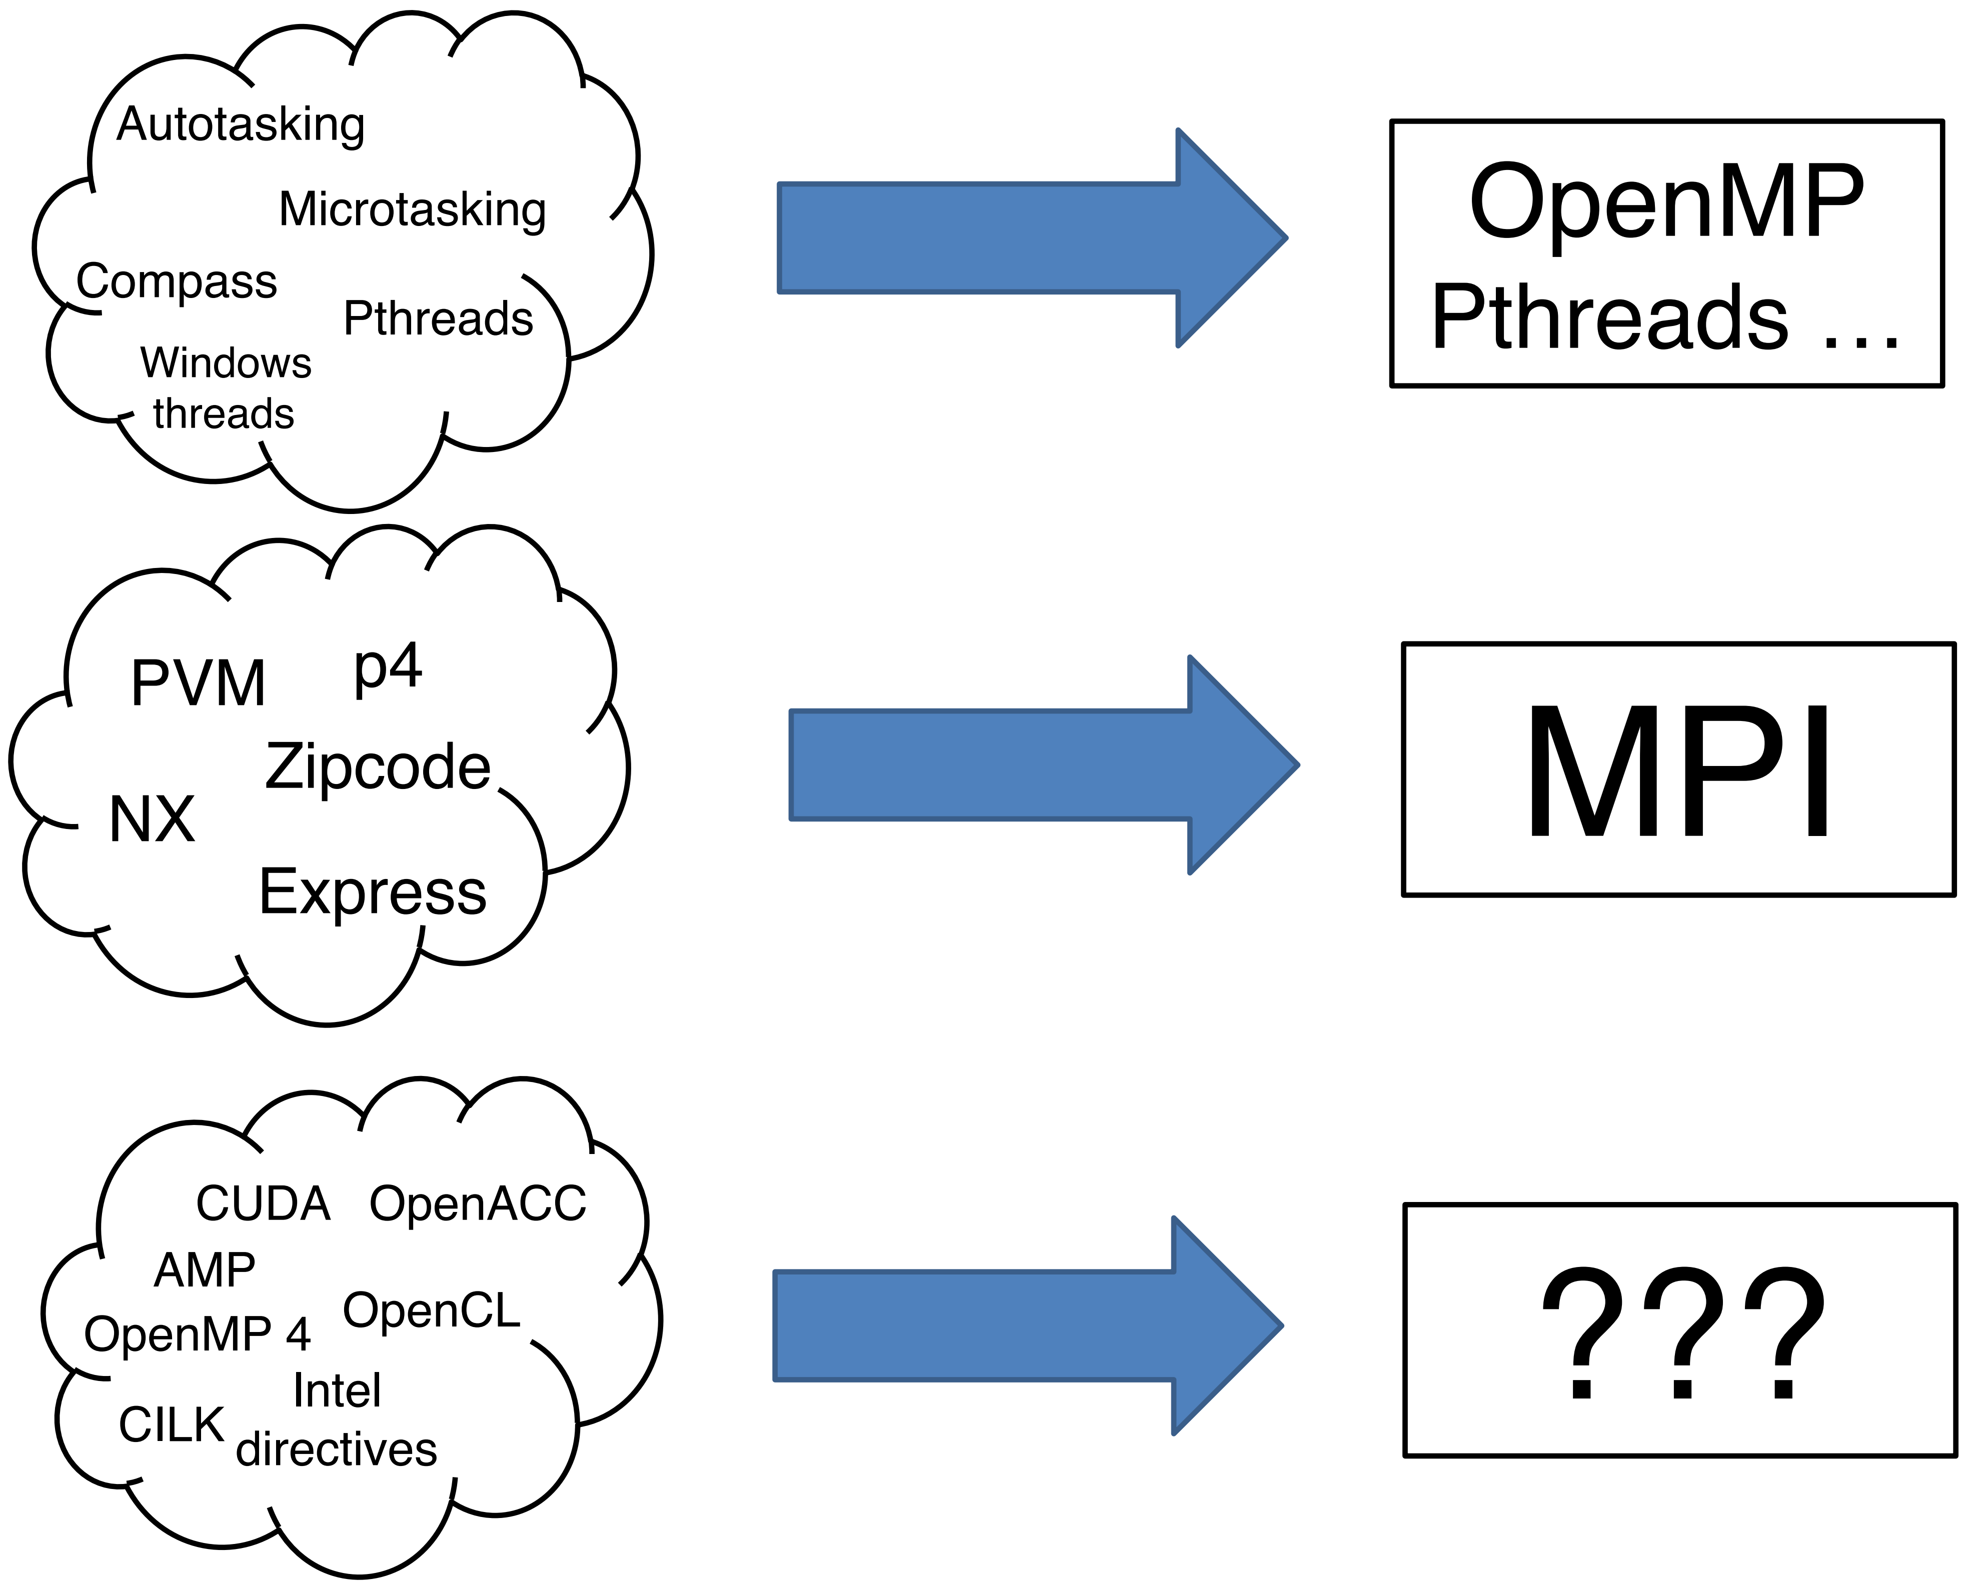

# Three primary user programs for access to LCF

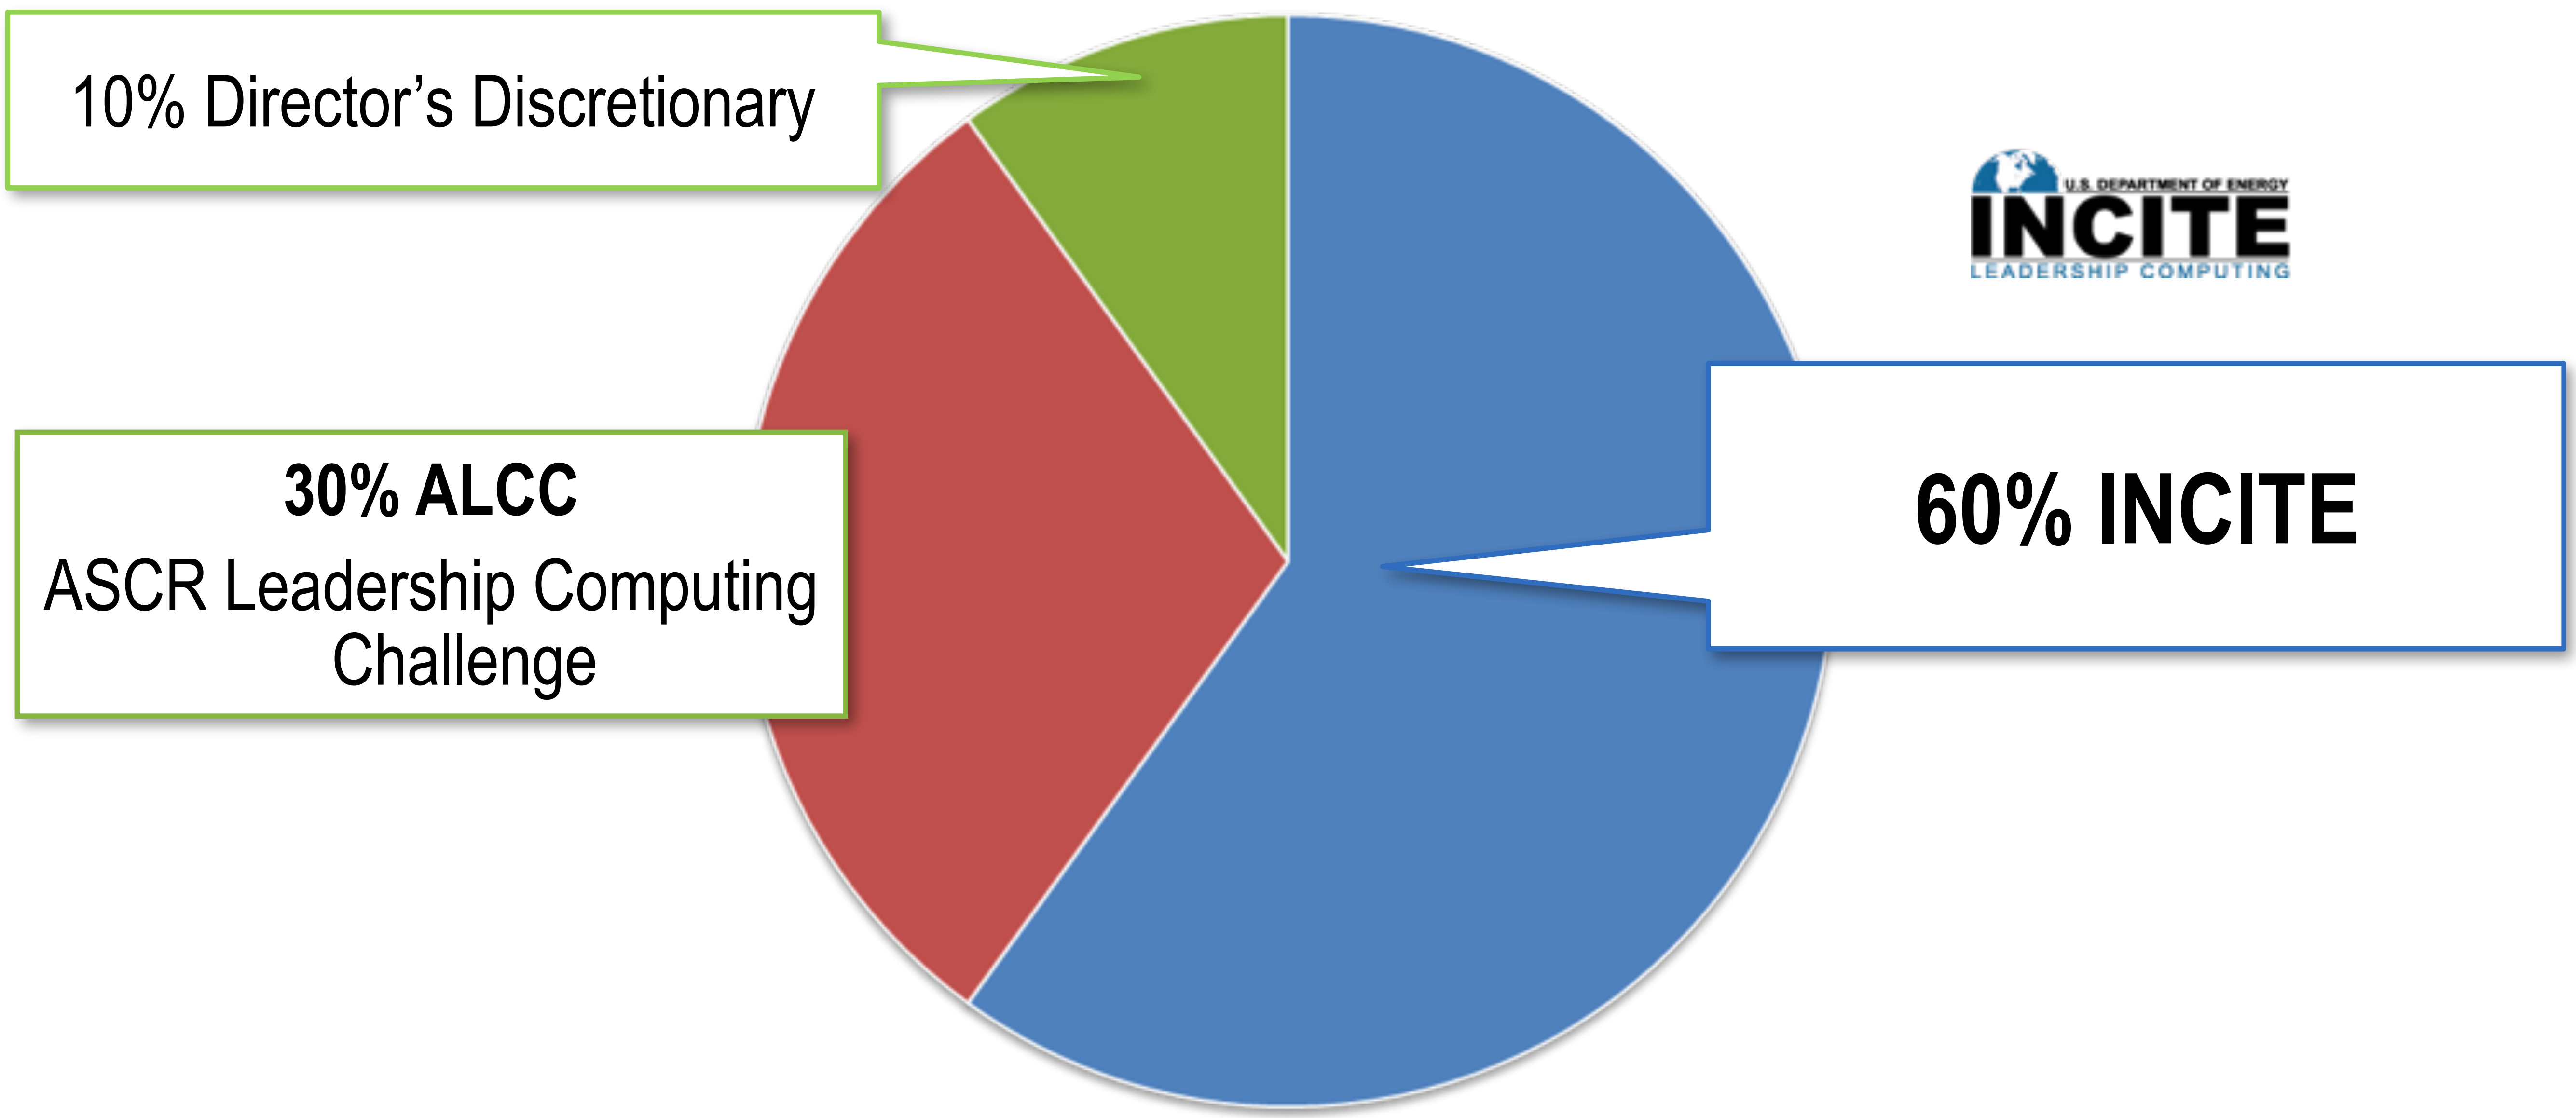


## Distribution of allocable hours

# LCFs support models

- “Two-pronged” support model is shared
- Specific organizational implementations differ slightly but user perspective is virtually identical

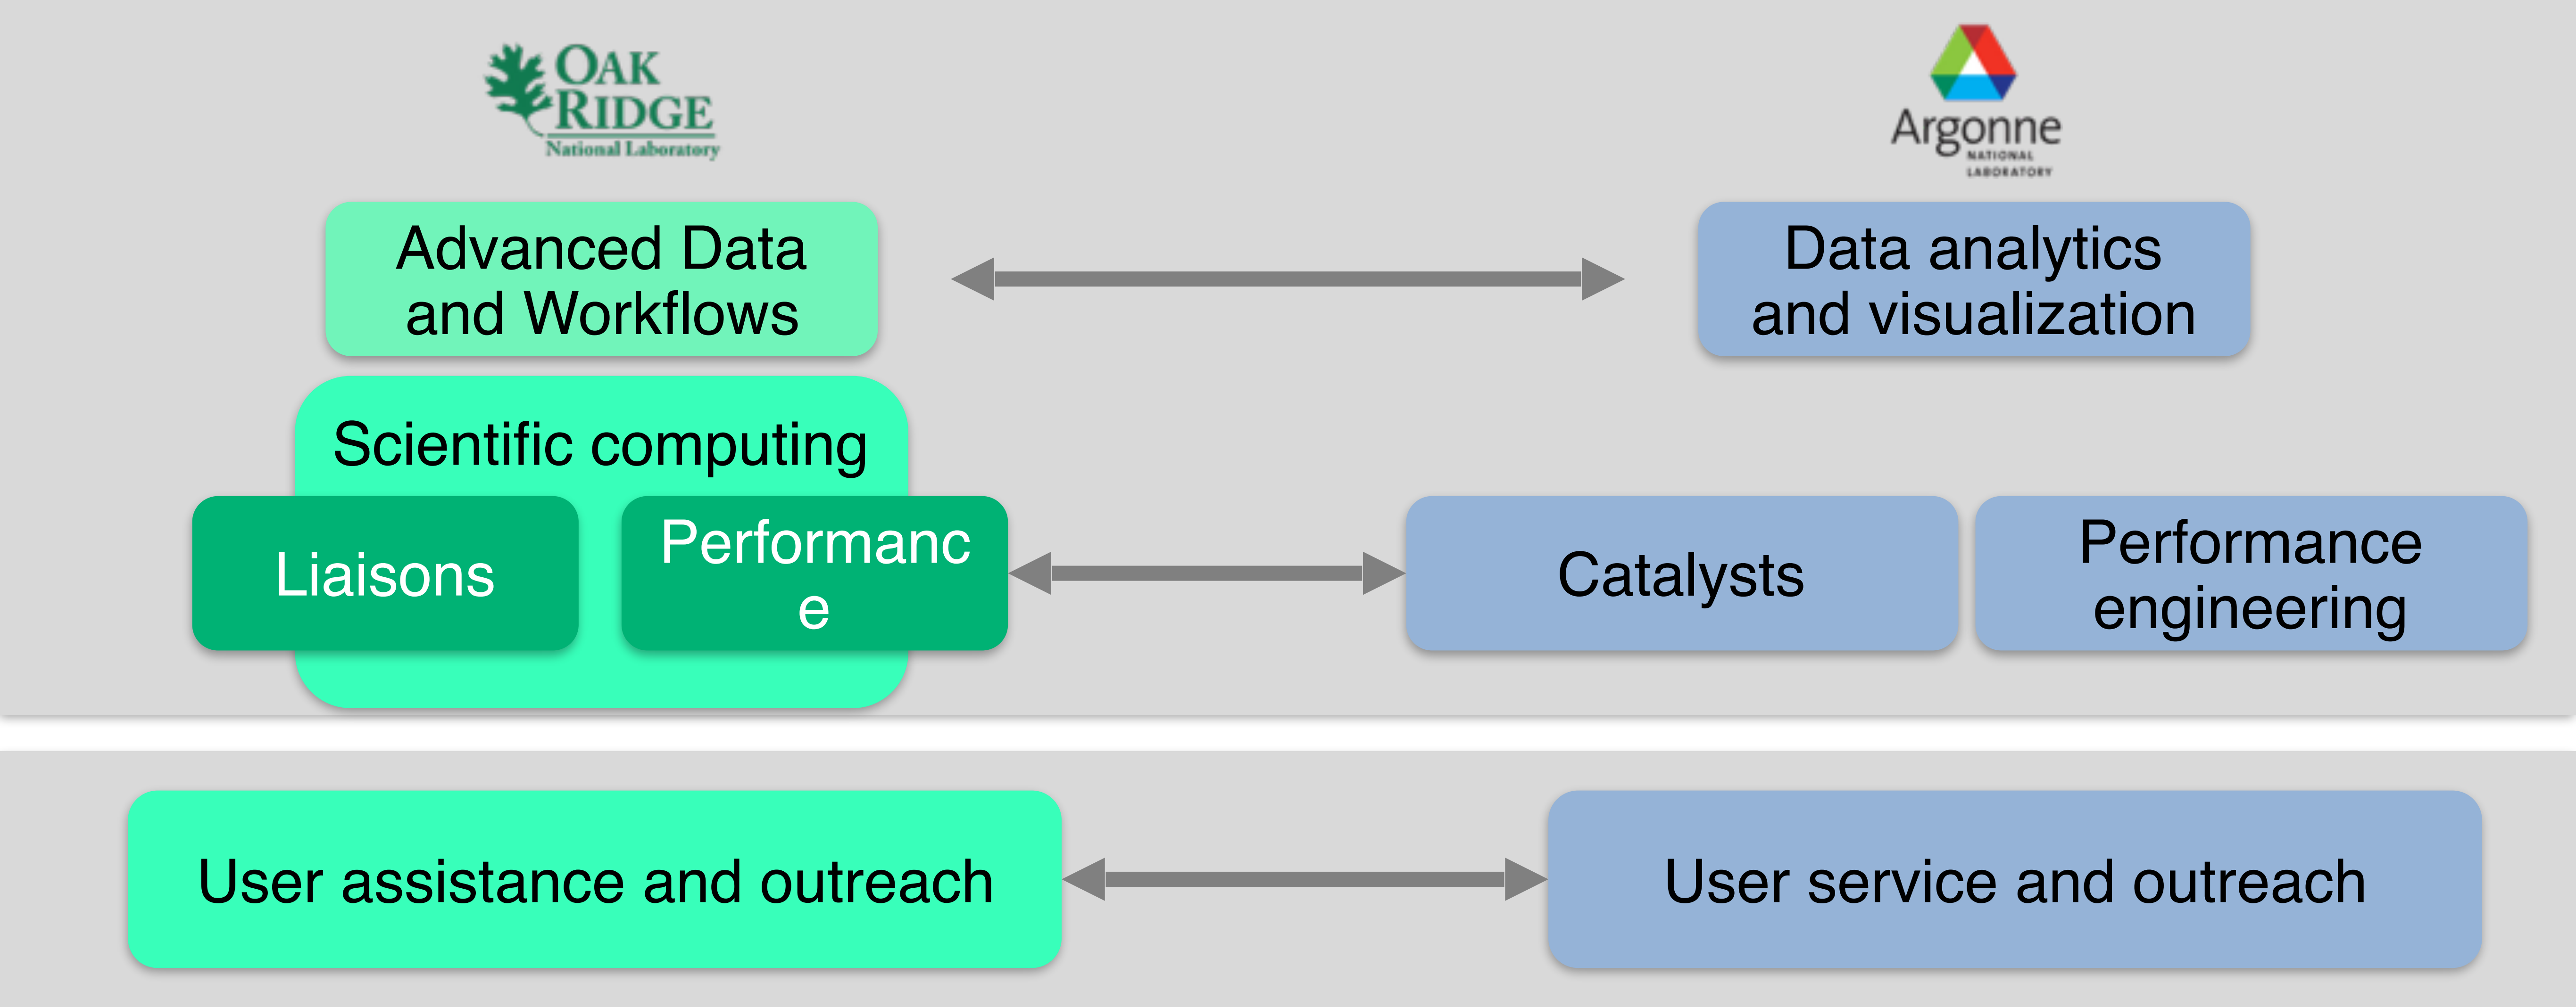

# Support basics

- User Assistance group provides “front-line” support for day-to-day computing issues
- SciComp Liaisons provide advanced algorithmic and implementation assistance
- Assistance in data analytics and workflow management, visualization, and performance engineering
can also provided for each project (both tasks are “housed” in ADW group at OLCF)

# BigPanda, ATLAS, and Titan

- ATLAS production is integrated with Titan
- Titan has already contributed a noticeable fraction of computing resources for MC simulations
- NYU and Manhattan College groups calculated the vector boson fusion channel for Higgs production and delivered more than 15 million fully simulated events on Titan, leading to an early physics results publication.

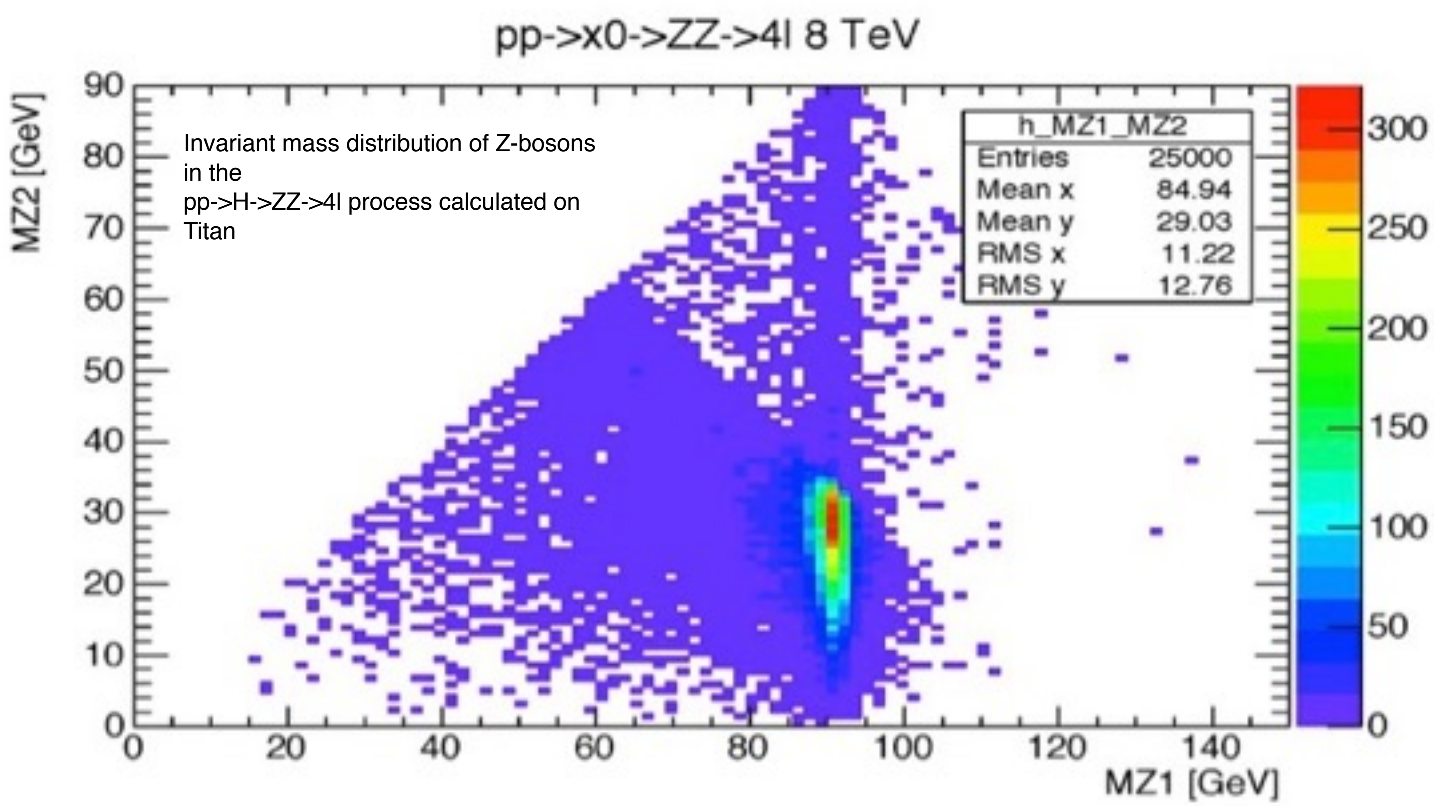


OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

# Titan Contribution to the ATLAS Scientific Program.

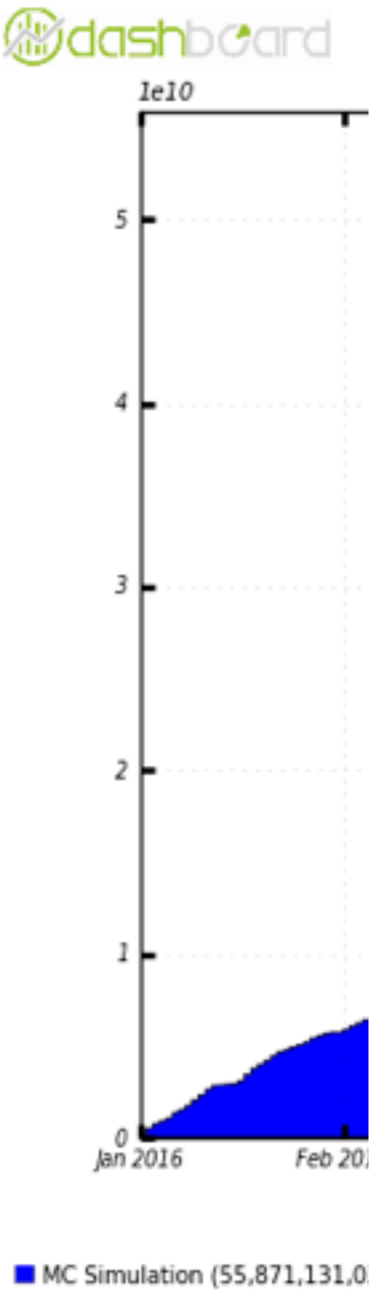


ATLAS Monte-Carlo Jobs monthly usage
Backfill mode (Jan-Aug)

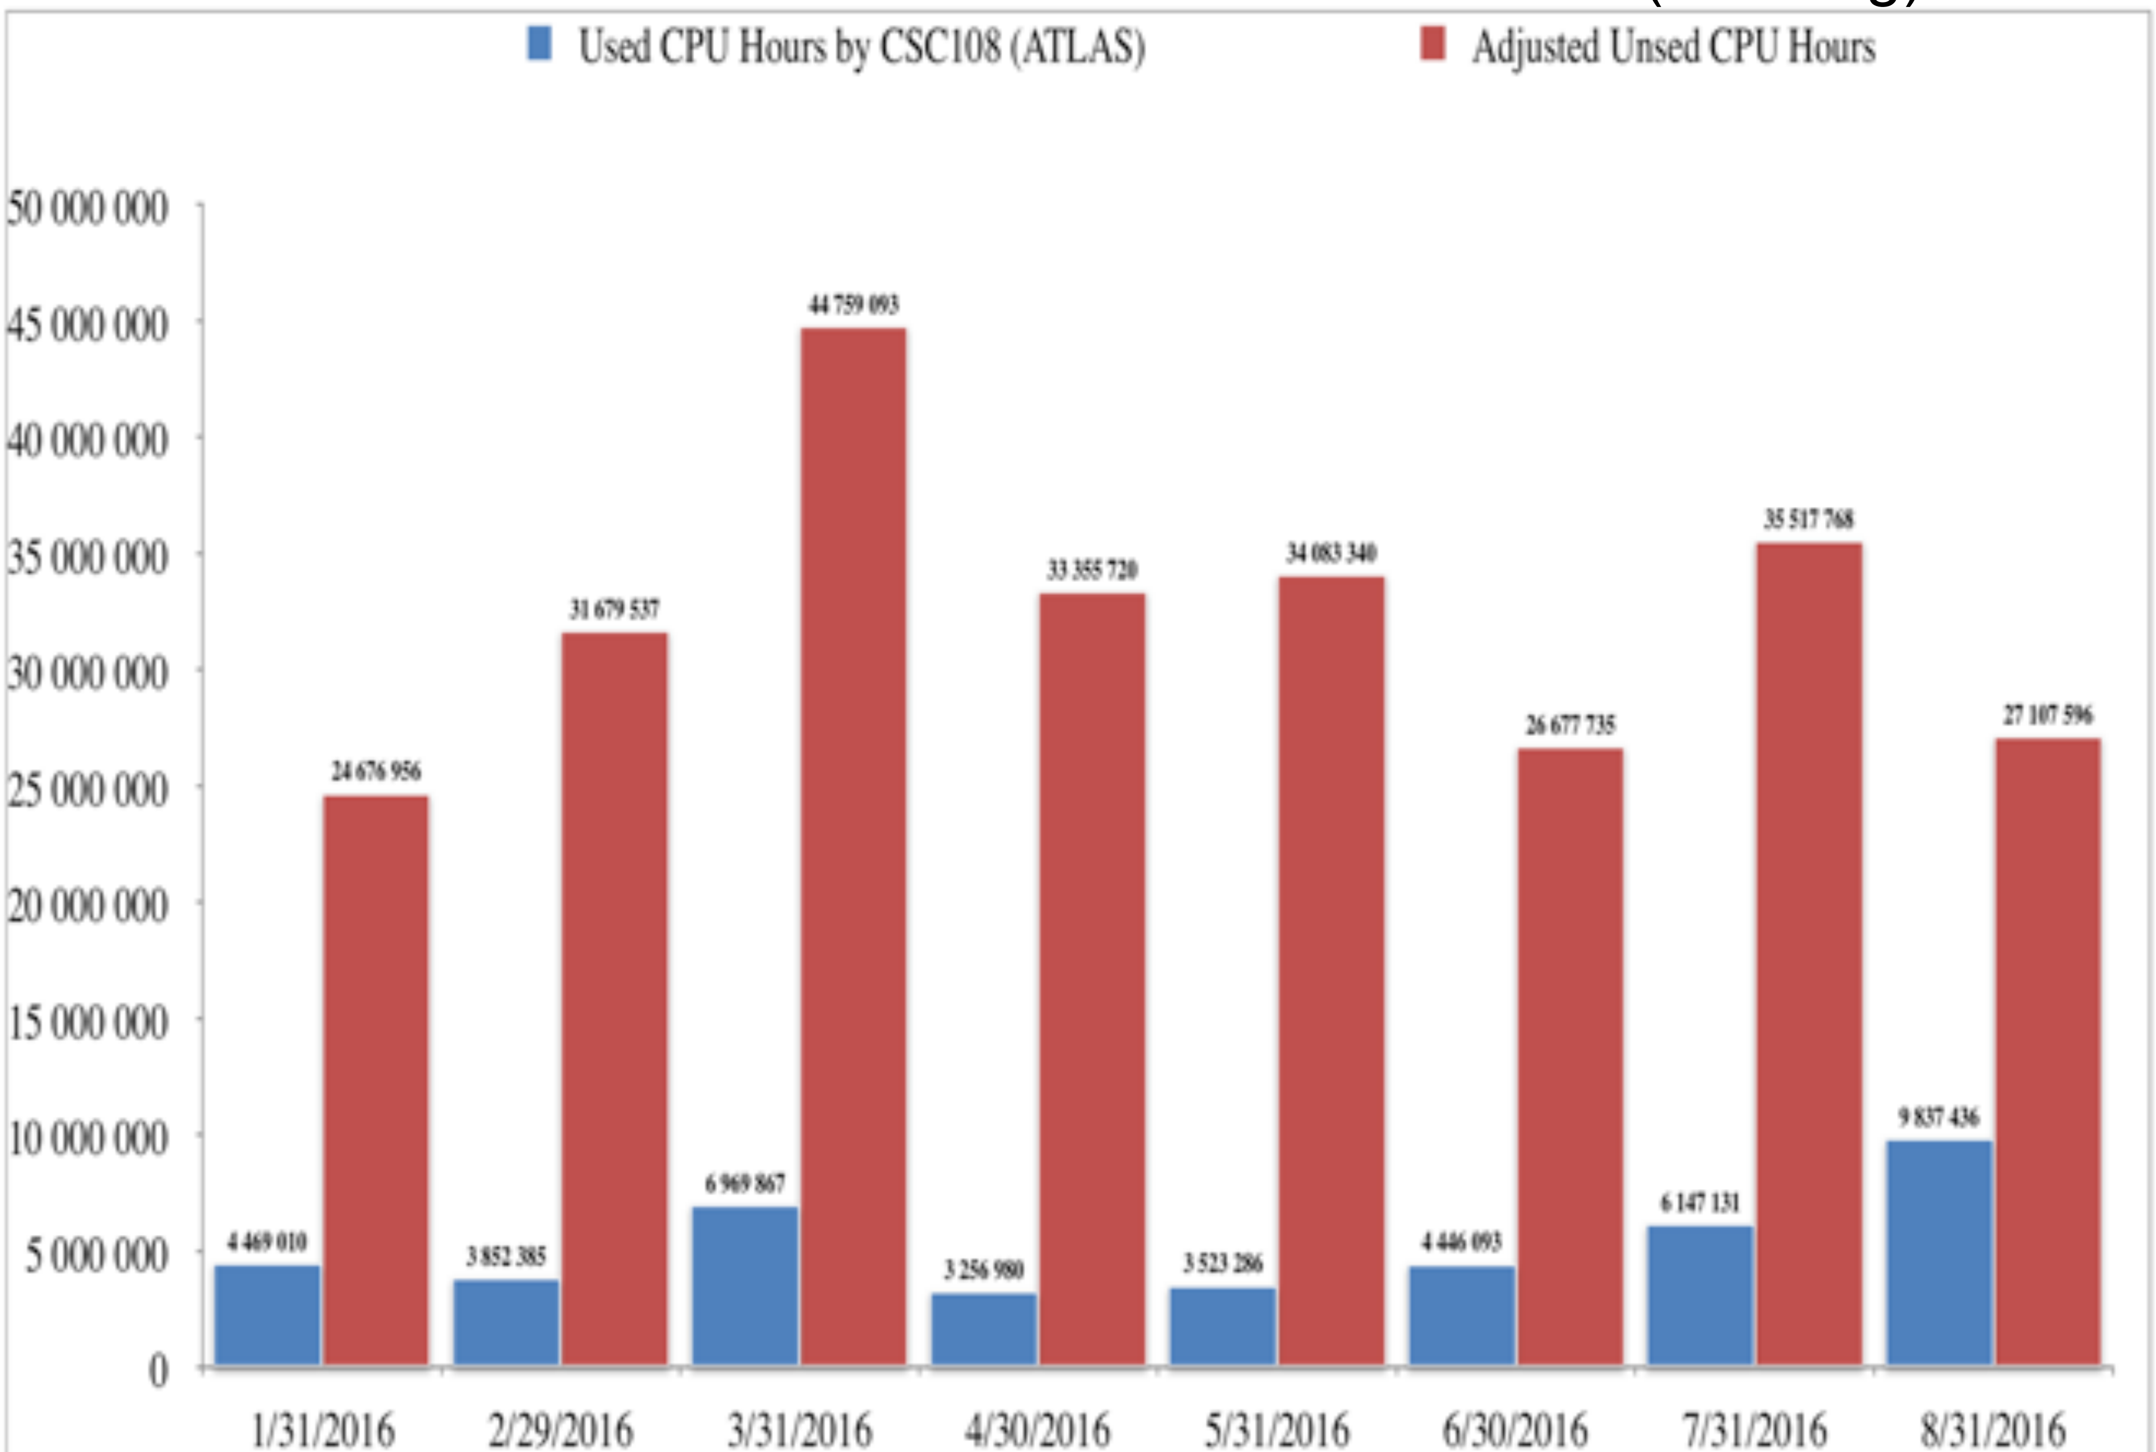


Alexei Klimentov

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

# Summary

- Jumping into the deep in of the pool with Titan was jarring, but is making the transition to Summit much easier (even pleasant!)
  - Exposing as much parallelism as possible — regardless of the new platform — is essential.
- Programming models are still evolving, but one has to start somewhere
- Several allocation programs available
- OLCF is a bespoke shop: We tailor solutions via collaborations with SciComp and ADW groups
- "On demand" computing is not easy for LCF's, but sheer scale can lead to some possibilities.

# Questions?

bronson@ornl.gov