Argonne NATIONAL LABORATORY

The Argonne Leadership Computing Facility

Scott Parker

U.S. DEPARTMENT OF ENERGY

Argonne and the Blue Gene

- **2005:**
  - Argonne accepts 1 rack (1024 nodes) of Blue Gene/L (5.6 TF)
- **2006:**
  - Argonne Leadership Computing Facility (ALCF) created
- **2008:**
  - ALCF accepts 40 racks (160k nodes) of Blue Gene/P (557 TF)
- **2009:**
  - ALCF approved for 10 petaflop system to be delivered in 2012
- **2012**
  - 10 PF Mira Blue Gene/Q hardware delivered to ALCF
- **2013:**
  - Mira in production

ALCF Systems

- ***Mira* – BG/Q system**
  - 49,152 nodes / 786,432 cores
  - 768 TB of memory
  - Peak flop rate: 10 PF
  - Linpack flop rate: 8.6 PF (#5 Top 500)
- ***Cetus & Vesta (T&D)* - BG/Q systems**
  - 4K & 2k nodes / 64k & 32k cores
  - 64 TB & 32 TB of memory
  - 820 TF & 410 TF peak flop rate
- ***Tukey* – Nvidia system**
  - 100 nodes / 1600 x86 cores/ 200 M2070 GPUs
  - 6.4 TB x86 memory / 1.2 TB GPU memory
  - Peak flop rate: 220 TF
- **Storage**
  - 30 PB capacity, 240 GB/s bw (GPFS)
  - Storage upgrade planned in 2015

ALCF Systems

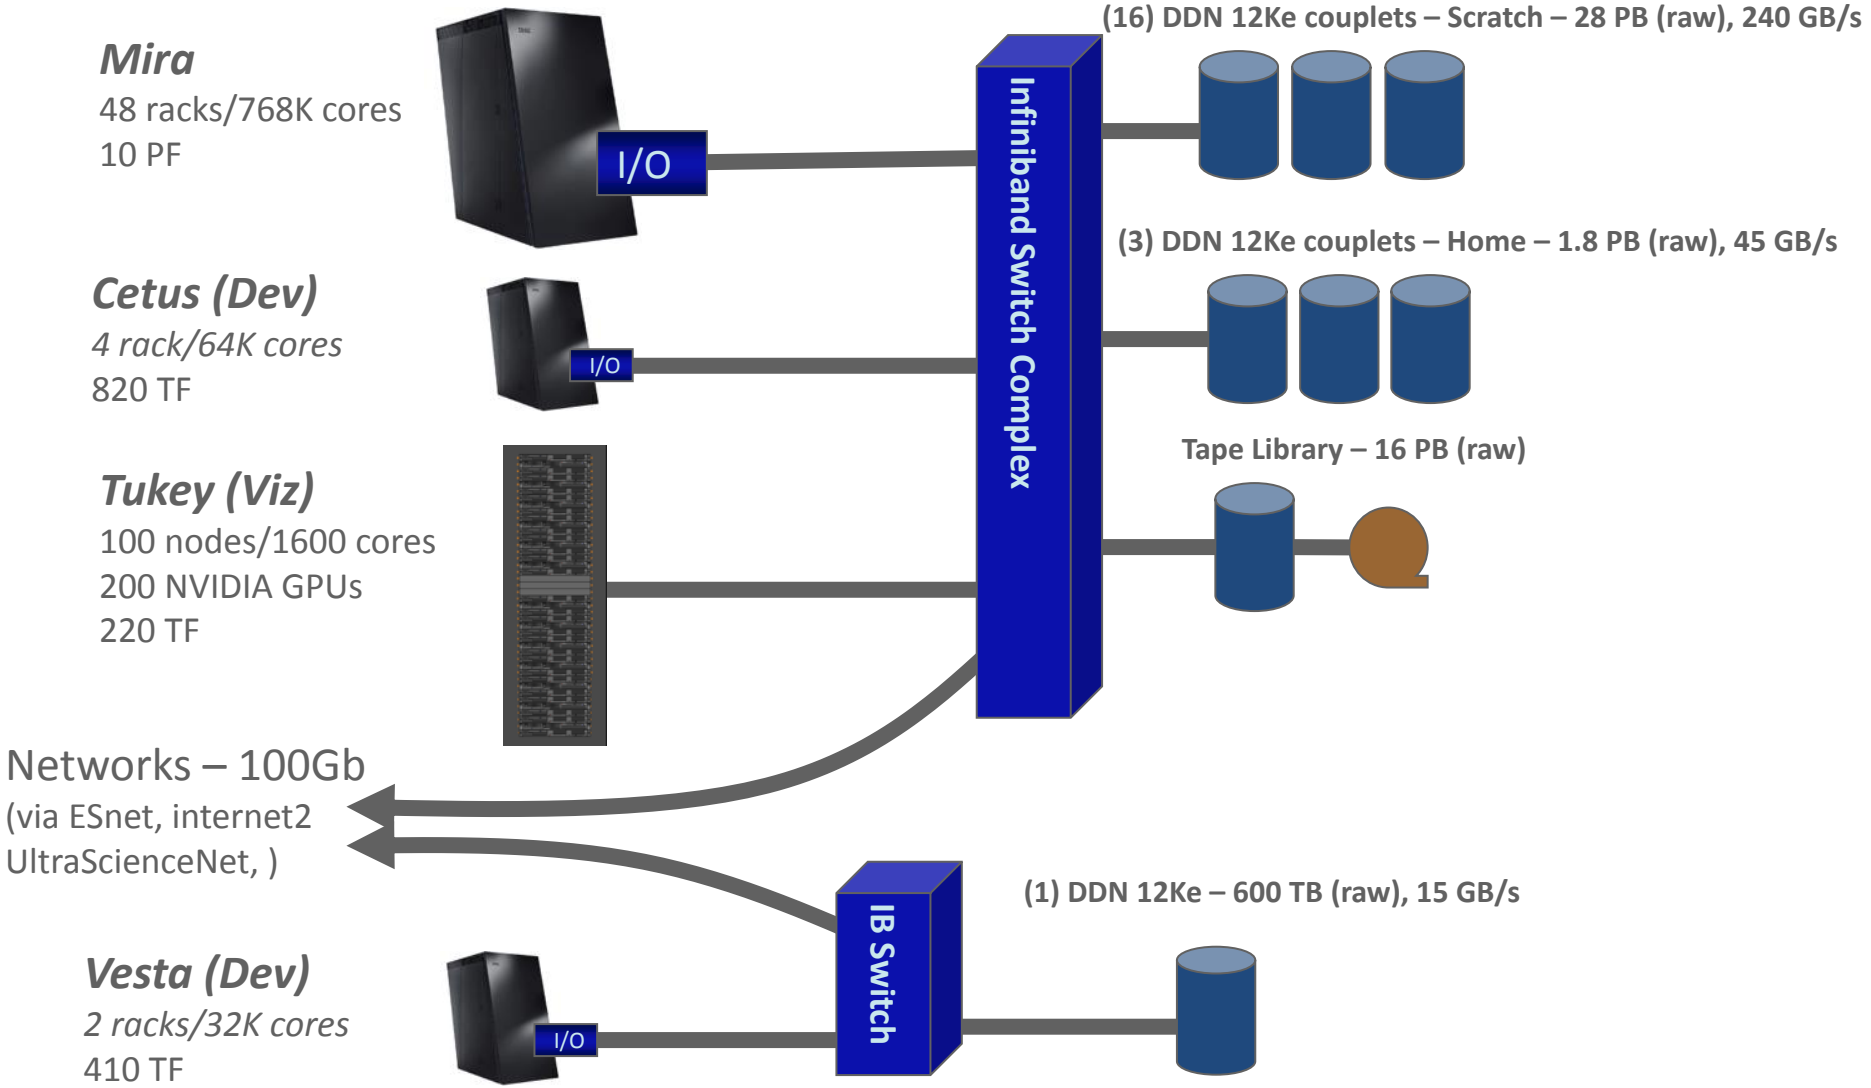

Blue Gene DNA

- **Leadership computing power**
  - Leading architecture since introduction, #1 half Top500 lists over last 10 years
- **Low speed, low power**
  - Embedded PowerPC core with custom SIMD floating point extensions
  - Low frequency (L – 700 MHz, P – 850 MHz, Q – 1.6 GHz)
- **Massive parallelism:**
  - Multi/Many core (L -2, P – 4, Q – 16)
  - Many aggregate cores (L – 208k, P – 288k, Q – 1.5M)
- **Fast communication network(s)**
  - Torus network (L & P – 3D, Q – 5D)
- **Balance:**
  - Processor, network, and memory speeds are well balanced
- **Minimal system overhead**
  - Simple lightweight OS (CNK) minimizes noise
- **Standard Programming Models**
  - Fortran, C, C++, & Python languages supported
  - Provides MPI, OpenMP, and Pthreads parallel programming models
- **System on a Chip (SoC) & Custom designed ASIC (Application Specific Integrated Circuit)**
  - All node components on one chip, except for memory
  - Reduces system complexity and power, improves price / performance
- **High Reliability:**
  - Sophisticated RAS (reliability, availability, and serviceability)
- **Dense packaging**
  - 1024 nodes per rack

ALCF Projects Span Many Domains

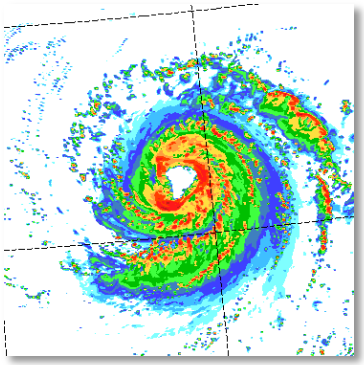

Climate

Predicting hurricane tracks to mitigate risks, hindcasting with climate model data to gauge impact of global change.

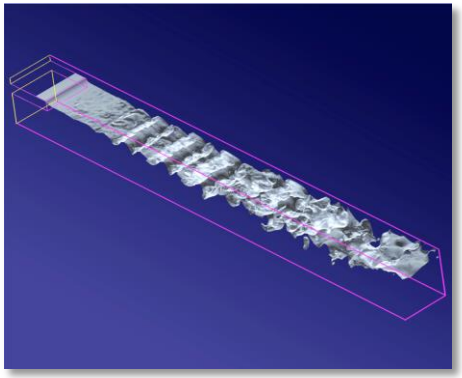

Engineering

Drag reduction by injecting air layer along ship hull

Fusion Energy

Control laser backscatter in inertial confinement fusion experiments at National Ignition Facility.

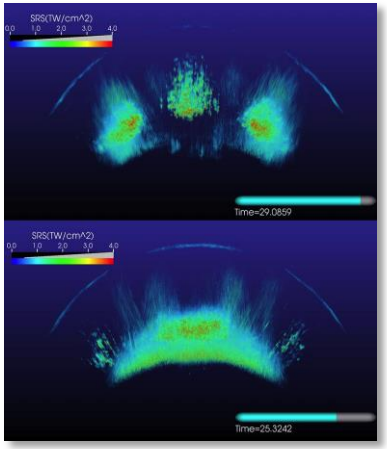

Nano Catalysts

Mapping out properties of gold nanoparticles to design catalysts for fuel cells and methane conversion.

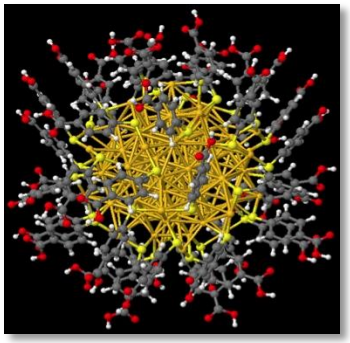

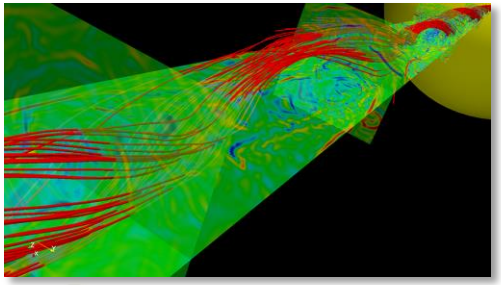

Space Plasmas

Investigate Alfvén waves as mechanism to heat solar corona.

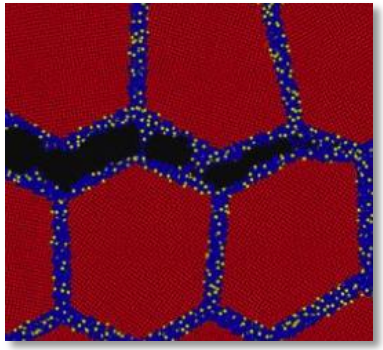

Materials Science

Molecular simulation of fracture dynamics in structural materials in next-generation nuclear reactors.

Allocation Programs at the LCFs

	60% INCITE		30% ALCC		10% Director's Discretionary	
Mission	High-risk, high-payoff science that requires LCF-scale resources		High-risk, high-payoff science aligned with DOE mission		Strategic LCF goals	
Call	1x/year – (Closes June)		1x/year – (Closes February)		Rolling	
Duration	1-3 years, yearly renewal		1 year		3m,6m,1 year	
Typical Size	30 - 40 projects	50M - 500M core-hours/yr.	5 - 10 projects	10M – 300+M core-hours/yr.	100s of projects	.5M – 10M core-hours
Review Process	Scientific Peer-Review	Computational Readiness	Scientific Peer-Review	Computational Readiness	Strategic impact and feasibility	
Managed By	INCITE management committee (ALCF & OLCF)		DOE Office of Science		LCF management	
Readiness	**High**		**Medium to High**		**Low to High**	
Availability	Open to all scientific researchers and organizations Over 5 billion core hours available in 2015					

Applying for a Discretionary Account

To apply for a discretionary account go to:

http://www.alcf.anl.gov/getting-started/apply-for-dd

Contact me:

Scott Parker (sparker@anl.gov)