

Computing Resources and Scientific Programs at Argonne Leadership Computing Facility

Wei Jiang

Leadership Computing Facility
Argonne National Laboratory

02/06/2015

DOE ASCR Computing Facilities

1. National Energy Research Scientific Computing Center (NERSC)



Principal provider of HPC services to **Office of Science programs**
Hopper: Cray XE6, 1.28 petaflops, 153,216 cores
Edison: Cray XC30, 2.39 petaflops, 124,608 cores
ALCC

2. Oak Ridge Leadership Computing Facility (OLCF)



Titan, Cray XT7, 27 petaflops
AMD Opteron CPU and NVIDIA Kepler GPU
18,688 CPU nodes & 18,688 GPUs
Computational science community: INCITE/ALCC

3. Argonne Leadership Computing Facility (ALCF)



Mira, IBM Blue Gene/Q, 10 petaflops
16 1.6GHZ PowerPC A2 core / Node, 4 hardware threads /
core, 49,152 nodes or 786,432 cores.
Computational science community: INCITE/ALCC

ALCF Computing Resources



Mira – the primary resource
48 racks, 786,432 cores
Hybrid MPI(PAMI)/OpenMP(POSIX)
Project/Size dependent job scheduler
Open science (any researcher can start from Director's Discretionary)



Cetus and Vesta: Test & Development Systems

Cetus – Debug/test machine for Mira
4 BGQ rack (4096 nodes)
Identical file system with Mira
Simple job scheduler

Vesta – Code porting/optimization/development
2 BGQ rack
Simple job scheduler

Resource Network

Mira

48 racks/768K cores
768 TB RAM
10 PF




Cetus (Dev)

4 rack/64K cores
64 TB RAM
838 TF



Tukey (Viz)

96 nodes/1536 cores
192 NVIDIA GPUs
6.1TB / 1.1TB GPU RAM
220 TF




Networks – 100Gb
(via ESnet, internet2
UltraScienceNet,)

Vesta (Dev)


2 racks/32K cores
32TB RAM
419 TF




GPFS file system




Mira-fs1
(6) DDN 12Ke couplets
7 PB (usable), **90 GB/s**



Mira-fs0
(16) DDN 12Ke couplets
20 PB (usable), **240 GB/s**
MHD/Rayleigh is here!



Mira-home
(3) DDN 12Ke couplets
1.1 PB (usable), **45 GB/s**



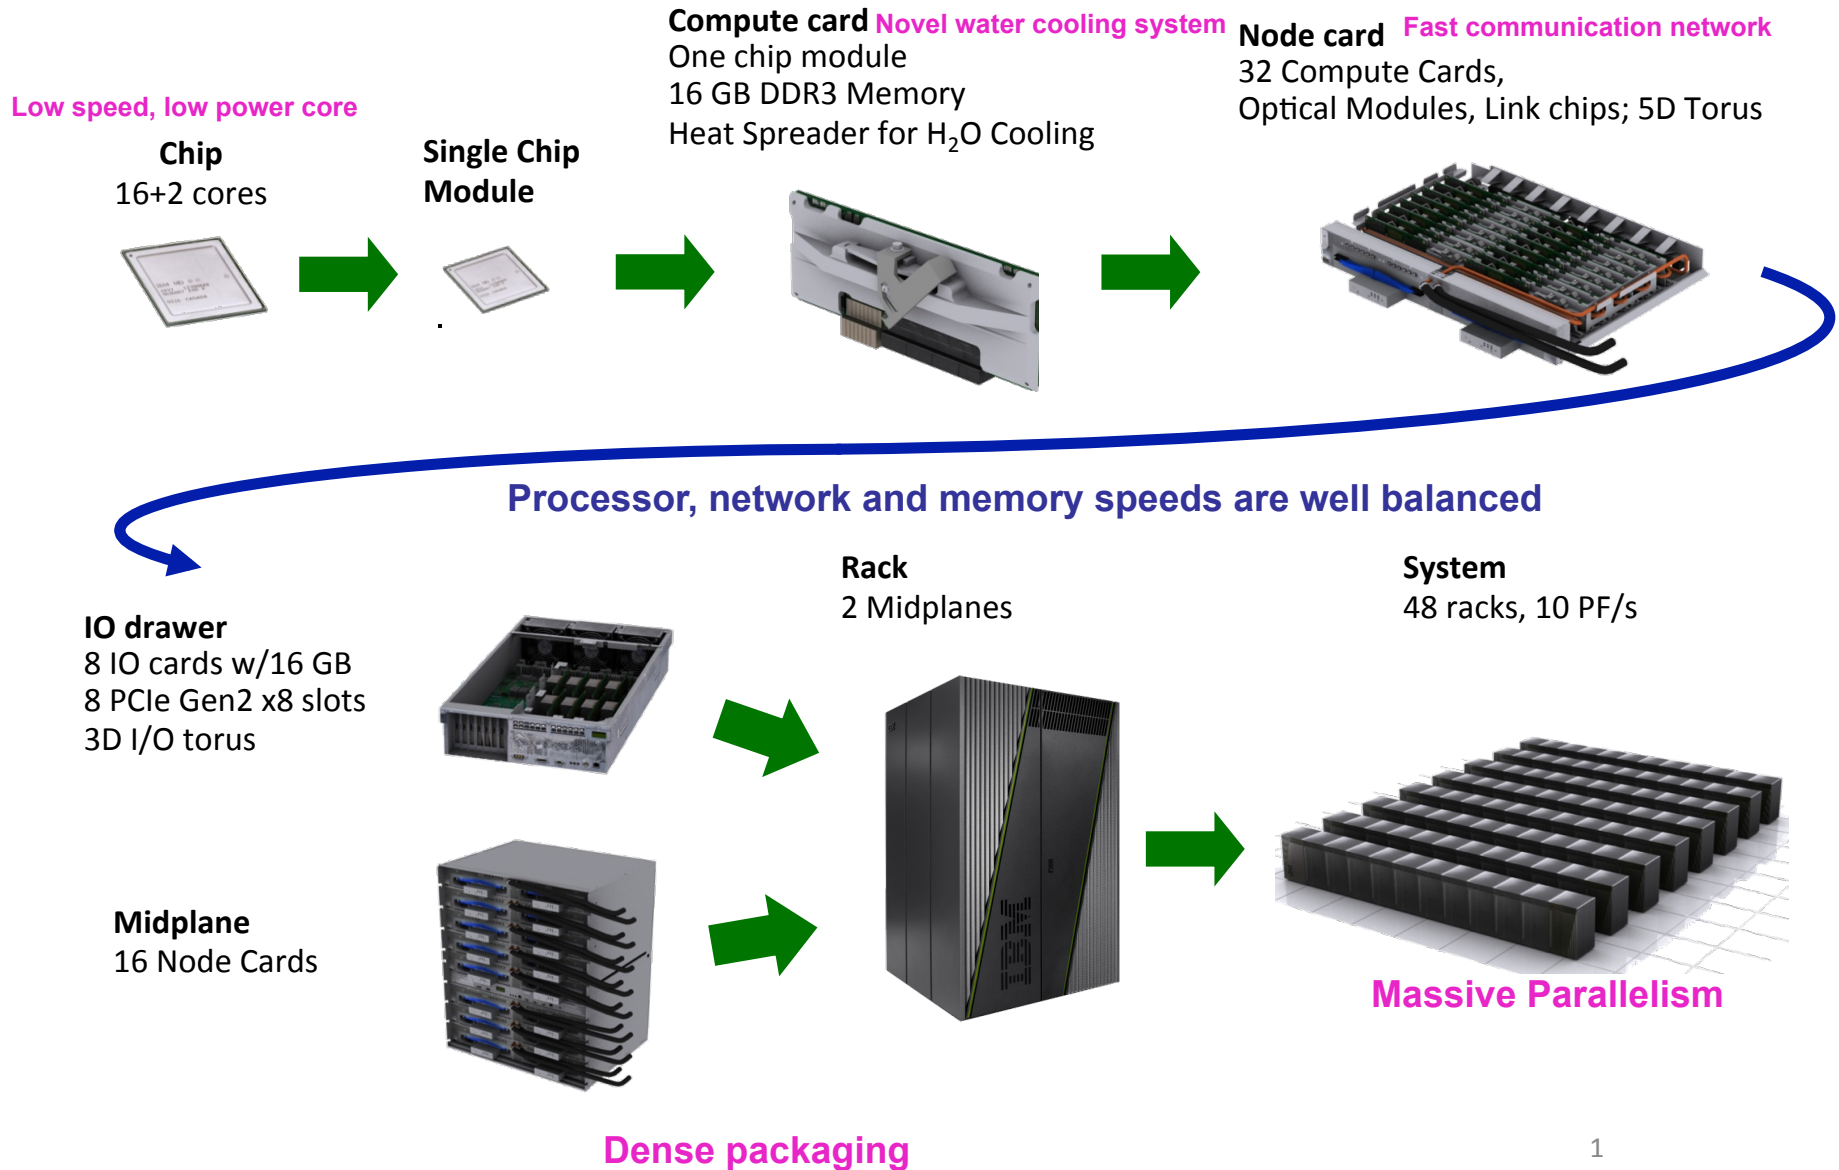
HPSS
(2) Spectralogic T-950
10,000 slot libraries
(48) LTO6 Tape Drives



(1) DDN 12Ke – 600 TB (raw), **15 GB/s**



Hardware Hierarchy of Blue Gene/Q



Application Software Stack

- Standard programming models:
 - Standard programming models
 - Fortran, C/C++, Python/Perl/Tcl.
 - MPI/OpenMP/Pthreads
- ALCF supports two sets of libraries:
 - IBM system and provided libraries: `/bgsys/drivers/ppcfloor`
 - glibc
 - MPI
 - PAMI (Parallel Active Messaging Interface)
 - Application development toolchains (compiler, linker, assembly, debugger etc)
 -
- Site supported libraries and programs: `/soft/libraries`
 - ESSL, PETSc, HDF5, netCDF, Parallel netCDF, Boost
 - ESSL is IBM's optimized Engineering and Scientific Subroutine library for BG/Q: BLAS, LAPACK, FFT, sort/search, interpolation, quadrature, random numbers, BLACS
 - Additional tuned libraries in `/soft/libraries/alcf` subdirectory
 - BLAS, CBLAS, FFTW2, FFTW3, LAPACK, METIS, PARMETIS, PARPACK, SCALAPACK, SILO, SZIP, ZLIB

Scientific Programs at ALCF

Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program (60%)

Open to researchers from academia, government labs, and industry, the INCITE program is the major means by which the scientific community gains access to some of the fastest supercomputers. The program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address “grand challenges” in science and engineering.

Award Size: high millions of compute hours **Award Duration:** One-three year

ASCR Leadership Computing Challenge (ALCC) Program (30%)

Open to scientists from the research community in academia and industry, the ALCC program allocates up to 30 percent of the computational resources at the ALCF, NERSC and Oak Ridge. Projects in the program are special situations of interest to the DOE, with an emphasis on high-risk, high-payoff simulations in areas directly related to the Department’s energy mission, national emergencies, or for broadening the community of researchers capable of using leadership computing resources

Award Size: low to high millions of compute hours **Award Duration:** One year

Director's Discretionary (DD) Program (10%)

Start-up time through the DD program is available to researchers in academia, industry, and other research institutions who can demonstrate a need for ALCF's leadership-class resources. These awards are primarily a "first step" for projects working toward an INCITE or ALCC allocation. DD awards are available year 'round and are usually between three and six months in duration. Award sizes generally range from the low tens of thousands to the low millions of compute hours.

Early Science Program (for machine upgrade)

The goals of the Early Science Program (ESP) were to prepare key applications for the architecture and scale of next generation machine and to solidify libraries and infrastructure that would pave the way for other future production applications.

Application Support Teams

Catalysts

Mission: provides key projects with an assigned expert to maximize and accelerate their research. Features include: full project lifecycle assistance; support in conjunction with ALCF hardware and software resources; tailored services for unique requirements of a given research initiative

The Catalyst Team:

Provides a "jump-start" in the use of ALCF resources

Aligns the availability of ALCF resources with the needs of the project team

Establishes a spirit of collaboration to maximize the value that ALCF can bring to our project partners

Performance Engineering

Mission: ensure the effective use of applications on ALCF systems and emerging systems by assessing and improving the algorithms used by the applications and techniques used to implement those algorithms. This team's efforts are driven by the needs of the applications.

ALCF Performance Engineers have extensive expertise and are involved in:

Application Scaling and single core optimizations

Performance modeling and projections

Performance tools and libraries collaborations

Performance benchmarks development

Programming models and communication software

Parallel I/O optimizations

Computer architectures

Standard organizations in programming languages and programming models

Data Analytics and Visualization Team

Analysis/Visualization Resource



Tukey – analysis and visualization cluster

Compute nodes: 96

Architecture: AMD X86_64

Processors: Two 2GHZ AMD Operon 6128 per node
(8 cores per CPU, 16 cores total)

GPUs: Two NVIDIA Tesla M2070 GPUs per node

Memory/node: 64 GB RAM per node, 6GB GPU
RAM per node

QDR infiniband interconnect

450GB local scratch space

Shares the same software environment, network and
file system with Mira

Cobalt job scheduler

Default 8000 core hours per Mira user

Visualization and Analysis Software

Data Analytics and Visualization Team

The Data Analytics and Visualization team has expertise in tools and methods for high-performance post processing of large datasets, interactive data exploration, batch visualization, and production visualization.

Team members help users with their visualization and analysis needs using ALCF high-performance resources and a suite of tools maintained for these purposes, including:

- Production tools for high performance visualization (ParaView, VisIt, VMD)
- Analysis tools (R, Octave)
- Presentation graphics (ImageMagick, ffmpeg)

Team members also have expertise and assist users with additional tools which are not generally available on ALCF resources, including:

- Analysis tools (MATLAB)
- Presentation graphics (PowerPoint, Keynote, Adobe Premiere)

The ALCF Data Analytics and Visualization team has strong connections to Argonne's Mathematics and Computer Science research and development in the area of visualization and analysis.

ParaView, VisIt and VMD on Tukey

ParaView:

Client/server mode (recommended): running Paraview client on local resource, and the ParaView server (pvserver) on the Tukey visualization nodes. ParaView should first be installed on local resource

DISPLAY environment variable needs to be set to the local X server

Start the ParaView Server: submit an interactive job, starting from a shell on a login node

Start ParaView Client: launch the client on local resource; need configure some server settings in the client.

Server Configuration: connect → Add Server → Configure Server → Connect

VTK format is preferred.

VisIt:

Local resource: download and install VisIt; download the Tukey host profile for VisIt;

Tukey login node: edit .soft_tukey to add VisIt and guarantee version match.

Running VisIt in Interactive Mode: start up on local resource; connect to Tukey via Open File; enter Cryptocard response; When a selected file is open, it will launch a job on Tukey.

Running VisIt in Batch mode:

VMD:


Local resource: X11 is installed; ssh -X tukey.alcf.anl.gov

Running VMD in login node:

Running VMD in interactive mode(recommended):

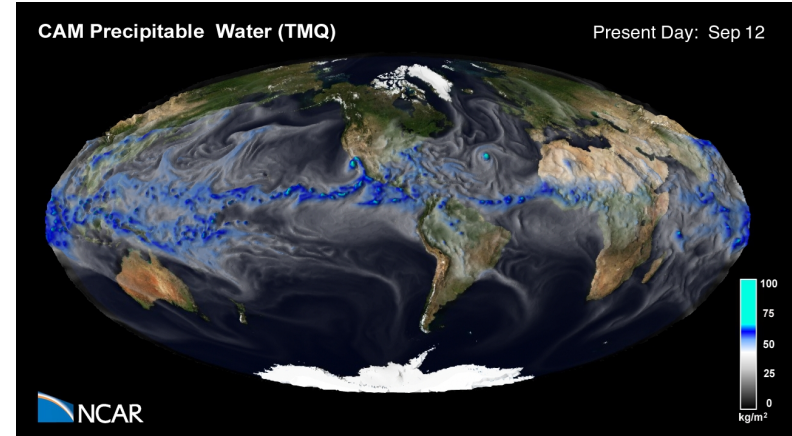
Visualizations of Extremely Large Datasets

ParaView




A Category 5 hurricane simulated by the CEM at 13km resolution
Accelerated Climate Modeling for Energy, 2015 INCITE project

ParaView




Present day simulation with a 1/4 degree horizontal resolution atmosphere
CESM Century-Scale Climate Experiments with a High-Resolution Atmosphere, 2015 INCITE project

ParaView



Autoignition of a turbulent hydrogen jet injected into hot oxygen
DNS/LES of Complex Turbulent Flows, 2015 INCITE project

VMD



Molecular details of the occlusion process in the ATP-driven calcium pump
Study of Large Conformational Changes in Biomolecular Machines, 2015 INCITE

World-Class Expertise and Project Lifecycle Assistance



Our people set us apart

To maximize your research, the ALCF has assembled a team of world-class experts in leadership computing to provide full project lifecycle assistance and support