The Argonne Leadership Computing Facility

An instrument of Change

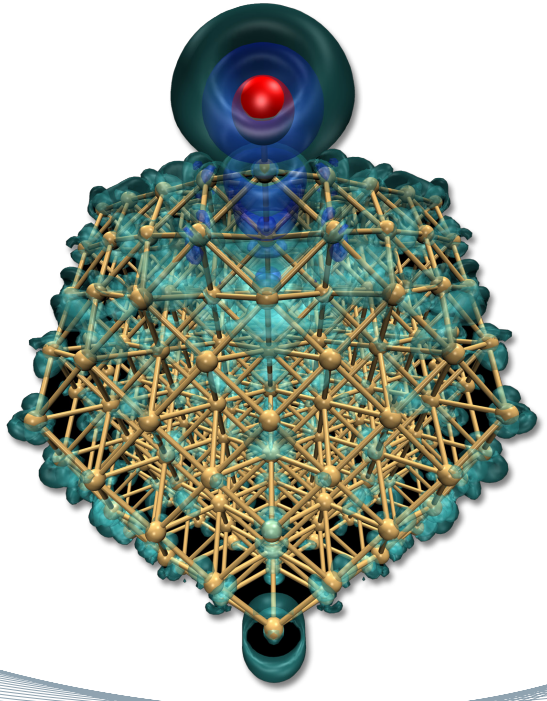

A Preeminent Global Resource

From Galileo's first telescope to the particle accelerators of today, historically, scientists have relied on instruments of exploration to understand the world around us. In this tradition, the Argonne Leadership Computing Facility (ALCF) was created and exists today as a preeminent global resource for breakthrough science and engineering research.

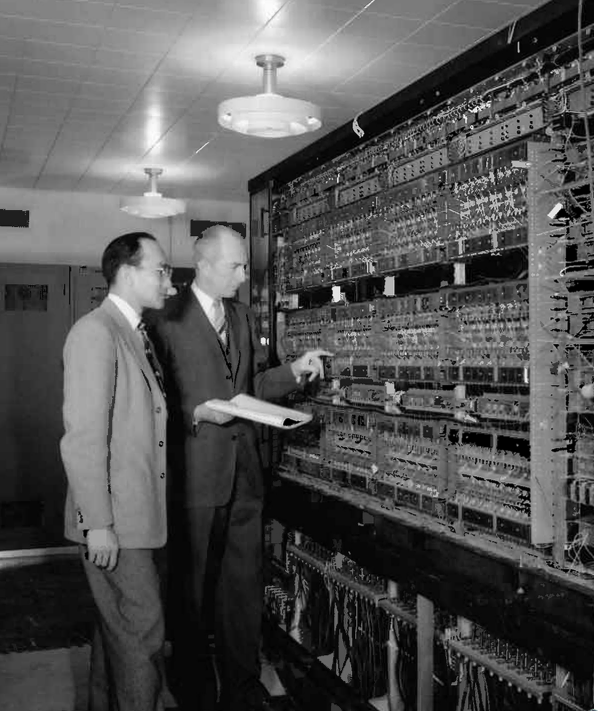

ALCF Continues a Tradition of Computing Innovation

In 1949, because the computers they needed weren't yet available commercially for purchase, Argonne physicists built their own. It was used to solve enormously complex mathematic problems. Thus began the Argonne tradition of innovation in providing leading computing resources to the scientific community—a tradition that continues today at the ALCF.

The seedbed for such groundbreaking software as MPI, PETSc, PVFS, GridFTP, and Cobalt, and as an essential partner in the development of some of the world's premier supercomputers, Argonne's historic stronghold at the forefront of scientific computing is undisputed.

The ALCF's team of system administrators, software developers, visualization specialists and computational scientists, collaborate with researchers to maximize and accelerate their discoveries.

Argonne NATIONAL LABORATORY

Argonne Leadership Computing Facility

ALCF wins EStar award

About the ALCF

High-performance computing is becoming increasingly important as more scientists and engineers use modeling and simulation to study complex chemical processes, exotic new materials, advanced energy networks, natural ecosystems and sophisticated energy technologies. One of the world's fastest computers for open science, Intrepid, Argonne's IBM Blue Gene/P, is capable of more than 500 trillion calculations per second. By 2012, Argonne will be home to an IBM Blue Gene/Q supercomputer capable of running programs at 10 quadrillion calculations per second. It will be able to do in one second what it would take every man, woman and child in the world to do if they performed a calculation every second for more than two weeks.

IBM Blue Gene Supercomputer recipient of the 2008 National Medal of Technology and Innovation

HPCwire Readers' Choice Awards 2009

Dr. Scott Parker leads a group of DOE Graduate Student Fellows on a tour of the ALCF's machine room. Here, Parker displays a tray of terabyte disk drives that make up part of the Intrepid, Argonne's IBM Blue Gene/P's massively parallel file system.

"Seeing the Blue Gene up close and interacting with our staff of experts provides an opportunity for researchers to understand how the Argonne Leadership Computing Facility can accelerate their research."

Accelerating Transitional Discovery

The ALCF is funded by the US Department of Energy's Office of Science and is one of Argonne's five user facilities. This year, the majority of the ALCF's nearly 1 billion compute hours will be used by Department of Energy-selected projects, including explorations into renewable energy, studies of the affects of global climate change, and efforts to unravel the origins of the universe.

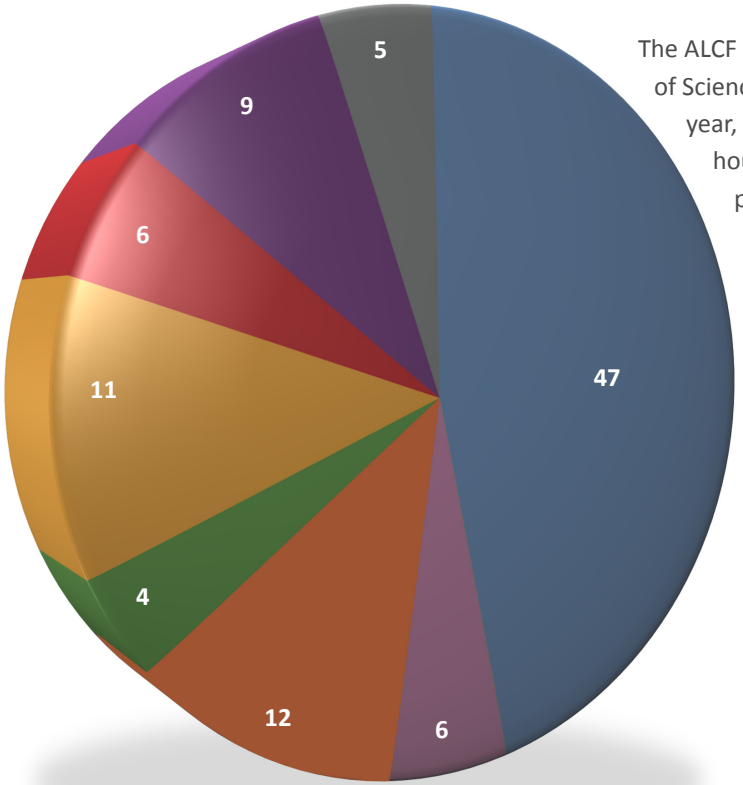

Percentage of compute hours used by scientific discipline
January–October 2010
Total hours used: 849 Million

- Physics
- Chemistry
- Biological Sciences
- Computer Science
- Earth Science
- Materials Science
- Engineering
- Energy Technologies

Access to the Resources of the ALCF Comes Through a Variety of Programs

60%	30%	10%	
Innovative and Novel Computational Impact on Theory and Experiment (INCITE)	**ASCR Leadership Computing Challenge Program (ALCC)**	**Early Science Program (ESP)**	**Discretionary Projects**
ALCF resources are available to researchers as part of the U.S. Department of Energy's INCITE program. Established in 2003, the program encompasses high-end computing resources at Argonne and other national laboratories. The INCITE program specifically seeks out computationally intensive, large-scale research projects with the potential to significantly advance key areas in science and engineering. The program encourages proposals from universities, other research institutions, and industry. It continues to expand, with current research applications in areas such as chemistry, combustion, astrophysics, genetics, materials science and turbulence.	Open to scientists from the research community in academia and industry, the ALCC program allocates resources to projects 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. Projects are awarded an ALCC allocation based on a peer review for scientific merit and computational readiness.	Allocations through the Early Science Program (ESP) provide researchers with preproduction hours (between system installation and full production) on the ALCF's next-generation, 10 petaflops IBM Blue Gene system. This early science period provides projects with a significant head start for adapting to the new machine and access to substantial computational time. During this shakedown period, users assist in identifying the root causes of any system instabilities, and work with ALCF staff to help develop solutions. More than four billion core hours are allocated through ESP.	Discretionary allocations are "start up" awards made to potential future INCITE projects. Projects must demonstrate a need for leadership-class resources. Awards may be made year round to industry, academia, laboratories and others, and are usually between three and six months in duration. The size of the award varies based on the application and its readiness/ability to scale; awards are generally from the low tens of thousands to the low millions of hours.

Catalyst Team

The Catalyst team provides key projects with an assigned expert, the "Catalyst," to maximize and accelerate research. Catalysts are computational scientists that are experts in their fields: computational chemistry, physics, fluid dynamics, astrophysics, etc., and actively participate in the research projects.

In addition to diagnosing problems using their knowledge of the scientific calculations being done, Catalysts also provide:

- Assistance with algorithm development and scaling efforts
- A single point of contact for key research initiatives
- Tailored services for unique project requirements
- Full project lifecycle assistance
- Value-added services and support in conjunction with ALCF resources

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.

- Production tools for high performance visualization (ParaView, VisIt)
- Analysis tools (R, MATLAB)
- Presentation graphics (PowerPoint, Keynote, Final Cut Pro)

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.

Operations Team

The ALCF Operations team consists of the Systems Group and the Advanced Integration Group.

The Systems Group is responsible for:

- Hardware maintenance
- Software maintenance
- Resolution of user tickets related to system issues
- Responding to requests for new software or versions of software
- Developing systems tools, particularly ones related to the unique system architectures and scale of ALCF resources

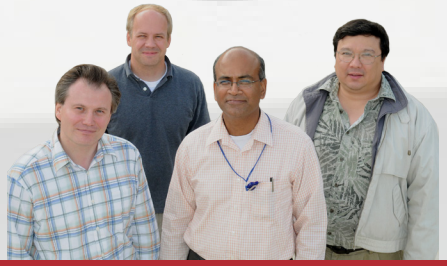

The Advanced Integration Group is responsible for:

- Ensuring the entire system software stack works together
- Assisting with I/O performance issues
- Bug fixes and feature requests for systems software

Performance Engineering Team

The mission of the Performance Engineering group is to help ALCF users achieve the best performance out of their applications. To this end, ALCF performance engineers work closely with the users in porting, tuning and parallelizing their applications on ALCF computers. They also assist in resolving performance-inhibiting I/O issues.

ALCF Performance Engineers have extensive experience in:

- Computer architectures
- Computational algorithms
- Porting, performance tuning and parallelizing of scientific applications and other software
- I/O

User Services and Outreach Team

The USO team provides frontline services and support to existing and potential ALCF users. The team also provides marketing and outreach to users, DOE and the broader community.

The USO provides the following services:

- Account/Project Administration
  - Account creation and administration
  - Project creation and allocation management
  - User access and foreign clearance processing
  - Setup of crytocard and ssh access
- Training and Education
  - Monthly INCITE/ALCC user call
  - Workshops and conferences
  - Website technical content
- Help Desk
  - Immediate answers to questions
  - Triage and troubleshoot problems
  - Escalation to appropriate resource
- User Outreach
  - Website content
  - Index of user publications
  - Facility tours
  - Public relations and media coordination
  - Displays, education materials

Leadership-Class Computing Resources

The Argonne Leadership Computing Facility systems are research machines open to Department of Energy INCITE projects and other select users.

Next-generation IBM Blue Gene

Preparation is underway for the installation of IBM's next-generation Blue Gene system. The 10 petaflops system will include more than 750,000 cores and more than 750,000 terabytes of memory.

Intrepid (IBM)

The ALCF state-of-the-art system gives the computational science community a computing capability dedicated to advancing knowledge towards solving humanity's most challenging scientific problems.

Blue Gene/P has a highly scalable torus network that can accommodate configurations with a petaflop of peak performance, as well as a high-performance collective network that minimizes the bottleneck common in simulations on large, parallel computers. Intrepid uses less power per teraflop than systems built around commodity microprocessors, resulting in greater energy efficiency and reduced operating costs. Blue Gene applications use common languages and standards-based MPI communications tools, so a wide range of science and engineering applications are straightforward to port, including those used by the computational science community for cutting-edge research in chemistry, combustion, astrophysics, genetics, materials science and turbulence.

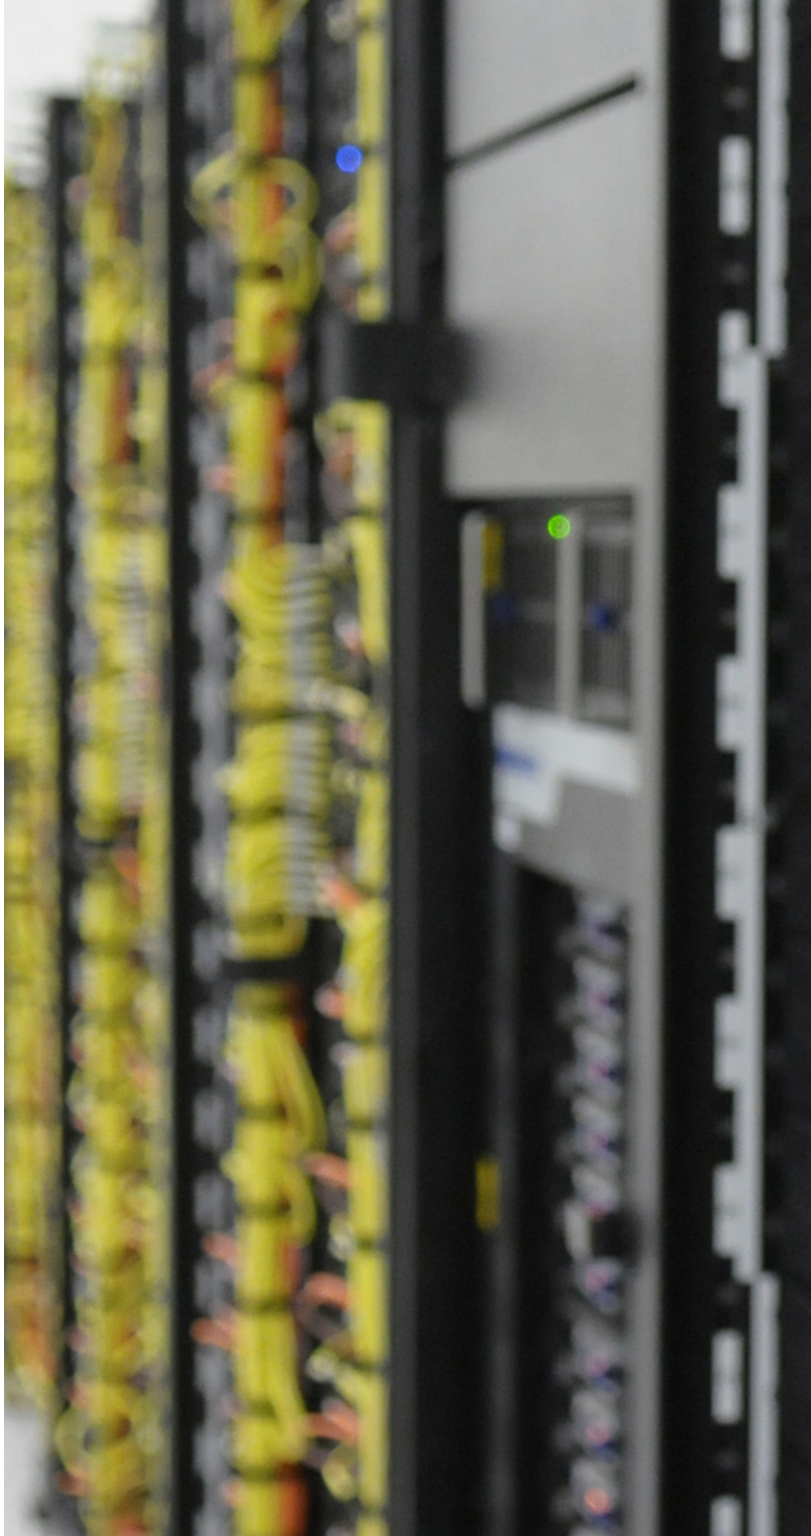

Intrepid (Blue Gene/P)—the ALCF's production machine for open science research

- 40,960 quad-core processors
- 163,840 cores
- 80 terabytes memory
- 557 teraflops

Surveyor (IBM Blue Gene/P)

ALCF's Surveyor is a Blue Gene system dedicated to tool and application porting, software testing and optimization, and systems software development. Surveyor has 1,024 quad-core nodes (4,096 processors) and two terabytes of memory. Peak performance is 13.9 teraflops.

Eureka (Visualizations and Data Analytics)

Most science applications that run on large-scale systems like Intrepid at the Argonne Leadership Computing Facility (ALCF) generate huge volumes of data that represent science results. Researchers analyze and explore this data output, and often convert the data to visual representations. To facilitate data analytics and visualization at the ALCF, researchers employ Eureka, one of the world's largest installations of NVIDIA Quadro Plex S4 external graphics processing units (GPUs). By using the NVIDIA visual computing system as the base graphics building block, Eureka enables breakthrough levels of productivity and capability in visualization and data analysis.

Eureka—the ALCF's visualization and data analytics solution

- 100 dual quad-core servers
- 200 Quadro FX5600 GPUs
- 111 teraflops
- more than 3.2 TB of RAM

Gadzooks (Visualization and Data Analytics)

Gadzooks is the ALCF's test and development system for visualization. It has four compute nodes, each with two 2.0 GHz quad-core Xeon servers with 32 GB RAM and eight NVIDIA Quadro FX5600 GPUs in two S4s.

Data Storage

The supercomputer's data systems consist of 640 I/O nodes that connect to 16 storage area networks (SANs) that control 7,680 disk drives with a total capacity of 7.6 petabytes of raw storage and a maximum aggregate transfer speed of 88 gigabytes per second. The ALCF uses two parallel file systems—PVFS and GPFS—to access the storage. The tape drives have built-in hardware compression allowing compression ratios of between 1.25:1 and 2:1, depending on the data, giving an effective capacity of 16-24 petabytes. An HPSS automated tape storage system provides archival storage with 16,000 tapes in two 10,000-slot libraries.

Networking

The Blue Gene/P uses five different networks for different communication operations. The 3D torus network is used for general-purpose, point-to-point message passing as well as for collective operations using irregular communication or large message sizes. Each node has six nearest neighbors. Each link provides a bandwidth of 425 MB/s per direction, for a total bidirectional bandwidth of 5.1 GB/s. Though each node has six bidirectional links on each node, there is only one shared DMA engine. The 3D torus network is also usable as a 3D mesh. The supercomputer connects to other research institutions using a total of 30 Gb/s of public network connectivity. Scientists can transfer datasets to and from other institutions over fast research networks such as the Energy Science Network (ESNet) and the Metropolitan Research and Education Network (MREN).

Theory and Computing Sciences Building

The Theory and Computing Sciences building, also known as TCS, provides the space and facilities to bring together over 600 researchers from across a wide range of computing and scientific disciplines.

Housing the ALCF and several other Argonne divisions, TCS was specifically designed to be an open and flexible workspace to encourage collaboration and the free flow of ideas. Construction of the more than 200,000 square-foot, seven-story-tall TCS building was completed in late summer 2009. It features a unique Zen garden, a vast library, as well as a multitude of open spaces and conference rooms.

Each year, hundreds tour the Supercomputing Support Facility at Argonne. Here, William Allcock, ALCF Director of Operations, explains the raw computing power behind the machines. Argonne's current supercomputer, Intrepid, is an IBM Blue Gene/P machine that is capable of more than 500 trillion calculations a second.

ALCF Users Speak Out

"If this resource were not available, there would have been no research. The project is such that it is impossible on other types of machines. This machine allows us to touch this problem for the first time, and prepare for the full-blown calculations on the new generation of machines."

Alexei Khokhlov
University of Chicago, Enrico Fermi Institute

"We wouldn't have been able to conduct a computational experiment of this size without our partnership with Argonne. Through this research, we experienced a dramatically exciting new approach to evaluate materials."

Tom Lange
Procter & Gamble

At Home at Argonne

With roughly 3,200 employees from nearly 60 nations (including about 1,000 scientists and engineers), Argonne National Laboratory hires top-quality staff who find personal and professional fulfillment in their work. The lab is located twenty-five miles southwest of Chicago—situated on 1,500 beautiful, wooded acres and surrounded by the Waterfall Glen Forest Preserve.

Argonne's surrounding communities offer housing and amenities to suit all tastes. The Chicago area boasts top-rated schools and universities and a superior transportation network of interstate highways, rail and bus lines. The laboratory's campus is about half an hour by freeway from two major international airports, O'Hare and Midway. Agronne's proximity to Chicago offers employees and their families access to world-class museums, an award-winning theatre scene, major music and dance companies, the world-renowned lakefront, boating and fishing, premier shopping and fashion, distinct and culturally diverse neighborhoods, more than 7000 restaurants and eateries, sports teams in all major professional leagues and more than 100 college teams.

ALCF Staff

Expert staff helps researchers maximize ALCF resources and achieve the best applications performance, accelerating key scientific discoveries and engineering breakthroughs for humanity.

On the cover
Electron density perturbation from carbon monoxide adsorption on a multi-hundred atom gold nanoparticle. The perturbation causes significant quantum size effects in CO catalysis on gold particles.

Jeffrey Greeley
Argonne National Laboratory

argonne

Leadership Computing Facility
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Argonne, IL 60439-4847

support@alcf.anl.gov

U.S. DEPARTMENT OF ENERGY

Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

CONTACT ▸ Argonne Leadership Computing Facility | www.alcf.anl.gov | (877) 737-8615

alcf-brochure-1110

Argonne NATIONAL LABORATORY
Argonne Leadership Computing Facility