ALCF Newsbytes

Argonne Leadership Computing Facility

Argonne National Laboratory

Volume 2, Issue 1 | Winter 2012

U.S. Department of Energy Awards ALCF Resources to Researchers Addressing Global Science Issues

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

Scientists at work on problems of global importance, including the quest for safe and renewable energy, curing diseases that ravage mankind, and understanding the affects of carbon emissions on the earth's climate, have a new tool in their arsenals thanks to recent awards of supercomputing time at the Argonne Leadership Computing Facility (ALCF)—one of the nation's premiere research facilities for world-transforming science.

In November, the Department of Energy (DOE) announced awards totaling 732 million compute hours on ALCF supercomputing resources. These hours will go to 31 projects through the DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. INCITE allocations are open to all scientists and are awarded based on a peer review to determine the project's potential real-world science impact and on its ability to fully utilize the extreme-scale compute resources of a leadership-class facility. Additional hours were made available to projects that will use DOE's Oak Ridge Leadership Computing Facility.

The ALCF provides world-class supercomputing resources and unsurpassed expertise in computational science to support breakthrough research in all scientific realms. This year's INCITE awardees include:

Biological Sciences

Andrew Binkowski, Argonne National Laboratory, received a 10 million hour award to conduct research that will allow a broader application of molecular simulation in therapeutic drug development.

George Karniadakis, Brown University, received a 50 million hour award to conduct multi-scale simulations for modeling blood flow in human brain blood vessels to better predict and understand cerebral aneurysms, sickle cell anemia, and malaria.

Chemistry

David Baker, University of Washington, was awarded 33 million hours to continue breakthrough research in protein structure calculation and design that may ultimately help scientists cure diseases and reduce carbon dioxide levels in the atmosphere.

continued ⏩

Table of Contents

Events of Interest

INCITE Proposal Writing Open Call and Webinars

The call for 2013 INCITE awards is April 11 through June 27.
Since its inception in 2004, the INCITE program has allocated nearly five billion processor hours on some of the world's fastest computers. Research teams from around the globe are eligible to compete for allocations that typically exceed 20 million hours. Click here to see a list of INCITE FAQs. For tips on writing your INCITE proposal, plan to participate in one of two interactive INCITE webinars, scheduled for March 26 and April 24. Watch your email from the ALCF for times.

Leap to Petascale Workshop Coming in May

This popular workshop is especially geared for users who are ready for large-scale runs. Plan to attend for hands-on assistance in tuning your applications. Tool and debugger developers will be on hand to lend their expertise. Registration information coming soon!

Argonne NATIONAL LABORATORY

U.S. Department of Energy Awards ALCF Resources to Researchers Addressing Global Science Issues (continued)

Alexei Khokhlov, The University of Chicago, received an award of 20 million hours to conduct simulations of the burning and detonation of hydrogen-oxygen mixtures for use in the design of safe hydrogen energy systems.

William Lester, Jr., ***UC Berkeley,*** received 4 million hours to model the potential use of ionic liquids for capturing carbon dioxide to further the development of clean energy systems.

Thierry Poinsot, European Center for Research and Advanced Training in Scientific Computation, received 10 million hours to continue simulations of two-phase combustion in gas turbines to facilitate the design of more fuel-efficient helicopter and airplane engines.

Donald Truhlar, University of Minnesota, was awarded 15 million hours to conduct large-scale electronic structure research of reactive systems relevant in the areas of nanotechnology, drug development and in the creation of clean, efficient fuels.

Computer Science

Ewing (Rusty) Lusk, Argonne National Laboratory, was awarded 5 million hours to continue work on enhancing the scalability and performance of major system software components to allow scientists to fully harness the power of high-end computing resources.

Ronald Minnich, Sandia National Laboratories, was awarded 10 million hours to build a software stack to address two of the major issues anticipated with exascale systems: fault management and adaptive runtime systems.

Patrick Worley, Oak Ridge National Laboratory, leads a team of researchers awarded 10 million hours to develop computer performance tools, middleware, and optimization strategies to maximize the usefulness of DOE leadership-class computing systems to the scientific community.

Earth Science

Thomas H. Jordan, Southern California Earthquake Center, leads an interdisciplinary team that was awarded 2 million hours at Argonne to construct the first physics-based seismic hazard model for California.

Roberto Paoli, CERFACS, was awarded 20 million hours for efforts to model global aviation affects on climate change, specifically the physics of the formation of contrail cirrus—ice clouds created by aircraft engine exhaust.

Warren Washington, National Center for Atmospheric Research, was awarded 30 million hours at Argonne for phase two of the Climate End Station—a coordinated Department of Energy (DOE) program for developing and deploying state-of-the-art coupled climate, chemistry and regional models in support of DOE's energy and carbon mission as well as for input into national and international energy policy considerations.

Energy Technologies

Mihai Anitescu, Argonne National Laboratory, was awarded 10 million hours to develop advanced optimization methods for the U.S. power grid under the uncertainty inherent in the high variability of renewable power sources as well as in transmission line contingencies that can unfold into cascading blackouts.

Paul Fischer, Argonne National Laboratory, was awarded 25 million hours to continue development of simulation capabilities for analyzing thermal transport in next-generation nuclear reactors capable of providing sustainable energy with a low carbon footprint.

Saivenkataraman Jayaraman, Sandia National Laboratories, was awarded 10 million hours to identify salt mixtures for use as heat-transfer fluids in solar-thermal energy facilities, making these power plants a viable, cost-effective energy source.

Umesh Paliath, General Electric Global Research, was awarded 45 million hours to use Large Eddy Simulation (LES) as an acoustic diagnostic and design tool to spur innovation in next-generation quieter aircraft engines and wind turbines, reducing cost and time-to-market while enhancing environmental friendliness.

Engineering

Said Elghobashi, University of California—Irvine, received 20 million hours to aid efforts to optimize global fuel usage in transportation by enhancing the understanding of the process of liquid fuel vaporization and mixing in turbulent flows.

continued

U.S. Department of Energy Awards ALCF Resources to Researchers Addressing Global Science Issues (continued)

James Glimm, Stony Brook University, received 35 million hour award to conduct chemistry combustion simulations within the engine of a scram jet with a new approach that have the potential to revolutionize the way we interpret Large-Eddy Simulations of turbulent combustion.

Kenneth Jansen, University of Colorado—Boulder, received 40 million hours to conduct simulations of active flow control on the vertical tail of commercial aircraft to facilitate next-generation, fuel-efficient designs.

Sanjiva Lele, Stanford University, received an allocation of 20 million hours to conduct direct numerical simulations to increase the scientific understanding of multi-material turbulent mixing in shock-accelerated flows—a fundamental phenomenon common to supernovae explosions, inertial confinement fusion and hypersonic propulsion systems.

Materials Science

Giulia Galli, University of California—Davis, received 25 million hours to conduct first-principles simulations of carbon-bearing systems to understand the properties of carbon in the deep Earth.

William George, National Institute of Standards and Technology, received 22 million hours to study the mechanisms that control the flow of large-particle and dense suspensions—a topic critical to industries such as cement and concrete, water treatment, and food processing.

Jeffrey Greeley, Argonne National Laboratory, received 10 million hours to perform calculations that will address fundamental, unresolved questions in nanonscience and provide insights into catalysts at the nanoscale—paving the way for groundbreaking, technological applications.

Priya Vashishta, University of Southern California, received a 45 million hour allocation to conduct simulations of the atomistic mechanisms of stress corrosion cracking of nickel-based alloys and silica glass—information essential for safe, reliable operation of nuclear reactors and nuclear-waste management.

Gregory Voth, The University of Chicago and Argonne National Laboratory, received 25 million hours to implement multiscale modeling methodologies to further research in chemistry and material science, and to conduct studies of charge transport to aid in the design of next-generation fuel cells and batteries.

Physics

Denise Hinkel, Lawrence Livermore National Laboratory, was awarded 63 million hours to conduct simulations of laser beam propagation in National Ignition Facility (NIF) targets, which will further scientific understanding of laser-plasma interactions and help attain the Department of Energy's goal of nuclear fusion at NIF.

Don Lamb, The University of Chicago, was awarded 40 million hours to continue a program of verification and validation of models of thermonuclear-powered supernovae to improve our understanding of these explosions and the properties of dark energy.

Paul Mackenzie, Fermi National Accelerator Laboratory, was awarded 50 million hours to continue research in lattice quantum chromodynamics and other strongly coupled field theories that will have a transformational impact on wide range of fundamental topics of importance in the study of high energy and nuclear physics.

Jean Carlos Perez, University of New Hampshire, was awarded 10 million hours to investigate Alfven waves and turbulence in the Sun's atmosphere, to help us understand the source of coronal heating and the origin of the solar wind, which remain critical unsolved problems in space physics.

James Vary, Iowa State University, and Steve Pieper, Argonne National Lab, were awarded 18 million hours to continue efforts towards developing a comprehensive description of all nuclei that will address several long-standing questions in nuclear theory relevant to applications in basic science, nuclear astrophysics and to the future of the nation's energy and security needs.

This year's allocations represent the largest amount of supercomputing time ever awarded under the INCITE program, reflecting both the growing sophistication of the field of computer modeling and simulation and the rapid expansion of supercomputing capabilities at DOE National Laboratories in recent years.

For more information about INCITE at the ALCF, visit: http://www.alcf.anl.gov/projects

The ALCF is Getting Q'd Up

First racks of Blue Gene/Q have arrived

Engineers are hard at work building the management infrastructure that will support Mira, the ALCF's 10 petaflops Blue Gene/Q system—and they don't have a moment to spare: The first two racks of the new system arrived at the datacenter on January 4.

These early racks were named for an asteroid and a constellation, respectively: Vesta, will be used for science application development and debugging, and Cetus will be used for system software development and testing.

Once Vesta and Cetus are installed and pass a series of clearance criteria, the first researchers to access them will be those selected in 2009 through the Early Science Program. In addition, key developers of essential computational tools for users will be given access to the system to complete the process of porting the tools for use on the new Blue Gene/Q architecture.

Said ALCF's Deputy Director Susan Coghlan, "With a system this innovative and powerful, we have a slew of researchers and collaborators eagerly vying for time on the new Q, even though full production could be a year away. We're very pleased with the installation progress to date and hope that it continues to run ahead of our scheduled milestones."

More about the Q

The third in IBM's family of Blue Gene supercomputers, the Blue Gene/Q operates at more than an order of magnitude faster than its predecessor, with a single one of Q's 16,384-processor racks more powerful than 16 racks of Blue Gene/P. Each of the Q's processors features 16 compute cores (up from four on the Blue Gene/P), plus a core allocated to operating system administrative functions.

With its small footprint and ultra-low power consumption (two to three times the computations-per watt of other machines its size), IBM Blue Gene/Q prototype systems have seized the first five spots on the Green500 list (November 11) which ranks the world's most energy-efficient supercomputers.

The first racks of Mira, the ALCF's 10-petaflop Blue Gene/Q system, have arrived at Argonne.

continued ⏩

The ALCF is Getting Q'd Up (continued)

Familiarity Breeds Ease of Use

The Blue Gene/Q offers users an open source and standards-based programming environment:

- Red Hat Linux distribution on service, front-end and I/O nodes
- Familiar Linux execution environment with support for most POSIX system calls
- Familiar programming models: MIP, OpenMPI, POSIX I/O
- Automatic SIMD (Single-Instruction Multiple-Data) FPU exploitation enabled by IBM XL (Fortran, C, C++) compilers

Applications Running on Blue Gene

Following is a partial list of applications developed on Blue Gene systems; many of which have ported successfully to Blue Gene/Q.

- CFD Alya System
- CFD (Flame) AVBP
- CFD dns3D
- CFD OpenFOAM
- CFD NEK5000, NEKTAR
- CFD OVERFLOW
- CFD Saturne
- CFD LBM
- MD Amber
- MD Dalton
- MD ddcMD
- MD LAMMPS
- MD MP2C
- MD NAMD
- MD Rosetta
- DFT GPAW
- DFT iGryd
- DFT KKRnano
- DFT ls3df
- DFT PARATEC
- DFT CPMD
- DFT QBOX
- DFT VASP
- Q Chem GAMESS
- Nucler Physics GFMC
- Neturonics SWEEP3D
- QCD CPS
- QCD MILC
- Plasma GTC
- Plasm GYRO (Tokamak)
- KAUST Stencil Code Ge
- BM: sppm, raptor, AMG, IRS, sphot
- BM: SPEC2006, SPEC openmp
- BM: NAS Parallel Benchmarks
- BM: RZG (AIMS, Gadget, GENE, GROMACS, NEMORB, Octopus, Vertex)
- Coulomb Solver – PEPC
- MPI PALLAS
- Mesh AMR
- PETSC
- MpiBlast-pio Biology
- RTM – Seismic Imaging
- Supernova Ia FLASH
- Ocean HYCOM
- Ocean POP
- Weather/Climate CAM
- Weather/Climate Held-Suarez Test
- Climate HOMME
- Weather/Climate WRF, CM1

Understanding Oil Plume Behavior in a Stratified and Rotating Ocean

A key challenge for scientists is understanding the flow patterns, mixing behavior, surface presence, and distribution of buoyant plumes in a stratified and rotating environment such as the ocean. Knowledge of this behavior is directly applicable to understanding the distribution within the Gulf of Mexico of the oil spilled by the Deepwater Horizon well. Studying the behavior of buoyant plumes, like those of the spill, could correlate the observations of oil on the ocean surface with the amount of oil that actually spilled.

Tamay Ozgokmen from the University of Miami, in collaboration with Argonne National Laboratory researchers, is doing the first-ever simulations of the buoyant plumes in a stratified and rotating environment. This project uses Nek5000, which the ALCF has helped to tune for Blue Gene/P. Nek5000 has demonstrated high performance (19% of peak at scale) and excellent computational efficiency. With it's high-order spectral element formulation, it requires a tenth of the computational resources (i.e., gridpoints and CPU time) as codes based on traditional discrectizations.

Initial results showed that it can be difficult for the plume to cross stratification barriers caused by pressure and temperature variations. Additional results showed behavior of the plume after the spill was shut off. Further studies will quantitatively explore the dispersion behavior and rotational behavior of the plume. The research will improve the ability of evaluating oil release at underwater drilling sites and perhaps mitigate consequences at the surface.

For more information, contact Tamay Ozgokmen, University of Miami (tozgokmen@rsmas.miami.edu).

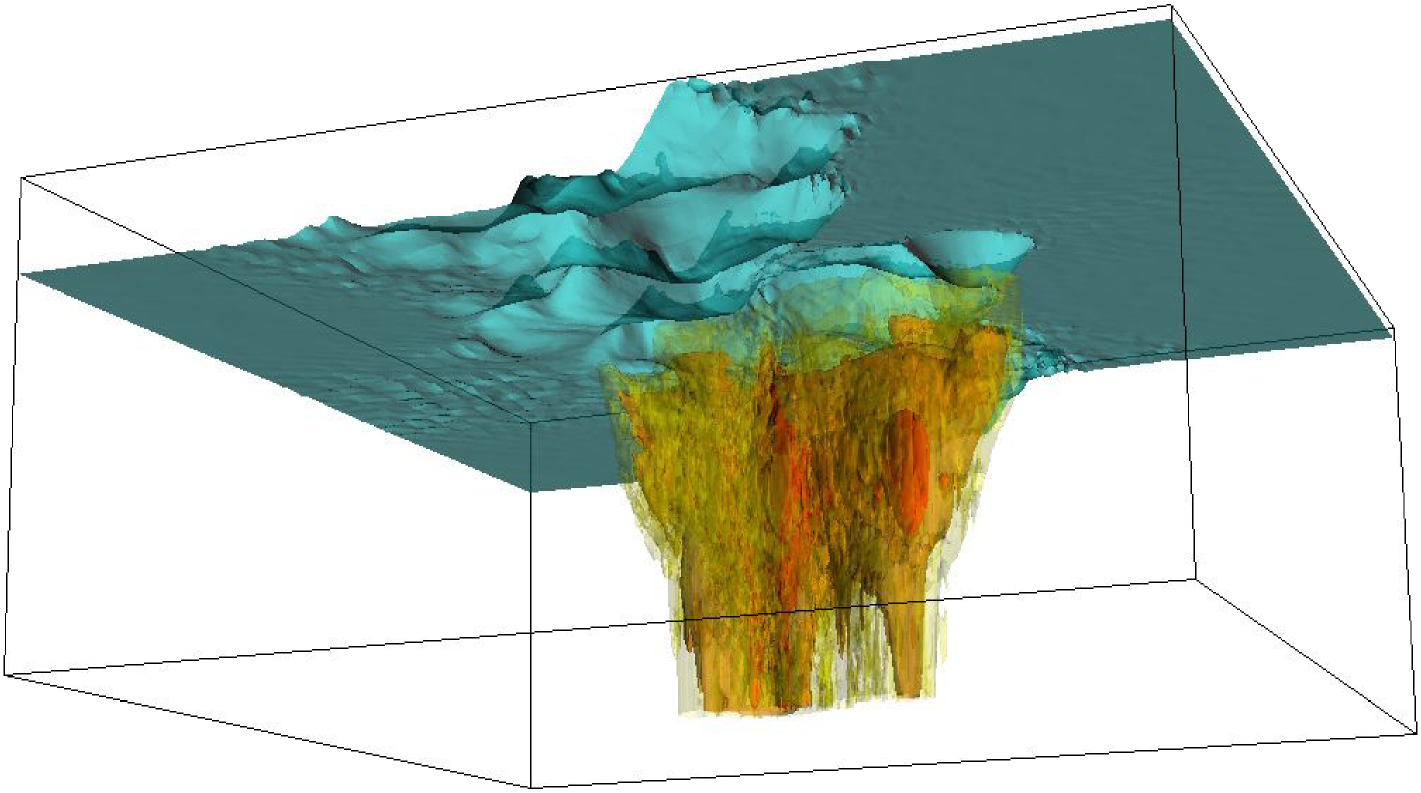

Side view of a buoyant plume interacting with ocean flows near the surface. Note the vertical coherence of the plume and inability to cross the base of the oceanic surface mixed layer. Image credit: Tamay Ozgokmen, University of Miami; Paul Fischer and Aleks Obabko, Argonne National Laboratory; and Hank Childs, Lawrence Livermore National Laboratory.

ExM Advances System Support for Extreme-Scale, Many-Task Applications

Exascale computers will enable and demand new problem-solving methods that involve many concurrent, interacting tasks, such as parameter sweeps, ensembles, and complex workflows. Running such "many-task" applications efficiently, reliably, and easily on extreme-scale computers is challenging. System software designed for today's mainstream, single-program multiple data (SPMD) computations does not typically scale to the demands of many-task applications.

At the ALCF, researchers from Argonne, the University of Chicago, and the University of British Columbia produce the advances in computer science and usable middleware needed to enable the efficient and reliable use of exascale computers for important existing applications and to facilitate the broader use of large-scale parallel computing by new application communities. The research is supported by the U.S. Department of Energy's (DOE) Office of Advanced Scientific Computing Research (ASCR) under the X-Stack program.

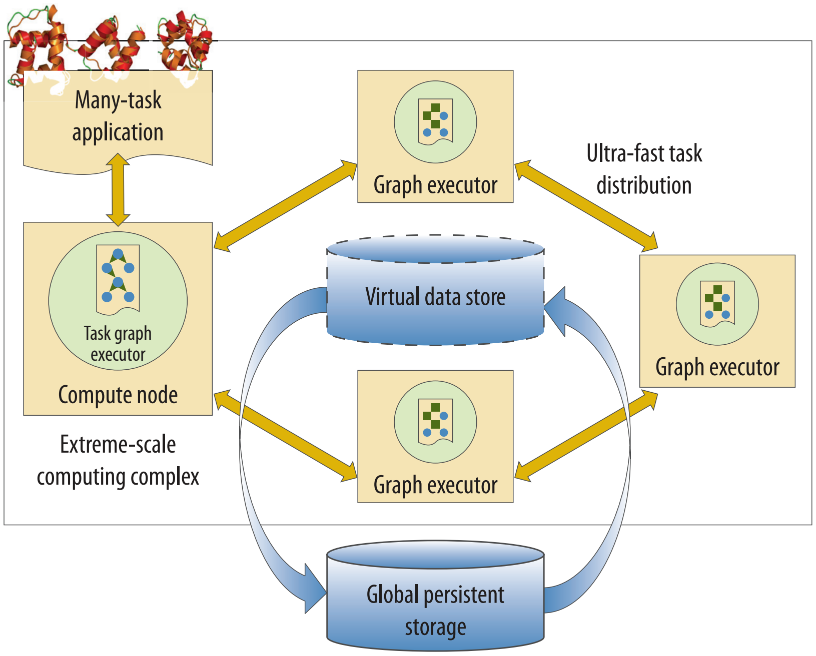

The ExM parallel data-flow program evaluator and virtual distributed data store are creating the task management and data exchange scalability needed to support exascale applications. Image credit: ExM project team.

To address these challenges, the ExM project is designing, developing, applying, and evaluating two new system software components: the ExM data store and the ExM parallel task evaluator. The data store utilizes both memory-based filesystems and distributed hash tables to facilitate the data exchange between tasks. The parallel task evaluator breaks down scalability barriers by permitting a significantly higher rate of task creation and management operations to support Swift's functional data-flow-based programming model. While the current Swift engine can launch approximately 700 "leaf" application tasks/second, it is limited to evaluating scripts on a single front-end host. Under the ExM model, script evaluation distributes itself among a theoretically unlimited number of compute nodes. In tests to date, the ExM "Turbine" prototype has achieved over 500,000 tasks/second on the relatively slow cores of a SiCortex 5832 system while operating on a complex distributed data structure. These early SiCortex results suggest that the ExM approach will scale well on DOE's Blue Gene and Cray XE leadership-class systems.The ExM/Swift implicitly parallel programming model is well suited for the upper-level logic of many prospective exascale applications ranging from climate model analysis to molecular biology to uncertainty quantification and extreme-scale ensemble studies.

Key accomplishments to date include the following: The Jets prototype has extended the many-task model to a hybrid many-parallel (MPI) task model. The Turbine prototype, using Argonne's Asynchronous Dynamic Load Balancing (ADLB) library, has shown encouraging scalability and has confirmed that exascale performance goals are reachable. MosaStore on the Blue Gene/P (BG/P) and other clusters is creating a model of a virtual data store. The AME "anyscale many-task engine" has explored integration and optimization of task dispatching and data movement, and has scaled on the ALCF Blue Gene/P to 16K cores. Researchers are already evaluating ExM tools on three science applications—earthquake simulation, astronomical image processing, and protein/RNA interaction. AME achieved a 73% reduction in traffic to the BG/P global filesystem for the Montage astronomy application. ALCF researchers are working with the ExM and Swift teams to execute molecular dynamics applications on the ExM Jets prototype, to create an interactive web-based interface for developing applications with the Swift programming mode, and supporting ExM execution on the Intrepid BG/P system using the ZeptoOS compute node kernel with Zoid IO forwarding.

For more details, visit the ExM project website (http://www.mcs.anl.gov/exm).

Spotlight on ALCF Staff: Robert Scott

This month, Robert Scott, technical support analyst at the ALCF gets caught in the Newsbytes spotlight.

NEWSBYTES: You're a part of the User Services and Outreach (USO) division here at the ALCF. What is the USO team's purpose?

ROBERT SCOTT: The USO team is often the first point of contact for new users, so from the very start, we help to shape their experience at the ALCF. We manage the account process, provide frontline user support through our Help Desk and facilitate ongoing user training through workshops. In addition, the USO shares the exciting breakthrough science underway at our facility through communications efforts, including the ALCF website.

NB: What's your role in the USO as a technical support analyst?

RS: My role is shared amongst three individuals on the Help Desk team. We answer calls and emails from users in need of assistance with accounts, projects, crypto card and ssh key access. We have the full ALCF team backing us for help with any specialized questions or issues that come in from our users.

NB: How has your work here been different than previous places you've worked?

RS: The atmosphere here is so engaging and positive. A casual conversation with a colleague can quickly turn into a moment where I learn something new. There's not a day that goes by where I don't pick up a useful nugget of knowledge. I also like the fact that management is approachable and helpful—they're always there to provide appropriate guidance.

NB: Tell us about an interesting project you're working on.

RS: I'm working on the new ALCF website which was built using Drupal technology. I attended a Drupal "camp" in Chicago, so it's been great to put my newly acquired skills to use on our new website.

NB: What do you like best about working at the ALCF?

RS: Outside of working with some of the most powerful, best-designed computers on the planet, we also get to collaborate with world-renowned scientists who are at the top of their fields.

For more information about the User Services and Outreach division at the ALCF, visit: http://www.alcf.anl.gov/user-services or email Robert Scott at scott@alcf.anl.gov.

Blue Gene/Q Dominates at SC11

Another SC is on the books for the ALCF. Each year, SC brings together the international community in high performance computing, networking, storage and analysis. The conference attracts a broad range of HPC stakeholders from across the scientific and technical computing community. Held November 12-18 in Seattle, this year's event set records in both attendance (11,455) and in the number of industry and research exhibits (336).

ALCF staff looks forward to SC as an annual venue for reconnecting with colleagues and making new connections. "SC offers a unique opportunity for researchers by bringing together the global HPC community where we can learn from each other and collaborate on ways to improve existing technology as well as develop visions for the future," said ALCF Division Director Mike Papka. "Much of what you see at SC today can be traced back to ideas developed over the years in partnerships forged at previous conferences. It's a perfect way for government, academia and industry to come together and accelerate discovery."

The Argonne SC11 booth was heavily trafficked, drawing attendees to a full schedule of talks and demonstrations given by leading researchers at work on ALCF systems, and by industry experts in HPC technology and applications.

Among the many highlights of the event for the ALCF was the unveiling of the IBM Blue Gene/Q—the system the ALCF will receive in 2012. With its unprecedented power efficiency and ultra capability, the Blue Gene/Q made a clean sweep of the major supercomputing rankings announced at SC11. The trifecta includes:

- Top500 – IBM Blue Gene/Q was named the most energy-efficient system with a rating of 2,029 megaflops per watt.
- Green500 – IBM Blue Gene/Q captured the top five positions on the Green500 list for superior energy-efficiency.
- Graph 500 – IBM Blue Gene/Q systems were ranked #1 and #7, demonstrating its superior performance in data-intensive computing.

Join the ALCF at SC12 in Salt Lake City, Utah.
For advance conference information, visit: http://sc12.supercomputing.org/.