FY 2006 Accomplishments

- ***NERSC to increase peak capacity by factor of 100.*** NERSC, through the competitive procurement process for NERSC-5, evaluated a number of vendors, using the NERSC Sustained System Performance (SSP) metric. The SSP metric was developed by NERSC to better gauge how well a proposed system will meet the needs of the Center's 2,500 users and measures sustained performance on a set of codes selected to accurately represent the Center's portfolio of challenging scientific applications. Based on the vendor evaluations, NERSC selected a 100 teraflop Cray Hood system. The system will consist of over 19,000 AMD Opteron 2.6-gigahertz processor cores, with two cores making up one node. Each node has 4 gigabytes of memory resulting in aggregate memory capacity of 39 terabytes. The NERSC upgrade will alleviate much of the current backlog of meritorious requests for high performance production computing resources.
- ***Leadership Computing Facility doubles XT3 capability to 54 teraflops.*** The LCF at ORNL upgraded the Cray XT3 supercomputer to increase the system's computing power to 54 teraflops. The LCF system is currently the largest high performance computing system in SC and will become a major computing resource for DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. The upgrade involved replacing all 5,212 single core processors with Cray's dual-core processors, doubling the memory and increasing the bisection bandwidth resulting in a supercomputer that consists of more than 10,400 processing cores and 21 terabytes of memory. The upgrade will allow the LCF to provide additional resources for computationally intensive, large- scale projects with the potential of high scientific impact such as the design of innovative nanomaterials, predictive simulations of fusion devices and the understanding of microbial molecular and cellular systems.
- ***ESnet moves forward with new architecture to deliver cost effective, high reliability, high performance networks for science.*** Energy Sciences Network (ESnet) is implementing a new architecture that is a double core ring with interconnected metropolitan area networks (MANs). The new architecture is designed to meet the increasing demand for network bandwidth, high reliability, and advanced network services as next-generation scientific instruments and supercomputers come on line. In FY 2006, the second and third MANs were completed in the Chicago and New York-Long Island Areas. They provide dual connectivity at 20 gigabits per second—which is 10 to 50 times the previous site bandwidths, depending on the site using the ring—while reducing the overall cost. These MANs connect ANL, Fermi National Accelerator Laboratory (FNAL), and BNL to the ESnet backbone to address critical requirements such as Large Hadron Collider (LHC) data. By increasing bandwidth to these sites, DOE advances research in areas such as climate change, genetics, renewable energy, nanotechnology, national security, and basic science in physics and chemistry through support for the large-scale science and large-scale SC collaborations nationwide.

- ***Terapaths Project enables experiments to use global networks effectively.*** The Terapaths project is enabling data-on-demand capability from scientific apparatus such as the Relativistic Heavy Ion Collider (RHIC) and data repositories such as CERN to computing centers used for data analysis. The Terapaths project is enabling data transfers with guarantees of speed and reliability that are crucial to applications with deadlines, expectations, and critical decision-making requirements. BNL needs to do RHIC production data transfers and LHC Monte Carlo challenges between BNL and the remote collaborators. The aggregate of their peak network requirements is beyond BNL capacity. The Terapaths project can modulate LHC data transfers to opportunistically utilize available bandwidth, ensuring that the RHIC production data transfer is not impacted.
- ***Lambda Station enables High Energy Physics researchers to prepare for LHC.*** Lambda Station, a high-capacity software-based network technology device, which combines advanced optical networks technologies with ultra high-speed data-intensive applications, has enabled researchers at FNAL, bracing for an onslaught of data from LHC experiments, to develop ultra fast connections for moving US-CMS (the U.S. part of the Compact Muon Solenoid detector) data between LHC Tier1 and Tier2 centers. Initial test results using Lambda Station technology enabled FNAL, the US-CMS Tier1 center to transfer data at line rate (10 gigabits per second) to Caltech, a US-CMS Tier2 center. This transfer rate represents a 15 fold increase in file system to file system transfer speed.
- ***The ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARS) enables global hybrid optical networks for science.*** OSCARS has developed and deployed a prototype service that enables on-demand provisioning of guaranteed bandwidth and secure virtual circuits within the ESnet production network. With the advent of service sensitive applications (such as remote controlled experiments, time constrained massive data transfers, distributed (Grid based) analysis of massive scientific data sets, etc.), it is essential to augment the services present in today's ESnet infrastructure. The default characteristic of the Internet today does not provide a user any service guarantees. There is neither the assurance that a packet will be delivered to its destination, nor any transport predictability (such as latency and jitter) when the packet is in transit. The concepts of dynamic provisioning and on-demand bandwidth assignment will provide predictable high bandwidth and/or low latency to high-impact science applications in a way that lets the science collaborations manage bandwidth as a predictable service as is currently done with high performance computing resources. In the past year, ESnet has successfully demonstrated that an end-user can provision circuits within ESnet (through a simple web-interface), and effectively obtain the required bandwidth. Further, most of SC's large-scale science is collaborative with the U.S. and international research and education community. In order to ensure that the new network services interoperate among DOE laboratories, universities, and other research institutions, the new services are being developed in close cooperation with Internet 2/Abilene, the U.S. research and education community, the DANTE (Delivering Advanced Network Technology to Europe) GÉANT network project, and the European research and education community.

- ***ASCR Distributed Network Environment software enables LHC researchers.*** One of the discoveries eagerly anticipated by particle physicists working on the world's next particle collider is that of supersymmetry, a predicted lost symmetry of nature. Physicists from the ATLAS experiment have used Open Science Grid (OSG) resources to show that there is a good possibility of discovering this long-sought symmetry with data collected during the first few months of LHC operation. OSG middleware and resources were combined with those from the LHC Computing Grid and the University of Wisconsin-Madison to simulate a vast number of particles interacting with the complex ATLAS detector. Comparing several of these simulations—some created accurately for the first time using the combined grid infrastructure—demonstrated that supersymmetry may be within reach. The OSG and the LHC Computing Grid are built on software developed by ASCR's distributed networking environment activity.
- ***Performance Tools, Modeling and Optimization*:** Computer science researchers significantly enhanced the performance of several SciDAC applications, including, doubling the performance of the materials quantum Monte Carlo code on the IBM Scalable Powerparallel supercomputer, and improvements to the scalaBLAST gene analysis software that enable it to address problems of a complexity impossible with earlier tools.
- ***Chemistry software enables computations 40 times more complex and 100 times faster than previous state of the art.*** The Full-configuration interaction (FCI) method developed by researchers at ORNL exactly solves the N-electron problem in a given 1-electron basis. This solution provides a critical benchmark for understanding and calibrating many-body calculations. The researchers have developed several new parallel-vector algorithms, each optimal for certain parameter values. The new code already enables computations over 40 times more complex and 100 times faster than previous work and is achieving 62% of peak speed on the ORNL Cray X1.
- ***Significant upgrade to DAKOTA, a state of the art optimization and uncertainty quantification package.*** Sandia National Laboratories have released Version 4.0 of the DAKOTA Optimization and Uncertainty Quantification Toolkit. DAKOTA is used throughout the DOE laboratories for design, surety, and validation studies. DAKOTA performs a multitude of iterative system analyses: it closes the design and analysis loop by using simulation codes as virtual prototypes to explore design scenarios and the effect of uncertainties. DAKOTA 4.0 provides a significant step forward in usability with a JAVA-based graphical user interface, support for algebraically-defined problems with AMPL, a library mode for seamless simulation integration, and upgrades to documentation and configuration management. DAKOTA is both a production code and a framework for innovative research. DAKOTA 4.0 deploys many new algorithms, such as the latest in probabilistic design capabilities, which are being applied to enable microsystem designs to perform well regardless of manufacturing variations. DAKOTA is open-source and has 3,000 registered installations from sites all over the world.
- ***Multiscale Analysis Reveals New Insights into Pattern Formation in Biological Systems.*** How do random mixtures of molecular components organize themselves into large-scale cellular structures? To answer this question, ANL researchers are

developing a multiscale mathematical framework to simulate the process of self-organization in biological systems. The focus is on microtubules, the building blocks of cellular structure. A major accomplishment has been the discovery of a previously unidentified cluster state. In addition to providing insights into biological systems, this work has potential application to the formation of transient patterns in populations of mobile bacteria.

- ***Simulations on Massively Parallel (32,000 Processor) Architectures Provide New Insights into Galaxies.*** ANL applied mathematicians have developed Nek5000, an advanced scientific simulation tool, that features high-order numerical discretizations and multigrid solvers capable of scaling to thousands of processors. Such extreme scalability is needed to simulate magnetohydrodynamics in complex domains. A major advance this year was the port of Nek5000 to the 32,000-processor Blue Gene platform at IBM Watson, where the researchers were able to identify an unexpected linear angular vector profile during their study of angular momentum transport in galaxies.
- ***Spectral Element Techniques Remove the Memory Bottleneck in Accelerator Simulations.*** The performance of accelerators is often limited by the Wakefield effect, which depends on the intensity and distribution of the electron bunch. Traditional finite-difference time-domain methods require about 10^{12} grids and 80 terabytes of memory for simulating the accelerator chamber. ANL computer scientists have developed a spectral element discontinuous Galerkin code that now makes the Wakefield problem completely tractable. Only 6×10^9 grids are needed to resolve a 500 gigahertz wavelength, with a memory requirement of only 400 gigabytes that is feasible on today's high performance computers.
- ***Toolkit Poised to Play Major Role in the Global Nuclear Energy Partnership (GNEP).*** Solvers from ANL's PETSc (Portable, Extensible Toolkit for Scientific Computing) system have been integrated into a neutronics code, to enable much larger simulations than previously possible. The neutronics code will be coupled with other codes by using a framework being developed for GNEP. This FY 2006 accomplishment will enable accurate, fully coupled, high-resolution simulations of fission reactors for the first time.
- ***Scientific Data Management.*** Ultrascale computing and high-throughput experimental technologies have enabled the production of scientific data about complex natural phenomena. However, answers to fundamental questions about the nature of those phenomena remain largely hidden in the massive quantities of produced data. One goal of the Scientific Data Management Center (ORNL; North Carolina State University; LBNL) is to provide a scalable high performance statistical data analysis framework to help scientists perform interactive analyses of these raw data to extract knowledge. They have developed an open source parallel statistical analysis package, called Parallel R, that lets scientists employ a wide range of statistical analysis routines on high performance shared and distributed memory architectures without having to deal with the intricacies of parallelizing these routines. Benefits have been demonstrated in Climate models and Computational Biology efforts.

- ***Simulating a Flame in Three Dimensions***. Researchers at LBNL have developed the most realistic direct numerical simulation combustion simulations to date. These laboratory-scale simulations of turbulent premixed methane combustion used 20 chemical species, 84 reactions, and no approximate models for fine scale (subgrid) turbulence or turbulence chemistry interactions. For the first time researchers have been able to compare full-field images of entire laboratory-scale turbulent flames with views of the same produced by computer simulations uncompromised by approximate fine scale turbulence models. The simulations captured with remarkable fidelity some major features of the experimental data, such as flame-generated outward deflection in the unburned gases, inward flow convergence, and a centerline flow acceleration in the burned gases. The simulation results were found to match the experimental results within a few percent. This agreement directly validated both the computational method and the chemical model of hydrocarbon reaction and transport kinetics in a turbulent flame. This advance has the potential to greatly increase our understanding of how fuels behave in the complicated environments inside turbulent flames. The work was featured on the July 19, 2005, cover of the Proceedings of the National Academy of Sciences.
- ***Supernova Combustion in the Distributed Burning Regime.*** Astrophysicists modeling entire stars need to know what is going on at scales their models can't resolve. For example, astrophysicists had predicted that in the late stage of a supernova explosion, the flame front burns further out from the center and the lower density of the star causes it to burn less vigorously and become more unstable due to increased turbulence and mixing of fuel and ash. This represents a different mode of combustion known as a "distributed burning regime." Using the Adaptive Mesh Refinement (AMR) combustion codes and running simulations for 300,000 processor hours at NERSC, the Supernova Science Center team, supported through SciDAC, was able to create the first-ever three-dimensional simulations of such an event.
- ***Scalable Linear Solvers.*** At the core of many DOE simulation codes is the need to solve huge linear systems on thousands of processors. Multigrid methods are called scalable or optimal methods because they can solve a linear system that scales linearly with problem size. This property makes it possible to solve ever larger problems on proportionally larger parallel machines in constant time. LLNL researchers are working on algebraic multigrid (AMG) approaches that are well suited for addressing a variety of problems that are challenging or impossible to solve using the classical multigrid approach. In 2006, a scalable solution of the definite Maxwell's equations was developed. The major difficulty with developing scalable solvers for these equations is the oscillatory, huge near-null space and the fact that such problems are often solved on unstructured grids. Previous attempts to construct AMG methods have had only partial success. This new auxiliary-space Maxwell solver (AMS) is an improved version of the Hiptmair/Xu solver and is the first provably scalable solver for the definite Maxwell's equations on quasi-uniform unstructured meshes that requires minimal additional information from the user.
- ***SciDAC computer management software selected by NNSA as standard.*** The Moab suite of cluster management software has been selected by the NNSA Advanced Simulation and Computing Program as their standard cluster workload and resource

management tool suite. The development of Moab was supported by the SciDAC Scalable Systems Software computer science project, and it is now also available and commercially supported by Cluster Resources, Inc. The Moab provides features necessary to run jobs on 100,000 processor systems.

- ***Unified Parallel C implements one-sided messages to improve scientific software performance on high performance computers.*** Partitioned Global Address space languages offer an alternative to message passing programming on high end systems. LBNL and University of California groups recently demonstrated that the languages can effectively leverage modern network hardware to provide performance that is faster than the two-sided send/receive models for some machines and computations. One example is the communication-intensive 3D Fast Fourier Transform (FFT), important to many scientific simulations. The compilers for these languages are highly portable and support interoperability with the Message Passing Interface, Fortran, and C/C++ languages.

- ***File system addresses critical bottleneck in high performance computer input/output.*** A critical problem in high performance computing is the speed at which data can be moved off the computer. In some cases the time required to move data off the computer is longer than the time the computer spends doing the calculation. Current commercial file systems may not be able to scale to petascale systems. The Argonne National Laboratory Parallel Virtual File System (PVFS) project for providing input/output (I/O) support to computational science has developed and deployed a new generation system, Parallel Virtual File System – version 2 (PVFS2). PVFS2 provides an innovative approach to parallel file systems by cooperatively working with user-level I/O software layers to achieve high performance. PVFS2 has been successfully ported to the IBM Blue Gene/L system. Also, new I/O interface standards have demonstrated their efficacy in PVFS2.

- ***Improving model accuracy and performance.*** The SciDAC Performance Engineering Research Center (PERC), led by LBNL, has developed statistical methodologies for modeling application performance as a function of its input parameters. These methods provide estimates not only of performance but also of the model accuracy, allowing the user to continue refining the model until sufficient accuracy is achieved. It has been demonstrated that these techniques can be applied to computer architecture sensitivity studies, reducing by two orders of magnitude the number of cycle accurate simulations required to determine the best choices for architectural parameters such as cache size.

- ***Identifying the stumbling blocks when scaling to larger systems.*** Rice University has developed a new performance analysis technique for identifying each program component whose performance limits the scalability of parallel programs and quantifying the impact on overall program performance and scalability. This enables improved pinpointing of the factors and quantifying the effects that cause the efficiency of parallel programs to drop as they are scaled to larger systems, a long-standing problem in computer science.

FY 2006 Awards

- ***2006 R&D 100 Awards.*** Two SciDAC supported tools, developed at LLNL, have been included in the R&D Magazine's "R&D 100 Awards" for 2006. Babel is a High-Performance Language Interoperability Tool that that enables software pieces written in different programming languages to seamlessly call each other. Sapphire is scientific data mining software; a software toolkit for the analysis of massive, complex datasets arising from scientific experiments, observations, and computer simulations.
- ***Lagrange Prize in Continuous Optimization***. An ANL researcher, as part of a collaboration with researchers at the Universities of Dundee and Namur, has won the Lagrange Prize in Continuous Optimization for seminal work on filter methods. The prize is awarded by the Mathematical Programming Society and the Society for Industrial and Applied Mathematics once every three years for original work published in the past six years. The researchers designed a mathematical filter algorithm that solves a problem of nonlinear programming and may have use in a range of engineering and design applications.
- ***American Society of Mechanical Engineers Timoshenko Medal.*** An LBNL researcher was awarded the 2005 Timoshenko Medal "for seminal contributions to nearly every area of solid and fluid mechanics, including fracture mechanics, turbulence, stratified flows, flames, flow in porous media, and the theory and application of intermediate asymptotics." The Timoshenko Medal was established in 1957 and is conferred in recognition of distinguished contributions to the field of applied mechanics. Instituted by the Applied Mechanics Division, it honors Stephen P. Timoshenko, world-renowned authority in the field, and it commemorates his contributions as author and teacher.
- ***IEEE Fernbach Award.*** An LBNL mathematician has been named recipient of the 2005 Sidney Fernbach Award. The award is given by the IEEE Computer Society for an outstanding contribution in the application of high performance computers using innovative approaches. The award was presented at the SC05 supercomputing conference.
- ***Dijkstra Prize in Distributed Computing.*** The Association for Computing Machinery awards the Edsger W. Dijkstra prize to the authors of outstanding papers in distributed computing. The 2005 Dijkstra prize was awarded to an ASCR funded researcher at Rice University and his co-author for their paper on algorithms for scalable synchronization on shared memory multiprocessors.
- ***Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program.*** INCITE awarded over 18 million hours of high performance computer time to 15 scientific teams from U.S. laboratories, academia, and industry to address critical problems in the design of more efficient aircraft and engines; understanding the molecular basis of Parkinson's Disease; advancing fusion as a future energy source; understanding of human and ecological processes affecting climate change; learning about how cell disruptions allow diseases and infections to occur; developing

stronger advanced materials and better understanding material properties; improving simulations of molecular collisions which can be used to study a wide range of scientific problems; developing of computing tools to improve computer visualizations and animations; understanding water and how light affects water in biological systems; computing the structure of proteins at the atomic level; understanding of the dark energy and dark matter thought to make up more than 90% of our universe; and simulations of particle accelerators used in scientific research.