High-End Computing

Research and Development

Definition of HEC R&D PCA

The activities funded under the HEC R&D PCA include research and development to optimize the performance of today's high-end computing systems and to develop future generations of high-end computing systems. These systems are critical to addressing Federal agency mission needs and in turn many of society's most challenging large-scale computational problems.

Most high-end systems are clusters of processor nodes developed for the commercial computing market for use in personal computers and Web servers, and are not specifically targeted for HEC computing. Although the "peak" performance of these processors has followed Moore's Law, increasing more than five orders of magnitude in speed in the last decade, the sustained performance of scientific applications measured as a fraction of the peak value has continued to decline. Two reasons for this decline are: (1) these HEC systems are hundreds of times larger than those used in commercial applications and require a highly specialized software infrastructure to use effectively, and (2) the current HEC systems are unbalanced in that they do not have the optimal ratios of system parameters such as computation rate versus memory bandwidth.

To remedy this situation, HEC R&D supports both hardware and software R&D specifically aimed at scientific and national security applications. HEC R&D focuses on teraflop- through petaflop-scale systems and computation. Research activities in this area seek fundamental, long-term advances in technologies to maintain and extend the U.S. lead in computing for generations to come. Current research focuses on advanced computing architectures, software technologies and tools for high-end computing, mass storage technologies, and molecular, nano-, optical, quantum, and superconducting technologies.

HEC R&D engages collaborative research teams from academic institutions, national laboratories, and industrial partners in the development of new architectures that are well suited for algorithms used in scientific applications. HEC R&D supports fundamental investigations in memory, interconnect, and storage technologies in order to improve system balance. HEC R&D involves research across the entire spectrum of software issues – operating systems, languages, compilers, libraries, development environments, and algorithms – necessary to allow scientific applications to use the hardware effectively and efficiently. In addition, HEC R&D supports system modeling and performance analysis, which enable researchers to better understand the interaction between the computational requirements of applications and the performance characteristics of any proposed new high-end architectures.

Broad Areas of HEC R&D Concern

- Hardware architecture, memory, and interconnects
- Power, cooling, and packaging
- I/O and storage
- Comprehensive system software environment
- Programming models and languages
- System modeling and performance analysis
- Reliability, availability, serviceability, and security

HEC R&D Agencies

NSF
DOE/SC
DARPA
DOE/NNSA
NSA
NIH
NOAA
NASA
NIST

Participating Agency
DoD/HPCMPO

HEC R&D PCA Budget Crosscut

FY 2004 estimate	FY 2005 Request
$355.1 M	**$317.5 M**

TECHNICAL GOALS

- Parallel architecture designed for scientific application requirements
- Parallel I/O and file systems for sustained high data throughput
- Systems scalable to hundreds of thousands of processors
- Reliable and fault-tolerant systems
- Improved programming model expressability and parallel compiler support
- Effective performance measurement, optimization tools
- Improved ease of use and time to solution

ILLUSTRATIVE TECHNICAL THRUSTS

- Parallel component architecture
- Parallel I/O and high-performance file systems
- Next-generation architecture
- Extreme-scale operating and runtime systems
- Global address space language
- Software infrastructure tools
- Development and runtime performance tools
- High-productivity computer systems
- Quantum computing

ROADMAPS FOR FEDERAL HEC R&D IN THE YEARS AHEAD

HARDWARE

Component	Near- to Mid-Term	Long-Term
Microarchitecture	Prototype microprocessors developed for HEC systems available	Innovative post-silicon technology optimized for HEC
Interconnect technologies	Interconnect technology based upon optical interconnect and electrical switches	All-optical interconnect technology for HEC
Memory	Memory systems developed for HEC needs. Accelerated introduction of PIMs	Revolutionary high-bandwidth memory at petaflop scale
Power, Cooling, and Packaging	Stacked 3-D memory and advanced cooling technologies address critical design limitations	Ability to address high-density packaging throughout the entire system
I/O and storage	Petaflop-scale file systems with RAS focused on HEC requirements	Revolutionary approaches to exascale "file systems"

SOFTWARE

Component	Near- to Mid-Term	Long-Term
Operating systems (OSs)	New research-quality HEC OSs that address scalability and reliability	Production-quality, fault-tolerant, scalable OSs
Languages, compilers, and libraries	Optimized for ease of development on selected HEC systems. Research-quality implementations of new HEC languages, supporting multiple levels of abstraction for optimization.	High-level algorithm-aware languages and compilers for automated portability across all classes of HEC systems
Software tools and development environments	Interoperable tools with improved ease of use across a wide range of systems. First research-quality IDEs available for HEC systems.	IDEs that support seamless transition from desktop to largest HEC systems
Algorithms	New multiscale algorithms suitable for HEC systems. Initial prototypes of architecture-independent parallel computations.	Automatic parallelization of algorithms for irregular and unbalanced scientific problems. Scaling up of parallel algorithms to enable detailed simulations of physical systems.

SYSTEMS

Component	Near- to Mid-Term	Long-Term
System architecture	Three or more systems capable of sustained petaflops (up to 100,000 processors or more) on wider range of applications. Programming much simpler at large scale. Emergence of adaptable self-tuning systems.	High-end systems capable of sustained 10 to 100 petaflops on majority of applications. Programmable by majority of scientists and engineers. Adaptable self-tuning systems commonplace.
System modeling and performance analysis	Accurate models/tools for HEC systems and applications. Tools and benchmarks provide better understanding of architecture/application interactions.	Models enable analysis, prediction of software behavior. Automated, intelligent performance and analysis tools and benchmarks widely available, easy to use.
Programming models	Research implementations of novel parallel computing models. "Non-heroic" programming: MPI follow-on for 1,,024 processors and robust DSM implementations (UPC, CAF, …) widespread and available for 1,024 processors.	Parallel computing models that effectively and efficiently match new or planned architectures with applications. Novel parallel computation paradigms foster new architectures and new programming language features.
Reliability, availability, and serviceability (RAS) + Security	Semi-automatic ability to run through faults. Enhanced prevention of intrusion and insider attack.	Self-awareness: reliability no longer requires user assistance. Systems will have verifiable multilevel secure environments.

"Federal Plan for High-End Computing" roadmaps outline mid-term and long-term R&D advances required to revitalize U.S. HEC leadership.

HEC PCAs: Coordination and Activities

HEC HIGHLIGHTS

The HEC Coordinating Group (CG) coordinates activities funded under both the HEC I&A and HEC R&D PCAs. The CG provides a forum in which Federal agency program managers coordinate and collaborate on HEC research programs and on implementation of Federal high-end computing activities. The HEC CG is charged with:

- Encouraging and facilitating interagency coordination and collaborations in HEC I&A and R&D programs
- Addressing requirements for high-end computing technology, software, infrastructure, and management by fostering Federal R&D efforts
- Providing mechanisms for cooperation in HEC R&D and user access among Federal agencies, government laboratories, academia, industry, application researchers, and others

The HEC CG held a Special Meeting on March 19, 2004, with presentations and program descriptions from ten Federal agencies: DARPA, DoD/HPCMPO, DOE/NNSA, DOE/SC, EPA, NASA, NIST, NOAA, NSA, and NSF. The meeting enabled the agencies to identify opportunities for collaboration and coordination, areas for HEC CG focus, and areas – such as open source software – whose scope is broader than the HEC CG's.

Multiagency Interests and Collaborations

The HEC agencies reported the following common interests, collaborations, and commitments across a wide range of topics, including:

- Acquisition coordination: DoD/HPCMPO, DOE/NNSA, NASA
- Air-quality modeling applications: DOE/SC, EPA, NOAA
- Applied research for end-to-end systems development: NASA, NOAA, NSF
- Benchmarking and performance modeling: DoD/HPCMPO, DOE/SC, NASA, NSA, NSF
- Connectivity and technology delivery to universities: NOAA, NSF
- Climate and weather applications: DOE/SC, NASA, NOAA, NSF
- Earth System Modeling Framework (ESMF): DOE/SC, NASA, NOAA, NSF/NCAR (details on next page)
- Grid demonstrations: DOE/SC, NOAA, NSF
- Hardware development: DARPA, DOE/NNSA, NSA
- HECRTF: DoD/HPCMPO, DoD/OSD, DOE/NNSA, DOE/SC, EPA, NASA, NIST, NOAA, NSA, NSF
- HEC-URA: DARPA, DOE/SC, NSA, NSF (details on next page)
- HPCS Phase II: DARPA, DOE/NNSA, DOE/SC, NASA, NSA, NSF
- HPCS productivity metrics: DARPA, DoD/HPCMPO, DOE/NNSA, DOE/SC, NASA, NSA, NSF
- Joint memo expecting open source software and Service Oriented Data Access (SODA) for work funded at DOE labs: DOE/NNSA, DOE/SC
- MOU among HPCS mission partners for joint planning, coordination of directions, and leveraging each other's activities: DARPA, DoD/OSD, DOE/NNSA, DOE/SC, NSA
- National Research Council study on the "Future of Supercomputing": DOE/SC, DOE/NNSA (report expected in 2004)
- Optical switches and interconnects: DARPA, DOE/NNSA, NSA
- Quantum information science: DARPA, NIST, NSA
- Reviews of SV2 [procurements?]: DOE/NNSA, NSA [and DOE/SC?]
- Quarterly reviews of Cray X1e/Black Widow R&D programs: DoD/HPCMPO, DOE/NNSA, DOE/SC, NRO, NASA, NSA
- Reviews of ASC White and ASC Q software environments: DOE/NNSA, DOE/SC
- Single photon sources: DARPA, NIST
- Spray cooling: DOE/NNSA, DOE/SC
- Standards: DoD/HPCMPO, NIST, NOAA, NSA
- Technology transfer from universities: DoD/HPCMPO, DOE/NNSA
- Testbeds: DoD/HPCMPO, DARPA, NIST, NOAA, NSA
- Unified Parallel Compiler (UPC): DOE/SC, NSA
- Weather research and forecast (WRF): NOAA, NSF/NCAR

MULTIAGENCY ESMF EFFORT

Another broad-based FY 2004 NITRD activity that involves collaboration among multiple Federal agencies and coordination across multiple CGs, including HEC and SDP, is the building of the Earth System Modeling Framework (ESMF). Initiated by NASA researchers to develop tools enabling broader scientific use of the agency's extraordinary Earth data collections, the effort has become a key R&D component of Federal climate research across many organizations.

The ESMF is a high-performance, flexible software infrastructure designed to increase ease of use, performance portability, interoperability, and reuse in climate, numerical weather prediction, data assimilation, and other Earth science applications. The ESMF is an architecture for composing multi-component applications and includes data structures and utilities for developing model components. The goal is to create a framework usable by individual researchers as well as major operational and research centers, and to engage the weather research community in its development.

The ESMF addresses the challenge of building increasingly interdisciplinary Earth systems models and the need to maximize the performance of the models on a variety of computer architectures, especially those using upwards of thousands of processors. The new structure allows physical, chemical, and biological scientists to focus on implementing their specific model components. Software engineers design and implement the associated infrastructure and superstructure, allowing for a seamless linkage of the various scientific components.

The following agency organizations collaborate:

- NSF-supported National Center for Atmospheric Research (NCAR)
- NOAA's Geophysical Fluid Dynamics Laboratory (GFDL)
- NOAA's National Center for Environmental Prediction (NCEP)
- DoD's Air Force Weather Agency (AFWA)
- DoD's High Performance Computing Modernization Program Office (HPCMPO)
- DoD's Naval Research Laboratory (NRL)
- DOE's Argonne National Laboratory (ANL) and Los Alamos National Laboratory (LANL)
- NASA's Goddard Space Flight Center (GSFC)
- NASA Goddard Global Modeling and Assimilation Office (GGMAO)
- NASA Goddard Institute for Space Studies (GISS)
- NASA Goddard Land Information Systems project (GLIS)

ESMF Version 2.0, the first version of the software usable in real applications, was released in June 2004. It includes software for (1) setting up hierarchical applications, (2) representing and manipulating modeling components, fields, bundles of fields, and grids, and (3) services such as time management and logging messages.

NEW HEC-URA

DARPA, DOE/SC, NSA, and NSF collaboratively developed the HEC-University Research Activity (HEC-URA) launched in FY 2004 as an outgrowth of the HECRTF effort. The initiative funds universities to conduct research in software specifically for high-end computing. The strategy is to invest in:

- Basic and applied research to refill the academic pipeline with new ideas and people
- Advanced development of component and subsystem technologies
- Engineering and prototype development
    - – Integration at the system level
    - – Development of new technologies
- Test and evaluation to reduce the risk for development, engineering, and government procurement

Research emphases include DOE/SC and DARPA efforts on runtime operating systems for extreme-scale scientific computations and the NSF and DARPA focus on languages, compilers, and libraries for high-end computing. The agencies will participate in annual topic selections, solicitations, peer review selections, PI meetings, and the overall execution of this activity. Selections will be in accordance with individual agency requirements. Both university and industry researchers will be invited to make proposals.

In its first year, the HEC-URA is focused on software R&D to build a critical mass in research teams, advance the field toward 2010/2015 HEC software, avoid duplication, share lessons learned, and develop links between basic research and advanced development and engineering.

HIGH-END COMPUTING REVITALIZATION TASK FORCE (HECRTF)

In FY 2004, the NITRD agencies and others are engaged in a major planning initiative to reinvigorate Federal leadership in the development of leading-edge supercomputing capabilities to serve the national interest.

The activity was launched in April 2003 with a charge from OSTP to all government agencies with a major stake in high-end computing. The charge directed the agencies, under the auspices of the NSTC, to develop "a plan to guide future Federal investments in high-end computing" that would lay out an overall investment strategy based on "the needs of important Federally funded applications of HEC." The plan was to include five-year roadmaps for investment in R&D in core HEC technologies, as well as approaches to addressing Federal HEC capability, capacity, and accessibility issues, and Federal HEC procurement.

The 63-member High-End Computing Revitalization Task Force (HECRTF), representing 12 departments and agencies, spent 15 months gathering information, soliciting expert opinion from scientists, IT researchers, industry, and Federal HEC users, and developing a comprehensive report addressing the OSTP charge.

On May 13, 2004, OSTP Director John H. Marburger released the "Federal Plan for High-End Computing" at a hearing of the House Science Committee on H.R. 4218, "the High Performance Computing Revitalization Act of 2004." In his testimony, Dr. Marburger commended the task force and called its report "an important first step" in the "sustained and coordinated effort" required to address the issues facing the Nation's high-end computing enterprise. He noted that results of the HECRTF effort were already beginning to emerge – for example, the new High-End Computing University Research Activity (HEC-URA) jointly developed by DARPA, DOE/SC, NSA, and NSF to fund academic investigations in key HEC areas. FY 2004 solicitations focus on HEC software tools and operating/runtime systems for extreme-scale scientific computation.

Beginning in June 2004, the HECRTF Implementation Committee, which was created to activate the Federal Plan's recommendations, began a series of activities that include: 1) a meeting to coordinate FY 2006 agency HEC budget requests; 2) a meeting to plan cooperative development and use of HEC benchmarks and metrics; and 3) planning for workshops on testbeds for computer science research and on file systems.

THE NEED FOR HEC REVITALIZATION: EXCERPTS FROM REPORT

The full "Federal Plan for High-End Computing" is available online at: http://www.nitrd.gov/pubs/2004_hecrtf/20040510_hecrtf.pdf

The United States has repeatedly demonstrated that leadership in science and technology is vital to leadership in national defense and national security, economic prosperity, and our overall standard of living. Today, progress in many branches of science and technology is highly dependent on breakthroughs made possible by high-end computing, and it follows that leadership in high-end computing is increasingly crucial to the nation.

***

In the past decade, computer modeling and simulation of physical phenomena and engineered systems have become widely recognized as the cost-effective "third pillar" of science and technology – sharing equal billing with theory and experiment. Simulations are performed on computing platforms ranging from simple workstations to very large and powerful systems known as high-end computers. High-end computers enable investigations heretofore impossible, which in turn have enabled scientific and technological advances of vast breadth and depth. High-end computing (HEC) thus has become an indispensable tool for carrying out Federal agency missions in science and technology.

***

[Recent agency studies] have revealed that current high-end computing resources, architectures, and software tools and environments do not meet current needs. Of equal concern, they observe that sustained investigations of alternative high-end systems have largely stopped, curtailing the supply of ideas and experts needed to design and develop future generations of high-end computing systems.

***

Revitalization of high-end computing is needed to refill the research pipeline with new ideas and highly trained people, support the development of robust and innovative systems, and lower industry and end-user risk by undertaking the test and evaluation of HEC systems and software technologies. This revitalization must support advancement across the innovation spectrum – from basic and applied research, to advanced development, to engineering and prototyping, to test and evaluation. Such a comprehensive approach is essential to the establishment of a sustainable R&D process that encourages the generation of competing innovations from the basic research phase, the development of early prototype HEC systems, and their evaluation on mission-critical test applications.

KEY RECOMMENDATIONS OF THE HECRTF REPORT

OVERARCHING GOALS

The HECRTF recommends that the Federal government and its private-sector partners carry out comprehensive, complementary, and synchronized actions over the next five years with these goals: 1) Make high-end computing easier and more productive to use; 2) Foster the development and innovation of new generations of HEC systems and technologies; 3) Effectively manage and coordinate Federal high-end computing; and 4) Make high-end computing readily available to Federal agencies that need it to fulfill their missions.

HEC R&D

To achieve goals 1 and 2, the plan's roadmaps propose an integrated, balanced, and robust program of basic and applied research, advanced development, engineering and prototype development, and test and evaluation activities in each of three core HEC technologies:

- **Hardware** (microarchitecture; memory; interconnect; power, cooling, and packaging; and I/O and storage)
- **Software** (operating systems; languages, compilers, and libraries; software tools and development environments; and algorithms)
- **Systems technology** (system architecture; system modeling and performance analysis; reliability, availability, serviceability, and security; and programming models)

As part of this research effort, the task force recommends that the government return to a strategy – the development of HEC systems specifically for HEC R&D – that had been used in the early 1990's. These "early access" systems, called Research and Evaluation Systems by the HECRTF, will enable testing of early prototypes and provide development platforms for new algorithms and computational techniques, software functionality and scalability testing, and system and application testbeds. "Evaluation projects that identify failed approaches save the government from acquiring systems that simply do not perform as expected," the report states. "They may also suggest more fruitful approaches to remove the sources of failure."

The HECRTF also recommends that its proposed R&D agenda be balanced and prioritized to provide a continuum of efforts across the four developmental stages "required to sustain a vigorous and healthy high-end computing environment." The table on page 16 provides an overview of the proposed R&D priorities over the next five years.

HEC RESOURCES

To achieve goal 4, the plan offers strategies to address three distinct aspects of Federal HEC resources:

- **Access –** Federal agencies with limited or no HEC resources could partner with large agencies under mutually beneficial multiyear HEC utilization agreements. The plan proposes establishing pilot partnerships with industry to work on HEC problems, with metrics to evaluate the success of these arrangements.
- **Availability –** Agencies such as DoD, NSA, DOE/NNSA, DOE/SC, EPA, NASA, NOAA, and NSF all have programs that provide HEC resources to their respective user communities. The plan calls for an increase in the resources available to Federal mission agencies to acquire and operate HEC resources.
- **Leadership –** To provide world-leading high-end computing capabilities for U.S. scientific research, the plan proposes that the Government develop what the HECRTF calls Leadership Systems, HEC systems that will offer 100 times the computational capability of what is available in the marketplace. These very advanced resources would be made available for a limited number of cutting-edge scientific applications, enabling the U.S. to be "first to market" with new science and technology capabilities. The Leadership machines, managed as a shared Federal research resource, would not only support research but drive HEC commercial innovation.

HEC PROCUREMENT

A significant barrier to achieving all of the plan's goals is presented by the technical complexity and demands of Federal procurement processes, which burden both agencies and industry vendors. The plan proposes testing more cost-effective, efficient approaches in interagency pilot projects on:

- **Performance measurement, or benchmarking –** agencies with similar applications will develop a single suite of benchmarks for use in a pilot joint procurement process.
- **Accurate assessment of the total cost of ownership (TCO) –** a multiagency team will evaluate TCO (acquisition, maintenance, personnel, external networking, grid and distributed computing) for several similar systems and develop "best practices" for determining TCO for comparing possible procurements.
- **Procurement consolidation –** Using these new tools, agencies will test a combined procurement solicitation and evaluate its effectiveness.

PLAN IMPLEMENTATION

The HECRTF offers one example of an interagency governance structure for the plan but notes that it is only a starting point for discussion.

HECRTF PLAN'S PROPOSED R&D PRIORITIES TO FY 2010

		CURRENT PROGRAMS*	INCREMENT COMPARED TO HEC R&D CURRENT PROGRAM				
		FY2004 ($ IN MILLIONS)	FY2006	FY2007	FY2008	FY2009	FY2010
HARDWARE	Basic and Applied Research	$5	▥	▥	▥	▥	█
	Advanced Development	$5	▥	▥	▥	▥	▥
	Engineering and Prototypes	$0		▒	▒	▥	▥
	Test and Evaluation	$2		▒	▒	▒	▒
	TOTAL	$12					
SOFTWARE	Basic and Applied Research	$33	▒	▒	▒	▥	▥
	Advanced Development	$21	▒	▒	▒	▥	▥
	Engineering and Prototypes	$15	▥	▥	▥	▥	▥
	Test and Evaluation	$2	▥	█	█	█	█
	Long-term Evolution and Support	$0	▥	▥	▥	▥	▥
	TOTAL	$71					
SYSTEMS	Basic and Applied Research	$4	▒	▒	▒	▥	▥
	Advanced Development	$40	▒	░	░	░	▒
	Engineering and Prototypes	$1	▥	█	█	█	█
	Test and Evaluation	$30	░	▒	▥	█	█
	TOTAL	$75					
TOTALS		$158[1]					

░ Modest redirection
▥ Moderate funding increment
▒ Modest funding increment
█ Robust funding increment

*Assumes no planning changes from FY 2004

This table provides an overview of the Federal HEC R&D investments recommended by the HECRTF to address the needs of Federally funded R&D over the next five years. The level of effort called for in the major HEC components – hardware, software, and systems – is shown for each of the following categories of R&D work, which the Task Force deemed vital to a robust HEC R&D program:

- **Basic and Applied Research:** Focus on the development of fundamental concepts in high-end computing, creating a continuous pipeline of new ideas and expertise.
- **Advanced Development:** Select and refine innovative technologies and architectures for potential integration into high-end systems.
- **Engineering and Prototype Development:** Perform integration and engineering required to build HEC systems and components.
- **Test and Evaluation:** Conduct testing and evaluation of HEC software and current and new generations of HEC systems at appropriate scale.

The HECRTF report notes that innovations in hardware and systems have a natural transition from research into industrial manufacturing but that software research and development will require a different strategy. Since major advances in HEC system software have generally occurred only when academe, industry, and government research laboratories have teamed to solve common problems (such as by developing the message passing standard, MPI, to resolve internodal communications issues), the Task Force concluded that the software revitalization strategy should also include significant government involvement to ensure long-term maintenance of critical HEC software infrastructure components. The table thus shows a Long-Term Evolution and Support category under software.

NOTE 1: The FY 2004 funding levels shown in the table represent aggregate investments by all the agencies that participated in developing the "Federal Plan for High-End Computing," not just those of NITRD agencies. In addition, the figures represent only HEC R&D activities as defined by the HECRTF plan. They do not include many related activities (such as networking, visualization, and application-specific software development) considered outside the scope of the plan.

Columns four through eight present investment levels recommended by the Task Force for FY 2006 through FY 2010 in each of the categories and activities. Modest funding increments above the FY 2004 level are shown in medium orange. Moderate funding increments above the FY 2004 level are shown in lined orange. Robust funding increments above the FY 2004 level are shown in solid orange. Modest funding redirections from the FY 2004 level are shown in pale orange.

HEC I&A AND R&D PROGRAMS BY AGENCY

SELECTED FY 2004 ACTIVITIES AND FY 2005 PLANS

HEC NSF HEC

High-end computing capabilities are important to work supported by all the NSF research directorates and offices including: Biological Sciences (BIO); Computer and Information Science and Engineering (CISE); Education and Human Resources (EHR); Engineering (ENG); Geosciences (GEO); Mathematical and Physical Sciences (MPS); and Social, Behavioral and Economic Sciences (SBE); and the Office of Polar Programs (OPP).

Directorates have complementary HEC investments. For example: R&D on computer architecture, networking, software, and cyberinfrastructure is funded by CISE; HEC devices are funded by MPS and ENG; mathematical algorithms are funded MPS and CISE; computational algorithms and libraries are funded by CISE, with some funding from MPS; science and engineering applications are developed primarily with funding from MPS, ENG, GEO, and BIO.

NSF HEC R&D Programs

NSF's largest investments in HEC resources are in the CISE Directorate. All CISE Divisions are responsible for HEC activities, as follows:

Computing and Communication Foundations (CCF) Division – has responsibility for:

- Formal and Mathematical Foundations – in particular algorithmic and computational science
- Foundations of Computing Processes and Artifacts – in particular high-end software, architecture, and design
- Emerging Models for Technology and Computation – includes biologically motivated computing models, quantum computing and communication, and nanotechnology-based computing and communication systems

Computer and Network Systems (CNS) Division – supports programs in computer and network systems, in particular distributed systems and next-generation software systems.

Information and Intelligent Systems (IIS) Division – supports programs in data-driven science including bioinformatics, geoinformatics, cognitive neuroscience, and other areas.

Shared Cyberinfrastructure (SCI) Division – supports programs in:

- Infrastructure development – creating, testing, and hardening next-generation deployed systems
- Infrastructure deployment – planning, construction, commissioning, and operations

The following illustrate ongoing HEC R&D:

- Chip Multiprocessor Architectures: on-chip shared-cache architectures, arithmetic logic unit (ALU) networks, shared memory architectures, and networks
- Scalable Multiprocessing Systems: system architecture (communication substrate) and (heterogeneous) ensembles of chip multiprocessors
- System-on-a-Chip building blocks and integrated functionality
- Dynamic and Static (compiler) Optimizations: memory systems (SMT/CMP support latency reduction, prefetching, and speculative execution [load/value prediction]), and verification and runtime fault tolerance
- Networking and storage

NSF also funds the following speculative research that could provide breakthrough HEC technologies:

- Nanotechnology that promises unprecedented scale-up in computing and communication through research in nano-architecture and design methodologies
- Post-silicon computing: development and application of technologies such as quantum, DNA, chemical, and other computing concepts to provide radically new models of computation, algorithms, and programming techniques.

NSF Cyberinfrastructure Framework

A key NSF-wide interest in FY 2004 is development of an overarching IT framework to support the communities performing research and using computing resources. Components include hardware (distributed resources for computation, communication, and storage), shared software cybertools (grid services and middleware, development tools, and libraries), domain-specific cybertools, and applications to make possible discovery and innovative education and training. The framework will consist of:

- Computation engines (supercomputers, clusters, workstations), including both capability and capacity hardware
- Mass storage (disk drives, tapes, and related technologies) with persistence
- Networking (including wireless, distributed, ubiquitous)
- Digital libraries and data bases
- Sensors and effectors
- Software (operating systems, middleware, domain-specific tools, and platforms for building applications)
- Services (education, training, consulting, user assistance)

R&D in every component of this infrastructure will enable HEC capabilities.

Support for the Supercomputing Needs of the Broad Scientific Community

NSF plays a unique role nationally in supporting supercomputing needs across the entire academic spectrum of scientific research and education. NSF's SCI Division provides support for and access to high-end computing infrastructure and research. NSF supports more than 22 high-performance computing systems that include both capacity and capability systems. In FY 2004, NSF supports supercomputing resources at the following national partnerships and leading-edge sites:

- The National Computational Science Alliance (the Alliance) led by the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign
- The National Partnership for Advanced Computational Infrastructure (NPACI) led by the San Diego Supercomputer Center (SDSC) at the University of California-San Diego
- The Pittsburgh Supercomputing Center (PSC) led by Carnegie Mellon University and the University of Pittsburgh, with its Terascale Computing System (TCS)

In FY 2003, NSF had 4,450 HEC users including 1,800 NPACI users, 1,200 NCSA users, and 1,450 PSC users – for an estimated total of 3,000 unique users across all the centers. These organizations participate in the following ongoing NSF infrastructure initiative:

Extensible Terascale Facility (ETF) – multiyear effort to build and deploy the world's largest, most comprehensive, distributed infrastructure (also known as the TeraGrid) for open scientific research. Four new TeraGrid sites, announced in September 2003, will add more scientific instruments, large datasets, computing power, and storage capacity to the system. The new sites are ORNL, Purdue University, Indiana University, and the University of Texas at Austin. They join NCSA, SDSC, ANL, CalTech, and PSC.

With the addition of the new sites, the TeraGrid will have more than 20 teraflops of computing power, be able to store and manage nearly one petabyte of data, and have high-resolution visualization environments and toolkits for grid computing. All components will be tightly integrated and connected through a network that operates at 40 gigabits per second.

NSF FY 2005 plans include the following:

HEC R&D:

- ***CISE emphasis area focused on high-end computing*** *– new initiative intended to stimulate research and education on a broad range of computing and communication topics, including architecture, design, software, and algorithms and systems for computational science and engineering. While the range of topics is broad, the research and education activity is to focus strictly on extreme-scale computing. One component of this emphasis area is:*
  - *The special NSF/DARPA activity "Software & Tools for High-End Computing" (ST-HEC), developed as a High End Computing University Research Activity (HEC-URA) – supporting high-end software tools for extreme-scale scientific computation, which are highly computation- and data-intensive and cannot be satisfied in today's typical cluster environment. The target hosts for these tools are high-end architectures, including systems comprising thousands to tens of thousands of processors.*
- ***Activities in planning or evaluation in all NSF directorates on cyberinfrastructure development*** *– will include research and education areas in computational science and engineering, such as high-performance computational algorithms and simulation-based engineering*

HEC I&A:

- *Expansion of the Extensible Terascale Facility (ETF) will continue, providing high-end computing and networking for colleges and universities*

- ***Cyber-Infrastructure Training, Education, Advancement, and Mentoring (CI-TEAM) –*** *new FY 2005 activity will support the preparation of a workforce with the appropriate knowledge and skills to inform the development, deployment, support, and widespread use of cyberinfrastructure. Goal is to support efforts that will prepare current and future generations of scientists, engineers, and educators to use, develop, and support cyberinfrastructure, and broaden participation of underrepresented groups and organizations in cyberinfrastructure activities.*
- ***Linked Environments for Atmospheric Discovery (LEAD)*** *– will create grid-computing environments for on-demand prediction of high-impact local weather such as thunderstorms and lake-effect snows.*
- ***International Virtual Data Grid Laboratory (iVDGL)*** *– collaboration of 16 institutions and more than 80 scientists is deploying a computational and data grid for four international high-energy physics experiments, including the Large Hadron Collider.*

HEC DOE/SC HEC

DOE/SC supports basic research that underpins DOE missions and constructs and operates for the U.S. scientific community large scientific facilities such as accelerators, synchrotron light sources, and neutron sources. Six offices oversee DOE/SC research: Basic Energy Sciences (materials, chemistry, and nanoscale science); Biological and Environmental Research (global climate, genomics research); Fusion Energy Sciences (plasma fusion energy); High Energy Physics (accelerator design, petascale experimental data analysis); Nuclear Physics (quantum chromodynamics [QCD], astrophysics); and Advanced Scientific Computing Research.

Role of Computational Science

Computational science is critical to the DOE/SC mission in energy production, novel materials, climate science, and biological systems, where:

- Systems are too complex for direct calculation and descriptive laws are absent
- Physical scales range over many orders of magnitude
- Several scientific disciplines are involved such as in combustion and materials science
- Multidisciplinary teams are required
- Experimental data may be costly to develop, insufficient, inadequate, or unavailable
- Large data files are shared among scientists worldwide

DOE/SC's major HEC activities are integrated in the following large-scale multiyear computer science initiative to accelerate advanced research in scientific fields of interest to DOE:

Scientific Discovery Through Advanced Computing (SciDAC) Program – provides terascale computing and associated information technologies to many scientific areas to foster breakthroughs through the use of simulation. Through collaborations among application scientists, mathematicians, and computer scientists, SciDAC is building community simulation models in plasma physics, climate prediction, combustion, and other application areas. State-of-the-art collaboration tools facilitate access to these simulation tools by the broader scientific community to bring simulation to a level of parity with theory and observation in the scientific enterprise.

DOE/SC HEC R&D Programs

Recent DOE/SC accomplishments in HEC R&D that help address DOE computational science needs include:

- PVM, a widely used early message passing model
- MPI, MPICH – THE standard message passing model in use today (MPICH is the reference implementation used by all commercial vendors to develop tuned versions)
- Global arrays – global memory model used by NWChem, which greatly simplified electronic structure programs in computational chemistry
- Co-Array Fortran – the first and only open source Co-Array Fortran compiler
- UPC – the first open source Unified Parallel C compiler
- SUNMOS/Puma/Cougar/Catamount – microkernel operating system used by ASC Red system and its 40 teraflops replacement being developed by Cray, Inc.
- Science Appliance – pseudo single system image Linux-

based operating system used by large-scale (10 teraflops) clusters

- OSCAR – in a partnership with industry, the most widely used open source toolkit for managing Linux clusters
- Low-cost parallel visualization – use of PC clusters to achieve high-performance visualization

SciDAC Integrated Software Infrastructure Centers (ISICs) – SciDAC interdisciplinary team activities that address specific technical challenges in the overall SciDAC mission. Three Applied Mathematics ISICs are developing new high-performance scalable numerical algorithms for core numerical components of scientific simulation, and will distribute those algorithms through portable scalable high-performance numerical libraries. These ISICs are:

- Algorithmic and Software Framework for Applied Partial Differential Equations (APDEC) – to provide new tools for (nearly) optimal performance solvers for nonlinear partial differential equations (PDEs) based on multilevel methods
- Terascale Optimal PDE Solvers (TOPS) – hybrid and adaptive mesh generation and high-order discretization techniques for representing complex, evolving domains
- Terascale Simulation Tools and Technologies (TSTT) – tools for the efficient solution of PDEs based on locally structured grids, hybrid particle/mesh simulations, and problems with multiple- length scales

Four Computer Science ISICs are working closely with SciDAC application teams and the math ISICs to develop a comprehensive, portable, and fully integrated suite of systems software and tools for the effective management and use of terascale computational resources by SciDAC applications. The Computer Science ISICs are:

- Center for Component Technology for Terascale Simulation Software ISIC – to address critical issues in high-performance component software technology
- High-End Computer System Performance: Science and Engineering ISIC – to understand relationships between application and architecture for improved sustained performance
- Scalable Systems Software ISIC – for scalable system software tools for improved management and use of systems with thousands of processors
- Scientific Data Management ISIC – for large-scale scientific data management

SciDAC National Collaboratory Software Environment Development Centers and Networking Research – DOE's investment in National Collaboratories includes SciDAC projects focused on creating collaboratory software environments to enable geographically separated scientists to effectively work together as a team and to facilitate remote access to both facilities and data.

DOE/SC Workshops

Workshops – some general, others highly focused – are a standard DOE/SC planning tool. FY 2004 workshops include:

- SciDAC
- Fast Operating Systems
- Multiscale Mathematics Workshop – to bridge the wide scale between modeling on continuum or atomistic basis; expected output includes a roadmap
- High Performance File Systems – such systems should support parallelism, be efficient, and be reasonably easy to use
- Distributed Visualization Architecture Workshops (DiVA)
- Data Management Workshop series – to address where file system ends and data-management system begins
- Raising the level of parallel programming abstraction
- Parallel tool infrastructure

DOE/SC HEC I&A Programs

National Energy Research Supercomputing Center (NERSC) – delivers high-end capability computing services and support to the entire DOE/SC research community including the other DOE laboratories and major universities performing work relevant to DOE missions. Provides the majority of resources and services supporting SciDAC. Serves 2,000 users working on about 700 projects – 35 percent university-based, 61 percent in national laboratories, 3 percent in industry, and 1 percent in other government laboratories.

NERSC's major computational resource is a10-teraflops IBM SP computer. Overall NERSC resources are integrated by a common high-performance file storage system that facilitates interdisciplinary collaborations.

Advanced Computing Research Testbeds (ACRT) – activity supports the acquisition, testing and evaluation of advanced computer hardware testbeds to assess the prospects for meeting future computational needs such as SciDAC and special-purpose applications. The ACRT provide two types of computer testbeds for evaluation – early systems and experimental systems. These research and evaluation (R&E) prototypes have been identified as a critical element in the HECRTF plan because they enable early partnerships with vendors to tailor architectures to scientific requirements. The results from these partnerships also play a key role in the choice of both high-performance production systems and potential leadership class systems governmentwide.

Scientific Application Pilot Programs (SAPP) – support targeted efforts to integrate new applied mathematics and computer science algorithms into SciDAC applications projects through strong liaisons between applications projects and ISICs. SAPP provides expertise in adapting high-performance algorithms to terascale computers and in adapting numerical modules to include application-specific knowledge.

Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Projects – new FY 2004 competitive program that supports a small number of computationally intensive, large-scale research projects that can make high-impact scientific advances through the use of a substantial allocation of computer time and data storage at NERSC. INCITE encourages proposals from universities and other research institutions.

In December 2003, DOE/SC selected three projects to receive a total of 4.9 million hours of supercomputing time – 10 percent of the total computing time available this year on NERSC's Seaborg system. The projects are expected to significantly advance understanding of the makeup of the universe, the chemical process by which plants convert sunlight to energy while removing carbon dioxide from the atmosphere, and the turbulent forces that affect everything from weather to industrial processes. The projects are:

- Thermonuclear Supernovae: Stellar Explosions in Three Dimensions
- Quantum Monte Carlo Study of Photoprotection via Carotenoids in Photosynthetic Centers
- Fluid Turbulence and Mixing at High Reynolds Number

DOE/SC FY 2005 plans include:

HEC I&A

- *Installation of a Linux Cluster with more than 5 Tflops at NERSC to support high-end capacity needs of science applications that do not scale to more than 256 processors*
- *Partnership with DOE Office of Fusion Energy to develop next generation of Tokamak simulation tools*
- *Initiative in multiscale mathematics to develop the mathematical underpinnings for the simulation of complex physical and biological systems*

HEC R&D

- *Continuation of partnership with DARPA in productivity metrics for high-end computers*
- *Continuation of HEC-URA efforts to develop software base for future high-end systems*
- *Complete evaluation of Cray X1 computer*

HEC NIH HEC

The scientific drivers behind requirements for supercomputers at NIH are recent and are increasing rapidly. Although biomedical supercomputing is still in an early stage of development, NIH grantees, taking the individual initiative to apply for computer time, already consume approximately one-third of the cycles provided by the NSF supercomputing centers at Illinois, Pittsburgh, and San Diego. NSF's planned TeraGrid will most likely receive similar use by the biomedical research community.

The largest portion of the current biomedical supercomputing use is for biomolecular modeling and simulation. This class of application will continue to grow rapidly. The areas of high-throughput bioinformatics, data integration, and modeling of complex biosystems will grow explosively as software linking high-throughput data gathering with automated analysis and modeling tools becomes widely available. This software is currently in early development, with NIH support through new grant programs instituted in the last few years. To realize the potential of imaging and medical informatics will also require increased computational resources.

NIH has in place a planning process for the next stage of growth by preparing an eight- to ten-year plan for bioinformatics and computational biology that outlines the software needs and challenges for the next decade of biomedical computing.

Principal new investment areas for NIH include:

- Expanded development of molecular and cellular modeling software
- Applied computer science, such as algorithm creation and optimization, creation of appropriate languages and software architectures applicable to the solution of biomedical problems
- R&D in biomedical computational science, by developing and deploying software designed to solve particular biomedical problems
- Providing infrastructure to serve the needs of the broad community of biomedical and behavioral researchers
- Enhancing the training for a new generation of biomedical researchers in appropriate computational tools and techniques
- Disseminating newly developed tools and techniques to the broader biomedical research community

In FY 2005, NIH plans to:

- *Continue developing long-term plan for bioinformatics and computational biology, including listed areas of new investment*

HEC DARPA HEC

In FY 2004, DARPA supports the following HEC programs:

High Productivity Computing Systems (HPCS) program – addresses all aspects of HEC systems (packaging, processor/memory interfaces, networks, operating systems, compilers, languages, and runtime systems). Goal of multiyear effort is to develop a new generation of economically viable high- productivity computational systems for the national security and industrial user community in the time frame of FY 2009 to FY 2010. Sustained petaflop performance and shortened development time are key targets. DARPA is working with DOE/SC and other agencies to understand how long it takes to develop solutions today. These systems will be designed to have the following impact:

- Performance (time to solution) – provide a 10-to-40-times performance speedup on critical national security applications
- Programmability (idea to first solution) – reduce the cost and time of developing applications solutions
- Portability (transparency) – insulate the research and operational application software from the system on which it is running
- Robustness (reliability) – apply all known techniques to protect against outside attacks, hardware faults, and programming errors

Target applications are intelligence and surveillance, reconnaissance, cryptanalysis, weapons analysis, airborne contaminant modeling, and biotechnology.

The HPCS program is structured to enable ongoing fielding of new high-performance computer technology through phased concept, research, development, and demonstration approach. The program is in its second phase, with productivity goals including:

- Execution (sustained performance) of 1 petaflop/second (scalable to over 4 petaflop/second)
- Development productivity gains of 10 times over today's systems

In phase 2, a productivity framework baselined to today's systems is used to evaluate vendors' emerging productivity techniques and will provide a reference for phase 3 evaluations of vendors' proposed designs. Phase 2 subsystem performance indicators are a 3.2 petabytes/second bisection bandwidth, 64,000 gigabyte updates per second (GUPS), 6.5 petabytes data streams bandwidth, and 2+ petaflops Linpack Fortran benchmark (an HPC Challenge) that augments the Top 500 benchmarks.

The NITRD agencies that provide additional funding for HPCS activities are DOE/SC, DOE/NNSA, NASA, NSA, and NSF; DoD's HPCMPO provides complementary support. Cray, Inc., IBM, and Sun Microsystems are phase 2 industry partners.

Council on Competitiveness HPC Initiative – HPCS-related co-funded by DARPA, DOE/SC, and DOE/NNSA has the goal of raising high productivity computing to be a major economic driver in the United States. Effort includes: an HPC Advisory Committee with representatives from the private sector; an annual private-sector user survey; and annual user meetings. The first conference, entitled "HPC: Supercharging U.S. Innovation and Competitiveness," was held July 13, 2004, in Washington, D.C.

Polymorphous Computing Architectures (PCA) Program– aims to develop the computing foundation for agile systems by establishing computing systems (chips, networks, and software) that will morph to changing missions, sensor configurations, and operational constraints during a mission or over the life of the platform. Response (morph) times may vary from seconds to months. The PCA objective is to provide:

- Processing diversity with a breadth of processing to address signal/data processing, information, knowledge, and intelligence, and uniform performance competitive with best-in-class capabilities
- Mission agility to provide retargeting within a mission in millisecond, mission-to-mission adaptation in days, and new threat scenario adaptation in months
- Architectural flexibility to provide high-efficiency, selectable virtual machines with two or more morph states, mission adaptation with N minor morph states, and portability/evolvability with two-layer morphware

The PCA hardware being developed includes:

- Monarch/MCHIP – 2005 prototype is projected to demonstrate 333Mhz, 64 GFLOPS, 64 GOPS, and 12 Mbytes EDRAM.
- Smart Memories – 2005 prototype is projected to demonstrate Brook Language based on modified commercial Tensilica CPU cores.
- TRIPS – 2005 prototype is projected to demonstrate 500 MHz, 16 Gflops peak floating point, and 16 GIPS peak integer. Production technology is projected to provide 4 Tflops peak floating point and 512 GIPS peak integer.
- RAW (early prototype) to provide chips and development boards (available now) and kernel evaluations with test results in 2004. The technology provides early morphware, software, and kernel performance evaluation. The morphware will be demonstrated in a development environment in which APIs, including Brook, will use high-level compilers to provide stream and thread virtual machine APIs to low-level compilers using PCA technology such as TRIPS, Monarch, Smart Memories, and RAW.

Networked Embedded Systems Technology (NEST) – will enable fine-grain fusion of physical and information processes. The quantitative target is to build dependable, real-time, distributed, embedded applications comprising 100 to 100,000 simple compute nodes. The nodes include physical and information system components coupled by sensors and actuators.

In FY 2005, work will continue on the following DARPA efforts:

- *High Productivity Computing Systems (HPCS)*
- *Polymorphous Computing Architectures (PCA)*

HEC NASA HEC

FY 2004 NASA accomplishments in high-end computing:

- Testbeds
  - – Built the first SSI Linux base supercomputer, the world's fastest machine on the Triad Benchmark
- Advanced Architectures
  - – Invested in new architecture streams and PIM working on languages and compilers and key application kernels for stream technology (this is tied into DARPA's PCA and morphware efforts)
- Grid Technology
  - – Continued work on NASA's Information Power Grid (IPG) demonstration in cooperation with Earth Science and Shuttle payload management
- Applications
  - – Developed advanced applications for aerospace, nanotechnology, and radiation biology
  - – Provided major advance in ocean modeling in the Estimating the Circulation and Climate of the Ocean (ECCO) code
  - – Major advance in aerospace vehicle modeling
- Programming environments
  - – Improved scalability of multilevel parallelism (MLP)
- Tools
  - – Workflow and database management programming environments
  - – Aero database/IPG Virtual Laboratory for workflow management and ensemble calculation database-management prototypes
  - – Hyperwall (for viewing large datasets) used in nanotechnology research, the Sloan Digital Sky Survey, and the Columbia Space Shuttle Mission accident investigation. This technology is now being transferred within NASA for general use.

Estimating the Circulation and Climate of the Ocean (ECCO) – project goal is to use NASA data and community code to improve Earth oceans modeling. Starting in August 2003, the project has:

- Developed serial #1 512 Altix supercomputer
- Ported and scaled ECCO code from 64 to 128 to 256 to 512 processors
- Determined correct environment variables, optimized code
- Assisted in revamping grid technique to improve prediction at poles
- Developed Kalman filter techniques at high resolution for data assimilation
- Developed archive for community access to data
- Improved resolution from 1 degree to 1/4 degree
- Developed unique visualization approaches

As a result of this work:

- Time to solution dropped from approximately one year to five days (about a 50-fold speedup)
- Resolution improved by a factor of four
- Code accuracy and physics improved
- Data assimilation at high resolution was enabled

Shuttle return-to-flight problems are now being ported to the Altix supercomputer for speedup by a factor of four or more for some turbo-machinery.

Supercomputing support for the Columbia Accident Investigation Board (CAIB) – NASA Ames supercomputing-related R&D assets and tools produced time-critical analyses for the CAIB. R&D and engineering teams modeled a free object in the flow around an airframe with automatic regridding at each time step. The environment in which this was accomplished entailed:

- State-of-the-art codes
  - – Codes honed by R&D experts in large-scale simulation, with their environment of programmer tools developed to minimize the effort of executing efficiently
- A modeling environment that included experts and tools (compilers, scaling, porting, and parallelization tools)
- Supercomputers, storage, and networks
  - – Codes and tools tuned to exploit the advanced hardware that was developed for large tightly-coupled computations
- Computation management by components of the AeroDB and the iLab
  - – An environment and tools for efficient executions of hundreds or thousands of large simulations and handling of results databases

- Advanced data analysis and visualization technologies to explore and understand the vast amounts of data produced by this simulation

NASA HEC Resources

NASA high-end computing resources include:

- Goddard Space Flight Center, where the foci are in Earth Science and Space Science
  - 1,392 processing element (PE) Compaq (2.2 Teraflops)
  - 640 PE SGI O3K (410 Gigaflops)
  - Sun QFS (340 Terabytes)
  - SGI DMF (370 Terabytes)
- Ames, where the foci are in Aeronautics and Earth Science
  - 512 PE SGI Altix (2.3 teraflops)
  - 1,024 PE SGI O3K (850 gigaflops)
  - SGI DMF (1,600 terabytes)

NASA FY 2005-FY 2007 HEC R&D plans include:

- *Collaborative decision systems to improve decision making*
- *Discovery systems to accelerate scientific progress*
- *Advanced networking and communications*
- *Advanced computing architectures and technologies*
- *Reliable software*
- *Adaptive embedded information systems*

HEC DOE/NNSA HEC

DOE/NNSA's Office of Advanced Simulation and Computing (ASC) provides the means to shift promptly from nuclear test-based methods to computer-based methods and thereby predict with confidence the behavior of nuclear weapons through comprehensive science-based simulations. To achieve this capability, ASC has programs in:

- Weapons codes and the science in those codes (advanced applications, materials and physics modeling, and verification and validation)
- Computational platforms and the centers that operate them
- Software infrastructure and supporting infrastructure (Distance computing [DisCom], Pathforward, Problem Solving Environments, and VIEWS)

ASC HEC R&D focuses on:

- Software-quality engineering oriented to processes to provide for longevity of codes
- Verification to ensure numerical solutions are accurate
- Validation to ensure that a problem is understood correctly
- Certification methodology – improvements through scientific methods
- Capability computing to meet the most demanding computational needs
- Capacity computing to meet current stockpile workloads
- Problem Solving Environments (PSEs) to create usable execution environments for ASC-scale applications
- Industry collaboration to accelerate key technologies for future systems
- Tracking requirements

ASC activities – Pathforward work includes:

- Optical switch (with NSA)
- Memory usage
- Lustre file system
- Spray cooling (with DOE/SC)

ASC HEC I&A platforms – HEC systems and software from Cray, HP, IBM, Intel, SGI, and Linux NetworX. ASC procures machines with differing technologies to support diverse requirements. These types include:

- Capability systems: Vendor-integrated SMP clusters used for critical Stockpile Stewardship Program (SSP) deliverables and for particularly demanding calculations. (These systems run a small number of large jobs.)
- Capacity systems: Linux HPC clusters that carry the largest workloads in terms of numbers of jobs supporting current weapons activities and as development vehicles for next-generation capability systems. (These systems run a large number of smaller jobs.)
- Disruptive technologies: The Blue Gene/L system, for example, is being procured as a highly scalable, low-

power, low-cost computer. It is the first high-performance computing system to incorporate all three attributes in a single system. ASC works with vendors and universities to push production environments to petaflop scales. ASC works with industry to develop new technologies such as optical switches, scalable visualization, memory correctness, spray cooling, and the Lustre file system.

DOE/NNSA FY 2005 HEC plans include:

- *Continue R&D in software-quality engineering; validation, verification, and certification of software; PSEs; and Pathforward topics*
- *Continue development of Blue Gene/L*

HEC NSA HEC

NSA has unclassified HEC R&D programs in the areas of:

- Architectures and systems
- High-speed switches and interconnects
- Programming environments
- Quantum information sciences
- Vendor partnerships

Cray X1e/Black Widow – under multiyear joint development effort, NSA supports Cray Inc. in extending its X1 NSA/ODDR&E/Cray development to build its next-generation scalable hybrid scalar-vector high-performance computer. The technology will result in an X1e upgrade in the market in late 2004. Black Widow system to be introduced in 2006 will provide outstanding global memory bandwidth across all processor configurations and will be scalable to hundreds of teraflops. Should be the most powerful commercially available system in the world at that time. Many agencies have regularly participated in the quarterly development reviews; several have serious plans for acquisition of these systems.

HPC Applications Benchmarking Study – project initiated by HPC users, the HPC User Forum, and IDC to study the performance of applications on leading supercomputers. The project will expand the understanding of the performance of a set of complete, difficult HPC applications on the largest computers around the world. Each application will have a specific data set and will represent an actual scientific problem. The results will be posted at the HPC User Forum Web site. This study's objectives are to compare leading large HPC capability class computers on a small set of large HPC problems, show major science accomplishments on these capability systems, explore the relative performance characteristics of the machines on these problems, and explore the scaling abilities of the computers. The HPC Users Forum effort is funded by NSA, DoD/HPCMPO and several other agencies.

Reconfigurable Computing Research and Development – an approach to general-purpose high-performance computing that takes advantage of field programmable gate array (FPGA) technologies and commercial platforms. This continuing effort is leveraging prototype results from academic, industrial, and government research and the continuing increase in the computational capability of commercial FPGA chips. NSA is conducting experiments and demonstrations of commercial reconfigurable computing machines using its applications and benchmarks. The principle difficulty being addressed is the need for software-level programming tools. There is substantial collaboration with the Air Force, NASA Langley, and the NRO in reconfigurable research and development.

Major FY 2004 achievements:

- Completed a highly successful demonstration on an NSA application of the C Compiler on a particular reconfigurable computer to achieve a 65x performance improvement over a modern 64-bit commodity microprocessor.
- Co-funded an NSA vendor partnership to develop an FPGA addition to a commercial architecture, with a prototype machine to be delivered in FY 2004.

Fine-grain Multithreaded Architectures and Execution Models – ongoing research program in collaboration with the University of Delaware to develop architectures, execution models and compilation techniques that can exploit multithreading to manage memory latency and provide load balancing.

Performance analysis of many real-world applications on today's large scale massively parallel processor systems indicates that often only a relatively small fraction of the potential performance of these machines is actually achieved in practice. The applications that exhibit these problems are generally characterized by memory access patterns that are dynamic and irregular, and hence difficult to partition efficiently, either by programmers or compilers. This situation frequently results in processors being "starved" for data due to the relatively high memory latencies. Fine-grain multithreaded models coupled with effective compilation strategies provide a promising and easy to use approach to reducing the impact of long, unpredictable memory latencies.

Work in the past year has resulted in an FPGA-based implementation of a hardware emulation platform and a compiler infrastructure that is enabling experimentation with thread scheduling and synchronization units.

Optical Technologies – Columbia University and Georgia Tech University commissioned by NSA to investigate high-end computer optical interconnect systems. Columbia's comparison of electrical and optical technologies showed conclusively that specifically designed all-optical switching can significantly outperform current and future electronic switches. The collaborative effort demonstrated that optical technologies can significantly improve bandwidth if implementation difficulties of all-optical switching can be overcome. Future efforts will investigate how current optical components might be integrated to form a demonstration system.

Support to DARPA HPCS Program, Phase 2 – NSA participation includes funding, technical guidance, and participation in program reviews.

Quantum Computing – major research program being conducted through NSA's ARDA affiliate. Collaboration involves DARPA and ARL. Efforts focus on quantum information sciences including investigation of selective materials, related quantum properties, and managing decoherence and entanglement issues. Funded research under this multiagency collaboration is underway at many universities, government laboratories, and other centers of excellence.

FY 2005 NSA plans in HEC include:

- *Continue R&D in all listed FY 2004 HEC R&D areas*
- *Increase funding support for Black Widow development consistent with its calendar-year 2006 introduction*
- *Maintain current research effort in all other efforts, including quantum computing.*

HEC NOAA HEC

NOAA uses HEC resources and cutting-edge applications to provide advanced products and services delivered to users. NOAA's HEC strategy is to:

- Develop skills, algorithms, and techniques to fully utilize scalable computing for improved environmental understanding and prediction
- Acquire and use high-performance scalable systems for research
- Optimize the use of HEC resources across all NOAA activities. This reflects NOAA's view of HEC as an enterprise resource.

NOAA's HEC I&A resources are located at the:

- Geophysical Fluid Dynamics Laboratory (GFDL) in Princeton, New Jersey. (In FY 2003, GFDL acquired a 1,408-processor SGI system.)
- Forecast Systems Laboratory (FSL) in Boulder, Colorado
- National Center for Environmental Prediction (NCEP), headquartered in Camp Springs, Maryland

NOAA expedites the development and use of improved weather and climate models by:

- Supporting advanced computing for NOAA environmental modeling research
- Developing software tools to optimize the use of scalable computing
- Infusing new knowledge through new talent such as postdoctoral students and contractors

NOAA Development Test Center (DTC) – establishing a collaborative environment in which university research and NOAA personnel can work together, focusing initially on development of the Weather Research and Forecasting (WRF) model that is designed as both an operational model and a research vehicle for the larger modeling community. Collaborators on the WRF include FNMOC, NAVO, AFWA, NCAR, and others. The goal is for Federal and university researchers working on specific problems to explore models, develop improvements, plug in to that part of the model, help the entire community, and increase speed going from basic research to applications. Potential benefits include frequent high-resolution analyses and forecasts produced in real time that are valuable to commercial aviation, civilian and military weather forecasting, the energy industry, regional air pollution prediction, and emergency preparedness.

NOAA is working cooperatively on the ESMF with NASA and others, and is coordinating with other agencies its research plans including the Climate Change Strategic Plan and the U.S. Weather Research Program. Developing grid technology with GFDL, PPPL, PU, FSL, PMEL, and CSU.

NOAA FY 2005 plans include the following:

- *Continue work on the ESMF for climate model development and extend it to all domains*
- *Explore the use of specific types of grid technology*
- *Develop a common weather research computing environment*
- *Apply the WRF modeling system standards and framework to NCEP's Mesoscale Modeling Systems, including the continental domain Meso Eta and nested Nonhydrostatic Meso Threats runs*

NIST

NIST works with industry and with educational and government organizations to make IT systems more usable, secure, scalable, and interoperable; to apply IT in specialized areas such as manufacturing and biotechnology; and to encourage private-sector companies to accelerate development of IT innovations. NIST also conducts fundamental research that facilitates measurement, testing, and the adoption of industry standards.

NIST HEC I&A

Activities and accomplishments include:

- Interoperable Message Passing Interface (IMPI) to support parallel computing research. The standard enables interoperability among implementations of Message Passing Interface (MPI) with no change to user source code.
- Java/Jini/Javaspaces-based screen-saver science system. Distributed-computing environment leverages java's portability to use otherwise unused cycles on any set of computing resources accessible over a local network (including desktop PCs, scientific workstations, and cluster compute nodes) to run any compute-intensive distributed Java program.
- Immersive visualization infrastructure for exploring scientific data sets drawn from both computer and laboratory experiments. Visualization projects include a three-wall Immersive Virtual Reality System that uses the DIVERSE open source software and locally built visualization software to display results at scale. The local software includes the glyph toolbox and techniques for immersive volume visualization, surface visualization, and interaction. All the software runs unchanged on Linux machines.
- Fundamental mathematical software tools to enable high-end computing applications. Examples include:
  - – PHAML, a parallel adaptive grid refinement multigrid code for PDEs
  - – Sparse Basic Linear Algebra Subroutines (BLAS)
  - – Template Numerical Toolkit (TNT) for object-oriented linear algebra
- HEC applications combining parallel computing with immersive visualization for basic sciences, such as physics, and other uses, such as building structure and material strength. Typical applications to which NIST applies these resources are:
  - – Nanostructure modeling and visualization (especially in nano-optics)
  - – Modeling cement and concrete, including modeling the flow of suspensions and fluid flow in complex geometries
  - – Computation of atomic properties
  - – Visualization of "smart gels" (which respond to specific physical properties such as temperature or pressure) to gain insight into the gelling mechanism
  - – Visualization of tissue engineering to study and optimize the growth of cells on scaffolding
- Applications-oriented problem-solving environments, such as Object-Oriented Finite Element Modeling of Material Microstructures (OOF) for modeling materials with complex microstructure and Object-Oriented MicroMagnetic Computing Framework (OOMMF) for micromagnetics modeling

NIST HEC R&D

Activities and accomplishments include:

- Research in quantum computing and secure quantum communications, as well as in measurement science aspects of nanotechnology, photonics, optoelectronics, and new chip designs and fabrication methods. Theoretical and experimental investigations include:.
  - – Algorithms and architectures for scalable quantum computing
  - – Demonstrations of quantum computing within specific physical systems such as ion traps, neutral atoms, and Josephson junctions
  - – High-speed testbed for quantum key distribution
  - – Single photon sources and detectors
  - – Quantum protocols

FY 2005 NIST plans include:

HEC I&A

- *Continue work in activities listed for FY 2004*
- *Pursue additional project in "Virtual Laboratories"*

HEC R&D

- *Continue work in activities listed for FY 2004*

HEC EPA HEC

EPA's mission is to protect human health and safeguard the environment through research, regulation, cooperation with state governments and industry, and enforcement. Areas of interest extend from groundwater to the stratosphere.

EPA HEC I&A programs are focused on tools to facilitate sound science using high-end computation, storage, and analysis. These programs enable relevant, high-quality, cutting-edge research in human health, ecology, pollution control and prevention, economics, and decision sciences. This facilitates appropriate characterization of scientific findings in the decision process. The HEC programs are performed in-house and as problem-driven research.

EPA is launching the Center of Excellence (CoE) for Environmental Computational Science to integrate cutting-edge science and emerging IT solutions to facilitate Federal and state-level partnerships and enhance the availability of scientific tools and data for environmental decision making. The CoE will enable collaboration from within and without the organization and will provide a flexible, dynamic computing and information infrastructure to ensure optimized yet secure access to EPA resources.

FY 2004 HEC programs include:

Multimedia Assessments and Applications (MAA) Framework – provides a foundation for research on how to structure compartmental modules and improve model integration and interchangeability. The MAA framework's objective is to provide software that supports composing, configuring, applying, linking, and evaluating complex systems of object-oriented models. It will improve EPA's ability to simulate the interaction between individual environmental media components (e.g., chemical fluxes, water cycle) and will enable distributed computation.

The MAA framework is tailored to multimedia models but is adaptable and generalized. It supports EPA programs such as the Chesapeake Bay Program, the Tampa Bay Region Atmospheric Chemistry Experiment, and the Office of Air Quality Planning and Standards. The framework is currently being tested by a number of clients.

Uncertainty Analysis (UA) Framework Development – developing tools to support the analysis of model sensitivities and the effects of input uncertainties on model predictions. Specific tasks of this EPA work are to:

- Construct a 400-node Intel-based supercomputing cluster called SuperMUSE, for Supercomputer for Model Uncertainty and Sensitivity Evaluation
- Develop platform-independent systems software for managing SuperMUSE
- Conduct uncertainty and sensitivity analyses of the Multimedia, Multipathway, Multireceptor Risk Assessment (3MRA) modeling system
- Develop advanced algorithmic software for advanced statistical sampling methods and global sensitivity analyses

Air Quality Modeling Applications (AQMA) – program aims to advance computational performance of the state-of-the-art Community Multiscale Air Quality (CMAQ) Chemistry-Transport Modeling System while maintaining modularity, portability, and single-source code. Efforts to improve CMAQ take into account both the typical Linux cluster in the States and also HEC deployments. Major areas of effort include algorithmic improvement, microprocessor tuning, and architecture assessment. Involves a phased deployment that enables the states, which are the key stakeholders, to participate in the development.

Grid deployment – goal is to provide phased deployment of an EPA-wide enterprise grid that will identify, develop, and integrate key technologies, align organizational policies such as security and networking, and field grid pilot systems to demonstrate success and benefits. Historically, agency researchers with high-end applications competed for time on EPA's high-performance computing resources located at the National Environmental Scientific Computing Center (NESC2). With the implementation of grid middleware, researchers will be able to tap unused processing capacity on local and remote clusters at the campus level or enterprise level. EPA's compute grid is being implemented in a phased approach with parallel development of both grid infrastructure and security policy. Pilot clusters have now been linked to demonstrate EPA compute grid services both internally and also to external partners.Ultimately, EPA researchers and trusted partners will be able to access a partner (or global) grid extending to organizations outside the agency.

Grid demonstration projects – part of EPA effort to combine state-of-the-science air quality modeling and observations to enable timely communication of meaningful environmental information, improve emission management decisions, and track progress in achieving air quality and public health goals. Technical collaborators in demonstration project on air quality include DOE/Sandia and NOAA; pilot partners include the state of New York and regional planning organizations. Phase I of the project includes delivery of an optimized CMAQ model and data sets to the client community in summer 2004 and eventually the running of CMAQ over the grid and at client sites.

In FY 2005, HEC I&A work will continue on the following EPA efforts:

- *EPA Center of Excellence for Environmental Computational Science (CoE)*
  - *Grid deployment and grid demonstration projects*
- *Air Quality Modeling Applications (AQMA)*

PARTICIPATING AGENCY

HEC DoD/HPCMPO HEC

The mission of DoD's HPCMPO is to deliver world-class commercial, high-end, high-performance computational capability to DoD's science and technology (S&T) and test and evaluation (T&E) communities, thereby facilitating the rapid application of advanced technology in superior warfighting capabilities. Development of future technologies supported by HPC includes: Micro Air Vehicles; Joint Strike Fighter; surveillance systems; smart weapons design; ocean modeling; parachute simulations; Unmanned Air Vehicle; and blast protection.

HPCMPO requirements are categorized in the following ten key Computational Technology Areas (CTAs):

- Computational Structural Mechanics (CSM)
- Computational Fluid Dynamics (CFD)
- Computational Chemistry and Materials Science (CCM)
- Computational Electromagnetics and Acoustics (CEA)
- Climate/Weather/Ocean Modeling and Simulation
- Signal/Image Processing (SIP)
- Forces Modeling and Simulation (FMS)
- Environmental Quality Modeling and Simulation (EQM)
- Computational Electronics and Nanoelectronics (CEN)
- Integrated Modeling and Test Environments (IMT)
- Other

HEC I&A

High Performance Computing Centers – include four major shared resource centers (MSRCs) and 16 distributed centers. MSRCs provide complete networked HPC environments, HPC compute engines, high-end data analysis and scientific visualization support, massive hierarchical storage, and proactive and in-depth user support and computational technology area expertise to nationwide user communities. The MSRCs are the Army Research Laboratory (ARL), Army Signal Command (ASC), ERDC, and NAVO.

Software Applications Support (SAS) – encompasses the following four components:

- Common HPC Software Support Initiative (CHSSI)
- HPC Software Applications Institutes (HSAIs)
- Programming Environment and Training (PET) program to enhance user productivity