

Trends in High-End Computing in United States Government

MINORU NOMURA

Information and Communications Research Unit



1 Introduction

Currently, United States Government research and development on information and communication technologies are conducted based on the Networking and Information Technology Research and Development (NITRD) program coordinated by the National Science and Technology Council (NSTC). Twelve federal agencies participate in the NITRD program consisting of seven Program Component Areas(PCAs). The PCAs related to High-end Computing (HEC) are High-end Computing Infrastructure and Applications (HEC I&A) and High-end Computing Research and Development (HEC R&D)^[1, 2].

In March 2003, the HEC Revitalization Task Force(“the Task Force”) was formed under NSTC as a special project of NITRD. The co-chairs are members from DoD/ODDR&E, the DOE/Office of Science, the National Coordination Office, and the Office of Science and Technology Policy. Other participating agencies include DARPA, DoD/HPCMP, DoD/Missile Defense Agency, DOE/NNSA, the EPA, NASA, the NIST, the NSA, the NSF, the OMB, and so on. The Task Force’s report lists 70 names.

The mission of the Task Force is to develop a plan for undertaking and sustaining a robust Federal high-end computing program to maintain US leadership in science and technology fields into the future. In May 2004, the Task Force released the Federal Plan for High-End Computing (“the HEC Plan”)^[3], which includes measures on HEC research and development,

HEC resources, and HEC system procurement over the coming 5 to 10 years. HEC is essential for science and technology development, national security, and international competitiveness. However, there is an awareness that the HEC (resources, architectures, and software tools and environments) used for government missions are not always meeting the computing needs of federal agencies.

The 108th US Congress debated at least three bills related to the HEC Plan during 2004. Of these, Congress passed the Department of Energy High-end Computing Revitalization Act of 2004 in November. Currently, the US Government strongly pushes HEC research and development, and its utilization as a strategy to increase national power.

The purpose of this article is first, to provide an outline of the HEC Plan in Section 2, and second, to discuss its most significant points in Section 3.

2 Outline of the “Federal Plan for High-end Computing”^[4]

In this section, the “Federal Plan for High-End Computing ” will be summarized.

The Task Force solicited input from leading applications scientists who utilize HEC to advance their research in various specialist fields. According to the survey, the estimates of additional capability needed to achieve the goals ranged from 100 to 1,000 times the current capability of today’s HEC resources. Table 1 shows “Science Challenges” and “Potential Outcomes with 100 to 1,000 Times Current Capability.”

Table 1 : Benefits of HEC to science and engineering

Area	Application	Science Challenge	Potential Outcome with 100 to 1,000 Times Current Capability
Physics	Astrophysics	Simulation of astrophysical environments such as stellar interiors and supernovae.	Yield understanding of the conditions leading to the origin of the heavy elements in the universe.
	High-Energy Physics	Achieve a detailed understanding of the effects of strong nuclear interactions so that the validity of the Standard Model can be tested to determine whether physics beyond the Standard Model occurs at extreme sub-nuclear distances.	Guide experiments to identify transition from quantum chromodynamics to quark-gluon plasma.
	Accelerator Physics	Accurate simulations of the performance of particle accelerators.	Optimize the design, technology, and cost of future accelerators, and use existing accelerators more effectively and efficiently.
	Nuclear Physics	Realistic simulations of the characteristics of the quark-gluon plasma.	By developing a quantitative understanding of the behavior of this new phase of nuclear matter, facilitate its experimental discovery in heavy ion collisions.
Nano-science	Catalyst Science / Nanoscale Science and Technology	Calculations of homogeneous and heterogeneous catalyst models in solution.	Reduce energy costs and emissions associated with chemicals manufacturing and processing. Meet Federally mandated NOx levels in automotive emissions.
	Nanoscale Science and Technology	Simulate the operation of nanoscale electronic devices of modest complexity.	Take miniaturization of electronic devices to a qualitatively new level enabling faster computers, drug delivery systems, and consumer and military electronics.
	Nanoscale Science and Technology	Simulate and predict mechanical and magnetic properties of simple nanostructured materials.	Enable the discovery and design of new advanced materials for a wide variety of applications potentially impacting a wide range of industries.
Aerospace	Simulation of Aerospace Vehicle in Flight	Simulate a full aerospace vehicle mission, such as a full aircraft in maneuver or an RLV in ascent or descent.	Reduce aerospace vehicle development time and improve performance, safety, and reliability.
	Full Liquid Rocket Engine Subsystems Simulation	Simulate full rocket engine subsystems during ascent including turbopump and combustion devices.	Provide capability for risk assessment during Earth-to-orbit and improve safety and reliability of space transportation systems.
	Aviation Systems Simulation	Execute high-fidelity airspace simulations and develop decision system and management tools for terminal area.	Provide capability for effectively managing national airspace and increase safety in terminal area.
Life Sciences	Structural and Systems Biology	Simulations of enzyme catalysis, protein folding, and transport of ions through cell membranes.	Provide ability to discover, design, and test pharmaceuticals for specific targets and to design and produce hydrogen and other energy feedstock more efficiently.
	Signal Transduction Pathways	Develop atomic-level computational models and simulations of complex biomolecules to explain and predict cell signal pathways and their disrupters.	Yield understanding of initiation of cancer and other diseases and their treatments on a molecular level, and the prediction of changes in the ability of microorganisms to influence natural biogeochemical cycles such as carbon cycling and global change.
National Security	Signals Intelligence	Model, simulate, and exploit foreign codes, ciphers, and complex communications systems.	Support U.S. policymakers, military commands, and combat forces with information critical to national security, force protection, and combat operations.
	Directed Energy	Advance the directed energy systems design process out of the scientific research realm into the engineering design realm.	Efficiently design next-generation directed energy offensive and defensive weapon systems. Change the design process from years to days.
	Signal & Image Processing & Automatic Target Recognition	Replace electromagnetic scattering field tests of actual targets with numerical simulations of virtual targets.	Design more stealthy aircraft, ships, and ground systems and create the ability to rapidly model new targets, enabling more rapid adaptation of fielded weapon systems' ability to target new enemy weapon systems.
	Integrated Modeling and Test of Weapon Systems	Model complex system interaction in real time with precision.	Replace many expensive, dangerous, and timeconsuming ground tests with virtual tests resulting in lower test costs and more rapid development of weapon systems.
Earth and Atmospheric Sciences	Climate Science	Resolve additional physical processes such as ocean eddies, land use patterns, and clouds in climate and weather prediction models.	Provide U.S. policymakers with leading-edge scientific data to support policy decisions. Improve understanding of climate change mechanisms and reduce uncertainty in the projections of climate change.
	Weather and Short-term Climate Prediction	Enable dynamical prediction of frequency and intensity of occurrence of hurricanes/typhoons and severe winter storms 90 days in advance.	Provide critical support to deployed naval, air, and land forces in local, regional, and global combat environments. Lives saved and economic losses avoided due to better severe weather prediction.
	Solid Earth Science	Improved statistical forecasting of earthquake hazards (fault-rupture probabilities and ground motion).	Provide prioritized retrofit strategies. Reduced loss of life and property. Damage mitigation.
	Space Science	Realistically simulate explosive events on the sun, the propagation of the energy and particles released in the event through the interplanetary medium, and their coupling to Earth's magnetosphere, ionosphere, and thermosphere.	Provide decision makers (both civilian and military) with status and accurate predictions of space weather events on time scales of hours to days.
Energy and Environment	Subsurface Contamination Science	Simulate the fate and transport of radionuclides and organic contaminants in the subsurface.	Predict contaminant movement in soils and groundwater and provide a basis for developing innovative technologies to remediate contaminated soils and groundwater.
	Magnetic Fusion Energy	Optimize balance between self-heating of plasma and heat leakage caused by electromagnetic turbulence.	Support U.S. decisions about future international fusion collaborations. Integrated simulations of burning plasma crucial for qualifying prospects for commercial fusion.
	Combustion Science	Understand interactions between combustion and turbulent fluctuations in burning fluid.	Understand detonation dynamics (for example, engine knock) in combustion systems. Solve the "soot" problem in diesel engines.

2-1 HEC: A strategic tool for science and technology leadership

(1) The case for HEC revitalization

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, 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.”

The HEC Plan states that this necessitates revitalization.

(2) Goals

- Make high-end computing easier and more productive to use
- Foster the development and innovation of new generations of high-end computing systems and technologies
- Effectively manage and coordinate federal high-end computing
- Make high-end computing readily available to federal agencies that need it to fulfill their missions

The HEC Plan states that in the course of making high-end computing easier and more productive to use, the most important thing for researchers is to minimize the time to solution from new idea to results. It also states that this should be the goal of research and development

in HEC systems. Figure 1 illustrates the elements of time to solution. Overall optimization requires the minimization not only of the calculations but also of each phase.

(3) Scope of the plan*1

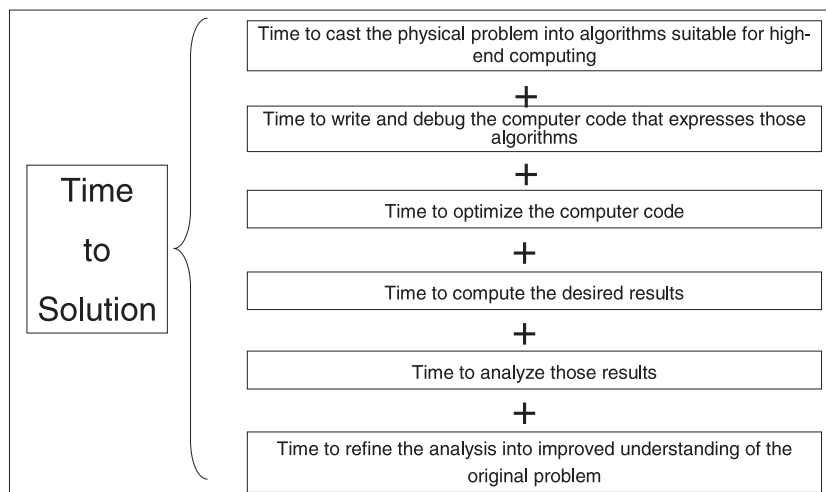
The HEC Plan includes a number of roadmaps outlining all the core technologies needed for high-end computers that might be manufactured within approximately 15 years. Key elements include:

- Core technology research and development in the hardware, software, and system technologies
- Capability, capacity, and accessibility strategies to assure that high-end computing resources are readily available to the science and engineering communities that need them
- Efficient procurement strategies that provide high-end computers that meet user requirements

Notes

*1 Visualization, networking, grid computing, general security issues, and application-specific software were considered outside the scope of this planning effort. Procurements of small-scale cluster systems also were not included in this planning activity.

Figure 1 : Time to solution



Source: Author’s compilation based on reference^[4]

Within the NITRD program, the annual HEC-related budget is approximately \$900 million, of which the activities considered by the Task Force represent about \$158 million (FY 2004). The Plan states that if these revitalization activities succeed, they will have a positive impact on the long-term activities of the entire \$2.6-billion government portfolio for high-end computing.

2-2 *Research and development*

The HEC plan states as follows.

In the R&D area, a gap between federal requirements on computational performance and requirements for commercial systems can be seen. Compared with the business computing market and the web-based commerce market, the HEC market is not large enough to divert computer industry attention. HEC procurements are approximately \$1 billion per year, while the server market by comparison is over \$50 billion per year. This is why industry concentrates on the server market, and the HEC systems it provides consist of very large collections of processors designed for the smaller systems required by that market. These massive multiprocessor systems have proven exceptionally difficult to program and achieving high levels of performance for some important classes of applications has been problematic.

Recently, processor performance is continually improving, and theoretical peak performance is rapidly rising. In multiprocessor systems, however, the increasing disparity between processor speed and memory bandwidth is constraining actual performance in real operating environments. Processor speed is growing approximately 40 percent per year, while memory speed is improving approximately 7 percent annually.

The HEC Plan states that the cluster-based systems on which recent HEC investment focuses in the USA are not well suited for all applications, and different architectures would be more suitable for some high-priority government applications. The HEC Plan also describes parallel efficiency as follows: "The current HEC focus on clustering hundreds of small nodes, each with a separate OS, results in poor parallel efficiency,

generally below 10% and sometimes lower than 1% of the peak on some applications."

Figure 2 shows the divergence between theoretical peak performance and sustained system performance (SSP) observed in major HEC centers.

(1) User requirements for HEC technology

The HEC Plan identifies the following primary challenges for effective use of high-end computing:

- Achieving high sustained performance on complex applications
- Building and maintaining complex software applications
- Managing dramatically increasing volumes of data, both input and output
- Integrating multiscale (space and time), multidisciplinary simulations

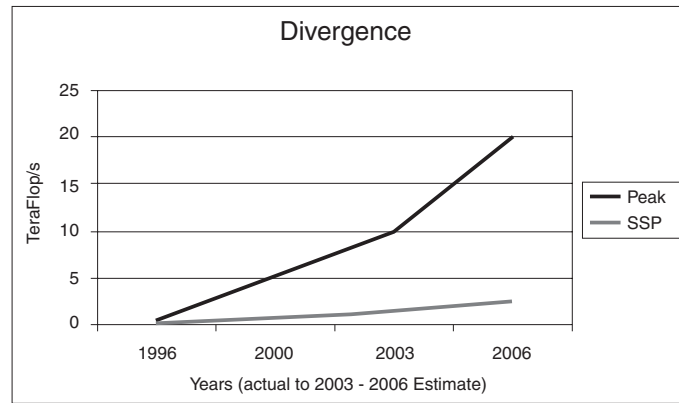
In addition, the HEC Plan identifies the following goals for future high-end computing systems:

- A 100-fold increase in sustained (as opposed to peak) performance (a level of performance required to solve a number of current scientific and technological problems)
- Ultra-fast processors and new algorithms, since not all problems can be easily parallelized
- Improvements in bandwidth and latency for both memory and communications fabric, which for many applications largely determine performance
- Architectures that can meet diverse application requirements

The HEC Plan also emphasizes the lack of software tools, programming models, and operating systems. It is stated that one could expect reasonable performance on up to 1,000 processors, but that one could not expect performance from systems of 100,000 processors (projected in the 2010 timeframe) without substantial improvements.

(2) HEC R&D Strategy

In order to respond to user requirements

Figure 2 : Divergence problem between theoretical performance and sustained performance

Peak : Theoretical peak performance
 SSP : Sustained system performance

in hardware, software and systems, the HEC Plan indicates (i) roadmaps, (ii) research and evaluation systems, and (iii) prioritization of HEC R&D investments for key technologies. They are outlined below.

(i) Roadmaps

The roadmaps include hardware, software, and systems. The roadmaps show two scenarios for the next decade.

First, the “current program” assumes no resource allocation changes from FY 2004.

Second, the “robust R&D program” indicates the probable scenario if new HEC system plans, execution, and system deployment are implemented in a timely manner.

Details are shown in Tables 2 through 7.

[Hardware Roadmap]

The “current program” and the “robust R&D program” are described in Table 2 and 3.

Under the “current program,” without additional research effort, there will probably be little progress beyond the next five years. Such improvement would depend primarily on industry-driven commercial-off-the-shelf (COTS) technology advances and the results of existing or past research investments. Furthermore, without significant technological breakthroughs, Moore’s Law will be coming to an end in the 2015 timeframe.

The “robust R&D plan” is a measure to push beyond the “current program.”

[Software Roadmap]

The “current program” and the “robust R&D

program” are described in Table 4 and 5.

The “current program” scenario depends on the Defense Advanced Research Projects Agency’s (DARPA) High-Productivity Computing Systems (HPCS) program for the release of new architectures in the next five years. Since the DARPA program ends in 2010, future improvement through the “current program” will be based mainly on those architectures. The “robust R&D program” is a measure to push beyond that scenario.

[System Roadmap]

The “current program” and the “robust R&D program” are described in Table 6 and 7.

The “current program” scenario is dependent upon existing research activities (including HPCS) and progress after the next five years will be difficult. The “robust R&D program” is a measure to push beyond that scenario.

(ii) Research and evaluation systems

Because proper development and evaluation are necessary for future large-scale systems with 10,000 to 100,000 processors to function, the Task Force recommends the procurement of research and evaluation systems as an essential HEC R&D strategy.

The “early access” systems called as research and evaluation systems by the Task Force enable early prototype testing and provide platforms necessary for the development of new algorithms and computational techniques. In addition, such systems are essential for the evaluation of the functionality and scalability of software. During software development, testing often

Table 2 : Hardware roadmap: Current program

	Near Term (within a year)	Mid Term(within five years)	Long-Term (within ten years)
Microarchitecture	COTS-driven microarchitecture	Multi-CPU cores per chip, memory bandwidth per CPU decreases	Moore's law end?
Interconnect technologies	Interconnect technology based upon electrical interconnect and electrical switches	Interconnect technology based upon electro-optical interconnect and electrical switches	Interconnect technology driven by telecom - expect moderate advances for HEC systems
Memory	Processor/memory performance gap addressed by caches, limits performance and ease of programming	Early COTS PIM-based and streaming technologies to address processor/memory gap	Evolutionary improvements; increased use of PIMs
Power, cooling, and packaging	Thermal/packaging–chip/system technologies limited by our ability to cool via air	Evolutionary improvements do not significantly advance our ability to develop high-end systems	System performance limited by “thermal wall”?
I/O and storage	I/O driven by COTS-based needs in areas of storage and links	Petaflop-scale file systems based upon COTS technologies, RAS issues will limit usability	Depends upon 3-D storage

Table 3 : Hardware roadmap: Robust R&D plan

	Near-to Mid-Term (within five years)	Long-Term (within ten years)
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”

PIM : Processor-In-Memory, RAS : Reliability, Availability, Serviceability
 COTS : Commercial-Off-The-Shelf

Table 4 : Software roadmap: Current program

	Near-Term (within a year)	Mid-Term(within five years)	Long-Term (within ten years)
Operating systems (OSs)	OSs adapted from desktops or servers. Fragile and do not scale over 1,024 processors	Early introduction of OSs that scale to 10,000 processors for at most two HPCS system architectures. Clusters remain a challenge.	Little progress is expected.
Languages, compilers, and libraries	Legacy languages and libraries (for example, Fortran, C, C++, and MPI). Compiler technology inadequate for achieving scalable parallelism.	Limited production quality compilers (for example, UPC and Co-Array Fortran [CAF]) for a few systems. MPI continues to dominate. Heroic programming required for computations on over 2,048 processors.	Limited additional improvements in programmability. Production-quality compilers for UPC and CAF widely available. Mostly incremental progress with compiler optimization and MPI implementation. No revolutionary advances in languages
Software tools and development environments	Wide variety of vendor specific or research-quality tools – limited integration, difficult to use, and little portability. No integrated development environments (IDEs) available for HEC systems.	Tool capability lags HEC systems (for example, debugging 250,000-processor jobs). IDE support for small-sscale (32-processor) systems only.	Gap between tool capabilities and ability to understand large systems widens. IDE support for mid-range shared memory systems
Algorithms	Efficient parallel algorithms for some problems (for example, dense linear algebra). Others require deep expert knowledge for efficient implementation.	Improved parallel algorithms for unstructured and sparse problems	Additional progress in mapping algorithms onto advanced architectures

Table 5 : Software roadmap: Robust R&D plan

	Near-to Mid-Term (within five years)	Long-Term (within ten years)
Operating systems (OSs)	New research-quality HEC OSs that address scalability and reliability	Production-quality, faulttolerant, 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 realistic simulations of physical systems.

Table 6 : System roadmap: Current program

	Near-Term (within a year)	Mid-Term (within five years)	Long-Term (within ten years)
System architecture	COTS-based systems from 10 to 100 Tflops peak (1,000 to 10,000 processors) with server-class operating systems – fragile and hard to program	At most two DARPA HPCS systems capable of sustained petaflops (up to 100,000 processors or more) on selected mission applications	Evolutionary improvements only beyond HPCS systems
System modeling and performance analysis	System modeling and performance analysis tools developed but ad hoc, incomplete, difficult to use, and not integrated	Accuracy improvements in models/tools for legacy systems and applications for use by experts. Modeling of HPCS systems faces complexity challenges.	Evolutionary improvements toward ease of use and integration with system
Programming models	Legacy parallel computing models limit ease of programming. Main model is message passing. “Non-heroic” programming practice: MPI at 64 to 256 and OpenMP at 16 to 128.	Minor progress in parallel computing models. “Non-heroic” programming: MPI-2 feasible for 128 to 512 processors and DSM implementations (UPC, CAF, ...) more widespread and available for 64 to 256 processors.	Incomplete implementation and acceptance of shared memory programming models (for example, UPC and CAF)
Reliability, availability, and serviceability (RAS) + Security	RAS achieved by defensive user actions (for example, checkpoint/restart) and rescheduling	Limited RAS solutions for up to 1,024-processor systems. Partial fault isolation and better profiling of user behavior to prevent inside attack.	RAS solutions for up to 10,000-processor systems. Some improvements in applications security

Table 7 : System roadmap: Robust R&D plan

	Near-to Mid-Term (within five years)	Long-Term (within ten years)
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 and prediction of software behavior. Automated and intelligent performance and analysis tools and benchmarks widely available and 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.

CAF: Co-Array Fortran, COTS: Commercial-Off-The-Shelf, DARPA: Defense Advanced Research Projects Agency, DSM: Distributed Shared Memory, HPCS: High Productivity Computing Systems, IDE: Integrated Development Environment, MPI: Message Passing Interface, OpenMP: Open specification for MultiProcessing, OS: Operating System, RAS: Reliability, Availability, Serviceability, UPC: Unified Parallel C

Figure 3 : Recommended priorities

		Current Program*	Increment compared to HEC R&D Current Program				
		FY 2004 (\$ in millions)	FY 2006	FY 2007	FY 2008	FY 2009	FY 2010
Hardware	a. Basic and Applied Research	\$5					
	b. Advanced Development	\$5					
	c. Engineering and Prototypes	\$0					
	d. Test and Evaluation	\$2					
Software	a. Basic and Applied Research	\$33					
	b. Advanced Development	\$21					
	c. Engineering and Prototypes	\$15					
	d. Test and Evaluation	\$2					
	e. Long-term Evolution and Support	\$0					
Systems	a. Basic and Applied Research	\$4					
	b. Advanced Development	\$40					
	c. Engineering and Prototypes	\$1					
	d. Test and Evaluation	\$30					
Total		\$158**					



* Assumes no planning changes from FY 2004.
 ** This total represents the aggregate investment across all agencies in HEC as defined in the scope of the plan section of the reference^[4].
 Source: Author's compilation based on reference^[4].

causes hardware breakdowns, interfering with application development. It is therefore necessary to separate software development testbeds from application development testbeds.

The performance information gained from extensive evaluations of research and evaluation systems is invaluable for the successful future procurement of HEC systems. If these evaluations have been able to identify failed approaches, the government will not acquire systems that do not perform as expected. In addition, such evaluations may also suggest more fruitful approaches through removal of the sources of failure.

(iii) Prioritization of HEC R&D investments

The HEC Plan examines prioritization after defining the four major stages in research and development.

- (a) Basic and Applied Research: Focus on the development of fundamental concepts in high-end computing through the continuous creation of new ideas and expertise.
- (b) Advanced Development: Select and refine

innovative technologies and architectures for potential integration into high-end systems.

- (c) Engineering and Prototypes: Perform the integration and engineering required to build HEC systems and components.
- (d) Test and Evaluation: Conduct testing and evaluation of HEC software as well as the current and new generations of HEC systems at appropriate scale.

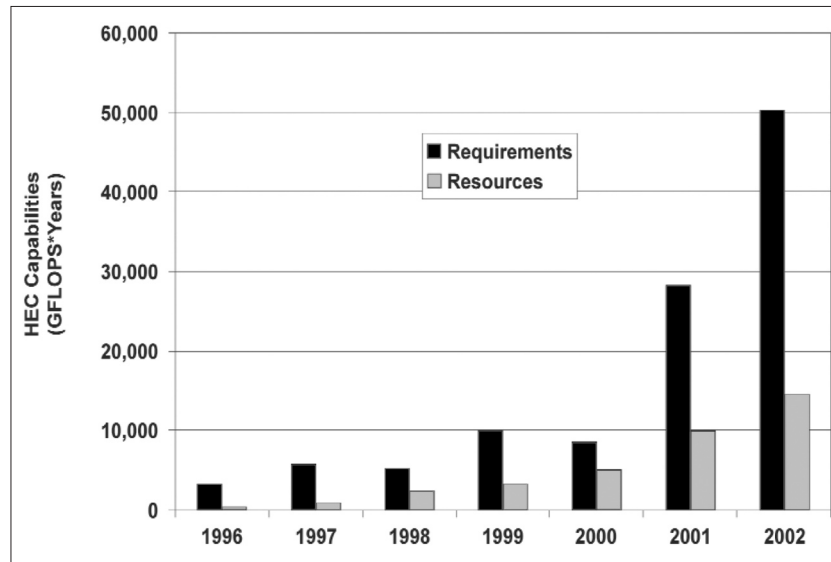
The HEC Plan also suggests that for long-term evolution and support, the government should maintain critical HEC software infrastructure over the long term.

Recommendations for R&D investment for each stage are described as shown in Figure 3. The Chart depicts the prioritization of each increment compared to the “current program.”

2-3 HEC resources

The Plan defines “HEC resources” as the acquisition, operation, and maintenance of HEC systems needed to carry out federal agency mission applications.

The HEC Plan indicates that since overall

Figure 4 : HEC requirements vs. available resources for the DoD High-Performance Computing Modernization Program

computing ability is not sufficient, robust investment in HEC resources is required. Some federal agencies borrow resources from other agencies because they cannot provide their users with sufficient computing capacity. The Plan also states that no civilian agency in the USA currently has access to leadership-class systems to provide true breakthrough capability for important computational problems.

(1) User requirements for HEC resources

Surveying the HEC requirements of a broad range of scientific disciplines across the federal government identified two classes of resource issues. The first is architectural availability, and the second is acquisition of HEC capacity. These issues are discussed in order below.

(i) Architectural availability

Today's HEC market is not producing products that satisfy the performance requirements of the most demanding scientific applications. Vendors provide excellent computers where commercial computing needs overlap with scientific needs. However, where scientific or defense needs do not overlap commercial IT needs, the products are insufficient.

(ii) Acquisition of HEC capacity

Federal need for HEC in science and technology is approximately triple current capacity and grows by about 80 percent annually. This trend of demand will strengthen as advanced application

usage and areas of application expand (Figure 4).

(2) Addressing HEC access, availability, and leadership

The Task Force proposes separate approaches to address the three distinct problems of accessibility, availability and leadership systems in HEC resources.

(i) Accessibility

Addressing the sharing of HEC resources

- Federal agencies whose researchers currently obtain HEC resources from other agencies should examine options for providing resources to users through cooperative agreements.
- Each federal agency should assess and make arrangements to provide for its own resource needs based on mission priority.

(ii) Availability

The Task Force calls for an increase in resources needed for mission execution.

- Federal agencies should examine the value of reallocating resources to cope with increasing demand for computing resources and increasing demand for already overburdened systems.
- Assessment and adjustment of the relative balance among research and engineering modes (theory, experiment, and computation) is needed for optimal resource reallocation.

(iii) Leadership systems

The Task Force proposes the development of so-called leadership systems in order to provide US science researchers with the world-leading HEC capabilities.

The goal of such systems is to provide computing capability at least 100 times greater than that currently available on the market. A limited set of scientific applications (perhaps 10 per year) would be selected and given substantial access to such systems. Much smaller time allocations could be available for a wider set of applications (perhaps 50 per year) for pilot experiments in preparation for full-scale runs in the future. By nature, Leadership Systems could be productive for several years, but they would need regular replacement with new leadership systems based on scientific needs and technologies emerging from research and development activities. The HEC Plan states that the results of core technical HEC R&D would be utilized for HEC systems at first, but over time those technologies could be applied to servers and finally to desktops. The Plan's recommendations are as follows.

- Provide leadership systems with leading-edge computing capability for highest-priority research problems.
- Federal agencies should manage leadership systems as national resources.
- Federal agencies should operate leadership systems as an open user facilities.
- Access to the system should be governed by a peer review process.

2-4 Procurement

Procurement of HEC systems is a very complex task. It thus requires approaches that reduce the burden on both the government and vendors. Ten years ago, it was common for an HEC system to have a service life of more than five years, but now average life span is about three to four years, necessitating shorter procurement periods.

The Plan proposes three interagency pilot projects (HEC benchmarks, TCO (total cost of ownership), procurement) to improve the efficiency of Federal HEC procurement practices.

A description of each project follows.

(i) HEC benchmarking pilot project

Sustained system performance is currently the only acceptable performance criterion for measuring procurement selection decisions. Other performance indicators, such as calculated peak performance and performance on a single benchmark such as LINPACK may be useful, but they should not be used as the basis for acquisition decisions. The HEC Plan concludes that benchmark performance on actual applications is the best indicator of a system's performance in an operational environment and makes the following recommendations.

- Selected agencies with similar HEC applications will develop a single suite of benchmarks based on their applications. This benchmark suite will be used at the pilot acquisition stage.
- Participating agencies use the benchmarking results, suitably weighted for their individual applications, instead of agency-specific benchmarks.

(ii) TCO pilot project

TCO includes all the financial aspects of providing HEC services, and comprises the following four major cost areas.

- Hardware
- System software
- Space, utilities, personnel, and extra-center communications (networking)
- User productivity (including costs of application software development)

The HEC Plan makes the following recommendations for this pilot project.

- A multiagency team would evaluate all elements of TCO (e.g., acquisition and maintenance, personnel, extra-center communications, and user productivity) across several similar systems and develop best practices for determining TCO.

(iii) Collaborative multiagency HEC procurement pilot project

Applying new techniques developed from the above two projects, participating agencies will develop a common method for procurement.

They will then evaluate its effectiveness. Evaluation criteria will include improved buying power, reduced overall labor costs, total procurement time, and ability to meet the needs of the participating agencies.

3 | Points of interest

Above I have outlined the “Federal Plan for High-End Computing.” The HEC program has many points of interest. I will describe some of them below along with relevant technology trends.

(1) Minimizing the time to solution

The HEC Plan frequently uses the term “time to solution.” This term indicates the time required for a researcher to obtain computation results, including the program development and testing periods as well as actual computation time. The Plan emphasizes “time to solution” and proposes making it the measure of the evolution of HEC. In addition, the Plan addresses time to solution is an important factor influencing costs over the whole life cycle of a HEC system, making it a baseline for all HEC activities, including HEC R&D, HEC resources, and HEC procurement.

(2) Emphasizing sustained system performance^{*2}

Regarding sustained system performance, the Task Force has discussed it in detail from the perspective of awareness of HEC systems, optimal

Note

^{*2} Information related to sustained system performance: Issues concerning the sustained system performance of US HEC systems are also discussed in a report of the National Research Council (NRC)^[8] written by US academic researchers. In addition, reference^[9] describes Japan’s Earth Simulator, which achieves high sustained system performance. See the references for details. (The content is not included in the HEC Plan, but is shown here for reference.)

R&D, and optimal procurement, and the HEC Plan strongly stresses it. The roadmap suggests improvement measures.

(3) Prioritization of R&D

The Plan’s recommendations for prioritization in HEC R&D investment are meaningful in that it demonstrates areas of emphasis. The FY 2004 budget related to the HEC Plan shows the relative weights of R&D areas. It is noteworthy that each total for software and for systems is greater than that for hardware. It is also noteworthy that the HEC Plan recommends many increases from the early stages, for “basic and applied research” and “advanced development” in hardware, for “engineering and prototypes”, “test and evaluation” and “long-term evolution and support” in software, and for “engineering and prototypes” in systems.

(4) Resource allocation for large, challenging problems: leadership systems

The HEC Plan states that leadership system facilities must be installed in response to large, challenging research problems that require the highest performance, and that they should be made accessible to researchers both from industry and from federal agencies. It touches on the fact that currently, in the USA as well, civilian agencies do not have access to leadership-class systems. This can be seen as urging improvement. In addition, the Plan describes the spinoffs from leadership system development as a goal. The Plan seems strongly aware that although at first the results of core technical R&D activities in HEC will be limited to HEC systems with federal missions, eventually those technologies will be applied to commercial products such as servers and finally desktops.

(5) Increased access to HEC

Regarding increased access to HEC, the HEC Plan reports rapidly increasing use in the National Institutes of Health (NIH) and increased access in industrial fields such as chemical, semiconductor, and materials sectors, where obtaining necessary data through experiments is difficult, time-consuming and/or expensive.

It is noteworthy that the Plan, at the front,

Table 8 : Comprehensive approach to R&D

Activity	Purpose	Performers
Basic and Applied Research	Refill the academic pipeline with new ideas and people	Academia and government labs
Advanced Development	Develop component and subsystem technologies	Mostly industry led, partnering with academia and government labs
Engineering and Prototypes	Integration at system level and development of Serial No. 1	Industry
Test and Evaluation	Reduce risk for development, engineering, and government procurement	Government labs and HEC centers

Source: Reference^[5]

describes very interesting scientific challenges and their potential outcomes in fields such as physics, nanoscience, aerospace, life sciences, national security, earth and atmospheric sciences, and energy environment. Over several pages, the Plan describes in detail the issues and HEC needs of climate and weather research, nanoscale science and technology, life sciences applications, and aerospace vehicle design and optimization. This illustrates that access to HEC will continue to increase.

(6) Procurement emphasizing TCO*³

Because TCO includes all financial aspects to provide the HEC service, the HEC Plan describes it as an element equal to benchmarks in determining system procurement. The Plan adds that time to solution, which drives costs during the HEC system life cycle, is an especially important factor.

The Plan also strongly notes user productivity as an element of TCO. In the roadmap, key issues for minimizing time to solution include ease of application software development such as important evolutions in compilers, and the programming environment with improved portability between HEC systems. Because the lifecycle of software is much longer than that of hardware, the optimized use of the huge accumulation of software assets that have been

developed and maintained over many years, and optimized portability of application assets regarding functionality and performance tuned for highly practical use are important issues.

(7) Practical performance measurement*⁴

Since reliable benchmarks that measure sustained system performance are an important element in determining procurement, the HEC Plan describes that federal agencies with similar applications develop and share benchmarks reflecting performance in actual operational environments. The Plan also states that research on “synthetic benchmarks” is being conducted with the support of DoD and the DOE to cope with cases where actual applications cannot be used.

(8) Comprehensive approach aiming for revitalization

The HEC plan states that the revitalization should be supported for the innovative development across the four major research stages: basic and applied research, advanced development, engineering and prototypes and test and evaluation. The HEC Plan calls it a comprehensive approach and also states that this approach is vital to the establishment of a sustainable R&D process.

Although the Plan does not give specific details, the Task Force presented the purposes

Note

*³ Information on user productivity: Refer to Reference^[3] describing HPCS activity in DARPA for a discussion of user productivity. (The content is not included in the HEC Plan, but is shown here for reference.)

*⁴ Information on benchmarks: See References^[10] and^[11] for current trends in benchmarks for measuring sustained system performance in real operational environments. (The content is not included in the HEC Plan, but is shown here for reference.)

and performers of the four major stages as shown in Table 8 at an international high performance computing, networking and storage conference (SC2004) on November 9, 2004. The table shows us how the approach is being promoted.

(9) Well-coordinated interagency plans from the user's point of view

The Task Force members who created the HEC Plan are listed in an appendix to it. All members belong to a user department of an HEC-related federal agency, and therefore the HEC Plan is based on a user perspective. In addition, the HEC Plan seems to have been generated by consistently considering the missions of federal agencies in HEC R&D, HEC resource and procurement.

4 | Conclusions

Regarding the HEC Plan, the 108th Congress debated at least three bills (HR4516, S2176, and HR4218) that included "Revitalization of High-End Computing" in their names during 2004. Of these, HR4516 the Department of Energy High-End Computing Revitalization Act of 2004^[6] ("the Revitalization Act") was approved in November^{*5}. During the course of the debates, HEC was described as having the ability to accelerate progress in fundamental sciences, as an essential component of national security and economic competitiveness, as having a ripple effect on industry, and as requiring the support of the Federal Government. In addition, Japan's Earth Simulator was cited several times as strongly demonstrating the necessity of HEC^[7]. In addition to amounts otherwise made available for HEC, the Revitalization Act is provided with authorization of appropriations of \$50 million

for FY2005, \$55 million for FY2006, and \$60 million for FY 2007, totaling \$165 million over the three-year appropriation period. The Department of Energy will use these funds for HEC research, HEC system development and procurement, the establishment of a software development center, and the transfer of HEC technology to the private sector.

The Revitalization Act specifies research of multiple architectures, research on software for HEC systems in collaboration with architecture development and the establishment of a high-end software development center. The Revitalization Act also specifies sustained access to HEC systems and to leadership systems by the research community in the USA. Therefore, the HEC environment for highly prioritized processes is enforced by providing the access to leadership systems for researchers in the United States industry, institutions of higher education, national laboratories, and other federal agencies.

So far, I have presented an outline of the HEC Plan and commented on points of interest. The Task Force has concludes that current HEC systems provided by industry are not always sufficient for the required performance of application used for Federal missions. Consequently, the Task Force compiled suggestions for HEC investment, aiming for the development of science and technology through cooperation among scientists, universities, industry, and federal agencies. In the HEC Plan, Federal agencies are identified as major users of HEC systems, and therefore Federal support is essential for HEC R&D to meet their requirements. Since Japan's Earth Simulator is regarded as an excellent system by the HEC-related personnel of the US Federal

Note

*5 The status of other bills: S2176, which mandates a five-year appropriation period and an \$800 million total budget, has nearly the same content as HR4516 and was debated in the Senate in March 2004. HR4218, High-End Computing Revitalization Act of 2004, is an amendment of the High-Performance Computing Act of 1991 and was received in the Senate after passing the House of Representatives in July 2004, and was referred to the Committee on Commerce, Science and Transportation. The HEC Plan covered in this article was presented in the deliberation of this bill at a hearing of the House Science Committee in May 2004. In addition, HR28 was presented in the 109th Congress in Jan. 2005.

Government, it will have a significant impact on deciding future R&D policies.

Currently, the US government is strongly promoting HEC-centric strategies in order to maintain its global leadership in science, engineering and technology, and is making every effort to maintain and succeed in technological capabilities that can generate a ripple effect through the pursuit of ultimate technologies related to HEC.

Acknowledgements

In writing this report, I received great help on gathering references and received valuable comments from Dr. Tetsuya Sato, Dr. Kunihiko Watanabe (Earth Simulator Center), Mr. Hisashi Nakamura (Research Organization for Information Science & Technology), Dr. Ryutaro Himeno, Mr. Yoshinari Fukui (Riken), Dr. Keiji Tani, Mr. Toshio Hirayama, Dr. Hiroshi Matsuoka (Japan Atomic Energy Research Institute), Prof. Chisachi Kato (University of Tokyo), Prof. Yoshiaki Nemoto, Prof. Hiroaki Kobayashi (Tohoku University), Prof. Hiroshi Kamimura (Tokyo University of Science), Dr. Satoshi Itoh (Toshiba) and other HEC specialists. I also received great assistance on compiling this report from Mr. Atsushi Ogasawara (National Institute of Advanced Industrial Science and Technology) and Dr. Akihiro Fujii (Science and Technology Foresight Center). I would like to express my sincere thanks to all of them.

Acronyms and full spellings

- CAF* Co-Array Fortran
- COTS* Commercial-Off-The-Shelf
- DARPA*
Defense Advanced Research Projects Agency
- DOE/NNSA*
Department of Energy/National Nuclear Security Administration
- DSM* Distributed Shared Memory
- EPA* Environmental Protection Agency
- HEC* High-end Computing
- HECRTF*
HEC Revitalization Task Force
- HPCC* HPC Challenge Benchmarks

- HPCMP*
High Performance Computing Modernization Program
- HPCS* High Productivity Computing Systems
- IDE* Integrated Development Environment
- MPI* Message Passing Interface
- NASA* National Aeronautics and Space Administration
- NIH* National Institutes of Health
- NIST* National Institute of Standards and Technology
- NITRD* Networking and Information Technology Research and Development
- NOAA* National Oceanic and Atmospheric Administration
- NSA* National Security Agency
- NSF* National Science Foundation
- NSTC* National Science and Technology Council
- ODDR & E*
Office of the Deputy Director Research and Engineering
- OMB* Office of Management and Budget
- OpenMP*
Open specification for MultiProcessing
- OS* Operating System
- OSTP* Office of Science and Technology Policy
- PIM* Processor-In-Memory
- RAS* Reliability, Availability, Serviceability
- TCO* Total Cost of Ownership
- UPC* Unified Parallel C

References

- [1] http://www.hpcc.gov/pubs/brochures/nitrd_20041029.pdf
- [2] FY 2004, 2005 "NITRD Blue Book"
- [3] "The NITRD Program: FY 2004 Interagency Coordination Report," Interagency Working Group on Information Technology Research and Development (IWG/ITR&D), October 2004: http://www.nitrd.gov/pubs/20041020_icr.pdf
- [4] "Federal Plan for High-End Computing: Report of the High-End Computing Revitalization Task Force" (HECRTF), May 10, 2004: http://www.nitrd.gov/pubs/2004_HECRTF/20040702_HECRTF.pdf
- [5] http://www.hpcc.gov/hecrtf-outreach/sc04/20041109_culhane.pdf
- [6] Department of Energy High-End Computing

- Revitalization Act of 2004, H.R. 4516
- [7] <http://thomas.loc.gov/cgi-bin/cpquery/T?&report=sr379&dbname=cp108&>
- [8] "Getting Up to Speed: The Future of Supercomputing" (2004): <http://books.nap.edu/books/0309095026/html/R1.html>
- [9] "Special report: Earth Simulator" Information Processing "Joho Shori", Feb. 2004 (in Japanese)
- [10] "S&T Topics", Science and Technology Trends, Feb. 2005
- [11] Kobayashi, et al. "Performance evaluation of SX system on HPC challenges", SENAC, Vol.38, No.1, pp. 5-28, 2005 (in Japanese)

(Original Japanese version: published in February 2005)
