

Update: NITRD Agencies Continue Implementation of *Federal Plan for High-End Computing*

Since the May 2004 release of the *Federal Plan for High-End Computing* by the White House Office of Science and Technology Policy (OSTP), the NITRD HEC agencies together have implemented many key recommendations of the *Federal Plan* and are providing ongoing support for these new efforts.

The HEC agencies collaborate through the HEC Interagency Working Group (HEC IWG), which coordinates high-end computing policy, strategies, and programs across NITRD member and participating agencies. Recent highlights of NITRD agencies' implementation of the *Federal Plan* include the following:

High-End Computing University Research Activity (HEC-URA): Begun in 2004 by NSF, DARPA, DOE/SC, and NSA, this research effort in operating systems, languages, compilers, libraries, software tools, and development environments has also been supported by DOE/NNSA. In 2006 and 2008, major emphases were file systems, storage, and I/O, including an annual national workshop to spur thinking among researchers about the R&D agenda in these topics; parallel language and programming environments; and HEC operating systems.

DARPA High-Productivity Computing System (HPCS) Program: Now in its final and prototyping phase (Phase III), this program begun in 2001 to pioneer a new generation of innovative, highly productive, economically viable HEC systems was opened by DARPA for participation by other agencies, becoming the first activity in which the HEC agencies could work collaboratively to implement the *Plan's* recommendations for development of new HEC requirements analyses, performance metrics, and assessment tools (such as means of calculating the total cost of ownership). Agencies involved in these activities included NSF, OSD, DOE/SC, NSA, NASA, and DOE/NNSA. DOE/SC and NSA are providing Phase III funding support.

Leadership Systems: The *Federal Plan* proposed that HEC agencies develop "leadership systems" – highest-capability computing platforms – and open them to U.S. scientists to engage in leading-edge computational science and engineering research. In FY 2009, the time on DOE/SC leadership systems made available to outside researchers through the agency's INCITE program will rise from 250 million hours to a billion hours. Also in FY 2009, NSF, DOE/SC, and DOE/NNSA will continue their acquisitions of prototype systems at or approaching the petascale level (see HEC I&A Budget Highlights on page 2).

Petascale Computing: As they move toward acquisition of petascale systems that may involve 100,000 or more processors, NSF, DOE/SC, and DOE/NNSA are collaborating in research activities focused on issues in petascale architectures, software, programming environments, and applications.

System Performance Assessment: Standard metrics used to assess the performance of high-end computers do not measure how a system will perform on the most demanding scientific applications. The HEC agencies are continuing their collaborative work to develop an interagency suite of HEC benchmarks that can accurately represent Federal advanced computing demands. Such a project was called for in the *Federal Plan* to improve the efficiency of the Federal HEC acquisitions process for both agencies and vendors; to provide a more common measure of performance across Federal HEC systems; and to promote interagency cooperation in acquisitions. Initial investments by DARPA and NSF generated the HPC Challenge (HPCC) suite of system benchmarks, composed of seven recognized computational kernels intended to span varied HEC architectures and memory-access patterns. OSD (HPCMP) is now leading a follow-on effort by the HEC agencies to develop a comparable suite of application benchmarks typifying the computational workloads of Federal HEC systems. In 2008, four applications used by multiple agencies and representing varied computational and programming characteristics will be tested on selected Federal platforms. Subsequent steps will include testing of additional applications, development of a Federal Web site for benchmark results and test-case datasets, and exploration of new overall system performance metrics.

Production Software Inventory: The HEC agencies are also developing a detailed inventory of production software installed on government high-end systems. In a preliminary 2007 survey, DOE/SC and DOE/NNSA collected data on the major categories of HEC software – operating and runtime systems, programming environments, and applications support – by type of software, function, and package name, for 30 systems managed by DoD, DOE/SC, NSA, NASA, NIST, and DOE/NNSA. An analysis of the survey results found that 20 software packages are used by many of the systems. However, there is "a long tail" of 135 software packages used by only one system. The survey found that, with a few exceptions, various forms of open-source software are in use across the types of software reported. Areas with substantial commonality include languages, MPI, tools, and Linux; specialized software was reported in system management and certain kinds of development tools. A key goal of this effort is to develop an approach for interagency cooperation in supporting and maintaining production-level open-source software for the Federal HEC infrastructure.