Advanced Scientific Computing Advisory Committee
November 21, 2014
Hyatt French Quarter New Orleans Hotel, New Orleans, La.

ASCAC members present:

Martin Berzins
Vinton Cerf (via telephone)
Barbara Chapman (afternoon only)
Jacqueline Chen
Silvia Crivelli
Jack Dongarra (afternoon only)
Roscoe Giles (Chair)
Sharon Glotzer (via telephone)
Gwendolyn Huntoon (via telephone)
Juan Meza (via telephone)
John Negele (via telephone)
Linda Petzold (via telephone)
Daniel Reed
Vivek Sarkar (via telephone)
Dean Williams (via telephone)

ASCAC members absent:

Anthony Hey

Also participating:

Katie Antypas**,** Services Department Head, National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory
Steven Binkley, Director, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Arthur Bland, Project Director, Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory
Keri Cagle, Oak Ridge Institute of Science and Education
Richard Carlson, Computer Scientist, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Christine Chalk, ASCAC Designated Federal Officer, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Sudip Dosanjh, Director, National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory
Al Geist, Chief Technology Officer, Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory
William Harrod, Director, Computational Science Research and Partnership Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Barbara Helland, Director, Facilities Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Paul Messina, Director of Science, Argonne Leadership Computing Facility, Argonne National Laboratory
Frederick O’Hara, ASCAC Recording Secretary, Oak Ridge Institute of Science and Education
Tjerk Straatsma, Scientific Computing Group Leader, Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory
Timothy Williams, Principal Project Specialist, Computational Science, Argonne National Laboratory

About 45 others were in attendance.

Morning Session

The meeting was called to order at 8:32 a.m. CST by the Chairman, **Roscoe Giles**.

Keri Cagle made safety and convenience announcements.

Giles welcomed the new members, Martin Berzins Silvia Crivelli, and Daniel Reed, and acknowledged the presence on the telephone of Vinton Cerf, Juan Meza, Linda Petzold, Gwendolyn Huntoon, John Negele, Dean Williams, Sharon Glotzer, and Vivek Sarkar. Giles announced that public commentary would be allowed after each presentation.

Steven Binkley was asked to present an update on the activities of the Office of Advanced Scientific Computing Research (ASCR).

The Government is in a continuing resolution. A budget for FY15 is under consideration by the House and Senate. It may last 6 months; if that happened, there would be impacts. No new measures could be started. There may be an omnibus bill in January.

The enacted FY14 appropriation for ASCR was $478.1 million; the FY15 President's request was $541 million, an increase of $62.9 million or 13.2%. The percentage increase is one of the largest proposed for an Office of Science (SC) program office. The Energy and Water committees strongly supported ASCR during their budget markups. The House Energy and Water Development Committee report called for funding ASCR at the full level called for in the President's request, specifically mentioning the need for "maintaining the nation's global leadership in computing technologies." The Senate Energy and Water Development Committee report called for funding ASCR at a level higher than that called for in the President's request. The Committee specifically recommended "$104,317,000 for the Oak Ridge Leadership Computing Facility [OLCF], $80,320,000 for the Argonne Leadership Computing Facility [ALCF], and $85,000,000 for the National Energy Research Scientific Computing Center [NERSC] facility at Lawrence Berkeley National Laboratory."

The FY15 budget priorities are to conduct R&D to produce exascale systems by 2022; to prepare today's applications to migrate to the exascale; and to acquire and operate more-capable computing systems.

The $91 million for the exascale maps onto Applied Mathematics, Computer Science, Computational Partnerships, and Research and Evaluation of Prototypes. The $9.9 million increase for the Big Data Initiative maps onto Applied Mathematics, Computer Science, Computational Partnerships, Next-Generation Networking for Science, and High-Performance Network Facilities. In addition, there is a Facilities increase of $30.4 million for preparing for next-generation machines.

The Computational Sciences Graduate Fellowship (CSGF) has no money in the funding request. There will be carryover money available and something in the budget. Patricia Dehmer has highlighted the CSGF to Congress. The Office of Management and Budget (OMB) passback on the FY16 budget request is due soon.

There is a backlog in Senate confirmations. Only a handful of DOE appointments have been approved. The Deputy Secretary for Science and the Director of SC are still pending.

In Office staffing matters, Christine Chalk has moved to the Facilities Division and taken over the responsibilities for the OLCF. A new hire, Robinson Pino will be working in High-Performance Computing Science and Applications. Carolyn Lauzon (a former American Association for the Advancement of Science Fellow in the Office) is now a federal employee and is the Argonne Leadership Computing Facility Program Manager. Ravi Kapoor has left government service to work in academia.

The Secretary of Energy Advisory Board (SEAB) has had an ongoing consideration of high-performance computing. On Aug. 10, 2014, it posted its report that looked at high-performance-computing needs in a broad context. Its key findings are

- A need exists for an exascale-class machine.

- Projectable technology development can enable one last “current-generation” machine.
- “Classical” high-end-simulation machines are already significantly impacted by many of the data-volume and architecture issues.
- Data centricity at the exascale is already important for DOE missions.
- A lot of common challenges and undergirding technologies span the computational needs.
- The factors that drive DOE’s historical role in leadership computing still exist and will continue.
- A broad and healthy ecosystem is critical to the development of exascale systems and beyond.
- It is timely to invest in science, technology, and human resources for the next generation of high-performance computing.

The SEAB report recommended that

- SC should lead the program and investment to deliver the next class of leading-edge machines by the middle of the next decade at the 1- to 10-exaflop level.
- This program should be executed by using a partnering mechanism with industry and academia.
- The National Strategic Computing Initiative (NSCI), a co-design process that jointly matures the technology base for complex modeling and simulation and for data centric computing, should be part of a jointly tasked effort among the agencies with the biggest stake in a balanced ecosystem.
- DOE should lead a cross-agency U.S. Government investment in “over-the-horizon” [beyond complementary metal oxide semiconductor (CMOS)] future high-performance computing technology.
- DOE should lead the USG efforts to invest in maintaining the health of the underlying balanced ecosystem.

The NSCI, launched in November 2013, has not been discussed very much. It is led by Patricia Falcone, Associate Director for National Security, and it is tasked to define the U.S. Government’s approach to high-performance computing for the next decade. It is to report in time to inform the FY16 budget deliberations. Membership includes the Office of Science and Technology Policy (OSTP), OMB, DOE, the intelligence community, the Department of Defense (DoD), the Department of Commerce (DOC), the National Aeronautics and Space Administration (NASA), the National Institutes of Health (NIH), and the National Science Foundation (NSF). Its purpose is to define how the federal government should proceed with high-performance computing in FY16 and beyond. It is in the process of finalizing its recommendations with a public announcement expected in the next 2 months. The SEAB Task Force was briefed by this Initiative.

Historically, ESnet [Energy Sciences Network] has served DOE facilities in the United States. It is a 100 Gb/sec network. During the past year, in preparation for the Large Hadron Collider’s coming back online after its upgrade, the interconnects with Europe have been upgraded to 100-Gb/sec circuits. Two such interconnects are already operational.

ASCR researcher Alexandre Chorin won the National Medal of Science for contributions to turbulence modeling. And the 2015 Society for Industrial and Applied Mathematics/Association for Computing Machinery Prize in Computational Science and Engineering was presented to the PETSc [Portable, Extensible Toolkit for Scientific Computation] team.

Giles reaffirmed the committee’s interest in the CSGF. Binkley replied that Dehmer has worked very hard on this. OMB plans to consolidate fellowship programs under the NSF with the CSGF as the model. Discussions with NSF are ongoing to corroborate their readiness and to go back to OMB.

Giles asserted the Committee’s interest in seeing the leadership of SC being confirmed. Binkley said that Dehmer has done a tremendous job carrying out not only her full-time job but also the job of Acting Director of SC. The nominees are still hanging in there, going out and visiting the national laboratories.

Chen asked Binkley to comment on the similarities between NSCI and the SEAB Task Force. Binkley answered that the SEAB Task Force is internal to DOE and answers only to DOE. The exascale is a topic of particular interest to the Secretary. NSCI is at the federal-government level; high-performance computing has been seen as a critical capability government-wide. NSCI was set up by OSTP to look at that interest broadly.

Steven Binkley was asked to present the new charge regarding DOE's Office of Scientific and Technical Information (OSTI), which is part of SC and that acts as the repository of scientific and technical information coming out of federal energy research. It is a separate, government-owned, government-operated facility. It used to be a paper-document storage mechanism. Things have changed dramatically since it was founded in 1947. OSTI has made the transition to electronic document storage and distribution. Further changes are coming down the road to make public access to federally funded research more facile. OSTI has been working out arrangements with various types of scientific publishers and federal agencies to make more information more available. SC operates OSTI for DOE. Dehmer is requesting ASCAC to establish a standing subcommittee to look at how OSTI operates; how it compares to other, similar facilities; and where the industry is going (i.e., at the ongoing increase in the amount of data and the interpretation of those data) and how OSTI should be involved in those trends.

OSTI's responsibilities to collect, preserve, and disseminate scientific and technical information were codified in the enabling legislation of DOE and its predecessor agencies and, more recently, was defined as a specific OSTI responsibility in the Energy Policy Act of 2005. External, independent advice is needed as OSTI transitions its products and services to methods appropriate to a new era of information gathering and sharing. Why should ASCAC supply this advice? SC funds 47% of federal basic research in the Physical Sciences. This charge has to do with data and information structure, and ASCAC is the most suited in SC to address these issues. The study seeks to determine:

- Are current OSTI products and services best in class and are they the most critical for the OSTI mission given the present constrained budget environment?
- Do OSTI products and services fulfill customer needs now?
- Are the OSTI products and services positioned to evolve to fulfill customer needs in the future?
- Are there new product types that OSTI should be anticipating?
- What is the national and international standing of OSTI with respect to similar organizations whether at other U.S. Federal Agencies, DOE Laboratories, or universities?

Giles opened the discussion of the charge to ASCAC members. Chen asked if the subcommittee would continue to exist beyond the initial charge. Chalk replied, yes. Giles asked if the subcommittee should have outside members and, if so, how one should represent the ongoing subcommittee to them. Chalk answered that it should have outside members who have been well-forewarned.

Giles noted that this subcommittee overlaps several other subcommittees that ASCAC has had (e.g., the subcommittee on big data).

Reed pointed out that the access issues regarding the results of open science overlap with the concerns of the NSF and of other agencies; however, classified data and information have different dynamics. Binkley responded that, in 1947, 90% of investment in the Atomic Energy Commission (AEC) was in national security. That number is much different today. NNSA's relevant budget today is less than half of SC's budget. And there are dual-use technologies (like accelerator-based technologies that were developed for the light sources and neutron sources that might be used to control fission reactions). In the next decade or two, the preponderance of DOE's research will be on the open-science side, and OSTI needs to understand what implications the broad sharing of those publications will have. Reed stated that the National Archives has a mix of classified and unclassified documents; their practices may inform this subcommittee and OSTI. Binkley commented that there is a large barrier to accessing National Archives and Records Administration (NARA) information if you are not part of the library community.

Crivelli noted that there are standards for storing data. Binkley added that data storage is being looked at across DOE. A lot of data are coming in for analysis during the next decade. The problem is just starting to be understood. Data storage and OSTI might intersect in the future. The different physical-science programs have different needs for data storage. Giles noted that this Committee could use a presentation from OSTI at the next meeting and will seek out subcommittee membership and leadership in the meantime.

Berzins noted that the reproducibility of data by researchers is another important issue.

William Harrod was asked to present an update on the exascale effort and to introduce the plan for achieving the exascale.

Everything that DOE is doing with the exascale is driven by applications in the natural sciences and national defense. The next generation of computing (i.e., a 1000-time improvement over the current generation of high-performance computers) has been talked about for 5 years. That jump cannot be made with current technology. Significant investment needs to be made to do scientific discovery, conduct engineering design, and pursue national prosperity.

There are four challenges faced: parallelism/concurrency, reliability/resiliency (checkpoint/restart is not an option), energy efficiency, and memory/storage (a cost issue as well as an efficiency issue). There are productivity issues in managing system complexity, portability of code, and generality for high-performance computing across the industry. Data movement across the whole system entails energy costs. System-design issues include scalability, efficiency, time to solution, and dependability at all levels of the design. Dependability includes security as well as reliability.

The strategy is to exploit the co-design process, as driven by the application workflow; to develop exascale software stacks; to partner with and find vendors to transition research to production speed; and to collaborate with other government agencies and other countries (as appropriate).

Exascale funding is being sought for application development (FY12 to FY15) and for science, engineering, and defense applications (FY16 to FY24), producing (through the vendors) a node prototype, a petascale prototype, an exascale prototype, and exascale delivery. A significant increase in budget is hoped for in FY16. Without that increase, it would not be worthwhile going forward.

The Fast Forward Program started funding vendors in 2012 to start designing the node, software, interconnect architectures, compilers, etc. The codesign practice has worked well. The Design Forward Program started in fall 2013 to see what the conceptual design would look like. Fast Forward 2 looked at memory technology that might apply to all vendors. It is just starting. Some vendors are building on what they have previously done. Others are looking at alternative processor designs.

There are three codesign centers in (1) materials, (2) advanced reactors, and (3) combustion. They have had notable successes (e.g., they have developed proxy applications, re-examined current algorithms, and explored programming models). Lessons they have learned include

- Interdisciplinary teams discover key issues,
- Personnel are in high demand,
- Horizontal integration is needed, and
- Computations and communications will be irregular.
- In addition, the system will have a deep-memory hierarchy, and code size can be reduced by a factor of 0.01 while getting 3-times-faster runtimes.

A lot of interest has been expressed by the Secretary of Energy, SEAB, OMB, Congress, OSTP, etc.

A preliminary version of a project design document has been developed, a red-team review of it has been conducted, an external review of it has been conducted, and it has been updated. New programs have been installed: Fast Forward 2, Design Forward 2, Resilience, Analytical Modeling for Extreme-Scale Computing Environments, Scientific Data Management, and Analysis and Visualization at the Extreme Scale.

ASCAC is being asked to review a document on the preliminary strategy. That document has limited descriptions of project management details and it has a preliminary list of proposed projects. The effort has three steps: R&D, application, and deployment. The document has a list of risks and potential mitigations.

In summary, exascale computing systems are essential for the scientific fields that will transform the 21st Century global economy: energy, biotechnology, nanotechnology, and materials science. The goal is to lead computational sciences and high-performance computing to continue to develop and deploy high-performance-computing hardware and software systems through exascale platforms. A DOE ASCAC subcommittee has identified the top 10 technology advancements that are critical to making productive, economically viable exascale systems. Innovations are required beyond the anticipated conventional practices and their incremental extensions.

Giles asked if the document were available. Chalk said that it will be posted on the web soon and will be fully public.

Steven Binkley was asked to present the charge concerning this document.

The charge letter is general. The main part is "to review the Department's draft preliminary conceptual design for the Exascale Computing Initiative. Specifically, we are looking for input from the community as to whether there are significant gaps in our plans or areas that need to be given priority or extra management attention. ASCAC should gather, to the extent possible, input from a broad cross-section of the stakeholders."

There are certain caveats. There are no funding levels cited in the document because that funding is under discussion by OMB. Also, the application software is being done differently by the National Nuclear Security Administration (NNSA) and SC. The NNSA effort is integrated into the Alliance for Computing at Extreme Scale (ACES; stockpile stewardship). The software developed has to be owned by the domain programs that the software applied to (and that has the funding allocated to it). The intent is that ASCR will partner closely with these offices. The climate program of the Office of Biological and Environmental Research (BER) has been very active during the past year, and ASCR's participation has become mature. The Office of Basic Energy Sciences (BES) is waiting to see if it gets FY15 funding for a next-generation materials code. ASCR's activity with BES is currently slim.

Another caveat is the acquisition strategy. The timeline calls for acquisition to start in 2018 or 2019. The plan does not go into much detail about how those acquisitions would be done.

The subcommittee's report would be due Mar. 15, 2015.

Giles asked if applications other than those specified would benefit. Harrod replied that the intent is to deploy technology as rapidly as possible. Technology developed in Fast Forward 1 is being used in the Collaboration of Oak Ridge, Argonne, and Livermore (CORAL) acquisition.

Cerf asked if there were a good enough definition of an exascale machine. Harrod replied, yes; and those definitions are being refined during the next year. It is hoped to release requests for proposals for Fast Forward designs later this year with those specifications and definitions.

Berzins noted that this represents a substantial shift in how programming is done and taught. There are some cultural and workforce changes that come with it. One must train many more people on the large petascale systems. There is considerable uncertainty in how these systems will be built. Harrod said that the seriousness of this issue is recognized; however, the problem has not been fleshed out completely.

The floor was opened to public comment. Richard Bower (Boston University) pointed out that the technology has to be marketable in smaller systems. Application people have to have access to the programming structure. Harrod said that the project is focused on building technology to go from the terascale to the exascale. This document represents a lot of effort. There has been a lot of discussion.

Chen asked if there were linkages in the document that look like Scientific Discovery through Advanced Computing (SciDAC) and that can bring up programs before the machine is operational.

Harrod said that there were no such linkages in the document and that that issue is something that needs to be expanded upon.

Giles said that he was seeking volunteers for the subcommittee that will respond to this charge.

A break was declared at 10:52 a.m. The meeting was called back into session at 11:01 a.m.

Roscoe Giles presented the Committee of Visitors (COV) Report on SciDAC. ASCAC was charged to assess the efficacy and quality of the SciDAC processes; the breadth and depth of portfolio elements; the national and international standing of the portfolio; and the degree to which the program is anticipating and addressing emerging challenges from high-performance computing and DOE missions. The bottom line is that SciDAC remains the gold standard for fostering interaction between disciplinary scientists and high-performance computing. The program managers (PMs) are to be commended on continuing the excellence of the SciDAC "brand."

The COV was conducted Oct. 6–7, 2014; meetings were held with program managers from BES, BER, Fusion Energy Sciences (FES), High-Energy Physics (HEP), and Nuclear Physics (NP) with access provided to all relevant documents.

The mission of the SciDAC Institutes is to provide intellectual resources in applied mathematics and computer science, expertise in algorithms and methods, and scientific software tools to advance scientific discovery through modeling and simulation. Science Application Partnerships (SAP) exploit leadership-class computing resources to advance scientific frontiers in an area of strategic importance to SC and link to the intellectual resources in applied mathematics and computer science, expertise in algorithms and methods, and scientific software tools at one, or more, SciDAC Institutes. The interesting question is how these efforts are sequenced. Information was provided by the science-area project offices on the selection of science areas, the setting up of institutes, and the establishment of partnerships. The COV found that this sequence worked well.

The COV found that the timing of the calls for institute proposals and the interrelated partner proposals is a challenge. Asking the PMs in the science areas to define their areas of interest, followed by the institute competition with knowledge of those areas, followed by the actual science program completion, was a good process.

The PMs are to be commended for having the courage to re-compete the data institute rather than accepting a suboptimal solution among the original proposals.

Process Documentation has much improved since the previous review in 2007.

Critical to the success of SciDAC-3 is the active role played by PMs in selecting projects and evaluating proposals in light of the highly complex and multidisciplinary mission needs. The COV commended the PMs.

The solicitation was complex and effective overall with cooperation among various PMs. The communication about the nature of the data institute was not as effective, necessitating a second call.

The requirements for pre-proposal and portfolio downselection varied widely across the offices, but the PMs are coping with those differences in a manner appropriate to the discipline.

The COV had several recommendations regarding solicitations and awards:

- Preserve decision documents, even for declined proposals, and provide summary feedback in the declination letter.
- Maintain coordination between science programs and ASCR priorities in timing decisions pertaining to future proposals.
- It is important that the PMs can impose the SciDAC priority filter over and above the peer reviewers, who (properly within their sphere) rank base on the traditional merits of quality and originality.

The COV found that projects are well monitored by PMs through frequent teleconferences.

The communication and interaction of PMs with the complex teams that are involved is essential. The level of interaction of the PMs with the teams is commendable.

The ability of PMs to travel to project meetings and important conferences is important but is currently insufficient. Current travel support is inadequate.

The program was adaptive to changing circumstances. For example, when one principal investigator (PI) became ill, there was an intervention that resulted in a two-PI arrangement that worked very well.

In regard to breadth and depth, the COV found that the aggressiveness of distributed memory scaling among the applications is much improved relative to previous SciDAC rounds. Now, software migration to many-core architectures is of strong interest to the many of the applications groups, as well. This a potential vulnerability in that the software may not be able to be maintained after the life of the program.

In breadth and depth, the COV recommended that

- The current balance among science-based algorithms and insights, mathematical/computational algorithms and high-performance computing be maintained.
- ASCR should pursue synergisms between SciDAC and Co-Design.
- In terms of demonstrating success for SciDAC collaborations, wide adoption of code developed by the institutes should be regarded as at least as meritorious as shared post-doctoral full-time equivalents (FTEs) in that it shows that the algorithmic and software technology has reached maturity.

In regard to emerging challenges, the COV found that the breadth and depth of the portfolio is appropriate. Summer schools and tutorial efforts by the Institutes to expand the science impact of their work are successful and important. The fact that a PI cannot count on an Innovative and Novel Computational Impact on Theory and Experiment (INCITE) allocation generates concern for SciDAC PMs and PIs. This represents a risk if allocations do not in fact meet the need. Also, INCITE is funded just once a year, placing timing limitations on SciDAC projects. However, there was no case of a project being stopped because of the lack of cycles.

In emerging challenges, the COV recommended that

- The balance between ASCR Leadership Computing Challenge (ALCC) and INCITE computing resources be tuned in light of SciDAC requirements.
- The COV strongly encourages the Institutes to expand outreach efforts in the out years of SciDAC-3 to reach a larger scientific community.

In regard to national and international standing, that COV found that SciDAC remains the gold standard for fostering interaction between disciplinary scientists and high-performance computing.

Berzins said that SciDAC does a great job rallying resources but is geared toward grand challenges, not what INCITE calls for. “Infrastructure” work should be included. Giles responded that that was part of the intent of the institutes. The COV focused on the process. The Data Institute still should have done more projects than it did.

Reed noted that the matching of computer allocations with needs is a problem across the research community, and that issue should be considered.

Giles called for the acceptance of the findings and recommendations, and a roll-call vote was taken. The vote was unanimously in favor with aye votes from Giles, Crivelli, Reed, Berzins, Chen, Williams, Negele, Cerf, Meza, Huntoon, Chapman, Petzold, Chapman, and Dongarra.

Sudip Dosanjh was asked to give an update on CORI.

NERSC’s latest machine is Edison, the first Cray petascale machine with Intel processors. It has a very high bandwidth with 64-GB double-data-rate third-generation (DDR3), 1600-MHz memory per node. It has exceptional performance.

Edison does not deploy accelerators or GPU. Disruptions in programming models are a challenge for NERSC because it has many users and many codes and because it does not select its users. There are

5000 users from around the world, 467 at different ASC institutions, supporting a very diverse workload, ranging from small to large scales.

NERSC is a primary computing facility for DOE/SC. DOE allocates the vast majority of the computing and storage resources of NERSC.

For CORI, there is a heightened awareness among application teams. Many codes are being adapted for next-generation systems. Technology changes (e.g., self-hosted many-core chips and tighter CPU/GPU integration) will make the transition easier. We must transition to energy-efficient architectures to meet the science needs of the users. NERSC works closely with users to meet their needs. In the future, it will be harder to meet those needs, largely because of energy usage. The users need to do both high-throughput (statistics, systematics, analysis, and uncertainty quantification) and larger physical systems with higher fidelity.

NERSC needs to transition to energy-efficient architectures through the use of lightweight or heterogeneous systems; many-core or hybrid is the only approach that crosses the exascale finish line. NERSC-8 (Cori) must provide a significant increase in computational capabilities over Hopper; have delivery in the 2015/2016 timeframe; provide high-bandwidth access to existing data stored by continuing research projects; and begin to transition users to more energy-efficient, many-core architectures.

ACES and NERSC formed a partnership for next-generation supercomputers. This was a very visible collaboration between ASCR and ASC to strengthen the impact on industry; to address challenges transitioning applications to advance many-core architectures with a broader coalition; and to act as a risk-mitigation strategy for NERSC-8 and Trinity systems by having a partner to work with on technical challenges in deploying and testing NERSC-8 and Trinity, two separate projects. The name Alliance for application Performance at the EXtreme scale (APEX) was adopted.

Cori will have 64 cabinets of Cray XC System with more than 9300 “Knights Landing” compute nodes with 64 to 128 GB of memory per node. It will also have more than 1900 Haswell compute nodes, an Aries interconnect, and a lot of the same software as Edison. Cori will have a Lustre file system, a nonvolatile random-access memory (NVRAM) burst buffer for I/O acceleration, significant Intel and Cray application transition support, and delivery in mid-2016 at Lawrence Berkeley National Laboratory (LBNL).

The Intel Knights Landing processor is a next-generation Xeon-Phi with more than 3 teraflops peak performance. It is a single-socket processor and it is self-hosted (not a co-processor and not an accelerator). It has more than 60 cores per processor with support for four hardware threads each, more cores than the current-generation Xeon-Phi. It uses the Intel Silvermont architecture enhanced for high-performance computing. It has a high-bandwidth, on-package memory with up to 16 GB/s capacity.

Cori will be installed in the new Computational Research and Theory (CRT) Facility for collaboration and more floor space. It will use natural air and water cooling with heat recovery to heat the building.

The big challenges are the applications. NERSC will initially focus on 20 codes. Currently, 10 codes make up 50% of the workload of Hopper and Edison. 25 codes make up 66% of the workload. Edison will be available until 2019/2020.

The NERSC Exascale Science Applications Program (NESAP) is being launched with strong support from vendors, developer workshops for third-party software, early engagement with code teams, leveraging of existing community efforts, a postdoc program, NERSC training and online modules, and early access to NERSC-8 technology. The 20 focus codes will be Tier-1 and Tier-2 codes from six SC offices. The selected codes are diverse in several dimensions: production codes versus up and coming, higher versus low readiness, and smaller community applications versus overlap with established ALCF/OLCF readiness codes. NESAP has already received recognition at Supercomputing 2014 for its collaboration with Intel and Cray to improve the performance of BerkeleyGW.

Extreme data science is important at all DOE facilities. NERSC is increasing data traffic exponentially. More than 1 petabyte of data is being imported each month now, and the first petabyte day is expected in 2020 as extreme data science plays a key role in scientific discovery. Examples of such discovery science are measurement of the theta-13 parameter of neutrinos, the discovery of more than 2000 supernovae in the past 5 years by the Palomar Transient Factory, trillions of measurements by the Plank satellite in mapping the cosmic microwave background radiation, and the materials project that has more than 5000 users.

NERSC currently deploys separate compute-intensive (Edison and Hopper) and data-intensive (Carver, PDSF, and Genepool) systems. Dramatically growing data sets require petascale+ computing for analysis. In addition, NERSC increasingly needs to couple large-scale simulations and data analyses.

Cori will employ several enhancements in handling data: data partition with large memory nodes and throughput-optimized processors, a burst buffer, and a larger disk system. The goals are to enable the analysis of large experimental data sets and *in situ* analysis coupled to petascale simulations. A data-intensive system will need a local disk, large memory nodes, massive serial jobs, complex workflows, a compute-gateway node, streaming of data from observational facilities, and an easy-to-customize environment.

The goal is to enable science that cannot be done on today's supercomputers. The impact will not be seen for quite a while.

Berzins asked how confident NERSC was in how the second-generation Intel Xenon 5 will perform. Dosanjh replied that there was a high degree of confidence because of its being self-hosted and having improved cores.

Chen asked if NVRAM burst capabilities were needed by users. Dosanjh answered that users will be able to deliberately use that capability. Eventually, it will be automatic for applications that use a lot of input/output.

The floor was opened for public comment. There being none, a break for lunch was declared at 11:59 a.m.

Afternoon Session

The meeting was called back into session at 1:01 p.m.

Tjerk Straatsma was asked to begin the review of the Applications Readiness Plan.

A meeting of ALCF, NERSC, and OLCF concluded that close collaboration on application readiness and performance portability will result in better programming activities, higher value for the ultimate users of the ported applications, and the avoidance of duplication. The outlined collaboration on application readiness and portability does not simply address the next-generation systems that will be coming to ALCF, NERSC, and OLCF but is also on the vendors' paths to exascale architectures.

Two tracks are viable for pre-exascale and exascale machines: many-core (Cori at NERSC) and hybrid multicore (Summit at OLCF). The path for the ALCF-3 is still to be announced.

The Application Developer Team wants to be heavily involved in application-readiness efforts because of its knowledge of the applications, which are working applications that are still under development. Optimizations will be included in the application releases. All programs will have an Early Science Project to demonstrate the application on real problems at scale and to shake down new system hardware and software. Vendor support is crucial; these are all brand-new systems. The team will provide access to multiple resources, including early hardware, and there will be joint training activities.

Portability is a crucial concern. There will be a limited number of codes worked on under three programs:

- NESAP at NERSC had its call for proposals in June 2014; 20 projects were selected; it is in partnership with NERSC; and it has eight postdoctoral fellows
- CAAR [the Center for Accelerated Application Readiness] at OLCF had its call for proposals in November 2014; eight projects are to be selected; it is in partnership with OLCF and IBM/NVIDIA; and it has eight postdoctoral associates.
- The ESP [Early Science Program] at ALCF still is in a planning and organization phase.

NESAP will work on 20 applications from six DOE program offices: ASCR (BoxLib AMR Framework and Chombo-crunch); HEP (WARP & IMPACT, MILC, and HACC); NP (MFDn, Chroma, and DWF/HISQ); BES (Quantum Espresso, BerkeleyGW, PARSEC, NWChem, and EMGeo); BER (Gromacs, Meraculous, MPAS-O, ACME, and CESM); and FES (M3D and XGC1). At Oak Ridge, six codes have previously been used: WL-LSMS, S3D, NRDF, LAMMPS, CAM-SE, and Denovo. The ALCF has done this before, too, with energy storage materials and catalysis (QMCPACK), biomolecular science (NAMD), high-speed combustion and detonation (ALLA/FTT), active aerodynamic flow control (PHASTA), and earthquake genesis (SORD).

The teams have a tentative timeline geared to the delivery dates for the new machines. The OLCF currently has the Titan and will have the Summit in 2018. NERSC currently has Edison and will have Cori in 2016. The ALCF currently has MIRA and will have the ALCF-3 in 2018.

Application portability among ALCF, NERSC, and OLCF is a critical concern to ASCR because application developers target a wide range of architectures, maintaining multiple code versions is difficult, porting to different architectures is time-consuming, many PIs have allocations on multiple resources, and applications far outlive any computer system. The primary task is exposing parallelism and data locality. Moving data is very expensive.

The three teams have met several times to keep abreast of what each team is working on and to address:

- How to coordinate the application readiness efforts,
- What guidance and tools can be provided, and
- What mechanisms and allocations can be provided?

Each team will have a Postdoctoral Associate Program to provide challenging scientific campaigns in critical science mission areas, experience in using ASCR computing capabilities, and training in software development and engineering practices for current and future massively parallel computer architectures. Basically, it is a question of having enough personnel to carry out portability objectives.

Katy Antypas continued the presentation with a discussion of portability.

Generally, an increase in portability comes with a decrease in performance. FASTMath developed two definitions of performance portability:

- The same piece of code runs on different architectures with “good” performance.
- A relatively small amount of effort is needed to make a change to get good performance within advertised tolerances across both current and future architectures.

The dominant programming model at centers is MPI-based. There are MPI only, MPI+OpenMP, and MPI+OpenACC/CUDA systems. The challenge is the deepening of memory hierarchies with fast memory and a slower memory that do not have the same size, same bandwidths, or same attributes. The Knight’s Landing architecture will use near memory as a cache.

Thread scaling and management is another challenge. An increasing number of threads that challenges scaling threaded performance will occur on each architecture, requiring different scheduling policies. It is important to sweep this architecture space.

One way to improve portability is to improve data locality and thread parallelism through refactoring loops, adding OpenMP, and improving vectorization. Another is to use portable libraries (taking the difficulty out of the application). The caveat is that libraries have to be supported and high performing.

The next strategy is to use MPI + OpenMP 4.0. OpenMP 4.0 has new capabilities that make it more attractive for use on or with accelerators by offloading data and computation to accelerators and by supporting more portable vectorization. Significant work is still necessary. All three centers are joining the OpenMP Standards Committee. A number of frameworks abstract architecture details away from the user and thus provide portability (e.g., KOKKOS or TIDA). A final strategy is to encourage portable and flexible programming rather than models that hinder portability (like CUDA). Good coding practices will help a lot, such as using parameters for the amount of threading and placements of threads, allowing data structures to be flexibly allocated to different memory spaces, and allowing task-level flexibility so work can be allocated on different computing elements (i.e., GPU and CPU).

Timothy Williams continued the presentation with a discussion of planning of joint activities.

The SC centers' roles include advocating portability, focusing application efforts, training developers, granting hardware access, and sharing best practices.

Developer training needs to target the center staff, application-readiness development teams, library and tool developers, and the general user community. The different topics that require training include multiplatform development, platform optimization, science-domain-specific sessions, and broader software engineering.

Upcoming training includes Common Application Readiness Approach at Oak Ridge National Laboratory (ORNL) in January 2015, the first of a series of sessions on Scientific Software Engineering at ANL in February 2015, two application portability workshops in spring 2015, ALCF-3 New Architecture Training at ANL in June 2015, and Extreme Scale Computing at ANL on Aug. 2-14, 2015.

Additional available training includes Intel Dungeon sessions (NERSC), vendor workshops (ORNL), NESAP training (NERSC), ESP hands-on workshops (ANL), hackathons (ORNL), ongoing scientific software engineering workshops (ANL), and OpenMP 4.0 training. Where appropriate, there will be shared training events. In all training, the leaders will advocate and instruct on portability.

The focusing application efforts will include Application Readiness Program coordination through joint participation in project review and selection with partnering by representatives from other centers including CORAL and APEX. ESP and CAAR proposal forms ask about proposals to other centers' programs. Other aids to coordination are annual meetings of ESP, NESAP, and CAAR project teams and the tools and libraries readiness efforts. ESP, NESAP, and CAAR will provide companion training or components of the tools and library efforts.

Granting hardware access will include Director's Discretionary or equivalent allocation at other centers as well as ESP, NESAP, and CAAR projects in the form of next-generation hardware simulators; access to next-generation hardware as early as possible (centers/vendors); large allocations of pre-production time for early science runs (ALCF-3, Cori, and Summit); and time on current systems for interim development.

A tools and libraries working group and a cross-laboratory training committee have been set up. Nondisclosure and export-control challenges will have to be managed.

Giles asked how the management of particular programs was shared among multiple centers. Straatsma replied that one center will help out even if the effort were being conducted at one of the other centers. It is desired that all users' programs be portable to all machines.

Berzins asserted that runtime systems should also be thought about and said that the things being done are exactly what needs to be done.

Arthur Bland was asked to present an update on CORAL activities.

The mission of the OLCF is to deploy and operate the computational resources required to tackle global challenges, providing world-leading computational and data resources and specialized services for the most computation- and data-intensive problems; providing a stable hardware/software path of increasing scale to maximize the productivity of users; providing the resources to investigate otherwise inaccessible systems at every scale; and with our partners, delivering transformational discoveries in all fields.

The science requires continual advancement of OLCF's computational capability of the Center. Jaguar was upgraded every year from 6,000 to 300,000 cores. The Center worked with its users and tool vendors to make that jump in scale. Titan is an upgrade of Jaguar with hybrid GPU/CPUs. Summit will have 5 times the performance of Titan with a 10% increase in power. Clock times have become stable; increases in performance have come largely from more-efficient algorithms. Titan has 18,688 compute nodes, each with a 1.45-TF peak, and a Cray Gemini 3-D Torus interconnect.

Scientific progress has been stupendous in fusion energy, turbo-machinery, nuclear reactors, liquid-crystal film stability, earthquake-hazard assessment, and climate science. There are two architecture paths for leadership computers: hybrid multicore like Titan (the fat-node path) and many cores (the skinny-node path). In future fat-node systems, data-movement issues are expected to be easier to deal with than they are with current systems. In future skinny-node systems, tens of thousands of nodes with millions of homogeneous cores with multiple levels of memory will be unlike prior generations in that they will likely be self-hosted.

Leadership computing capabilities are needed from 2018 through 2022. CORAL was formed by grouping the three Labs who would be acquiring Leadership computers in the same timeframe (2017) to develop one request for proposals and draft statement of work.

For risk reduction and to meet a broad set of requirements, two architectural paths will be selected, and Oak Ridge and Argonne must choose different architectures. Once selected, a multi-year lab-awardee relationship will be used to co-design the computers. Both R&D contracts are jointly managed by the three national laboratories. Each national laboratory will manage and negotiate its own build contract.

A buying team met to select the set of two proposals that provided the best value to the government, evaluating the DOE mission requirements being met, technical proposal excellence, feasibility of schedule and performance, diversity of platforms, overall price, supplier attributes, and risk evaluation. Ultimately, the request for proposals resulted in the non-recurring engineering (NRE) contract, the OCLF contract (IBM, NVIDIA, and Mellanox), the Lawrence Livermore National Laboratory contract (IBM, NVIDIA, and Mellanox), and the upcoming ALCF contract.

IBM owns the OCLF contract. They will use multiple IBM POWER9 CPUs and multiple NVIDIA Tesla® GPUs using the NVIDIA Volta™ architecture. The CPUs and GPUs are completely connected with high-speed NVLink™ with 3400 nodes (fewer but much fatter) with large coherent memory (more than 512 GB). An additional 800 GB of NVRAM, which can be configured as either a burst buffer or as extended memory, will produce more than 40-TF peak performance per node. The interconnect is a dual-rail Mellanox® EDR-IB full, non-blocking fat tree. The memory is IBM Elastic Storage with 1-TB/s I/O and a 120-PB disk capacity with burst buffers smoothing out the transfer of data. NVLINK is their GPU high-speed interconnect with 5 to 12 times PCI-E Gen3 bandwidth with planned support for POWER® CPUs.

Summit software will be Linux®. The programming environment will be as much like Titan's as possible to ease portability. It will use OpenMP and OpenACC, Libraries, and debugging routines.

Summit will thus have 5 to 10 times the application performance of Titan.

Dongarra noted that the energy use was impressive and asked about the risk-mitigation plan. Bland replied that there were two architecture paths to back up each other and that the project had held several discussions with others building such machines to provide risk mitigation.

Berzins stated that the bandwidth of the GPUs is the challenge and asked if there were a way to cope with all the links that would be needed. Geist responded that, with multiple GPUs per link, the node has complete connectivity to avoid a bottleneck. Bland added that one can stream all the double-data-rate (DDR) memory to the GPU. Berzins noted that it required a balanced architecture and asked about the software stack. Bland replied that IBM and Cray have mature software stacks. CORAL will use early users and early science projects to flesh out an understanding of the software.

Chen asked if the GASNET and PGAS languages would be supported on Summit. Bland said that he did not know.

Richard Carlson was asked to present an update on Next-Generation Networking for Science (NGNS).

The goal of the program is to research, develop, test and deploy advanced network and middleware technologies critical to DOE's science mission. A unique focus is the end-to-end aspect of the high-performance, high-capacity network and middleware technologies needed to provide secure access to distributed science facilities, high-performance-computing resources, and large-scale scientific collaborations.

The NGNS budget trend has been upward. Funding used funding opportunity announcements (FOAs) and kept the awards going for 3 years; increases in funding are being used to start new projects. NGNS program management uses regular calls and interactions with PIs, workshops to gather new requirements, and annual PI meetings that foster cross-fertilization.

Accomplishments have included

- A PanDA [Production and Distributed Analysis] workload-management system funded for Titan. It is hoped to make this a production system rather than a one-off.
- A multicore-aware data transfer middleware has been developed with 80-Gbps sustained transfer from disk to disk.
- A workflow-management system that can map the workflow and make experiments repeatable from the data produced.
- A 2-year design study has been funded to create a roadmap for distributed computing in the future.

Two projects were funded, one with cofounding from BER to deal with climate change to deliver scientific productivity; outreach establishes a new, holistic perspective for the broader scientific community. The other was in analytical modeling for distributed computing to produce workflow-component analytical models, integrated end-to-end models, experimental validation and verification, and simple-to-use tools.

In analytical modeling for extreme-scale computing environments, four large projects were started in FY14:

- IPPD: integrated end-to-end performance prediction and diagnosis for extreme scientific workflows,
- Panorama: predictive modeling and diagnostic monitoring of extreme science workflow,
- Ramses: robust analytical models for science at extreme scales, and
- X-Swap: extreme-scale scientific workflow analysis and prediction

Activities that have been proposed include workshops on workflow science (Science Automation) and on exploring optical networking from computers to networks. It is hoped to release FOAs this year on a software-defined-network-enabled terabit optical network for extreme-scale science and on dynamic distributed-resource management.

In conclusion, the NGNS program continues to support both applied and basic research in networking and scientific collaborations. Existing programs are being augmented by new areas of research on software productivity and analytical modeling of scientific workflows. Future programs are being developed to address challenges emerging from current work.

Giles noted that DOE had developed these programs and asked what the mechanism was to link them. Carlson replied that a workflow scientist was needed to analyze the infrastructure and to validate the workflow.

The floor was opened to public comment. There was none.

Paul Messina was asked to present an update on INCITE

The Department of Energy High-End Computing Revitalization Act of 2004 established the Leadership Computing Facilities and called on them to provide access on a competitive, merit-reviewed basis to researchers in U.S. industry, institutions of higher education, national laboratories, and other federal agencies.

Sixty percent of leadership-class computing (5.8 billion core-hours in 2015) is allocated to the INCITE, which promotes transformational advances in science and technology through large allocations of computer time, supporting resources, and data storage at the Argonne and Oak Ridge leadership computing facilities for computationally intensive, projects. The INCITE program has Science Panel reviewers to make recommendations on resource allocations. INCITE has two major criteria: the merit criterion, which judges if the research campaign has the potential for significant domain and/or community impact, and the computational leadership criterion, which seeks computationally intensive runs that cannot be done anywhere else.

Anyone can apply to INCITE in a broad array of domains. No number of hours is designated for a particular science area. The program does have a yearly cycle that proceeds from solicitations through a computational readiness review, a panel peer review, selection, awards, access processing, and access and monitoring. From 2007 to 2010, the hours requested were twice those allocated; from 2011 to 2014, the requests were three times the allocations; this despite that fact that Monte Carlo calculations that used to take 1 year on a cluster can now be done in 3 hours on the Titan.

In combustion research, the S3D speedup on the Titan enabled the researchers to conduct calculations with the world’s largest-ever Reynolds number, 13,050.

In a study of high-Reynolds-number turbulent channel flow, direct numerical simulations on Mira led to a more complete understanding of wall-bounded turbulence (e.g., energy loss in pipelines). The simulations used two-thirds of Mira’s capabilities, and the results have been made available to the research community.

A study of the production of hydrogen on demand from the splitting of water demonstrated orders-of-magnitude acceleration of the reaction rate and higher yield by alloying aluminum nanoparticles with lithium. The reaction rates and yield are high enough for industrial use.

The number of projects awarded INCITE hours has slacked off as the size of projects has increased to match the performance characteristics of the machines. For 2015, the total requests were for about 15 billion core-hours, an increase of 1 billion core-hours from the previous year; 56 projects were awarded, of which 30 were renewals. The acceptance rates were 27% for nonrenewal submittals and 91% for renewals. With the new proposals, 48% of the PIs had never before led an INCITE proposal, and 23% of nonrenewal projects awarded time were led by new PIs. INCITE actively engages with new research teams through outreach, such as workshops, email distributions, and individual networking.

For 2015, Titan has 30 projects using an average of 75 million core-hours; Mira has 37 projects using an average of 96.5 million core-hours. Some projects have dual awards on both machines.

In the 2015 cycle, 97 science experts participated in the panel review. Of the reviewers, more than 60% are society fellows, NSF or DOE Early Career awardees, laboratory fellows, National Academy members, or department chairs; 48% participated in the 2014 INCITE review. There was a large number of foreign reviewers to avoid conflicts of interest. The 2015 INCITE panel got good reviews from applicants.

Thirty percent of the time on LCF machines is going to INCITE proposals.

Dongarra asked what percentage of cycles was going to non-DOE-funded people. Messina said that the program asks what funding the PIs get, but it does not vet the replies. Of 56 PIs, 35 said that DOE was their source of funding; 14 said NSF; 6 said foreign sources; 4 said industry was their source; 4 are unknown; 3 are from NASA; and 3 are from other agencies: NNSA, Office of Naval Research, and the DoD.

Chapman asked how many foreign applications were received and how much time was given them. Messina did not know the answer to that question but offered to find out and get the numbers to her.

Crivelli asked, in general, how many of the results are “transformational.” Messina said that the centers list the results and the resultant publications. The answer to the question resides on the LCF websites. Giles noted that ASCAC also reviews the LCFs and asks that same question of the centers.

Berzins asked if the times cited include the GPUs. Messina replied that yes, they do, and they reflect that effect with multiplication factors.

Barbara Chapman was asked to present the final report from the ASCAC Subcommittee on Workforce Development Needs.

The charge was to identify disciplines that need more workforce training at grad student or postdoc level for DOE Office of Science mission needs, including disciplines not well represented in academic curricula, disciplines in high demand, disciplines where the DOE labs may play a role in providing needed workforce development, and programs needed at the graduate student or postdoctoral levels, Letter report on findings and recommendations by June 30.

Computing disciplines needed for DOE missions do not fit neatly into traditional academic categories, so the term “computing sciences” was coined for this report. However it is characterized, though, human resources data typically are not available for this range of subjects.

Scientific computation is at the core of much of Office of Science R&D, and ASCR provides critical technology and expertise that enable DOE’s world leadership in scientific computation. Maintaining a sufficient workforce is this area is critical.

The national laboratories were surveyed on areas where they experience recruitment/retention difficulties. They reported a low number of qualified applicants, many of whom are foreign nationals; long times to fill positions; and spending significant efforts on recruiting. They reported a wide range of competencies were required to carry out their missions. A standardized method for gathering this information would be useful. Analysis of the data from the national laboratories showed a strong demand for MS- and PhD-level computing sciences positions; a longer time to fill positions than industry takes; and attrition rates that compare favorably with industry’s.

Many of the jobs in this area require a multidisciplinary education. A study of academic curricula showed that interdisciplinary computational science and engineering (CS&E) studies are emerging, but they cannot impart the complexities of the field and do not provide the full skillset needed by DOE. There has been an insufficient number of graduates, and the number of graduates has been flat or declining in relevant fields. Interdisciplinary data-science education is beginning to emerge but is not likely to satisfy demand.

In terms of workforce demand, there is a huge demand for graduates with computing expertise, the retirement of the current workforce is expected to grow, the national laboratories cannot compete with industry compensation, awareness among graduates about career opportunities at the national laboratories is low, and conference travel restrictions impede recruitment.

The lack of diversity in U.S. graduates in computing sciences is a major factor in the national shortage. Participation by African American and Hispanic graduates is very low, and the percentage of females is declining. Foreign nationals make up 58% of computer sciences graduates.

The DOE CSGF program was established in 1991. It was jointly funded by ASCR and NNSA to train graduate students to meet national workforce needs in computational sciences. It is extremely competitive and provides practical work experiences at DOE national laboratories. Its participants are consistently highly qualified. Multiple reviews attest to the success of the CSGF program.

The national laboratories also put significant training and outreach efforts into attracting, training, and retaining workers in computing sciences. They engage graduate students and postdocs through such programs as summer internships and university subcontracts. Many staff come from these programs.

The national laboratories face many challenges in maintaining the workforce. They may want to work together to make recruiting more successful. Career paths need to be re-examined to offer competitive choices and to provide a more attractive workplace. The current funding model makes it hard for young staff to establish themselves. Opportunities for employees to grow professionally should be considered. Resources should be provided to address work/life balance. The national laboratories should adapt to the shorter-term commitment (higher workforce turnover) that is becoming common in this field. Mid-career entry to the national laboratories should be facilitated, and the national laboratories should engage in education.

The Subcommittee's major findings were that

- The multidisciplinary national laboratories face workforce recruiting and retention challenges in computing sciences.
- Insufficient educational opportunities are available at academic institutions.
- There is a growing demand for graduates in Computing Sciences that far exceeds the supply.
- The exemplary DOE CSGF program has been deemed highly effective in multiple reviews.
- The DOE national laboratories have individually developed measures to help recruitment and retention, yet they need help at the national level.

The Subcommittee developed the following recommendations:

- Investment should be made to preserve and increase the DOE CSGF program for more high-quality students, particularly students from underrepresented populations and demographics.
- New fellowship programs, modeled after the CSGF program, should be established for research opportunities in enabling technologies in the computing sciences.
- Support for local laboratory programs should be expanded, and greater inter-laboratory sharing of information about locally successful programs and workforce-related data should be encouraged.
- A DOE-funded Computing Leadership graduate curriculum advisory group should be established to spearhead participation in efforts within the Association for Computing Machinery (ACM), Computing Research Association, and NSF to develop and annually publish competencies of DOE need at the graduate and undergraduate level.
- Working with ACM's Special Interest Group on High-Performance Computing, NSF, and other organizations, DOE should provide a rich repository of mission-oriented learning materials and engagement opportunities to attract and guide individuals towards careers in areas of DOE need.
- Working with other agencies and organizations, certificate programs should be established to address the need for competency certification.
- Attractiveness of DOE opportunities should be improved with continued relocation assistance, ongoing professional development and position rotation, and a sabbatical program for DOE employees.
- Awareness of DOE opportunities should be increased.
- Working with other agencies, DOE should develop a strategic plan with programs and incentives to pro-actively recruit, mentor, and sustain the involvement of significantly more women, minorities, people with disabilities, and other underrepresented populations.

Berzins said that these were all well-laid-out recommendations, but it all comes down to more money. If one does not compete with salaries, one does not get the workers. Chapman responded that there is a passion for science that is reflected at the national laboratories that does not exist elsewhere. One has to tell prospective employees about these opportunities. Crivelli said that students have commented how one can work on a variety of things at the national laboratories and not be focused on one narrow topic. Chan added that not all students are focused on money; they are influenced by a variety of motivations.

Cerf noted that he worked for Google. He had two points: (1) If one is not going to learn new things, one will not stick around. (2) Students need to be enticed into a computing career early-on, maybe in middle school.

The floor was opened for public comment. There being none, the meeting was adjourned at 3:56 p.m.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
Dec. 9, 2014

ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE

Meeting Minutes
November 21, 2014

I hereby certify that these minutes constitute a reasonably comprehensive and accurate record of the meeting of the Advanced Scientific Computing Advisory Committee held on November 21, 2014.

Roscoe C. Giles

Roscoe Giles, Chair

1-9-2015

Date