Draft Minutes
Advanced Scientific Computing Advisory Committee (ASCAC)
March 24, 2015
American Geophysical Union
Washington, D.C.

ASCAC Members Present

Martin Berzins
Vinton Cerf
Jacqueline Chen
Silvia Crivelli
John Dolbow
Jack Dongarra
Roscoe Giles (Chair)
Anthony Hey (via telephone)
Juan Meza (via telephone)
John Negele (via telephone)
Linda Petzold
Daniel Reed
Dean Williams

ASCAC Members Absent

Barbara Chapman
Sharon Glotzer
Gwendolyn Huntoon
Vivek Sarkar

Also Participating

Christine Chalk, Designated Federal Official for the ASCAC
Gregory Bell, Director, Scientific Networking Division, Lawrence Berkeley National Laboratory
John Steve Binkley, Associate Director, Office of Advanced Scientific Computing, Office of Science, USDOE
Laura Biven, Science and Technology Advisor, Office of the Deputy Director for Science Programs, Office of Science, USDOE
Keri Cagle, Oak Ridge Institute for Science and Energy
Jody Crisp, Oak Ridge Institute for Science and Energy
Melinda Comfort, Attorney-Advisor, Office of the General Counsel, USDOE
Mattan Erez, Electrical and Computer Engineering Department, University of Texas at Austin
Robert French, High-Performance Computing Support Specialist, Oak Ridge National Laboratory
Carolyn Lauzon, Program Manager, Leadership Computing Challenges, Office of Advanced Scientific Computing, Office of Science, USDOE
Steven Lee, Program Manager, SciDAC Institutes, Office of Advanced Scientific Computing, Office of Science, USDOE
Jeffrey Nichols, Associate Laboratory Director, Computing and Computational Sciences Directorate, Oak Ridge National Laboratory
Lucy Nowell, Program Manager, Data and Visualization, Office of Advanced Scientific Computing, Office of Science, USDOE
Frederick O'Hara, ASCAC Recording Secretary
Franklin Orr, Under Secretary of Energy for Science and Energy, USDOE
Douglas Ray, Assistant Laboratory Director, Fundamental and Computational Sciences Directorate, Pacific Northwest National Laboratory

Morning Session

Before the meeting started, **Melinda Comfort**, Office of the General Counsel, USDOE, presented an ethics briefing to the Committee members present at the meeting who were not federal employees and one new member was sworn in.

The meeting was called to order at 9:01 a.m. by the Chairman, **Roscoe Giles**. He noted that the FY16 budget request had many references to Advanced Scientific Computing Advisory Committee (ASCAC) recommendations. Giles had testified before Congress on next-generation supercomputing and was well received.

Jody Crisp made safety and convenience announcements.

Franklin Orr was introduced and immediately thanked Roscoe Giles for his 15 years of service on the Committee, to Patricia Dehmer for her leadership over the years within the Office of Science (SC), and to Stephen Brinkley for his management of the Office of Advanced Scientific Computing (ASCR).

The nation and the world are facing an energy and climate challenge. The future climate needs to be projected, and that cannot be done without high-performance computing. Energy needs to be transformed to energy services through the production, distribution, and use of energy. More-efficient energy transformations, distribution, and uses are needed. President Obama recognizes this. He has set some greenhouse-gas reduction goals. There is a lot to do, and new tools need to be invested in to achieve these reductions. Secretary Moniz organized the Department of Energy (DOE) in three parts: the nuclear deterrent, cleanup of the environment, and science and energy. Within Science and Energy, the applied and fundamental science programs need to be linked. High-performance computing permeates all that is done by these programs.

There are 17 national laboratories (13 in Science and Energy); applied programs; crosscutting initiatives (the energy–water nexus, subsurface engineering, supercritical carbon dioxide, cybersecurity, grid modernization, and exascale computing). The *Quadrennial Energy Review* is being published to assess the state-of-the-art and to see what can be done.

The FY16 budget is currently being defended on the Hill. SC is the biggest supporter of physical science in the United States at $5.3 billion, of which $627 million is for ASCR, a 15% increase from last year's request (which does not map one-to-one with the Congressional appropriations). This seems to indicate that many agree that investment is needed in advanced computing.

Computing facilities are an important part of SC. Last year saw two initiatives:

- a five-times improvement in computing performance at the Oak Ridge Leadership Computing Facility (OLCF) with only a 10% increase in electrical power consumption
- $100 million for extreme computing (in the form of FastForward 2)

ASCAC has been important by virtue of the excellent advice that it has provided in years past. Huge progress has been made in high-performance computing: 15 years ago, the National Nuclear Security Administration (NNSA) needed a supercomputer to simulate nuclear tests. ASCR had a $125 million budget. Then the Accelerated Strategic Computing Initiative (ASCI) took place. Today ASCR has a $620 million budget, an increase of almost a factor of 5. Petascale computing, the National Energy Research Supercomputing Center, the Energy Sciences Network (ESnet), and two leadership computing facilities have been tremendous accomplishments. All this led to high-performance computing's becoming a major contributor to SC's science efforts. The program Scientific Discovery through Advanced Computing (SciDAC) has been operating since 2004 helping SC do science.

The Department looks forward to ASCAC's future advice in modeling and simulation for materials development, combustion, and the full range of energy sciences. Those sciences include climate change, which requires subgrid physical processes and higher resolution, which in turn require bigger, faster computers.

Cerf asked whether the argument will be made that investment in scientific computing will pay off in the commercial sector. Orr replied, yes. That argument was made before Congress during the week before this meeting. The government pays vendors to develop advanced computer technology that meets DOE mission needs but for which there is no commercial market, as yet. Cerf asked if the Office still worked with the National Information Technology R&D (NITRD) Program. Binkley replied, yes. Cerf asked if anything was being learned about fusion from high-performance computing. Orr replied that fusion is an important example where R&D can only be pushed forward by modeling and simulation by high-performance computing. Binkley added that solving the fundamental problems in fusion requires sophisticated physics. Past advances have been the result of modeling and simulation. At ITER [formerly the International Thermonuclear Experimental Reactor], plasma instability is one problem that needs more research and larger simulations. Orr added that the research effort is trying to produce high gradients (e.g., in temperature), which have great instabilities, requiring high-performance computing. On the security side, the need is to compute the performance of complex systems. Giles stated that, for SC, a crucial element is the ability of ASCR personnel to work with personnel on the application side. Orr pointed out that that is why the cross-cutting initiatives are being set up, to engender such links and collaborations.

Giles pointed out that developing the workforce is important across the government. Orr responded that there are efforts going on to educate the workforce, but not enough such efforts.

John Steve Binkley was asked to update the Committee on the activities of the Office of Advanced Scientific Computing. He welcomed Daniel Hitchcock back to ASCAC as a visitor and acknowledged Hitchcock's contributions to the Office's accomplishments over the years before his retirement.

SC's user facilities serve 31,000 researchers each year, of which 25% use the computing facilities that ASCR provides. ASCR's capabilities include mathematics research, computer science research, SciDAC partnerships, exascale computing deployment in the next decade, facilities, and a postdoctoral program. ASCR has a strong focus on the exascale, facilities, and dealing with large amounts of scientific data.

In the FY16 budget request to Congress, SC is asking for a 5.3% increase versus the FY15 enacted appropriation. For ASCR, that number is 14.8%, the biggest increase among SC offices; the additional funding is mostly for the exascale program. For ASCR's research programs, funding is projected to be relatively flat for FY15 to FY16; most of the increase is in facilities budget and for exascale initiative.

The exascale computing initiative is one of the crosscuts mentioned by Orr in the previous presentation and is being conducted in partnership with the NNSA. The FY15 budget of $91 million for ASCR's portion is being increased to $177.9 million in FY16, an uplift of $86.9 million; an additional $64 million is coming from the NNSA budget.

In regard to the DOE Computational Science Graduate Fellowship (CSGF), the Office of Management and Budget (OMB) tried to consolidate such programs and do away with the CSGF. But Sec. Moniz wrote to the President's Science Advisor in the Office of Science and Technology Policy, recommending a reconsideration of that budget cut. Now the CSGF has a $10 million budget request being submitted to Congress.

In personnel matters, Marc Kastner has withdrawn his candidacy for the directorship of SC to accept a position running the Science Philanthropy Alliance. Sandy Landsberg is departing DOE to go to the Department of Defense High-Performance Computing Modernization Program; a nationwide search will look for a replacement. Melea Baker is retiring from federal service after 30 years. Robert Lindsay, a program manager in the Research Division, is retiring from federal service. Ceren Susut-Bennett is on a 4-month detail at the National Science Foundation.

The Secretary of Energy Advisory Board (SEAB) looked at exascale computing and filed its report in August 2014; Sec. Moniz has asked three follow-up questions, which will be addressed at the next SEAB meeting on March 31, 2015. The first question was, can the Task Force provide an additional level of granularity on the allocations against major technology areas and their timing that regulate success in getting to the 1- to 10-exascale range in a decadal timeframe? The Task Force envisions that a decade-long program will be required to achieve exascale computing; estimates that the funding profile for exascale will extend through 2024, with peaks in the range about $350 million a year; and notes that the cost of individual exascale computers will be in the range of $200 million to $250 million. The second question was, could the Task Force provide suggestions for what the Department could undertake to expand industrial high-end, high-performance computer use? The Task Force suggested creating and supporting an easy-to-navigate DOE portal describing all publicly available computing resources and programs; continuing to support competitive programs that provide access to leading-edge high-performance computing at DOE; support programs leading to the commercialization of new or matured codes; and being a key partner with the university community, the national academic accreditation bodies, and the private sector in enhancing engineering and science degree programs. The third question asked for further thoughts on how a beyond-exascale research program might be structured. The Task Force identified three promising areas for advancing high-performance computing: quantum computing, superconducting circuits, and neuromorphic computing. It suggested that (1) DOE invest to maintain and strengthen the computational ecosystem, including working with universities, which would allow DOE to understand what already is under way while focusing on more-advanced elements of over-the-horizon computing, including software development; and (2) combining the path to the exascale with investments to sustain the advanced computing ecosystem and to look over the horizon with funding of $100 million to $150 million per year, including $20 million or $25 million per year to enable DOE to stay abreast of developments being sponsored by others.

Computational capacity is based on the needs of the research community, identified through the formal requirements-gathering process. This process needs to be revisited. The purpose of this process is to ensure the ability of ASCR facilities to support SC mission science in the exascale regime. This process is being mapped out and socialized with the domain researchers. A series of workshops will be held, one per domain program, between June 2015 and September 2016. The platform acquisitions have been firmed up in the Exascale Computing Initiative timeline: CORAL [Collaboration of Oar Ridge, Argonne, and Livermore] at the beginning of FY18, APEX [Autonomic Performance Environment for eXascale] at the end of FY20, and exascale at the beginning of FY24.

In the near term for the Exascale Computing Initiative (FY15), ASCR will carry out reviews; initiate rapid requirements assessments, seeking broader application of high-performance computing across the government; seek CD-0 [Critical Decision Zero] approval; prepare the FY17 Exascale Computing Initiative budget request; conduct a second external agency review of the initiative; have a red team review of the CD-1 [Critical Decision One] package; complete the rapid assessments needed for the FY17 budget; submit the FY17 budget to OMB; and release the FY16 funding opportunity announcements (FOAs).

The Office is on the cusp of announcing the Argonne National Laboratory (ANL) component of the CORAL acquisitions for upgrading the leadership computers. In the Oak Ridge National Laboratory (ORNL) Summit computer, a CORAL rack has 779 TF, operating at 55 kW; the system as a whole has about 200 of these racks, standing on a 6 ft^2 footprint. The NERSC-8 Cori delivery to the National Energy Research Scientific Computing Center (NERSC) is expected in 2016. It is a 64-cabinet Cray XC system that will have 10 times the Hopper sustained performance with about 9300 Knights Landing compute nodes, about 1900 Haswell compute nodes, an Aries interconnect, a Lustre file system, and a nonvolatile random access memory (NVRAM) burst buffer. [The system is named after Gerty Cori, the first American woman to be awarded a Nobel Prize in science.

The Lawrence Berkeley National Laboratory (LBNL) Computational Research and Theory Building is being constructed with University of California at Berkeley funds with probable occupancy in May 2015.

The question arises, what comes after the exascale? The fundamental limits (11 nm) of the lithographic features of complementary metal oxide semiconductors (CMOSs) are being approached. There are lots of issues about where chips will be made. Technology options include quantum computing, neuromorphic computing, probabilistic computing, and other technologies. There have been many recent ASCR workshops on these topics. The workshop on quantum computing noted that quantum computing technologies are making rapid progress, and algorithms are coming for quantum-computing calculations. Several community activities have addressed the usefulness of quantum computing and quantum devices. Quantum computing is different because it requires a variety of disciplines: computer science, high-energy physics, and materials science. There is a confluence of knowledge between quantum computing and particle physics. A lot of work is going on about sensors, cryptography, communications, networking, metrology/measurement, and accurate timekeeping. There is also work being done on new materials for quantum devices, cubit manipulation, error correction, system integration, and software.

Cerf pointed out that Google has a joint program with the National Aeronautics and Space Administration’s Ames Laboratory on quantum computing and the development of algorithms. He asked if graphic processing units (GPUs) have any further future in high-performance computing. Binkley replied, yes. They are easier to program and provide an adequate speed-up of 3 to 5 times. Cerf pointed out that, at the National Ignition Facility (NIF), industry is brought in and provided computational services that industry could not afford on its own. He asked if DOE could do that. Binkley answered, yes. The LCFs already allocate time to industry on a competitive basis. There are several pathways for industry to gain access to these world-class machines. Cerf noted that prototyping needs to be iterative. Binkley agreed and pointed out that ASCR projects prototype at the node level and then at successively higher levels.

Dongarra asked how much of the budget will go to quantum computing and other next-generation computers. Binkley replied, $5 million, depending on how Congress reacts to the President’s request. The major funders of quantum computing are in the Department of Defense, which provides $15 to 18 million. Also, several national laboratories have laboratory-directed research and development (LDRD) investments in quantum computing.

The floor was opened to public comment. There being none, a break was declared at 10:36 a.m. The meeting was called back into session at 10:47 a.m.

Gregory Bell was invited to give an update on ESnet.

In 1986, DOE wanted ESnet to be a big local area network, performing as if there were only one user. ESnet is now a high-speed national network, optimized for DOE science missions. It connects 40 laboratories, plants, and facilities with more than 100 university and facility networks. It just completed a high-speed European connection. Its budget was $32.6 million in FY14, and it employs 42 full-time-equivalent employees. It is older than the commercial Internet, and it is growing twice as fast. It is the DOE user facility that serves all the others. Its fiber assets and access to spectrum are shared with Internet2, but ESnet is designed for different goals than the general Internet. It is developed for elephant, not mouse, flows. With elephant flows, almost lossless networks are required.

With ESnet, scientific progress is completely unconstrained by the physical location of instruments, people, computational resources, or data. The science data transferred each month has seen exponential growth since 1990. These data come from the Spallation Neutron Source (SNS), fusion facilities, light sources, and small devices routinely operated across the Office of Basic Energy Sciences (BES) network. By 2020, everyone will have a computer in his or her pocket, and scientists will have a sensor and computer.

Today, 80% of ESnet traffic originates or terminates outside the DOE complex. The Large Hadron Collider (LHC) is now moving to a federated model; ESnet is expected to support this evolving data model in which 30 nations are participating. There are dozens of universities lined up with 100-GB connections. This network is now spilling outside the Office of High Energy Physics (HEP) facilities (e.g., to light sources). An example of super facilities that are connected is the connection between the Linac Coherent Light Source (LCLS) and NERSC. The data flow from a single LCLS detector triples the network utilization for a major high-performance computing center. It would be good to automate the setting up of the workflow.

A 340-Gbps extension has been built between the United States and Europe. The immediate driver is more capacity for LHC Run 2. However, the extension supports all DOE missions, reducing barriers to European collaborators and instruments. It will accommodate ITER and light sources in Europe. The optical fiber of the Internet crosses land and the ocean floors and parallels the Eastern Telegraph Company's 1858 system and its general connections. There are many things that break underwater cables, so there are four high-speed parallel cables.

ESnet operations focus on simplicity, automation, and cooperation. It is decreasing monitored hosts and increasing auto-patched hosts while decreasing operating system versions. Its portal has now been greatly enhanced, allowing one to dial into any facility to see the amount of traffic.

Many honors were garnered by the ESnet5 for its software, testbed work, and technology. Staffing remains lean despite all these honors. The organization is highly diverse (half of ESnet managers are women), except in network engineering, where an effort is being made to increase diversity. It sponsored two early career women from national laboratories to attend an important annual conference for network engineers. It is also sponsoring the Grace Hopper Conference for Women's Career Development.

ESnet's vision and strategic goals guide impacts. The vision is that discovery is to be unconstrained by geography. The strategic goals are to improve networking practices globally; to provide information and tools for optimal network use; and to perform R&D in pioneer architectures, protocols, and applications.

SC invests nearly $1 billion per year in university research, so campus networks matter to DOE. A DMZ [demilitarized zone] is a firewall configuration that adds an extra level of security to a local-area network. Super DMZ as a service would be essentially a super facility on demand. Therefore, Science DMZ has been developed. It is a network-design pattern for data-intensive science; it had its origin in ESnet and NERSC. Three components are required: a friction-free network path, dedicated systems and data-transfer nodes, and performance-measurement nodes called perfSONAR. Science DMZ is now recognized as a best practice. The National Science Foundation (NSF) is investing $60 million to promote adoption by U.S. universities, and more than 120 universities in the United States have deployed this DOE architecture. In addition, the U.S. Department of Agriculture and the National Institutes of Health are supported, and the National Aeronautics and Space Administration (NASA) and the National Oceanic and Atmospheric Administration (NOAA) are investigating the possibility of joining. Australian and Canadian universities are following suit.

The evolution of Science DMZ as a regional cyberinfrastructure platform was recently announced. The Pacific Research Platform Initiative is the first large-scale effort to coordinate and integrate Science DMZs. All major California research and education institutions are now participating. The next step will be national integration and then international integration.

ESnet has had major impacts on the broader networking community:

- A “fasterdata” knowledgebase is being developed.
- Sample campus and regional cyberinfrastructure plans are being posted on the web.
- The production software OSCARS has been adopted by more than 40 networks in universities.
- perfSONAR has had 1200 deployments worldwide (with other, unregistered users).
- With NSF support, science DMZ has been deployed on more than 120 campuses in the United States alone.
- ESnet is frequently among the first customers for new technology (e.g., the 100-GB transatlantic network and the new 400-GB terrestrial lines).
- ESnet has demonstrated the first software-defined networking (SDN) network for optical transport, which resulted in a new product for Infinera.
- It was the first public customer for disruptive white-box networking gear.
- It has worked closely with industry to develop a packet-processing pipeline useful to DOE science missions.
- Its testbed has been available to researchers and industry since 2011 and has been used to prototype ESnet6.
- It is engaged in a federation of other testbeds and has a platform for collaborations and demonstrations.

Cerf pointed out that there are latency problems with the network (e.g., between Los Angeles and San Diego) on the order of tens of milliseconds and asked how ESnet was going to deal with such latency. Bell replied that there is no way to avoid latency. One needs to identify local problems and deal with them. Cerf said that one can use a hollow fiber in which the light would travel through a vacuum at twice the speed that it travels through the glass.

Reed asked Bell if he could say something about spectrum management at the transition from public to private networks. Bell admitted that ESnet did not have expertise in this area. The data are converted to fiber transmission as soon as possible. Reed replied that ESnet needs to talk to the Federal Communications Commission about licensed and unlicensed spectrum management.

Giles asked if ESnet’s planning was well integrated with new facilities. Bell replied that ESnet works through the DOE offices, conducting requirement reviews every 3 years. This process gives ESnet a good idea of what the offices’ needs will be.

The floor was opened to public comment. Douglas Ray asked about the likelihood of Europe and Japan adopting Science DMZ. Bell replied that the Large Hadron Collider (LHC) has had an integrating effect in Europe and will probably have a similar effect on Japan.

Giles asked if ESnet had any thoughts about the Square Kilometer Array (SKA), which involves countries that are not necessarily parts of the networking framework. Bell replied that it is an amazing project. DOE and NSF are not officially participating now but may in the future. This global-scale radio telescope uses 10 million 10-Gb links to bring data from the detectors back to Jodrell Bank Observatory. ESnet hopes to keep informed about what the South African Research Network (SARN) is doing until ESnet becomes directly involved. Dongarra interjected that SARN is at 1 EB per hour in data generation.

Daniel Reed was asked for an update on the exascale.

- A subcommittee has been formed.
- Meetings are being scheduled.
- A teleconference is being scheduled to discuss proposals in more detail.
- A final report is due in September for informing the FY17 budget.

Giles asked if there were going to be a June red-team review of exascale activities. Binkley replied that the appropriate personnel need to get together and talk through how the two reviews (the red team's and the Subcommittee's) will cooperate.

A break for lunch was declared at 11:45 a.m.

Afternoon Session

The meeting was called back into session at 1:33 p.m.

Anthony Hey was asked to present an update on the ASCAC Subcommittee associated with the Office of Scientific and Technical Information (OSTI).

ASCAC received a charge to establish a standing subcommittee to advise SC about "matters associated with the DOE Office of Science and Technical Information (OSTI)," which was established in 1947 to collect, preserve, and disseminate scientific and technical information from the R&D activities of DOE and its predecessor agencies. SC is looking for external, independent advice as OSTI transitions its products and services to methods appropriate to the new era of information gathering and sharing. The subcommittee is to examine four issues:

- Are OSTI's products and services best in class, and do they fulfill the most critical needs?
- Do they meet customer needs now?
- Are they positioned to meet customer needs in the future?
- What is the national and international standing of OSTI, and in what areas should OSTI be a clear leader to fulfill its DOE-mandated responsibilities?

The subcommittee has been formed, and its report is due at the summer ASCAC meeting.

DOE has launched the Public Access Gateway for Energy and Science (PAGES), a web-based portal that will provide free public access to DOE-funded, accepted peer-reviewed manuscripts or published scientific journal articles within 12 months of publication. As it grows in content, PAGES will include access to DOE-funded authors' accepted manuscripts (hosted primarily by DOE's national laboratories and grantee institutions) in addition to the public-access offerings of publishers. For publisher-hosted content, DOE is collaborating with the publisher consortium CHORUS, the Clearinghouse for the Open Research of the United States.

The open-access policy relies on publisher CHORUS infrastructure. The backup consists of copies of publications and repositories, but deposit is not mandated. It only allows fair use, not text mining or translation. It is a possible alternative to the Shared Access Research Ecosystem (SHARE), a higher-education and research-community initiative to ensure the preservation of, access to, and reuse of research outputs. SHARE was developed by the Association of Research Libraries, the Association of American Universities, and the Association of Public and Land-Grant Universities.

Binkley noted that the open-access policy was set up to keep a “dark copy,” the deposition of which is mandated. Biven added that deposition is mandated under the terms of grants made to R&D funding awardees.

The Office of Science and Technology Policy (OSTP) also issued a memo that refers to R&D data. Now the archiving of data is the responsibility of the program offices. The linkage of publications and the underlying data is a challenge.

Brian Hitson of OSTI made a presentation to the Subcommittee during the week before this meeting. It covered the history, authorization, and structure of the workflow for gathering publications from national laboratories and researchers. He pointed out that DOE had to comply with copyright law under the OSTP memo and that OSTI was waiting for SHARE to reach maturity before instituting full cooperation.

Hey hoped to have the full Subcommittee visit OSTI in Oak Ridge and prepare its report by the summer ASCAC Meeting. OSTI holds some classified weapons research (about 10% of its holdings), and Dean Williams, who has the necessary clearance, will be responsible for reviewing that portion of OSTI’s operations.

Cerf noted that arXiv and CHORUS overlap in subject matter coverage. Hey replied that CHORUS has the published version and arXkiv has preprints. Hitson noted that if arXiv hosts the final accepted manuscript, then a DOE-funded author can provide OSTI the arXiv URL to the accepted manuscript along with other metadata, and OSTI will harvest that accepted manuscript for dark archiving purposes. A preprint is not a substitute for the final, peer-reviewed accepted manuscript. If it is in CHORUS, OSTI will harvest the final preprint from arXiv for its dark archive.

Cerf asked what relationship DOE has with the Research Data Alliance (RDA). Hey replied that he could not point to any great connection. RDA is in an early stage of development, and he was on its Council. Biven added that no DOE people are involved with the management of RDA, but DOE is talking with the NSF about managerial participation. Nowell further added that DOE has some senior laboratory people serving on technical committees and tasked to keep abreast of RDA to let the Office know when it needs to take action.

Crivelli noted that publications are straightforward to deal with but that data are a different matter. Hey responded that there can be huge amounts of data; one does not want to save all the data; the question is how to choose what should be saved. The grant holder is currently being asked to choose what to save. It is very complicated. RDA is trying to grapple with this issue. Giles asked if the Subcommittee’s report would cover this issue. Hey replied, yes, especially the linking of the publications and supporting data. Another issue that needs to be dealt with is the permanence of links.

Crivelli noted that linking data to the software with which it was been created has been talked about. At this point Hey stated that the conference center in the United Kingdom where he was attending a conference was being closed and he had to leave. Giles said that this issue could be addressed later. Cerf pointed out that Mahadev Satyanarayanan at Carnegie Mellon University has been doing good work on rendering data as part of Project OLIVE [Open Library of Images for Virtualized Execution]. It would be wise if the effort were extended. Preserving the original software for rendering may not work over the long term.

Lucy Nowell was asked to report on data management, analysis, and visualization.

The Computer Science Division of ASCR announced an FOA on Scientific Data Management, Analysis, and Visualization [SDMAV] at the Extreme Scale. The FOA was informed by three ASCR reports:

- The *Data Crosscutting Requirements Review*, which identified the research challenges of maximizing data's analytic value while minimizing power and memory costs, maintaining data provenance, using input/output and memory to support the reuse and repurposing of data, comparing simulation data to observational data, richly representing scientific semantics, generating concurrency and masking latency, and mitigating hardware failures and faults.
- The *ASCAC Data Subcommittee Report*, which discussed the natural synergies among the challenges facing data-intensive science and exascale computing.
- *Scientific Discovery at the Exascale*, which found that the disruptive changes posed by a progressive movement toward the exascale and high-performance computing threatened to derail the scientific discovery process because today's success in extracting knowledge from large simulations will not be generally applicable to the exascale.

The FOA addressed five themes:

- Usability and user interface design
- In situ methods for data management, analysis, and visualization
- Design of in situ workflows to support data management, processing, analysis, and visualization
- New approaches to scalable interactive visual analytic environments
- The need to develop proxy applications or workflows and/or simulations for data management, analysis, and visualization software to support the co-design of extreme-scale systems

The review criteria were scientific and/or technical merit, appropriateness, competency of applicant's personnel and adequacy of the proposed resources, reasonableness and appropriateness of the proposed budget, and relevance to the mission of ASCR. The overall budget was approximately $4 million per year for 3 years with the potential to increase it to $7 million per year, depending on the proposal quality. Each funded project could have a total annual budget of from $150,000 for a small project in a single institution to $1.5 million for a project spanning a larger portion of the scope of research in multiple institutions. An award to non-national-laboratory applicants could not exceed $500,000 per year for 3 years if partnered with a national laboratory or $350,000 per year if not so partnered.

Seventy-four unique pre-proposals were received; 39 projects/pre-proposals were encouraged; 35 encouraged projects' proposals were received and reviewed; and 9 projects were recommended for award (3 university led and 6 national-laboratory led). The total funding awarded was $24.852 million spanning 3 years. About two-thirds went to national laboratories, one-quarter to universities, and the remainder to industry. The funded projects were

- Usable Data Abstractions for Next-Generation Scientific Workflows
- Optimizing the Energy Usage and Cognitive Value of Extreme Scale Data Analysis Approaches
- Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery

- A Unified Data-Driven Approach for Programming and In Situ Analysis and Visualization
- XVis: Visualization for the Extreme-Scale Scientific-Computation Ecosystem
- High-Performance Decoupling of Tightly Coupled Data Flows
- In Situ Indexing and Query Processing of Adaptive-Mesh-Refinement Data
- Performance Understanding and Analysis for Excess Scale Data Management Workflows
- Extreme-Scale Distribution-Based Data Analysis

In other activities, the Division held a day-long event, Storage Systems and I/O (SSIO) Summit, on September 15, 2014, in Rockville, Md., to see what to do to reinvigorate research in this area. The workshop found that in situ data analysis is important; new technology is complicating the storage hierarchy; a coherent view and management methods of storage resources are needed; SSIO designs are hindered by their isolation from system-level resource management, monitoring, and workflow systems; application/system SSIO behavior is not well understood; results validation may change the role of SSIO systems; new programming models and systems drive new persistence mechanisms; increasingly complex data abstractions are being sought by users; and community access to data on applications and systems is needed. The workshop identified a series of near-term research priorities in SSIO architectures, metadata, name spaces, provenance, science-based support, and SSIO understanding.

The Division also sponsored a Data Council Meeting on September 16-17, 2014, in Rockville, Md., which focused on the exascale preliminary planning documents on data management and data analysis and visualization. The workshop devolved community concurrence with the Exascale Preliminary Program Design Document, took steps toward defining the data stack for the exascale, and made recommendations on coordination across the portfolio.

A follow-up series of workshops was held on SSIO:

- Science Requirements Review on December 8 focused on scientific-use cases for SSIO and patterns of data movement
- Crosscutting Computer Science Review on December 9 focused on identifying dependencies across areas and coordination needed to see that needs are met
- Storage Systems and I/O Workshop on December 10-11, which was charged to assess the state-of-the-art, the research needed to address identified requirements, and approaches to reinvigorating the field
- The Burst Buffers Workshop on December 11 focused on DOE requirements and how to address them in future vendor requests for proposals

The Division sponsored a data/visualization principal investigator (PI) meeting on January 13-15, 2015, in Walnut Creek, Calif., with about 80 participants from 20 projects. It examined what technologies could be provided to the exascale in such topics as data movement, programming models, data reduction, run time, power reduction, and burst buffers.

The Division is currently planning for anticipated announcements in SSIO; providing coordination to ensure that SDMAV dependencies are met; holding a workshop on workflows on April 20-21, 2015, to address end-to-end workflows and the need for an in situ workflow management system; organizing workshops on requirements for experimental and observational science; and planning for the 2016 SDMAV PI meeting.

There are so many reports of this topic that it is a challenge to review all of them.

Chen asked how one can unify the analysis, input/output, solvers, etc. in one system. Nowell said that that is why an in situ workflow-management system is needed. This is a new area. The workflows are currently embedded in the codes or use a Python script. There are more questions than answers now. A workshop on provenance may also be needed. There are several workshops sponsored by DOE and NSF coming up.

Berzins asked if there were any examples of provenance models for a multiyear project. Nowell noted that one research project spends more money on a provenance plan than ASCR spends on its entire multiyear plan. One needs to identify what provenance to capture and how to use it.

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

Mattan Erez was asked to describe his Early Career Award research on balancing correctness, performance, and efficiency with containment-domain resilience.

The constraints that one needs to deal with are power/energy, time, cost and money, and correctness of results. Resilience is a big challenge for DOE computations because, with an exascale system, something bad happens about every minute (vs every year in a commercial system), and the trend in time to failure is downward (shorter and shorter).

The baseline process for controlling such faults and errors is checkpoint-restart, which is not good enough for large systems and codes. Efficiency drops off quickly, and correctness is at risk. As the overhead goes up, energy becomes problematic. The cost of resilience lies in preparation, detection, mitigation, and implementation.

There are three paths to addressing resilience: software, hardware, and algorithms. Hardware can fix the problem [with chipkill-correct and chipkill-detect memory or with pipeline silent data corruption (SDC) reduction] but is costly. Algorithmic detection uses iterative converging algorithms, redundant information, and probabilistic methods, but different applications and scales require different techniques, and one does not want to write a program over and over to accommodate those requirements.

Instead, one needs to adapt or co-tune. Containment domains elevate resilience to first-class extraction, and one can develop program-structure abstractions, composable resilient program components, regimented development flow, and supporting tools and mechanisms. One ends up with containment domains that are abstract resilience constructs that span system layers, have hierarchical and distributed operation for locality, are scalable to large systems with high energy efficiency, are heterogeneous to match disparate error/failure effects, are proportional and effectively balanced among costs, have tunable resilience specialized to an application or system, and are analyzable and auto-tuned.

Containment domains embed resilience within an application. Semantics are simple because erroneous data are never communicated outside the domain, and each domain provides a recovery mechanism. A containment domain preserves data, computes, detects faults, and recovers from errors.

The focus of this work is on the programming and execution model support. Containment domains manage preservation, restoration, and re-execution. Specific policies can be written by the user. Containment domain abstraction is amenable to analysis and auto-tuning.

In using containment domains, one annotates the code, profiles and extrapolates a containment domain tree, supplies machine characteristics, analyzes and auto-tunes, refines trade-offs, repeats, and then executes and monitors how it operates in terms of efficiency. Containment-domain annotations express intent for the domain with a containment-domain hierarchy for scoping and consistency, preservation directives exploiting locality, correctness abstractions, recovery customization, and a debug/test interface.

State preservation and restoration application programming interfaces (APIs) can be hierarchical (limited to a given level) and proportional (preserving only when preservation is worth it). Correctness abstractions can cover detectors, requirements, and recovery.

Faults and failures can result from application crashes, process crashes, process unresponsiveness, failed communications, hardware failures, lost resources, wrong values, and degraded resources. These errors are detected by system-provided detectors and user-specified detectors, producing a consistent and unified reporting and analysis system. There are APIs for debugging, testing, and tools.

Progress has been made on the execution model. Building systems has been found to be hard and tricky. There has been a limited release of a single-node runtime; Message Passing Interface (MPI) runtime is very close. The code is open source and will soon be available on Bitbucket. Useful collaborations are already in progress with FastForward 2, DEGAS [Dynamic Exascale Global Address Space], Legion and Swarm, Stanford's PSAAP [Predictive Science Academic Alliance Program], and LBNL's TOORSES [Towards Optimal Order Resilient Solvers at Extreme Scale]. Autotuned containment domains perform well in several codes that have been tested and have improved energy efficiency at scale.

Containment domains should be used when errors may corrupt results and failures kill applications; when the error rate becomes very high, requiring detection and recovery; when the failure rate becomes very high, requiring specialized preservation and recovery; and when problems require a general, analyzable, and opposable solution.

Giles stated that most people still think of computers as being perfect, but they are not. He asked if the mechanisms that detect faults are perfect. Erez replied that detection is built-in on the hardware side, but verification is also built-in to the software to make sure data are within expected bounds. Rerun can be undertaken. Multiple faults can point to a problem that needs to be addressed. Giles said that containment domains reduce intrinsic errors that can be trapped in a domain. Erez added that they also allow correction.

Chen asked what tuning and mapping were. Erez replied that tasked base models can have hierarchies that raise efficiency. One can also detect errors and make detection tunable (i.e., deciding what one does and does not need to make copies and maps of). If one keeps a copy, that copy and the map of it can provide a recovery point that the recovery code can jump to. Autotuning and mapping go together and work well for programs that are not too dynamic. There is not yet a process for programs that are highly dynamic.

Cerf asked if Erez were familiar with sharding and suggested that it might be interesting to compare containment domains with sharding. Erez said that he was familiar with sharding but pointed out that one does not want to distribute errors to all the neighbors, as occurs with sharding, and one does not want to make so many copies as to overwhelm the system. Rather, containment domains seek an intermediate solution. Cerf said that Erez equated "intent" with "preservation of data." Erez added: also responses to errors. Cerf asked that he think about other assertions about the program (like intent), as well.

Petzold agreed that this approach would work for academic codes but wondered whether it would work with industrial-strength codes. She asked if one could put limits on domains. Erez said that that was a great idea but that it was beyond the research team right now.

Steven Lee was asked to report on the SciDAC-3 Committee of Visitors (COV).

A COV was charged to look at processes used by SciDAC-3 and to assess the breadth and depth of the portfolio elements. The COV report was delivered November 21, 2014, and the response by the Office was posted January 5, 2015. The major comment by the COV was that

“SciDAC remains the gold standard for fostering interaction between disciplinary scientists and [high-performance computing]. The program managers are to be commended on continuing the excellence of the SciDAC ‘brand’.” The COV made eight recommendations.

Recommendation 1 was about preserving decision documents and providing summary feedback in any declination letter. ASCR agrees with this recommendation. The Portfolio Analysis and Management System (PAMS) has recently been developed and employed to support and document the complete research-funding process for SC research programs, including SciDAC. Decision documents for declined proposals are in PAMS.

Recommendation 2 spoke to the importance of the ability of program managers to impose a SciDAC priority filter over and above the peer reviewers. ASCR agrees with this recommendation. The overall quality of the SciDAC program relies on the careful management of the solicitation, review, and selection process relative to each science discipline.

Recommendation 3 said that there should be coordination between the science programs and ASCR priorities in timing decisions pertaining to future proposals. ASCR agrees with this recommendation. Close coordination and communication among SciDAC program managers has been essential in managing this complex program and will be maintained

Recommendation 4 recommended maintaining or creating an appropriately balanced emphasis on science-based algorithms and insights, mathematical/computational algorithms, and high-performance computing. ASCR agrees with this recommendation. The SciDAC program will continue to balance its portfolio of high-performance algorithms and software to address the strategic research priorities of SC.

Recommendation 5 urged ASCR to pursue synergisms between SciDAC and co-design. ASCR agrees with this recommendation. Scalability and architecture awareness are primary characteristics of SciDAC-3 software and science applications. ASCR is always looking for ways to prepare SciDAC for future architectures and to benefit from leveraging results from ASCR research projects.

Recommendation 6 suggested that wide adoption in the field of codes developed by the institutes should be regarded at least as meritorious as shared postdoctoral funding in that it shows that the algorithmic and software technology has reached maturity. ASCR agrees with this recommendation. The wide adoption of codes produced by SciDAC projects continues to be one of ASCR’s success stories.

Recommendation 7 strongly encouraged the institutes to expand outreach efforts in the out years of SciDAC-3 to reach a larger scientific community. ASCR agrees with this recommendation. The SciDAC institutes are actively involved in expanding outreach to the wider computational science community through annual summer schools, extensive tutorials, and research-project collaborations.

Recommendation 8 cautioned ASCR to be attentive that the balance between the ASCR Leadership Computing Challenge (ALCC) and the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program computing resources is tuned in light of SciDAC requirements. ASCR agrees with this recommendation. Sufficient access to advanced scientific computing resources is essential to the successes of the SciDAC program. ASCR can address this risk when considering its computing resource allocation policies in FY16.

Giles stated that there were exciting things going on in using advanced computing to advance science.

Chen asked what the mechanisms of interaction were among institutes. Lee replied that ASCR has not been able to leverage what is coming out of the codesign centers. Chen followed up by asking about the future. Lee responded that SciDAC-4 was being planned, but the future of the codesign centers was unknown. Binkley added that the codesign centers probably will not be ended but may be refocused. Cerf noted that William Kahan of the University of California at Berkeley has been teaching accuracy for 40 years and asked if anyone in SciDAC looked at such assessments of accuracy. Petzold said that that is what numerical analysts do. Giles concurred that that is done outside SciDAC.

Lee continued his report. The COV voiced 11 findings:

- The timing of the calls for institute proposals and the interrelated partnership proposals is a challenge. Asking the program managers 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 program managers are to be commended for having the courage to re-compete the Data Institute rather than accepting a sub-optimal solution among the original proposals.
- Process documentation has much improved since the last review in 2007.
- Projects are well monitored by program managers through frequent teleconferences.
- PI meetings are an excellent mechanism for oversight.
- The program managers seem to be able to work together very effectively in supporting the projects. Negotiations among program managers were essential, and positive solutions were readily achieved.
- The communication and interaction of program managers with the complex teams that are involved is essential. The level of interaction of the program managers with the teams is commendable.
- The ability of program managers to travel to project meetings and conferences is important but is currently insufficient. Current travel support is inadequate.
- The program was adaptive to changing circumstances. For example, when one PI became ill, there was an intervention that resulted in a two-PI arrangement that worked very well.
- SciDAC remains the gold standard for fostering interaction between disciplinary scientists and high-performance computing. The program managers are to be commended on continuing the excellence of the SciDAC "brand."
- Informal conversations of the reviewers with overseas colleagues indicate that SciDAC is seen as a model program, which foreigners wish could be replicated in their home countries.

The COV wrote a good report, and the recommendations and findings were very helpful.

Giles asked if there were any idea of the timescale for SciDAC-4. Binkley replied that it is an FY17 start. The Office is just assembling its FY17 budget request, and that request will be going to OMB in September. Planning has to be done in close cooperation with the domain offices and has to be done in June. Giles cautioned that there should not be a big gap between SciDAC-3 and SciDAC-4.

Binkley thanked ASCAC for the SciDAC Subcommittee's COV and pointed out the need for a COV of DOE's networking efforts. Giles said that the new request for a networking COV was a standard COV request, and that he would appoint a subcommittee chair in a few days.

A break was declared at 4:16 p.m. The meeting was reconvened at 4:32 p.m.

Carolyn Lauzon and **Robert French** were called upon to describe and demonstrate the Tiny Titan, a desktop demonstration of how parallel computing works.

The Tiny Titan was developed to aid the OLCF's outreach to students and the general public when it was determined anecdotally that K–12 students found the ORNL computing tour boring.

The OCLF’s tradeshow/exhibit outreach was limited to posters and fliers. In general, lay audiences do not understand what it means for computers to collaborate because they have just used one computer at a time and did not have any frame of reference outside high-performance computing. To explain the importance of high-performance computing, one needs to convey how supercomputers break tasks apart, share information with neighboring computers, and work together to speed up solving a problem.

Tiny Titan is interactive; fun for kids; explains high-performance-computing concepts like decomposition, communication, and scaling; and is visually engaging. This tiny supercomputer that was created at ORNL could visit schools and was actually used at the Tennessee Governor’s Schools. It was written up in *Popular Science*. It reached about 150 grade schoolers at ORNL’s “take your child to work” day. It was used in more than 100 tours of the OLCF since June 2014. It has also been displayed and used at the American Museum of Science and Energy. It is portable and has made elementary school, high school, and community college visits.

SC needed a kid-friendly activity for a booth at a children’s science day event. Tiny Titan was the perfect solution.

Tiny Titan was a team build that included members from ASCR, HEP, IT Support, and Office of Nuclear Safety. In an office setting, Tiny Titan is great for engaging visitors and federal agents and contractors. Other opportunities to share Tiny Titan were the May 2015 Science Bowl, in a DOE museum in the Forestall Building, as a do-it-yourself classroom tool, and other outreach activities conducted by SC personnel.

Cerf asked what the parts cost. French replied, about a thousand dollars.

Committee members were offered the opportunity to try out the Tiny Titan, and French led them through the various capabilities of the machine and explained what concepts are demonstrated by the exhibition.

Giles suggested that a Raspberry Pi 2 could bring the cost down to about $100 rather than $1000.

Dolbow asked if the plans were available online and suggested that a Java applet be developed that would emulate a virtual copy of the Tiny Titan.

Cerf suggested programming this machine for Angry Birds or something similar and putting it up on BitHub. French pointed out that one benefit in exhibiting the Tiny Titan is that the task demonstrated is very explicit. One can point to a light or to a box as a representative of data, memory, or communication.

The floor was opened to public comment. Nichols stated that SciDAC-4 in 2017 may be an opportunity for ASCR to connect with BES and the Office of Biological and Environmental Research (BER) and be part of the funding proposed in the President’s request, which has special language designating $19 million for climate modeling and exascale computing and $12 million for computer materials for use in the exascale.

There being no further public comment, the meeting was adjourned at 5:07 p.m.

Respectfully submitted,
Frederick M. O’Hara, Jr.
Recording Secretary
April 6, 2015