Draft ASCAC Facilities Letter

2-28-2013

Introduction

This letter is ASCAC's response to the charge of December 20, 2012 to review and advise the Office of Science with respect to ASCR's plans for new and upgraded major facilities.

ASCAC formed a subcommittee to respond to the charge. The subcommittee included representatives of the current and future user community, experts in scientific computing, and people experienced with comparable facilities outside of the DOE Office of Science.

The subcommittee members were:

Tom Bettge, National Center for Atmospheric Research	Dr. Vincent Chan, General Atomics
Dr. Jackie Chen, Sandia National Laboratories	Dr. Thom H. Dunning, National Center for Supercomputing Applications
Dr. Timothy C. Germann, Los Alamos National Laboratory	Dr. Roscoe Giles (chair), Boston University
Dr. Andrew B. White, (Los Alamos National Laboratory, retired)	

The subcommittee met with the acting ASCR AD, Barbara Helland, and representatives of the major ASCR facilities at their strategic planning meeting on January 30, 2013. This provided an excellent context in which the subcommittee could discuss and review the facilities plans. The subcommittee discussed and prepared its report in several teleconference calls and by email.

Facilities

ASCR computing, networking, storage, software and applications support are key underpinnings of the activities of the Office of Science. The impact of ASCR facilities is realized in the successes of research programs of all the offices in SC. It is important to consider the proper balance between these underpinnings, and realize it changes over time. Identifying application and technology drivers is crucial, and ASCR facilities staff has considerable experience and expertise in this area.

Because of the unique and rapidly changing role of computing and data in all areas of science, we believe that investment in this area is critical to the overall mission of the Office of Science and to DOE and to the nation. The facilities we are commenting on here represent the minimum necessary to support our needs and do not explicitly incorporate a full scale commitment to exascale computing development and deployment as envisioned in ASCAC's Fall 2010 report.

The three major facilities brought to our attention by the ASCR AD reflect a balanced roadmap for upgrading existing ASCR computing capabilities to meet the expected and emerging needs of DOE and the nation's scientists:

1. Upgrading the production computing facility at NERSC, which supports more than 600 projects sponsored by the DOE Office of Science Program Offices.
2. Upgrading the Leadership Computing Facility (LCF) at ANL and ORNL, which advances the frontier of computational science and discovery for the nation and the world.
3. Increasing the network bandwidth of ESnet, which enables the large data flows needed for DOE computing, experiments and analysis in an expanding national and international collaborative environment.

In addition, to meet the emerging critical need to support and develop large-scale data science, we propose adding a fourth facility to the portfolio:

4. A Virtual Data Facility (VDF). This multi-site facility would add huge data storage and analysis resources to the existing ASCR facilities to alleviate the data challenges to all SC programs, and is being considered by the ASCR facilities leaders.

Findings

The table below summarizes our findings; a further discussion of each facility and justification for their categorization are provided in the following sections. Given the rapid pace of technology change, we feel it necessary to distinguish near-term (within 5 years) and far-term (towards the end of the 10-year timeframe covered by this charge) readiness levels, as described below.

Facility	Impact	Readiness (2014-2017)	Readiness (2018-2024)
NERSC	A	A	B
LCF	A	A	B
ESnet	A	A	B+
VDF	A	B+	C

(The classifications used are those described in the charge letter,
Impact: A="absolutely central", B="important", C="lower priority", D="don't know enough yet".
Readiness A="ready to initiate construction", B="significant science/engineering challenges to resolve before initiating construction", C="mission and technical requirements not yet fully defined")

Impact

Each of the four facilities has a key role to play in a balanced ecosystem of DOE high performance computing, and each in its own way contributes in an essential way to DOE's ability to contribute to world leading science. We agree with the AD's assessment that the facilities she identified (1-3) are in the highest "(a): absolutely central" category. We also believe that the proposed Virtual Data Facility is in this category.

1. NERSC is the main engine that supports the breadth of scientific computing for the Office of Science. It provides a broad user base with advanced technology and applications support and is the vehicle by which cutting edge computing technologies enter production.
2. The LCF pioneers the application of extreme scale systems. It helps develop and use the most advanced computing systems for the open science community, including industry, and also works intensively with key user teams to enable breakthrough computations. The lessons learned about large scale computing systems and user support inform NERSC and others about how to broaden and extend the impact of advanced scientific computing to the wider community.
3. ESnet provides the key data linkage for instruments, people, and computational resources. The projected data growth in the next decade is exponential and in some cases faster than Moore's Law. ESnet has a leadership role in delivering highly resilient data transport optimized for large-scale science. Upgrading to 400 Gb/s on the backbone will have a large impact in addressing this challenge.
4. The VDF will provide an integrated focus on data science across all SC computational and experimental facilities. The ASCAC report on data science and exascale computing notes the emerging impact of "big data" on computation, experiment and science as a whole. Key to DOE's leadership in computing is the development of data science at a scale commensurate with the needs of modern experiment, theory, and computation. This is the challenge VDF directly addresses.

Timelines and Readiness

The committee believes that all existing ASCR facilities – NERSC, LCF and ESnet - are ready to upgrade their facilities in the near-term (2014-2017). That is, there are no significant scientific or engineering challenges as yet unresolved. Specifically, the NERSC CRT building is under construction and scheduled for completion in 2015. The CD (Critical Decision) process is underway for power upgrades for LCF to accommodate their next generation systems.

Beyond the near term, there is considerable uncertainty in the performance of systems for a given footprint, power envelope and cost. Therefore, we have divided the 2014-2024 report period into the near term (2014-2017) and long term (2018-2024). The out-year uncertainty is manageable if there is a significant, robust exascale program *addressing issues in* hardware, software and applications.

Additional elements required for effective facilities

Effective computing facilities are comprised of hardware, software, and applications development and support. Support for applications must come from all offices of SC, particularly in light of the ongoing fundamental change in computing and programming, led by the LCF and soon to be embraced by NERSC. Hardware lifecycles are short (3 years) and predictable within known technologies - shorter than the decadal horizon of this report. Application development and support has a significantly more complex and nuanced timeline - starting with early adopters with significant support at the LCF and then developing to include a broader community, including NERSC. It is important to consider all these components in thinking about the timeline.

Broad applications support must take on a new aspect for the future of these facilities. The LCF has very successfully implemented a relatively small collection of important applications on the next generation of energy-efficient, SIMD systems. However, significant additional domain-specific support to migrate the broad range of DOE applications relying on NERSC systems is required for future success. This is a responsibility of SC as a whole, not only of ASCR.

Facilities Details

LCF: History of impact and leadership nationally

The Leadership Computing Facility (LCF) at Argonne National Laboratory (ALCF) and Oak Ridge National Laboratory (OLCF) were established in 2004 with a mission to provide the world's most advanced computational resources to the open science community. Through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) and Advanced Leadership Computing Challenge (ALCC) programs, computational resources are provided to scientists from the research community in industry, academia, and national laboratories. LCF users have continuously achieved numerous wide-ranging research accomplishments and technological innovations. The full breadth and impact of science productivity cannot be adequately discussed in a few paragraphs. But as one measure, OLCF scientists published more than 300 research articles in 2012 alone, where their work is recognized through peer-review publications in high-impact journals such as Science, Nature, and The Proceedings of the National Academy of Sciences (PNAS).

Specific high-impact scientific achievements from the past year include some of largest nuclear structure studies ever performed and the world's largest high-resolution cosmology simulation, modeling over one trillion particles. At OLCF, exploration of the nuclear landscape carried out by INCITE researchers and their theoretical prediction of isotopes was featured in Nature in 2012. James Vary and collaborators answered one of the fundamental questions of nuclear structure physics by exploring the limits of nuclear stability by showing there are approximately 7,000 possible combinations of protons and neutrons allowed in bound nuclei with up to 120 protons, providing fundamental insight into theoretical constraints on isotopes. At ALCF, an Outer Rim cosmology simulation which was 15 times larger than the largest simulation previously carried out in the US is providing extremely valuable results for ongoing and upcoming DOE-funded sky surveys, such as the Dark Energy Survey and the Large Synoptic Survey Telescope, as well as setting new standards for computational performance, achieving 69.2% of peak performance (13.94 Pflop/s) on Sequoia. Researchers working under an INCITE climate end station project were the first to definitively show carbon dioxide as the major driver of planetary warming by producing a more comprehensive global paleoclimate proxy dataset coupled with the simulation of the Earth system's energy transport mechanisms during the last deglaciation, and in a separate modeling projection quantified the mechanisms driving sea-level rise. Other work included the development of a fully self-consistent microscopic theory that describes inhomogeneous supernova core-collapse including the transformation of matter from a neutron-rich heavy nuclei and a free neutron and electron gas to a homogeneous neutron, proton, and electron liquid. Other scientists and engineers, working on the detailed simulation of nuclear reactors, demonstrated that fuel-rod acceleration, velocity and displacements - using the fluid forces computed with large-eddy simulation - can be predicted with a high degree of accuracy. Other representative achievements by OLCF users include generation of spectroscopic and photometric data for more accurate distance measurements, necessary for planning future NSF and DOE missions such as the Large Synoptic Survey Telescope and the Palomar Transient Factory; calculations to enable the experimental verification of Bose glass; and screening of 2 million compounds against a targeted receptor in a matter of days, as opposed to the months that would be required for computing clusters, creating a vast library of molecular compounds that can be used for future screenings of potential drug candidates.

LCF staff have also worked with industry under INCITE, ALCC, and Director's Discretionary allocations. For example, Ramgen has partnered with OLCF to transform design approaches for shockwave-based turbomachinery, which has the potential to reduce the capital cost of CO2 compression for carbon capture and sequestration (CCS) by 50 percent. The collaborative work, involving all elements of the design framework, enabled the use of intelligent optimization techniques where ensemble simulations of varying design parameters are combined into a single simulation on Jaguar/Titan, capable of utilizing more than 240,000 cores. Ramgen's CEO and Director Doug Jewett characterized the outcome by saying "The use of Jaguar has cut the projected time from concept to a commercial product by at least two years and the cost by over $4 million."

A key factor in the successful acceleration of such scientific discoveries and engineering breakthroughs by the INCITE and ALCC user communities are the assistance in porting, tuning, and scaling applications to run on extreme-scale systems provided by the LCF. For instance, each INCITE project is assigned a computational scientist (a "catalyst" at ALCF, or a "liaison" at OLCF) to assist in scaling and improving application performance in that science area. Without these experts, using these large systems would be difficult to use, and scientists less productive.

The LCF has provided the world's best example of interactions with the high performance computing industry. The systems deployed over the last several years (TITAN and MIRA) are the result of continuing, enlightened collaborative efforts that began with the LCF program itself. IBM (collaboration on MIRA) and Cray (collaboration on TITAN) form the backbone of the U.S. high performance computing industry and this is in significant part due to their collaborations with DOE.

Readiness

Both Argonne and Oak Ridge have signed CD-0 (Critical Decision-mission need) documentation for power up-grades to their existing buildings. There are no significant unresolved scientific or engineering challenges. These upgrades will enable each center to field projected 2016 systems, even at the upper bound presented. In addition, LCF has a strong, significant partnership with LLNL. This team has recently completed a request for information (RFI) relative to a joint SC-NNSA procurement in 2016 timeframe. This, together with the excellent development and support staff already at the LCF, leads to our *short-term readiness* grade of (a)*, ready to begin.*

Beyond this point in time, there is increasing uncertainty as to the characteristics (power, footprint, cost) of advanced computing systems, and therefore also of the ability of the LCF to field the most aggressive systems. *Our long-term readiness grade is* (b)*,* since there are significant scientific/engineering challenges to resolve which are beyond the control of ASCR. The workforce has been excellent, providing an exceptional collaboration with important DOE applications. Further, the LCF workforce serves as a model for other facilities in effectively deploying advanced, extreme scale computing systems. However, this excellent track record may erode if we cannot retain an expert development and support staff.

NERSC

History of Impact in service to SC mission

From the time NERSC was established at Livermore in 1974, it has grown extensively in the number of scientific disciplines supported. Today NERSC is the main engine that supports production scientific computing within DOE Office of Science. In 2012 over 600 projects benefited from the high performance computing environment at NERSC, including fusion energy, materials science, lattice QCD, chemistry, climate science, earth science, astrophysics, biosciences, accelerator science, combustion, nuclear physics, engineering, and math and computer science. The impact of NERSC is highly visible – over 1500 peer reviewed journal publications are produced each year. We note only a few of the many widely recognized breakthroughs and/or discoveries: Nobel Prize awards in 2007 and 2012 from scientific simulations at NERSC; Supernova 2011fe was caught within hours of its explosion in 2011, and telescopes around the world were rapidly redirected to it; and the new approach developed by MIT researchers to desalinate sea water using sheets of grapheme, a one-atom-thick form of the element carbon - Smithsonian Magazine's fifth "Surprising Scientific Milestone of 2012".

In addition to decades of provisioning state-of-the-art supercomputing cycles for discovery and simulation, over the past decade NERSC has been at the forefront of recognizing the need for data intensive computing and analysis. Indeed, data import at NERSC has overtaken data export, and NERSC has met this extreme data challenge by providing the necessary computational and storage resources for data-intensive science. Experts in high performance computing, computer systems engineering, data, storage and networking provide an environment to maximize the productivity of NERSC users. By working directly with users via use cases, staff has been able to effectively prioritize and maintain a balanced resource facility.

Gathering and implementing user requirements is a long-standing strength of NERSC. NERSC has almost 5,000 users affiliated with Office of Science programs, i.e. BES, HEP, NP, FES, BER and ASCR. A cornerstone of this relationship is extensive, triennial reviews of computing requirements for each SC program office.

Readiness

The NERSC CRT facility is scheduled to come online in 2015. This new facility will have upwards of 20,000 ft^2 available and initially provisioned at about 12 MW of available power. These are certainly sufficient to field NERSC-8, and our evaluation of the readiness in the 2014-2017 time frame from a hardware perspective is (a) ready to begin.

However, application readiness is a significant challenge, because NERSC-8 will be on a different technology trajectory than Hopper and Edison. That is, this system will be more closely aligned with those already in place at the LCF, presenting similar programming and performance difficulties to many of NERSC's 600 applications that the handful of LCF application codes have already faced. This experience from LCF will help NERSC in their plan to address this challenge, building upon their extremely successful user support model. This will require strengthening its workforce. In addition, it will have to deal with two factors that are not completely under their control: (1) a successful exascale hardware, and software and application R&D effort and (2) significant domain-specific support from the other programs in SC (e.g. BES, BER).

ESnet

Impact: a, national and international and essential

In the past 25 years, DOE Office of Science has provided leading edge network connectivity for scientific discoveries through ESnet, a national network that connects 40 labs and facilities with > 100 networks. This has greatly facilitated the collaborative interactions of DOE funded scientists in geographically distributed locations, with other collaborative agencies and commercial enterprises, and with major international experiments. For LHC alone, there is a dedicated 'Overlay Network' that includes 30 networks and over 40 institutions.

Data traffic is growing exponentially driven by two factors. (1) The growth in experimental data. For example, the ALS at Berkeley Lab is a 24-hour operational facility and 45 beams are available to users in excess of 4000 hours per year. It programs will produce up to 5 Gbps of data that will need to be streamed to remote supercomputers for processing and the resulting images returned to ALS in near real-time. (2) The growth in LCF capabilities that enable more realistic simulations of grand challenge problems. One of the biggest producer and user is the climate community. NCAR and LBNL will distribute data from the CCSM model, and ORNL will distribute observational and other data. The data will be replicated worldwide for analysis and validation, and bandwidth requirements will likely approach 100 Gbps in the next few years.

To meet the challenge, ESnet with the help of ARRA funding has successfully upgraded its network to deliver data at 100 Gbps, making it one of the fastest systems in the work. This breakthrough will not only make sharing of information between labs more efficient and pave the way for new discoveries, it also holds the potential for driving innovations that find their way into the commercial sector. The proposed 400 Gbps upgrade is a logical next step and is on the strategic path to a Tbps network.

Readiness

The success of the proposed ESnet upgrade depends on international collaboration in cases where connections between multiple countries are involved. Issues regarding compatibility of systems, data transfer policies and cyber security will have to be addressed. ESnet will also have to correct historical understaffing, exacerbated by recent waves of retirement, and loss of staff to the commercial sector. The plan is excellent for the first 7 years. There could still be glitches in out years depending on technology advances not yet known (although less so than LCF's and expectations of exascale technologies).

Grade =2+ or 3

Virtual Data Facility

Description (cross reference facilities report suggestions here)

Large-scale, experimental and observational user facilities are a unique aspect of the Office of Science portfolio. Today, these facilities produce data at a prodigious rate, e.g. the limit for a large-scale DOE facility is about 10 PB/year. That figure is certain to escalate exponentially in the future; in fact, Peter Denes of LBNL estimates that detector data rates will increase 20-fold over the period from 2010-2015 and sequencers about 50-fold. The ability to effectively capture, store, filter, analyze, curate and archive data across all SC facilities is critical to the science mission of DOE. **Impact = 3 (essential)**

The subcommittee believes that (1) a "big data" storage and analysis facility with common interfaces and workflows will be necessary and that (2) building on present ASCR facilities, at least in the short term, will provide both early successes, such as NERSC's work with JGI, and considerable economies. In addition, there is often considerable synergy between analysis and visualization of large computational and observational data sets.

The subcommittee participated in a discussion of a proposed Virtual Data Facility at the January 30-th ASCR strategic planning meeting. This facility would up-grade NERSC, LCF and ESnet resources to provide "big data" storage, processing and networking capabilities. We believe that such a facility, built on existing infrastructure, would provide a valuable, initial capability. However, the software, tool and workflow infrastructure necessary to make this useful across a variety of DOE data sources is formidable and relatively unexplored. *Thus, our overall near-term readiness = 2.*

In the long-term, the facility readiness requirements are unknown. Studies such as that underway by the ASCAC Subcommittee on Synergistic Challenges in Data-Intensive Science and Exascale Computing will provide important guidance for this long-term vision. In addition, early experience with the proposed virtual facility will provide valuable insight. *Over overall long-term readiness = 1*.

To: **Chairs of the Office of Science Federal Advisory Committees:**
Professor Roscoe C. Giles, ASCAC
Professor John C. Hemminger, BESAC
Professor Gary Stacey, BERAC
Professor Martin Greenwald, FESAC
Professor Andrew J. Lankford, HEPAP
Dr. Donald Geesaman, NSAC

From: **W. F. Brinkman**
Director, Office of Science

I am writing to present a new charge to each of the Office of Science Federal Advisory Committees. I would like each Advisory Committee to help us with an important task—the prioritization of proposed scientific user facilities for the Office of Science. To meet a very compressed timetable, **we will need your final report by March 22, 2013**.

This charge derives from Administration efforts to improve the efficiency, effectiveness, and accountability of government programs and requirements of the Government Performance and Results Modernization Act of 2010. In order to improve the agency's performance, and in compliance with this Act, DOE has established several Priority Goals, including the following goal for the Office of Science:

> *Goal Statement: Prioritization of scientific facilities to ensure optimal benefit from Federal investments. By September 30, 2013, formulate a 10-year prioritization of scientific facilities across the Office of Science based on (1) the ability of the facility to contribute to world-leading science, (2) the readiness of the facility for construction, and (3) an estimated construction and operations cost of the facility.*

To accomplish this goal, DOE will undertake the following steps. We will need your help with step #2, as described below.

1. *The DOE/SC Associate Directors will create a list of proposed new scientific user facilities or major upgrades to existing scientific user facilities that could contribute to world leading science in their respective programs from 2014 to 2024 (the timeframe covered by this goal).*

   This step is complete. The Associate Directors have developed material describing the nature of a number of proposed new or upgraded facilities, the scientific justification for the facility or upgrade, and the various inputs from the scientific community that provided motivation for the proposal. Additionally, the Associate Directors have provided assessments of their existing scientific user facilities to contribute to world-leading science through 2024. The Associate Directors will be in touch with their respective FACA chairs shortly to submit this material directly to you.

2. *The information developed by the DOE/SC Associate Directors will be used by the DOE/SC as the basis for engagement with the DOE/SC Federal Advisory Committees and others to seek advice and input on new or upgraded scientific user facilities necessary to position the DOE/SC at the forefront of scientific discovery. The Federal Advisory Committees will seek additional outside input as necessary. In particular, for programs that have a significant existing or potential user base outside of the DOE/SC, the Federal Advisory Committees will be encouraged to seek input from the broader scientific community and existing facility user committees.*

   In order for your Advisory Committee to execute step #2, I suggest that you empanel a subcommittee to review the list of existing and proposed facilities provided to you by the program Associate Director, subtracting from or adding to the list as you feel appropriate. To address the concerns of the broad facilities user community, the subcommittees should include representatives of the broad, multi-disciplinary community that stands to benefit from these facilities, including representatives whose research is supported by other Federal agencies. In its deliberations, the subcommittees should reference relevant planning documents and decadal studies. If you wish to add facilities or upgrades, please consider only those that require a minimum investment of $100 million. More detailed instructions for the report are given below.

3. *Finally, with input from the DOE/SC Federal Advisory Committees and other stakeholders, the DOE/SC Director will prioritize the proposed new scientific user facilities and major upgrades across scientific disciplines according to his/her assessment of the scientific promise, the readiness of the facility to proceed to construction, and the cost of construction and operation. In making this prioritization, the DOE/SC Director will consider the resource needs for research support and for robust operation of existing facilities and will engage leaders of other relevant agencies and the Administration to ensure priorities are coordinated with related investments by other agencies and reflect cross-agency needs where appropriate.*

Please provide me with a short letter report that assigns each of the facilities to a category and provides a short justification for that categorization in the following two areas, but do not rank order the facilities:

1. *The ability of the facility to contribute to world-leading science in the next decade (2014 – 2024)*. Please include both existing and proposed facilities/upgrades and consider, for example, the extent to which the proposed or existing facility or upgrade would answer the most important scientific questions; whether there are other ways or other facilities that would be able to answer these questions; whether the facility would contribute to many or few areas of research and especially whether the facility will address needs of the broad community of users including those supported by other Federal agencies; whether construction of the facility will create new synergies within a field or among fields of research; and what level of demand exists within the (sometimes many)

scientific communities that use the facility. **Please place each facility or upgrade in one of four categories: (a) absolutely central; (b) important; (c) lower priority; and (d) don't know enough yet.**

2. *The readiness of the facility for construction.* For proposed facilities and major upgrades, please consider, for example, whether the concept of the facility has been formally studied; the level of confidence that the technical challenges involved in building the facility can be met; the sufficiency of R&D performed to-date to assure technical feasibility of the facility; and the extent to which the cost to build and operate the facility is understood. **Please place each facility in one of three categories: (a) ready to initiate construction; (b) significant scientific/engineering challenges to resolve before initiating construction; and (c) mission and technical requirements not yet fully defined.**

Each SC program Associate Director will contact the Chair of his or her Federal Advisory Committee to discuss and coordinate the logistics of executing this charge. We realize that the six SC programs will require somewhat different approaches, in part based on recent and future community planning activities. In addition, if you would like to discuss the charge further, please feel free to contact Pat Dehmer (patricia.dehmer@science.doe.gov). Thank you for your help with this important task.