

ASCAC Facilities Statement

March 21, 2013

Introduction

This 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, retired (Los Alamos National Laboratory)	

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. Although ASCR is renowned for fielding some of the most powerful computing available at any given time, the real impact of ASCR facilities is realized in the successes of the research programs of all the offices in SC—basic energy sciences, biological & environment science, fusion energy science, high energy physics, nuclear physics and ASCR itself. 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 the nation. The facilities we comment on here represent the minimum necessary to support the needs of the Office of Science, DOE and the nation 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 and networking 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 the data storage and analysis resources to the existing ASCR facilities to address the data challenges to all SC programs, and is being considered by the ASCR facilities leaders.

Findings

The table following summarizes our findings; a further discussion of each facility and justification for their categorization are provided in subsequent 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 in the table.

Facility	Impact	Readiness	
		(2014-2017)	(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 range of scientists and engineers with advanced technology and applications support and is the vehicle by which cutting edge computing technologies enter production. NERSC helps emerging fields of science and engineering take advantage of supercomputing and, at the same time, provides the production resources needed by all of the programs in the Office of Science.
2. The LCF must continue to address the most challenging compute- and data-intensive problems in the national research portfolio. It helps develop and use the most advanced computing systems for the open science community, including industry, and also works intensively with key science 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 services 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 capability for 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 and the CD (Critical Decision) process is underway for power upgrades for LCF to accommodate its 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 far-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 scientific computing expertise, including applications development and support. Support for applications development and support must come from all offices of SC, particularly in light of the ongoing fundamental change in computing and programming, pioneered by the LCF and soon to be embraced by NERSC. Hardware lifecycles are short (3-4 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 who embrace technology advances, often with significant support from the LCF, and then expanding to include a broader community, including NERSC. It is important to consider all these components in thinking about the timeline.

Broad support for the scientific and engineering applications of the future must be an integral part of 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 petascale computing 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 national leadership (impact = A, absolutely central)

The Leadership Computing Facility (LCF) at Argonne National Laboratory (ALCF) and Oak Ridge National Laboratory (OLCF) was 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. The LCF is a national resource as well as a DOE resource. Scientists and engineers using the LCF have 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, more than 500 peer-reviewed research articles based directly upon LCF projects were published in 2012 alone, including several 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 the 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, a trillion-particle Outer Rim cosmology simulation, which was 15 times larger than the largest simulation previously carried out in the US, is providing invaluable results for ongoing and upcoming DOE-funded sky surveys, such as the Dark Energy Survey and the Large Synoptic Survey Telescope. In addition, this simulation is setting new standards for computational performance, achieving 69.2% of peak performance (13.94 Pflop/s) on Sequoia. Researchers working on a climate end station INCITE 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:

- 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.
- Detailed simulation of nuclear reactors, which 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.
- 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.
- 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 is the assistance in porting, optimizing, and scaling applications to run on the 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 optimizing application performance in that science area. Without these experts, these large systems would be difficult to use, and scientists less productive.

The LCF has provided the world's best examples of interaction 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; this is in significant part due to their collaborations with DOE.

Readiness (near-term = A, far-term = B)

Both Argonne and Oak Ridge have signed CD-0 (Critical Decision-mission need) documentation for power upgrades to their existing buildings. There are no significant unresolved scientific or engineering challenges. These upgrades will enable each center to field the projected 2016 system, even at the upper bound expected for energy consumption. In addition, the 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 the 2016 timeframe. This, together with the excellent development and support staff already at the LCF, leads to our near-term readiness grade of A.

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 far-term readiness grade is B, since there are significant scientific/engineering challenges to resolve which are beyond the control of ASCR. The workforce at the LCF has been excellent, supporting exceptional collaborations with teams working on 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 service to Office of Science mission (impact = A, absolutely central)

From the time NERSC was established at Livermore in 1974 to provide computing resources to the fusion energy community, it has grown extensively in the number of scientific disciplines that it supports. Today NERSC is the main computing resource 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, mathematics and computer science. The impact of NERSC is highly visible – over 1500 peer reviewed journal publications are produced each year. We note here 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 graphene, 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 capabilities 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. These include scientists and engineers from DOE's laboratories, other federal agencies, the academic community and more. A cornerstone of this relationship is extensive, triennial reviews of computing requirements for each SC program office.

Readiness (near-term = A, far-term = B)

The NERSC CRT facility is scheduled to come online in 2015. This new facility will have upwards of 20,000 ft² of machine room space available, initially provisioned with 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 many of NERSC's 600 applications with difficulties in programming and performance similar to those that the handful of LCF application codes have already faced. The 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 the NERSC workforce. In addition, NERSC 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

History of a national and international resource (impact = A, absolutely central)

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 more than 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. First is 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. Projects at the ALS will produce up to 5 Gb/s of data that will need to be streamed to remote supercomputers for processing and the resulting images returned to ALS in near real-time. Second is the growth in LCF capabilities that enable more realistic simulations of grand challenge problems. One of the biggest producers as well as users of data 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 Gb/s 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 Gb/s, making it one of the fastest systems in the world. This breakthrough will not only make sharing of information between DOE's laboratories 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 Gb/s upgrade is a logical next step and is on the strategic path to a Tb/s network.

Readiness (near-term = A, far-term = B+)

The success of the proposed ESnet upgrade depends on international collaboration in cases where connections among 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, leading to our near-term readiness score of A. There could still be glitches in out years depending on technology advances that are not yet quantifiable, although less so than LCF's expectations of exascale technologies, leading to a B+ far-term readiness.

Virtual Data Facility

A new Department of Energy facility (impact = A, absolutely central)

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. 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. The impact of this virtual facility would be absolutely central to DOE's science mission.

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

Readiness (near-term = B+, far-term = C)

The subcommittee participated in a discussion of a proposed Virtual Data Facility at the January 30th ASCR strategic planning meeting. This facility would upgrade NERSC, LCF and ESnet resources to provide coordinated storage, archival, analysis and networking capabilities for extremely large data sets. We believe that such a facility, built on existing infrastructure at LCF (OLCF and ALCF) and NERSC and connected via upgraded ESnet resources, would provide a valuable, initial capability for the Office of Science. 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 = B+.

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