Workshop Report
Rethinking NSF's Computational Ecosystem for 21st Century Science and Engineering

Workshop Website: https://uiowa.edu/nsfcyberinfrastructure
Workshop Report: https://www.uiowa.edu/nsfcyberinfrastructure/report.pdf

This report summarizes the discussions from a workshop convened at NSF on May 30-31, 2018 in Alexandria, VA. The overarching objective of the workshop was to rethink the nature and composition of the NSF-supported computational ecosystem given changing application requirements and resources and technology landscapes. The workshop included roughly 50 participants, drawn from high-performance computing (HPC) centers, campus computing facilities, cloud service providers (academic and commercial), and distributed resource providers. Participants spanned both large research institutions and smaller universities.

Organized by Daniel Reed (University of Utah, chair), David Lifka (Cornell University), David Swanson (University of Nebraska), Rommie Amaro (UCSD), and Nancy Wilkins-Diehr (UCSD/SDSC), the workshop was motivated by the following observations. First, there have been dramatic changes in the number and nature of applications using NSF-funded resources, as well as their resource needs. As a result, there are new demands on the type (e.g., data centric) and location (e.g., close to the data or the users) of the resources as well as new usage modes (e.g., on-demand and elastic). Second, there have been dramatic changes in the landscape of technologies, resources, and delivery mechanisms, spanning large scientific instruments, ubiquitous sensors, and cloud services, among others.

Executive Summary

The workshop presenters and discussion participants both highlighted the shifts in cyberinfrastructure needs and expectations, with a broad base of applications and resource needs that are much richer and more diverse than that provided by high-end clusters and batch-oriented applications. Simply put, the community's needs are evolving rapidly, emphasizing the need to support a wide range of cyberinfrastructure capabilities, not just leading edge high-performance computing platforms.

1. The demand for cyberinfrastructure capabilities is dramatically expanding and diversifying and substantially exceeds what NSF has currently provisioned for the national research community.
2. The high end (big data and big compute) and the long tail (laboratory and campus) of computationally-driven research are no longer distinguishable and are deeply intertwined, with complex workflows, data-driven computations, and interactive explorations now common at all scales.
3. Workforce training and availability remain challenges in campus cyberinfrastructure sustainability and support, limiting research uptake of extant and new cyberinfrastructure. This should incorporate NSF's unique role in broadening participation.
4. Machine learning and big data analytics are now an essential part of the scientific discovery process, complementary to and increasingly integrated with computational modeling. Simulation and modeling, measurement and instrumentation, and data analysis are deeply interdependent, as an integral part of the scientific process.
5. Disciplinary and cross-disciplinary collaborations require converged instrument, data management, and computational capabilities. Put another way, cyberinfrastructure is an ecosystem in rapid change that is composed from many distinct pieces.
6. Incentives are needed for university investment to complement NSF's and to increase streamlined access to both, as well as commercial cloud services. NSF cannot be expected to fund the entirety of

national cyberinfrastructure, just as it did not fund the entire Internet. Likewise, NSF should not replicate what is well-supported by other agencies.

7. Under current NSF budgets, there are neither adequate resources to meet the demand for midrange, campus hardware, services, and training (e.g., via MRI awards), nor to provision high-end hardware, services, and support comparable to those deployed by the EU, Japan, or China, or the Department of Energy. NSF needs a clearly articulated strategy and appropriately matched resources.

In turn, these themes suggest several potential directions:

1. Clearly and unambiguously delineate solicitations and investments in cyberinfrastructure innovation versus cyberinfrastructure operations, while recognizing the former ultimately informs the latter.
2. Interoperability and sustainability – across scientific instruments and domains – remain elusive goals, limiting opportunities for reproducibility, collaboration, and discovery. NSF might consider a funding model that requires collaboration across MREFC projects to drive interoperability.
3. New and creative kinds of partnerships – public-private and interagency – are necessary to sustain national research competitiveness and NSF leadership. NSF cannot do this by itself.
4. Absent new funds, NSF must make difficult, strategic choices about its cyberinfrastructure, recognizing this may mean ceding leadership in some scientific domains. At present, NSF is investing in too many things, and it should focus limited resources on those things it does best. Simply put, the agency needs a clear, long-term strategy, derived from principles that are clearly articulated and understood.

Paraphrasing one workshop participant, NSF should not be constrained by the status quo, but should look to the future. In that spirit, others pointed to NSF's catalytic contributions to the Internet's formation and wondered what today's analogs were. Examples might include (a) building a set of national data archive anchor points and fund campus connections, just as the Internet was built based on peering points, (b) creating innovative partnerships with public cloud providers that include services the national community, or (c) connecting NSF's cyberinfrastructure strategy directly to its MREFC projects or Big Ideas.

Workshop Agenda

The workshop featured two opening plenaries that addressed the wide dynamic range of cyberinfrastructure needs and opportunities:

Rick Stevens, Argonne National Laboratory, *Big Data, Big Science*
Ilkay Altintas, San Diego Supercomputing Center, *The Long Tail, Research and Discovery for All*

These plenaries were then followed by a series of panel discussions shaped by four framing questions:

1. Looking forward, what are the essential components of, and relationships contributing to the NSF-supported computational ecosystem? What are the key application drivers (and usage modes) that motivate these components?
2. What are the appropriate models for, and relationships contributing to resource management, operation and delivery?
3. What are the opportunities, usage modes, and use cases for using future cyberinfrastructure components in an integrated and holistic manner?
4. What the opportunities and use cases for integration with complementary investments by other agencies in the US and internationally?

Panel members and panel moderators were drawn from the attendees, who shared their perspectives and ideas related to the four framing questions. In addition, the workshop included breakout discussions in smaller groups to address the questions and offer suggestions for strategic directions and implementation approaches.

Appendix A contains the detailed workshop agenda, including speakers and panel participants. Appendix B includes the complete list of non-NSF attendees.

The discussion below includes a summary of each panelist's presentation, prepared by the panelist, as well as a summary of the breakout discussions. As such, each represents the perspective of one individual. Slides from the presentations, as well as this workshop summary, can be found on the workshop website: https://uiowa.edu/nsfcyberinfrastructure.

Panel 1: Future Ecosystem Models and Application Drivers
Looking forward, what are the essential components of, and relationships contributing to the NSF-supported computational ecosystem? What are the key application drivers (and usage modes) that motivate these components?

Abani Patra (SUNY at Buffalo)
Modeling, forecasting, and decision making have been transformed by advances in available computing, data and AI related methodology. The linear path from observation through modeling, analysis and prediction has been disrupted and replaced by a process in which *automated data driven modeling* allows for analysis and prediction from partially specified models with iterative improvement as additional data becomes available. Consequently, the nature and demand for digital infrastructure (computing, data and networking) that facilitated model evaluation at higher and higher fidelities of well-founded models needs to change.

The new paradigm needs to embed computing in data handling and disrupt the old process with *in situ* modeling. Whereas current approaches that use ensembles for quantifying uncertainty in parameters and input data for well-defined models have led to many insights, the next advance, where the model is a dynamic and temporary construct, will require very different computing hardware and access mechanisms. The future of applications will be workflow centric and require the digital infrastructure to quickly process both large amounts of data and computing using such data.

Similarly, computing has been dominantly parallel for 30+ years, but scientific thinking has been sequential for 400+ years, where well defined hypotheses have been systematically tested and incrementally improved. The challenge now is to enable and promote "concurrent thinking" where many hypotheses must be simultaneously explored using automation. Exploration of a larger set of hypotheses will reduce time to discovery and innovation by leveraging many ideas in a concurrent thinking environment with the explorations quickly seeking to converge to the best solutions. Collaborative efforts supported on community platforms will be essential for promoting such concurrent thinking.

Don Krieger (University of Pittsburgh)
The Krieger research group seeks to better understand, diagnose, and treat traumatic brain injury by extracting high fidelity neuroelectric traces from noninvasive brain recordings. The tools target grid computing resources, using compute cycles opportunistically which would otherwise be unused. research results demonstrate characterization of changes in brain function with unprecedented detail, useful

measures of neuroelectric activity for the first time from the fiber bundles which conduct information from one brain region to another (the white matter) ,and high-resolution, high-fidelity measures of ongoing brain function which yield clinically relevant information for each individual. These efforts demonstrate the value of opportunistic use of high performance computing resources, and the importance of providing support for a high-risk effort which required more than five years to mature.

Shaowen Wang (University of Illinois)

Novel and replicable tools, methods, and infrastructure are urgently needed to harness exponentially increasing data velocity, volume and variety for urban discovery and innovation. Specifically, three types of interrelated capabilities represent significant challenges and opportunities: dynamically urban sensing instruments based on Internet of Things approaches, computation with high-performance cloud and edge computing for handling geospatial and urban big data and enabling analytics, and spatiotemporal analytics with advanced cyberGIS and machine learning. Integrating these three components to provide a suite of data-driven capabilities will catalyze new approaches and insights into scientific challenges related to urbanization, ranging from emergency management to environmental sustainability and flood mitigation, and from the nexus of food, energy, and water to urban safety, health and air quality.

What are the key application drivers (and usage modes) that motivate these components? Urban areas have ecological, environmental, and socio-economic impacts far beyond their proportion of the population. More than half of the world's population currently lives in urban areas (80% in the U.S.) and this proportion is projected to climb up to as high as 70% by 2050. Whereas urban areas aim to provide physical and socio-economic infrastructure designed to provide a higher quality of life, the implementation and management of urban infrastructure has not been without flaws. Indeed, urban challenges include rising energy demand and corresponding increases in greenhouse gas emissions, increase in air and water pollutants, heat emissions, development of hotspots known as urban heat islands, changing microclimates, uneven rainfall patterns and related water issues, urban blight, and crime. This rapidly increasing 'stress' affects the sustainability of cities and renders urban ecosystems increasingly vulnerable to natural and anthropogenic hazards.

Unless societies strive to make current and future urban design, operations and resource management decisions to reduce urban stress based on scientific understanding from a variety of disciplines, we will continue to threaten human health, strain energy resources, and reduce economic productivity. To this end, complex and massive geospatial data is increasingly collected for understanding and tackling such challenges, motivating many geospatial and urban observatories that could play essential roles in resolving these challenges through science, engineering, and policy innovations. However, such observatories require innovative cyberGIS and cyberinfrastructure approaches to integrating heterogeneous and voluminous geospatial big data with associated analytics and modeling capabilities for a variety of scientific problem-solving and decision-making purposes.

Frank Wuerthwein (UCSD)

My input is based on my experience as an experimental particle physicist, working on the CMS experiment at the Large Hadron Collider (LHC), but also more generally on experience supporting research computing for science with instruments of all kinds, ranging from the LHC to advanced LIGO, IceCube, and a host of medium size experiments such as XENON1T, VERITAS, the South Pole Telescope, Simons Array, GlueX, Mu2e, Minos, Minerva, and DES. In addition to these physics and astronomy instruments, those in the life sciences increasingly produce substantial data volumes, leading to similar processing and data management requirements. As a result, we can broadly talk about deriving scientific insight from data taken with instruments of all kinds as a key application driver for the future. This includes large instrument collaborations across multiple institutions, funding agencies, and nations but also individual PIs or small groups that either operate their own instruments, or buy time on somebody else's instrument, locally, or at another institution.

The common usage model is to archive the data produced by the instrument(s) and their corresponding simulations in "passive storage," transfer the archived data to "active storage" for processing, and potentially store the processing results again in passive storage for later (re)use. The differentiation between active and passive storage is motivated by price. Exabyte scale distributed archives as required for the LHC are substantially cheaper when implemented as passive rather than active storage, making the split cost-effective. Price consciousness is mandated by the scale of the data for the largest individual instruments, as well as the collective of all the more modest data producers. For the latter, there is an additional challenge of coordination across a cottage industry of petabyte scale data producers that lack an option for economies of scale that experiments at the LHC have. At any given time, the amount of active storage needed for processing is a small fraction of the total archival storage required to preserve all previously accumulated data.

Science is a team sport, and large instruments tend to require large multi-institutional (and multi-national) teams to design, build, and operate the instruments, and to derive science from the data collected. The resulting scientific collaborations benefit greatly from integrating resources (human, software, and hardware) that participants commit as "in kind" contributions. More often than not, this integration is interagency and international.

Distributed High Throughput Computing (DHTC) is the ideal paradigm for meeting the research computing needs of these instrument collaborations. DHTC is unique in its ability to integrate compute and storage resources subject to different policies in a dynamic and heterogeneous environment. Policies may span the full range from preferred to fair share to opportunistic. Resources may be operated by universities and regional, national, and international organizations, including network providers, supercomputing centers, and commercial clouds. By easily supporting an ecosystem of resources and policies DHTC is also an ideal sharing paradigm for data. An obvious use case is multi-messenger astronomy where individual scientists want to access data from many different instruments from within a single framework in order to facilitate correlation across instruments. Each instrument collaboration manages a data life cycle that includes a transition from data being private to the collaboration to becoming public to a larger audience, or even completely open to all. DHTC as a paradigm is built from the ground up on the federation of resources. Tools from within the DHTC ecosystem are thus excellent building blocks to build upon to support data and resource sharing across a wide range of scales and science domains.

Rommie Amaro (UC San Diego)
Key application drivers in biomedical and environmental sciences now (already) draw heavily from data-rich experimental approaches that must be integrated to develop highly detailed physical models of systems with many components and that vary in time and spatial dimensions. Investments in the experimental side of these science areas is already producing the types of datasets that are, and will continue to, benefit from integrative approaches to enable an understanding of emergent behavior of complex systems. Multiple cyberinfrastructure frameworks must be supported to enable this science.

First, large massively parallel machines will continue to be a requirement to simulate large-scale atomistic based modeling and prediction of whole systems (e.g., viruses, cells, sea spray aerosol particles, etc.). For these large systems (i.e., atom counts approaching a billion or more) a machine with fast interprocessor communication is key (e.g., NCSA Blue Waters, DOE Summit, etc.). These machines are envisioned to be part of the flexible and capable cyberinfrastructure of the present and future.

At the same time, simulations of single proteins/ion channels, aka traditional drug targets or proteins known to nucleate ice, are benefitting from novel data analysis methods that combine statistical methods, such as Markov state modeling, with rigorous theoretical underpinnings (statistical mechanics). The "resolution revolution" in cryoEM will drive this area of science already and over the next 5-10 years. The broad application of these new methods not only to molecules, but to supramolecular complexes, will require researchers to have access to many hundreds or thousands of GPU cards to gain good statistics. Such studies will revolutionize our understanding of biological phenomena such as allostery, for example, which is potentially key to developing more efficacious therapies for a wide array of disease areas.

To run computational molecular sciences, as described above, in rigorous and reproducible ways, the cyberinfrastructure of the future must enable interoperability with different levels of workflows and workflow frameworks. Storage of the resulting data and/or in some instances analysis on the fly without storage will be a challenge for researchers. A networked data "superfacility" is likely to be a continuing need (hubs sites that can enable researchers to analyze data across sites and by researchers remotely [without requiring individual labs to pull data locally]).

Finally, the coordinated development and support of a flexible cyberinfrastructure will require ongoing participation and discussions among all the national research funding agencies. As there is no "one ideal solution" to the many needs of all research groups, we must continue to have broad discussions about these topics and, as much as possible, insist that interagency communication is facilitated and required.

Panel 1 Breakout Summary

The intent of the breakout session was to move beyond the panel presentations, which largely focused on current and near-term future (within 5 years) applications and instead think about the future. Multi-messenger astronomy, citizen science and adaptive/real time processing were areas that were put forth for discussion.

Breakout participants agreed that an increase in the computational resources is needed for improved accuracy and resolution, but also saw different usage scenarios than the steady-state, queue-based loads of the past. The participants noted a strong need for data integration, preservation and annotation, hybrid, interactive explorations involving simulation and experimentation and real time feedback. One example is in earthquake engineering, where building destruction is simulated and 1-2 second turn around in computation is needed to interact with shake tables. Many current NSF-funded machines are well-suited to steady state load, but this is not how some science is done in many fields today.

In addition, breakout participant and panelist Don Krieger contributed several observations about high-end machines. Don observed that as machines increase in size and support more users, the ability of some users to slow down all other users' jobs increases. Network and disk I/O are the most common vulnerabilities on these machines and a fractional expenditure and attention to design for these components will help all get better use from these large systems. Don mentioned that the use of deferred/kill at will queues could extract greater efficiencies from large machines without impacting jobs in other queues. Don also observes that data movement and data archival services require significant planning and often significant effort to deal with the problems. For big science projects, these issues are well thought out by team members. But for small groups and individual scientists, these can be entirely misunderstood and mishandled. It might make sense to require a formal data flow plan for allocation requests and to provide review, feedback, and consultation by an expert in the specific system.

Ian Foster observed that artificial intelligence will drive new classes of applications based on both large-scale use of “conventional” AI methods to scientific problems (“HPC for AI”) (e.g., to process large genomic, imaging, simulation, linguistic, environmental datasets). AI will also be used to drive or add value to simulation (“AI for HPC”) (e.g., adaptive learning to decide what simulations to perform next, deep learning to generate surrogate models for large and complex simulations). Ian also saw the need to create and operate “learning systems:” computer systems that can automatically ingest, integrate, learn models from, train models with, perform inference on models, retrain models for individuals and communities. He also noted that data may not be in one place.

Composable services (learning, data classification) will play a role, but the type of infrastructure is subject to debate? How are services discovered? How can the community bring persistent services into the scientific workflow, what role will gateways play and how can OAC help with all of this? These questions were posed, but not answered.

Participants also compared investments in networking, which were originally funded by the NSF and later picked up by campuses. Is there a comparable need for OAC to support individual researcher’s data, for example through a program for storage? NSF could provide seed funding with different costs for active vs passive storage. OAC could encourage data-oriented research and perhaps also work with cloud and CI providers to enable “affordable” data archive services. Some in the breakout wondered about the role of campus libraries as a possible transition plan.

Returning to the non-standard computing discussion, participants recommended rethinking the metrics of success for CI, moving beyond utilization as the primary measure of success and evaluating different costs for different usage models.

Participants also discussed the type of automation needed to support applications, for example enabling closed-loop simulations and different type of schedulers (hierarchical scheduling). Future simulation tools will likely embed techniques to learn the simulation environment and adapt to availability. Science gateways can help standardize the interactions with cyberinfrastructure but can restrict users as well. Gateways of the future will need to be flexible. There is a tension between automation versus customizable solutions.

Researchers are increasingly using (and generating) large data collections that are too large to download. Associated computational services are needed co-resident with the data. A new type of cyberinfrastructure is required to host and automate such services, one that will support workflows and resource provisioning and management. OAC can support research in this area and services that can process, analyze and visualize data at the source.

In a related vein, the breakout group considered provisioning data flows, where processing is included in the workflow. Delivering data to applications in many computing sites is an example. The group, however, worried about cost models. Rob Gardner pointed out the importance of computing at the edge, service capabilities at the edge, or in the network, such as distributed gateway components. The group agreed that we need to examine the data and application lifecycle in a holistic way, in a very heterogeneous environment, and consider new modalities of provisioning resources. One proposed idea was to focus on pathfinder or grand challenge applications.

Breakout members believed that both CI staff and the user community need retraining to support these new modes of science, which rely on heterogeneous CI, summarized as submit locally/compute globally. Modernizing OAC's training resources, for example those deployed via XSEDE was suggested. The group also discussed how to make scientists aware of existing capabilities and engage them in the evolution of NSF-funded CI. Interdisciplinary workforce development and education on the use of public clouds for the long tail is also necessary.

The group also touched on the importance of reproducibility and how can CI help, in particular through recent advances in containerized software. Finally, the group discussed how scientists, infrastructure providers and the NSF can all be successful, given that they have in some sense conflicting goals as far as time-to-solution and utilization. Don Krieger left the group with an observation from keynote speaker Rick Stevens, who pointed out that creative ideas will sometimes fail. This can be true of some OAC investments in CI, as well as future directions recommended in this workshop, but to penalize innovation in favor of more pedestrian machines was seen as detrimental by some.

Panel 2: Future Resource Management, Operation, and Delivery Summary

What are the appropriate models for, and relationships contributing to resource management, operation and delivery?

Dan Fay (Microsoft)

For future cyberinfrastructure (CI), there is a need to move from thinking just about raw computations, storage, and bandwidth, but to look at how scientific research can be supported via services that abstract the complexity of managing the infrastructure and hardware for the scientists. Due to the commoditization and affordability, one of the challenges for sciences nowadays is that the computational demand will always exceed the available resources – due to simulations, modeling, analytics, etc.

To help scientists focus on the research challenges and their core expertise of generating scientific insight they should look at leveraging cloud resources and services. One of the benefits of utilizing Cloud Vendors services is they operate at such scale and volumes they are able to provide access to the latest hardware (CPUs, GPUs, etc.) as well as make investments into new technologies (FGPAs, TPUs, Quantum, etc.). NSF and other agencies should look to see how to partner with cloud vendors to create a sustainable model with the aim of having scientist focus on their core expertise, doing the science and not managing infrastructure.

Tom Furlani (SUNY at Buffalo)

The current NSF model of supporting campus-based cyberinfrastructure (CI) systems, national capacity class CI systems, and national capability class CI systems provides the nation's academic community with a diverse CI ecosystem to support its research needs. While demand will always exceed capacity, given the critical role that simulation science and data driven science play in science and engineering as well as the U.S. economy, additional funding is needed to support acquisition at all levels of the NSF CI ecosystem. No matter the service model employed (cloud, national systems, campus systems) to support simulation and data-driven science, local staff with expertise in computational science and data analytics are needed to work with faculty and students in order for them to fully leverage existing resources. This is especially true for researchers in the domain sciences where there is not a long history of utilizing high performance computing.

In terms of campus CI, continued investment by NSF in moderate sized campus-based CI systems is desirable. The arguments in favor of this strategy are several. First, many of today's scientific application codes do not scale well beyond a few hundred cores or the problems being addressed by researchers are not large enough for the application to run at larger core counts. Second, many of the large national facilities are designed/optimized for the most demanding computational problems with expensive interconnects and file systems designed for large I/O rates, making smaller jobs comprising tens to hundreds of processors not a cost-effective use of these resources. Third, the operating costs (power, cooling, and staffing) of campus systems are typically borne by the host institution thereby maximizing NSF's investment. Finally, campus CI plays a crucial role in the training of a next generation who will lead the creation of the new HPC tools and methodologies.

Sharon Broude Geva (Michigan)

Broude Geva's remarks included the need to address various policies. These included a sharing policy (e.g. to allow OSG-style sharing explicitly in guidelines), data management policies that are of elevated importance with clearer guidelines for NSF panel review, and institutional mechanisms and approvals that would be needed to use awarded cloud credits. It is possible that an institution would not have a mechanism already in place to allow a researcher to conduct an analysis of HIPAA-restricted data on a commercial cloud resource, or would not want to, or be able to, sign such an agreement. This obstacle could be removed by ensuring that the terms of use of awarded cloud credits for a research project be the same as the those of allocations on government-funded national computational resources.

This is just an example of a broader issue that arises when researchers, regulated by the legal standards and agreements of their institution, use computational resources that are not owned and certified by their institution. This is very different from the use of on-premise campus resources (that are owned, certified and regulated by the researcher's institution) or even national computational resources (that are owned, certified *and* regulated by the national provider).

Broude Geva also emphasized that the Branscomb pyramid[1] could be modified to represent the current landscape and needs. The levels in the current form of this pyramid would include the same levels as Branscomb, with the addition of commercial cloud at the peak, but would no longer reflect congruence between the levels. The levels are not congruent but complementary, and no longer represent an "on-ramp" picture where the assumption is researchers start at the broad base and progress to the tip. This last comment echoes Ilkay Altinas' remarks concerning the need for composable systems which include multiple aspects of cyberinfrastructure even for initial investigations.

David Lifka (Cornell)

David Lifka discussed the Aristotle Cloud Federation project, an NSF Data Infrastructure Data Infrastructure Building Blocks (DIBBs) Award (No. 1541215).[2] Aristotle provides a model that allows institutions to share local hybrid clouds with each other providing elasticity and access to diverse research data at the various partner site as well as special software and hardware configurations before having to spend money on public cloud resources. Aristotle includes an allocation and accounting systems so partner institutions can earn credits on each other sites when users other partner institutions run on their local resources. It also includes a portal that makes it easy to see the availability of resources on partner sites, NSF resources and eventually public clouds resources allowing researchers to make decisions on

---

[1] *From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing*. NSF Blue Ribbon Panel on High Performance Computing. Arlington, Va.: NSF, August 1993.

[2] See http://federatedcloud.org/ for more information.

where to run based on availability and cost. The ultimate goal of Aristotle is to improve the time-to-science for researchers. Lifka noted that this may be an interesting model for the NSF community to share their resources while earning credits on other institutions resources. This allows institutions to provide their researchers with a broad range of resources without having to invest in all of them locally.

Panel 2 Breakout Summary

As a group, there was consensus that NSF must fund things that cannot or will not be funded any other way. The reality of finite funding is that choices will need to be made between competing priorities (e.g., campus or national infrastructure). This can be mitigated by increasing the total sum available, of course, but clear alternatives remain. Much good research is done even in small universities, many of which have been uniquely well-served by the MRI program. Several participants lobbied for a cyberinfrastructure-specific MRI, to alleviate the difficulties of obtaining a limited submission opportunity in a campus environment. Regardless of current programmatic details, the breakout group was unanimous in believing that NSF must have the flexibility to respond dynamically to a rapidly changing landscape. As highlighted in the recent National Academies cyberinfrastructure report, scientists are utilizing “all types and capabilities of systems, from large numbers of single-commodity-nodes to jobs requiring thousands of cores”[3] as well as a wide assortment of data and networking needs. A survey of research computing professionals and NSF directorates concerning how to allocate funding may be instructive to help define the community’s current research priorities and the facilities needed to fulfill them.

The group also observed that it is non-trivial to measure the complete cost/benefit ratio of NSF investments; transparency is needed. Impact can be estimated, but workforce development, diversity, and other aspects of NSF’s mission are not as easily quantified. In particular, there is not a comprehensive method to do accounting of all resources funded by NSF. This could lead to unanticipated side effects if campus awards are curtailed or eliminated. The millions of jobs run there would have to go somewhere else, most likely NSF run centers. NSF may want to consider rewarding campuses for their investment. Especially when a proposal competes to house expensive equipment, campuses much be prepared to well utilize the funding awarded. In the end, NSF funds more research than they fund computing, resulting in oversubscription of XSEDE resources, which is exacerbated by computing demands from other agencies.

Initial debates about resource management and delivery options focused on expert personnel as a critical component of successful cyberinfrastructure delivery. Several examples such as Campus Champions (CC) or XSEDE’s ECSS were described as critical to scientific advance but insufficient in numbers to meet demand. Regionally tasked staff might help to alleviate this shortfall. Benefits could include greater use of cloud or national resources if there was a local expert to help researchers with initial utilization. Along these lines, it was mentioned that the NSF CC* programs changed campus culture, spurring local networking expertise. A similar program to promote workforce development to incentivize local computational and data scientists could, for instance, result in the integration of otherwise isolated clusters on campuses with national resources. These key personnel, ranging from ECSS experts and developers to CCs, are often in careers that need professionalization.

As stated above, workforce development is an NSF distinctive, with diversity and participation by the underserved uniquely served by required broader impact components of every proposal. This success has not been broadly replicated concerning other required elements such as facilities statements or data management plans. While these are seen as valuable, it is suggested further guidance or requirements could lead to improvements in the ability of reviewers to evaluate site readiness and ongoing

---

[3] National Academies, *Future Directions for NSF’s Advanced Computing Infrastructure*, 2016

responsibilities and commitments. A “resource sharing” document was suggested that would describe how NSF funded resources would be shared with the broader community. This may go beyond providing login access to individuals to becoming part of a larger shared platform such as that provided by OSG.

Finally, the complementary missions of innovation and production infrastructure were noted. More clarity on deliverable requirements for RFPs that address the latter was expressed as a need. There is a natural tension between what holds promise for the future and what is currently working reliably to support computational science. Commercial providers, who are often innovating on hardware at a faster rate than national centers, hold advantages here. Clouds can limit the required initial investment needed for rapidly evolving technology. Tightly coupled machines will continue to require a periodic huge investment, while distributed high-throughput computing may be purchased incrementally. Coincidentally, the former is decreasing in importance for many research projects.

Panel 3: Integration and Holistic Operations
What are the opportunities, usage modes, and use cases for using future cyberinfrastructure components in an integrated and holistic manner?

Alan Blatecky (RTI)
Blatecky highlighted four possible strategies for NSF and the community to consider.

Strategy 1: NSF no longer establishes leadership HPC Facilities to provide compute resources; NSF partners with DOE for exascale science; the majority of existing compute services provided by NSF should be migrated to clouds, consortiums and campuses rather than continue to support these users with leadership machines.

Although Moore’s Law has driven the increased capabilities of Leadership-class supercomputers over the last three decades, the cost of these leadership-class machines has also continued to rise to point where NSF no longer has enough funding (~$600M) to establish and operate the next generation exascale machine. However, DOE has announced plans to fund and acquire the next generation leadership exascale machine, so NSF should not compete with DOE to establish a next generation leadership machine, but should instead explore a partnership to with DOE to support NSF Exascale science.

Strategy 2: HPC resources should no longer managed as a separate Cyberinfrastructure capability; compute support should be included in research proposals along with data, software and personnel.

Computing support capabilities and services have dramatically changed since NSF established the supercomputing centers in the mid-1980s. Many of the compute capabilities required for science were not available in 1984, so NSF established the centers. However, the rapid growth of clouds, along with the tremendous increase in capability of affordable single racks of equipment, means that there are now alternative options beyond NSF resources for researchers who need compute capabilities. It is time for NSF to review the usefulness of providing general compute resources. For example, when networking services and capabilities began to be widely supported by industry in the early 1990s, NSF stopped supporting NSFnet, and let the networking marketplace provide those services. Compute resources have evolved to a similar point; NSF should focus on providing those CI resources that are not generally available in the marketplace. It is worth noting that there are several companies (Amazon, Google, Microsoft) that have established (and are continuing to establish) massive compute and data centers to provide compute and storage services. The cost of many of these centers cost more than $1B each and contain millions of nodes.

Strategy 3: Use the NSF Big Ten Ideas to drive the acquisition and development of advanced CI capabilities; unique compute architectures, data structures, workforce, software development.

One of the biggest challenges in the development and support of cyberinfrastructure over the next two decades will be to find ways to effectively build and deploy integrated cyberinfrastructure that supports sharing and interoperability. What makes this very challenging, is that it is not possible to develop a common universal cyberinfrastructure architecture that will meet all the needs of research and education. Cyberinfrastructure can only be effectively deployed when it is developed to address specific implementations and applications. This means that the impetus for development must come from domains and researchers who are doing research or trying to solve one or more problems. This approach in turn creates a range of problems and issues that need to be addressed, including reinvention, duplication, redundancy, lack of scale, lack of interoperability and little incentive to work on integration or sharing of resources and capabilities as these efforts do not directly accrue to or directly benefit the research being done. One of the most effective ways to NSF to insure that the future cyberinfrastructure meets the needs of science, is for OAC to engage and partner with all the Directorates and scientific domains, and since NSF has already identified *10 Big Ten Ideas in Science*,[4] the new cyberinfrastructure effort should leverage that vision and strategy.

Strategy 4: NSF to provide leadership in the deployment and integration of emerging capabilities including AI, IoT, G5, Blockchain, P4 networks, Quantum computing

While a confluence of factors will continue to drive the development of cyberinfrastructure in general, NSF should focus on developing those components or capabilities of cyberinfrastructure that specifically support engineering and research. This will include basic foundational cyberinfrastructure research leading to new capabilities such as new computer architectures and infrastructure; challenges associated with embedding Artificial Intelligence into the conduct of science; new network architectures as fiber and wireless capabilities eliminate issues of geographic location; proliferation of IOT devices that can be used or modified to support science; coordination with other international efforts.

A related challenge is to address research issues associated with a "human-in-the-loop" cyberinfrastructure that goes beyond human interfaces and explores AI/symbiotic interfaces as devices and intelligence become handmaidens of human efforts, where automatic awareness, assistance capabilities, and bots will do things without our conscious or proactive involvement. This will also include the establishment and support of new "tribes" or research communications who conduct science with a set of tools and at a scale almost unimaginable today.

Robert Ricci (Utah)

To create a cyberinfrastructure (CI) that is integrated across different facilities, and domains of science, it is critical that the impetus for federation and consolidation come from the bottom up, rather than the top down. The people in the best position to understand which features can be usefully combined – either across facilities, or by consolidation – are the researchers who use them. Not all users of cyberinfrastructure will necessarily try to push beyond the boundaries of existing CI, or combine elements from multiple Cis in one workflow, but there will always be a vanguard who will. It is from this vanguard that we can get the best information on what future CI should look like.

---

[4] 10 Big Ideas for Future NSF Investments, https://www.nsf.gov/about/congress/reports/nsf_big_ideas.pdf, 2016

It is imperative for CI to empower "power users" to build integrations ad specializations atop our CI, and for CI providers to pay careful attention to what they build. This is akin to the idea of "desire paths": looking at the walking paths people actually take, not just those that have been designed for them. Finding the things that users do that are not exactly in line with the intention of a CI facility is exactly how we will learn what future facilities should look like. This can only happen if we give users the freedom to use our CI in ways we had not anticipated; otherwise, we will never really learn how to build future CI that is not just a small iteration on the last one.

Dan Stanzione (TACC)
Much of our current data on what our CI consumers need is biased by how we allocate our systems — if we limit the scope of what we allocate, we limit the scope of the problems as we well. We have lots of evidence that our current XSEDE resources run problems that fit well on XSEDE — and that different offerings draw different kinds of computational science — there are all kinds of users out there, operating at all kinds of scales — including the very largest — but science can also happen on a single node. We face a number of challenges moving forward, not the least of which is our crisis in building scientific software. We ask users to build software that is conservative and does not change, so we can be sure the answers are right, and at the same time ask them to be patched to the latest version for security.

Although commercial software manages divergent demands by scaling up the budget as it becomes more successful, we do not have the same model for research software. We also face a crisis in reproducibility, that ultimately leads to a crisis in credibility — for computational science, and for science more broadly. We have viewed gateways as a "web interface" to our CI, but they may mean much more. It's science services on the web — and may have capabilities beyond the system the gateway interfaces too — including data management, metadata collection, collaboration, etc.

With so many challenges in computational engineering and science, and an entire NSF division devoted to CI, it's time we think of CI as a discipline. If we have a discipline, we need to reflect on our first principles, on the knowledge our practitioners should possess, and on basic research within CI — which means CI research can be basic and not applied (i.e. not tied to generating a specific output for a user in another discipline).

We need to also embrace the cloud — the cloud spends a great deal on "IT infrastructure," but just because the money is there doesn't mean they solve our problems — but it does provide an opportunity for leverage. When "Beowulf clusters" came to HPC, they were successful not because they ignored the commodity markets, but because they embraced the strengths of the markets to build a custom solution for science. The cloud can be similar — is the "as-a-service" model for software what we need to embrace from the cloud for science? Let us think about the cloud experience in our CI. Applications should run, and be composable and consumable in many modes. There are many non-trivial challenges to create the cloud experience from the components the cloud is driving costs down on — for instance, our users expect to see the same data everywhere, presented as a file system. How we encapsulate jobs matter, network costs matter, data locality matters. But this is the challenge to tackle, not run from.

A vision for the future — we should have a lot of good hardware resources — scale, but not entirely homogeneous. HPC, virtual machines (VMs), interactive and batch, a variety of storage across the lifecycle. We should seek to build on top of these resources a data-rich, service-rich cyberinfrastructure. And on top of that we need an integrated ecosystem of professional practitioners. A future CI will let you compute what you want, store what you want, pick from a rich set of services (analysis, data quality, vis, learning), and have support for reproducibility, quality assurance, uncertainty quantification, security, and compliance.

Craig Stewart (Indiana)

Scientists and researchers adapt to resources available in their environment, just like any other humans. There is no more reason to believe that the recent rise in the percentage of jobs running on XSEDE as single-node jobs represents what people need than there is to believe that this represents scientists adapting to an environment in which there are too few large systems. Indeed, there are many anecdotal reports that suggest that:

- The overall capacity of computational systems provided by the NSF for use by the national research community is inadequate to meet current demand
- The availability of large-scale parallel computing resources in particular is well below the needs of high priority government-funded research; and in this area in particular investment by individual campuses is difficult because of the high price of such systems and generally low number of users qualified to use such systems on any given campus.
- There is a growing diversity in the kind of resources that the national research community needs – as evidenced by:
  - the tremendous accomplishments reported by users of Blue Waters,
  - the popularity of Jetstream (and its success in attracting new users – 80% of the users of Jetstream have not previously used any NSF-funded cyberinfrastructure)
  - the oversubscription of requests relative to available resources ranging from Stampede 2 to Comet and Bridges

The days of the NSF supercomputer center are gone, and so are the days when the demand for advanced cyberinfrastructure from the national research community was as simple as it was in the late 1980s. At that time, there were a suite of agreed upon grand challenge problems that could be addressed with the fastest available supercomputers, and only such supercomputers. In those days there was also a small community of people capable of using such systems.

The NSF simply cannot pay for all of the advanced cyberinfrastructure needs of the entire US higher education community. Thus, Stewart recommended the following:

- Greater investment and more balanced investment across a variety of different types of computing systems. In the early 1980s there were maybe 4 architectures (tops, counting generously) that were relevant to the mission of the NSF supercomputer centers. Now there are major differences in types of advanced cyberinfrastructure (leadership class parallel systems; everyday HPC clusters; HTC; cloud) and processors (IBM Power, Intel, a resurgent AMD, ARM, and NVidia to name a few). The NSF should treat its investment in CI facilities like stocking a hardware store. People are going to come in to find tools, use the tools they need, and for any given client (or researcher) the tools nee3ded will be a small fraction of those made available.
- Greater investment in high-end highly scalable computer systems – systems in capability that will fall between the current Blue Waters system approaching what the DOE describes as "leadership class" supercomputers. Given the diversity and uncertainty in the processor market today, the NSF should invest in at least two systems that are substantially larger than the current Blue Waters system.
- Greater investment in cloud systems designed to support research (like the current Jetstream system)

- The NSF should change its approach to allocations of computational systems, including the following changes:
    - There should be one allocation process for use of all NSF-funded computational resources ranging from the "large scale" MRI awards to the successor to Blue Waters
    - NSF should offer multi-year allocations of computational resources to researchers funded by the NSF (and where possible to researchers funded by other federal funding agencies). This does not have to mean that NSF guarantees that a specific system last n years, but that it commits to provide and support certain types of systems for n years. (The minimum for effective long term allocations is probably 3 years, which is also probably the effective maximum). Without that, researchers and research centers will always have reason to justify the wasteful spending of federal funds on small-scale, inefficient, and likely improperly secured cybersecurity dedicated to specific projects.
    - Coupled with the above, the NSF should create a small set (3 or 3) CI centers that provide "everyday" advanced CI systems for use by FFRDCs, MMURFs, and other research projects. Stop creating a cyberinfrastructure center on a per-project center and do it at scale, do it well, and provide excellent security. Consolidated CI resources should, if implemented well decrease cost, increase assurance of good cybersecurity, and improve quality of CI services delivered to NSF-funded researchers. This could be done in a way that took advantage of backfill systems to run single processor or high throughput jobs.
    - Limit the ability to apply for use of NSF-funded CI resources to researchers who have a current federal grant award. ("No resources without prior peer review of the scientific or engineering research" – with the caveat that some could argue that some federal funding via EAGER and other mechanisms does not constitute full peer review)
    - Reward local campuses for local investment in cyberinfrastructure. The ONLY way to build a national open research cyberinfrastructure that meets the needs of researchers working at institutions of higher education is to have a larger share of the national CI fabric purchased and paid for by individual colleges and universities. One way to incent this is to give "credit" or "bonus points" to researchers who apply for resources on NSF-funded cyberinfrastructure and which come from campuses that invest heavily in local cyberinfrastructure. [One would want to weight this properly by campus size, etc., and put in factors to preserve good access for researchers from MSIs and financially -challenged institutions in EPSCoR states). Similarly, reward researchers from campuses that make strong and long-term investments I storage and link good storage system capabilities with technologies to make data collected with NSF funding actually discoverable and practically useful for the rest of the US research community.
- Invest more in software that enables researchers to make effective use of distributed cyberinfrastructure

Panel 3 Breakout Summary

Based on the panel three breakout question, the group offered the following recommendations:

- An NSF-wide approach to CI is needed. This could start with a common platform across Major Research and Equipment and Facilities Construction (MREFC) projects. This may be a promising opportunity to include science leaders from multiple domains and NSF should derive holistic benefits for a broad array of science. By reducing costs for CI through a common, shared platform, the NSF would be able to reinvest more money into the instruments and or science. This is then likely to affect a broad swath of science projects at different scales. *NSF should have a solicitation that requires multiple MREFCs to collaborate on CI.*
- Specific areas that need to be understood and addressed include integration, heterogeneity, and sustainability, as well as diversity and broader impact from a user and community perspective.

- Systems need to be viewed as composable.
- *Researchers need dynamic capabilities to have access to all resources and capabilities required by a modern workflow. This is possible with commercial cloud services, but not on current academic systems.*
- The community must create a holistic approach to data and data sharing. This could be tied to existing resources, but data need to be accessible from other facilities. Metadata creation and extension are also critical to be supported/required. Centers can provide storage in short term, but the long-term need is for a model to store data in perpetuity. NSF should explore both medium and long-term approaches. MREFC facilities understand this, and there are potential synergies with their activities.
- The current scope of cyberinfrastructure (CI) training is insufficient, as it does not train CI professionals. We need clearer career paths and formal recognitions, as well as pipelines that lead to leadership.
- Should NSF reward and/or create incentives for campus investment? This fits naturally with professional development. CI sharing could be included via NSF's broader impacts requirements. Data management plans must be taken seriously, and investment in data should include the Investigation of different models for data curation.

Panel 4: Interagency and International Collaboration Opportunities

What the opportunities and use cases for integration with complementary investments by other agencies in the US and internationally?

Ian Foster (Argonne)

Foster's talk focused on the important topic of software cyberinfrastructure. He argued that the future of scientific software CI must encompass, as is already broadly the case in industry, cloud-hosted platform services to which scientists and scientific applications can outsource time-consuming tasks relating, for example, to data and computation management. He noted that the establishment of such services provided a wonderful opportunity for interagency and international collaboration, as such services both require broad adoption to be successful and, when broadly adopted, provide a natural foundation for interoperation.

Foster used the example of the Globus service to illustrate principles and approaches that have been demonstrated to work effectively. This cloud-hosted research data management service is operated by the University of Chicago for the research community. It provides Web, REST, and command line access to powerful security and data management functions that orchestrate data transfer, sharing, publication, and other activities on and among storage system endpoints running Globus Connect software. With more than 10,000 active endpoints and 100,000 registered users, Globus has become an important element of the national cyberinfrastructure for many institutions and scientists. A subscription model allows US and international universities and national laboratories to contribute to its long-term sustainability and continued development.

He concluded by noting the following points:

- High-quality services that "just work" can have a transformative impact on science by eliminating barriers that frequently hinder innovative thinking. But they are expensive to build and sustain. Interagency and international collaboration can help "cross the chasm" to affordable services.
- The power of combining public cloud and science resources: Science often involves the use of specialized resources, from instruments and data repositories to supercomputers. Artful

integration with cloud services (as in Globus) can greatly enhance the accessibility and usability of those specialized resources.

- High-quality, persistent platform services can slash the cost of developing scientific applications and also promote interoperability. The Globus platform demonstrates this power at work.
- The need for sustainability so that projects can build with confidence: Platform services will only be used if projects (especially long-lived science facilities) can count on them persisting into the future. An unfortunate truth is that few science software projects have a workable sustainability model; the result is often reinvention.
- Subscriptions as a sustainability model that works: Globus demonstrates that subscriptions can be used to sustain useful services, in a manner that provides both close engagement with customers and the positive returns to scale (the more users, the more funds) that are essential to viability.

Barbara Helland (Department of Energy)

The Department of Energy's (DOE) Office of Science's (SC) mission is to deliver scientific discoveries and major scientific tools to transform our understanding of nature and advance the energy, economic and national security of the United States. SC is the Nation's largest Federal sponsor of basic research in the physical sciences and the lead Federal agency supporting fundamental scientific research for our Nation's energy future. SC accomplishes its mission and advances national goals by supporting:

- The frontiers of science—exploring nature's mysteries from the study of fundamental subatomic particles, atoms, and molecules that are the building blocks of the materials of our universe and everything in it to the DNA, proteins, and cells that are the building blocks of life. Each of the programs in SC supports research probing the most fundamental disciplinary questions.
- The 21st Century tools of science—providing the nation's researchers with 26 state-of-the-art national scientific user facilities - the most advanced tools of modern science - propelling the U.S. to the forefront of science, technology development and deployment through innovation.
- Science for energy and the environment—paving the knowledge foundation to spur discoveries and innovations for advancing the Department's mission in energy and environment. SC supports a wide range of funding modalities from single principal investigators to large team-based activities to engage in fundamental research on energy production, conversion, storage, transmission, and use, and on our understanding of the earth systems.

From the beginning, DOE has recognized the role of advanced computing and networking in addressing science and engineering challenges. DOE, through the Office of Science's Advanced Scientific Computing Research (ASCR) program, one of SC's six program offices, provides the science and technology community, access to world-class supercomputers and the tools to use them for science and engineering. ASCR accomplishes this by advancing research in applied mathematics, computer science, and advanced networking; developing and maintaining world-class computing and network facilities for science; developing exascale and future generations of computing hardware and software tools for science and engineering, in partnership with the research and U.S. industry; and training the future generation computational science workforce. The following sections describes each of these components of ASCR's portfolio.

ASCR Research. In the past, advances in computer science and applied mathematics often contributed even more than Moore's Law in increasing the performance of scientific codes on new architectures. But this is changing, increasingly ASCR needs to be agile to adapt to a rapidly changing computing landscape while still be focused on the development of future technologies such as quantum or neuromorphic computing. Emerging trends are pointing to a future that is increasing *instrumented* with data from sensors, satellites, drones, offline repositories; *interconnected* through the Internet of Things, composable infrastructure and heterogeneous resources; and *automated and accelerated* through the development of faster and flexible pathways that support real-time analysis of data couple with machine learning to gain scientific insights. To meet SC's and the Nation's future high performance computing (HPC) needs, ASCR's research program must focus on following challenges.

- *Extreme Heterogeneity with Post-Moore Technologies.* Basic research is needed in new algorithms, new software stacks, programming models to support the heterogeneous systems of the future which may include quantum and neuromorphic accelerators
- *Adaptive Machine Learning, Modeling, and Simulation for Complex Systems.* Algorithms and tools must be developed that support automated decision making from intelligent operating systems, in situ workflow management and improved resilience. In addition, basic research in uncertainty quantification and artificial intelligence is necessary to build statistically and mathematically rigorous foundations for advances in science domain-specific areas.
- *Data Tsunami.* With the rapid growth in data from SC user facilities, new investments are needed in software and a coordinated infrastructure to accelerate scientific discovery by addressing challenges and opportunities associated with research data management, analysis, and reuse.

In addition, ASCR must continue to accelerate progress in scientific computing through the Scientific Discovery through Advanced Computing (SciDAC) programs. SciDAC connects ASCR applied mathematicians and computer scientists with scientists in other SC disciplines to address two challenges: to broaden the community and thus the impact of high performance computing, particularly to address the Department's missions, and to ensure that progress at the frontiers of science is enhanced by advances in computational technology, most pressingly, the emergence of the heterogeneous and many-core architectures and machine learning techniques. It also serves to connect ASCR's Research program with ASCR's computing facilities.

ASCR Facilities. ASCR and its predecessor organizations have been providing high performance resources for the scientific community for over four decades, first through the high performance production computing at NERSC at LBNL and most recently, the Leadership Computing Facilities (LCFs) at ORNL (OLCF) and ANL (ALCF). NERSC delivers high-end capacity computing services for the SC research community, supporting over 7,000 computational scientists in about 700 projects annually. In contrast, the LCFs deliver the highest computational capability to the national and international researchers, including U.S. industry and normally each have approximately 1,000 users and 100 projects.

ASCR's computing facilities and the other SC research facilities generate many petabytes of data each year. Moving data to where it is needed requires advanced scientific networks and related technologies provided through ESnet, SC's high performance network user facility, delivering highly reliable data transport capabilities optimized for the requirements of large-scale science. ESnet currently maintains one of the fastest and most reliable science networks in the world with a 100 gigabit per second (Gbps) "backbone" network that spans the continental United States and the Atlantic Ocean. ESnet interconnects the DOE's national laboratory system, dozens of other DOE sites, and ~200 research and commercial networks around the world—enabling tens of thousands of scientists at DOE laboratories and academic institutions across the country to transfer vast data streams and access remote research resources in real time. ESnet's traffic continues to grow exponentially—roughly 66% each year since 1990—a rate more than double the commercial internet. As a user facility, ESnet engages directly in efforts to improve end-to-end network performance between DOE facilities and U.S. universities. ESnet is recognized as a global leader in innovative network design and operations, and is heavily engaged in planning a complete upgrade of its backbone network in order to create new and transformative opportunities for experimental and observational science. When completed, ESnet-6 will increase its capacity and capability to handle the expected exponential growth in network traffic; improve reliability and resiliency, including cyber resiliency; and provide flexible support for increasing complex workflows.

The facilities regularly gather requirements from the other SC research programs through a robust process to inform upgrade plans. These requirements activities are also vital to planning and to prioritizing ASCR research directions and inform the community of new computing trends, especially as the computing industry moves toward exascale computing. Because of the growth in data from simulations and experiments, ASCR's computing facilities are exploring new technologies for deploying shared archival storage. As a first step they released a released a request for information to survey the landscape.

The OLCF's deployment of the 200 petaflop Summit system is the result of past requirements reviews. Although designed to meet the nation's leadership expected computational science needs in the 2018-2022 timeframe, Summit is also capable of over 3 billion mixed precision calculations per second, or 3.3 exaops, for certain scientific applications, primarily because of the inclusion of graphical processing units (GPU). Summit represents the first step in the computational paradigm shift toward driving transformational science through the integration of artificial intelligence, enabling researchers to apply techniques like machine learning and deep learning. The NERSC-9 upgrade is expected to provide similar capabilities for SC researchers. ASCR's first exascale system, which is expected to be deployed at the ALCF in 2021, is being designed from its inception to support classical computational science simulations as well as data analytics and artificial intelligence investigations or a combination of all three methods.

Exascale Computing Initiative. Maximizing the benefits of U.S. leadership in computing in the coming decades will require an effective national response to increasing demands for computing capabilities and performance, emerging technological challenges and opportunities, and competition with other nations. As one of the leading Federal agencies in the National Strategic Computing Initiative (NSCI), DOE will sustain and enhance its support for HPC research, development, and deployment as part of a coordinated Federal strategy. Specifically, SC and the DOE National Nuclear Security Administration (NNSA) are partnering on the Exascale Computing Initiative (ECI) to overcome key exascale challenges in parallelism, energy efficiency, and reliability, leading to deployment of a diverse set of exascale systems in the 2021-2022 timeframe. The ECI's goal for an exascale-capable system is a fifty-fold increase in sustained performance over today's Titan system at Oak Ridge National Laboratory (ORNL), with applications that address next-generation science, national security, engineering, and increased convergence between exascale and large-data analytic computing. The two components of the ECI include the planning, site preparations, and non-recurring engineering (NRE) at the Leadership Computing Facilities (LCF) to prepare for deployment of at least one exascale system in 2021, and the Exascale Computing Project (ECP), which includes only the R&D activities required to develop exascale-capable computers.

Exascale Computing Project (ECP). To create an exascale ecosystem with critical methods and tools required for efficient and effective use of exascale systems, ECP has identified the following goals:

- Deliver mission results with exascale-ready applications, addressing currently intractable exascale challenge problems of strategic importance and national interest
- Deliver a sustainable software product suite required by exascale applications and platforms, sustainable into the future
- Deploy integrated ECP products on targeted systems at DOE HPC Facilities (pre-exascale and exascale)

- Transition results from the PathForward partnerships with vendors into Facility NRE to enhance capabilities of delivered exascale systems

To address the its goals, ECP has created three technical focus areas:

The Application Development (AD) effort develops and enhances the predictive capability of applications critical to the DOE, including the science, energy, and national security mission space through a co-design process to integrate of appropriate software and hardware.

Software Technology (ST) spans low-level operational software to high-level applications software development environments that will result in a comprehensive, coherent software stack that enables application developers to productively write highly parallel applications that effectively target diverse exascale architectures. Software Development kits (SDK) will be the key delivery vehicle for software products with a direct line of sight to application development projects. All SDK releases will be delivered as “build from source” via Spack with a focus on ensuring that the software compiles robustly on all platforms of interest to ECP (including testbeds at the facilities)

The Hardware and Integration (HI) focus area is centered on the integrated delivery of specific outcomes (ECP Key Performance Parameters, or KPPs) and products (e.g., science as enabled by applications, software, and hardware innovations) on targeted systems at leading DOE computing facilities. As a result, HI manages the interactions between ECP and the computing facilities. HI also includes the PathForward partnerships with vendors and Outreach and Training to share ECP accomplishments with the greater computational science community.

Computational Science Graduate Fellowship (CSGF). Recruiting, developing, and sustaining a workforce with top-notch computational science skills is necessary for sustaining national progress through scientific and technological discovery and innovation. Accordingly, the DOE Computational Graduate Fellowship (CSGF) program looks to develop the kinds of talent necessary to meet future workforce needs.

Started in 1991, the DOE Computational Science Graduate Fellowship (CSGF) was designed to broadly train advanced computational scientists. Because computational science is multi-disciplinary, fellows must follow plan of study that transcends the bounds of traditional academic disciplines with course work in applied mathematics, computer science as well as a scientific discipline. To provide a real world experience, the fellows must also participate in at least one twelve week practicum with a computational scientist at DOE National Laboratory. In addition to fully funding the fellows’ tuition and fees, CSGF also provides a yearly stipend of $36,000.

To ensure that CSGF is meeting its program goals, it undergoes program reviews and longitudinal studies. The most recent longitudinal study of CSGF, conducted in 2017, interviewed alumni dating back to the beginning of the program and found that “One of the benefits we confirmed was the substantial proportion of DOE CSGF alumni who transition directly to permanent employment at DOE laboratories. Meanwhile, most other alumni found employment in highly specialized positions in U.S. industry or in academia, where they are in a position to contribute to the development of the next generation of scientists. A substantial proportion of alumni overall have also achieved leadership positions within these various settings, thereby helping to spread the influence of the program beyond the government agency that sponsors the fellowships.” In addition, the study noted that “the fact that alumni have received a large number of professional awards and patents and published research in prestigious journals at an impressive and compounding rate serves as evidence of their contributions to the broader scientific community through the generation of new knowledge and innovations."

In summary, DOE's ASCR program gives the science and technology community, including U.S. industry, access to world-class supercomputers and the tools to use them for science and engineering. ASCR accomplishes this through coordinated activities to develop, operate and maintain world-class computing and network facilities for science; and advance research in applied mathematics, computer science, and advanced networking, prepare for exascale and future computer generations and train a computational workforce.

Miron Livny (Wisconsin)
Livny's input was based on his experience as the PI and Technical Director of the Open Science Grid (OSG), The Director of the UW-Madison Center for High Throughput Computing (CHTC), and a researcher in the area of High Throughput Computing (HTC) that I pioneered. In the past 12 month, OSG delivered more than 1.6B core hours and CHTC delivered more than 400M core hours. The HTCondor distributed job and resource management system is widely deployed by academic and industrial organization across the world.

Almost all aspects of research cross organizational. It is common practice today that researchers from different organizations to collaborate on addressing a science question. These organizations may be in different countries. Collaboration are founded on mutual trust and involves sharing of ideas and in many cases also involves sharing of resources. It is therefore critical and natural that the research computing infrastructure (also referred to as Cyberinfrastructure) that powers scientific discovery in a growing number of domains will facilitate effective computing across organizational boundaries. This is especially critical for HTC applications that can harness dynamic and heterogeneous computing resources.

All aspects of the research computing eco-system – planning, implementation, and operation – need to recognize the existence of boundaries and develop processes, procedures and technologies that work across them. It is common practice to point at "silos" and "stovepipes" as factors that hinder progress. The solution is not in eliminating (breaking) them but in developing means to interconnect these organizational structures as in many cases these organizational structures are there for a reason and facilitate local autonomy in planning, execution and operation. The boundaries provided by organizational structures constitute spears of trusts.

The science community has developed the organizational techniques, social skills and software technologies to facilitate international and cross agencies collaborative use of research computing. The research community offers people who know how to work in such environments, the national and international networking infrastructure offers the means needed to carry data and applications across boundaries and software tools that knows how to provision and use resources that are operated by different organizations located in different countries are available and deployed. These organizations include commercial providers of processing and storage capacities.

It is important to note that in the same way that funding agencies constitute an intra-national organization structure, these agencies have an internal structure that offers opportunities for collaboration, integration and joint investments. In the case of NSF, crossing the boundaries between directorates, divisions and projects to facilitate collaboration and sharing is an opportunity. Materializing these collaboration opportunities will require an effort to establish mutual trust between the organizations and to develop joint value proposition and metrics to monitor progress. This requires a commitment to sustained effort that tracks the evolution of the research computing eco-system and adjusts to changes in computational technologies and business models. Current indicators point at an acceleration in the rate of change in technologies. This makes multiyear decision and implementation process difficult.

Padma Raghavan (Vanderbilt)

The "big data revolution" in science and society is well underway. NSF and OAC should lead the global research community's response to shape and seize this opportunity. It is imperative that we not only rapidly harness scientific breakthroughs that will not be possible otherwise but also develop the ecosystem of innovation that can fuel sustained US global leadership in research and education long into the future.

NSF and OAC have pioneered an ecosystem of scientific discovery based on computational modeling which often demand large numbers of calculations (or operations) per unit data. In addition to the fact there is a growing gap between demand and supply of NSF hosted CI, the scientific data deluge will be truly disruptive with direct and indirect effects[5]. For example, while there will be direct effects related to the challenges of harnessing knowledge throughout the long scientific data life cycle there will also be indirect effects from how we collaborate and partner across agencies and universities to develop the CI of the future in which the costs of simply *moving* data (as opposed to operating on it) will dominate[6].

A major opportunity is for NSF to partner with other agencies that bring complementary strengths to collaboratively drive-up returns on investments. For example, with DOE's commitment to Exascale, which brings with efficiencies of scale in both acquisition and operation of CI, NSF should explore models to support needs of NSF funded scientific applications that require coupled parallel high-performance computing at the petascale and beyond. As there are only about 20 or so of such scientific applications that have continued to be scaled-up from the terascale to the petascale, a focused effort to benchmark and determine demand could be illuminating. This would allow NSF to direct resources to training and education and more importantly, to advance research in innovative CI through prototype systems with new memory architectures or software systems for data-intensive high-throughput workflows.

NSF should consider new models for "CI as a service" by removing barriers for the academic use of "infrastructure as a service" (cloud computing). NSF should develop commercial-federal-university partnerships to drive down costs and drive up researcher productivity with sustainable CI through models for "platform as a service" with toolkits and standards for specific domains/communities, or "software as a service" with subscription models for sustainability.

NSF should develop a mid-scale CI roadmap for data-driven discovery and innovation. The scientific data deluge from instruments such as the Cryo-EM, fMRI can easily top 100's of terabytes per day/week. Such volumes of data demand that some CI be close to the data source for initial pre-processing with seamless access to remote clouds for later workflows. The current CI network has many university-scale nodes and a few leadership class nodes with gaps in the middle. The opportunity is to develop a midscale CI program with new funding that is part of the overall NSF road for midscale scientific instrumentation.

While in the longer term we will need new programs with additional funding for midscale CI, in the near term, there is a critical need for re-balancing the NSF CI portfolio of resources. Scientific breakthroughs

---

[6] P. Kogge and J. Shalf, "Exascale Computing Trends: Adjusting to the New Normal for Computer Architecture," *Computing in Science & Engineering*, 2013

will increasingly depend on workflows that are data-intensive with data acquisition through instruments external to the CI. Strategic re-balancing and prioritizing is needed to advance these workflows which are woefully underserved by the current NSF CI resources.

Panel 4 Breakout Summary

Panel members noted that the list of potential federal and private collaborators was quite clear, though identifying mechanisms for collaboration was more challenging. As one discussant put it, science is uniform, but the infrastructure across agencies is not. Consequently, science suffers when researchers must themselves bridge this gap.

In that spirit, the lessons learned from the Globus project and the Open Science Grid provide some insights. These included a focus on services with economies of scale; enabling applications that span agencies, public cloud providers, and institutional resources; and facilitating international collaboration. All of these depend on sustainability, a prerequisite for building applications and workflows with long lifetimes. In turn, sustainability triggered a discussion of subscription service models and cost recovery, particularly the need to build subscription costs into institutional budgets.

Several participants raised the need for exemplar use cases to motivate collaborations and support. Absent an understanding of common needs, infrastructure developers and providers may build capabilities that either do not meet pressing needs or are less suited for collaborative science, particularly that spanning agencies and countries. In turn, this led to a discussion of interoperable campus infrastructure nodes, regional hubs, and national facilities, with an emphasis on data sharing and associated challenges.

Finally, the discussion returned to NSF’s cyberinfrastructure needs and challenges and how NSF might collaborate with and leverage investment by other Federal agencies and cloud providers. Several possibilities were discussed, including NSF negotiating with cloud providers to reduce prices for academic purchase or NSF purchasing cloud services and allocating them via peer review. In either case, the focus would be use of the unique hardware and software attributes of the cloud for scientific discovery rather than simply a low cost computing engine.

Appendix A: Workshop Agenda

Day 1 – Wednesday, May 30, 2018
7:30-8:30 **Registration**
8:30-9:15 **Welcome**
Erwin Gianchandani, Acting Assistant Director for CISE

Framing and Workshop Goals

Manish Parashar, Office Director, Office of Advanced Cyberinfrastructure/CISE
Dan Reed, University of Iowa/Utah

9:15-10:15 Scene Setting Plenaries

Rick Stevens, ANL, *Big Data, Big Science*
Ilkay Altintas, SDSC, *The Long Tail, Research and Discovery for All*

10:15-10:45 Break

10:45-12:00 Panel 1: Future Ecosystem Models and Application Drivers
Looking forward, what are the essential components of, and relationships contributing to the NSF-supported computational ecosystem? What are the key application drivers (and usage modes) that motivate these components?

Moderator: Rommie Amaro and Nancy Wilkins-Diehr
Panelists: Abani Patra, Don Krieger, Shaowen Wang, Frank Wuerthwien

12:00-1:00 **Lunch**

1:00-2:30 **Panel 2: Future Resource Management, Operation, and Delivery**
What are the appropriate models for, and relationships contributing to resource management, operation and delivery?

Moderator: David Swanson
Panelists: Dan Fay, Tom Furlani, Sharon Broude Geva, David Lifka
2:30-3:00 **Break**

3:00-4:30 **Breakouts**

Panel 1 Discussion: (Leads: Ewa Deelman**)**
Panel 2 Discussion: (Leads: Dana Brunson**)**

4:30-5:00 **Breakout Reports**

Dinner (On Your Own)

Day 2 – Thursday, May 31, 2018
8:15-8:30 **Reflections on Day One**

8:30-10:00 **Panel 3: Integration and Holistic Operations**
What are the opportunities, usage modes, and use cases for using future cyberinfrastructure components in an integrated and holistic manner?

Moderator: David Lifka
Panelists: Alan Blatecky, Robert Ricci, Dan Stanzione, Craig Stewart

10:00-10:30 **Break**

10:30-12:00 **Panel 4: Interagency and International Collaboration Opportunities**
What the opportunities and use cases for integration with complementary investments by other agencies in the US and internationally?

Moderator: Dan Reed
Panelists: Ian Foster, Barbara Helland, Miron Livny, Padma Raghavan

12:00-1:00 **Lunch**

1:00-2:30 **Breakouts**

Panel 3 Discussion: (Leads: Wendy Huntoon**)**
Panel 4 Discussion: (Leads: Gwen Jacobs**)**

2:30-3:00 **Breakout Reports**

3:00-3:30 **Workshop Summary**

Appendix B: Workshop Attendees

First Name	Last Name	Institution
Ilkay	Altintas	UCSD
Rommie	Amaro	UCSD
Karan	Bhatia	Google
Alan	Blatecky	RTI
Jim	Bottum	Internet
Dana	Brunson	Oklahoma State University
Paola	Buitrago	University of Pittsburgh
Aaron	Culich	UC-Berkeley
Ewa	Deelman	USC-ISI
Peter	Elmer	CERN
Doug	Jennewein	University of South Dakota
Dan	Fay	Microsoft
Ian	Foster	University of Chicago
Thomas	Furlani	SUNY at Buffalo
Rob	Gardner	University of Chicago
Sharon	Broude Geva	University of Michigan
David	Hancock	Indiana University
Barbara	Helland	Department of Energy
Wendy	Huntoon	KIMBER
Gwen	Jacobs	University of Hawaii
Klara	Jelinkova	Rice University
William	Kramer	University of Illinois
Don	Krieger	University of Pittsburgh
David	Lifka	Cornell University
Miron	Livny	University of Wisconsin
Amit	Majumdar	UCSD
Lauren	Michael	University of Wisconsin
Henry	Neeman	University of Oklahoma
Nick	Nystorm	University of Pittsburgh
Sanjay	Padhi	Amazon
Sudhakar	Pamidighantam	Indiana University
Abani	Patra	SUNY at Buffalo
Gregory	Peterson	University of Kentucky
Howard	Pfeffer	Internet2
Padma	Raghavan	Vanderbilt
Dan	Reed	University of Iowa/Utah
Robert	Ricci	University of Utah
Dane	Skow	North Dakota State University
Xiaohui Carol	Song	Purdue University

Dan	Stanzione	University of Texas
Craig	Stewart	Indiana University
Shawn	Strande	UCSD
David	Swanson	University of Nebraska
James	Taylor	Johns Hopkins University
John	Towns	University of Illinois
John	Vidale	USC
Shaowen	Wang	University of Illinois
Nancy	Wilkins-Diehr	UCSD
Frank	Wuerthwein	UCSD
P.K	Yeung	Georgia Tech