ACTI

What's Next for Campus Cyberinfrastructure?

ACTI Responds to the NSF ACCI Reports

July 2012

ADVANCED CORE TECHNOLOGIES INITIATIVE

EDUCAUSE

Authors and Contributors

Special thanks go to the following ACTI-CCI authors and contributors to this report.

Guy T. Almes, *ACTI-CCI Chair*
Director, Academy for Advanced Telecommunications and Learning Technologies
Texas A&M University

Vijay K. Agarwala
Senior Director, Research Computing and Cyberinfrastructure
The Pennsylvania State University

Asbed Bedrossian
Director of Enterprise Middleware Applications
University of Southern California

Rajendra Bose
Manager, Research Computing Services
Columbia University

Jill B. Gemmill
Executive Director, Cyberinfrastructure Technology Integration
Clemson University

Richard Greenfield
Strategic Research Analyst
University of Alaska

Thomas Hauser
Director of Research Computing
University of Colorado, Boulder

Curtis W. Hillegas
Director of Research Computing
Princeton University

Sally Jackson
Professor of Communication
University of Illinois at Urbana-Champaign

Richard Knepper
Manager, Core Services
Indiana University Bloomington

Timothy Lance
President and Board Chair
NYSERNet, Inc.

Clifford Lynch
Executive Director
Coalition for Networked Information

Donald F. McMullen
Cyberinfrastructure Strategist
Great Plains Network

Hideko J. Mills
Manager, IT Research Infrastructure (former)
University of Wisconsin–Madison

Gregory E. Monaco
Director for Research and Cyberinfrastructure Initiatives/Great Plains Network
Kansas State University

Michael R. Mundrane
Deputy CIO
University of California, Berkeley

Kim Owen
Advanced Applications Outreach Coordinator
North Dakota State University

James Pepin
Chief Technology Officer
Clemson University

David Stack
Deputy CIO
University of Wisconsin–Milwaukee

Table of Contents

1. Introduction 1
2. Grand Challenges 3
    2.1. Synopsis of Task Force Report 3
    2.2. Elements to Be Reinforced from a Campus Perspective 5
    2.3. Elements to Be Constructively Criticized from a Campus Perspective 7
    2.4. Implications for Campus Cyberinfrastructure Leadership 8
3. Campus Bridging 9
    3.1. Synopsis of Task Force Report 9
    3.2. Elements to Be Reinforced from a Campus Perspective 10
    3.3. Elements to Be Constructively Criticized from a Campus Perspective 12
    3.4. Implications for Campus Cyberinfrastructure Leadership 13
4. Cyberlearning and Workforce Development (CLWD) 14
    4.1. Synopsis of Task Force Report 14
    4.2. Elements to Be Reinforced from a Campus Perspective 15
    4.3. Elements to Be Constructively Criticized from a Campus Perspective 16
    4.4. Implications for Campus Cyberinfrastructure Leadership 17
5. Data and Visualization 18
    5.1. Synopsis of Task Force Report 18
    5.2. Elements to Be Reinforced from a Campus Perspective 21
    5.3. Elements to Be Constructively Criticized from a Campus Perspective 22
    5.4. Implications for Campus Cyberinfrastructure Leadership 25
6. High Performance Computing (HPC) 26
    6.1. Synopsis of Task Force Report 26
    6.2. Elements to Be Reinforced from a Campus Perspective 27
    6.3. Elements to Be Constructively Criticized from a Campus Perspective 27
    6.4. Implications for Campus Cyberinfrastructure Leadership 28
7. Software for Science and Engineering 28
    7.1. Synopsis of the NSF Task Force Report 28
    7.2. Elements to Be Reinforced from a Campus Perspective 31
    7.3. Elements to Be Constructively Criticized from a Campus Perspective 32
    7.4. Implications for Campus Cyberinfrastructure Leadership 34
8. Conclusion 37

Appendix A: ACTI-CCI Working Group Members 39

Appendix B: NSF ACCI Task Forces Chairs and Members 41

1. Introduction

The mission of the EDUCAUSE Advanced Core Technologies Initiative Campus Cyberinfrastructure (ACTI-CCI) Working Group is to help higher education institutions develop institutional strategies and plan resource deployments in the emerging and evolving technological landscape.[1] The mission extends to helping users at these institutions harness the power and capabilities of new integrated IT tools and systems for education and research. ACTI-CCI activities include sponsoring workshops, conferences, white papers, and other documents on these topics, as well as include interactions and close cooperation with federal funding agencies and others to ensure that grants for research and educational activities target key components and essential cyberinfrastructure tools, methods, and technologies.

In February 2009, a joint workshop of ACTI-CCI and the Coalition for Academic Scientific Computation (CASC) issued a report and recommendations that addressed the challenges and strategies for developing a coherent cyberinfrastructure from local campus to national facilities.[2] The joint report observed that extremely large computing clusters, such as those at federally funded centers, will provide and support excellent scalability for only a very few software applications. The report then noted the proliferation of 1,000–2,000 core clusters on many campuses. The report concluded that it is not only practical but also *optimal* to solve a large number of computational problems at the campus level.

The joint report immediately preceded the formation of six task forces by the National Science Foundation's Advisory Committee for Cyberinfrastructure (NSF ACCI), which were charged with investigating long-term cyberinfrastructure issues. Members of the ACCI task forces included hundreds of representatives from the education and research communities. The task forces solicited broad input from these communities to inform their work. In turn, institutions and individuals provided feedback and participated in the work of the task forces. Since the task forces' final reports were published in April 2011, they have been widely circulated and discussed at campuses and conferences.[3]

This ACTI-CCI white paper is a broad-based response to the findings of the ACCI task forces from the campus perspective. ACTI comprises more than 64 member institutions that include both higher education and regional networking organizations. Sixty-four representatives from 35 of these institutions serve on the Campus Cyberinfrastructure Working Group that is responsible for producing this work (see Appendix A: ACTI-CCI Working Group Members). With this white paper, ACTI-CCI aims to provide the higher education community with a

---

[1] See www.educause.edu/acti/working-group/acti-campus-cyberinfrastructure-acti-cci-working-group.

[2] EDUCAUSE Campus Cyberinfrastructure Working Group and Coalition for Academic Scientific Computation, *Developing a Coherent Cyberinfrastructure from Local Campus to National Facilities: Challenges and Strategies* (February 2009), http://www.educause.edu/Resources/DevelopingaCoherentCyberinfras/169441.

[3] The reports of the NSF-wide ACCI task forces are available at http://www.nsf.gov/od/oci/taskforces/.

thorough and thoughtful reflection on each of the six reports and an analysis of the role of campus cyberinfrastructure in each of these areas. Therefore, specific recommendations for campus leaders are included in each section. The ACTI-CCI also intends the analyses in this document to be of value to the ACCI and the NSF Office of Cyberinfrastructure (OCI) as it administers the Cyberinfrastructure Framework for the 21st Century Science and Engineering (CIF21) portfolio of programs.

This white paper includes six chapters that correspond to the ACCI task forces' reports. Each chapter has four components, including a synopsis of the task force report, a discussion of elements that warrant reinforcing from the campus perspective, a constructive critique of elements from the campus perspective, and a discussion of the implications of the report for the leadership of campus cyberinfrastructure.

In developing this white paper, the ACTI-CCI has concluded that campus cyberinfrastructure cannot be ignored when planning and developing the national cyberinfrastructure because many of the underlying services and activities that compose the national cyberinfrastructure are campus based.[4] Some of the examples discussed in greater detail in this document include the following:

- Advances in cyberinfrastructure are often conceived and developed at campuses.
- Campuses train the next generation of cyberinfrastructure-ready investigators and update the skills of current researchers.
- Distributed grid environments draw on campus computational resources, including cycles that would otherwise be lost.
- Campuses provide the champions who encourage the use of cyberinfrastructure, including outreach to new communities and disciplines.
- Campuses are multidisciplinary environments that serve as models for interdisciplinary collaboration at the national level.
- Campuses have committed to the ongoing curation of numerous data sets and collections.
- Visualization requires resources provided by campuses.
- Campuses are indispensable partners in determining access control and trust relationships between researchers and national centers.
- Much research software, especially open-source software, is written at campuses.

---

[4] This conclusion is consistent with the strategy outlined in the recent report from the National Academies, *Research Universities and the Future of America: Ten Breakthrough Actions Vital to Our Nation's Prosperity and Security*, available at http://www.nap.edu/catalog.php?record_id=13396.

2. Grand Challenges

2.1. Synopsis of Task Force Report

The overarching claim of the task force's final report is that cyber science and engineering (CS&E)—understood as computational science and engineering, including data-intensive computing—has emerged as a new yet central discipline in addressing the grand challenges of modern science and engineering.[5] This discipline, with a strong intellectual core in mathematics, computer science, and engineering, emphasizes collaborations among the traditional disciplines and must now be recognized as a strategically important discipline by the NSF and by related federal research agencies.

The report describes, in some detail, several specific grand challenges from discipline areas relevant to NSF's many divisions, and outlines the nature of the difficulties. While enhanced cyberinfrastructure (including the usual array of computing, storage, visualization, networking, software, and operational infrastructure) is required to meet these challenges, the report places emphasis on advanced computational methods and algorithms.

The Grand Challenges Task Force faced a delicate task. On the one hand, they repeated and strengthened the case, first articulated in the early 1980s, that computational science complements and strengthens theory and experimentation, the two "classical pillars" of science. On the other hand, they work to keep computational science and engineering unified with data-intensive approaches, to avoid the splintering or confusion that might result if data-intensive science were to be viewed as a totally separate "fourth paradigm."

Thus, by articulating a strong case for CS&E as an intellectually serious and innovative new discipline of strategic relevance to a broad array of grand challenges in the science and engineering areas important to NSF—and an enterprise requiring vigorous investment in cyberinfrastructure—the task force has produced a report that can claim to explain the relevance and importance of all six NSF ACCI task force reports.

While the descriptions of each of the specific grand challenges have value, the selection of these specific science and engineering problems as "grand challenges" is somewhat arbitrary. More interesting is the attention focused on six areas where the attention of NSF is directed:

---

[5] *National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Grand Challenges: Final Report, March 2011*, http://www.nsf.gov/od/oci/taskforces/TaskForceReport_GrandChallenges.pdf.

- **Advanced Computational Methods and Algorithms**
  The task force argues that computational methods "are often taken for granted due to past successes and their largely hidden role in powering CS&E applications."[6] They then proceed to make specific points about the need for advanced approaches, for example, to multiscale and multiphysics modeling, advanced discretization methods, scalable solvers, and uncertainty quantification. They note the value of addressing the dynamic data-driven application system (DDDAS) concept.
- **High Performance Computing (HPC) for Grand Challenge Problems**
  The task force supports the call for innovation in and provisioning of very high-end HPC systems. Of particular note is their comment that exascale systems will need to have data moving capabilities to match their computational capabilities. Similarly, they emphasize that system software—both operating systems and a "reinventing" of the Message Passing Interface (MPI) libraries to support exascale application parallelism—will be as critical as their hardware advances.
- **Software Infrastructure for Grand Challenge Communities**
  The task force calls for investment in applications, including data analysis and visualization applications, and in the staffing and software maintenance approaches needed to keep application software healthy over time.
- **Data and Visualization**
  The task force calls for a "data infrastructure" to deal with the "scale and complexity" of modern scientific data. They note, for example, that simply running HPC models at remote (e.g., TeraGrid, now XSEDE) sites and then copying the output data back for local data analysis and visualization fails since the growth in the size of such output data has dramatically outpaced growth in end-to-end network performance. This method calls both for improved remote visualization techniques and for improved access to remote data, including being able to mount remote TeraGrid file systems. It is of such importance that we treat it later in the context of the issues raised by the Data and Visualization Task Force (see Section 5).
- **Education, Training, and Workforce Development in CS&E**
  In contrast with the emphases of the Cyberlearning and Workforce Development Task Force (see Section 4), the focus here is specifically on preparing students to "design algorithms and write software for modern architectures." The challenge is not only to educate computational specialists, but to also educate new generations of computational scientists within the various disciplines. Curriculum that balances domain topics with mathematical/computational skills, software engineering skills, and the ability to exercise these skills over the coming range of data and computational scales will be difficult to achieve within current university graduate and undergraduate educational structures.

---

[6] Ibid., xvii.

- **Grand Challenge Communities and Virtual Organizations**
  The task force notes the value of the virtual organization (VO) idea in promoting effective harnessing of infrastructure and in supporting flexible forms of collaboration. They call for better understanding the nature of VOs and of collaboration and for the application of emerging best practices in VOs as a means of strengthening research.

2.2. Elements to Be Reinforced from a Campus Perspective

The most notable element of the report is the very articulate case it makes for cyber science and engineering as a field. The natural alliance between those engaged in CS&E as an academic discipline and those contributing to operational cyberinfrastructure, including at the campus level, is strategically important to the future health of American research universities and the broad scientific and engineering enterprise. Achieving the grand challenges of 21st-century science and engineering will require powerful computational, data, networking, and visualization infrastructure, but applying the best of such infrastructure in a brute force manner will not achieve success. Unhappily, for example, there are any number of application codes used by computational scientists that do not scale well beyond some number, often 1,000 or less, of computing cores. For such applications, the benefits of Moore's law will be very limited. A healthy CS&E research community is needed to overcome this problem.

The report makes an excellent case for improved curriculum relevant to CS&E. The curriculum issues addressed are excellent; they are focused on a healthy growing CS&E community and are essential if the grand challenges are to be met. Just as increasingly parallel cyberinfrastructure cannot be effectively harnessed in merely brute force ways, so also effective use of cyberinfrastructure in the cause of science and engineering cannot be accomplished unless young scientists and engineers in data- and computationally-intensive areas are provided improved instruction in these areas in relevant departments. This improved curriculum must address both students who see themselves as "computational scientists" and those primarily motivated by a specific discipline, and it must engage students well before their postgraduate studies.

The Grand Challenges report makes a number of points about data and visualization with greater clarity and forcefulness (in some cases) than does the Data and Visualization report. This is specifically true about issues relating to large-scale data and the need for data infrastructure. We support these points and comment in greater detail in Data and Visualization.

The report notes that high-end (e.g., exascale) cyberinfrastructure can no longer focus only on delivering double-precision floating point divides, but that organizing and moving large amounts of data, using effective data management, and organizing whole workflows will be needed. We support this emphasis on data movement capabilities. We also agree that the problem of moving large data in a geographically distributed environment is increasingly hard and not one that we can ignore. If scientists and engineers are to be able to make effective use of the variety of high-value cyberinfrastructure resources distributed at national centers and campuses, then addressing this problem successfully will be key to any comprehensive national cyberinfrastructure plan.

As the report mentions, techniques such as ensembles of independent yet similar model runs are very important, particularly in dealing with problems with uncertainty in parameters or in initial conditions. We emphasize a key consequence of this: that multiple parallel "large" systems are useful in addition to a few "very large" systems.[7]

The report makes cogent points about the need for improved software and an improved software maintenance life cycle. We support these points and emphasize that these developments will strengthen the value of cyberinfrastructure resources, both at national centers and on our campuses.

We also concur with the positive points made about the value of the InCommon federated authentication infrastructure and its application, e.g., to provide remote access to TeraGrid/XSEDE file systems. We note that an early project of XSEDE will be to create a geographically distributed file system, accessible at both campus and national center resources and containing data stored at both campuses and national centers. Effective use of InCommon and related technologies will be key to crafting secure and easy-to-understand approaches to controlled sharing of data in such systems.

We also share the concern that, so far, the NSF DataNet project has yielded few usable tools.[8] We hope this will soon change. The innovations emerging from DataNet and related efforts will strengthen the value to scientists of cyberinfrastructure resources on campus and at national centers and will lower barriers to the aggressive concerted use of campus and national center resources.

We further concur with stressing that virtual organizations could be a valuable way to integrate users across multiple facilities. Improvements to software and operational tools to support VOs will, again, strengthen the effectiveness of resources both on campus and at national centers.

The report articulates a case for continued advances in high-end (e.g., exascale) cyberinfrastructure. These advances are highly relevant to existing grand challenges and have historically benefited campus cyberinfrastructure. The need for appropriate investments in

---

[7] While the concepts "large" and "very large" will change over time, having parallel large systems where multiple jobs can execute concurrently is essential now and will remain important, even as very large systems continue to increase in relative size and become more affordable.

[8] The Sustainable Digital Data Preservation and Access Network Partners (DataNet) derives from the NSF Office of Cyberinfrastructure (see http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503141). A handful of projects have been funded under this program, the largest of which is DataOne (http://www.dataone.org/).

high-end infrastructure remains valid even in times of limited resources. We fully support this perspective and these recommendations.

2.3. Elements to Be Constructively Criticized from a Campus Perspective

One significant oversight of this report is that it does not discuss campus cyberinfrastructure in any detail. We note, positively, that the task force members are highly respected computational scientists who have no difficulty getting large allocations on the high-end NSF-supported (e.g., TeraGrid, now XSEDE) centers. We believe, however, that campus cyberinfrastructure is relevant to the other points made in the Grand Challenges report in several ways.

Computational science depends not only on high-end systems capable of very large, tightly coupled (e.g., MPI) model runs but also on the highly distributed fabric of instruments, data management systems, data-analysis systems, and visualization systems. The instruments (which comprise the ultimate source of data) and the optic nerves of the scientists (involved in viewing visualizations and sometimes steering computations) are, necessarily, highly geographically distributed. Much of this distributed infrastructure is on campuses. Planning for the combined infrastructure that will be relevant to scientists and engineers who are pursuing these grand challenges must consider national centers, campus cyberinfrastructure, and other elements.

Although the report does not explicitly point to the campus in its discussion of cloud resources, many of the points raised in this discussion could also be applied to typical campus resources. For example, in the section on HPC and the challenges of exascale computing, the task force reports that a "balanced provision of extreme and moderate scale computing is clearly needed as advocated in the three tier structure of [the 2007 NSF cyberinfrastructure vision document]" and that the "intermediate and lower end tracks must clearly include these mechanisms for modest scale HPC provisioning."[9] While applicable to commercial cloud resources, they are also applicable to campus cyberinfrastructure, including campus resources based on virtual machine technology and other campus resources.

The report's section on data infrastructure can be read in a way that assumes that network performance will inevitably be too weak to support the needs of CS&E users. This reading can lead to a conceptual framework in which CS&E users must place their large data in one location and keep it in that one place for the duration of a given project. This would, however, make it so that a user is unable to use both national and campus resources effectively, as well as create an unfortunate kind of "customer lock in" for the user who is only able to use one specific national resource in the cyberinfrastructure context. This rather dismal prospect must be addressed vigorously by cyberinfrastructure leaders, including at TeraGrid/XSEDE centers, at the

---

[9] *Grand Challenges Final Report*, 40.

campuses, and at network organizations such as Internet2 and ESnet. NSF must press these different sets of leaders to work in concert to prevent embracing this pessimistic view.

In the extreme case, whenever a return to "sneakernet" (the movement of large data sets via the physical transport of storage units) occurs, not only does data movement become cumbersome, but it also becomes very difficult to make use of modern techniques of access control, metadata, data presentation, and the data mining of multiple data sets.

2.4. Implications for Campus Cyberinfrastructure Leadership

The Grand Challenges report reminds us that, in addition to cyberinfrastructure "users" and cyberinfrastructure "providers," there are the computational scientists working to make cyberinfrastructure more valuable, the so-called bridging people. Campus CI leaders need to talk with them and support them. Cross-disciplinary efforts may also show that CI users in one disciplinary area can serve as bridging people in another. Cross-fertilization happens not only at the disciplinary level but also at the CI-support level.

The emphasis on data infrastructure presents campus cyberinfrastructure leadership with challenges and opportunities.

- If campus CI leaders can work effectively with the leadership of national cyberinfrastructure centers such as XSEDE and with networking leadership, then we can work toward a future in which healthy XSEDE and campus cyberinfrastructure resources are connected by powerful networks that, together, present users with a distributed data infrastructure of high flexibility, functionality, reliability, and performance. In this future, data sets seem light, the networks seem fast, and the effective use of multiple geographically distributed resources seems natural.
- If, on the other hand, this conversation fails to happen, we may be stuck with a future in which XSEDE and campus cyberinfrastructure resources are only weakly connected and thus appear, in the context of data-intensive science, as islands. In this future, data sets seem bulky and heavy, the networks seem slow or even highly viscous, and serious barriers to the effective use of multiple geographically distributed resources diminish the value of each of those resources.

This issue is closely related to those addressed in the Campus Bridging report in Section 3.

3. Campus Bridging

3.1. Synopsis of Task Force Report

The Campus Bridging Task Force emphasizes that the goal of campus bridging is to enable the seamless integration of a scientist's or engineer's personal cyberinfrastructure; the scientist's campus cyberinfrastructure; the cyberinfrastructure at other campuses; and regional, national, or international cyberinfrastructure.[10]

The report acknowledges the tremendous diversity of CI resources and the resulting opportunities of harnessing these resources and also the corresponding challenges in achieving effective integration.

The report focuses on several areas where the NSF should take leadership. It states that NSF should:

- Encourage the use of the InCommon global federated system in the services that NSF deploys and supports.
- Establish a blueprint for national cyberinfrastructure. This blueprint must use criteria appropriate for research infrastructure, include a national cyberinfrastructure software roadmap, and it must continue to fund campus cyberinfrastructure (e.g., through the existing MRI program).
- Create a new program funding high-speed connections from campuses to the nearest landing point for a national network backbone to support rapid data movement of large scientific data sets.
- Fund national facilities for at least short-term storage and management of data to support collaboration, scientific workflows, and remote visualization.
- Continue research, development, and delivery of new networking technologies.
- Support the evolution and maturation of cyberinfrastructure through careful analyses of needs and outcomes. This process will include collecting disciplinary community requirements and planning long-term cyberinfrastructure software roadmaps to support disciplinary community research objectives. All studies of cyberinfrastructure needs and outcomes should be published in open, refereed, scholarly literature.

The report also makes several strategic recommendations to campus leadership and the U.S. higher education community:

- Institutions of higher education should lead efforts to fund and invest in university-specific, state-centric, and regional cyberinfrastructure to create local benefits in research

---

[10] *National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Campus Bridging: Final Report, March 2011,* http://www.nsf.gov/od/oci/taskforces/TaskForceReport_CampusBridging.pdf.

accomplishments and economic development and to aid the global competitiveness of the United States.

- Every institution of higher education should have a cyberinfrastructure plan, developed and endorsed at the highest level of its governance.
- Institutions of higher education should adopt criteria for tenure and promotion that reward contributions such as data sets, online scholarly services, and relevant forms of software. Accomplishing this goal will require new forms of peer review.

For commercial cloud providers, the report recommends close collaboration to remove barriers for the use of their infrastructure by the U.S. research community.

3.2. Elements to Be Reinforced from a Campus Perspective

It is noteworthy that the areas of federated access and networking were among the task force report's strongest emphases, and that these areas are the most important from the campus perspective. Federated access and networking are crucial to allow utilization of resources based on existing researcher credentials and access. These, taken together, address the first barriers faced by science users as they approach the combined cyberinfrastructure.

The first recommendation to NSF concerned InCommon and related technologies and infrastructures as a means of promoting "[e]ffective, efficient federated identity management and authentication." We concur that this is foundational to campus bridging. Deploying and exploiting InCommon and these related technologies/infrastructures will address several campus bridging concerns, including:

- The ability of science users to have a consistent identity on their campus that can then be recognized on all cyberinfrastructure resources, whether deployed at the departmental, campus, or federal level.
- The ability of collaborations and virtual organizations to use these identities (and groups of these identities) to control access to resources and to information.
- The ability of cyberinfrastructure resource providers to support science users from several campuses in a controlled, yet flexible, manner.

In summary, rather than continue the current pattern of science users creating new user IDs at each resource, the community should leverage existing campus identities. Greater use of InCommon and improved software tools for exploiting InCommon will also provide a welcome counterexample to the general rule that convenience and security are in tension with each other. Improved convenience will actually strengthen security by using evolving InCommon tools that support both strong assurance that a given identity relates to a specific person with specific attributes, and tools that support strong (e.g., two-factor) authentication.

We also note that the increasingly sophisticated campus identity/authentication infrastructures that are federated via InCommon are themselves an important part of campus cyberinfrastructure. InCommon makes this distributed campus-level infrastructure crucially relevant to the national cyberinfrastructure area.

It will be important to track the success of technologies such as the NSF-funded CILogon service, which bridges from InCommon identities to grid certificates, and the increasingly secure forms of InCommon authentication mechanisms. Also, it should be noted that the success of these efforts will depend on establishing and maintaining trust relationships among resource providers (both on campus and at remote sites), campus identity/authentication infrastructure providers, and users.

Similarly, we reinforce the value of the report's call for NSF to initiate an appropriate successor to the historic "connections" program.[11] It is important, at this point, to note the similarities and the differences between the current situation and the situation during the mid- to late 1980s in which the original connections grants were made. In both cases there is a clear understanding that effective networking is key to the effective exploitation of the wide variety of cyberinfrastructure resources by university faculty, staff, and students. In the 1980s, however, the focus was very much on access *from* the campus *to* the new NSF supercomputer centers. In the present situation, this focus is enriched by the value of campus cyberinfrastructure resources, with particular emphasis on instruments (e.g., genome sequencers) and computational and storage resources that serve an increasingly data-intensive science and engineering world. We face two unhappy trends: (1) campus network staff are increasingly concerned with responding to challenges in the areas of security, privacy, and network management, and (2) the size of relevant scientific data sets is increasing more rapidly than the end-to-end performance of commonly deployed university networks. We note, therefore, the emphasis in the task force report recommendations on support for architectures and engineering appropriate for the "rapid movement of large scientific data sets." From a campus cyberinfrastructure perspective, we stress the emphasis, present in the report, that this movement is *among* cyberinfrastructure resources (located a various universities and labs) rather than merely *between* a campus (user) and a national center (resource).

In responding to this recommendation, we also point out the, perhaps obvious, importance of any NSF networking initiatives to dovetail with networking plans for XSEDE with university network initiatives such as Internet2 and with agency network initiatives such as ESnet.

---

[11] A few months before this report was finalized, NSF had announced a funding program responsive to this task force recommendation. This initiative is very welcome, though any substantive comment on the details of this program is outside the scope of this report.

We also point out the obvious but important tie-in between the Campus Bridging report recommendations on network infrastructure and the sections of the Grand Challenges report on the challenges posed by large data (particularly Section 6.3, "The Need for a Data Infrastructure").

An overarching idea in both reports, particularly from our working group perspective, is the importance of cyberinfrastructure leaders at the national (e.g., XSEDE), regional, and campus levels to communicate and work with each other in the interests of science and of science users. In this context, we welcome the early indications from the first months of the XSEDE project of a healthy awareness of this, including inclusion of campus and regional cyberinfrastructure resources in their architecture, and we look forward to working with XSEDE in making this a positive reality for science users.

Approaches, standards, and best practices for managing research outputs (which include primary data) comprise an area very much within the scope of campus bridging.

3.3. Elements to Be Constructively Criticized from a Campus Perspective

The TeraGrid/XSEDE Campus Champions program is not mentioned in the text of the task force report, and this is an unfortunate omission. Just as campus cyberinfrastructure leaders call for XSEDE leadership to take campus cyberinfrastructure more seriously, campus cyberinfrastructure leadership must also devote part of their limited resources to helping science users on campus make the most effective use of the XSEDE program and of its resources.

It would have been appropriate to discuss opportunities and challenges associated with the Open Science Grid, which has long provided a constructive way for campus cluster, storage, and network resources to contribute directly to the national grid infrastructure. One problem area for the campus bridging vision is that the Open Science Grid and the TeraGrid (and now XSEDE) have historically adopted differing approaches to the use of grid technologies, including X.509 certificates, which makes it difficult for campus resources to interoperate with these two grid efforts. We hope that technical infrastructures such as InCommon, Globus Online, and the Global Federated File System (GFFS) are developing with the goal of interoperability with each other and with campus cyberinfrastructure resources.[12]

The grand challenge problems, and the CI tools needed to solve them, are bigger than any one discipline, institution, or sector. A symbiotic partnership of academe, government, and industry is critical. The task force report fails to mention interaction with and contributions from

---

[12] We concur with the Campus Bridging report's recommended adoption of InCommon and mention Globus Online and GFFS as examples of relevant tools rather than by way of specifically recommending them. While we are happy to see XSEDE exploring these technologies, commenting substantively on specific XSEDE technology choices is outside the scope of this report.

corporate research, but as the corporate community applies greater resources than the entire NSF budget, the NSF is at risk of being relegated to a peripheral role.

Finally, we note that what is sometimes referred to as a cohesive campus cyberinfrastructure is often in reality a complex ecology of cyberinfrastructure resources on our campuses, each managed by such diverse entities as central IT organizations, departments, and individual research projects.

3.4. Implications for Campus Cyberinfrastructure Leadership

The task force's report makes a number of observations and recommendations that Campus Cyberinfrastructure Leadership should consider.

The first is very direct: "Institutions of higher education should lead efforts to fund and invest in university-specific, state-centric, and regional cyberinfrastructure." We underscore and expand on the importance of this commitment. Our universities should indeed fund campus cyberinfrastructure, including computing, storage, visualization, data, personnel, and networking assets on campus. During a time when our universities are under financial pressure, this is not always easy, but steady growth and stewardship of campus cyberinfrastructure during these lean years will be critical for the future of science and engineering research on our campuses. They should also work to develop partnerships with government and industry because the problems for which CI tools are most critical require multisector cooperation and coordination.[13]

Campus leaders should also support healthy state- and regional-level cyberinfrastructure. The most common example, and one that all universities must pay attention to, is the state- and/or regional-level advanced computer networks that connect each campus to the backbones of national university networks such as Internet2 and agency networks such as ESnet. Too often campuses keep their regional networks on "bread-and-water rations" with the minimum resources needed for lowest-common-denominator needs. One outcome from this that threatens a number of regional networks is a business model inordinately dependent on resale of commodity networking. Our science user community, especially those engaged in data-intensive science, needs effective, sustainable, high-speed, wide-area, end-to-end performance, which can only be achieved with healthy regional networks and by well-crafted connections to them (including intentional use of ScienceDMZ and perfSONAR technologies).

Similarly, the Campus Champions program of the TeraGrid, now reinforced in the XSEDE effort, provides an excellent opportunity for the campus to support its science users in their use

---

[13] See, for example, the High Performance Computing Consortium (HPC^2) in New York (http://hpc2.org/) and a new broad, multistate collaboration, the New York Genome Center (http://www.nygenome.org/).

of cyberinfrastructure and also to create and enrich communications and a supportive relationship between XSEDE leaders and campus cyberinfrastructure efforts.

The second recommendation is more subtle: "Every institution of higher education should have a strategic plan, developed and endorsed at the highest level of its governance, for the establishment of a coherent cyberinfrastructure." There must be sufficient engagement, both by high levels of the administration, but also by campus IT leaders and by leaders among the computational and data-intensive scientists on campus and with their counterparts in government and industry, to ensure strong buy-in from all key stakeholders. The detailed recommendations found in the Campus Bridging workshop report are worth reading.[14] Not only are such plans of great value to campuses, but so is the participatory process of doing the planning. Accomplishing this goal, however, requires time, patience, and intense sharing of ideas and listening to others.

The final strategic recommendation is also important: "Institutions of higher education should adopt criteria for tenure and promotion that reward the range of contributions involved in the production of digital artifacts of scholarship." We note that effecting such change will require communication among campus cyberinfrastructure leaders, leaders of departments and other academic units, and the faculty to achieve, and we appreciate the difficulty of bringing it about. As with the importance of working with the entire university community to build a healthy strategic cyberinfrastructure plan, however, this conversation is important to have if the potential contribution of cyberinfrastructure-enabled science and engineering is to be achieved.

4. Cyberlearning and Workforce Development (CLWD)

4.1. Synopsis of Task Force Report

The charge for the Cyberlearning and Workforce Development Task Force originated from NSF's *Cyberinfrastructure Vision for 21st Century Discovery*, which identifies many workforce development goals, including fostering the broad use of cyberinfrastructure-enabled learning and research environments, supporting the development of new professions needed for CI-enabled opportunities, and promoting a fuller utilization of the potential U.S. workforce.[15]

---

[14] Patrick Dreher et al., eds., *Campus Bridging: Campus Leadership Engagement in Building a Coherent Campus Cyberinfrastructure Workshop Report* (2011), https://pti.iu.edu/sites/default/files/cbtf_cio_sm.pdf.

[15] National Science Foundation, Cyberinfrastructure Council, *Cyberinfrastructure Vision for 21st Century Discovery* (NSF 07-28), March 2007, http://www.nsf.gov/pubs/2007/nsf0728/index.jsp.

The CLWD final report does not address which technical skills are needed or how to train people in technical fields; rather, it concerns how to adjust our most fundamental ideas about education and workforce preparation for the 21st-century global economy.[16]

The report advocates revolutionary change (e.g., reinventing the education system around cyberinfrastructure and creating new academic-industry-government cyberinfrastructure institutes), while acknowledging the need for evolutionary improvements (e.g., adding new topics to the curriculum, such as data-intensive science, and changing delivery methods for learning material).

The report focuses on the idea of the "Continuous Collaborative Computational Cloud," or C4, which stresses the evolution of society toward greater and more continuous connectivity, a more thorough embedding of information resources into ordinary life, and a greater reliance on social interaction as a source of both learning and acting. The claim is that C4 is driving "a fundamental reorganization in the process of the creation and intergenerational transmission of knowledge."[17] C4 is linked to the topics of pervasive and ubiquitous computing, which are now critical areas of research and teaching.

The task force urges an NSF commitment to the fuller utilization of human capacity—more people participating, with educational opportunity across the spectrum of individual capability and across demographic segmentation. The report further stresses that the NSF needs to invest in the discovery of educational methods that are proven effective across all demographics and the entire spectrum of individual ability.

4.2. Elements to Be Reinforced from a Campus Perspective

A central premise of the report is that the nation (which includes those of us in the CCI community), with guidance from the NSF and corporate research leaders, should be aiming to educate everyone: that is, preparing everyone in the world to contribute to the global economy that has been created through information and computing technologies. The ACTI-CCI group broadly agrees with the report's contention that the growth of research/education has created not only the opportunity but also the imperative to integrate cyberinfrastructure into teaching and learning. Cyberinfrastructure challenges our ideas of education and highlights the need to assess what we are currently doing.

The concept of C4 must be informed by developments such as the Amazon Elastic Compute Cloud and similar services that have far greater resources at their disposal and a customer base

---

[16] *National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Cyberlearning and Workforce Development: Final Report, March 2011*, http://www.nsf.gov/od/oci/taskforces/TaskForceReport_Learning.pdf. A 2008 NSF report covers related topics and provides an alternative perspective: NSF Task Force on Cyberlearning, *Fostering Learning in the Networked World: The Cyberlearning Opportunity and Challenge* (June 24, 2008), http://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf08204

[17] *CLWD Final Report*, 33.

for sustainability. These CI tools that may have germinated in academe were, in fact, created independent of it and have greater capability to become C4 than an NSF-directed project.

There was general agreement among ACTI-CCI members with the report's view that educators need to adjust their teaching to reflect how people do science and engineering today, emphasizing multidisciplinary areas and the important area of continuing professional development as technology changes. The challenge of educating everyone is a scientific one that must be met through research on learning. Revolutionary change as called for in the task force report probably includes very fundamental change in the structure of educational institutions.

4.3. Elements to Be Constructively Criticized from a Campus Perspective

Many ACTI-CCI members were unfamiliar or uncomfortable with the C4 terminology. Some within ACTI-CCI feel that C4 will be mistaken for a design proposal rather than a description of a visible direction of ongoing change. Comments were made that C4 is an attempt to "impose" a design on something that is evolving on its own, and some felt the C4 graphic in the report is not very helpful in clarifying either the overall sense of C4 or the implications of C4 for our overall education system. Further, some members felt that the report neglected to fully appreciate ongoing work in educational technology and related work on standards for content and tools.

Another concern from our working group is that K–12 teachers are not supported well in learning about or using C4-related technologies. In discussions, ACTI members questioned where educational content is coming from and how it will be delivered to K–12 classrooms (which often lack the necessary technology), as well as how classroom time will be used and whether there will be a shift away from the lecture model.

At the October 18, 2011, ACTI meeting, ACTI-CCI members expressed the concern that social issues overshadow technical problems in education. Experienced educators in the fields of math and science underscored that educators must respond to existing programs (e.g., No Child Left Behind, etc.) that cannot be ignored while pursuing the use of advanced technology in the classroom. Others commented that a decline in classroom performance is linked to social rather than technical problems. For example, one instructor observed the trend that his recent classes in math had a harder time grasping abstract concepts than those in previous years.

The report, in advocating revolutionary change, will also require the NSF itself to consider organizational barriers to revolutionary change. We note that the report recommendations are predicated on the NSF embracing the very bold proposals for supporting learning science more comprehensively and "merging the technological and cognitive communities."

4.4. Implications for Campus Cyberinfrastructure Leadership

Assuming our economic "sector" (education) is about to go through a long period of highly disruptive change, we believe the CCI community needs to consider how IT support services are organized now on campus and what changes may be needed in the near future. Two specific areas of report have implications for those of us who are part of the CCI community on university campuses.

The first area concerns data-intensive science and the future directions of computational science, as well as the impact these new programs will have if enacted on IT and research staff hiring, retention, and professional development.[18] A more involved IT infrastructure will be needed on campuses to support the report's vision. Staff and faculty with computational skills will need to be recruited and hired on campuses, and new organizational support models will be needed. In addition, we believe cultural change is needed within the computer science discipline. Many computer science departments focus on research publications and abstract problems as opposed to applied, disciplinary examples and problems. Furthermore, university administrators need to appreciate the computational aspects of interdisciplinary work and understand the infrastructure, including staffing needs, required for such work.[19]

The second area involves the campus bridging discussions and the call to create new academic-industry-government cyberinfrastructure institutes.[20] The creation of such institutes will certainly change the landscape of research activity for those of us in academia. The structure of the current NSF XSEDE program versus the previous TeraGrid project may be a preview of the types of changes called for in this area.

---

[18] Ibid., chapter one.

[19] See, for example, the National Academies study under way under the auspices of its Board on Research Data and Information. More information on the study, "Future Career Opportunities and Educational Requirements for Digital Curation," can be found at http://sites.nationalacademies.org/PGA/brdi/PGA_069853.

[20] *CLWD Final Report*, chapter four.

5. Data and Visualization

5.1. Synopsis of Task Force Report

5.1.1. Main Points from the Data and Visualization and Grand Challenges Reports

In the executive summary of the Data and Visualization report, the authors identify six main areas of "challenges and opportunities that will require focused and sustained investment with clear intent and purpose."[21] They are, in brief:

1. **Infrastructure Delivery**: Deliver data infrastructure and services that are essential research assets associated with software and visualization tools and that require specific budget allocations to establish and maintain.
2. **Culture and Sociological Change**: Develop academic rewards for those who maintain research data sets, services, or software.
3. **Roles and Responsibilities**: Develop a shared data stewardship model involving everyone who uses, stores, or provides information. A trust relationship model is needed re: ownership and interdependencies of data.
4. **Economic Value and Sustainability**: Develop realistic cost models addressing institutional/national repositories and services.
5. **Data Management Guidelines**: Identify and share best practices for data management.
6. **Ethics, Privacy, and Intellectual Property**: Invest in training the research community in privacy-preserving data access.

The task force concludes its summary by stating that they have "not identified any recommendations for specific investments in visualization." The report does, however, recommend the findings of the Software for Science and Engineering Task Force report, which in turn references the 2006 *NIH-NSF Visualization Research Challenges Report*.[22] Because the recommendations regarding visualization in the 2006 NIH-NSF report appear to address our primary visualization concerns, we summarize that report's key observations and recommendations:[23]

- Visualization hardware capabilities are advancing consistent with Moore's law. The current constraint has remained in display technology with a limited number of pixels per display screen as compared to 600 DPI printing technology.

---

[21] *National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Data and Visualization: Final Report, March 2011*, 7, http://www.nsf.gov/od/oci/taskforces/TaskForceReport_Data.pdf.

[22] See the *National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Software for Science and Engineering: Final Report, March 2011*, http://www.nsf.gov/od/oci/taskforces/TaskForceReport_Software.pdf; *NIH/NSF Visualization Research Challenges Report* (January 2006), http://vgtc.org/wpmu/techcom/national-initiatives/nihnsf-visualization-research-challenges-report-january-2006.

[23] *NIH/NSF Report*, 24.

- NIH and NSF should make coordinated investments in visualization.
- NIH and NSF should review and create new funding/policy changes designed to encourage domain scientists to partner with visualization researchers, with the aim of making visualization expertise as integral to biological sciences as statistical expertise.
- NIH and NSF should develop a "coordinated and sustained national investment in a spectrum of centralized and distributed research programs to promote foundational, transitional, and applied visualization research in support of science, medicine, [and] business"—emphasizing foundational research and collaboration with domain specialists.
- NIH and NSF should create and maintain curated data collections and open-source software and provide long-term funding for these.
- More research is needed in how and why visualization works; this investigation would explore visual representation's design space, develop new interaction approaches, and exploit novel display hardware.
- The long-term sustainability of software should be addressed in addition to long-term management of data.

Unfortunately, the 2006 recommendations are not available to the Data and Visualization report reader, thereby leaving the (perhaps incorrect) impression that all investment should be directed to big data challenges, not visualization.

To better understand the NSF ACCI perspective on data and visualization, we also need to consider additional recommendations in this area made in the Grand Challenges report. That report pointed out that digital scientific data and visualization have strategic value (they are transforming science) to the Grand Challenges communities, and data science and data infrastructure are major components of research infrastructure. Data and visualization recommendations from the Grand Challenges report focused on infrastructure and education:[24]

- Fund research on data management, network infrastructure, data analysis, and data visualization to manage the pipeline from field instrument to large-scale data analysis to end-user visualization and to the public and policy makers.
- Support data-intensive computing.
- Build next-generation data scientists able to work in multidisciplinary teams.
- Build data-curation professionals to support metadata, indexing, and access who can work in multidisciplinary teams with domain scientists.

In the section on data and visualization, the Grand Challenges report makes a detailed and powerful case for the need for a data infrastructure. As data sets become larger and more complex, the role of storage, movement, and management of data become more difficult and

---

[24] *Grand Challenges Final Report*, 62.

more important. One source of the problem is that, while the leaders of the early NSF supercomputer centers and NSFNET were aware of the interrelatedness of their efforts, computing and networking leaders during the PACI and TeraGrid eras often lacked that shared sense. The significant and well-engineered networking effort within the TeraGrid, for example, was focused on connecting TeraGrid resources to each other rather than working with others to strengthen end-to-end performance between TeraGrid resources and their users. Given that the Grand Challenges Task Force consisted of computational science experts rather than networking experts, they naturally reflected the consequences without drawing attention to the causes. Thus the report's observation that "Large scale data easily generated in a few days may [take] weeks to transfer back to home institutions."[25] Given the seriousness of the report's call for an improved data infrastructure and the difficulties of achieving such an infrastructure, several points can be made relevant to a campus cyberinfrastructure perspective:

- Improved high-speed networks are needed with end-to-end performance adequate for the needs of data-intensive science. Since the writing of the report, major national networks, including Internet2, ESnet, and many of the leading regional optical networks are moving to 100 Gb/s link speeds, with the capability of several parallel such links on each optical fiber in wide-area networks. These networks need to be harnessed to support the data-moving needs cited in the report.
- To exploit high-speed networks and the InCommon federated authentication infrastructure, appropriate middleware is needed to "facilitate migration, data integrity, caching." Given the right architectures, middleware, and networks, campus infrastructure will better be able to contribute to the data infrastructure called for.
- The report calls for enhanced data access approaches, including the ability "to mount remote disk resources from the TeraGrid locally." We concur with the report, and are glad to see the XSEDE activity in this area.[26]

If the report's call for an improved data infrastructure is read in the context of the combined XSEDE plus campus cyberinfrastructure used by university researchers, the resulting opportunities and challenges are evident.

5.1.2. Comments on the Consistency of the Two Reports

The reports of the NSF ACCI task forces were never intended to be viewed as a single coherent work product. They were used by NSF as input to the Cyberinfrastructure Framework for the 21st Century Science and Engineering (CIF21) report. Similar to how CIF21 tied together the reports of the ACCI task forces for NSF, the following is an attempt to identify how, taken

---

[25] Ibid., 57.

[26] XSEDE, the Extreme Science and Engineering Discovery Environment, aims to support access both to XSEDE storage resources from the campus and to campus storage resources from XSEDE service providers; see http://www.xsede.org/.

together, the Grand Challenges and Data/Vis reports might provide recommendations for the area of data and visualization.

While not inconsistent with each other, the Grand Challenges and Data and Visualization reports emphasize very different areas. For example, whereas the Data/Vis report states that cultural change regarding data is essential within the scientific community, this topic is barely mentioned in the Grand Challenges report. The Data/Vis report also identifies privacy, distributed access control, and intellectual property policy as important and critical research issues, while the Grand Challenges report treats these topics as solved issues and instead points to the need to disseminate best practices in these areas to practitioners. Finally, where the Data/Vis report reduces the role of NSF to that of a catalyst—without addressing where funding for community data and visualization should come from—the Grand Challenges report specifically calls for funding for these two areas.

Overall, the Grand Challenges report makes a more forceful case that visualizing data has large gaps that need to be addressed, particularly for large, complicated data sets; in contrast, the Data/Vis report is silent on the concept that exascale/extremely large data sets may be an important, special problem. The Grand Challenges report, however, does not identify any grand challenge specific to data or visualization (see Section 5.3.2).

5.2. Elements to Be Reinforced from a Campus Perspective

Four emphases in the Data and Visualization report coincide with rapidly growing needs we see from a campus perspective.

5.2.1. Data Infrastructure Investment and Service Models

In pointing out the value of direct investment in data services and visualization tools the report states: "The NSF has the opportunity to provide leadership in developing sustainable service models... that are a prerequisite for contemporary science.... [T]he intent is to flag that data management and software tools (e.g., for visualization) are a necessary investment for breakthrough science and not a hidden cost."[27] Campuses are in extremely early stages of addressing the data deluge. IT organizations and/or libraries frequently face expectations from faculty and administrators regarding the availability of unlimited, archival quality, searchable (usable) storage and typically have no specific plan for funding such needs. Developing scalable best practices in curation, metadata standardization and use, data preservation, and data access will be essential as data grows in the years ahead.

---

[27] *Data and Visualization Final Report*, 23.

5.2.2. Need for Culture/Sociological Changes

The task force remarks on the need for campus cultural change resonate with our campuses' need for IT and data domain specialists who can work in multidisciplinary teams. This change will advance more rapidly when domain scientists are rewarded academically and professionally for applied, collaborative work, whether that work occurs in a tenured faculty role or professional staff role. There is a diversity of opinion about large-scale data integration; however, researchers will need to adjust to NSF's expectations if they want to continue to pursue that funding stream.

5.2.3. Need for Shared Stewardship Models and Sustainability

As the report indicates, there is a need for data stewardship models that identify "data owners" at various stages of the data life cycle while recognizing that data ownership is a role/responsibility that may transition over time. In addition, data preservation also requires stewardship, shared responsibilities, and a shared cost model.

5.2.4. Provide Best Practices

Finally, the report correctly recognizes that campuses need proven methods for dealing with data collections and learning how to build sustainable collections, including determining what to keep, how long to keep it, and whether institutional investments can leverage or be leveraged at a domain or regional level or on another basis. Community and common practice building that combines funding and propagation of best practices from NSF, the Department of Education, information sciences, library sciences, and other federal and private sources would be beneficial for campuses.

5.3. Elements to Be Constructively Criticized from a Campus Perspective

5.3.1. Funding Support

The chief shortcomings of the report lie in its omissions. The overall focus of the Data and Visualization report is on data management. In our experience, there is also a need to identify and report on challenges and frontiers in visualization research and the application of visualization within disciplines.

By not including any specific recommendations for future investments or programs in data visualization, the report gives the impression—whether intended or not—that there is no need to provide coordinated funding for visualization research and for visualization within disciplines. We believe that it will be essential to commit funding to the areas identified in the 2006 *NIH-NSF Visualization Research Challenges Report* and listed in Section 5.1.1 in order to address the complete spectrum of big data challenges.

5.3.2. Lack of Progress in Addressing Visualization Challenges

As noted above, recommendations from the NIH-NSF 2006 report on visualization challenges (summarized in Section 5.1.1) are excellent, yet little progress has been made—they remain as challenges in the 2011 task force report. This is not a critique of the Data/Vis report, but does show an apparent lack of progress in solving challenges from five years ago and speaks to the need for additional investment and emphasis in support of progress.

Neither the Data and Visualization nor the Grand Challenges report identifies any specific "grand challenge areas" for either data management or visualization. Because big challenges frequently drive innovation and global change, identifying these might make identifying funding priorities and organizational models clearer.

5.3.3. Stakeholder Representation

One question that arises when reviewing the task force is whether its members sufficiently represented the broad data and visualization communities. Task force membership appears to have drawn from a distinguished set of data experts, but disproportionately from outside the United States (especially from the United Kingdom) and from industry (especially Microsoft). As a result, the report may not have received adequate input from the visualization, DataNet, and library communities that engage in this area. For example, the report failed to include any discussion of federating data, which is an area that frequently arises when this topic is discussed broadly.

5.3.4. User Education

We believe there should be more emphasis on user education to make visualization relevant to the broadest range of users. This recommendation appears in the Grand Challenges report but not in the Data and Visualization report. Our campus experience leads us to believe that more people would use visualization, or would increase the sophistication of tools they are currently using, if they knew what was available and had assistance using the tools. Another factor, specific for big data, is that we haven't yet worked with big data sets enough to know how visualization can help. There isn't enough of a sense yet within the community about the new techniques that will be needed for visualizing big data.

5.3.5. Visualization as a Communication Tool

Visualization is especially important to make data "speak" to students and the general public. The importance of visualization in making science more understandable through education and outreach should be emphasized.

5.3.6. Partnerships

The data deluge requires that NSF develop cross-agency and industrial partnerships. We recommend that more consideration be given to the role of industry and other federal agencies in solving problems identified in the Data/Vis report.

5.3.7. Access Control—and Trust Relationships

Data management is not a solved problem. Access control (referenced as one aspect of privacy and IP in the 2006 report) and trust relationships remain challenging. The growing amount of digital data requires more complicated access-control–supported embargoed data and complex interorganizational trust relationships. InCommon provides a good start but does not adequately address some long-term needs. The task force report relegates this to education in best practices, and may underestimate the research challenges that remain.

There are some serious challenges in declaring and using trust relationships in access control across multiple administrative domains. Data sets will last for long periods of time (decades or centuries), and therefore the access and control problem will also be long-lasting. Solutions may also endure for long periods of time (decades or centuries) and will be critical enabling factors for the construction and operation of large-scale federated data collections.

5.3.8. Big Data Analytics

We believe that it is time to adopt the use of a term that is broader than "data and visualization" that captures dealing with the data itself, data mining, and machine learning. Data analytics is as important as simulation for advancing science. A new forum may be needed for the entire community to discuss questions related to data (e.g., visualization, data mining, algorithms, analytic methods, and anonymization) and the ways in which these new techniques and developments relate to more traditional practices and research issues in data management. This seems to be consistent with the multiagency "big data" initiative that the Obama administration announced in the spring of 2012, which includes investments in these areas.[28]

---

[28] For more information on the big data research and development initiative, see the White House Office of Science and Technology Policy website at http://www.whitehouse.gov/administration/eop/ostp; access the archived webcast of the March 29, 2012, announcement at http://www.nsf.gov/news/news_videos.jsp?cntn_id=123607&media_id=72174&org=NSF.

5.3.9. Models for Sustainable Collections

Models for sustaining collections of irreplaceable data generated by large collaborations are needed. One key cross-cutting area is that of models for sustaining collections of irreplaceable data generated by large collaborations. This is primarily a social and economic problem, but it directly impacts the availability of data and requires innovation in both technology and organizational models of scientific collaborations operating across many institutions in many countries. The time and money already spent to develop such models for experimental data from the Large Hadron Collider and for the Orion/Neptune ocean observatory indicate the complexity of data management for working experiments without addressing the long-term needs to preserve these data.

5.3.10. End-to-End Infrastructure Issues

End-to-end infrastructure issues underlie planning for big data analytics. Big data is an end-to-end issue related to networks, particularly network requirements for linking users, data repositories, computing elements and visualization facilities in a productive way; this challenge was not mentioned in the Data/Vis report but is vitally important. For example, being able to predict the performance of an analytical pipeline that draws on data and compute elements from many locations will make researchers more efficient in designing analyses and computational experiments. Although cloud computing and storage models have great potential in research, scaling data-intensive computing into data and compute clouds to get a desired level of performance remains largely a mystery.

5.4. Implications for Campus Cyberinfrastructure Leadership

The campus data infrastructure will be a foundational component for big data analytics, and campuses will expect that IT and libraries provide the investment.

The campus bridging model may help provide visualization with best practices for how a campus can provide visualization support and expertise in collaboration with project-specific participants and national resources (such as the TACC Visualization center).

Campus IT organizations should seek to discover what visualization resources are being asked for and identify the technical, financial, and staffing challenges in providing visualization support. Campuses can help identify the use of visualization for on-campus education and also the points of strain.

XSEDE cyberinfrastructure should coordinate with campus cyberinfrastructure with respect to data and visualization. Various design choices for visualization have different trade-offs; we recognize that there is no "one size fits all." Yet some choices, if widely adopted, would enable broader support for users. For example, if XSEDE clusters were supported together with visualization at the server side, there might be broader acceptance of consistent technology.

6. High Performance Computing (HPC)

6.1. Synopsis of Task Force Report

The High Performance Computing Task Force was chaired by Thomas Zacharia from the University of Tennessee and Oak Ridge National Laboratory and co-chaired by Jim Kinter from the Center for Ocean-Land-Atmosphere Studies. In addition to the chairs, the task force included 15 representatives from universities, computing manufacturers, and government laboratories, plus Rob Pennington, who served as the NSF liaison.

The task force was charged to focus on six areas:

1. Ensuring access
2. Application development and support
3. Computer science and engineering
4. Integration of research and education
5. Training
6. Policy implementation

The task force addressed these areas through a series of three workshops: "Sustainability for the National Computational Cyberinfrastructure," "Application Drivers for Exascale Computing and Data Cyberinfrastructure," and "Broader Engagement and Workforce Development."[29] In its final report, the task force made four overarching recommendations to the NSF:[30]

1. Develop a sustainable model by 2015–16
2. Invest now for exascale access in 2018–20
3. Broaden outreach
4. Continue the process for gathering community input

The detailed discussion in the report focused on three major areas:

1. Cyberinfrastructure sustainability
2. Exascale computing
3. Broader engagement

---

[29] See the Community Input on the Future of High-Performance Computing Workshop Series website at http://www.nics.tennessee.edu/workshop.

[30] *National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on High Performance Computing: Final Report, March 2011*, http://www.nsf.gov/od/oci/taskforces/TaskForceReport_HPC.pdf.

6.2. Elements to Be Reinforced from a Campus Perspective

The report provides a clear vision for the highest end and national centers. Specific goals were provided for both the near-term hundreds of petaFLOP (2015–16) and long-term exaFLOP (2018–20) capabilities. An important aspect of this vision is the reinforcement of the idea that to make exascale efforts successful, the NSF should commit to stable and sustained funding for HPC centers. This recommendation was given in conjunction with encouragement for HPC centers to develop long-term relationships with vendors.

The report also provides strong support for engaging a broader range of user communities in HPC and reaching out to disciplines that are not currently using HPC extensively in their research or delving deeper into disciplines that already are important, specifically in cases where scholarly advances can be achieved through the use of HPC.

6.3. Elements to Be Constructively Criticized from a Campus Perspective

This report could have reached deeper and taken greater risk when considering recommendations—particularly in the longer term.

Specifically, this report focused almost entirely on what have been considered Tier 1 and Tier 2 centers without acknowledging the role and continued need for growth in smaller regional and campus HPC centers—both as feeders for the larger national centers and as production engines where a large amount of scientific activity and discovery take place. The scaling of many applications does not go beyond 1,000 or 2,000 cores, and while in the past campus and regional centers were not able to handle models that large, many are now quite capable of running them. By neglecting these centers, the report misses a very large opportunity to grow the utility of HPC.

The report also could have paid more attention to the growing importance of big data and the growing challenge of data mobility as data sets grow faster than networking infrastructure. This area was identified as a growing concern, but there is a need to develop strategies to address the challenges presented by the large data sets that will need to be the input and products of petascale and exascale systems.

Computing architectures and programming models are in a period of rapid change and development. While the report acknowledges that there is uncertainty in this area, we need a strong vision for how academia, industry, and computing vendors can work together to drive the future and ensure that the HPC community has a clear path to leverage new and improving technologies.

6.4. Implications for Campus Cyberinfrastructure Leadership

Campus and regional HPC centers have served an important and enduring role in the development and sustainability of HPC in academic research. The talent and projects that germinate at and are cultivated by these smaller centers grow to be the next generation of key users and leaders at the national centers. A plan that does not include these smaller centers will be counterproductive to developing a strong national cyberinfrastructure.

Computing architectures are entering a phase of transformational change. A clear plan is needed to engage computing manufactures with academia, national laboratories, and industry to seize and steer the resulting opportunities for the advancement of HPC in research.

With changing architectures, programming techniques and system practices will transform. Continued engagement among the key stakeholders listed above will be necessary to ensure that knowledge and expertise in these areas keep up with available hardware architectures. The community needs to remain particularly aware of the fact that staffing varies widely among the constituencies, and we have to make sure that researchers and computing professionals (ranging from graduate students or even undergraduates through seasoned professional programmers and engineers) have access to the tools and training they need to leverage the new architectures.

Campuses should look for opportunities to actively engage with the NSF to help shape the future of HPC.

Campuses should engage in education, outreach, and training programs sponsored by the NSF, and they should provide feedback on the effectiveness of those programs.

7. Software for Science and Engineering

7.1. Synopsis of the NSF Task Force Report

Software has been identified as a "critical and pervasive component of the cyberinfrastructure for science and engineering."[31] This is entirely consistent with numerous other reports both contemporary and historical. It is a critical success factor that is emphasized in the final report from the Software for Science and Engineering Task Force, while recognizing that, though "software is *a*, if not *the*, 'grand challenge' of cyberinfrastructure… [it] is historically among the least coordinated and systematically funded components of cyberinfrastructure."[32] Because software is, in a sense, virtual instead of tangible like hardware, the life cycle does not constitute an easily structured period of use after which it can be comfortably decommissioned. A return to original data for the purpose of verification or repurposing often requires access to original

---

[31] *Software for Science and Engineering Final Report*, 4.

[32] Ibid., 4.

software. Continued use and emphasis on software almost invariably requires some level of ongoing commitment and development. Despite these facts, it was identified in the report that access to funding for the long-term maintenance and/or evolution of software was difficult to obtain from funding agencies. The challenge in obtaining stable resources in concert with the monotonically increasing complexity of software efforts combine to seriously undermine this critical support element of science and research. The task force established a goal to provide NSF with recommendations on how to "meet the demand on software to deliver ubiquitous, reliable, and easily accessible computer and data services."[33]

The National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Software for Science and Engineering was established in June 2009 and received a charge that asked it to review, assess, and make recommendations associated with three specific areas:

- *Identify specific needs and opportunities across the spectrum of scientific software infrastructure*. Characterize the specific needs and analyze technical gaps and opportunities for NSF to meet those needs through individual and systemic approaches.
- *Design responsive approaches*. Develop initiatives and programs led (or co-led) by NSF to grow, develop, and sustain the software infrastructure needed to support NSF's mission of transformative research and innovation leading to scientific leadership and technological competitiveness.
- *Address issues of institutional barriers*. Anticipate, analyze, and address both institutional and exogenous barriers to NSF's promotion of such an infrastructure.

In their research, the task force noted tension between "organically" developed community codes and the need for learning and adopting rigorous software engineering practices and validation for sustainability. Communities of interest pursue their independent interests despite the absence of sufficient numbers of adequately trained software designers and implementers. The results naturally vary significantly in quality and sustainability. In addition to this, the task force pointed out the importance of compilers and run-time systems that are able to:

- Leverage complexity in multicore systems.
- Generate good code for a variety of CPUs, memory configurations, and co-processors, such as MICs and GPGPUs.
- Understand and optimize execution based on actual performance information.

In their report, the task force recommended that NSF:

---

[33] Ibid., 4.

- Develop a multilevel (e.g., individual, team, and institute), long-term program of support of scientific software elements, including support of extreme scale data and simulation and NSF's major research equipment and facilities (MREFC) projects.
- Promote verification, validation, sustainability, and reproducibility through software developed with federal support.
- Develop a consistent policy on open-source software that promotes scientific discovery and encourages innovation.
- Support software collaborations among all of its divisions, related federal agencies, and private industry.
- Obtain community input on software priorities.

Specifically, code reuse and hiding hardware-dependent complexity are identified as problems arising in computer-intensive science that still need special attention (although linkages to and interactions with computing architecture programs within the NSF were not mentioned).

- NSF should support the development of portable systems through such things as automatic code generation and auto-tuning approaches.
- NSF should encourage close communication between chip designers, system builders, and software developers through appropriate collaborative research grants.

Other recommendations include:

- NSF should support standards development in both application-specific data formats and generic requirements for multiscale, multimodel integration.
- NSF should support the development of new numerical libraries and the sustainability of existing ones, which will provide a bootstrap for old and new application codes.
- Support is needed to enable collaboration with industry computer vendors and software developers.

Finally, the report also addressed software as it relates to data, federation, and collaboration. Big data problems now exist at the laboratory level, where data from one experiment can be in the gigabyte to terabyte range (and more) when data are aggregated across experiments. Historical paradigms and approaches for analysis have not scaled sufficiently with the increased size of even casual data sets. To support interpretation, findability, and cross-domain reuse, internal and external annotations for data files should be based on ontologies developed and supported at the domain level.

7.2. Elements to Be Reinforced from a Campus Perspective

Overall, this is an impressive and intelligently written report that addresses the three key charges through exploration and makes recommendations in each category. It recognizes and acknowledges that software is a critical and pervasive component of the cyberinfrastructure for science and engineering, but that the NSF does not adequately support its evolution and the corresponding life-cycle costs. This issue is compounded by steadily increasing code complexity that "could exceed the capacity of the national centers for software development and support." The task force was sympathetic to the idea that it is a challenge for software to address extreme scale (for both data and computation), but there is widespread belief that the current state of software and programming is not scaling well to multicore 1,000-processor systems.

In addition, the report articulated a need for generic data modeling and integration software. "Software is required to assist with data modeling, with the representation and exchange of semantic information, and with the mechanics of large-scale data integration." Although these tools currently exist, they have yet to be fully adopted by computational scientists and technologists, and the NSF could play a role in helping encourage discovery and implementation. The report identified an action item to generate a catalog of sustainable software. Further, the report recommended that the "NSF should support the development of new and sustainability of existing numerical libraries. This will provide a bootstrap for old and new application codes."

The section of the document on software evolution introduces a corresponding taxonomy and acknowledges the importance of developing ontologies and vocabularies for data formats. The task force presents a number of interesting ideas about organizing the evolution of software and data—a clear recognition that massive growth in data-set size has huge implications. Portability was strongly advocated, with recommendations that the "NSF should support development of portable systems through such things as automatic code generation and auto-tuning approaches" and that the "NSF should encourage close communication between chip designers, system builders, and software developers through appropriate collaborative research grants." They recommend an emphasis on best practices and that the "NSF should support standards development in both application specific data formats, and generic requirements for multi-scale, multi-model integration."

The Grand Challenges report notes "The Message Passing Interface (MPI) based programming model based on an inherently flat architecture … will need to be reinvented to meet application challenges…." Unfortunately, it also acknowledges that current programmers lack sufficient skills to address these issues broadly. While key simulation codes grow organically, they are fundamentally unsustainable either because of the increasing complexity or through the misallocation of resources through directly duplicative efforts. This can be mitigated through the establishment of guidelines and standards, but the report rightly contrasts the formulation of standards against the desirability of innovation. It establishes the NSF as having a clear role in supporting both the creation of standards as well as providing incentives and rewards for innovation.

It was acknowledged that software infrastructure would be driven by applications in four key areas of science—astrophysics, atmospheric science, evolutionary biology, and chemical separations—but it also identified an absence of systems and period processes for broadly determining software requirements or use. An interesting recommendation was that software could automatically measure how often it was used so that the NSF and others could make effective recommendations based on that information.

7.3. Elements to Be Constructively Criticized from a Campus Perspective

Software infrastructure is not limited to the support of astrophysics, atmospheric science, evolutionary biology, and chemical separations. Each of these key areas has unique needs, but many of their challenges also exist in other areas. There is an opportunity to distinguish clearly between general software challenges that also apply to the four key areas and specific challenges for each area.

While there was a recommendation to encourage collaboration between hardware and software vendors, the key suggested approach was through grants. There are other options. The NSF could develop an efficiency score based on energy use, for example, and have teams compete for a prize or an award or a certification based on that. The creation of efficient, capability-balanced nodes that use the minimum power to achieve their results would require significant engineering effort that combines both hardware and software, but that would be an interesting alternative to the current approach of designing around fast, high-wattage CPUs in order to maximize single node overall performance. So while it is laudable to encourage collaboration, it is suggested that this collaboration also be focused in the direction of deliverables relevant to researchers. Concerning data management and metadata, the empowerment of libraries or other groups to establish clear metadata standards for specific research domains would greatly facilitate the ability to work across and among specific communities. General data formats should be advocated over specific data formats, and we believe this should be an immediate and high priority. The long-term preservation of results requires the incorporation of metadata early in the process or it tends not to happen. The opportunity cost of delay is potentially profound.

Portability is a generic proxy for a more specific vision of workflow and dataflow from "office-to lab-to campus-to external-to center." While the report acknowledges the need for portability, and recommends that portability be encouraged and supported, a stronger statement would be more clearly supportive of research and researchers at all levels. Context is relevant, and the impact of barriers to movement depends greatly on the phase of research. The application of portable or global file systems is a specific approach that attempts to resolve portability issues and to encourage moving data to research. Working with ever-larger data sets, however, will increasingly emphasize moving research to data instead, and this has the potential for far greater compatibility issues, depending on the approach.

While the report acknowledges the creative tension between standards and innovation, it does not provide sufficient insight into how to achieve an appropriate balance between the two. The creation of software standards can take place either through directed projects, indirectly through the peer-review process that emphasizes one activity over another, or organically via market forces. Regardless of process, these standards naturally have the potential to facilitate reuse and extension. What is important to acknowledge is that standards are, in general, at cross-purposes to innovation. Standards contribute to stability and predictability within the area they attempt to define by minimizing variability in order to maximize translatability. Innovation can result, but it is often orthogonal to the standard's intent. This creative dynamic exists along any typical maturity curve.

The current MPI is adequate, but it requires too much effort on behalf of developers to achieve efficient results. There has been progress, but the current state of compilers and cluster operations neither automates optimization sufficiently nor adequately abstracts low-level complexity away from the developer. The report did not stress the importance of fault-tolerant and self-validating software, which will certainly be required for very large scale computation and to deal with the existing and increasingly complexity of software implementations. This challenge is intensifying within the context of MapReduce parallelization. This production approach is an evolving research technique that must become much more readily accessible to research developers because of its potential to vastly improve big data analysis. As a National Research Council committee stated in 2008, "The increases in code complexity could exceed the capacity of the national centers for software development and support."[34] This is true along the entirety of typical research workflow: model creation, calculation, analysis, visualization, curation. EDUCAUSE and ACTI-CCI could emphasize the need for focused software engineering education programs for undergraduate students that encourage interest in research applications and in graduate training programs as professional development.

Many programmers lack skills, but it is simply not practical to bring all developers up the complexity curve to ensure broad productivity in the current environments. This was true in the past and is becoming increasingly true as software development complexity inexorably

---

[34] Committee on the Potential Impact of High-End Computing on Illustrative Fields of Science and Engineering, National Research Council, *The Potential Impact of High-End Capability Computing on Four Illustrative Fields of Science and Engineering* (Washington, DC: National Academies Press, 2008), 59, http://www.nap.edu/catalog.php?record_id=12451.

increases.[35] The challenge in the software development space today is that the current set of technologies that must be mastered and integrated to be truly productive has become quite large and is getting larger. These technologies present a hurdle for serious developers but are potentially insurmountable for more casual developers. Many research teams do not have dedicated, full-time professional developers and rely instead on individuals that serve in multiple roles. While some can be quite proficient, developers will always be of varying skill. The capability of research team developers can be increased by making either the developer or their tools more capable. We believe that the former is a path of rapidly diminishing returns while the latter offers better long-term possibilities.

Power usage can be addressed through a number of means, not merely by having software interact with hardware to power down elements. In fact, having researcher software micromanage the power state of nodes would add yet another complexity element to their code and is in direct contrast to other observations mentioned here. Optimizing hardware and system software specifically for scale and power consumption might be a better approach and is more broadly translatable (at the system level by systems engineers and administrators) instead of within the software of individual researchers.

Understanding the demographic impact of any specific software product is key to setting relative priorities between competing options. An accounting model that tracks installations or launches is one way to facilitate a data-driven process. The application of gaming theory to solicit directed feedback and leverage the wisdom of crowds could be an alternative approach. Both methods have the potential to provide data to the process. What is not explicit in either approach is a deterministic mechanism for ensuring that market-like dynamics will produce good results over time.

The recommendation for the NSF to support numerical libraries could go further with a vision that includes permanent support for standardized general libraries that would otherwise languish. Individual compiler vendors would support the general language libraries because they would not be able to sell the compiler otherwise. This is not so for science-related libraries. There is typically not a sufficient market to make these maintainable. The NSF could form permanent programs to support and evolve basic science libraries.

7.4. Implications for Campus Cyberinfrastructure Leadership

7.4.1. Open Source

Commercial software, unless it is general, often cannot justify the development and support effort for specialized communities without imposing extremely high marginal costs. This

---

[35] Software sustainability is a problem that suffers from neglect generally. The report, *Cyberinfrastructure Software Sustainability and Reusability: Report from an NSF-funded Workshop*, speaks directly to this issue (https://scholarworks.iu.edu/dspace/handle/2022/6701).

constitutes a high barrier to entry and has the perverse impact of excluding communities of interest that are most in need of access to specialized software. It is not simply a matter of research communities partnering with commercial providers to provide additional, targeted functionality. Their respective interests are generally misaligned: the community is typically concerned with supporting research in general, while the vendor is often trying to protect intellectual property and preserve future profits.

Open source, while having significantly lower licensing costs, is not automatically the most cost-effective approach to solving software needs. Well-funded research activities that produce effective software solutions that are released to the public provide important and valuable contributions to their respective communities. The software itself is essentially a highly subsidized byproduct of the research activity. This does not, however, constitute an automatic market or a commitment to ongoing support and evolution. Communities that support open source through research generally do so in an asymmetric fashion, with most persistent endeavors anchored by a strong, core group of contributors. This dynamic, coupled with the increasing complexity of software, means that there is a risk of the most capable contributors changing focus and exiting an activity that cannot be reasonably absorbed by those that remain.

Despite this risk, open-source software is still being investigated as a realistic alternative to commercial software. Open source does not have a profit emphasis driving decisions and has the potential to permit subsequent developers the option of directly leveraging accessible code to permit meeting specific use cases. For example, general open-source geographic information systems (GIS) options have continued to be competitive with commercial options and currently provide much of the same or even better functionality than commercial packages. In the GIS use case, resources are sufficient and developers with the required skill can be brought to bear. Yet even with more specialized or obscure use cases, it is possible for a core group of motivated developers to advance open-source options. The increasing complexity of code development and the commensurate associated costs, however, make narrow use open-source projects increasingly difficult, and activities have been trending in the alternative direction of broader markets to drive interest and effort. This is not dissimilar to commercial efforts and may produce similar outcomes.

7.4.2. Software Discussion

The prior state of high-performance research was driven by tightly coupled hardware and operating systems. This, in conjunction with optimizing compilers, enabled researchers to readily develop software that performed quite well for them under most circumstances. The performance complexity was largely subsumed by highly skilled software developers focused on these specific hardware/software couplings. Current research has increasingly moved in the direction of leveraging more general infrastructure that is not highly optimized at either the hardware or software level and that often only exhibits reasonable overall performance.

An approach to dealing with the performance of specific elements of the research environment is to scale the research more broadly. The ability to deploy research on the vastly increasing amounts of inexpensive and cheap hardware partially ameliorates the situation, but existing research problems, software, and techniques generally do not scale to arbitrarily large environments. The evolution of highly optimized hardware allowed systems engineers to optimize performance at the machine level. This can continue, in some sense, at well-supported national centers. Campuses, acting both individually and in close coordination with each other, will continue to represent pathways to national centers and, as such, should not be ignored. Unfortunately, campuses have vastly different economic realities. On campus, what cannot be addressed in hardware must be addressed in software: generic hardware and available operating system kernels are the norm. Researchers are finding that hand optimizing their software for commodity hardware in a campus environment achieves mixed results.

Likewise, the hardware performance curve is scaling core-wise and not speed-wise. The ability to rapidly optimize multithread or multiprocess analysis is demonstrably more difficult than the same activity performed on a sequential environment. Additionally, there is a locality complication because data exchange between respective cores differs from data exchange between respective CPU chips or between respective compute nodes. The future state for traditional sequential research is also multithreaded, at least at the core level, in order to compensate for CPU speeds that are not increasing. Software needs to automatically or at least trivially scale node-wise and core-wise to support traditional independent jobs.

Progress in software that is explicitly research domain-specific is helpful but is highly likely to be unsustainable over time. Common aspects of research that can be modeled and better supported have a much better chance of being sustained if their impact is broad. That which serves broadly will typically not serve with great specificity. Respective research domains will certainly need to extend broad software tools and techniques in order to provide for their domain specific needs. Well-supported general research tools can be funded in a targeted fashion and still have broad impact. This is true regardless of business model, with both open-source and commercial endeavors being far more likely to persist if they serve research needs broadly.

7.4.3. Need for New Paradigms

Large numbers of researchers have made significant contributions to software development using current methods. These approaches and methodologies have sustained research communities through an historical software-driven period of growth, but it has become increasingly clear that they are reaching the point of diminishing returns. The tendency toward increased complexity within any existing set of methods or paradigms is both normal and natural. Software development in support of research has followed this trend, and it has now become impractical for research teams to have developers or individuals serving in that role with sufficient skill to effectively produce custom software in support of their needs. This problem can be partially addressed by improving the abilities of research teams, but this will not break software development out of the current complexity spiral. We believe the traditional approach to software development does not address this issue.

Even within the context of hardware clustering, there now exist two entirely distinct paradigms. The traditional perspective of high-performance computing underscores the challenge of providing a balanced set of nodes to deliver significant compute (either numerical or graphical) capabilities and fast, deterministic interprocess communication. The existing MPI-based approach on today's standard hardware configurations using existing software is not scaling well enough for many researchers. Likewise, big data and its corresponding need to traverse permanent storage for repeated analysis suggest a very different approach with data layout, bandwidth to storage, and structured process management representing the significant challenges. In this use case, interprocess communication tends not to be synchronization bound and so represents a small fraction of the overall effort. Big data research today often fits more naturally in the MapReduce solution space. This represents a new approach, but to a different class of research problems.

The computer science curriculum currently falls short of providing sufficient skills for either of these increasingly complex approaches. While there is nothing inherently wrong with pedagogy that begins with single-processer serial use cases to introduce basic concepts, it does not go far enough to adequately provide for current and future researcher needs. To best prepare future researchers, the current curriculum needs to be refreshed to reflect future-looking skills that rely heavily on the decomposition of problem domains and the corresponding system and programming challenges that this introduces. Current and future researcher needs will not necessarily drive the market, however, and it is not clear that computer science programs will naturally evolve in a direction to address them.

8. Conclusion

ACTI-CCI acknowledges the considerable efforts of the six ACCI task forces and their volunteer members in developing the reports under significant time constraints while simultaneously soliciting and including the feedback from hundreds of representatives from the education and research communities. The task force reports are valuable contributions to the national dialogue on developing a comprehensive plan for a national cyberinfrastructure that aids researchers in making wise and informed use of both federal and local resources. ACTI-CCI commends the

NSF for commissioning these reports and encourages the agency to draw on their recommendations when planning future programs.

Higher education leaders are encouraged to appreciate the need for and the value of investing in local cyberinfrastructure resources, not as an alternative but as a complement to federal initiatives. Effectively bridging local and national resources is a necessary condition for the development of a national big data effort as announced by the Office of Science and Technology Policy (OSTP) on March 29, 2012. Campus-level efforts will entail both organizational and cultural changes because their central IT organizations are often organized to focus support on administrative and academic computing. Although campus research centers can assist, they cannot by themselves build a seamless cyberinfrastructure with peer institutions or national centers because they are embedded within their respective campus infrastructures.

Powerful visualization tools are available in the commercial and open-source markets, but there is no widely known default solution available for researchers to get started and reasonably expect assistance from hallway colleagues or help desk staff. Therefore, campus IT organizations must undertake the time-intensive process of engaging with research groups one at a time until the institution has an overall understanding of the types of visualization that are needed or until a limited set of breakthrough tools become available that researchers can adopt on their own.

Even if they are not yet aware of it themselves, there are researchers at every institution whose scholarship would benefit from more powerful computational tools. These researchers are unlikely to spontaneously jump from using only the capabilities of a desktop PC to running computational jobs at a national center without a measure of local support, which is why the bridging of personnel, outreach, and training is just as important as the bridging and networking of hardware resources. This bridging will only be accomplished by the national centers reaching out to the campuses and vice versa with the goal of meeting midway.

The greatest need for investment in human capital is in the realm of software. Just as learning to shift with a manual transmission is no longer considered a prerequisite to having a driver's license, educational curricula must be updated to reflect the comprehensive paradigm shift away from single-processor, serial architectures. In addition, many midcareer researchers need to retool to take advantage of new programming paradigms. At most institutions, making these adaptations will require spending considerable time in academic committee work and training, which are necessary though not often glamorous undertakings. Meanwhile, maintaining a sustainable software environment will continue to be a nonautomated, labor-intensive task into the foreseeable future.

ACTI-CCI welcomes the opportunity to engage in active dialogue with the NSF regarding the future of the national cyberinfrastructure at both the federal and local levels.

Appendix A: ACTI-CCI Working Group Members

The following were members of the EDUCAUSE Advanced Core Technologies Initiative Campus Cyberinfrastructure Working Group at the time this report was published.

David Ackerman
Associate Vice President, .edu Services
New York University

Vijay K. Agarwala
Senior Director, Research Computing and Cyberinfrastructure
The Pennsylvania State University

Guy T. Almes, *ACTI-CCI Chair*
Director, Academy for Advanced Telecommunications and Learning Technologies
Texas A&M University

Greg Anderson
Retired Senior Director
University of Chicago

Katarzyna Azzara
Technology Enterprise Architecture Director
Columbia University

Charles Bartel
Director, Global IT Services
Carnegie Mellon University

Asbed Bedrossian
Director of Enterprise Middleware Applications
University of Southern California

Rajendra Bose
Manager, Research Computing Services
Columbia University

James R. Bottum
Vice Provost for Computing and IT and CIO
Clemson University

Perry Brunelli
Director of Network Services
University of Wisconsin–Madison

Duncan Buell
Professor, Computer Science
University of South Carolina

Douglas Carlson
Associate Vice President, Communications and Computing Services
New York University

David Crass
Director of Cyberinfrastructure
University of Wisconsin–Milwaukee

Alan Crosswell
Associate Vice President and Chief Technologist
Columbia University

James Cuff
Director of Research Computing and Chief Technology Architect
Harvard University

James Davis
CIO
Iowa State University

Wilson Dillaway
Infrastructure Planning and Security Officer
Tufts University

Thomas Dodds
CIO and Vice President for Information Technologies
Cornell University

Michael Fary
Enterprise Data Architect
University of Chicago

Steve Fleagle
CIO and Associate Vice President
The University of Iowa

Jill B. Gemmill
Executive Director, Cyberinfrastructure Technology Integration
Clemson University

Rick Golden
Assistant CIO
University of Nebraska

Richard Greenfield
Strategic Research Analyst
University of Alaska

Ted Hanss
Chief Information Officer, Medical School
University of Michigan–Ann Arbor

Thomas Hauser
Director of Research Computing
University of Colorado, Boulder

Curtis W. Hillegas
Director of Research Computing
Princeton University

William F. Hogue
Vice President for Information Technology and CIO
University of South Carolina

Sally Jackson
Professor of Communication
University of Illinois at Urbana-Champaign

Kurt J. Jeschke
Manager, Network Engineering
The Pennsylvania State University

James A. Jokl
Associate Vice President and Chief Enterprise Architect
University of Virginia

Paul Killey
Executive Director of IT, College of Engineering
University of Michigan–Ann Arbor

Kenneth J. Klingenstein
Senior Director, Middleware and Security
Internet2

Richard Knepper
Manager, Core Services
Indiana University Bloomington

Boyd Knosp
Associate Dean, Information Technology
Carver College of Medicine
The University of Iowa

John E. Kolb
Vice President and CIO
Rensselaer Polytechnic Institute

Steve Krogull
Director, Systems Engineering and Operations
University of Wisconsin–Madison

Timothy Lance
President and Board Chair
NYSERNet, Inc.

Betty Leydon
Vice President for Information Technology and CIO
Princeton University

David Lifka
Director, Center for Advanced Computing
Cornell University

Clifford Lynch
Executive Director
Coalition for Networked Information

Donald F. McMullen
Cyberinfrastructure Strategist
Great Plains Network

Hideko J. Mills
Manager, IT Research Infrastructure (former)
University of Wisconsin–Madison

Gregory E. Monaco
Director for Research and Cyberinfrastructure Initiatives/Great Plains Network
Kansas State University

Randy Monroe
Associate Director, Network Services
Carnegie Mellon University

Kevin Morooney
Vice Provost for Information Technology/CIO
The Pennsylvania State University

Michael R. Mundrane
Deputy CIO
University of California, Berkeley

Anita Nikolich
Executive Director, Infrastructure
University of Chicago

Michael Norman
Director, San Diego Supercomputer Center
University of California, San Diego

Kim Owen
Advanced Applications Outreach Coordinator
North Dakota State University

James Pepin
Chief Technology Officer
Clemson University

Valerie E. Polichar
Infrastructure Analyst/Technical Projects Coordinator
University of California, San Diego

Jerry Protheroe
Manager, Research Services
The University of Iowa

Lynn Rohrs
Manager, eSystems and HPC Services
New York University

Brian Seiler
Network Analyst
University of Illinois at Urbana-Champaign

Larry Smarr
Director, Calit2
University of California, San Diego

Michael K. Smeltzer
Director of Networking
University of Illinois at Urbana-Champaign

Oren Sreebny
Senior Director, Emerging Technologies and Communications
University of Chicago

David Stack
Deputy CIO
University of Wisconsin–Milwaukee

Charles Thompson
Assistant Dean and Director, Engineering IT Shared Services
University of Illinois at Urbana-Champaign

Scott Valcourt
Director, Strategic Technology
University of New Hampshire

Shelton Waggener
Senior Vice President
Internet2

David H. Walker
Consultant

Steven Worona
Speaker/Consultant
Chestnut Hill Policy Center

Tom Zeller
Senior Technology Analyst
Indiana University Bloomington

Appendix B: NSF ACCI Task Forces Chairs and Members

The following lists the chairs and co-chairs for the NSF ACCI task forces mentioned in this report and identifies if and when members of the ACTI-CCI Working Group participated as members on the various task forces.

Task Force on Grand Challenges

- Chair: J. Tinsley Oden (University of Texas at Austin)
- Co-chairs: Omar Ghattas (University of Texas at Austin) and John Leslie King (University of Michigan)
- ACTI-CCI members: None

Task Force on Campus Bridging

- Chair: Craig Stewart (Indiana University)
- Co-chairs: Guy Almes (Texas A&M University) and Jim Bottum (Clemson University)
- ACTI-CCI members: Guy Almes (Texas A&M University), Jim Bottum (Clemson University), Clifford Lynch (CNI), Michael Mundrane (University of California, Berkeley), Jim Pepin (Clemson University)

Task Force on Cyberlearning and Workforce Development

- Co-chairs: Geoffrey C. Fox (Indiana University) and Alex Ramirez (Hispanic Association of Colleges and Universities)
- ACTI-CCI member: Michael Mundrane (University of California, Berkeley)

Task Force on Data and Visualization

- Chair: Not identified
- ACTI-CCI member: Guy Almes (Texas A & M University)

Task Force on High Performance Computing

- Chair: Thomas Zacharia (University of Tennessee/Oak Ridge National Laboratory)
- Co-chair: Jim Kinter (Center for Ocean-Land-Atmosphere Studies)
- ACTI-CCI members: None

Task Force on Software for Science and Engineering

- Co-chairs: David Keyes (KAUST and Columbia) and Valerie Taylor (Texas A&M University)
- ACTI-CCI members: None