

NSF-Sponsored Workshop Report:
**Sustainable Funding and Business Models for
Academic Cyberinfrastructure Facilities**

November 2010

*Final report for the National Science Foundation-sponsored workshop
held May 3-5, 2010 at Cornell University*

Table of Contents

Preface 3

Workshop Organizing Committee..... 4

Executive Summary 5

1.0 Introduction 7

2.0 Workshop Objectives and Participation 9

3.0 Organizational Models 10

4.0 Regional Organizational Models 11

5.0 Requirements for Resources and Services 12

6.0 Funding Models and Budget Sustainability 15

 6.1 Understanding Costs and Potential for Recovery 17

 6.2 Additional Motivating Factors 18

 6.3 Common Strategies, Models, Opportunities and Challenges 19

 6.3.1 Centralization of CI Resources and Services 19

 6.3.2 University Funding Models 21

 6.3.3 External Funding Models 21

 6.3.4 Cost Recovery Models 22

7.0 Staffing and Succession Planning 24

8.0 Industry and Vendor Relations 26

9.0 Metrics of Success and Return on Investment (ROI) 27

 9.1 Quantitative Metrics of Success 27

 9.2 Qualitative Metrics of Success 28

 9.3 New Challenges 28

10.0 Conclusions 30

Citations 31

Appendix A: Workshop Announcement 34

Appendix B: Workshop Agenda 36

Appendix C: Terminology 40

Appendix D: Workshop Participants (On-Site Participation) 41

Appendix E: Workshop Participants (Web-Based Participation) 48

Appendix F: Workshop Presentations and Breakout Findings 51

Appendix G: Workshop Position Papers 52

Appendix H: Related Presentations, Papers, Web Sites 56

Acknowledgements 57

Preface

This report summarizes the observations and recommendations from the National Science Foundation-sponsored workshop, “Sustainable Funding and Business Models for High Performance Computing Centers,” held May 3-5, 2010 at Cornell University, with additional support from Dell and Intel. Workshop participants, attending both in person and virtually via WebEx, were asked to submit position papers discussing the challenges that they face in funding and managing academic research computing facilities. The organizing committee accepted 28 position papers, which are available online at the workshop website: <http://www.cac.cornell.edu/SRCC>. 87 senior HPC and cyberinfrastructure (CI) experts from across the nation, as well as representatives from industry and Dr. Jennifer Schopf from the NSF, attended the workshop; 32 additional professionals participated via WebEx.

The workshop served as an open forum for identifying and understanding the wide variety of models used by directors to organize, fund, and manage academic cyberinfrastructure facilities. An ancillary but equally important outcome of the workshop was the degree of transparency and collegiality displayed by the participants while discussing the benefits and challenges of the models that they ascribe to or aspire to. By openly sharing their personal experiences and knowledge, insights were gained which through this report should provide value not only to institutions facing the challenges of establishing new CI facilities, but to more established facilities who are increasingly called on to justify the significant expenses of CI staff and infrastructure and the resulting return on investment.

Workshop Organizing Committee

Stanley C. Ahalt, Ph.D.
Director, Renaissance Computing Institute
ahalt@renci.org

Amy Apon, Ph.D.
Director, Arkansas High Performance Computing Center, University of Arkansas
aapon@uark.edu

David Lifka, Ph.D.
Director, Cornell University Center for Advanced Computing and Director Research Computing, Weill
Cornell Medical College
lifka@cac.cornell.edu

Henry Neeman, Ph.D.
Director, OU Supercomputing Center for Education and Research, University of Oklahoma
hneeman@ou.edu

Executive Summary

On May 3-5, 2010 the National Science Foundation (NSF) sponsored a workshop entitled “Sustainable Funding and Business Models for High Performance Computing (HPC) Centers” at Cornell University. A distinguished group of scientists, engineers, and technologists representing cyberinfrastructure (CI) facilities of all sizes and scope gathered to discuss models for providing and sustaining HPC resources and services. Attendees included directors and CIOs from national centers; departmental, college-level and central IT; and, research groups, as well as vice provosts and directors from research administration.

Those assembled for this workshop were acutely aware of the critical role that CI facilities play in sustaining and accelerating progress in numerous research disciplines, thereby promoting the discovery of new fundamental knowledge while simultaneously spurring practical innovations. The disciplines that are profoundly impacted include those that require sophisticated modeling, simulations, or analytic processes in order to understand and manipulate complex physical or sociological models and data that are otherwise incomprehensible. Examples include weather and climate modeling, molecular design for advanced materials and pharmaceuticals, financial modeling, structural analysis, cryptography, and the spread of disease. Many of these disciplines are now confronting, and benefiting from, new sources of observational data, exacerbating the need for center-level economies of scale for computation, storage, analysis and visualization.

This report summarizes the observations and findings of the workshop. Workshop participants were encouraged, prior to the workshop, to submit position papers discussing the challenges that they face in funding and managing academic research computing facilities. 28 position papers were accepted and may be accessed at the Sustainable Research Computing Centers wiki at <http://www.cac.cornell.edu/SRCC>.

At the national level, the NSF and the Department of Energy support formidable national HPC centers that provide a moderate number of national users with world-class computing. By contrast, a substantial number of scientific and engineering researchers depend upon departmental, campus, or regional/state research computing resources to fulfill their fundamental science and engineering computational requirements and to educate the students that are critically needed if we are to “weather the storm” and compete for quality jobs in the evolving global economy [1][2]. In some cases, local resources are also used by researchers to transition their research to the better-equipped and/or large-scale national facilities.

While workshop participants represented a broad spectrum of cyberinfrastructure facilities, ranging from the largest national centers to very small facilities just being formed, the primary focus of the workshop was on small to medium-sized CI facilities. The recent economic downturn has presented significant funding and organizational challenges to these facilities, calling into question their long term sustainability.

The papers and the subsequent workshop discussions identified and documented a variety of models used to organize, fund, and manage academic HPC and cyberinfrastructure facilities. One tangible outcome of the workshop was the collective realization of the profound challenges faced by many facilities, as well as the significant benefits that can be derived by different models of CI facility governance and operation. *Consequently, this report is not only informative for those creating new CI facilities for research, but also provides key insights into the efficacy of extant facilities, and supplies justifications for long-established facilities.*

The body of the report addresses a range of issues at some length, including:

- Organizational models and staffing
- Funding models
- Industry and vendor relationships
- Succession planning
- Metrics of success and return on investment.

Each of these topics is discussed from the significantly varying perspective of the many workshop participants, and the report thus captures a breadth of opinions that have not, heretofore, been assembled in a single report. The participants did reach a consensus on the importance of clearly stating, and endorsing, the fundamental precepts of the CI community, which are:

- Computational science is the third pillar of science, complementing experimental and theoretical science.
- Support for advanced research computing is essential, and CI resources need to be ubiquitous and sustained.
- Computational resources enable researchers to stay at the forefront of their disciplines.
- The amount of data that is being acquired and generated is increasing dramatically, and resources must be provided to manage and exploit this “data tsunami.”
- Disciplines that require computational resources are increasing rapidly, while, simultaneously, computationally-based research is becoming increasingly interdisciplinary and collaborative.

The conclusions and recommendations from the workshop are:

- **Broadening the CI Base** – The health and growth of computational science is critical to our nation’s competitiveness. While there is understandably a significant amount of attention and energy focused at the top of the Branscomb Pyramid [3], *the base or foundation of the computational pyramid must continue to develop and expand in order to both underlie and accelerate our scientific progress and to produce the next generation of researchers and a US workforce equipped to effectively bring these innovations to bear on our global competitiveness.*
- **Toward Sustainability** – Because computational science and CI are essential infrastructure components of any academic institution that has research as a fundamental part of its mission, *sustained support for computational science is essential and should involve a partnership of national funding agencies, institutions of higher education, and industry.* Notably, *the model of support that is appropriate for each specific institution requires strategic vision and leadership with substantial input from a diversity of administrators, faculty and researchers.*
- **Continued Collaboration** – Organizations such as the Coalition for Academic Scientific Computation (CASC), Southeastern Universities Research Association (SURAgriid), and the Great Plains Network (GPN) provide the community with an opportunity to share best practices, to disseminate results, and to collectively support continued investments in computational science at all levels of US academic institutions. *By working together, the HPC and CI communities best serve the mutually reinforcing goals of (1) sustaining the entire computational pyramid while (2) generating economic growth via breakthroughs in science and engineering.*

Policy and funding decisions that dis-incent collective community behavior, and that thereby impede shared improvement, are harmful, and should be avoided.

1.0 Introduction

High Performance Computing (HPC) continues to become an increasingly critical resource for an expanding spectrum of research disciplines. Both the National Science Foundation (NSF) and the Department of Energy (DOE) have created and support a powerful set of national cyberinfrastructure facilities that provide select national users with access to state-of-the-art computing capabilities. These facilities include both the NSF Track 1 and Track 2 facilities that are either already online or will be coming online soon, as well as the DOE HPC centers, including the DOE Leadership Class Facilities. The petascale Computational Science and Engineering applications that run at these facilities model a class of phenomena that are difficult or impossible to measure by any other means. The availability of tier-1 facilities such as these enable scientists and engineers to accelerate time to discovery, create new knowledge, and spur innovation.

National resources provide formidable computing capabilities to key researchers that work on extraordinarily complex problems. Yet, the consensus among participants in this NSF Workshop is that the vast majority of scientific and engineering researchers continue to rely on departmental, campus, or regional/state research computing resources. A recent Campus Bridging survey, which will be appearing in report form soon, supports this hypothesis, and we believe this can be shown to be true if appropriate surveys of the entire HPC ecosystem are conducted. Departmental, campus and regional resources are used to fulfill fundamental science and engineering computational requirements, and to educate the students that are critically needed if we are to “weather the storm” from both a competitive and a national security perspective. More local resources are also used by some researchers to prepare their software for eventual migration to the national facilities.

To satisfy these requirements, many universities have been focusing on identifying economies of scale, creating second and third tier CI facilities that provide HPC resources to their research communities in the most cost-effective and sustainable ways possible. However, the recent economic downturn is creating challenges in sustaining these facilities. Second and third tier facilities are faced with major challenges in funding, organizational structure, and long-term sustainability. Though we recognize that the first and second tier facilities funded by the NSF and those serving academic partners through the DOE may face budget pressures, the focus of this workshop is on unit, institutional and regional CI facilities and the budget challenges they may face in the coming years as the NSF transitions from the TeraGrid to a new model of funding, creating even more competition for funding. The identification of suitable sustainability models for cyberinfrastructure facilities is more important than ever. Resource sharing among tier-2 and tier-3 CI facilities, for example, is one approach to satisfying generic computing problems that do not require the highest level computing systems and can help bring the power of cyberinfrastructure to broader communities [4]. We believe that the survival and expansion of second and third tier CI facilities is crucial to national efforts to advance science and engineering discovery and is essential if we are to increase the number of US students with computational and data analysis skills.

Academic institutions take a wide variety of approaches to research computing. Some universities and university systems consider research computing a strategic investment and have attempted to provide sustained support for significant research computers, including sizeable parallel clusters, which are typically housed in formally recognized centers. Other universities view research computing as a tactical need, and may provide only intermittent funding for research computing for smaller, informal facilities. In either case, these research computing facilities are struggling to understand how best to organize, manage, fund, and utilize their hardware and staff.

Industry standard computing solutions provide a low cost entry into HPC hardware, but there are significant hidden costs, including:

- Building renovations, including space, power and cooling
- Administrative staff to install, maintain and support computational resources and research users
- Infrastructure requirements such as disk storage, backup, networks, and visualization
- Consulting staff who are specialists in complex domains such as weather and climate modeling, molecular design for advanced materials and pharmaceuticals, financial modeling, structural analysis, cryptography, and the spread of disease
- Consulting staff adept in supporting the scaling and optimization of research codes and the training of students and post-docs, as well as assisting researchers in identifying and leveraging national and regional resources and funding opportunities.

Our national research computing ecosystem must be sustained and expanded, lest our ability to compete at every level, including the most elite levels, be compromised. This workshop offered a unique opportunity to begin a dialogue with colleagues in leadership positions at academic institutions across the nation on CI facility requirements, challenges, experiences and solutions. This report summarizes the findings and recommendations of this workshop, both to raise awareness and to encourage continued open and collaborative discussions and solutions. It is the result of a productive workshop which led to a shared understanding of organizational, policy, funding, and management models that result in *sustainable cyberinfrastructure facilities*. An ancillary, but equally important outcome, is the degree of transparency across the extant facilities, which will provide evidentiary justification for cyberinfrastructure facilities that are struggling to become established and are increasingly called on to justify the significant expenses, and the resulting return on investment (ROI), that naturally occur as facilities become established.

2.0 Workshop Objectives and Participation

The objective of this workshop was to provide a forum for an open discussion among Center Directors, campus Chief Information Officers, and Research Officers on the topic of Sustainable Funding and Business Models for Academic Cyberinfrastructure (CI) Facilities. Eighty-seven academic HPC and cyberinfrastructure experts from across the country, as well as representatives from industry along with Dr. Jennifer Schopf from the National Science Foundation (NSF), participated in the workshop held May 3-5, 2010 at Cornell University. An additional thirty-two participants accessed the workshop presentations and findings of the breakout sessions via WebEx (www.webex.com). Appendix D and Appendix E list the workshop participants, both on-site and Web-based.

All participants were strongly encouraged to submit position papers covering any or all of the proposed workshop discussion topics, including:

- Organizational models and staffing
- Funding models
- Industry and vendor relationships
- Succession planning
- Metrics of success and return on investment.

Appendix G provides links to 28 accepted workshop position papers. Appendix H provides links to other useful papers and publications.

Invited workshop presentations and breakout sessions were designed to stimulate participation and allow those in attendance to focus on and provide detailed input and feedback on all topics. Appendix F provides links to the workshop presentations and summary slides from the breakout sessions.

3.0 Organizational Models

In order to establish a foundation for comparing institutional models for research computing and cyberinfrastructure support, an obvious place to begin was to develop an understanding of the various reporting structures, institutional leadership advisory boards, and interactions with key users of the facilities. Virtually all workshop participants represented institutions with one of the following four organizational structures:

1. A director reporting to the Chief Information Officer (CIO) of the university, as part of the overall campus IT mission
2. A director reporting to the Vice Provost/President/Chancellor for Research (VP/CR), as part of the overall campus research mission
3. A director reporting to the Provost/Chancellor as part of the overall campus infrastructure mission
4. A director reporting to one or more Deans of heavily invested colleges, often in conjunction with the CIO or VP/CR, as part of a focused research mission for specific college(s).

Cyberinfrastructure facility directors, whether they are faculty or staff, must be skillful leaders. CI facility directors may be either tenured/tenure-track faculty members or non-tenure-track research staff. Directors who are faculty are often engaged in personal research that requires computational resources and services and, therefore, are well suited to justify the importance of these services to the administration of an institution. Directors who are non-faculty research staff typically understand the service mission of a CI facility and, therefore, are well suited to make service their primary focus since they do not have the same teaching and research pressures as tenured/tenure-track faculty (albeit there are other pressures surely).

Faculty Advisory Committees and other types of oversight boards can be useful to CI directors. Faculty Advisory Committees typically perform the following functions:

1. Recommend strategic directions
2. Identify new requirements
3. Promote the requirements of researchers
4. Provide input on allocation and policy decisions.

Other types of oversight boards, which often also include faculty members, may include members from industry as well as colleagues from outside institutions. Oversight boards typically provide advice on one or more of the following areas:

1. User issues
2. Administration
3. Funding
4. Technical direction
5. Strategic opportunities and partnerships.

4.0 Regional Organizational Models

As the need for computational and data analysis capabilities grows and expands to new fields, funding facilities (space, power and cooling) and hiring staff with the appropriate skills and experience to run cyberinfrastructure resources is becoming more and more challenging. In order to address these challenges, some institutions are choosing to form regional partnerships as a means of cost and expertise sharing.

Shared data centers support growing research needs from participating members with the flexibility for expansion through phased deployments. Establishing regional data centers also provides the opportunity to leverage green technologies for power and cooling. The University of California institutions [5], the New Jersey Institute for Modeling and Visualization [6], and the Massachusetts Green HPC Academic Research Computing Facility (GHPCC) [7] are three such recent state-supported efforts.

Other groups are forming regional models that leverage grid technology to share resources and expertise. Like the shared data center model, this model provides research capabilities for institutions who independently could only afford to offer resources and services at a much smaller scale. SURAgriid [8] and the Great Plains Network [9] are good examples of such regional collaborations.

SURAgriid's Strategic Plan recognizes the value of regional engagement and collaboration: "The overall intent of the SURAgriid strategic plan is to provide researchers with a cyberinfrastructure strategy that extends local resources through the shared resources of SURAgriid and in doing so provides a gateway to national (and international) infrastructures, and establishes SURAgriid as an integral component of each SURAgriid member's infrastructure solution for competing in the 21st century. This implies a collaborative effort of the SURAgriid community to articulate a core set of standards, conventions and policies that supports the integration of our member's campus CI resources into a regional whole, under the banner of a regional Virtual Organization [10]." The Coalition for Academic Scientific Computation (CASC) also includes the importance of community engagement in its mission statement, which emphasizes the facilitation of "information exchange within the academic scientific computation and communication community [11]."

Regional facilities can be a catalyst for economic, educational and workforce development, as well as an effective way for individuals and organizations to focus a strategically targeted fraction of their effort on a larger community-shared set of CI resources and services. By doing so, they are also providing a potential path for researchers at their institutions to scale research from campus to regional or national resources.

5.0 Requirements for Resources and Services

The research computing infrastructure, HPC systems, and cyberinfrastructure resource requirements of every institution's researchers, students and faculty are unique. The first step in providing support for academic research computing is to understand user requirements, i.e., what are the services and resources that users will use and/or that the institution sees as strategic and, therefore, necessary to provide. Developing services that meet user needs requires that an inventory of existing resources and services currently available at the institution, even if provided by other organizations at the institution, be performed, as well as a detailed cost analysis of providing these services.

Crucial to a successful analysis is a full accounting of costs, including the hidden costs, involved with each service. For example, providing support for an HPC system requires not only space, power, cooling, networking and computing equipment, but also staff support for running the system, and for helping researchers use the system. Staff members require office space, phones, personal computers, printers, training, travel, and benefits. All costs must be taken into consideration in order to reveal the total cost of operating an organization's CI resources. Knowing all costs is extremely beneficial in developing a sustainable funding model. Once user requirements, the costs of the services required, and the amount and sources of funding are identified and understood, negotiations with an institution's administration for resources and support can begin in earnest. Despite the desire on the part of some users to maintain an existing resource or service or to establish a new one, if adequate institutional support and/or external funding are not available or if users are unwilling or unable to pay for it, such a resource or service will be difficult or impossible to sustain locally.

To be successful in negotiating for CI facility resources and support, a mission statement that resonates with campus faculty and researchers, and that is clearly aligned with the goals of the institution, is essential. In today's challenging economic environment, administrative management will carefully weigh each CI investment based on cost, breadth of impact, strategic potential and alignment with the institution's mission and goals. Providing data and identifying leading faculty and researchers who will support that data will help this process.

It is noted that in the same way that NSF-funded national centers used a shared CI model in order to provide high-end (terascale and petascale) supercomputing resources, a community of local HPC facilities can certainly look to collaborative and shared strategies to attain economies of scale required to sustain their services – whether those services are computation, storage, consulting (or others as noted below). Indeed, concepts of virtual organizations and cloud computing are very informative as to how a future HPC ecosystem might sustain individual resources and services.

During the workshop, a broad range of activities, resources and services that CI facilities provide were discussed. Note: every institution offers a unique configuration of some or all of these activities.

- **Consulting** – Providing professional technical support for the effective use of local, regional and national cyberinfrastructure resources. This activity could include (a) ensuring that researchers can access resources from their own workstations, (b) facilitating multidisciplinary research, (c) supporting data analysis, possibly including statistical analysis, (d) providing scientific application expertise, and (e) supporting existing and emerging technologies.

- **Computing** – Providing computational resources locally, regionally, or nationally, depending upon mission and funding, as a production service. Providing this service with the right level of technology capabilities requires understanding the users whom a CI facility will be serving. Some hardware options can be expensive and may not be necessary for most or all targeted research. Part of providing computing resources is keeping resources reasonably current and identifying opportunities to deploy new types of resources. North Carolina State University’s Virtual Computer Lab (VCL) cloud computing environment is a good example of adapting a service offering to meet the changing needs of users [12].
- **Data Storage** – Providing data storage and backup services for local, regional, and national users, based on mission and funding. Data storage is a multifaceted service, and the exploding volume of data being produced by scientific instruments, distributed data sensors and computer simulations make this a growing challenge. Data storage involves providing high performance or parallel file systems for large-scale computational simulations, reliable storage for accumulated research data, and backup solutions for data that is difficult or cost prohibitive to recreate. The NSF, National Institutes of Health, and other federal funding agencies have announced plans to soon begin requiring data management plans as part of new proposals [13]. This will require all institutions to revisit their data storage strategies and implementations, as it will impact how datasets are created, formats that are used, metadata solutions, methods for tracking provenance, and, in some cases, long-term curation.
- **Networking** – Providing various levels of network connectivity at the local, regional, and national scale based on mission and funding. Networking is essential for accessing local and remote cyberinfrastructure resources, moving data to and from CI resources, and performing visualizations and analyses of data at remote resources where the volume of data makes transfer to local resources impractical.
- **Visualization** – An important component of data analysis is visualization. As the volume of data produced in research continues to grow at rapid rates, using visualization to analyze that data continues to grow in importance. Visualization resources range from workstation tools to dedicated visualization clusters with graphic accelerators to specialized installations that support three-dimensional immersive graphics at extremely high resolutions.
- **Education and Training** – Providing various levels of education and training based on mission and funding. This is an extremely important part of any academic institution’s research mission and essential in developing a workforce equipped to compete globally. Training involves helping researchers and students learn computational and data analysis skills, including programming, parallel programming, use of computational resources (local and remote), numerical analysis, algorithm development, debugging, and performance optimization. Training is typically offered as workshops (hours to days in duration) or as academic courses (half or full semester in duration) that provide an in-depth understanding of complex topics.
- **Software Development** – The development of software tools, libraries, and techniques to improve the usability of local, regional, and national cyberinfrastructure resources is based on mission and funding. This typically involves research and development efforts that focus on the latest, often leading-edge, CI resources in order to ensure optimal utilization by researchers. Depending on the scope and mission of a CI facility, in-house software development can range from a mission critical service to an unaffordable luxury.

- **Virtual Organizations** – As the pervasiveness of regional and national cyberinfrastructure resources increases, the need for appropriate infrastructure and tools to facilitate collaboration is becoming more important. Economies of scale may be instituted by providing efficient and reliable services around the systems, networking, data storage, and software tools required by virtual organizations.
- **Outreach Activities** – This set of activities focuses on reaching and supporting new users of cyberinfrastructure resources and broadening impact at the local, regional, and national scale, based on the CI facility’s mission and funding. At the national and regional levels, this includes activities such as the TeraGrid Campus Champions program, the Open Science Grid, the Great Plains Network or SURAgriid. At the institutional and regional level, this involves activities such as introductory or “getting started” workshops, open houses, presentations at neighboring or collaborating institutions, or support for getting new researcher projects underway. Feedback from this particular workshop, for instance, indicated that it was highly valued as an outreach forum for HPC and CI leaders.
- **Economic Development** – This set of activities is focused on the sharing of information, technologies, and services through corporate partnership agreements at the local, regional, and national scale. Companies seek to gain a competitive advantage through these agreements and depend upon colleges and universities to develop an intelligent workforce with the drive and skills to compete.

6.0 Funding Models and Budget Sustainability

Workshop participants identified three commonalities in successful cyberinfrastructure facilities: (1) an organizational model and reporting structure that is compatible with its institution's mission, (2) a portfolio of resources and services based on current and emerging requirements of its research community, and (3) a funding model that is commensurate with the scope of its mission, whether local, regional, or national. Developing a sustainable funding model that enables CI facilities to retain a skilled and proficient technical staff while providing a current computational infrastructure was a common goal of all workshop participants and, as such, was one of the most popular topics of the workshop. This concern was further heightened by organizational budget pressures resulting from the recent downturn in national and state economies.

Participants expressed the importance of frequency and clarity in conveying the fundamental assumptions of the CI community to institutional administrators. These tenets include:

- Computational science is the third pillar of science, complementing experimental and theoretical science. This has been widely cited in the scientific literature and acknowledged in Congressional testimony and Federal reports [14].
- Support for advanced research computing is no longer a luxury; it is a fundamental tool and, as such, computational resources need to be as ubiquitous as networks, phones, and utilities.
- Access to computational resources is a necessity for many researchers in order to stay at the forefront of their disciplines. Further, the amount of data researchers are acquiring, generating and needing to analyze is growing rapidly. Providing resources to store this data, along with the hardware, software, and experienced staff to assist in data mining, analysis and visualization, is essential. As more and more knowledge is created in or converted to digital form every day, data will be used not only to enhance research discovery, but as an important part of the education delivered by classroom instructors or through discipline-specific Science Gateways.
- The number of disciplines that require computational resources is increasing rapidly. More and more researchers in the social sciences, economics, and the humanities are embracing cyberinfrastructure resources and services as required tools for analysis and discovery. CI use by science and engineering fields, such as astronomy, bioengineering, chemistry, environmental engineering, and nanoscience, is also growing, driving the need for rapid access to CI facilities with varying levels of scalability in order to answer questions that until now were intractable.
- Contemporary computationally-based research is becoming increasingly interdisciplinary and collaborative in nature. Professional staffs adept at developing CI software, tools, technologies and techniques are necessary in order to bridge the gap between disciplines and to turn what has been described as “mountains of data” into knowledge.

A variety of funding and budget models were shared during the workshop. It was evident that no single solution works for everyone, and that every model will require modification over time. Dr. Eric Sills, Director for Advanced Computing at North Carolina State University, captured this concept in his position paper:

"Sustainability evokes the feeling of perpetual motion - start it and it sustains itself - but sustainability actually requires nearly continuous ongoing work, adaptation, and adjustment."

Essential for success are a solid understanding of an organization's computational and data analysis requirements, a clear mission statement that addresses these requirements, and an institutional commitment to develop, maintain and support a sustainable funding model. Flexibility and adaptability are required in order to anticipate and react to ever-changing research requirements and technologies.

Most sustainable funding models include the following qualities:

- **A Strong Value Proposition** – The resources and services required by researchers should be provided in a highly efficient and effective manner. Research requirements need to be carefully analyzed in order to define the services that will most likely enable and accelerate research success. The number of researchers, students, and/or departments that require specific resources and services should be quantified; this is an important step in securing institutional commitment and support.
- **Transparency** – Sharing the cost basis for specific resources and services is essential in order to gain understanding and trust. CI directors need to demonstrate that facility costs are similar to, or better than, what researcher costs would be if they performed the work themselves and/or with graduate student labor, factoring in that professional CI services should provide better quality, availability, utility, and economies of scale for the institution as a whole. In virtually all cases the CI facility will receive some level of direct funding support from the institution. In the case that users are charged directly for the use of resources and services, institutional funding support can be applied to chargeback rates in order to keep costs down for the end user. A well-informed research community, knowledgeable of the true costs of CI, will be better able and willing to support their infrastructure providers in articulating and justifying the need for CI funding.
- **Fairness** – Ensure that generally available resources and services are available in an equitable manner to all intended users of the facility, i.e., at the same access level and the same cost. This is essential in order to serve a broad, rather than narrow, user community. A broad and loyal user community will reduce risk for the CI facility and can increase partnership, joint proposal, and service opportunities.
- **Economies of Scale** – By identifying resources and services that are in wide demand, economies of scale solutions may be implemented that reduce overall institutional costs. This may increase the value proposition of the CI facility by reducing institutional redundancies and maximizing resource usage. Beyond the local institution, economies of scale are accessible through virtual organizations or collaborative partnerships. Intra-organization economies of scale can provide CI resource value to the local organization while, concurrently, contributing resource value to external entities.
- **Base Funding** – Organizations interested in establishing a cost recovery model need to define the mission of the CI facility, what resources and services it will provide, as well as their associated costs. Next, which costs can and should be recovered from users, versus those costs that are institutionally accepted as required core infrastructure, need to be clarified. The appropriate level of base funding provided by the institution to the CI facility may then be rationally established.

6.1 Understanding Costs and Potential for Recovery

There are four major costs involved in operating an academic cyberinfrastructure facility. Covering these costs is the primary objective of any sustainable budget model.

- **Staff** – While workshop participants agreed that staff are the essential resource that make a CI facility a unique and valued resource, many participants stated that they are short-staffed due to insufficient staff funding. The need for multi-skilled staff, in particular, is critical. Staff require not only advanced computational skills, but, in many cases, expertise in one or more scientific domains. Funding, recruiting, training, and retaining staff with the requisite experience and expertise is difficult. CI facility staffing requirements are extensive. Relatively high salaries coupled with required overhead expenses, i.e., vacation, training, etc. make cost recovery for CI staff time difficult.
- **Facilities** – The amount of data center space, power, cooling, and office space required to provide professionally operated, maintained and supported CI resources and services is substantial. The power, heat and space density of current computational and data storage resources continues to increase. Facilities that can handle this kind of density are expensive to build and, even with proper design and planning, will be out-of-date and will require significant updates every 10 to 15 years. Depending upon the institution, these facilities may, or may not, be covered partially or completely by indirect funding.
- **Hardware Resources** – As scientific problems addressed by researchers scale upward in terms of complexity, so too do the computational resource requirements, in terms of number of processors and cores, high-speed interconnects, memory, disk storage, network connectivity and bandwidth, and visualization. The challenge is not only the one-time cost of acquiring hardware resources, but also the recurring cost of maintaining them over their service lifetime and, ultimately, replacing them with new technologies at appropriate intervals, based on performance and utility needs and relative consumption of space, power and cooling.
- **Software Resources** – Software and tools are necessary for the operation of computational and data resources (e.g., scheduling software, deployment, patching and monitoring tools, support for parallel file systems, etc.) and for the researcher’s effective use of the resources (e.g., mathematical libraries, parallel programming libraries, specialized applications, compilers, debuggers, performance tuning tools, etc.). There are costs and trade-offs associated with commercial, open source, public domain, and custom software. As several workshop participants commented, when it comes to software, “there’s no free lunch.” Commercial software has licensing and maintenance costs, much like hardware. Open source, public domain software and custom development of research applications require an investment in staff time for the development and maintenance of the software. The true cost of staff support, development effort, and maintenance for all types of software is not negligible and must be carefully considered in light of the institution’s overall mission and budget.

6.2 Additional Motivating Factors

There are additional motivating factors for creating a sustainable budget model for local cyberinfrastructure facilities that are strategic in nature. Determining the right level of funding to support these efforts is crucial and requires a clear understanding of the needs of local researchers, both now and in the future.

- **Supporting Local Research** – The advanced computing skills of faculty and research staff typically fall into one of three categories: (1) users with little to no experience that, therefore, do not necessarily know how advanced computing can help them accelerate their research, (2) experienced users who require relatively straightforward resources, such as high-throughput clusters and small-to-large scale HPC clusters, (3) specialized users who can take advantage of national extreme-scale resources. Most faculty and researchers fall into categories 1 and 2. As trends toward increasingly sophisticated simulation tools, global collaborations, and access to rapidly growing data sets/collections are required by more and more disciplines, local resources will need to grow in capability. Local researchers and their students represent a clear opportunity to broaden the impact and expand the use of CI resources in order to enhance our nation's competitiveness. These researchers, however, need access to more local and regional CI facility resources and staff in order to expand their CI use and, most importantly, engage more undergraduate and graduate students in parallel computing and data-intensive scientific discovery.
- **Supporting Training** – The amount of effort required by faculty and students to learn the necessary computational skills to effectively use advanced CI facilities is substantial. As more and more institutions make a commitment to support computational science and interdisciplinary research in the form of "Computational Science" degrees and/or certifications, such as those funded by NSF CISE Pathways to Revitalized Undergraduate Computing Education (CPATH) awards, the ability to access computational resources for training purposes should increase.
- **Gateways to National Cyberinfrastructure** – It is becoming increasingly important for local institutional CI facilities to be well connected to, if not seamlessly integrated with, regional and national resources. Local institutional resources and services are in a position to provide an "on ramp" to large-scale national resources for researchers who require access to more capacity or capability than can be reasonably provided at the campus level. Local researchers who require access to national resources also may need an appropriate level of local staff support and infrastructure, such as software and tools, to make timely and effective use of national resources.
- **Utility Support** – As more and more disciplines require computational and data analysis resources, supporting new researchers who have little or no intrinsic interest in the inner workings and complexities of the resources will require local staff support and the availability of user-friendly interfaces that enable users to access resources as a seamless and ubiquitous utility. Workshop participants noted that many researchers who need dedicated access to their own private resources see little value in managing these resources since their core focus is on research rather than computational support. Having local staff and CI facilities where these resources can be installed, managed, and maintained in a professional manner, with optimal utility, is becoming more and more important.

- **Economies of Scale and Scope** – New energy efficient computer systems and virtualization technology, along with enterprise-class storage solutions, are enabling new economies of scale that make centralized resources and services increasingly attractive in comparison to highly distributed cyberinfrastructure resources spread throughout an academic institution. Furthermore, the existence of a centralized highly-skilled staff to support the use of these resources is far more efficient than trying to do so at multiple departmental or individual college levels.
- **Federal Cyberinfrastructure Funding Opportunities** – Academic institutions that provide their CI facilities with sustained funding for a limited but consistent number of core staff and/or resources may be in a better position to leverage those resources and expertise in order to effectively compete for federal research grants. By providing sustained funding for some level of core cyberinfrastructure, institutions are more likely to develop a CI staff that is highly proficient in operating and supporting the resources required for campus research and more likely to identify regional opportunities for collaboration.

6.3 Common Strategies, Models, Opportunities and Challenges

Several common strategies for sustainable budget and funding models were discussed; some are well established, others are just emerging. While it is clear that no one solution fits all, there are lessons to be learned from each that may be applied à la carte to the development of a successful model.

Many institutions use a blended funding model that combines funds from internal university sources, external state, federal, and industry sources, and cost recovery (chargebacks) for services provided. Internal university funding is usually required to start and sustain a CI facility, and to attract subsequent investment by faculty researchers. Funding for CI staff may eventually be offset by research grants. Many institutions cover physical CI facility infrastructure support, such as machine room power, cooling and floor space, with indirect funds. New equipment purchases may be covered by internal university funds, external research grants, or some combination of internal and external funds (as long as they are properly accounted for).

6.3.1 Centralization of CI Resources and Services

Centralization was a common strategy that many of the workshop participants were working towards in hopes of saving money by providing operational efficiencies and economies of scale and scope. Of course, while centralization can benefit one's local institution, "centralization" models may also span institutions, and accrue similar benefits.

- **Benefits of Centralized CI**

Facilities – Increased efficiency in the use of space, power, cooling, and more focused and consolidated long-term planning. Lower operating costs by eliminating the need for less efficient distributed facilities to house computational infrastructure. A well-run centralized data center can improve advanced computing quality (security, stability, and continuity) by providing professional systems administration and maintenance. Sharing of facilities and resources is emerging as an important component of many green initiatives.

Staffing – A core staff supporting centralized resources enables an institution to attract and retain higher quality faculty and CI staff with deeper skills in critical areas such as parallel computing, scientific applications, visualization, and data storage/analysis.

Economies of Scale, Scope and Cost Sharing – A cost sharing model that allows faculty and researchers to contribute to a well-run, cost-efficient enterprise CI facility enables everyone to get more for their research dollar. Condominium clusters and enterprise storage solutions are two good examples. Several forms of condominium clusters were described by workshop participants, each providing a different level of buy-in with associated benefits [15][16]. In its simplest form, research groups contribute funds to a centrally-managed CI facility that purchases compute cores for a large cluster that is shared by researchers. The “condo” approach provides researchers with a much bigger and better run resource than they could purchase and maintain independently and generates valuable economies of scale for the institution in the areas of facilities, power consumption, and staffing. In addition, when researchers work together and pool funds, there is an opportunity for increased bargaining power with Original Equipment Manufacturers (OEMs) and Independent Software Vendors (ISVs). Interdisciplinary research opportunities may also be more likely to develop.

Enhanced Research Support – Professionally run CI resources enable faculty and researchers to focus on their research rather than on the management of their own computing infrastructure. The establishment and availability of some level of general-purpose computational resources allows faculty and researchers to explore the value of advanced computing for their research without requiring a large initial investment for their own infrastructure and the staff to run it.

- **Challenges of Centralized CI**

Costs – Funding a large-scale centralized data center can be difficult. These facilities typically have a high cost per square foot, and attracting sponsored funding is difficult for buildings that are not designed for teaching or education. There is a perception that libraries are a necessary core infrastructure investment while cyberinfrastructure facilities are an expense. This perception, coupled with the difficulty in grasping the degree of impact that the digitization of the vast majority of knowledge will have on research and education in the coming years, contributes to an underfunding of CI facilities by US colleges and universities, and their supporters.

Access Control – Providing researchers access to the resources in a CI facility involves special considerations, both for the physical access to facilities and for administrative access to the computational infrastructure. Access requirements are often in conflict with the basic principles of operating a secure and stable production resource.

Strategic Oversight and Policy Decisions – Appropriate faculty and researchers should be identified to solicit feedback on sensitive issues such as queuing policies, access priorities, and the specific types of heterogeneous computing resources required to effectively serve the institution’s research community. It is important to ensure that key stakeholders have a say in the operation of a CI facility and its resources.

6.3.2 University Funding Models

Some institutions fund cyberinfrastructure facilities, resources and services completely or at a very high level. These institutions view CI as a necessary and critical part of campus infrastructure much like administrative IT, the library, and networking. Some of these institutions fund CI entirely from core internal budget, or with indirect funds from research grants. Other institutions have formed a "Partner's Program" model where faculties leverage base university funding to expand a large central resource rather than buying their own. In this model, the institution typically provides some amount of base level funding support, and cost sharing by researchers is used to make up the difference. Inter-institutional funding models for multi-campus or state-wide systems are possible through consortium or collaborative agreements.

- **Benefits of University-Funded CI**

Efficiency – Base funding for CI reduces individual department costs by eliminating the need to build and support their own resources and optimizes institutional CI operations and maintenance.

Strategic Advantages – The goal of institutional funding is typically to provide a strategic advantage for its faculty, researchers and students. Providing access to cyberinfrastructure resources and services to those who may not yet have funding to explore new areas of research may yield innovation and breakthroughs otherwise not possible. In addition, undergraduate and graduate students at these institutions gain valuable experience in computational science, which is rapidly becoming integral to research in most disciplines, from the traditional sciences to the social sciences and the humanities.

- **Challenges of University-Funded CI**

Sustainability – How will institutions develop a business model that enables them to sustain the staff, computational resources and services on an ongoing basis, especially during economic downturns?

Motivation – If resources and services are free to faculty researchers, is there adequate motivation for faculty to compete for grants that support their computational requirements at some level?

6.3.3 External Funding Models

Institutions that receive much of their CI funding from external sources such as federal grants and industry are typically able to focus on very large and even extreme scale resources and services that are not financially feasible for institutions running under local funding models. NSF-funded centers such as the TeraGrid resource providers are good examples of these types of facilities. In order for the TeraGrid resource providers to successfully compete for external funding, sizable investments at the University and/or state level have been required.

- **Benefits of Externally Funded CI**

Efficiency – To provide extreme scale resources intended to support select world-class research and enhance competitiveness, there are efficiencies that can be leveraged at the federal level by

supporting a limited number of centers with skilled staffs and extreme scale resources.

Innovation – By pushing the limits of computational scale and performance, these centers produce innovations in software, tools, and affiliated technologies. This has a positive effect not only on research disciplines whose applications run on these resources, but also on the field of computer science and, more broadly, computational science.

National Competitiveness – Industrial outreach and collaboration are important metrics of success for nationally funded facilities. Technologies that are developed through the pursuit of extreme scale and performance find their way into capabilities that industry can use to develop new or better products and services.

- **Challenges of Externally Funded CI**

Funding – During economic downturns, federal and especially state support funding (e.g., legislative line items) is limited and therefore competition is much higher. In addition, federal funds to support institutional resources are increasingly more difficult to secure, as the NSF and other agencies appear to be focused primarily on the extreme scale.

Sustainability – How do institutions that rely heavily on externally-funded projects sustain their staff expertise beyond their center’s immediate funding horizon? At the extreme scale, most national scale centers operate with a constant push toward bigger, faster, and higher performance resources. How do national resources fund hardware refreshes at a proper, i.e., competitive, pace?

6.3.4 Cost Recovery Models

The ability and willingness of research teams to pay for centralized CI computational resources or staff consulting services are important factors to consider in deciding whether to move to a cost recovery model. Institutions with considerable external and/or internal funding per faculty member are typically vastly better positioned to implement cost recovery approaches for the uncovered costs than those organizations with lower levels of research funding and, consequently, higher recovery needs.

Researchers operate under different measures of productivity and reward structures, i.e., the number of publications produced, the number of students mentored and graduated, and the number and scientific impact of computationally-enabled discoveries. For modestly funded researchers, the value proposition of paying funds directly into a central CI facility may be difficult to justify with respect to their particular reward structure – possibly to the point where their incentives favor abandoning computationally-intensive research rather than paying service fees for it. On the other hand, if fees for centrally accessed CI computational or staff resources are low enough relative to the productivity gains enabled, some selective use of centralized services (or emerging technologies such as cloud computing) may make sense, even for a modestly funded research group.

Well-funded research teams may already be near or at the maximum practical number of members that the team leadership can reasonably mentor, and so productivity is less likely to be improved by increasing the size of the team than by providing current team members with additional resources, including (and in some cases especially) CI resources, in which case the value proposition of CI facility service fees can be vastly more justifiable.

If an institution decides to implement a cost recovery model, the costs for access to resources and services are covered in part or whole by a fee-for-service. These costs are best kept transparent, so that the value proposition of a professional, centralized service is readily apparent and thereby discourages faculty from constantly building their own one-off systems. The cost of using a centralized resource should not exceed the cost of faculty deploying their own resource in their department or lab. Hopefully, centralized resources are cost competitive with graduate student labor by providing superior service, though meeting this price point has implications with respect to institutional support and subsidies.

There are benefits and challenges in implementing a cost recovery model:

- **Benefits of a Cost Recovery Model**

Steady-State Funding – If faculty researchers are well served, and if they have sufficient research budgets to cover such costs, they will likely see value in and subsequently be willing to pay for resources and services. The more satisfied and well-funded researchers that a particular CI facility supports, the better the cost recovery model will work for that facility. Steady-state funding from the institution enables CI facilities to continually “right-size” their CI offerings based on demand. Using a cost recovery model also provides a transparent mechanism for an institution to monitor the impact of its financial support or subsidy [17].

Positive Incentives – Given a cost recovery model where resources are not provided for “free,” faculty and researchers may be more motivated to write proposals and win grants to cover the costs of computational resources and services. This may have a positive financial impact on the researchers, the cyberinfrastructure facility, and the institution’s overall research portfolio.

Economies of Scale, Scope and Cost Sharing – By contributing research funds toward well-run CI facility resources and professional services, the whole is greater than the sum of the parts. Researchers have access to staff and computational resources when they need them and more resources for peak computing times than they could fund on their own.

- **Challenges of a Cost Recovery Model**

Demand and Resistance – Cost recovery models assume researcher support and demand for CI facility resources and services, as well as an ability to pay. Getting started under a cost recovery model can be challenging, especially for institutions moving to a cost recovery model from one that was formerly heavily or completely subsidized by the university, i.e., where the resources were “free.” Overcoming this change takes full-time CI leadership and hard work in order to identify what researchers really want and what, if anything, they can afford and are willing to pay for. The CI facility must provide a strong value proposition to both the institution and the CI users.

Innovation – One concern is that a CI facility operated in a pure service mode will fall behind the technology curve and lose value to its researchers. If the facility is unable or unwilling to adapt over time, this is a legitimate concern. The counter argument is that a CI facility operating under a cost recovery model is more motivated than ever to ensure that it provides resources and services that researchers demand, lest it will lose value, become obsolete, and no longer be required.

7.0 Staffing and Succession Planning

The number and variety of staff at a CI facility depends on the type and level of services and resources supported as well as the number of computational researchers, both expert and novice, that are supported. Generally, each type of technology needs some level of expertise in the facility, whether it is a cluster resource, storage, file system, scheduler, specialized network, or security infrastructure. For small CI facilities, acquiring the variety of skills necessary to deliver all technologies can be challenging, and these facilities may want to choose to keep technology and vendor choices limited so that the support staff can manage the systems more effectively. One way to reduce the number of staff required and still maintain relatively complex HPC technologies at a CI facility is to buy commercial products with support where they are available. For example, commercial services and products for cluster installation, storage and file systems, and scheduling software are readily available. Essentially, the benefit of this approach – leveraging external expertise and capabilities – has much in common with emerging cloud computing models.

A CI facility that aims to provide highly available services with 24-hour uptime must have a large enough staff for someone to be on call at all times. If there are not enough funds to provide this level of support, then users must realize that a major failure or user issue that occurs late at night or on the weekend may not be serviced until the next working day. Of course, models of shared support – as in a consortium or virtual organization – can offset this staffing requirement by rolling support to another site based on an agreed upon schedule.

Staff can generally be categorized as “inward facing” – servicing the systems and resources of a CI facility, or “outward facing” – providing user support and analysis services. Some staff have the skills to meet both inward-facing and outward-facing tasks; this flexibility is especially valuable at small CI facilities.

The transition of a CI facility from one director to another can be disruptive, and may be a substantial setback. In many cases, CI facilities are driven by the personality of the director, and when this individual leaves, the vision and persistence of the CI facility may be threatened. Some methods that can help to alleviate the impact of the loss of a CI facility director are: (1) engage staff in the operational decisions of the facility prior to the director leaving, (2) ensure that university administration value the importance of the ongoing mission of the CI facility through regular reports, engagement, and communication, (3) ensure that a funding model is in place for continued operation of the facility, (4) ensure that the “hero” users of the institution will lobby the institution to sustain the operation of the facility, (5) create faculty experts in various aspects of CI technologies, facility operation, and in the authoring of proposals that support the resources of the facility, (6) if feasible, make recommendations for the succession director, and (7) actively participate in the greater cyberinfrastructure community and use best practices to better prepare for changes in your local organization.

Changes in senior university administration personnel were identified as another area of concern. Several facility directors stated that their CI initiatives were largely supported by one or two senior officials at their institution who viewed their activities as strategic. When these administrative positions turn over, there is no guarantee that new officials will have the same vision or appreciation for CI initiatives. Directors are challenged with educating their university administration broadly on the importance of CI to their institution and providing them regular updates with metrics of success that is in alignment with the

mission and priorities of the institution, in order to ensure that CI becomes an integral part of the fabric of the institution rather than the strategy of a single administrator. Community development of materials and publications that provide CI blueprints and demonstrate ROI, cost avoidance and cost savings are needed by the CI community. The National Science Foundation and organizations such as CASC, SURAGrid, and GPN can help in this regard.

8.0 Industry and Vendor Relations

Providing advanced computing services based on technologies that provide optimal performance and economies of scale increases the relevance of academic CI resource providers to industry and vendors. The rate at which new computing technologies are developed and existing technologies improve is accelerating each year. It is part of what makes computational science so exciting and so challenging at the same time. Researchers are always anxious to take advantage of technologies that will allow them to get better results faster, but they balance this desire against how much effort they or their research group must invest versus the relative payoff. This forces cyberinfrastructure service providers to constantly keep abreast of new technologies and to rapidly adopt only those with promise, because adoption often requires time-consuming testing and implementation. Having a sustainable recovery model promotes careful decision-making processes when it comes to evaluating new technologies and implementing “right-sized” solutions. This level of experience and expertise makes cyberinfrastructure providers attractive places for industry to partner with. Industry is excited about the potential of new technologies but cannot always invest the required time and resources at the same level that academic CI providers can. Through industry partnerships and, in special cases, technology transfer agreements, academic cyberinfrastructure providers can leverage their intellectual investments in new technology research.

Advanced computing infrastructure vendor relations are a special class of industry relationships, in that vendors are not only interested in seeing their solutions deployed in academic institutions that provide innovative, cost-effective solutions, but are also interested in working with CI facilities as a technology development partner. Forming a meaningful technical relationship with numerous vendors provides the leaders of CI facilities with technology roadmaps that are essential for strategic planning purposes. Further, vendor partnerships often lead to access to early release hardware and software for testing and performance evaluation purposes. As these partnerships mature, vendors can learn how to tailor their products, both current and future, to meet new research requirements, thus improving their ability to compete. These types of strategic partnerships also motivate vendors to provide more aggressive product pricing to their academic partners to help them be more competitive in grant competitions.

While industry partnerships should be considered by all institutions, it should be noted that not all institutions can engage in these partnerships or accept corporate funding due to the source of their primary funding.

9.0 Metrics of Success and Return on Investment (ROI)

Justifying technology and staff expenditures is an important issue for academic CI facilities. Institutions often look to make cuts in IT services first. In order to secure and sustain institutional CI support, it is helpful for CI directors to identify metrics of success and clearly and effectively communicate ROI to senior administrators on a regular basis.

Workshop participants differentiated between quantitative and qualitative metrics of success. It was also noted that the definition of success depends largely upon the audience that the metric findings are intended for. Quantitative metrics are measurable data that typically have straightforward collection methods, e.g., system accounting data, consulting logs (how consulting time was spent), sponsored program data such as those measured by the University at Buffalo's "Metrics on Demand" tool [18], and lists of grants and publications enabled. Qualitative metrics tend to be areas that intuitively sound compelling and believable, but generating statistical data to support them is a challenge. Customer satisfaction testimonials and internal and external committee reviews are common examples.

Workshop participants expressed an interest in developing more compelling quantitative metrics and accounting methods for CI. This is a "New Challenge Area" that needs additional attention, discussion and community collaboration.

9.1 Quantitative Metrics of Success

Workshop participants identified the following quantitative metrics of success:

- **Service Metrics** – These are typically based on standard accounting data. Examples include the number of user accounts, the percentage of campus researchers served, the number of departments supported, computing resource utilization and demand, and research data stored, served or transferred. Measurements are usually based on the fiscal year and show both accumulated numbers and new usage for that fiscal year as a means of showing growth.
- **"Science Driver" Metrics** – Communicate how an academic CI facility supports science at its institution. Examples include number of presentations and papers published as a result of having access to services and resources, the amount of staff time spent participating in, supporting, or enabling multidisciplinary research, and courses or workshops offered. Details for courses and workshops often include whether they are offered for academic credit, number of courses, workshops, or modules that are available, and the number of users and/or students that have taken advantage of them.
- **Funding Metrics** – The number of grants, awards and funding that can be attributed to having access to the services and resources provided by a CI facility. Examples include funds generated through an approved cost recovery model; new funds from grant proposals submitted and awarded, including awards such as the NSF CAREER award; external funding (federal funding agencies and industry) specifically for the CI facility or its staff and researchers; researcher participation in supported resources providing economies of scale such as condominium clusters or centralized research data storage; and, the number of jobs created and retained.

- **Intellectual Property Metrics** – The number of patents, copyrights, start-up companies enabled and industry agreements established or industry gifts given, based on having access to the services and resources provided by the CI facility. The depth of a CI facility relationship with a particular company can positively impact university-wide development, increasing the likelihood of, for example, gifts for new academic buildings or facilities, equipment donations, alumni giving by employees, etc.
- **Outreach Metrics** – Support for activities that broaden impact and reach underrepresented groups. These metrics are important in order to measure and improve upon the impact of projects on these communities. The establishment of activities that other researchers can leverage helps build and maintain credibility. Examples include support for NSF Research Experiences for Undergraduates (REUs) and frameworks for education and training such as the “Virtual Workshops” developed and delivered by the Cornell University Center for Advanced Computing for the NSF and other federal agencies [19].

9.2 Qualitative Metrics of Success

Workshop participants identified the following qualitative metrics of success:

- **Economic Development** – Again, based on funding and mission, this is the ability to act as a local, regional, or national resource in order to support industry by providing access to services and resources that make industry more competitive. As research computing becomes more prevalent in large commercial enterprises, this is becoming a more difficult ROI argument for industry; however, there is a growing opportunity with small and mid-size businesses, many of whom are embracing HPC and parallel computing for the first time.
- **Researcher Satisfaction** – Due to the availability of resources and services provided by CI facilities, many researchers and students are more than willing to make positive statements such as: "My productivity has increased significantly," "I have more time to do research and not worry about running my cluster," "I have more publications, or "I have more time to focus on my research and will graduate earlier." While this type of enthusiasm is essential for continued institutional support, it can be difficult to quantify, particularly in terms of cost savings or cost avoidance.
- **Strategic Metrics** – These metrics should communicate a cyberinfrastructure facility’s relevance and importance to its home and/or partner institutions. Examples include the impact on faculty and student recruitment and retention, the integration with regional and national resources such as TeraGrid and Open Science Grid, and partnering on large-scale national cyberinfrastructure proposals.

9.3 New Challenges

Workshop participants noted several areas where the methods of collecting data to provide new and potentially more meaningful metrics of success are needed:

- **Cost Savings and Cost Avoidance Metrics** – Measuring how much money is saved or establishing a dollar value for costs avoided by an institution due the availability of a CI facility

are metrics that can play an important role in securing ongoing institutional funding support. An example is the creation of a centralized data center. Intuitively it seems obvious that a centralized data center with optimal staffing, space, power and cooling for research computing should provide a huge cost savings. However, it can be difficult to provide an actual dollar amount for money saved or costs avoided by the existence of such a facility versus many distributed private clusters installed across a campus in facilities not designed for research computing equipment.

- **Institution Budget Metrics** – This is an institution’s understanding of the relative importance of a CI facility as critical core infrastructure and budgeting for it as such. Comparisons to other critical core infrastructure such as libraries, core facilities providing access to instrumentation (e.g., mass spectrometers, gene sequencers or clean rooms), and administrative IT are common, but are difficult to compare without considering the overall mission, metrics of success, and priorities of the institution. The growing and accelerating role of computation and data curation, mining, and analysis in research and education is not always understood or welcomed by university administrators. The value of CI must be effectively communicated to administrators, many of whom are not computational scientists.

10.0 Conclusions and Recommendations

This report describes the many ideas, strategies, models and experiences of the participants of the NSF workshop on sustainable funding and business models that are in use or under consideration at academic cyberinfrastructure facilities across the nation. There are many lessons learned in both the report and the accompanying 28 position papers listed in Appendix G. This report is not intended to promote any one specific CI funding or business model, but is offered as a summary for institutions that are reevaluating their funding strategies or starting a CI facility from scratch. Hopefully, the collegiality and openness that was exhibited by the 87 participants at the workshop and 32 WebEx participants that led to this report is only the beginning of continued discussions and sharing of experiences that will help broaden and strengthen computational science at all interested institutions. The Web site Sustainable Research Computing Centers (SRCC) at <http://www.cac.cornell.edu/SRCC> includes this report, links to presentations and position papers, and information on a SRCC LinkedIn social networking group that is available to facilitate further discussions. The conclusions and recommendations of this workshop are:

- **Broadening the CI Base** – The health and growth of computational science is critical to our nation’s competitiveness. The Branscomb Pyramid has been an accepted model for the computational science ecosystem since 1993, when it was described in the National Science Board Report 93-205 [20]. A significant amount of attention and energy is often focused at the top of the pyramid, as the excitement of extreme scale and performance is something everyone can appreciate. However, *the base or foundation of the computational pyramid must continue to develop and expand in order to produce the next generation of researchers and a US workforce equipped to effectively bring these innovations to bear on our global competitiveness.* The findings of this workshop will hopefully help more institutions play a meaningful role in a national cyberinfrastructure in which growing participation is crucial. Increased geographic participation through the development of regional models and the provisioning for adequate training were singled out by the workshop participants as two important needs.
- **Toward Sustainability** – Computational science has established itself as the third pillar of science complementing theory and experimentation. Data-intensive scientific discovery is emerging as the fourth paradigm. Because computational science and CI are essential infrastructure components of any academic institution that has research as a fundamental part of its mission, *sustained support for computational science is essential and should involve a partnership of national funding agencies, institutions of higher education, and industry.* Notably, *the model of support that is appropriate for each specific institution requires strategic vision and leadership with substantial input from a diversity of administrators, faculty and researchers.* Clearly, there is no “one-size-fits-all” solution. Strong institutional commitment through base funding is essential. State and federal funding through legislation and grants combined with various cost sharing mechanisms and recovery models that offer compelling value propositions by offering economies of scale are necessary to cover the remaining costs.
- **Continued Collaboration** – Organizations such as the Coalition for Academic Scientific Computation, Southeastern Universities Research Association, and the Great Plains Network provide the CI community an opportunity to continue discussions and sharing that started as a result of this workshop. *Support of computational science at all levels of US academic institutions will generate additional opportunities for collaboration, innovation, and, ultimately, the ability to compete globally and generate new economic growth.*

Policy and funding decisions that dis-incent collective community behavior, and that thereby impede shared improvement are harmful, and should be avoided.

Citations

- [1] Committee on Prospering in the Global Economy of the 21st Century: An Academy for American Science and Technology, National Academy of Sciences, National Academy of Engineering, Institute of Medicine. (2007). *Rising above the gathering storm: energizing and employing America for a brighter economic future*. Retrieved from: http://www.nap.edu/catalog.php?record_id=11463.
- [2] Members of the 2005 “Rising Above the Gathering Storm” Committee; Prepared by the Presidents of the National Academy of Sciences, National Academy of Engineering, and Institute of Medicine. (2010). *Rising above the gathering storm, revisited: rapidly approaching category 5*. Retrieved from: http://www.nap.edu/catalog.php?record_id=12999.
- [3] National Science Board – NSF Blue Ribbon Panel on High Performance Computing Report. (1993). *From desktop to teraflop: exploiting the U.S. lead in high performance computing*. Retrieved from: <http://www.nsf.gov/pubs/stis1993/nsb93205/nsb93205.txt>.
- [4] Hey, T., Tansley, S., Tolle, K. (2009). *The fourth paradigm: data-intensive scientific discovery*. Microsoft Research, pg. 252.
- [5] Lim, R., Jung, G., Welcome, T., Moore, R., Papadopolous, P., Mori, W.B. & Alvarez, R. (2010). *University of California shared research computing services pilot program*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/22.UC-SRCS_Position_Paper.pdf.
- [6] Joiner, D. (2010). *The New Jersey Institute for Modeling and Visualization*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/9.joiner_position_paper.pdf.
- [7] Adrion, R., Blank, K., Hill, C., Kurose, J. & Ruckenstein, A. (2010). *A Massachusetts green high-performance academic research computing facility*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/19.GHPCC_cornell_white_paper_final.pdf.
- [8] Crane, G., Robinson, J.-P. & Smith, P. (2010). *Enabling and sustaining campus-to-campus cyberinfrastructure*. Southeastern Universities Research Association paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/8.SURA-Campus-to-Campus-CI-Paper.pdf>.
- [9] McMullen, R. & Monaco, G. (2010). *The role of regional cyberinfrastructure in sustaining campus high performance computing centers*. Position paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/26.TheRoleofRegionalCyberinfrastructureinSustainingCampusHighPerformanceComputingCenters.pdf>.

- [10] Vandenberg, A., Sachon, M., Austin, S.F., Connolly, J., Crane, G., Hale, J., Johnson, S., Katz, D., Robinson, J.-P., Smith, P., Wong, V., Akli, L., Wang, D., Yafchak, M.F. (2008). *SURAGrid strategic plan 2008-2012*. Retrieved from:
<http://www.sura.org/programs/docs/SURAGridSPlan.pdf>.
- [11] The Coalition for Academic Scientific Computation home page mission statement. Retrieved from:
<http://casc.org/>.
- [12] Sills, E. (2010). *North Carolina State University model for providing campus high performance computing services*. Position paper. Retrieved from:
<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/15.NCStateHPCModel.pdf>.
- [13] National Science Foundation Press Release 10-077 (2010). Scientists seeking NSF funding will soon be required to submit data management plans. Retrieved from:
http://www.nsf.gov/news/news_summ.jsp?cntn_id=116928&org=NSF&from=news
- [14] National Coordination Office for Information Technology Research and Development, Executive Office of the President of the United States. (2005). *Report to the president on computational science: ensuing America's competitiveness*. (2005) pgs 12-14. Retrieved from:
<http://www.nitrd.gov/pitac/reports/index.html>.
- [15] Pummill, J., Brunson, D. & Apon, A. (2010). *Community funding models for computational resources*. Position paper. Retrieved from:
http://www.cac.cornell.edu/~lifka/Downloads/SRCC/14.pummill_brunson_apon.pdf.
- [16] Hillegas, C. (2010). *Princeton University position paper for the NSF Workshop – Sustainable Funding and Business Models for High Performance Computing Centers*. Retrieved from:
<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/25.PrincetonUniversity-NSFSustainableHPCCenters.pdf>.
- [17] Lifka, D., Alvord, R., Mehringer, S. & Redfern, P. (2010). *Overview of the Cornell University Center for Advanced Computing sustainable funding model*. White paper. Retrieved from:
<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/10.Overview-of-Cornell-CAC-Sustainable-Funding-Model.pdf>
- [18] Furlani, T. (2010). *Metrics of success and return on investment*. University of Buffalo Center for Computational Research position paper. Retrieved from:
http://www.cac.cornell.edu/~lifka/Downloads/SRCC/4.nsfunsolicited2009sustainctrs_whitepaper_furlani_20091212.pdf.
- [19] Cornell University Center for Advanced Computing Education and Outreach. (2010). Retrieved from: <http://www.cac.cornell.edu/education/train.aspx>.

[20] National Science Board – NSF Blue Ribbon Panel on High Performance Computing Report. (1993). *From desktop to teraflop: exploiting the U.S. lead in high performance computing*. Retrieved from: <http://www.nsf.gov/pubs/stis1993/nsb93205/nsb93205.txt>.

Appendix A: Workshop Announcement

National Science Foundation Workshop on Sustainable Funding and Business Models for High Performance Computing Centers
May 3 – May 5, 2010 at
Cornell University, Ithaca, NY

Applications to register and position papers are sought for the NSF-sponsored Workshop on Sustainable Funding and Business Models for High Performance Computing Centers. To apply for registration, please go to https://mw1.osc.edu/srcce/index.php/Main_Page and follow the links to register.

The purpose of the workshop is to provide a forum for an open discussion among High Performance Computing (HPC) center directors, campus information officers and campus research officers on the topic of sustainable funding of, and business models for, research computing centers. The discussion will yield a shared understanding of organizational models, funding models, management models and training models that result in sustainable funding for research computing centers.

Participants in the workshop will be better prepared to elucidate and champion the need for established research computing centers, and they will have the necessary data to explain how and why such centers must be established and can be sustained. Further, this workshop will prepare higher education institutions located in economically disadvantaged areas of the country with models for successful research computing centers that, if created and sustained, can markedly impact local economies. Additionally, by developing and sharing institutionally-siloed knowledge across diverse centers, this workshop will facilitate the establishment and implementation of similar centers elsewhere, and will strengthen and enrich broader learning communities. Finally, by promoting sustained research computing centers, this workshop will help to ensure early exposure to advanced computational concepts for all science and engineering students.

Up to seventy-five invited leaders in the operation and organizational administration of sustainable funding for HPC centers will participate on-site. In addition, WebEx conferencing of the meeting will reach additional participants. Broad engagement of the research computing community is sought, to ensure adequate representation from various stakeholders and also to ensure meaningful participation by all during the event.

Submission of position papers from the academic research computing community is strongly encouraged. The position paper process is intended to serve two purposes: (1) to solicit input from the larger community; (2) to serve as a mechanism for individuals to be selected to participate on-site in the workshop. Position papers are limited to 3 three pages and must be submitted by March 15, 2010. A review panel will review the papers and use them as the basis for deciding who will be invited to participate on-site.

Cornell University is hosting this NSF-sponsored workshop Monday, May 3 - Wednesday, May 5, 2010 in Ithaca NY. The workshop will include (a) an informal reception at 6pm on Mon May 3 at the Cornell Johnson Museum of Art and (b) an evening dinner cruise on Cayuga Lake on Tuesday, May 4. The workshop will conclude at 12:00 noon on Wednesday, May 5.

The organizing committee, along with an invited group of participants, will generate a complete report on the findings of the workshop. The report will also be posted on the on the CASC website and submitted to EDUCAUSE for publication.

Please feel free to contact members of the organizing committee by email if you have any additional questions, concerns or suggestions.

Organizing Committee

Stanley C. Ahalt, Ph.D.
Director, Renaissance Computing Institute
ahalt@renci.org

Amy Apon, Ph.D.
Director, Arkansas High Performance Computing Center, University of Arkansas
aapon@uark.edu

David Lifka, Ph.D.
Director, Cornell University Center for Advanced Computing
lifka@cac.cornell.edu

Henry Neeman, Ph.D.
Director, OU Supercomputing Center for Education and Research, University of Oklahoma
hneeman@ou.edu

Appendix B: Workshop Agenda

**Cornell University
May 3-5, 2010
Agenda**

Monday, May 3, 2010

- 3:00pm Afternoon Check-in at Statler Hotel & Marriott Executive Education Center, Cornell
6:00pm - 8:00pm Informal Reception, Herbert F. Johnson Museum of Art – sponsored by Intel (10 min. walk from Statler)

Tuesday, May 4, 2010

- 7:45am - 8:30am Continental Breakfast, 1st floor Statler Foyer with extra room in Yale/Princeton
8:30am – 8:40am Welcome to Cornell – Robert Buhrman, 1st floor Statler Amphitheater
8:40am - 9:00am Overview, Goals and Brief Introductions by Participants – Stan Ahalt
9:00am - 9:45am “The Cornell Center for Advanced Computing Sustainability Model” – Dave Lifka
9:45am - 10:30am "Bridging Campuses to National Cyberinfrastructure – Overview of OCI Sustainable Center Activities" - Jennifer Schopf
- 10:30am -11:00am Break
- 11:00am - 11:45am "The Penn State Sustainability Model" – Vijay Agarwala
11:45am - 12:00pm Afternoon agenda discussion and breakout planning – Stan Ahalt
- 12:00pm - 1:15pm Lunch at Carrier Grand Ballroom, 2nd floor Statler – sponsored by Dell
- 1:15pm - 3:00pm Breakout Sessions/Leads:

Organizational Models & Staffing – Stan Ahalt
Funding Models – Dave Lifka
Industry & Vendor Relationships – Amy Apon
Succession Planning – Henry Neeman
Metrics of Success and Return on Investment – Vijay Agarwala
- Breakout room capacities: Amphitheater – 92; Yale/Princeton – 44; Columbia – 20; Dartmouth –20; Harvard – 14
- 3:00pm - 3:30pm Break
- 3:30pm - 4:45pm Reports from the breakout sessions
- 5:15pm Meet in front of Statler Hotel to board Ithaca Limousine buses to Dinner Cruise
6:00pm – 9:00pm MV Columbia Dinner Cruise on Cayuga Lake. Boat departs Pier at 708 W. Buffalo St. at 6:00pm – sponsored by Dell
9:00pm Buses return to Statler Hotel

*For complete workshop information, visit the Sustainable Research Computing Centers wiki:
<http://www.cac.cornell.edu/SRCC>*

Wednesday, May 5, 2010

- 7:45am - 8:30am Continental Breakfast, Statler Foyer with extra room in Yale/Princeton
- 8:30am - 8:45am Welcome and Agenda Review – Stan Ahalt, Statler Amphitheater
- 8:45am - 9:45am Federal Funding Opportunities and Strategies for Tier-2 and Tier-3 Research Computing Centers - Jim Bottum
- 9:45am - 10:15am Open Discussion on the need for Collaboration and Advocacy – Henry Neeman
- 10:15am - 10:30am Break
- 10:30am - 11:30 Panel Discussion on Industry & Vendor Relationships – Moderator: Dave Lifka;
Panelists: David Barkai, Tim Carroll, Loren Dean, Ed Turkel
- 11:30am - 12:00pm Wrap up including identification of areas of consensus or lack thereof and report planning - Stan Ahalt & Dave Lifka
- 12:00pm Adjourn and Box Lunches available in Yale/Princeton room
- 12:00pm - 1:00pm Organizing committee generate report writing assignments and deadlines – Harvard room

Speakers/Panelists

Vijay Agarwala, Director, Research Computing and Cyberinfrastructure, Penn State University, vijay@psu.edu

Stanley C. Ahalt, Ph. D., Director, Renaissance Computing Institute, ahalt@renci.org

Amy W. Apon, Ph. D., Director, Arkansas High Performance Computing Center, University of Arkansas, aapon@uark.edu

David Barkai, Ph.D., HPC Computational Architect, Intel Corporation, david.barkai@intel.com

Jim Bottum, Ph.D., Vice Provost and Chief Information Officer for Computing and Information Technology, Clemson University, jb@clemson.edu

Robert Buhrman, Ph.D., Senior Vice Provost for Research, Cornell University, rab8@cornell.edu

Timothy Carroll, Senior Manager of HPC, Dell, tim_carroll@dell.com

Loren Dean, Director of Engineering, MATLAB Products, MathWorks, loren.dean@mathworks.com

David Lifka, Ph. D., Director, Cornell University Center for Advanced Computing, lifka@cac.cornell.edu

Henry Neeman, Ph. D., Director, OU Supercomputing Center for Education & Research, University of Oklahoma, hneeman@ou.edu

Jennifer Schopf, Ph.D., Program Officer, National Science Foundation, jschopf@nsf.gov

Ed Turkel, Manager, Business Development, Scalable Computing & Infrastructure Organization, Hewlett Packard Company, Ed.Turkel@hp.com

Workshop Discussion Topics

The following six topics were the focus of the workshop with a particular focus on the needs and goals of second and third tier high performance computing centers.

1. Organizational Models and Staffing

Currently a number of such models exist. Centers are in place as separate entities subsidized by a consortium of individual universities. They may also exist as a part of a larger Information Technology operation on a campus, as a division within an institution's research administration structure, as a research center associated with one or several colleges within a university, or in various hybrid forms. Leaders representing each area will present an overview of these organizational models and the advantages and disadvantages of each.

2. Funding Models

As central subsidies for centers decline, various fee-for-service models are being put into place. The mix of services and fee structures range across a number of categories from maintenance and management of computing resources to consultation with major research projects, to a package of fees for services. Centers are also increasingly competing for extramural funds for both research and industrial contracts. We will discuss examples of each of these funding models and the markets or situations in which they appear to be most successful.

3. Vendor Relationships

Smaller centers do not have the buying power of the major centers and thus are less likely to receive the pricing and array of options available to those larger entities. Strategies that emerge from this situation range from creating strong ties with a single vendor to encouraging long-term and better support to developing local expertise and expending staff resources on assembling heterogeneous systems. We will discuss various strategies and problems associated with vendor relationships, as well as the potential for regional and/or national cooperation that might lead to a broader set of options.

4. Succession Planning

Many centers have limited staff and thus potentially face major problems as key leaders or key staff retire or take positions elsewhere. With a very limited pool of expertise in high performance computing, such transitions can lead to the demise of a center unless actions are taken to anticipate possible changes and to provide a succession plan that will work. At the same time, many centers are being asked to transition from one organizational model to another. Such transitions pose similar problems, as staff may resent the changes and thus may move to alternative jobs. These issues will be discussed and potential approaches to their solution will be discussed.

5. Metrics of Success and Return on Investment

As budgets become tighter, centers are increasingly asked to justify their return on investment. Metrics are therefore becoming an increasingly important aspect relating to the survival of HPC centers. Approaches to defining metrics of success such as return on investment, gathering and maintaining the

necessary data such as resource usage and usability, depth and breadth of impact, and effective means of presenting them to key decision-makers will be discussed.

6. Industry Relationships

As both industry and academic centers are pressed by budget limitations, there are opportunities for joint projects with and services to industry that could become an important aspect of center activities. Examples of industry partnerships, services and service models, and the challenges of developing an industrial customer base will be addressed at the workshop.

Appendix C: Terminology

- **Cyberinfrastructure (CI) Facilities** – Centers or a centralized group or organization within an academic institution that provide research computing resources and services. This term is meant to be more inclusive than “Centers” because many workshop participants who provide research computing services and resources at their institution are not part of a “Center,” but are a group within Central IT or another organization.
- **Core Facilities** – A group that provides research infrastructure typically under a fee-for-service model in academic institution. Traditional core facilities typically provide access to expensive instrumentation or facilities such as gene sequencers or clean rooms for nano-fabrication.

Appendix D: Workshop Participants (On-Site Participation)

1. Agarwala, Vijay
Director of Research and Cyberinfrastructure
The Pennsylvania State University
2. Ahalt, Stanley
Director
Renaissance Computing Institute
3. Allen, Gabrielle
Associate Professor, Computer Science
Louisiana State University
4. Apon, Amy
Director, Arkansas HPC Center
University of Arkansas
5. Athanasoulis, Marcos
Director of Research IT
Harvard Medical School
6. Atkins, Daniel
Associate Vice President for
Research Infrastructure
University of Michigan
7. Bangalore, Purushotham
Associate Professor and Director
Collaborative Computing Laboratory
University of Alabama at Birmingham
8. Bansil, Arun
Professor, Physics
Northeastern University
9. Barkai, David
HPC Computational Architect
Intel
10. Bose, Rajendra
Manager, Research Computing Services
Columbia University
11. Bottum, James
Vice Provost and CIO
Clemson University
12. Bozylinski, Garrett
CIO
University of Rhode Island

-
- | | |
|--|---------------------------|
| 13. Brenner, Paul
Associate Director, HPC Group | University of Notre Dame |
| 14. Bresnahan, Glenn
Director, Scientific Computing & Visualization | Boston University |
| 15. Buhrman, Robert
Senior Vice Provost for Research | Cornell University |
| 16. Campbell, John
Associate Vice President,
Center for Advanced Computing | Purdue University |
| 17. Carlson, Doug
Assistant Vice President,
Communications & Computing | New York University |
| 18. Carroll, Timothy
Senior Manager of HPC | Dell |
| 19. Clebsch, William
Associate Vice President, IT Services | Stanford University |
| 20. Connolly, John
Director, Center for Computational Sciences | University of Kentucky |
| 21. Crane, Gary
Director, IT Initiatives | Southeastern Universities |
| 22. Crosswell, Alan
Associate Vice President &
Chief Technologist | Columbia University |
| 23. Crowley, Kate
Director of IT Finance | University of Rochester |
| 24. Dean, Loren
Director of Engineering, MATLAB | The MathWorks |
| 25. Deumens, Erik
Director, HPC Center | University of Florida |

26. Devins, Robert
HPC Researcher
University of Vermont
27. Dougherty, Maureen
Supervising Systems Programmer
University of Southern California
28. Fratkin, Susan
Washington Liaison
Coalition for Academic Scientific Computation
29. Fredericksen, Eric
Associate Vice Provost,
University IT – Academic & Research
University of Rochester
30. Fronczak, Christine
HPC Marketing Programs Manager
Dell
31. Fujimoto, Richard
Chair, School of Computational
Science & Engineering
Georgia Institute of Technology
32. Furlani, Thomas
Director, Center for Computational Research
University at Buffalo
33. Gemmill, Jill
Executive Director, Cyberinfrastructure
Technology Integration
Clemson University
34. Goldiez, Brian
Deputy Director,
Institution of Simulation & Training
University of Central Florida
35. Hamilton, Victoria
Director, Research Initiatives
Columbia University
36. Hargitai, Joseph
Faculty Technology Specialist
New York University
37. Hauser, Thomas
Associate Director, Research Computing
Northwestern University

-
- | | |
|--|---------------------------------------|
| 38. Hawkins, Ronald
Dept. Director, Industry Relations | San Diego Supercomputer Center |
| 39. Hillegas, Curt
Director, TIGRESS HPC Center | Princeton University |
| 40. Johnsson, Lennart
Director, Texas Learning &
Computation Center | University of Houston |
| 41. Katz, Daniel
Senior Computational Researcher | University of Chicago |
| 42. Khanna, Gurcharan
Director of Research Computing | Rochester Institute of Technology |
| 43. Klingenberg, Alys
Assistant Director, Finance & Operations | University of Rochester |
| 44. Kohlmeyer, Axel
Associate Director, Institute for Computational
Molecular Science | Temple University |
| 45. Krishnamurty, Ashok
Interim, Co-Executive Director | Ohio Supercomputer Center |
| 46. Labate, Bill
Director, Academic Technology Services | University of California, Los Angeles |
| 47. Lance, Timothy
President | NYSERNet |
| 48. Lifka, David
Director, Cornell Center for Advanced
Computing & Director Research
Computing, Weill Cornell Medical College | Cornell University |
| 49. Lim, Ramon
Director of Strategic IT Projects
Office of the President | University of California |

50. Liu, Honggao
Director of HPC
Louisiana State University
51. Lombardi, Julian
Assistant Vice President
Office of Information Technology
Duke University
52. Majchrzak, Daniel
Director of Research Computing
University of South Florida
53. Marinshaw, Ruth
Assistant Vice Chancellor for
Research Computing
University of North Carolina-Chapel Hill
54. Marler, Bryan
Director, High Performance Computing
Hewlett-Packard
55. McMullen, Donald
Director of Research Computing
University of Kansas
56. Mehringer, Susan
Assistant Director, Consulting
Cornell Center for Advanced Computing
Cornell University
57. Miller, Therese
Chief Operating Office
Pervasive Technology Institute
Indiana University
58. Monaco, Greg
Director for Research and Cyberinfrastructure
Great Plains Network (GPN)
59. Moore, Richard
Deputy Director
San Diego Supercomputer Center
60. Nabrzyski, Jaroslaw
Director, Center for Research Computing
University of Notre Dame
61. Neeman, Henry
Director OU Supercomputing Center
for Education & Research
University of Oklahoma

-
- | | |
|--|-------------------------------------|
| 62. Odegard, Jan
Executive Director, Ken Kennedy
Institute for Information Technology | Rice University |
| 63. Papadopoulos, Philip
Program Director UC Computing Systems | San Diego Supercomputer Center |
| 64. Pepin, James
CTO and Director of HPC | Clemson University |
| 65. Redfern, Paul
Assistant Director, Strategic Partnerships
Cornell Center for Advanced Computing | Cornell University |
| 66. Reese, Philip
Research Computing Strategist | Stanford University |
| 67. Ricks, Matthew
Executive Director Computing Services | Stanford University |
| 68. Robinson, John-Paul
System Programmer Lead | University of Alabama at Birmingham |
| 69. Roeber, Ronald
Director, Research Computing &
Data Management | University of Nebraska-Lincoln |
| 70. Rohrs, Lynn
Manager, Research Computing Services | New York University |
| 71. Schopf, Jennifer
Program Officer | National Science Foundation |
| 72. Siedow, James
Vice Provost for Research | Duke University |
| 73. Sills, Eric
Director for Advanced Computing | North Carolina State University |
| 74. Slattery, Paul
Dean of Research | University of Rochester |

-
- | | |
|--|----------------------------------|
| 75. Smith, Philip
Senior Director, HPC Center | Texas Tech University |
| 76. Spadaro, John
Director of Technical Architecture | Brown University |
| 77. Stampalia, Jacqueline
Associate Director of Research | Rensselaer Polytechnic Institute |
| 78. Stanzione, Dan
Deputy Director | Texas Advanced Computing Center |
| 79. Steinhardt, Sam
Executive Director, Business Services | Stanford University |
| 80. Topham, David
Assoc. Prof., Microbiology & Immunology | University of Rochester |
| 81. Tsinoremas, Nicholas
Director, Miami Center for Computational
Sciences | University of Miami |
| 82. Turkel, Ed
Manager, SCI Business Development | Hewlett-Packard |
| 83. Vandenberg, Art
Account Representative,
Information Systems & Technology | Georgia State University |
| 84. Warnes, Gregory
Director, Center for Research Computing | University of Rochester |
| 85. Wilgenbusch, James
Director of HPC Facility | Florida State University |
| 86. Wroblewski, William
Executive Director, Infrastructure Service | University of Michigan |
| 87. Yetsko, Mary
Workshop Coordinator
Cornell Center for Advanced Computing | Cornell University |

Appendix E: Workshop Participants (Web-Based Participation)

1. Alvord, Resa
Assistant Director, Systems
Cornell Center for Advanced Computing
Cornell University
2. Andresen, Daniel
Associate Professor, Department of
Computing and Information Sciences
Kansas State University
3. Brauner, Shane
Director, IT
Schrodinger
4. Brunson, Dana
Senior Systems Engineer – HPCC
Oklahoma State University
5. Chourasia, Amit
Senior Visualization Specialist
San Diego Supercomputer Center/
University of Californian, San Diego
6. Combariza, Jaime
Associate Director of Research Computing
Dartmouth College
7. Dreher, Patrick
Chief Domain Scientist, Cloud Computing
Renaissance Computing Institute
8. Duffy, Edward
Associate Researcher, Cyberinfrastructure
Technology Integration
Clemson University
9. Greve, Ron
Assistant Director of Research
South Dakota State University
10. Grooms, Jerry
HPC Systems Manager
University of Kentucky
11. Hare, Tracey
Administrative Coordinator
Clemson University
12. Hart, Brian
Manager, CSSD Business and
Resource Management
University of Pittsburgh

13. Hellman, Rebecca
Business Manager & Acting Director
Center for High Performance Computing
North Dakota State University
14. Heslop, Janet
IT Associate Director
Cornell University
15. Johnson, Del
Director, Research Computing
South Dakota State University
16. Joiner, David
Assistant Professor
Kean University
17. Jung, Gary
Manager, High Performance Computing
Services
Lawrence Berkeley National Laboratory
18. Meacham, Steve
Senior Staff Associate
National Science Foundation
19. Mendan, Joseph
Financial Reporting Specialist/CIMS
Rochester Institute of Technology
20. Mort, Brendan
HPC Consultant & Computational Scientist
University of Rochester
21. Nordquist, Russell
HPC System Administrator
University of Rochester
22. O'Leary, Patrick
Director, Center for Advanced Modeling and
Simulation
Idaho National Laboratory
23. Poole, Michael
Associate Vice President and
Chief Technology Officer
University of Massachusetts
24. Pormann, John
Director of Scalable Computing Support Center
Duke University

25. Rupp, Bill
Research and Development Analyst
University of Pittsburgh
26. Schopf, Paul
Dean of Research & Computing
George Mason University
27. Sugimura, Tak
Director, Hawaii Institute for
Molecular Sciences
University of Hawaii
28. Swanson, David
Director, Holland Computing Center
University of Nebraska-Lincoln
29. Taylor, Jackie
Director of College Partnerships
Rochester Institute of Technology
30. Teig von Hoffman, Jennifer
Assistant Director, Scientific Computing and
Visualization
Boston University
31. von Oehsen, James
Director of Computational Science
Clemson University
32. Walsh, Kevin
Graduate Student
University of California, San Diego

Appendix F: Workshop Presentations and Breakout Findings

Cornell University Center for Advanced Computing Sustainability Model – David Lifka

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/Lifka.pdf>

Bridging Campuses to National Cyberinfrastructure: Overview of OCI Sustainable Center Activities – Jennifer Schopf

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/Schopf.pdf>

Penn State Sustainability Model – Vijay Agarwala

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/Agarwala.pdf>

Sustainability for HPC Centers, A Macro View – Jim Bottum

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/Bottum.pdf>

Open Discussion on the need for Collaboration and Advocacy – Henry Neeman

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/Neeman.pdf>

Dan Atkins' Principles – Dan Atkins

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/AtkinsPrinciples.pdf>

Breakout Findings: Organizational Models, Staffing & Succession Planning

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/Breakout1.pdf>

Breakout Findings: Funding Models, Industry & Vendor Relationships

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/Breakout2.pdf>

Breakout Findings: Metrics of Success and Return on Investment

<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/Breakout3.pdf>

Appendix G: Workshop Position Papers

1. Allen, G. & Katz, D.S. (2009). *Computational science, infrastructure and interdisciplinary research on university campuses: experiences and lessons from the Center for Computation and Technology*. Louisiana State University position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/1.AllenKatz_WhitePaper.pdf.
2. Goldiez, B., Tafur, S. & Palaniappian, R. (2009). *HPC sustainability plan and position paper*. University of Central Florida position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/2.HPC-Sustainability-Plan-and-Position-Paper_v2.pdf.
3. Smith, P. (2009). *HPC sustainability (metrics)*. Texas Tech High Performance Computing Center position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/3.HPC_Sustainability_metrics.pdf.
4. Furlani, T. (2010). *Metrics of success and return on investment*. University of Buffalo Center for Computational Research position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/4.nsfunsolicited2009sustainctrs_whitepaper_furlani_20091212.pdf.
5. Kennedy, C.L. & Sheldon, P. (2010). *The Vanderbilt Advanced Computing Center for Research and Education*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/6.nsfunsolicited2009sustainctrs_whitepaper_sheldon_20100113.pdf.
6. Dougherty, M. (2010). *High-performance computing sustainability at the University of Southern California*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/7.USC-High_Performance_Computing_Sustainability.pdf.
7. Crane, G., Robinson, J.-P. & Smith, P. (2010). *Enabling and sustaining campus-to-campus cyberinfrastructure*. Southeastern Universities Research Association position paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/8.SURA-Campus-to-Campus-CI-Paper.pdf>.
8. Joiner, D. (2010). *The New Jersey Institute for Modeling and Visualization*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/9.joiner_position_paper.pdf.

9. Beck, S.D. (2010). *Computational research across the curriculum: addressing the challenges of sustainable HPC at the academy*. Louisiana State University Center for Computational & Technology position paper. Retrieved from:
<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/10.ComputationXCurriculum.pdf>.
10. Lifka, D., Alvord, R., Mehringer, S. & Redfern, P. (2010). *Overview of the Cornell University Center for Advanced Computing sustainable funding model*. White paper. Retrieved from:
<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/10.Overview-of-Cornell-CAC-Sustainable-Funding-Model.pdf>.
11. Atlas, S. R. (2010). *White paper submitted to the NSF Workshop on High Performance Computing Center Sustainability*. Retrieved from:
http://www.cac.cornell.edu/~lifka/Downloads/SRCC/11.AtlasCARC-SustainabilityConference2010WhitePaper_012410_final.pdf.
12. Bose, R., Crosswell, A. Hamilton, V. & Mesa, N. (2010). *Piloting sustainable HPC for research at Columbia*. Position paper. Retrieved from:
http://www.cac.cornell.edu/~lifka/Downloads/SRCC/12.Columbia_NSF_sustDC_position_paper_3pg.pdf.
13. Devins, R. & Elledge, A. (2010). *A user-centric approach to HPC sustainability*. Vermont Advanced Computing Center position paper. Retrieved from:
http://www.cac.cornell.edu/~lifka/Downloads/SRCC/13.UVM_SRCC_Whitepaper.pdf.
14. Pummill, J., Brunson, D. & Apon, A. (2010). *Community funding models for computational resources*. Position paper. Retrieved from:
http://www.cac.cornell.edu/~lifka/Downloads/SRCC/14.pummill_brunson_apon.pdf.
15. Sills, E. (2010). *North Carolina State University model for providing campus high performance computing services*. Position paper. Retrieved from:
<http://www.cac.cornell.edu/~lifka/Downloads/SRCC/15.NCStateHPCModel.pdf>.
16. Liu, H. (2010). *The sustainability of high performance computing at Louisiana State University*. Position paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/16.LSU-HPC-Sustainability.pdf>.
17. Wilgenbusch, J. (2010). *Sustainable research computing at Florida State University*. Position paper. Retrieved from:
http://www.cac.cornell.edu/~lifka/Downloads/SRCC/17.FSU_HPC_NSF_sustain.pdf.

18. Deumens, E. & Taylor, C. (2010). *Experiences with sustainability at the University of Florida HPC Center*. White paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/18.UF-HPCC-sustainability.pdf>.
19. Adrion, R., Blank, K., Hill, C., Kurose, J. & Ruckenstein, A. (2010). *A Massachusetts green high-performance academic research computing facility*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/19.GHPCC_cornell_white_paper_final.pdf.
20. Papadopoulos, P., Moore, R. & Hawkins, R. (2010). *The Triton Resource at UCSD – implementing sustainable HPC at the University of California*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/20.SDSC_position_paper_final_03122010.pdf.
21. Clebsch, B., Steinhardt, S., Ricks, M. & Reese, P. (2010). *Culture change and research computing at Stanford*. Position paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/21.Stanford-SRCC-NSFWorkshop.pdf>.
22. Lim, R., Jung, G., Welcome, T., Moore, R., Papadopolous, P., Mori, W.B. & Alvarez, R. (2010). *University of California shared research computing services pilot program*. Position paper. Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/22.UC-SRCS_Position_Paper.pdf.
23. Mori, W., Davis, J. & Labate, B. (2010). *The UCLA shared cluster hosting program, providing sustainable HPC Services at UCLA*. Position paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/23.UCLASustHPC.pdf>.
24. Krishnamurthy, A., Gordon, S., Kelley, K., Wohlever, K. & Chalker, A. (2010.) *Data center needs for a sustainable HPC center: issues facing the Ohio Supercomputer Center*. Position paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/24.100315NSFSustainableCenters.pdf>.
25. Hillegas, C. (2010). *Princeton University position paper for the NSF Workshop – Sustainable Funding and Business Models for High Performance Computing Centers*. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/25.PrincetonUniversity-NSFSustainableHPCCenters.pdf>.
26. McMullen, R. & Monaco, G. (2010). *The role of regional cyberinfrastructure in sustaining campus high performance computing centers*. Position paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/26.TheRoleofRegionalCyberinfrastructureinSustainingCampusHighPerformanceComputingCenters.pdf>.

27. Combariza, J.E. (2010). *Research computing cyberinfrastructure at Dartmouth College*. Position paper. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/27.RCI-DartmouthCollege.pdf>.
28. Modi, J., Schwenker, A., Sinha, R. & Tolly, M. (2010). *Survey of CASC members on the role of industrial partnership programs in high performance computing center sustainability*. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/28.CASC-Survey-Industrial-Partnerships.pdf>.

Appendix H: Related Papers, Presentations, Web Sites

1. Voloudakis, J. (2010). *A systemic model for IT economic sustainability*. EDUCAUSE Center for Applied Research, bulletin 1. Retrieved from:
<http://www.educause.edu/Resources/ASystemicModelforITEconomicSus/194795>.
2. Berman, F. (2008). *Making research and education cyberinfrastructure real*. EDUCAUSE Review, vol. 43, no.4. Retrieved from: <http://www.cac.cornell.edu/~lifka/Downloads/SRCC/EducauseCI-FINAL08.pdf>.
3. Lifka, D. (2010). *High Performance Computing Center Sustainability NSF Workshop*. Invited presentation at the Coalition for Academic Scientific Computation (CASC) Fall 2009 meeting. Retrieved from: <http://www.cac.cornell.edu/SRCC>.
4. Education Advisory Board (2009). *University business executive roundtable agenda - national best practice initiatives 2009-2010*. Retrieved from:
http://www.educationadvisoryboard.com/uber_council_agenda.asp.

Acknowledgements

This workshop was funded in part by the National Science Foundation Office of Cyberinfrastructure grant number 0944039. The Organizing Committee gratefully acknowledges the National Science Foundation, Dell Inc. and Intel Corporation for their support of this workshop. The Committee would also like to thank everyone who took time to attend the workshop for their open and collegial participation and the Cornell University Center for Advanced Computing for hosting the workshop.