High Performance Academic/Research Computing Assessment

Supporting the Needs of the Research and Academic Communities
Review and Recommendations

Executive Summary

Report Prepared for UAF Chancellor Brian Rogers
By the University of Alaska Office of Information Technology

Steve Smith, Chief Information Technology Officer
Karl Kowalski, Executive Director User Services
James Durkee, Executive Director Technology Oversight

February 15, 2011

Executive Summary

In October 2010, UAF Chancellor, Brian Rogers, tasked the Office of Information Technology with preparing an assessment of the research and academic demands for computational science and high performance computing (HPC) at UAF. This study evaluated the demands for computing capacity, storage, and support staffing to meet current and projected research and academic demands for high performance computing and thus makes recommendations for its ongoing support.

The process engaged UAF, UAA and UAS researchers, faculty, Deans and Institute Directors and key experts from across the country.

Key Observations

Executive leadership at the University is engaged and committed to take action to preserve critical high performance computing capabilities and to implement a plan to take Alaska into the future.

The University of Alaska has a unique proposition as a potential partner for institutions outside of the arctic. Alaska's physical location in the arctic, the experience of its faculty, students, and staff in working in arctic environments, and its deep expertise in the many disciplines related to the arctic make Alaska a very attractive partner for other universities with established or developing programs of arctic study.

People, especially PhD-level scientists with the specialized HPC skills to partner with faculty to effectively leverage computational techniques in their research, are the most critical component of an overall HPC strategy.

The Arctic Region Supercomputing Center has built a well-respected reputation in the HPC community and the brand may be worth preserving. However, the ARSC model represents a particular type of HPC support and comes with some local baggage. If the University chooses to broaden its efforts beyond this model, rebranding is advisable.

The University of Alaska has an advantage should it choose to focus on data-intensive computing due to the presence of experts in data management. The generation (through simulation) and collection (through remote sensing and other automated technologies) of massive data sets, and the associated data intensive computing, are hallmarks of modern HPC.

The University needs to maintain a core of central HPC compute resources (both clusters and high throughput Condor-style grids) to handle funded projects that need to buy services, pre-funding exploratory work, unfunded research (to the extent the university chooses to support it), and research data management.

The investment in high performance computing resources should be viewed as a system resource. Engagement and partnership with the other UA campuses and with State and local agencies will greatly enhance UAF’s position as the HPC center for the State.

Finally, any coordinated approach needs to leverage the extensive and very high end resources available at little or no cost through the TeraGrid and the national labs.

Near Term Recommendations (0-12 Months)

UAF should immediately begin a broad communications campaign designed to communicate clearly to all interested members of the University community the current state of affairs, the commitment to excellence in HPC being made by the University, the plans for the near term, and how people can get involved and remain informed as long term plans are developed.

UAF should articulate the role of computational science/high performance computing in its mission as a research University. This should be reflected in UAF Vision 2017 and UAF's strategic academic and research master plans. If it is important, the University should say so.

UAF should refocus ARSC on a carefully chosen set of HPC services and partnerships that focus on the competitive strengths that Alaska enjoys, and place it within the organizational structure of UAF. The science and computational support should report to the Vice Chancellor for Research. The technical and operational support should lie with the Office of Information Technology. This organizational structure leverages the operational expertise and vendor relationships that exist in IT and enables the scientific research community to focus on research and computation.

UAF should maintain HPC operational staff at least at the present levels in order to provide support for the existing and planned systems.

UAF should maintain and support high performance computing at a base level of 30% of the operational costs and adopt a model to continue sustainable support. The first transitional year will require a higher initial level of support.

UAF should seek commitment by all Institute Directors and Deans to support (through actions and finance) high performance computing for research and academics.

UAF should seek to maintain its presence as the statewide center for high performance computing, ensuring engagement with UAA's research community, UAS, the State of Alaska, other research organizations in the State (e.g. BASC), and establish relationships that are mutually beneficial.

UAF should establish and support a strong and broadly-represented faculty governance group to direct investment in high performance computing over time.

UAF should implement a plan to fund, recruit and retain PhD-level scientists with specialized HPC skills to partner with faculty on their research. Such a plan should include provisions for consultation with faculty during proposal preparation, buyout of fractions of the HPC specialists’ time over extended periods when grants are funded, and bridge the HPC specialists between grants.

UAF should decommission the Midnight cluster. It is at the end of its useful service life and provides little or no capacity not available more efficiently on the new PACMAN cluster. Near and long term needs and expectations of the University related to the costs of power for HPC systems should be discussed openly as plans are made for new acquisitions for end-of-life systems. While this system is owned outright, UAF must carefully consider whether any revenues obtained from the sale of compute cycles outside UAF is cost effective given the operational costs.

Medium Term Recommendations (1-3 Years)

UAF should consider the way it recruits and hires scientists. It should seek scientists with skills in numerical, quantitative, computational science. This will serve to broaden the computational base and bring expertise to the research units in which they reside.

UAF should seek the partnership of other universities, research institutes and national labs with compatible research interests and commitment to high performance computing.

UAF should seek external funding to maintain, upgrade and replace the PACMAN cluster and related HPC infrastructure as they age.

UAF should pursue additional network capacity connecting Alaska with the rest of the world. As success in HPC grows, individual projects will require dedicated network resources beyond the University's existing capacity.

Long Term Recommendations (3+ Years)

UAF should continue to leverage grant resources such as EPSCOR RII to fund and sustain HPC infrastructure.

UAF should pursue long-term center support through grants such as NSF Industry & University Cooperative Research Program (I/UCRC) and other forms of industrial collaboration.

UAF should continue to support PhD level scientists with specialized HPC skills over the long term.

It important that the University of Alaska recognize that high performance computing is a system resource. HPC equipment and infrastructure are expensive to implement and support, duplication of service within UA does not make sense. Collaboration with the other campuses and their faculty will also be key in building a strong and sustainable HPC platform. Through the implementation of these recommendations and through the establishment of a transparent, reliable HPC organizational and funding model, greater participation and increased sustainability will result throughout the UA system.

High Performance Academic/Research Computing Assessment

Supporting the Needs of the Research and Academic Communities
Review and Recommendations

Report Prepared for UAF Chancellor Brian Rogers
By the University of Alaska Office of Information Technology

Steve Smith, Chief Information Technology Officer
Karl Kowalski, Executive Director User Services
James Durkee, Executive Director Technology Oversight

February 15, 2011

"If [physical] infrastructure is required for an industrial economy, then we could say that cyberinfrastructure is required by a knowledge economy"

Atkins Report, "Revolutionizing Science and Engineering through Cyberinfrastructure"

Table of Contents

Executive Overview 1
What is Computational Science? 4
National Trends........ 5
History of High Performance Computing at UAF 7
    Arctic Region Supercomputing Center........ 7
    Alaska Satellite Facility........ 9
    Institute of Arctic Biology 9
    Life Science Informatics 9
    Department of Chemistry........10
UAF Research Goals 11
OIT Interview and Community Engagement Results 12
    Hardware/software 13
    Storage........13
    People:........14
Discussion of Financial and Organizational models 16
    Financial Models 16
    Campus-Funded........16
        Benefits: 17
        Challenges of Campus-Funded: 17
    Externally-Funded 17
        Benefits of Externally-Funded........17
        Challenges of Exterally-Funded........17
    Cost Recovery (recharge) Model........18
        Benefits of Cost Recovery Model 18
        Challenges of Cost Recovery Model........18
Organizational Models 18
    Central IT Model........18

Research Unit Model 19
Shared Responsibility Model 19
Metrics of Success and Return on Investment 19
Quantitative Metrics of Success 19
Qualitative Metrics of Success 20
New Challenges 21
Findings & Recommendations 22
Near Term Recommendations (0-12 Months) 22
Medium Term Recommendations (1-3 Years) 24
Long Term Recommendations (3+ Years) 24
Summary 25
Endnotes 26
Appendices .....
Report from the OIT HPC Advisory Committee ..... A
High Performance Academic Research Computing Survey ..... B
High Performance Computing (HPC) Consultant Visit Agenda C
NSF Sustainable Funding and Business Models for Academic Cyberinfrastructure Facilities ..... D
Resources Used in Preparation of this Report ..... E
Chancellor Roger's Memo on High Performance Academic/Research Computing Assessment ..... F

Executive Overview

This report reviews the research High Performance Computing (HPC) and Computational Science (CS) needs for the University of Alaska and makes short term and long term recommendations. It is based on interviews with University of Alaska Fairbanks (UAF) Institute Directors, Deans, faculty and researchers, University of Alaska Anchorage (UAA) researchers, a faculty survey, technology leaders from other research Universities, and the work of three external reviewers.

The assigned overarching goal is to enable research and academic advancement across the disciplines while getting as much value as possible from every dollar spent on high performance computing resources. The University of Alaska is rich in computing resources but can be more effective by adopting best practices that support greater coordination and cooperation at all levels. This will require UAF to plan, share, leverage existing resources and expertise, design funding models that promote sustainable resources, rely on user governance and implement recommendations of the research computing needs analysis.

Since 1993, through Department of Defense funding, the Arctic Region Supercomputing Center (ARSC) provided computational resources for researchers at the University of Alaska and the Department of Defense as a DoD Supercomputing Resource Center (DSRC). This funding is scheduled to end in May 2011. Researchers and departments both support and depend on ARSC while others express frustration and concerns. In spite of these differing views, it is clear that the University has benefitted from ARSC's computational resources. More importantly UAF has benefited from the staff that supports computational science for academic researchers. The loss of DoD funding for ARSC raises the question, *"What is needed to continue computational support for research and academics at the University of Alaska?".*

This report reviews and makes recommendations in several areas:

1) Demand for High Performance Computing/Computational Science at UAF and throughout the UA system
2) Key resources required to support HPC
3) HPC organizational structure
4) Requirements for sustainable funding
5) Metrics

Demand. Computational Science is recognized as the third pillar of science along the traditional areas of theory and experimentation. UA researchers identified computational science as an important part of their science and research in both the survey and individual interviews. Seventy-four percent (74%) of survey respondents reported using some level of high performance computing ranging from local clusters to national supercomputing resources.

Resources. There is clear consensus from UA researchers and directors that the most valuable resources are the individuals, computational experts, able to translate the science to computer models for their research. These are typically individuals with a PhD in a discipline who are also fluent in programming, modeling, visualization and debugging. Access to high performance computers is also important, but not as critical as the people. Current resources provided by the PACMAN and MRI-R2 grants, as well as available local, departmental clusters, appear adequate for the research demand through the FY11 & FY12 fiscal cycles.

Organization. Multiple organizational models were reviewed and it is recommended to split scientific and computational science support from technical and operational support for hardware and software. Computational scientists should report to the Vice Chancellor for Research and technical and operational staff should reside in the central IT operations (currently Office of Information Technology). While this only addresses the organization of centralized resources, it is important to note that many departmentally based resources exist as well. It is our hope that through the establishment of a transparent, reliable HPC organizational and funding model that greater participation will result.

Sustainable Funding. It is clear from recent events surrounding ARSC that UA should not rely too heavily on any single funding source. A combination of base funding from the campuses, support from individual research projects, larger grants, such as the NSF EPSCoR RII grants, and public and private partnerships are needed to diversify and stabilize UA's HPC position. Future grants should include line items for staff and equipment. UAF is data rich in high quality, arctic systems research. It will need to leverage the valuable data collected and develop long term relationships with researchers at other institutions interested in arctic research.

Going forward, it important that the University of Alaska recognize that high performance computing is a system resource. HPC equipment and infrastructure are expensive to implement and support and duplication of service within UA does not make sense. Collaboration with the other campuses and their faculty are key in building a strong and sustainable HPC platform.

Metrics. Given fiscal pressures facing the University and increasing accountability, it will be strategically important that the University establishes key metrics for evaluating the success of its HPC program. These should be both quantitative and qualitative.

Introduction

The University of Alaska Fairbanks has an outstanding reputation and tradition as a research leader. UAF is close to the top 100 in annual federal obligations for research averaging over $120 million dollars per year. [i] As competition for research funding becomes increasingly competitive, UAF's research institutes must make important strategic decisions in order to continue their work. Having an up-to-date understanding of how researchers do their work is critical in making informed decisions about supporting their information technology needs.
In October, 2010, UAF Chancellor Brian Rogers asked the Office of Information Technology (OIT) to conduct an assessment of the high performance computational needs of the University's research and academic communities.

Between October 2010 and January 2011, OIT conducted a series of interviews with UAF Institute Directors, Deans, faculty and researchers, University of Alaska Anchorage (UAA) researchers, and technology leaders from other research Universities across the nation. OIT also reviewed high performance computing at other universities and engaged three reviewers to assist with the analysis.

The goal of the study was to understand how faculty and researchers currently use technology in support of their research, how they see that changing in the future, what support structures they currently use and anticipate using, and to identify organizational and financial models that will enable the University to sustain that level of support into the future. Over the course of the analysis, over 50 research professionals were interviewed or participated in face-to-face discussions and over 170 participated in an online survey. The results indicate a strong need for support and resources for computational science within our research community. The difficult task was identifying to what level and with what financially sustainable model.

The report from the OIT HPC Advisory Committee (Appendix A) captures the current state of high performance computing at UAF and makes some key observations and targeted recommendations.

The focus on IT needs stems from the recognition that technology influences all aspects of the research process from collecting, storing, analyzing, and sharing data to communicating and collaborating with research colleagues. Across disciplines, the advancement of knowledge is becoming increasingly dependent on technology.[i]

This study came on the heels of a National Science Foundation (NSF) sponsored workshop on sustainable funding and business models for academic cyberinfrastructure. The final report and recommendations issued by this group in November 2010 provided many examples of successful business and financial models to explore. (Appendix D)

This assessment also looks at computational resources currently available and their impact on research at UAF. This includes the role that the Arctic Region Supercomputing Center (ARSC) has played in computational science at UAF over the last decade. As Department of Defense funding is reduced, it is important to celebrate the successes that ARSC has had and the role the remaining resources will play in continuing to support research at UAF

What is Computational Science?

In their "Initiative to Support Computational Science at the University of Virginia", the U.Va. Task Force on Information Technology Infrastructure characterized computational science and its relationship to research and IT. The following is excerpted with permission from that report.[ii]

What is computational science? Science is the process by which we observe the world, infer general principles that systematize those observations, and then deduce the observable consequences of those principles. These activities involve the collection and analysis of information.

Computational science is an emerging discipline and, as such, one whose definition is itself evolving. The traditional definition of computational science is that it undertakes an investigative approach to the understanding of systems (physical, biological, mechanical, etc.) through the use of mathematical models that are solved on high performance computers. Computational science grew out of necessity; the development of the modern computer was driven, in large part, by the need to solve complex equations in science and engineering. Today computational science includes not only simulation and numerical modeling, but also the powerful new ways in which data generated from experiments and observation can be manipulated and probed to gain scientific insight. These activities are often referred to as data mining and analysis. The broadest definition of computational science would encompass the development and application of problem-solving methodologies that rely on advanced computational resources, including processor, storage, network, and display technologies.

> *My research would not be possible without access to high performance computer. –UAF Researcher*

It is helpful to state what computational science is not. First, it is not computer science. Computer science is a specific discipline: the science of the computer itself. Computer science is often involved in important and ground-breaking computational science research, but there no requirement that all computational science directly involve computer science. Second, all uses of a computer in service of science are not necessarily computational science. The computer has undoubtedly become an indispensable tool in such activities as scholarly writing, literature research, production and manipulation of graphics and images, and experimental control and data acquisition. These activities are not normally regarded as computational science, even though they may require a significant amount of technical expertise. Third, computational science is not confined to supercomputing, i.e., the use of the most advanced and largest processors. The purpose of computational science is the science, and its procedures follow the usual course of scientific inquiry. Research begins with a question, a conjecture, or a *hypothesis to be tested and ends with a moment of discovery that constitutes at least a partial answer. The intermediate steps generally include framing the question in terms of a model; this most often means a determination of the equations governing the model. The scientist then looks for previously unrealized implications of this model. In computational science the governing equations must be put in terms amenable to solution by a numerical algorithm. This algorithm must then be converted to a machine-executable form through programming. This application must then be placed on a suitable hardware platform and, ideally, implemented in an appropriate and efficient manner for this platform. After testing and validation, the resulting program is used to execute a series of numerical experiments that have the potential to probe the question. The outcomes of these experiments are often themselves rather complex, requiring a comparable computational effort to analyze. Finally, out of this analysis one aims for new insights into the system under study. Computational science can be so effective in facilitating the exploration of data, that the outcomes often close the loop by assisting in the formulation of new questions for further research. Many of these research steps are familiar; they are the same steps taken in any theoretical or experimental science. But computational science requires additional procedures and new areas of expertise. Discipline-trained scientists will not necessarily possess all of the requisite knowledge and, given the pace at which the technology evolves, may find it difficult even to maintain an appropriate level of expertise. Continual advanced training and research collaborations are essential.*

National Trends

Twenty-five years ago computational science was severely hampered by limited resources. The computers that were available to the academic community were far from the state of the art, even for the era. They were located centrally and employed a timesharing model that greatly increased the difficulty of use. Supercomputers existed, but they were primarily located within government labs and agencies, employed for programmatic efforts where the recognized importance of computation could easily justify the high costs of these machines.

The importance of computational science to fundamental research was acknowledged in a significant way in 1985 with the establishment of the National Science Foundation (NSF) supercomputing centers, dedicated to making state of the art computational resources available to a broad community of investigators. Concurrently, another important trend was underway, namely the rise of the desktop computer and the Internet. During the next decade the capabilities of the supercomputers located at national facilities increased rapidly, driving the state of the art and enabling ever more complex and sophisticated modeling by leading researchers. At the same time the widen availability of powerful desktop workstations began to make computational science more ubiquitous, increasing the number of practitioners within the scientific disciplines.

The most recent decade has seen the desktop workstation become increasingly widespread and powerful. A typical processor now has substantially greater capability than that of the Cray computers installed in the first supercomputing centers 20 years ago. Similarly, the capabilities of the Internet have increased to the point that 1 Gigabit bandwidth to the desktop is not *uncommon. The facilities at the national centers have also seen remarkable improvements. The NSF centers now allocate roughly 100 million service units (SUs) per year, where the service unit is one standard processor hour. Leading research groups are routinely awarded more than one million SUs per year in support of their work, an amount of computing power substantially greater than that available to the entire national community only a short time ago.*

The rapid pace of development continues at the national level. The NSF has established the national TeraGrid, which uses a dedicated high-speed network to link together an impressive array of computational resources located throughout the country. The NSF's plan for the next five years calls for the development of petascale computing capabilities, that is, processors capable of delivering 10^15 floating point operations per second and systems that can manipulate datasets containing on the order of 10^15 bytes.[iii] Other federal agencies, e.g., DOE, NASA, and NIH, are also developing substantial computational capacity independent of the NSF initiatives. These agencies are making this investment because they recognize that ground breaking research will be produced with these facilities.

For a long time the emphasis has been on the hardware and the growth of the raw processor power. But this focus overlooks areas of equal or greater importance, such as the software that runs on these systems and the people who develop and use that software. Million-fold increases in processor speed have not been matched by commensurate improvements in these other areas. Programming tools, languages, and techniques have scarcely evolved; researchers continue to program in outdated languages or employ legacy codes that have not been updated since the punch-card era. They employ storage and analysis techniques that were adequate for megabytes of data but are completely ill-suited to the petabytes to come. While visualization techniques are well-developed as long as only (x,y) plots or two dimensional contours are required, datasets now often span an N-dimensional parameter space and may also include time dependent evolution information. Both the software and the hardware required to display and interact with such data are still relatively undeveloped. Thus the requirements for computational science include not only discipline specific knowledge, but also up-to-date knowledge in areas such as algorithms, programming, optimization, data management, and visualization; a single researcher cannot hope to master all of these.

> *We really need more human expertise in operating and developing research software. –UAA Researcher*

Dealing with these complexities will require the collaboration of many skilled people, but such people are generally not widely available, and the training to develop such people is lacking. Graduate and undergraduate curricula remained centered on the research methodologies employed when the members of the faculty were themselves educated, despite the revolutionary changes seen in the last few decades.

Twenty years ago researchers were clearly limited by the available hardware; they could easily devise projects that would exceed the capabilities of the fastest computers. Now the relation is inverted: the capabilities of the most powerful computers easily exceed the average researcher's ability to utilize them in a meaningful way.

Thus despite all of the progress of the past two decades, there remain significant concerns regarding the development of computational science. These concerns are quite clearly spelled out by the President's Information Technology Advisory Committee (PITAC) in their June 2005 report "Computational Science: Ensuring America's Competitiveness." This report details both those steps that must be taken at the national level and those required within universities and research communities. The principal finding of the PITAC report is as follows:

Computational Science is now indispensable to the solution of complex problems in every sector, from traditional science and engineering domains to such key areas as national security, public health, and economic innovation. Advances in computing and connectivity make it possible to develop computational models and capture and analyze unprecedented amounts of experimental and observational data to address problems previously deemed intractable or beyond imagination. Yet, despite the great opportunities and needs, universities and the Federal government have not effectively recognized the strategic significance of computational science in either their organizational structures or their research and educational planning. These inadequacies compromise U.S. scientific leadership, economic competitiveness and national security.[iv]

History of High Performance Computing at UAF

There are many resources available throughout the UAF campus designed to meet the needs of researchers. Some are localized within departments or Institutes. Others are centralized. A brief description of some major resources follows. This by no means is meant to be exhaustive or all inclusive, but to be illustrative of the types of resources needed and available through the campus.

Arctic Region Supercomputing Center

Since 1993, The Arctic Region Supercomputing Center has served the computational needs of many University of Alaska departments and institutes and the Department of Defense as a DoD Supercomputing Resource Center (DSRC) within the DoD's High Performance Computing Modernization Program (HPCMP). ARSC computational systems and resources include a wide range of high performance computing, storage and visualization technology.

Specialists at ARSC have provided expertise in massively parallel supercomputing, parallel vector supercomputing, code optimization, storage, networking and visualization. ARSC's relationship with the University of Alaska facilitates collaborative research opportunities for academic and government scientists. ARSC actively participates in a joint-faculty program supporting University of Alaska Fairbanks (UAF) researchers in several departments across campus. ARSC staff provide on-site training and support UAF classes throughout the year.

Recent changes in the Department of Defense funding have necessitated a downsizing of ARSC's operations. Diversification of funding and mission has enabled ARSC to secure grant funding for the PACMAN project and additional funding through an MIR grant. Currently, a core staff of 9 technical, operational and administrative staff will support the HPC resources PACMAN, Midnight and implementation of the MIR-R2 grant.

ARSC doesn't cost me money, but their software is really stale. For more advanced stuff it's really cantankerous.
–UAF Researcher

PACMAN is a Penguin Computing cluster comprised of Opteron processors. It has 2032 compute cores and a 109 terabytes (TB) Panasas version 12 file system. PACMAN is an acronym for the Pacific Area Climate Monitoring and Analysis Network. PACMAN is appropriate for solving large compute and memory intensive parallel jobs. The multi-core nodes on PACMAN make the system appropriate for both OpenMP and MPI applications. PACMAN runs Red Hat Enterprise Linux 5.5 and the Torque/Moab batch system scheduler.

Midnight is a Sun cluster comprised of Opteron processors. It has 1584 compute cores and a 68 TB Lustre file system. Midnight is appropriate for solving large compute and memory intensive parallel jobs. While it's possible for a single job to span both the X2200 and X4600 nodes, the vast majority of jobs will be run on one node type or the other. The X2200 and X4600 nodes all are available to jobs parallelized with MPI or a hybrid MPI/OpenMP model. Better performance for MPI jobs will generally be achieved on the X2200's. The X4600's are useful for large shared-memory challenges.

The ARSC was one reason I chose to work here. My research program was targeting the power available in the ARSC to apply bioinformatic methods to linguistic data to test hypotheses on the nature of the relation between Athabaskan in North America and Yeneseic in Siberia, of the relatedness between Mesoamerican and Californian Indigenous languages, and on the multiple routes of migration of the original populations of the Americas. Millions of calculations are necessary for this research which requires supercomputers
–UAF Researcher.

Midnight runs SuSE Linux Enterprise Server 9.3 and the PBS batch system scheduler. It is attached to ARSC's Academic mass storage system including robotic tape library.

Bigdipper is a Sun SPARC Enterprise T5440 Server that acts as the long term storage server for ARSC academic computational resources. Bigdipper provides permanent mass storage for user data at ARSC. The system is connected to tape drives in the STK SL8500 automated tape library. Together with the library, bigdipper has enormous capacity, theoretically more than seven petabytes (7,000 TB). This storage capacity is complemented by bigdipper's 125 TB of local disk cache, which allows users to keep large amounts of data staged from tape for long periods of time. Users have a small, quota-limited home directory and an archive directory. Other file systems are available for special purpose applications and groups.

Alaska Satellite Facility

The ASF processing system consists of 30 data nodes with dual-quad core Intel Xeon processors used to support ASF. The nodes supply data processing capability and a cloud of 1.5 petabytes (PB) of network attached storage (NAS). The nodes run CentOS 5 linux and use the xfs file system to allow for large volume size (between 40 and 90 TB). The raw data is stored in a robotic tape library, processed on one of the nodes, and the resultant products stored on NAS for search and download via the web.

Institute of Arctic Biology

Aeshna, the Spatial Ecology Lab cluster, contains 12 nodes; 7 Penguin Altus 1300 nodes each with two Opteron CPUs (14 cores), and 5 Dell PowerEdge 1435SC nodes each with two Quad core CPUs (40 cores), for a total of 54 cores (52 compute cores and 2 on the head node that hosts the clustering operating system). They are running Red Hat Linux with a custom-built clustering system written and maintained by Shawn Houston. Attached to the cluster is a Silicon Mechanics large storage unit (Storform iS707, soon upgraded to R515), with capacity for 48 TB, but currently populated with 38 TB of hard disks, 30TB formatted capacity.

This cluster is used almost exclusively for running the Terrestrial Ecosystem Model (TEM), a custom-written model that takes the existing spatial distribution of ecosystem types (e.g upland white spruce, lowland black spruce bogs) along with various climate scenarios as inputs to predict the future distribution of those ecosystem types across the landscape. A typical run might be for the entire boreal forest in North America on a 1 km grid, under current CO_2 concentration and 2x CO_2, with a model run of 200 years.

Life Science Informatics

The Life Science Informatics Primary Compute Cluster consists of a 264 CPU Core, with 1.1TB RAM, 4 node cluster with (1) 3 node IBM x3950 M2 as single SMP and (3) IBM x3850 X5. The Login Cluster is a 32 Core, 256GB RAM, 16 node cluster with (16) Penguin Altus 1300 dual Opteron CPUs. The GenePattern Cluster is a 12 CPU Core, 48GB RAM, 6 (8) node cluster, (1) G4 node as file server, (1) Xserve XRaid with 1.1TB disk, (1) G5 head node not used for computation, and (6) Xserve dual G5s.

Storage consists of the following distributed as needed to individual clusters/systems or used for disk-to-disk backup:(1) NexSAN SATABoy with 14TB of disk,(1) NexSAN STABeast with 27TB of disk, (1) FalconStor NSS HC650 and (2) DEU with 32TB disk, (1) FalconStor NSS HC620 and (2) DEU with 32TB disk.

Life Science Informatics is the Bioinformatics Core for the University of Alaska (http://biotech.inbre.alaska.edu/) with systems used by almost 300 undergraduate students, graduate students, professors, researchers, and their collaborators from all around the globe for Life Science research. There are about 100 software packages available on the Primary Compute Cluster and more on the Login Cluster. Researchers include phylogeneticists, evolutionists, geneticists, ecologists, mathematicians, psychologists, and provide easy to use web based access to our systems and storage in order to facilitate more science faster.

Department of Chemistry

The chemistry department WebMO servers are based on a Intel S5000PSL server motherboard with two 2.5 GHz quad-core Xeon processors (L5420). Each has 32 GB of 667-MHz double-buffered memory and three 1-TB hard drives, which are connected through an Adaptec RAID card as "Raid 1+hot spare". This is a high $/GB setup (1000GB/$500) but it is reliable. These are housed in an AntecTitan650 box with a total of 7 fans.

WebMO is a java-based web front-end for Ab Initio computational chemistry software. With funding from the UAF Technology Advisory Board, John Keller of this department built two servers and installed Gaussian ab initio computational chemistry software. This hardware-software combination is a good deal more powerful in terms of speed and versatility than HyperChem. HyperChem has its place as a relatively easy to use tool for learning the basics of computational molecular modeling; however it is slow carrying out ab initio and density functional theory calculations. The latter are generally more accurate, albeit slower, than molecular mechanics or semi-empirical methods.

Both Gaussian 03 and 09, as well as other packages such as NWChem, are also installed on the Midnight supercomputer at the Arctic Region Supercomputing Center (ARSC). This WebMO installation is designed for use by any students or faculty at UAF, the State of Alaska, or beyond (within reason) for teaching or research. Large scale computational projects should still be run at ARSC or other dedicated supercomputer centers.

I maintain my own cluster and manage my own storage for research and student use. This takes some of my time, but gives me a lot of flexibility: if I need to eject all student jobs, I can. –UAF Researcher

During the 2010 spring semester, over 2000 jobs were submitted to WebMO, mainly from students in Chem 623 Molecular Modeling course at UAF. The systems had a 99.99% uptime performance during that time. Calculations this year involved optimization and vibrational analysis of transition metal complexes, cyclodextrins, and Diels-Alder reaction transition states.

While this system fills a particular computational niche, it does come with a disclaimer: *This system should be considered in testing mode. It may suffer unannounced outages, including hard disk failure or network glitches, which could prevent you from connecting to the server.*

UAF Research Goals

Research is integral to UAF's goals, mission and its future. This is clearly articulated not only in the UAF 2010 Strategic Plan, but in the Vision 2017 Task Force report as well.
As a research-intensive university, UAF seeks to maintain international prominence in research and scholarship with emphasis on the circumpolar North. The schools, colleges and research institutes headquartered on the Fairbanks campus focus on baccalaureate and graduate programs, basic and applied research, and research outreach. These programs derive much of their strength from the extensive research and scholarship conducted by the faculty.[v] The goals of the research program are to:

- Increase research programs that address the Arctic and its indigenous people
- Focus on participation and outcomes associated with the International Polar Year
- Expand and improve both applied and collaborative research ventures
- Increase the proportion of students and faculty engaged in research and scholarly activities
- Document and disseminate indigenous knowledge
- Increase, promote and monitor undergraduate research opportunities, activities and accomplishments
- Provide competitive stipends for graduate assistantships and increase the number of graduate assistantships and postdoctoral fellowships
- Increase revitalized research space in existing facilities on West Ridge

In 2017 UAF will be a place where research and scholarship are fully integrated in the undergraduate and graduate experience. As core functions of UAF, research and scholarship require prominence in its vision for UAF and the highest priority for future investments. These core missions must permeate its purpose, its plans and their implementation.[vi]

The recommendations presented in Vision 2017 are based upon these basic assumptions:

- UAF should serve the people and state of Alaska by addressing the problems and needs of its community, the state and nation.
- UAF should build upon its unique strengths, expertise and location to address local, national and international problems and needs of the state and nation for which the university is particularly suited.

- UAF scholarship and research should focus on areas that will yield economic benefits and address Alaska's and the nation's intellectual needs while being mindful of our responsibility to also address issues of global importance.

Comprising a strategy for UAF's future, the recommendations outlined in Vision 2017 link assessment, investment and measurement by evaluating its current abilities and demand for services, and the need to invest in tomorrow's infrastructure and measure performance against expectations and need. The world is changing and UAF must be responsive to the needs of society in the areas where it has recognized expertise and can make valuable contributions.

Just as the world is changing, so is the nature of research. Computational science has become the third pillar of science complementing experimental and theoretical science. This has been widely cited in scientific literature and acknowledged in Congressional testimony and Federal reports.[vii] High Performance Computing (HPC) continues to become an increasingly critical resource for an expanding spectrum of research disciplines.[viii] Therefore it is vital that UAF strategically address and plan for computational science resources and infrastructure to support this trend in the context of its research goals.

> *Any university that is serious about biological research must provide computational support that is not itemized by the nickel and dime on grants. Funding agencies like NSF consider this part of core facilities that should come out of overhead and/or state funds, just like heat and light. –UAF Researcher*

OIT Interview and Community Engagement Results

Over the course of 8 weeks, OIT conducted a series of interviews with Deans, Directors, key stakeholders and HPC leaders from other Universities across the country. In addition, OIT conducted online survey of faculty, researchers and administrators from across the UA system.The survey contained a series of questions, both multiple choice and free response that prompted participants to identify their current and future computing needs as it relates to their research. Respondents were well distributed across schools, colleges and institutes and across disciplines. In total 187 individuals participated in the survey: 172 from UAF, 11 from

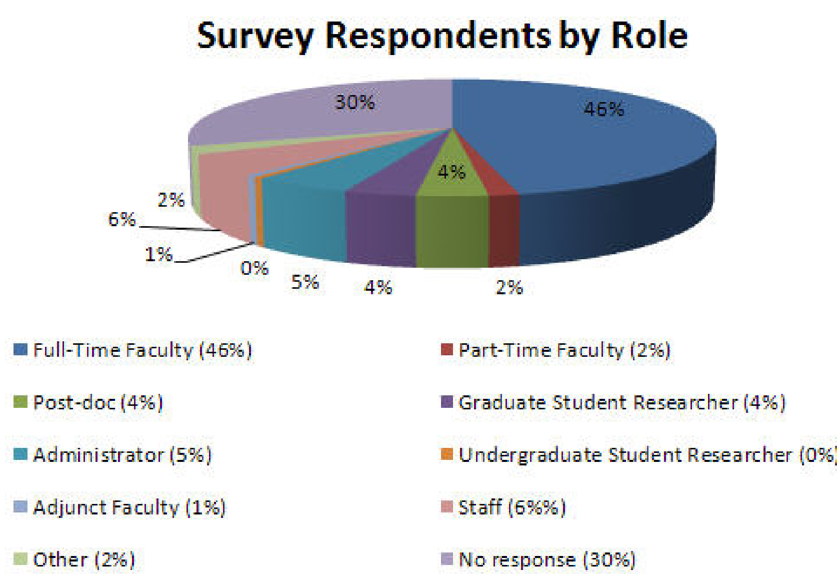

UAA, 2 from UAS, and 2 not affiliated. (see Appendix B for complete survey results)
Some interesting facts, trends and analysis were identified and fall into three main categories: Hardware/software, storage, and people.

Hardware/software

Respondents were asked to identify how their current computational research requirements were being met. Each was able to select more than one option if applicable. Overwhelmingly, ninety-seven percent (97%) of respondents utilize a desktop computer and seventy-four percent (74%) use some sort of high performance computing resources (departmental cluster, national labs, or ARSC.)

> *My research is especially focused upon modeling wildfire smoke, volcanic ash, and ice fog. About 300,000 CPU hours per year are used for air quality and ice fog modeling.*
> *–UAF Researcher*

Those that utilize high performance computing resources expressed the need for local hardware resources, indicating that granting agencies do not look favorably toward funding hardware resources when national centers are available, but having those resources locally gives a competitive advantage when applying for grants.

Another concern expressed was software licensing. Much of the software required for modeling and visualization is specialized by field and is often expensive. This limits the number of licenses that departments and institutes can afford and thus its widespread use. Many individuals stated that it would be desirable to have a centrally funded suite of software that would be available for both research and instruction.

Storage

As computational needs were discussed with interviewees, it became apparent that in addition to computing hardware and software, there was a tremendous need for electronic storage. Most researchers and directors saw this need growing as quantities of and complexity of research data, modeling results and visualization output increased. Additionally, new NSF requirements for data management and access have necessitated the need to focus on accessible storage and portals for access to data.

> *The major needs in my lab revolve around the storage and analysis of very large DNA sequence datasets. The greatest need is for customized support to provide scripting, install complicated software, and provide web-based solutions.*
> *–UAF Researcher*

As the chart below depicts, both short term and long term storage needs are anticipated to increase over time. The HPC consultants that were engaged as part of this study characterized UAF as “data rich” and recommended that UAF take advantage of its vast wealth in Arctic data by aggregating data across disciplines and making it more accessible to researchers around the world.

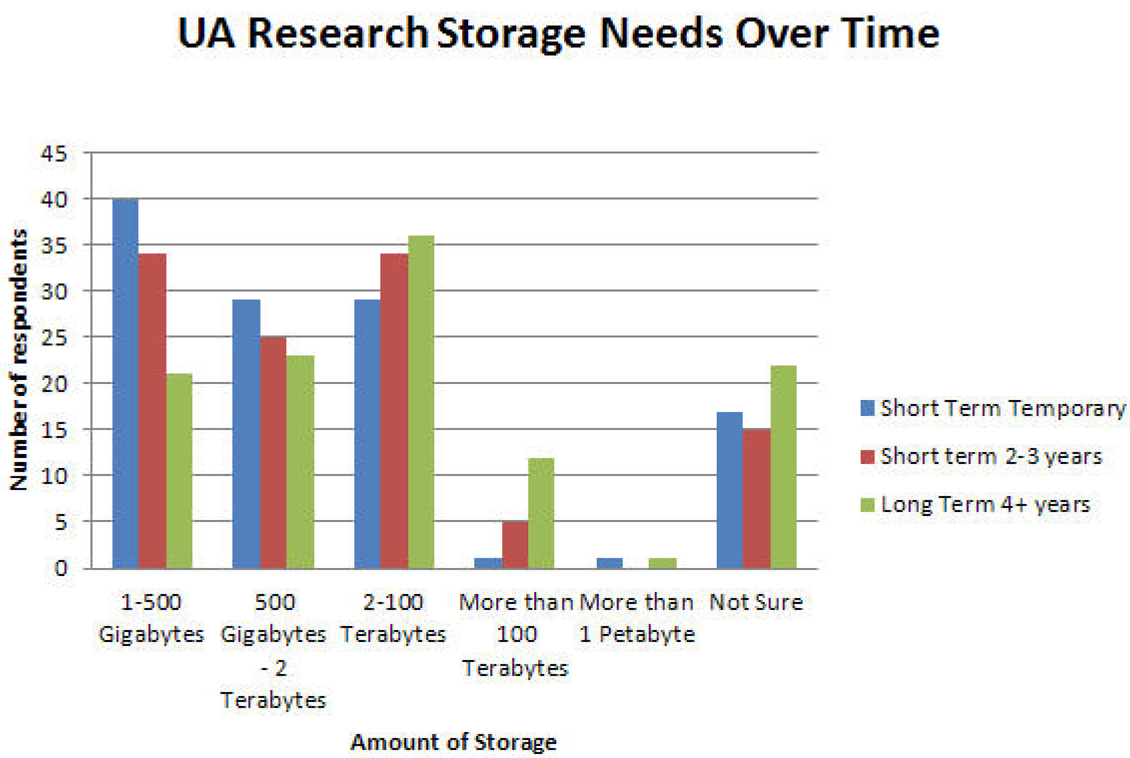

People:

A message that was repeated again and again through interviews and survey narratives was the importance of people. Many researchers indicated that support for bridging science and computation was very important to have as a local resource. The vast majority commented that given a choice between having local hardware or local support, they would choose local support. In interviews, the research community stressed that vital to their success was the PhD level computational scientists that were able to assist them in coding, compiling, paralleling and debugging of their code.

> *Hardware is important. Far more important are people serving computation and software development.*
> *–UAF Researcher*

While a few institutes had computational expertise within their units, much of UAF’s computational science support came from the scientists at ARSC. In the current ARSC reorganization, seven such support staff and scientists remain, but are currently term funded by grant resources. Concern exists over the sustainability of this support beyond the current grant cycle.

Consultant Observations

As part of this study, OIT interviewed HPC leaders at many other Universities and invited three onsite to assist in the recommendation process. Amy Apon, Director of the Arkansas Supercomputing Center, Gerry McCartney, Oesterle Professor and CIO for Purdue University and Michael McPherson, Vice-Provost and Deputy CIO at the University of Virginia visited UAF January 12-14, 2011 and engaged a wide range of stakeholders. (See Appendix C)

The consultants observed that withdrawal of DoD funding for ARSC is perceived as more of a crisis by researchers than it really is. There may be a need to improve the communication among stakeholders about the status of needs and resources. It appears to the consultants that the current and planned (and already funded) build-out of the NSF-funded PACMAN cluster and the additional MRI-R2 grant should meet all current and short term UAF HPC needs. The current newest HPC system is less than a year old, and plans are in place for expansion of resources in the next year. With a typical lifetime of an HPC system of three to five years, these resources provide a foundation and buffer of time for planning additional acquisitions. Faculty who have grants or have submitted proposals based on available compute resources at ARSC do not appear to be in danger of losing the necessary resources. Likewise, with the interim funding made available by UAF to retain a core of ARSC technical staff, operations and basic consultation on the existing systems appears to be secure for now.

People, especially PhD-level scientists with the specialized HPC skills to partner with faculty to effectively leverage computational techniques in their research, are the most critical component of an overall HPC strategy. UAF has taken action to retain seven such individuals from ARSC on soft money through at least the end of the calendar year.

Executive leadership at the University is engaged and committed to taking action to preserve critical HPC capabilities and to implement a plan to take Alaska into the future.
The University of Alaska has a unique proposition as a potential partner for institutions outside of the arctic. Alaska's physical location in the arctic, the experience of its faculty, students, and staff in working in arctic environments, and its deep expertise in the many disciplines related to the arctic make Alaska a very attractive partner for other universities with established or developing programs of arctic study.

> *I work with scientists in Europe and the lower 48 in creating and analyzing large genomics datasets. –UAA Researcher*

Nationally, ARSC is a well-respected organization with a brand that may be worth preserving. Locally, the reputation of ARSC is also quite good. When weighing the choice of keeping the name or rebranding, it should be kept in mind that there is some local baggage attached to the name ARSC; it conjures a particular kind of HPC support, one in which compute cycles are the scarce resource. If Alaska chooses to broaden its efforts beyond just proving cycles then rebranding may be advisable to help reinforce the change.

The generation (through simulation) and collection (through remote sensing and other automated technologies) of massive data sets, and the associated data intensive computing, are a hallmark of modern HPC. Alaska has an advantage should it choose to focus on data-intensive computing due to the presence in ARSC of experts in data management currently working on the DOD data management project.

HPC at a research university is a layered infrastructure. The university needs to maintain a core of central HPC compute resources, both clusters and high throughput Condor-style grids, to handle funded projects that need to buy services, prefunding exploratory work, unfunded research (to the extent the university chooses to support it), and research data management. Some number of researchers will have needs sufficiently specialized or intense that it will make sense for them to buy and operate their own HPC equipment, but that should be a minority of researchers. Finally, any coordinated approach needs to leverage the extensive and very high-end resources available at little or no cost through the TeraGrid and the national labs.

Discussion of Financial and Organizational models

The NSF Sponsored workshop on Sustainable Funding and Business Models for Academic Cyberinfrastructure Facilities outlined in their findings key factors of successful and sustainable financial and organizational models that UAF should consider. The report tackles many of the same issues the University of Alaska faces: centralization (or not), funding, staffing, organization and governance. (see Appendix D for complete report)

Financial Models

Funding large scale computing centers can be difficult. These facilities are often expensive to operate, especially so in Alaska where energy expenses can be 3-4 times that of comparable facilities in the Lower 48. External agencies often do not want to contribute costs to central computing facilities that are not directly within teaching or research buildings. In addition, many Universities perceive computing facilities as an expense while they view libraries and classroom labs as necessary core infrastructure. As one UAF researcher put it, "*The University has no problem funding a chemistry or engineering lab. Well, my lab is the computer*," referring to research that requires high performance computing for modeling and visualization.

Three types of funding models emerge: campus-funded, externally-funded, and recovery models. A summary overview of these models from the NSF report is included below.

Campus-Funded

Some institutions fund high performance computing facilities, resources and services completely or at a very high level. These institutions view high performance computing as a necessary and critical part of campus infrastructure much like administrative IT, the library, and networking. Some of these institutions fund HPC entirely from core internal budget, or with indirect funds from research grants.

Benefits:

Efficiency-Base funding for HPC reduces individual department costs by eliminating the need to build, support and maintain their own resources.
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 Campus-Funded:

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?

Externally-Funded

Institutions that receive much of their HPC 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.

Benefits of Externally-Funded

Efficiency – To provide extreme scale resources intended to support select world-class research and enhance competitiveness.
Innovation – By pushing the limits of computational scale and performance, these centers produce innovations in software, tools, and affiliated technologies
National Competiveness – 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 Exterally-Funded

Funding – During economic downturns, federal and especially state support funding (e.g., legislative line items) is limited and therefore competition is much higher. Funding can be eliminated as recent ARSC experience shows.
Sustainability – How do institutions that rely heavily on externally-funded projects sustain their staff expertise beyond their center's immediate funding horizon?

Cost Recovery (recharge) Model

The ability and willingness of research teams to pay for centralized HPC computational resources or staff consulting services are important factors to consider in deciding whether to move to a cost recovery model.

Benefits of 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

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.

Economies of Scale, Scope and Cost Sharing – By contributing research funds toward well-run HPC facility resources and professional services, the whole is greater than the sum of the parts.

Challenges of Cost Recovery Model

Demand and Resistance – Cost recovery models assume researcher support and demand for HPC facility resources and services, as well as an ability to pay.

Innovation – One concern is that a HPC facility operated in a pure service mode will fall behind the technology curve and lose value to its researchers.

Many institutions that participated in the workshop 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. In addition, many were working towards a centralized model of HPC support to leverage facilities, staffing and economies of scale.

In our interviews with other Universities, the funding trend was between 20-50% institutional base funding with the remaining costs being funded through a combination of external grants and cost recovery (recharge).

Organizational Models

There are numerous examples of successful HPC organizations at other research universities. These organizational models tend to fall into three basic approaches: a reporting structure through the central IT department, through the VP for Research or a combination of both.

Central IT Model

In organizations where high performance computing reports through the central IT unit, there tends to be a singular focus on operations and maintenance of equipment, software and support. Scientific computation tends to be embedded within the schools or colleges and therefore not part of the IT structure.

Research Unit Model

Universities with HPC that reports through the Vice Chancellor for research tended to be highly focused on both computational science and on equipment. They had large staffs both scientific and technical in order to maintain support for research. This is the model that ARSC operated under.

Shared Responsibility Model

The vast majority of universities interviewed had a split organizational model. The central IT unit was responsible for maintaining HPC hardware, software, network and backups. The research unit was responsible for providing a core of PhD level computational scientists that were shared amongst the research institutes. These scientists provided support for coding, debugging, modeling and visualization. These universities also had HPC advisory councils that governed the HPC resources, setting operational windows, cycle allocations and upgrade paths.

Just having a central system for IT consultation, design, and implementation would be helpful. Right now we are forced to do much of the design independently of OIT, then ask them to help us with the problems. –UAF Researcher

Universities with the shared model had close working relationships with their central IT organizations and leveraged the economies of scale that each could bring to the relationship.

Metrics of Success and Return on Investment

Given fiscal pressures facing the University and increasing accountability, it will be strategically important that the University establishes key metrics for evaluating the success of its HPC program. Justifying technology and staff expenditures is essential as IT services of any kind are areas where institutions often look to make cuts.

The NSF workshop on Sustainable Funding and Business Models for Academic Cyberinfrastructure articulated three key areas in evaluating metrics: quantitative, qualitative and new challenges. (Appendix D)

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 HPC 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 HPC 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 HPC 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 HPC facility. The depth of a HPC 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.[ix]

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 HPC 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.

New Challenges

NSF 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 HPC 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 HPC 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 HPC must be effectively communicated to administrators, many of whom are not computational scientists.

These metrics should be clearly communicated to the user community and the University at large. The University of Buffalo has created a data portal called Metrics on Demand (UBMoD) UBMoD is an open source data warehouse and web portal for mining statistical data from resource managers commonly found in high-performance computing environments. It was developed at the Center for Computational Research at the University at Buffalo, SUNY.

UBMoD presents resource utilization (CPU cycles consumed, total jobs, average wait time, etc) for individual users, research groups, departments, and decanal units. A role-based authentication scheme can be set up to display data in the format best suited to the end-user (for example, as a student, research group leader, department chair, dean, etc.). The interface provides a dashboard overview of resource consumption along with fine-grained control over the time period and resources that are displayed. The information is presented in easy to understand interactive charts, graphs, and tables.[x]

Recommendations

These recommendations are based on interviews, surveys, consultant visits and examination of other university model. UAF should consider taking the following near term, midterm and long term actions as they relate to support of high performance computing for UAF and the UA system.

Near Term Recommendations (0-12 Months)

UAF should immediately begin a broad communications campaign designed to communicate clearly to all interested members of the University community the current state of affairs, the commitment to excellence in HPC being made by the University, the plans for the near term, and how people can get involved and remain informed as long-term plans are developed.

UAF should articulate the role of computational science/high performance computing in its mission as a research University. This should be reflected in UAF Vision 2017 and UAF's strategic academic and research master plans. If it is important, the University should say so.

UAF should refocus ARSC on a carefully chosen set of HPC services and partnerships that focus on the competitive strengths that Alaska enjoys, and place it within the organizational structure of UAF. Based on interviews, consultant recommendations and examination of successful HPC models at other institutions, we recommend a shared responsibility model. The science and the PhD level computational support should report to the Vice Chancellor for Research. This would include the Chief Scientist and seven (7) support staff (including PhD level computational scientists) currently on soft money that supports HPC. The technical and operational support should lie with the central IT organizations (Office of Information Technology). This would include one (1) network specialist, three (3) systems/programming specialists, two (2) storage experts, and two (2) user services experts. This organizational structure leverages the operational expertise and vendor relationships that exist in IT and enables the scientific research community to focus on research and computation.

There is one position with fiscal expertise. We recommend this individual be integrated into either the VC for Research or the OIT business office. Integrating into an existing business office will both provide backup support and will allow this individual to provide additional support for developing grants and building relationships between researchers and their counterparts at other institutions.

UAF should maintain HPC technical and operational staff at least at the present levels in order to provide support for the existing and planned systems and transition them to the Office of Information Technology over the next four to six months. This will leverage current expertise, allow for cross training and provide greater coverage and support for HPC resources.

UAF should maintain and support high performance computing at a base level of 30% of the operational costs and adopt a sustainability model to continue. The first transitional year will require a higher initial level of support. This percentage was determined based on average funding levels of other universities which ranged from 20%-50% central base funding.

UAF should seek commitment by all Institute Directors and Deans to support (through actions and finance) high performance computing for research and academics.

UAF should seek to maintain its presence as the statewide center for high performance computing, ensuring engagement with UAA's research community, UAS, the State of Alaska, other research organizations in the State (e.g. BASC), and establish relationships that are mutually beneficial.

UAF should establish and support a strong and broadly-represented faculty governance group to direct investment in high performance computing over time. This group would assist the Office of the Vice Chancellor of Research and the Office of Information Technology in recommending strategic direction, identifying new requirements, promoting the requirements of researchers and providing input on allocation and policy decisions. This should be a system wide group representing all UA Campuses.

UAF should implement a plan to fund, recruit and retain PhD-level scientists with specialized HPC skills to partner with faculty on their research. Such a plan must include provisions for consultation with faculty during proposal preparation, buyout of fractions of the HPC specialists' time over extended periods when grants are funded, and bridging the HPC specialists between grants.

UAF should begin to identify a funding stream to move a core of these PhD level scientists from soft funding into the central research unit to be shared by all needing computational science support. The technical staff should continue to be funded centrally.

UAF should decommission the Midnight cluster. It is at the end of its useful service life and provides little or no capacity not available more efficiently on the new PACMAN cluster. The consultants and ARSC Staff indicate all non DoD work now done on Midnight can be handled by PACMAN and the MRI hardware. Near and long term needs and expectations of the University related to the costs of power for HPC systems should be discussed openly as plans are made for new acquisitions for end-of-life systems. While this system is owned outright, UAF must carefully consider whether any revenues obtained from the potential sale of compute cycles outside UAF is cost effective given the operational costs.

Medium Term Recommendations (1-3 Years)

UAF should change the way it recruits and hires scientists. It should seek scientists with skills in numerical, quantitative, computational science. This will serve to broaden the computational base and bring expertise to the research units in which they reside.

UAF should seek the partnership of other universities, research institutes and national labs with compatible research interests and commitment to high performance computing.

UAF should seek external funding to maintain, upgrade and replace the PACMAN cluster and related HPC infrastructure as they age.

UAF should pursue additional network capacity connecting Alaska with the rest of the world. As success in HPC grows, individual projects will require dedicated network resources beyond the Universities existing capacity.

Long Term Recommendations (3+ Years)

UAF should continue to leverage grant resources such as EPSCOR RII to fund and sustain HPC infrastructure.

UAF should pursue long-term center support through grants such as NSF Industry & University Cooperative Research Program (I/UCRC) and other forms of industrial collaboration.

UAF should continue to support PhD level scientists with specialized HPC skills over the long term.

Summary

It is clear from the results of this assessment that support for computational science at UA is essential in fulfilling the academic and research missions the University. It is equally clear that the most important components of that support are the human resources which provide PhD level science and programming support bridging science and computation, modeling and visualization.

There is overwhelming desire for UAF to continue to maintain and fund HPC resources at a level that will continue to support existing research and continue to provide faculty and researchers with a competitive advantage when applying for grants.

There are pros and cons of centralized versus distributed HPC and what users expressed is desire for transparent, easy to access, easy to use computational resources with clear, consistent procedures for acquisition regardless of where those resources lie. Establishing a system that is inclusive of those principles will guarantee satisfaction and increase participation in a centralized HPC resource.

Adoption of the recommendations contained within this report should lay the foundation for a solid, reliable, sustainable high performance computing resource to support both research and academic high performance computing.

These recommendations leverage current resources in the most economic manner available. They take advantage of existing expertise and focus technical and research operations within appropriate University structures. This provides greater reliability and support base within in each unit.

Implementation will require close collaboration between central IT operations, the research community and the HPC faculty governance group. This will bring positive cultural change to both operations.

It is important that the University of Alaska recognize that high performance computing is a system resource. HPC equipment and infrastructure are expensive to implement and support. Duplication of service within UA does not make sense. Collaboration with the other main campuses and their faculty are key in building a strong and sustainable HPC platform. Through the implementation of these recommendations and through the establishment of a transparent, reliable HPC organizational and funding model, greater participation and increased sustainability will result throughout the UA system.

Endnotes

[i] Lane, Cara, Jancie Fournier, and Tom Lewis, with Gred Koester, Ed Lazowska, Erik Lundberg, Henry Lyle III, and William Washington. " Scientific Advances and Information Technology: Meeting Researchers' Needs in a New Era of Discover" (Case Study 3, 2010). Boulder, CO: Educause Center for Applied Research, 2010, available from http://www.educause.edu/ecar

[ii] University of Virginia, Task Force on Information Technology Infrastructure Supporting Research in Science and Engineering, (2006) "An Initiative to Support Computational Science at the University of Virginia." Retreived from: http://www.virginia.edu/vpcio/documents/ITRTF_report_9JULY06.pdf

[iii] National Science Foundation. (2007) Cyberinfrastrucure Vison for 21st Century Discovery. Retrieved from: http://www.nsf.gov/pubs/2007/nsf0728/nsf0728.pdf

[iv] President's Information Technology Advisory Committee (PITAC) Report,"Computational Science: Ensuring America's Competitiveness," June 2005, retrieved from: http://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf

[v] University of Alaska Fairbanks. (2006). University of Alaska Fairbanks Strategic Plan 2010. Retrieved from: http://www.uaf.edu/strategic/F_StrategicPlan06-2_2.pdf

[vi] University of Alaska Fairbanks (2008). UAF Vision 2017 Task Force Report: Shaping the future of UAF. Retrieved from: http://www.uaf.edu/vision2017/final/VTF_finalreport.pdf

[vii] 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: ensuring America's competitiveness.* (2005) pgs 12-14. Retrieved from: http://www.nitrd.gov/pitac/reports/index.html

[viii] National Science Foundation. (2010). Sustainable Funding and Business Models for Academic Cyberinfrastructure Facilities. http://www.cac.cornell.edu/~lifka/Downloads/SRCC/NSF-sponsored-Sustainable-Cyberinfrastructure-Report.pdf

[ix] 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 .

[x] University of Buffalo, Metrics on Demand. http://ubmod.sourceforge.net/

Appendix A

Report from the OIT HPC Advisory Committee

REPORT FROM THE OIT HPC ADVISORY COMMITTEE

Date: 29 January 2011
From: Gerry McCartney, Oesterle Professor of IT and University CIO, Purdue University
Michael R. McPherson, Associate VP and Deputy CIO, University of Virginia
Amy Apon, Director, Arkansas HPC Center, Professor, University of Arkansas
To: Steve Smith, Chief Information Technology Officer, University of Alaska

EXECUTIVE OVERVIEW

The University of Alaska is presented with both a challenge and a unique opportunity. Although the pain of extensive layoffs at the Arctic Region Supercomputing Center (ARSC) is very real and undeniable, the sudden departure of the majority funding partner provides the opportunity to remake the high-performance computing (HPC) environment for the entire Alaska system with relatively few constraints. The University has a faculty and technology staff who are experienced with the power of HPC, a leadership that is committed to maintaining excellence in this critical area, a core of people and facilities retained from ARSC that form a solid foundation for a new plan, and a unique position and expertise in the arctic that makes it an attractive partner for discipline-based HPC collaborations. In our opinion, the University of Alaska is well-positioned to leverage the past successes of ARSC to create a new HPC strategy that is sustainable and that meets the needs of the System in the present and in the future.

OBSERVATIONS

From our study of materials provided to us and conversations with leadership, faculty, and staff on-site, we make the following observations:

- Withdrawal of DOD funding for ARSC is perceived as more of a crisis by researchers than it really is. There may be a need to improve the communication among stakeholders about the status of needs and resources. It appears to us that the current and planned (and already funded) build-out of the NSF-funded Pacman cluster and the additional MRI grant should meet all current and short term UAF HPC needs. The current newest HPC system is less than a year old, and plans are in place for expansion of resources in the next year. With a typical lifetime of an HPC system of three to five years, these resources provide a foundation and buffer of time for planning additional acquisitions. Faculty who have grants or have submitted proposals based on available compute resources at ARSC do not appear to be in danger of losing the necessary resources. Likewise, with the interim funding made available by UAF to retain a core of ARSC technical staff, operations and basic consultation on the existing systems appears to be secure for now.

- People, especially PhD-level scientists with the specialized HPC skills to partner with faculty to effectively leverage computational techniques in their research, are the most critical component of an overall HPC strategy. Alaska has taken action to retain seven such individuals from ARSC on soft money through at least the end of the calendar year
- Executive leadership at the University is engaged and committed to taking action to preserve critical HPC capabilities and to implement a plan to take Alaska into the future.
- The University of Alaska has a unique proposition as a potential partner for institutions outside of the arctic. Alaska's physical location in the arctic, the experience of its faculty, students, and staff in working in arctic environments, and its deep expertise in the many disciplines related to the arctic make Alaska a very attractive partner for other universities with established or developing programs of arctic study.
- Nationally, ARSC is a well-respected organization with a brand worth preserving. Locally, the reputation of ARSC is also quite good. When weighing the choice of keeping the name or rebranding, it should be kept in mind that there is some local baggage attached to the name ARSC; it conjurs a particular kind of HPC support, one in which compute cycles are the scarce resource. If Alaska chooses to broaden its efforts beyond just proving cycles then rebranding may be advisable to help reinforce the change.
- The generation (through simulation) and collection (through remote sensing and other automated technologies) of massive data sets, and the associated data-intensive computing, are a hallmark of modern HPC. Alaska has an advantage should it choose to focus on data-intensive computing due to the presence in ARSC of experts in data management currently working on the DOD data management project.
- HPC at a research university is a layered infrastructure. The university needs to maintain a core of central HPC compute resources (both clusters and high-throughput Condor-style grids) to handle funded projects that need to buy services, pre-funding exploratory work, unfunded research (to the extent the university chooses to support it), and research data management. Some number of researchers will have needs sufficiently specialized or intense that it will make sense for them to buy and operate their own HPC equipment, but that should be a minority of researchers. Finally, any coordinated approach needs to leverage the extensive and very high-end resources available at little or no cost through the TeraGrid and the national labs.

NEAR TERM RECOMMENDATIONS (0-12 MONTHS)

1. Choose the basic approach the System wishes to pursue, one that best fits with the culture of the System. There is no single model that works everywhere. We present options here, and recommend Option #2a or #2b, based on what we heard and on our experience.

   #1. Continuation of ARSC as a standalone supercomputing center, but rather funded entirely or primarily by the University instead of with the previous DOD contract. We believe it will be more cost-effective to leverage the resources of the university in areas of staffing, expertise, and physical facility, and that a closer integration of the center with the University organizational structure will be a better solution.

   #2. Refocus ARSC on a carefully-chosen set of HPC services and partnerships that focus on the competitive strengths that Alaska enjoys, and place it within the organizational structure of UAF. Here the committee offers two options that can be effective and acknowledge that both of these are successful in many academic HPC centers. The selection of these will depend on local UAF culture and preferences.

   #2a. The science portion of the HPC organization logically reports to the Vice Chancellor for Research. It is possible to place the operational responsibility for the HPC organization with the Chief Information Technology Officer. This organizational structure leverages the operational expertise and vendor relationships that exist in IT. While the staff would enjoy the community of shared expertise, we recommend that HPC operational staff be maintained at least at the present levels in order to provide support for the existing and planned systems.

   #2b. The HPC organization, including both the science portion and the operational responsibility, reports to the Vice Chancellor for Research. The model recognizes the large difference in culture between operational support for HPC and operational support for enterprise IT systems. The user community of HPC systems is typically much smaller than the user community of the university's enterprise systems, and many users are much more familiar with the actual hardware configuration and optimal use of the system. Because of the often leading edge nature of HPC systems, even with systems that are designated as "production" as compared to "experimental", academic HPC systems may be much more fragile than enterprise system. In a successful center a sophisticated user community works in partnership with the staff to keep the systems running well. Unlike with traditional enterprise computing, in an academic HPC center, users of HPC systems frequently communicate with HPC operational staff to help identify hardware errors, to help identify bugs in the software stack that may include bugs in commercial as well as open-source production infrastructure software, and to get help on how to optimally use the system. Users may also suggest helpful operational policies in the use of the file systems, schedulers, and allocations. Depending on the local culture at UAF, this type of close relationship between HPC users and HPC operational staff may be best supported under the Vice Chancellor for Research.

Whether #2a or #2b is selected, the success of HPC will depend on a strong partnership between IT and the Vice Chancellor for Research, and University resources such as data center facilities, IT expertise for such things as campus networking, security, and vendor relationships will need to be leveraged by the HPC center.

2. Develop and implement a long-term funding model for HPC that fits with the culture and business practices of the University of Alaska. Such a funding model should include a plan for ongoing demonstration of the return on investment in HPC. As with organizational structure, there is no single funding model that will work everywhere, and a successful sustainable funding model will need to include a suite of funding sources and strategies. Sources of funding may include:

   - State support (legislative, line-item)
   - Campus research office line item
   - Campus research office grant overhead, or through startup and match funds
   - Office of Information Technology line item
   - Single project support from agencies, including NSF, DOD, ...
   - Long-term center grants from agencies, such as NSF Science and Technology Center grants, or I/UCRC grants, and others
   - Collaborative grants with faculty, as PI, as co-PI, and including line items for staff or equipment in faculty-led grants with no senior role
   - Industry collaborations
   - NSF EPSCoR RII grants

   For an excellent discussion of the possibilities and trade-offs of both structural and funding models see the report of the NSF-sponsored workshop "Sustainable Funding and Business Models for Academic Cyberinfrastructure Facilities” at http://www.cac.cornell.edu/SRCC/.

3. Immediately begin a broad communications campaign designed to communicate clearly to all interested members of the University community the current state of affairs, the commitment to excellence in HPC being made by the University, the

plans for the near-term, and how people can get involved and remain informed as long-term plans are developed.

4. Continue to maintain and support a strong and broadly-representative faculty governance activity to direct investment in HPC over the long term, such as the already established User Advocacy Group, as described in the MRI-R2 proposal.

5. Implement a plan to fund, recruit, and retain PhD-level scientists with specialized HPC skills to partner with faculty on their research. Such a plan must include provisions for consultation with faculty during proposal preparation, buyout of fractions of the HPC specialists' time over extended periods when grants are funded, and bridging the HPC specialists between grants.

6. Decommission the Midnight cluster. The Midnight cluster is at the end of its useful service life and provides little or no capability not available more efficiently with the new Pacman cluster. The cost of operating Midnight ($500K or more per year in electricity alone) can be much better directed to other priorities. Power is expensive in Alaska. The near and long-term needs and expectations of the University related to the cost of power for the HPC systems should be discussed openly as plans are made for new acquisitions and end-of-life for systems.

7. Investigate tuning of building, campus, System, and external networks to facilitate large data flows and address faculty concerns about network performance as a limiting factor in their HPC work.

MEDIUM TERM RECOMMENDATIONS (1-3 YEARS)

1. Seek a handful of partners (other universities, research institutes, national labs, etc.) with compatible research interests and commitment to HPC. The opportunity exists for partners to specialize in areas of natural advantage (for example, Alaska specializes in curation of massive data and in hands-on arctic work; a partner with deep expertise in HPC operations and favorable electric rates specializes in provision of compute cycles). The relationship with universities in Hawaii through the Pacman project is an excellent example of the type of collaboration that is possible.

2. Seek external funding to maintain, upgrade, and replace the Pacman cluster and related HPC infrastructure as they age. The current MRI grant with PI Newby is funded at a level that indicates high regard by NSF. We think that Alaska's chances of maintaining this level of campus infrastructure through grant funding such as MRI are good.

3. Pursue additional network capacity connecting Alaska with the rest of the world, particularly access to either owned fiber, IRUs, or carrier pricing models that allow

the System access to additional dedicated lambdas (10Gbps or higher) at or near marginal cost. As success in HPC grows, you will likely find that individual projects will require dedicated network resources beyond your existing capacity.

LONG TERM RECOMMENDATIONS (3+ YEARS)

1. The State of Alaska is to be commended for the foresight in leveraging the existing EPSCoR RII grant for the acquisition of HPC resources. Over the long term this source of funding can be a key component of sustainable HPC infrastructure.
2. Consider pursuing long-term center support through grants such as NSF I/UCRC, and with other forms of industrial collaboration.
3. As stated in the observations, PhD-level scientists with the specialized HPC skills to partner with faculty to effectively leverage computational techniques in their research are critical to the success of the HPC strategy. These staff should be supported over the long term, and increased, if at all possible.

Appendix B

High Performance Academic Research Computing Survey

UAF Research Computing Needs Survey

Survey Name: Research Computing Resources Survey
Response Status: Partial & Completed
Filter: None
Feb 09, 2011 7:50:11 PM

1. What is your primary role?

	Number of Response(s)	Response Ratio
Full-time faculty	80	47.0%
Part-time faculty	3	1.7%
Post-doc	6	3.5%
Graduate Student Researcher	7	4.1%
Administration (Dean, Director, etc)	8	4.7%
Undergraduate Student Researcher	1	<1%
Adjunct Faculty	1	<1%
Staff	7	4.1%
Other	3	1.7%
No Responses	54	31.7%
Total	170	100%
16 Comment(s)		

2. Please select your School, College, Institute Affiliation. (You may select more than one)

	Number of Response(s)	Response Ratio
College of Engineering & Mines	14	12.2%
College of Liberal Arts	26	22.8%
College of Natural Sciences & Mathematics	36	31.5%

College of Rural & Community Development	3	2.6%
Graduate School	1	<1%
School of Natural Resources & Agricultural Sciences	3	2.6%
School of Education	1	<1%
School of Fisheries & Ocean Sciences	9	7.8%
School of Management	2	1.7%
Arctic Region Supercomputing Center	3	2.6%
Geophysical Institute	18	15.7%
Institute of Arctic Biology	5	4.3%
International Arctic Research Center	9	7.8%
Institute of Marine Science	3	2.6%
Institute of Northern Engineering	6	5.2%
Museum of the North	3	2.6%
Rasmuson Library	1	<1%
Alaska Satellite Facility	0	0.0%
Other	4	3.5%
Total	114	100%
2 Comment(s)		

3. How are your computational research requirements currently being met?

	Number of Response(s)	Response Ratio
Desktop computer (local personal computer, workstation)	101	88.5%
Faculty/Research Group Cluster	10	8.7%
Department-owned cluster	8	7.0%
ARSC	48	42.1%
LSI/Biosciences Cluster	7	6.1%
National Systems (National Labs, TeraGrid)	6	5.2%
Open Science Grid	0	0.0%
N/A	0	0.0%
Other	6	5.2%
Total	114	100%
16 Comment(s)		

4. If your computational research needs are not being met, please specify the unmet needs now and into the future (RAM, CPU, bandwidth, etc requirements) What is the driving need for this growth? Please enter "N/A" if this question is not relevant to your research needs.

44 Response(s)

5. What operating system do you typically use for your scientific research? (check all that apply)

	Number of Response(s)	Response Ratio
Linux (any flavor, SUSE, RedHat, SLES, CNK/SLES, etc)	56	49.1%
Windows	72	63.1%
OSX (Apple)	51	44.7%
BSD/BSDi (Berkeley)	1	<1%
UNIX/AIX	19	16.6%
Cell OS	1	<1%
UNICOS/lc	0	0.0%
Windows HPC	1	<1%
CentOS	3	2.6%
Solaris/ Open Solaris	3	2.6%
in-house OS, programmed own	3	2.6%
Other	1	<1%
Total	114	100%

7 Comment(s)

6. If computing services were provided on campus, which of the following would enhance your research productivity? (Select all that apply.)

	Number of Response(s)	Response Ratio
Support to locate department/researcher purchased servers in a central	33	30.5%
Support to run a department/researcher purchased cluster in a central	17	15.7%
Provide "Condo Cluster" services- you purchase cluster and add to other	18	16.6%
Provide "Computing Cycles" (campus procures a large cluster and allocates	32	29.6%
Commercial offerings/outsourcing	2	1.8%
on-demand, large memory, data-intensive supercomputing	32	29.6%
N/A	28	25.9%
Other	7	6.4%
Total	108	100%
16 Comment(s)		

7. Would you or your laboratory benefit from the availability of basic training courses in using operating systems and programming languages?

	Number of Response(s)	Response Ratio
YES	66	38.8%
NO	36	21.1%
N/A	11	6.4%
No Responses	57	33.5%
Total	170	100%
20 Comment(s)		

8. What major commercial or open-source software are you using?

	Number of Response(s)	Response Ratio
SPSS	15	14.0%

SAS	15	14.0%
NCSS/PASS	0	0.0%
MATLAB	45	42.0%
SRB	0	0.0%
iRODS	1	<1%
BLAST	4	3.7%
OpenSees	0	0.0%
Abaqus	4	3.7%
Ansys	3	2.8%
Fluent	3	2.8%
N/A	16	14.9%
Other	45	42.0%
Total	107	100%
25 Comment(s)		

9. What other software packages would enhance your research productivity and what is preventing you from using them now?

39 Response(s)

10. What are your greatest data management needs? (Select as many as apply)

	Number of Response(s)	Response Ratio
Storage capacity, more short term (1-3 years) storage for research data	60	53.5%
Ability to process and manage large quantities of research data	46	41.0%
Data management software (Oracle, MySQL, SRB/iRODS)	11	9.8%
Data analysis software (SPSS, SAS, Cognos)	30	26.7%
Transferring experimental data to storage facility	16	14.2%
Transferring data from storage to desktop or cluster	29	25.8%
Sharing your data collection with colleagues (via web or resources such as	35	31.2%
Access to national or community repositories (examples? PDB? NVO?	7	6.2%

Data Backup	51	45.5%
Long term data preservation of large data sets	35	31.2%
Metadata creation for large data sets for archival purposes	13	11.6%
Long term access via a common repository	17	15.1%
Data/format compatibility	16	14.2%
Meeting data privacy/security requirements (FISMA, HIPAA)	6	5.3%
Live online access to large quantities of live data	11	9.8%
N/A	11	9.8%
Other	4	3.5%
Total	112	100%
9 Comment(s)		

11. Are your current research needs for data storage being met?

	Number of Response(s)	Response Ratio
YES	92	54.1%
NO	18	10.5%
No Responses	60	35.2%
Total	170	100%
24 Comment(s)		

12. What do you anticipate your research data storage requirements will be for temporary datasets? (Please select only one option)

	Number of Response(s)	Response Ratio
1- 500 Gigabytes	36	21.1%
500 Gigabytes - 2 Terabytes	28	16.4%
2-100 Terabytes	29	17.0%
More than 100 Terabytes	1	<1%

More than 1 Petabyte	1	<1%
Not Sure	13	7.6%
N/A	4	2.3%
Other	0	0.0%
No Responses	58	34.1%
Total	170	100%
7 Comment(s)		

13. What do you anticipate your short term (2-3 years) research data storage requirements will be ? (Please select only one option)

	Number of Response(s)	Response Ratio
1- 500 Gigabytes	30	17.6%
500 Gigabytes - 2 Terabytes	23	13.5%
2-100 Terabytes	33	19.4%
More than 100 Terabytes	5	2.9%
More than 1 Petabyte	0	0.0%
Not Sure	13	7.6%
N/A	4	2.3%
Other	0	0.0%
No Responses	62	36.4%
Total	170	100%
5 Comment(s)		

14. What do you anticipate your long term (4 or more years) research data storage requirements will be? (Please select only one option)

	Number of Response(s)	Response Ratio
1- 500 Gigabytes	19	11.1%
500 Gigabytes - 2 Terabytes	21	12.3%

2-100 Terabytes	33	19.4%
More than 100 Terabytes	12	7.0%
More than 1 Petabyte	1	<1%
Not Sure	20	11.7%
N/A	5	2.9%
Other	0	0.0%
No Responses	59	34.7%
Total	170	100%
6 Comment(s)		

15. Does the current campus network (LAN and WAN) meet your current needs?

	Number of Response(s)	Response Ratio
YES	95	55.8%
NO	18	10.5%
No Responses	57	33.5%
Total	170	100%

16. If no, please be as specific as possible and describe your particular area of need, including peak and sustained data rates if possible.

17 Response(s)

17. Does your current or near-term future (next 2-3 yrs) research require connectivity to any national laboratories, research centers or international collaborations?

	Number of Response(s)	Response Ratio
YES	40	23.5%
NO	72	42.3%
No Responses	58	34.1%
Total	170	100%

18. If "YES," please describe.

37 Response(s)

19. What types of cyberinfrastructure, resources or system administration would be most useful to you, your group, lab, department, or "unit"? This could be in the form of deployment of a cluster, providing robust storage, providing long-term data management, machine room space, and others. (Select all that apply.)

	Number of Response(s)	Response Ratio
IT system administration	35	31.8%
Cluster system administration	29	26.3%
Machine Room Rack Space and/or Co-Location Facility	14	12.7%
Compute Resources	37	33.6%
Access to supercomputer cycles	40	36.3%
Storage Resources	53	48.1%
Networking Resources	31	28.1%
Archival planning Resources	17	15.4%
Metadata creation Resources	12	10.9%
N/A	24	21.8%
Other	5	4.5%
Total	110	100%

14 Comment(s)

20. What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department, organized research unit or other identifiable campus “unit”? This could be in the form of programming and staff support, managing a software stack, and others. What kind of advanced expertise would be of value (visualization, parallelization, database design)? (Select all that apply.)

	Number of Response(s)	Response Ratio
Interface/Portal development (GUI, Web-based)	23	21.2%
Database and/or data management support (e.g., Schema design,	24	22.2%
Scientific programming/Modeling	54	50.0%
Visualization (Scientific, Medical, etc.)	40	37.0%
Managing a Software Stack (builds, revision control)	17	15.7%
Statistical support (e.g., survey design, analysis)	25	23.1%
Software tailoring (e.g., porting code, scripting)	25	23.1%
Software parallelization/ optimization for clusters	38	35.1%
Technical Documentation	11	10.1%
Virtualization	6	5.5%
N/A	19	17.5%
Other	5	4.6%
Total	108	100%
11 Comment(s)		

21. Which of the following user support options would you be most interested in? (Check all that apply.)

	Number of Response(s)	Response Ratio
Helpdesk support	72	70.5%
desktop computer support and troubleshooting	45	44.1%
software diagnosis	23	22.5%
analysis of job failures	26	25.4%
training in programming, modeling, cluster or supercomputer use	63	61.7%
Other	6	5.8%
Total	102	100%
12 Comment(s)		

22. Please rank the following research computing funding options.

Top number is the count of respondents selecting the option. Bottom % is percent of the total respondents selecting the option.	**unreasonable option/ would not work**	**would be difficult, but might work**	**Could work pretty easily**	**option, would work just**
Pay for compute cycles and storage as they are used	32 36%	34 38%	15 17%	8 9%
Buy-in for a block of individual or departmental time	28 33%	34 40%	15 18%	8 9%
Include compute time and storage as a line item in my grant	17 19%	32 36%	29 32%	12 13%
fully centrally supported by campus or state funds (may result in central pull-back to cover	16 18%	16 18%	28 31%	31 34%
condominium style-my grant will add clusters and storage to the 'condo'	24 30%	32 40%	22 27%	3 4%
Split funding: 50% central, 50% customer/researcher supported	22 26%	34 40%	27 31%	3 3%
seek corporate users to offset cost to researchers	33 38%	24 28%	14 16%	16 18%
centrally supported through externally-generated overhead	19 22%	18 21%	27 31%	23 26%
phase-in costs, pay out of overhead, then work to line item in grant	19 24%	37 46%	20 25%	4 5%

33 Comment(s)

23. Can you suggest an innovative funding option that would work for your research program and sustain the infrastructure and support?

36 Response(s)

24. What would attract you to change from local/departmental clusters/storage to other providers (on or off campus)?

	Number of Response(s)	Response Ratio
reduced cost	43	40.1%
improved security	28	26.1%
central administration of hardware	25	23.3%
leverage economies of scale	21	19.6%
spend more time on my research and less on adminstration of my hardware	48	44.8%
expanded access to expertise	30	28.0%
improvements in up time, performance and reliability	33	30.8%
ease of use	38	35.5%
N/A	24	22.4%
Other	2	1.8%
Total	107	100%
16 Comment(s)		

25. What are the pros and cons of the services you currently use?

48 Response(s)

26. Please enter any additional comments you would like to add about your research computing needs.

23 Response(s)

27. (optional) Please enter your contact information below.

First Name	44
Last Name	44
Email Address	44

Constant Contact Survey Results

Survey Name: UAF Research Computing Resources Survey
Response Status: Partial & Completed
Filter: None
Feb 09, 2011 7:50:11 PM

1. What is your primary role? - Other responses		
	Answer	**Respondent**
	retired math faculty	
	administrative faculty	
	faculty	

1. What is your primary role? - Comments		
	Answer	**Respondent**
	I am involved in experimental and computational research related to alternative energy, space sciences, physical properties of snow and ice, and water resources. I utilize HPC for constructing new models for analysis of landed missions on mars and the moon. I also utilize commercial multi-physics programs to solve problems related to snow physical properties. I currently utilize both the computational and scientific expertise at ARSC in constructing my space sciences modeling capabilities.	
	As Dept Chair, I am also the recipient of faculty complaints about technology, especially classroom technology	
	I have to bring in 59% of my salary.	
	But will be transitioning to full time faculty in January.	
	Emeritus prof, working part time	
	Newly hired	
	Assistant Professor Department of Computer Science UAF	
	Also part-tiem director	
	Also, 49% appt as director of Alaska Native Language Archive.	
	I do not myself use the supercomputing facilities, but I know that the availability of those facilities has made UAF a much stronger institution overall. We need to do everything possible to keep them. Thanks.	

	Research administration for College, and also □ data management for College: AUP data, PBB data,□ program review, accreditation data, strategic planning scenarios for Dean and Dept. Chairs, and so forth.	
	I'm a beginning Assistant Professor of English and the Director of Composition.	
	Maintain the two Photo and Video editing labs for the Journalism Department.	
	I was a graduate student until a few months ago at UAF. I am now a postdoc in France, but still use ARSC supercomputers to work with UAF faculty.	
	I am working with parallel CIPL2-MPI numerical model on ARSC HPCC.	
	Alaska Tsunami Mapping Program,□ A developer of tsunami numerical codes	
2. Please select your School, College, Institute Affiliation. (You may select more than one) - Other responses		
	Answer	**Respondent**
	ESTES	
	Chancellor's Office	
	Chukchi Campus	
	Geographic Information Network of Alaska	
2. Please select your School, College, Institute Affiliation. (You may select more than one) - Comments		
	Answer	**Respondent**
	As I said, I started using ARSC machines while I was in College of Natural Sciences & Mathematics, but now I am in ITER Organization, France.	
	Geophysics/ Permafrost	
3. How are your computational research requirements currently being met? - Other responses		
	Answer	**Respondent**
	web server to host open access data	
	other academic clusters	
	Cloud, colleagues	
	resources as other institutions	
	OIT/ Provost Reserach funds for software	
	2 Computer Labs	
3. How are your computational research requirements currently being met? - Comments		
	Answer	**Respondent**
	All of my primary HPC use is done at ARSC.	
	personal purchases	
	All my research projects rely on CPU time available at ARSC. I nearly lost a project because of the article in the News Miner on Friday last week.	

	For me, other chemistry faculty, and chemistry students, the need is almost completely in the area of computational chemistry. This means using HyperChem on local PCs, or Gaussian on a linux server. For the latter software, most jobs currently are accommodated by large fast multi-core node. ARSC seems to be losing its Gaussian license, and Chemistry has recently been funded by TAB to purchase a departmental license.	
	I currently have research accounts on the HPC centers at Caltech, Princeton, and Harvard.	
	I maintain the UAF CS Bioinformatics Powerwall, an array of 10 nvidia GeForce GTX 280 graphics cards, connected with Core2 Duo processors and gigabit ethernet. Software is g++, Ubuntu Linux, OpenMPI, and my own homebrew "MPIglut" and "cudaMPI" software.□ □ I also maintain NetRun, a web-accessible programming environment used by CS and CE students. It's hooked into the powerwall to run MPI and CUDA jobs.	
	My needs are currently met, but the LSI/Biosciences cluster has been on the verge of disappearing for years now and its fate hangs in the balance. It is a missed opportunity that UAF admin has not provided a line-item base budget for this important program or sought to grow it more vigorously.	
	I stopped to use any supercomputing service when I realized that ARSC started to build research capacity in my own research field.	
	Primarily via desktop workstation and software, but I pull a lot of what we need from PAIR reports, Banner modules, QMenu/QAdhoc, EndNote,have created our own purposes database in Access/Excel, and various other UAF sources.	
	Our art/film courses have been using the ARSC lab for 3D graphics, animation, and video production in addition to labs within CLA departments. Loosing ARSC will put more pressure on our departmental labs, whose funding is provided by TAB grants or departmental funds. Oversight and maintenance of some of these labs is provided by faculty -without OIT or staff support.	
	I use a desktop computer, but all my needs are not currently met.	
	We have two computer labs with high-end Mac computers in them that provide our students with essential access to Final Cut and Adobe editing suites. We also provide access to our own file server for storage of projects.	

	About 300,000 CPU hours for air quality and ice fog modeling. Combined models for emissions and numerical weather models are running at ARSC on quasi-operational mode to support the public community, customer like the Air Force, DEC, the Borough,, EPA, etc.. My research is especially focused upon modeling wildfire smoke, volcanic ash, and ice fog.	
	99% ARSC, 50,000+ CPU hours annually 1% Desktop	
	I will need to use ARSC computers starting in February for the reminder of my thesis.	
	Desktop, laptop, smart-phone and cloud computing.	
4. If your computational research needs are not being met, please specify the unmet needs now and into the future (RAM, CPU,		
	Answer	**Respondent**
	n/a	
	N/A	
	Need continuous purchase of annual license (ABAQUS) as well as technical support of Lynux machines.	
	The current ARSC resources are satisfactory, however as codes continue to grow and become more complex this will obviously put higher strain on the systems I use.	
	My current needs are being met. If ARSC is defunded, it is not certain that my current or future needs will be funded.	
	N/A	
	I need about 200,000h of CPU time each year. I cannot get it from other computer centers like NCAR as my grants are through EPA, DEC, FNSB. These agencies also do not provide any support to buy CPU time. The number of CPU hours needed will increase as the model resolution will increase and the complexity of processes considered in my modeling will increase.	

	If ARSC would continue with a Gaussian installation, that would be good. Whether it is worth the $20K+ to purchase a supercomputer license (vs. $5K for a departmental license), is a good question. However, due to the onerous security and account policies that surround the ARSC installation, even if this software were available, it would still be usable only for research purposes by graduate students and above. These computers are not easily accessed by students probably because of their DOD affiliation. In contrast, the Chemistry Department's WebMO installation (http://chem.uaf.edu/keller/WebMO/) is free to all, with certain limitations. We use this for undergraduate, even HS, chem computing. WebMO is a JAVA based graphical frontend for Gaussian and other computational chemistry packages which are installed on the Chemistry Department's small linux cluster. This is a model that could be adopted, or adapted elsewhere.	
	N/A for research, but for basic internet access and ability to use streaming video for classes we could use a whole lot more bandwidth and faster connection...	
	N/A	
	There is concern that the Biosciences Cluster and the IT support for it are not funded in a manner that will ensure their longevity. This cluster is critical to faculty conducting phylogenetic research who often need 2-4 CPUs per analysis which can take days to weeks to complete.	
	I don't know specifically about research needs, but the entire KuC campus internet is almost as slow as dial-up. it can take several minutes to log on in the morning - and using video conferencing from our desktops is painful. our usage often exceeds capacity.	
	n/a	
	n/a so far	
	My computational requirements are currently being met by ARSC and departmental computer support.	
	N/A	
	More Web-based resources are needed for the ARSC systems. Data portals and computational portals.	
	Low-tech but moderate capacity and bandwidth shared server space that can be accessed both by University and non-University collaborators -- ideally not charged back to the department, which is cash-strapped and can't afford it.	
	N/A	
	N/A	

	Current need is for approximately 20TB storage, with sufficient bandwidth to provide data sharing to other major repositories. This storage need will grow by approx 20-30 TB over the next 5 yrs.	
	Could make use of additional computational resources for one current and one pending project.	
	So far I my personal computer has been adequate for my computational needs. However I am currently pursuing a research project that requires a great deal of computational time (specifically running simulations that involve a lots of numerical integration). I was hoping to use the facilities at ARSC for this purpose.	
	Bandwidth going in and out of the Duckering building is lower than anywhere else on campus and no one can tell us why. Can you explain.	
	Occasionally in my research I need access to more ram, but that hasn't happened in a while, so at the moment, N/A.	
	The major needs in my lab revolve around the storage and analysis of very large DNA sequence datasets. Life Sciences Informatics has done a fantatistic job of meeting these needs, with the limited resources that have been available to LSI. That being said, with ever growing datasets and changing/novel analyses, we continue to need greater storage and clusters with larger RAM, more nodes and faster cycles. Some of our analyses take months to run on the current LSI systems, and some analyses have been impossible due to RAM limitations. However, the greatest need is for customized support to provide scripting, install complicated software, and provide web-based solutions. LSI provides all of these services, but LSI is overworked and understaffed.	
	While my laboratory computer is meeting some of my current needs, my research program was targeting the power available in the ARSC to apply bioinformatic methods to linguistic data to test hypotheses on the nature of the relation between Athabaskan in North America and Yeneseic in Siberia, of the relatedness between Mesoamerican and Californian Indigenous languages, and on the multiple routes of migration of the original populations of the Americas. Millions of calculations are necessary for this research which requires supercomputers. The ARSC was one reason I chose to work here. I have additional needs for the archive storage of volumes of multimedia data on endangered languages that are not currently met by UAF. I was surprised to learn that UAF is unlike most other research universities in not making network storage space available for faculty and research groups.	

	So far data needs are being met, but frequently the desktop machine slows down (when on UAF network) hampering ability to download data needed, or upload data reports requested. In my Research Administration role, transmitting large proposal packages via email (e.g., PureEdge packages such as for US Dept. of Education) often exceed mail message attachment capacity. We're just now trying that OSP share drive as a remedy. Some of the problems experienced are directly related to which internet browser I use for the task: Internet Explorer works better for anything I have to do with Federal portals (Grants.gov, Fastlane, Era Commons etc...) but is much slower and can hangup occasionally. Mozilla Firefox is at least faster when doing anything with data on UA network.	
	n/a	
	HD Video (the broadcast standard as of 2009) requires very high bandwith, Terra bytes of RAID server space, and software licensing that OIT currently doesn't have as site license (namely Final Cut Pro Studio). Currently all of our needs are being funded through TAB grants, research work, and student fees through the department, but it is a constant struggle to keep up with the growing demands.	
	I need video editing software that is not available through OIT downloads. I would like to have space beyond the hard drive of my office computer and my portable drives to back up my audio and video data.	
	Desktop computer is now 5 years old, overburdened with UAF system software that has caused it to get slower and slower. My research involves statistical analysis of very large datasets by social science standards. To run a complex multivariate model can take several hours. It is not unusual to return in the morning and find that the computer has rebooted for routine maintenance and the work has been lost. There are no funds to replace my computer unless I get a research grant, as the assumption in CLA is that we do not need computers with significant data processing capacity. It is increasingly difficult to obtain a research grant as time passes during which my productivity is compromised. There is no institutional support for backing up my data, which sits on my local hard drive and is backed up on a portable hard drive purchased with personal funds. Last time I checked, the charge for a small amount of networked disk space was OUTRAGEOUS, but this resource would be very useful to me.	

	n/a	
	NVivo Software, which organizes and interprets qualitative data such as transcripts, surveys, questionaires, and interviews.	
	Our current arrangement requires that students either check-out or purchase an external hard drive for their video projects because our server is not set up to provide incremental backups of all student work. We would, however, benefit greatly from a storage solution that may provide between 5-10TB worth of space for archiving of raw footage, project files and finished video projects along with incremental backups that would facilitate recovery from accidental deletion. As well, the potential to remotely access these files, rather than having to download them every time for quick access, would be very helpful. Currently, students that save on the server have to move files locally for quick access & then reupload.	
	In the future (next 3-4 years) I anticipate needing assess to ARSC facilities for a very large scale simulation of societal networks (complex systems), built on the basis of agent-based modeling of human agents involved in dyadic communication.	
	in house technical support staff are unable/unwilling to provide adequate assitance to support staff and students in use of several software packages. i.e. links to programs appear broken so students cannot even start the programs, sometimes these links open the wrong program, sometimes students must change codes using the command line in ordere to be able to use programs.	
	My computational needs are currently met by ARSC supercomputers.	
	N/A	
	My "workstation" is very outdated (it's a salvage machine), and the ARSC worstations have painfully outdated software on them.	
	N/A	
	I mostly need lots of ram. 8 - 10 GB per processor would be most desirable for manipulating large grids in programs currently written to use a single processor.	
	It would be nice if MATLAB can be available on pacman.	
	Most of my computing needs relate to software. I have one-time licenses for analytical and word-processing software, but it would be helpful if the University could buy site licenses for them so they can be easily updated: Atlas.ti, SPSS, OmniFocus, OmniOutline, MS Office. I also run Parallels on my Mac in order to run Atlas and SPSS. The driving need is just to keep my existing software updated.	
5. What operating system do you typically use for your scientific research? (check all that apply) - Other responses		
	Answer	**Respondent**

	ChromeOS	
5. What operating system do you typically use for your scientific research? (check all that apply) - Comments		
	Answer	**Respondent**
	Also are attempting to pilot a database for alumni/financial aid and grant reporting purposes.	
	Linux is the operating system used by ARSC supercomputers.	
	I don't do scientific research. I do artistic collaborations.	
	Since I use a Mac, sometimes I can't open/access PC files and documents sent to me.	
	I have checked only those of which I am aware. There may be others out there.	
	wound like toi se linux but without decent technical support this is more hassle than it's worth	
	Primarily Linux/Unix/Aix, some OSX (increasing) and a little windows for OS specific modeling.	
6. If computing services were provided on campus, which of the following would enhance your research productivity? (Select all		
	Answer	**Respondent**
	Supercomputer is needed like those at ARSC	
	Don't ignore the GPU: CUDA-MPI hybrid programming can give ridiculous speedups.	
	LSI	
	undergraduate and graduate researchers would also have access	
	Storage Space	
	more licenses, software infrastructure	
	supercomputer like mightligh in ARSC	
6. If computing services were provided on campus, which of the following would enhance your research productivity? (Select all		
	Answer	**Respondent**
	I don't think there is a one solution approach.	
	I need ARSC's supercomputers & the support provided by the ARSC consultants. I will be non-competitive in getting funding if I have to ask for CPU support or cannot show that the CPU time is available. I can ask about 100K/y per proposal. Anything higher makes the proposal uncompetitive. Having to pay for CPU-time, means cutting my support or my students' support. This means less graduate students, education & research. I currently support 4 graduate, 1 undergraduate students, 59% of my salary.	
	My needs are fully met by the hardware on my desktop.	

	User-built clusters is NOT a good strategy for computational science. A centralized facility is much more efficient overall, since the hardware and procedures are uniform for all. How to devise an efficient and fair queue and priority system for such a centralized facility is the main questions. (ARSC currently does well with this.)	
	We would prefer to manage our own machines within the research group.	
	ARSC provides the technical support and the hardware for supercomputing on campus, plus the training of students, faculty and staff in the use of these systems. This service is indispensable for providing the educational and research opportunities for Alaskan students and researchers in projects requiring high-end super computing.	
	None of the above- every single time the U has tried to cluster things and offer me savings, I have lost years of data, received pitiful amounts of storage space, had software eliminated that I was actively using, and endured years of refusal to upgrade software to the point I had to develop my own server. □ Put money into more help desk personnel especially ones why are Mac saavy.	
	The priorities would be from the bottom option to the top, i.e., supercomputing would be the top choice.	
	Provide for the current need better and effectively not attempt to bite off more than what you can provide.	
	None of the options above are sufficiently rapid, flexible and user-oriented to support my research. Long-term, central institutional support of LSI is what is critically needed.	
	It was difficult to check a box here. □ □ Qualitative research is interdisciplinary. NVivo is relevant to education, for example, a department that is not in CLA. □ □ Is there a choice here in which the innovation could bring researchers from different disciplines together? (If so, I'd like the check that box.)	
	We have the hardware. I mean, any old computer lab can be turned into a cluster. However the software isn't there. Moreover, when software *is* there, you need more than one license available to start clustering.	
	I think ARSC does a perfect job on this. I believe having department owned and run clusters will only isolate researchers at UAF- where interdisciplinary research is extremely important for UAF's growth as a research university. Plus, not all departments have the same amont of funding. So, supercomputing should be free to academic users as in ARSC.	

	Definitively not commercial	
	I imagine electrical work is expensive, housing all the servers in one place (ideally) would be more efficient for the electrical labor.	
	The option I selected would only be marginally useful to me, but is the most useful of the alternatives.	
7. Would you or your laboratory benefit from the availability of basic training courses in using operating systems and		
	Answer	**Respondent**
	Maybe.	
	These should be available through university computer science course offerings and special courses should not be needed. If this refers to offering courses through the CS then yes, but that is an academic issue. Specialized course offerings would have such a small base that these	
	Python programming□ R for statistics	
	We always made use of the classes offered by ARSC to get new graduate students up to speed. I use the ARSC computer lab in 2 of my classes.	
	I have run into situations where my productivity could have been increased if I had some programming skills. Especially now that SPSS works with Python.	
	Chemistry computing applications are not used at the OS level: they are graphical or text-based input. The necessary training is in ab initio quantum chemistry theory and practice, which is available in the chemistry department.	
	I feel I have the skills I need.	
	Based on my experience, much of the support for direct projects comes either from other group members (including the advisor), or from detailed online documentation, or from email feedback from dedicated support. Classes are beneficial, but not as much as dedicated consult support.	
	Training in supercomputer use is now provided by ARSC. Training in matlab programming is currently provided as part of departmental instruction in physical oceanography. To my knowledge, no other training is provided by SFOS/IMS.	
	I'd be happy to help train faculty and grad students in GPU programming. I teach CUDA to my graduate students and upper-division undergraduates in CS.	
	Ongoing training in Matlab, IDL, KML, Python, Linux/unix, Comsol, fortran	
	Especially for student researchers and student assistants	
	I tried taking those html classes for webpage design... If it takes me four hours to get dreamweaver to turn the page, I'm better off outsourcing the work.	
	Basic and advanced training would be helpful, especially in Applescript and other efficiency tools	

	Probably, but it would depend on the classes.	
	It couldn't hurt.	
	I think basic engineering classes should be restructured to require software earlier, and ALSO TEACH THE SOFTWARE. Which software isn't too important, though the engineering department uses primarily matlab, with a touch of solidworks and mathematica. I personally like python, but one can't have everything amirite?	
	Pretty much always.	
	A course in "programming for dummies" would be great.	
	This might be useful, even though I think it is hard to offer courses to people with such different backgrounds and skill levels.	
8. What major commercial or open-source software are you using? - Other responses		
	Answer	**Respondent**
	STATA, Eviews	
	Mathematica	
	SYSTAT	
	GIS software: GRASS, ARCINFO, stats: statview, systat	
	COMSOL	
	petsc, mpi, netcdf, python, idv	
	ArcGIS, Statistica,R for statistics	
	DMOL, Gaussian, Python, in-house sfotware	
	R	
	Gaussian; Tinker; NWChem;	
	R	
	GMT	
	Mathematica, various programing lang.	
	imageJ and other imaging software	
	GOCAD, CUBIT, SPECFEM3D, Mathematica	
	MrBayes, Garli, PAUP*, BEAST, Bayesphylogenies, BEST	
	comsol	
	CUDA, g++, Mathematica	
	Lahey Fortran, IDL	
	see comment	
	OpenFOAM, VisIt	
	R	
	R	
	Mathematica, GAP, Geometer's Sketchpad	
	COMSOL	

	Comsol	
	Microsoft Office PRO suite	
	Final Cut Studio Pro (7)	
	LiveCode	
	Praat, R	
	IDL	
	R, various genetics stats programs	
	agent-based modeling software	
	Eclipse, IRAP,	
	Atlas.ti	
	IDL	
	COMSOL, python, perl, octave, node.js, haskell, GNU Scientific Library, ...,	
	VISIT	
	ENVI/IDL, GMT, GRASS GIS, VisIt, Octave	
	Python, Fortran 77	
	ArcGIS, WaSiM-ETH, ORIGIN	
	fortran, idv, ucar, netcdf	
	FEniCs	
	Atlas.ti	
	Fortran	

8. What major commercial or open-source software are you using? - Comments

	Answer	Respondent
	It would be better to replace ABAQUS with ANSYS if the latter price is cheaper. Right now only ABAQUS is available.	
	But I can no longer afford to purchase this on my own. So I have large datasets that are no longer usable to me.	
	All the good stuff produced by the national computing labs and climate centers is open source, but local support (=people who know how to use/install it) for it is very helpful.	
	ArcGIS, Statistica,R for statistics	
	I use NCL, IDL, FORTRAN, C++, and various libraries available on the ARSC supercomputers	
	ncks fortran NCO totalview debugger	

	GOCAD is commercial (Paradigm) for 3D geological models, and I only use it through a Harvard license. CUBIT is at ARSC and is for 3D meshing. SPECFEM3D is Fortran90 software for 3D seismic wavefield simulations in complex media.	
	Mthematica	
	We use a number of commercial statistical, visualization and computational packages for a PC including Surfer, Sigma Plot, Statistica, Corel. ARSC provides fortran compilers, NCL, netcdf libraries, MPI software for multiple processor applications.	
	We write our own analysis software.	
	Matlab is essential, including symbolic math	
	I am working with private software like ORIGIN, DEC-FORTRAN-compiler, WINEDIT, MATHTYPE, etc.	
	Vulcan, IRAP, Eclipse, CMG, Rockware, Whittle, ARCMAP, ERDAS, many more	
	My lab uses about half of the 50+ programs available through LSI. I also use another ~10 programs on my personal computer or available through other public portals. Genomic and population genetic sequence analyses require numerous software programs.	
	Open Source: Bayes Phylogenies Bayes Traits Mr. Bayes Mesquite Paup SplitsTree QuickTree ELAN Arbil Audacity PRAAT R QGIS	
	Through the Provost's office we purchased 3 years worth of 20 site licenses of this software for UAF. Licences have been installed on over 45 machines and 15 departments including the Rasmuson Library's laptops, UA Museum, and Statewide's Public Affairs Office. The software license expires in 6 months, and I have no idea where we will find another $16K to retain our licenses.	
	I checked only those of which I was aware.	

	some of my graduate students also use matlab, but IDL is the main commercial software used	
	Expect to of the agent-based modeling in the next 3-4 years.	
	IDL the following software packages are used occasionally: ENVI ArcGIS	
	Too many to list.	
	...BLAS/LAPACK/ATLAS, qgis	
	Python with Numpy. It would be great if whatever computing platform comes if Numpy could be included with Python. Plus related extensions like MatPlotLib & MatPlotLib BaseMap http://numpy.scipy.org/ http://matplotlib.sourceforge.net/gallery.html http://matplotlib.sourceforge.net/basemap/doc/html/index.html	
	PGI and Pathscale compiler, TotalView debugger is extremely important for us. If there is no efficient debugger at an HPC center, then we cannot develop numerical codes and will try to find another HPC center. TotalView is the MUST for us.	
	For qualitative research, Atlas.ti and similar software are important.	
9. What other software packages would enhance your research productivity and what is preventing you from using them now? -		
	Answer	**Respondent**
	I need to use econometric software in order to enhance my research productivity, but so far those software has been purchased through my personal expenditure and school/university-wide support would be helpful, particularly in updating my software and research skills without finanacial restraints.	
	ANSYS is better for use than ABAQUS. However, it would be perfect if we have both software available. Right now the budget insufficiency is the mere reason to prevent me from using ANSYS.	
	campus licence for ARCGIS	
	nVivo--cannot afford to purchase on my own	
	N/A. ARSC installs/supports the software we need. They respond to software requests.	
	GIS, the high license fee	
	With the current level of computer support ARSC provides I am able to meet my research needs.	

	I am averse to learning new software unless there is a REALLY compelling reason to invest the time.	
	See # 4 - Currently dealing with questions of FERPA and security of local database.	
	GOCAD would enhance my productivity. It coses $5000 for a university license per year, which is not affordable to me at this time. Nor are there other users demanding it, as far as I know.	
	SAS would enhance my research productivity -- the cost and hassle of the license is prohibitive.	
	SFOS/IMS provides site licenses for some products such as Microsoft Office, Windows and Matlab. Other software licenses must be paid for and renewed or updated through research grants. For financial reasons, ARSC will discontinue support for its path fortran compiler which I currently use for the ROMS model and post processing software. I will therefore have to reconfigure the software for a new compiler. These tasks are not always trivial and can be very time consuming; time which must be charged to research grants since there is no other funding option.	
	Good FORTRAN compilers and good debugging tools.	
	YAMS and TetMesh are geometry processing packages that would help some of my model generation work, but they're too expensive for me to maintain a license.	
	Would be useful to have a VPN solution that would only route University data over the VPN, rather than all data -- in other words a VPN bridge on the local machine rather than simply a repeater.	
	None	
	FLUENT,the general purpose industry-standard CFD code. No licenses anymore on ARSC.	
	None.	
	Mathematica	
	A web interface to phred/phrap/consed. This set of Unix programs is very complicated to install and are not user friendly. These could be made available via the LSI web portal, but there has been insufficient technical support to make this happen.	
	ArcGIS which has a prohibitively high license fee	
	Additional modules for SPSS beyond the basic program are needed, including Categories and Regression, among others.	
	Photoshop... the keyclient doesn't work from my home machine, even over my DSL line; I have to be on-campus on a networked machine to prepare photos. Likewise, if I try to work on photos from my laptop while traveling to conferences or field sites.	

	We have a site license for almost all of the Adobe Master Production Suite. Adding the Full Production Suite (including After Effects) would be excellent.	
	A full-featured video editing program	
	Qualitative data analysis software. Too expensive.	
	LISREL; it's expensive	
	I'm starting to use dreamweaver; very useful.	
	NVivo. Cost.	
	Final Cut Server - having a huge backup of archived material for our News TV Reporting students to grab when archived material is needed for news reports. For longer projects, having the ability to distribute the rendering process to a server and then down to our own machines would be beneficial. Lack of infrastructure doesn't provide us with a fast enough connection to our own server to allow for fast distribution of rendering.	
	IDL, problem is the money to keep the licenses updated	
	We need an update to SPSS at UAF--the available version is 2 releases old. I'd like to have systat available or if SPSS is updated, we need more modules built in (there are many stats tests that do not come with the program). We need to update to Windows 7. I have this on my personal computers but not at work. I have problems accessing Sequencher. I am trying to resolve them but it isn't a quick fix. Dragonspeak. All of these are cheaper if the uni buys a site license rather than individuals buying software	
	NCL, NetCDF libraries	

	ARSC's basic libraries are ancient. It's nearly impossible to compile anything from source on the workstations as it is without building your own libC, etc. It's kind of a shitty situation. If I could easily compile things, I wouldn't care that ARSC doesn't have X obscure thing. Oh, and the security settings with ARSC make using github no worky. I work best with a remote git repository.	
	N/A	
	high resolution atmospheric models	
	N/A	
	Enabled MPI Input/Output, Parallel NetCDF	
	Project management software, idea-tracking software, e.g., OmniFocus, OmniOutline, etc. I do use them, but don't have the ability to update them regularly.	
10. What are your greatest data management needs? (Select as many as apply) - Other responses		
	Answer	**Respondent**
	secured long term storage of information	
	physical storage of paper-based data	
	Long uptimes for analyses that take weeks or monhts	
	git	
10. What are your greatest data management needs? (Select as many as apply) - Comments		
	Answer	**Respondent**
	Very difficult to extract longitudinal data from Banner system	
	Long term storage of digital information: e.g. images, movies, data files - increasing demand by funding agencies - requirement	
	I need to process long datasets all at the same time, e.g. two years of atmospheric and chemical data where each day has more than 2GB.	
	Local storage is sufficient.	
	ARSC provides the data storage capacity for large quantities of model data (10s to 100s of TB). Data can easily be moved from storage to working disks for post-processing. Moving large amounts of data from a remote site to a local computer may be very time consuming.	
	Our programs are growing and changing and our needs for computational capacity and data management will increase in the future.	

	Live online access to large data sets is necessary to make model output data available. We also would need long-term safe data storage for climate research purpose.	
	A stable compute environment managed by folks with a lot of experience with HPC.	
	Realistically, I just need git to version control a few megs of files. On the other hand, live, real-time data is always awesome, and it's a shame that we don't see this more in science.	
11. Are your current research needs for data storage being met? - Comments		
	Answer	**Respondent**
	secured long term storage facility is needed	
	I'm a qualitative researcher, and must keep physical paper materials for years. I need physical, secure storage space.	
	It would be nice if the data would go offline and we would not have to batch-stage them so often	
	At ARSC	
	No particular needs	
	I have a 2 Tbyte disc in my iMac and 8 Tbytes in a Drobo box on my desk. This supports comprehensive back up of the computer disc.	
	mostly	
	ARSC provides the data storage.	
	I use ARSC heavily and my research depends on their computing, support and storage capabilities.	
	Yes by ARSC	
	only through ARSC	
	Yes, via ARSC	
	We don't appear to have any campus resources available for data backup/storage made available to faculty.	
	Via ARSC...so not sure if they will be in a few months.	
	They are currently being met by LSI. However, as next generation sequencing datasets come online, the current LSI systems will be under strain.	
	I have need to archive multimedia data on a network that has an adequate back-up routine.	
	with a buffalo drive and zipdrives.	
	This is only because currently we're not creating an archive of our student's prior material - if we had additional space and better access, we wouldn't be meeting our data storage needs.	

	The /scratch and /archive situation is awkward enough, but then ARSC also just broke everything when switching up archive servers. >_< I would mostly prefer automatic rsync between /scratch and /archive.	
	I use two ftp sites provided by the GI but it would be better to organize the data storage and distribution.	
	However, FY12 and FY13 continued downsizing of ARSC will be a factor.	
	I am depended on the ARSC resources and will continue to be depended on such resources when I start as a UAF assistant professor this spring.	
	I am currently only working with small synthetic data sets, but this will change in the near future.	
	Yes, although the major limitation is in the limits on my time of keeping my cluster operating and safely backed up.	
12. What do you anticipate your research data storage requirements will be for temporary datasets? (Please select only one		
	Answer	**Respondent**
	[No Responses]	
12. What do you anticipate your research data storage requirements will be for temporary datasets? (Please select only one		
	Answer	**Respondent**
	HD video recording (hundreds of hours)	
	I have fully adequate resources here.	
	A single seismic wavefield simulation produces volumetric fields that require a few GB of storage.	
	One 2TB drive is $100. Charging more than $100/TB/year will result in very few users of the storage service.	
	Presently the output from the computation of the Indian Ocean Tsunami require about 500 GB.	
	in a RAID configuration for backup.	
	I don't really have data storage needs that large; I'm in the humanities, so I generally back up data myself on an external hard drive.	
13. What do you anticipate your short term (2-3 years) research data storage requirements will be ? (Please select only one		
	Answer	**Respondent**
	[No Responses]	
13. What do you anticipate your short term (2-3 years) research data storage requirements will be ? (Please select only one		
	Answer	**Respondent**
	HD video recording	
	There is a trend to consider more species, which means more data have to be stored.	
	Not an issue.	

	I need approximately 10 TB for a 2-year research topic.	
	see above	
14. What do you anticipate your long term (4 or more years) research data storage requirements will be? (Please select only one		
	Answer	**Respondent**
	[No Responses]	
14. What do you anticipate your long term (4 or more years) research data storage requirements will be? (Please select only one		
	Answer	**Respondent**
	HD video recordings archive	
	Not an issue.	
	I would lean on the low end (10 TB) of the 2-100 TB range. The key point is that the computations should be REPEATABLE, so there should be no absolutely essential need to save simulation results.	
	As more modeling data is generated, more than 100 TB may be required.	
	Understand, I will never use university shared storage. Not reliable.	
	see above above	
16. If no, please be as specific as possible and describe your particular area of need, including peak and sustained data rates if		
	Answer	**Respondent**
	I am not on campus	
	na	
	Because the department servers and firewalls are decentralized this creates problems accessing datasets and information from the various locations on campus. The network is slow sometimes as well, although I know I "should" be getting gigabit speeds.	
	The rates vary by one order of magnitude for some reason. At some times (off-peak), the download rate could be 5 MB per second; at other times (during the day) 0.3 MB per second. This is not a major problem, however.	
	There is a 10Mbit connection from 7th floor of the Elvey Building to ARSC. From there, I cannot transfer ~1TB radar data collected in the field to and from ARSC in a reasonable ammount of time (less than a few days). Instead, it is easier to hand-cary a hard drive to ARSC where I can transfer it to long-term storage.	
	WHAT'S A lan? WHAT'S A wan??	
	UAF's firewall prevents me from connecting to compute servers on most networks (no way to initiate contact with a NAT'd server), and the few networks that have real IPs typically require bizarre port remapping hacks to access ordinary ports such as SSH (port 22).	
	Peak rate of approx 17 Mbps is insufficient for moving large quantities of data to ARSC.	

	to slow, unreliable	
	MOST Of the time, as long as I physically work on UAF machine... not well when off-site. Other aggravations as noted in comments above.	
	We need to be able to stream HD Video to multiple computers from a server, while still being on the main network. Our lab currently only has 2 LAN dataports hardwired into the lab that get shared through a splitter hub. This then requires us to have desktop machines with wireless in order to work with higher bandwidth materials.	
	A dedicated line (or at least faster) to our server and storage spaces.	
	Wireless hardly works in my office and I often have visitor scholars and collaborators from campus who need wifi access. □ □ The library has a convoluted log in system for quick seach--I log in and have to log in a second time before it works--this is at work and then at home (for the second login--at work I don't need the first login)	
	The VPN is a proprietary piece of shit. Please replace it with something openVPN can handle. Generally, the easier it is for me to ssh files back and forth, the better, and right now it's a massive pain in the ass.	
	N/A	
	Internet speed in my office at IARC is extremely very low.	
	Department-to-Department networks need up-grades to sustain 10-to-100 GB/s transfer rates.	

18. If "YES," please describe. - Responses

	Answer	Respondent
	I collaborate with researchers from other countries. Is this what you mean?	
	Currently working with the department of defense, Argonne National Lab and Stanford Linear Accelorator.	
	connection to online data providers like www.iobis.org - we provide data to those through our own data node housed in sfos.	
	Much of my research involve teams residing at NASA centers. At some point, as our computer code become operational, we will likely interact with NASA supercomputing facilities.	
	This is primarily an issue of knowledge transfer. The national labs and climate centers develop software we use. (E.g. PETSc at Argonne and NetCDF at UCAR.) Local technical support and a local user base require "connectivity", including travel.	
	The models I use have been developed at NCAR and EPA. Thus, I must be able to download updates. Moreover, I collaborate with NCAR and EPA.	

	frequent transfer of data sets between UAF and Argonne Nat. Lab	
	work closely with NOAA folks in Seattle. At present data sharing with them is very difficult and usually entails mailing hard drives.	
	I work closely with people at ORNL and to a lesser degree with people at PPPL, NERSC and LLNL	
	I collaborate with several non-U.S. citizens on code developments for HPC. We have access to other clusters (Caltech, Harvard, Princeton) for developmental purposes. It is essential that they have access to ARSC for the same reasons. Each clusters is somewhat different, and it takes a substantial effort to fully test.	
	LANL will serve us unclassified data from a customer.	
	We have a collaborative project requiring data transfer between ARSC and the NOAA facility in Seattle, and to a colleague in Chile. We anticipate the need to make data sets of 10s of TB available to these people.	
	We will need to download the next IPCC dataset from the storage facility to analyze it.	
	I do collaborative research with faculty at Illinois (UIUC) and Washington (UW Seattle). It would be beautiful if ordinary pedestrian TCP/IP traffic from these institutions wasn't filtered to the point of being unusable.	
	Joint research with Texas A&M require transfering input and output data for Supercomputing.	
	can not predict at this point in time, depends upon funded project mix	
	Collaborators in lower 48, connections to national lab	
	Need to be able to exchange datasets with collaborators at National Labs.	
	I exchange large sequence datasets with a variety of vendors and collaborators, including DOE Sandia labs, DOE JGI and the Broad Institute.	
	Connections to Language Archives of the Max Planck Institute (Netherlands/Germany) and the AILLA (Archive of the Indigenous Languages of Latin America, UTexas, Austin).	
	Collaboration on drafting papers and reports with international colleagues; communications with students doing international field research.	
	Collaboration with Ohio State University over next five years.	
	As the Research Administrator for the College, we have numerous international research projects of varying sizes. We currently have active research collaborative projects with researchers in: Canada, Russia, Japan, New Zealand, Eygpt, Argentina, New Guinea, and anticipate others.	
	I routinely need to use ftp sites to upload data sets.	

	In the sense that we need to be able to share large pieces of raw audio and video data with high resolution, which cannot be passed around as attachments.	
	Animal genetics work collaborative with USDA/ARS in Montana.	
	uploading to genbank data management and dissemination for nsf grant	
	Connectivity to other climate centers, the NOAA National Weather Service, and to laboratories such as Sandia National Lab or the DOE Atmospheric Research Measurement Program sites at the North Slope of Alaska.	
	Products of the collaborations between my department and international institutes (e.g. Japanese) will require data transfer of the size (10s-100s giga-bytes) between UAF and the international counterparts. Majority of these products are mainly the outputs from the global climate/earth system models, and/or satellite data.	
	I connect to ARSC computers from ITER Organization France. I also collaborate with Oak Ridge National Lab.	
	Argonne for example	
	Sustain high-speed connectivity is needed for: 1) transfer TB datasets from NASA DAACs 2) transfer TB information datasets to "Cloud" centers (national and university sponsored) and 3) perform high-resolution graphics rendering over internet ssh tunnel with remote login at high performance computing centers (NCSA University of Illinois Urbana-Champaign for example). Number 3 anticipates disestablishment of ARSC in FY12/13. We are budgeting and planning for this contingency.	
	DOE's ARM network; probably Japanese network of some type.	
	I am depended upon supercomputers such as ARSC in my work.	
	NCAR, LANL and GFDL/NOAA	
	Pending research proposal will require video conferencing facilities.	
	Nothing out of the ordinary -- only standard networking protocols (ftp/sftp, ssh).	
19. What types of cyberinfrastructure, resources or system administration would be most useful to you, your group, lab,		
	Answer	**Respondent**
	more mathematica licenses	
	Computer science/ physics resource staff	
	ARSC	
	desktop support, smart classrooms	

	Hostile takeover of CEM Technical Services	
19. What types of cyberinfrastructure, resources or system administration would be most useful to you, your group, lab,		
	Answer	**Respondent**
	I am very happy with the computational support and service that ARSC provided. I need that this support continues to carry out my research and education. I need free CPU time in sufficient manner.	
	In the long run, installation and physical maintenance of a computational chemistry cluster in an OIT "farm" would be ideal due to severe space limitations in the Reichardt Building. We could use linux remote desktop to manage the WebMO application from Reichardt.	
	In fact, our most pressing need is to maintain the independence of our department computer lab in the face of OIT push for homogenization and integration of computer labs on campus.	
	I'm strongly in favor of a centralized HPC center with support staff.	
	Not sure what this means relative to what we already have - is the question asking what we need that we already have? or just what we don't already have?	
	PLEASE LEAVE OUR DEPARTMENTAL COMPUTING FACILITIES ALONE!!!!!!!!!!!!!!!!!!!	
	Ask SFOS IT director	
	I'd definitely consider trading out cycles on my 10 GPU, 20 CPU cluster, since it does have low average utilization. Resolving the security issues while still allowing real work to get done would be tricky, though.	
	Networking trumps all.	
	Do not know what Metadata means.	
	Need additional, functional, supported smart classrooms so as to teach research methodologies, hold seminars/research meetings, and so forth.	
	Just having a central system for IT consultation, design, and implementation would be helpful. Right now we are forced to do much of the design independently of OIT, then ask them to help us with the problems.	
	We *all* know that those guys are incompetent.	
	Not sure about this question	
20. What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department,		
	Answer	**Respondent**
	GIS lab	
	hardware and software support	
	What the consultants did at ARSC is what I need	
	desktop support, smart classrooms	

	FCP Server?	
20. What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department,		
	Answer	**Respondent**
	Our in house codes could especially benefit from the expertise of someone savvy in parallelization and optimization.	
	free or low-cost software, e.g. SPSS and nVivo, plus tech support for use	
	Python programming, R for stats	
	I do not need that anyone does some software stuff for me. I do that myself. I need people who help me if the model crashes due to software problems/incompatibility or stupid behavior of students. I need people who are better at UNIX, perl, etc. than I am in the case I have problems installing a new model version.	
	The type of support that has been available from the ARSC support staff has been wonderful. As a major user, I work and have worked with the support staff "consultants" at a number of the large national super computer centers (NERSC, Pittsburgh Super Computing Center, Univ of Ill Super computing Center, ORNL) and I can say without reservation that ARSC consultants are the best in the business in their willingness and ability to help.	
	PLEASE LEAVE OUR DEPARTMENTAL COMPUTING FACILITIES ALONE!!!!!!!!!!!!!!!!!!	
	Funding agencies are increasingly demanding that researchers seeking support (and the institutions they work for) have a credible data management plan to archive the data, provide metadata and provide web access to the data. The university will be under increasing pressure to provide these resources to remain competitive in acquiring grants.	
	Simple troubleshooting for day to day desktop computers.	
	I have my own programmer who does scientific programming, visualization, web design, and database architecture.	
	as noted above.	
	Engineers mostly need to be able to crunch MATLAB code, and be able to make sweet graphs. I think for larger problems we could also use some kinda parallelization action. I think my peers also lag a bit in terms of simple know-how. Also, virtualization is always useful, at least when it's not so hard to do. See: http://www.stackvm.com	
21. Which of the following user support options would you be most interested in? (Check all that apply.) - Other responses		
	Answer	**Respondent**
	software purchasing	

	ARSC consultants	
	N/A	
	smart classrooms	
	lab design consultation and matenience	
	In-house Tech already provided	
21. Which of the following user support options would you be most interested in? (Check all that apply.) - Comments		
	Answer	**Respondent**
	The process for purchasing software for use in my and my students' research is totally obscure, and there seems to be no OIT support for, e.g., handling licensing issues---or if there is such support it is not clear who to talk to or whether such requests would be handled politely.	
	I am very happy with the support that the ARSC consultants provide and with their quick response. I would not want to miss that service as it saved me a lot of time in the past. ARSC and their consultants' support was and is essential for my research success and for me in securing funding.	
	THe HPC course that I have been teaching with ARSC staff (Tom Logan) "Core Skills in Computational Physics" has been a very useful course for the students in a number of disciplines.	
	Dedicated support (via email) is as important as having the computing cluster itself.	
	I need helpdesk support when youse guys arbitrarily make my password non-effective and I'm not here to do anything about it.	
	Particular with respect to supercomputing.	
	I mostly need help in program development, parallelization of the models. Help and knowledge of scientifical math libraries for solving large set of equations.	
	high end support in the selected areas	
	Get a couple of help desk people certified in apple systems	
	I was engaged in supercomputing for several years, I need none of them.	
	I'd be interested in learning useful computing things. However, they should be offered on the East end of campus every once in a while instead of on West Ridge.	
	The GI already offers the needed support.	
22. Please rank the following research computing funding options. - Comments		
	Answer	**Respondent**
	I don't have a grant. Am I locked out of access to computing power?	
	Recent proposal success with NASA was facilitated by access to ARSC facilities and the ability for ARSC to provide partial support for computational aspects.	

	Most of these decisions should be made centrally, not the ways suggested above. If the decision is spread across all PIs an extraordinary amount of time will be wasted. Administrators have a job to do.	
	My research computing are exclusively on desktop (Windows XP) computers in my lab	
	The uppermost amount I can ask on a grant is 100K. If I have to pay for CPU time I cannot support as many graduate students & undergraduate students as I currently do. That means less education & research. The research institutes should give some of their overhead to cover the supercomputer center. It is wrong to have to pay for CPU time from grants. Agencies assume coverage by overhead. We could not get OIT support due to CRC but had to pay for CRC. This is wrong as overhead already overed OIT.	
	condominium style would have to have a central support staff for this to work. I am concerned that the price for cycles will make proposals non-competitive.	
	What is "research computing"? Is it different from "teaching computing" ?	
	I realize this is an evaluation of the need for ASRC. I have seen presentations on their need for funding a couple times. I find their "business" model completely unrealistic. I am opposed to funding ASRC out of overhead. It is not like the library.	
	N/A to our local database	
	Since I have and can get cycles at a number of other HPC facilities, Charging my grants would not work.	
	Excellent list of possibilities. If partly supported in proposals, then it is important that the cost is low enough not to jeopardize the success of the proposal. ANY PERSON at UAF should be able to at least try some things out for free -- this is the "If you build it, they will come" philosophy.	
	A base level of free access needs to be provided for suitability testing, training, and student educational needs.	
	LEAVE OUR DEPARTMENTAL COMPUTING FACILITIES ALONE!!!! Do whatever you want to those who desire a central facility.	
	When given a choice, many funding agencies will support those projects where computational resources are provided by the academic institution or agency requesting the funds rather than charged to the grant. Corporations are driven by the profit motive and are unlikely to be a reliable source of support. On-campus support facilities provide academic and research opportunities that remote facilities cannot supply. I don't know what "Buy-in for block..." or "Condominium style.." refers to.	

	My grants can support computing personnel at UAF but my community (CLIMATE) already has many computing resources available to us for free when a NSF or DOE grant if funded.	
	Do I pay more just because my programs need a lot more supercomputer hours to run effectively?	
	The arts are poor. Our needs are miniscule compared to many who will be answering this survey. We don't have a secure grant funding base -- and at current funding levels our departmental funding could not cover costs.	
	I would STRONGLY OPPOSE overhead being taken from grants obtained by researchers who do not use ARSC to fund ARSC. This would mean that one set of reserachers are subsidizing another set. UNFAIR and UNIMAGINABLE.	
	grant-funding could work for particular aspects of my projects, but there is still a need for an institutional commitment to data storage.	
	Sounds like you are trying to justify the use of the super computer by requiring that departments be charged a fee whether they have a use or not. If the users cannot pay for it, get rid of it. Don't tax those of us who will never use it.	
	Meanwhile, it is rather difficult to get funds for supercomputing. Thus, only the fully centrally supported by campus or state funds would be helpful. A computer center should only be working as a service center.	
	Do not understand the last option (phase-in costs..).	
	"Overhead" is the same as other unrestricted funds; I see no difference in that option vs. other central support options, unless the "Overhead" came from only one or a few units.	
	Costs of computing support would exceed currently generated ICR; the tail would wag the dog in that those programs with justifiable research computing needs and comparably fewer external funding opportunities or any that generate overhead would be disadvantaged (e.g, arts, humanities, social sciences). Seeking corporate users might be politically difficult (is that UAF mission) and who's use would take precedence? Isn't that essentially the mess we got into with ARSC?	
	an important option was left out in this query, "might work under certain conditions"	
	not really applicable in the humanities	
	Most of these, I believe, don't apply to us.	
	Some proposal funding agencies do not support costs for CPU hours, since it is assume to be included in the overhead at some competing agencies.	

	Computer resources NEED to be provided by the University and cannot depends on external funding. This in turn will enhance the probability for external funding of research. Funding organizations needs to see an increase compromise of the university supporting high performance computing.	
	I don't see the attraction for corporate users. I don't imagine UAF can compete with Amazon / IBM etc. in the cloud space. Some sort of phase-in of costs would seem to me to be most workable. Step 1: get a handle on computer usage (e.g. look at ARSC stats) Step 2: set up a new department or something Step 3: allow researchers to continue to work but have them write CPU time or something into these new proposals. So when they get funded there will be money to pay for the department	
	Sudden change to any direct pay or buy-in option will kill all projects that have not allocated funds so far and cannot re-allocate such funds under their current grant. This will mostly adversely affect graduate students. Phase-in of pay option over 3 to 5 years will ensure that future grant proposals will include possible cost of HPC resources.	
	Grants in the social sciences (outside of Psychology) tend to be shorter term, so researchers in CLA are likely to have periods without grant support which may interrupt their ability to pay for infrastructure.	
	If I tried to put $$ for supercomputing services in an NSF grant, NSF will tell me to apply for time at one of the centers it already supports. I think any pay by the use scheme is doomed to fail. The campus should provide the level of computing support it can afford from centralized funds, without imposing new taxes on the research units.	
23. Can you suggest an innovative funding option that would work for your research program and sustain the infrastructure and		
	Answer	**Respondent**

	Win the lottery? I think you have to play to the strengths of the institution, which would undoubtedly mean focussing on the academic support leveraged across multiple states or regions. With the correct connections and support this could allow for more reasonable assessments of where the $ could come from, but it would have to be functional to the point that users (in whatever state) felt like they have "someone" on their side. I think research will fund that no matter what...in science you have enough problems to hack at conceptually that dealing with resource issues is a bigger headache. Having a facility and infrastructure to make problems you don't need more transparent would be something folks could go for. I know that is what I would like.	
	provide part of overhead to faculty to allow for updating outdated equipment (e.g. 11 year old desktop computer)	
	Including major computing support requirements in proposals as a direct line item can be difficult since those are facilities type costs are strongly frowned upon. It is preferred by agencies for the organization to already have those resources (thus, smaller computational costs are acceptable). It is easier to have part of the OH from a grant to support large-scale computational efforts.	
	cut administrative bloat, centralize operations as much as possible, and shift funding	
	Propose an Arctic Science Computing Center with a focus on: * some supercomputing resources, esp. green computing * open source software development in support of climate and cold regions infrastructure * close integration with observations (e.g. portals)	
	Slim the administration and use the money saved for research. Reduce ARSC to just the service component. Currently they are also doing research that was paid by DoD. Sell some CPU time to other users in town, e.g. GEVA, FNSB, the city, NOAA, NWS, etc.	
	I think the university should be funding the computational resource requirements of the university much more than they are. They are an essential resource for many and the university should recognize it as such.	

	If HPC is incorporated in meaningful ways into undergrad and grad curricula, it can be paid for by fees and tuition dollars. Taking advantage of rapidly declining #/cpu cycle is important, and so probably a 2-year hardware replacement schedule should be factored in. Ultimately this will result in monetary savings. On the other hand, if these facilities continue to be "walled off" as "research" as they currently are at ARSC, they will be much harder to pay for, and justify.	
	I am happy with what I have.	
	Department of Education Title III funding	
	I line item from the Alaska Legislature to support the valuable resource that local HPC facilities are.	
	ICR to the grant holder to be spent on in-house services. That way you pay for what you usenot what you don't use.	
	I think the options in 22 cover what I could imagine. Perhaps you could have a model that would reward researchers (with ARSC CPU-hours or support) who also sustain CPU-grants for national centers. This would encourage the largest users to ALSO seek computational options outside of ARSC.	
	no	
	I think the line item in grants is the only way I know of	
	Innovative?: leave my research program alone. Leave our departmental computer lab and facilities alone.	
	I think that resources such as ARSC should be fully supported by the state. Corporations are subject to the whims of the market and commitments by non-profits can changed with changing economic and social demands. ARSC should be a public facility open to Alaskan educational and research institutions who can show a credible need for the resources.	
	See my comment above.	
	The massive supercomputers ought to be centrally supported by State and University.	
	not innovative, but effective and fair --- ICR allocation to HPC unit on top of core funding from State funds provided to the U. We researchers then pay with grant or departmental funds for specialized/dedicated services	
	UAF admin should seek line-item funding from the Alaska legislature for these initiatives.	
	I already have funding for my program. I don't use the supercomputer, never will use the supercomputer, and have no intention of paying for the supercomputer	

	The portion of research overhead generated from the campuses but given to Statewide could be redirected, at least in part, to support computational needs for research. This may result in a higher overhead rate assessed by the federal government (via Navy) since the overhead will be utilized closer to its intended purpose. But this will also effectively increase the amount of budget in proposals which may not be desirable from PI's point of view. However, since computing needs are usually considered as part of a university's infrastructure and are not something to be included in a grant, this may be a reasonable trade off. Many research projects conducted at UAF have mid-term or long-term benefits to multinational companies who are involved in the extraction of AK's natural resources. Portions of the tax coming from these companies could be dedicated, using, say, a formula, to support the computational needs. Many of these companies do hire our graduates.	
	NO	
	Any university that is serious about biological research must provide computational support that is not itemized by the nickle and dime on grants. Funding agencies like NSF consider this part of core facilities that should come out of overhead and/or state funds, just like heat and light. Specialized requirements that serve only a single lab and involve considerable personnel time, such as custom scripting, should be explicitly budgeted in grants. But system administration and general support for clusters, like those provided by LSI, should come out of grant overhead and state funds.	
	I'm trying...	
	Dollars spent on increasing efficiency--especially helping users to script repetitive tasks--would ultimately save time and money. A campus scripting specialist who would both train and write scripts (without charge) would pay for him/herself in improved efficiencies across the university in both research and administration.	
	I like the idea of corporate or federal/state agency users and the idea of then having buy-in to the system to enhance possibilities of future support.	
	n/a	

	At the University of Massachusetts Amherst, there was a "Technology Fellows" program. The fellows were made up of Graduate Teaching Assistants currently teaching College Writing. Fellows would meet formally for a year, reading scholarship on writing and technology. The nature of the program was to conduct inquiries regarding writing pedagogy and technology, and then to engage these inquiries through systematic teacher research. All projects made significant contributions to the local teaching program, as local teachers were the primary audience for the program. The Fellows carried an additional stipend on top of their University Graduate Assistantship. Fellows were also provided with new hardware and any software they needed in order to reflect on the relationship between technology and writing pedagogy and to design cutting edge classroom application was provided and available for the next group of fellows.	
	Haven't had time to investigate options.	
	Not innovative: Reduce ARSC to a service center, which is partly financed by research projects, and partly by the University (-through overhead?). A balance between customer support and UA funding needs to be established. The ARSC computing expenses might be reduced by more innovative energy-reducing measures.	
	It sounds like you guys are already considering something AWS-esque; That would be my major contribution. I think that, with the DoD pulling out, that computing resources from ARSC should become more available to students and less stale in terms of software. At least, I hope so. There is of course the massive staff cut to counteract the decrease in bullshit red tape.	
	Getting a Governmental or non Governmental (like, industrial) agency to fund the research and super computing at ARSC will be excellent.	
	Program managers don't want to fund local computer resources. On the other hand they see very positive when the university support the proposal-project with local resources. -This need to be provided by the University and i think it is strictly related to the global research vision of the university in the long term.	
	We are investigating several options at this time. We wish not to comment further at this time.	
24. What would attract you to change from local/departmental clusters/storage to other providers (on or off campus)? - Other		
	Answer	**Respondent**

	NOTHING WOULD ATTRACT ME	
	absolutely nothing	
24. What would attract you to change from local/departmental clusters/storage to other providers (on or off campus)? -		
	Answer	**Respondent**
	I prefer control rather than being dependent on other providers...minimizes time waiting for someone to show up who may not solve the problem.	
	The support ARSC provided is perfect.	
	We need independent hardware to support particular software we need for teaching. Becoming part of the Borg, while possibly cheaper, would prevent us from teaching effectively.	
	Both the in-house and off-campus models do NOT satisfy the above options. Only the centralized institute-wide option can achieve these.	
	Huge, centralized = yucko There would be lots of promises made about advantages, but there would actually be no accountability and the promises would be like campaign promises. Leave our departmental facilities ALONE!	
	I have been satisfied with the services received from ARSC to date. An off-campus facility would be counterproductive to my research and instructional efforts so I would oppose it. Granting agencies do not want us spending their research dollars on computer maintenance or software debugging, particularly since most PIs lack the training to service high end computers or diagnose software failures quickly and efficiently. ARSC provides those services fee of charge to the research project.	
	Access to much larger systems	
	Ability to share with wide variety of collaborators.	
	One size does not fit all.	
	Up-time is essential.	
	Absolutely nothing. I have been burned too many times by capricious decision making at the computer center. Us little guys always get lost in the shuffle. The university has cost me thousands of dollars in my past attempts to join clusters, and I will not be part of any attempt to do it again.	
	Off-campus will require extreme scrutiny to avoid ITAR problems and preserve data confidentiality.	
	A cluster and storage infrastructure managed by folks who know what they are doing, with a long track record of doing it right.	
	I don't like the ARSC kerberos system. I would like to be able to set up passwordless SSH to a cluster.	

	Definitively not off campus. This will end up killing any possibility of new research development.	
	The other provider would have to have demonstrated a record of effective and responsive system administration before I would consider it.	
25. What are the pros and cons of the services you currently use? - Responses		
	Answer	**Respondent**
	I am currently (and have been for 5+ years) a heavy user of ARSC's supercomputers most recently Midnight. I have found ARSC support to be outstanding and would like to see ARSC maintained. The simplest and least expensive way would seem to be to add a line item to grant requests.	
	I like having the control. I hate having the headache.	
	Pros - low cost & high capabilities of ARSC that allow for more competitive proposals and the ability to tackle more challenging (risky) problems than would otherwise be possible. Good support by knowledgeable people, easy and timely access to staff, facilities, and computational resources. Cons - I haven't had problems.	
	ARSC has been extremely useful because direct PI use of grant money for hardware is painful. Flexibility of supercomputing resources is just as important as peak performance.	
	I am in control and frustration is minimal.	
	ARSC consultants are highly efficient, the supercomputer has hardly any down times, ARSC training for new users is excellent, CPU time is free which helps me securing funding, their facilities helps me to teach my classes in a modern way so students are well prepared for high paying, high demand jobs in all fields of atmospheric sciences and climate research.	
	Very happy with ARSC services provided to date.	
	Pro: full control Con: requires local person with expertise and desire to make time for administering the cluster.	

	PRO: As stated above, the type of support that has been available from the ARSC support staff has been wonderful. As a major user, I work and have worked with the support staff "consultants" at a number of the large national super computer centers (NERSC, Pittsburgh Super Computing Center, Univ of Ill Super computing Center, ORNL) and I can say without reservation that ARSC consultants are the best in the business in their willingness and ability to help.	
	ARSC is a world-class computational center with rapid and accurate feedback from its support staff. At present it is not easy to become a user, especially if you are a non-US citizen. This should not be the case for an academic computing center but is understandable for a DoD center.	
	expensive	
	Pros: economies of scale, central support of hardware and software toolchains, low costs. Cons: Variable and uncertain software licenses (i.e. UAF MATLAB site-license) cause problems every around February.	
	Pros: It works Cons: none	
	I get good help from sfos	
	ARSC is an on-campus computing facility with the expertise to diagnose and correct hardware or software failures quickly and efficiently. If a problem arises, I discuss the issues in person with ARSC personnel to correct the problem with a minimum of time on my part; I am being supported by the funding agencies to perform research, not debug computer systems and software. ARSC has provided invaluable educational support for faculty, staff and students to make super computing systems a reality for University of Alaska and its disappearance would be a serious blow to the credibility of the university as a serious institution of higher education and research in the 21-st century. Because ARSC is a shared resource, the amount of resources available at any given time can be limited, however, I have had no problem to date receiving the support and computational resources required to complete my research.	

	I maintain my own cluster, and manage my own storage, for research and student use. This takes some of my time for sysadmin work, but gives me a lot of flexibility: if I need to eject all student jobs so I can run a benchmark for a paper, I can. I use the GPU for most of my research: compared to classical processors such as used by ARSC, the floating-point performance (per second, per dollar, or per watt) is incredibly good.	
	The ARSC systems are pretty slow to be upgraded (i.e., 'midnight' has a very old software stack). Can't a more aggressive upgrade path be found?	
	PROS: No incremental cost to department. CONS: No inter or intra sharing or backup of the very moderate amount of data we wish to use for collaborations.	
	Pros: Easy access to hardware and skilled, helpful personnel.	
	Presently I develop program for my research on Unix and PC machines. After this introductory step. I move (Fortran) programs to Supercomputer. Running programs and analyzing output is 24hr job. ARSC people have been essential in program development, data analysis and visuallization.	
	con -- doing it myself or in my group is a huge barrier to doing it at all, technology changes faster than we can keep up with it or at least takes effort that could be spent better on the research rather than the tool I'm using pro --- it's mine, and therefore can tailor it to my particular optimization	
	There has been no central financial support for the LSI/Biosciences cluster, which has now existed for about six or seven years. It needs a line item in the university budget and a commitment to hire faculty and programming staff with relevant expertise. OIT services seem minimal and disorganized. I had better computational support when I was undergraduate at another university 20 years ago.	
	Con: current solutions to long-term data storage are ad-hoc and thus jeopardize long-term survivability of important data.	

	I hae a good system. I would love to have a real Apple tech on board so I could call someone knowledgeable whenever issues with my desktop system and server come up. The help desk people have been superb, but it is always a learning curve for them to deal with MACOS.	
	After stopping to use ARSC-computers I transferred all software required to Desktop computers.	
	ARSC is really the best we have - extremely professional, competent and user friendly. Only pros and no cons.	
	I use only desktop support, somewhat infrequently (once to several times a year). I have no particular problems with that.	
	I really try hard in not using any of your services, I do not get good service and when I ask about a problem your people are quick to say it isn't them without checking to see if there really is a problem or not.	
	They are either provided by unverified could services or are payed for by me out of pocket.	
	N/A	
	ARSC systems have always been too hard to use as much as I would have liked. Support staff were helpful, but not as helpful as they could have been. Some consultants were generous with their time, others were not interested in helping me get my job done at all. Machine interface, security, data management policies have never been optimal for me. I have my own linux machine but it's a pain to administrate and I don't get as much help with that as I used to.	
	LSI is GREAT. It just needs more support.	
	Con: Doesn't always work Con: Expensive to install/update/troubleshoot Con: Time consuming to install/update/troubleshoot Con: Decentralized Con: Insecure Con: Limited accessibility Pros: Better than nothing	
	The problem with central storage has been that Computing Services has been ineffective in communicating to users about impending changes and not much concerned about consequences to users. We lost a whole database, and the switch to Roxen and demise of the faculty server cost a tremendous loss of web presence and google searchability--important for the dissemination of research.	

	In the past, I relied on audio recorded transcripts. After I typed the transcripts, I printed reams of paper, reviewed them, highlighting sections, eventually using scissors to cut the data into pieces in order to discover categories. This laborious process is the nature of qualitative research, where an inductive process leads the researcher to make a claim about what the data means. NVivo is designed to work for these same results, but saves resources such as paper and time for the researcher and the institution.	
	Pros: low cost. Cons: doesn't provide the incremental and archival backups that we need.	
	ARSC provides excellent computing resources and help if needed. I am afraid that the service will not be as efficient as it is now, particularly when reading through all the expensive options for a researcher that you suggest in the survey. the competition for NSF awards will get even worse if we need to add computer support/time there.	
	very little local support at no/low cost	
	N/A	
	I exclusively use the services of the Arctic Region Supercomputing Center (ARSC). I am very happy with the ARSC services. ARSC consulting support is world class. Cons: At times there are waiting times for submitting jobs.	
	ARSC doesn't cost me money, but their sofware is really stale. COMSOL has a shallow learning curve early on, but for more advanced stuff it's really cantankerous. It's hard to ad-hoc parallelize with ARSCputers because of the Kerberos thing. I like that I have admin on my office box, but I wish it was newer. The VPN would be way more handy if it didn't require a custom kernel module in Linux that didn't even compile without sketch user-built patches. Cisco sucks, mmkay?	
	As far as ARSC services are concerned - currently its only pros.	
	Currently is good needs better access high speed, license of Fluent, ARSC is great.	
	We have no comment at this time.	
	ARSC provide excellent customer support and it meets all my research needs.	
	ARSC is the perfect resource to smoothly conduct my research in reasonable time.	
	All of my data is backed up on an external hard drive, which is probably good enough for my current databases but in collaborating across universities, it would be good to have more centralized data storage.	

	PROS: Scaled to fit my needs. Systems/processes can be fully customized for my needs. CONS: When computer problems occur, they immediately become an emergency that sucks up my time.	

26. Please enter any additional comments you would like to add about your research computing needs. - Responses

	Answer	Respondent
	I didn't understand most of these questions, which perhaps is a consequence of the fact that most of my computing is done on my personal desktop (well, laptop) computer. My biggest computing need is access to software packages (in particular, Mathematica). I do not feel that it is easy to get access to these packages, especially long-term or for graduate students.	
	Thanks for asking us.	
	ARSC's service component is essential for my research and my teaching and my ability to secure funding. If ARSC and free CPU time fall away, I will not be able to support as many students as I do and I will not be able to teach and educate my students the modeling skills they need for employment in their field after graduation.	
	My research would not be possible without access to high performance computer. I have multi year grants that were awarded under the assumption that I would have access to cycles at ARSC. I am concerned that changes to ARSC will jeopardize by ability to meet my commitments.	
	see above for comments re "research" vs "teaching".	

	We have done (and published) a fair amount of fundamental research which has been facilitated by our HPC facilities here on campus. This has been something I have written into a number of proposals (that have been funded to the tune of over $2 million over the last 10 years). In at least a few cases the reviewers specifically mentioned the availability of these resources as a plus. That said, I have access to a huge number of cycles (with much poorer support) at centers ranging from NERSC and ORNL to the Spanish super computing center and so paying (beyond the overhead) for the cycles just would not make sense for my funding agencies. Please contact me if you have any questions.	
	I was recently hired by UAF and would not have accepted the job without the high-performance computing presence of ARSC. UAF has a golden opportunity to lead academic HPC with the likes of Harvard, Princeton, and Caltech, who have comparable HPC centers. Requirements for a successful HPC center: 1. HPC cluster contains 2000+ cores, sufficient disk space, and dedicated support staff to facilitate scientific tasks. 2. Cluster operates at full capacity with a fair scheduler, so that the maximal amount of good science can be performed. 3. The cluster is upgraded on a regular basis. 4. The university promotes the center as an integral scientific component of the university. 5. The university helps to support the center at some financial level.	
	Leave our departmental computing facilities ALONE	

	ARSC is currently critical to three grants I am participating in. These grants have brought overhead recovery to the department and provided research and educational opportunities that would not have been possible without ARSC. This research was funded with the understanding that ARSC would provide the computational and data storage services required to complete the research objectives. If the university is unable to sustain the critical infrastructure necessary to fulfill its research commitments as outlined in our proposals, it will become increasingly difficult for us to obtain support for future research projects.□ In my experience, the research and educational opportunities of students can be limited because they lack the programming skills to manage large volumes of data; skills which are growing in importance as ever larger data sets are being generated by remote sensing. Information in these data sets can be critical to research, management and public policy challenges.	
	ARSC is very effective, and I hope they can become better supported through central campus resources. I have brought in $millions in external funding that would not have been granted without ARSC's availability.	
	Hardware â□" a supercomputer is important as a tool for fast and massive computations. By far more important are people serving computation and software development. The methods of computations are in permanent development. Without people who help us with quick access and implementation of new methods we will not be able to make any essential progress. The cooperation with ARSC personnel was the main driving engine of our scientific progress. The specialized skill in program parallelization, solving large system of equations will disaper with those who worked in ARSC. We will need to spend money and time to eventually educate new personnel and it will take long time for them to gain experience.	
	UAF needs to make a basic commitment to data storage through creation of an institutional repository. This is a primary need for a major research university.	
	Some research projects on campus contain sensitive data that may require approval from the funding agency prior to their release to the public. Policies and their implementations at ARSC have been able to ensure the integrity of the data. In that regard, it's not clear whether outsourced management, take Google as an example (although not for high performance computing), can provide the same level of support. This aspect needs to be considered carefully during the decision making process.	

	Very interseting approach you mention acedemic needs but never address any of the acedemic needs I see or come across every day. You are unwilling to support our infrastructure needs by supplying more switches which should of been installed already and increasing the local connection speed from 10Mbps to 100Mbps. I have a hard time in supporting your needs.	
	I used ARSC for data storage of log files because OIT could not accommodate keeping these files long term (as in over 1-2 years). I need long term storage of files for my longitudinal research.	
	N/A	
	There is a right way to do it, and a wrong way. So far, we have only been able to afford the wrong way, which doesn't set up a great model for students who may want to know the right way to do it. The "computer science" aspect of video editing, data management, special effects, 3D animation etc. is an art unto itself, and requires a much bigger structure of support than what can be provided on a departmental level.	
	I am in my first semester as Director of Composition. First Year Composition, a sequence that consists of two academic writing classes, which are required of every student enrolled at UAF, has a vibrant teaching community. Teachers, largely graduate students and adjunct faculty, are in need of more support for their own research on teaching and curricular design. However, these teachers are sharing ideas and innovations right now by word of mouth and a box in the copy room that says "place useful handouts here." One of my goals (research and administrative) is to continue to foster a culture of learning and teaching at UAF. One way of achieving this goal is to point out the need for a program-run database, qualitative research software, and training and innovating professional development opportunities for these teachers.	
	our experience with OIT (all their service) is not very good ... actually quite bad, our experience with ARSC is very good. please do not make the supercomputer facility controlled by OIT	
	The large scale social network/complex systems research I am planning is simply not possible without access to large scale computing resources capable of accomplishing millions of iterations of a very large scale network, over hundreds of runs as parameters are varied systematically.	

	I believe ARSC perfectly caters to the supercomputing needs of the university. The consultants and specialists at ARSC are exceptionally talented and helpful. I know because I would not have a PhD if they or ARSC weren't there. They contributed, and still do, to every step of my research. I believe it's hard to replace ARSC. Free academic computing time, good machines, exceptionally helpful and knowledgeable staff, good data storage facilities - combining all of these is very difficult if not close to impossible. UAF can never be the same without ARSC.	
	This wouldn't affect me but the biggest hit will probably be to single-person researchers (no team) probably tenure track faculty who aren't super successful yet i.e. not bringing in much external research dollars or... older lazy faculty who don't chase research dollars but still do some research. These people aren't bringing in money but still need the computing resources for getting results and then writing papers. Not super fair for the successful people to be supporting them all the time. More generally, it seems tricky, coming up with a plan can't take a long time. The longer the gap between ARSC shutting down and the next thing coming online I imagine researchers will probably transition to offsite resources. Someone working with LANL might start using their computers, a person working independently may find it simple enough to use amazon's services. I don't know, feels like a decision is pretty time-sensitive because once people find other alternatives...	
	I like the mightnight type of supercomputer in ARSC managed by a team with help desk.	

Constant Contact Survey Results

Survey Name: Research Computing Resources Survey II
Response Status: Partial & Completed
Filter: None
Feb 09, 2011 7:52:14 PM

What is your primary role?

	Number of Response(s)	Response Ratio
Full-time faculty	7	36.8%
Part-time faculty	0	0.0%
Post-doc	1	5.2%
Graduate Student Researcher	0	0.0%
Administration (Dean, Director, etc)	2	10.5%
Undergraduate Student Researcher	0	0.0%
Adjunct Faculty	0	0.0%
Staff	4	21.0%
Other	1	5.2%
No Responses	4	21.0%
Total	19	100%
0 Comment(s)		

What is your primary campus affiliation?

	Number of Response(s)	Response Ratio
UAA	11	57.8%
UAF	2	10.5%
UAS	1	5.2%

Other	1	5.2%
No Responses	4	21.0%
Total	19	100%
0 Comment(s)		

Please select your School, College, Institute Affiliation. (You may select more than one)

	Number of Response(s)	Response Ratio
College of Engineering & Mines	0	0.0%
College of Liberal Arts	0	0.0%
College of Natural Sciences & Mathematics	0	0.0%
College of Rural & Community Development	0	0.0%
Graduate School	0	0.0%
School of Natural Resources & Agricultural Sciences	0	0.0%
School of Education	0	0.0%
School of Fisheries & Ocean Sciences	0	0.0%
School of Management	0	0.0%
Arctic Region Supercomputing Center	0	0.0%
Geophysical Institute	0	0.0%
Institute of Arctic Biology	1	33.3%
International Arctic Research Center	0	0.0%
Institute of Marine Science	0	0.0%
Institute of Northern Engineering	0	0.0%
Museum of the North	0	0.0%
Rasmuson Library	0	0.0%
Alaska Satellite Facility	0	0.0%
Other	2	66.6%
Total	3	100%
0 Comment(s)		

Please select your School, College, Institute Affiliation. (You may select more than one)

	Number of Response(s)	Response Ratio
College of Arts & Sciences	4	36.3%
College of Business and Public Policy	0	0.0%
College of Education	1	9.0%
College of Health and Social Welfare	3	27.2%
Community and Technical College	0	0.0%
School of Engineering	1	9.0%
School of Nursing	0	0.0%
Graduate School	0	0.0%
University Honors College	0	0.0%
Insitute of Social and Economic Research	0	0.0%
Environment and Natural Resources Institute	0	0.0%
Institute of Circumpolar Health Studies	0	0.0%
Justice Center	0	0.0%
Center for Human Development	0	0.0%
Center for Behavioral Health Research	0	0.0%
Alaska Center for Rural Health	0	0.0%
Biomedical Program	0	0.0%
Alaska Small Business Development Center	0	0.0%
Other	2	18.1%
Total	11	100%
0 Comment(s)		

Please select your School, College, Institute Affiliation. (You may select more than one)

	Number of Response(s)	Response Ratio
College of Arts & Sciences	1	100.0%

School of Education	0	0.0%
School of Management	0	0.0%
Career Education	0	0.0%
Other	0	0.0%
Total	1	100%
0 Comment(s)		

How are your computational research requirements currently being met?

	Number of Response(s)	Response Ratio
Desktop computer (local personal computer, workstation)	10	90.9%
Faculty/Research Group Cluster	1	9.0%
Department-owned cluster	2	18.1%
ARSC	0	0.0%
LSI/Biosciences Cluster	2	18.1%
National Systems (National Labs, TeraGrid)	0	0.0%
Open Science Grid	0	0.0%
N/A	1	9.0%
Other	2	18.1%
Total	11	100%
2 Comment(s)		

If your computational research needs are not being met, please specify the unmet needs now and into the future (RAM, CPU, bandwidth, etc requirements) What is the driving need for this growth? Please enter "N/A" if this question is not relevant to your research needs.

6 Response(s)

What operating system do you typically use for your scientific research? (check all that apply)

	Number of Response(s)	Response Ratio
Linux (any flavor, SUSE, RedHat, SLES, CNK/SLES, etc)	3	27.2%
Windows	9	81.8%
OSX (Apple)	2	18.1%
BSD/BSDi (Berkeley)	0	0.0%
UNIX/AIX	0	0.0%
Cell OS	0	0.0%
UNICOS/lc	0	0.0%
Windows HPC	0	0.0%
CentOS	0	0.0%
Solaris/ Open Solaris	0	0.0%
in-house OS, programmed own	0	0.0%
Other	0	0.0%
Total	11	100%
1 Comment(s)		

If computing services were provided on campus, which of the following would enhance your research productivity? (Select all that apply.)

	Number of Response(s)	Response Ratio
Support to locate department/researcher purchased servers in a central	6	60.0%
Support to run a department/researcher purchased cluster in a central	1	10.0%
Provide "Condo Cluster" services- you purchase cluster and add to other	1	10.0%
Provide "Computing Cycles" (campus procures a large cluster and allocates	4	40.0%
Commercial offerings/outsourcing	1	10.0%
on-demand, large memory, data-intensive supercomputing	6	60.0%
N/A	1	10.0%
Other	0	0.0%

Total	10	100%
2 Comment(s)		

Would you or your laboratory benefit from the availability of basic training courses in using operating systems and programming languages?

	Number of Response(s)	Response Ratio
YES	5	26.3%
NO	4	21.0%
N/A	2	10.5%
No Responses	8	42.1%
Total	19	100%
1 Comment(s)		

What major commercial or open-source software are you using?

	Number of Response(s)	Response Ratio
SPSS	3	27.2%
SAS	3	27.2%
NCSS/PASS	0	0.0%
MATLAB	4	36.3%
SRB	0	0.0%
iRODS	0	0.0%
BLAST	1	9.0%
OpenSees	1	9.0%
Abaqus	1	9.0%
Ansys	1	9.0%
Fluent	0	0.0%
N/A	1	9.0%

Other	3	27.2%
Total	11	100%
1 Comment(s)		

What other software packages would enhance your research productivity and what is preventing you from using them now?

3 Response(s)

What are your greatest data management needs? (Select as many as apply)

	Number of Response(s)	Response Ratio
Storage capacity, more short term (1-3 years) storage for research data	4	36.3%
Ability to process and manage large quantities of research data	6	54.5%
Data management software (Oracle, MySQL, SRB/iRODS)	1	9.0%
Data analysis software (SPSS, SAS, Cognos)	3	27.2%
Transferring experimental data to storage facility	0	0.0%
Transferring data from storage to desktop or cluster	0	0.0%
Sharing your data collection with colleagues (via web or resources such as	4	36.3%
Access to national or community repositories (examples? PDB? NVO?	1	9.0%
Data Backup	4	36.3%
Long term data preservation of large data sets	2	18.1%
Metadata creation for large data sets for archival purposes	2	18.1%
Long term access via a common repository	2	18.1%
Data/format compatibility	0	0.0%
Meeting data privacy/security requirements (FISMA, HIPAA)	0	0.0%
Live online access to large quantities of live data	1	9.0%
N/A	1	9.0%
Other	0	0.0%
Total	11	100%
0 Comment(s)		

Are your current research needs for data storage being met?

	Number of Response(s)	Response Ratio
YES	7	36.8%
NO	4	21.0%
No Responses	8	42.1%
Total	19	100%
0 Comment(s)		

What do you anticipate your research data storage requirements will be for temporary datasets? (Please select only one option)

	Number of Response(s)	Response Ratio
1- 500 Gigabytes	4	21.0%
500 Gigabytes - 2 Terabytes	1	5.2%
2-100 Terabytes	0	0.0%
More than 100 Terabytes	0	0.0%
More than 1 Petabyte	0	0.0%
Not Sure	4	21.0%
N/A	0	0.0%
Other	0	0.0%
No Responses	10	52.6%
Total	19	100%
1 Comment(s)		

What do you anticipate your short term (2-3 years) research data storage requirements will be ? (Please select only one option)

	Number of Response(s)	Response Ratio
1- 500 Gigabytes	4	21.0%
500 Gigabytes - 2 Terabytes	2	10.5%
2-100 Terabytes	1	5.2%
More than 100 Terabytes	0	0.0%
More than 1 Petabyte	0	0.0%
Not Sure	2	10.5%
N/A	0	0.0%
Other	0	0.0%
No Responses	10	52.6%
Total	19	100%
1 Comment(s)		

What do you anticipate your long term (4 or more years) research data storage requirements will be? (Please select only one option)

	Number of Response(s)	Response Ratio
1- 500 Gigabytes	2	10.5%
500 Gigabytes - 2 Terabytes	2	10.5%
2-100 Terabytes	3	15.7%
More than 100 Terabytes	0	0.0%
More than 1 Petabyte	0	0.0%
Not Sure	2	10.5%
N/A	0	0.0%
Other	0	0.0%
No Responses	10	52.6%
Total	19	100%
1 Comment(s)		

Does the current campus network (LAN and WAN) meet your current needs?

	Number of Response(s)	Response Ratio
YES	9	47.3%
NO	2	10.5%
No Responses	8	42.1%
Total	19	100%

If no, please be as specific as possible and describe your particular area of need, including peak and sustained data rates if possible.

1 Response(s)

Does your current or near-term future (next 2-3 yrs) research require connectivity to any national laboratories, research centers or international collaborations?

	Number of Response(s)	Response Ratio
YES	1	5.2%
NO	10	52.6%
No Responses	8	42.1%
Total	19	100%

If "YES," please describe.

1 Response(s)

What types of cyberinfrastructure, resources or system administration would be most useful to you, your group, lab, department, or “unit”? This could be in the form of deployment of a cluster, providing robust storage, providing long-term data management, machine room space, and others. (Select all that apply.)

	Number of Response(s)	Response Ratio
IT system administration	3	27.2%
Cluster system administration	1	9.0%
Machine Room Rack Space and/or Co-Location Facility	2	18.1%
Compute Resources	5	45.4%
Access to supercomputer cycles	5	45.4%
Storage Resources	5	45.4%
Networking Resources	3	27.2%
Archival planning Resources	2	18.1%
Metadata creation Resources	1	9.0%
N/A	1	9.0%
Other	0	0.0%
Total	11	100%

1 Comment(s)

What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department, organized research unit or other identifiable campus “unit”? This could be in the form of programming and staff support, managing a software stack, and others. What kind of advanced expertise would be of value (visualization, parallelization, database design)? (Select all that apply.)

	Number of Response(s)	Response Ratio
Interface/Portal development (GUI, Web-based)	4	40.0%

Database and/or data management support (e.g., Schema design,	2	20.0%
Scientific programming/Modeling	3	30.0%
Visualization (Scientific, Medical, etc.)	2	20.0%
Managing a Software Stack (builds, revision control)	0	0.0%
Statistical support (e.g., survey design, analysis)	4	40.0%
Software tailoring (e.g., porting code, scripting)	3	30.0%
Software parallelization/ optimization for clusters	3	30.0%
Technical Documentation	1	10.0%
Virtualization	1	10.0%
N/A	1	10.0%
Other	0	0.0%
Total	10	100%
0 Comment(s)		

Which of the following user support options would you be most interested in? (Check all that apply.)

	Number of Response(s)	Response Ratio
Helpdesk support	7	63.6%
desktop computer support and troubleshooting	6	54.5%
software diagnosis	1	9.0%
analysis of job failures	1	9.0%
training in programming, modeling, cluster or supercomputer use	6	54.5%
Other	0	0.0%
Total	11	100%
0 Comment(s)		

Please rank the following research computing funding options.

Top number is the count of respondents selecting the option. Bottom % is percent of the total respondents selecting the option.	unreasonable option/ would not work	would be difficult, but might work	Could work pretty easily	option, would work just
Pay for compute cycles and storage as they are used	2 22%	4 44%	3 33%	0 0%
Buy-in for a block of individual or departmental time	1 11%	5 56%	3 33%	0 0%
Include compute time and storage as a line item in my grant	0 0%	2 22%	6 67%	1 11%
fully centrally supported by campus or state funds (may result in central pull-back to cover	0 0%	1 13%	5 63%	2 25%
condominium style-my grant will add clusters and storage to the 'condo'	2 25%	3 38%	3 38%	0 0%
Split funding: 50% central, 50% customer/researcher supported	0 0%	2 25%	5 63%	1 13%
seek corporate users to offset cost to researchers	1 13%	4 50%	3 38%	0 0%
centrally supported through externally-generated overhead	1 13%	1 13%	4 50%	2 25%
phase-in costs, pay out of overhead, then work to line item in grant	1 13%	2 25%	2 25%	3 38%

1 Comment(s)

Can you suggest an innovative funding option that would work for your research program and sustain the infrastructure and support?

0 Response(s)

What would attract you to change from local/departmental clusters/storage to other providers (on or off campus)?

	Number of Response(s)	Response Ratio

reduced cost	7	63.6%
improved security	4	36.3%
central administration of hardware	6	54.5%
leverage economies of scale	3	27.2%
spend more time on my research and less on adminstration of my hardware	5	45.4%
expanded access to expertise	6	54.5%
improvements in up time, performance and reliability	2	18.1%
ease of use	5	45.4%
N/A	1	9.0%
Other	0	0.0%
Total	11	100%
2 Comment(s)		

What are the pros and cons of the services you currently use?

4 Response(s)

Please enter any additional comments you would like to add about your research computing needs.

2 Response(s)

(optional) Please enter your contact information below.

First Name	4
Last Name	4
Email Address	4

UAA-UAS Research Computing Needs Survey Results

Survey Name: Research Computing Resources Survey II
Response Status: Partial & Completed
Filter: None
Feb 09, 2011 7:52:14 PM

What is your primary role? - Other responses		
	Answer	**Respondent**
	Researcher/editor	
What is your primary role? - Comments		
	Answer	**Respondent**
	[No Responses]	
What is your primary campus affiliation? - Other responses		
	Answer	**Respondent**
	SW	
What is your primary campus affiliation? - Comments		
	Answer	**Respondent**
	[No Responses]	
Please select your School, College, Institute Affiliation. (You may select more than one) - Other responses		
	Answer	**Respondent**
	SW	
	CRS	
Please select your School, College, Institute Affiliation. (You may select more than one) - Comments		
	Answer	**Respondent**
	[No Responses]	
Please select your School, College, Institute Affiliation. (You may select more than one) - Other responses		
	Answer	**Respondent**
	DSS	
	Office of Research and Graduate Studies	
Please select your School, College, Institute Affiliation. (You may select more than one) - Comments		
	Answer	**Respondent**
	[No Responses]	
Please select your School, College, Institute Affiliation. (You may select more than one) - Other responses		
	Answer	**Respondent**
	[No Responses]	

Please select your School, College, Institute Affiliation. (You may select more than one) - Comments		
	Answer	**Respondent**
	[No Responses]	

How are your computational research requirements currently being met? - Other responses		
	Answer	**Respondent**
	Outside collaborations	
	work laptop to work from home	

How are your computational research requirements currently being met? - Comments		
	Answer	**Respondent**
	We have used ARSC in the past for running compute-intensive tasks (research using genetic algorithms). The enrollment/access process was a bit cumbersome. We are currently using local GPU-based workstations for our needs.	
	We can run smaller-scale evolutionary computation runs on modern desktop computers. For larger runs, it was helpful to have had access to ARSC resources.	

If your computational research needs are not being met, please specify the unmet needs now and into the future (RAM, CPU,		
	Answer	**Respondent**
	N/A	

	Hardware for data visualization; "true" multiprocessing would be nicer for some problems than the GPU-based processing (CUDA) we are doing now. We really need more human expertise in operating and developing research software.	
	With the broad range of applications for evolutionary optimization of image compression algorithms, my research can only benefit from the availability of massively parallel computational resources. Without such access, we are left attempting to use GPUs or run on a classroom full of PCs to get work accomplished.	
	I am unable to receive large data files via email	

	Basically, the computational resources I need to carry on my research are not available. While we could stumble along doing a limited number of small-scale runs on desktop computers, what we really need are massively parallel supercomputing resources that support distributed computation of image compression and reconstruction transforms for fitness evaluation of thousands of candidate solutions over the course of each evolutionary computing run, and we need to conduct a much larger number of runs.	
	We are using workstations to meet our soil liquefaction analyis. One analysis is taking a fiarly high end pc/workstation (16GB ram and quadcore cpu) about one day to complete. We wish we can cut the computation time to a few hrs.	
What operating system do you typically use for your scientific research? (check all that apply) - Other responses		
	Answer	**Respondent**
	[No Responses]	
What operating system do you typically use for your scientific research? (check all that apply) - Comments		
	Answer	**Respondent**

	Many of our software packages runs only under windows.	
If computing services were provided on campus, which of the following would enhance your research productivity? (Select all that		
	Answer	**Respondent**
	[No Responses]	
If computing services were provided on campus, which of the following would enhance your research productivity? (Select all that		
	Answer	**Respondent**
	Human support for research development	
	I don't know much about how universities should run computing services, server space and control, etc.	
Would you or your laboratory benefit from the availability of basic training courses in using operating systems and programming		
	Answer	**Respondent**
	R, C languages	
What major commercial or open-source software are you using? - Other responses		
	Answer	**Respondent**
	IPA	
	Atlas T/I	
	ArcGIS	
What major commercial or open-source software are you using? - Comments		
	Answer	**Respondent**
	We use R mostly	
What other software packages would enhance your research productivity and what is preventing you from using them now? -		
	Answer	**Respondent**
	Geneious - cost of a license	

	We need the MATLAB Genetic Algorithms and Direct Search toolbox, as well as the Wavelet toolbox. It would be helpful if MATLAB would make it easy to distribute fitness evaluation across several CPUs or GPUs.	
	Windows based GUI	
What are your greatest data management needs? (Select as many as apply) - Other responses		
	Answer	**Respondent**
	[No Responses]	
What are your greatest data management needs? (Select as many as apply) - Comments		
	Answer	**Respondent**
	[No Responses]	
Are your current research needs for data storage being met? - Comments		
	Answer	**Respondent**
	[No Responses]	
What do you anticipate your research data storage requirements will be for temporary datasets? (Please select only one option) -		
	Answer	**Respondent**
	[No Responses]	
What do you anticipate your research data storage requirements will be for temporary datasets? (Please select only one option) -		
	Answer	**Respondent**
	More is better	
What do you anticipate your short term (2-3 years) research data storage requirements will be ? (Please select only one option) -		
	Answer	**Respondent**
	[No Responses]	
What do you anticipate your short term (2-3 years) research data storage requirements will be ? (Please select only one option) -		
	Answer	**Respondent**
	More is better	
What do you anticipate your long term (4 or more years) research data storage requirements will be? (Please select only one		
	Answer	**Respondent**
	[No Responses]	
What do you anticipate your long term (4 or more years) research data storage requirements will be? (Please select only one		
	Answer	**Respondent**
	More is better	

If no, please be as specific as possible and describe your particular area of need, including peak and sustained data rates if		
	Answer	**Respondent**
	I can't download large data files. Sometimes the network is down and I can't use basic internet-based programs like UA Online, Blackboard, etc.	

If "YES," please describe. - Responses		
	Answer	**Respondent**
	I work with scientists in Europe and the lower 48 in creating and analyzing large genomics datasets.	

What types of cyberinfrastructure, resources or system administration would be most useful to you, your group, lab, department,		
	Answer	**Respondent**
	[No Responses]	

What types of cyberinfrastructure, resources or system administration would be most useful to you, your group, lab, department,		
	Answer	**Respondent**
	You're assuming that basic computer users with storage needs know what these teche terms refer to	

What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department,		
	Answer	**Respondent**
	[No Responses]	

What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department,		
	Answer	**Respondent**
	[No Responses]	

Which of the following user support options would you be most interested in? (Check all that apply.) - Other responses		
	Answer	**Respondent**
	[No Responses]	

Which of the following user support options would you be most interested in? (Check all that apply.) - Comments		
	Answer	**Respondent**
	[No Responses]	

Please rank the following research computing funding options. - Comments		
	Answer	**Respondent**

	most of these questions are not relevant to social science research	
Can you suggest an innovative funding option that would work for your research program and sustain the infrastructure and		
	Answer	**Respondent**
	[No Responses]	
What would attract you to change from local/departmental clusters/storage to other providers (on or off campus)? - Other		
	Answer	**Respondent**
	[No Responses]	
What would attract you to change from local/departmental clusters/storage to other providers (on or off campus)? - Comments		
	Answer	**Respondent**
	Still need easy access to the data and sometimes the hardware	
	Useful help and advice in data analysis and storage.	
What are the pros and cons of the services you currently use? - Responses		
	Answer	**Respondent**

	There were two major drawbacks to working on ARSC systems during my recent research: 1. We were not allowed to utilize more than a small number of processors at a time. This limitation made it difficult to collect data for statistical validation of the approach, and hampered our ability to "push the envelope", i.e., to determine an upper-bound on the amount of improvement our approach could obtain compared to the state-of-the-art. 2. Most every time ARSC upgraded to a new version of MATLAB, there were problems. Our code usually had to be modified. Often the new version did not work well with the toolboxes we used. In some ways it would have been better (for us) if ARSC provided a more stable computing set-up.	
	Located far away.	

	My research could truly use as many CPUs as possible. We've typically been given access to no more than a few. We have had to aim at a relatively moving target, and with each MATLAB modification, incompatibilities between toolboxes, lack of support for parallelization, etc. have become problems. Some of these problems can be blamed upon MATLAB.	
	Pros: solely owned by the research group; no hassle;□ □ Cons: high cost; too much time to manage hardware and software, etc.; lack of security	

Please enter any additional comments you would like to add about your research computing needs. - Responses

	Answer	Respondent
	I like having access to our X drive that allows me to safely store secure files on a campus/department network drive, yet access them easily from home.	
	Once somebody figures out a really good way to distribute computations effectively across a large number of GPUs, my need for supercomputing resources will be greatly reduced. GPUs have the potential to outperform supercomputers, and at a fraction of the cost. Major vendors are currently missing a market.	

Appendix C

High Performance Computing (HPC) Consultant Bios and Visit Agenda

UAF Chancellor, Brian Rogers, has tasked the Office of Information Technology with preparing an assessment of the research and academic demands for computational science and high performance computing at UAF. This assessment will evaluate the demands for computing capacity, storage, and support staffing to meet current and projected research and academic demands for high performance computing and make recommendations for its ongoing support.

This assessment will help identify faculty research cyberinfrastructure needs, including software programs, compute cycles, data storage, data management, archiving, networks, facilities space, expertise and programming. To assist the University, we have engaged the knowledge and expertise of the following leaders in research computing:

Amy Apon is Professor of Computer Science and Computer Engineering at the University of Arkansas, Fayetteville and Director of the Arkansas High Performance Computing Center. She holds a Ph.D. in Computer Science from Vanderbilt University, with a research emphasis in the performance analysis and evaluation of parallel and distributed systems. She also has an M.S. in Computer Science, and M.A. in Mathematics, and a B.S.Ed. in Mathematics Education from the University of Missouri-Columbia. Amy served as the co-chair for the NSF Workshop on Sustainable Funding Business Models for High Performance Computing Centers and is the 2011 elected Chair of the Coalition for Academic Scientific Computation (CASC), an organization of more than 60 of the nation's most forward thinking universities and computing centers.

Gerry McCartney has served as Purdue University's chief information officer since June 2006. Under McCartney's leadership, Purdue has developed the nation's largest campus-serving cyberinfrastructure for research, currently with two supercomputers listed in the internationally known Top 500 list. Also during his tenure, Purdue has developed some of the nation's most advanced learning and classroom technologies, including Signals and Hotseat. He was named Purdue's Oesterle Professor of Information Technology in 2009. In 2010 McCartney provided oversight to a campus-wide restructuring of the information technology resources used by the nearly 15,000 faculty and staff on campus. He earned his doctorate in sociology and anthropology from Purdue in 1996 after receiving diplomas in advanced programming and systems analysis from Trinity College in Dublin, Ireland, in 1982 and 1984, respectively. He received his bachelor's and master's degrees in 1981 and 1982 from NUI Maynooth in Ireland

Mike McPherson has served as University of Virginia's Associate Vice President and Deputy Chief Information Officer since July 2006. Prior to his appointment, he served in a number of other information technology leadership roles in higher education, including Director of Information Technology for the College of Literature, Science, & the Arts; Special Counselor to the Provost (both at the University of Michigan); and Interim President and CEO of Merit Network, Inc., the nation's longest-running regional research and education network. He holds a Bachelor of Arts in Multidisciplinary Humanities (Physics, History of Art, Theatre), Michigan State University, 1989 and has completed some graduate work in the area of astrophysics. Mike helped lead the University of Virginia in developing a ground up model for sustainable research computing.

For Questions and Comments please contact:

Jim Durkee, Executive Director Technology Oversight — or — Karl Kowalski, Executive Director User Services

jrdurkee@alaska.edu (907) 450-8315 — kekowalski@alaska.edu (907) 450-8440

OFFICE OF
Information Technology

High Performance Computing (HPC) Consultant Site Visit Agenda
January 12 – 14, 2011
University of Alaska Fairbanks

Wednesday, January 12, 2011

10 -11:30 am	**Orientation and Tour** *Butrovich 102b*
11:45 – 1:15 pm	**Deans and Directors Luncheon** *IARC 501*
1:15 – 2:30 pm	**Walking Tour of West Ridge Research Facilities**
3 – 4 pm	**Video Conference with UAA and UAS Executive Leadership** *UAF/SW Butrovich 212b*
4 – 5 pm	**Audio Conference with Mark Myers, UAF Vice Chancellor of Research** *Butrovich 102d*

Thursday, January 13, 2011

7 – 7:45 am	**Breakfast Meeting** *Princess Riverside Lodge*
9 – 10 am	**President Patrick Gamble** *President's Conference Room*
10:15 – 11:15 am	**Chancellor Brian Rogers** *Chancellor's Office*
12:00 – 1:15 pm	**Faculty Luncheon** *Wood Center Rooms C/D*
1:30 – 2:30 pm	**UAF Cabinet** *Chancellor's Conference Room*
2:30- 3:30 pm	**Audio Conference with Greg Newby, ARSC** *Chancellor's Conference Room*
3:45 – 4:15 pm	**Frank Williams, Director of ARSC** *Butrovich 102d*
7 – 8:30 pm	**Dinner with Chancellor**

Friday, January 14, 2011

7:30 – 8:30 am	**Breakfast with Dan Julius**
8:45 am – 12 pm	**Draft Writing/Debrief** *Butrovich 102d*
12pm – 1pm	**Lunch**

Appendix D

NSF Sustainable Funding and Business Models for Academic Cyberinfrastructure Facilities

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 (SURAgrid), 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 s*ustainable 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. SURAgrid [8] and the Great Plains Network [9] are good examples of such regional collaborations.

SURAgrid's Strategic Plan recognizes the value of regional engagement and collaboration: "The overall intent of the SURAgrid strategic plan is to provide researchers with a cyberinfrastructure strategy that extends local resources through the shared resources of SURAgrid and in doing so provides a gateway to national (and international) infrastructures, and establishes SURAgrid as an integral component of each SURAgrid member's infrastructure solution for competing in the 21st century. This implies a collaborative effort of the SURAgrid 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 SURAgrid. 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 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 then 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 Competiveness – 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/srcc/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 — The Pennsylvania State University
   Director of Research and Cyberinfrastructure

2. Ahalt, Stanley — Renaissance Computing Institute
   Director

3. Allen, Gabrielle — Louisiana State University
   Associate Professor, Computer Science

4. Apon, Amy — University of Arkansas
   Director, Arkansas HPC Center

5. Athanasoulis, Marcos — Harvard Medical School
   Director of Research IT

6. Atkins, Daniel — University of Michigan
   Associate Vice President for
   Research Infrastructure

7. Bangalore, Purushotham — University of Alabama at Birmingham
   Associate Professor and Director
   Collaborative Computing Laboratory

8. Bansil, Arun — Northeastern University
   Professor, Physics

9. Barkai, David — Intel
   HPC Computational Architect

10. Bose, Rajendra — Columbia University
    Manager, Research Computing Services

11. Bottum, James — Clemson University
    Vice Provost and CIO

12. Bozylinkski, Garrett — University of Rhode Island
    CIO

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 — Louisiana State University
    Director of HPC

51. Lombardi, Julian — Duke University
    Assistant Vice President
    Office of Information Technology

52. Majchrzak, Daniel — University of South Florida
    Director of Research Computing

53. Marinshaw, Ruth — University of North Carolina-Chapel Hill
    Assistant Vice Chancellor for
    Research Computing

54. Marler, Bryan — Hewlett-Packard
    Director, High Performance Computing

55. McMullen, Donald — University of Kansas
    Director of Research Computing

56. Mehringer, Susan — Cornell University
    Assistant Director, Consulting
    Cornell Center for Advanced Computing

57. Miller, Therese — Indiana University
    Chief Operating Office
    Pervasive Technology Institute

58. Monaco, Greg — Great Plains Network (GPN)
    Director for Research and Cyberinfrastructure

59. Moore, Richard — San Diego Supercomputer Center
    Deputy Director

60. Nabrzyski, Jaroslaw — University of Notre Dame
    Director, Center for Research Computing

61. Neeman, Henry — University of Oklahoma
    Director OU Supercomputing Center
    for Education & Research

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. Wrobleski, 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 — North Dakota State University
Business Manager & Acting Director
Center for High Performance Computing

14. Heslop, Janet — Cornell University
IT Associate Director

15. Johnson, Del — South Dakota State University
Director, Research Computing

16. Joiner, David — Kean University
Assistant Professor

17. Jung, Gary — Lawrence Berkeley National Laboratory
Manager, High Performance Computing Services

18. Meacham, Steve — National Science Foundation
Senior Staff Associate

19. Mendan, Joseph — Rochester Institute of Technology
Financial Reporting Specialist/CIMS

20. Mort, Brendan — University of Rochester
HPC Consultant & Computational Scientist

21. Nordquist, Russell — University of Rochester
HPC System Administrator

22. O'Leary, Patrick — Idaho National Laboratory
Director, Center for Advanced Modeling and Simulation

23. Poole, Michael — University of Massachusetts
Associate Vice President and
Chief Technology Officer

24. Pormann, John — Duke University
Director of Scalable Computing Support Center

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-Industial-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.

Appendix E

Resources Used in Preparation of this Report

Appendix E

Resources used in preparation of this report

Personal Interviews
Brian Barnes, Director, UAF Institute of Arctic Biology
Jock Irons, Institute of Arctic Biology
Larry Hinzman, Director, UAF International Arctic Research Center
Roger Smith, Director, UAF Geophysical Institute
Carl Tape, Assistant Professor of Geophysics, UAF Geophysical Institute
Dan White, Institute of Northern Engineering
Terry Whitledge, Director, Institute of Marine Science
Frank Williams, Director, Arctic Region Supercomputing Center
Paul Layer, Dean College of Natural Sciences and Mathematics
Nichole Molders, Department Chair, professor, Atmospheric Sciences
David Newman, professor, Space Physics and Aeronomy
Carol Lewis, Dean, School of Natural Resources & Agricultural Sciences
Doug Goering, Dean, UAF College of Engineering & Mines
Virgil (Buck) Sharpton, Director, Geographic Information Network of Alaska
Barbara Horner-Miller, Assistant Director, Arctic Region Supercomputing Center
Virginia Bedford, Assistant Director, Arctic Region Supercomputing Center
Greg Newby, Chief Scientist, Arctic Region Supercomputing Center
Scott Arko, Deputy Director, Alaska Satellite Facility
Paul McCarthy, Interim Dean, Rasmusson Library
Jon Genetti, Associate Professor, Computer Science
Fred Smits, Executive Director, Infrastructure Services, UA Statewide Office of Information Technology
Dale Denny, Manager, Data Center Operations, UA Statewide Office of Information Technology
Dave Lifka, Director, Cornell Center for Academic Computing
Amy Apon, Director, Arkansas Supercomputing Center
Vijay Argawala, Senior Director of Research Computing and Cyberinfrastructure, Penn State
Jason Holt, IT Services, Penn State
Michael Phen, IT Services Penn State
Michael McPherson, Deputy CIO and Vice-Provost, University of Virginia
Gerry McCartney, Oesterle Professor, CIO, Purdue University
Willliam Whitson, Directory of Research Support, Purdue University
Brian Rogers, Chancellor, University of Alaska Fairbanks
Mark Myers, Vice-chancellor for Research, University of Alaska Fairbanks
Patrick Gamble, President, University of Alaska
Thomas Hauser, Director Research Computing, University of Colorado at Boulder
Craig Stewart, Executive Director, Pervasive Technology Institute, Indiana University
Gary Miksik, Project Manager, Pervasive Technology Institute, Indiana University

Cherri Pancake, Professor, School of Electrical Engineering and Computer Science, Oregon State
Erik Deumens, Directory of Computing, University of Florida
Timothy Thomas, Deputy Director for Advanced Research Computing, University of New Mexico
Terry Herdman, Associate Vice President for Research Computing, Virginia Tech
Nicholas Polys, Director of Visual Computing, Virginia Tech
James Myers, Director, Computational Center for Nanotechnology Innovations, Rensselaer Polytechnic Institute
Jacqueline Stampalia, Associate Director Research Computing, Computational Center for Nanotechnology Innovations, Rensselaer Polytechnic Institute
Kelly Gaither, Director Data and Information Analysis, Texas Advanced Computing Center
Aaron Dubrow, External Relations, Texas Advanced Computing Center
Dan Stanzoine, Deputy Director, Texas Advanced Computing Center
Jerry Perez, Research Associate, Texas Tech University
Maureen Dougherty, Director, High Performance Computing and Communications Information Technology Services, University of Southern California
John van der Zwaag, Research Associate, Mississippi State University
Matthew Jones, Associate Director and Lead Computational Scientist, University at Buffalo
James Bottum, Vice Provost for Computing and Information Technology, Clemson University
Ralph Roskies, Scientific Director, Pittsburgh Supercomputing Center

Forums:

Deans and Directors Luncheon

Kevin Berry, School of Management
Anita Hartmann, College of Liberal Arts
Jim Long, International Arctic Research Center
Steve Sparrow, School of Natural Resources and Agriculture Science
Doug Goering, College of Engineering and Mines
Paul Layer, College of Natural Sciences and Mathematics
Dan White, College of Engineering and Mines
Paul McCarthy, Rasmusson Library
Nettie Labelle-Hamer, Alaska Satellite Facility/Geophysical Institute
Steve Smith, Office of Information Technology
Karl Kowalski, Office of Information Technology
Jim Durkee, Office of Information Technology
Julie Larweth, Office of Information Technology

Faculty Luncheon

Ronald Daanen, Institute of Northern Engineering
Chris Fallen, Arctic Region Supercomputing Center
Georgina Gibson, International Arctic Research Center

Roger Hansen, Geophysical Institute
John Haverlack, School of Fisheries and Ocean Sciences
Shawn Houston, Institute of Arctic Biology
Jock Irons, Institute of Arctic Biology
John Keller, College of Natural Sciences and Mathematics/Chemistry
Constantine Khroulev, Geophysical Institute
Zygmunt Kowalik, School of Fisheries and Ocean Sciences
Jonah Lee, College of Engineering and Mines
James Long, International Arctic Research Center
Andres Lopez, School of Fisheries and Ocean Sciences/MOTN
Don Morton, Arctic Region Supercomputing Center
Chung-Sang Ng, Geophysical Institute
Oralee Nudson, Arctic Region Supercomputing Center
Jacques Philip, Center for Alaska Native Health Research/IAB
Elena Suleimani, Geophysical Institute
Carl Tape, Geophysical Institute
Peter Webley, Geophysical Institute
Essam Yesin, Geophysical Institute
Xiandong Zhang, International Arctic Research Center
Vladimir Alexeev, International Arctic Research Center
Rui Han, International Arctic Research Center
Robert Torgerson, Arctic Region Supercomputing Center
Uma Bhatt, College of Natural Sciences and Mathematics/GI
Nicole Moelders, GI/DAS/CSM
Jenny Hutchings, International Arctic Research Center
Sergei Maurits, Arctic Region Supercomputing Center
Kenneth Coyle, Institute of Marine Science
Steve Smith, Office of Information Technology
Karl Kowalski, Office of Information Technology
Jim Durkee, Office of Information Technology

UAA, UAS Videoconference:

Rich Whitney, CIO, Vice-Provost, University of Alaska Anchorage
Doug Causey, Chair Biological Sciences, University of Alaska Anchorage
Kim Peterson Ph.D., Associate Dean for Research, University of Alaska Anchorage
Kenrick Mock, UAA Associate Professor of Computer Science
Shannon Atkinson, professor, UAF School of Fisheries and Ocean Sciences
Tony Gharrett, professor, UAF School of Fisheries and Ocean Sciences

Reports

NSF, (2010) Sustainable Funding and Business Models for Academic Cyberinfrastructure Facilities, Lane, Cara, Janice Fournier, and Tom Lewis, with Greg Koester, Ed Lazowska, Erik Lundberg, Henry Lyle III, and William Washington, “Scientific Advances and Information Technology: Meeting Researchers’ Needs in a New Era of Discovery,” ECAR Case Study 3, 2010

U.Va. Task Force on Information Technology Infrastructure, (2006), “AN INITIATIVE TO SUPPORT COMPUTATIONAL SCIENCE AT THE UNIVERSITY OF VIRGINIA,” Retrieved from: http://www.virginia.edu/vpcio/documents/ITRTF_report_9JULY06.pdf

Jokl, Jim, Terry Lockard, Tim Sigmon, Martha Sites, Diane Walker, (2006), “A Technology Framework for Enabling Scholarship and Research at the University of Virgina. Retrieved from: http://www.itc.virginia.edu/oit/org/Technology-Framework.pdf

Michaels, George, Reagan Moore, James Huesmann, Jeff Yauney, (2006)“University of Alaska Information Technology Review Needs and Assessment Document, Available at: https://sites.google.com/a/alaska.edu/ua-research-computing-needs/home/files

NSF, Cyberinfrastructure Council, (2007), “Cyberinfrastructure Vison for 21st Century Discovery,” Retrieved from: http://www.nsf.gov/pubs/2007/nsf0728/index.jsp

President’s Information Technology Advisory Committee, (2005). “Computational Science: Ensuring America’s Competitiveness,” Retrieved from: http://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf

Erik Deumens, Erik, Charles Taylor, (2010), Experiences with sustainability at the UF HPC Center,” High Performance Computing Center, University of Florida, Retrieved from: http://www.cac.cornell.edu/~lifka/Downloads/SRCC/18.UF-HPCC-sustainability.pdf

Dunning, Thom Jr., (2010) “Congressional Testimony:Subcommittee on Research and Science Education Committee on Science and Technology, U.S. House of Representatives, Hearing on “The State of Research Infrastructure at U.S. Universities” Retrieved from: http://www.casc.org/papers/THDJr_CongressionalTestimony_Feb10.pdf

Hacker, Thomas, Bradley Wheeler, (20017), “Making Research Cyberinfrastructure a Strategic Choice,” Educause Quarterly. Retrieved from: http://www.educause.edu/EDUCAUSE+Quarterly/EDUCAUSEQuarterlyMagazineVolum/MakingResearchCyberinfrastruct/157439

Stewart, Craig, Stephen Simms, Beth Plale, Matthew Link, David Hancock, Geoffrey Fox, (2010) “What is Research Cyberinfrastructure?” Indiana University, personal copy provided by Craig Stewart, available at: http://portal.acm.org/citation.cfm?id=1878347

State of Alaska, (2010), “Enterprise IT Strategic Plan,” Retrieved from: http://doa.alaska.gov/eib/EnterpriseStrategicPlan2010_2012Draft.pdf

National Oceanic and Atmospheric Administration, (2008), "High Performance Computing Strategic Plan for FY2011-2015," retrieved from: http://www.cio.noaa.gov/HPCC/pdfs/HPC_Strategic_Plan.pdf

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

Committee on Institutional Cooperation, (2010), A research Cyberinfrastructure Strategy for the CIC: Advice to the Provosts from the Chief Information Officers, Retreived from: http://www.cic.net/Libraries/Technology/2010Report_-_6_21reduced.sflb.ashx

University of California San Diego, (2009), Blueprint for the Digital University, UCSD research cyberinfrastructure, Retreived from: http://research.ucsd.edu/documents/rcidt/RCIDTReportFinal2009.pdf

National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure, (2003), Revolutionizing Science and Engineering Through Cyberinfrastructure, Retreived from: http://www.nsf.gov/od/oci/reports/atkins.pdf

Voloudakis, J. (2010). A systemic model for IT economic sustainability. EDUCAUSE Center for Applied Research, bulletin 1.

Berman, F. (2008). Making research and education cyberinfrastructure real. EDUCAUSE Review, vol. 43, no.4.

Lifka, D. (2010). High Performance Computing Center Sustainability NSF Workshop. Invited presentation at the Coalition for Academic Scientific Computation (CASC) Fall 2009 meeting.

Education Advisory Board (2009). University business executive roundtable agenda - national best practice initiatives 2009-2010.

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

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 .
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 .

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

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

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

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

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 .

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 .

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 .

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 .

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 .

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

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

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

.

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 .

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 .

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 .

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

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 .

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

.

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 .

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 .

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

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 .

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 .

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 .

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-Industial-Partnerships.pdf

Appendix F

Chancellor Roger's Memo on High Performance Academic/Research Computing Assessment

UNIVERSITY OF ALASKA FAIRBANKS

Office of the Chancellor

320 Signers' Hall, P.O. Box 757500, Fairbanks, Alaska 99775-7500
www.uaf.edu

Brian Rogers, Chancellor
(907) 474-7112
(907) 474-6795 fax
uaf.chancellor@alaska.edu

Bob Shefchik, Executive Officer
(907) 474-7489
(907) 474-7475 fax
rrshefchik@alaska.edu

MEMORANDUM

Date: October 22, 2010

To: Steve Smith, Chief Information Technology Officer

From: Brian Rogers, Chancellor

RE: High Performance Academic/Research Computing Assessment

Per this memo I am directing the Office of Information Technology to undertake an assessment of the research and academic demands for high performance computing at UAF. This assessment should evaluate the demands for computing capacity, storage space, and support staffing necessary to meet current and projected research and academic demands for high performance computing.

UAF presently has significant assets currently available through NSF PACMAN award, the follow-up MRI award, and the Arctic Region Super Computing Center. Past efforts at assessing the supercomputing program have focused more on what it will take to run and maintain existing assets rather than what UAF needs to have in place to meet the service demands of the research and academic programs. This effort should to come at the question from the service demand side of the issue.

This project will be run and managed by OIT. Input from researchers and staff will be critical in determining the needs based upon the wide array of work conducted at UAF. Funding will be made available from UAF to fund external costs incurred by OIT in accomplishment of this task. This task is on a tight time line, as decisions on hardware, staffing, and organizational structure could need to be made in time for a potential change June 1, 2011. Hence, I would like the review accomplished and recommendations provided by February 15, 2011.

Bob Shefchik will be available to facilitate timely responses to any information or data requests that OIT will need to accomplish this task.

BDR

cc: Bob Shefchik