XSEDE

eXtreme Science and Engineering Discovery Environment

20012 Q1: January 1, 2011, through March 31, 2011

<u>XSEDE Senior Management Team</u>

John Towns (NCSA)	PI and Project Director
Jay Boisseau (TACC)	Co-PI and Users Services Director
Ralph Roskies (PSC)	Co-PI and Extended Collaborative Support Service-Projects Director
Nancy Wilkins-Diehr (SDSC)	Co-PI and Extended Collaborative Support Service-Communities Director
Robert Harrison (UT-NICS)	Interim Co-PI
Victor Hazlewood (UT-NICS)	Interim Operations Director
Scott Lathrop (Shodor)	Education and Outreach Director
Kathlyn Boudwin (ORNL/NICS)	Manager of Project Management
Janet Brown (PSC)	Manager of Systems and Software Engineering
Kurt Wallnau (SEI)	Manager of Software Development and Integration
TBD	Chair of the User Advisory Committee
Carol Song (PU)	Interim Chair of the XD Service Providers Forum

XSEDE QUARTERLY REPORT

Contents

1 Overview 10

1.1 Project Context 10

1.2 Project Highlights 12

2 Science and Engineering Highlights 15

2.1 Advanced Scientific Computing: Studying Transportable Explosives with Uintah (Martin Berzins, University of Utah) 15

2.2 Atmospheric Sciences: High-Performance Computing for Ensemble-Based Cloud-Resolving Hurricane Analysis, Prediction, and Predictability (Fuqing Zhang, Pennsylvania State University) 16

2.3 Atmospheric Sciences: Studying Magnetic Reconnection (Homa Karimabadi, University of California-San Diego; and William Daughton, Los AlamosNational Laboratory) 16

2.4 Chemical, Thermal Systems: Numerical Modeling of Dense Particulate Flows in Fluidized Bed Reactors (Olivier Desjardins, Cornell University) 17

2.5 Climate Science: Simulating Regional Climate Change, a California Case Study (Alex Hall, UCLA) 18

2.6 Genomics: Identification of Lignocellulosic Enzymes through Sequencing and Assembly of Metagenome Soil Enrichment (Mostafa Elshahed, Rolf Prade & Brian Couger, Oklahoma State University) 19

2.7 Organic Chemistry: Tuning the Optical and Electronic Properties of Conjugated Polymers (Aimée Tomlinson, North Georgia College & State University) 20

2.8 Systematic and Population Biology: Large-scale Phylogenetic Estimation (Tandy Warnow, The University of Texas at Austin) 21

3 XSEDE Project Office 1.1 23

3.1 Overview 23

3.2 Project Management and Reporting 1.1.1 24

3.3 Systems and Software Engineering 1.1.2 24

3.4 Architecture and Design 1.1.3 25

3.5 External Relations 1.1.4 28

3.6 Industry Relations 1.1.5 29

3.7 Software Development and Integration 1.1.6 29

4 XSEDE Operations 1.2 31

4.1 Overview 31

4.2 Security 1.2.1 31

4.3 Data Services 1.2.2 32

4.4 XSEDEnet 1.2.3 33

4.5 Software Support 1.2.4 ... 35
4.6 Accounting and Account Management 1.2.5 ... 36
4.7 Systems Operational Support 1.2.6 ... 36
5 User Services 1.3 ... 41
5.1 Overview ... 41
5.2 Training 1.3.1 ... 41
5.3 User Information & Interfaces 1.3.2 ... 41
5.4 User Engagement 1.3.3 ... 42
5.5 Allocations 1.3.4 ... 43
6 Extended Collaborative Support Service – Projects 1.4 ... 45
6.1 Overview ... 45
6.2 Extended Research Teams Support 1.4.1 ... 46
6.3 Novel and Innovative Projects 1.4.2 ... 48
7 Extended Collaborative Support Service – Communities 1.5 ... 50
7.1 Overview ... 50
7.2 Extended Support for Community Codes 1.5.1 ... 50
7.3 Extended Science Gateways Support 1.5.2 ... 51
7.4 Extended EOT Support 1.5.3 ... 56
8 Education and Outreach 1.6 ... 57
8.1 Overview ... 57
8.2 Education 1.6.1 ... 57
8.3 Outreach 1.6.2 ... 61
8.4 Community Requirements 1.6.3 ... 64
8.5 E&O Infrastructure 1.6.4 ... 67
8.6 Campus Bridging 1.6.5 ... 67
9 TAIS/Audit Services 1.7 ... 69
10 TAIS/Technology Insertion 1.8 ... 76
10.1 Highlights ... 76
10.2 XD – TIS ... 76
11 ExTENCI ... 79
11.1 Overview ... 79
12 XD Service Provider Reports ... 83
12.1 Overview ... 84
13 FutureGrid Service Provider Quarterly ... 85
13.1 Executive Summary ... 85

13.2 Science Highlights 85
13.3 User-facing Activities 87
13.4 Security 89
13.5 Education, Outreach, and Training Activities 89
13.6 SP Collaborations 91
13.7 SP-Specific Activities 91
13.8 Publications 91
13.9 Metrics 92
14 Indiana University - Service Provider Report 100
14.1 Executive Summary 100
14.2 Science Highlights 101
14.3 User-facing Activities 102
14.4 Security 105
14.5 Education, Outreach, and Training Activities 105
14.6 SP-Specific Activities 106
14.7 Publications 108
14.8 Metrics 108
15 LONI/LSU - Service Provider Quarterly Report 110
15.1 Executive Summary 110
15.2 Science Highlights 110
15.3 User-facing Activities 110
15.4 Security 111
15.5 Education, Outreach, and Training Activities 111
15.6 SP Collaborations 111
15.7 SP-Specific Activities 111
15.8 Publications 111
15.9 Metrics 111
16 NCSA - Service Provider Quarterly Report 113
16.1 Executive Summary 113
16.2 Science Highlights 113
16.3 User-facing Activities 114
16.4 Security 115
16.5 Education, Outreach, and Training Activities 115
16.6 SP Collaborations 115
16.7 SP-Specific Activities 115

16.8 Publications 116
16.9 Metrics 117
17 NICS - Service Provider Quarterly Report 127
17.1 Executive Summary 127
17.2 Science Highlights 128
17.3 User-facing Activities 133
17.4 Security 134
17.5 Education, Outreach, and Training Activities 134
17.6 SP Collaborations 135
17.7 SP-Specific Activities 136
17.8 Publications 137
17.9 Metrics 139
18 Pittsburgh Supercomputing Center - Service Provider Quarterly Report 150
18.1 Executive Summary 150
18.2 Science Highlights 153
18.3 User-facing Activities 154
18.4 Security 156
18.5 Education, Outreach, and Training Activities 156
18.6 SP Collaborations 157
18.7 SP-Specific Activities 158
18.8 Publications 158
18.9 Metrics 159
19 Purdue University - Service Provider Quarterly Report 166
19.1 Executive Summary 166
19.2 Science Highlights 167
19.3 User-facing Activities 170
19.4 Security 171
19.5 Education, Outreach, and Training Activities 171
19.6 SP Collaborations 172
19.7 SP-Specific Activities 172
19.8 Publications 173
19.9 Metrics 173
20 San Diego Supercomputer Center (SDSC) Service Provider Quarterly Report 195
20.1 Executive Summary 195
20.2 Science Highlights 197

20.3 User-facing Activities ... 197
20.4 Security ... 199
20.5 Education, Outreach, and Training Activities ... 199
20.6 SP Collaborations ... 202
20.7 SP-Specific Activities ... 204
20.8 Publications ... 205
20.9 Metrics ... 206
21 TACC - Service Provider Quarterly Report ... 226
21.1 Executive Summary ... 226
21.2 Science Highlights ... 229
21.3 User-facing Activities ... 234
21.4 Security ... 235
21.5 Education, Outreach, and Training Activities ... 235
21.6 SP Collaborations ... 241
21.7 SP-Specific Activities ... 241
21.8 Publications ... 244
21.9 Metrics ... 244
A XSEDE Project Milestones Update ... 267
B XSEDE Schedule with Progress Update ... 268
C XSEDE Risk Register ... 287
D XSEDE Change Control Report ... 292
E Metrics ... 293
E.1 XSEDE Resource and Service Usage Metrics ... 293
E.2 XSEDE Program Metrics ... 303
F XSEDE Publications Listing ... 314
F.1 XSEDE Staff Publications ... 314
F.2 Publications from XSEDE Users ... 314
ASC070031 ... 314
AST020001 ... 315
AST090005 ... 315
AST090105 ... 315
AST090107 ... 316
AST100014 ... 316
AST110016 ... 316
ATM050012 ... 317

ATM050014 317
ATM70007 317
ATM110006 317
CHE040014 318
CHE090047 318
CHE100005 319
CHE100033 319
CHE100039 319
CHE100123 319
CHE110011 319
CHE120020 319
CHE120021 320
CHE120025 320
CTS030025 320
CTS050054 320
CTS070066 321
CTS100024 321
CTS100055 321
CTS110027 321
CTS110029 321
CTS120015 322
DBS110003 322
DMR090071 322
DMR100033 322
DMR110037 323
DMS060014 323
EAR100012 324
EAR100021 324
EAR120003 324
IRI120005 324
MCA01S007 324
MCA01S018 325
MCA01S026 325
MCA03S007 326
MCA08X014 327

MCA08X013 .... **327**
MCA08X004 .... **327**
MCA08X003 .... **327**
MCA08X002 .... **328**
MCA06N007 .... **328**
MCA06N038 .... **328**
MCA07S005 .... **329**
MCA07S017 .... **329**
MCA08X032 .... **330**
MCA08X034 .... **330**
MCA93SO13 .... **331**
MCA94P014 (could not paste pubs list from PDF) **332**
MCA94P018 .... **333**
MCA94P030 .... **333**
MCA95C014 .... **334**
MCA98N017 .... **334**
MCA98N020 .... **336**
MCA99S007 .... **336**
MCB040002 .... **336**
MCB060006 .... **336**
MCB070038 .... **338**
MCB070118 .... **340**
MCB080013 .... **340**
MCB080071 .... **340**
MCB090003 .... **340**
MCB090108 .... **341**
MCB090131 .... **341**
MCB100049 .... **341**
MCB100065 .... **341**
MCB100146 .... **341**
MCB110056 .... **342**
MCB120047 .... **342**
MSS080019 .... **343**
PHY100035 .... **343**
PHY110010 .... **343**

PHY110015 .... **343**
PHY110016 .... **343**
PHY110018 .... **344**
PHY120016 .... **344**
G XSEDE Citation Report **345**
H XSEDE EOT Event Details **346**

1 Overview

The Extreme Science and Engineering Discovery Environment (XSEDE) is the most advanced, powerful, and robust collection of integrated digital resources and services in the world. It is an integrated cyberinfrastructure ecosystem with singular interfaces for allocations, support, and other key services that researchers can use to interactively share computing resources, data, and expertise.

1.1 Project Context

Scientists, engineers, social scientists, and humanities experts around the world—many of them at colleges and universities—use advanced digital resources and services every day. Computational technologies and resources such as supercomputers, visualization systems, storage systems and collections of data, software, and networks are critical to the success of those researchers, who use them to advance our understanding of our world, and to make our lives healthier, safer, and better. XSEDE integrates these resources and services, makes them easier to use, and helps more people use them. XSEDE currently supports 16 supercomputers and high-end visualization and data analysis resources across the country.

Digital services, meanwhile, provide users with seamless integration to NSF's high-performance computing and data resources. XSEDE's integrated, comprehensive suite of advanced digital services is developing and implementing tools, methods, and policies to federate with other high-end facilities and with campus-based resources, serving as the foundation for a national cyberinfrastructure ecosystem. Common authentication and trust mechanisms, global namespace and filesystems, remote job submission and monitoring, and file transfer services are examples of XSEDE's advanced digital services. XSEDE's distributed systems architecture allows open development for future digital services and enhancements.

XSEDE also provides the expertise to ensure that researchers can effectively use the supercomputers and tools. Those include:

- Extended Collaborative Support that includes teaming with individual research groups or with research communities to extend their capabilities.
- An advanced hardware and software architecture rooted in user requirements and hardened by systems engineering that allows for individualized user experiences, consistent and enduring software interfaces, improved data management, and ways for campus resources to be transparently integrated into the overall XSEDE infrastructure.
- The XSEDE User Portal, a web interface that allows users to monitor and access XSEDE resources, manage jobs on those resources, report issues, and analyze and visualize results.
- Coordinated allocations of NSF's high-end resources and digital services.
- A powerful and extensible network, in which each XSEDE service provider is connected to a Chicago-based hub at 10 gigabits per second and has a second 10 gigabit-per-second connection to another national research and education network.
- Specialized community-provided services that serve a particular function and allow for rapid innovation and experimentation.
- Advanced cybersecurity to ensure that XSEDE resources and services provide confidentiality, integrity and availability of information
- Training, Education, and Outreach efforts that expand the scope and scale of activities to foster greater community participation in XSEDE-based projects through curriculum development, live and web-based training offerings, outreach at professional society meetings, and engagement of under-represented faculty and students.
- Advanced support for novel and innovative projects.

- A fellowship program that brings Campus Champions working closely with Extended Collaborative Support Service staff on user identified challenges for up to a year.
- The Technology Insertion Service, which allows researchers to recommend technologies for inclusion in the XSEDE infrastructure and enables the XSEDE team to evaluate those technologies and incorporate them where appropriate.

1.1.1 Communities Served

The national, and global, user community that relies on XSEDE for HPC resources has grown tremendously. XSEDE continued to see increased HPC resource user numbers in Q1 2012, with 1,063 new users added in the quarter. The number of open individual accounts reached 6,313, and the number of non-gateway individuals charging jobs was 2,165. An additional 1,039 users submitted jobs via science gateways—that is, nearly 32% of XSEDE users submitting jobs worked through gateways; but this represents a decrease of over 350 XSEDE users submitting jobs via gateways over the previous quarter. As a result, we initiated a thorough analysis of the gateway usage statistics which is ongoing at the time of writing this report. Counting current individual accounts and gateway users together, the XSEDE community numbered 7,352 users.

Further details can be found Appendix E.

1.1.2 XSEDE's Integrated, Distributed Environment

XSEDE is taking on the difficult but necessary task of documenting a clearly specified architectural design for its distributed systems architecture. Given the nature of the end game of the proposal competition that ultimately resulted in the XSEDE award, the project has had to substantially redesign the architecture originally proposed in order to incorporate innovative and important elements of the previously competing proposal. While this has been difficult and has led to some confusion, the project is making progress in this area and will begin to produce design documents that specify the architecture in detail during the coming months.

1.1.3 Project Governance

The XSEDE project has established an organizational structure and governance that promotes efficient and effective project performance. As this is a distributed project involving 17 partner institutions and with many other stakeholders including NSF, and thousands of users, it was necessary to establish a governance model that balances efficiency and inclusiveness. The XSEDE governance model has strong central management to provide rapid response to issues and opportunities, delegation and decentralization of decision-making authority, openness to genuine stakeholder participation, and improved professional project management practices including formal risk management and change control.

The XSEDE governance model is geared towards inclusion of, and responsiveness to, users, service providers, and the NSF scientific community. The various stakeholders have input through three distinct advisory bodies, which have direct access to the XSEDE Project Director and the XSEDE senior management team through regularly scheduled meetings. In order to remain well informed of the requirements of the user community, XSEDE leadership receives advice and counsel from the User Advisory Committee, the XD Service Providers Forum, the XSEDE Advisory Board, and the TEOS Advisory Committee. These advisory committees are intimately involved with XSEDE management in guiding the project towards optimal operations, service, and support for users.

The XSEDE project is managed by a senior management team consisting of the PI/Project Director as chair, the co-PIs and key leaders of major areas of the XSEDE project, the Chair of the User Advisory Committee, and the Chair of the XD Service Providers Forum. This team is constituted from those responsible for the day-to-day operation of the project and is the highest-level management body in the organization. In order to be responsive to both the user community and the set of Service Providers with whom we will collaborate, the chairs of the User Advisory Committee and the XD Service Providers Forum are members of this team.

1.2 Project Highlights

At the close of the third quarter of the first year of XSEDE, the project has begun to really get traction across a wide range of activities, and in particular areas that were new to XSEDE vs. TeraGrid. In many areas we have already met many of our annual goals although we are not through the year. Overall progress is quite good with some exceptions noted in the details in various sections of this report.

We have made significant progress in our project governance. This quarter we completed the lengthy process of formally defining the XD Service Providers Forum by working closely with a set of founding Service Providers. Though this resulted in something a bit different than originally planned, what has been created is of broader scope and far more inclusive. The *XD Service Providers Forum: Charter, Membership, and Governance* document can be found on the Project Documents page[1]. In addition we have formed the XSEDE Advisory Board (XAB) and the User Advisory Committee (UAC). Both will have initial meetings in the coming quarter. Representatives from all of these will also be determined to participate in our governance. The TEOS Advisory Committee has also been formed and we have developed a document describing the role and responsibilities of the TEOS Advisory Committee. These activities put in place the multiple close engagements of stakeholders in the governance of XSEDE that we planned from the start.

XSEDE continues its high level of support of researchers with our Extended Collaborative Support Services staff engaged in engaged in over 80 active projects during this past quarter covering a variety a areas. The requests for such support continue to be very strong via the allocation request process. Of particular note, during this quarter we have our first example of executing our flexible hires strategy to bring on board expertise to assist prospective users that XSEDE does not already have—in this case in the area of Digital Humanities. We are currently looking to identify similar assistance with High Throughput Computing (e.g. Condor or OSG) experience.

Through these engagements we continue to support important science and engineering advances; we provide highlights on 8 additional success supported by XSEDE across a variety of domains in Section 2 of this report.

Also of note, the Campus Champion Fellows program application process is now complete and we will be opening the application process shortly. This is complemented by the Campus Champions program is showing strong signs of community building among the Champions. There are now 112 campuses and 153 Champions. In the current quarter, the group has initiated over a dozen conversations addressing challenges they have found in performing their duties as Champions and in supporting their campus users. What is striking is that the Champions have begun to effective support one another with less reliance on XSEDE staff to answer questions and provides support marking an important transition for this community.

XSEDE's Allocations team continues to coordinate peer-review of XSEDE's over-requested computing resources, and began discussing storage allocations for implementation. Requests totaling 807M SUs were reviewed at the March 2012 XRAC meeting with only 397M SUs available. XSEDE Allocations staff and XSEDE site representatives are faced with a difficult

[1] https://www.xsede.org/web/guest/publications

situation of continued requests for supplements, transfers, and startups but no significant pool of SUs to satisfy the user community requests!

XSEDE's online presence continues to grow and be enhanced. Usage of the XSEDE web site and XSEDE User Portal (XUP) have continued to increase as new technical documentation was published and new capabilities were added to the XUP. The XUP now offers improved allocations management and user profiles capabilities.

Training continues to be very effective reaching over 2,700 users and potential users in this quarter period via in-person, online, and webcast training workshops. In addition, we are collaborating with the Blue Waters project to develop this year's petascale symposium that had been started by Blue Waters and TeraGrid. It will consist of 4 invited talks and at least 8 contributed papers and will be held in Chicago just preceding the XSEDE12 conference.

In our growing efforts in Campus Bridging Team, we engaged in two major activities during this quarter: definition of use cases for the Architecture and Design team and a pilot project for testing the Global Federated File System (GFFS) software on campuses. A set of seven use cases for campus bridging activities have been documented covering the use of resources by campus researchers, data movement between researchers and XSEDE resources, and making local resources available to the larger community via XSEDE in exchange for other resources and are being used to drive the architectural design process.

This quarter has seen significant progress in our ongoing efforts to redefine the distributed systems architecture for XSEDE. There has been a particular focus on developing the architecture definition to support our campus bridging efforts. We have also successfully delivered all three major new capabilities (GO-Data, EMS, GFFS) from the Software Development and Integration team to Operations for acceptance testing; issues raised for two of these components (GO-Data, EMS) have resulted in improvements to those services for the providers. These will progress through the operational readiness and acceptance testing in the next few months.

Throughout the quarter we have successfully put more of our planning and management processes in place and as more of them are available we have increasingly been making use of them to improve how the project is managed and how we execute our work. This quarter also has seen the first publicly available documents emerge on the xsede.org website. These are all on the web site linked from the Project Documents page[2]. We will be posting many more such documents as they become available.

In the area of XSEDE Operations we have finalized and received approval by the International Grid Task Force's TAGPMA certification body[3] of our initial plans for an XSEDE Certificate Authority and enhancements to security that include the completion of security awareness training materials and plans to leverage the REN-ISAC Security Event Service. Testing of a new network-monitoring tool has been completed and transitioned operational monitoring to the XSEDE Operations Center. This will leverage perfSONAR Services at PSC, IU, and NCAR installed to perform regular throughput test between sites. We have defined the criteria for selecting a technology for the XSEDE-wide file system which will be used in the next quarter to make a technology decision and begin planning for deployment of this capability. Operational readiness reviews and acceptance tests were completed for the Globus Online file transfer service and Execution Management Services.

---

[2] https://www.xsede.org/web/guest/project-documents

[3] http://www.tagpma.org/

Overall, we achieved over 97% uptime for all production central services while fielding over 2,600 tickets of which nearly 40% were closed within 2 business days. The User Engagements team has completed its evaluation and recommendation of a new ticket system. Expected to be approved in Q2CY2012, the new system will further improve user issue analysis, and thus drive improvements to all of the other user services activities (web documentation, XUP, training, allocations, etc.) as well as XSEDE resources.

As part of the education program, we have been working on several efforts to create educational materials and present educational opportunities for faculty and students. We have also been planning for a number of workshops for faculty this spring and summer. Two one-day workshops providing an overview of computational science educational resources and tools for community college faculty in Nevada are confirmed for college sites in Las Vegas and Elko and are being co-sponsored by the Nevada EPSCoR office.

Over the past quarter, we have continued work with the campuses interested in formal computational science programs. We have been engaging with the campuses through audio conferences and campus visits with work continuing on the development of model competencies for both undergraduate and graduate programs in computational science.

As might be expected, the XSEDE Project Office has been the driver for Year 2 planning activities throughout the quarter coordinating the effort from the Project Management Team. This will be an ongoing effort throughout the following Quarter also.

2 Science and Engineering Highlights

2.1 Advanced Scientific Computing: Studying Transportable Explosives with Uintah (Martin Berzins, University of Utah)

Figure 2.1 An array of 64 explosive cylinders ignited in the lower left corner undergoing deflagration (left) and detonation (center) with an individual container fragmenting (right). *(Credit: M. Berzins and C. Wright, University of Utah)*

Led by the University of Utah's Martin Berzins and Charles Wight, a team of scientists is using the processing power of Kraken, the Cray XT5 at the National Institute for Computational Sciences, to simulate burning and detonation processes in transportable explosives. Recent simulations have led the team to believe that the key to preventing detonations is to pack the trucks transporting the explosives in such a way that temperatures and pressures are regulated so as to avoid a chain reaction explosion to detonation scenario. The difference in burning, or deflagration, and detonation is a crucial one. In deflagration, the transfer of heat to each individual device strictly limits the rate of burning. However, detonation is a phenomenon that occurs as the result of a shock wave, a catalyst that moves 1 million times faster than deflagration and in the end is approximately 1 million times more devastating. Nearly all of current knowledge surrounding explosions has evolved from experiments performed on single devices, or at most three related devices. Berzins and his team are pioneers in the field of explosive arrays, or the simultaneous explosion and resulting detonation of multiple devices in a single area. Thanks to Kraken, the team has developed models for the transition from deflagration to detonation, and they quickly reaching the point where they can begin to apply these models to the full problem. They are currently in the process of scaling up the models for explosions and detonations. That scaling up means simulating whole truckloads of explosives from the initial burning process to the resulting detonation. For now the team is only simulating a few explosives at a time, a daunting task even with Kraken's 100,000-plus computing cores. This computing power comes in handy especially when dealing with detonations, where the team believes that the phenomenon results ultimately from very fine-scale behaviors. This modeling across scales, from the micrometer inside explosives to the crater-sized destruction created from the explosion, is what the team refers to as "continuum level science." Due to the relationships between the different scales, this type of problem specifically requires a system such as Kraken, with the necessary processing power for simulating explosions and detonations from their microscopic beginnings to their larger cataclysmic finales. The University of Utah team used 3.5 million processor hours on Kraken last year alone. The team previously used the NSF's PetaApps program to get their simulations to their current fidelity. Their weapon of choice in this quest is Uintah, a general purpose scientific computing and engineering software application. Developed by the Department of Energy, the Department of Defense, and the National Institutes of Health, Uintah is an all-purpose tool that is currently putting out proverbial fires in areas from fluid flow to heat transfer to biochemistry to materials. In Berzin's case, it is excelling in every sense of the word. "Uintah is scaling very, very well, up to our biggest runs of 100,000 cores," he said. "We've run our software on almost all of Kraken in a scalable way with good results with adaptive meshes, where all the action is, where there are high gradients of pressure and temperature." Those gradients require enormous computing power and are a major reason why Berzins and his team require a computer as powerful as Kraken to continue to move forward.

2.2 Atmospheric Sciences: High-Performance Computing for Ensemble-Based Cloud-Resolving Hurricane Analysis, Prediction, and Predictability (Fuqing Zhang, Pennsylvania State University)

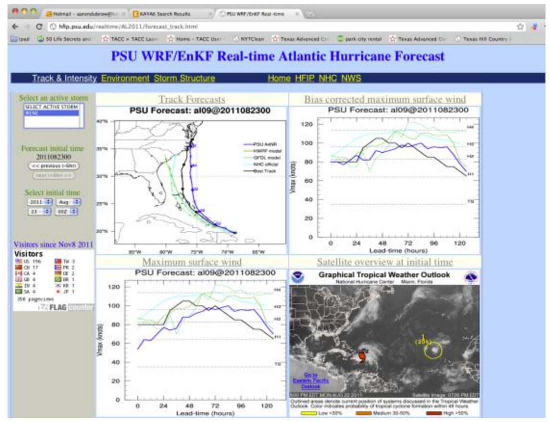

Figure 2.2 Zhang's website shows his group's real-time Atlantic Hurricane Forecast (created using Ranger) plotted alongside comparable model forecasts, satellite images of the storm, and other pertinent information.

When Hurricane Irene swept through New England in August 2011, predictions of its path were on target, but it arrived significantly weaker than originally forecast, leading to a larger evacuation than was required. Forecasting how hurricanes form, intensify, or dissipate is different and more challenging than predicting their paths. Over the course of several years, researchers from Penn State and the National Oceanographic and Atmospheric Administration (NOAA) used the Ranger supercomputer at the Texas Advanced Computing Center to develop and test a new, high-resolution hurricane forecasting system that improves on today's operational methods in several important ways. First, the system adopts a higher-resolution state-of-the-art computer model (4.5km grid spacing, compared to those of 9km) capable of more explicitly resolving a hurricane's inner-core dynamics and structure. Second, it ingests airborne Doppler radar taken by planes flying through the hurricane — information that NOAA has been collecting for 30 years but that has yet to be put into operational models. Third, the system applies a new data assimilation method that improves how the system ingests the Doppler radar data. Using Ranger, Zhang's system forecasted the track and intensity of every major storm in the Atlantic throughout the 2011 hurricane season, sharing the results of his simulations on his personal research website at http://hfip.psu.edu/realtime/AL2011/forecast_track.html. Zhang's forecasts were shown to improve intensity predictions by an average of 20 to 40 percent over the National Hurricane Center's (NHC) official forecasts for storms that have the airborne Doppler radar data. The prediction system is one of a handful the NHC is assessing to become part of the operational forecasting system used in emergency situations.

2.3 Atmospheric Sciences: Studying Magnetic Reconnection (Homa Karimabadi, University of California-San Diego; and William Daughton, Los AlamosNational Laboratory)

Earth's magnetic dipole creates a type of shield known as the magnetosphere, which protects humans from radiation unleashed by the sun in the form of an ionized gas known as plasma. Unfortunately, this radiation can occasionally sneak past Earth's defenses in a process known as magnetic reconnection, creating the northern lights and wreaking havoc on electronics and our daily lives. "When a storm goes off on the sun, we can't really predict the extent of damage that it will cause here on Earth. It is critical that we develop this predictive capability," said Dr. Homa Karimabadi, a space physicist at the University of California–San Diego (UCSD). Karimabadi's team, in close collaboration with Dr. William Daughton at Los Alamos National Laboratory, is currently using the Oak Ridge Leadership Computing Facility's (OLCF's) Cray XT5 Jaguar supercomputer to better understand how plasma spewed from the sun interacts with Earth's magnetosphere, investigating exactly what happens when the two meet and precisely what gets through and why. Initially the team used the National Science Foundation-funded, University of

Tennessee-managed Cray XT5 system known as Kraken to model the magnetic reconnection process that is the major cause of space weather, research that was published in the April 10, 2011, edition of *Nature Physics*. "One of the surprising outcomes of our research is the ubiquity and nature of turbulence in the magnetosphere," said Karimabadi. "This is important, since turbulence implies more efficient mixing of the plasma and fields, and after all, space weather arises because the plasma and fields emanating from the sun can penetrate and mix with the plasma and fields of Earth's magnetosphere." Extracting meaningful knowledge from these simulations poses its own challenges. A single run can generate more than 200 terabytes of multi-dimensional, highly complex data. To meet this challenge, Karimabadi has been working with Burlen Loring of Lawrence Berkeley National Laboratory on the development of specialized visualization techniques. "By color coding the magnetic field lines in our visualizations of the solar wind and the magnetosphere, we can track the level of mixing," said Karimabadi. Thanks to the speed and power of systems such as Kraken and Jaguar, this complex mixing phenomenon is finally becoming understandable. Because of limits on computing strength, earlier global modeling of Earth's magnetosphere was based on magnetohydrodynamics simulations, which neglected critical effects in the reconnection process. Essentially the plasma was treated as a fluid, or a single body. Karimabadi says, "With petascale computing, we can now perform 3D global particle simulations of the magnetosphere that treat the ions as particles, but the electrons are kept as a fluid with resolution that was well out of reach until recently." Jaguar and Kraken are fantastic analogous telescopes, and when coupled with Karimabadi's use of two individual codes, H3D and VPIC, the sun's stormy effect on our planet becomes clearer. While the two applications are indeed different (H3D is a global code developed at UCSD that treats the ions as particles and the electrons as a fluid, while VPIC, developed at Los Alamos National Laboratory, defines both electrons and ions as particles), they are complementary. H3D gives the team a look at the big picture, while VPIC serves as the magnifying glass, revealing all the gritty, noisy details. "Our codes are very efficient at using all the available computational power of Jaguar and Kraken," said Karimabadi, adding that as soon as the fastest hardware comes out, his software is ready to take full advantage.

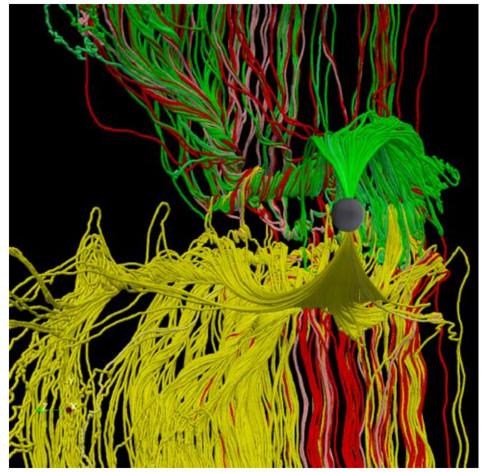

Figure 2.3 The complex entanglement and mixing of the field lines results from the formation of a magnetic flux rope generated by the magnetic reconnection process. *(Credit: H. Karimabadi and B. Loring)*

2.4 Chemical, Thermal Systems: Numerical Modeling of Dense Particulate Flows in Fluidized Bed Reactors (Olivier Desjardins, Cornell University)

A Cornell University-based team led by Olivier Desjardins has used both the Cray XT5 Kraken and SGI Altix 1000 UltraViolet Nautilus systems at the National Institute for Computational Sciences (NICS) to model the movement of particles in fluidized bed reactors (FBR). These simulations seek to understand and predict the multiphase flow inside FBRs, while ultimately improving the overall efficiency for thermochemical conversion of biomass (plant material) into liquid fuel. FBRs work by passing a fluid (either a gas or a liquid) through a grainy solid material at high velocity, which causes the solid to behave as though it were a fluid. This type of reactor has been used since the 1920s, primarily by oil and petrochemical industries to increase production of fuel in the United States. While FBRs are still used to produce fuels, they also can be used to produce industrial polymers like rubber and vinyl, and are even found at water and waste treatment utilities. Desjardins' team coupled its in-house computational fluid dynamics code NGA with a Lagrangian particle-tracking scheme in order to simulate dense particle flows in a lab-scale fluidized bed reactor. NGA scaled to more than 49,000 cores on Kraken. The team simulated a 4-inch FBR located at the National Renewable Energy Laboratory at half-scale. This simulation contained 17 million grid cells and nearly 16 million particles and required approximately 4 million CPU hours to run long enough to reach a statistically steady state. The simulation was run using more than 8,000 cores on Kraken for approximately 20 days, creating one of the largest particle-laden simulations to date. Particles tend to cluster and fall along the walls of vertical risers (a component of a FBR), binding surrounding particles as they pass and decreasing the efficiency of thermochemical conversion. Understanding the gas-solid interactions is key to improving the overall process efficiency. The team performed small-scale simulations to reproduce experiments found in literature. Periodic simulations of a section of a lab-scale riser were conducted using 19,000 grid cells and 300,000 particles on 288 cores. Steady state particle velocities were found to disagree with experimental results, indicating a need to simulate the full-scale reactor, which will be the team's next step with this project. Desjardins also has used Kraken to perform direct numerical simulations of gas-liquid annular flows in horizontal pipes. Annular space refers to the space surrounding one cylindrical object placed inside another. This "pipe within a pipe" scenario often is used to increase efficiency in a number of industries, including oil and gas. It's critical to be able to predict performance of annular flow to ensure efficient performance. "One of the fundamental aspects of our project is the large size of the simulations that we run, and the ability to run high core counts is essential to our progress," explained Desjardins. "This makes access to Kraken a vital aspect of our research. Nautilus has also been a great resource for visualizing our data." The team used a number of other XSEDE systems at the Texas Advanced Computing Center to conduct their research — the Dell Linux cluster Lonestar, the Sun Constellation Linux cluster Ranger, and the Sun visualization cluster Spur.

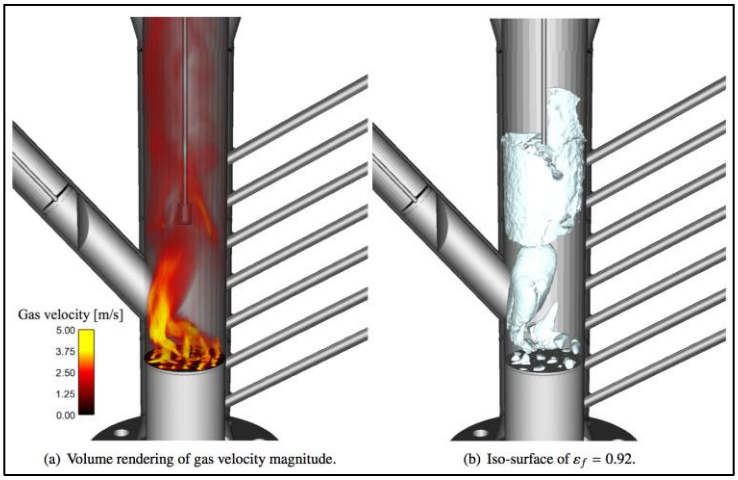

Figure 2.4 Instantaneous snapshot of the simulation shown within the reactor, with volumetric rendering of the gas velocity on the left, iso-surface of gas volume fraction on the right. Gas velocity is greatest inside the bubbles. (*Credit: J. Capecelatro and O. Desjardins, Cornell University*)

2.5 Climate Science: Simulating Regional Climate Change, a California Case Study (Alex Hall, UCLA)

Alex Hall, a professor in UCLA's Institute of the Environment and Sustainability, studies how climate change can have regional effects, and his work focuses on factors not well accounted for in large-scale global climate models. In the Los Angeles area, this means taking into account how the coastal Pacific Ocean affects conditions over the Sierra Mountains not far inland — natural features that, among other effects, shape winds known as the Santa Anas that have fueled some of the most furious wildfires to occur in densely populated areas. Hall works with a coalition of government, universities, and private concerns whose aim is to develop a Climate Action Plan that considers some of the “What if?” questions involved with climate change in the Los Angeles region. Since the Santa Ana phenomenon isn’t represented in coarse resolution global models, Hall seeks to recover that phenomenon by regionalizing the simulation at high resolution. Recent results from modeling by Hall and colleagues, reported in *Climate Dynamics* (2011), show that the connection between sea-surface temperature and winds at high resolution must be taken into account in considering the regional climate-change effects. Their work used the Pittsburgh Supercomputing Center’s Blacklight (also the National Center for Supercomputing Applications’ Ember, and before that PSC’s Pople) to solve a series of challenges in running computational experiments that integrate the effects of the coastal ocean with inland topography. At a resolution of two kilometers (compared to the roughly 100-kilometer resolution of global models), they capture the interplay between surface temperatures of the ocean and inland mountains in creating wind and rain conditions. They have reproduced California rain events as far back as data is available. In ongoing work, they expect to investigate other uncertainties associated with Los Angeles regional climate-change scenarios, including not only the Santa Ana winds, but also critical questions associated with availability of water resources.

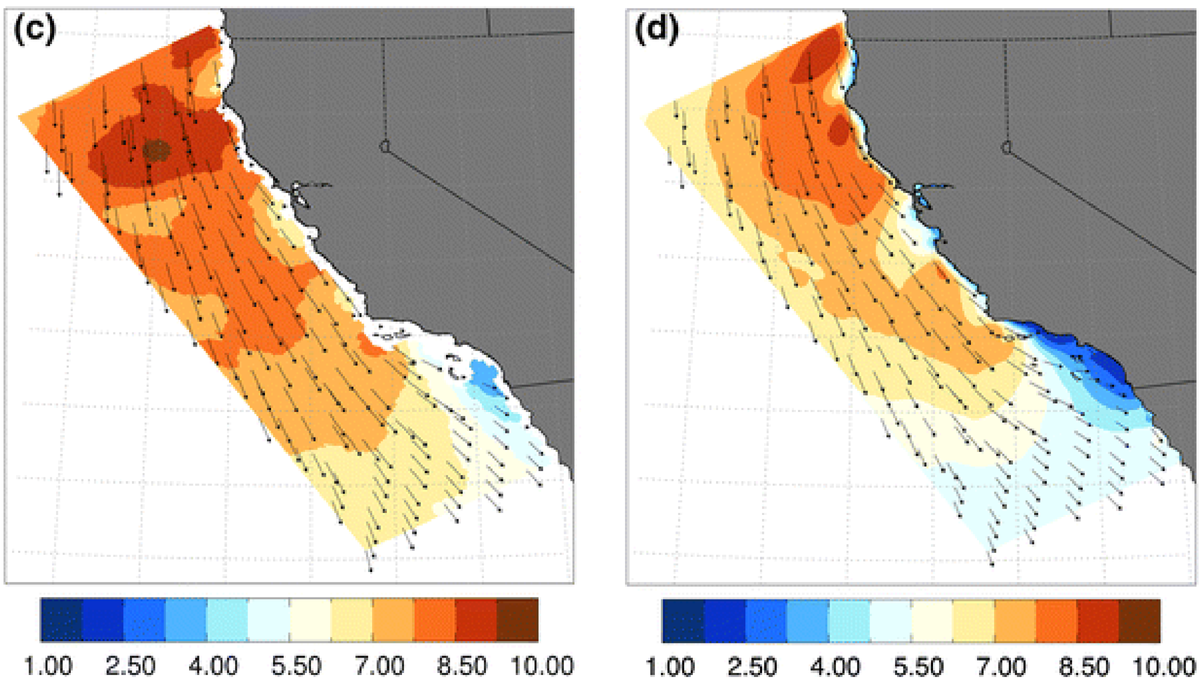

Figure 2.5 Southern California Wind in the summer of 2002. The first graphic (c) shows mean wind-speed direction (arrows) and magnitude (decreasing from red to blue) measured by satellite compared to the results of Hall’s coupled ocean and atmosphere model (d).

2.6 Genomics: Identification of Lignocellulosic Enzymes through Sequencing and Assembly of Metagenome Soil Enrichment (Mostafa Elshahed, Rolf Prade & Brian Couger, Oklahoma State University)

Metagenomics allows the study of microbial communities like those present in this stream receiving acid drainage from surface coal mining. Taking advantage of the Pittsburgh Supercomputing Center’s Blacklight and its shared-memory architecture for genomic data-analysis (the advantages of which were highlighted in an article in *GenomeWeb*, February 1, 2012), these researchers (with computation handled by Couger) completed what they believe to be the largest “metagenomics” assembly to date. Metagenomics is the study of genetic material recovered directly from environmental samples. This project focused on a soil sample from a sugar-cane field in Brazil, enriched in a bioreactor. The work aims at identifying previously unknown enyzmes that may be effective in breaking down non-feed stock lignocellulosic plants, such as switchgrass, wheat straw, and others, that have the potential to produce biofuel with a much higher ratio of energy per quantity of input crop than feed stocks such as corn. These non-feed stock materials, however, are recalcitrant to degradation and require costly pre-treatment to liberate fermentable material, which economically limits their suitability for industrial-level biofuel production. From their enriched soil sample, cultured for eight weeks, the researchers sequenced DNA (via a next-generation Illumina sequencer) that yielded 1.5 billion pair-data reads of 100 basepairs each, approximately 300 gigabases in total. Using Blacklight's shared-memory, and assembly software called Velvet — made available on Blacklight as a pre-compiled module by XSEDE consultant Phil Blood of PSC — the full assembly came out to 3.5 terabytes of RAM. Because of the large amount of RAM required to do the assembly, it couldn't be accomplished on any existing distributed-memory system, says Couger, and Blacklight, he believes, is the only existing architecture on which this problem was tractable. Analysis of the assembled metagenome is still under way and indicates more than 8,000 gene candidates related to glycoside hydrolase, a category of enzymes that can degrade plant cell walls. For future work, this project has also received an allocation on the National Center for Supercomputing Applications' Forge system.

Figure 2.6 Metagenomics allows the study of microbial communities like those present in this stream receiving acid drainage from surface coal mining.

2.7 Organic Chemistry: Tuning the Optical and Electronic Properties of Conjugated Polymers (Aimée Tomlinson, North Georgia College & State University)

Conjugated polymers (CPs), organic macromolecules with a backbone of alternating double and single bonds, have, in recent years, gained research attention related to their potential to replace traditional inorganic semi-conductor materials in applications such as field-effect transistors, light-emitting diodes, and photovoltaic cells. The 2000 Nobel Prize in chemistry recognized Alan Heeger, Alan MacDiarmid and Hideki Shirakawa for the discovery and development of these conductive polymers. CPs could reduce costs of device fabrication, through solution-based techniques, and offer the ability to fine tune the energy levels for specific applications through chemical synthesis. Tomlinson used the Pittsburgh Supercomputing Center's Blacklight (100,000 service units) and also used the National Center for Supercomputing Applications' Ember and San Diego Supercomputer Center's Trestles for extensive density-functional computations (using GAUSSIAN09) on a series of new *trans*-benzobisoxazoles(BBO) compounds. XSEDE consultant Marcela Madrid of PSC worked closely with Tomlinson and helped to develop scaling data to support her XRAC proposal. BBOs are promising for the development of new polymers because of their structure, which facilitates efficient charge transport, and they have been used to synthesize CPs. Investigating how varying the strength of electron donor and acceptor components of the CP backbone facilitates manipulating electronic properties of the compounds, Tomlinson's computations showed reasonable correlation between computational findings and experiment. Her work, published in *The Journal of Organic Chemistry* (2011), reinforces the promise of CP materials by showing that their electronic properties can be varied in predictable ways. Tomlinson benefitted from access to Blacklight's shared-memory architecture, since GAUSSIAN requires shared memory; Tomlinson's GAUSSIAN jobs typically used approximately 90 percent of the 128 Gbytes available on 16 cores of Blacklight.

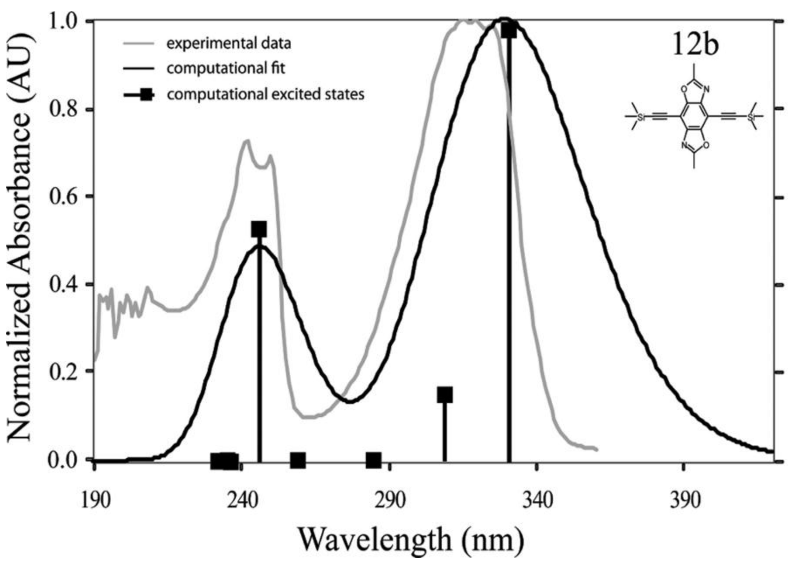

Figure 2.7 Comparison of the experimental UV-visual spectrum of a model Benzobisoxazoles compound (12b) with the predicted excited states.

2.8 Systematic and Population Biology: Large-scale Phylogenetic Estimation (Tandy Warnow, The University of Texas at Austin)

Phylogenetics is the branch of life science that studies the evolutionary relationships among organisms based on genetic evidence. Tandy Warnow, a computer science professor at The University of Texas at Austin, used the Ranger supercomputer to develop and test smarter, faster, and more accurate genetic alignment and tree-building algorithms and applied them to some of the largest biological datasets ever created. Among her projects, she and her team collaborated with evolutionary biologists from the Smithsonian to resolve the evolutionary history of flightless birds (known as *ratites*) and to answer the question: How did so many similar species get to the far-flung corners of the Earth? Smithsonian scientist Michael Braun discovered through DNA analysis that an ancient family of birds found in South American, the tinamou, is one of the most closely related groups to emus and ostriches — and they can fly! This fact,

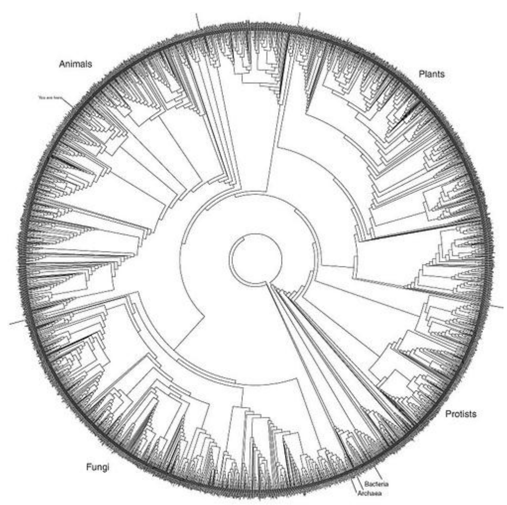

Figure 2.8 This phylogenetic tree, created by David Hillis, Derreck Zwickil and Robin Gutell, depicts the evolutionary relationships of about 3,000 species throughout the Tree of Life. Less than 1 percent of known species are depicted.

combined with the lack of skeletal evidence for flightless birds before the time of continental breakup, led to a re-conceptualization of the ratite branch of the avian tree. Warnow worked with Braun to reanalyze his controversial findings using her software. Their study confirmed the evolutionary relationship that Braun found. By improving the quality of the avian tree of life through computational methods, a new history emerged.

3 XSEDE Project Office 1.1

3.1 Overview

This quarter has seen significant progress in our ongoing efforts to redefine the distributed systems architecture for XSEDE. As noted below, we have produced a publically available high-level description of that architecture and the hard work of defining the details of the architecture has begun with a focus on the project's efforts in campus bridging. This activity has helped to refine the processes we will be using going forward as the Architecture and Design Team tackles additional areas.

We have also successfully delivered all three major new capabilities (GO-Data, EMS, GFFS) from the Software Development and Integration team to Operations for acceptance testing; issues raised for two of these components (GO-Data, EMS) have resulted in improvements.

Throughout the quarter we have successfully put more of our planning and management processes in place and as more of them are available we have increasingly been making use of them to improve how the project is managed and how we execute our work.

The Sciforma software has been acquired and initial staff training and software configuration has been completed. We have also installed and tested a web client to provide convenient and cost effective access to the DOORS Requirements database and have made this available to a limited number of XSEDE staff members and updated the content against the requirements baseline. SDI initiated an ongoing open planning activity to identify and prioritize technical enhancements to operational software, and introduced a significantly more agile approach to planning, launching and tracking software development and integration activities within SDI and on cross-WBS activities with substantial SDI involvement (e.g., Campus Bridging pilots).

This quarter also has seen the first publicly available documents emerge on the xsede.org website. Many of these come from the XSEDE Project Office and we have made the following documents available:

- Project Definition documents
    - Original Project Summary
- Policy documents
    - Service Provider Definition
- Architecture and Design documents
    - Level 1 and Level 2 Decomposition
- Engineering documents
    - System Requirements Specification
- Production Environment and Operations documents
    - Software and Service Baseline
    - Technical Security Baseline

These are all on the web site linked from the Project Documents page at: https://www.xsede.org/web/guest/project-documents.

In a related effort we have been slowly making progress on establishing a digital object repository for XSEDE project objects to allow long-term preservation and sharing of these objects.

As might be expected, the XSEDE Project Office has been the driver for Year 2 planning activities throughout the quarter coordinating the effort from the Project Management Team. This will be an ongoing effort throughout the following Quarter also.

Another significant step forward has been the formation of the XSEDE Advisory Board (XAB). As a key participant in our planning process, we selected and invited 10 members from the community to participate in this board which will provide guidance to help XSEDE achieve the maximum impact across diverse scientific disciplines and communities. The XAB membership consists of:

- Kevin Franklin, Illinois
- Phil Maechling, SCEC
- Rob Leland, Sandia
- Karin Remington, NIH
- Valerie Taylor, TAMU
- Jerry Ostriker, Princeton
- Theresa Windus, Iowa State
- Tom Lange, Procter & Gamble
- Greg Voth, UChicago

This group will be joined by the Chair of the SP Forum along with two additional SP Forum representatives, and the Chair of the User Advisory Committee once each of these individuals has been identified.

3.2 Project Management and Reporting 1.1.1

The project management team has accomplished the following tasks for the quarter:

Planning:

The process and infrastructure for the Year 2 Program Plan was developed, and the process was initiated in preparation for the first annual review of the XSEDE project.

Risk Management:

The risk register was updated with current status of identified risks, and was updated with known retired risks.

Project Management Software Tool:

The PM software tool Sciforma completed the purchase process and was installed on XSEDE servers. A requirements and configuration session was held March 13-14, and training for the project managers was held March 19 – 23. Following the configuration changes we requested and acceptance testing, we are anticipating a rollout for use by the project in late April.

Reporting:

The second quarterly report was published and submitted to NSF via Fastlane.

Quarterly Meeting:

The third XSEDE Quarterly meeting was planned and held March 6/7 in Austin following the Resource Allocation Committee (RAC) meeting. We will follow and co-locate with the RAC meeting whenever feasible to minimum travel expenses and maximize travel efficiency because a subset of the RAC meeting attendees also attend the XSEDE Quarterly meetings.

3.3 Systems and Software Engineering 1.1.2

During the quarter ending 03/31/11, the Systems and Software Engineering Team has accomplished the following:

Requirements Management

We have installed and successfully tested a web client that provides more convenient and cost effective access to the Doors Requirements database and have made this available to a limited number of XSEDE staff members. We also worked on developing “read only” access to this system, so that we can provide safe access to the requirements database to a broader group of XSEDE staff members.

The consolidated, approved, list of requirements have been cleaned up and entered into the DOORs database. Inconsistencies have been removed; all requirements now have unique identifiers and form a cohesive set.

The approved baseline XSEDE requirements provide admirable overall goals for the project. However, these requirements were developed during the proposal phase of the project and were based on user surveys and interactions that took place during 2010 or earlier, so may not reflect current user needs. During this quarter SSE started working with XSEDE staff to help them update and more fully develop their requirements to reflect current activity and trends. This will also enable us to gain better insight into progress made towards satisfying project requirements.

XSEDE Consolidated Glossary of Terms

XSEDE brings together staff from many very specialized areas such as system engineering, cyber-infrastructure, project management, as well as many technical specialties such as system design, operations, security, and networking. However, each of these specialties has its own unique vocabulary, which can lead to confusion for staff members who are not familiar with that specialized area. To try to help staff understand the different terminologies, we have started a consolidated Glossary of Terms to define terms and acronyms in use across the project as an aid to XSEDE staff members. Going forward, this is intended to be a shared document, updateable by anyone on the project who wants to contribute definitions and information from their area.

XSEDE Digital Object Repository

S&SE staff investigated how best to preserve the collected works of XSEDE (key documents, web pages, videos, etc.) for future use. We evaluated several different systems (IU Scholarworks, University of Illinois IDEALS, Cornell University Fedora Commons, ACM and IEEE Digital Libraries, OCLC.) using a set of required characteristics that any suitable candidates systems must have (ease of use, accessibility, cost effectiveness, ability to handle a wide variety of media and formats, security, longevity, etc.) We selected the UI IDEALS system as the best candidate system because it met all of our needs, offers a natural tie-in with the XSEDE because of the connection with UI, and because the IDEALS staff is interested in working with XSEDE. IDEALS staff have set up a work area for XSEDE and we have begun testing this system.

Risk Review

The 3 risks for Systems and Software Engineering remain the same: architecture obsolescence, hindered deployment and operation of XSEDE due to architectural ambiguity, and lack of qualified system personnel. Unfortunately, we once again triggered the third risk - lack of qualified system personnel. At the end of the last quarter a key SSE staff member left the project, but we were able to replace him this quarter with a new staff member. Unfortunately, just as this new person was coming up to speed on the project, and starting to make valuable contributions in the area of requirements analysis, he was let go as a result of administrative decisions at his home institution. At the end of the quarter we were looking into options for trying to replace this individual.

3.4 Architecture and Design 1.1.3

The A&D team (Felix Bachmann, Ian Foster, Andrew Grimshaw, David Lifka, Morris Reidel, & Steve Tuecke) delivered three important documents this quarter:

1. XSEDE architecture Level 1 and 2 Decomposition

https://www.xsede.org/documents/10157/281380/XSEDE+Architecture+2012-03-03+Level-1-and-2+ver+1.pdf

2. Campus Bridging Use Cases

   https://www.xsede.org/c/wiki/get_page_attachment?p_l_id=57086&nodeId=51230&title=XSEDE+Architecture+and+Design&fileName=XSEDE+Architecture+and+Design%2fXSEDE_Campus_Bridging_Use_Cases_2012_feb_2.pdf

3. Campus Bridging Requirements cross-matrix with Use Cases

   https://www.xsede.org/c/wiki/get_page_attachment?p_l_id=57086&nodeId=51230&title=XSEDE+Architecture+and+Design&fileName=XSEDE+Architecture+and+Design%2f2012-2-18-SRS-Reqs-121911-cross-matrix-with+campus+bridging+use+cases-rbs.xlsx

The "XSEDE architecture Level 1 and 2 Decomposition" document provides all interested parties a high-level view and scope of the XSEE architecture. New versions of this document are intended to be released quarterly that include increment updates, changes and additions based on priorities set by XSEDE leadership and important stakeholders. It will also include level 3 decomposition information for each area as they are completed and approved by all necessary stakeholders.

XSEDE leadership has identified the following areas that require architectural documentation at a level 3 decomposition level of detail. The ordered prioritization of these areas is reviewed and updated as necessary on a quarterly basis. The current list, point stakeholders and order of work is as follows:

1. Campus Bridging – Rich Knepper & Craig Stewart
2. Security – Randal Butler & Jim Basney
3. Science Gateways – Suresh Marru & Nancy Wilkin-Diehr
4. Integration and Interoperability – TBD
5. Inter-process communication – TBD
6. Replication and Fault Tolerance

The A&D has been using the following process to create level 3 decomposition documentation for each area:

1. Collect and document stakeholder requirements
2. A&D Team documents how the architecture addresses stakeholder requirements, including areas of debate/options
3. Stakeholders review A&D Team proposed solution and decide to:
   a. Accept, which entails one of:
      i. Choose a single solution
      ii. Agree that multiple solutions are required and must be tracked
   b. Request further consideration/information by A&D Team
   c. Escalate to XSEDE leadership in the event of an unresolved discrepancy
4. Repeat above steps until A&D Team solution is accepted by the stakeholders and approved by XSEDE leadership
5. Completed documentation is integrated into a public-facing XSEDE Architecture document which provides a Level 1 and Level 2 decomposition plus Level 3 decompositions for each area's use cases and scenarios as the architects document them.
6. The next area to be addressed is identified by XSEDE leadership and stakeholders and the process begins again

The "Campus Bridging Use Cases" and the "Campus Bridging Requirements cross-matrix with Use Cases" documents are the result of the first step of the process above. The architects worked in very close collaboration with Rich Knepper and Craig Stewart from the Campus Bridging area. Thanks to excellent support and guidance from Felix Bachmann this collaboration provided everyone with a better understanding of the architecture documentation process and with the Campus Bridging use cases and scenarios.

These document will be used as examples for the stakeholder leads in other areas so that they can begin working on the corresponding documents for their areas while the architects work on steps 2-4 in the level 3 documentation process for Campus Bridging.

David Lifka continues to work identify other parallel activities that can begin in a staggered fashion to improve the efficiency and total time required to complete the entire first pass (decomposition levels 1-3 for all areas) of the XSEDE architecture document.

Parallel Activities

In addition to documenting the XSEDE architecture, the architects have also been involved in parallel activities that will help them understand the technical concerns and perspectives of the SSE, SD&I, Security, SYS-OPS & leadership teams. Some of these activities include:

- Production Globus Data Transfer Service/GO-Data
- Beta Execution Management Service/EMS (Genesis/Unicore)
- Beta Global Federated File System/GFFS (Genesis)

Current Schedule

The current schedule for completing the level 3 decompositions for the following areas is as follows:

1. Campus Bridging – Program Year 1 (due to be completed by June Quarterly Meeting)
   - Lead Stakeholders: Rich Knepper & Craig Stewart
2. Science Gateways – Program Year 2 (7/1/2012 – 8/9/2012)
   - Lead Stakeholders: Suresh Marru & Nancy Wilkins-Diehr
3. Computing – Program Year 2 (8/9/2012 – 11/15/2012)
   - Lead Stakeholders: TBD
4. Big Data – Program Year 2 (11/15/2012 – 3/7/2013)
   - Lead Stakeholders: TBD
5. Connecting Instrumentation – Program Year 2 (3/7/2013 – 6/13/2013)
   - Lead Stakeholders: TBD
6. Collaboration – Program Year 3
   - Lead Stakeholders: TBD

Completion of a level 3 decomposition of each area will result in:

1. A use cases document
2. A use cases/SRS requirements matrix document
3. A new release of the Public Facing XSEDE Architecture Document that includes the additional detail from of the use cases and SRS requirements matrix documents.

3.5 External Relations 1.1.4

3.5.1 Accomplishments

SDSC communications staff promoted Gordon, a unique supercomputer at the San Diego Supercomputer Center (SDSC) and a key XSEDE resource, which went into production during the first quarter.

Promoted the following events on XSEDE website, Facebook, Twitter and through email lists: XSEDE12 CFP and updates, Summer School on Astroinformatics at SDSC, XSEDE- and PRACE-sponsored HPC summer school in Dublin, Blue Waters-XSEDE Extreme Scaling Workshop, GlobusWORLD 2012, 14th International Congress of Quantum Chemistry, and Second International Workshop on Advances in HPC Earth Sciences.

Edited content, and designed and produced websites for HPC summer school and Blue Waters-XSEDE workshop, and created online registration and payment system for BW-XSEDE workshop.

For XSEDE12: edited and added content to the website; created letterhead, badge template, registration page graphics, CFP promotional poster, and expanded timeline; edited numerous documents/proposals; followed up on process of engaging with ACM for digitalization of conference proceedings.

Sent confirmations and recorded all applications for HPC summer school.

Created online registration and connected to payment system for quarterly meeting.

Edited and compiled 2011 science and EOT highlights for NSF.

Collaborated with ECSS and EOT staff to produce application and project list for Campus Champions Fellows Program and for Student Engagement Program, promote ECS Symposium series, and provide communications assistance for Campus Champions' online and print materials.

Created first draft of general-use XSEDE slide deck and poster.

Created job description for new designer/communicator/programmer FTE.

Created Publications and Security pages in website.

Participated in document repository discussions and investigated an option (currently testing that solution).

Continued to produce two monthly newsletters, one focused internally and one externally.

Continued to track media hits.

Edited, posted and fielded questions about Digital Humanities Specialist position.

Continued working on strategic communications plan.

Developed Program Plan for Year 2.

3.5.2 Media Hits

XSEDE was included in more than 45 news items on XSEDE-related topics, as tracked during the quarter. An Excel file containing specific media hits has been submitted for inclusion in the Appendix.

3.5.3 Challenges

XSEDE Communications became involved in a greater number of projects this quarter -- especially those related to the ramp-up of several XSEDE events -- while down two positions; one is the loss of WBS Level 3 Manager Bill Bell, who took another job at the University of Illinois, and the other is a communications/design/programmer FTE, for which the hiring process recently began.

3.6 Industry Relations 1.1.5

Due to a loss in staffing, little has been done in this area of the project as yet. Some effort has been spent in identifying a new lead for this area. The only deliverable was the ongoing outreach to industry with our training and other offerings. This has been happening to a small extent but has not been well organized.

3.7 Software Development and Integration 1.1.6

The Software Development and Integration (SD&I) team made substantial progress in several areas. All three major configuration items under development (GO-Data, EMS, GFFS) were delivered to Operations for acceptance test; issues raised for two of these components (GO-Data, EMS) have resulted in improvements scheduled for PY1 4Q, while acceptance testing of the third component (GFFS) is being scheduled.

In addition to completing its first delivery of components to Operations, SDI also initiated an ongoing open planning activity to identify and prioritize technical enhancements to operational software, and introduced a significantly more agile approach to planning, launching and tracking software development and integration activities within SDI and on cross-WBS activities with substantial SDI involvement (e.g., Campus Bridging pilots).

A good start has been made on introducing Wiki-automation of technical project management activities such as tracking of milestones, schedules, and deliverables. SDI defined conventions for using the Wiki that allow us to track different kinds of engineering activities, the requirements they address, the components they modify, the staff working on these activities and their budgeted effort, milestones and deliverables.

In previous reports, we indicated that a continued challenge to SDI has been incomplete technical documentation and insufficient community-held software engineering practices. The conventions alluded to, above, are a start at consolidating and institutionalizing our understanding of different kinds of engineering activities undertaken by SDI (and more generally, by XSEDE), the practices that different kinds of activity entail, and different kinds of coordination these activities require.

SDI has also initiated practices to incrementally improve the quality of the XSEDE technical baseline documentation thru project management tasks associated with projects being launched in PY1 4Q. Specifically, two project management tasks involve improving project traceability to baseline requirements and production components.

Specific SDI 3Q accomplishments include:

- All configuration items (GO-Data, GFFS, EMS) were delivered to Operations for acceptance testing; security issues raised for GO-Data and EMS have resulted in scheduled enhancements to these components as well as the XSEDE security baseline.
- SDI planning identified more than 90 engineering activities to introduce new capabilities or enhance operational capabilities; these were consolidated and prioritized (internal SDI) by technical and effort feasibility and by technical dependency.
- Of the proposed activities, fourteen (14) were selected for possible launch in Y1-4Q 2012; the remaining activities were partitioned into those that might be, with additional planning, feasible for Y2-1Q (including topics

- Y1-4Q activities span a range of topics, including: specific technical enhancements to EMS and GO-Data; defining consolidated XSEDE test practices and architecture; delivering a standalone version of Unicore-6 (disaggregated from Genesis-II); substantial improvements to XSEDE system information registries, and dedicated support for upcoming campus bridging requirements.

4 XSEDE Operations 1.2

4.1 Overview

The Operations group consists of ~30 FTEs and is responsible for implementing, delivering, maintaining, and evolving an integrated cyberinfrastructure capability of unprecedented scale that incorporates a wide range of digital capabilities to support the national scientific and engineering research effort. The Operations group follows the XSEDE Project Management methodologies detailed in the Project Execution Plan by allocating and coordinating staff in accordance with the XSEDE work breakdown structure (WBS), scheduling tasks through the XSEDE project scheduling process, and identifying and reviewing risk on an ongoing basis. Operations staff is subdivided into six teams based on the WBS:

1.2.1 Security
1.2.2 Data Services
1.2.3 XSEDEnet (Networking)
1.2.4 Software Testing and Deployment (formerly Software Support)
1.2.5 Accounting and Accounts Management
1.2.6 Systems Operational Support

All project schedule items for Operations WBS 1.2 were updated for the baseline project schedule and are reflected in Appendix B. The Operations group completed all tasks scheduled for this quarter. Progress is being made on all other tasks scheduled for Year 1. All risk items for the group's activities were reviewed and are listed in Appendix C. No risks were triggered during this quarter.

Highlights within the Operations group for this quarter include the approval of initial plans for an XSEDE Certificate Authority by the governing body and enhancements to security that include the completion of security awareness training materials and plans to leverage the REN-ISAC Security Event Service. The Data Services groups completed a requirements matrix and identified the evaluation criteria for selecting a technology for the XSEDE-wide file system. The XSEDEnet group completed testing of a new network-monitoring tool and transitioned operational monitoring to the XSEDE Operations Center. The Networking group also installed perfSONAR Services at PSC, IU, and NCAR to perform regular throughput test between sites. The Software Deployment and Testing group conducted operational readiness reviews and acceptance tests for Globus Online file transfer service and Execution Management Services this quarter. The Systems Operations Support group continued to insure availability of central services, achieving over 97% uptime for all central services. The XSEDE Operations Center fielded 2,651 tickets of which 39% were closed within 2 business days. The following sections further detail the activities of the six Operations teams during this reporting period.

4.2 Security 1.2.1

The Security Operations team continued to insure confidentiality, integrity, and availability of data and services through day-to-day security operations and incident response. During this quarter the incident response team handled approximately five compromised user accounts. Also the IR team coordinated with a former TG site when their head node was briefly compromised, but no XSEDE accounts or resources were affected. There were no vulnerabilities this period requiring an XSEDE wide notification. In addition, the Security group engaged in the following efforts to enhance the security of the XSEDE infrastructure:

XSEDE Certificate Authority: The requirements and high-level plans for an XSEDE CA were gathered, documented and published on the XSEDE staff page. The concepts for this CA were presented to the International Grid Task Force, TAGPMA certification body, who provided feedback and approval to move forward.

A draft use case document has been produced, but campus-bridging requirements are still being explored. Future effort will involve documenting and comparing implementation options against the gathered requirements. In addition Open Science Grid and XSEDE management will meet on plans June 7 & 8th to discuss collaboration and future plans.

XSEDE OTP Service: An XSEDE one time password (OTP) evaluation is underway and the current project plan is to have a basic interoperable OTP implementation available to XSEDE SPs by July 2012.

Coordinated Log Analysis: During this quarter the Security group established plan to leverage the REN-ISAC Security Event Service. Each of the XSEDE Service providers will send their security incident information to REN-ISAC for analysis and correlation. These logs will be correlated with a larger pool of logs from other REN-ISAC members and the correlated data (SES data) will be fed back to the SPs and used to configure filtering and intrusion detection systems so sites are better capable to identifying attacks and probes in progress.

XSEDE Federation Risk Assessment: An XSEDE risk and threat analysis is ongoing. Documents detailing this effort are available on the staff WIKI under the security projects section. Policy updating and development work is planned to begin next quarter.

Security Training Program: Training materials have been completed and supplied to the XSEDE training staff. The Security group plans to offer the training at XSEDE12 as part of a new user tutorial.

Architecture and SD&I: Members of the Security Operations team continue to be heavily involved discussions about the SD&I Configuration Items: Detailed discussions about the XSEDE security architecture have occurred between the Architecture team and Security team. Specific discussions are now underway with the Genesis II team concerning how to advance their security implementations. Security Operations Readiness reviews of Globus OnLine and EMS were completed this quarter, and follow-up discussions with the development teams are underway to address the issues raised.

4.3 Data Services 1.2.2

Data Services activities for this quarter focused on maintaining the existing operational infrastructure, including GridFTP servers, archive resources, and wide-area file systems, along with migrating and improving documentation for data resources in XSEDE.

The XSEDE-Wide File System (XWFS) assessment process proceeded throughout the reporting period and is on schedule. A full requirements matrix has been completed listing all high-level and derived requirements and comparing each technology to those requirements, and benchmarks have been defined along with further evaluation criteria. The XWFS evaluation is documented on the XSEDE staff wiki under the Operations Data Services section including the meeting minutes, the XWFS requirements, benchmark scripts, and other evaluation information.

A wide-ranging discussion was begun on Data Services metrics and reporting criteria, and a year 2 effort has been defined to improve the reporting framework for existing components in the Data Services area as well as new components coming from SD&I in Increment #1. Preparation began for deployment of the Global Federated File System and Globus Online components with minimal disruption to production resource utilization.

4.4 XSEDEnet 1.2.3

Intermapper was selected as the network-monitoring tool for XSEDENet. Following configuration and testing, operational monitoring was transitioned to the XOC in Feb 2012. Work continues to enhance the reliability of the monitoring and reduce the number of false positives reported. The high noise ratio was caused by the polling method required by NLR and the IU GRNOC. Direct access to the SNMP data is not available.

During the quarter, hardware was received and configured to host perfSONAR services at Service Provider sites. Hosts were shipped to all sites. XSEDE perfSONAR Services are available at PSC, IU, and NCAR. These hosts are configured to perform throughput test between XSEDE sites on a regular basis.

With the news that the NLR Network Operations Center (NOC) will transition organizations from Indiana University to the Corporation for Education Network Initiatives in California (CENIC) in April 2012, investigation into the impact on XSEDE operations was initiated.

During the quarter, connectivity was established between the Galaxy project at Penn State and XSEDE via PSC and Three Rivers Optical Exchange (3ROX). The current XSEDEnet topology is shown below.

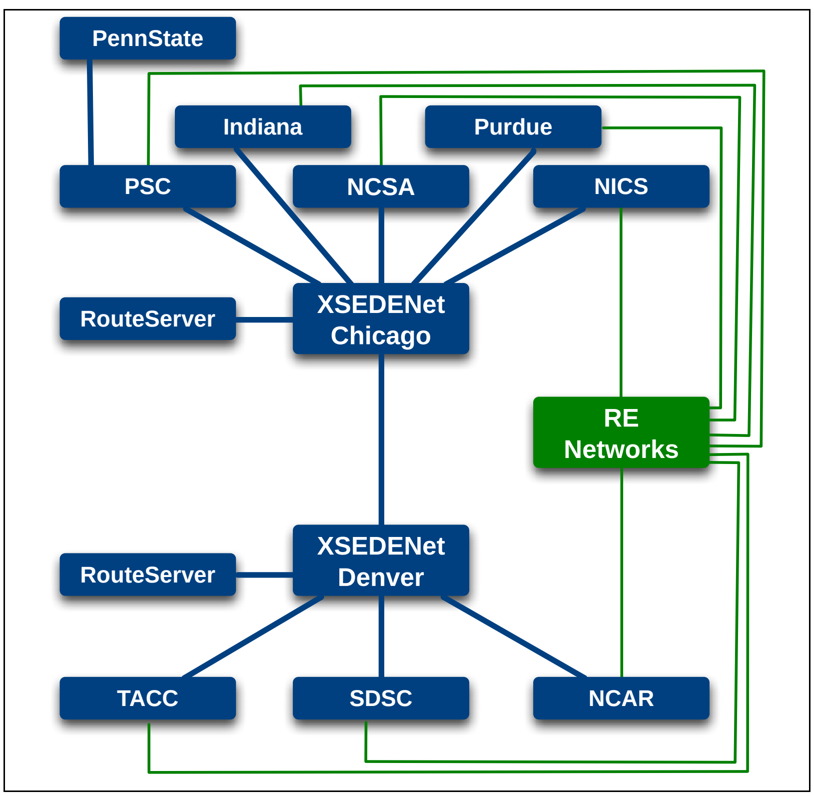

Traffic utilization of the Chicago-Denver link is shown below. The first graph shows the peak bandwidth across the link for the period. The second graph shows link utilization as a percentage.

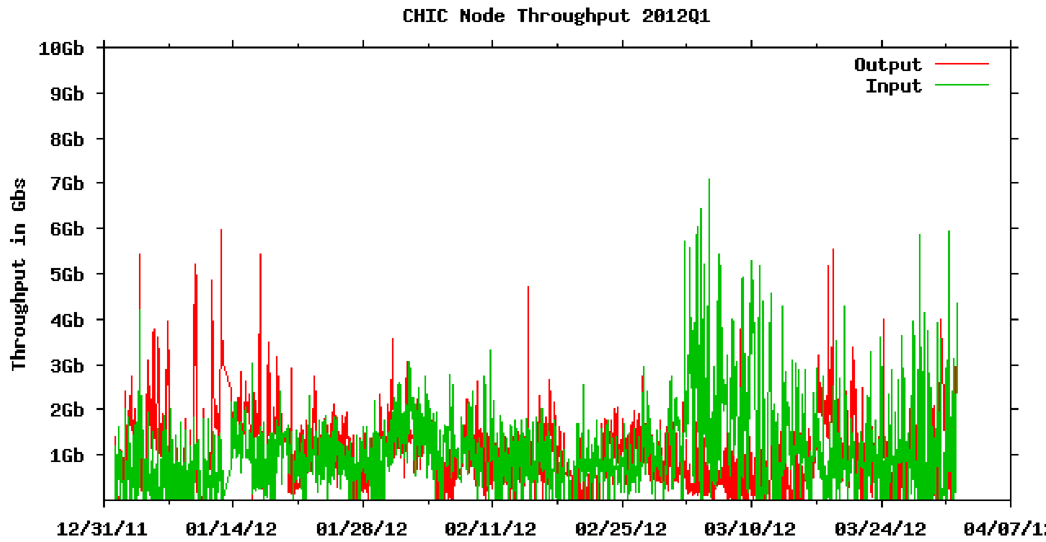

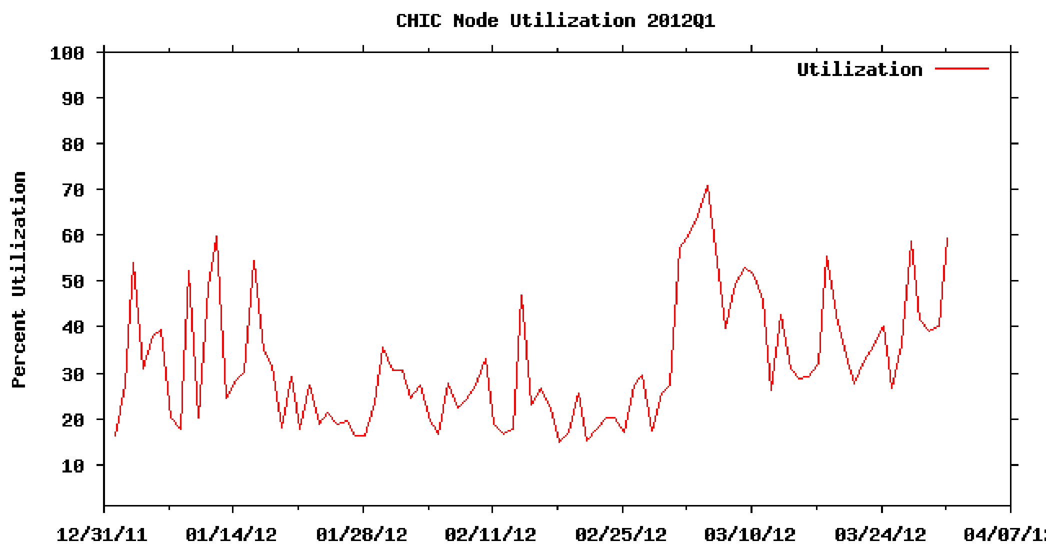

Traffic offered load into XSEDEnet links is shown below:

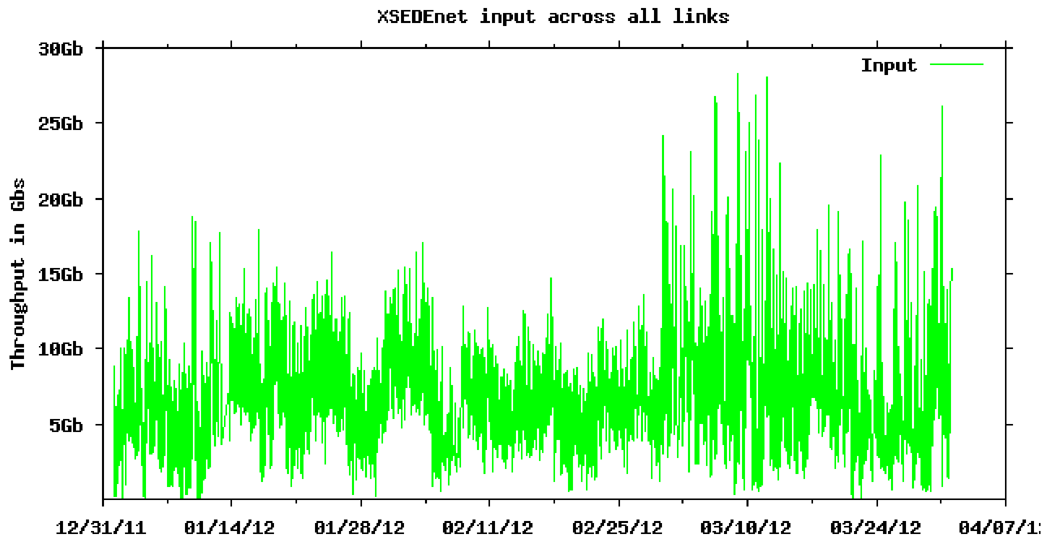

Below is the list of outages for the quarter by site.

Site	Up	Downtime
CHIC	99.87%	2:12:42
DENV	99.87%	2:18:13
FutureGrid	99.98%	0:17:16
IU	99.99%	0:11:39
NCAR	99.97%	0:32:20
NCSA	99.99%	0:13:46
NICS	99.87%	2:12:00
PSC	99.93%	1:16:48
Purdue	99.98%	0:16:03
SDSC	94.84%	3days 17:7:29
TACC	99.04%	16:32:59

4.5 Software Support 1.2.4

The Software Deployment and Testing group conducted operational readiness reviews (ORRs) and acceptance tests for two Configuration Items (CIs) this quarter. The first of these was for the Globus Online file transfer service, which passed its ORR and its functionality tests but had to be rejected due to security concerns about how the service caches XSEDE credentials. A work plan for resolving these issues has been compiled and will be completed during the next quarter.

The second configuration item tested this quarter was Execution Management Services (EMS), which includes UNICORE services and clients as well as job submission components of Genesis

II. This CI passed its ORR and is currently in acceptance testing, with a beta deployment planned to follow.

A test plan for the Global Federated File System (GFFS) was also completed this quarter. Operational readiness review and acceptance testing on GFFS is expected to start once the EMS acceptance test is completed next quarter, with a beta deployment planned to follow.

4.6 Accounting and Account Management 1.2.5

A focus for the Accounting and Account Management group in this quarter was to significantly improve user account request turnaround time. This work included completion of the decoupling of the Proposal Submission system (POPS) database from NCSA's database, implementing new person tracking capabilities in the POPS database, changing the routing of account requests through POPS when received from the XSEDE User Portal, and alterations to the POPS Account Management Information Exchange (AMIE) implementation. With this release completed (actual release date of 4/9/12), all account requests now reach the XSEDE Central Database (XDCDB) within two of hours of being submitted, and reach the relevant site(s) within three hours. More changes and improvements were made to POPS during this period that include:

- The portal login of the person making a submission is now captured and sent in the resulting AMIE packet
- A portal login field was added to the PI and co-PI pages, with the requirement that the PI portal login be entered for a successful submission, and any entries are verified against the XDCDB
- Work began to change the "Organization" field for PIs and co-PIs from a simple text field to a pulldown, with autocomplete (necessary for the account request work mentioned above)

Due to an extended health-related leave of absence of a key member of the Allocations working group, several members of the A&AM group filled in for this absence. Theses A&AM staff members worked closely with the Allocations workgroup, and with members of the XSEDE Helpdesk to cover this work. The work included the processing of proposal submissions (transfers, extensions, supplementals), account request reconciliation and processing, processing of tickets, and various other tasks.

4.7 Systems Operational Support 1.2.6

4.7.1 Overview

The Systems Operational Support (SysOps) group is responsible for operating the XSEDE Operation Center and administering all XSEDE centralized services. During the reporting period, effort was focused on several scheduled tasks with high importance placed on 1.2.6.2, 1.2.6.3, 1.2.6.4, 1.2.6.11, and 1.2.6.13. Several planned and unplanned outages occurred, but data integrity was maintained and proper failover procedures were exercised when appropriate. As such, no central services experienced any less than 97% uptime for the period.

4.7.2 XSEDE Operations Center

During the reporting period the XSEDE Operations Center (XOC) fielded 2,651 tickets. Among these 2,381 were submitted via email to help@xsede.org, 254 were submitted via the

XSEDE User Portal, and 16 were submitted via phone to the XOC. There were 1,044 tickets closed within 2 business days, which equates to 39% for the reporting period. This is up from 35% last quarter. There were a total of 2,263 tickets responded to within 24 hours, which equates to 85% for the reporting period. This metric is up from 72% from last quarter. The following graph shows the ticket breakdown (opened/closed) to each major resolution center:

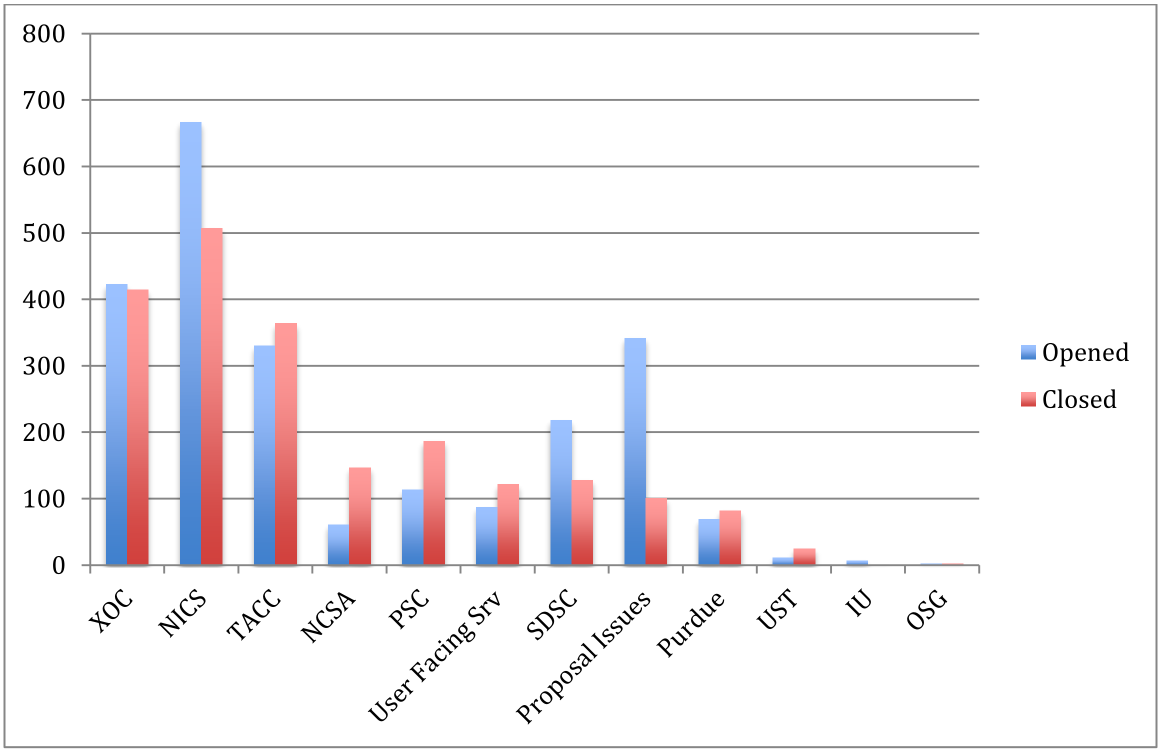

The tickets can also be further broken into 7 distinct problem categories. The below pie chart shows this breakdown:

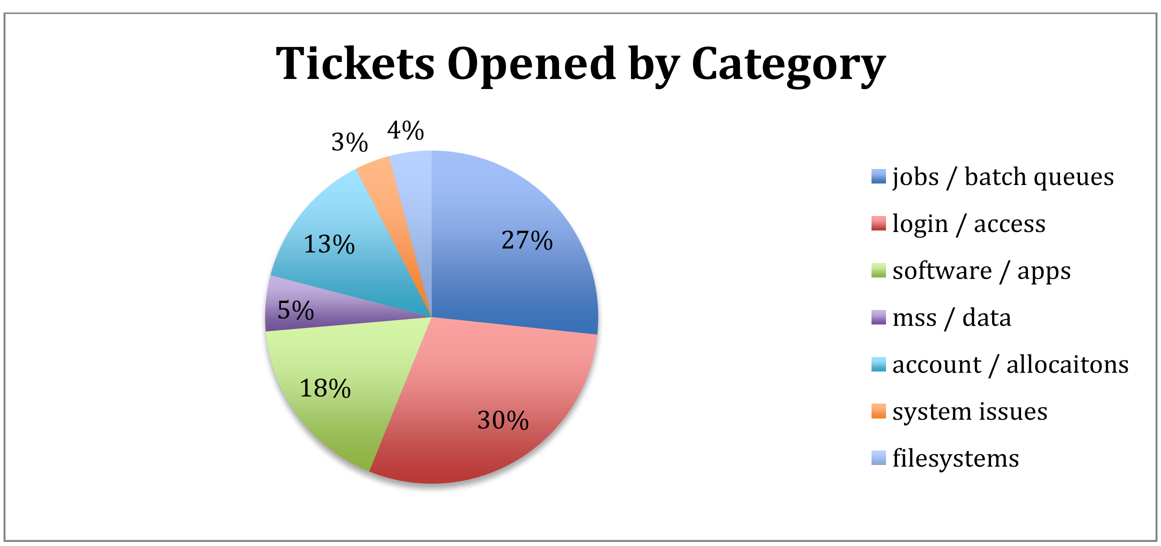

This pie chart represents a significant portion of the 2,651 tickets but does not represent the entire range. Tickets largely fit in the previous 7 categories but there are other categories that are not significant enough the visually represent.

4.7.3 *Central Services*

During the reporting period we experienced several outages both planned and unplanned that affected various central services. Many of these outages were the result of individual servers or sites experiencing intermittent technical difficulties. The most significant outage, this quarter, was with the Virtual Machine (VM) environment at Indiana University. This incident was the result of a large VM migration that experienced the perfect storm of bad events. Operations staff were able to revert to previous backups for these VMs. Indiana University Operations Management and XSEDE Operations Management continue to develop procedures and service level agreements to ensure continued stability. It is worth noting that up until this point the VM environment at IU was extremely reliable as they consistently maintained 99% uptime.

No significant power events occurred during the reporting period. The following table describes each service that experienced an outage, the corresponding downtime/uptime, and the nature of the outage:

Service	Percentage of Uptime (Number of hours)	Nature of the Outage
Information Services	**97.4%** (56 hours)	Unplanned
Karnak	**97.4%** (57.5 hours)	Unplanned
IIS Metrics	**97.4%** (56 hours)	Unplanned
XSEDE User Portal	**99.97%** (45 minutes)	Unplanned
TG Wiki	**97.4%** (56 hours)	Unplanned
Globus Listener	**97.4%** (56 hours)	Unplanned
Source Repository	**97.4%** (56 hours)	Unplanned
Build and Test	**97.4%** (56 hours)	Unplanned
Software Distribution	**97.4%** (56 hours)	Unplanned
Bugzilla	**97.4%** (56 hours)	Unplanned
User Profile Service	**97.4%** (56 hours)	Unplanned
Kerberos Backup	**99.99%** (5 mins)	Unplanned
XDCDB Backup	**99.99%** (30 mins)	Planned
AMIE Backup	**99.99%** (30 mins)	Planned
RDR	**99.77%** (5 hours)	Planned
Speedpage	**99.91%** (2 hours)	Planned

Portal/WWW backup	**99.6%** (8.5 hours)	Planned
Certificate Authority (not in production)	**99.6%** (8.5 hours)	Planned
Kerberos Backup	**99.91%** (2 hours)	Planned

The remaining services did not experience an outage during the reporting period.

4.7.4 INCA

At the time of this report, the Inca deployment was executing 963 tests for XSEDE software and services. Of these, 120 of these tests were running for six central XSEDE services: Inca, Information Services, Karnak, MyProxy, User Portal, and the XDCDB. The table below shows the definition of an outage for each service and the uptime percentages as detected by Inca. All services fall within acceptable limits of their high availability service definition.

Service	Definition of outage	Uptime (Details of outages)
Inca	Inca status pages are unavailable or not able to fetch data from the database (i.e., test details page fails to load). Tests every 5 mins.	**100%** (One partial outage for 16 hours – test results not stored)
Information Services	Information Web pages are unavailable. Tests every 15 mins.	**99.996%** (Two outages for 45 minutes of downtime)
Karnak	Karnak front page fails to load. Tests every 30 mins.	**99.93%** (Two outages for a total of 1.5 hours of downtime)
MyProxy	MyProxy server does not respond to credential query check. Tests every hour.	**100%** (No outages detected)
User Portal	Portal front home page fails to load correctly. Tests every 30 mins.	**99.97%** (One outage for a total of 30 minutes of downtime)
XDCDB	Connection to database refused or slow (using check_postgres.pl script). Tests every 5 mins.	**100%** (No outages detected)

4.7.5 Syslog Monitoring Project

Simulation of historical events and real-time monitoring setup was completed using Ranger's (TACC) Syslogs. This was accomplished by importing the Syslogs into a database. The import of this data into Ganglia was also started. Solicitations with TACC and NICS are underway to obtain Lustre logs for continued work. Real time log analysis and prediction is still the end goal and we are one step closer to realizing that goal.

5 User Services 1.3

5.1 Overview

XSEDE User Services activities continue to increase in quality and quantity, providing more technical information, training, and user capabilities each quarter to a growing user community and additional constituents. Usage of the XSEDE web site and XSEDE User Portal (XUP) also continued to increase as new technical documentation was published and new capabilities were added to the XUP. The XUP has become the one-stop source of information and a persistent information and services tool for most XSEDE users, and now offers improved allocations management and user profiles capabilities (with more enhancements in progress). Training reached thousands of users and potential users in the reporting period via in-person, online, and webcast training workshops. Allocations continues to conduct peer-review of XSEDE's over-requested computing resources, and began discussing storage allocations for implementation. Finally, User Engagement has completed its evaluation and recommendation of a new ticket system. Expected to be approved in Q2CY2012, the new system will further improve user issue analysis, and thus drive improvements to all of the other user services activities (web documentation, XUP, training, allocations, etc.) as well as XSEDE resources. User Engagement also drafted the first XSEDE user survey, to be released in Q2CY2012. In summary, User Services activities are operating smoothly through the third project quarter and poised for completion of some major milestones by the end of the project year.

5.2 Training 1.3.1

The XSEDE training efforts continued apace in the reporting period. More than 20 training events took place with well over 700 attendees, and online training drew more than 2,000 visitors to the 36 complete online modules.

The new data management plan course debuted in webcast form, and several courses were adapted to online seminars presented to the XSEDE scholars group. Planning continued for a number of events at minority serving institutions, with a session at Florida A&M taking place in Q2CY2012, and events in Tennessee and North Carolina set for Q3CY2012. After a workshop in February, XSEDE training content is now embedded in iPlant training courses targeted at the plant biology community as well. Joint training with the National Center for Genome Analysis will commence in Q2CY2012.

Overall, in both delivery and in course development, XSEDE remains on track to deliver on the milestones for program year 1.

5.3 User Information & Interfaces 1.3.2

Both the XSEDE web site and XSEDE User Portal (XUP) released new features and surpassed expected deliverables for this quarter. The XSEDE web site continued to expand its list of user guides by releasing a software guide for Globus Online and a user guide for the Open Science Grid. In addition, the web site created a new getting started guide to introduce new users how to get started on acquiring an allocation and starting to use XSEDE resources and a new file transfer page with dynamic GridFTP endpoints. The web site continued to improve documentation and navigation based on both XSEDE staff and user feedback. The team also released new user news categories that focus more on services oriented news instead of institution specific and a new downtime calendar that uses the new user news information.

This quarter also included new capabilities in allocation management in XUP. These include new burn rate usage meters for users to see how fast they are using their allocations. Also PIs and Users have allocation usage graphs they can view as well. PIs can view the SU usage of the project as a whole and for individual users in their project. Users can now also view their own individual SU usage on resources in a project. The team also released a new user profile page that enables users to upload an avatar image that is used across the XSEDE site, e.g., when posting user forum items. New profile enhancements are in progress for the next quarter.

The overall usage of both the web site and user portal continues to increase with almost 2.5 million hits on the web site and over 1.3 million on the user portal. Furthermore, almost four thousand XSEDE users logged in to the user portal and out of 2,216 users running jobs on XSEDE over 70% of them are logging in to the user portal. Lastly, over 30,000 portal file transfers resulted in over 2.5 TB of data transferred via the portal file manager and an increase across the board for all XUP services.

The Knowledgebase (KB) team expanded and improved the KB articles by adding 73 new KB items and updated existing articles to insure accuracy. There are currently 478 articles in the KB and there were 68,619 hits on the KB within the reporting period.

5.4 User Engagement 1.3.3

XSEDE User Engagement is organized as two working teams: Feedback and Consulting.

5.4.1 Feedback

The Feedback team completed all required feedback activities and required reporting, other than the development and implementation of the annual user survey. Significant progress was made on the survey during this quarter, and the development and implementation of the user survey is expected to be complete early next quarter. The survey is expected to be administered early enough next quarter to allow for the completion of the required analysis and reporting by the end of the quarter. No usability panels were requested; thus, no usability test plans were required this quarter.

The Feedback team conducted a focus group entitled “Petascale and Beyond” at the SIAM Conference on Parallel Processing for Scientific Computing in February, 2012. The intent of the focus group was to determine the needs, opinions, and overall satisfaction of petascale users. Topics of discussion included XSEDE resources and services, the allocations process, support, and the direction these users expect to see supercomputing head in the next two years. A report documenting this event is posted to the XSEDE staff site for reference.

The Feedback team completed data mining on the tickets submitted to help@xsede.org during the previous quarter. The data mining revealed that the majority of tickets were related to routine operational issues such as:

- jobs/batch issues,
- login/access issues,
- software/applications issues, and
- systems issues.

In addition to these issues, individual XSEDE sites also received tickets related to archival storage and file system issues. Tickets associated with chemistry software and efficient data transfer continue to be submitted. A new XSEDE document entitled *Data Transfers and Management* (https://www.xsede.org/data-transfers) that was deployed by XSEDE User Services in the reporting period should reduce the number of tickets related to data transfer in the future; such tickets will be monitored over the next few quarters to confirm the impact. Both NICS and

NCSA received a number of GPU-related tickets associated with their local resources. None of the issues and trends observed this quarter are considered significantly problematic. A report documenting the results of the quarterly data mining activities is posted to the XSEDE staff site for reference.

5.4.2 Consulting

The Consulting team conducted all required consulting activities but failed to meet the milestones associated with the selection, deployment, and release of a new/improved XSEDE ticket system. A cost-risk-benefits analysis for the adoption of the recommended ticket system was completed near the end of the quarter, and the associated report is expected to be submitted to the XSEDE Senior Management Team early next quarter. The selection and deployment of a new/improved ticket system are expected next quarter.

5.5 Allocations 1.3.4

The table below shows the overall allocations management activity handled by POPS and the allocations staff for the reporting period. Note that for transfers, the table shows only the positive side of the transaction to show the total transfer level; there is a corresponding negative amount, adjusted for resource exchange rates.

POPS Requests and Awards

	Research				Startup				Education				Campus Champions			
	# Req	SU Req	# Awd	SUs Awd	# Req	SU Req	# Awd	SUs Awd	# Req	SU Req	# Awd	SUs Awd	# Req	SU Req	# Awd	SUs Awd
New	48	198,995,720	44	54,778,458	149	22,272,682	133	15,601,409	9	1,696,014	9	1,706,013	18	12,181,120	17	10,592,017
Prog. Report	4	16,313,212	4	14,402,325	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
Renewal	74	592,391,322	72	327,640,399	19	2,960,117	16	1,935,087	3	945,019	2	40,001	9	6,959,131	9	5,288,017
Advance	31	25,997,280	28	9,213,840	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
Justification	2	3,500,002	1	1,710,000	0	0	0	0	0	0	0	0	0	0	0	0
Supplemental	23	42,483,955	19	19,127,441	13	3,154,001	10	683,002	2	100,000	1	20,000	4	350,000	4	350,000
Transfer	97	32,189,200	83	21,948,391	40	2,944,273	37	1,527,278	3	69,284	2	56,000	1	2,600,000	1	190,000
Extension	83	n/a	67	n/a	42	n/a	37	n/a	1	n/a	1	n/a	0	n/a	0	n/a

The March quarterly allocations meeting, XRAC, was planned and held in Austin, TX. The next two XRAC meetings have been scheduled for Knoxville/Oak Ridge, TN, in June 2012 and Champaign/Urbana, IL, in September 2012.

Requests totaling 807M SUs were reviewed at the March 2012 XRAC meeting, 200M SUs more than the previous quarter. Reviewer recommendations totaled 450M SUs, and 397M SUs were available. A large number of moves were needed to bring recommended awards in line with the individual resource available limits. However, the XSEDE Allocations staff and XSEDE site representatives are faced with a difficult situation of continued requests for supplements, transfers, and startups but no significant pool of SUs to satisfy the user community requests!

The XSEDE Allocations staff received 425 tickets within the reporting period, more than double the previous quarter. Most, if not all, were addressed and a high rate of user satisfaction achieved.

Lastly, XSEDE Allocations staff along with the XSEDE Operations/accounting group, led by Steve Quinn, are working on improving reporting of not only PI awards to their respective program officer but also to the entire XSEDE community via quarterly award announcements in conjunction with the XSEDE ER team.

6 Extended Collaborative Support Service – Projects 1.4

6.1 Overview

The Extended Collaborative Support Service (ECSS) pairs members of the XSEDE user community with expert staff members for an extended period to work together to solve challenging science and engineering problems through the application of cyberinfrastructure. In depth staff support, lasting weeks to up to a year in length can be requested at any time through the XSEDE allocations process. Expertise is available in a wide range of areas, from performance analysis and petascale optimization to the development of community gateways and work and data flow systems. We divided ESCC efforts in two, one designated as Projects, headed by Ralph Roskies; the other, designated as Communities, headed by Nancy Wilkins-Diehr. These groups have very close interactions, with common Project Management support. The ECSS Projects and Communities L2 and L3 managers have a one-hour call every two weeks at which outstanding issues are discussed. Other interaction is largely by e-mail. ECSS staff also participate in reviewing adaptive proposals associated with XRAC meetings. 30 ECSS staff reviewed adaptive proposals this quarter and 2 served as full XRAC members.

An ESRT project is a collaborative effort between an XSEDE user group and one or more ECSS staff members, whose goal is to enhance the research group's capability to transform knowledge using XD resources and related technologies. The mission of the Novel and Innovative Projects (NIP) team is to provide proactive, sustained efforts to jump-start XSEDE projects by non-traditional (to HPC/CI) users. Activities may range from initial contact to the development and execution of successful projects, including those that receive extended collaborative support. The scope of NIP includes disciplines whose practitioners have rarely availed themselves of HPC/CI resources in the past. It also includes demographic diversity, such as researchers and educators based at MSIs and EPSCoR institutions, and SBIR recipients.

The NIP team often identifies disciplines where the current ECSS team does not have the requisite expertise to assist prospective users. For this, we budgeted for flexible hires for a year at a time. This quarter we identified such an area (Digital Humanities), and made our first flexible hire (Alan Craig from NCSA). He will work in this area at 49% effort. In other areas (Genomics, Databases, Gordon optimization), we were able to reassign staff at an SP site into the regular ECSS team. We are trying to do something similar for people with High Throughput Computing (e.g. Condor or OSG) experience.

We have already most of our annual goals although we are not through the year. These can be seen the L3 milestones found in the individual L3 reports. ESRT has a large distributed staff and many different projects with different start/end dates. The project management staff is excellent and has helped tremendously, but we eagerly await the project management software (which was finally acquired this quarter) to replace email and cumbersome spreadsheets. We have speeded up the assignment of personnel to projects, by calling on our excellent project managers (Karla Gendler and Natalie Henriques). See their report below. ECSS (represented by Ralph Roskies, Sergiu Sanielevici and Nick Nystrom) has collaborated with Blue Waters (led by Bill Kramer and Scott Lathrop) to develop this year's petascale symposium that had been started by Blue Waters and TeraGrid. It will consist of 4 invited talks and at least 8 contributed papers. It will be held in Chicago just preceding the XSEDE12 conference.

PM report: Over the last three months, ECSS's Project Managers (Karla Gendler and Natalie Henriques) have continued to maintain spreadsheets of project requests, active projects, project assignments and staffing in lieu of formal project management software. They have streamlined the process for receiving start-up and supplemental requests such that they make the initial assignment of L3 manager thus reducing time to contact PIs and additional email. They review and track workplans, entering quarterly objectives into the spreadsheets. They manage and attend ECSS meetings and XSEDE PM meetings, posting notes and action items to the ECSS wiki once the meeting has concluded. They provide ECSS information to the XSEDE PM office and relay information from the PM team to ECSS. They also maintain the ECSS wiki and mailing lists. They actively participate on the committee in charge of purchasing the PPM software and testing the PPM software. Gendler has taken an active leadership role to ensure ECSS needs are met in regards to the PPM software. It has been purchased and as such, Gendler attended NCSA for a week to participate in a requirements gathering session with Sciforma and both Gendler and Henriques attended a week-long training session on how to use the PPM software. They coordinated the gathering of information for this report and have published all of the information to the wiki.

6.2 Extended Research Teams Support 1.4.1

An ESRT project is a collaborative effort between an XSEDE user group and one or more ECSS staff members, whose goal is to enhance the research group's capability to transform knowledge using XD resources and related technologies. Typical ESRT projects have a duration of several months up to one year and include the optimization and scaling of application codes to use 100,000 nodes or more per job; aggregating petabyte databases from distributed heterogeneous sources and mining them interactively; or helping to discover and adapt the best work and dataflow solution for simulation projects that generate ~100 TB of persistent data per 24-hour run. The first year of the XSEDE ESRT program is also managing projects transitioned from the TeraGrid ASTA program, and all of the ASTA projects will be completed by the end of June 2012.

A request for ESRT support is made by the principal investigator (PI) of a research team via the XSEDE resource allocation process. If the request is recommended by the reviewers and suitable to be an ESRT project, and if staff resources are available, a statement of work for up to one year will be developed by in collaboration by the PI, the ESRT team leader, and the ESRT manager and project manager. The work plan will include staff assignments from the pool of available advanced support experts who have the necessary skills. The ESRT team leader, working with the ECSS project manager, will be responsible for project tracking and reporting and for requesting additional resources or assistance from XSEDE management as needed.

Metrics that describe the state of the ESRT program are provided in Table 1 and Table 2, which t quantify ESRT requests and projects for this quarter.

Table 1 ESRT project metrics for this quarter

Metric	**XRAC**	**Startups/Edu**
Number of requests	19	29
Number of requests recommended for support	17	11
Number of projects initiated	11	8
Number of work plans completed	0	1
Number of work plans completed for previous quarter(s)	2	5

Table 2 ESRT project breakdown

Metric	XRAC	Startup/Edu	TG ASTA
Number of projects active	27	17	12
Number of projects completed	0	0	11

There are 13.28 FTEs assigned to ESRT from NCSA, NICS, PSC, SDSC, and TACC.

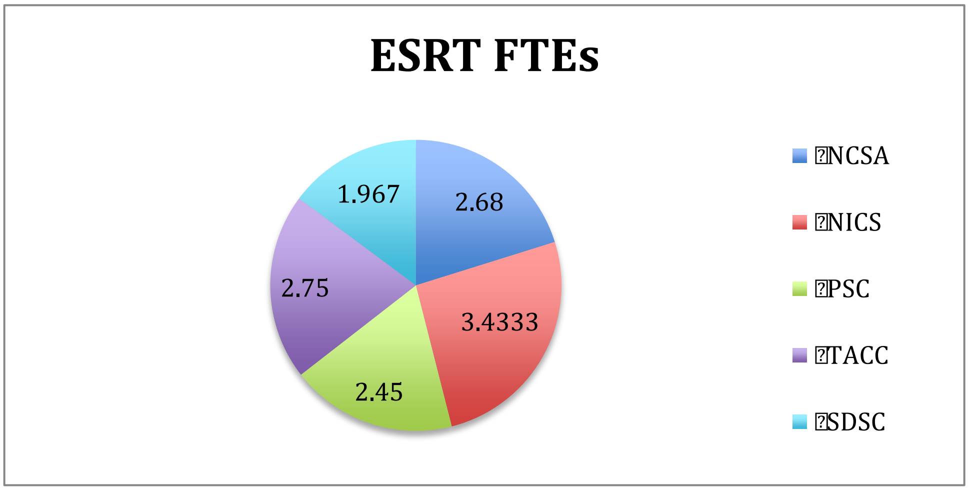

There has been one main challenge with managing the ESRT program - the management of projects and people. ESRT has a large distributed staff and many different projects with different start/end dates. The project management staff is excellent and has helped tremendously, but we eagerly await the project management software to replace email and spreadsheets.

The following sections highlight a few projects just to provide some examples of the kinds of work that is being done in the ESRT program.

6.2.1 Processing Census Data (McHenry, UIUC)

ECSS Project Team: Jay Alameda (NCSA), Rui Liu (NCSA), Yan Liu (NCSA), and Ye Fan (NCSA)

For this project, there are six deliverables defined in the work plan. By the end of March 2012, the first three deliverables have been achieved. Specifically, the Census code has been installed on four HPC clusters (Ember@NCSA, Blacklight@PSC, Lonestar and Ranger@TACC). The code has been thoroughly benchmarked in both non-HPC and HPC environments using two performance profiling libraries - hprof and psjprof. One computing bottleneck caused by the code was identified and communicated to the project developers. This discovery resulted in 20% reduction on CPU time. There were profiling results differences observed across machines: one difference interpreted as the difference of processor architecture while a second is still under investigation. Data I/O experiments were conducted on Ember in order to study the I/O patterns of Census data processing. Consequently, a data transfer solution was then prototyped and recommended to the project team to significantly improve the data transfer performance among raw data storage system, shared file systems on HPC, and local disks on HPC nodes.

Given current project progress, it is expected that the project team will request continued support from ECSS in order to eventually support the production runs on 1940's U.S. census data. The ECSS team will work closely with them to define new objectives in the scope of user support.

6.3 Novel and Innovative Projects 1.4.2

The mission of the Novel and Innovative Projects (NIP) team is to provide proactive, sustained efforts to jump-start XSEDE projects by non-traditional (to HPC/CI) users. Activities may range from initial contact to the development and execution of successful projects, including those that receive extended collaborative support. The scope of NIP includes disciplines whose practitioners have rarely availed themselves of HPC/CI resources in the past. It also includes demographic diversity, such as researchers and educators based at MSIs and EPSCoR institutions, and SBIR recipients. Bringing these communities to XSEDE is expected to lead to the consideration of applications and programming modes that have not been the focus of HPC in the past, such as those necessary for data analytics and informatics, and of innovative technologies such as streaming from instruments, mobile clients, and the integration and mining of distributed, heterogeneous databases. We expect that the implementation of campus bridging processes and technologies will be particularly important for these communities.

Biweekly teleconferences and the use of the project wiki and email list have been successful in catalyzing communications among team members, who have often benefitted from each other's contacts and expertise and helped each other.

We began to address the main challenge identified in previous quarters: balancing our ability to *attract* novel projects with our ability to help them *succeed* in becoming efficient users of XSEDE resources. One aspect of this is matching the expertise of ECSS staff to the characteristic requirements of NIP, which often differ from those of traditional HPC applications. To do this, we completed the first ECSS flexible contract hiring process and recruited Alan Craig of I-CHASS to become our *Digital Humanities* expert as of April 1. We also rebalanced our staff to add expertise in wide area data management and in database technology.

Another strategy is to *consolidate and leverage* our methods for eliciting XSEDE-ready projects from target communities. We have had considerable success doing this in the Genomics community, where our previously reported work with the Broad Institute has generated 13 new Startup allocations on Blacklight alone. We have now initiated a similar effort with the GIS community, leveraging the *CyberGIS* project led by NIP member Shaowen Wang (NCSA). We are exploring the development of projects of mutual interest with RENCI and with the arts and humanities program at TACC, as well as with I-CHASS.

Consolidation is also an effective strategy for ensuring the success of the novel and innovative projects we help to launch. Thus, we have determined that the most efficient way to support all the various research groups who want to use the *Trinity* and *Allpaths-LG* codes on Blacklight and possibly on other XSEDE platforms is by means of the ECS project *Collaboration with the Broad Institute - Genomics Community Capabilities*. A number of these groups had requested ECS at the individual research team level, and our NIP genomics experts are in touch with the individual groups regarding their specific research objectives and usage modes. But the first priority is to work with the developers to make the codes as efficient as possible on the XSEDE platforms and to document a general set of best practices for all users.

The NIP-initiated ECS projects *Interactive Large Scale Media Analytics* and *Virtual Worlds Exploratorium Research on Dynamic Networks in Massive Multiplayer Online Games* also have the potential to impact wider communities if their work plans are successfully executed, and are therefore classified as Community Capability projects.

The NIP-initiated ECS projects *An Inverse Stiffness Mapping Approach to the Early Detection of Breast Cancer*, *Computational Explorations in to Population Level Disease Spread using Agent Based Modeling*, *Large Scale Semi-Supervised Information Extraction (Restructuring NELL for Parallel Performance)*, *Large-shared-memory supercomputing for game-theoretic analysis with fine-grained abstractions, and novel tree search algorithms*, *and* *Processing Census Data* are being executed in with the research teams that requested them.

6.3.1 *We continue to face the **challenge** to sustain the development of non-traditional projects and their successful execution using XSEDE services and resources. To address them, we will continue to pursue and refine the strategies discussed above. We will make sure we leverage our strengthened relationships with the Genomics, Digital Humanities, GIS, and Computational Economics communities. Planned NIP outreach activities, including the May 15-17 tutorial at the Great Lakes Regional Bioinformatics conference in Ann Arbor, MI (http://www.iscb.org/glbio2012) and the July 27-28 tutorial at the Chicago-Argonne Initiative for Computational Economics (http://ice.uchicago.edu/) are expected to promote these efforts. We will continue to monitor the need to rebalance the ECSS staff expertise pool and, if necessary, to recruit additional contract positions to meet the support requirements we discover as we interact with new communities.*

7 Extended Collaborative Support Service – Communities 1.5

7.1 Overview

There was a significant focus on bioinformatics this quarter in ECSS-Communities, with several projects focusing on the Galaxy framework beginning in earnest and porting work beginning for oft requested community codes Trinity and AllPaths-LG.

The gateway team continues to work with XSEDE's Architecture and Design (A&D) and Software Development and Integration (SD&I) as the architecture definitions evolve. ESTEO continues novel interactions with the community including staff hosting of a professor and student from from Instituto Politechnico Nacional in Mexico City.

The ECSS Symposium continues each month and is open to the public to highlight work going on in ECSS projects and allow ECSS staff to learn from one another. Speakers this quarter discussed the use of Condor for scientific applications, how to use Nautilus for visualization and analysis, experiences developing software tools for the scientific community through the Apache Foundation, GPU accelerated AMD simulations, a general purpose workbench for science gateway development (GPSI) and GPU-based radiation treatment. These talks have been very well received and attract both the user community and staff.

The Campus Champion Fellows program application process is now complete. Applicants can choose from 23 projects where they can gain expertise in many different areas of cyberinfrastructure, including sequencing analysis, performance analysis, hybrid parallelization, I/O patterns and workflows.

7.2 Extended Support for Community Codes 1.5.1

ESCC has met its initial goal of supporting at least 10 projects per year with the acceptance of 10 completed work plans and ongoing development of workplans for three more projects. Over the past quarter, 3 requests have arrived via the XRAC process and 5 requests via Startup proposals. Out of these, 5 were approved to move forward as projects. A few requests were determined to be short term in nature and were referred to local XSEDE staff to be handled directly.

ESCC also plans to work with the newly formed User Advisory Committee to take recommendations for Community Codes and report on current initiatives. In anticipation of this effort, ESCC has initiated a project to quantify which community codes are being used across the XSEDE centers to better support our users. This will allow the team to focus on codes with the most impact with limited support.

As an example of this, ESCC has created another internal project to support efforts to port and optimize commonly used genomics tools, Trinity and AllPaths-LG, to XSEDE HPC architectures. After seeing multiple requests for assistance in porting/optimizing these codes over the past two quarters, requests have been consolidated into one overarching project. This provides the advantage of tracking one project with a common goal that will support multiple research groups. The PI of the combined project is a member of the ESCC staff that is in communication with the research teams who are dependent upon this software.

The challenges mentioned in the last quarterly, namely support for bioinformatics researchers and project management are being met. The new project management tool should be available before the end of the next quarter, and a number of FTE's have been reassigned to support bioinformatics. The primary challenge going forward will be engaging with the teams who've developed software under NSF programs such as PetaApps and SI2 to determine what software has been developed that we can install and support. Efforts will be coordinated with the XD TIS program.

Below are summaries of the progress of two active ECSS projects as an example of the types of support being provided.

7.2.1 ESCC projects.

7.2.1.1 Interactive Large Scale Media Analytics (PI Virginia Kuhn, USC)

ECSS Team: Ritu Arora (TACC), David Bock (NCSA), Kai Lin (SDSC)

This project was initiated via a startup request, and is the result of a NIP outreach effort by Dora Cai (NCSA) and is being led by Ritu Arora (TACC). This is the second project in ESCC to require database support.

The goal of this project is to use HPC resources to enable interactive, real-time query of the multi-media database for automatic metadata extraction, video analytics and large scale tagging using a single end-user interface with appropriate data visualizations. Due to the large size of the databases, 5 TB, and computational intensity of the analytics, HPC resources are required. The final objective of the ESCC project is to create a database on Gordon that may be queried in real time.

The lead consultant is working with the PI to develop a metadata schema for video analytics that will work well on HPC resources. A test database has been setup on Corral (TACC resource) to investigate the development of a Resource Description Framework (RDF) on the database to allow the Medici web-interface (from NCSA) talk to the database seamlessly. The compute-intensive steps of the video-analytics pipeline will be performed on Gordon. There is also an effort by another ECSS consultant, David Bock (NCSA) to develop a visualization tool to explore the database. Kai Lin (SDSC) is providing assistance for Gordon related issues.

7.2.1.2 Distributed and Loosely Coupled Parallel Molecular Simulations Using the SAGA API (PI Ronald Levy, Rutgers University)

ECSS Team: Yaakoub El Khamra (TACC), Matt McKenzie (NICS)

This project was initiated via an XRAC request and spans three centers, TACC, NICS, and SDSC. The PI proposed to use the SAGA (Simple API for Grid Applications) framework to enable molecular dynamics simulations using NAMD, IMPACT and AMBER to be coupled across multiple systems. ESCC assistance was required to set up and harden SAGA and the associated tools that sit on top of SAGA on Ranger, Lonestar, Kraken, and Trestles. ECSS staff also assisted in developing the workflow for the coupled applications.

SAGA has been deployed on Lonestar, Ranger, Kraken and work is in progress on Trestles. SAGA is in a production phase on single machines, work continues on Lonestar and Ranger to support distributed simulations. Kraken will be added next.

Several tutorials have also been held about the tool, and a submission has been made for a BigJob/SAGA tutorial at XSEDE12 to showcase the tools and their use. Work continues to support, debug and troubleshoot operations for users of SAGA.

7.3 Extended Science Gateways Support 1.5.2

The mission of the Extended Support for Science Gateways (ES-SGW) is to provide assistance to researchers wishing to access XSEDE resources through web portals also called as science gateways and/or data repositories. The group assists both new and advanced groups and has experience in the use of web technologies, grid software, fault tolerance, complex workflows, and security and accounting aspects of the program.

This quarters much of the effort continued to analyze and sustaining the successes of the predecessor TeraGrid gateway program. Significant time was expended in preparation of work plans and scoping the gateway projects that are slightly daunting. The gateway group continued pro-active participation in interacting with XSEDE A&D and SD&I teams in assistance with gateway architecture requirements, assisting with testing and operational plans.

Support for existing projects is taking more time than anticipated and the group is planning exit strategies for gateway support and transitioning of maintenance tasks such as monitoring of grid services to XSEDE operations. Transitioning of these tasks will allow assignments of gateway staff to new activities.

Starting January 2012, the gateway group is starting to step up even more for XSEDE wide activities including requirement gathering, testing and normalizing current usage. Gateway staff will need to work with different XSEDE teams to maximize gateway contributions with minimal time commitment.

Gateways are heavily vulnerable to XSEDE security and architectural changes. If architecture is defined early, it will greatly help the EC-SGW group to scope support projects. Currently, completion of some workplans is delayed because of the difficulty in estimating effort to complete the projects. A project management tool will assist very much in activity tracking with multiple parallel involvements and especially with short-term XSEDE-wide interactions.

Effort is being this quarter to complete workplans for work initiated in the TeraGrid program that have been extended. A total of 10 Science Gateway projects are ongoing.

7.3.1 ECSS SGW projects.

7.3.1.1 OLAM (PI Craig Mattocks, University of Miami)

ECSS team: Yu Ma (IU)

The objectives of the OLAM extended support project is to deploy the new global icosahedral Ocean Land Atmosphere Model (OLAM) Earth System Model on XSEDE computational resources in order to simulate current and future climate scenarios, high- impact storm events (hurricanes) and local circulations; develop an OLAM web portal science gateway for teaching graduate-level meteorology, climate, and predictability courses at the University of Miami, and generate regional climate change projections that can be used to guide water management decisions in South Florida.

In this quarter the support focused on building gateway infrastructure for classroom teaching of OLAM model at University of Miami. Specific activities include building a portal interface to specify model domain, a user interface to configure temporal information and model input configuration and a wizard style interface to launch the model on specified grids.

Next quarter will focus on integrating the gateway/portal with back end services to launch, monitor jobs to XSEDE resources and build interface to download/visualize model outputs.

7.3.1.2 Dark Energy Survey Simulation Working Group (PI Andrey Kravtsov, University of Chicago)

ECSS team: Raminder Singh (IU)

The Dark Energy Survey (DES) is an upcoming international experiment that aims to constrain the properties of dark energy and dark matter in the universe using a deep, 5000-square degree survey in five optical bands. The production of a suite of large N-body simulations will allow the development of a Blind Cosmology Challenge (BCC) process for the DES collaboration. The gateway extended support is focusing on developing workflow solutions for (N-body production): 2LPT + L-GADGET codes. The workflow will prepare input files for these tasks, submit 2LPT initial conditions (ICs) code, verify correct execution, archive IC's, submit L-gadget simulation code, verify correct execution, archive output, re-submit and verify until code terminates successfully.

This quarter focused on:

- Created the workflow for L-gadget application on Ranger
- Configuring grid job submissions to launch custom libraries. This task involved co-ordination with Globus team and TACC system administrators.
- Deployed the workflow services on IU Quarry gateway hosting machine
- Documented and demonstrated the workflow creation and execution steps to the project team.
- Analyzed the performance of RNN code. The analysis revealed that this code is not suffering from a single bottleneck, but that almost all time is spent in 8 equally time consuming routines. Therefore an optimization is difficult. There is also some considerable load imbalance, which is expected in a code that performs tree-searching. Lastly, memory access within the tree search can be improved by ordering the output within Gadget. Code that would do this does already exist and a simple test will reveal if the compute time is better spent in preordering the output in Gadget or suffering through non-optimal memory access in RNN.

Optimization findings in detail:

- On a smaller scale serial performance was checked and there was no single bottleneck, but a whole array of routines that accounted for 10-15% of the computing; and none of them offered good optimization opportunities. At production scale there is quite some load imbalance, but that is to be expected in any tree-based code and remedies are difficult.
- To improve memory access the data may be written out in a more organized fashion by Gadget. The code to do this is already there. The time in Gadget would increase, but the time in the tree-search would decrease. So simple tests will show whether this is worth it.

Next quarter will focus on:

- Providing automatic submission of L-gadget code in case a simulation exceeds the maximum queue walltimes
- Providing data movement steps to persistent storage at TACC Ranch or IU HPSS
- Test workflow with production like data and release the workflow for production runs.

7.3.1.3 CIPRES (PI Mark Miller, UCSD)

ECSS team: Terri Schwartz (SDSC), Wayne Pfeiffer (SDSC)

CIPRES gateway focused this quarter on installation and availability via the CIPRES gateway of enhanced versions of heavily used phylogenetics codes. The team also focused on availability on one or more XSEDE resources of bash script for doing comprehensive phylogenetic analyses of large data sets using RAxML-Light, Parsimonator, and RAxML. Numerous enhancements were made to the bash script for doing comprehensive phylogenetic analyses with RAxML-Light on

Trestles. Christian Goll of the Heidelberg Institute for Theoretical Studies wrote the initial version of the script the previous quarter in consultation with Pfeiffer and Alexis Stamatakis, also at HITS. The recent enhancements are as follows:

- Input random number seeds now give reproducible results.
- Bootstrap searches use the -D option, which allows them to run faster than regular searches.
- A "usebest" option was added to start the regular searches from the best bootstrap searches.
- An option to handle partitioned alignments was added.
- It is possible now to do only bootstrap searches or only regular searches. In this case, a comprehensive analysis is not done.
- A test for "bootstopping" convergence is always done whenever bootstrap searches are done. This is an after-the-fact test. That is, the number of bootstrap searches is specified at the outset and is not increased regardless of whether the convergence test is passed.
- A Perl script was written to summarize usage of the many jobs spawned by the bash script and was incorporated by Goll in the bash script.
- An option was added to pack multiple executions of Parsimonator and RAxML-Light within a single node. This allows efficient use of the normal queue instead of the shared queue on Trestles and can be used on other computers that do not have a shared queue.

CIPRES support also focused on additional benchmark runs of the BEAST phylogenetics code and its associated BEAGLE library using two more data sets from users. Redid some previous runs and updated the CIPRES Web page with the new benchmark results, all obtained on Trestles.

Based upon these results, one of the rules for selecting the number of BEAST and BEAGLE threads when running via the CIPRES gateway was revised. The rules still call for BEAST jobs to use 8 threads total in the shared queue on Trestles. The change is to use at most 4 BEAST threads with 2 BEAGLE threads each, even for data sets with large numbers of partitions. The CIPRES interface has been updsted accordingly.

In response to user-reported problems with BEAST 1.6.1, BEAST 1.6.2 was installed and tested on Trestles and put it into production in the gateway. This fixed one problem, but not another, which involves the interface between BEAST and the BEAGLE library. The latter problem is reportedly fixed in BEAST 1.7, which should be available soon.

Other CIPRES activities include availability on one or more XSEDE resources of a bash script for doing comprehensive phylogenetic analyses of large datasets using RAxML-Light, Parsimonator, and RAxML. The foundations were laid for exposing RAxML-Light for the CIPRES Science Gateway and the iPlant discovery environment. This work involved creation of infrastructure, including:

- Preliminary investigation into making RAxML-Light script sufficiently robust to run from the portal
- "build and test" environments with Apache, Tomcat and Shibboleth for multiple developers
- Consultation with iPlant developers and implementation of the Shibboleth authentication system required by iPlant. This work is in progress, leading to single sign-on capability for iPlant users.
- Evaluation of Apache Airavata's workflow system for running RAxML-Light using the CIPRES Portal. Airavata's workflow manager uses GRAM, which hasn't been

dependable enough for CIPRES in the past. Use of Airvata workflow manager to run RAxML-Light in the iPlant DE will be investigated once RAxML-Light is working within the iPlant DE.

- Installation and availability via the CIPRES gateway of enhanced versions of heavily used phylogenetics codes.
- Coordination of the updating of existing codes on the Gateway, and release of new patches for bugs that were discovered in individual interfaces.

7.3.1.4 Multi-scale Computational Models (PI Mohammad Mofrad, UC Berkeley)

ECSS team: John Lockman (TACC)

The effort focused on understanding the project requirements issues related to Condor jobs. The team assisted with simple Condor jobs using the Condor standard universe and resolved issues with Condor submissions.

7.3.1.5 Galaxy Deployment for the Scripps Research Institute (PI Rathi Thiagarajan, TSRI)

ECSS team: Suresh Marru (IU)

The project focused on assisting the researchers at TSRI to manage, map, and analyze the terabytes of data generated from gene sequencing experiments. The project plans to sequence multiple complete human methylomes by using specialized mapping software required such as Bismark, BS_Seeker. These software programs map the raw reads to four separate human reference genomes using Galaxy and Bowtie.

This quarter focused on setting up a local instance of Galaxy on the Quarry gateway hosting service, installing and maintaining the needed applications on IU Mason cluster leveraged from the National Center for Genome Analysis Support (NCGAS) project at IU. The team helped moving data from TSRI to IU Data capacitor and loading it into Galaxy.

Next quarter will focus on engaging the users in using the Mason cluster to first filter raw data using a locally running version of Galaxy and then map these filtered reads to the reference human genome using a third party mapper such as Bowtie.

7.3.1.6 CESM- The collaborative CESM-ESG integration project (PI Don Middleton, NCAR)

ECSS team: Don Middleton (NCAR), Lan Zhao (Purdue), Christopher Thompson (Purdue)

The integrated Community Earth System Model (CESM) and Earth System Grid (ESG) portal project focused on completing the critical features in the original project plan to support various user needs and enhance the user experience. Specific work activities include:

- Support for large files to handle through conventional HTTP protocols and improved input uploading and output downloading functions.
- Upgraded to latest version of CESM model
- Testing of a pre-release version of the CESM model against the ESG network for the purposes of greater metadata generation and handling.

Next quarter will focus on:

- Developer transition.
- Preparation for CESM model release, testing and validating model cases, building additional ESG publication options into the system for users such as the ability to unpublish previously published datasets,

- Supporting a graduate classroom use-case for the portal, and integration of post-processing codes upon datasets published to the ESG network.

7.4 Extended EOT Support 1.5.3

ESTEO work for the period includes an intense focus on preparing materials for upcoming conferences. XSEDE12 and SC12 are targets of several tutorials and presentations by our staff. Phil Blood has been coordinating the technical program for the conference and getting all of the various program chairs working together in anticipation of reviewing the proposed papers, posters, tutorials, and presentations. Some of the upcoming tutorials will cover:

- Accelerator programming: OpenACC and intermediate CUDA
- Introduction to XSEDE
- Introduction to Unix/Linux
- Eclipse IDE and the Parallel Tools Platform

During the period, staff continued to provide outreach and training activities even as they prepare new material for large events later in the year. For example, in January, Marcela Madrid presented our XSEDE New User training webinar and she's also working on the Accelerator programming tutorial listed above. On the outreach side, Amit Majumdar hosted a professor and student from Instituto Politechnico Nacional in Mexico City at SDSC for a week and had great interaction with them -- showing them the HPC ropes.

At NICS/RDAV, an intensive HPC workshop was run Mar. 19-22 that covered many aspects of HPC: parallel programming and I/O, GPU programming, Numerical libraries, and various other topics covering tools and debugging techniques. Most of the XSEDE staff based at NICS took part in the workshop.

Jay Alameda presented half and full day tutorials on the Eclipse Parallel Tools Platform (Eclipse PTP) to a visiting team from KISTI and also to the UCAR software engineering assembly in Boulder.

We also continue to review web tutorials and provide new content. The multicore tutorial in ci-tutor is under revision now and will be released in the next few days. It has been revised after a thorough review from Marcela Madrid and Galen Arnold. Sandie Kappes in the EOT group is busy reworking the updated material and designing new quizzes and questions to go with it. Amit Chourasia reviewed our online visualization tutorials.

The full list of training courses for the period is now available online at https://www.xsede.org/web/xup/course-calendar.

8 Education and Outreach 1.6

8.1 Overview

Overall, the Education and Outreach teams are on track with their planned activities. There activities have made slower progress than expected (student program, Rice faculty council, and formation of TEOS Advisory Committee) but efforts are underway to catch-up in these areas to be on track for the project year.

During this quarter, 17 new education accounts and 154 new startup accounts were allocated. XSEDE continues to engage new communities in using XSEDE's resources and services.

8.2 Education 1.6.1

Education activities for the third quarter continued all of year one activities: 1) conducting educational programs for faculty and students, 2) development of competencies for computational science programs, 3) initiating collaborations with universities interested in starting or expanding formal computational science degree and certificate programs 4) Creation of computational science materials for pre-service science and math educators.

8.2.1 Educational Programs for Faculty and Students

As part of the education program, we have been working on several efforts to create educational materials and present educational opportunities for faculty and students. James Demmel from the University of California, Berkeley began teaching his CS267 Parallel computing course[4]. The video lectures for that course are being captured for later conversion into an online offering with the assistance of the XSEDE staff at Cornell University. In addition, we have been planning for the conversion of class exercises and instructions so that they will run on XSEDE systems and we can pilot both an online and live, distributed version of the class in year 2 of the XSEDE program.

We have also been planning for a number of workshops for faculty this spring and summer. We have confirmed two one-day workshops providing an overview of computational science educational resources and tools for community college faculty in Nevada as college sites in Las Vegas and Elko. Those workshops are being co-sponsored by the Nevada EPSCoR office.

We also have commitments to run five summer workshops on a variety of topics:

- Computational Chemistry for Chemistry Educators (Pennsylvania) June 11-15
- Computational Thinking from a Parallel Perspective: (Louisiana) July 8-14
- Computational Biology for Biology Educators (Maryland) July 16-18
- Intro to Parallel Programming and Cluster Computing: (Oklahoma) Aug 5-11
- Data Analysis and Visualization (New Jersey) TBD

The workshops will provide priority registration to the faculty at universities that are currently working with XSEDE on the establishment of formal programs. Participant costs will be subsidized by the XSEDE education program.

Planning is also underway for the second offering of the parallel programming summer bootcamp at University of California, Berkeley. Last year that workshop attracted the participation of several hundred people from around the world. The XSEDE education program has also agreed to help sponsor a computational physics developers workshop to help build examples that can be incorporated into undergraduate physics classes.

[4] http://www.cs.berkeley.edu/~knight/cs267/resources.html

8.2.2 *Development of Competencies*

Work continued on the development of model competencies for both undergraduate and graduate programs in computational science. The current version of the competencies can be seen at http://www.rrscs.org/competencies. Over the past few months, we have reviewed and revised the undergraduate competencies with assistance from some of the faculty at several of the institutions with which we are working on formal programs. The graduate level core competencies have also been reviewed and now awaiting the addition of domain specific competencies in various science and engineering domains. The core competencies at both levels are now sufficient to help guide the initiation of new computational science programs. As we work with institutions on their implementation of such programs, we will document how the competencies are being used to guide course development and incorporate any newly suggested changes or variations in how the competencies are being used. We will also begin to compile a set of course descriptions and syllabi needed to define the courses that are used in those programs and make those available to other institutions interested in related course development.

8.2.3 *Initiating University Collaborations*

Over the past quarter, we have continued work with the campuses interested in formal computational science programs. We have been engaging with the campuses through audio conferences and campus visits. These sessions have included a description of the model programs and competencies, discussion of the current expertise on each campus, and establishment of working groups to help create or enhance formal computational science programs. Table 1 summarizes the activities for the current quarter with each institution.

Table 1: Institutions with Interests in Computational Science Programs

Institution	Status and Interests
Clark Atlanta	Visit made in March 2012. Working committee established. Program development confirmed as a part of the institutions strategic plan.
Clemson University	Several conference calls to discuss arrangements for a visit and organization to look at potential programs. Visit on hold until Autumn 2012.
Finger Lakes Community College	Presentation made at workshop in Austin to introduce XSEDE and possible collaboration on an NSF project to introduce research into the community college curriculum.
Kean University	Expanded program endorsed by campus administrators. Campus visit upcoming on April 26.
Montgomery Community College	Conference calls indicated most interest in computational biology. Hosting one of the summer workshops to begin faculty professional development
Norfolk State University	Interested in graduate programs in computational science. Visit forthcoming April 17, 2012.
North Carolina	Made visits to North Carolina Asheville, North Carolina State, University of North Carolina, and General Administration of the NC System. All potentially interested in statewide effort.

	Series of video conferences with the faculty have been scheduled.
North Carolina Central	Interested in undergraduate major spanning the sciences. Could be integrated into statewide effort.
Southern University	Interested in both undergraduate and graduate concentrations and new graduate degree. Visit completed on April 2, 2012 and working committee formed to start program preparation.
Stockton College	Expanded program endorsed by campus administrators. Campus visit upcoming on April 27.
The Ohio State University	Working on integration of modeling and simulation into pre-service science teacher education program

8.2.4 Other Activities

A number of visits were made this quarter to engage faculty and students in computational science activities. Theses are summarized in Table 2.

Table 2: Visits Made	
Location and Date	Description
Jan 20: East Carolina University	2 hour intro to modeling workshop, 6 faculty attended
Feb 1-3: CI-TEAM/CE21 meeting, NSF	Panel and then Parallel Thinking session lead, 9 faculty attended
Feb 6-7: Richmond Public Schools	Computational science and mathematics
Feb 9-16: Dublin City University, Ireland	Computational Science: Dynamic Explorations, funded by int'l grant to Shodor, partially funded by Noyce
Feb 19-20: VMI campus visit	Faculty consults, department meeting, Dean's Seminar for cadets
March 1: North Carolina School of Science and Math	Innovations in Computational Science
March 2-11: Visit to Monash University	2 half-day workshops (one for teachers from John Monash School, one for undergrads), colloquium, individual faculty visits, 3 meetings of curriculum committee putting together new masters
March 23-24: North Carolina Academy of Science	Computational Thinking: Dynamic Explorations for a Changing World workshop for 35 (25 college faculty, 5 high school teachers, 5 college students)

We also continued our work with the pre-service teachers program at The Ohio State University.

During winter quarter, Dr. Karen Irving and a graduate student continued to work on models and modeling as topics in science education pre-service teacher education instruction and for students in middle and high school classroom. Syllabi for new semester iterations of our science methods classes explicitly include modeling as topics for pre-service teachers. Progress has been made on developing lessons for pre-service teachers related to the water cycle. We are currently working on producing a set of lessons related to modeling in each of four science disciplinary fields (chemistry, physics, earth & space, and life science).

Our final activity was the creation of a preliminary, unified ontology for computational science education based on several existing schema and related efforts for a computational science education roadmap. The resulting scheme will be used to create the metadata for computational science training and education materials and translated into an online resource for the XSEDE community. We expect a test metadata schema to be completed in April and then tested against some initial materials. We will then be able to design the relevant infrastructure and procedures to create a server to distribute the information and the processes for receiving and reviewing materials.

Planning for the 3rd annual EU-US HPC summer school has progressed. There were 236 applications from graduate students and postdocs, of which 25 from the US and 35 from Europe were selected. The selection of presenters and topics is nearly complete. The event will occur June 24-28, 2012 in Dublin, Ireland. Additional information is available at https://www.xsede.org/web/summerschool12.

Planning for the 6th annual Extreme Scaling Workshop continues. The workshop will be held July 15-16 in Chicago. The call for submissions closes in April, after which the final agenda will be posted and registrations for the workshop will open. Details about the workshop are available at: https://www.xsede.org/web/xscale.

8.2.5 Challenges

The main challenge continues to be is the slow pace of change for programs at the universities. Faced with very tight budgets at all higher education institutions, the prospects for creating new programs may become even slower. Nevertheless, an interdisciplinary program like computational science is being well-received on the campuses and makes it likely for us to succeed in establishing new programs in the long run.

These are all activities that promoted computational science education:

- Jan 20: East Carolina University, 2 hour intro to modeling workshop, 6 faculty attended
- Feb 1-3: CI-TEAM/CE21 meeting, NSF, panel and then Parallel Thinking session lead, 9 faculty attended
- Feb 6-7: Richmond Public Schools, Computational science and mathematics ***18 AP teachers
- Feb 9-16: Computational Science: Dynamic Explorations, in partnership with Dublin City University, funded by int'l grant to Shodor, partially funded by Noyce *** 23 faculty or teachers, 5 American pre-service teachers
- Feb 19-20: VMI campus visit, faculty consults, department meeting, Dean's Seminar for cadets ***60 cadets, 8 at faculty meeting, 4 individual consults
- March 1: Innovations in Computational Science: North Carolina School of Science and Math ***2 teachers, 14 students
- March 2-11: Visit to Monash University: 2 half-day workshops (one for teachers from John Monash School ***15 attendees, one for undergrads *** 10 attendees), colloquium

***20 local attendees and 5 r more at remote sites, individual faculty visits, 3 meetings of curriculum committee putting together new masters

- March 23-24: Computational Thinking: Dynamic Explorations for a Changing World, at North Carolina Academy of Science, workshop for 35 (25 college faculty, 5 high school teachers, 5 college students)

8.3 Outreach 1.6.2

8.3.1 Underrepresented Engagement

Accomplishments:

SURA is deepening its involvement with MSIs as planned. In addition to participating in the Emerging Researchers Network Conference, there have been campus visits to:

- Clark Atlanta University: This visit resulted in a follow-up visit with the XSEDE Education group to discuss how to integrate computational science into their curriculum. CAU has also begun the process to become a Campus Champion site.
- Norfolk State University: Following up with Education group
- North Carolina Central University: Following up with Education Group
- Southern University: Following up with Education Group
- Florida A&M University
- Bowie State University

The Minority Researchers Community now has a mailing list set up and are holding group calls once or twice per month to address topics of mutual interest.

From Rice University's XSEDE efforts:

Following a presentation by Richard Tapia to the XSEDE Leadership team at the Quarterly meeting in March, 2012 in Austin, TX, the Minority Faculty Council has been formed. The roles this group of academic leaders are expected to play include serving as a resource for the XSEDE partnership, giving information and advice about broadening participation, i.e. an advising capacity with "bite" and as mentors for students, with the goal of raising retention among minority students. Richard's talk to XSEDE Leadership is available on the XSEDE staff wiki (www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Minority+Faculty+Council) Members are:

- Raquell Holmes – University of Connecticut & Boston University
- Roscoe Giles – Boston University
- Leticia Velasquez – University of Texas at El Paso
- Jeanine Cook – New Mexico State University
- Juan Vargas – Microsoft
- Dorian Arnold – University of New Mexico
- Legand Burge – Howard University

One additional member may be added, from a Tribal College. The members will conduct a face-to-face meeting at XSEDE'12.

The first year of the XSEDE Scholars Program provided five webinars for their students:

- Juan Vargas: Innovation in Cloud Computing
- Maytal Dahan: Introduction to XSEDE User Portal
- Dan Stanzione: Parallel Programming with MPI, part I
- Dan Stanzione: Parallel Programming with MPI, part II

- Roger Moye: What is the XSEDE Campus Champions program; Intro to HPC; How to submit jobs on an Xsede machine

To facilitate full participation in these sessions, all of the XSEDE Scholars set up portal accounts and were added to an XSEDE Education allocation, giving them access to basic XSEDE resources.

In addition to these successful programs, NSF requested a “Highlights” article from the Empowering Leadership Alliance. Since there is much overlap between the goals and offerings of the XSEDE Scholars Program, and the ELA, the article ended up highlighting features of both programs. The article is on the XSEDE staff wiki at https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/XSEDE+Scholars.

Challenges

- Faculty from MSIs have limited time to devote to writing applications for allocations (teaching workload, lack of student support at their institutions). This makes it much more challenging to meet their needs for acquiring and utilizing XSEDE resources.
- Based on feedback from the evaluation of XSEDE Scholars, training, and Campus Champions, there are indicators that our newer constituencies are less advanced than we expect for the content we are providing. We need to identify the deficiencies and determine how XSEDE can/should address them if we are going to successfully grow the XSWEDE user base and HPC workforce.

8.3.2 Speakers Bureau

Accomplishments

XSEDE was asked, by the Clemson and University of Kentucky Campus Champions, to provide an orientation speaker for the Clemson EPSCoR Student Engagement program. The enthusiasm and interest in these students was evident in the questions they had about supercomputing and computation science, as well as their enjoyment and participation in the Clemson machine room and automotive design center tours. All of the 83 students have been added to the Student Engagement email list and have received information about the summer projects. A number of the early applicants have cited the orientation session as their source of information about the availability of the opportunity.

In addition to the Clemson request, XSEDE was also present at the following events:

- American Physics Society meeting, Boston, MA
- Emerging Researchers Network annual meeting, Atlanta, GA
- National Society of Black Engineers, Pittsburgh, PA

128 participants at these events expressed interest in further contact with XSEDE. The students have been added to the mailing list and some have already applied for the Student Engagement projects. Researchers and educators have been handed off to NIP and TEOS Education for appropriate follow-up.

Challenges

The Speakers Bureau planning process needs input from other XSEDE areas with respect to which events to target, especially in traditional fields of science. Support for PY1 events has been weak outside of Outreach team members.

8.3.3 *Student Engagement*

Accomplishments

In addition to helping start up the Clemson EPSCoR Student Engagement program, the XSEDE team has begun identifying projects and recruiting students for summer opportunities. Twenty-two projects were submitted from seven XSEDE sites. 161 applications were submitted to the program, with 126 students completing the process (by supplying resumes and transcripts or other proof of student status). An initial scan of where the students are from and how they heard of the program seems to indicate that there has been a lot of cross-group participation happening, to raise interest and awareness among students. More detailed analysis will be conducted after the placement process has been completed.

Applicants' information is being distributed to the supervisors of the projects they expressed interested in. Interviews will be conducted by teleconferencing and decisions should be made shortly thereafter. The program is expected to begin with a synchronous orientation session, based at the three lead Student Engagement sites (PSC, NICS, TACC), with remote participation by supervisors and students at the other sites. The projects are scheduled to run through the first week of August, to allow the students time to return to their homes/campuses in time to prepare for the fall semester.

Challenges

The portal interface (forms) for project and student applications are inadequate and push too much processing back on the Student Engagement team. The Student Engagement team is investigating alternatives for future years.

8.3.4 *Campus Champions*

Accomplishments

The Campus Champions program is showing strong signs of community building among the Champions. There are now 112 campuses and 153 Champions. In the current quarter, the group has initiated over a dozen conversations addressing challenges they have found in performing their duties as Champions and in supporting their campus users. Some examples:

- Clemson EPSCoR Student Engagement Program: drawing on XSEDE Student Engagement and Campus Champions for speakers and mentors
- Training: initiated a survey of Champion training needs. Cross-listing with what is currently available (or should be) in the XSEDE training suite versus what specific needs Champions have for advocating and supporting their campuses effectively. This has already resulted in a "New Champions Training" webinar, which will be repeated on a regular basis as the Champions program continues to grow.
- Creation of two "working groups" to address what the Champions have identified as their most critical needs: outreach to their campuses and providing effective user assistance. While these groups are just getting started, it is a clear sign of the health of the community that they are initiatives from the Champions and not driven by an externally perceived need.

Cross-agency and international opportunities:

- The Statement of Work for the Champions/OSG liaison role was finally completed and agreed upon by both XSEDE and OSG.
- HPC Wales has been in touch with the Champions program as they consider ways to begin a similar program for their community.

- The MARC program (from NIH), through an award to PSC, is working with the Champion at Tennessee State University to investigate the possibility of developing an internship program similar to MARC.

The Campus Champions Fellowship program experienced many of the same startup issues that the Student Engagement program had. 19 projects were approved by their PIs to be offered to the Champions. The application for participation opened in April, 2012 and Champions will have one month (until May 6) to apply. A review team is in the process of being constituted. The first two Fellows will be placed in the fourth quarter of the first plan year.

Challenges

Due to financial constraints at Purdue University, the leader of the Campus Champions program, Kay Hunt, may not be able to continue in her role beyond June 30. Her non-XSEDE funding is being cut, and she may need to leave XSEDE to find a position that will provide fulltime support. We are working to resolve this before the risk of losing Kay escalates further.

8.4 Community Requirements 1.6.3

8.4.1 Highlights

The TEOS Community requirements activities remain on track/target as stated last quarter. In January 2012, Edee Wiziecki and Scott Lathrop with input from Level 3 Managers produced a document describing the role and responsibilities of the TEOS Advisory Committee members to ensure that both internal TEOS members and the external Advisory Committee members understood the role of the AC (document posted on Wiki).

The TEOS Level 3 Managers with input from their partners compiled a list of potential Advisory Committee members. Names of potential members were discussed among Managers during several phone meetings in January to ensure the Committee represented members experienced in areas that the TEOS activities targeted. The following members were invited and accepted to serve:

- Carlos Castillo-Chavez, Arizona State University (*)
- Roscoe Giles, Boston University (*)
- Gwen Jacobs, Montana State University (*)
- Rubin Landau, Oregon State University (**)
- Wilf Pinfold, Intel (**)
- Nora Sabelli, SRI (**)
- Valerie Taylor, Texas A&M University (*, ***)
    - * - serving 3 years
    - ** - serving 2 years
    - *** - member of the XAB

Valerie Taylor agreed to be the liaison to the XSEDE Advisory Board (XAB), as she serves on both committees. The first teleconference meeting of the TEOS AC was held on March 14, 2012 with a follow-up call on March 26, 2012. The purpose of the first call was to update the TEOS AC members on the activities conducted thus far and to receive input from the AC members on our direction, etc. Each of the TEOS Managers presented an overview of their area of concentration and AC members provided their input. An additional teleconference meeting is scheduled for April 24, 2012 for interested AC members to discuss the external evaluator matrix.

In February, the CR Manager oversaw the production of a community requirements survey with input from TEOS Managers and based upon prior feedback from the community to receive comments on current planned activities and to gauge the interests of the community on other activities. It was disseminated using Survey Monkey to several large partner listservs, and to a large list of names associated with TeraGrid EOT (see survey on Wiki). There have been over 70 responses to the survey. The deadline for completion is April 15, 2012.

Next Quarter Focus: The first face-to-face meeting of the TEOS Advisory Committee and Level 3 Managers will be held in Chicago on May 22, 2012. We are currently planning the logistics and agenda. The CR Manager will oversee the analysis of the survey results to be disseminated in June 2012.

Risks: **(1)** There are no new risks.

8.4.2 Evaluation

The External Evaluation team has continued to work through the plans they have developed and documented in the TEOS Evaluation Matrix, that is posted on the wiki.

8.4.2.1 Training

Evaluators have drafted a suite of tools for the training evaluation. This suite includes a Post Training Session Survey, Semi Annual Training Survey, Power User Interview Protocol, and a Non PI Power User Interview Protocol. The Post Training Session Survey will be administered via the XSEDE User Portal after users complete online training sessions. This survey will capture participant satisfaction with the training session and the added value of the session to their current work or research. Users will be reminded to complete surveys for past training sessions when they log in to register for additional classes. The Semi Annual Training Session Survey will be administered via the XSEDE User Portal to users who have participated in training during the last 6 months. The semiannual form will be used to track participants longitudinally to determine the impact and outcomes of training programs. Evaluators are using the Power User Interview Protocol to conduct interviews with PIs who have active allocations of 5 million hours or more. With the permission of their PI, evaluators will also interview members of Power Users' labs using the Non PI Power User Interview Protocol to uncover how XSEDE resources are introduced, adopted, and distributed in the project's most active labs.

The evaluation is currently in the final stages of review. We expect this suite to be live and in use by June 2012. Evaluators will also utilize the longitudinal tracking system to further investigate the impact training activities have on research and education outcomes amongst users.

8.4.2.2 Education

Evaluators have been compiling Dr. Steven Gordon's completed Observation Guides to ultimately build case studies on HPC program implementation in higher education. Recently, Dr. Gordon completed a 3 day workshop at Southern University where attendees, including faculty and administrators, are interested in implementing an undergraduate concentration in HPC. Evaluators will follow up with participants from this workshop through a web based survey and phone interviews in April 2012.

8.4.2.3 Outreach

XSEDE Scholars Program: The evaluation has collected and analyzed data from four of the five total webinars that the XSEDE Scholars Program (XSP) has sponsored. Data was not gathered from the first webinar on November 3, 2011, since the program evaluation plan had not yet been developed. Interim Reports for each survey administered have been generated and submitted to program coordinators, Laura McGinnis, Scott Lathrop, and each webinar presenter. In general, the scholars have found the webinars to be useful. Attendance for the webinars however has been steadily decreasing due to a class or research schedule conflict, as reported by the scholars. To address this, program coordinators broadcasted the webinar given on March 21, 2012 at a different time. Evaluators are compiling formal recommendations based on their findings to inform year 2 program plans and scholar selection.

Maytal Dahan has also completed generating the longitudinal tracking system to monitor the scholars post participation. Items of interest include training sessions taken and affiliation with any allocations. Testing of this system will begin at the end of April 2012.

SURA: Evaluators are working with Linda Akli to monitor the impact of the minority research community as well as the spring regional events. Lorna Rivera will attend the May 7-8 workshop at Vanderbilt University to conduct observations, interviews, and administer surveys.

Campus Champions: The Campus Champions Training Needs Survey was active on February 3 to February 19, 2012 and had a 49% (73/149) response rate. The survey was conducted with the purpose of identifying areas where Champions need tailored materials. Over 50% of Champions consider themselves novice to low experienced as Campus Champions, reinforcing the need for tailored training materials to be developed. Based on the evaluation results, the following recommendations were made to program staff: Provide a link to a comprehensive list of tailored real time training resources on the Campus Champions website; and new Champions should be matched with experienced Champions at program induction. To view the detailed interim report, please contact Laura McGinnis.

Evaluators are attending calls for the newly developed Outreach and User Assistance working groups. Lorna Rivera will attend their first face – to – face meeting on May 23 – 24 at Notre Dame.

Student Engagement Program: Evaluators have developed a short demographic survey to be administered to all program applicants prior to notification of acceptance into the program. This is expected to go live in mid April 2012. Lorna Rivera will attend the kickoff event at PSC May 30 – 31 (tentative dates) to conduct observations and gather baseline data on students and supervisors.

8.4.2.4 Campus Bridging

Evaluators have updated the TEOS Evaluation Matrix to include plans for the GFFS Early Adopters Program. We are currently working with Craig Stewart and Rich Knepper of IU to formalize the evaluation plans. To date the plans include a pre implementation survey to capture PI expectations, on-site observations, a post implementation survey and a 1 year follow up to compare their current research activities to pre-GFFS implementation. Areas of interest include how results vary according to Carnegie classification, number of students, faculty, allocations, and Campus Champion association.

*8.4.2.5 **Major Findings***

Training: Power User PIs lack understanding of XSEDE resources and deployment in their labs. To address this evaluators will conduct interviews with lab managers, post docs, and graduate students in Power User labs with permission from PIs.

Outreach: The level of HPC knowledge that XSEDE Scholars possess is much lower than what XSEDE staff expected. This gap in knowledge has been reflected through student report and a decrease in student webinar attendance. Evaluators recommend this be addressed through the year 2 selection process by selecting students with a deeper knowledge of HPC, or maintaining the current selection criteria while modifying the webinar content to be more introductory.

8.4.2.6 Plans for Next Quarter

Evaluators will increase data collection during the fourth quarter of year 1. Many of the Outreach and Campus Bridging activities are expected to occur during this quarter.

Evaluators also plan to be involved in the TEOS Advisory Committee by regularly reporting findings and presenting to the Committee. Lorna Rivera will attend and present at the May 22 meeting in Chicago.

The annual EU-US Summer School will be incorporated into the evaluation in order to capture the longitudinal impact of the school. Evaluators are currently developing an evaluation plan for the June 2012 Summer School in Dublin.

XSEDE12 will be evaluated by Lizanne DeStefano and Lorna Rivera. We have received UIUC IRB approval for amendments submitted to the original XSEDE protocol. Particular attention will be paid to the Student Program affiliated with the conference. A meeting to discuss the evaluation plan with stakeholders will be scheduled in April.

8.5 E&O Infrastructure 1.6.4

The Education and Outreach Blog was grown by 52 entries, nearly all events targeted at educators and students. Events featured in the Blog originate with XSEDE TEOS members, and the scrounging work done by Ange Mason and Jim Ferguson to find events of interest to share with the community.

Updates to the main XSEDE.org web site continue to be done on a request basis to the XSEDE external relations lead. External relations has promised to allow direct access to the public web pages for editing by TEOS staff, specifically the E&O Infrastructure staff, but so far has not allowed that access. We are told this is pending the completion of the External Relations strategic plan. There are a few E&O pages that need timely updates, which will be done in April 2012, and others that would like an overhaul and do not want to do this by request but instead by working with other TEOS staff.

Facilities provided by the XSEDE Portal for use in registration of students for E&O programs in the student engagement and education areas have proven inadequate for the job. Scott Lathrop and Jim Ferguson have been made aware of the problems and required functionality needed of the Portal, and these requirements are being communicated to the Portal team to see if they can be met. If they cannot be met in a reasonable amount of time, alternative (non-Portal) will be identified and developed for use by the TEOS teams.

A new space on the XSEDE portal was requested and created for the campus bridging pilot projects. Members of these projects from different non-partner universities will have access to these spaces along with XSEDE staff who are working on various aspects of the campus bridging effort. The space will only be visible to the participants to start, but is likely to be a public resource eventually, after the pilots have helped identify issues and solutions to problems that will inevitably arise.

8.6 Campus Bridging 1.6.5

The Campus Bridging Team engaged in two major activities during this quarter: definition of use cases for the Architecture and Design team and a pilot project for testing the Global Federated File System (GFFS) software on campuses.

In cooperation with the Architecture and Design and S,D & I teams, the CB team created seven use cases for campus bridging activities and began the construction of requirements and quality attributes for fulfilling these use cases. These use cases cover use of resources by campus researchers, data movement between researchers and XSEDE resources, and making local resources available to the larger community via XSEDE in exchange for other resources. The use cases include:

- InCommon Authentication
- Economies of scale in training and usability
- Long term remote interactive graphic session
- Use of data resources from campus on XSEDE, or from XSEDE at a campus
- Support for distributed workflows spanning XSEDE and campus-based data, computational, and/or visualization resources
- Shared use of computational facilities mediated or facilitated by XSEDE
- Access to "____ as a Service" mediated or facilitated by XSEDE

The GFFS Pilot program identified 2 friendly-user sites for initial testing (Indiana University and Louisiana Tech) and 4 pilot sites, out of 17 total proposals for participation in the GFFS program. XSEDE Operations team has identified resources at each of the XSEDE SP's which will accommodate GFFS testing. The pilot sites are:

- CUNY – PI: Paul Muzio, Champion: none
- KU – PI: Thorbjorn Axelsson, Champion: McMullen
- Miami – PI: Joel Zysman, Champion: none
- TAMU – PI: Guy Almes, Champion: will designate

Work has started with the friendly-user sites to identify requirements for interoperating with XSEDE SP resources. XSEDE will be working with the Galaxy gateway at IU and with Tom Bishop's research team at Louisiana Tech. The pilot project teams have had an initial kickoff meeting and will be meeting bi-weekly in order to move project work forward. The XSEDE User Portal will create a space in which the Campus Bridging pilot teams can work and store documents, similar to the Campus Champions space in the XSEDE portal.

XSEDE Campus Bridging will participate in a panel at the Internet2 Spring Member's meeting April 24th on Campus Bridging initiatives in XSEDE, in the Open Science Grid, at IU, and at Cornell.

9 **TAIS/Audit Services 1.7**

In this TAS quarterly progress report we discuss continued software development, meetings, technical presentations and technical progress made toward scheduled software releases.

9.1 Technical Progress

XDMoD User Interface:

During the last quarter, a major update of XDMoD (XDMoD 1.6) was released by TAS. XDMoD 1.6 (https://xdmod.ccr.buffalo.edu) is the first version providing access to all XSEDE users with an active allocation as well as a public view with substantial functionality. Appendix 1 below shows the current XDMoD usage. While the usage is substantial, we anticipate increased usage in the future as more users become aware of its availability and capability. Major new features of version 1.6 include: an update of the Usage Explorer and the Application Kernel Explorer, a greatly enhanced public view available to everyone and full implementation of the XDMoD roles (Program Officer, Principal Investigator, User, Center Director and Public). Each of these new features will be described in detail below.

Usage Explorer:

The Usage Explorer has been improved both in capability and usability. Particularly noteworthy is the ability to plot multiple dependent variables (y-axes), as shown in Figure 1. Figure 1 shows a screen shot of the upgraded Usage Explorer. Displayed in the window is a plot that shows total XSEDE usage in XD SU's charged for the first quarter of 2012 by job size on the primary (left hand side) axis and average wait time per job on the secondary (right hand side) axis. The plot was created by selecting "Job Size" from the "Groupings" heading and "Total XD SUs Charged" and "Avg Wait Hours Per Job" from the "Metrics" heading – three clicks total. Under the "Chart Options" heading, the grouping and metric(s) that are plotted are displayed. The actual data upon which the plot is created is displayed in the window below the plot, which can be hidden from view if desired. The plot can also be made available for the custom report generator by dragging and dropping the chart to the "Report Generator" tab or by clicking the box that reads "Available For Report". It can also be exported in either PNG (portable network graphics) or EPS (encapsulated postscript) format. The data itself can be exported in either CSV (comma separated values) or XML (extensible markup language) format. As shown in Figure 2., the data can be filtered in a variety of ways to display only a desired subset of the data. For example, the plot in Figure 2 was filtered to display only the "NICS-KRAKEN" data. Taken in its entirety, the Usage Explorer provides a powerful and flexible interface to facilitate analysis of the XSEDE usage data.

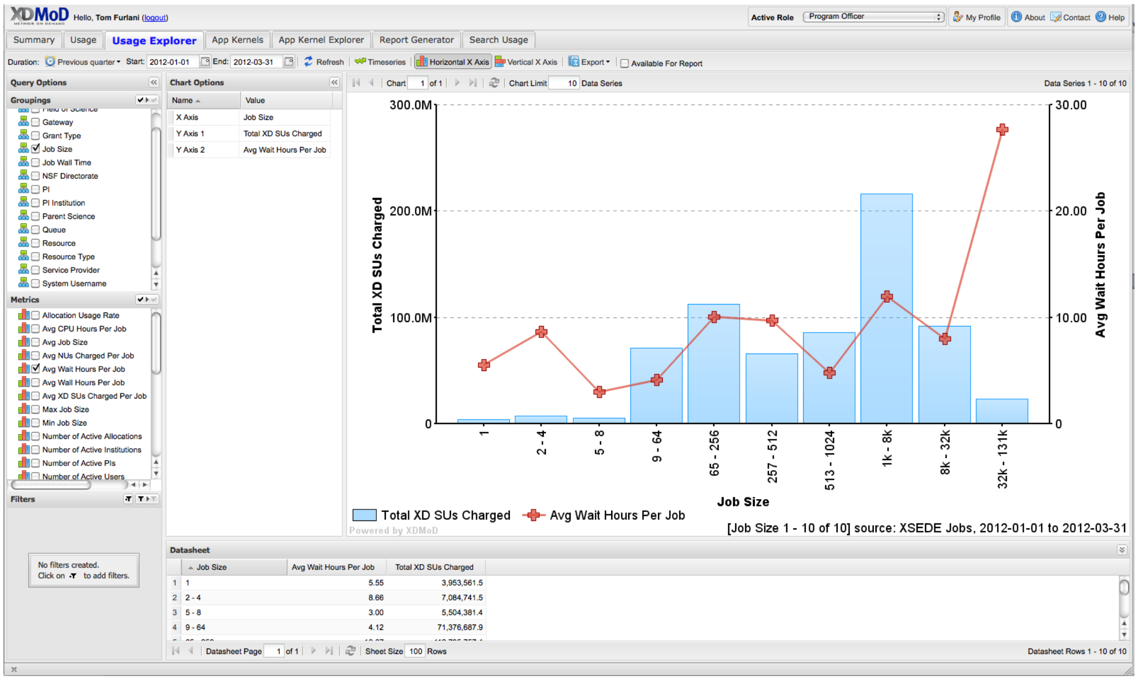

Figure 1. The XDMoD Usage Explorer. A plot of the total XSEDE usage for the first quarter of 2012 is shown on the primary axis (left hand axis) and the average wait hours per job is shown on the secondary axis (right hand axis).

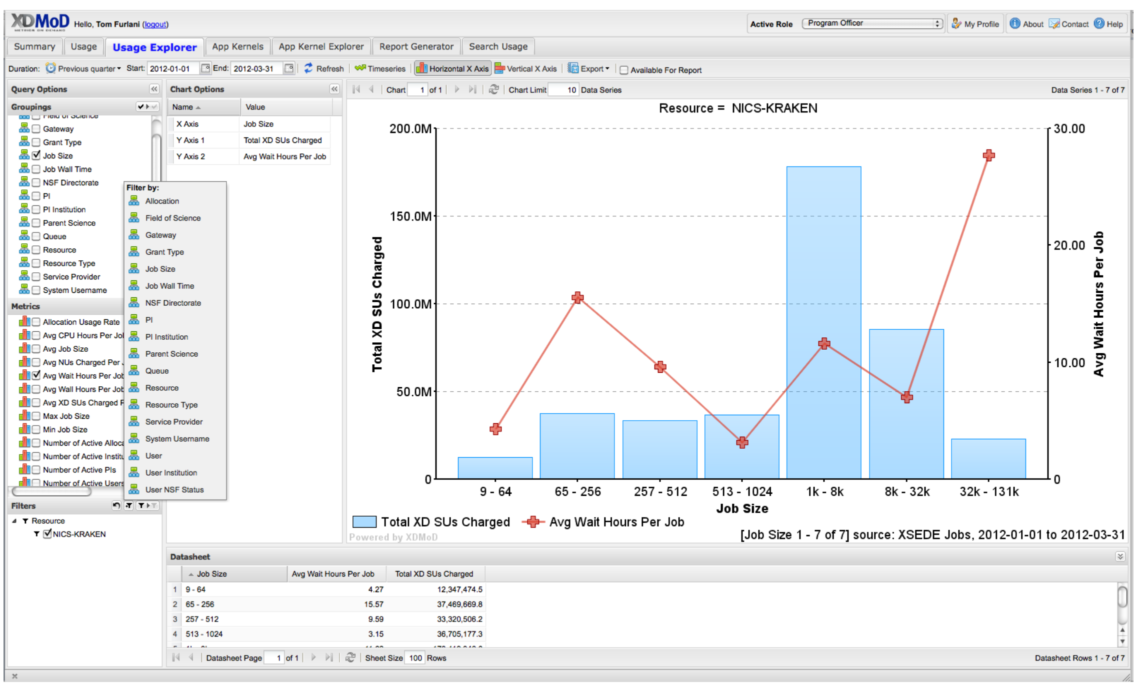

Figure 2. The XDMoD Usage Explorer. A plot of the total usage for the first quarter of 2012 is shown on the primary axis (left hand axis) and the average wait hours per job is shown on the secondary axis (right hand axis). The data have been filtered to show only NICS-Kraken usage.

Application Kernel Explorer:
The Applications Kernel Explorer has been similarly improved to make it more powerful and easier to use. Figure 3 shows a screen shot of the XDMoD App Kernel Explorer tab. The resource selected isedge.ccr (CCR's largest production cluster). The application kernel is HPCC (high performance computer challenge), specifically the FFT floating point performance on 8, 16, 32, 64 and 128 cores. Plots such as this can be made by simply selecting a resource, an application kernel metric and a number of cores. This plot can be exported as a PNG or EPS file and the data can be exported as a CSV or XML file.

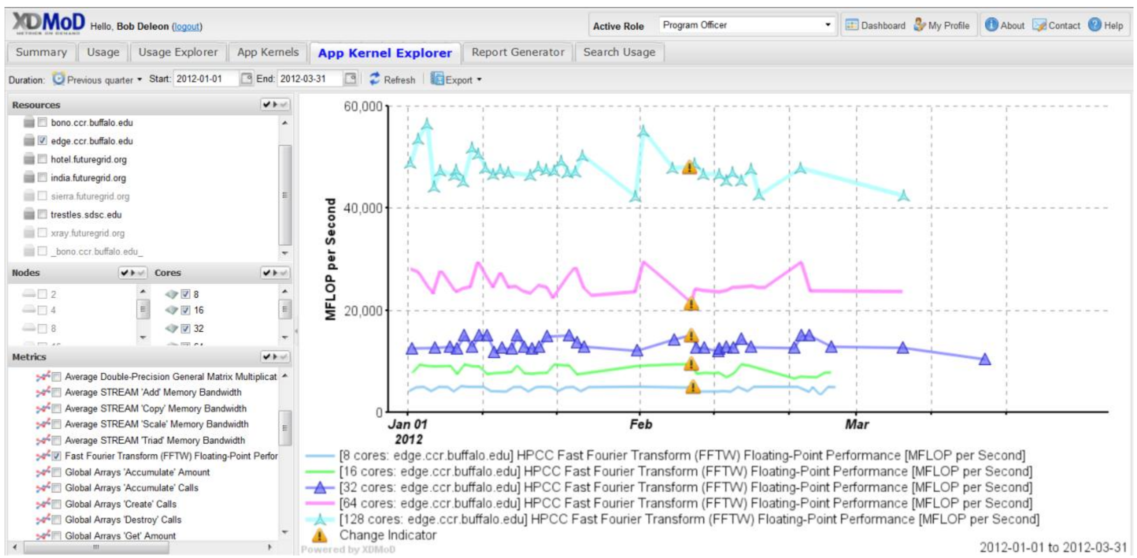

Figure 3. The Applications Kernel Usage Explorer. Shown is a plot of the 2012 first quarter data on CCR's large production cluster (ccr.edge.buffalo.edu) for the HPCC application kernel FFT floating point performance on 8,16, 32, 64 and 128 cores.

Expanded XDMoD Public Site:
The XDMoD public site has been expanded so that non-XSEDE users can get general information about the XSEDE usage and view the capabilities of XDMoD itself. Figure 4. shows a screen shot of the new public view. The example chosen for Figure 4 is a plot of the total XSEDE usage for the first quarter of 2012. As shown on the chart tree on the left hand side in Figure 4, a large variety of metrics are available even to the public. In addition by clicking on the "Scope" selector box (upper left hand corner) data for individual service providers can be selected as shown in Figure 5 (note that TACC is selected in the Scope selector box).

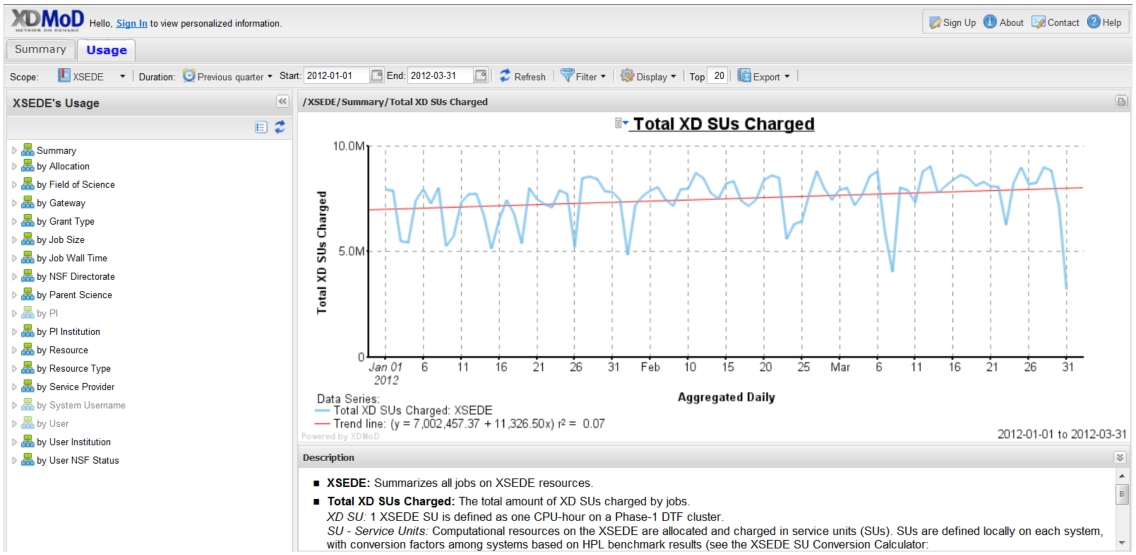

Figure 4. Screen shot from the Public view of XDMoD. Total XSEDE usage for the first quarter of 2012 is shown in the plot.

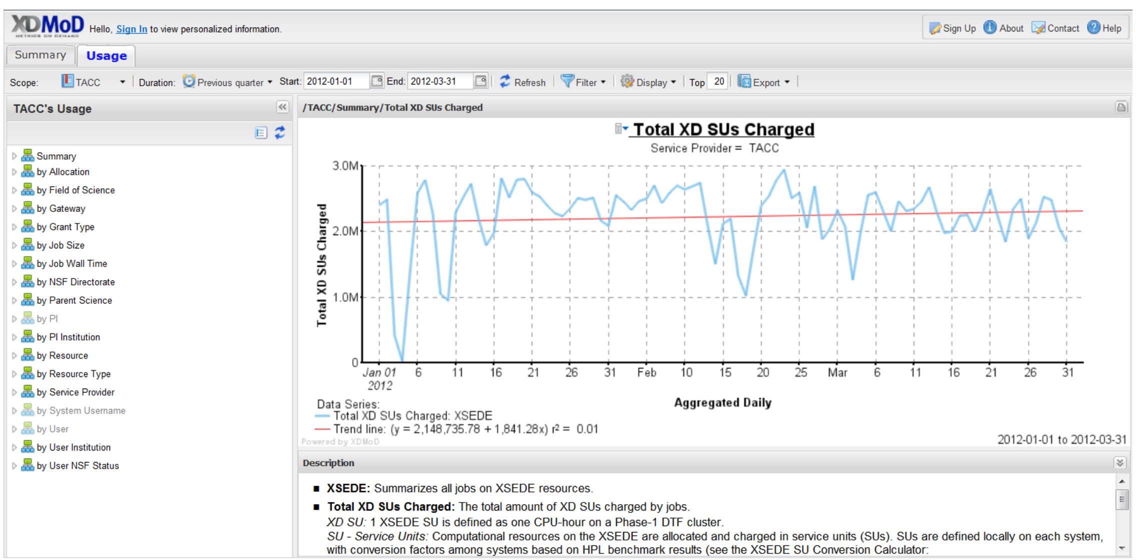

Figure 5. Screen shot from the Public view of XDMoD. Total TACC usage for the first quarter of 2012 is shown in the plot.

Implementation of the XDMoD role system:

It has always been intended that XDMoD would be organized by user role in order to display the information most pertinent to each XDMoD user. With the release of XDMoD 1.6 now all of the original roles are supported including: Program Officer, Center Director, Principal Investigator and User in addition to the Public role.

Application Kernel Development:

Last quarter, in preparation for wider deployment of the application kernels throughout XSEDE resources, an allocation request (Jones – PI) was made by the TAS team for Blacklight, Forge, Kraken, Lonestar4, Ranger and Trestles. The TAS Application Kernel audit process is designed to run from a user's perspective and accordingly an XSEDE user allocation is required. During the past quarter this allocation request was approved. Work is currently in progress on deploying the application kernels on these machines. We anticipate that the application kernels will be running on a number of additional XSEDE resources by the end of the current quarter.

Support of Planning for Advanced Cyberinfrastructure (ACI):

TAS has supported NSF in a strategy planning process. One of the most important aspects of the TAS program and XDMoD in particular is providing the capability of supplying historical usage data which can serve as the basis for planning future cyberinfrastructure. The TAS team worked closely with NSF to provide required data for the planning document:

A Data History of TeraGrid/XSEDE Usage: Defining a Strategy for Advanced CyberInfrastructure(ACI)

B. I. Schneider

Office of Cyberinfrastructure, National Science Foundation, Arlington, Virginia 22230

(Dated: March 11, 2012)

Incorporation of POPS data into XDMoD:

As part of the process of providing data for the ACI planning document described above, a great deal of allocation information was mined from the POPS data base, for example see Figure 6. This effort was lead by Indiana University (Gregor von Laszewski and Fugang Wang) who developed software to access a portion of the POPS data base. We are presently in the process of establishing the mechanisms to incorporate access to all of the POPS data within XDMoD. This will provide access to all of the allocation data through the XDMoD portal.

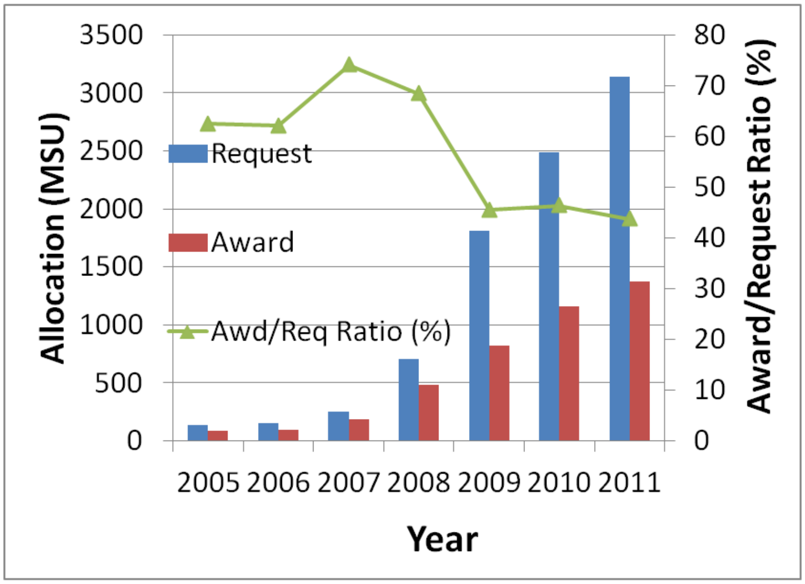

Figure 6. Historical allocation data mined from the POPS data base..

9.2 <u>Meetings, Events, Publications and Presentations</u>

Meetings:

Weekly meetings of the TAS Working Group are held every Wednesday afternoon. UB project members attend in person and the Indiana University subcontractors attend by web and audio conference.

9.3 TAS Appendices

Appendix 1. XDMoD Current Usage

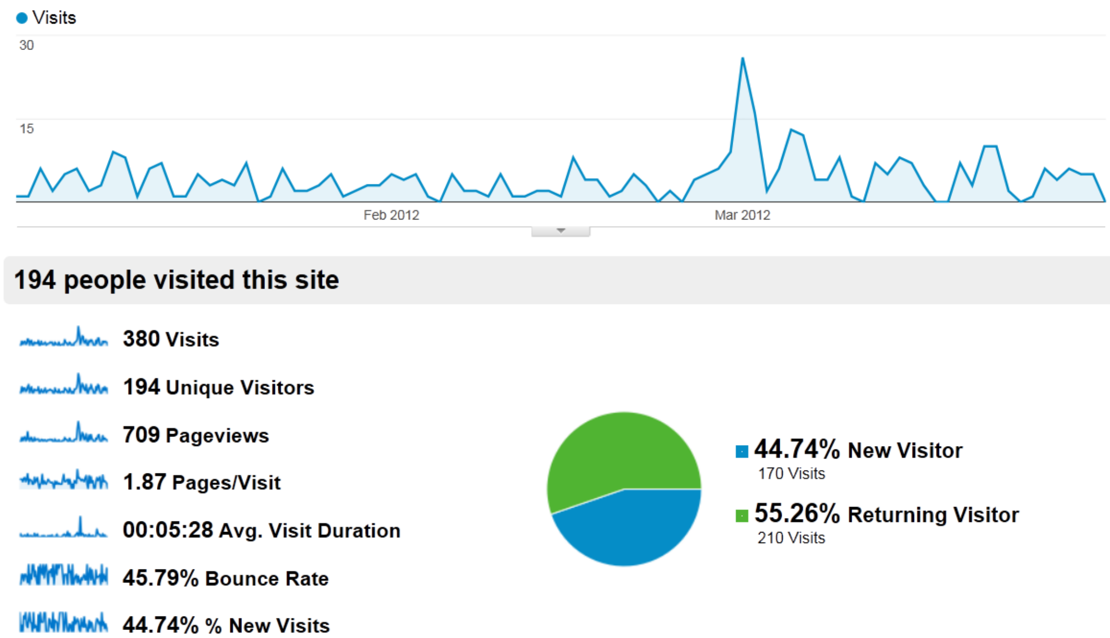

Figure A1. Google analytics overview of XDMoD usage for Jan 1, 2012 to Mar 31, 2012.

10 TAIS/Technology Insertion 1.8

10.1 Highlights

XTED version 0.7 was released providing support for managing user requests. Evaluations have been started on workflow management systems.

10.2 XD – TIS

10.2.1 Technology Identification and Tracking

Update

Development for iteration 0.7 of XTED was completed. Version 0.7 provides features supporting User Request feedback management. As this information is collected in XTED, it will be used to determine technologies selected for evaluation. Version 0.7 is currently undergoing quality assurance review and will be pushed to production upon approval.

Version 0.7 Use Cases

1. Create User Request - Allows an Authenticated User to request a particular capability. This is a very generic 'catch-all' for soliciting feedback from the community and covers things such as documenting user forum feedback.
2. Manage User Request - TIS Staff manages a User Request, ex. updating content.
3. View Requests - The primary listing of all 'Request' related records viewable by TIS Staff.

10.2.2 Technology Evaluation Laboratory

The Technology Evaluation Laboratory (TEL) is composed of a combination of reserved resources housed at TACC, the FutureGrid systems, and Service Provider (SP) systems. These systems represent the HPC computing environment found in XD and will facilitate the Technology Evaluation Process (TEP) conducted by the TIS team. The test system housed at TACC has been configured with 6- Dell 710 nodes and can be accessed via ssh at the public address: tis.futuregrid.tacc.utexas.edu.

In the last quarter dedicated TEL nodes were integrated into the FutureGrid and the operating system was upgraded to CentOS 6.2 to support the latest in virtualization technologies. The team has helped the FutureGrid group in providing base virtual machine images for testing purposes which can also be used by the TEP group. New work was completed evaluating OpenStack cloud services on Indiana's FutureGrid system, India. We will continue to work to get OpenStack available on other test systems including dedicated nodes TEL in the coming quarter. Documentation of the uses of this hardware is ongoing.

10.2.3 Technology Evaluation Process

Summary: We published the final report for the Globus Online evaluation and started 2 more evaluations, Pegasus Workflow Management System and Unicore for Workflow management. NICS identified a replacement for Victor, Michael Campfield, and he has started with work on the Pegasus evaluation.

The second evaluation, Globus Online was finalized and the report was published.

We found that Globus Online met the requirements and should be recommended for use. We will publish the final report next quarter.

The Fourth evaluation, Pegasus Workflow Management System, was started.

Peter Enstrom is leading this one, and we expect it will take longer as we develop requirements, use cases and test plans for Workflow Management systems

The Fifth evaluation, Unicore for Workflow Management, was started.

Robert Budden is leading this one. We expect to make use of prior work done on the Pegasus evaluation to speed this one along.

Gensis II, GFFS was not ready for reevaluation during this time.

Goals

Our goal for the next quarter is to finish the fourth and fifth evaluations and to start the sixth.

10.2.4 Technology Insertion Support

The YNT team continues to work with the TIS evaluation teams to develop procedures and instrumentation in the evaluation process to support documenting deployment object components including validation tests and usage examples. Additional useful metadata from the evaluation process, such as time, expertise and effort required to perform each evaluation are being collected during evaluations to inform prospective implementers regarding deployment effort and expertise required. The YNT team continues to review existing deployment models in XSEDE and other CI projects, for guidance in establishing appropriate YNT procedures.

To facilitate transition of recommended and supported technologies in XSEDE, the YNT team continues to observe current deployment testing activities of XSEDE Software Development and Integration (SD&I) processes associated with XSEDE Configuration Items for execution management services (EMS), global federated file system (GFFS), and Globus Online data movement. During this quarter, the focus of SD&I on these configuration items has been on evaluation by XSEDE Operations and Security groups prior to production deployment. The experience and information gathered will guide preparation of deployment materials and documentation in ways that can readily be utilized by XSEDE SDI Configuration Item teams, as well as XSEDE Operations integration test teams.

10.2.5 Usage, Satisfaction, and Impact Evaluation

Led by the TIS user advocate, the Usage, Satisfaction and Impact Evaluation team is charged with engaging the user community in the evaluation and testing process, by various methods including focus group teleconferences aimed at finding out what functional requirements are not being met, and with feeding these requirements into the TIS identification and tracking, evaluation and testing activities.

Pending widespread public announcement that XTED is available to developers who wish to register candidate technologies for TIS evaluation, we have explained the process to several interested groups including RENCI and University College London. We have also begun discussions with RENCI about leveraging their insights into the requirements of their user base, focusing on bioinformatics, digital humanities, and the needs of educators. We are now exploring the development of a requirements gathering collaboration with I-CHASS (focused on their projects in the humanities, arts and social sciences). We are also leveraging partner institutions' work in the NIH funded *Minority Access to Research Careers (MARC) Bioinformatics* program, to discover the technology needs that are specific to researchers whose labs are on minority school campuses.

To enter the requirements we gather, in a format that is useful for TIS evaluations, we are beta testing version 0.7 of XTED. We will refer salient requirements to the User Requirements Evaluation Process (UREP) run by the XSEDE Systems Engineering team. Requirements that are validated by the UREP will be preferentially used in the TIS evaluation process.

We continue to work with the YNT team to developed an evaluation priority metric to guide selection and to establish priorities among technology evaluation candidates (see 9.2.4 above), stressing the need for the user requirements to be given appropriate weight in this metric.

11 ExTENCI

XSEDE and Open Science Grid have been working together to support science that can make use of both cyberinfrastructure platforms.

11.1 Overview

The Extending Science Through Enhanced National Cyberinfrastructure (ExTENCI) Project is a joint Open Science Grid (OSG) and TeraGrid project, funded by OCI. The PIs are Paul Avery (U. Florida), Ralph Roskies (PSC), and Daniel S. Katz (U. Chicago).

The goal of ExTENCI is to develop and provide production quality enhancements to the national cyberinfrastructure that will enable specific science applications to more easily use both OSG and XSEDE/TeraGrid or broaden access to a capability to both XSEDE/TeraGrid and OSG users.

ExTENCI has four primary areas of work, each of which is discussed below.

11.1.1 ExTENCI – Distributed File System (Lustre-WAN)

Activities on the Distributed File System are the following:

- Incorporated a new release of Lustre, 2.1.5.4, into the VM packages we distribute
- All Lustre clients were upgraded from 2.0.6.2 to 2.1.5.4 in the hopes of increasing stability.
- Conducted Lustre clients and server version interoperability tests
- The University of Southern Florida, a Tier 3 site, installed the Lustre client. The USF client was added to the PSC monitoring tool, wanpage.psc.edu.
- FNAL: deployed an Object Storage Server (OSS). This is the first distributed OSS outside the UF site. The total storage for the ExTENCI Lustre WAN is now 60TB. Benchmarks were run to confirm the performance to the new storage would be similar to performance of a local file system.
- Performed multiple client tests of the ATLAS and LQCD application using the Lustre WAN. Benchmarks results were documented on the project wiki.
- PKINIT configured and testing started on PSC VM test servers
- Packaged a VirtualBox VM Lustre image and tested it at the University of Florida
- Began investigating the PAM/LDAP federated authentication
- Presentation about ExTENCI Wide Area Lustre approved for the April Lustre User Group Meeting
- A paper for the XSEDE 12 meeting was written and submitted.

11.1.2 ExTENCI – Virtual Machines

The Clemson team has been working on the use of the OpenFlow protocol in its cloud computing resources (OneCloud). The aim is to provide advanced networking services for users, in order to configure VM instances and make them appear as seamlessly integrated as possible with the user's own resources. Initial work has been on "Elastic IP" and "Security Groups", services named by Amazon for the leasing of public IP address, its dynamic/late binding association with running instances and dynamic firewall. These two functionalities are now in production in Clemson's OneCloud. A user is able to start VMs, lease Clemson public IPs and associate them on the fly with the instances. Documentation and screencasts are available (http://sites.google.com/site/cuonecloud/). This work has been tested by Massimo Canonico (FutureGrid user) and Owen Synge from DESY (Germany). Mike Wilde's group working on the use of SWIFT in SCEC is also testing this resource, and STAR is also testing these two new functionalities. Software developed has been contributed back into the OpenNebula 3.4, released early April. Security Groups are being included in the code base and release is planned for OpenNebula 3.5. This work is also described in a blog entry on the OpenNebula web site (http://blog.opennebula.org/?p=2695). Clemson presented this work with Elastic IP and Security Groups at the First GENI Research and Educational Experiment Workshop in Los Angeles, CA in conjunction with GEC13. Another paper has been accepted at VTDC workshop during HPDC in June 2012. Further work in the use of Openflow in the cloud will result in more advanced network services, such as dynamic load balancing, virtual private cloud and efficient data transfer between sites.

As one of the Year 2 deliverables, the Purdue team is developing Cloud Dashboard, a tool for easier access to Wispy through a graphical user interface. Using the Dashboard, users can start, stop and interact with the VMs running on Wispy in a browser. An X11 version has been completed, and the team is adapting it to an online version that will work in a web browser. The Purdue team automated the authentication process on Wispy to make it easier for users to access. This authentication process has been integrated into the Cloud Dashboard. The team also successfully integrated the U Wisconsin's VM authoring and management tools (the early February version) into the Dashboard and tested it with Wispy. During a visit by the ExTENCI project manager on Feb. 14, the Purdue team demonstrated the prototype Cloud Dashboard using the integrated Wisconsin's VM tool for staging VM images into Wispy. The team gave an overview of several nanoHUB applications that run on Wispy and the comparison with running on other resources (comparable performance for these applications). Although the team demonstrated and shared some initial usage and charts of the CMS VMs running at Purdue, a more general solution for running the CMS VMs will be the use of the GlideinWMS factory (this work is on hold until the Fermilab GlideinWMS team is available to proceed). The Purdue team also worked with Clemson to get BOINC CernVMs (LHC work) running successfully on Wispy.

Purdue's Wispy system was moved to a new data center in February as the previous data center housing Wispy was closed due to a power and cooling savings plan for IT on the West Lafayette campus of Purdue University. In the near future, Wispy will undergo a hardware upgrade to servers based on Intel's Sandy Bridge series of processors.

The Wisconsin team completed support for staging and de-staging on Purdue's Wispy and Clemson's OneCloud services in its VM management tool, and started support for control operations on these services. It also incorporated useful feedback from Purdue's experience using the tool. The team had FutureGrid's Eucalyptus, Nimbus, and OpenStack services approved as the final component of the corresponding work item. Wisconsin also helped to define FutureGrid's requirements for a formal ExTENCI project definition. The OpenStack integration has progressed to being approximately half-completed, and preliminary work was started on the other two services. Additionally, the team has completed a draft of the cloud information services report as defined in WBS 5.1, and is in process of collecting feedback on the report from the VM team. The team has also begun to develop a survey of stakeholders to provide more focused requirements and motivating scenarios to guide further development of the cloud information service.

The ExTENCI VM team has recently received a request from the STAR project which needs massive resource for 5k job slots for 4 months for a near real-time data production (Run 12 data). These jobs will also need large memory and large local storage per job. Purdue is coordinating the discussion with FutureGrid and Clemson about resource availability. A new FutureGrid project for ExTENCI has been created to help with the investigation, and it is available for any ExTENCI activity to use.

11.1.3 <u>ExTENCI – Workflow & Client Tools</u>

Hardware and software tasks during this period:

SCEC:

- *Task I1.4* (Port the currently sequential Broadband code to Swift in order to parallelize it and to enable it to run efficiently on both TeraGrid and Open Science Grid. Run the "hanging wall" benchmark in one day. Coordinate with SCEC to continue working with latest version of Broadband code), to start 8/1/2011, to finish 3/31/12, 60% complete.

  We have installed the latest release 11.2.2 of the SCEC Broadband package on Ranger, along with its pre-requisite supporting packages and libraries. We addressed issues in unit testing in consultation with SCEC. The acceptance tests were completed on Ranger. A number of tests show failures due to divergence from the defined tolerance values. The tests were then successfully run after applying broadband patches supplied by SCEC. The stages of the Hanging Wall challenge problem, and its inputs, outputs and degrees of parallelism have been defined with SCEC. Hanging wall simulations were defined in Swift and successfully tested on Ranger. 16 simulations have been successfully run in parallel using Swift.

- *Task I1.5* (Build version of parallel Broadband code that used VMs to run on some TG and OSG systems and potentially Amazon EC2), to start 4/1/2012, to finish 7/31/12, 25% complete.

  Broadband has been installed within a VM on FutureGrid. The Broadband unit tests are now passing, and work is proceeding on the acceptance tests. Swift workflows have been tested on FutureGrid using resource management mechanisms provided by Nimbus.

Additional applications:

- *Task I2.3* Run an application with some workflow segments on XSEDE and other segments on OSG. Measure workflow behavior and efficiency. Started 2/1/2011, to finish 3/31/2012.

  70% complete. Candidate workflows for this milestone are: (a) collaboration graph analysis (PI: Andrey Rheztsy) or (b) LAMP protein modeling (PI: Greg Voth). Both are NSF-funded. (Revised 1/6/2012 to select apps that have live science needs and multi-grid requirements).

  Application (a), the SciColSim collaboration graph analysis workflow has been ported to Swift, tested on a Cray system (including science runs), and is now being tested on OSG and XSEDE (Ranger). The main application code involved has been converted to run serially or on OpenMP.

  For application (b) we have switched to focus on an application for the Reichman lab at Columbia. In addition, for the NSF project "RDCEP" a Swift workflow with 120K tasks using the DSSAT land use application is being tested in production on the UChicago Cray system and being moved to OSG.

11.1.4 <u>ExTENCI – Job Submission Paradigms</u>

Evaluation / implementation of the P* model on OSG. Advances in BigJob deployment and documentation on XSEDE and OSG. Integrated, interoperable runs on XSEDE and OSG using

BFAST (Next-Generation Sequence code). Significant sized runs on XSEDE and OSG but with many data movement bottlenecks.

Documentation and hardening of the SAGA Condor Adaptor to make it work with OSG glideinWMS. Investigations into how data-intensive applications can be supported programmatically on OSG (with operating data-set of 5GB per task), as well as on OSG and XSEDE simultaneously.

Integration with Cactus Gateway and documentation capabilities provided to the Gateway are underway.

12 XD Service Provider Reports

2012 Q1: January 1, 2012, through Mach 31, 2012

XD Service Provider Forum Chair

Carol Song	Purdue University (PU)

XD Service Provider Principal Investigators

Sean Ahern	U Tennessee – National Institute for Computational Science (UT-NICS)
Jay Boisseau	Texas Advanced Computing Center (TACC)
Geoffrey Fox	Indiana University (IU)
Kelly Gaither	Texas Advanced Computing Center (TACC)
Michael Levine	Pittsburgh Supercomputing Center (PSC)
Honggao Liu	Louisiana Optical Network Initiative/Louisiana State University (LONI/LSU)
Miron Livny	University of Wisconsin, Madison (Open Science Grid)
Richard Loft	National Center for Atmospheric Research (NCAR)
Richard Moore	San Diego Supercomputer Center (SDSC)
Michael Norman	San Diego Supercomputing Center (SDSC)
Carol Song	Purdue University (PU)
Craig Stewart	Indiana University (IU)
John Towns	National Center for Supercomputing Applications (NCSA)
Jeffery Vetter	Georgia Tech
Thomas Zacharia	U Tennessee – National Institute for Computational Science (UT-NICS)

12.1 Overview

The Service Provider Forum (SPF) has finalized the Service Provider Forum Charter document and voted unanimously to ratify the charter. The document is now available publicly at the XSEDE web site (https://www.xsede.org/web/guest/project-documents). The SPF Chair is working to put this and other SPF public documents, such as the membership application form, on the SPF's public web page (www.xsede.org/web/sp-forum).

The Service Provider Forum (SPF) has formally organized and elected its officials according to the SPF Charter. The election results are as follows:

Chair: Carol Song, Purdue University

Vice Chair: Dave Hancock, Indiana University

Representatives to the XSEDE Advisory Board:

Dick Glassbrook, Georgia Tech

Miron Livny, University of Wisconsin, Madison

During this quarter, the SPF discussed and provided input to a number of XSEDE documents, including the SP letter of intent, the XSEDE award letter to NSF Program Directors, the baseline software and security documents, as well as candidate nominations for the XSEDE Advisory Board (XAB) and the User Advisory Committee (UAC).

13 FutureGrid Service Provider Quarterly

13.1 Executive Summary

- A new GPU Cluster (*delta*) with 16 nodes is operational for early users. Each GPU has 192GB RAM.

- Sixteen (16) new project requests were approved this quarter. See *1.9.3 SP-specific Metrics* for project statistics

- A new portal front page design was implemented, featuring FutureGrid Services, Support, News, Projects, and links to social media. Also highlighted are middleware experimentation, education and training, data-intensive applications, and more.
- A new SLURM-based tool for creating virtual clusters was developed. This tool distinguishes itself from other tools via a queuing system based on SLURM and configures the cluster in such a way that jobs can be submitted to the SLURM management node. This tool is used to conduct benchmarks for dynamic provisioning.

13.1.1 Resource Description

FG Hardware Systems

Name	System type	# Nodes	# CPUs	# Cores	TFLOPS	Total RAM (GB)	Secondary Storage (TB)	Site
india	IBM iDataPlex	128	256	1024	11	3072	335	IU
hotel	IBM iDataPlex	84	168	672	7	2016	120	UC
sierra	IBM iDataPlex	84	168	672	7	2688	96	SDSC
foxtrot	IBM iDataPlex	32	64	256	3	768	0	UF
alamo	Dell PowerEdge	96	192	768	8	1152	30	TACC
xray	Cray XT5m	1	168	672	6	1344	335	IU
bravo	HP Proliant	16	32	128	1.7	3072	192	IU
delta	SuperMicro	16	32	192	TBD	3072	144	IU
Total		457	1080	4384	43.7	17184	1252	

FG Storage Systems

System Type	Capacity (TB)	File System	Site
DDN 9550 (Data Capacitor)	339 shared with IU + 16 TB dedicated	Lustre	IU
DDN 6620	120	GPFS	UC
SunFire x4170	96	ZFS	SDSC
Dell MD3000	30	NFS	TACC
IBM dx360 M3	24	NFS	UF

13.2 Science Highlights

Sergey Blagodurov
School of Computing Science
Simon Fraser University
Burnaby BC, CA

Optimizing Shared Resource Contention in HPC Clusters

Abstract

Contention for shared resources in HPC clusters occurs when jobs are concurrently executing on the same multicore node (there is a contention for allocated CPU time, shared caches, memory bus, memory controllers, etc.) and when jobs are concurrently accessing cluster interconnects as their processes communicate data between each other. The cluster network also has to be used by the cluster scheduler in a virtualized environment to migrate job virtual machines across the nodes. We argue that contention for cluster shared resources incurs severe degradation to workload performance and stability and hence must be addressed. We also found that the state-of-the-art HPC cluster schedulers are not contention-aware. The goal of this work is the design, implementation and evaluation of a scheduling framework that optimizes shared resource contention in a virtualized HPC cluster environment.

Intellectual Merit

The proposed research demonstrates how the shared resource contention in HPC clusters can be addressed via contention-aware scheduling of HPC jobs. The proposed framework is comprised of a novel scheduling algorithm and a set of Open Source software that includes the original code and patches to the widely-used tools in the field. The solution (a) allows an online monitoring of the cluster workload and (b) provides a way to make and enforce contention-aware scheduling decisions on practice.

Broader Impacts

This research suggests a way to upgrade the HPC infrastructure used by U.S. academic institutions, industry and government. The goal of the upgrade is a better performance for general cluster workload.

Results

Link to problem outline, the framework, and preliminary results:
http://www.sfu.ca/~sba70/files/ClusterScheduling.pdf

Dave Hancock
Research Technologies
University Information Technology Services
Indiana University

Comparison of Network Emulation Methods

Abstract

Dedicated network impairment devices are an accepted method to emulate network latency and loss, but are expensive and not available to most users. This project will compare the performance of a dedicated network impairment device with that of performing network emulation within the Linux kernel to simulate the parameters when utilizing TCP.

Intellectual Merit

Results

The experiment consisted of host-to-host Iperf TCP performance while increasing parallel streams and inducing RTT latency utilizing FutureGrid's Spirent XGEM Network Impairments device. The hosts were two IBM x3650's with Broadcom NetExtreme II BCM57710 NIC's. RedHat release 5.5 Linux distribution was installed on each host, keeping stock kernel tuning in place. An Ethernet (eth0) interface on each host was connected back-to-back while the second Ethernet (eth1) passed through the Spirent XGEM and Nexus 7018 using an untagged VLAN, as illustrated in this diagram:

The direct host-to-host link saw an average delay of .040 ms while the path through the XGEM (.004 ms) and Nexus (.026 ms) was .080 ms.

Dual five minute unidirectional TCP Iperf tests were conducted, one each across the direct and switched path.

This testing will be used as validation of WAN impairment of TCP and LNET over a single 100 gigabit Ethernet circuit being conducted in Germany through a partnership with the Technische Universität Dresden. **Broader Impacts** Results will be widely disseminated through a joint paper with Technische Universität Dresden and presented to Internet2 members in April.	Tests were initiated independently and occurred at approximately the same start time with a deviation of +/- 3 seconds initiation. Results were gathered for each direct (D) and switched (S) test. Follow-up tests were executed increasing the number of parallel streams Iperf (command line option -P) could transmit. The number of streams included single, sixteen, thirty-two, sixty-four and ninety-six. Delay was added via the Spirent at increments of default (.080 ms), 4.00 ms, 8.00 ms, 16.00 ms, 32.00 ms, 64.00 ms, 96.00 ms and 128.00 ms RTT. The matrix yielded forty data points. Additionally the experiments were repeated utilizing two different kernel tuning profiles, increasing the data points to 80 and 120. The data points and graph (only switched path) show that as delay increased overall TCP performance increased as the number of parallel threads were increased.

13.3 User-facing Activities

13.3.1 System Activities

- IU iDataPlex ("*india*"). System operational for Eucalyptus, OpenStack, and HPC users.
- IU Cray ("*xray*"). System operational for production HPC users.
- IU HP ("*bravo*"). System operational for production users. 50% of the nodes being used for testing with the Network Impairment Device. HDFS available on remaining nodes.
- IU GPU ("*delta*"). All 16 nodes of this new cluster are operational for early users and integrated into the *india* scheduler. Each GPU has 192GB RAM.
- SDSC iDataPlex ("*sierra*"). System operational for production Eucalyptus, Nimbus, and HPC users.
- UC iDataPlex ("*hotel*"). System operational for production Nimbus and HPC users. Deployment plan for Genesis II is in progress. Conducted a massive overhaul of logging and monitoring systems
- UF iDataPlex ("*foxtrot*"). System operational for production Nimbus users.
- TACC Dell ("*alamo*"). System operational for production Nimbus and HPC users. Five (5) nodes are provisioned for XSEDE TIS testing. A new Bright Computing cluster management system has been installed.

All system outages are posted at https://portal.futuregrid.org/outages_all.

13.3.2 Services Activities (specific services are underlined in each activity below)

Cloud Services and Support Software

Eucalyptus. The commercial version of Eucalyptus (Version 3.0) was installed on the gravel test cluster, and assistance from the Eucalyptus support team on network settings in on-going. Configuration details will be documented at https://wiki.futuregrid.org/index.php/Euca3.

Nimbus. Final release of Nimbus Infrastructure 2.9. The major additions in this release are support for Availability Zones, configurable EC2 multi-core instances, more robust support for

LANTorrent and new administration tools which allow administrators to easily control VMs running on their cloud. The administrators can also choose to give more information to the user, e.g., allow them to inspect on what physical machines their virtual machines are running. In addition, the release also includes bugfixes and additions to documentation. With the exception of Availability Zones, all new features and enhancements were driven by FutureGrid.

OpenStack. While the Cactus version of OpenStack is running in "production mode" on *india*, testing on the next version, Essex, has begun.

SLURM. A new tool to easily create virtual clusters was developed. This tool distinguishes itself from other tools via a queuing system based on SLURM and configures the cluster in such a way that jobs can be submitted to the SLURM management node. The cluster also includes OpenMPI. The tool has been uploaded to PyPI and is easily installed with pip. This tool is used to conduct benchmarks for dynamic provisioning. The tool has been developed with S3-compliant tools and APIs. Although running only on OpenStack today, the objective is to be able to run the tool on Eucalyptus, Nimbus, and OpenNebula.

Inca. A new Inca test to verify ViNe capabilities between *sierra* and *foxtrot* was deployed. The test is based on the FutureGrid tutorial, "Connecting VMs in private networks via ViNe overlay." A VM with a public IP address is initiated on *sierra* and a VM with a private IP address is initiated on *foxtrot*. The test verifies that the *sierra* VM can connect to the *foxtrot* VM, and then shuts them both down. This test currently runs once per day and the results are available on the Inca page in the portal. A new Inca test was also deployed to verify that email is being sent from the ticketing system. This ensures that our communication to the end user who originally generated a ticket is properly being c communicated with. The test runs once a day and is configured to send email to the FutureGrid system administrators if a failure is detected.

ViNe. ViNe development activities focused n testing ViNe Version 2 (ViNe with management enhancements). While basic functionality of ViNe2 was verified, it was identified that only simple scenarios of overlay network deployment and management is supported. These include: overlay topology in which most of the load concentrates in a single node; difficulty in supporting the currently deployed ViNe on *sierra* and *foxtrot*; and limited auto configuration capability. It was decided that the ViNe management code will be reqorked in order to enable the full capability of ViNe's infrastructure. Target delivery by May 1st. Additional activities focused on improving the ViNe Central Server (VCS), specifically the procedures to generate default configuration values and parameters. Basic front-end interfaces have been developed to allow FutureGrid users to directly interact with VCS to check the status of ViNe overlays and to issue ViNe requests (e.g. creation of a new overlay).

Experiment Management

Experiment Harness. Continued modification of the TeraGrid/XSEDE glue2 software for use on FutureGrid. Adding support for providing respurce information in JSON.

Pegasus. Pegasus is now visible on the FutureGrid portal, complete with documentation in the User Manual (https://portal.futuregrid.org/manual/pegasus) and an updated tutorial (http://pegasus.isi.edu/futuregrid/tutorials).

Image Management. Improvements made to the image management tools to support different kernels. Users may now select the kernel they want for their images.

HPC Services

UNICORE and Genesis. Delivered ~23,000 CPU hours on *india* and ~16,000 hours on *sierra* via the cross campus grid (XCG) to an Economics application. This is the first in a series of runs where the expectation is the delivery of ~400,000 hours to the user on XCG. This user will be able to use the same mechanism to use his XSEDE allocation on TACC's Ranger system, once the XSEDE EMS goes operational

Performance

Vampir. Completed an OpenMPI 1.5.4 build, which includes VampirTrace 5.8.4 and is configured with an installation of PAPI 2.4.0 as well, all to be deployed on the *bravo* cluster.

PAPI. The PAPI VMware component was completed, including network performance testing and preparation for bringing these capabilities to FutureGrid VMs.

FutureGrid Portal

A new portal front page design was implemented, featuring FutureGrid Services, Support, News, Projects, and links to social media. Also highlighted are middleware experimentation, education and training, data-intensive applications, and more.

Additional redesign efforts are on-going in the areas of:
- Support
- Getting Started
- Hardware, Networking Status
- Projects, including a "featured" project and a new content-rich approach to each project

The IU KnowledgeBase was updated to properly synchronize KB content with the portal and is in final testing. Once released, this will allow the use of default Drupal content searching (and other services, such as comments), for value add to the KB content, as well as to collect and direct feedback on KB content.

13.4 Security

No security issues occurred during this period.

13.5 Education, Outreach, and Training Activities

Agreement with XSEDE and University of Michigan (Virtual Summer) was reached to offer a Summer School on Science Clouds July 30 thru August 3, 2012.

The FutureGrid Project Challenge was announced. The FPC is open to both students and non-students and provides an opportunity to showcase their skills and knowledge by utilizing FutureGrid computing resources in the execution of applications. Eight awards are envisioned, four for students and four for non-students, to be based on criteria that include the following:
- Interoperability
- Scalability
- Contribution to Education

- Research (innovation, quality of papers, new software, algorithms, insightful performance measurements, etc.)

The award will be based on FutureGrid work and its analysis or on ideas and software developed with major use of FutureGrid. First place award for students is a trip to SC12 (up to a $1,000), where a demonstration will be provided. All other awards are $200 cash.

The first formal FutureGrid User Survey was finalized and sent out to all FutureGrid portal account holders. Preliminary results are being reviewed.

Dialog was initiated with XSEDE regarding the establishment of User Forums on the XSEDE portal for Map/Reduce and Science Cloud Applications groups.

A concentrated review effort on all FutureGrid tutorial materials was completed. This has now become a "continuous improvement" process, as there are times when a tutorial on, say, Eucalyptus, will work fine on one FG resource, but not on another.

Events this quarter:

Type	Title	Location	Date(s)	Hours	Number of Participants	Number of Under-represented people	Method
Indiana University							
Presentation	Advances in Clouds and their Application to Data Intensive Problems	University of Southern California, Los Angeles, CA	02/24/2012	1.0	50	Unknown	Synchronous
Type	**Title**	**Location**	**Date(s)**	**Hours**	**Number of Participants**	**Number of Under-represented people**	**Method**
University of Chicago							
Presentation	Nimbus Platform: Infrastructure Outsourcing for Science	Trends in High-Performance Distributed Computing, Vrije Universiteit, Amsterdam, NL	03/14/2012	0.5	20-50	Unknown	Synchronous
Presentation	Building an Outsourcing Ecosystem for Science	Computer Science Seminar, Northern Illinois University, Dekalb, IL	03/22/2012	0.75	20-50	Unknown	Synchronous

University of Texas at Austin							
Presentation	FutureGrid: An Experimental High-Performance Grid Test-bed	TACC, Austin, TX	03/21/2012	1.0	~20	Unknown	Synchronous

13.6 SP Collaborations

- The Eighth Open Cirrus Summit and Federated Clouds Workshop, September 21, 2012
  Focus on test beds, including Future Grid, Open Cloud Consortium, and others. The cloud-computing test beds considered for this workshop will present and share research into the design, provisioning, and management of services at a global, multi-datacenter global scale. Through this venue, a collaborative community around test beds will be fostered, providing ways to share tools, lessons and best practices, and ways to benchmark and compare alternative approaches to service management at datacenter scale.
- European Middleware Initiative (EMI)
  The European Middleware Initiative (EMI) is a joint project of the middleware consortia gLite, ARC, UNICORE and dCache. Funded by the European Commission, the major goal of this project is to harmonize a set of EMI products out of the different middleware components by providing streamlined EMI releases. The aim of this Future Grid project is to set up a permanent test bed of EMI releases for exploration by US partners in order to disseminate the activities happening in Europe. Training material will be provided so that the EMI products can be tested and evaluated by potential US stakeholders using FutureGrid for a hands-on experience with European middleware.

13.7 SP-Specific Activities

See 1.3.2 Services Activities.

13.8 Publications

Javier Diaz, Gregor von Laszewski, Fugang Wang, Geoffrey C. Fox, “Abstract Image Management and Universal Image Registration for Cloud and HPC Infrastructures,” Technical Report, 3/2012.

Kate Keahey, Frederic Desprez, Geoffrey Fox, Emmanuel Jeannot, Kate Keahey, Michael Kozuch, David Margery, Pierre Neyron, Lucas Nussbaum, Christian Perez, Olivier Richard, Warren Smith, Jens Vöckler, Gregor von Laszewski Supporting Experimental Computer Science Published as ANL MCS Technical Memo 326, Argonne, Argonne National Laboratory, 03/2012 summarizing Support for Experimental Computer Science Workshop at SC11 Seattle, WA, 11/18/2011.

Zhenhua Guo, Geoffrey Fox, Mo Zhou, "Investigation of Data Locality and Fairness in MapReduce," Technical Report, 3/4/2012.

Yang Ruan, Zhenhua Guo, Yuduo Zhou, Judy Qiu, Geoffrey Fox, "HyMR: a Hybrid MapReduce Workflow System," Technical Report, 2/28/2012.

Gregor von Laszewski, "MapReduce for Scientists using FutureGrid," Technical Report, 2/28/2012.

Zhenhua Guo, Geoffrey Fox, and Mo Zhou, "Improving Resource Utilization in MapReduce," Technical Report, 2/18/2012.

Geoffrey C. Fox, Supun Kamburugamuve, Ryan Hartman, "Architecture and Measured Characteristics of a Cloud Based Internet of Things," Technical Report, 2/14/2012.

Thilina Gunarathne, Bingjing Zhang, Tak-Lon Wu, Judy Qiu, "Scalable Parallel Scientific Computing Using Twister4Azure," Technical Report, 2/11/2012.

Zhenhua Guo, Geoffrey Fox, "Improving MapReduce Performance in Heterogeneous Network Environments and Resource Utilization," Technical Report, 1/29/2012.

Geoffrey Fox and Dennis Gannon, "Cloud Programming Paradigms for Technical Computing Applications," Technical Report, 1/17/2012.

Yunhee Kang and Geoffrey Fox, "Performance Analysis of Twister based MapReduce Applications on Virtualization System in FutureGrid," Advanced Science Letters to appear 2012.

13.9 Metrics

13.9.1 Standard systems metrics

FutureGrid #Jobs run by (78) users (Jan-Mar 2012) - largest to smallest

Area *(page 1 of 2)*	User	# Jobs	Avg. Job Size (cpus)	Avg. Wait Time (h)	Wall Time (d)	Avg. Mem (MB)
University of Virginia - Genesis	xcguser	47159	1	1.1	1554.1	283.3
University of Buffalo	charngda	27938	37.2	0.7	3374.4	144.7
University of Virginia - UNICORE	unicore	1104	10.4	2.8	541.1	100.9
SDSC - INCA	inca	641	10	0.4	155.9	15.5
LSU - SAGA	pmantha	617	32	0.1	281.6	93.1
Indiana University	jychoi	341	15.6	1.7	615.2	2128.2
Indiana University	jdiaz	225	1	0	1.1	9.6
Indiana University (graduate class)	qismail	142	14.3	0	59.1	161.8
Indiana University (graduate class)	sbpatil	131	11.9	0.1	23.8	30.8
Indiana University (graduate class)	sahshah	121	6.2	0.2	16	25.9
Indiana University	ktanaka	113	8.5	0	5.3	24.8
Indiana University (graduate class)	aurankar	100	5.2	0.1	46.6	38.1
LSU (graduate class)	sivak2606	84	6.7	1.3	13.7	9
Indiana University (graduate class)	vuppala	79	7	0.2	14.1	49.1
LSU - SAGA	luckow	73	76.3	0.1	133.9	414.6

University of Utah (Project #149)	jasonkwan	70	7.8	0.2	80.1	13991.9
Indiana University (graduate class)	aalthnia	64	12.4	0.3	28.5	59.5
LSU - SAGA	marksant	63	16.9	0	7.6	141.9
University of Piemonte Orientale (class)	sguazt	62	1	0	47.1	1725.7
Indiana University	taklwu	60	12.8	0.1	32.5	62
Indiana University (graduate class)	hbharani	57	13.6	0	28.4	81.5
Indiana University (graduate class)	ralekar	49	16.7	0	30.6	39
Indiana University (graduate class)	cdesai	47	8.1	0	21.6	458.6
Indiana University (graduate class)	jshrotri	46	15.3	0	14	237
Indiana University (graduate class)	hsavla	44	13.7	0	20.8	87.2
Indiana University (graduate class)	bhasjais	43	7.6	0.2	6	22.6
LSU - SAGA	oweidner	41	7.7	0	1.9	2247.1
Indiana University (graduate class)	amkazi	36	11.7	0	4.3	22.6
LSU - SAGA	ssarip1	34	38.4	0	7.4	20.7
Indiana University (graduate class)	vdkhadke	34	13.1	0.1	19.6	33.7
Indiana University (graduate class)	granade	27	8.7	0	2.5	23.1
Indiana University (graduate class)	adherang	25	12.2	0	11.6	23.1
Indiana University (graduate class)	retemple	25	12.8	0	4.6	34.9
Indiana University (graduate class)	mmasrani	23	6.6	0	1.3	107.8
Indiana University (graduate class)	zhangqua	23	15.3	0.1	1.9	189
Indiana University (graduate class)	pvishwak	22	14.5	0.1	14.6	410.5
Indiana University (graduate class)	pransund	17	8.6	0.4	4.8	21.2
Indiana University (graduate class)	shenhan	16	3.2	0	0.6	20.7
Indiana University (graduate class)	venubang	16	9.4	0.5	10.8	22.2
Indiana University (graduate class)	vs3	14	11.7	0.4	4.2	17.1
Indiana University (graduate class)	yichfeng	12	11.8	0.7	21.3	303.3
Indiana University (graduate class)	banerjea	11	9.9	0	1.2	60.4
Simon Fraser University	blagodurov	10	7.6	0	0.1	6
Area *(page 2 of 2)*	User	# Jobs	Avg. Job Size (cpus)	Avg. Wait Time (h)	Wall Time (d)	Avg. Mem (MB)
Indiana University (graduate class)	minglu	10	11.5	0.9	10.1	38.1
Indiana University (graduate class)	bypatel	10	14.6	0	0.8	22.4
Indiana University (graduate class)	ilajogai	10	11.5	0	3.9	496.1
Indiana University (graduate class)	snagde	9	6	1.7	3.1	12.8
Indiana University (graduate class)	ritekavu	9	9.3	0.2	2.2	32.4
Indiana University	ajyounge	9	4.9	0	220.6	29148.8
Indiana University (graduate class)	abtodd	8	9.4	0	10.5	170.9
Indiana University (graduate class)	psharia	8	6.6	0	0.2	13.5
Indiana University (graduate class)	rjhoyle	8	12.5	0	0.8	18.2
Indiana University (graduate class)	shdubey	8	8.4	0	1	22.3
Indiana University	sharif	7	32.6	0.2	0	0.6
University of Utah (Project #149)	tkakule	7	8	0	2.8	370.5
Indiana University (graduate class)	dkmuchha	6	9.7	0	1.3	20.9
Indiana University (graduate class)	quzhou	5	7	1.4	0	0
Indiana University	skarwa	4	16	0	0.9	94
Indiana University	gpike	4	130.3	0	0.1	1.1
Indiana University	lihui	3	8	0	10.7	12.9
University of Florida (graduate class)	ssaripel	3	1	0.9	0	6.3
Indiana University (graduate class)	vjoy	3	11	0	1.2	22.1
Indiana University	yuduo	3	64	0	670.3	57.1
Indiana University (graduate class)	chaikhad	3	8.3	0	0.8	601.5

University of Utah (Project #149)	diarey	2	8	0	7.4	18640.1
Indiana University (graduate class)	nvkulkar	2	8.5	0	0.1	20.7
Indiana University (graduate class)	yuangao	2	16	0	3.5	14.3
Indiana University (graduate class)	sannandi	2	8.5	0	2.1	19.8
TACC	wsmith	2	8	0	0	0
Indiana University (graduate class)	bsaha	2	8.5	0	0.1	35
University of Florida (graduate class)	shailesh2088	2	32	0	0.1	180.9
Indiana University (graduate class)	pnewaska	1	1	0	0.1	18.2
Indiana University (graduate class)	gomoon	1	2	0	0	0
Indiana University	rpkonz	1	1	0	0	21.1
TACC	dgignac	1	1	0	0	0
Indiana University (graduate class)	ainausaf	1	1	0	0.1	17
University of Central Florida (Project #191)	laurenba	1	16	1.3	0.2	23.2

13.9.2 Standard systems metrics (continued)

Top 25 (Average) Memory Users - Largest to Smallest

Area	User	# Jobs	Avg. Job Size (cpus)	Avg. Wait Time (h)	Wall Time (d)	Avg. Mem (MB)
Indiana University (ScaleMP)	ajyounge	9	4.9	0	220.6	29148.8
University of Utah (Project #149)	diarey	2	8	0	7.4	18640.1
University of Utah (Project #149)	jasonkwan	70	7.8	0.2	80.1	13991.9
LSU - SAGA	oweidner	41	7.7	0	1.9	2247.1
Indiana University	jychoi	341	15.6	1.7	615.2	2128.2
University of Piemonte Orientale (class)	sguazt	62	1	0	47.1	1725.7
Indiana University (graduate class)	chaikhad	3	8.3	0	0.8	601.5
Indiana University (graduate class)	ilajogai	10	11.5	0	3.9	496.1
Indiana University (graduate class)	cdesai	47	8.1	0	21.6	458.6
LSU - SAGA	luckow	73	76.3	0.1	133.9	414.6
Indiana University (graduate class)	pvishwak	22	14.5	0.1	14.6	410.5
University of Utah (Project #149)	tkakule	7	8	0	2.8	370.5
Indiana University (graduate class)	yichfeng	12	11.8	0.7	21.3	303.3
University of Virginia - Genesis	xcguser	47159	1	1.1	1554.1	283.3
Indiana University (graduate class)	jshrotri	46	15.3	0	14	237
Indiana University (graduate class)	zhangqua	23	15.3	0.1	1.9	189
University of Florida (graduate class)	shailesh2088	2	32	0	0.1	180.9
Indiana University (graduate class)	abtodd	8	9.4	0	10.5	170.9
Indiana University (graduate class)	qismail	142	14.3	0	59.1	161.8
University of Buffalo	charngda	27938	37.2	0.7	3374.4	144.7
LSU - SAGA	marksant	63	16.9	0	7.6	141.9
Indiana University (graduate class)	mmasrani	23	6.6	0	1.3	107.8
University of Virginia - UNICORE	unicore	1104	10.4	2.8	541.1	100.9
Indiana University	skarwa	4	16	0	0.9	94
LSU - SAGA	pmantha	617	32	0.1	281.6	93.1

Top 15 (Average) Job Size Users - Largest to Smallest

Area	User	# Jobs	Avg. Job Size (cpus)	Avg. Wait Time (h)	Wall Time (d)	Avg. Mem (MB)
LSU - SAGA	luckow	73	76.3	0.1	133.9	414.6
Indiana University	yuduo	3	64	0	670.3	57.1
LSU - SAGA	ssarip1	34	38.4	0	7.4	20.7
University of Buffalo	charngda	27938	37.2	0.7	3374.4	144.7
Indiana University	sharif	7	32.6	0.2	0	0.6
University of Florida (graduate class)	shailesh2088	2	32	0	0.1	180.9
LSU - SAGA	pmantha	617	32	0.1	281.6	93.1
LSU - SAGA	marksant	63	16.9	0	7.6	141.9
Indiana University (graduate class)	ralekar	49	16.7	0	30.6	39
Indiana University	skarwa	4	16	0	0.9	94
Indiana University (graduate class)	yuangao	2	16	0	3.5	14.3
University of Central Florida (Project #191)	laurenba	1	16	1.3	0.2	23.2
Indiana University	jychoi	341	15.6	1.7	615.2	2128.2

13.9.3 Standard User Assistance Metrics

RT Ticket System

Created tickets in period, grouped by status:

Status	Tickets
deleted	0
new	4
open	21
rejected	0
resolved	184
stalled	1
Total	212

Resolved tickets (184) in period, grouped by queue (category):

a) 017 FutureGrid account requests
b) 037 Eucalyptus issues
c) 009 Nimbus issues
d) 005 Portal issues
e) 100 General issues
f) 007 *hotel* issues
g) 001 *alamo* issues
h) 005 User Support issues
i) 003 Systems issues

New/Open tickets (25) in period, grouped by queue (category):

a) 005 FutureGrid account requests
b) 001 Eucalyptus issues
c) 002 Nimbus issues
d) 003 Portal issues
e) 006 General issues
f) 007 User Support issues
    i) 003 external to users
    ii) 004 internal to project team
g) 001 Systems issues

13.9.4 SP-specific Metrics

Knowledge Base:

- Total number of documents available in the FutureGrid KB at the end of the 1st quarter = 105
- Number of documents added = 16
- Number of documents modified = 28
- Total number of documents retrieved = 20,671
- Total number of documents (retrieved minus bots) = 9,853

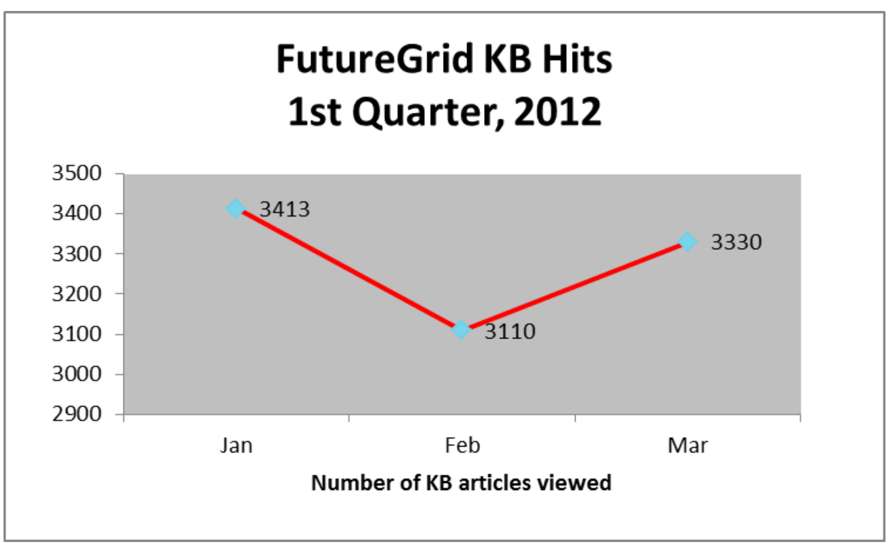

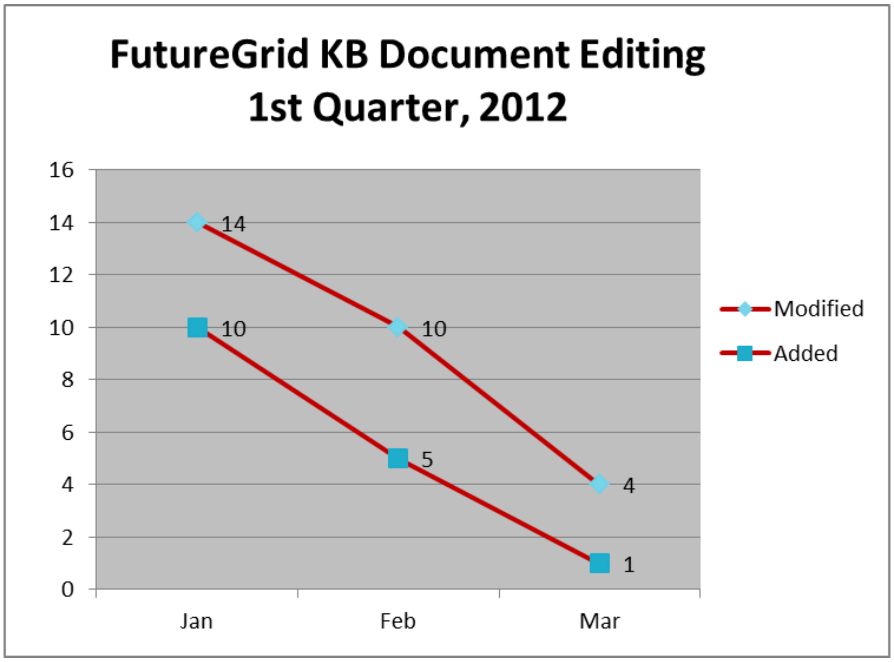
FutureGrid KB Document Editing
1st Quarter, 2012
16
14
12
10
8
6
4
2
0
Jan
Feb
Mar
14
10
10
5
4
1
Modified
Added

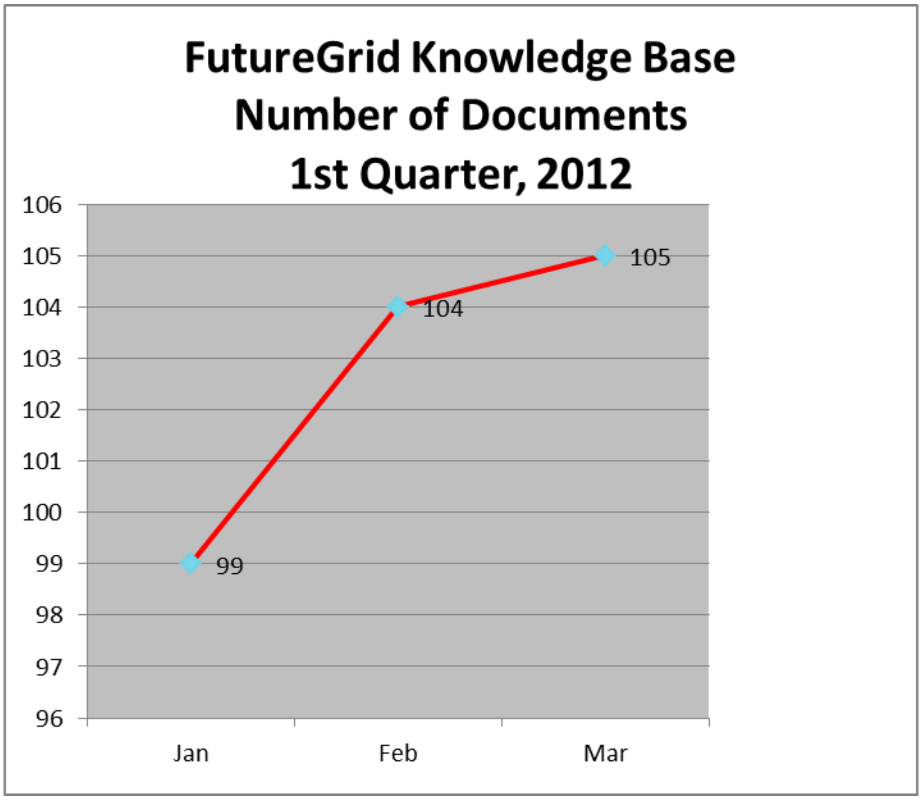

Projects:

- FutureGrid project count to date: 199
- Sixteen (16) new projects added this quarter
- Categorization of projects to date:

a) Project Status:

Active Projects: 183(92%)
Completed Projects: 10(5%)
Pending Projects: 2(1%)
Denied Projects: 4(2%)

b) Project Orientation:

Research Projects: 172(86.4%)
Education Projects: 24(12.1%)
Industry Projects: 2(1%)
Government Projects: 1(0.5%)

c) Project Primary Discipline:

Computer Science (401): 165(82.9%)
Biology (603): 9(4.5%)
Industrial/Manufacturing Engineering (108): 3(1.5%)
Not Assigned: 6(3%)
Genetics (610): 2(1%)
Physics (203): 1(0.5%)
Aerospace Engineering (101): 1(0.5%)
Statistics (403): 1(0.5%)
Engineering, n.e.c. (114): 2(1%)
Biosciences, n.e.c. (617): 1(0.5%)
Biophysics (605): 1(0.5%)
Economics (903): 1(0.5%)
Electrical and Related Engineering (106): 2(1%)
Pathology (613): 1(0.5%)
Civil and Related Engineering (105): 1(0.5%)
Biochemistry (602): 1(0.5%)
Atmospheric Sciences (301): 1(0.5%)

d) Project Service Request/Wishlist:

High Performance Computing Environment: 96(48.2%)
Eucalyptus: 101(50.8%)
Nimbus: 110(55.3%)
Hadoop: 70(35.2%)
Twister: 34(17.1%)
MapReduce: 66(33.2%)
OpenNebula: 30(15.1%)
Genesis II: 29(14.6%)
XSEDE Software Stack: 44(22.1%)
Unicore 6: 16(8%)
gLite: 17(8.5%)
OpenStack: 31(15.6%)
Globus: 13(6.5%)
Vampir: 8(4%)
Pegasus: 9(4.5%)
PAPI: 7(3.5%)

14 Indiana University - Service Provider Report

14.1 Executive Summary

During the current quarter, IU expanded and made more robust the services it offers via Quarry, and made significant strides forward in operational activities. This past quarter was a time of particular productivity as regards campus bridging activities. Four pilot projects were selected for the campus bridging pilots, and IU representatives worked at length with A&D, SD&I, and Operations teams to develop XSEDE use cases and quality attribute specifications. In addition, IU made significant progress on the first annual XSEDE survey and in preparations for the inaugural XSEDE12 conference.

14.1.1 Resource Description

Indiana University SP Systems

Quarry (Virtual Machines) - The XSEDE science gateway hosting environment provides hosting services for XSEDE researchers running science gateways, making high performance computing more accessible. Two VMs for XSEDE projects were added this quarter, bringing the total to 48. We have completed the first phase of hardware installation for the gateway hosting refresh. Software installation has also been completed and the initial migration preparations have begun. All new VM's starting Q4 are allocated in the new environment.

Data Capacitor - The Data Capacitor is a high speed/high bandwidth Lustre storage system for research computing that serves all IU campuses. It is comprised of 6 servers (2 MDS, 4 OSS). The total usable storage capacity is 427TB (1TB = 2**40 bytes).

HPSS - HPSS forms the core of IU's Scholarly Data Archive, a tape based archival data storage service. It is comprised of 23 IBM x3650 M2 servers, two TS3500 tape libraries, a total of 48 Jaguar4 tape drives, 500TB disk cache, and 15PB of tape. Data is replicated between the Indianapolis and Bloomington campus. While TeraGrid allocations will be provisioned until July 2012, currently HPSS is not accepting new XSEDE allocations.

Rockhopper - IU announced a new and novel approach to systems delivery and support during the current reporting period. Rockhopper is a collaborative effort between Penguin Computing, IU, the University of Virginia, the University of California Berkeley, and the University of Michigan to provide supercomputing "cluster on demand" services in a secure US facility. Researchers at US institutions of higher education and federally-funded research centers can purchase computing time from Penguin Computing, and receive access via high-speed national research networks operated by IU. It takes just minutes to go from holding a credit card in one's hands and filling out a web form to being computing on Rockhopper (the system itself is owned by Penguin; cycles on Rockhopper are purchased from Penguin). Rockhopper is a 4.4 TFLOPS system based on AMD processors.

14.2 Science Highlights

Realizing the Universe for the Dark Energy Survey (DES) Using XSEDE Support

The Dark Energy Survey (DES) is an upcoming international experiment that aims to constrain the properties of dark energy and dark matter in the universe using a deep, 5000-square degree survey of cosmic structure traced by galaxies. To support this science, the DES Simulation Working Group is generating expectations for galaxy yields in various cosmologies. Analysis of these simulated catalogs offers a quality assurance capability for cosmological and astrophysical analysis of upcoming DES telescope data. Our large, multi-staged computations are a natural fit for workflow control atop XSEDE resources.

We have completed the first of the Blind Cosmology Challenge (BCC) N-body computations, and are preparing to add galaxies and weak-lensing signals to the data. We are prepared to run the next (and all future) N-body simulations in the Apache Airavata workflow environment, and plan to fold in the galaxy and weak-lensing post-processing stages in the next semester. Using a workflow environment has reduced sources of error and inefficiency associated with an individual running every step of the process. Development of do-while construct in workflow helps us overcome queue wall-time limits of XSEDE resources (TACC Ranger).

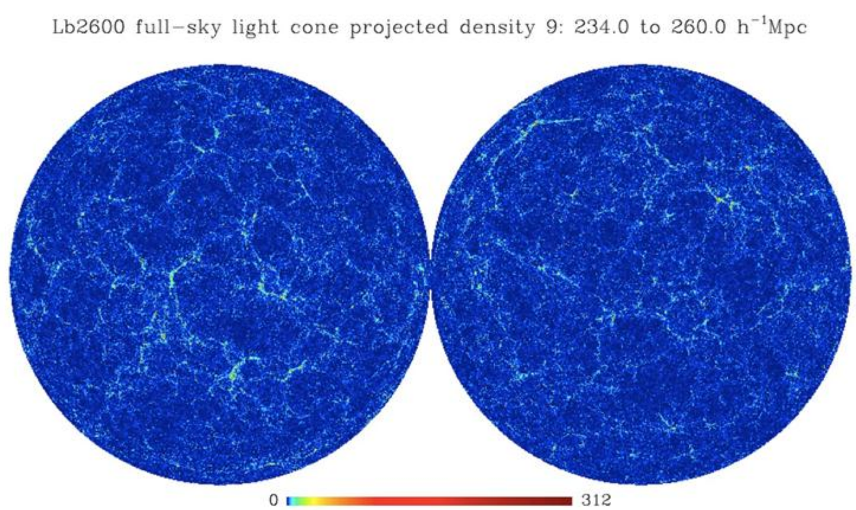

Fig. 1 The density of dark matter in a thin radial slice as seen by a synthetic observer located in the 8 billion light-year computational volume. Image courtesy Matthew Becker, University of Chicago.

Fig. 2: A synthetic 2x3 arcmin DES sky image showing galaxies, stars, and observational artifacts.
Courtesy Huan Lin, FNAL.

DES will be the first project to combine four different statistical methods to probe the properties of dark matter and dark energy. Each method has associated sources of systematic error. Our simulations will provide support for analyzing systematics in the three methods associated with large-scale structure. We are coordinating a Blind Cosmology Challenge, in which a variety of sky realizations in different cosmologies are analyzed, in a blind manner, by DES science teams. The BCC process depends on a number of computationally intensive steps, execution of which can be greatly facilitated in a workflow environment supported by XSEDE. The resultant catalogs offer stringent tests of science pipelines' ability to robustly derive cosmological parameters with design accuracy.

Principal Investigators: Prof. August Evrard (University of Michigan), Prof. Andrey Kravtsov (University of Chicago).

14.3 User-facing Activities

14.3.1 System Activities

Startup Gateway projects:

IU contributions to Extended Collaborative Support for Science Gateways

Gateway Leadership: IU gateways group contributed to the ongoing gateway co-ordination and developing work plans and executing them for requests made through ECSS in gateway related area. The activities include the community engagement and co-ordination of gateway support personal across all the sites.

Specifically this quarter, the IU group has lead the gateway requirements sessions, a collaborative requirement writing exercise involving all funded science gateway ECSS staffers. The group will formalize these requirements into a cohesive document for the consumption of the XSEDE architecture and design team.

Dark Energy Survey Simulation Working Group: The Dark Energy Survey (DES) is an upcoming international experiment that aims to constrain the properties of dark energy and dark matter in the universe using a deep, 5000-square degree survey in five optical bands. The production of a suite of large N-body simulations will allow the development of a Blind Cosmology Challenge (BCC) process for the DES collaboration. The IU gateway extended support is focusing on developing workflow solutions for N-body production process executing 2LPT and L-GADGET codes. The workflow will prepare input files for these tasks, submit 2LPT initial conditions (ICs) code, verify correct execution, archive IC's, submit L-gadget simulation code, verify correct execution, archive output, re-submit and verify until code terminates successfully.

The simulation group has developed python code that manages directories, code compilation and parameter files for the first few steps of the overall production process (power spectrum generation, N-body initial condition generation, and N-body calculation) to minimize errors that could be introduced by setting up each step manually. This python program is included as a first step in an Apache Airavata scientific workflow framework that wraps the code and is capable of managing all of the jobs that need to run for the production process. The workflow has been extensively tested on small simulations and is ready to be used for a full-scale production simulation with the service units extended from last year's allocation.

Next quarter will focus on: running the production workflow for N-body simulation on TACC Ranger cluster; migrating the DES code and Workflow to SDSC Trestles; Explore file transfer mechanisms to Archival systems and post processing environment; and providing automatic processing steps within the workflow.

***Ocean Land and Atmospheric Model (OLAM)*:** The IU gateways group is continuing its extended collaborative support to the OLAM project. During this quarter, the Ocean Land Atmosphere Model (OLAM) Earth System Model was deployed and XSEDE computational resources in order to simulate current and future climate scenarios, high- impact storm events (hurricanes) and local circulations. The OLAM web portal science gateway is being developed for teaching graduate-level meteorology, climate, and predictability courses at the University of Miami, and generates regional climate change projections that can be used to guide water management decisions in South Florida.

In this quarter the support focused on build gateway infrastructure for classroom teaching of OLAM model at University of Miami. Specific activities include: building a portal interface to specify model domain, user interface to configure temporal information and model input configuration and a wizard style interface to launch the model on specified grids. Due to a problem with the model code, on PI's request, the gateway group has also supported the WRF and WRF Pre-processing codes on SDSC Trestles.

Next quarter will focus on integrating the gateway/portal with back end services to launch, monitor jobs to XSEDE resources and build interface to download/visualize model outputs for both WRF and OLAM models.

Startup Gateway projects:

Einstein Genome Gateway: The Einstein Genome Gateway is targeted to provide large scale computational interface between XSEDE compute and storage resources and the national community with large-scale computational needs in genomics and related fields.

The IU science gateways group is recently assigned to initiate the ECSS support on this project and had initial discussions with the project team. Quarry gateway hosting infrastructure is setup as a development testbed and targeted production platform.
Next quarter will focus on job management and workflow abstractions to seamless migrate from Einstein medical school's computational resources to XSEDE.

Petascale computations in mineral physics with the Quantum ESPRESSO (VLAB): The IU gateways group is helping the VLAB project team to build workflow infrastructure to perform first principles calculations of unprecedented magnitude and scope in mineral physics. The project will optimize the execution of the Quantum ESPRESSO software for materials simulations in XSEDE systems.

The IU gateways group is helping the VLab developer Pedro da Silveira in using Apache Airavata workflow APIs in submitting and monitoring jobs to XSEDE resources.

14.3.2 Services Activities

Extended Collaborative Support for Science Gateways

The IU ECSS team worked with DES simulation group to design and develop a workflow environment to run their N-body production process in an efficient manner. The simulation group has developed python code that manages directories, code compilation and parameter files for the first few steps of the overall production process (power spectrum generation, N-body initial condition generation, and N-body calculation) to minimize errors that could be introduced by setting up each step manually. This python program is included as a first step in an Apache Airavata scientific workflow framework that wraps the code and is capable of managing all of the jobs that need to run for the production process. Figure 1 below shows the graphical interface for running a fully featured suite of cosmology simulations. The workflow has been extensively tested on small simulations and is ready to be used for a full-scale production simulation with the service units extended from last year's allocation. We are delaying running the production level simulation to completion because, upon starting the workflow for full production on March 29th, we found inconsistencies with some previous work. We want to clean up this previous work before moving on to the new work that will be run entirely by the workflow. With the workflow in place on Ranger, we can potentially burn ~75k CPU hours per day in full production, (3 jobs, 1024 cores each, running for 24 hours).The work has demanded some new features be added to the Apache Airavata framework, which will be made generally available through future software releases. The development and results from the first production run have led to a paper being prepared for the XSEDE12 conference.

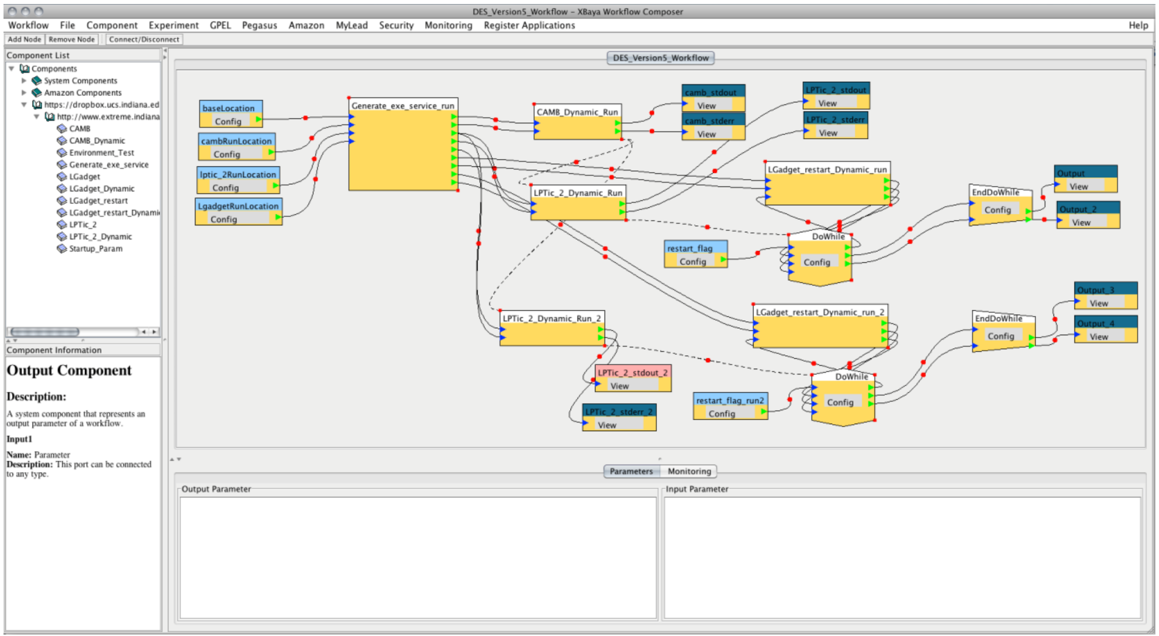

Plans for next quarter:

- Submit XSEDE paper before April 25rd.
- Run the production workflow for N-body simulation on Ranger.
- Migrate DES code and workflow to use SDSC Trestles
- Explore file transfer mechanisms to Archival systems and post processing environments
- Automate post processing steps into workflow.

14.4 Security

No changes.

14.5 Education, Outreach, and Training Activities

Campus Bridging

The Campus Bridging Team engaged in two major activities during this quarter: definition of use cases for the Architecture and Design team and a pilot project for testing the Global Federated File System (GFFS) software on campuses.

In cooperation with the Architecture and Design team, the CB team created seven use cases for campus bridging activities and began the construction of requirements and quality attributes for fulfilling these use cases. These use cases cover use of resources by campus researchers, data movement between researchers and XSEDE resources, and making local resources available to the larger community via XSEDE in exchange for other resources.

The GFFS Pilot program identified 2 friendly-user sites for initial testing (Indiana University and Louisiana Tech) and 4 pilot sites, out of 17 total proposals for participation in the GFFS program. XSEDE Operations team has identified resources at each of the XSEDE SP's which will accommodate GFFS testing. Work has started with the friendly-user sites to identify requirements for interoperating with XSEDE SP resources. XSEDE will be working with the Galaxy gateway at IU and with Tom Bishop's research team at Louisiana Tech. The pilot project teams have had an initial kickoff meeting and will be meeting bi-weekly in order to move project work forward. The XSEDE User Portal will create a space in which the Campus Bridging pilot teams can work and store documents, similar to the Campus Champions space in the XSEDE portal.

XSEDE Campus Bridging will participate in a panel at the Internet2 Spring Member's meeting April 24th on Campus Bridging initiatives in XSEDE, in the Open Science Grid, at IU, and at Cornell.

1.1 SP Collaborations

No new collaborations during this reporting period.

14.6 SP-Specific Activities

XSEDE12 Conference

Planning continues for the XSEDE 12 Conference, with registration set to open on May 14. The Call for Papers has been distributed widely, with a due date of April 25. Sponsorship solicitation letters and prospectus have been sent to a list of about 30 potential contributors, with a goal of raising $150,000. A site visit to the InterContinental Hotel in Chicago was made to assess overall logistics with further follow-up on a few specific issues. Richard Tapia has accepted an invitation to be the keynote speaker at the conference. The XSEDE Student Program is being chaired by Jenett Tillotson from IU. She hopes to expand the program this year to include a "student day" for local area high school students.

XSEDE Annual User Survey

Work began on the 2012 XSEDE User Satisfaction Survey in mid-February, with the aim of defining a new process for evaluating how users perceive and use XSEDE services (systems, software, support, etc.), what they recommend for changes and improvements to XSEDE services, and how those perceptions and inputs change over time.

Drawing on lessons learned and data gathered from previous TeraGrid surveys, Indiana University (IU) proposed the use of a new, comprehensive survey instrument that would be sent to two populations. The survey instrument would be similar, in parts, to past surveys, but diverge based upon the population being surveyed. The two populations to be surveyed are:

- XSEDE users: categorized as PIs, postdoctoral fellows, research staff, graduate students, and undergraduate students conducting work on XSEDE resources using XSEDE allocations.

- Researchers who are federally funded PIs but not XSEDE users: Researchers funded as a PI by the NSF, DOE and NIH between 2006-2010, inclusive, but who do not have accounts or allocations on XSEDE. This category would also include potential users of XSEDE resources who do not have federal awards, but who have appropriate uses and applications of XSEDE resources. This would include campus personnel, faculty and students who have cyberinfrastructure research and/or education needs that would benefit from the use of XSEDE resources and services. (This information will also provide data of potential use in discussions among NSF, NIH and DOE).

Working with the XSEDE project's User Engagement team (and various service owners) and beginning with the questions used in past surveys, a new survey instrument has been developed and is presently in final draft form; the final draft is expected to go before XSEDE leadership (and the XSEDE User forum) for review and approval during the week of April 9. The survey (in its current draft form) has been submitted to the IU Institutional Review Board, and will be administered by the IU Center for Survey Research. Upon approval by the IU IRB and XSEDE leadership, the survey will be deployed for not less than one month. A full report will be published as a technical report in advance of the June NSF review of XSEDE.

XSEDE Ops Failover

IU GlobalNOC, continued work on plans for failover of the XSEDE Operations Center. Several critical areas are pending implementation in order to move forward with failover plan. Stephen McNally is to give further guidance when issues in these areas have been resolved. These areas include:

- Phones – completed XSEDE NOC greeting on IU phone system. - Pending Toll Free provider redirecting the number and process.
- Email – completed setup of Email acceptance for failover on IU side. Pending XSEDE/NCSA setup, policy and procedures, documentation before we continue. Continuing to communicate with them – **pending XSEDE.**
- Ticketing – current ticketing system is non-portable. Replacements being considered are RT and needs a decision and implementation. **Pending XSEDE.**
- Monitoring – currently SNAPP monitoring from Pittsburgh. Steve looking at how to get GRNOC access to current monitoring. **Pending XSEDE.**
- Process and Procedure Documentation – XSEDE currently looking at centralizing all XSEDE documentation to "staff wiki" – **Pending XSEDE.**
- Training - needs to be done for GRNOC staff – whether XSEDE personnel come here or subset go to University of Illinois ? Recommend whatever is chosen – it is a video that can be used across shifts new/different personnel. – **Pending XSEDE.**
- Fail-over Documentation - Steve to work on documenting the fail-over plan when's, how-to's, etc… – **Pending XSEDE.**

XSEDE Networking

During the reporting period the GlobalNOC completed the following tasks related to XSEDE:

- Transition of IU and Purdue XSEDE connections to Monon100.
- Represent IU in the Network-Ops call with other XSEDE Participants.

14.7 Publications

No publications or presentations for this reporting period.

14.8 Metrics

14.8.1 Standard systems metrics

No systems metrics to report for Indiana University for this quarter.

14.8.2 Standard User Assistance Metrics

XSEDE Knowledge Base

Summary statistics for the XSEDE Knowledge Base:

- Total number of documents available in the XSEDE KB at the end of the 1st quarter = 478
- Number of documents added = 73
- Number of documents modified = 37
- Total number of documents retrieved = 108,076
- Total number of documents (retrieved minus bots) = 68,619

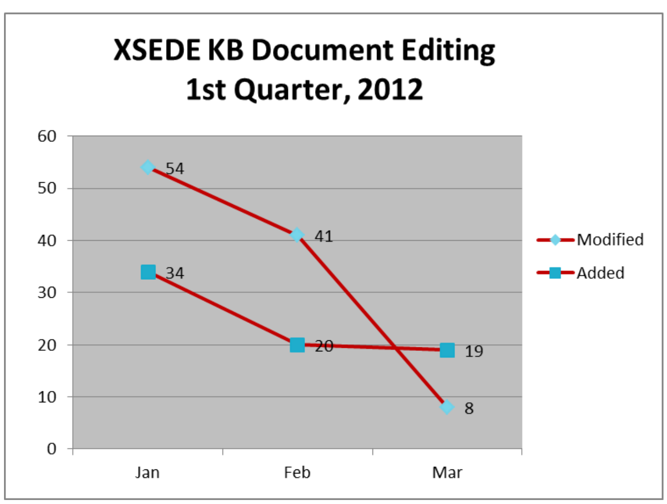

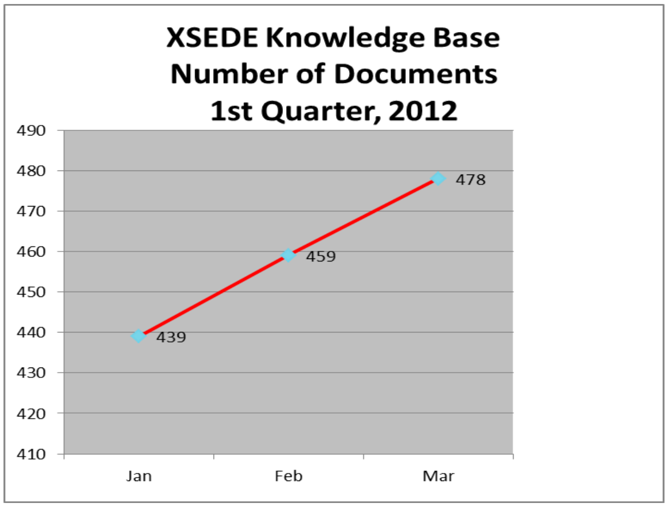

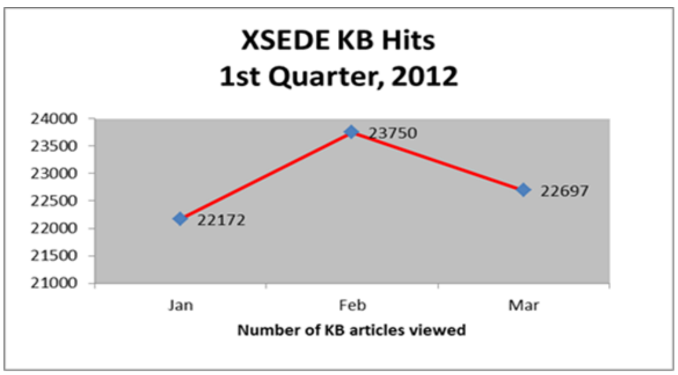

14.8.3 SP-specific Metrics

Quarry Virtual Machines

There was an outage of 56 hours during February 2012 of the Quarry Gateway Hosting environment that affected 10 virtual machines.

15 LONI/LSU - Service Provider Quarterly Report

15.1 Executive Summary

Provide short, high-level self-contained description of significant activities by the SP institution during the reporting period.

15.1.1 *Resource Description*

QueenBee is a 50.7 TFlops Peak Performance, 680 node Dell server cluster. Each node is equipped with 2 quad-core Intel 2.33GHz Xeon 64-bit processors, and 8 GB RAM. SDR (10Gb/s) InfiniBand is used for the interconnection fabric, and 10Gb/s Ethernet connects the system to the LONI optical network. 400TB of Lustre-based storage provides workspace for parallel jobs. The operating system is currently Red Hat Enterprise Linux (RHEL) v4. The system is allowing existing TeraGrid allocations to continue running until exhausted, but no new allocations are currently being accepted. The system is approaching 5 years of age and is out of warranty.

15.2 Science Highlights

Focused descriptions of science accomplishments by users of SP resources and services. *(Science highlights that span multiple SP institutions and/or use XSEDE services would be preferentially used for XSEDE-wide science highlights.)*

Each highlight answers the following questions, and optionally includes a representative image:

What is the science problem that the PI is attempting to simulate or solve?
What were the scientific results and what is their impact?
What activities happened on the SP resources, and how did the SP support the work beyond providing the resources?

15.3 User-facing Activities

(If multiple awards at the SP institution, provide sub-sections for each award and for common infrastructure/services)

15.3.1 *System Activities*

What activities have happened with respect to the SP resources and supporting infrastructure (filesystems, networking, etc.) that impact users positively or negatively, and if negatively what is being done to address it? Describe new deployments (or decommissioning), significant configuration or performance changes, any significant problems or successes, and key system statistics like uptime, availability and utilization.

15.3.2 *Services Activities*

What activities have happened with SP-provided services that impact users positively or negatively, and if negatively what is being done to address it? Includes user support procedures/activities, scheduler/queue options, software packages, portal work, developed software, etc. *(Report on SP-funded services; XSEDE-funded services hosted/delivered at the SP site would be preferentially reported in the general XSEDE report.)*

15.4 Security

Describe any change in security procedures and all security incidents/responses.

15.5 Education, Outreach, and Training Activities

(There may be overlap between SP-funded activities reported here and XSEDE-funded activities reported in the main report. If gray, err on the side of redundancy in both locations.) Details all EOT SP-funded or SP-based activities, including classes, tutorials, workshops, training, etc. For training and workshops, include location, number of attendees, and a brief description of the goals of the training. Where appropriate, describes feedback we've received on our EOT activities.

1. Training
2. Education
3. Conferences/workshops
4. Presentations

15.6 SP Collaborations

Collaborations with non-XSEDE institutions/projects, including international collaborations (ExTENCI comes to mind but if it's reported in the XSEDE whole, no need to report here). *(Focus is on SP-based collaborations; XSEDE-based or multi-SP collaborations would preferentially be reported in the general XSEDE report.)*

15.7 SP-Specific Activities

(If multiple awards at the SP institution, provide sub-sections for each award and for common infrastructure/services) A free format section for the SP to highlight activities of their choosing that do not fall into other defined sections of this report. This could include a description of innovative work, relevant technologies or other developments that has not yet impacted users but is relevant to current or future SP or XSEDE activities.

15.8 Publications

All publications and presentations by SP-funded staff members and publications that utilized SP resources in the prior reporting period in bibliography format. *(There may be overlap between SP-funded publications reported here and XSEDE-funded publications reported in the main report. If gray, err on the side of redundancy in both locations.)*

15.9 Metrics

If multiple awards at the SP institution, provide sub-sections for each award

15.9.1 Standard systems metrics

XDMod team will provide tables of metrics. The originally proposed list:

1. Job summaries by core count or memory footprint
2. Machine usage by scientific discipline
3. Machine usage by project
4. Top 20 users, sorted my machine use
5. Other standard metrics that the SP-Forum and NSF agree to.

15.9.2 Standard User Assistance Metrics

(These metrics are for front-line user support only. More hands-on user support (AUSS-like) may be discussed in Science Highlights or Services Activities.) *(As with common system metrics above, a standard set of metrics will be generated by XSEDE staff (?) for each SP that uses the XSEDE ticketing system. If a different ticketing system is used locally, the SP must produce the same common metrics.)*

1. Tickets opened during report period.
2. Tickets closed during report period.
3. Distribution of response time on tickets during the reporting period.

Ticket metrics for this section will depend on what will be provided by the XSEDE central.

15.9.3 SP-specific Metrics

(if multiple awards at the SP institution, provide sub-sections for each award and for common infrastructure/services)

The SP may include additional metrics specific to their resources and services. There should be reasonable continuity in the reported metrics/format from report to report.

The types of information included here for various SPs may be integrated into the XDMoD for future reports.

16 NCSA - Service Provider Quarterly Report

16.1 Executive Summary

NCSA’s *Forge* GPU computing system is retiring September 30, 2012. Users were sent a notice to prepare them for the transition to other systems. NCSA is determining if it can maintain access to the tape archive beyond September 30, 2012.

16.1.1 *Resource Description*

Forge consists of 36 Dell PowerEdge C6145 quad-socket nodes with dual 8-core AMD Magny-Cours 6136 processors and 64 GB of memory. Each node supports 6 NVIDIA Fermi M2070 GPUs.

NCSA's hierarchical archival storage system is available for permanent storage of data. Access is via FTP- and SSH-based transfer clients, including GridFTP clients. NCSA's mass storage now holds more than six petabytes of data and has the capacity to archive ten petabytes of data.

16.2 Science Highlights

Computational Mathematics: Hybrid Lagrangian/Eulerian rotorcraft wake simulations using GPGPU (Christopher Stone, Computational Science & Engineering LLC and , and Earl Duque, Intelligent Light & Penn State)

A big problem with helicopters is the interaction between the vortex wake coming off the rotor blades and how those vortex structures affect the lift behavior and drag behavior of the rotor blades themselves. The disturbed air may further change as it passes over the fuselage of the vehicle, creating new wake structures that hit the tail. This causes a condition called tail shake that results in vehicle control problems—such as the pilot can’t land.

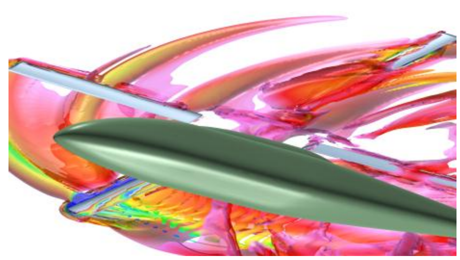

Figure 1. A visualization of the wakes coming off a helicopter rotor. The simulation was conducted with a new coupled Eulerian and vortex particle method code. The visualization was then created using FieldView. Image courtesy of Intelligent

That’s why Earl Duque, manager of the Applied Research Group (ARG) at Intelligent Light, together with Christopher Stone, founder of Computational Science & Engineering LLC, used NCSA’s Forge and now-retired Lincoln to develop and test a simulation software that would more accurately predict rotor wake interactions. The ARG was a participant in a five-year project working on CFD research for rotorcraft as part of the Army’s Vertical Lift Research Center of Excellence based out of Penn State.

Ability to accurately model wakes and predict tail shake will dramatically reduce the number of wind tunnel or flight tests required, which are expensive. For their project, Stone fused standard grid-based finite-difference (Eulerian) CFD software with the vortex particle method. Grid-based methods are pretty good for modeling the flow close to the fuselage body or the rotor wing, but are not very accurate farther away from the body. If the CFD used is not accurate enough, the vortex that’s striking the tail rotor won’t have the right physics thus causing the tail rotor shake in flight tests. With the hybrid method, he and Duque were hoping to more accurately predict how a vortex would flow back and strike the tail rotor, providing improved predictions of the rotorcraft's handling.

Stone, who was the first to utilize GPUs in the hybrid code, reports a magnitude difference in code performance using NCSA's GPU machines. But more importantly, he says, is that they had the computing power to attempt simulations that were not previously possible. They simulated loads on the rotor and fuselage in real-world flight conditions and then used Intelligent Light's FieldView software to visualize the predicted rotor wakes interacting with a helicopter fuselage. Their original theory that the hybrid method is superior to predict the rotor wake far down stream was upheld with this project, but they are still working on the validation.

16.3 User-facing Activities

16.3.1 System Activities

Filesystems and Storage:

NCSA is continuing deployment of the new tape archive.

Networking:

The NCSA XSEDEnet circuit continues to see use since coming online. The report below shows the peak utilization the circuit has seen in the 1/1/2012 – 3/31/2012 timeframe.

Network Utilization expressed in percent of a 10Gb circuit utilized:

Interfaces Reporting Tool - Utilization

Tue 10 Apr 2012, 11:10 (America/Chicago)

Sun Jan 1 00:00 2012 to Sun Apr 1 00:00 2012

Device	Interface	Rx Util	Tx Util
exit-east	xe-0/0/2	37.8	65.2

1 record

The NCSA network supporting our SP resources is redundant and fault tolerant. As a result we have been able to sustain a network service uptime of 100% to our SP resources. This measurement is over the 1/1/2012 – 3/31/2012 period in which NCSA has been connected to XSEDEnet.

Ping Availability

Tue 10 Apr 2012, 11:22 (America/Chicago)

Sun Jan 1 00:00 2012 to Sun Apr 1 00:00 2012 America/Chicago

Totals (SLA Transitions: 0, SLA Outage Time: 0secs, Transitions: 0, Outage Time: 0secs) Averages (SLA Availability: 100.00%, Availability 100.00%)

Device	SLA Availability	SLA Transitions	SLA Downtime	Availability	Transitions	Downtime
npcf-hpc-1	100.00	0	0d 00h 00m 00s	100.00	0	0d 00h 00m 00s
exit-east	100.00	0	0d 00h 00m 00s	100.00	0	0d 00h 00m 00s

2 devices displayed

NCSA continues to manage the DNS services for XSEDE which includes the teragrid.org and xsede.org domains. We delegate subdomains to other SP sites and make DNS adds and deletions as requested by project groups. We setup and mange the dynamic dns system. We also acquire certificates for the various web services being deployed on XSEDE.

Systems:

NCSA's new GPU-based computational resource *Forge* transitioned from friendly user to production this quarter, replacing the previous system *Lincoln*. The system is described above.

16.3.2 Services Activities

Continued support of NCSA's Dell NVIDIA Cluster Forge and the archival system MSS.

16.4 Security

The NCSA security team did not have any incidents involving XSEDE users or resources for the first quarter 2012. There were also no changes to the security monitoring during the quarter.

16.5 Education, Outreach, and Training Activities

NCSA provides online training via CI-Tutor. Course usage statistics are captured in the following table:

Course Title	**Jan-12**	**Feb-12**	**Mar-12**	**Total**
Access Grid Tutorials	5	4	3	12
BigSim: Simulating PetaFLOPS Supercomputers	6	5	1	12
Debugging Serial and Parallel Codes	7	12	9	28
Getting Started on theTeraGrid	9	3	4	16
Intermediate MPI	22	22	13	57
Introduction to MPI	104	116	134	354
Introduction to Multi-core Performance	12	9	30	51
Introduction to OpenMP	32	64	33	129
Introduction to Performance Tools	6	10	9	25
Introduction to Visualization	5	13	6	24
Multilevel Parallel Programming	17	8	11	36
Parallel Computing Explained	56	52	26	134
Parallel Numerical Libraries	13	9	6	28
Performance Tuning for Clusters	12	7	5	24
Tuning Applications for High Performance Networks	3	3	2	8
Using the Lustre File System	5	14	4	23
XSEDE Cybersecurity	5	5	2	12
Total	320	356	299	975

16.6 SP Collaborations

Nothing to report.

16.7 SP-Specific Activities

Continuing work on Residual Based Turbulence Models for Large Eddy Simulation on Unstructured Tetrahedra Meshes with PI A. Masud, UIUC. Involves visualizing high Reynolds flow around a moving airfoil. The computational aspect focuses on developing novel numerical methods that capture multiscale physics on unstructured grids. Visualization focused on animations that show the flow evolution over the periodic vertical movement of the airfoil.

Continuing work on Supernovae simulations with PI Dean Townsley, University of Alabama under charge number (TG-AST090086). Involves extensive work using custom rendering system to generate a variety of different volume visualization representations. Included below are some examples of such volume representations showing the "flam" variable and associated colormaps

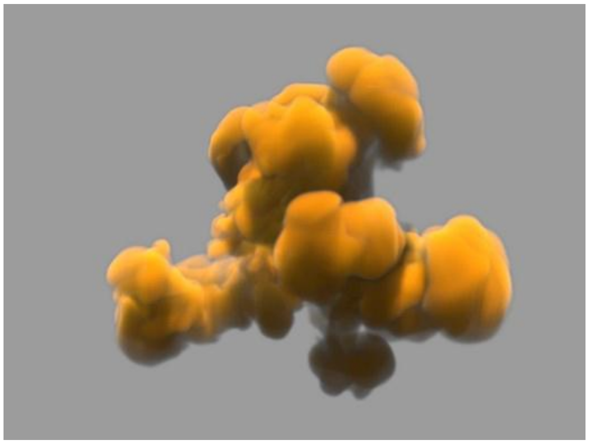

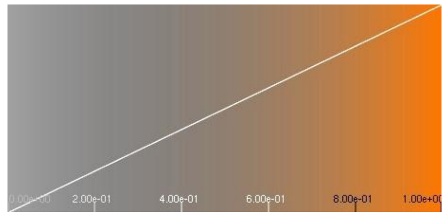

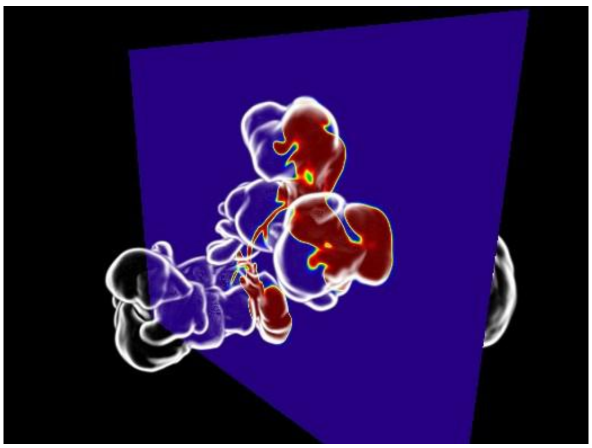

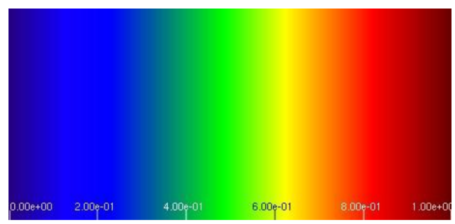

Initial work on Interactive Large Scale Media Analytics with PI Virginia Kuhn. Involves developing visualization representations to analyze underlying video databases associated with the tagging system. Investigating visual analytics for media data.

16.8 Publications

Papers

None to report this quarter.

16.9 Metrics

NCSA Ticket resolution times by category from XSEDE Ticket System

Time to Resolution	account issues	file systems	grid software	jobs/batch queues	login/access issues	mss/data issues	network issues	software/apps	system issues	other
0-1 hr										
1-24 hr										
1-7 d						1			1	1
1-2 wk						1			14	
> 2 wk						1			13	
Still Open			1	1		3			8	

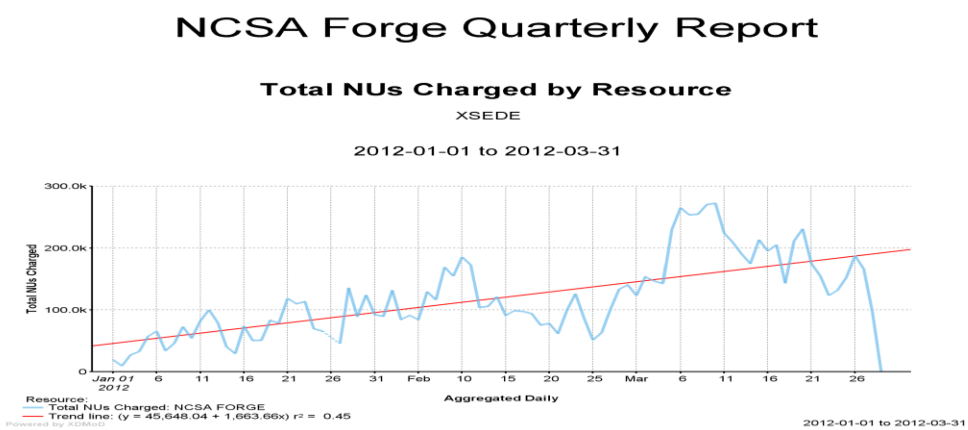

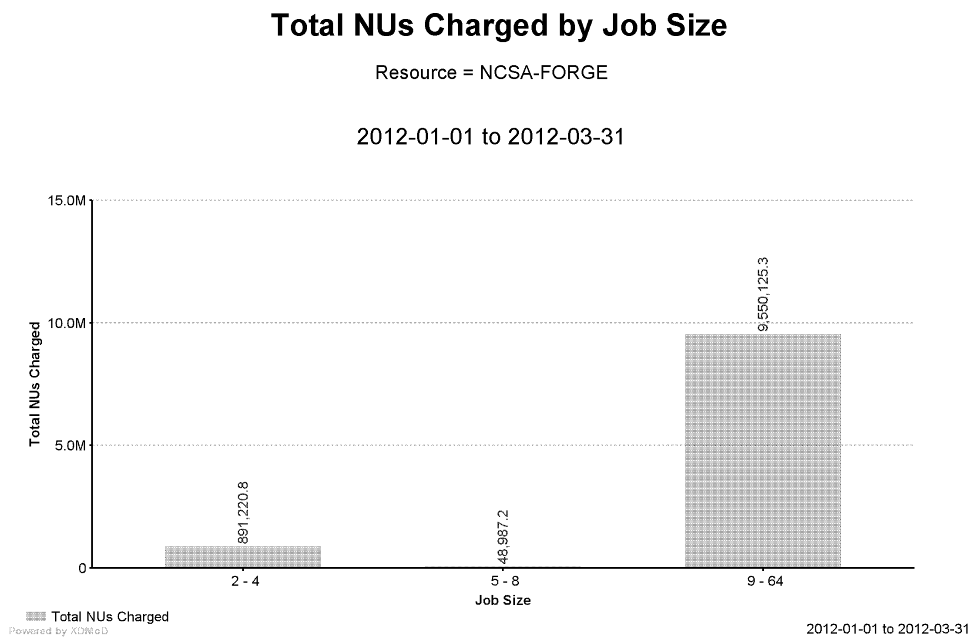

Total NUs Charged by Job Size
Resource = NCSA-FORGE
2012-01-01 to 2012-03-31
15.0M
10.0M
5.0M
0
Total NUs Charged
891,220.8
48,987.2
9,550,125.3
2 - 4
5 - 8
9 - 64
Job Size
Total NUs Charged
Powered by XDMoD
2012-01-01 to 2012-03-31

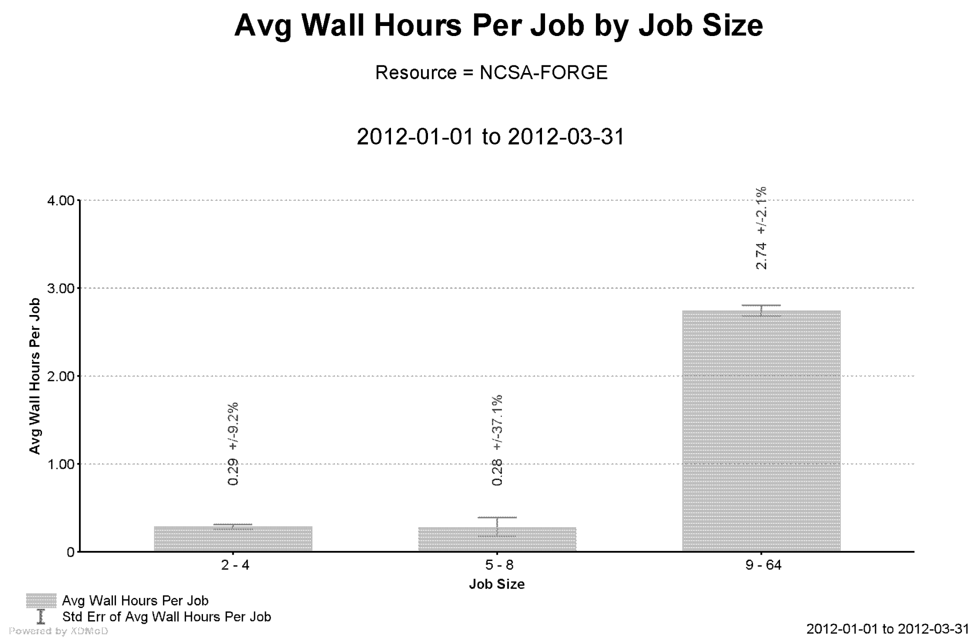
Avg Wall Hours Per Job by Job Size
Resource = NCSA-FORGE
2012-01-01 to 2012-03-31
4.00
3.00
2.00
1.00
0
Avg Wall Hours Per Job
0.29 +/-9.2%
0.28 +/-37.1%
2.74 +/-2.1%
2 - 4
5 - 8
9 - 64
Job Size
Avg Wall Hours Per Job
Std Err of Avg Wall Hours Per Job
Powered by XDMoD
2012-01-01 to 2012-03-31

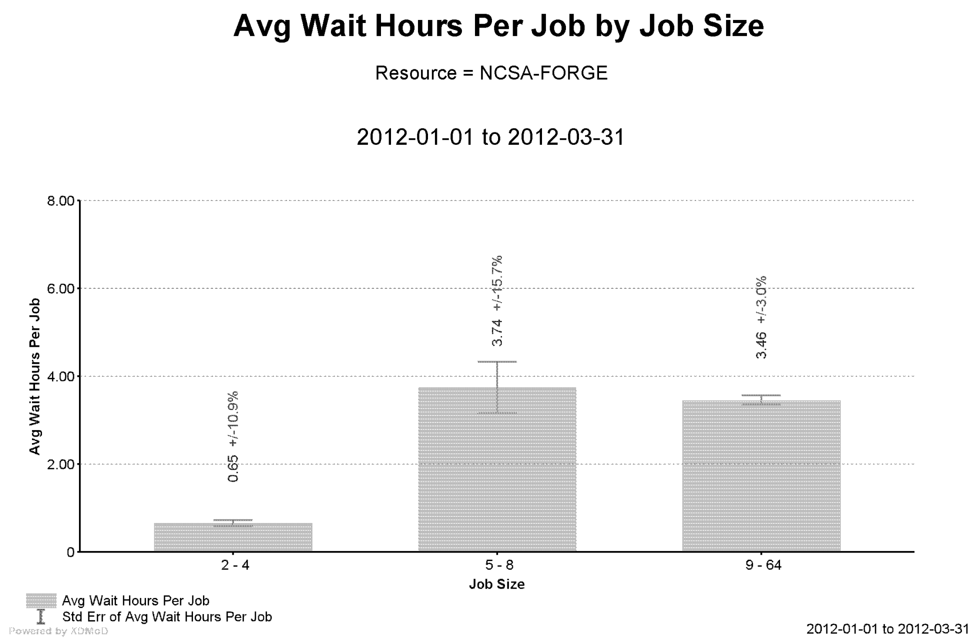
Avg Wait Hours Per Job by Job Size
Resource = NCSA-FORGE
2012-01-01 to 2012-03-31
Avg Wait Hours Per Job
8.00
6.00
4.00
2.00
0
0.65 +/-10.9%
3.74 +/-15.7%
3.46 +/-3.0%
2 - 4
5 - 8
9 - 64
Job Size
Avg Wait Hours Per Job
Std Err of Avg Wait Hours Per Job
Powered by XDMoD
2012-01-01 to 2012-03-31

User Expansion Factor by Job Size

Resource = NCSA-FORGE

2012-01-01 to 2012-03-31

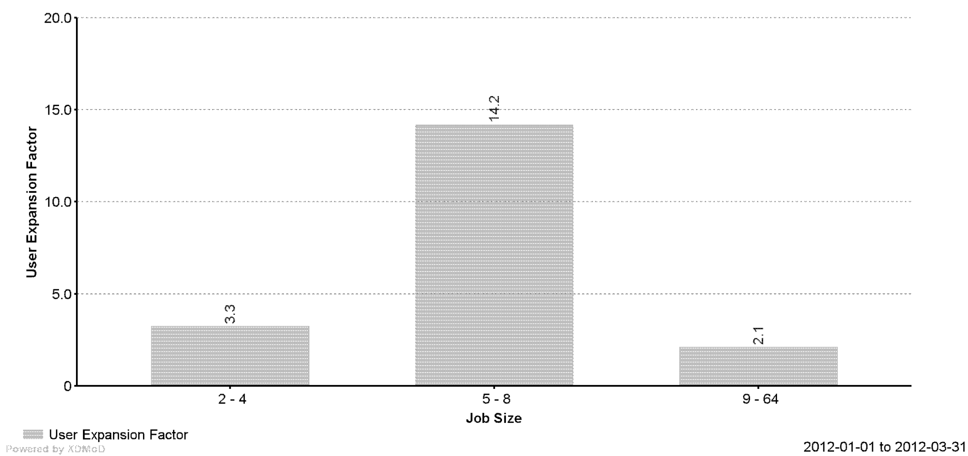

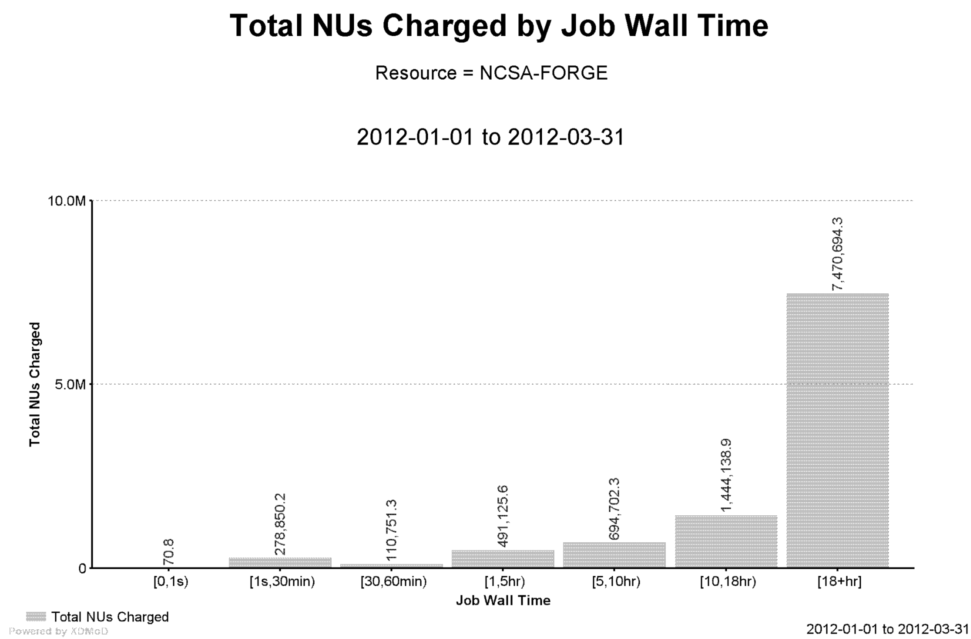

Total NUs Charged by Job Wall Time
Resource = NCSA-FORGE
2012-01-01 to 2012-03-31
10.0M
5.0M
0
Total NUs Charged
70.8
278,850.2
110,751.3
491,125.6
694,702.3
1,444,138.9
7,470,694.3
[0,1s)
[1s,30min)
[30,60min)
[1,5hr)
[5,10hr)
[10,18hr)
[18+hr]
Job Wall Time
Total NUs Charged
Powered by XDMoD
2012-01-01 to 2012-03-31

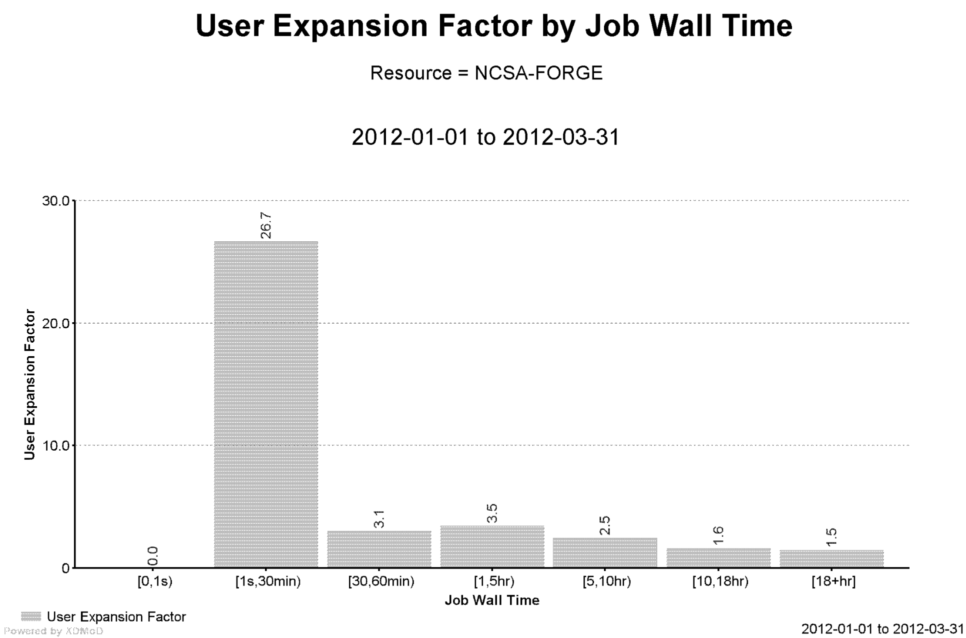
User Expansion Factor by Job Wall Time
Resource = NCSA-FORGE
2012-01-01 to 2012-03-31
User Expansion Factor
30.0
20.0
10.0
0
0.0
26.7
3.1
3.5
2.5
1.6
1.5
[0,1s)
[1s,30min)
[30,60min)
[1,5hr)
[5,10hr)
[10,18hr)
[18+hr]
Job Wall Time
User Expansion Factor
Powered by XDMoD
2012-01-01 to 2012-03-31

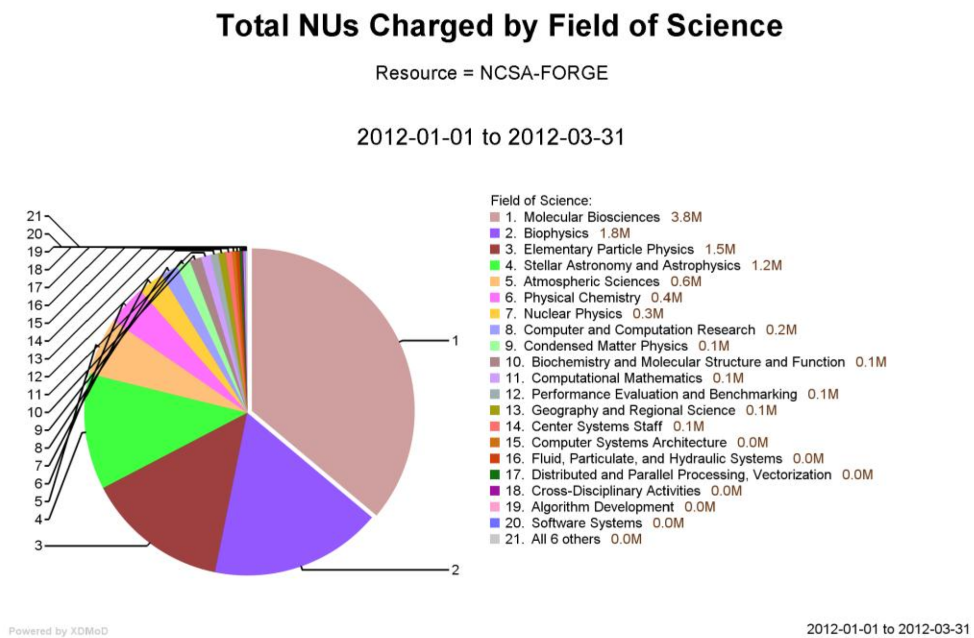
Total NUs Charged by Field of Science
Resource = NCSA-FORGE
2012-01-01 to 2012-03-31
Field of Science:
1. Molecular Biosciences 3.8M
2. Biophysics 1.8M
3. Elementary Particle Physics 1.5M
4. Stellar Astronomy and Astrophysics 1.2M
5. Atmospheric Sciences 0.6M
6. Physical Chemistry 0.4M
7. Nuclear Physics 0.3M
8. Computer and Computation Research 0.2M
9. Condensed Matter Physics 0.1M
10. Biochemistry and Molecular Structure and Function 0.1M
11. Computational Mathematics 0.1M
12. Performance Evaluation and Benchmarking 0.1M
13. Geography and Regional Science 0.1M
14. Center Systems Staff 0.1M
15. Computer Systems Architecture 0.0M
16. Fluid, Particulate, and Hydraulic Systems 0.0M
17. Distributed and Parallel Processing, Vectorization 0.0M
18. Cross-Disciplinary Activities 0.0M
19. Algorithm Development 0.0M
20. Software Systems 0.0M
21. All 6 others 0.0M
Powered by XDMoD
2012-01-01 to 2012-03-31

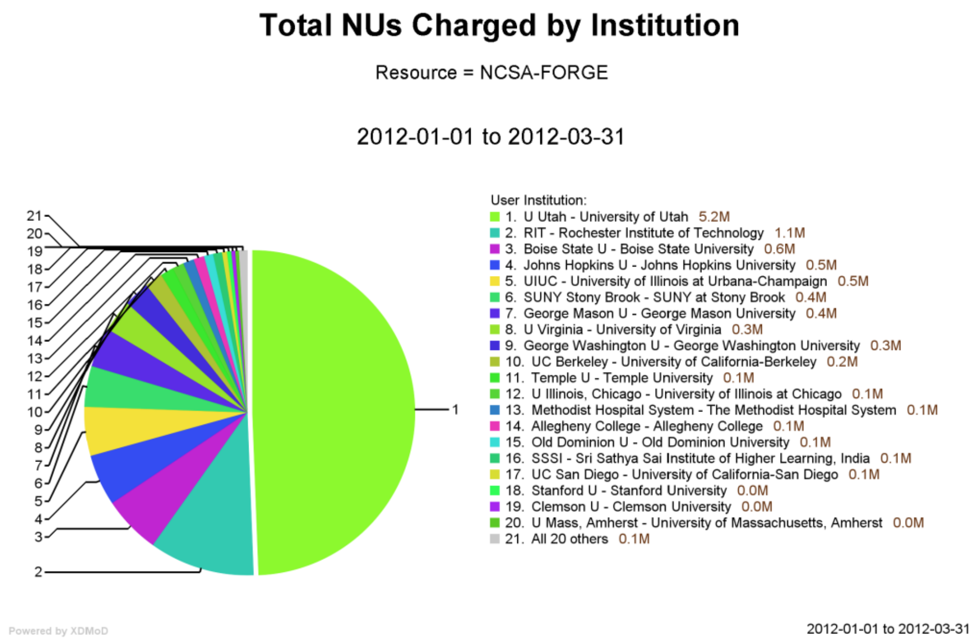
Total NUs Charged by Institution
Resource = NCSA-FORGE
2012-01-01 to 2012-03-31
User Institution:
1. U Utah - University of Utah 5.2M
2. RIT - Rochester Institute of Technology 1.1M
3. Boise State U - Boise State University 0.6M
4. Johns Hopkins U - Johns Hopkins University 0.5M
5. UIUC - University of Illinois at Urbana-Champaign 0.5M
6. SUNY Stony Brook - SUNY at Stony Brook 0.4M
7. George Mason U - George Mason University 0.4M
8. U Virginia - University of Virginia 0.3M
9. George Washington U - George Washington University 0.3M
10. UC Berkeley - University of California-Berkeley 0.2M
11. Temple U - Temple University 0.1M
12. U Illinois, Chicago - University of Illinois at Chicago 0.1M
13. Methodist Hospital System - The Methodist Hospital System 0.1M
14. Allegheny College - Allegheny College 0.1M
15. Old Dominion U - Old Dominion University 0.1M
16. SSSI - Sri Sathya Sai Institute of Higher Learning, India 0.1M
17. UC San Diego - University of California-San Diego 0.1M
18. Stanford U - Stanford University 0.0M
19. Clemson U - Clemson University 0.0M
20. U Mass, Amherst - University of Massachusetts, Amherst 0.0M
21. All 20 others 0.1M
Powered by XDMoD
2012-01-01 to 2012-03-31

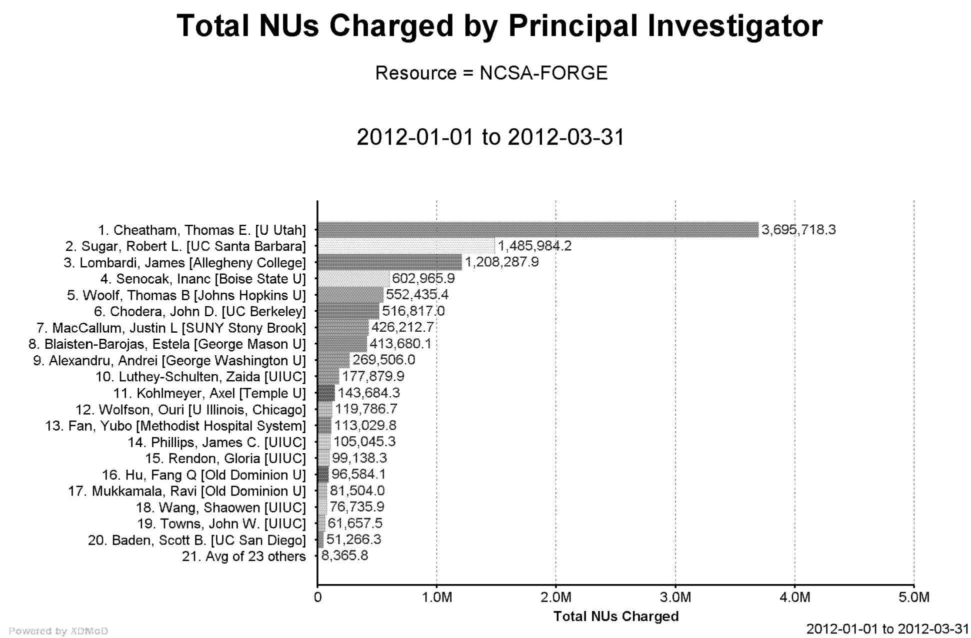

Total NUs Charged by Principal Investigator
Resource = NCSA-FORGE
2012-01-01 to 2012-03-31
1. Cheatham, Thomas E. [U Utah]
2. Sugar, Robert L. [UC Santa Barbara]
3. Lombardi, James [Allegheny College]
4. Senocak, Inanc [Boise State U]
5. Woolf, Thomas B [Johns Hopkins U]
6. Chodera, John D. [UC Berkeley]
7. MacCallum, Justin L [SUNY Stony Brook]
8. Blaisten-Barojas, Estela [George Mason U]
9. Alexandru, Andrei [George Washington U]
10. Luthey-Schulten, Zaida [UIUC]
11. Kohlmeyer, Axel [Temple U]
12. Wolfson, Ouri [U Illinois, Chicago]
13. Fan, Yubo [Methodist Hospital System]
14. Phillips, James C. [UIUC]
15. Rendon, Gloria [UIUC]
16. Hu, Fang Q [Old Dominion U]
17. Mukkamala, Ravi [Old Dominion U]
18. Wang, Shaowen [UIUC]
19. Towns, John W. [UIUC]
20. Baden, Scott B. [UC San Diego]
21. Avg of 23 others
3,695,718.3
1,485,984.2
1,208,287.9
602,965.9
552,435.4
516,817.0
426,212.7
413,680.1
269,506.0
177,879.9
143,684.3
119,786.7
113,029.8
105,045.3
99,138.3
96,584.1
81,504.0
76,735.9
61,657.5
51,266.3
8,365.8
0
1.0M
2.0M
3.0M
4.0M
5.0M
Total NUs Charged
Powered by XDMoD
2012-01-01 to 2012-03-31

17 NICS - Service Provider Quarterly Report

17.1 Executive Summary

In 2009, the National Institute for Computational Sciences (NICS) delivered the first academic petaflop computer to the NSF community—a Cray XT5 called Kraken. By the end of 2010, systems at NICS were delivering more than 70% of all NSF compute cycles. This quarter, Kraken sustained a utilization of 93% and a 99% uptime.

The addition of the SGI Altix, called Nautilus, and the Remote Data and Visualization (RDAV) center serves to broaden the services provided by NICS to the NSF community and increases the potential for breakthrough science (Section 17.2). RDAV's purpose is to aid in the significant challenge of transforming large-scale data into knowledge and insight by providing scientists with well-engineered and well-supported remote visualization, analysis, and scientific workflow technologies. Nautilus provided a 98% uptime for the quarter and 50% utilization.

There were 667 help tickets opened during the quarter and 507 closed with a mean-time-to-resolution of just over 67 hours (Table 5).

17.1.1 *Resource Description*

NICS currently has two NSF funded computational resources: Kraken and Nautilus. These systems share a Network File System (NFS) that contains user directories, project directories and software directories. One-time password tokens provide secure access to both the computational and storage resources at NICS.

17.1.1.1 Kraken

Kraken is a Cray XT5 consisting of 9,408 compute nodes, each containing two 6-core AMD Istanbul Opteron processors and 16 GB of on-node memory. The result is 112,896 compute cores that deliver 1.17 PF at peak performance with 147 TB total memory. Communications take place over the Cray SeaStar2+ interconnect. A parallel Lustre file system provides 3.3 PB (raw) of short-term data storage.

17.1.1.2 Nautilus

Nautilus, an SGI Altix UV 1000 system, is the centerpiece of NICS Remote Data and Visualization (RDAV) Center that is also located at ORNL. It has 1024 cores (Intel Nehalem EX processors), 4 TB of global shared memory, and 8 GPUs in a single system image yielding 8.2 TF at peak performance. A parallel Lustre file system provides 427 TB (raw) of short-term data storage.

17.1.1.3 HPSS Archival Storage

The High Performance Storage System (HPSS), developed and operated by ORNL, is capable of archiving hundreds of petabytes of data and can be accessed by all major leadership computing platforms. Incoming data is written to disk and later migrated to tape for long term archiving. This hierarchical infrastructure provides high-performance data transfers while leveraging cost effective tape technologies. Robotic tape libraries provide tape storage. The center has four SL8500 tape libraries holding up to 10,000 cartridges each. The libraries house a total of 24 T10K-A tape drives (500 GB cartridges, uncompressed), 60 T-10K-B tape drives (1 terabyte cartridges, uncompressed), and 20 T10K-C tape drives (5 terabyte cartridges, uncompressed). Each T10K-A and T10K-B drive has a bandwidth of 120 MB/s. Each T10K-C tape drive has a bandwidth of 240

MB/s. Disk storage is provided by DDN storage arrays with nearly a petabyte of capacity and over 12 GB/s of bandwidth. This infrastructure has allowed the archival system to scale to meet increasingly demanding capacity and bandwidth requirements with more than 9 PB of NICS data stored as of March 2012.

17.2 Science Highlights

17.2.1 Kraken

17.2.1.1 Advanced Scientific Computing: Studying Transportable Explosives with Uintah (Martin Berzins, University of Utah)

See main body of XSEDE quarterly report, 2.1.

17.2.1.2 Atmospheric Sciences: Studying Magnetic Reconnection (Homa Karimabadi, University of California-San Diego; and William Daughton, Los Alamos National Laboratory)

See main body of XSEDE quarterly report, 2.3.

17.2.1.3 Chemical, Thermal Systems: Numerical Modeling of Dense Particulate Flows in Fluidized Bed Reactors (Olivier Desjardins, Cornell University)

See main body of XSEDE quarterly report, 2.4.

17.2.1.4 *Environmental Biology: UT's Kraken supercomputer helps track white-nose syndrome (Thomas Hallam and Jeffrey Nichols, University of Tennessee/Oak Ridge National Laboratory)*

Caitlin Rockett, crockett@utk.edu

In February 2006, a recreational caver documented the North American emergence of white-nose syndrome (WNS), a fungus-driven disease that is rapidly killing cave-dwelling bats. In six years, WNS has killed at least a million bats and has spread from New York to 15 other U.S. states.

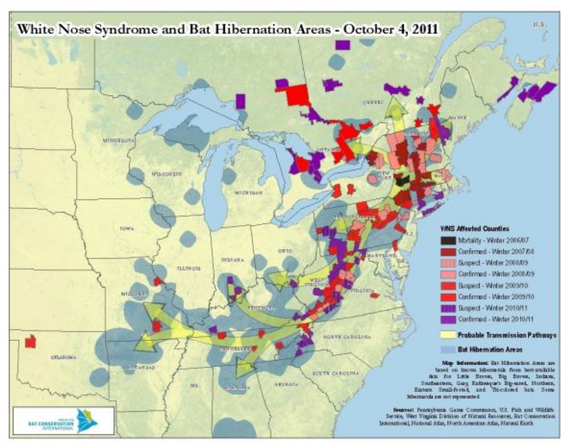

Figure: A map depicting areas of bat hibernation and cases of white-nose syndrome (WNS) from 2006-2011. Solid purple areas are the most recently confirmed cases of WNS, while the yellow arrows show probable transmission pathways (Courtesy: T. Hallam and J. Nichols, University of Tennessee/Oak Ridge National Laboratory).

Using the Cray XT5 known as Kraken at the National Institute for Computational Sciences (NICS), University of Tennessee (UT) recent professor emeritus Thomas Hallam and his colleagues have been exploring how WNS will spread in the continental United States and how to treat it.

WNS is transmitted via direct contact with the fungus *Geomyces destructans.* The ultimate outcome of WNS in North American bats is death—wing deterioration and loss of body mass make it impossible for bats to survive. The fungus produces a fuzzy white growth on the bat's snout, wings, and ears, and sometimes penetrates into the bat's hair follicles and sebaceous glands, generating ulcers on the skin. Of the nine species of bats that have thus far been affected by WNS, all are cave-dwelling, hibernating species.

According to Hallam, the spread of WNS has slowed considerably in the past two to three years—the disease has been found no further south than Tennessee and has yet to surface farther west than Ohio (although the fungus has been located on single bats in Missouri and Oklahoma, it has not caused the mass mortality associated with the disease). Observations and the team's simulations indicate that temperature has something to do with slowing down the spread of WNS.

The team used 6,000 processors per simulation on Kraken to predict how WNS will spread over the next 15 years—each year resulted in one simulation. While the Southern U.S. saw moderated spread, nationwide spread is still predicted despite being slowed by the Great Plains and the Rocky Mountains. "Our simulations show white nose reaching the edge of the plains in 2014, then slowing down in 2015 and expanding rapidly once it reaches Colorado," explained Jeff Nichols, a computational scientist at ORNL and a code developer for Hallam's team.

However, there is a more positive result from the simulations—there is a cave-rich corridor across Southern Illinois that has not been affected by WNS. Hallam and his team have hopes that the caves of this area can be treated with fungicides. If bats with WNS continue to spread west, this corridor acts as a transitional zone between the East and the Midwest. Once the bats reach these caves, the fungicides would treat the disease. Again, deeper research is necessary; no antifungal treatments have undergone substantial testing, and biologists must be sure that any treatments used will not pose a threat to the natural ecosystem of a cave.

Hallam and his colleagues have much more to examine about this perplexing disease. Their future plans involve strengthening their models by incorporating internal cave temperatures, adding more caves, and integrating Canadian data.

17.2.1.5 *Advanced Scientific Computing: Highly Scalable Parallel Docking to Speedup Drug Discovery (Yuri Peterson and Bhanu Rekepalli, Medical University of South Carolina/National Institute for Computational Sciences)*

Caitlin Rockett; crockett@utk.edu

Drug discovery starts with a process known as high-throughput screening (HTS) where huge public libraries of chemical compounds are tested for their effectiveness against a biological target—like a protein—that is known to play a key role in the disease. Laboratory experimental testing is tedious and expensive, but high-performance computers are now able to reduce the time and expenses involved in narrowing down a subset of chemicals to be tested. However, having computational power is only half of the solution; researchers need codes that can scale to these large machines.

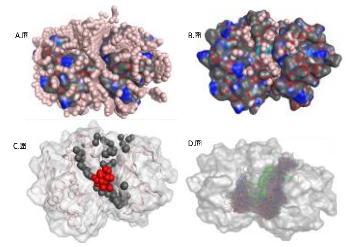

Figure 1. The docking oncoprotein target preparation workflow for HP-D. A. Generation of all spheres for drug placement from showsphere. B. Manual selection of Pocket 2 spheres, C. Addition of spheres to finalize docking site. D. Display of results of random 5000 compounds from zin clean leads on RGS4 for quality control before large scale Kraken HP-D.

Figure: The docking oncoprotein target preparation workflow for HP-D. A. Generation of all spheres for drug placement from showsphere. B. Manual selection of Pocket 2 spheres, C. Addition of spheres to finalize docking site. D. Display of results of random 5000 compounds from zinc clean leads on RGS4 for quality control before large scale Kraken HP-D (Courtesy: Y. Peterson and B. Rekepalli, Medical University of South Carolina/National Institute for Computational Sciences.

Yuri Peterson of the Medical University of South Carolina (MUSC) and Bhanu Rekepalli of NICS are using NICS' flagship system "Kraken" to improve the efficiency of two commonly used codes in drug discovery, Dock6 and AutoDock. Both codes work to simulate how candidate drug molecules will dock (interact) with proteins associated with diseases.

Peterson's group at MUSC was running Dock6 on a small cluster called CBRC (Computational Biology Resource Center) located on campus. The group was using the application to study how 1.3 million compounds dock with seven target proteins that have been identified in ovarian cancer. The team also wanted to change the conformation of the proteins, meaning each compound would be docked to a few variations of each of the seven proteins.

"Assuming their cluster was running 24 hours a day, seven days a week, 365 days a year, it would still take them a few years to achieve that," said Rekepalli. As such, Rekepalli's team began to optimize Dock6 to be run on Kraken. The researchers decided to split the work between Kraken and CBRC, using Kraken to run flexible docking while the cluster ran rigid docking simulations. Flexible docking is anywhere from three to five times more computationally intensive (requiring many processors) than rigid docking because it allows the compound to change shape (as it would in nature) during the docking process.

Kraken was able to dock all 1.3 million compounds to each of the seven proteins in around 40 minutes using only 8,000 of the machine's 112,896 cores—it took three months to complete this task on the CBRC cluster. The scientists also changed the conformation of the proteins (four conformations for each protein), leading to a total of around 36 million compound-protein interactions, all completed within a matter of days. Since these initial test runs, the researchers have done testing with Dock6 up to 36,000 cores on Kraken and results are looking consistent as they scale higher.

Once the researchers have scaled Dock6 to the entirety of Kraken they will run different conformations of each protein until they have narrowed the millions of compounds down to a few hundred. This subset will undergo molecular dynamics (MD) simulations to more deeply investigate the atomic interactions taking place during docking

AutoDock will undergo the same optimization and scaling techniques in order to achieve full scaling on Kraken. While Dock6 only allows the compound to be flexible, AutoDock allows malleability of both compound and protein, leading to a more natural representation of the actual process. Again, millions of compounds will be narrowed down to hundreds and then submitted to MD simulations. By comparing results from AutoDock and Dock6, Rekepalli and Peterson hope that this thorough process eliminates statistical errors and leads to a stronger subset of compounds for experimental testing.

17.2.2 Nautilus

17.2.2.1 Atmospheric Sciences: Seeing is Believing: Nautilus helps paint a clearer picture of the data from the climate science endeavor Project Athena (Lawrence Marx, Center for Ocean-Land-Atmosphere Studies)

See main body of XSEDE quarterly report, 2.2.

17.2.2.2 Chemical, Thermal Systems: Numerical Modeling of Dense Particulate Flows in Fluidized Bed Reactors (Olivier Desjardins, Cornell University)

Caitlin Rockett; crockett@utk.edu

A Cornell University-based team led by Olivier Desjardins has used both the Cray XT5 “Kraken” and SGI UltraViolet “Nautilus” systems at NICS to model the movement of particles in fluidized bed reactors (FBR). These simulations seek to understand and predict the multiphase flow inside FBRs, while ultimately improving the overall efficiency for thermochemical conversion of biomass (plant material) into liquid fuel.

FBRs work by passing a fluid (either a gas or a liquid) through a grainy solid material at high velocity, which causes the solid to behave as though it were a fluid. This type of reactor has been used since the 1920s, primarily by oil and petrochemical industries to increase production of fuel in the United States. While FBRs are still used to produce fuels, they can also be used to produce industrial polymers like rubber and vinyl, and are even found at water and waste treatment utilities.

Desjardins team coupled their in-house computational fluid dynamics code NGA with a Lagrangian particle-tracking scheme in order to simulate dense particle flows in a lab-scale fluidized bed reactor. Parallel performance of NGA was assessed on Kraken and scaled to over 49,000 cores.

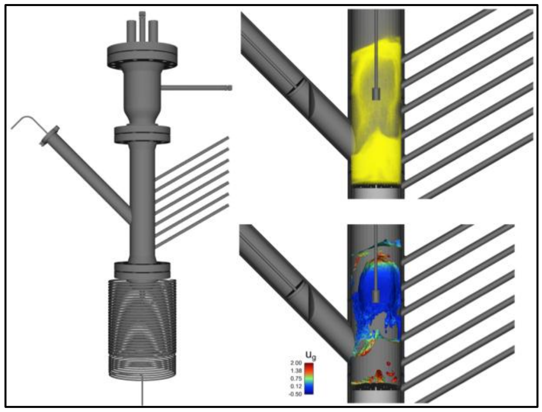

Figure: Simulation of NREL's 4-inch fluidized bed reactor. Reactor geometry (left), Lagrangian particles (top right), iso-surface of gas volume fraction colored by gas velocity (bottom right). The simulation was conducted on NICS Kraken using 8112 cores for approximately 20 days.

The team simulated a 4-inch FBR located at the National Renewable Energy Laboratory at half-scale. This simulation contained 17 million grid cells and nearly 16 million particles and required approximately 4 million CPU hours to run long enough to reach a statistically steady state. The simulation was run using over 8,000 cores on Kraken for approximately 20 days, creating one of the largest particle laden simulations to date.

Particles tend to cluster and fall along the walls of vertical risers (a component of a FBR), binding surrounding particles as they pass and decreasing the efficiency of thermochemical conversion. Understanding the gas-solid interactions is key in improving the overall process efficiency.

The team performed small-scale simulations to reproduce experiments found in literature. Periodic simulations of a section of a lab-scale riser were conducted using 19,000 grid cells and 300,000 particles on 288 cores. Steady state particle velocities were found to disagree with experimental results, indicating a need to simulate the full-scale reactor, which will be the team's next step with this project.

Desjardin has also used Kraken to perform direct numerical simulations of gas-liquid annular flows in horizontal pipes. Annular space refers to the space surrounding one cylindrical object placed inside another. This "pipe within a pipe" scenario is often used to increase efficiency in a number of industries, including oil and gas. It's critical to be able to predict performance of annular flow to insure efficient performance.

"One of the fundamental aspects of our project is the large size of the simulations that we run, and the ability to run high core counts is essential to our progress," explained Desjardins. "This makes access to Kraken a vital aspect of our research. Nautilus has also been a great resource for visualizing our data."

17.3 User-facing Activities

17.3.1 System Activities

17.3.1.1 Kraken

Availability

Kraken had an overall system availability of 97% for this quarter with 69 total hours of downtime. Downtime for the quarter consisted of 67 hours of scheduled downtime and 2 hours of unscheduled downtime (

Table 3).

Table 3: Summary of maintenance Stats for Kraken in Q3PY1.

Maintenance Stats -- Kraken Cray XT5	
Number of planned reboots	6
Number of unplanned reboots	3
Total reboots	9
Job failures due to system faults	712
Total time in period	2183 hrs (100%)
Scheduled Downtime	67 hours (3%)
Unscheduled Downtime	2 hours (0%)
Total Downtime	69 hours (3%)
Total time available to users (total-downtime)	2114 hours (97%)
% System Utilization	93%

Internet2 Connectivity Improvement

The maximum bandwidth of NICS' network link to Internet2 was upgraded from 8 gigabit/s to 10 gigabit/s in February. This has removed the 1 gigabit/s per data stream limitation, and provided better network performance for remote users.

17.3.1.2 Nautilus

Availability

Nautilus had an overall system availability of 98% for this quarter. Downtime for the quarter consisted of 3 hours of scheduled downtime and 28.75 hours of unscheduled downtime (Table 4).

Table 4: Summary of maintenance Stats for Nautilus in Q3.

Maintenance Stats -- Nautilus SGI UV1000	
Number of planned reboots	1
Number of unplanned reboots	11
Total reboots	12
Total time in period	2114 hrs (100%)
Scheduled Downtime	3 hours (0%)
Unscheduled Downtime	28.75 hours (1%)
Total Downtime	31.75 hours (2%)
Total time available to users (total-downtime)	2085 hours (98%)
% System Utilization	50%

17.3.2

17.3.3 Services Activities

17.3.3.1 Kraken

Software Packages
NICS currently supports 893 unique application builds on Kraken that include pre-compiled binaries and builds with PGI, GNU, Pathscale, Cray, and Intel compilers. These builds include 293 unique versions and 141 unique applications and libraries.

Environment
In March, NICS' staff performed an OS upgrade on Kraken, upgrading the Cray Linux Environment (CLE) from version 2.2 to version 3.1. This upgrade also changed the OS on the system's service nodes from SLES 10.1 to SLES 11. The upgrade also changed the version of Lustre on the system from 1.6.5 to 1.8.4. To prepare for the upgrade, the Meta Data Target was migrated to a new RAID array. Since the upgrade, Lustre performance has improved in all tests done.

17.3.3.2 Nautilus

Software Packages
NICS' staff currently supports 293 unique application builds on Nautilus that include pre-compiled binaries and builds with PGI, GNU, and Intel compilers. These builds include 175 unique versions and 103 unique applications and libraries. No new packages were installed during the quarter, but upgrades were applied to R, MATLAB, Visit, and Paraview.

Environment
No changes were made to the environment on Nautilus in this quarter.

17.4 Security

No security incidents or enhancements occurred during the quarter.

17.5 Education, Outreach, and Training Activities

Type	Title	Location	Hours	Date(s)	Number of Participants	Underrepresented Community

Mini- Symposium at SIAM PP12	Large Scientific Applications Performing at Petascale Level and Beyond	Savannah, GA	4	02/14	25	unknown
Workshop/Tutorial	Keeneland workshop (2+ days) Multiple talks, hands-on sessions	Atlanta, GA	16	2/20-21	30	5
Presentation	'Understanding the Encoding of JPEG Images	Univ of Tennessee Knoxville, Mathematics Dept Junior Colloquium	1	1/19	35	5
Presentation	'High Performance Computing in Engineering	Univ of Tennessee Knoxville, Society of Women Engineers	1	2/9	25	23
Conference presentation	'Computer Art and the Letter X'	Atlanta, GA, Gathering 4 Gardner 10	1	3/29	400	40
Conference Exhibit	American Physical Society March Meeting	Boston, MA	20	2/28-3/1	4000	unknown
Workshop/Tutorial	NICS/RDAV HPC Spring Workshop (4 full days)Mulitple talks, hands-on sessions	Univ of Tennessee Knoxville	28	03/19-22	40	8
Seminar	Scientific Computing in Practice at NSF, DoE, and DoD HPC Centers	University of Kentucky, Lexington	2	2/15	20	3
Invited Talk	Challenges to Scientific Understanding at the Exascale	Engineer Research and Development Center, Vicksburg, MS	1	3/15	20 – 30	unknown

17.6 SP Collaborations

EPSCOR

The Experimental Program to Stimulate Competitive Research, or EPSCoR, establishes partnerships with government, higher education and industry that are designed to effect lasting improvements in a state's or region's research infrastructure, R&D capacity and hence, its national R&D competitiveness. NICS participates along with researchers from twenty-seven other states and the Commonwealth of Puerto Rico. The partnership is based on existing and planned collaborations in the advanced materials and systems biology domains where computational science is driving new approaches and insights. The collaborative team has proposed to build cyberinfrastructure (CI) linked, community specific knowledge environments that embody the desktop to XSEDE ecosystem by using campus-based CI at a regional research institution as an essential bridge for connecting faculty investigators to national resources such as the XSEDE.

Keeneland

The Keeneland Project is a five-year, $12 million NSF Track 2D award that has enabled the Georgia Institute of Technology, NICS, and the Oak Ridge National Laboratory to deploy a small, experimental, high-performance computing system consisting of an HP system with NVIDIA Tesla accelerators attached. The project team has been using this initial system to develop scientific libraries and programming tools to facilitate the development of science and engineering research applications, while also providing consulting support to researchers who wish to develop applications for the system. NICS' staff provides administration and support for the Keeneland project.

In 2012, the project will upgrade the heterogeneous system to a larger and more powerful system based on a next-generation platform and NVIDIA accelerators. It is anticipated that the final system will have a peak performance of roughly 2 petaflops. The project will then operate the upgraded system as a XSEDE resource for two years.

17.7 SP-Specific Activities

AACE

In 2011 the Joint Institute for Computational Sciences (JICS) established the Application Acceleration Center of Excellence (AACE) in partnership with NICS, industry leading vendors, and academic institutions. The center's objectives are:

- Accelerate NSF projects toward exascale with state-of-the-art heterogeneous architectures
- Spur development of new algorithms and codes optimized for accelerator-based architectures
- Disseminate fundamental knowledge
- Facilitate effective exchange of expertise and cross-disciplinary collaboration

NICS is assisting NSF users with their transition to the Intel MIC architecture by researching parallelization techniques on the Intel MIC platform and by porting key NSF applications (already millions of lines of code) to the Intel MIC architecture in advance of its commercial release. NICS will also offer training on the Intel MIC architecture following its commercial debut.

"NICS has provided insight to Intel regarding the technology requirements of the scientific computing community," added Joe Curley, director of marketing for Intel's Technical Computing Group. "The impact of our partnership can be seen in the focus of our Intel MIC 'Knights Ferry' software development platform on extending well understood, high-level, standard programming languages and models."

NICS continues to engage a variety of technology providers to determine the role their technologies will play in the quest for exascale, while at the same time providing valued input for product development to meet the needs of the NSF research community.

Industrial Partnerships

NICS currently provides expertise and computational resources to two industrial partners. The goal of the partnerships is to speed innovations to market through application of leading edge simulation capabilities. The benefit to the industrial partners is a condensed design cycle and reduced prototyping and manufacturing costs.

17.8 Publications

17.8.1 User Publications

Submitted

1. Brown, W. M., Kohlmeyer, A., Plimpton, S. J., and Tharrington, A. N. “Implementing Molecular Dynamics on Hybrid High Performance Computers - Particle-Particle Particle-Mesh,” *Journal of Chemical Physics*, 2011.
2. Q. Zhao, Y. Lam, M. Kheirabadi, C. Xu, K. N. Houk, C.E. Schafmeister. “Hydrophobic Substituent Effects on Proline Catalysis of Aldol Reactions in Water” Journal of Organic Chemistry (2012).
3. J. Nordhaus, T. Brandt, A. Burrows, & A. Almgren. “The Hydrodynamic Origin of Neutron Star Kicks" M.N.R.A.S., 2011 (arXiv:1112.3342)

Accepted

1. Rempel,M.,“Numerical sunspot models: Robustness of photospheric velocity and magnetic field structure”, Astrophysical Journal, 2012.
2. Mittal, R., Erath, B., and Plesniak, M., "Fluid Dynamics of Human Phonation and Speech," Annual Review of Fluid Mechanics. To appear 2013.
3. Bhardwaj, R. and Mittal, R., “Benchmarking a coupled Immersed-Boundary-Finite-Element Solver for large scale flow-induced deformation”, AIAA Journal. To appear, 2012.
4. Zheng, X., Seo, J. H., Vedula, V., Abraham, T., and Mittal, R., “Computational Modeling and Analysis of Intracardiac Flows in Simple Models of the Left Ventricle”, European Journal of Mechanics- B/Fluids. To appear, 2012.
5. A. Burrows, “Multi-Dimensional Core-Collapse Supernova Explosions," the proceedings of the 11th International Symposium on the Origin of Matter and Evolution of Galaxies (OMEG11)), held at RIKEN, Tokyo, Japan on Nov. 14 - Nov. 17, 2011 (AIP Conference series).
6. Dan Willenbring, Lu T. Liu, Yan Xu, Pei Tang, "Isoflurane Alters the Structure and Dynamics of GLIC", 9th Annual Safar Symposium, Pittsburgh, PA 2011 (Awarded: Best Poster)
7. Dan Willenbring, Lu T. Liu, Yan Xu, Pei Tang, "Isoflurane Perturbs Intersubunit Interactions in GLIC", 58th Annual meeting of the Association of University Anesthesiologists, Philadelphia, PA 2011
8. Dan Willenbring, Lu T. Liu, Yan Xu, Pei Tang, "Binding of Isoflurane to GLIC Alters the Structure and Dynamics of the Protein", 55th Annual meeting of the Biophysical Society, Baltimore, MD 2011
9. J. Terasaki and J. Engel, “Testing Skyrme Energy-Density Functionals with the QRPA in Low-Lying States of Rare-Earth Nuclei”, Phys. Rev. C 84 (2011) 014332 (11 pages).
10. Borovikov, S.N., Pogorelov, N.V., Burlaga, L.F., Richardson, J.D., “Plasma near the Heliosheath: Observations and Modeling”, Astrophys. J., 728, L21 (2011)

11. Borovikov, S.N., Pogorelov, N.V. "The Effects of the Sun Rotation on the Solar Wind Propagation", Proceedings of the Numerical Modeling of Space Plasma Flows ASTRONUM-2010, ASP conference series, vol. 444, pp 105-110 (2011).
12. Pogorelov, N. V., Borovikov, S. N., Heerikhuisen, J., Kim, T., Kryukov, I. A.,Zank, G. P., in Proceedings of the Numerical Modeling of Space Plasma Flows ASTRONUM-2010, ASP conference series, vol. 444, pp 130-135 (2011).
13. Heerikhuisen, J., Pogorelov, N.V., An Estimate of the Nearby Interstellar Magnetic Field Using Neutral Atoms Astrophys. J., 738, 29 (2011).
14. Florinski, V., Ferreira, S.E.S., Pogorelov, N.V., Galactic cosmic rays in the outer heliosphere: theory and Models, Space Science Reviews, 03/2011, doi:10.1007/s11214-011-9756-1 (2011).
15. Pogorelov, N.V., Heerikhuisen, J., Zank, G.P., Borovikov, S.N., Frisch, P.C., McComas, D.J., Interstellar Boundary Explorer measurements and magnetic field in the vicinity of the heliopause, Astrophys. J., 742, 104 (2011).
16. Borovikov, S.N., Pogorelov, N.V., Ebert, R.W., Solar rotation effects on the heliosheath flow near solar minima, Astrophys. J., in press (2012)
17. H.B. Guo, A. Johs, J.M. Parks, J.C. Smith, "Mercury Detoxification by Bacteria: Simulations of Transcription Activation and Mercury-Carbon Bond Cleavage, in Modeling of Molecular Properties", Wiley-VCH, in Press.
18. B. Mostofian, J.C. Smith, X. Cheng, "The Solvation Structures of Cellulose Microfibrils in Ionic Liquids", Interdisciplinary Sciences - Computational Life Sciences, in press

Published

1. L. Van Roekel, B. Fox-Kemper, P. P. Sullivan, P. E. Hamlington, and S. R. Haney. The form and orientation of Langmuir cells for misaligned winds and waves. Journal of Geophysical Research-Oceans, 2012.
2. L. Cavaleri, B. Fox-Kemper, and M. Hemer. Wind waves in the coupled climate system. Bulletin of the American Meteorological Society, 2012.
3. A. Webb and B. Fox-Kemper. Wave spectral moments and Stokes drift estimation. Ocean Modelling, 40(3-4):273-288, 2012.
4. B. Fox-Kemper, G. Danabasoglu, R. Ferrari, S. M. Griffies, R. W. Hallberg, M. M. Holland, M. E. Maltrud, S. Peacock, and B. L. Samuels. Parameterization of mixed layer eddies. III: Implementation and impact in global ocean climate simulations. Ocean Modelling, 39:61-78, 2012.
5. W.-S. Cho, M. Luisier, D. Mohata, S. Datta, D. Pawlik, S. L. Rommel, and G. Klimeck, "Full band atomistic modeling of homo-junction InGaAs band-to-band tunneling diodes including band gap narrowing", App. Phys. Lett. 100, 063504 (2012).

17.8.2 Staff Publications

Submitted

1. Simmerman, S., et al. "Exploring Similarities Among Many Species Distributions," IEEE Computer Graphics and Applications Special Issue.

2. Rekepalli, B., et al. “HSPp-BLAST: Highly Scalable Parallel psi-BLAST for Very Large-scale Position Specific Iterative Sequence Searches,” ISCA 4th International Conference on Bioinformatics and Computational Biology (BICoB 2012).
3. Rekepalli, B., et al., “Defining Dark Matter and Grey Matter of the Protein Sequence Space,” Bioinformatics Journal 2012.
4. Rekepalli, B., et al., “The Three Dimensional Reconstruction of Large Micro-CT Data Sets for the Production of Bio-CAD Blueprints of Tissue Structures,” CAD 2012.
5. Huang, J., et al., “Advancing Sustainable Energy Systems Through Computational and Theoretical Chemistry/ Physics,” in Sustainable Energy Systems2011.
6. Kasper Kristensen, Ida-Marie Høyvik, Branislav Jansik, Poul Jørgensen, Thomas Kjærgaard, Simen Reine, and Jacek Jakowski, "Pushing the frontiers of accurate quantum mechanical calculations"submitted to Nature Chemistry

Accepted

1. Patel, P., et al., “OpenMP-style Parallelism in Data-Centered Multicore Computing with R, at the 17th ACM SIGPLAN 2012 Symposium on Principles and Practice of Parallel Programming (PPoPP),” New Orleans, LA, February 2012.
2. J. Meredith, S. Ahern, D. Pugmire, R. Sisneros, ?EAVL: The Extreme-scale Analysis and Visualization Library,? Eurographics Symposium on Parallel Graphics and Visualization (EGPGV), 2012

Published

1. Robert C Ehemann, Predrag S Krstic, Jonny Dadras, Paul R C Kent, Jacek Jakowski, "Detection of hydrogen using graphene", Nanoscale Research Letters 2012, 7:198 doi:10.1186/1556-276X-7-198
2. J. Meredith, R. Sisneros, D. Pugmire, S. Ahern, ?A Distributed Data-Parallel Framework for Analysis and Visualization Algorithm Development,? Fifth Workshop on General Purpose Processing on Graphics Processing Units (GPGPU), ACM Architectural Support for Programming Languages and Operating Systems, 2012

17.9 Metrics

17.9.1 Standard systems metrics

The following subsections contain system metrics for NICS’ resources that are allocated through XSEDE: Kraken and Nautilus.

Note that job wait times and job expansion factors as reported by XDMoD are skewed by user specified job dependencies. NICS has implemented an “effective queue time” metric to eliminate the influence of job dependencies on these statistics. The effective queue time is a measure of the wait time incurred only once a job is eligible to run and is not a factor of individual workflows. In the future job wait times and expansion factors will be reported based on effective queue times.

Another issue with wait time and expansion factor by job size, as currently reported, is that the job size bins overlap multiple scheduling queues at NICS, and thereby, overlap multiple scheduling policies. This too will be corrected in future reporting. Also note that the error bars associated with the mean values in Figure 3, Figure 4, Figure 13 and Figure 14 represent the standard deviation of the sampled mean which is the standard deviation divided by the square root of N, where N is the sample size.

17.9.2 Kraken

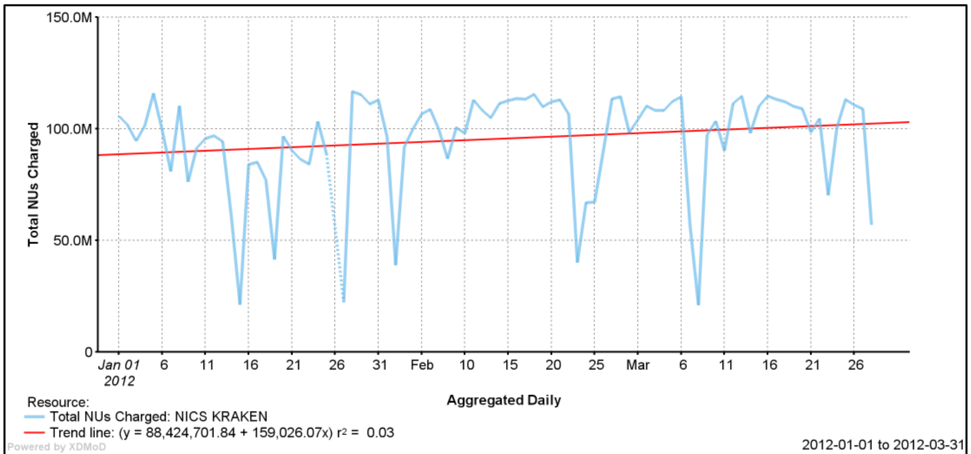

Figure 1: Daily resource consumption in Mega-normalized units (1e^{6}) charged on Kraken for Q3PY1.

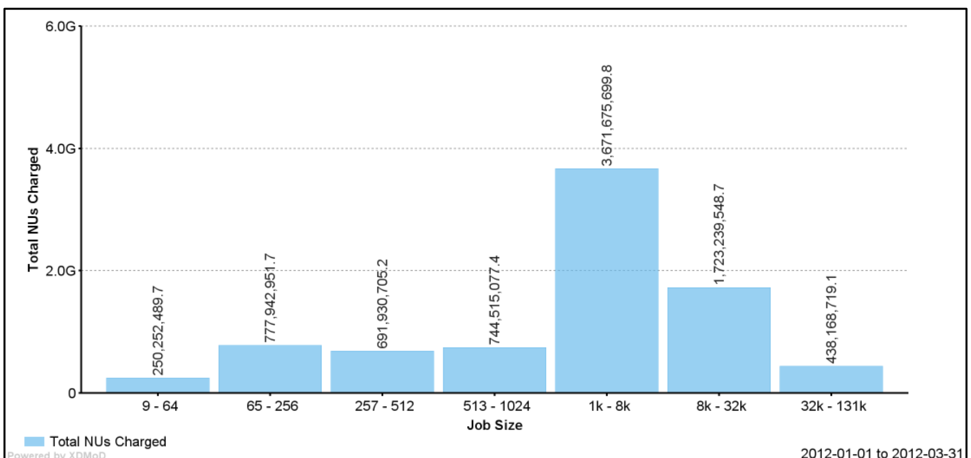

Figure 2: Total resource consumption in Giga-normalized units (1e^{9}) by job size for Kraken in Q3PY1.

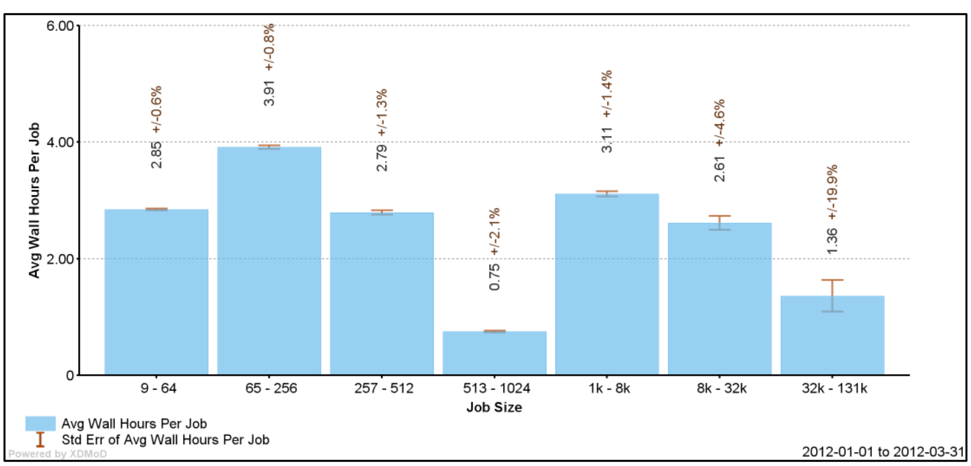

Figure 3: Average wall hours by job size on Kraken in Q3PY1.

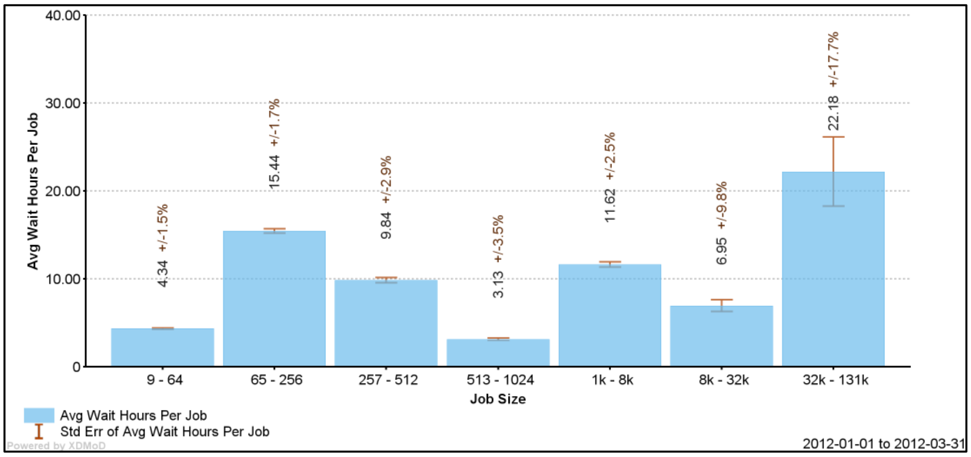

Figure 4: Average wait hours by job size on Kraken in Q3PY1.

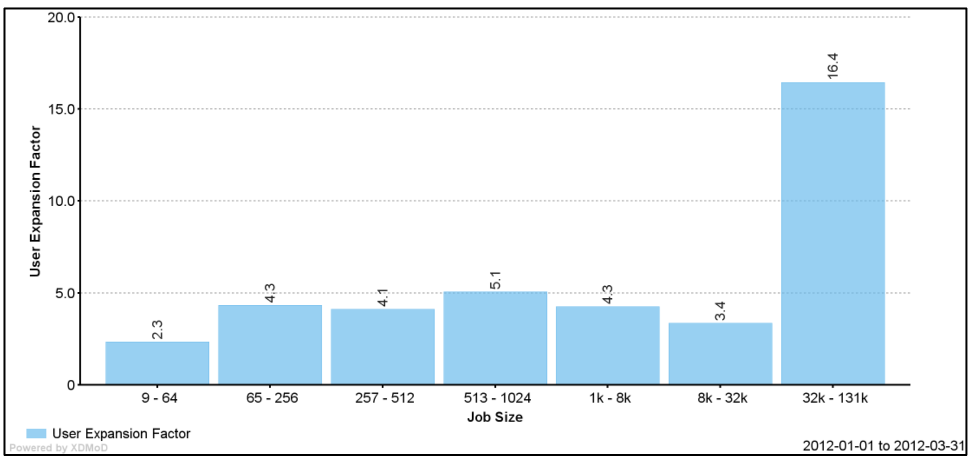

Figure 5: Expansion factor by job size for Kraken in Q3PY1.

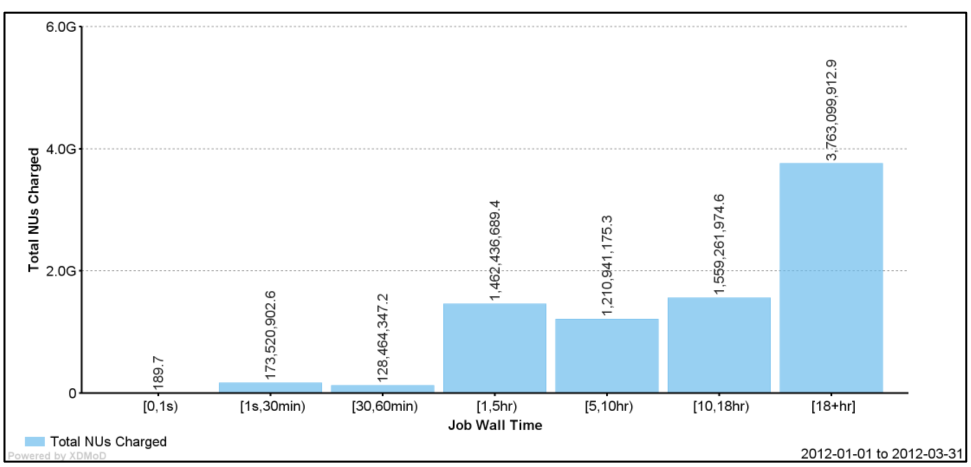

Figure 6: Total resource consumption in Giga-normalized units ($1e^9$) by wall time for Kraken in Q3PY1.

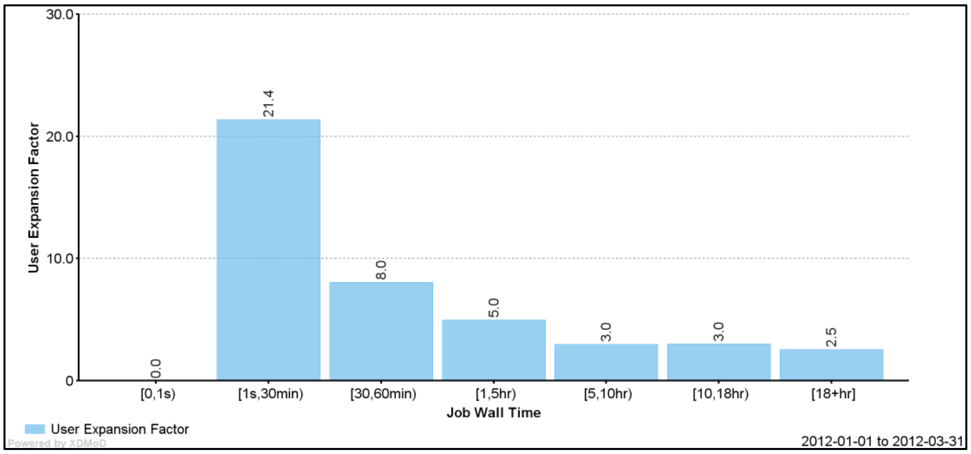

Figure 7: Expansion factor by wall time for Kraken in Q3PY1.

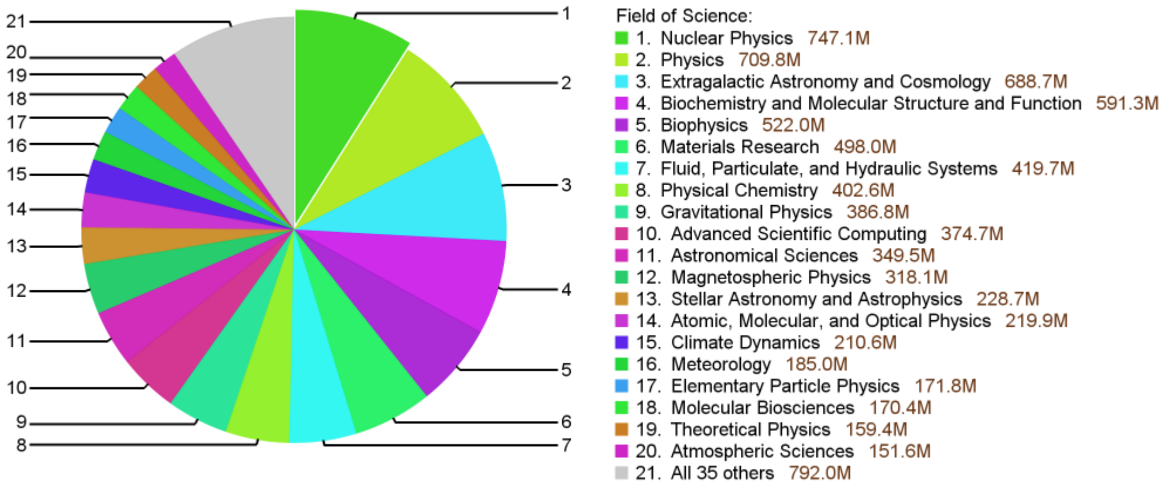

Figure 8: Resource consumption in mega-normalized units ($1e^6$) by field of science for Kraken in Q2PY1.

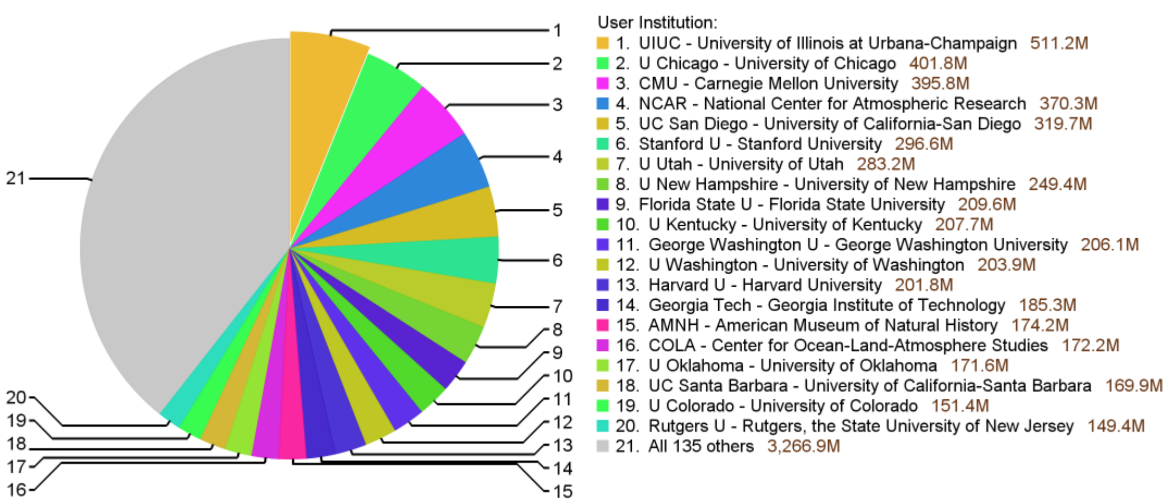

Figure 9: Resource consumption in mega-normalized units ($1e^6$) by institution in mega-normalized units for Kraken in Q3PY1.

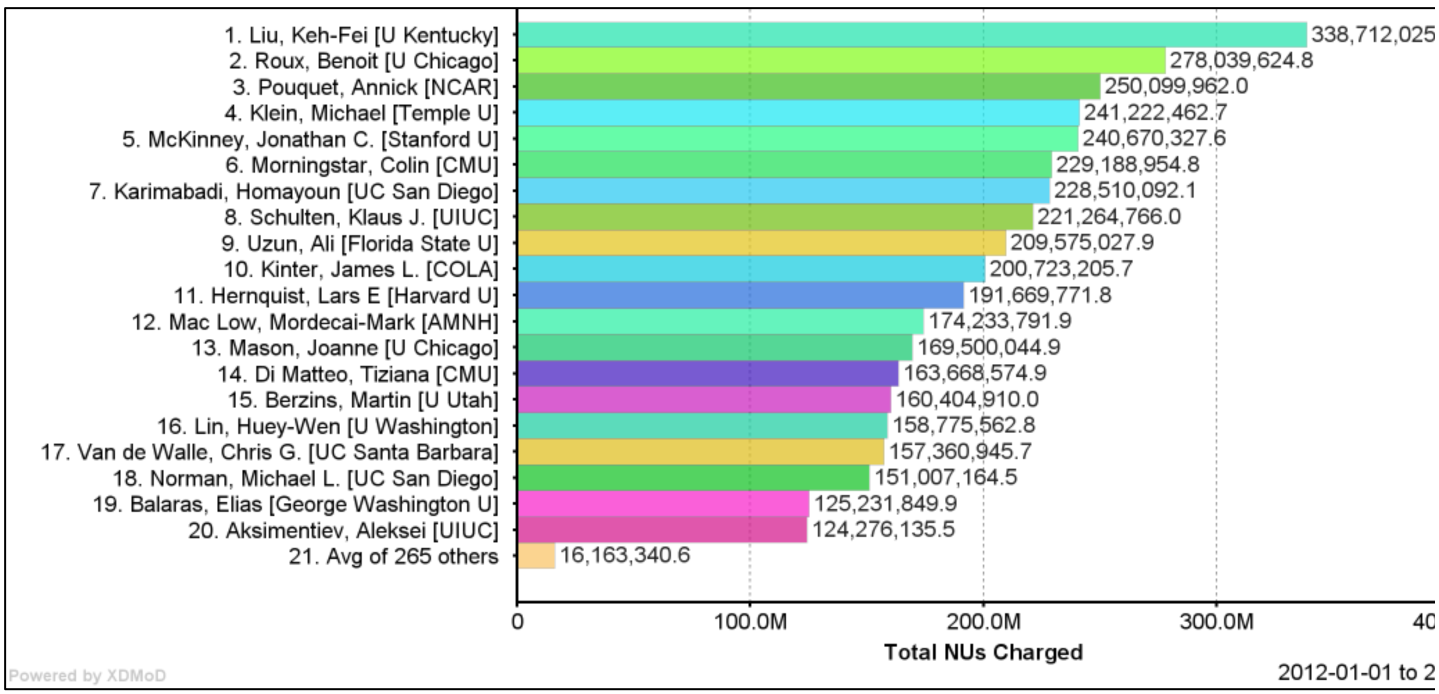

Figure 10: Resource consumption by PI in Mega-normalized units for Kraken in Q3PY1.

17.9.2.1 Nautilus

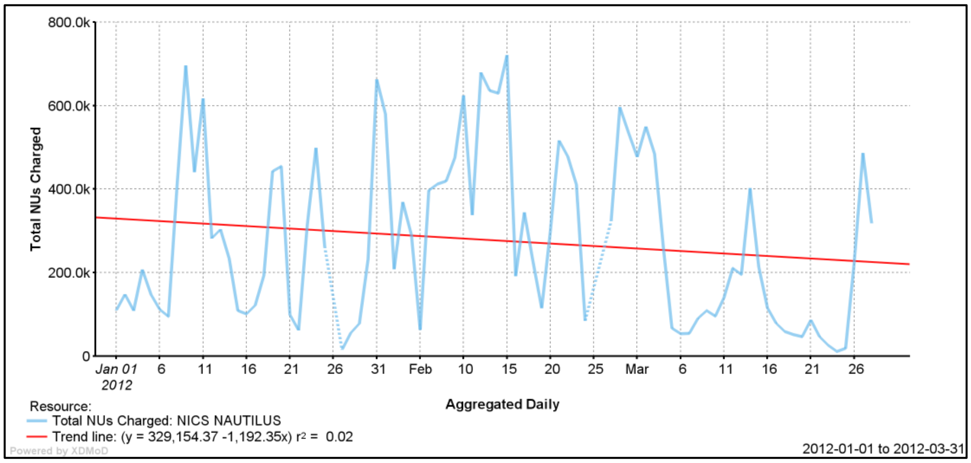

Figure 11: Daily resource consumption in kilo-normalized units ($1e^3$) charged on Nautilus in Q3PY1.

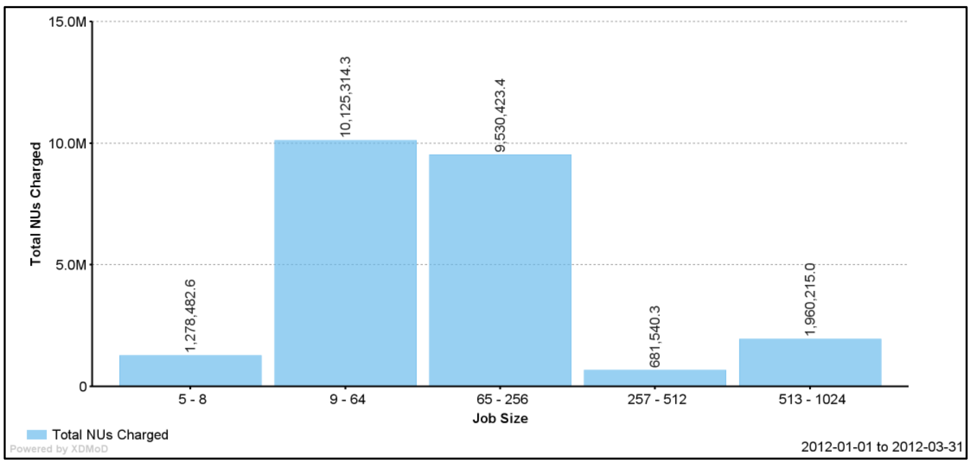

Figure 12: Total resource consumption in mega-normalized units ($1e^6$) by job size for Nautilus in Q3PY1.

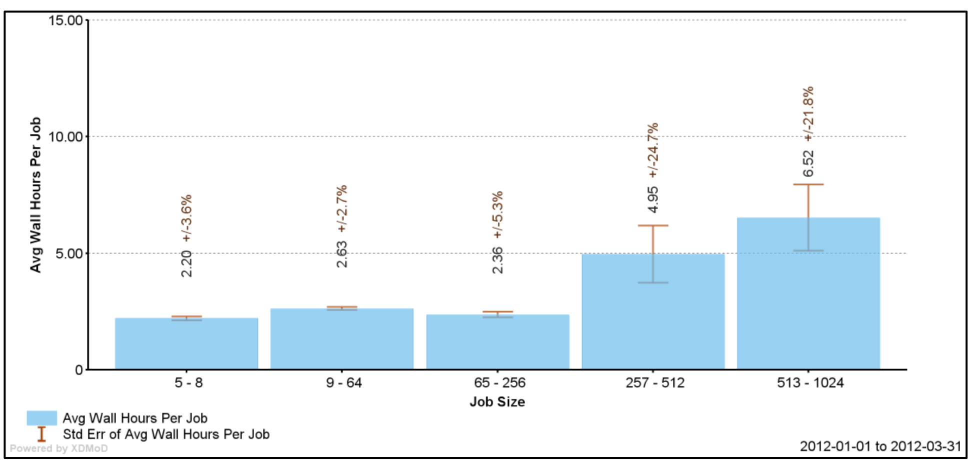

Figure 13: Average wall time in hours by job size on Nautilus in Q3PY1.

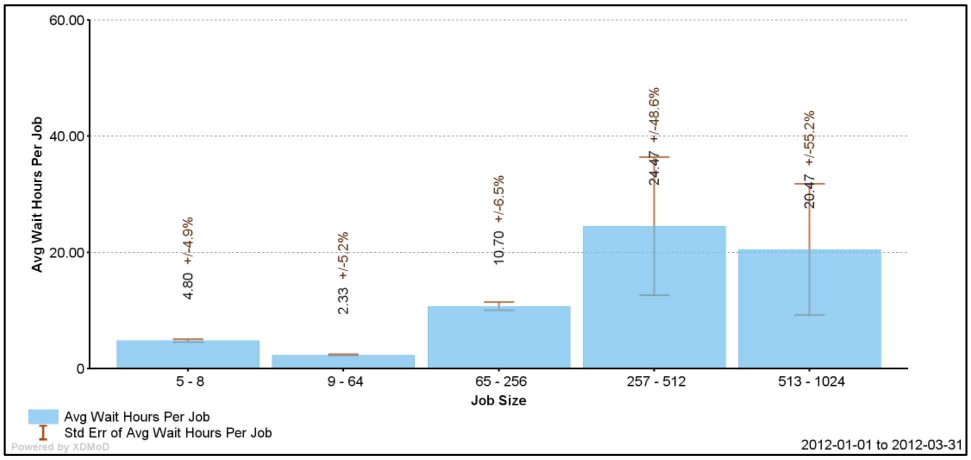

Figure 14: Average wait time in hours by job size for Nautilus in Q3PY1.

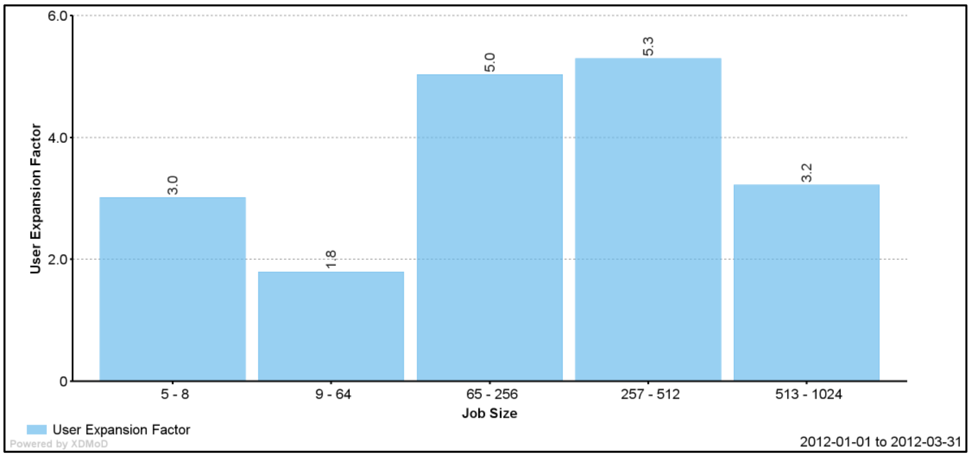

Figure 15: Expansion factor by job size for Nautilus in Q3PY1.

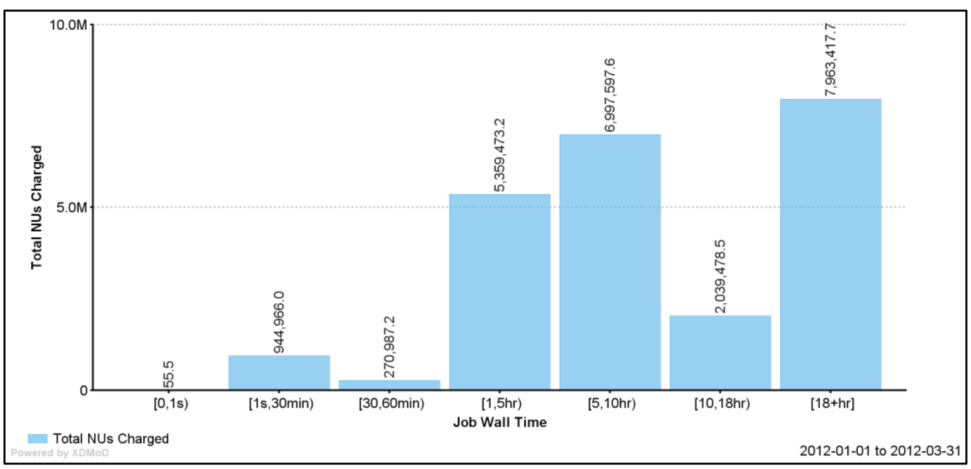

Figure 16: Total resource consumption in mega-normalized units (1e6) by wall time for Nautilus in Q3PY1.

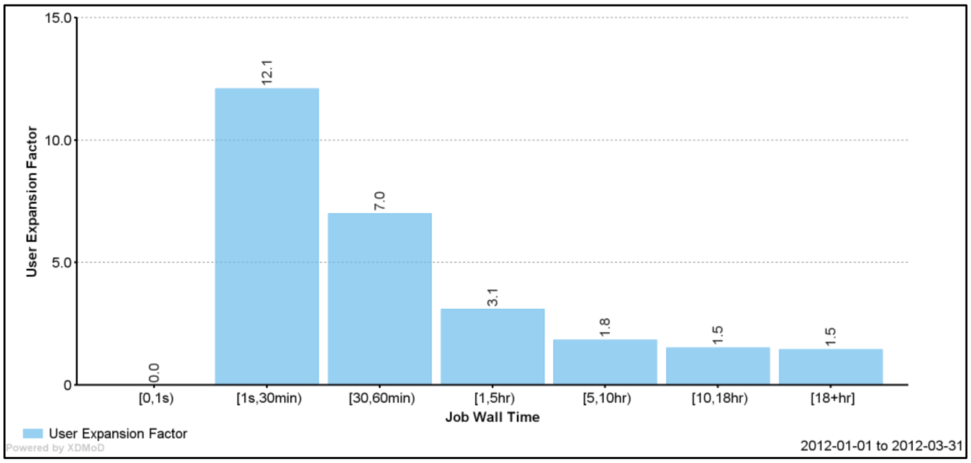

Figure 17: Expansion factor by wall time for Nautilus in Q3PY1.

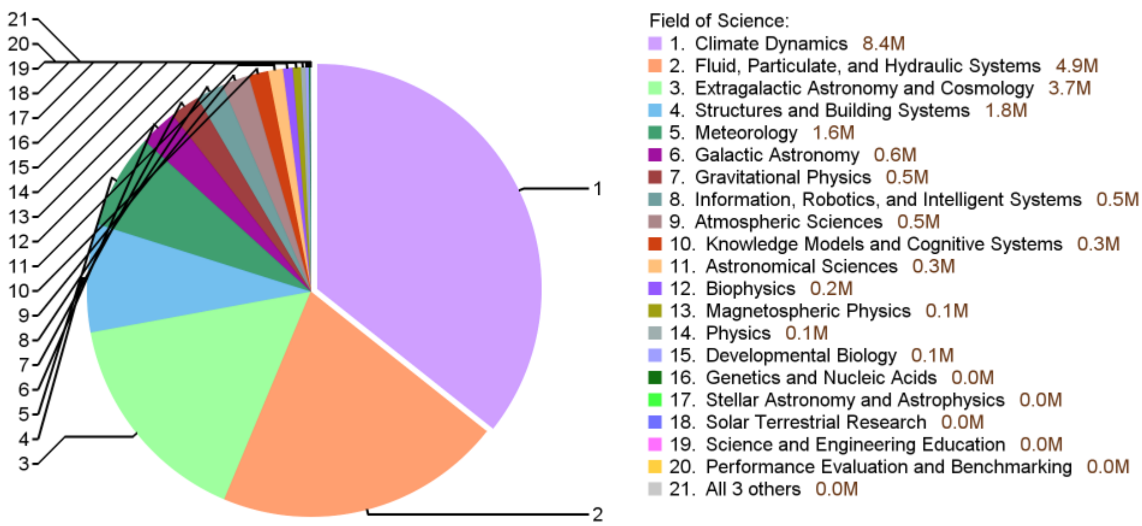

Figure 18: Resource consumption by scientific domain in mega-normalized units for Nautilus in Q3PY1.

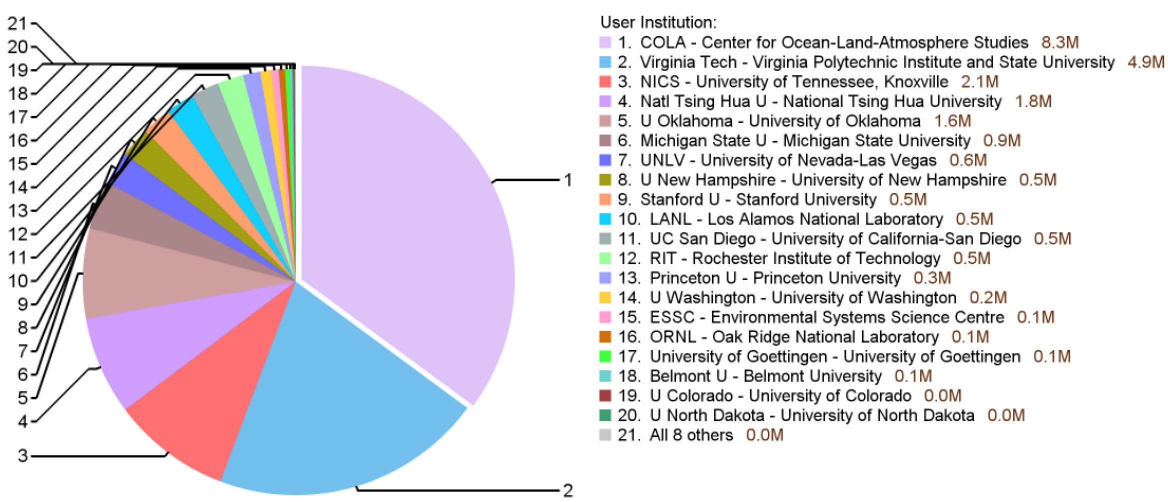

Figure 19: Resource consumption by institute in mega-normalized units for Nautilus in Q3PY1.

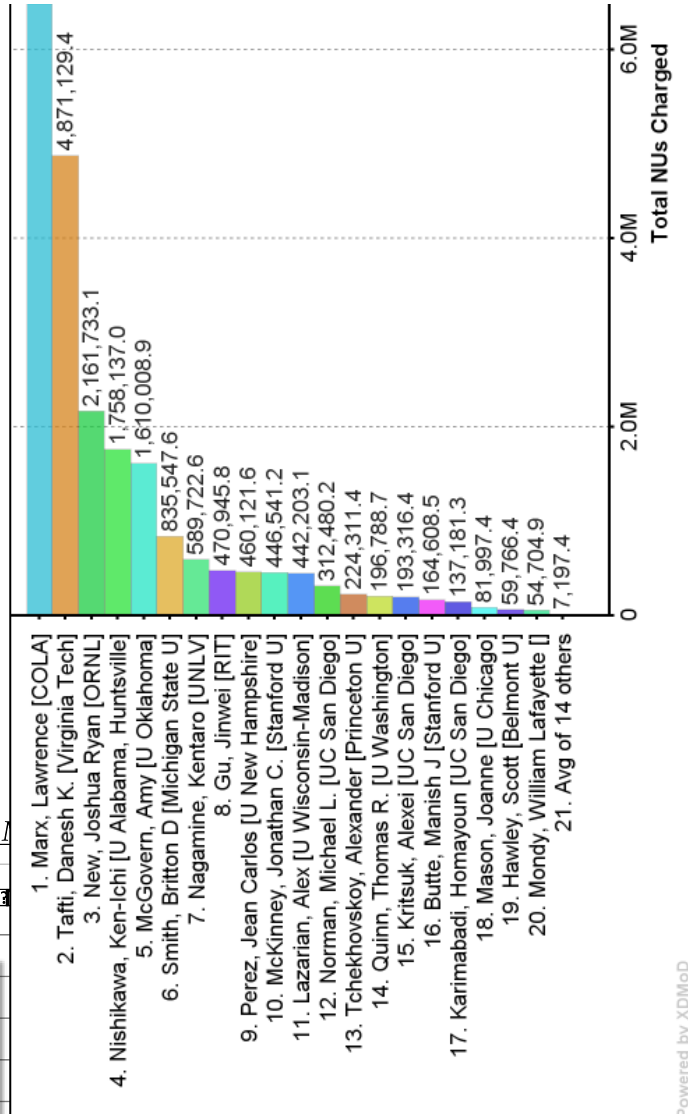

Figure 20: Resource consumption by PI in mega-normalized units for Nautilus in Q3PY1

17.9.3 *Standard User Assistance I*

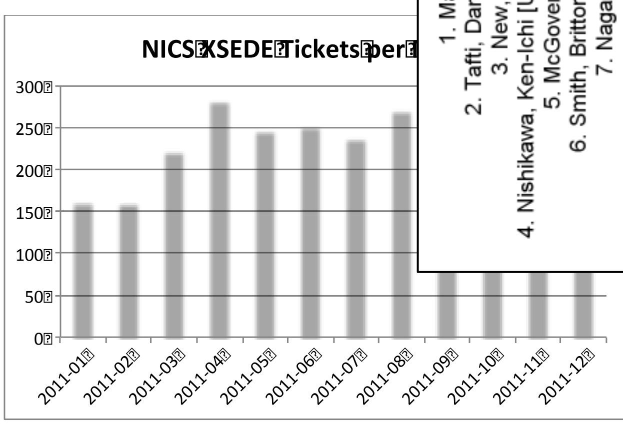

Figure 21: New XSEDE tickets opened by month.

ed

ese

ues

r the quarter with most tickets being resolved within 24 hours.

Table 5: Ticket resolution times by category for Q3PY1.

	account issues	filesystems	grid software	jobs / batch queues	login / access issues	mss / data issues	network issues	other	software / apps	system issues	Grand Total
0-1 Hour	1	3	1	8	42	-	1	2	1	4	**63**
1-24 Hours	5	16	1	54	85	3	-	8	18	3	**193**
1-7 Days	15	10	3	65	58	16	2	5	44	4	**222**
1-2 Weeks	1	2	1	17	23	5	-	1	21	5	**76**
> 2 Weeks	-	2	-	33	20	9	2	1	31	2	**100**
Still Open	-	-	1	3	1	3	-	-	3	2	**13**
Grand Total	**22**	**33**	**7**	**180**	**229**	**36**	**5**	**17**	**118**	**20**	**667**

17.9.4 SP-specific Metrics

NICS' resources provided roughly 60% of computational cycles that were delivered to the NSF community in this quarter (Figure 23), and NSF charges accounts for 98% (**Error! Reference source not found.**) of the total charges on these resources.

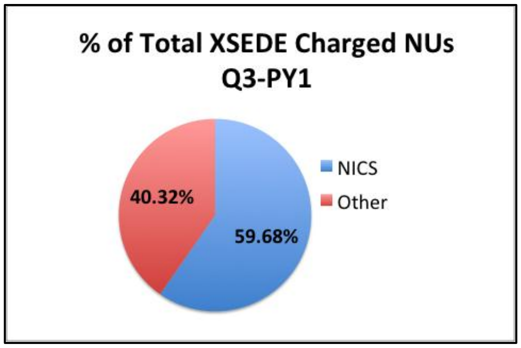

Figure 22: NICS as a percentage of total charged Giga-normalized units (1e9) for Q3PY1.

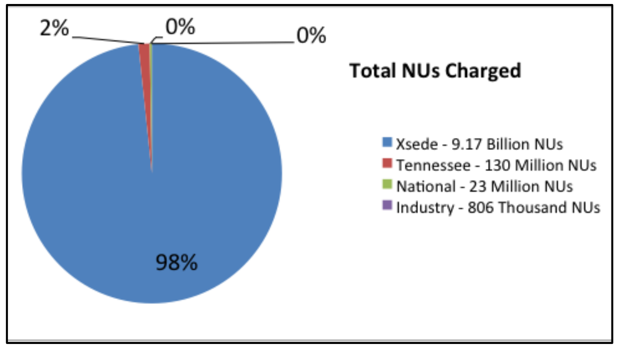

Figure 23: XSEDE charges as a percentage of total charges on NICS resources in Q3PY1.

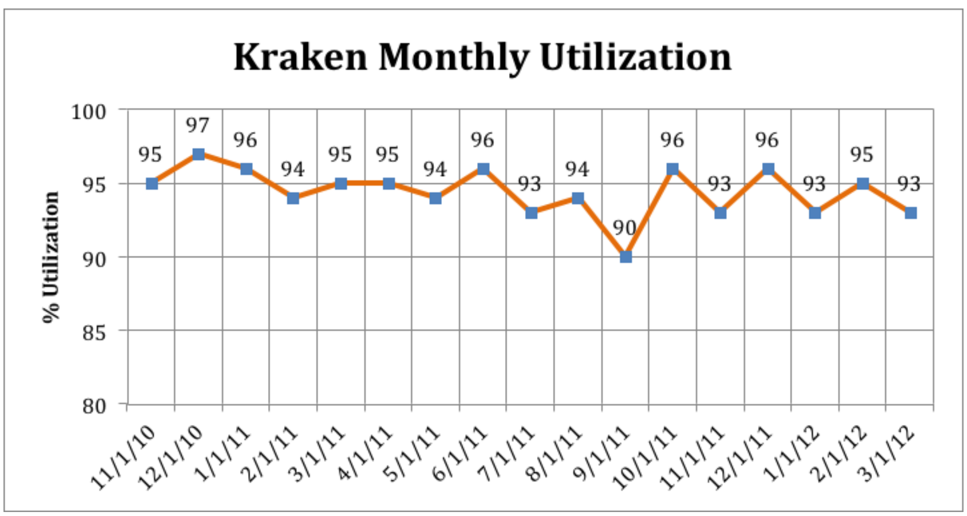

Figure 24: Total monthly utilization for Kraken for the past year.

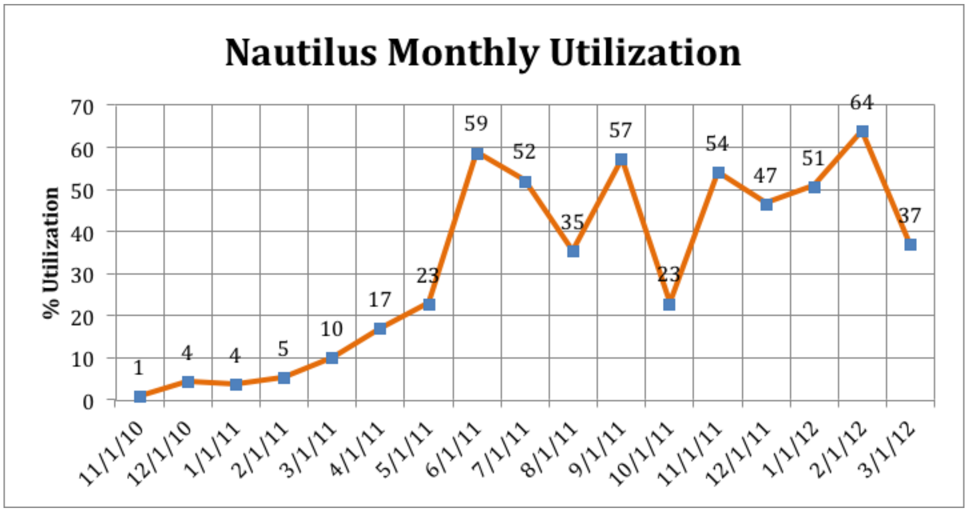

Figure 25: Total monthly utilization for Nautilus for the past year.

18 Pittsburgh Supercomputing Center - Service Provider Quarterly Report

18.1 Executive Summary

The Pittsburgh Supercomputing Center (PSC) operates and supports two powerful and unique resources for the national research community. *Blacklight,* an SGI Altix UV 1000 acquired with the assistance of an NSF grant and operated as an XSEDE resource, is the world's largest shared-memory system, providing two partitions of 16TB each. *Anton* is a special-purpose computer for molecular dynamics calculations, performing up to 100 times faster than conventional supercomputers. It is on loan, at no cost, from its designer, D.E. Shaw Research (DESRES) with operational funding provided by NIH. It is available to the national community through a peer-review committee convened by the National Research Council. It is the only such system outside of DESRES. These systems, as well as other smaller clusters, are supported at PSC by a central file system, archival storage and extensive LAN, MAN and WAN infrastructure. PSC has completed the transition from its disk-tape archive system to a new disk-only archive called *arc* that provides much faster access to files in the archive. *Arc's* building-block architecture enables *arc* to scale well beyond its initial deployment of four petabytes.

In addition to scientific achievements presented in the XSEDE report, PSC resources enabled significant progress in the study of how varying preparedness levels for dealing with pandemics among the states affects the United States as a whole.

In networking, PSC staff worked with the Galaxy Genomics Portal at Penn State to enable connectivity between Galaxy and XSEDE via PSC's Three Rivers Optical Exchange (3ROX) so that Galaxy can use XSEDE as a backend. PSC staff also worked with Penn State to tune the network transfers over this new connection and helped them with their perfSONAR platform. 3ROX also established a 10-Gbps link with Drexel University. This 3ROX/Drexel partnership provides network services to a wide range of sites in western Pennsylvania and West Virginia along with Drexel and its affiliated research sites and the 14 Pennsylvania State System of Higher Education universities.

Users continue to find *Blacklight's* large shared memory to be very valuable, particularly in genomics. For instance, Rebecca Duncan of the University of Miami said, "My research in functional and comparative genomics involves a lot of [memory intensive] computation … Although my university has a high performance computer, it did not meet the memory requirements needed for some steps of my assembly pipeline. *Blacklight* enabled me to overcome the difficulties …" Users also continue to praise PSC's well-regarded user support functions.

PSC engaged in a broad range of Education, Outreach and Training activities. These included a training workshop at Lehigh University, a workshop to encourage girls to study STEM, development of a short course on phylogenetics, continuation of an international student intern program, delivery of professional development sessions for science teachers at a local high school, mentoring a high school robotics team, and participation on an advisory committee for a project at Tennessee State University.

PSC collaborates with two groups at Harvard University on studying the wiring diagram of the mouse visual cortex. The Harvard teams obtain high resolution images of extremely thins sections of the visual cortex, and the PSC team assembles the images for viewing and analysis. Among the issues that must be handled is the transmission of the high volumes of data from Harvard to PSC over the Internet.

18.1.1 Resource Description

PSC provides a range of computing and storage platforms for the national science community.

For applications requiring very large shared memory, high-productivity programming models, and/or moderate parallelism with a high-performance system-wide interconnect, PSC operates *Blacklight*, an SGI UV 1000 cc-NUMA shared-memory system comprising 256 blades. Each blade shares 128GB of local memory, and holds two Intel Xeon X7560 (Nehalem) eight-core processors, for a total of 4096 cores and 32 TB across the whole system. Each core has a clock rate of 2.27 GHz, supports two hardware threads and can perform nine Gflop/s for a total system floating point capability of 37 Tflop/s. *Up to 16 TB of this memory is accessible as a single memory space to a shared-memory program.* Message-passing and PGAS programs can access all 32 TB on the system. *Blacklight* is part of the National Science Foundation XSEDE integrated national system of cyberinfrastructure.

Additionally, PSC has an SGI Altix 4700 systems called *Salk*, smaller than *Blacklight*, that is also targeted at applications requiring large shared memory, high-productivity programming models, or moderate parallelism with a high-performance, system-wide interconnect. *Salk* is administered for the NIH-funded National Resource for Biomedical Supercomputing (NRBSC) and offers 144 Intel Itanium Montvale processors providing a peak aggregate speed of 0.96 Tflop/s with 288 GB shared memory. This system supports advanced programming languages and models including UPC and Star-P.

PSC operates an *Anton* special-purpose supercomputer for molecular dynamics (MD) simulation. Designed by D. E. Shaw Research (DESRES) and provided to PSC without cost by DESRES, it is available for non-commercial research use by universities and other non-profit institutions. The *Anton* is hosted by PSC and available to the national biomedical community with funding from NIH's National Institute of General Medical Sciences. This is the only *Anton* computer operated outside DESRES. The *Anton* computer is supplemented by a high performance file storage system for simulation trajectories and an analysis cluster (*Kollman*). Each of the four nodes in the analysis cluster consists of two Intel Westmere six-core processors and 96 GB of memory. The high-performance file storage system consists of a 500-TB Lustre file system. The file system and the analysis cluster nodes are interconnected over Quad Data Rate (QDR) InfiniBand.

The originally-planned duration of production use of *Anton* at NRBSC/PSC was twelve months. However, because of the overwhelming scientific and operational success during the initial period, NRBSC/PSC and DESRES agreed to extend the project for another nine months beginning on October 4, 2011 and lasting until June 2012.

PSC operates several Linux clusters for biomedical and computer science research as well as several high-end servers and powerful workstations for development, analysis, and visualization tasks.

The production workload on all of the PSC computing platforms is managed by PBS/Torque. Several scheduler policy modules used include a locally-developed module, *Simon*, and the Maui scheduler.

All of the PSC computing platforms have access to *Brashear,* PSC's shared, central file system using the Lustre file system architecture. It comprises eight storage nodes and 350 TB of direct-attached disks, forming a large I/O cluster globally accessible within the PSC site. Access to the file system is provided by InfiniBand, 10-Gigabit Ethernet and 1-Gigabit Ethernet. Each node in the I/O cluster is a Lustre Object Storage Server (OSS) hosting multiple Object Storage Targets (OSTs).

PSC is a partner in the Lustre *Albedo* Wide Area File System project along with five other XSEDE sites. PSC is taking the lead by managing the metadata service for the *Albedo* file system in addition to providing a portion of its bulk object storage. In addition to *Albedo*, PSC has made available a test file system running pre-release Lustre software incorporating advanced security with Kerberos authentication. *Albedo* is mounted on *Blacklight*.

To support the file repository needs of its users, PSC has deployed a file archival system named *arc*. *Arc* is a disk-based storage system designed with dual strategies of long term archiving and rapid access to data. Each building block in *arc* has one petabyte of useable disk storage, which is managed by the ZFS file system and the PSC-developed SLASH2 replicating distributed file system. ZFS and SLASH2 provide multiple layers of robust data integrity checking to protect user data against data corruption. This building-block architecture will enable *arc* to scale well beyond its initial deployment of four petabytes.

Users can access *arc* from within PSC using the familiar PSC file archiving utility, *far*. From outside PSC, users can employ a variety of well known file transfer methods such as scp and gridftp. These transfers are handled by a series of dedicated data transfer servers, which are collectively addressed as data.psc.edu.

PSC network facilities consist of production and research LAN, MAN, and WAN infrastructures. The network technology used for the LAN infrastructure consists of switched Ethernet ranging in speeds from 100 Mb/s to 10 Gb/s. The PSC MAN infrastructure is DWDM-based, supporting multiple 10-Gigabit Ethernet waves. Network connections include a 10-Gigabit Ethernet connection to the XSEDE backbone network via National LambdaRail, a 10-Gigabit Ethernet connection to National LambdaRail's FrameNet service, a nationwide Layer2 network, and a 10-Gigabit Ethernet connection between PSC's offices and its remote supercomputing machine room at 4350 Northern Pike (formerly known as Westinghouse Energy Center).

WAN connections are provided by the Three Rivers Optical Exchange (3ROX), a regional network aggregation point that provides high-speed commodity and research network access, primarily to sites in Western and Central Pennsylvania and West Virginia. 3ROX is based at Carnegie Mellon University and is operated and managed by the Pittsburgh Supercomputing Center. While the primary focus of 3ROX is to provide cost-effective, high-capacity, state-of-the-art network connectivity to the university community, this infrastructure also provides well-defined network services to both community (K-12, government) and commercial entities in Western Pennsylvania. University member sites currently include Carnegie Mellon University, the Pennsylvania State University, the Pittsburgh Supercomputing Center, the University of Pittsburgh, Norfolk State University, and West Virginia University. The 3ROX commodity Internet component consists of multiple high-performance WAN connections to major Internet service providers including Gigabit Ethernet connections to Cogent and Global Crossing. The research component of 3ROX includes a 10-Gigabit Ethernet connection, with 5 Gb/s of bandwidth, to the Internet2 network. In addition to the Internet2 connection, 3ROX also has a 10-Gigabit Ethernet connection to National LambdaRail's PacketNet service, a nationwide routed IP network. Explicit routing is used to maintain the acceptable use policies associated with the various production and research network infrastructures.

18.2 Science Highlights

In addition to major science accomplishments that are highlighted in the XSEDE report, we present selected others specific to PSC.

18.2.1 *Systematic and Population Biology: Preparedness for Pandemics: Does Variation Among States Affect the Nation as a Whole? (Margaret A. Potter and Shawn T. Brown, University of Pittsburgh - Graduate School of Public Health)*

Since states' public health systems differ as to pandemic preparedness, this study explored whether such heterogeneity among states could affect the nation's overall influenza rate. The Centers for Disease Control and Prevention produced a uniform set of scores on a 100-point scale from its 2008 national evaluation of state preparedness to distribute materiel from the Strategic National Stockpile (SNS). This study used these SNS scores to represent each state's relative preparedness to distribute influenza vaccine in a timely manner and assumed that "optimal" vaccine distribution would reach at least 35% of the state's population within 4 weeks. The scores were used to determine the timing of vaccine distribution for each state: each 10-point decrement of score below 90 added an additional delay increment (varied between 1 to 10 days) to the

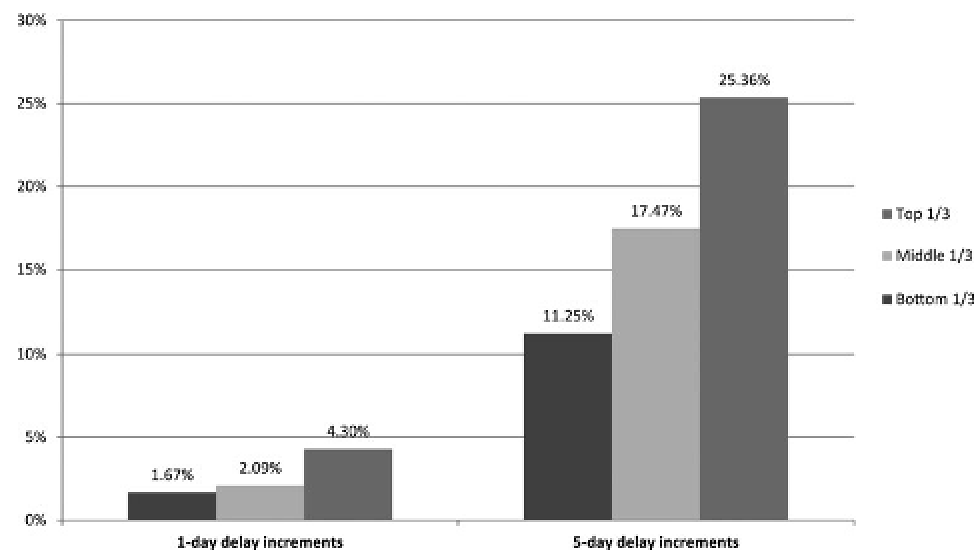

Comparison of Percent-Change in Influenza Attack Rates from Optimal Scenario to SNS-Modified Scenario among States Grouped by SNS Score

distribution time. A large-scale agent-based computational model simulated an influenza pandemic in the US population. In this synthetic population each individual or agent had an assigned household, age, workplace or school destination, daily commute, and domestic intercity air travel patterns. Simulations compared influenza case rates both nationally and at the state level under 3 scenarios: no vaccine distribution (baseline), optimal vaccine distribution in all states, and vaccine distribution time modified according to state-specific SNS score. Between optimal and SNS-modified scenarios, attack rates rose not only in low-scoring states but also in high-scoring states, demonstrating an interstate spread of infections. Influenza rates were sensitive to variation of the SNS-modified scenario (delay increments of 1 day versus 5 days), but the interstate effect remained (see figure).

This study required the use of a nation-wide agent-based model originally developed by Neil Ferguson at Imperial College and modified by Shawn Brown (U Pitt / PSC) for the study and to be run on PSC resources. The model has quite an intensive shared memory requirement (74 GB of RAM per run) with 10 independent runs being needed to run simultaneously for each experiment, and so large shared-memory resources of *Blacklight* were instrumental in completion of the work. The study represents how real world measurements of emergency preparedness can be used in pandemic planning. The work was funded by the University of Pittsburgh NIH MIDAS National Center of Excellence and the University of Pittsburgh Centers for Disease Control Public Health Adaptive Systems Study.

18.3 User-facing Activities

18.3.1 System Activities

HPN-SSH: FreeBSD announced the inclusion of the PSC-developed High Performance SSH (HPN-SSH) as one of the highlighted features included in its FreeBSD 9.0 release. See http://www.freebsd.org/releases/9.0R/announce.html for a list of additional features.

10-Gbps Link with Drexel University: An announcement of a 10-Gbps link between PSC's Three Rivers Optical Exchange and Drexel University was released on January 30. This 3ROX/Drexel partnership provides network services to universities, research sites and K-12 schools in western Pennsylvania and West Virginia along with Drexel and its affiliated research sites and the 14 Pennsylvania State System of Higher Education universities. It also allows for further connectivity with the KINBER/PennREN cross-state network that is currently being constructed. See http://www.psc.edu/publicinfo/news/2012/013112_3rox_drexel.php for additional information.

18.3.2 Services Activities

Galaxy Genomics Portal: Galaxy, the open, web-based platform for data intensive biomedical research at Pennsylvania State University, has a project with XSEDE to use XSEDE as a backend. This past quarter, in order to facilitate connectivity between Galaxy and XSEDE-based resources, a direct connection between Penn State and the XSEDE backbone was completed via PSC's Three Rivers Optical Exchange, allowing this ECSS project to go forward. As part of the network connection, PSC's Joe Lappa and Kathy Benninger worked with Penn State to tune the network transfers over this new connection and helped them get the Penn State perfSONAR platform working with the XSEDE perfSONAR platform.

In addition, Phil Blood of PSC has been involved from the start in developing this project, and we added Josephine Palencia to the team to provide WAN data management expertise.

Large Memory Usage on *Blacklight*: In previous quarterly reports, PSC has reported user comments on how they have effectively used *Blacklight's* large shared memory in their research, particularly in genomics applications. PSC has received a number of additional comments from users about *Blacklight's* usefulness, a sampling of which follows:

- Ronald Worthington (Southern Illinois University, Edwardsville): My research goal is to discover new, rare genetic variations in individual genomes, and that requires fast computing nodes and large memory resources, since the individual genetic data files are of gigabyte size. My campus does not have a computing cluster available to all faculty, and my lab does not have workstations with memory above 8 GB. So I am unable to use the open source software needed for my genomics studies without *Blacklight*.
- Brendan O'Connor (Carnegie Mellon University): I've been using *Blacklight* to train large-scale latent semantic models of natural language. It is well-known that one needs large amounts of data to do this effectively -- millions to billions of words at the very

least. However, most researchers using large textual datasets tend to use very simple algorithms since the computational burden becomes high – I've done this myself in previous work. By contrast, by using *Blacklight*, we can use dozens to hundreds of cores to parallelize our algorithms (I'm using MPI/Python/C) and therefore take advantage of much more powerful Bayesian latent variable models that are rarely applied to these large-scale datasets. It is exciting!

- Kevin Guise**,** Mount Sinai School of Medicine: "Advances in high density neuronal recordings increase the need for access to sufficiently powerful computational resources for the analysis of neuronal signals … our current analysis required simulations of processes that would have taken weeks - or even months - to complete using the resources available to us. Use of *Blacklight* allows us to perform these analyses in days and therefore broadens the range of hypotheses we can reasonably test regarding how different regions of the brain interact with each other in order to support memory-guided behaviors."

- Panagiotis Foteinos, Old Dominion University: "Image to mesh conversion (I2M) is a fundamental step for finite element simulations in many (bio-)engineering applications and for the visualization of complex multi-labeled medical images. Advances in medical imaging have rendered the acquisition of large high resolution possible. *Blacklight* gives us the opportunity to fine tune our parallel I2M algorithm on a large number of cores, and hence be able to produce high resolution meshes within seconds …"

- Mark Lawson, U. of Virginia: "Recently our group has taken on a lot of projects that require *de nov*o transcriptome assembly which is a process that requires a large amount of memory [but] … our local resources are not sufficient. So we were very happy to see an implementation of Trinity on *Blacklight* as it allowed us to finally assemble these reads."

- Rebecca Duncan, U. of Miami: "My research in functional and comparative genomics involves a lot of [memory intensive] computation, including *de novo* transcriptome assembly without a reference genome … Although my university has a high performance computer, it did not meet the memory requirements needed for some steps of my assembly pipeline. *Blacklight* enabled me to overcome the difficulties …"

- Matthew Woitaszek, NCAR, who has an educational allocation for a course on HPC that he teaches at the University of Colorado Boulder: "From the instructional perspective, *Blacklight* was great! The SGI UV series is wonderful for exploring shared-memory programming at larger scales than the 16-core node."

GenomeWeb: *Blacklight* has been highlighted on GenomeWeb. The lead sentence of the article says, "When it comes to high-performance computing systems suited to genomics, researchers might be hard pressed to do better than Pittsburgh Supercomputing Center's *Blacklight* supercomputer system." The article says that large shared-memory architectures like *Blacklight* might be the way forward for genomic data analysis on high-performance computing systems.

For the full article, see http://psc.edu/publicinfo/pdf/GenomeWeb_BioInformatics_021312.pdf.

As reported previously, Matthew Dublin of GenomeWeb had posted a "Highlight from SC11" that contains an interview with PSC's Phil Blood. In the interview Blood discusses *Blacklight* and the value of *Blacklight's* large shared memory in genome assemblies. (See http://www.genomeweb.com//node/1005326?hq_e=el&hq_m=1167682&hq_l=4&hq_v=92dcb774c6.)

GenomeWeb LLC is an independent, privately-held online and print publisher based in New York. GenomeWeb's editorial mission is to serve readers with exclusive, in-depth coverage of the technology, institutions, and scientists that make up the worldwide research enterprise of molecular biology.

18.4 Security

PSC had no security incidents during the reporting period.

As the use of devices such as mobile phones and tablets becomes the daily norm, PSC staff members are increasingly using these mobile devices for access to PSC data and resources. Therefore, PSC's security team is developing security tips and recommendations that one can use for personal mobile devices. The first guide of security tips for Android devices has been completed. PSC also plans on having a similar guide for Apple's iOS devices.

18.5 Education, Outreach, and Training Activities

International Student Intern Program: Rodrigo (Rod) Sellanes was the newest participant in PSC's highly successful international collaboration for student internships with UCEMA (a university in Argentina <http://www.ucema.edu.ar/foreign-students/about-ucema>). Sellanes worked in PSC's Systems & Operations group during January and February on PSC's new file archive project. In particular he looked at system log message collection and reporting.

Girls Math/Science Partnership Program - Tour Your Future: On Saturday, 28-Jan-2012, PSC hosted 10 middle school and high school students for an afternoon of learning about the role supercomputing plays in science and the skills and technology these resources require. The speakers for the session were Kathy Benninger, Marcela Madrid, and Laura McGinnis, with Deb Nigra helping shepherd the group, especially during the hands-on activities.

In the first session, the girls received an introduction to PSC and then learned about what computational thinking is and why it's a critical skill for the future workforce. Next they did a hands-on exercise that demonstrated fundamental ideas of supercomputing, such as serial versus parallel processing, latency, collision management and bottlenecks.

The second session was an introduction to the systems and networking that make science happen. Benninger provided an introduction to the overall facilities that PSC provides and gave a tour of the on-site machine room, which showed a good sample of the hardware involved.

Finally, the girls met with Madrid, who talked about her professional journey from graduate student working with NMR to computational scientist. Using hemoglobin as an example, she then showed the girls how to load the structural data from the Protein Data Bank and manipulate it to see how its structure affects drug design.

A lot of material was covered in the two-hour session, which gave the girls and the staff a very positive sense of the future of science and technology.

Teacher Development: PSC and Oakland Catholic High School (OCHS) received a grant from the Allegheny Intermediate Unit to conduct professional development workshops for high school STEM teachers at OCHS. Pallavi Ishwad of PSC conducts the workshops which cover BEST and CMIST.

BEST (Better Educators of Science for Tomorrow) is a program of PSC's NIH-funded National Resource for Biomedical Supercomputing (NRBSC) that prepares teachers to refocus their teaching strategies towards exposing and encouraging students to become aware of emerging and exciting biomedical careers. It provides a high school level bioinformatics curriculum and ongoing support for teachers. CMIST (Computational Modules in Science Teaching) brings innovative science tutorials into secondary school classrooms, focusing on integrative computational biology, physiology, and biophysics. In contrast to other teaching tools, CMIST modules include high quality, biologically realistic 3-D animations produced with cutting-edge simulation and visualization software developed at NRBSC.

The schedule for delivered and planned workshops at OCHS follows:

- First workshop, 12_Dec-2011, Introduction to CMIST and BEST
- Second workshop, 20-Jan-2012, CMIST - Molecular Transport in Cells
- Third workshop, 08-Feb-2012, CMIST - Enzymes - Structure and Function
- Fourth workshop, 15-Mar-2012, First Semester of a high school curriculum in Bioinformatics
- Fifth workshop, 25-Apr-2012, Second Semester of a high school curriculum in Bioinformatics

Short Course on Phylogenetics: Alex Ropelewski and Hugh Nicholas worked with Dr. Dannie Durand of Carnegie Mellon University in helping her prepare for teaching a CMU-accredited short course on phylogenetics for molecular biologists. They supplied a well understood and already tested dataset for use in the course as well as a video lecture about interpreting multiple sequence alignment results from several popular programs. In return Durand created and gave to them a set of homework exercises based on the dataset they supplied. She also provided them with a responsive and questioning audience of her graduate students in front of whom to film the video lecture on interpreting multiple sequence alignment results. The presence of an active, questioning audience greatly improved the quality of the video lecture.

Mentoring Robotics Team: PSC's Chad Vizino mentors a robotics team called the Terabytes. Recently, the team won the Inspire Award at the Ohio state championship, qualifying it for the FIRST Tech Challenge World Championship from April 25-28 in St. Louis, one of 128 teams to compete. (No championship was held in the Pittsburgh area, allowing the team to compete in Ohio.) More information is given in the news article listed below:
http://www.post-gazette.com/stories/local/neighborhoods-east/students-set-for-robotics-contest-628757

Training at Lehigh University: PSC has had a long standing relationship with Lehigh University, and we have participated in their annual high performance computing events since 2007. This year, on Thursday, March 22, John Urbanic gave a half day session to about 35 students on OpenMP and OpenACC in conjunction with their HPC Symposium, formerly known as the Lehigh HPC Day. See http://hpcsymposium.lehigh.edu/2012/. Urbanic reported that the session went extremely well and he even had a request to present a workshop at Rutgers University on OpenACC.

18.6 SP Collaborations

Tennessee State University (TSU): TSU has been awarded a grant from NSF that provides funds for students studying bioinformatics as well as for bioinformatics lab and computer equipment. The PI's are Ali Sekman (Computer Science) and Terry Johnson (Biology). PSC's Hugh Nicholas and Alex Ropelewski serve on the External Advisory Committee for this grant along with Dr. Bhanu Rekepalli from Oak Ridge, and they recently attended the first annual meeting of the Advisory Committee at TSU. While there, they met with Dr. Keith Hargrove, Dean of the College of Engineering at TSU, and several other administrators and researchers. They also gave several guest lectures.

Pallavi Ishwad of PSC accompanied Nicholas and Ropelewski and talked with the outreach people at Tennessee State University (TSU) and also at North Carolina Agricultural and Technical State University (NCA&T) in order to discuss joint outreach projects. There was interest at both NCA&T and TSU in using some of the PSC-developed bioinformatics materials in their summer outreach programs. There was also interest in collaborating on additional outreach efforts.

Harvard University: For the past few years, PSC and the Center for Brain Science at Harvard University have been collaborating on a project to map the neuron-to-neuron wiring diagram of the mouse visual cortex. The Harvard team captures both in vivo functional (epi-fluorescence two-photon calcium) and high-resolution serial section transmission electron microscopy (ssTEM) image sets to correlate neural firing patterns with detailed ultrastructure. TEM acquisition uses a customized multiple camera array to maximize capture rate. The TEM image data, up to a TByte per day, are transmitted to PSC where the PSC team processes and assembles the images for viewing and analysis using PSC-developed software. Art Wetzel and Greg Hood of PSC have aligned the initial substack of Wei-Chung Allen Lee's (Harvard) new mouse visual cortex dataset for a poster at the March 2012 Cold Spring Harbor meeting on Neuronal Circuits. This is part of an ongoing capture series that will reach ~125 TBytes by the end of 2012.

Wetzel, Hood and PSC networking staff members have also been working with a second Harvard team to increase the bandwidth for a scanning electron microscopy (SEM) collaboration that will study different areas of mouse brain and other organisms.

18.7 SP-Specific Activities

User Appreciation: In response to proactive support provided by Marcela Madrid, Kenno Vanommeslaeghe of Alex MacKerell's group at the University of Maryland at Baltimore wrote, "… based on our past and present experiences, if there would be a prize for best supercomputer user support, I'm sure PSC would win it."

Kudos for PSC's Networking Group: Regarding the problem of standard SSH small window sizes slowing down transmission of files over high performance networks, bloggers at http://blogs.techworld.com/power-consciousness/2012/03/maximising-the-use-of-our-network-resources/index.htm said, in reference to PSC-developed HPN-SSH, that "... the solution is available and we need to look no further than the experts at the University of Pittsburgh Supercomputer Center [sic]" in a discussion about maximizing the use of their network resources.

18.8 Publications

Lasinski ME, Romero NA, Brown ST, Blaudeau JP. "Recent performance improvements to the DFT and TDDFT in GAMESS". Journal of computational chemistry. 2012 Mar 15; 33 (7):723-31. PMID: 22241553.

Lee BY, Assi TM, Rajgopal J, Norman BA, Chen SI, Brown ST, Slayton RB, Kone S, Kenea H, Welling JS, Connor DL, Wateska AR, Jana A, Wiringa AE, Van Panhuis WG, Burke DS. "Impact of introducing the pneumococcal and rotavirus vaccines into the routine immunization program in Niger". American journal of public health. 2012 Feb; 102 (2):269-76. PMID: 21940923.

L. McGinnis, K, Benninger, M. Madrid, D. Nigra, C. Begandy, "Tour Your Future", Girls Math/Science Partnership, January 28, 2012.

L. McGinnis, “XSEDE: A National CI-Ecosystem for Scientific Discovery”, Clemson EPSCoR Student Engagement Program, February 22-23, 2012.

L. McGinnis, “Think Like a Computer”, Carnegie Science Center SciTech Festival, March 6-9, 2012.

L. McGinnis, “Think Like a Computer - for Teachers”, PRSEF Faculty Workshop, March 30, 2012.

Margaret A. Potter, JD, MS; Shawn T. Brown, PhD; Bruce Y.Lee, MD, MBA; John Grefenstette, PhD; Christopher R. Keane, ScD; Chyongchiou J.Lin, PhD; Sandra C. Quinn, PhD; Samuel Stebbins, MD, MPH; Patricia M.Sweeney, JD, RN, MPH; Donald S. Burke, MD, “Preparedness for Pandemics: Does Variation Among States Affect the Nation as a Whole?”, J Public Health Management Practice,2012,18(3),233-240.

Szarecka A, Lesnock KR, Ramirez-Mondragon CA, Nicholas HB Jr, Wymore T. “The Class D beta-lactamase family: residues governing the maintenance and diversity of function.” *Protein Eng Des Sel*. 2011 Oct;24(10):801-9 Epub 2011 Aug 22. PubMed PMID: 21859796; PubMed Central PMCID: PMC3170078.

E. Wierzbinski, A. De Leon, X. Yin, A. Balaeff, K. Davis, S. Rapireddy, R. Venkatramani, S. Keinan, D. Ly, M. Madrid, D. Beratan, C. Achim, D. Waldeck, “The Effect of Backbone Flexibility on Charge Transfer Rates in Peptide Nucleic Acid Duplexes”, submitted to Journal of the American Chemical Society, March 2012.

18.9 Metrics

18.9.1 Standard User Assistance Metrics

Numbers below and the numbers in Table 1 refer to all tickets handled by the PSC help desk in PSC’s local ticket system.

Tickets Created between 2012-01-01 and 2012-03-31: 421

Tickets Resolved between 2012-01-01 and 2012-03-31: 438

Table 1: Distribution of times to resolution for the 417 tickets that were created as well as resolved between 2012-01-01 and 2012-03-31.

Ticket Type	0–1 hr	1–24 hr	1–7 Days	1–2 weeks	> 2 weeks
System issues	9	16	31	1	1
Software/apps	4	22	79	12	7
Other	1	9	6	6	0
Login/access	2	1	3	0	1
Jobs/batch queues	6	50	70	6	2
File systems	5	7	17	5	0
Account issues	2	6	28	2	0

Numbers in Table 2 refer to tickets relating to PSC that were handled in the central XSEDE Ticket System.

Table 2: Distribution of times to resolution

Time to Resolution	account issues	file systems	grid software	jobs/batch queues	login/access issues	mss/data issues	network issues	software/apps	system issues	other
0-1 hr		1			1	2	1			
1-24 hr	1		1							
1-7 d	3	2		15	9	3		8		
1-2 wk	2			15	9			10	1	1
> 2 wk		1		10	5	1		10		1
Still Open	1									

18.9.2 PSC-specific Metrics

Key system statistics for *Blacklight* are shown in Table 3.

Table 3: Operational Statistics - *Blacklight*		
Number of unplanned outages	17	
Number of planned outages	4	
Total outages	21	
Number of jobs failures due to system faults	279	
Total time* in period (hours)	4368	100%
Scheduled Downtime (hours)	34.00	1%
Unscheduled Downtime (hours)	236.52	5%
Total Downtime (hours)	270.52	6%
Total time available to users (total-downtime)	4097.48	94%
% System Utilization	79.3%	

* On *Blacklight* a node is half the machine. Time values listed are expressed in node hours.

18.9.3 Standard system Metrics

The following charts of standard system metrics for *Blacklight* were provided by the Technology Audit Services team.

PSC Blacklight Quarterly Report

Total NUs Charged by Resource

XSEDE

2012-01-01 to 2012-03-31

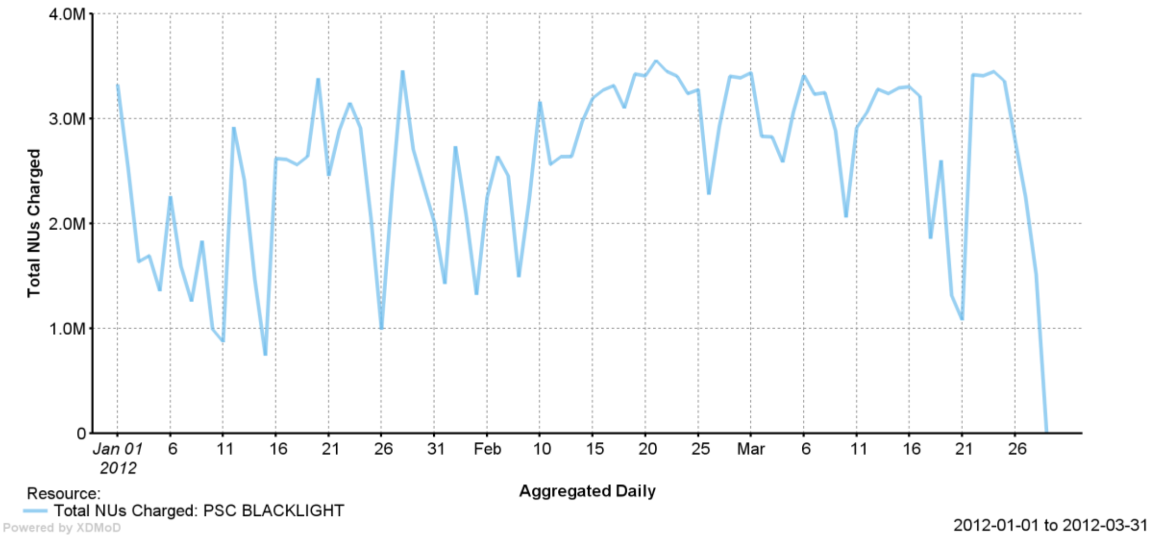

Total NUs Charged by Job Size

Resource = PSC-BLACKLIGHT

2012-01-01 to 2012-03-31

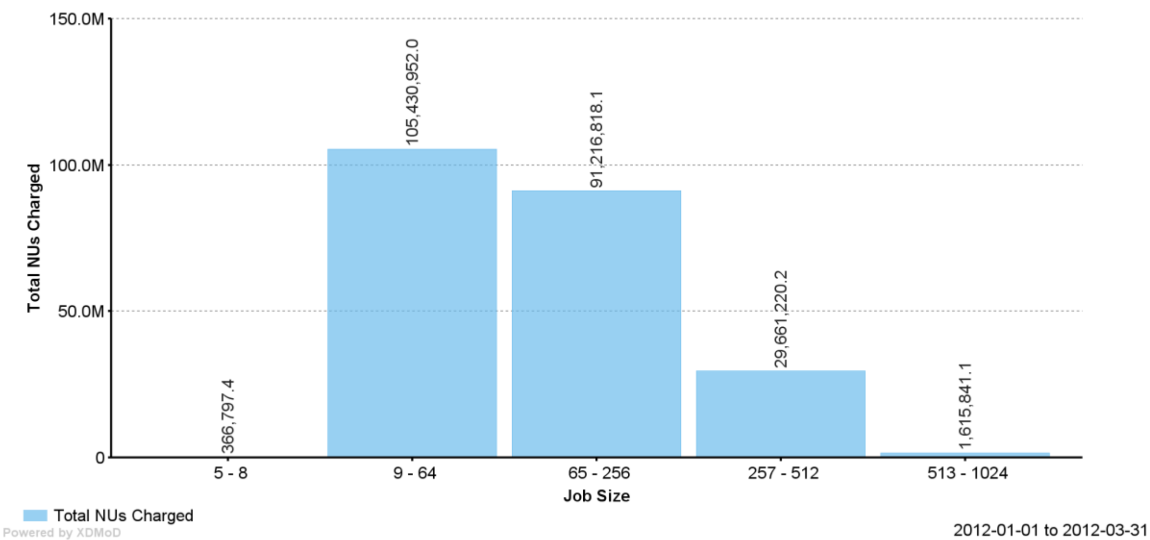

Avg Wall Hours Per Job by Job Size

Resource = PSC-BLACKLIGHT

2012-01-01 to 2012-03-31

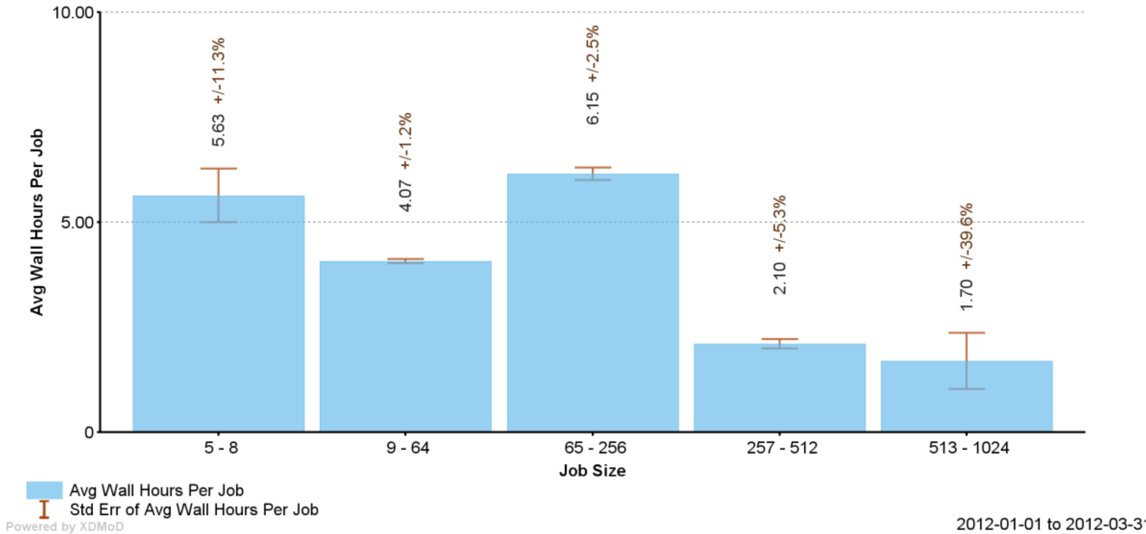

Avg Wait Hours Per Job by Job Size

Resource = PSC-BLACKLIGHT

2012-01-01 to 2012-03-31

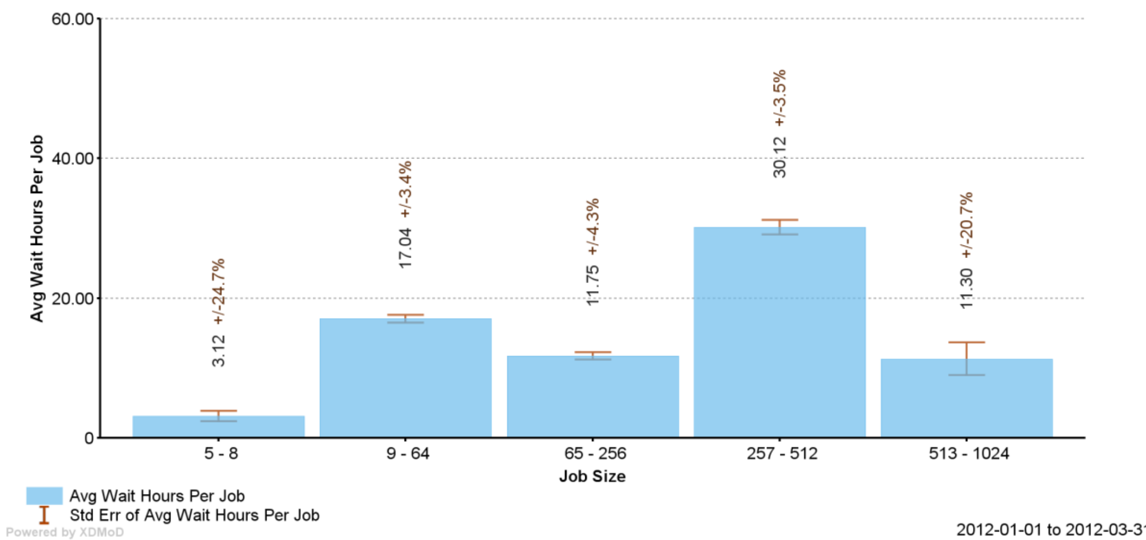

User Expansion Factor by Job Size

Resource = PSC-BLACKLIGHT

2012-01-01 to 2012-03-31

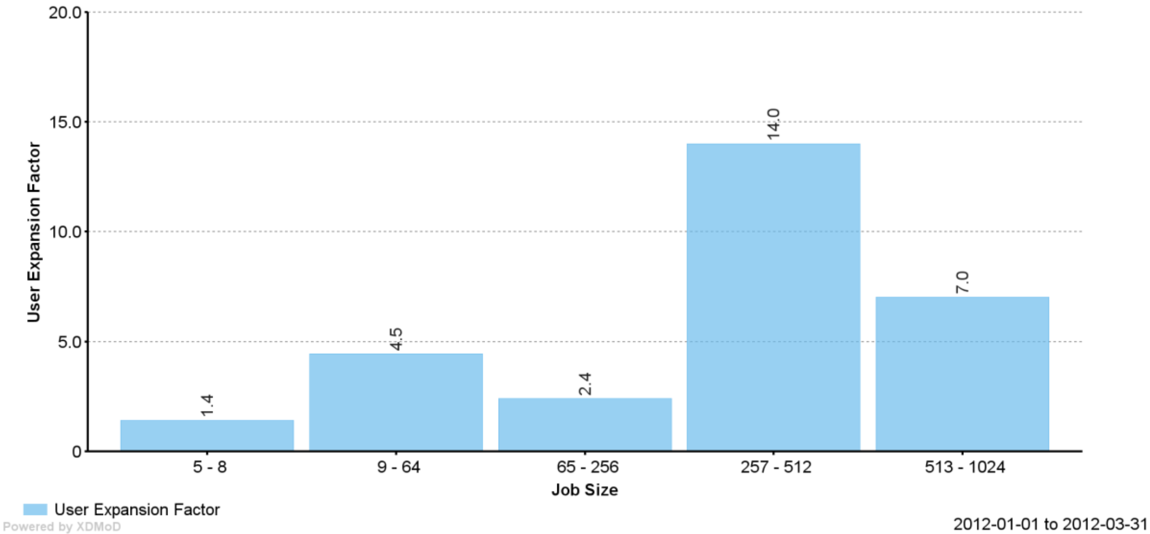

Total NUs Charged by Job Wall Time

Resource = PSC-BLACKLIGHT

2012-01-01 to 2012-03-31

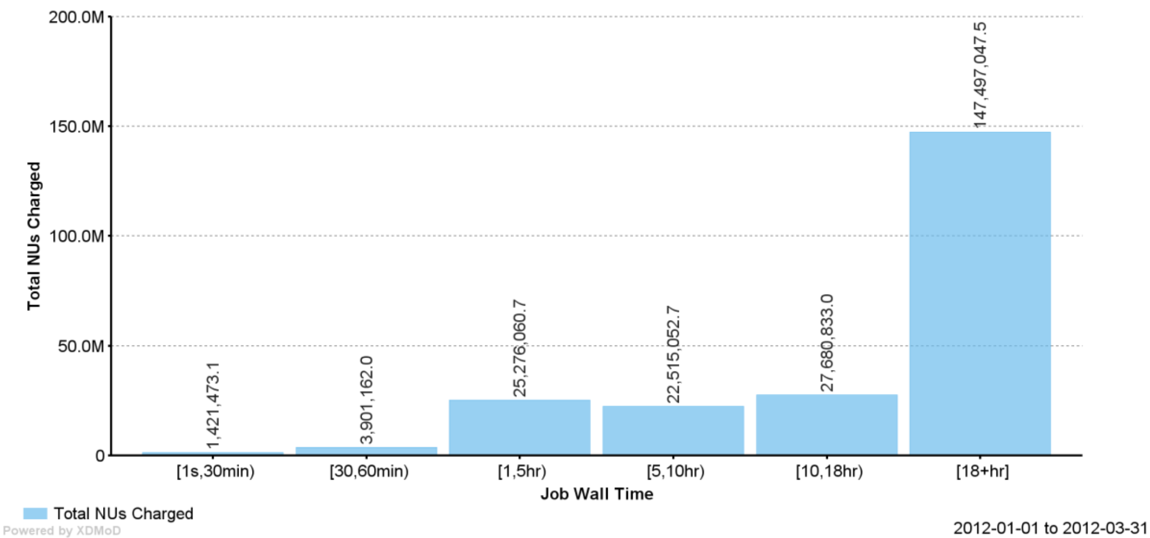

User Expansion Factor by Job Wall Time

Resource = PSC-BLACKLIGHT -- Service Provider = PSC

2012-01-01 to 2012-03-31

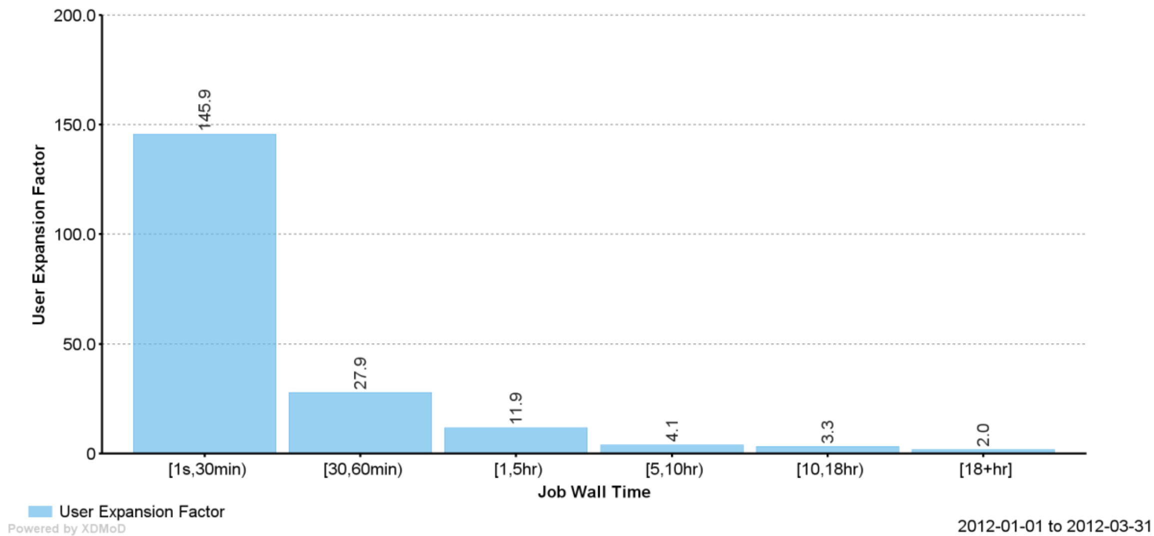

Total NUs Charged by Field of Science

Resource = PSC-BLACKLIGHT

2012-01-01 to 2012-03-31

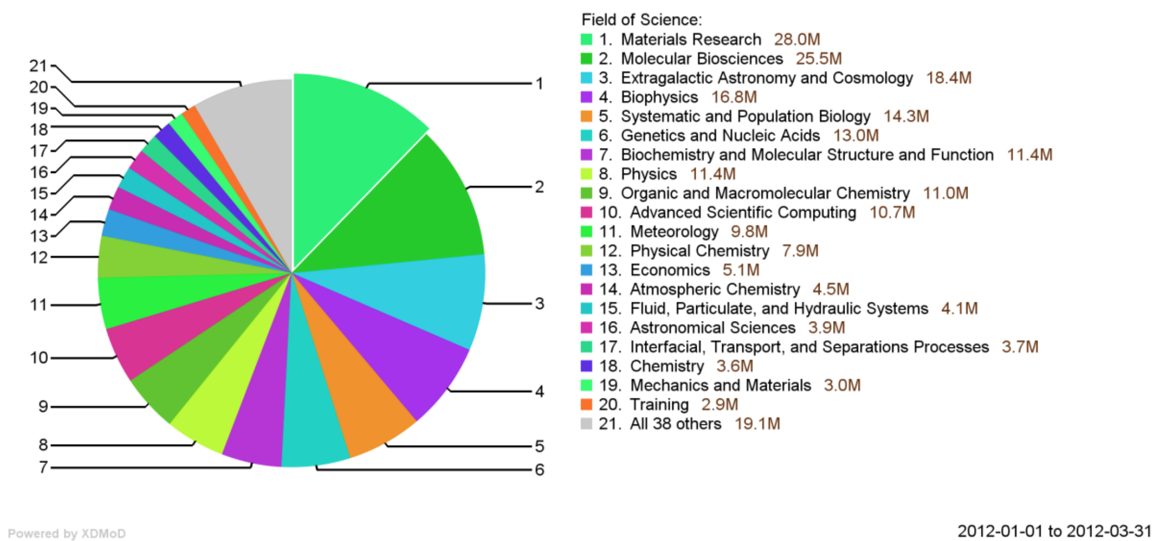

Total NUs Charged by Institution

Resource = PSC-BLACKLIGHT

2012-01-01 to 2012-03-31

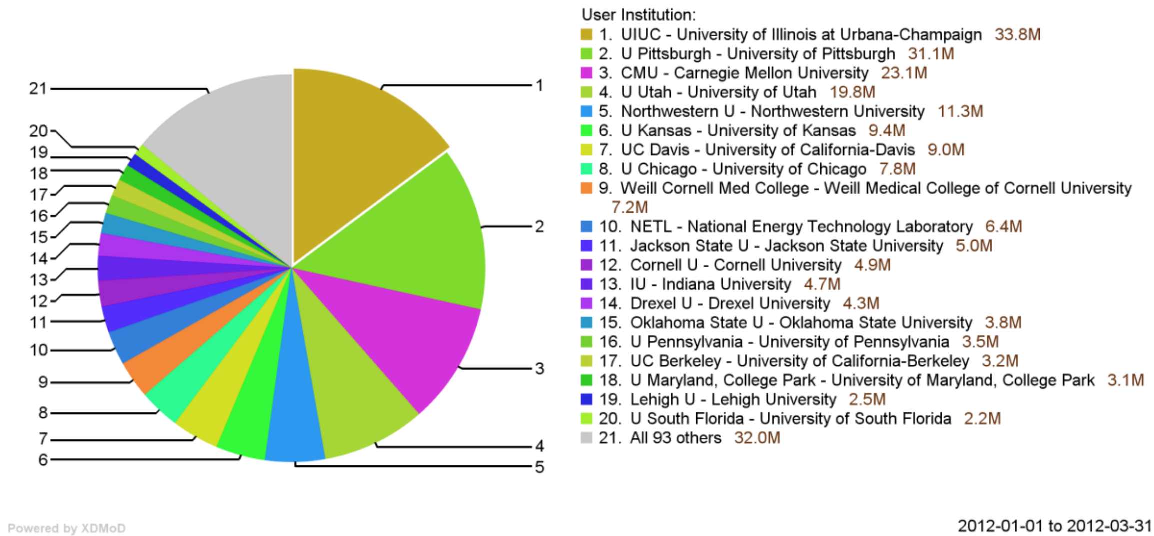

Total NUs Charged by Principal Investigator

Resource = PSC-BLACKLIGHT

2012-01-01 to 2012-03-31

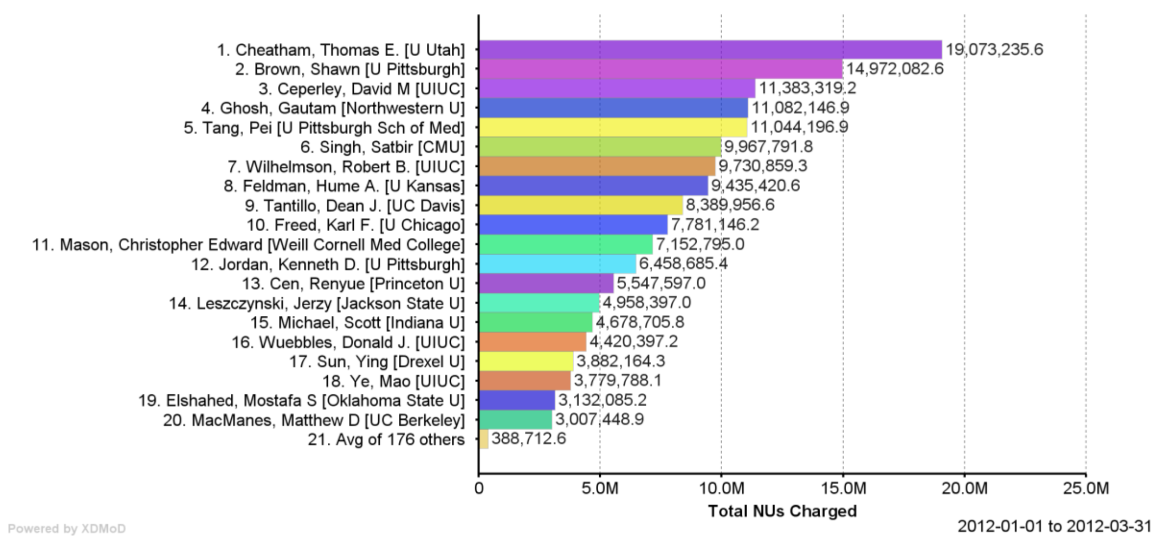

19 Purdue University - Service Provider Quarterly Report

19.1 Executive Summary

Purdue continues to be a service provider (SP) to the XSEDE project as TeraGrid transitioned to XSEDE in July 2011, providing an HPC cluster (Steele), a high-throughput computing resource (the Purdue Condor pool), and a cloud resource (Wispy) to XSEDE users. The SP operates the systems and provides helpdesk and user support, as well as participate in XSEDE wide training and outreach activities. These activities are funded by the NSF awards #0503992, #0932251.

Purdue's roles in XSEDE also include the following (funded by the XSEDE award):

- ECSS staff to provide expert support to user requests for advanced technical assistance.
- Lead for the XSEDE Campus Champion Program to further expand the XSEDE user base by reaching out to campuses, providing training and outreach activities, and getting feedback from campus users to improve XSEDE services.
- An XSEDE – OSG liaison to bridge the efforts of the two cyberinfrastructure projects by providing assistance and guidance to users in utilizing the OSG resources and identifying and organizing training and outreach activities to broaden user base.

These activities and impacts are reported in the overall XSEDE report sections.

Purdue continues to support a number of science gateways that utilize XSEDE data and computational resources, bridge OSG computation high-throughput HPC (HTHPC) jobs to XSEDE resource, and develop, deploy virtualization tools to support scientific users of cloud resources as part of the joint TeraGrid-OSG project ExTENCI.

19.1.1 <u>*Resource Description*</u>

The Steele Cluster

The Steele cluster consists of 893 dual quad-core Dell 1950 compute nodes, running Red Hat Enterprise Linux, version 4. Each node thus has 8 64-bit Intel 2.33 GHz E5410 CPUs and either 16 GB or 32 GB of RAM. They are interconnected with either Gigabit Ethernet or InfiniBand. The machine offers access to the 120 TB scratch space. Steele's peak performance is rated at 66.59 TFLOPS. Steele cluster is well suited for a wide range of both serial and parallel jobs. Steele replaced the Purdue Lear cluster and was made available to TG users in May 2008. Its projected useful lifetime is through July 2013. In October 2009, Purdue RP has increased the TG *dedicated* portion of Steele from 22 nodes to nearly 200 nodes (1600 cores), which significantly reduced job wait time. Steele has no effective runtime limit on XSEDE jobs. Additionally, XSEDE users may leverage the larger Steele cluster by utilizing the standby queues with no node limit but subject to 4 or 8 hour runtime limits.

Condor Pool

The Purdue Condor pool is a shared resource among the resource owners (academic users at Purdue) and XSEDE/OSG users. Consisting of more than 40,000 cores, the Condor pool is an opportunistic resource which allows Condor jobs access to machines that are not being used by their owners. The Purdue Condor pool is designed for high-throughput computing, and is excellent for parameter sweeps, Monte Carlo simulation, or most any serial application. Also, some classes of parallel jobs (master-worker) may be run effectively in Condor. 30% of all Condor-usable cycles are available to XSEDE users at a minimum level of service. On average the Purdue Condor pool is able to provide up to 9 million CPU hours to XSEDE users per year.

The Purdue Condor resource, recently named DiaGrid, has expanded tremendously from a total of 7700 CPUs at the end of 2007 to its current size of more than 40,000 CPUs (system information shown in Table 1). The Purdue Condor pool consists of nodes from 10 institutions, including Purdue's West Lafayette campus, University of Wisconsin-Madison, University of Nebraska-Lincoln, University of Louisville, Indiana University, University of Notre Dame, Indiana State University, IUPUI Fort Wayne, Purdue North Central and Purdue Calumet campuses. Memory on most of the compute nodes is 1 GB, 2 GB and 4 GB per core, while a small number of nodes have larger memory (e.g., 10GB per core). With a total of over 177 TFLOPS available, the Purdue Condor pool can provide large numbers of cycles in a short amount of time. All shared areas and software packages available on Steele are available on Condor. Available to TeraGrid users since 2006, the Condor pool is self-renewing as old machines in the pool are retired and new ones added over time.

Table 1: Purdue Condor pool information as of October 2011

System Information	Cores	Total Memory	Local Interconnect	Processor Speed
X86_64 LINUX	39,866	105 TB	10 Gb or 1 Gb Ethernet	Various (2.1, 2.33, 2.5, 3.2 GHz)
INTEL & X86_64 WINDOWS	409	1.09 TB	1 Gb or 100Mb Ethernet	Various (2.13, 2,66, 3.6GHz)
INTEL LINUX	44	66 GB	1 Gb or 100Mb Ethernet	various
Total	**40,319**	**106 TB**		

Wispy

Purdue's *Wispy* is a special XSEDE resource, a cloud computing platform for research and education use. Wispy consists of 32 64bit, 4-core HP DL140 connected via 1 Gigabit Ethernet network with the capacity of supporting 128 VMs. Wispy runs KVM and the Nimbus cloud software. It provides users with the capability of packaging their applications and operating systems completely inside the Virtual Machine (VM) images, submitting these VMs to run in Wispy with up to 4 CPUs and 16GB of memory each, and have full control over the execution environment. Current usage includes small, instant, on-demand clusters for various tasks and running complicated or prepackaged applications on additional hardware resources.

19.2 Science Highlights

Computational Studies of Protein/Peptide Interactions in Biological Membranes
PI: Wonpil Im, University of Kansas (field of science: biophysics)

Wonpil Im's lab at the University of Kansas is exploring the dynamics, orientation, and interaction of various proteins and protein-like peptides in membrane environments. The goal is to systematically investigate properties of biological membranes, such as those that protect cells, and to better understand the protein and peptide structures and dynamics in these membranes. In particular, Im's lab has recently developed a method of understanding both orientation and dynamics of transmembrane helices of membrane proteins, structures that can act as channels through membranes both for beneficial purposes and with negative consequences, from solid-state nuclear magnetic resonance (NMR) observations. These membrane constituent interactions form the basis for many vital cellular processes. A better understanding of these processes can promote development of new pharmaceutical therapies at the molecular level.

In addition, Im's lab has made considerable efforts in developing a community-wide resource, CHARMM-GUI (http://www.charmm-gui.org), a web-based graphical user interface to facilitate and standardize the usage of common and advanced simulation techniques in the widely used CHARMM and NAMD simulation programs. CHARMM-GUI is now in use by many researchers for building various complex membrane systems on the Web. Recently, features were added to help users perform glycoprotein simulations and ion channel Brownian dynamics, important to understanding fundamental roles of glycoproteins, which are often involved in cell-to-cell interactions, and ion channels, pore-forming proteins that form, and regulate traffic through, routes across membranes. This group reports a number of publications (listed in Section 1.8) that are related to this work.

The calculations for this research require extensive computational resources hence the support from XSEDE is indispensable. The work can involve numerous single-node parallel jobs running for lengthy periods, as much as 720 hours, which is well supported by the Steele cluster XSEDE resource at Purdue. In this quarter, this group has used more than 2.2 million CPU processor hours on Steele.

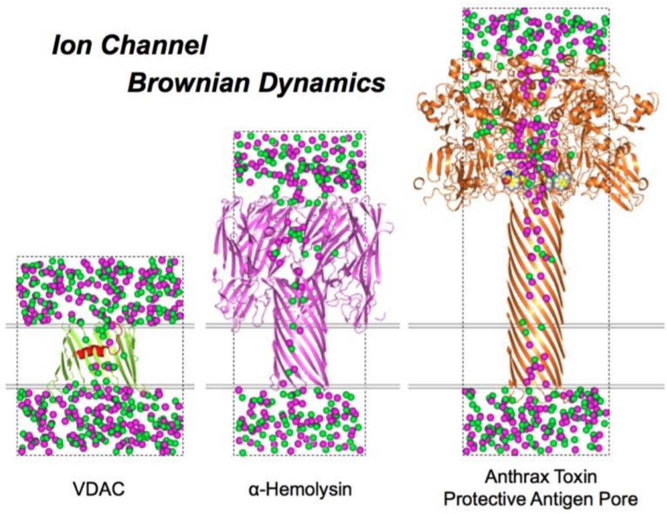

Snapshots of ion trajectories from the BD simulations of (A) hVDAC1, (B) α-HL, and (C) PA (PROF) bathed in 1.0 M KCl symmetric solutions.

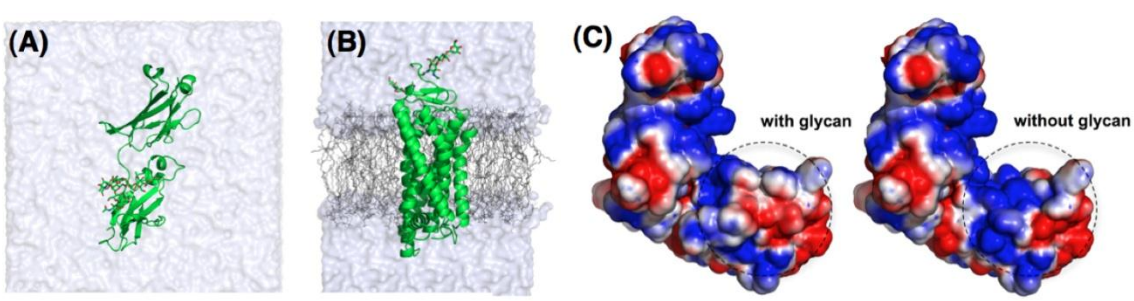

Images of glycoprotein simulation systems (A) in solution with constant region of immunoglobulin-1 (IgG-1) (PDB: 1L6X) and (B) in a bilayer with bovine rhodopsin (PDB:1GZM). (C) The electrostatic potential surface of PDB:1L6X with and without glycan.

Sigma Delta Pulse Processing
PI: Martin McCormick, MIT (field of science: Circuits and Signal Processing).

Computing devices based on continuous-time digital circuits could offer a number of advantages in reliability, robustness, speed, lower power consumption, and simplification that complements mixed analog-digital designs. The technology might be applied in tiny devices like hearing aids (to improve sound filtering, for example); in low-power flexible sensors and other flexible devices; in less-error-prone high-bandwidth communications; and in more quickly solving computationally demanding inference problems. These circuits, in essence, function something like analog circuits, which use continuous time voltages and currents, by processing a stream or pulse of digital bits, as opposed to classical digital processing with its discrete time and discrete value characteristics. Massachusetts Institute of Technology graduate student Martin McCormick is studying an alternative paradigm for designing computers using continuous-time digital circuits with help from the Steele cluster XSEDE resource at Purdue. McCormick uses a low-density parity-check (LDPC) coder simulation, which simulates a method of transmitting a message over a noisy transmission channel, to study various aspects of pulse processing with continuous-time digital circuits, among them the performance floor for the technology. Recent work suggests that this approach matches the performance of ideal analog decoders that are not physically realizable. These simulations involve numerous iterations. The availability of Steele allows McCormick to reduce his time to results from several days to less than a day.

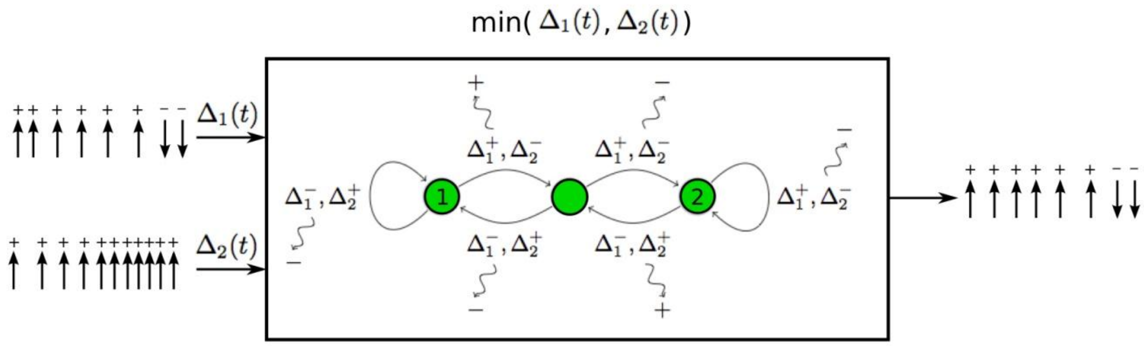

An example of the kind of pulse domain operations which can be handled by continuous-time digital circuits.

Windows HPC Cluster helping with data intensive computing
PI: Brian Pijanowski, Forestry and Natural Resources, Purdue University

Purdue SP staff has recently deployed a Windows based cluster supercomputer now available to the Purdue campuses. Forestry and Natural Resources Professor Bryan Pijanowski's research looks at how people and land use play into the state of ecosystems, including issues like projecting loss of prime farmland to urbanization over the next 50 years and the potential effects of climate change on fish habitats. His lab examines such questions with computer modeling, among other techniques, sometimes on a national scale and at almost-unheard-of resolutions as fine as every 30 meters. The work can involve hundreds of gigabytes to terabytes of data to begin with, trillions of data points and software that runs — in some cases will only run — in Microsoft Windows. In the past, that could mean

With the new Windows HPC cluster Prof. Bryan Pijanowski can look at the whole world at almost-unheard-of 30-meter resolution.(Photo courtesy of Purdue ITaP) running a model for months on a desktop computer. But the new cluster supercomputer using the high-performance computing version of Windows has reduced the time it takes to get results to a few hours and is enabling more detailed and accurate modeling. “We could actually run the entire world at 30-meter resolution if we wanted,” Pijanowski says.

While similar in functionality to Purdue’s other community cluster supercomputers, the Windows HPC cluster offers potential advantages for some users and developers. Users get an entry point to high-performance computing with the familiar Windows user interface. The Windows HPC cluster makes high-performance computing available to researchers who rely on Windows-only software like ArcGIS and Microsoft Excel. In addition, the new cluster can run Windows HPC versions of research software common on Purdue’s Linux community clusters, such as MATLAB and SAS.

19.3 User-facing Activities

19.3.1 System Activities

The Steele cluster had a scheduled maintenance during the university’s spring break. There was also an unplanned outage during March that had some impact on user jobs that had to be restarted. The Wispy resource’s hardware has been relocated to a different data center during a planned maintenance on February 28. In the near future Wispy will undergo a hardware upgrade to servers based on the Intel’s Sandy Bridge series of processors. The table below summarizes the monthly SP resource availability.

	% Uptime (monthly)		
Month	**Condor**	**Steele**	**Wispy**
Jan 2012	100%	100%	100%
Feb 2012	100%	100%	96%
Mar 2012	100%	97%	100%

Other systems:

Purdue SP deployed its fifth and most powerful community cluster, Carter, which began full operation at the end of March 2012 (project lead Ed Lee shown in the picture). Carter was built through a partnership among Purdue, HP, Intel and Mellanox using upcoming technologies. Carter was ranked 54th on the latest international Top500.org list of the world's most powerful supercomputers. The new cluster features HP compute nodes with two eight-core Intel Xeon E-5 “Sandy Bridge” processors, 16 cores per node, 32 gigabytes of RAM, a 500 GB system disk and 56 gigabits per second InfiniBand interconnects. Though Carter is dedicated for campus use, its nodes are in the Purdue Condor pool, available to XSEDE and OSG users.

Purdue has recently deployed a new cluster that runs Windows HPC Server 2008 R2, the latest version of Microsoft's operating system for high-performance computing. Windows HPC is an alternative to Linux-based clustering software packages, with similar features such as cluster management tools, a job scheduler and a message passing interface (MPI) library. This cluster is enabling many faculty, many in earth sciences, whose research requires data intensive computing on a Windows platform (see science highlight in Section 1.2).

19.3.2 Services Activities

The SP user support staff worked with many XSEDE users during this quarter. The following highlights the SP support for a number of community gateways that use the XSEDE resources.

In the past quarter Purdue's efforts focused on wrapping up the CESM gateway as a TeraGrid project and transitioning it into XSEDE. The original objective, set out by the TG GIG project, of developing a prototype of an integrated climate modeling environment that supports end-to-end model execution on TG/XD, metadata generation, data/metadata publishing to ESG and output analysis. The prototype has been completed, including the installation of the latest release of CESM plus a new model baseline with significantly enhanced metadata capability, a CESM gateway integrated with the ESG data and metadata publishing and advanced data analysis (using LAS – Live Access Server) and integrated data transfer capability for users to transfer large input and output files related to climate model simulations. The integrated workflow will be demonstrated at an upcoming XSEDE science gateway meeting in May. The SP has set up a production site and a development site in preparation for the formal release of the CESM gateway.

The SP staff helped fixed several problems in the C4E4 gateway caused by changes on the backend cluster system. The gateway was successfully used in a graduate class ABE 522 (Ecohydrology). Students ran SWAT auto calibration simulations which entails long running jobs that would not be feasible on regular lab computers. Positive feedbacks were received at the end of the semester.

The SP also worked with a NSF CI-TEAM project at Purdue to enable not only computation intensive models using XSEDE resources but also sharing of input models (SWAT in this application). The staff helped the WaterHUB science gateway create an GIS-enabled online tool named SWATShare which allows users to upload, share and edit their SWAT models, run the models on XSEDE, download the output file and visualize the results online. This type of application requests are more frequently seen by our staff in recent years. There is a rapidly growing need for collaborations and sharing of tools, data and configurations through the use of cyberinfrastructure beyond the traditional HPC applications. SWATShare was used in a graduate level class CE549 (Computational Watershed Hydrology) recently.

19.4 Security

Purdue disabled one account on January 12, 2012 due to compromise at Texas Tech University.

19.5 Education, Outreach, and Training Activities

19.5.1 EOT Events

Type	Title	Location	Date(s)	Hours	Number of Particip ants	Number of Under-represent ed people	Method
Training	New Champions Training	Virtual	Jan. 25, 2012	1			Presentatio n

	By Kim Dillman						
Meeting	Collaboration with the Wales HPC Campus Champion leadership By Kim Dillman	Virtual					presentation
Symposium	Using Condor by Ben Cotton	ECSS Symposium	Jan. 10, 2012	30 min.			presentation

19.5.2 Education

Purdue SP continues to recruit students to join the SP staff team and train them in HPC system administration, software maintenance, and scientific application support. The SP has funded one undergraduate student, Wesley Weber, this quarter to assist with system administration tasks. Wesley is a first year computer engineering student. He handles hardware fixes and front line support for the SP staff. We also have several graduate students to support the operations and users of the SP's XSEDE resources, as well as application support to improve services and capabilities, e.g., enabling online climate model runs using CESM, soil water assessment model (SWAT) and land data assimilation model (HRLDAS) on XSEDE resources.

The SP staff has been supporting two classes in using XSEDE science gateways: a graduate class on ecohydrology (Purdue ABE 522) and a Civil Engineering graduate course (CE549 - Computational Watershed Hydrology), during the 2011-2012 school year. The staff worked with the instructors before the class began and throughout the semesters to ensure that their specific requirements are met by the gateways and that problems and issues resolved promptly. Among the lessons learned are that the need for data sharing and data transfer often varies from class to class, and timely solutions are critical during the semesters in order for the students to complete their assignments.

19.6 SP Collaborations

Purdue RP continues to collaborate with OSG, for example, to support the OSG MPI (high-throughput HPC) jobs on Steele. In this quarter, 3174 OSG MPI jobs ran and consumed 4397 processor-core hours on Steele. Various OSG VO serial jobs consumed a total of approximately 3.75 million processor hours in Purdue's Condor pool during this quarter.

Purdue SP is contributing to and coordinating the virtualization project within the ExTENCI project. The detailed report on Purdue activities related to ExTENCI is included in the overall XSEDE report.

19.7 SP-Specific Activities

The SP staff has investigated ways to make Wispy access easier for end users. Purdue has been collaborating with U of Wisconsin, Fermilab, and Clemson, etc on virtualization technologies to help XSEDE and OSG users. Wispy is being used in this project as a testbed for running scientific applications in virtual machines on it. The SP staff provided access to system staff at other sites and, through this process, we have identified issues and ways to improve access by automating the authentication steps. The SP is also developing tools for easier access to and control of virtual machines that run on Wispy. This part of the work is now integrated into the ExTENCI project.

19.8 Publications

Publications by researchers who utilized SP resources (PI's name in **boldface** text**):**

1. T. Kim, K.I. Lee, P. Morris, R.W. Pastor, O.S. Andersen, and **W. Im**, Influence of Hydrophobic Mismatch on Structures and Dynamics of Gramicidin A and Lipid Bilayers. *Biophys. J. 102:1551-1560* (2012).
2. X. Cheng and **W. Im**, NMR Observable-Based Structure Refinement of DAP12-NKG2C Activating Immunoreceptor Complex in Explicit Membranes. *Biophys. J. 102:L27-L29* (2012).
3. S. Park, T. Kim, and **W. Im**, Transmembrane Helix Assembly with Window Exchange Umbrella Sampling. *Phys. Rev. Lett. 108:108102* (2012).
4. K.I. Lee, S. Jo, H. Rui, B. Egwolf, B. Roux, R.W. Pastor, and **W. Im**, Web Interface for Brownian Dynamics Simulation of Ion Transport and Its Applications to Beta-Barrel Pores. *J. Comput. Chem.* 33:331-339 (2012). *[cover]*
5. T. Kim, S. Jo, and **W. Im**, Solid-State NMR Ensemble Dynamics as a Mediator Between Experiment and Simulation. *Biophys. J.* 100:2922-2928 (2011).
6. S. Jo and **W. Im**, Transmembrane Helix Orientation and Dynamics: Insights from Ensemble Dynamics with Solid-State NMR Observables. *Biophys. J.* 100:2913-2921 (2011).
7. K.C. Song, P.W. Livanec, J.B. Klauda, K. Kuczera, R.C. Dunn, and **W. Im**, Orientation of Fluorescent Lipid Analog BODIPY-PC to Probe Lipid Membrane Properties: A Comparison of Simulation with Experiment. *J. Phys. Chem. B* 115:6157–6165 (2011).
8. H. Rui, R. Kumar, and **W. Im**, Membrane Tension, Lipid Adaptation, Conformational Changes, and Energetics in MscL Gating. *Biophys. J.* 101:671-679 (2011).
9. K.I. Lee, H. Rui, R.W. Pastor, and **W. Im**, Brownian Dynamics Simulations of Ion Transport through VDAC, Voltage Dependent Anion Channel. *Biophys. J.* 100:611-619 (2011).
10. H. Rui, K.I. Lee, R.W. Pastor, and **W. Im**, Molecular Dynamics Studies of Ion Permeation in VDAC, Voltage Dependent Anion Channel. *Biophys. J.* 100:602-610 (2011).

Staff publications

G. J. Bowen, J. B. West, L. Zhao, G. Takahashi, C. C. Miller, and T. Zhang, "Cyberinfrastructure for Environmental Isotope Analysis and Modeling", AGU Eos newspaper.

19.9 Metrics

19.9.1 Standard systems metrics

See system metrics charts below.

19.9.2 Standard User Assistance Metrics

Category	Tickets Received	Tickets Closed	Activities

account issues	1	1	Shell change requests
csa requests	0	0	
filesystems	7	6	Project space request, missing files, quota increase, albedo questions, shared space
gateways	0	0	
gpfs-wan	0	0	
grid software	3	1	GridFTP monitoring, certificate information, GridFTP client request
inca test reports	5	3	MDS stale providers and failures
jobs / batch queues	25	23	Wait time to job start, number of CPUs per core questions, condor errors, redirecting output from jobs, tg_debug queue unusable, running R jobs, queues available
login / access issues	14	14	default shell change, login questions, issues connecting to Steele, submitting gromacs jobs
mss / data issues	2	2	Albedo questions, quota increase
network issues	0	0	
other	3	3	shell change requests and quota increase requests
refund request	0	0	
reservation request	0	0	
security	0	0	
software / apps	16	14	mopac, charmm, namd, hdf5, gfortran, matlab, phoenix and rosetta, VASP
system issues	6	4	CA failures, stale MDS providers
workshops	0	0	
TOTAL	82	71	

Ticket resolution times by category from XSEDE Ticket System

Time to Resolution	account issues	file systems	grid software	jobs/batch queues	login/access issues	mss/data issues	network issues	software/apps	system issues	other
0-1 hr										
1-24 hr		1	1	2	2	1		1		1
1-7 d		2		8	4			3	5	4
1-2 wk		1		10	5			4		1
> 2 wk		1	1	5	1			2		1
Still Open				1				1		

19.9.3 SP-specific Metrics

None.

PURDUE Steele Quarterly Report

Total NUs Charged by Resource

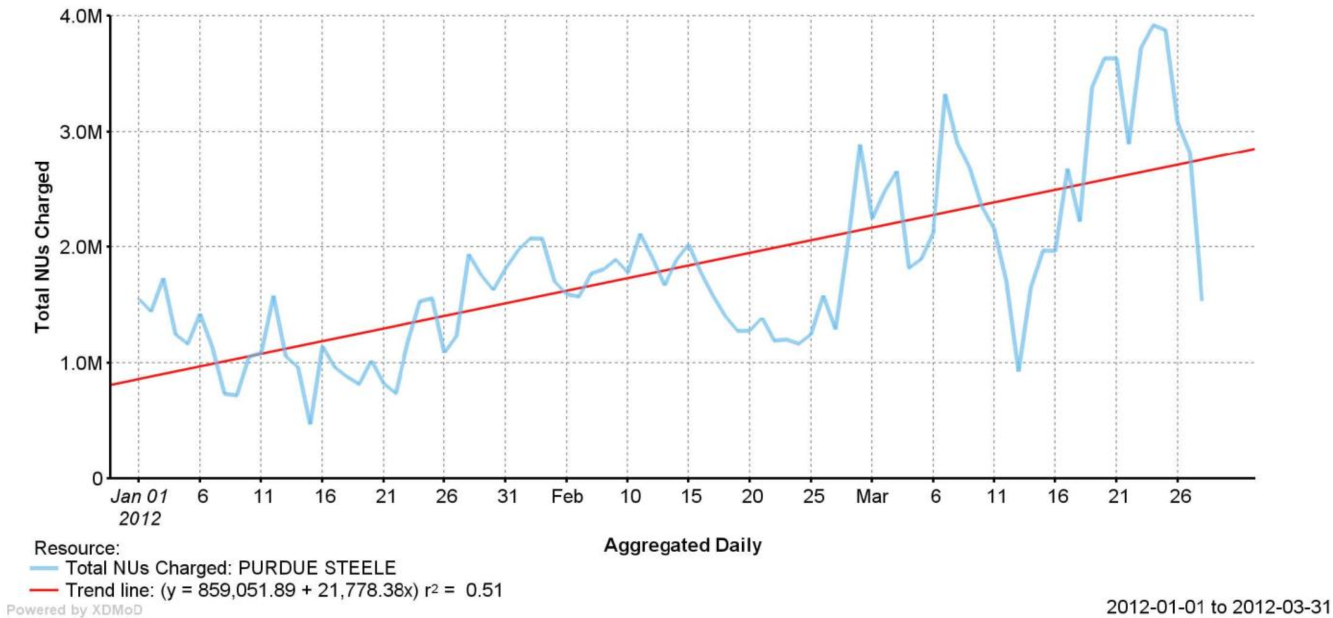

XSEDE 2012-01-01 to 2012-03-31

Total NUs Charged by Job Size

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

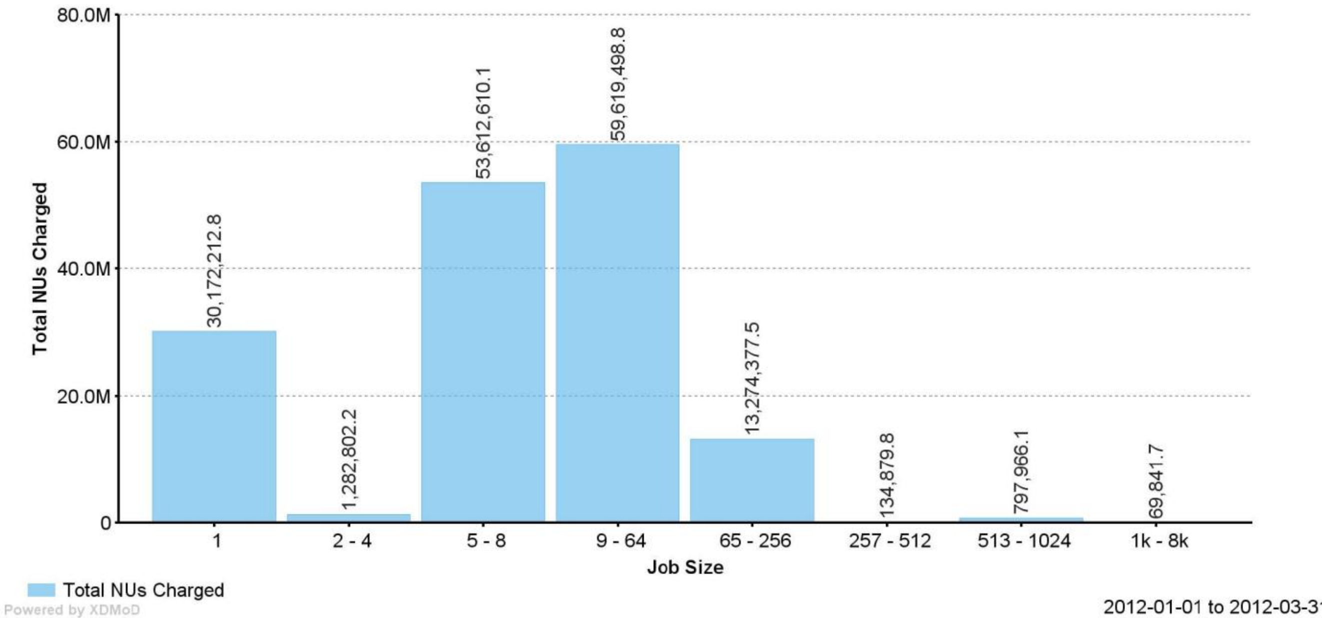

Avg Wall Hours Per Job by Job Size

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

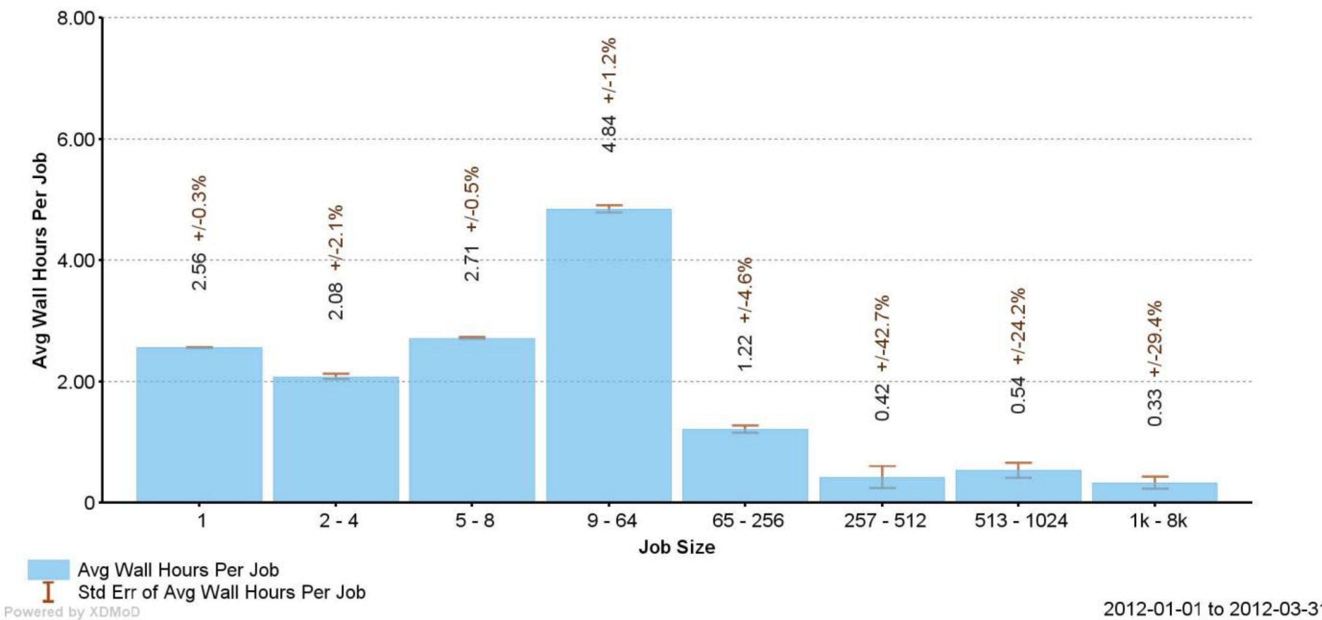

Avg Wait Hours Per Job by Job Size

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

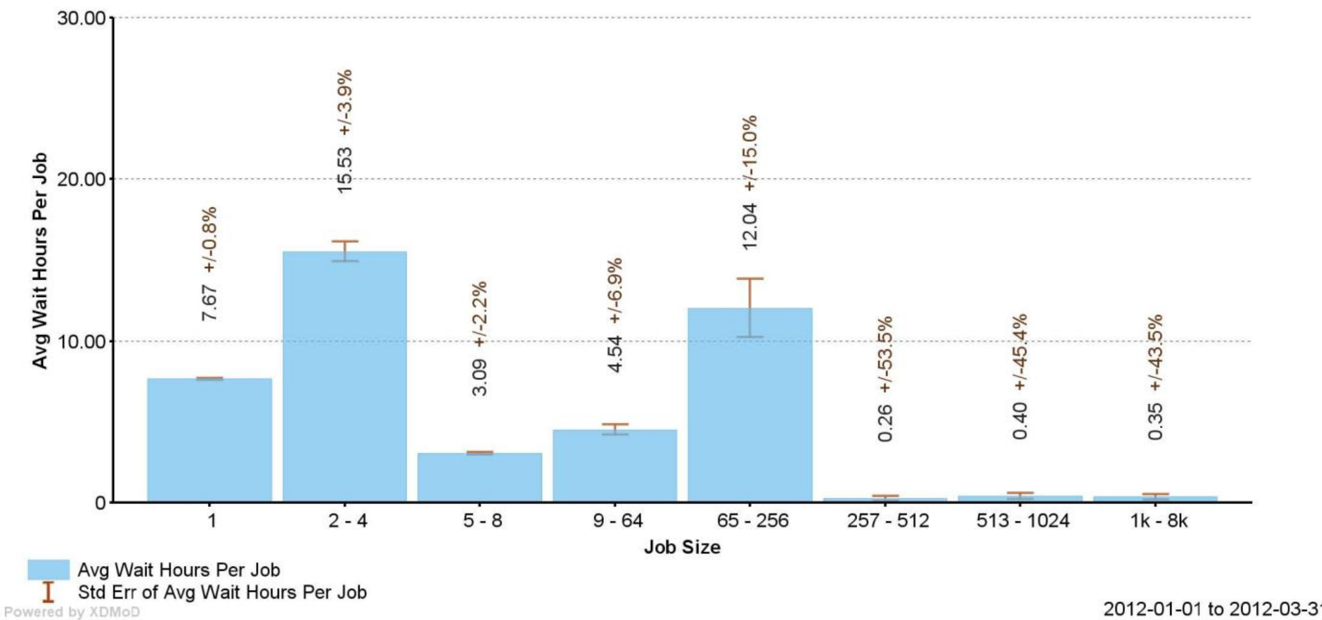

User Expansion Factor by Job Size

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

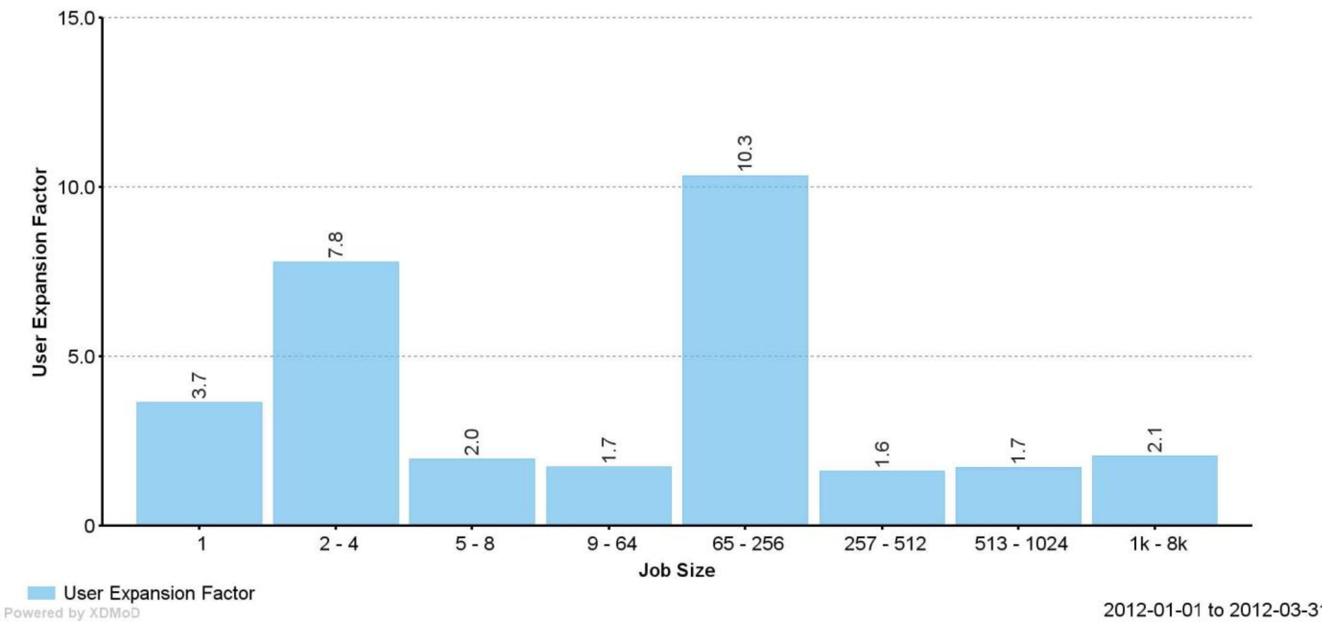

Total NUs Charged by Job Wall Time

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

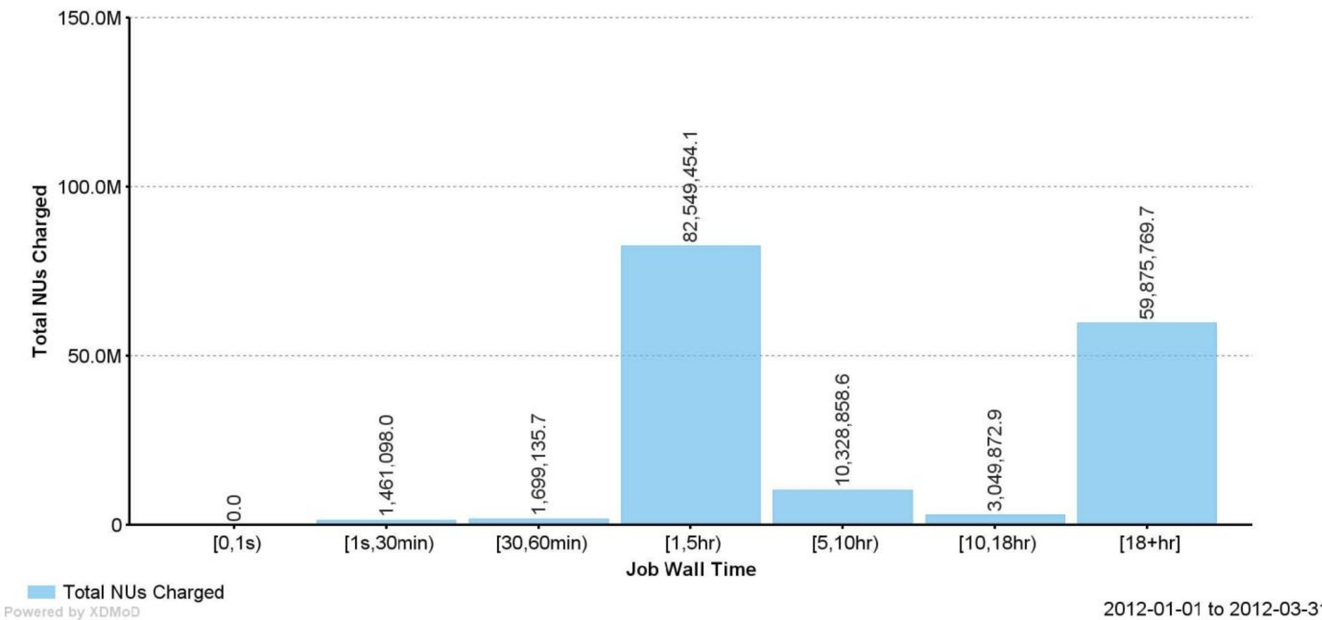

User Expansion Factor by Job Wall Time

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

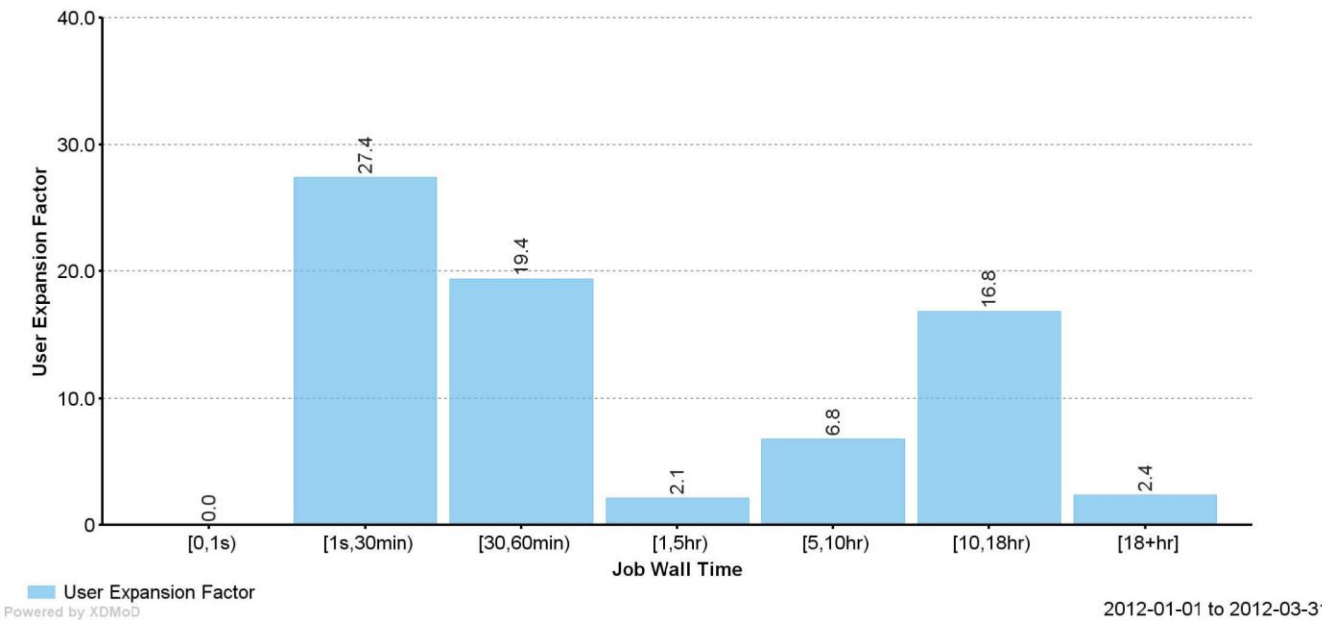

Total NUs Charged by Field of Science

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

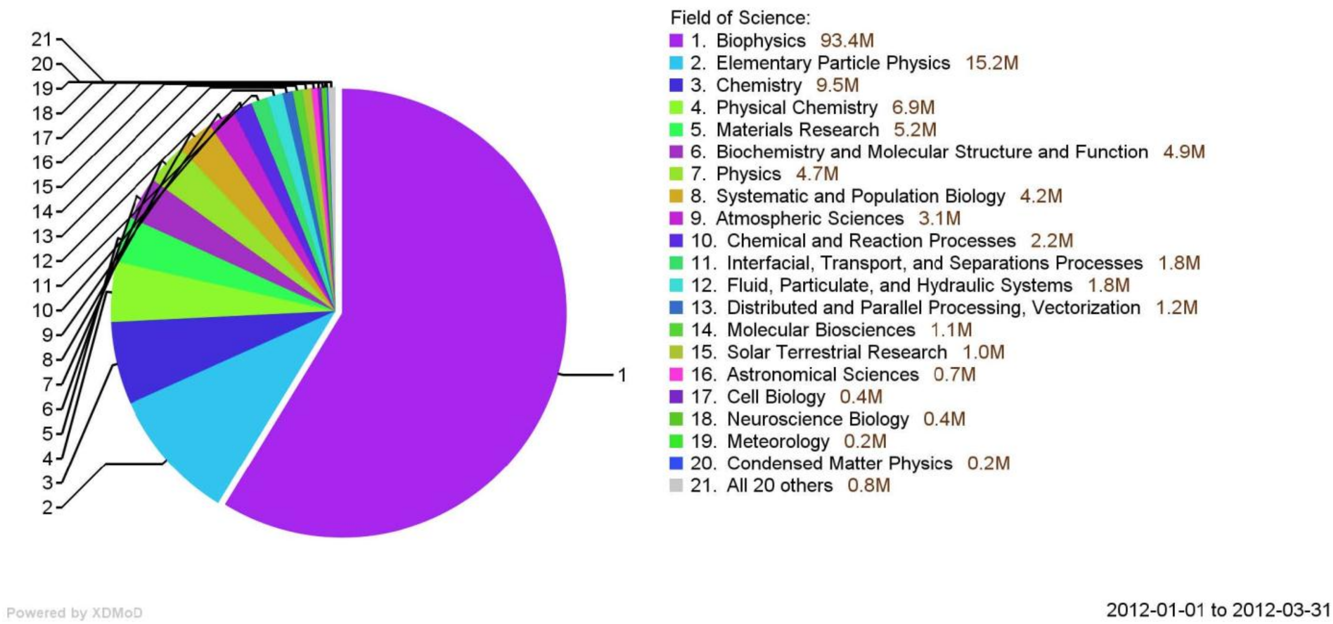

Total NUs Charged by Institution

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

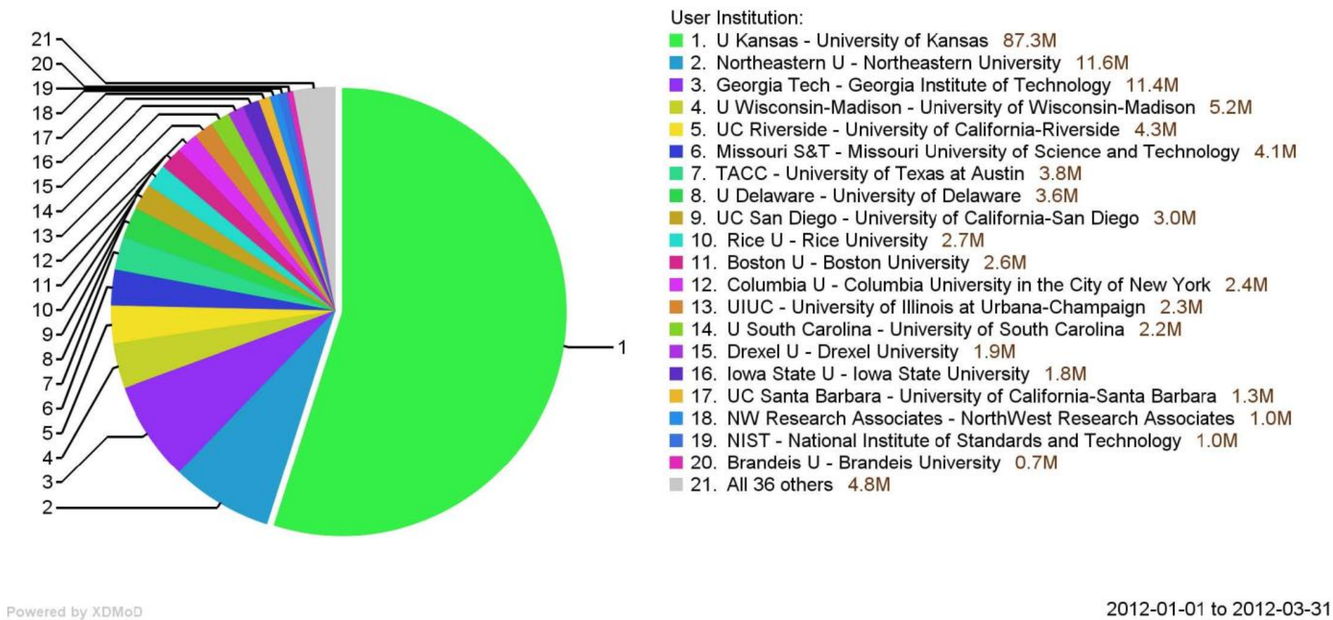

Total NUs Charged by Principal Investigator

Resource = PURDUE-STEELE 2012-01-01 to 2012-03-31

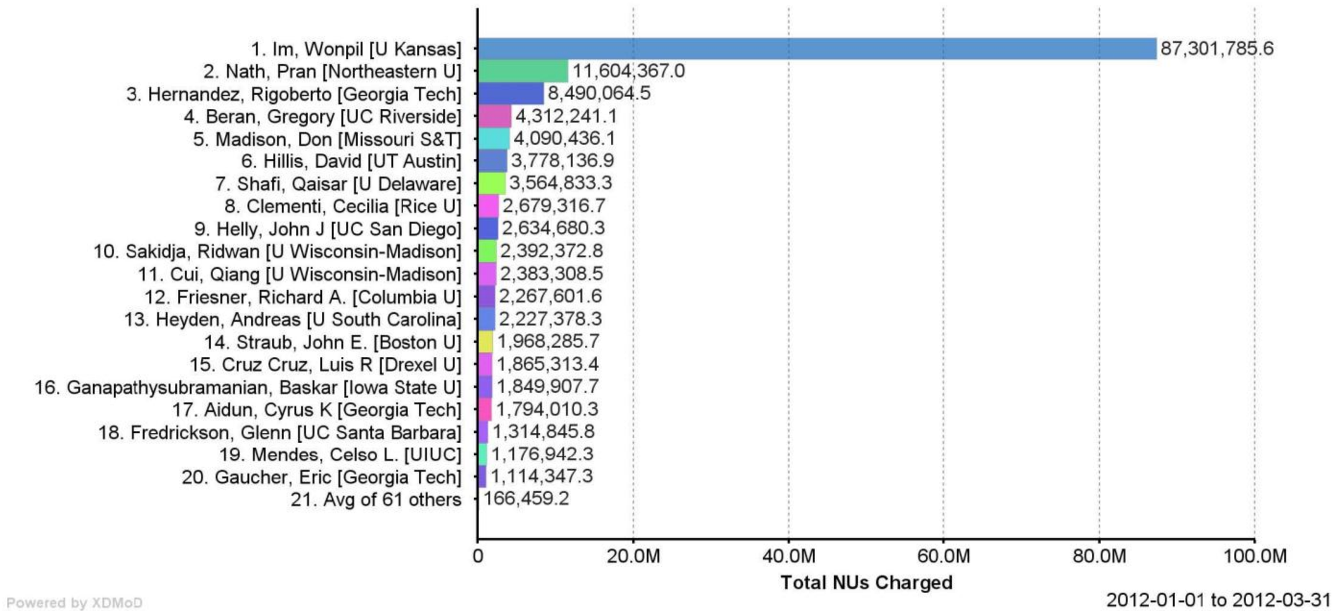

PURDUE Condor Quarterly Report

Total NUs Charged by Resource

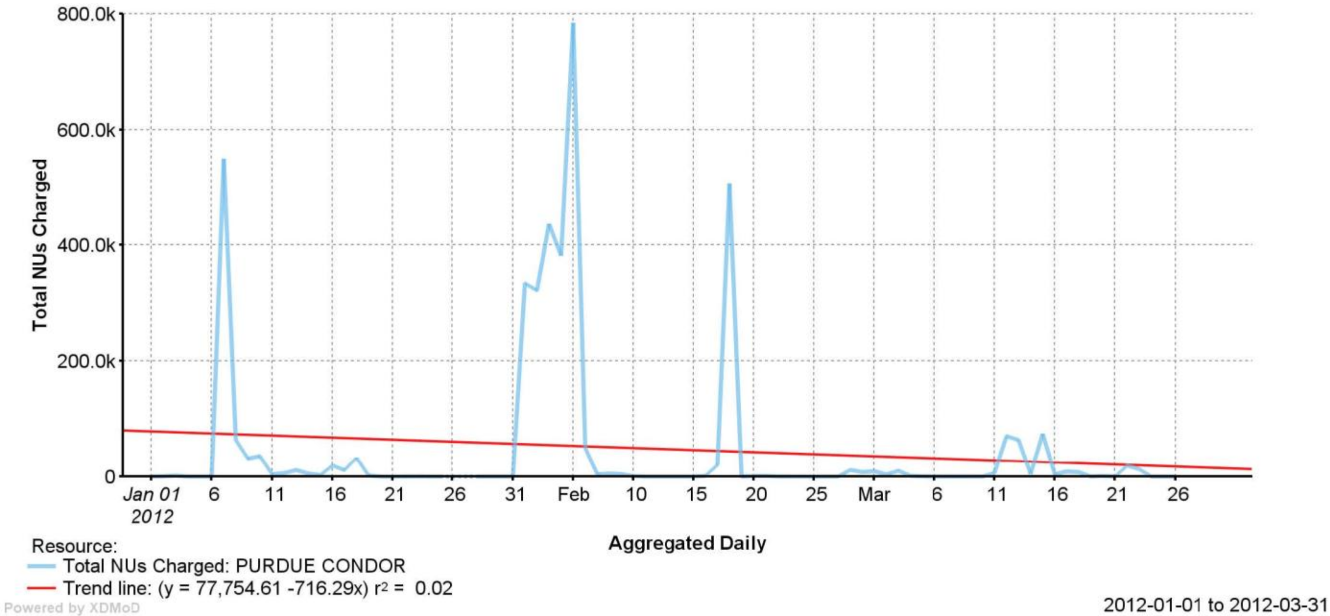

XSEDE 2012-01-01 to 2012-03-31

Total NUs Charged by Job Size

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

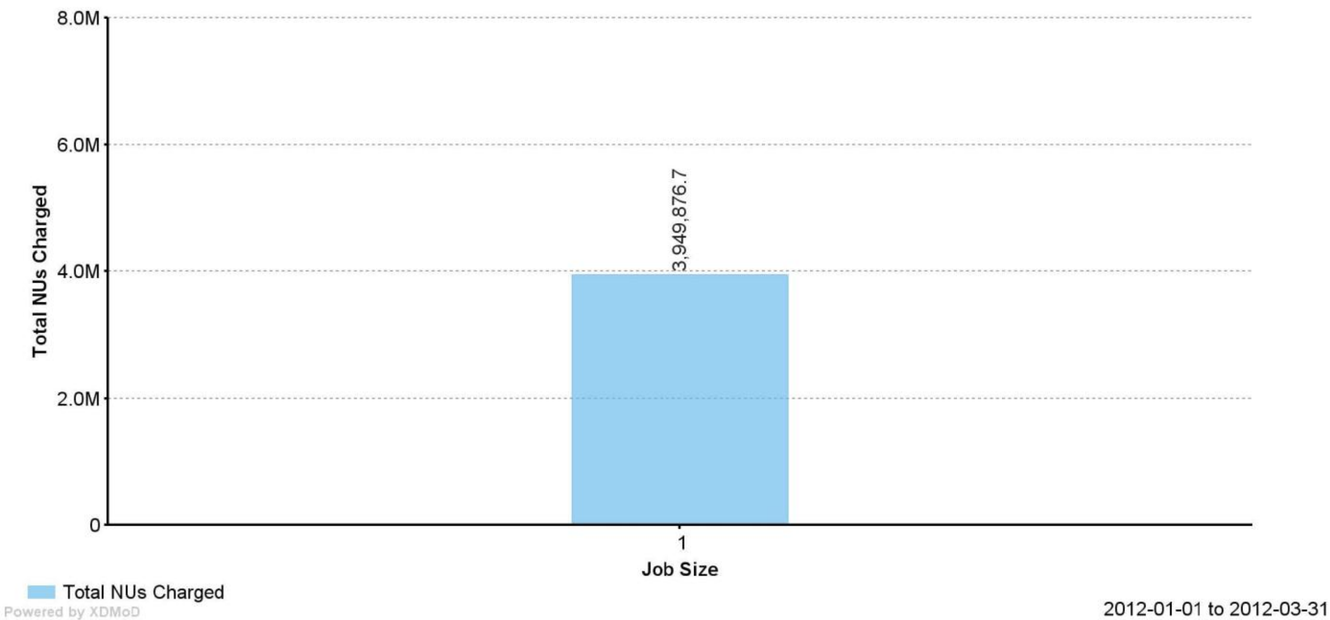

Avg Wall Hours Per Job by Job Size

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

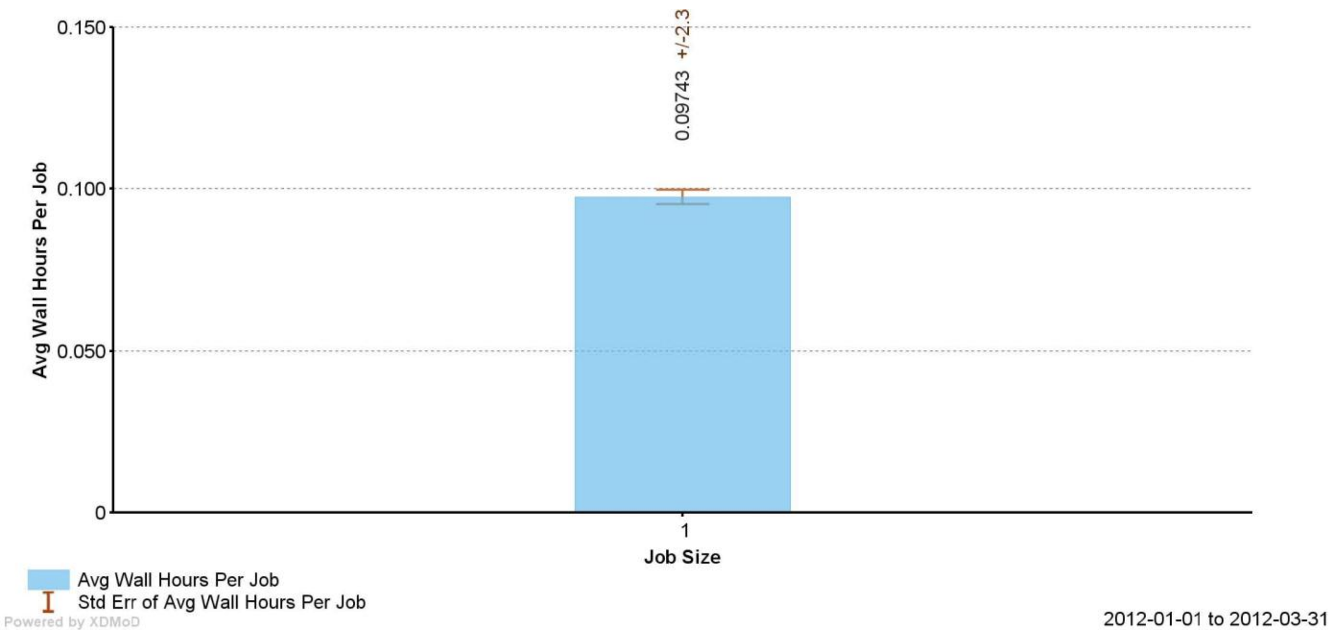

Avg Wait Hours Per Job by Job Size

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

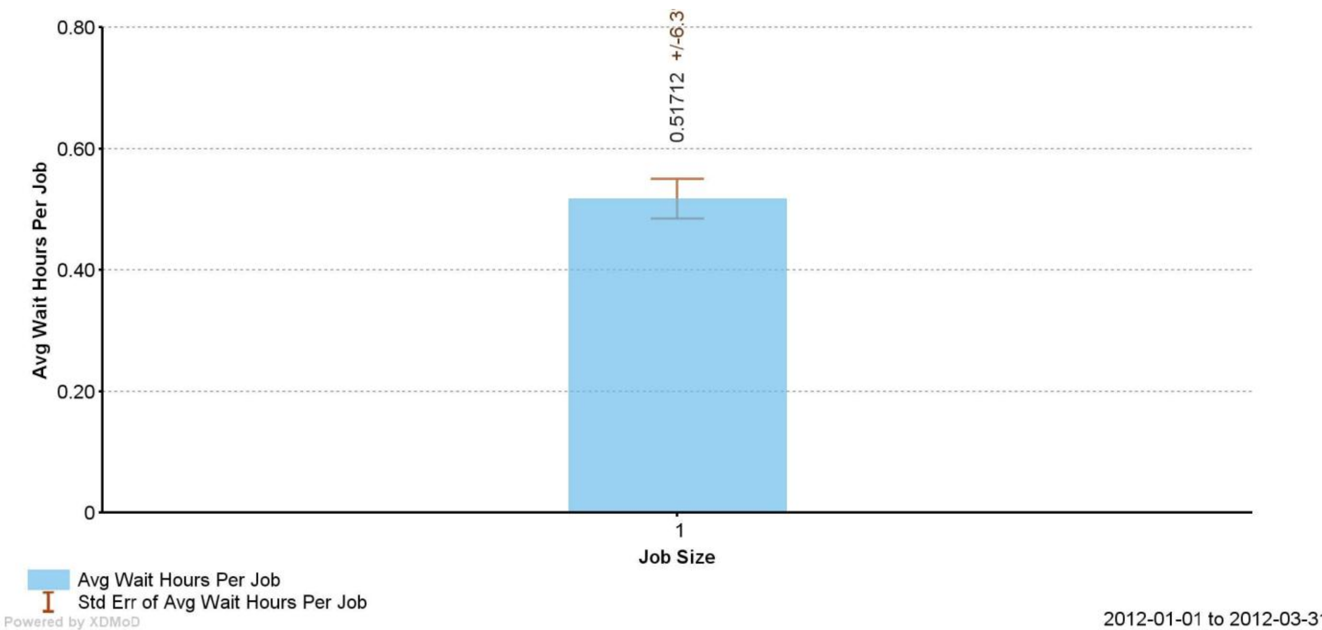

User Expansion Factor by Job Size

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

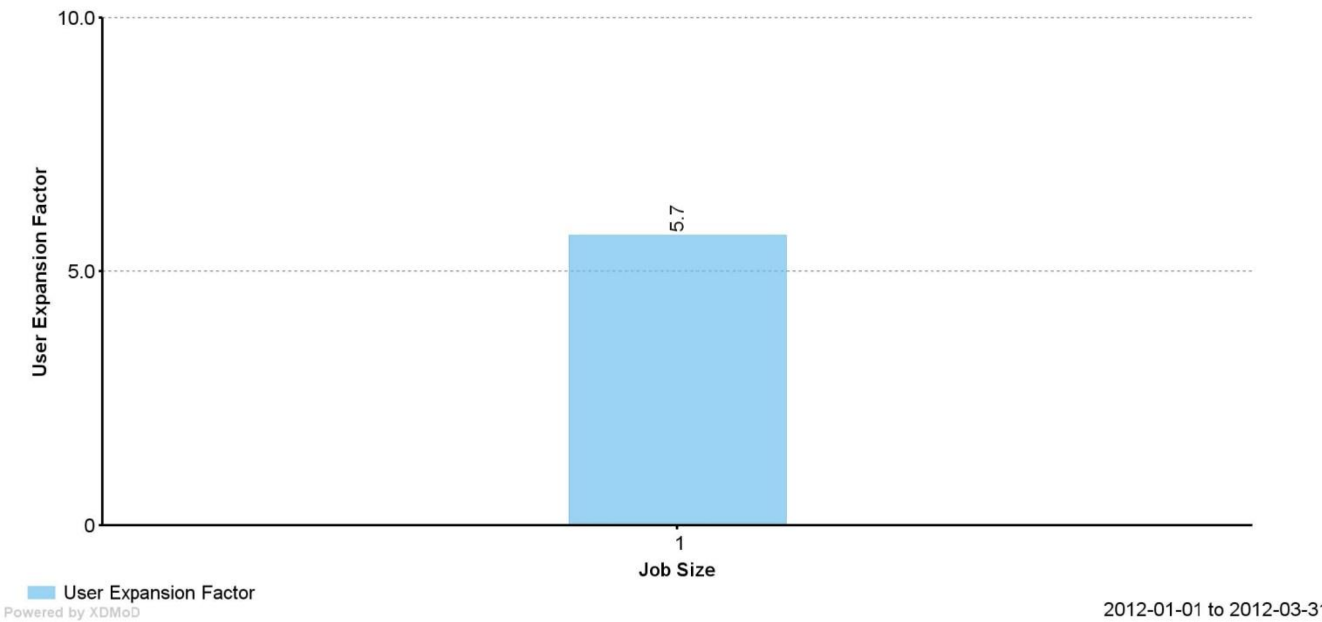

Total NUs Charged by Job Wall Time

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

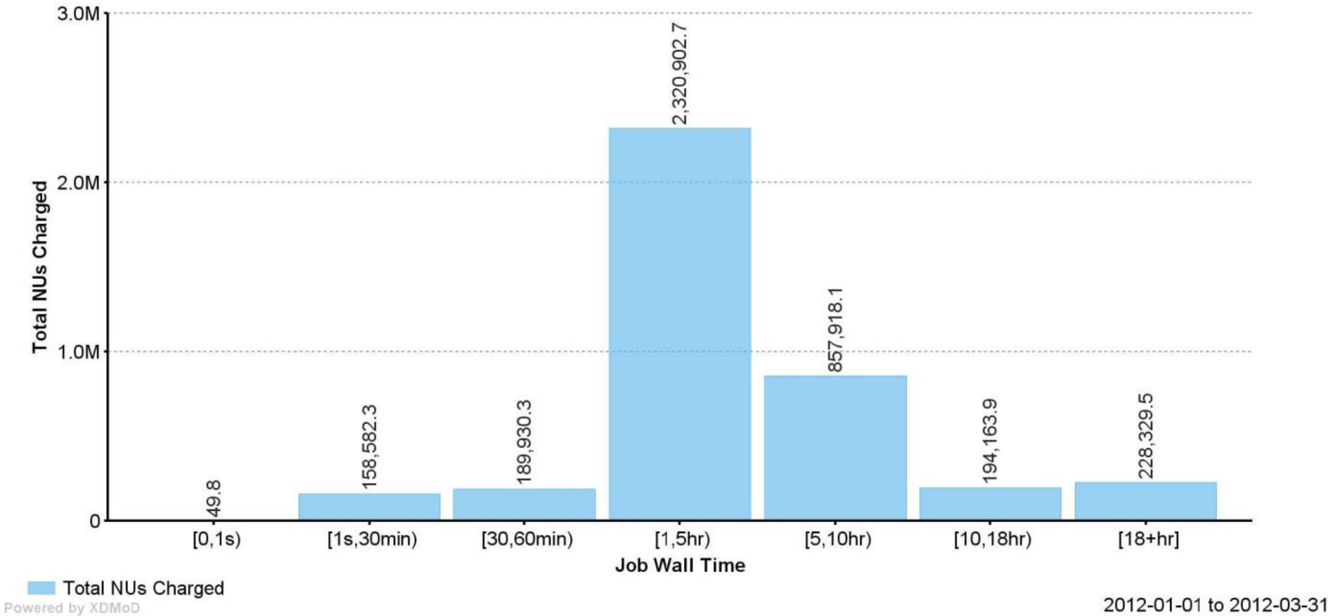

User Expansion Factor by Job Wall Time

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

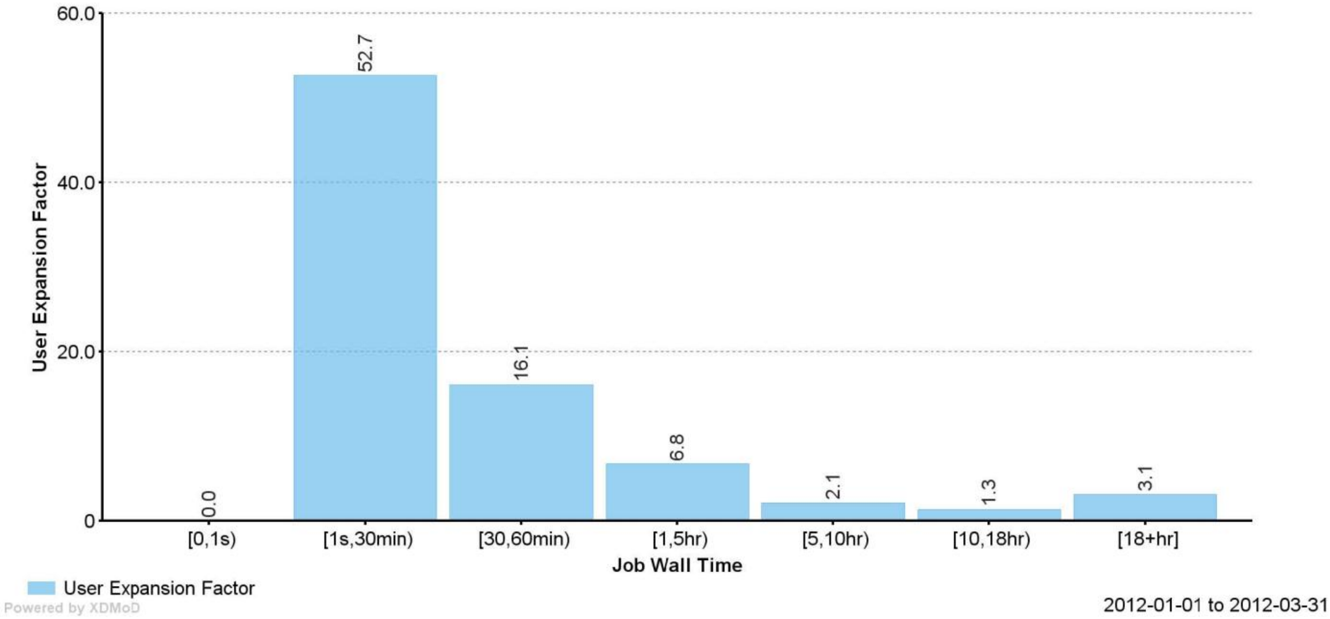

Total NUs Charged by Field of Science

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

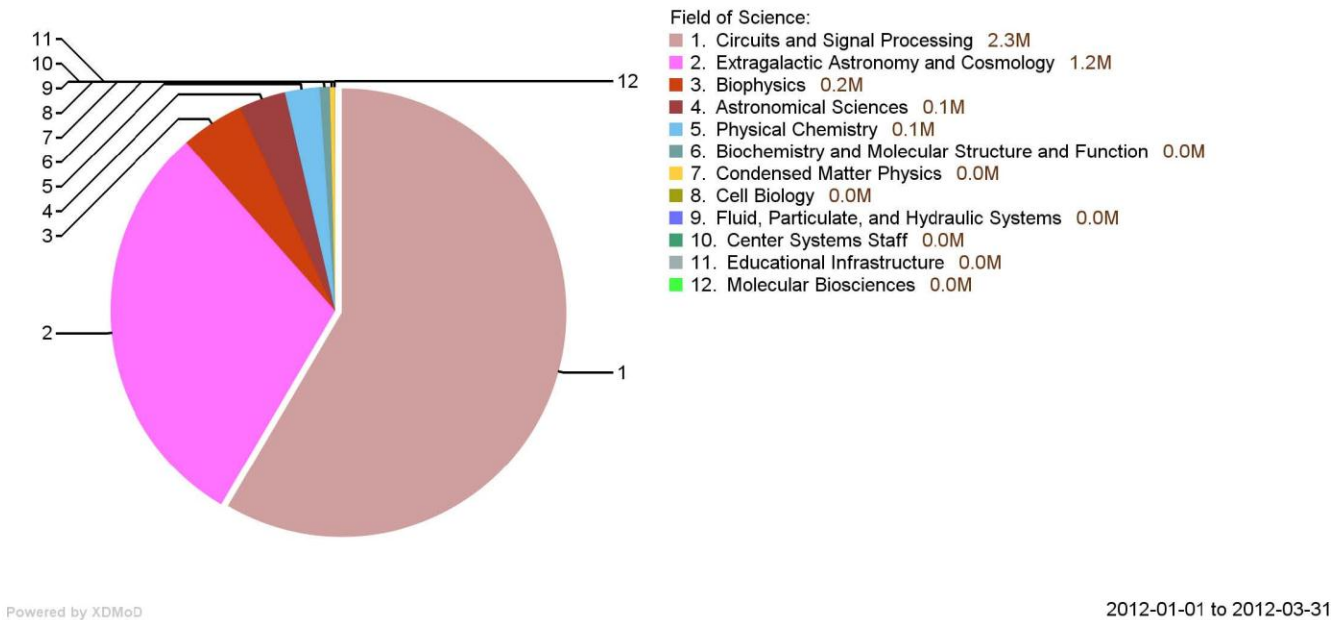

Total NUs Charged by Institution

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

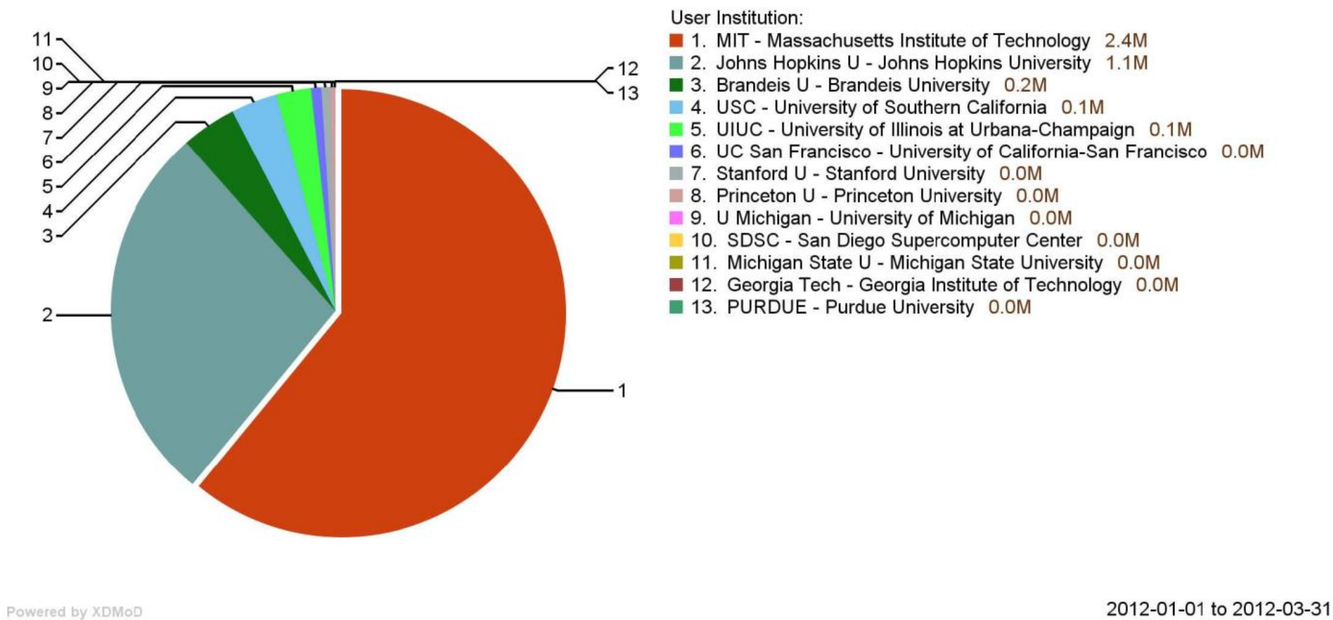

Total NUs Charged by Principal Investigator

Resource = PURDUE-CONDOR 2012-01-01 to 2012-03-31

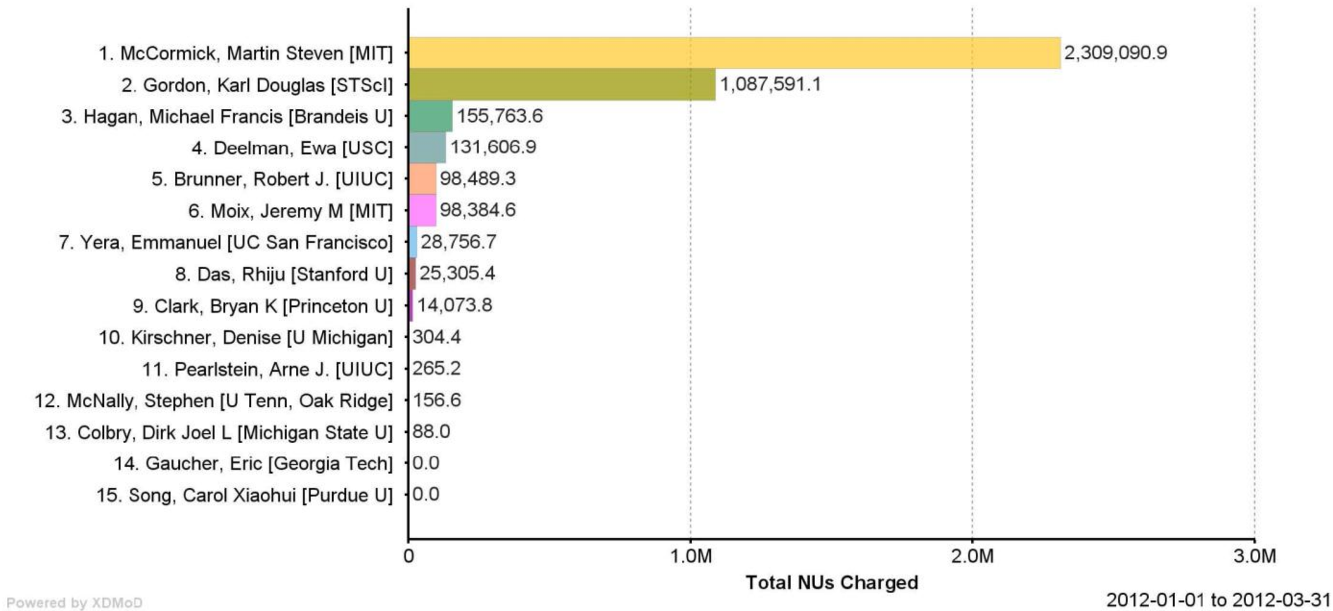

20 San Diego Supercomputer Center (SDSC) Service Provider Quarterly Report

20.1 Executive Summary

After passing a rigorous set of acceptance tests and a 20-day reliability test, *Gordon* was officially accepted on February 24, 2012 and has now been added to the XSEDE resource portfolio as the first system designed specifically for data intensive computing.

The interest in *Gordon* is high, and we are encouraging allocation requests that can make use of its unique architecture features. Eleven projects were awarded just over 11M SU on *Gordon* at the December 2011 XRAC meeting and another 15 awards totaling 20M SUs were made at the March 2012 XRAC. There have also been a number of supplemental and transfer requests. Allocated users were given access to *Gordon* on February 25, and following a brief friendly user period, job accounting became active on March 5. Although only in production for a short period of time during the subject reporting period, XDMoD statistics show that 6,877 jobs were run for a total of 6.8M SUs. Thus far the system has been performing very well with users taking advantage of the flash memory and large memory vSMP "supernodes". In addition, 3 users were awarded startup allocations for dedicated use of Gordon's I/O nodes. This is a prototype allocations model for projects in the areas of data mining and database applications. Two of these projects also received Extended Collaborated Service and Support (ECSS).

Dash has now been decommissioned as an XSEDE resource. Users with active allocations on *Dash* were invited to transfer their allocations to *Gordon*. Over its roughly 18 month life, *Dash* was both an allocable system in the TeraGrid program that let users explore the data intensive features that would be in *Gordon*, and a critical element in the *Gordon* project, that allowed the project team to address risk areas and develop mitigation strategies. This strategy of early prototyping of new technologies is an excellent model that should be considered for future projects where there is a high degree of innovation in the target system. This is especially true in the rapidly developing landscape of data intensive computing where SPs may have more insight in the application challenges than the vendors and are therefore in a better position to integrate the hardware and software elements of the design.

Trestles is now in its second year of production and is successfully supporting the modest-scale/gateway user community. The system is over-requested in the allocation process and there are a significant number of requests for transfers and supplements. This past quarter has seen the system utilization ramp up from the targeted 70% level to now well over 90%; with this increase in utilization, queue wait times and expansion factors, while still good relative to most fully-allocated XSEDE systems, are increasing beyond desired levels. We are analyzing usage data and user job patterns to optimize utilization while maintaining low expansion factors, and to understand the impact of job mix and scheduler tuning that could be applied not only to *Trestles*, but also to other XSEDE production systems. Results of this analysis will be submitted for presentation at XSEDE'12.

SDSC is continuing its evolutionary deployment of the *Data Oasis* Lustre parallel file system. As part of the Gordon acceptance testing, the integrated 4PB system, which supports *Gordon, Trestles* and SDSC's *Triton*, demonstrated 100 GB/s aggregate bandwidth. During the first quarter, Gordon users have had access to a ~2PB/50GBps scratch file system, while Trestles users accessed the 800TB Phase 1 system. During the upcoming quarter, Trestles users will see a new more robust scratch system, and 400 TB of "project storage" will be deployed for non-purged allocated parallel file system storage accessible by both Gordon and Trestles users. An additional 2 PB of project storage will be deployed next year.

We are on track to decommission the GPFS-WAN disk storage system and SDSC's tape archive by June 2012, and are working with users to delete or migrate their data to other resources.

This reporting period saw a large number of effective education, outreach and training activities that engaged users, students and new user communities with conferences, summer workshops and education programs.

20.1.1 Resource Description

Gordon is an Appro-integrated cluster that was co-designed by SDSC and Appro to address the challenges of data intensive computing, which can be summarized simply as, *moving data efficiently through the memory hierarchy*.

When Gordon entered production in February 2012 it became the third largest system in the XSEDE program as measured by Flops and by memory. However, Gordon was designed primarily as a data intensive system and has special features that distinguish it from other XSEDE resources. These features include:

- 1,024, dual socket compute nodes based on Intel's Xeon E5 processor (code named Sandy Bridge). *Gordon* is the first system in XSEDE to use the Sandy Bridge. Each node has two 8-cores, 2.6 GHz processors, with 64 GB of DDR3 memory (64 TB for the full system). The L1, and L2 cache is 32kB, and 256kB, respectively, and there is a large 20 MB shared L3 cache. The nodes use a PCI-Gen3, Intel motherboard. *Gordon* has a peak performance of 341 Tflop/s, with a 286 Tflop/s performance as measured by Linpack.
- 64 flash-based I/O nodes, each with 16, 300 GB Intel 710 SSD's (300 TB of flash for the full system). The drives deliver excellent performance for sequential I/O, but more importantly for data intensive computing, their random performance as measure by input/output operations (IOPS) is over 35K per drive, or 35M for the full system.
- Large memory SMP nodes via ScaleMP's vSMP Foundation software. The minimum size of a supernode 1TB. Nodes of up to 4 TB have been created, with larger sizes possible. For every 16 compute nodes in a supernode, there is a I/O node that also provides 4.8 TB of scratch space.
- A dual rail, QDR, 3D torus network. This is a 4x4x4 torus of switches, with 16 compute notes and one I/O node on each switch. This is a non-proprietary interconnect based on Mellanox infrastructure. As deployed, one rail serves user MPI traffic, while the other is used for I/O to both the Data Oasis Lustre-based file system and the flash drives.
- A high performance, Lustre-based file system called Data Oasis, which has measured performance of 100 GB/s. There is currently 2 PB deployed for Gordon, with an additional 2 PB coming in one year.

The Trestles HPC system is targeted to modest-scale and gateway users with the objective of achieving high-scientific productivity by maintaining short turnaround times, and having responsive scheduling policies and a robust software applications environment. It is an Appro Linux cluster with a peak performance of 100 TFlop/s, allocated 100% to XSEDE users. There are 324 quad-socket compute nodes, with 8-core 2.4 GHz AMD Magny-Cours processors, for a total of 32 cores per node and 10,368 total cores for the system. Each node has 64 GB of DDR3 RAM, with a theoretical memory bandwidth of 171 GB/s. The compute nodes are connected via QDR InfiniBand interconnect, with a fat tree topology, and each link capable of 8 GB/s (bidrectional). *Trestles* uses the *Data Oasis* Lustre parallel file system.

Data Oasis is a shared Lustre parallel file system across multiple HPC systems – *Gordon*, *Trestles,* and SDSC's *Triton* HPC system. Its current configuration and near-term evolution are discussed above. It is primarily funded by SDSC and partially by NSF.

GPFS-WAN (Global Parallel File System-Wide Area Network) is a disk-based storage system that can be mounted by HPC systems at SDSC and across the XSEDE network. GPFS-WAN has ~615TB (usable) disk storage, managed by 16 IBM p575 (Power5) nodes. GPFS-WAN is currently used for XSEDE storage allocations and for testing and evaluation of an XSEDE wide area file system. The resource is no longer available for new allocations, and only previous allocations are being honored until decommissioning, currently scheduled for mid-2012.

Tape archival storage at SDSC consists of IBM Jaguar 3 drives and six Powderhorn silos with a capacity of ~30K cartridges, with data managed under both HPSS and SAM-QFS. SDSC is no longer receiving funding for archival storage and this resource is not offered to *Gordon or Trestles* users. The system is read-only, with ~3PB of TeraGrid/XSEDE legacy data. We have consolidated to a single system by migrating HPSS data to SAM-QFS and, and have initiated disposition of the legacy data to alternative storage with a planned decommissioning June 2012.

20.2 Science Highlights

Gordon officially went into production during the first quarter of 2012. The flash memory-based system is now helping a diverse group of researchers throughout the XSEDE community delve into a wide range of data-intensive projects, from analyzing stock market trades to climate prediction to finding new drugs for diseases. Please see **SDSC's *Gordon* Supercomputer Ready for Researchers** (http://www.sdsc.edu/News%20Items/PR030512_gordon.html) to read more about selected research project highlights.

SDSC also announced during the first quarter plans to host a two-week summer school to help the next generation of astronomers manage the ever-increasing amount of data generated by new instruments, digital sky surveys, and simulations. Participants will be issued user accounts on *Gordon*, which will contain several astronomical data sets and simulations. The July 9-20 course is being held in in conjunction with the University of California's High-Performance AstroComputing Center (UC-HiPACC). Please see **SDSC to Host Summer School on AstroInformatics** (http://www.sdsc.edu/News%20Items/PR030112_astroinformatics.html) for the full story.

SDSC is remaking itself into a center of expertise on all aspects of "big data" research, including genomics, one of the fastest growing areas of scientific study. In addition to new XSEDE resources such as *Gordon*, SDSC has four tiers of specialized data storage (http://www.sdsc.edu/services/StorageBackup.html), which is crucial for genomics and other researchers who need to sift through massive amounts of data. Please see **SDSC's "Big Data" Expertise Aiding Genomics Research** (http://www.sdsc.edu/News%20Items/PR032012_xgen.html) for the full story.

20.3 User-facing Activities

20.3.1 System Activities

The *Gordon* system was brought into production during the reporting period. System management on *Gordon* is handled using Rocks 5.4.3, which configures custom installations on the nodes (compute, IO, and login) based on the CentOS 5.6 Linux distribution. Because the default kernel package for CentOS 5.6 does not support AVX, a custom package based on the 2.6.34.7 Linux kernel is used. For the InfiniBand HCA drivers and subnet manager, the Mellanox release of the 1.5.3 OpenFabrics Enterprise Distribution (OFED) is used, which supports *Gordon*'s ConnectX-3 HCAs. The custom kernel needed for AVX conflicted with OFED in the support of iSER, which provides RDMA access to the SSDs as block devices from the Gordon IO nodes. Mellanox developed a custom version of their OFED release for SDSC, integrating a later version of the iSER code, which supports our kernel. We are now in the process of deploying the updated system configuration with RDMA access to the IO node SSDs. A major upgrade to the Rocks is also underway which will provide the full AVX and iSER support required for *Gordon*.

In March, we carried out the first major *Gordon* software and environment update, providing many new applications to the users. We also utilized nodes from the decommissioned *Dash* cluster to establish a *Gordon* staging and development cluster for applications and systems development, and to isolate configuration changes.

We continue to work very closely with ScaleMP on vSMP upgrades and performance improvements. Notably, we are now testing a version that will significantly improve Lustre performance from an SMP node. SDSD's partnership with ScaleMP and success in deploying vSMP on Gordon has led to additional requests for information from other organizations interested in use vSMP.

The *Trestles* system was stable throughout the reporting period and its utilization is now regularly beyond the target utilization of ~70%, often exceeding 90%. During the reporting period, *Trestles* operated without significant interruption. Due to minor hardware problems or user errors, 76K core-hours (0.3%) were lost.

As previously noted, the *Dash* system was decommissioned during the reporting period.

We have redesigned and are now procuring servers and infrastructure to replace the existing data movers. This will improve transfer rates from all XSEDE sites, and position us for planned transfers of data to support data intensive project, such as those with Blue Waters. We will report on the deployment and benchmarking of this resource in the next quarterly report.

20.3.2 Services Activities

Eleven projects were awarded just over 11M SU on *Gordon* at the December 2011 XRAC meeting. Awardees were: University of Utah; Purdue University; University of Colorado; University of Miami; University of South Florida; UC San Diego; University of Maryland; University of Texas Southwestern Medical Center; University of Oklahoma; University of Illinois at Urbana-Champagne; and the University of Wisconsin Eau Claire. Another 15 awards totaling 20M SUs were made at the March XRAC to these institutions: UC Santa Barbara; University of Texas Health Science Center, San Antonio; Columbia University; Boston University; UC Los Angeles; MIT; Yale University; UC San Diego; University of Washington; Carnegie-Mellon University; University of Chicago (2); Wake Forest University; North Georgia College & State University; and NYU.

SDSC saw over 500 support tickets created and resolved between January 1 and March 31, 2011. These tickets included account questions, file system issues, software requests, globus support, code support, password resets, code optimizations and debugging, allocation refunds/problems, albedo setup, software support, licensing queries, and resource availability.

During the quarter, *Gordon* was opened up for production use by XSEDE allocated users. Over 30 applications were installed as part of the deployment, with new ones coming weekly. Typical *Gordon* support included fielding questions and providing information on how to appropriately submit jobs to the virtual shared memory (vSMP) node, updating documentation, and software install requests. In particular, Gaussian, VASP, and NWCHEM were installed and tested on both the regular compute and vSMP nodes. Support was also provided to users regarding appropriate use of the dual-rail torus architecture of *Gordon*. User services staff members were also involved in testing and performance evaluation of gridftp services to the *Gordon's* scratch filesystem.

In support of *Trestles* users, user services staff worked proactively to modify user scripts to run out of local scratch space (which is flash memory) instead of the *Data Oasis* Lustre filesystem when appropriate (e.g. Gaussian jobs, jobs with many cores writing to a single file). This was very useful in enhancing the performance of the applications codes while lowering the load on the Lustre filesystem. Additionally, user services staff also worked with users to modify job I/O patterns to make better use of the *Data Oasis* filesystem, as well as optimizing code runs to achieve the best performance possible out of *Trestles*.

SDSC staff catered to numerous software requests on *Trestles* including: *bioinformatics codes* – tophat, Columbus, lammps, vasp, amber, namd, visit, and hdf5, as well as many user developed codes that needed help on being ported to the system.

The SDSC support team also handled a large (1 million SU) allocation transfer from *Ranger* to *Trestles* for Dr. Alan Aspuru-Guzik's group researching a computational design of organic semiconductors. The group had approximately 2 months to finish usage of the newly transferred service units, and the scheduler was modified to allow the group to run on 10% of the cores (1024) throughout this period. The research was successfully completed a week ahead of schedule and a science story has been submitted to XSEDE that details the results.

Amit Chourasia worked on installation and integration of of VisIt software on *Gordon* system. Users will be able to submit jobs to Gordon through VisIt application remotely. Documentation was also created for end users to get started.

20.4 Security

No security incidents occurred during this reporting period. However, a moderate effort was required to stay abreast with XSEDE-wide security topics and minor changes to operational logistics.

The effort to develop support for password authentication to XSEDE's Kerberos realm has been successfully completed and the software deployed on *Gordon.* As a result of the effort, SDSC now offers the users of *Gordon* a flexible variety of authentication methods. A fair number of those users have adopted the use of their XSEDE portal password for authentication.

With *Gordon* deployed and in production, we will shift focus to *Trestles* in the coming months. We will perform a security assessment of a proposed software upgrade and hope to apply the software developed and lessons learned during the *Gordon* deployment to *Trestles.*

20.5 Education, Outreach, and Training Activities

20.5.1 Training

Training during the first quarter of 2012 was minimal, as support for early *Gordon* users consumed the bulk of the support team's efforts. However, data from early user support provided valuable feedback to identify topics for a strategic and coordinated training program for *Gordon*, to be launched in the second quarter of this year. At this time, two workshops are planned for May 2012, and the SDSC Summer Institute planning is underway. Whereas last year's institute focused on *Gordon*, this year's will include both *Gordon* and *Trestles*, and will include a component on science gateways. Instructional presentations within the context of conferences are be listed below in the "Presentations" section.

20.5.2 Education

SDSC K-14 Teacher (T) and Student (S) Education and Community Outreach

SDSC hosted 11 workshops, courses, and events for students, K-14 educators, and the general public. The programs varied from a two-hour class to multiple-meeting workshops, to an all-day family activity for 50,000 event attendees. Dates, titles, participant numbers (with gender breakdowns), sites, and audiences (T = teachers; S = students) are listed in the table below.

Date	**Workshop Title**	Participants (F, M)	**Location**	**T/S/C**
Jan. 7	TeacherTECH Science Olympiad Coaching Event	21 (0, 21)	SDSC	T
Jan. 21	SMART Team Hands-on Lab Activities	22 (7, 14)	TSRI	S
Jan. 25	American Teacher documentary screening	90 (55, 35)	SDSC	T
Feb. 4	SMART Team Modeling Work Session	62 (23, 39)	SSB 107	S
March 3	Sea Perch High School Competition	18 (11, 7)	SDSC	T
March 3	Exploring Torrey Pines State Park: A Geology Excursion	17 (5, 12)	Torrey Pines	S
March 10	San Diego Mayor's Cyber Cup Challenge	58 (1, 57)	SDSC	S
March 10	SMART Team Final Presentations	75 (32, 43)	TSRI	T/S/C
March 17	SMART Team Science Festival Event	48 (21, 27)	TSRI	T/S/C
March 24	San Diego Science Festival	50,000	SD Petco Park	C
March 31	Learn MIT's Scratch! Let's Get Programming!	16 (7, 9)	SDSC	S

SDSC's TEOS team launched an NSF-funded Computing Education for a 21st Century Workforce (CE 21) program project in January 2012. The CE21 project, entitled "Computing Principles for ALL Students' Success" (ComPASS) was granted to a partnership led by SDSC, working with San Diego State University, the UCSD Computer Science and Engineering department, and the Computer Science Teachers Association's San Diego chapter. The new project supports teacher professional development to implement a new computer science principles course at the college and pre-college levels. The course is being introduced to pave the way for a new AP (Advanced Placement) test in computer science; one that will focus on the foundation principles of computational thinking. Geared for ALL students, not just computer science majors, this course was taught at five pilot sites across the country in 2010-2011, one at UCSD, taught by Dr. Beth Simon. Undergraduates at UCSD (freshmen) and SDSU (freshmen and seniors planning to become teachers) are now offered the course on basic principles of computing as part of general education (required of all freshmen non-CS majors within UCSD's Sixth College). At the same time, in-service professional development courses for high school and community college teachers in the region will prepare them to teach their own students the same content, but tailored for their students' learning pace and school schedules. The project will be carefully evaluated to determine if this course and its associated pedagogy are successful in reaching the cognitive and attitudinal learning objectives that support broadening participation in computing by women and minorities.

SDSC and UCSD worked closely with four high schools in the region, continuing their work started in the last quarter (Computing Principles for All Students Success (ComPASS) pre-project pilot), meeting weekly to assess progress, and helping the leadership team of teachers who volunteered for the pre-project pilot. Teacher feedback from the fall's pilot program proved extremely valuable in helping SDSC and UCSD project leadership understand the challenges faced by the high school teachers, and the modifications to the college-level course that are necessary for high schools.

During the first quarter of 2012, SDSC TEOS staff worked with the project's curriculum developer to outline an approach for teaching the pedagogical techniques for the new CS Principles course. Evaluation criteria for assessing the effectiveness of the pedagogical training curriculum and educator training were outlined in partnership with the project's education researcher, Dr. Alexander (Sasha) Chizhik, San Diego State University.

SDSC received a $10,000 grant from UCSD California Space Grant Consortium program to support a TeacherTECH series focused on earth and space sciences.

20.5.3 Outreach and Training Presentations

SDSC created outreach materials for the San Diego Science Festival, held annually for the past 4 years, and subsequent to its launch, catalyzing the National Science Festival in Washington D.C. This year, SDSC hosted a booth featuring 3D visualizations from computational science, and offering elementary age children (30% of the 50,000 attendees) an original activity book that followed SDSC Education mascots Chip, Mina, and Max (left) as they explored the world of supercomputing.

This quarter, SDSC announced its annual Research Experience for High School Students (REHS) program featuring 12 SDSC research opportunities for high school students 16 years of age and older. The application deadline was March 31, 2012, and SDSC Education staff are currently reviewing the more than 200 applications. Also announced during the first quarter of 2012 and coordinated by SDSC's Education Manager was another internship program on campus, Career Experiences for High School Students (CEHS), a UCSD campus-wide internship program featuring 11 internship opportunities for high school students 16 years of age and older.

Six SDSC staff served as judges at the annual Greater San Diego Science and Engineering Festival, where they awarded first place and honorable mention awards in three areas of computing, and a first place award in computational science. All awardees have been invited to submit their posters to XSEDE 12, and at least one has reported back that she has submitted her poster. Several of the students who garnered awards had been introduced to computational science and engineering through interactions with SDSC either as interns, or through the on-campus STEM leadership program, COSMOS, which includes an introduction to SDSC.

Posts to the XSEDE blog during the reporting period totaled 63 postings, comprised of a combination of single news posts, the XSEDE weekly Newsroom posts and the monthly XSEDE Education Blog Spot post. We also updated summer internship opportunities on the HPC University site.

The SDSC Education, TeacherTECH, and StudentTECH web sites continue to see the highest activity of all SDSC-specific sites. This semester the TeacherTECH Community Portal is seeing some of the highest activity on record, particularly from a half-dozen or so high school and community college courses. The portal hosts more than 50 active courses in computational STEM topics. The StudentTECH site has seen a lot of activity for its internships and upcoming workshops. The education site itself continues to see much activity related to the instructional 'gems' within its resources. SDSC Education sites are currently being moved to the Cascade CMS, to be consistent with the SDSC main web site.

20.6 SP Collaborations

1.1.1 Collaborations with SP XSEDE Users

SDSC approved a 1M SU allocation transfer from *Ranger* to *Trestles* for Dr. Alan Aspuru-Guzik's group to carry out research on the computational design of organic semiconductors. The group had approximately two months to finish using the newly transferred service units, and their workflow consisted of running thousands of independent two-week, one-core jobs. We created a special Quality of Service (QoS) to restrict the maximum number of running jobs by a user to 1024, thereby enabling a high throughput without the user monopolizing the machine. This allowed the user to run on 1024 cores, about 10% of *Trestles*, while keeping the rest of the system available for other users. The job submission filter was modified to direct long jobs by allocation account of that user to the special QoS. Some tuning was necessary, as the user changed behavior twice during running. The first change was from one-core jobs to two-core jobs. This doubled the amount of total running cores to 2048, around 20% of the resource. In response, we decreased the cap to 512 jobs. The second change was a reduction in duration under the long job limit to 30 hours. This allowed an unlimited number of short jobs to run for that user. We subsequently extended the QoS to all jobs from that user. The research was successfully completed a week ahead of schedule and a science story has been submitted to XSEDE that details the results.

Robert Sinkovits worked with Mao Ye's group (Project: The Lost World: Systematic Truncation of Current Financial Data and Its Impact on Research and Policy Making) to help them make the transition to *Gordon*. A significant amount of data had to be moved from PSC and Mahidhar Tatineni was instrumental in working through data mover issues and assisting with the transfers (as noted above, we will soon be deploying new data movers to assist with projects like this). The software suite was easily ported to *Gordon* and one of the key programs used in the preprocessing of the stock market transaction data was optimized to run from 1.5x to 3.7x faster. The latter speedup was obtained by modifying the code to write binary data rather than CSV (comma separated value) text files. At this point only the more modest set of changes were implemented since writing binary files would require extensive modifications throughout the software suite. Work will begin on the transition to using binary data after a planned set of production runs are completed in April.

Pietro Cicotti continued his collaboration with Michela Taufer (Project: MapReduce clustering on large datasets using SSDs and virtual shared memory) to periodically discuss progress and results, and occasionally meet with the PI and her student, Boyu Zhang. In the past quarter performance bottlenecks that interfered with optimal scaling were identified. Porting the application from Hadoop to MapReduce is expected to improve scaling and the MapReduce library was installed on *Gordon* in anticipation of this work. The results of this collaboration are expected to have a positive impact on other *Gordon* users who want to run Hadoop application and take advantage of *Gordon's* flash drives

Shawn Strande, Mahidhar Tatineni, and Robert Sinkovits met with Francesco Paesani's lab (Project: Computational modeling of atmospheric chemistry: from clusters to aerosols) in March.

The goal of this meeting was for the *Gordon* team and the Paesani lab to learn more about each other's work and develop a plan for making best use of SDSC resources. We initiated a collaboration to improve the performance of the complicated three-body water potentials that are used extensively throughout the simulations. Gains of 1.3x have been realized so far and the Paesani lab has transferred the modified software to the original developers who distribute precompiled binaries to the broader computational chemistry community.

Now that all *Gordon* hardware is available and deployed, work has started on the three dedicated *Gordon* I/O node projects.

- Michael Baitaluk: IntegromeDB: an integrated system and biological search engine
- Phil Bourne and Andreas Prlic: RCSB Protein Data Bank
- Greg Quinn: Real time generation of frame-based protein structure molecular animations for display on mobile devices

A number of preliminary steps have been taken, including: configuration of user environments, database installations (DB2, PostgreSQL, MySQL), account creations, resolution of security issues, enabling direct public access to I/O nodes. The project PIs have been contacted and work should begin in earnest in April.

1.1.2 Collaborations with External Partners

Visualization Collaborations

Amit Chourasia participated in following collaborative activity

- Continued collaborated with Darcy Ogden at SIO/UCSD on visualization of volcano simulations. Visualization results are available at the following website http://visservices.sdsc.edu/projects/volcano.
- Continued collaboration with SCEC on visualization of stress tensors.
- New collaborative work with John Helly on visualization of Climate simulation data

Gordon 3D Torus

Gordon's dual rail, 3D torus network is unique among XSEDE systems. To further leverage the capabilities for data intensive computing, SDSC is collaborating with D. K. Panda's group at Ohio State University to test and deploy new versions of the MVAPICH2 library that can take advantage of this topology. In addition to the support already provided by MVAPICH2 (support for the 3D torus with appropriate Service Levels, integrated multi-rail communications support, flexible selection of InfiniBand adapter and ports, and others) the new version will provide support for network fault resilience for multi-rail, optimized collectives for 3D-torus topologies, and support for upcoming MPI-3 standard and features.

Benchmarking of bioinformatics codes on Gordon

Wayne Pfeiffer installed five bioinformatics codes on Gordon and ran benchmarks to compare their performance on Gordon with that obtained on other computers at SDSC. The codes considered were two assemblers -- SOAPdenovo and Velvet -- and three phylogenetics codes -- RAxML, RAxML-Light, and Parsimonator. The speedup from Trestles to Gordon was:

- Approximately 1.5 for SOAPdenovo, Velvet, and RAxML
- 1.75 for RAxML-Light, and

- 2.5 for Parsimonator

Wayne Pfeiffer made benchmark runs of the Velvet assembler using vSMP on Gordon for a full Daphnia genome. The run times were then compared to those on Blacklight at PSC and on Triton PDAF at UCSD obtained by Mahidhar Tatineni and Pfeiffer, respectively.

For the Daphnia assembly, which uses about 140 GB of memory for the graph step, but only 43 GB for the hash step, Gordon is 1.15x faster than Blacklight and 1.20x faster than Triton PDAF. This comparison demonstrates that vSMP on Gordon is needed only for the graph step and works well.

20.7 SP-Specific Activities

A number of new capabilities were deployed in the SDSC scheduling system to meet specialized needs on *Trestles* and *Gordon*. Topology scheduling was deployed to efficiently schedule compute nodes connected in a 3D torus. This allows users to specify the degree of connectivity required for their jobs, with those jobs being assigned to compute nodes according to relative switch topology. To serve multiple users of limited software licenses, a method for incorporating licenses into the scheduling process, along with core and memory allocation, was deployed. To allow for preventative maintenance and the occasional unplanned unavailability of individual scheduling servers, failover capability, through the NFS filesystem, was added to the batch system. This allows a pool of scheduling servers to run, with any individual server able to continue scheduling if any of the others become unavailable. Instances of the scheduler run on each of the *Trestles* and *Gordon* login nodes.

XSEDE software stack-related activities included a compilation of openssh with both Kerberos and GSI authentication mechanisms. This NCSA guide proved very useful:

http://grid.ncsa.illinois.edu/gssapi-mechglue/openssh/

MDS info services, WSRF job and load information services, and GLUE2 resource information services were deployed on *Gordon* using the XSEDE Baseline Software document. These software packages were repackaged to conform to site security and systems policies.

We have written a paper describing the capabilities and implementations of on-demand urgent computing and user-settable reservations on *Trestles*. This paper was accepted to the Urgent Computing Workshop in Omaha in June 2012, as part of the International Conference on Computational Science, ICCS 2012.

SDSC has coordinated with NSF and had ongoing communications with users about decommissioning its tape archive by June 2012. The archive contains legacy data from earlier generations of SDSC HPC resources, and users are being asked to delete or migrate their data to other resources by June 2012.

We are also communicating with users about decommissioning the GPFS-WAN file system mid-2012. This has been an allocated resource and nearly all allocations will have expired by the time the system is decommissioned in June 2012; nevertheless it represents further shrinkage of the space available for online storage (along with IU's Data Capacitor) and the only remaining allocated resource of this type will soon be *Albedo*.

20.8 Publications

Y. Bazilevs, A.L. Marsden, F. Lanza di Scalea, A. Majumdar, and M. Tatineni, "Toward a Computational Steering Framework for Large-Scale Composite Structures Based on Continually and Dynamically Injected Sensor Data," International Conference on Computational Science, Omaha, Nebraska, June 4-6, 2012.

Y. Cui, K. B. Olsen, J. Zhou, P. Small, A. Chourasia, S. M. Day, P. Maechling and T. H. Jordan (2012). Development and Optimizations of a SCEC Community Anelastic Wave Propagation Platform for Multicore Systems and GPU-based Accelerators. To be presented at Seismological Society of America 2012 Meeting, San Diego, Apr 2012

Dodson, M.L., Walker, R.C., Lloyd, R.S., "Carbinolamine Formation and Dehydration in a DNA Repair Enzyme Active Site", PLoS ONE 2012, 7(2): e31377. doi:10.1371/journal.pone.0031377

Goetz, Andreas W; Williamson, Mark J; Xu, Dong; Poole, Duncan; Grand, Scott Le; Walker, Ross C. "Routine microsecond molecular dynamics simulations with amber - part i: Generalized born", 2012. Journal of Chemical Theory and Computation., 2012, in press. DOI:10.1021/ct200909j

Nurisso, A., Daina, A., Walker, R.C., "A Practical Introduction to Molecular Dynamics Simulations: Applications to Homology Modeling", ed. A. Orry, Springer, Methods in Molecular Biology Series, 2012, vol 857, ISBN 978-1-61779-587-9. (Book chapter)

S. Seale, J. H. Steidl, L. H. Seale and A. Chourasia (2012). Visualizing Structural Response and Site Amplification Using Earthquake Data Recorded at the NEES@UCSB Field Site. To be presented at Seismological Society of America 2012 Meeting, San Diego, Apr 2012.

K. Yoshimoto, D. Choi, R. Moore, A. Majumdar, E. Hocks, "Implementations of Urgent Computing on Production HPC Systems", International Conference on Computational Science, ICCS 2012.

20.9 Metrics

20.9.1 Standard systems metrics

SDSC Gordon Quarterly Report

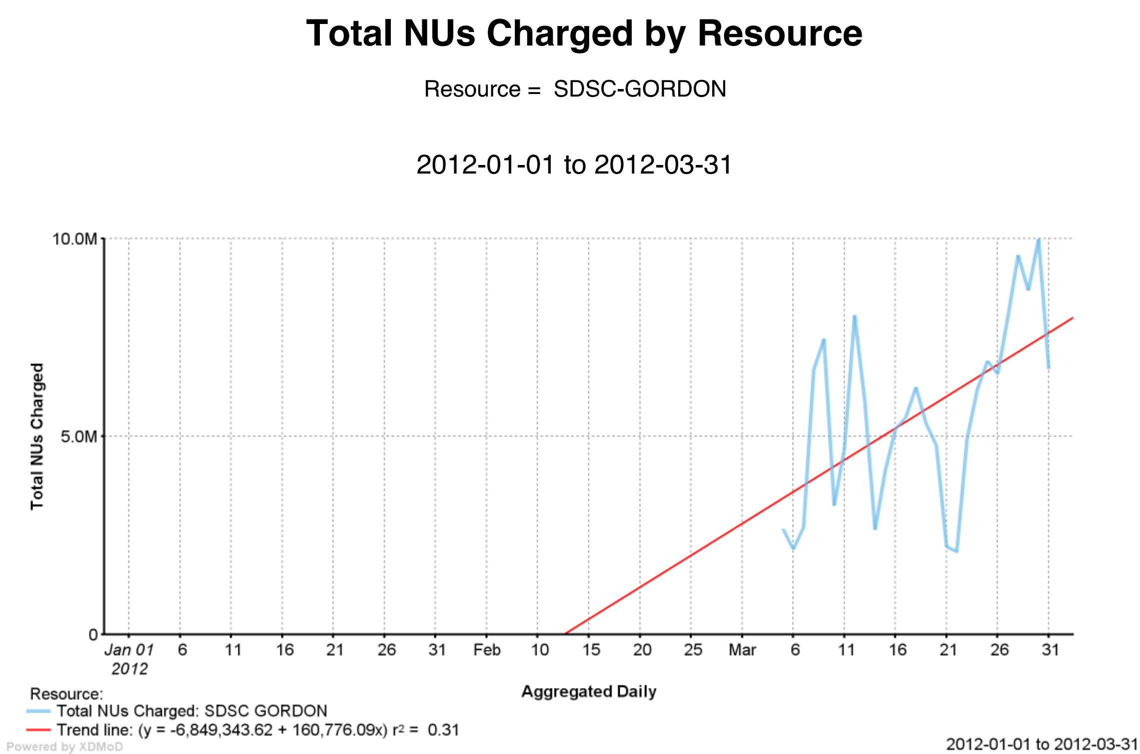

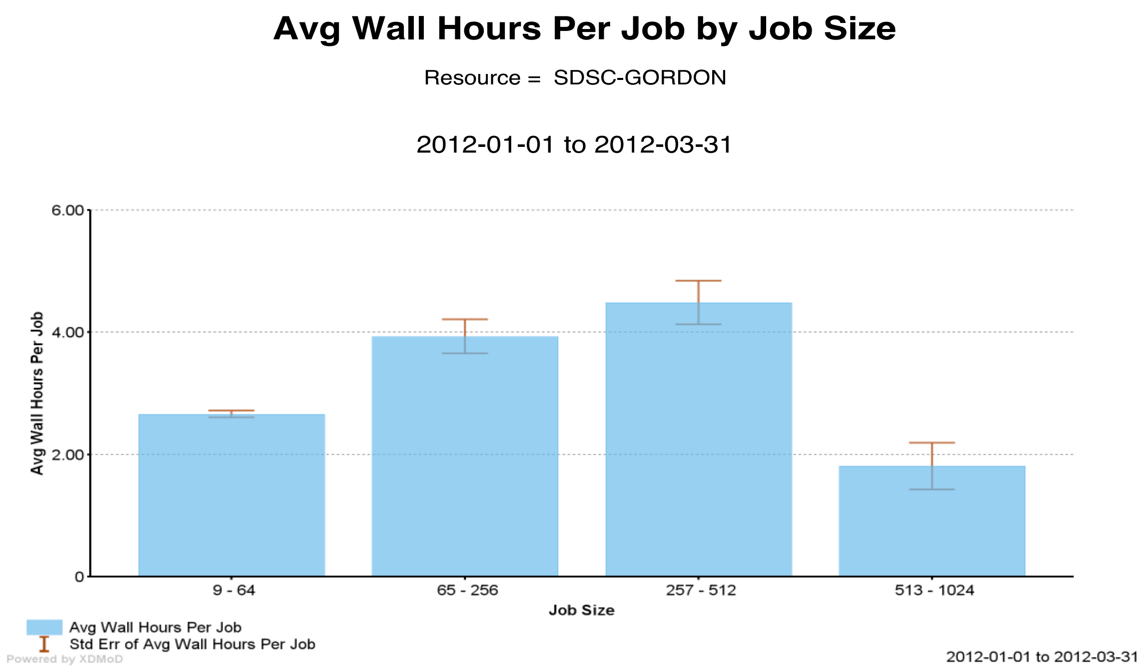
Avg Wall Hours Per Job by Job Size
Resource = SDSC-GORDON
2012-01-01 to 2012-03-31
6.00
4.00
2.00
0
Avg Wall Hours Per Job
9 - 64
65 - 256
257 - 512
513 - 1024
Job Size
Avg Wall Hours Per Job
Std Err of Avg Wall Hours Per Job
Powered by XDMoD
2012-01-01 to 2012-03-31

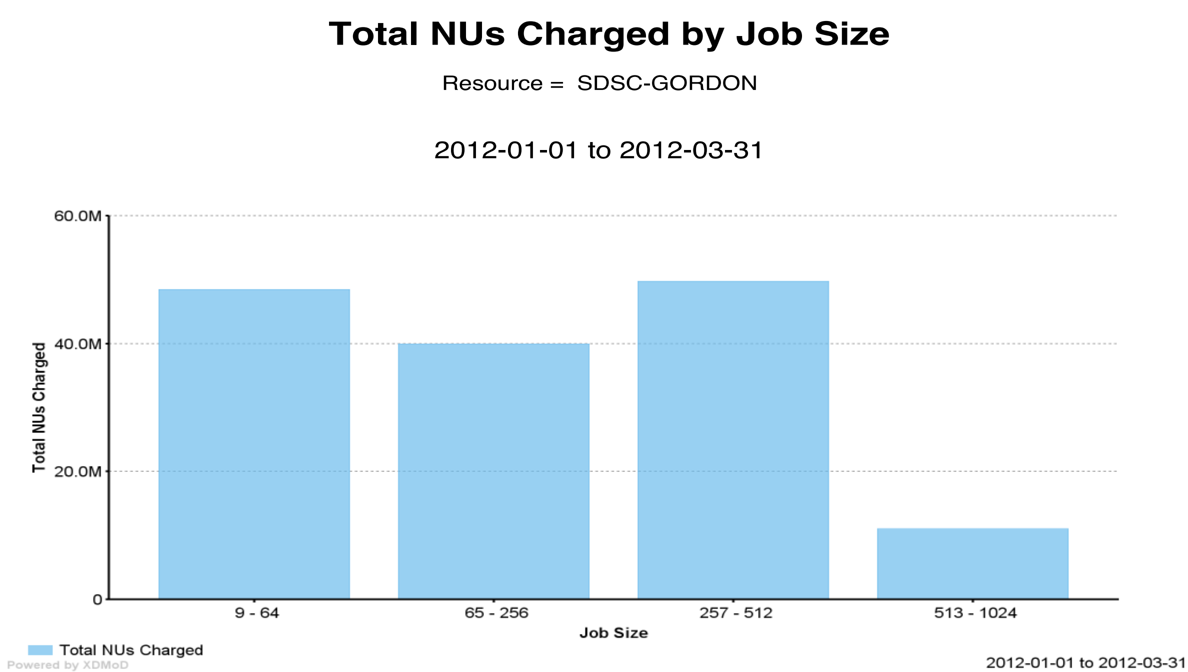
Total NUs Charged by Job Size
Resource = SDSC-GORDON
2012-01-01 to 2012-03-31
60.0M
40.0M
20.0M
0
Total NUs Charged
9 - 64
65 - 256
257 - 512
513 - 1024
Job Size
Total NUs Charged
Powered by XDMoD
2012-01-01 to 2012-03-31

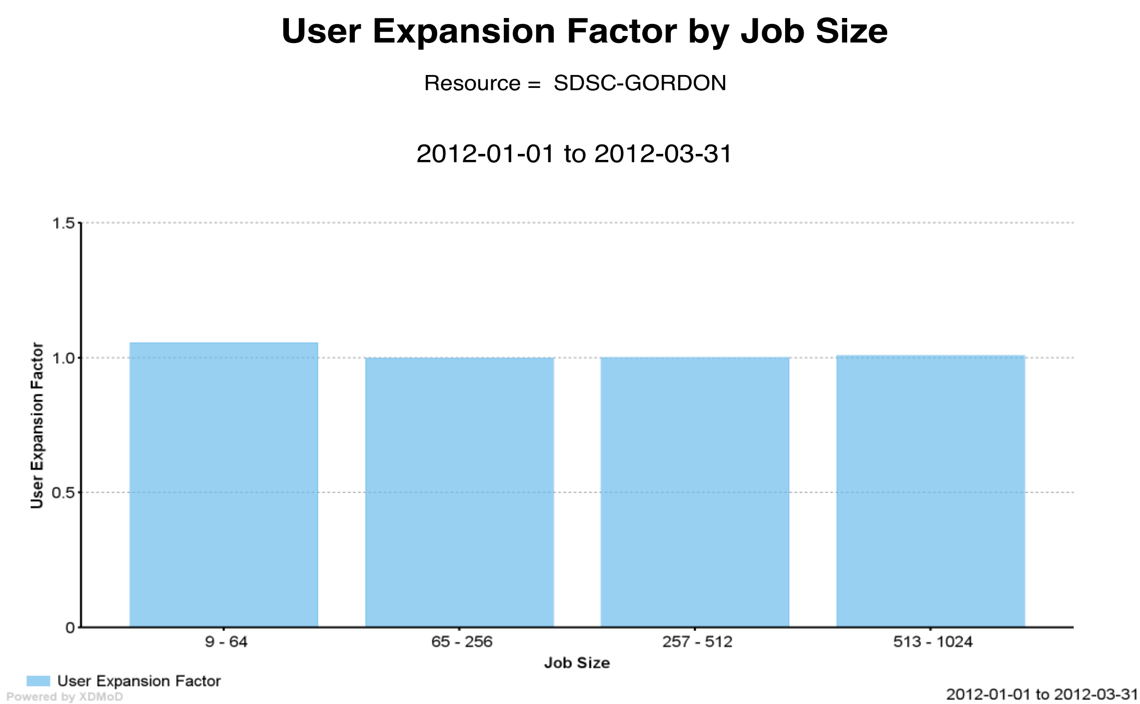

User Expansion Factor by Job Size
Resource = SDSC-GORDON
2012-01-01 to 2012-03-31
1.5
1.0
0.5
0
User Expansion Factor
9 - 64
65 - 256
257 - 512
513 - 1024
Job Size
User Expansion Factor
Powered by XDMoD
2012-01-01 to 2012-03-31

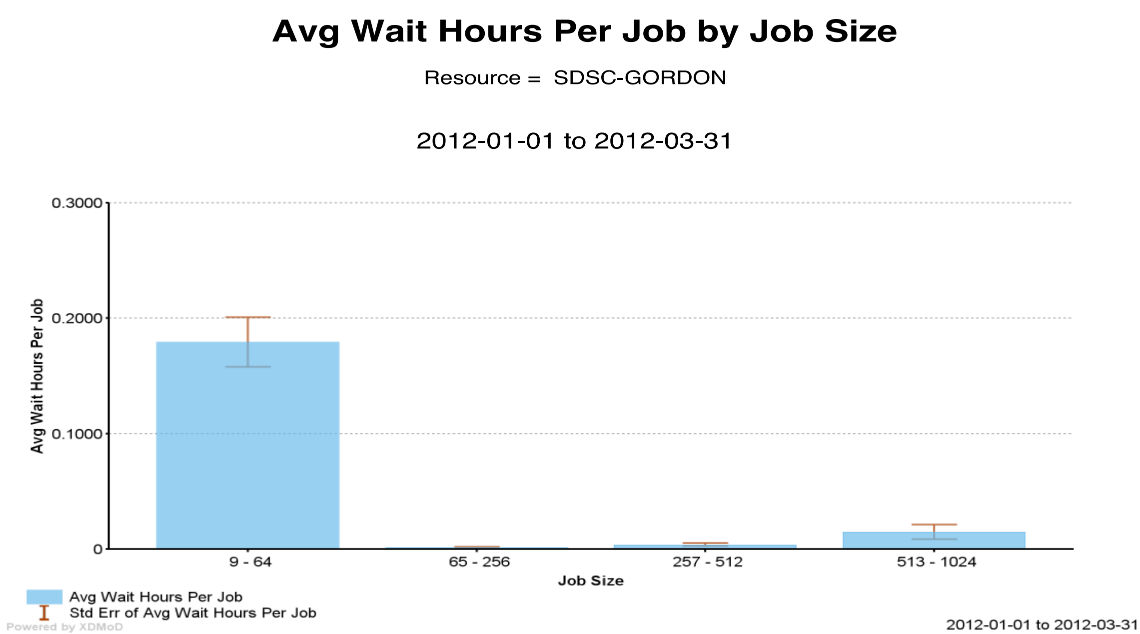

Avg Wait Hours Per Job by Job Size
Resource = SDSC-GORDON
2012-01-01 to 2012-03-31
0.3000
0.2000
0.1000
0
Avg Wait Hours Per Job
9 - 64
65 - 256
257 - 512
513 - 1024
Job Size
Avg Wait Hours Per Job
Std Err of Avg Wait Hours Per Job
Powered by XDMoD
2012-01-01 to 2012-03-31

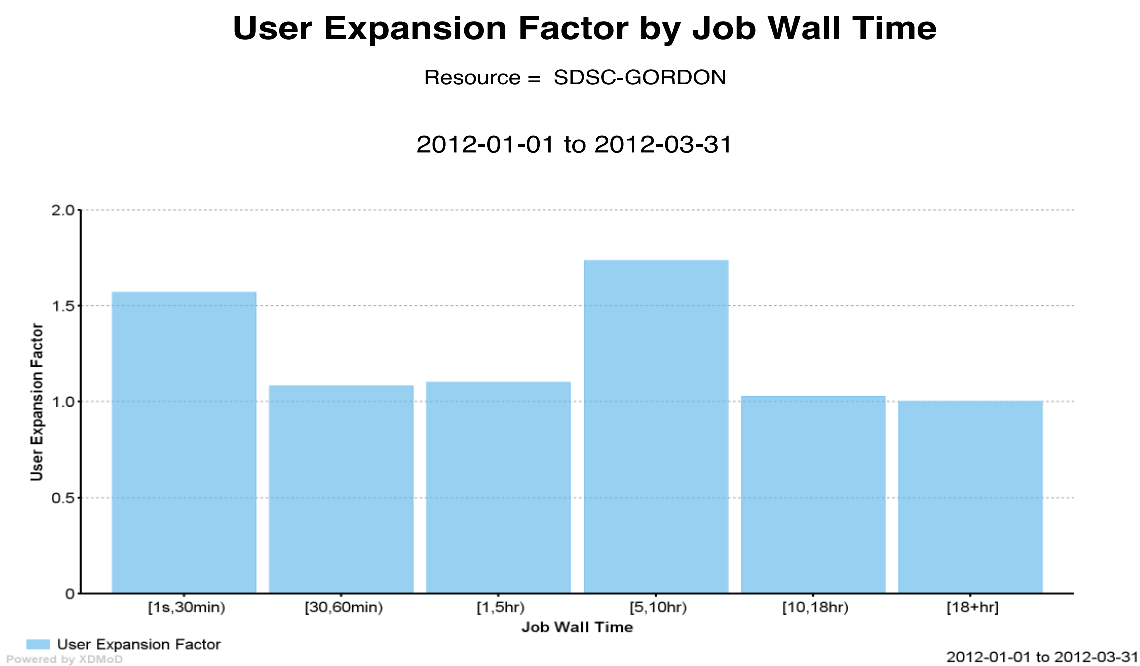

User Expansion Factor by Job Wall Time
Resource = SDSC-GORDON
2012-01-01 to 2012-03-31
2.0
1.5
1.0
0.5
0
User Expansion Factor
[1s,30min)
[30,60min)
[1,5hr)
[5,10hr)
[10,18hr)
[18+hr]
Job Wall Time
User Expansion Factor
Powered by XDMoD
2012-01-01 to 2012-03-31

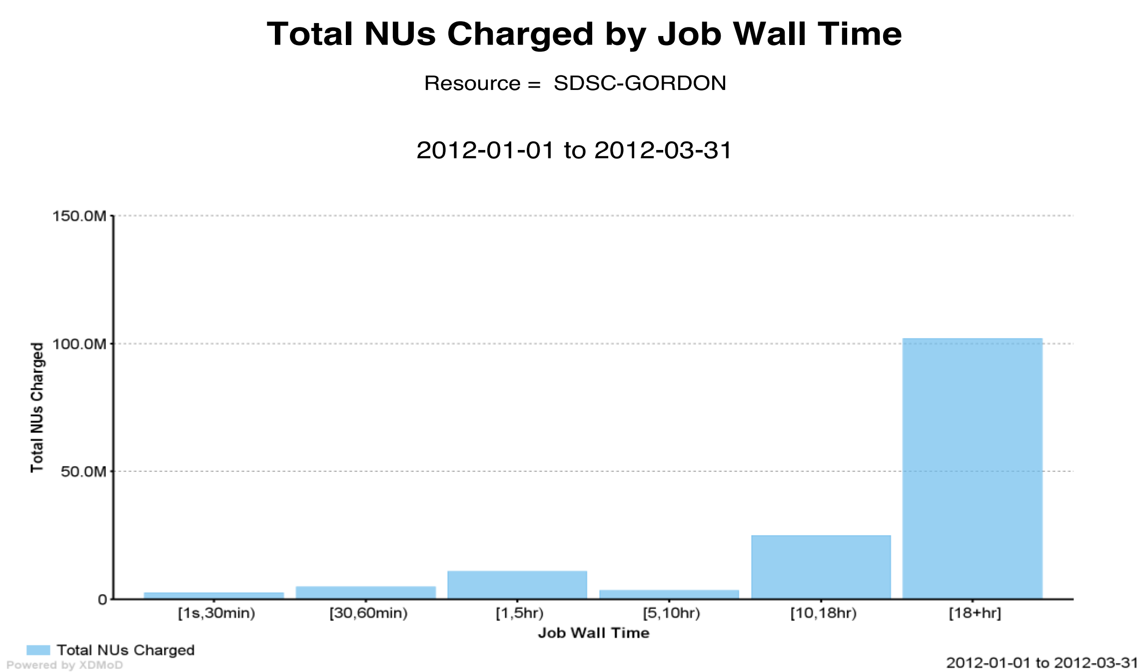

Total NUs Charged by Job Wall Time
Resource = SDSC-GORDON
2012-01-01 to 2012-03-31
150.0M
100.0M
50.0M
0
Total NUs Charged
[1s,30min)
[30,60min)
[1,5hr)
[5,10hr)
[10,18hr)
[18+hr]
Job Wall Time
Total NUs Charged
Powered by XDMoD
2012-01-01 to 2012-03-31

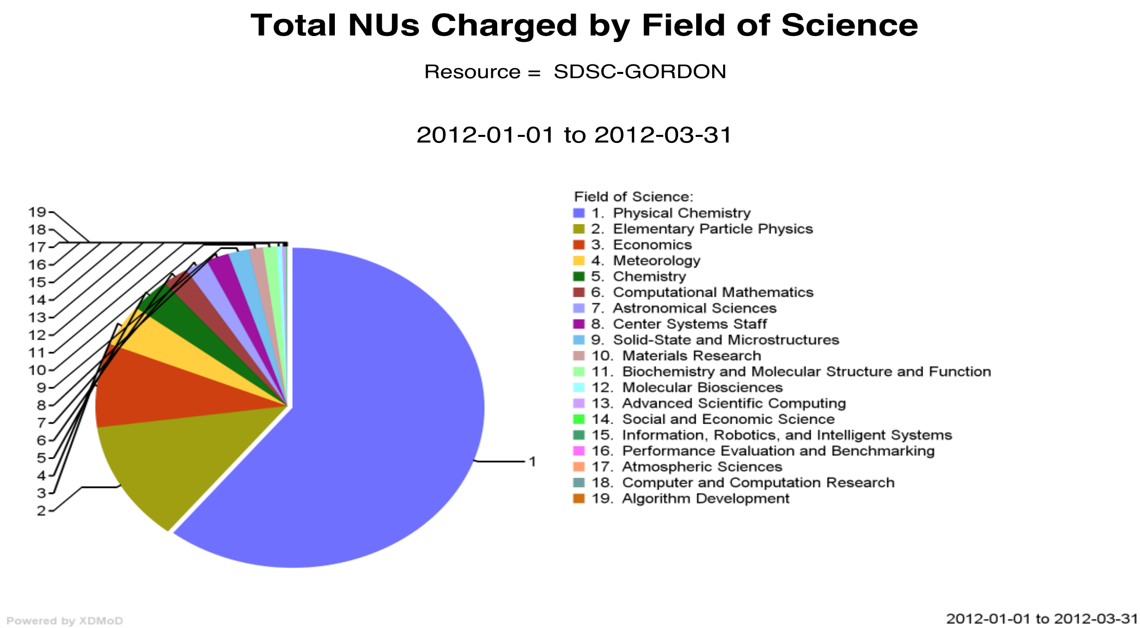
Total NUs Charged by Field of Science
Resource = SDSC-GORDON
2012-01-01 to 2012-03-31
Field of Science:
1. Physical Chemistry
2. Elementary Particle Physics
3. Economics
4. Meteorology
5. Chemistry
6. Computational Mathematics
7. Astronomical Sciences
8. Center Systems Staff
9. Solid-State and Microstructures
10. Materials Research
11. Biochemistry and Molecular Structure and Function
12. Molecular Biosciences
13. Advanced Scientific Computing
14. Social and Economic Science
15. Information, Robotics, and Intelligent Systems
16. Performance Evaluation and Benchmarking
17. Atmospheric Sciences
18. Computer and Computation Research
19. Algorithm Development
Powered by XDMoD
2012-01-01 to 2012-03-31

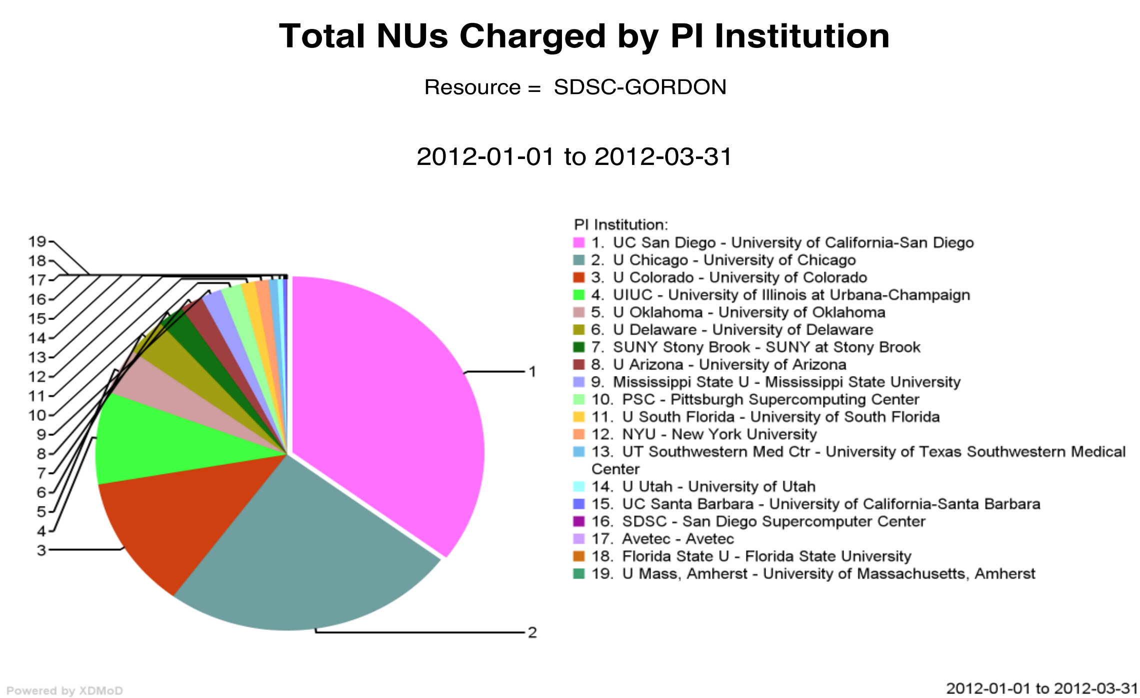
Total NUs Charged by PI Institution
Resource = SDSC-GORDON
2012-01-01 to 2012-03-31
PI Institution:
1. UC San Diego - University of California-San Diego
2. U Chicago - University of Chicago
3. U Colorado - University of Colorado
4. UIUC - University of Illinois at Urbana-Champaign
5. U Oklahoma - University of Oklahoma
6. U Delaware - University of Delaware
7. SUNY Stony Brook - SUNY at Stony Brook
8. U Arizona - University of Arizona
9. Mississippi State U - Mississippi State University
10. PSC - Pittsburgh Supercomputing Center
11. U South Florida - University of South Florida
12. NYU - New York University
13. UT Southwestern Med Ctr - University of Texas Southwestern Medical Center
14. U Utah - University of Utah
15. UC Santa Barbara - University of California-Santa Barbara
16. SDSC - San Diego Supercomputer Center
17. Avetec - Avetec
18. Florida State U - Florida State University
19. U Mass, Amherst - University of Massachusetts, Amherst
Powered by XDMoD
2012-01-01 to 2012-03-31

Total NUs Charged by PI

Resource = SDSC-GORDON

2012-01-01 to 2012-03-31

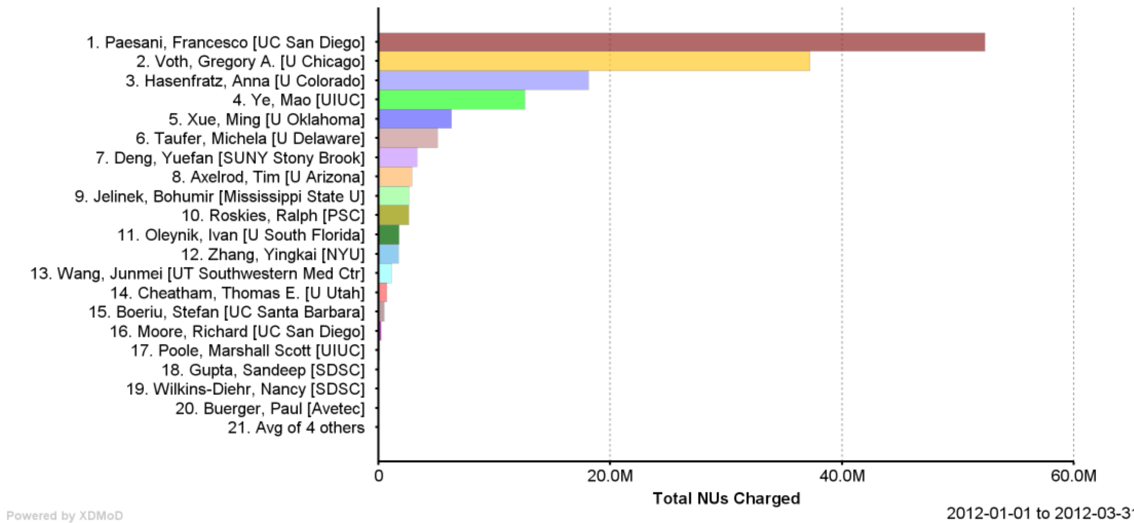

SDSC Trestles Quarterly Report

Total NUs Charged by Resource

XSEDE

2012-01-01 to 2012-03-31

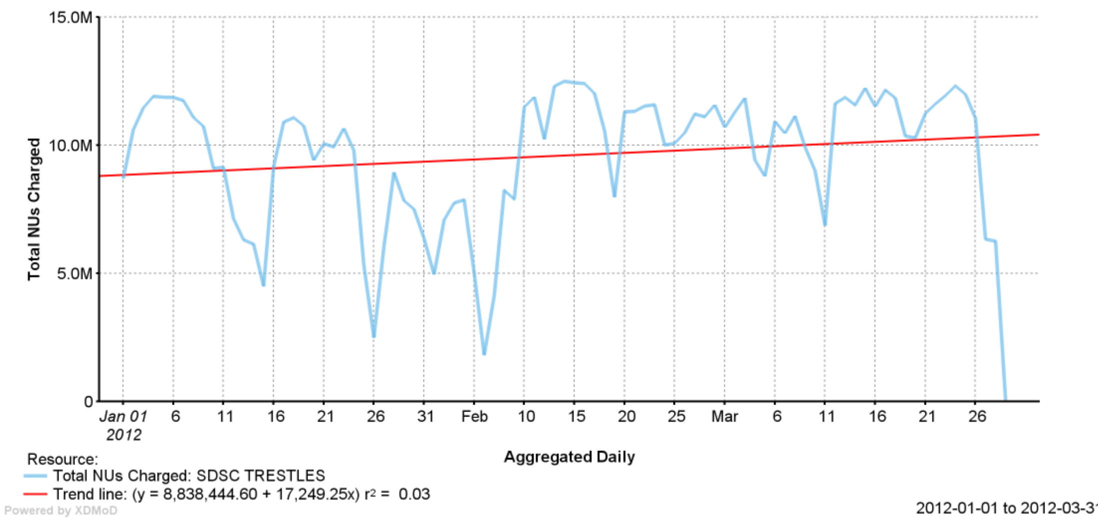

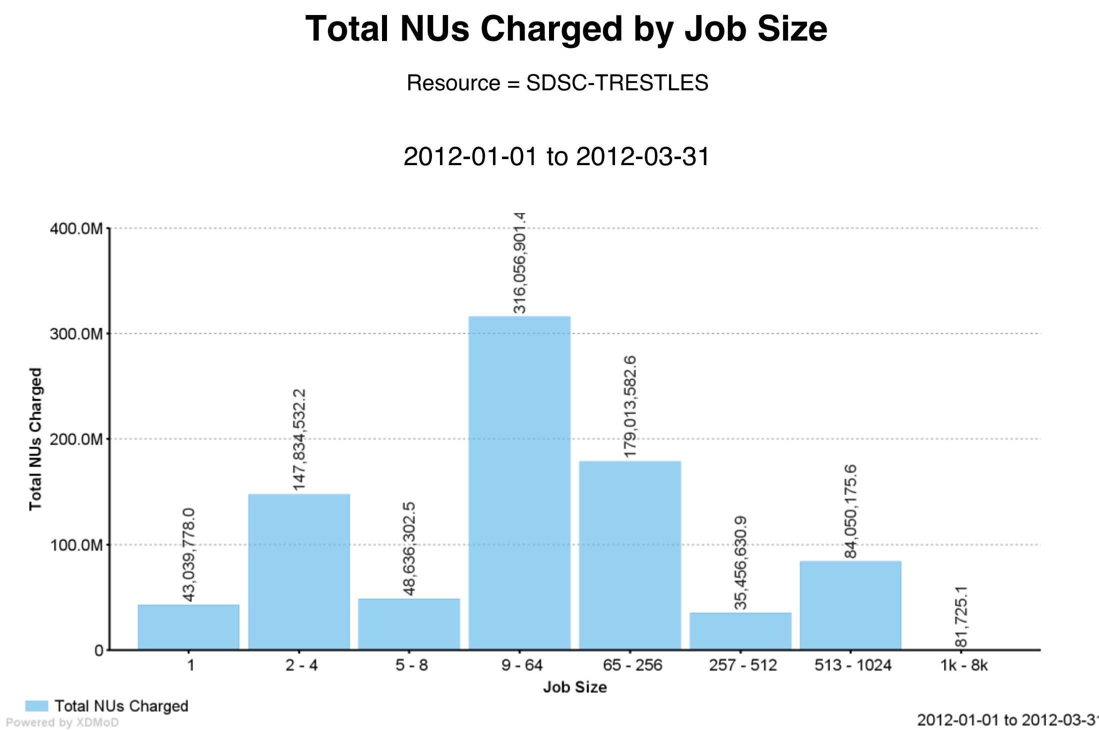
Total NUs Charged by Job Size
Resource = SDSC-TRESTLES
2012-01-01 to 2012-03-31
Total NUs Charged
400.0M
300.0M
200.0M
100.0M
0
43,039,778.0
147,834,532.2
48,636,302.5
316,056,901.4
179,013,582.6
35,456,630.9
84,050,175.6
81,725.1
1
2 - 4
5 - 8
9 - 64
65 - 256
257 - 512
513 - 1024
1k - 8k
Job Size
Total NUs Charged
Powered by XDMoD
2012-01-01 to 2012-03-31

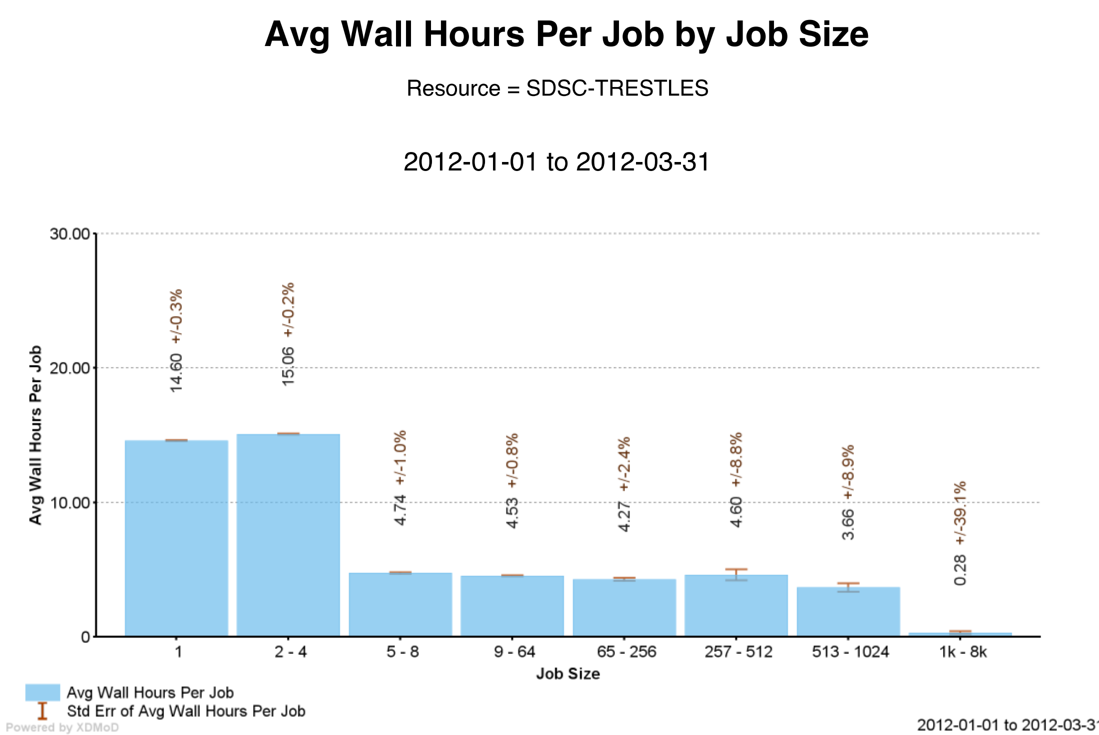
Avg Wall Hours Per Job by Job Size
Resource = SDSC-TRESTLES
2012-01-01 to 2012-03-31
Avg Wall Hours Per Job
30.00
20.00
10.00
0
14.60 +/-0.3%
15.06 +/-0.2%
4.74 +/-1.0%
4.53 +/-0.8%
4.27 +/-2.4%
4.60 +/-8.8%
3.66 +/-8.9%
0.28 +/-39.1%
1
2 - 4
5 - 8
9 - 64
65 - 256
257 - 512
513 - 1024
1k - 8k
Job Size
Avg Wall Hours Per Job
Std Err of Avg Wall Hours Per Job
Powered by XDMoD
2012-01-01 to 2012-03-31

Avg Wait Hours Per Job by Job Size

Resource = SDSC-TRESTLES

2012-01-01 to 2012-03-31

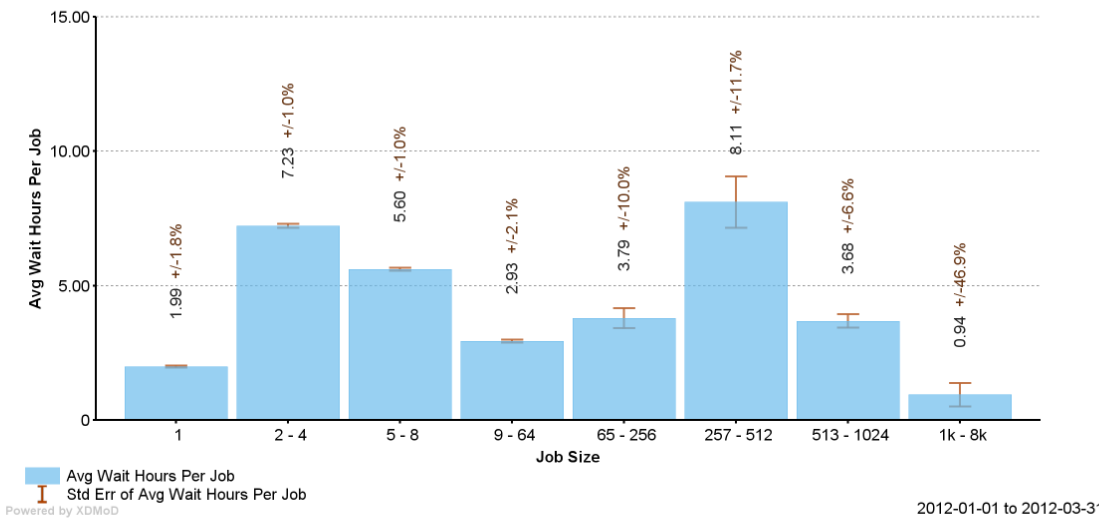

User Expansion Factor by Job Size

Resource = SDSC-TRESTLES

2012-01-01 to 2012-03-31

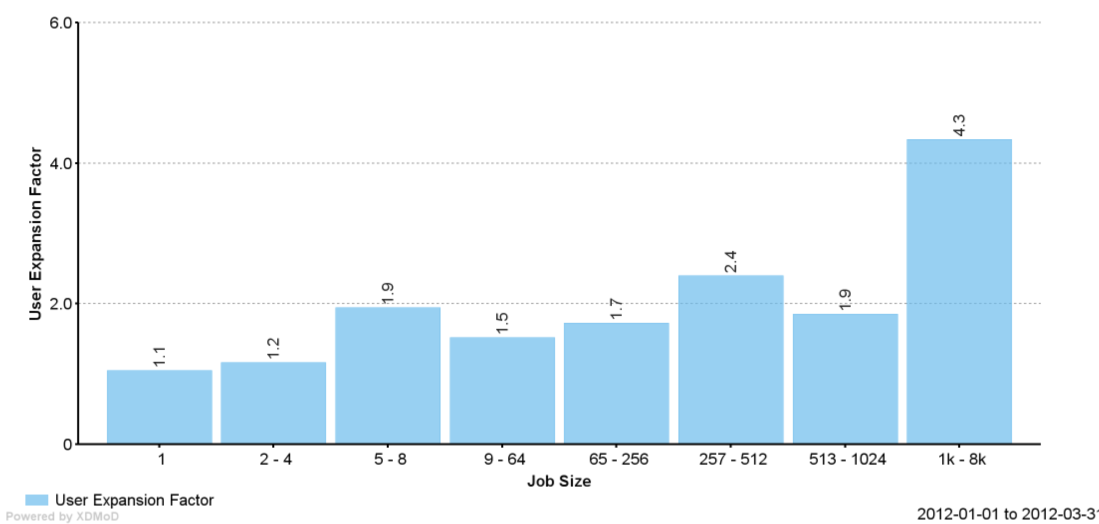

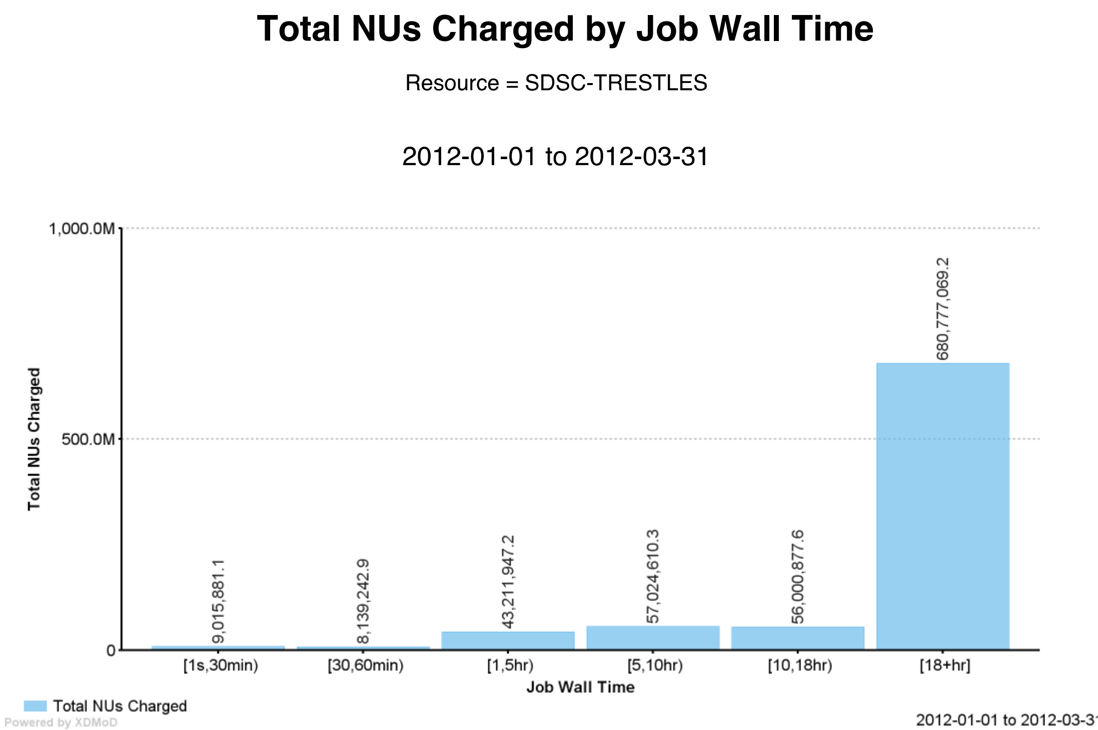
Total NUs Charged by Job Wall Time
Resource = SDSC-TRESTLES
2012-01-01 to 2012-03-31
Total NUs Charged
1,000.0M
500.0M
0
9,015,881.1
8,139,242.9
43,211,947.2
57,024,610.3
56,000,877.6
680,777,069.2
[1s,30min)
[30,60min)
[1,5hr)
[5,10hr)
[10,18hr)
[18+hr]
Job Wall Time
Total NUs Charged
Powered by XDMoD
2012-01-01 to 2012-03-31

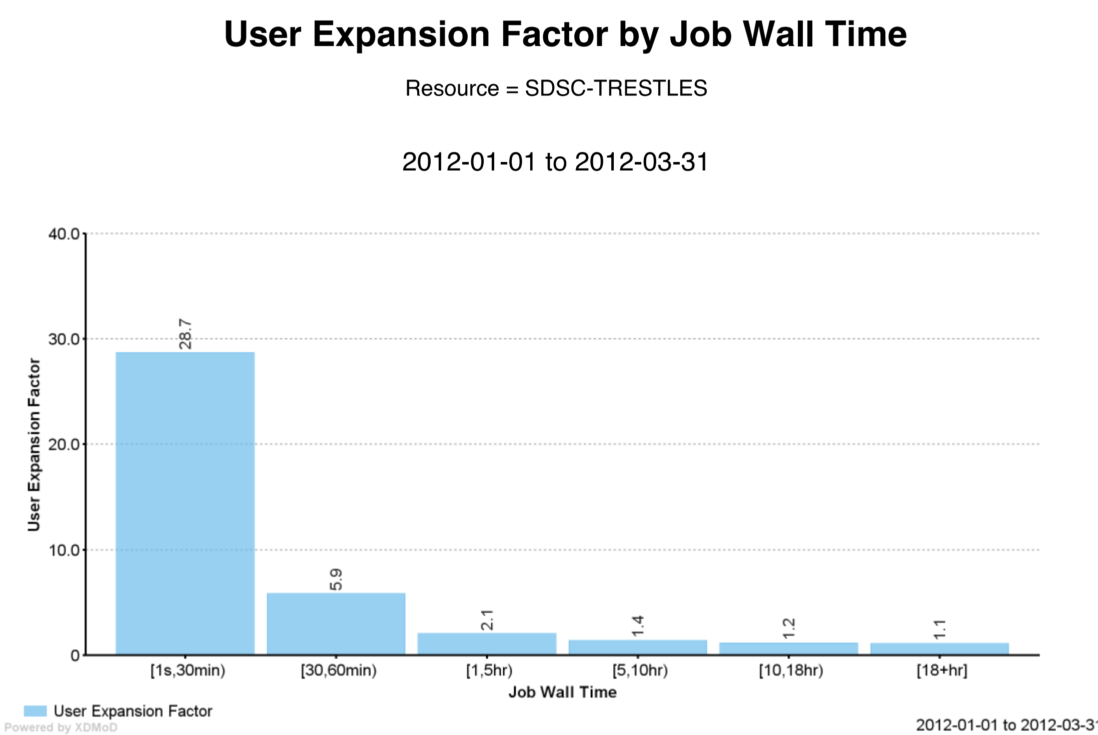
User Expansion Factor by Job Wall Time
Resource = SDSC-TRESTLES
2012-01-01 to 2012-03-31
User Expansion Factor
40.0
30.0
20.0
10.0
0
28.7
5.9
2.1
1.4
1.2
1.1
[1s,30min)
[30,60min)
[1,5hr)
[5,10hr)
[10,18hr)
[18+hr]
Job Wall Time
User Expansion Factor
Powered by XDMoD
2012-01-01 to 2012-03-31

Total NUs Charged by Field of Science

Resource = SDSC-TRESTLES

2012-01-01 to 2012-03-31

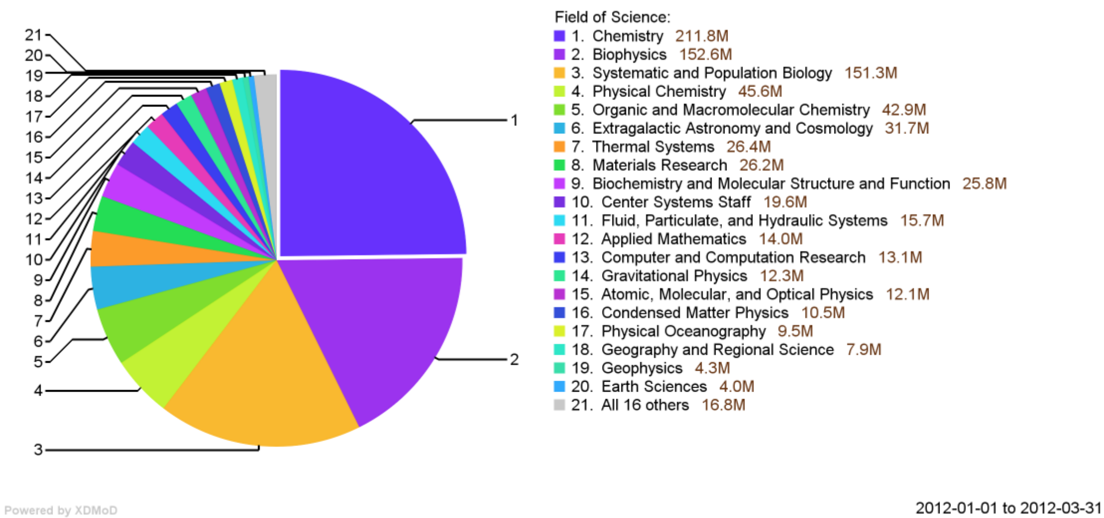

Total NUs Charged by Institution

Resource = SDSC-TRESTLES

2012-01-01 to 2012-03-31

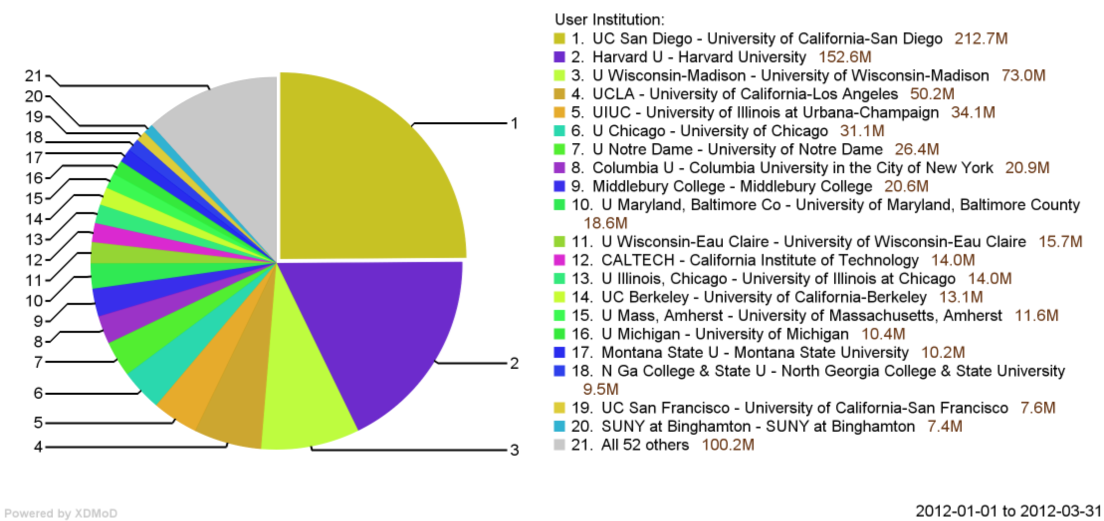

Total NUs Charged by Principal Investigator

Resource = SDSC-TRESTLES

2012-01-01 to 2012-03-31

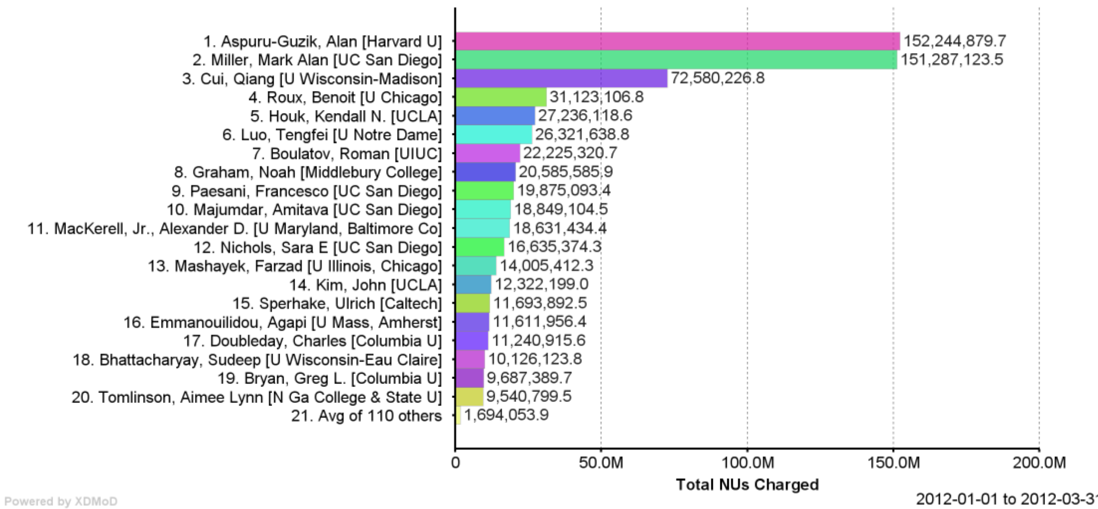

SDSC Dash Quarterly Report

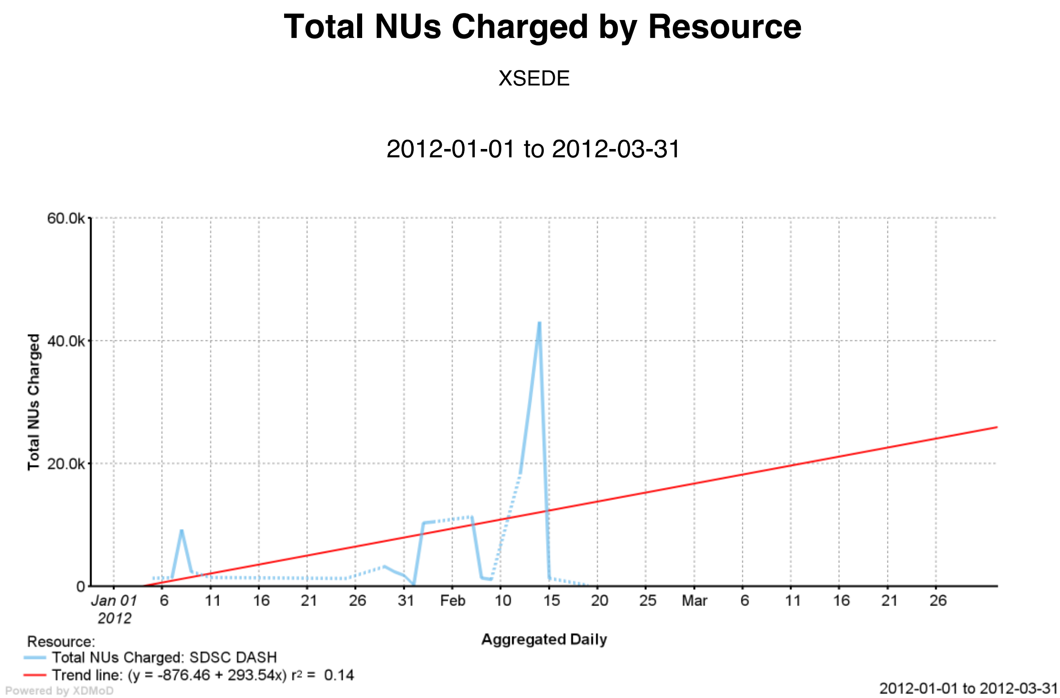

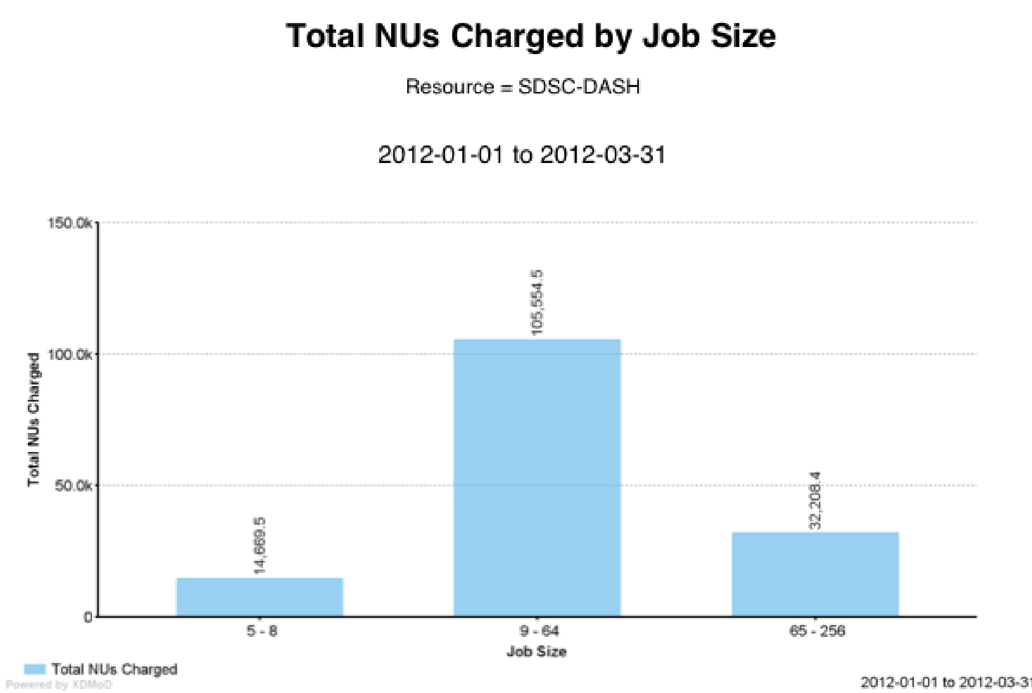

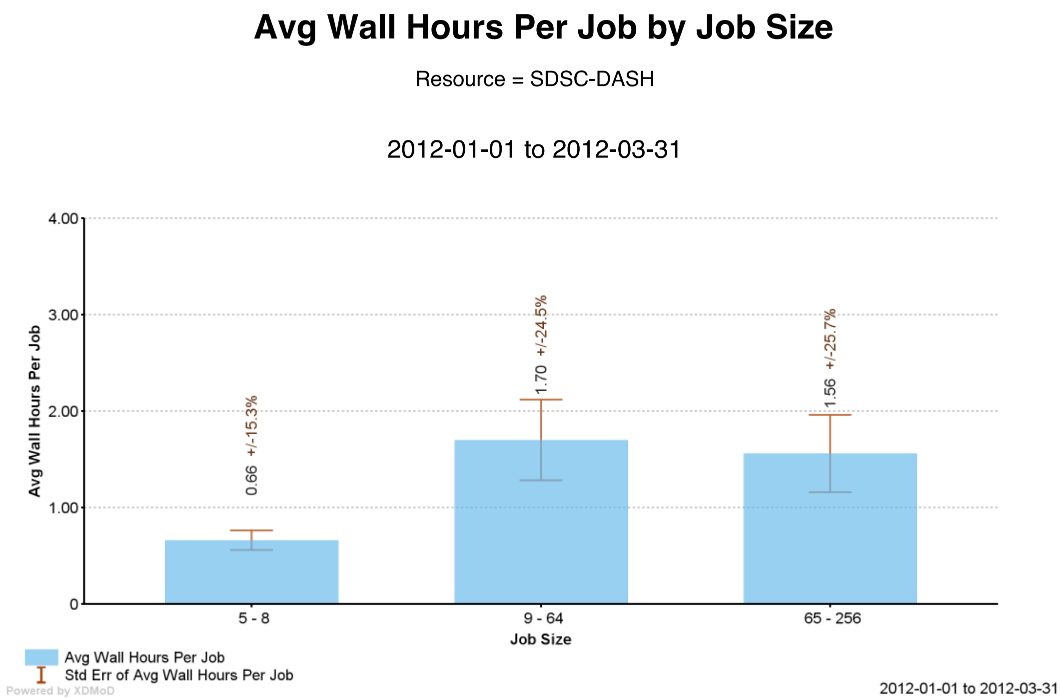

Avg Wall Hours Per Job by Job Size
Resource = SDSC-DASH
2012-01-01 to 2012-03-31
Avg Wall Hours Per Job
4.00
3.00
2.00
1.00
0
0.66 +/-15.3%
1.70 +/-24.5%
1.56 +/-25.7%
5 - 8
9 - 64
65 - 256
Job Size
Avg Wall Hours Per Job
Std Err of Avg Wall Hours Per Job
Powered by XDMoD
2012-01-01 to 2012-03-31

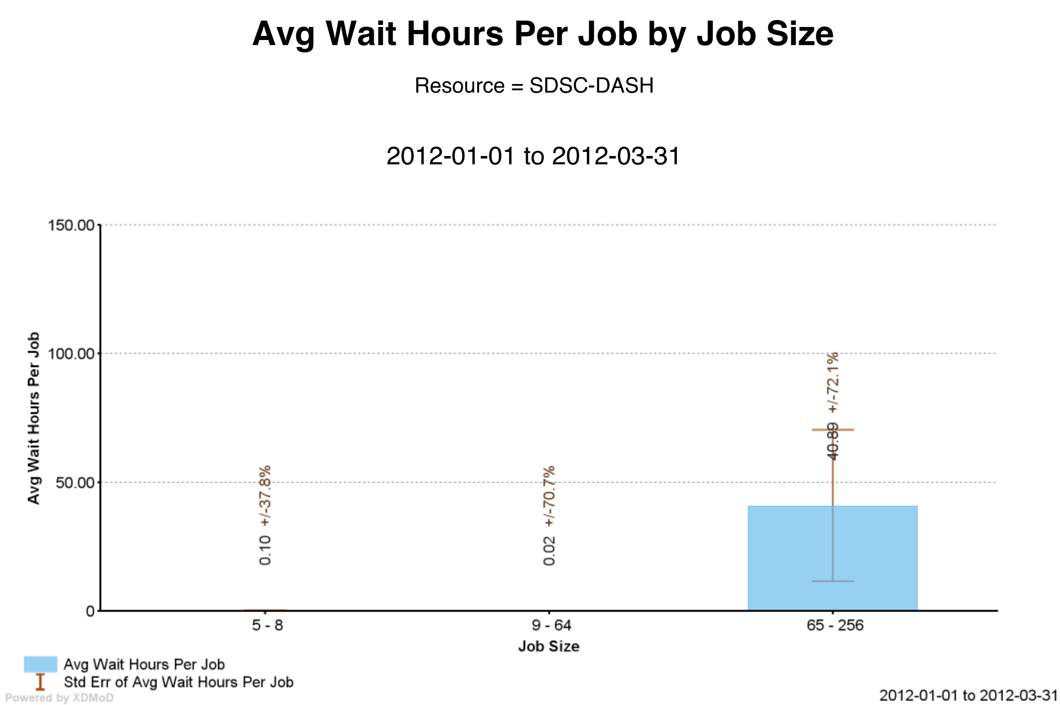

Avg Wait Hours Per Job by Job Size
Resource = SDSC-DASH
2012-01-01 to 2012-03-31
Avg Wait Hours Per Job
150.00
100.00
50.00
0
0.10 +/-37.8%
0.02 +/-70.7%
40.89 +/-72.1%
5 - 8
9 - 64
65 - 256
Job Size
Avg Wait Hours Per Job
Std Err of Avg Wait Hours Per Job
Powered by XDMoD
2012-01-01 to 2012-03-31

Total NUs Charged by Field of Science

Resource = SDSC-DASH

2012-01-01 to 2012-03-31

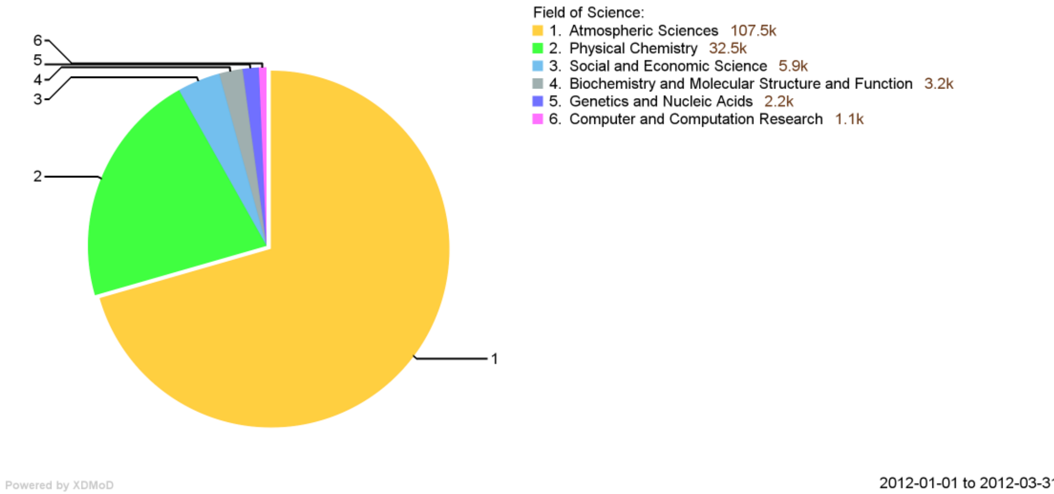

Total NUs Charged by Institution

Resource = SDSC-DASH

2012-01-01 to 2012-03-31

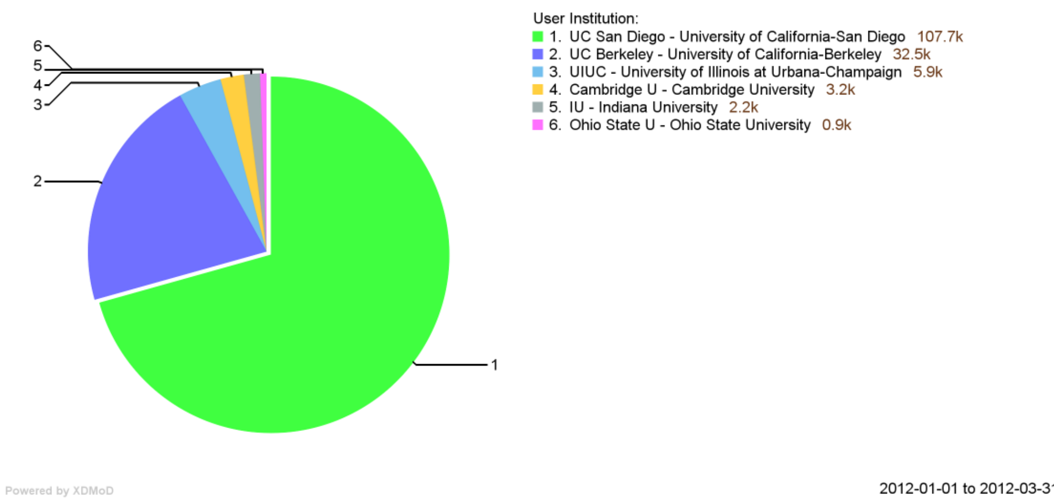

Total NUs Charged by Principal Investigator

Resource = SDSC-DASH

2012-01-01 to 2012-03-31

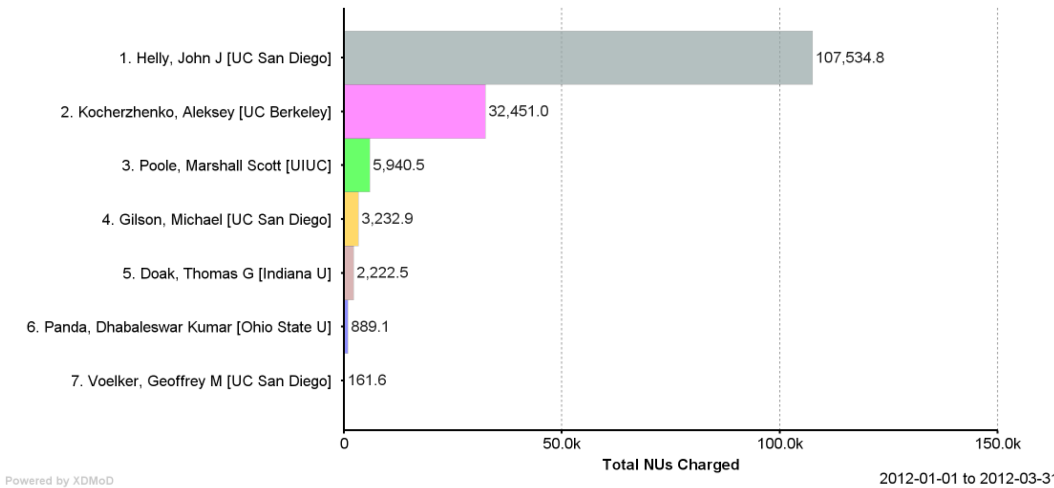

User Expansion Factor by Job Size

Resource = SDSC-DASH

2012-01-01 to 2012-03-31

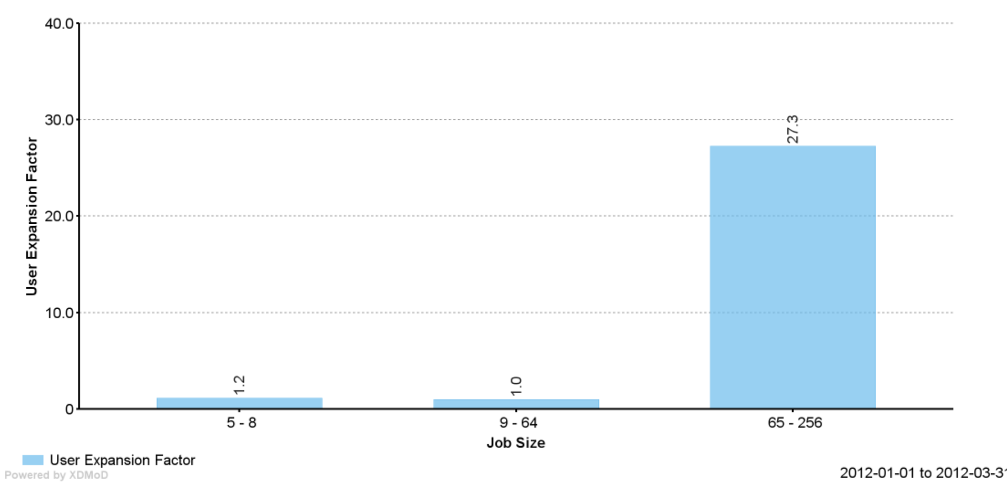

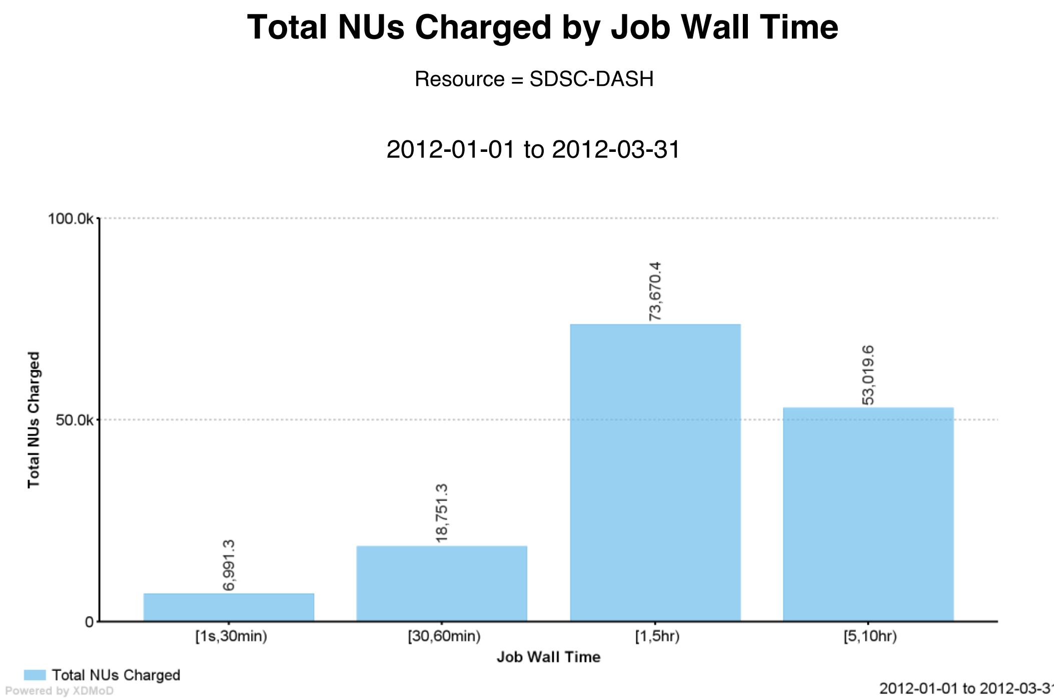

20.9.2 Standard User Assistance Metrics

Between the two tickets systems used to support Trestles and Gordon (NCSA's XSEDE ticket system and SDSC's local ticket system) a total of 545 tickets were created between January 1 and March 31, 2012. 501 of those tickets were closed, leaving 44 tickets that we are still working to resolve. The average time to close the tickets across the two systems was 5.8 days, with a median time of 3.1 days. The XSEDE ticket system statistics are provided in the table below (SDSC local ticket stats not shown.)

SDSC Ticket resolution times by category from XSEDE Ticket System

Time to Resolution	account issues	file systems	grid software	jobs/batch queues	login/access issues	mss/data issues	network issues	software/apps	system issues	other
0-1 hr				4				2		
1-24 hr	4	7		15	7			5	2	2
1-7 d	3	8	1	29	19	5		17	2	3
1-2 wk	2	1		12	8	2		4		2
> 2 wk		1		14	3	2		13		1
Still Open	1		1	2	1	4		9		

20.9.3 <u>SP-specific Metrics (Trestles)</u>

Trestles Metrics:

Number of Jobs run by Job Size (Cores)

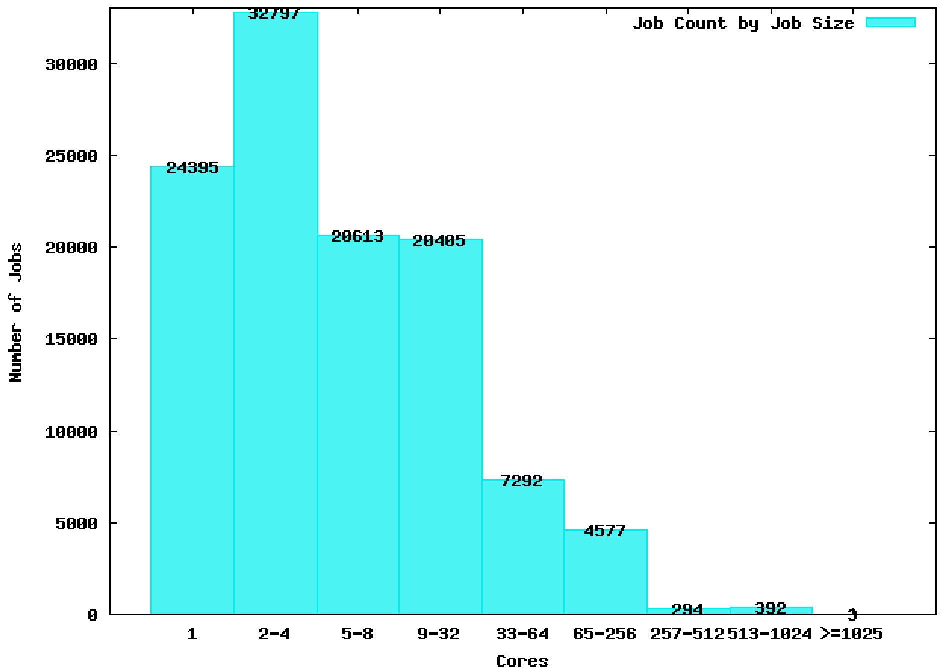

Average System Wait Time (Hours) by Job Size (Cores)

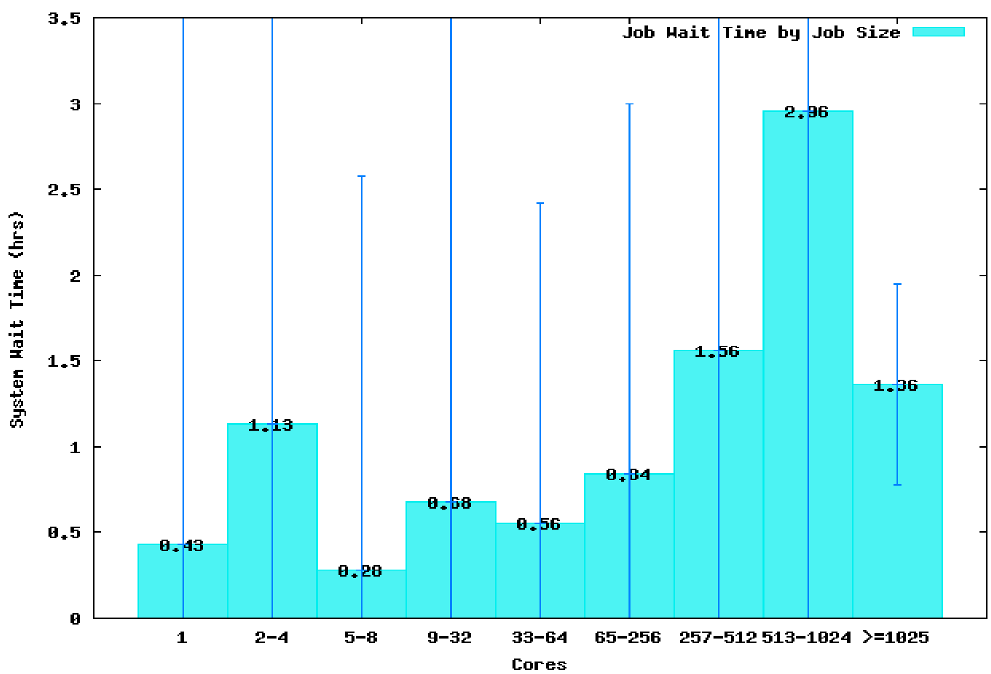

Average Requested Wall Time (Hours) by Job Size (Cores)

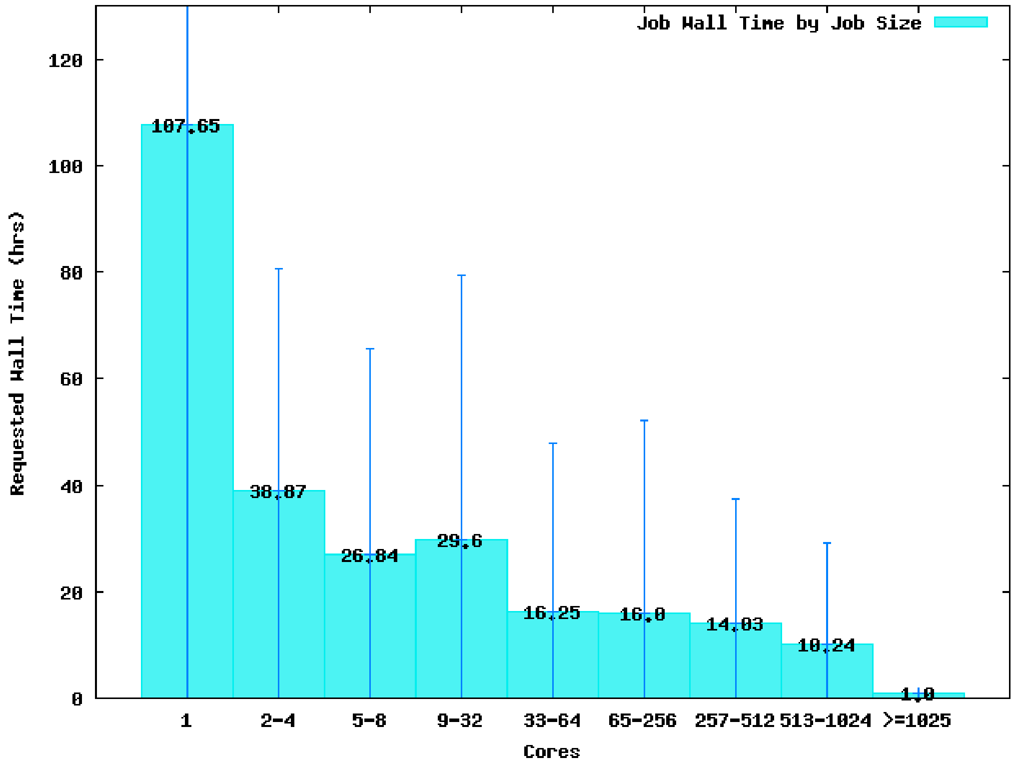

Average "Scheduler" Expansion Factor by Job Size (Cores)

(Requested Wall Time + System Wait Time)/Requested Wall Time

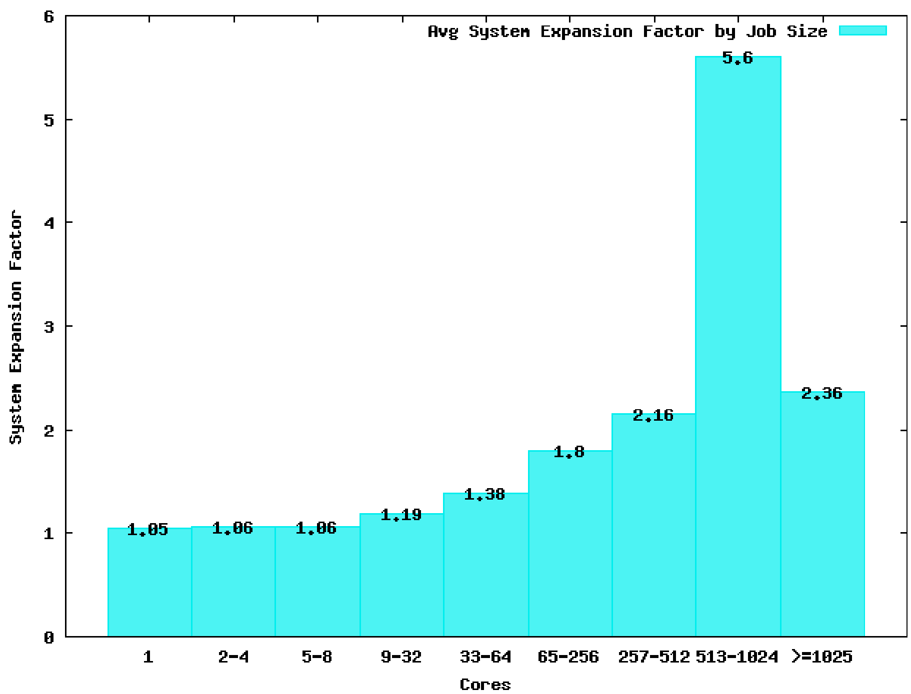

Average Scheduler Expansion Factor by Requested Wall Time (Hours)

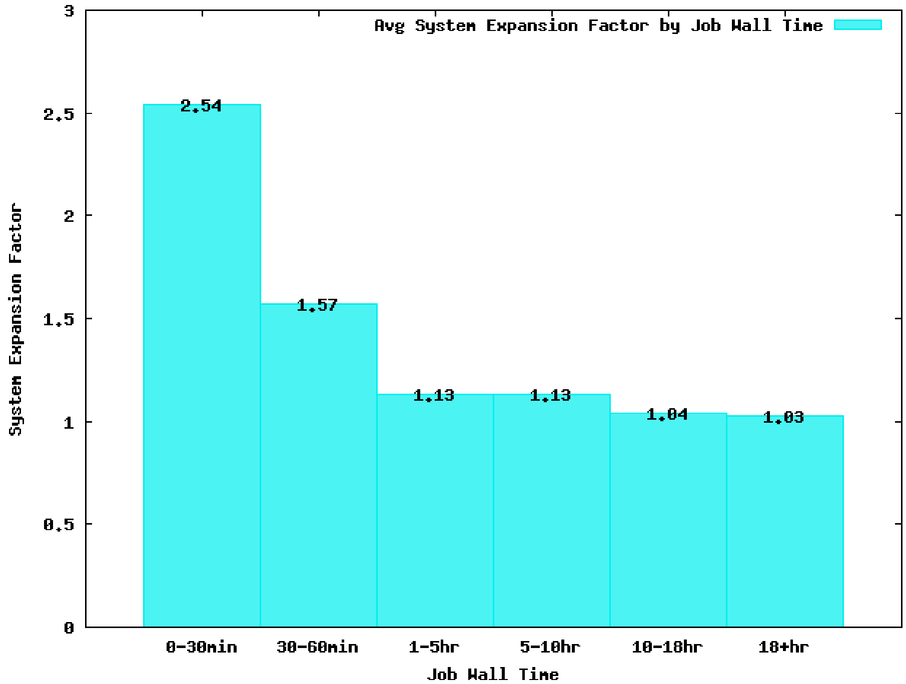

21 TACC - Service Provider Quarterly Report

21.1 Executive Summary

The Texas Advanced Computing Center (TACC) at The University of Texas at Austin (UT Austin) develops and deploys an integrated infrastructure of advanced computing resources to enhance the research and education activities of the faculty, staff, and students at UT Austin, and in Texas and across the US through its involvement in various state and national programs, including the NSF funded eXtreme Digital Resources for Science and Engineering (XD) project. This infrastructure includes high performance computing (HPC) systems, advanced scientific visualization (SciVis) systems, data servers and storage/archival systems, grid computing servers, IT systems, high-bandwidth networks, and a comprehensive software environment comprising applications, tools, libraries, databases, and grid software. TACC services include technical documentation, consulting, and training in HPC, SciVis, and grid computing.

TACC staff continued to contribute to the success of the XSEDE project to date. Significant effort was expended in the User Services User Engagement, User Interaction & Interfaces, and Training activities, XSEDE web site, and the Extended Collaborative Support Services area.

TACC continues to take a leadership position in XSEDE training efforts by offering 17 training workshops attended by 640 students during the reporting period. The Cornell University team continued to add training modules to and improve the Ranger virtual workshop.

Work continues on the expansion of the TACC computing center to accommodate the new Stampede cluster.

21.1.1 Resource Description

21.1.1.1 Sun Constellation Linux Cluster (Ranger)

The TACC Sun Constellation Cluster contains 62,976 cores (2.3 GHz) within 3,936 Sun Constellation blades (nodes), an X4600 Rocks master node, 4 X4600 user login nodes, 4 X4600 user gridftp nodes, 4 X4600 data movers, 2 X4600 nodes dedicated to supporting the SGE batch system, 2 X4600 external management service nodes, 2 X4600 InfiniBand subnet management nodes, an X4100 software build node, and 6 X4600 metadata server nodes to support the Lustre parallel file systems. Multiple work and home file systems are configured from 1.7 PB of storage managed by the Lustre parallel file system management software. Two Sun Data Center 3456 switches are the core of an InfiniBand fabric through which all components are connected. The basic configuration is as follows:

- 3936 Sun Constellation Blade Servers, each with
  - four quad-core 2.0 GHz processors
  - 32 GB of Memory
  - 8 GB flash drive
- 1.7 PB of storage managed by the Lustre Parallel File System software
- InfiniBand Interconnect

21.1.1.2 Dell Westmere Linux Cluster (Lonestar)[5]

The TACC Dell Westmere Cluster contains 22,656 compute cores (3.33 GHz) within 1,888 Dell PowerEdge M610 compute blades (nodes), 15 PowerEdge R610 compute-I/O server-nodes, and 2

[5] Began production on February 1, 2011.

PowerEdge M160 login/management nodes. Each compute node has 24 GB of memory, and the login/development nodes each have 24 GB. 14 large memory (1TB) nodes are available for high-throughput computing and applications that require access to a shared-memory architecture and 8 GPU nodes are configured for visualization and applications that can take advantage of the computational speed of the GPUs. The system storage includes a 421 TB parallel WORK Lustre file system, a 841 TB parallel SCRATCH Lustre file system, and 275 TB of local compute-node disk space (146GB/node). A QDR InfiniBand switch fabric interconnects the nodes (I/O and compute) via a fat-tree topology, with a point-to-point bandwidth of 40Gb/sec. The basic configuration is as follows.

- 1888 Dell PowerEdge M610 Blade Servers, each with
  - Dual Intel Westmere 6-core, 3.33 GHz processors
  - 24 GB of Memory
  - 146 GB of Local Disk
- 14 Dell PowerEdge R910 servers, each with
  - Four 6-core, 2.0 GHz Intel Xeon processors
  - 1 TB of Memory
  - 292 GB of Local Disk
- 8 Dell PowerEdge C6100 servers, each with
  - Two NVIDIA M2070 GPUs
  - Two 6-core, Intel Xeon X5670 2.93 GHz processors
  - 24 GB of Memory
  - 146 GB of Local Disk
  - 16-lane PCI Express to Dell C410x PCI expansion box housing the NVIDIA GPUs
- 421 TB Lustre Parallel File System (WORK)
- 841 TB Lustre Parallel File System (SCRATCH)
- QDR InfiniBand Interconnect

21.1.1.3 DELL/NVIDIA Visualization and Data Analysis Cluster (Longhorn)

The TACC DELL/NVIDIA Visualization & Data Analysis Cluster, Longhorn, is a hybrid CPU/GPU system designed for remote, interactive visualization and data analysis. In addition, Longhorn supports production, compute-intensive calculations on both the CPUs and GPUs via off-hour queues. The large, per-node memory is intended to support serial and parallel visualization and analysis applications that take advantage of large memories, multiple computing cores, and multiple graphics processors. Longhorn is an ideal companion resource for working with large data sets created on Ranger, since Longhorn can directly access Ranger's Lustre parallel file system through a 10 GigE network link.

The system consists of 256 dual-socket nodes, each with significant computing and graphics capability. Total system resources include 2048 compute cores (Nehalem quad-core), 512 GPUs (128 NVIDIA Quadro Plex S4s, each containing 4 NVIDIA FX 5800s), 13.5 TB of distributed memory and a 210 TB global file system. Longhorn configuration details can be found below.

128 NVIDIA Quadro Plex S4s, each with

- 4 NVIDIA FX 5800 GPUs
- 16GB Graphics Memory (4GB per GPU)

- 2 independent graphics busses, one per GPU pair

240 Dell R610 Compute Nodes, each with

- 2 Intel Nehalem quad-core processors (8 cores) @ 2.53 GHz
- 48GB RAM
- 73GB local disk
- connected to 2 dedicated NVIDIA FX 5800 GPUs via Quadro Plex graphics bus

16 Dell R710 Compute Nodes, each with

- 2 Intel Nehalem quad-core processors (8 cores) @ 2.53 GHz
- 144GB RAM
- 73GB local disk
- connected to 2 dedicated NVIDIA FX 5800 GPUs via Quadro Plex graphics bus

Mellanox QDR InfiniBand Interconnect

14 Dell PowerVault MD1000 Direct Attached Storage Arrays (210TB global file system, managed by the Lustre Parallel File System)

21.1.1.4 Terascale Sun Visualization Cluster (Spur)

TACC's Terascale Sun Visualization Cluster contains 128 compute cores, 1 TB aggregate memory and 32 GPUs. Spur acts not only as a powerful stand-alone visualization system: it also enables researchers to perform visualization tasks on Ranger-produced data without migrating to another file system and to integrate simulations and rendering tasks on a single network fabric. The cluster consists of the following hardware:

- 1 Sun Fire X4600 server with 2 NVIDIA Quadro Plex model 4. The X4600 contains 8 dual-core CPUs (16 cores total) and 256GB of RAM. Each Quadro Plex model 4 contains 2 NVIDIA Quadro FX5600 GPUs;
- 1 Sun X4400 servers, with 4 quad-core CPUs (16 cores total) and 128GB of RAM, connected to 2 NVIDIA Quadro Plex model 4. Each Quadro Plex model 4 contains 2 NVIDIA Quadro FX 5600 GPUs;
- 6 Sun X4400 servers, each with 4 quad-core CPUs (16 cores total) and 128GB of RAM, and each connected to an NVIDIA Quadro Plex S4. Each Quadro Plex S4 contains 4 NVIDIA Quadro FX 5600 GPUs; and
- Total system capability: 128 cores, 1TB aggregate memory, 32 GPUs.

Because Spur shares Ranger's interconnect fabric and file systems, researchers will be able to easily transition between HPC runs to generate and visualize data. Furthermore, visualization software is able to harness both the rendering power of the graphics hardware and the compute power of Ranger to enable the analysis of terascale and larger data sets.

21.2 Science Highlights

XSEDE Q1 2012 Science Highlights:

Atmospheric Sciences (ATM)

"High-Performance Computing for Ensemble-based Cloud-Resolving Hurricane Analysis, Prediction and Predictability"

PI: Fuqing Zhang, Pennsylvania State University

Upgrading the Hurricane Forecast
TACC's Ranger supercomputer helps researchers predict storm intensity with greater accuracy

By Aaron Dubrow (aarondubrow@hotmail.com)

When Hurricane Irene swept through New England in August 2011, predictions of its path were on target, but it arrived significantly weaker than originally forecast, leading to a larger evacuation than was required. Forecasting how hurricanes form, intensify, or dissipate is different and more challenging than predicting its path.

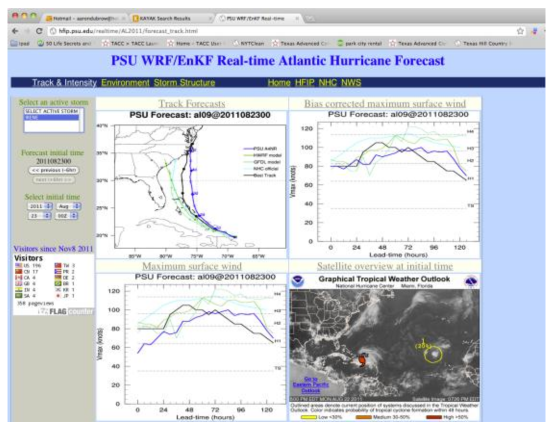

Figure 1. *Zhang's* [*website*]() *shows his group's real-time Atlantic Hurricane Forecast (created using Ranger) plotted alongside comparable model forecasts, satellite images of the storm, and other pertinent information.*

Over the course of several years, researchers from Penn State and the National Oceanographic and Atmospheric Administration (NOAA) used the *Ranger* supercomputer at the Texas Advanced Computing Center to develop and test a new, high-resolution hurricane forecasting system that improves on today's operational methods in several important ways.

First, the system adopts a higher-resolution state-of-the-art computer model (4.5km grid spacing, compared to those of 9km) capable of more explicitly resolving a hurricane's inner-core dynamics and structure. Second, it ingests airborne Doppler radar taken by planes flying through the hurricane — information that NOAA has been collecting this data for 30 years, but that has yet to be put into operational models. Third, the system applies a new data assimilation method that improves how the system ingests the Doppler radar data.

Using *Ranger*, Zhang's system forecasted the track and intensity of every major storm in the Atlantic throughout the 2011 hurricane season, sharing the results of his simulations on his personal research website. Zhang's forecasts were shown to improve intensity predictions by an average of 20 to 40 percent over the National Hurricane Center's (NHC) official forecasts for storms that have the airborne Doppler radar data.

The prediction system is one of a handful the NHC is assessing to become part of the operational forecasting system used in emergency situations.

--

Education

"Education Visualization Laboratory (EdVisLab)"

PI: Kelly Gaither, Texas Advanced Computing Center

Scientific Visualization for Educators
Students, faculty and staff to benefit from new visualization technologies

By Aaron Dubrow (aarondubrow@hotmail.com)

Members of the College of Education (COE) and the Texas Advanced Computing Center (TACC) gathered in early February 2012 for the grand opening of the Education Visualization Laboratory (EdVisLab), part of the COE's Learning Technology Center. The lab represents a new generation of tools and technologies that will enhance researchers' capacity to represent and interpret large datasets. It also represents the non-traditional collaborations into which HPC centers are increasingly participating.

Figure 2. *TACC research associate Brandt Westing (left) showcases the high-resolution capabilities of a College of Education dataset.*

Staff in TACC's Visualization and Data Analysis group designed and deployed the new EdVisLab. The lab is equipped with a tiled array of 15, 30-inch high-resolution monitors, as well as a 73-inch, stereoscopic 3D display all operated by human-computer interaction (HCI) interfaces. The HCI interfaces allow researchers and students to engage the display without the use of a mouse or controller.

The EdVisLab is being reserved for research studies at the University of Texas that benefit from the incorporation of visualizations of large data sets. For instance, Taylor Martin, curriculum and instruction professor within the College of Education, is using virtual manipulatives and online educational games to study how children learn mathematics.

The data generated from these sources is so immense that tracking patterns and trends with traditional analytical techniques has been a challenge. Martin's group and TACC are currently collaborating in the EdVisLab using various visualization techniques and data aggregation methods to create visual representations of student learning processes. So far, preliminary visualizations have succeeded in summarizing results and correlating parameters to help researchers navigate, grasp trends, and deduce conclusions from the data.

--

Physical Chemistry

“Direct Dynamics Simulations of Many-Atom Complex Molecular Systems”

PI: William Hase, Texas Tech University

Collision Chemistry

Texas Tech researcher uses TACC supercomputers to explore ion implantation and applications for molecular recognition in medicine

By Aaron Dubrow (aarondubrow@hotmail.com)

Supercomputers allow chemists to study the dynamics of individual molecules when they collide with a level of detail and granularity not available through laboratory methods. William Hase, professor of Chemistry at Texas Tech University, used the *Ranger* and *Lonestar* supercomputers to better understand the nature of atomic collisions related to biological and industrial applications.

A recent project simulated peptide ions colliding with surfaces. These collisions are being used to template an organic surface with a specific type of peptide on it. Because each peptide molecule has properties that make it interact strongly with a small subset of other types of molecules, researchers believe molecular templates may one day allow scientists to design small molecules that can control cells during disease or slow aging.

Collaborating with Julie Laskin, a laboratory fellow at the Pacific Northwest National Laboratory, they are learning how readily molecules can bind to the surface, or how well they are trapped in the surfaces. The information helps the experimentalists tune their experiments, suggests new areas of study, and provides insights that can be compared with experimental data.

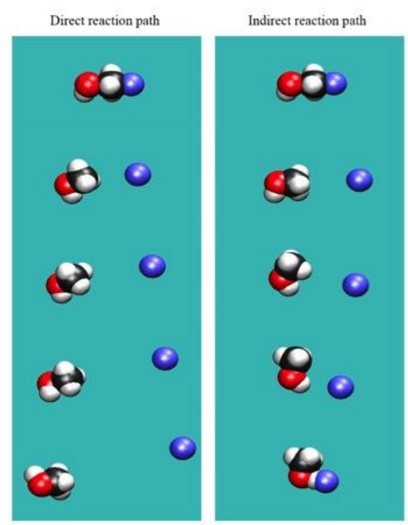

Figure 3. *The two pathways for motion from the [HO---CH3---F]- central barrier to the CH3OH + F- reaction products. Most of the trajectories follow the direct dissociation path. A small amount ~10%, form the CH3OH---F- hydrogen-bonded intermediate and follow an indirect path.*

--

Systematic and Population Biology

“Large scale phylogenetic estimation”

PI: Tandy Warnow, The University of Texas at Austin

A Tree of Life Grows in Texas
Scientists use *Ranger* supercomputer to create a more accurate evolutionary history

By Aaron Dubrow (aarondubrow@tacc.utexas.edu)

Phylogenetics is the branch of life science that studies the evolutionary relationships among organisms based on genetic evidence. Tandy Warnow, a computer science professor at The University of Texas at Austin, used the *Ranger* supercomputer to develop and test smarter, faster, and more accurate genetic alignment and tree-building algorithms and applied them to some of the largest biological datasets ever created.

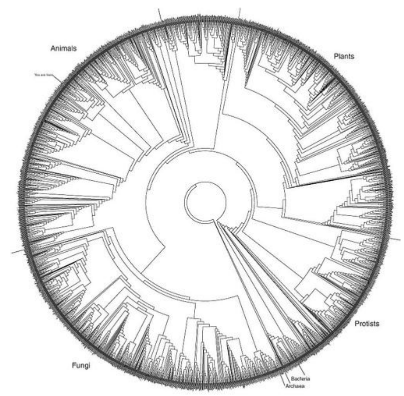

Figure 4. *This phylogenetic tree, created by David Hillis, Derreck Zwickil and Robin Gutell, depicts the evolutionary relationships of about 3,000 species throughout the Tree of Life. Less than 1 percent of known species are depicted.*

Among her projects, she and her team collaborated with evolutionary biologists from the Smithsonian to resolve the evolutionary history of flightless birds (known as *ratites*) and to answer the question: how did so many similar species get to the far-flung corners of the Earth?

Smithsonian scientist Michael Braun discovered through DNA analysis that an ancient family of birds found in South American, the tinamou, is one of the most closely related groups to emus and ostriches — and they can fly! This fact, combined with the lack of skeletal evidence for flightless birds before the time of continental breakup, led to a re-conceptualization of the ratite branch of the avian tree. Warnow worked with Braun to reanalyze his controversial findings using her software. Their study confirmed the evolutionary relationship that Braun found. By improving the quality of the avian tree of life through computational methods, a new history emerged.

--

Institutional Infrastructure

“Data Management and Collections Projects”

PI: Maria Esteva, Texas Advanced Computing Center

Attention Researchers: Protect Your Data!
TACC, The University of Texas Libraries and Information Technology Services partner to offer data management resources to researchers

By Aaron Dubrow (aarondubrow@tacc.utexas.edu)

Figure 5. *Example of museum specimens of Dionda argentosa collected in Texas in 1893 and maintained in the Smithsonian Museum. Over 124,000 records such as this are stored on TACC's servers, along with various associated images, that are being used for research including development of species distribution models.*

Government funding agencies, private foundations, and universities are asking faculty researchers to pay extra attention to their data, but for large and complex data sets this can be a challenge. To address the growing need for data management plans, TACC partnered with the Libraries at The University of Texas at Austin and UT Information Technology Services to develop a data management ecosystem through which faculty can manage data related to their research projects.

TACC already hosts many collections (created by the Texas Natural Science Center and other museums and natural history centers nationwide) on the *Corral* system, which assures performance, web access and security. As part of the coordinated strategy, UT Libraries acts as the final resting place for datasets up to one gigabyte, including publications and papers. TACC will house larger (terabyte and petabyte-sized) datasets and collections that require complex architectures and functionalities, such as geographic information system (GIS) data and relational database management system services. ITS will provide hardware, co-location, network access, web services and technical assistance.

The management of digital data is an evolving field. Many aspects of the project, such as licensing, privacy, acceptable file formats, and the creation of metadata to assist in the organization and retrieval of information, are still being discussed and TACC staff is taking a lead in determining the most efficient solutions.

Throughout the year, TACC will be holding training sessions for interested researchers. To learn more about the project, explore the Data Management at UT page on the UT Libraries website.

--

Neuroscience

Medial temporal lobe function in human memory

PI: Allison Preston, The University of Texas at Austin

The Cognitive Neuroscience of Memory

By Faith Singer-Villalobos (faith@tacc.utexas.edu)

Dr. Alison Preston, an assistant professor in the Department of Psychology at The University of Texas at Austin and a recipient of the National Science Foundation's CAREER Award explores how the brain supports memory and how memory influences the decisions people make. Her research has many societal implications, including a major influence on the criminal justice system, and how people record events in their life. Much of the work done in Dr. Preston's lab combines laboratory experimentation with data-driven computation.

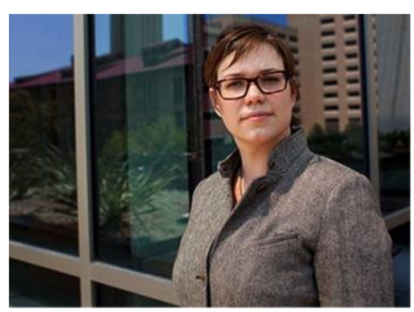

Figure 5. *Alison Preston is an assistant professor in the Department of Psychology and Section of Neurobiology at The University of Texas at Austin.*

For example, her research involves putting people in a magnetic resonance imaging (MRI) machine and taking thousands of pictures of their brains while they are doing various memory tasks, such as learning a long list of words for 45 minutes. Her job is to mine that very rich data set and relate it to various aspects of the task—whether someone is likely to remember words or not remember them, for example.

Preston uses TACC's HPC systems, specifically Lonestar 4 (an XSEDE resource), to perform the data analysis for the brain imaging studies to try to understand how these patterns of brain response relate to human behavior.

By combining state-of-the-art MRI techniques with advanced statistical approaches to data analysis (e.g., multivariate statistics, structural equation modeling, computational modeling), Preston's research aims to extend beyond standard theories of memory to focus on a number of cognitive functions.

Preston said, "While many things we do can be done on a standard personal computer, it would be very slow going to do so. Advanced computing, like the resources available at TACC, allows us to do things much faster, which has a direct influence on the rate of scientific discovery. In our case, advanced computing can cut months, or sometimes years, from our analysis time."

21.3 User-facing Activities

(If multiple awards at the SP institution, provide sub-sections for each award and for common infrastructure/services)

21.3.1 System Activities

Longhorn continues to be a workhorse for remote interactive visualization and large-scale data analysis activities. TACC staff continue to monitor usage of the system and the Longhorn user portal and reach out to new communities, including the social sciences and the humanities.

The capacity of the Ranch storage facility was increased by 10PB with the addition of 2,000 tape cartridges.

Ranger availability was 97% during the reporting period. While still above the required availability, it is below Ranger's lifetime 98%+ average. One 12 hour preventative maintenance was taken, and a 2 day outage was required to bypass chilling lines as part of the construction project to bypass Stampede. Construction should be able to continue from this point forward without additional impacts to Ranger. During the reporting period, Ranger successfully completed it's originally scheduled production life. NSF has granted a one year extension on Ranger operations, and the extension year began February 5th. During this period, operations will continue unchanged, and security upgrades/patches and application upgrades will continue, but no major upgrades to the software environment will happen.

21.3.2 Services Activities

All Systems

Dr. George Khelashvili and Niklaus Johner at Weill Cornell Medical College requested support of Replica-Exchange Molecular Dynamics (REMD) on TACC resources. REMD is a technique used to enhance sampling relative to a standard molecular dynamics simulation by allowing systems of similar potential energies to sample conformations at different temperatures. By doing so, energy barriers on the potential energy surface might be overcome, allowing for the exploration of new conformational space. TACC support staff modified NAMD provided Tcl sever and socket connections scripts and made them suitable for Ranger and Lonestar. The users were very cooperative and provided important feedback to test the computational procedures for their applications. During the test, multiple replicas on a single node were found to be extremely inefficient due to the fact that multiple MPI tasks occupied the same set of cores. The problem was solved by properly setting the CPU binding affinity environment variables in MVAPICH and MVAPICH2 with the assistance of Yaakoub El Khamra and John Cazes. The changes have been updated in the latest release of ibwrapper. In addition, Niklaus found that the REMD scripts worked fine for 8 replicas but not 16 replicas. By going through a number of tests, we found that the event-driven loop in the implementation was unresponsive with more than 14 replicas using Tcl/Tk version 8.4. A newer version of Tcl/Tk(8.5) was built for the users and the calculations are now being carried out routinely by the research team.

TACC staff provided custom installations of R for the following research teams: Matthew Moskwik (UT Austin), Karthik Ram (UC Berkeley), Russell Poldrack (UT Austin), and David Romps (UC Berkeley).

TACC staff installed 8 custom modules on Ranger for XSEDE PI Stephanie Stenzel’s research group at the University of Michigan.

TACC staff worked with XSEDE user Artem Baskin (Univ. of Illinois at Chicago) to modify the current installation of Siesta to enable more efficient post-processing. This work included building, testing, and installing additional Siesta utilities for data transformation and informing users of input methods that result in significant speedup.

Longhorn

The TACC visualization group continues to provide high quality user support for the Longhorn XD Visualization project, providing at least one visualization person every weekday for hands on help ranging from startup on Longhorn to more in depth application support. Additionally, the Longhorn XD Visualization project provides more in depth user support to work on longer term projects that emerge as users become more proficient on the resources. In general, this more in depth support involved the development of tools and techniques to support user needs.

21.4 Security

Describe any change in security procedures and all security incidents/responses.

There are no changes in security procedures or security incidents to report within this reporting period.

21.5 Education, Outreach, and Training Activities

21.5.1 Outreach

XSEDE activity in Q1Y2012 ramped up with participation in monthly All TEOS meetings led by Scott Lathrop NCSC, bi-weekly Outreach meetings led by Laura McGinnis at PSC, and bi-weekly Underrepresented Outreach meetings led by Linda Akli at SURA. The meetings provided an opportunity for information and clarification of activities across TEOS and illuminated opportunities for collaboration within the existing planned activities.

Planning for the first cohort of XSEDE Student Engagement participants has been underway in Q1Y2012, via conference calls with Laura McGinnis of PSC and Jim Ferguson of NICS. The deadline for applications was April 2 with selection of participants expected by April 19. XSEDE members submitted twenty-two projects for student engagement. Twenty two students will be selected to attend orientation at TACC, NICS, or PSC on May 30, work on their project remotely or on sight for six weeks, present their works in progress in July at XSEDE 12, and spend the remaining 2-3 weeks to prepare their projects for completion by transition to fall student engagement participants.

Recruiting for the first cohort of student engagement participants included promoting the program to students and faculty at; the National Society of Hispanic Physicists/National Society of Black Physicists Conference in Austin, Texas in Q4Y2011, the NSF/AAAS Emerging Researchers National Conference in STEM in Atlanta, Georgia February 23-25; the National Society of Black Engineers Convention in Pittsburgh, Pennsylvania March 28-April 1.

Activities with the Underrepresented Outreach group included meeting with Computer Science and Electrical Engineering faculty at Clark Atlanta University before the Emerging Researchers National Conference. The visit laid the groundwork for collaborating with Clark Atlanta on infusing scientific computation into their undergraduate curriculum, starting with the development of a visualization lab. TACC VisLab associates provided input on the selection of equipment for the Clarke Atlanta lab, based on their experience with developing a small but powerful tiled display for the College of Education at the University of Texas at Austin. The Clarke Atlanta meeting also resulted in the finalization of a visualization training session by TACC in Atlanta on April 9.

In conjunction with the TEOS meetings and the input provided by Kay Hunt at Purdue, a campus champion is in the works for Austin Community College (ACC) a two-year college in Austin, Texas. Lorraine Brill is an adjunct faculty member at ACC in computer science and has attended multiple TACC training workshops. Dr. Brill will collaborate with TACC to develop a scientific computing certificate for ACC. Steve Gordon at OSC has offered to assist in the development of the certificate through the XSEDE Education group. Dr. Brill will serve as a liaison between TACC and ACC to increase the awareness of and participation in scientific computation by ACC students.

Finally, a conference call with Dan Stanzione and Steve Gordon at OSC laid the groundwork for OSC providing consultative services to institutions preparing to offer scientific computation to their students. The meeting led to Dr. Stanzione presenting the services of XSEDE TEOS in Lost Pines, Texas to affiliates of the $3.3 million National Science Foundation initiative to incorporate research in the teaching of biology at community colleges. The initiative is led by Finger Lakes Community College in upstate NY.

The Austin Forum on Science, Technology & Society

In this quarter, TACC hosted a total of 3 monthly Austin Forum events with invited speakers from areas of interest focused on science and technology. The goal of The Austin Forum on Science, Technology & Society is to engage and educate the local community about the numerous ways in which science and technology enhance the quality of their everyday life, as well as the health, prosperity and security of the nation. One hour is devoted to a presentation and Q&A discussion between the speaker and guests. Ample time for networking is offered, both preceding and following the speaker presentation. A total of 743 people attended The Austin Forum.

TACC Facility Tours

From K-12 and higher education groups, TACC conducted facility tours impacting 271 people; 50% were under-represented. An overview of XSEDE and TACC were given at each event.

The participation of TACC in the annual Explore UT Open House allowed for the integration of interactive activities during tours of the Visualization Laboratory (Vislab) for 400 people, including K-12 students, parents, and teachers.

This quarter, TACC also created and implemented two new educational workshops to enrich the facility-tour experience:

1) Cybersecurity-Creating Strong Passwords- Launched at Explore UT, TACC facilitated a workshop to increase awareness of Internet safety, while teaching deep concepts of how computer hardware, software, and networks operate.

2) Tree of Life- Using iPlant's web interface, the Ranger supercomputer, and software, high school students looked at a piece of the genome for fruit, and figured out which fruits are most closely related.

The following table lists TACC outreach activities during the reporting period.

Type	Title	Location	Date(s)	Number of Participants	Number of Under-represented people
The Austin Forum on Science, Technology & Society	"The Predictive Brain: How Past Memories Influence Future Decisions" with Dr. Alison Preston	AT&T Conference Center	1/11/12	360	Not tracked.
Vislab Tour	St. Dominic Savio Catholic School (HS juniors and seniors)	Vislab	1/27/12	12	5
Vislab Tour	College of Nat Sci Deans Scholars Honors HS Recruits	Vislab	1/30/12	12	8
Vislab Tour	VIP: World Bank/Science, Tech, Innovation Prog Officer	Vislab	1/31/12	2	1
Vislab Tour	UT Women in Engineering FIG (freshmen)	Vislab	2/1/12	15	15
Vislab Tour	T-STEM Learning Network Specialist	Vislab	2/1/12	1	1
The Austin Forum on Science, Technology &	"A Preview to SXSWi -- Social Media & Non-Profits" with Hugh Forrest, David Neff, and Heidi	AT&T Conference Center	2/7/12	298	Not tracked.

Society	Adams				
Machine RoomTour	Andrews High School	PRC	2/10/12	4	0
Speaking Engagement	Kelly Gaither: UT WEP (graduate and PhD students)	UT	2/16/12	25	25
Vislab Tour	Prospective Graduate Students, UT Aerospace Engineering	Vislab	2/17/12	15	3
Vislab Tour	College of Nat Sci Deans Scholars Honors HS Recruits	Vislab	2/20/12	5	3
Vislab Tour	VIP: Intelligent Communities	Vislab	2/27/12	2	1
Vislab Tour	CS Ambassadors, Recruit w/Family	Vislab	2/29/12	4	2
Machine Room Tour	VIP: Ultra Electronics	PRC	3/1/12	1	0
Vislab Tour	Astronomy Graduate Student Recruitment	Vislab	3/1/12	14	3
Machine Room Tour	Mechanical Engineering Graduate Student Recruitment	PRC	3/2/12	8	Not tracked.
Vislab Tours	Explore UT	Vislab	3/3/12	400	Not tracked.
Vislab Tour	SITE Conference Participants (20 educators)	Vislab	3/6/12	20	Not tracked.
Vislab Tour	VIP: Verizon	Vislab	3/6/12	2	1
The Austin Forum on Science, Technology & Society	Teen Researcher: Discovering A New Treatment for Ovarian Cancer" with Shree Bose and Dr. Alakananda Basu	AT&T Conference Center	3/6/12	85	Not tracked.
Tours/Worksho p	Leander High School	Vislab/PR C	3/20/12	56	33
Vislab/Machin e Room Tours	VIP: Missouri S&T	Vislab/PR C	3/21/12	2	1
Vislab Tour	UT GLUE (undergraduate engineering women)	Vislab	3/27/12	28	28
Vislab Tour	Vislab: Ischool advisory council	Vislab	3/28/12	3	1
Vislab Tour	Jewish Academy	Vislab	3/29/12	15	6
Vislab Tour	UT College of Ed (undergrad students)	Vislab	3/29/12	25	19

21.5.2 Education

Scientific Computing Curriculum and Courses

TACC's Spring 2012 courses continued to attract students interested in acquiring scientific computation skills. Reflecting this continued interest was the addition of Scientific and Technical Computing to the Spring course offerings, making it available Fall and Spring. Spring enrollment totals were 36 graduate students and 45 undergraduates. TACC staff taught: Introduction to Scientific Programming (4 graduate students and 27 undergraduates), Parallel Computing for Science & Engineering (16 graduates and 11 undergraduates) and Scientific and Technical Computing (16 graduates and 7 undergraduates). Students were assigned accounts on Ranger.

The courses are offered in the Flawn Academic Center in a customized classroom housing both lecture space and a computer instruction laboratory. The classroom customization was made possible through a partnership with Chevron to increase instruction in scientific computing. Fall 2012 is targeted for a pilot webcasting of instruction from the lab to a higher education institution in the state of Texas without scientific computation course offerings.

In addition, the partnership with the UT Austin Division of Statistics and Scientific Computation (DSSC) entered its third year. DSSC offers an Undergraduate Certificate and a Graduate Portfolio Program in Scientific Computation. Four of the scientific computing classes taught by TACC fulfill requirements for the certificate and portfolio. Documentation of students completing the certificate program appears on their transcript as a notation. Five students have completed the program to date. As of Spring 2012, 36 undergraduates and 10 graduate students are enrolled in the program. DSSC's scientific computing course descriptions are online at:

http://www.tacc.utexas.edu/education/academic-courses

21.5.3 Training

TACC conducted 16 training workshops and staff members at the Cornell Center for Advanced Computing conducted one training workshop related to the Ranger project during the reporting period. A total of 640 students attended the workshops either in person or via webcast. The following table lists the date, title, location, and attendance for each workshop.

Date	Class	Attendance	Underrepresented Community	Location
1/10/2012	Bring your own code workshop	27	4	UT Southwestern Medical Center
1/17/2012	Introduction to PETSc	38	3	TACC, webcast
1/19-20/2012	Data Analysis on Ranger and Lonestar	31	6	Cornell Center for Advanced Computing
1/20/2012	Intro. to SciVis Workshop	17	6	Old Dominion Univ. (SURA)
1/26/2012	Linux/Unix Basics	59	11	TACC, webcast
2/2/2012	Data Management Planning and Execution	57	16	TACC, webcast
2/6-7/2012	Intro. to Parallel Computing on Ranger and Lonestar	55	12	TACC, webcast
2/17/2012	Bring your own code workshop	16	4	UT Health Science Center at Houston
2/20/2012	Advanced PETSc	36	3	TACC,

				webcast
2/23/2012	Intro. to Scientific Visualization	29	2	TACC
2/24/2012	Visualization Workshop	12	NA	UT at San Antonio
2/25/2012	PerfExpert Tutorial	7	NA	New Orleans
3/5/2012	Fortran 90/95/2003 for HPC	57	8	TACC, webcast
3/19/2012	C Programming Basics	58	10	TACC, webcast
3/26/2012	C++ Programming Basics	68	12	TACC, webcast
3/29/2012	Writing a Data Management Plan	48	11	TACC, webcast
3/30/2012	Visualization Workshop	25	13	Florida A&M Univ. (SURA)

The Virtual Workshop provides users access to twenty-two training modules with new modules under development and existing modules being reviewed for updates. Users who are logged in to the XSEDE portal can pass-through to the Virtual Workshop, or they can use guest registration.

Available Modules

An Introduction to Linux
An Introduction to C Programming
An Introduction to Fortran Programming
Python on Ranger and Lonestar
Balancing Scripts and Compiled Code in Scientific Applications
MATLAB Programming
Parallel Programming Concepts and High-Performance Computing
Ranger Environment
Message Passing Interface (MPI)
MPI Point-to-Point Communications
MPI Collective Communications
MPI One-Sided Communication
MPI Advanced Topics
OpenMP
Hybrid Programming with OpenMP and MPI
Profiling and Debugging
Optimization and Scalability Series – Part 1: Planning for Parallel (major update completed)
Computational Steering
Large Data Visualization
ParaView
VisIt
Using Databases (major update completed)

Modules Under Development
Advanced Batch
Allocations
Distributed Debugging Tool (DDT)
EnVision

Multi-node Map Reduce (Module planning and material selection done)
Parallel I/O
Profiling with mpiP
R
Data Transfers

Modules Being Updated
MPI Collective Communications
Large Data Visualization
ParaView (Minor updates done, still in work)
VisIt
Python
Ranger Environment

Table: Virtual Workshop Usage

	Page Loads	Unique Visitors	First Time Visitors	Returning Visitors
Q1 '11	4,456	920	730	190
Q2 '11	16,281	2,988	2,509	479
July – Sept 2011	9,208	2,905	2,457	448
Oct – Dec 2011	10,068	3,615	3,019	596
Jan – Mar 2012	16,800	5,318	4,249	1069

Note: the Q2 '11 numbers were a result of high activity after an online news release on the Virtual Workshop was sent out.

21.6 SP Collaborations

The Longhorn XD Vis team continues to work with P.K. Yeung and Diego Donzis to visualize their large scale turbulent flow. Additionally, a paper was co-authored by the team and submitted to IEEE CG&A to detail the crucial role that visualization and data analysis play in analyzing and understanding turbulent flow simulations at 4096^3 cells per time slice (= 68 billion cells) and 17 time slices (= 1 trillion total cells). The visualization techniques presented in this paper allowed us to investigate the dynamics of intense events individually or as they form clusters. Understanding the geometrical and dynamical descriptions of these intense events allow scientists to get closer to a more complete understanding of turbulent flows and more accurate models for engineering applications. The paper has been accepted and will appear in IEEE CG&A in 2012.

21.7 SP-Specific Activities

(If multiple awards at the SP institution, provide sub-sections for each award and for common infrastructure/services) A free format section for the SP to highlight activities of their choosing that do not fall into other defined sections of this report. This could include a description of innovative work, relevant technologies or other developments that has not yet impacted users but is relevant to current or future SP or XSEDE activities.

Longhorn

As a part of the Longhorn XD Vis project, TACC collaborates with NCAR. Version 2.1.0 of VAPOR was released on January 25, 2012, and over 1000 copies of the software have been downloaded to date. Many of the new capabilities of 2.1, such as hedgehog plots and display of 3D scene files, were developed with the direct support of this award.

Work on re-factoring *vaporgui's* internal architecture to facilitate 3rd party extensibility continued in the reporting period. Progress was made on integrating a new object-oriented internal data model into vaporgui. The new data model provides first class support for missing data values, and provides an abstract representation of many commonly used computation grids (e.g. AMR, staggered, stretched, terrain following, etc.). This quarter VAPOR's statistical operators (e.g. histogram), and hedge-hog plot visualizer were converted to the new data model. Conversion of VAPOR's flow integrators is well underway and expected to be completed in Q2CY2012. The two remaining data operators – a 2D data probe and a Python-based variable calculator – should also be converted in Q2.

Work continued on PIOVDC: a parallel API for reading and writing data to a VAPOR Data Collection (VDC) from an MPI-based simulation code. The API, which is an extension to NCAR's Parallel IO library, will enable modelers to write wavelet-encoded data directly, bypassing the need for a post-processing translation step, an important step toward our extreme data handling goals. This quarter a GNU *configure* based build environment was completed, which will facilitate building the library on different platforms. The library was ported to, and benchmarked on, Cray XT5, IBM Power6, and a commodity Linux/Intel cluster. A number of performance optimizations were made and the library is now delivering IO bandwidths approaching raw MPI-IO speeds. Finally, the NCAR team began working with turbulence modelers to integrate the code into the Geophysical High-Order Suite for Turbulence (GHOST) model. The turbulence researchers are proposing a 6144^3 run this summer, on which plans call for using PIOVDC for output. Performance tuning and porting will continue in Q2.

The XD Visualization team engaged in support of the following research projects.

- Lee Lior, NYU used Longhorn for a Molecular dynamics study of DNA mutations. Using Longhorn, Lee has investigated the contribution of local sequence context to the successful incorporation (or mis-incorporation) of certain nucleotides by human DNA polymerase iota. In addition, she has investigated the treatment of particular adenine compounds by human DNA polymerase iota, with special focus on the structural features that allow unique bypass of adenine compound lesions. These studies shed light on the structural methods by which DNA mutations occur in a critical human polymerase.
- Chuan Xiao, UTEP used Longhorn for a Cryo-EM reconstruction of the giant marine virus CroV. Chuan has used Longhorn to process electron microscopic images of the Cafeteria roenbergensis (CroV) marie virus, integrating the 2D images into high-resolution 3D maps of the virus structure. The better understanding of the CroV virus enabled by this project will contribute to the understanding of its evolutionary path and of its impact on ocean ecology.
- Katelyn White, UC Santa Cruz used Spur for a Mangetohydrodynamic study of the solar interior. Using Spur, Katelyn has produced visualizations of the project's high resolution simulations of the processes found in the solar interior to better understand the production of the magnetic activity observed at the solar surface.
- Greg McIvor, UT-Austin used Spur to compute an accurate simulation of black hole mergers. Greg has used Spur to visualize high-resolution black hole merger simulations. These simulations are higher-resolution than prior efforts, and the resulting visualizations will provide more detail to reveal details in the resulting simulated waveforms.

XD Vis PI Gaither co-taught a Visualization and Data Analysis for Science and Engineering course at The University of Texas at Austin. This course had 16 total students made up of 8 undergraduate students and 8 graduate students. All students were required to run on Longhorn and do at least one programming assignment on the resource. Additionally, approximately half of the students developed their semester long project on Longhorn, with applications ranging from Cosmology to Fluid Mechanics. All of the students finished the course with a proficiency in using remote visualization resources, and more importantly, developing applications on these resources as well.

As part of the Longhorn XD Vis effort, NCAR overhauled the VAPOR user documentation and rewrote it using Drupal. The new documentation is available from http://docs.vapor.ucar.edu/docs.

In support of the Longhorn XD Visualization and Data Analysis (VDA) Services program, SURA is leading an effort to:

- identify established and emerging computational science programs at MSIs;
- promote the use of VDA services to researchers from under-represented groups;
- and coordinate VDA training for researchers from MSIs and under-represented groups.

Sixteen participants from Norfolk State University, Old Dominion University (ODU), University of Virginia Alliance for Computational Science and Engineering (UVACSE), and the Virginia Institute for Marine Science at the College of William and Mary attended the January 20th workshop at ODU. Academic disciplines represented include computer science, engineering, marine science, political science, and geography. Participants from academic research computing support at ODU and UVACSE attended to learn about the potential of ParaView and ViSit to enhance visualization services offered at their institutions.

Twenty-four participants attended the January 30th workshop at Florida A&M University. The participants included fifteen from Florida A & M University and nine from Florida State University. Academic disciplines represented include computer science, physics, library science, mathematics, and earth, ocean and atmospheric sciences. Florida A & M University is a doctoral granting HBCU with PhD programs in biomedical engineering, chemical engineering, civil engineering, electrical engineering, mechanical engineering, industrial engineering, pharmaceutical sciences, physics, and environmental science.

Registration is open for the April 9th workshop hosted by Clark Atlanta University. This workshop is an inaugural event for Clark Atlanta University's new visualization lab which will provide visualization services to the campus and will include some of the same software tools (ViSit and ParaView) that are offered on Longhorn. Clark Atlanta is a doctoral degree granting HBCU with PhD programs in the Biological Sciences and Chemistry and also offers masters programs in Computer Science and Physics.

In February, SURA assisted Kofi Nyarko, Associate Professor and Director of the Engineering Visualization Research Laboratory at Morgan State University's School of Engineering to obtain startup allocations on Longhorn and add user accounts. Dr. Nyarko and some of his students attended the June 2011 "Introduction to Scientific Visualization using Longhorn" at SURA. Currently, Dr. Nyarko has several students experimenting with EnVision on Longhorn. SURA is also providing guidance to Mark Jack, Assistant Professor of Physics at Florida A&M University, who is preparing an XSEDE Research Allocation Request for Ranger and Longhorn.

The next workshop will be held on April 9th at Clark Atlanta University in Atlanta, Georgia.

SURA will continue to provide support to Mark Jack of Florida A&M University and Dr. Nyarko of Morgan State University.

21.8 Publications

All publications and presentations by SP-funded staff members and publications that utilized SP resources in the prior reporting period in bibliography format. *(There may be overlap between SP-funded publications reported here and XSEDE-funded publications reported in the main report. If gray, err on the side of redundancy in both locations.)*

Jeong B., Navratil P., Gaither K., Abram G., Johnson G. "Configurable Data Prefetching Scheme for Interactive Visualization of Large-Scale Volume Data" *Proceedings of Visualization and Data Analysis (VDA) 2012*

Navratil P., Fussell D., Lin C., Childs H. "Dynamic Scheduling for Large-Scale Distributed-Memory Ray Tracing" *Proceedings of Eurographics Symposium on Parallel Graphics and Visualization (EGPGV) 2012*

Navratil P. "General-Purpose GPU Programming: Hardware Considerations" *Guest lecture for Dr. Calvin Lin's (UT at Austin) parallel programming class (CS 380P)*

Navratil P., Johnson G., Westing B. "Visualization Clusters: from Tiled Displays to Remote Visualization" *IEEE Cluster 2011*

Navratil P. "Visualization Clusters: from Tiled Displays to Remote Visualization" *Invited talk at Microsoft Research*

Navratil P. "Texas A&M GPU Workshop" *Invited talk at Texas A&M University*

Terrel A., Mandli K. "ManyClaw: Slicing and dicing Riemann solvers for next generation highly parallel architectures" *TACC-Intel Highly Parallel Computing Symposium*

Arora R., Bangalore P., Mernik M. "Formal and Practical Aspects of Domain-Specific Languages: Recent Developments" *Authored book chapter*

Arora R., Bangalore P., Mernik M. "A FRAMEWORK FOR SEMI-AUTOMATIC EXPLICIT PARALLELIZATION" *Authored book chapter*

Marker B., Terrel A., Poulson J., Batory D., van de Geijn R. "Mechanizing the Expert Dense Linear Algebra Developer" *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*

21.9 Metrics

21.9.1 Standard systems metrics

Ranger

Total Normalized Service Units (NUs)

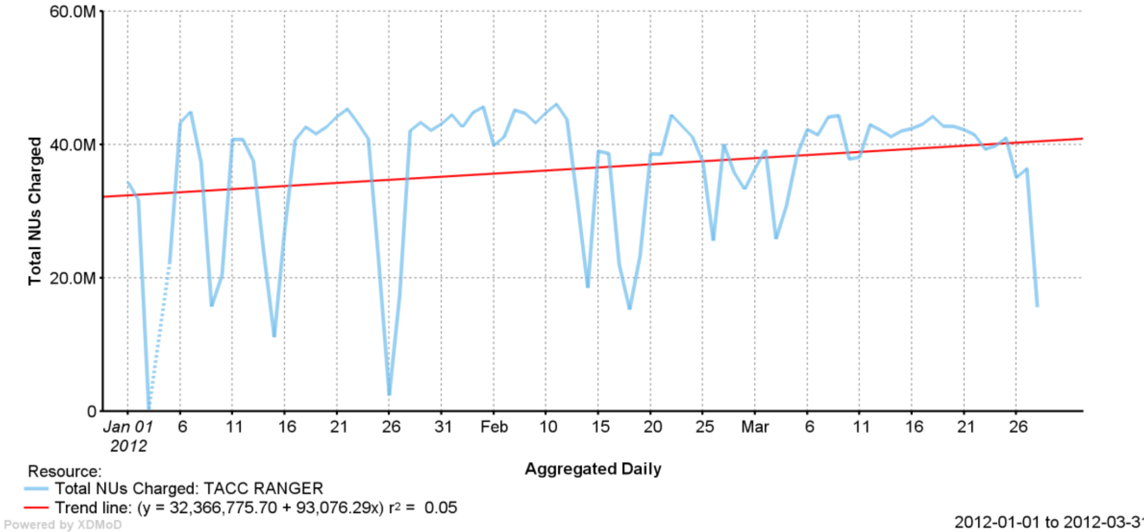

Total NUs Charged by Job Size

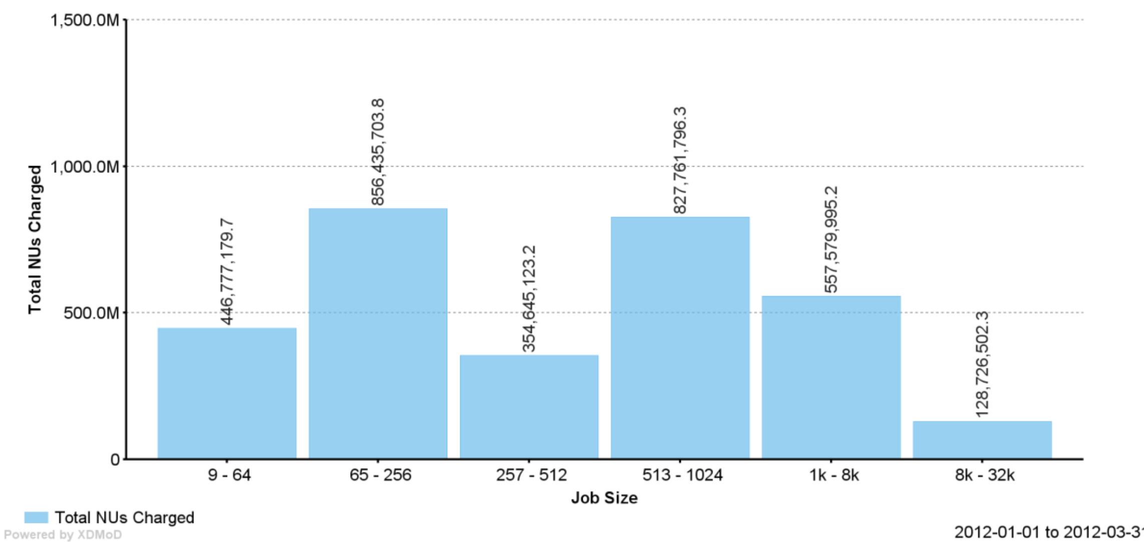

Avg Wall Hours Per Job by Job Size

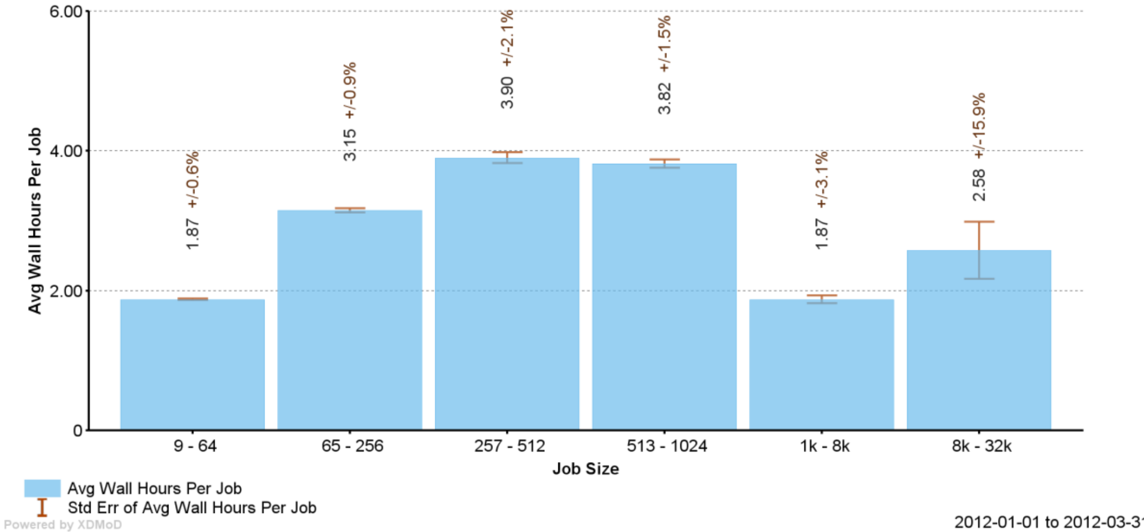

Avg Wait Hours Per Job by Job Size

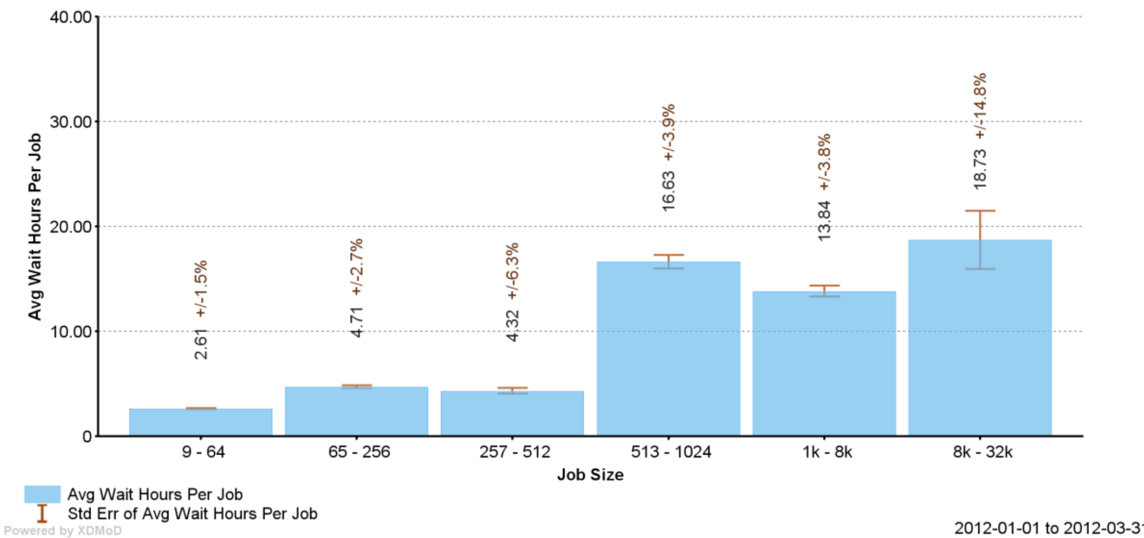

User Expansion Factor by Job Size

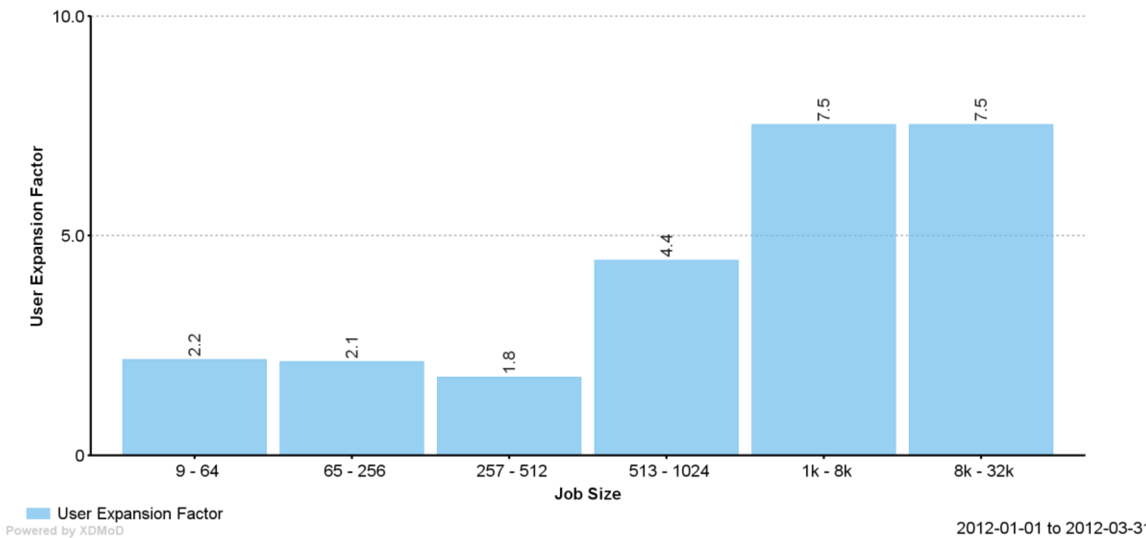

Total NUs Charged by Job Wall Time

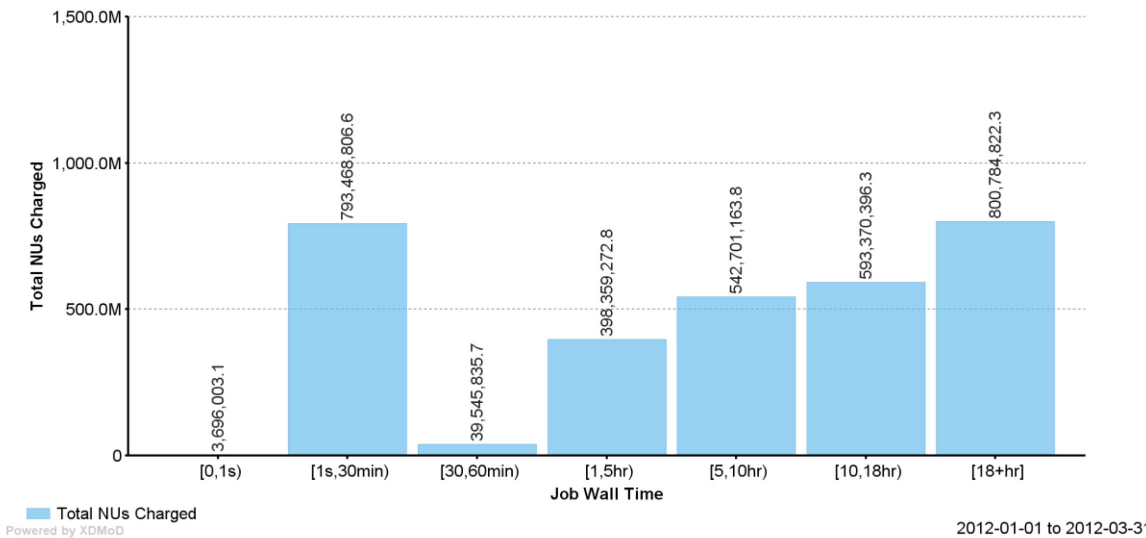

User Expansion Factor by Job Wall Time

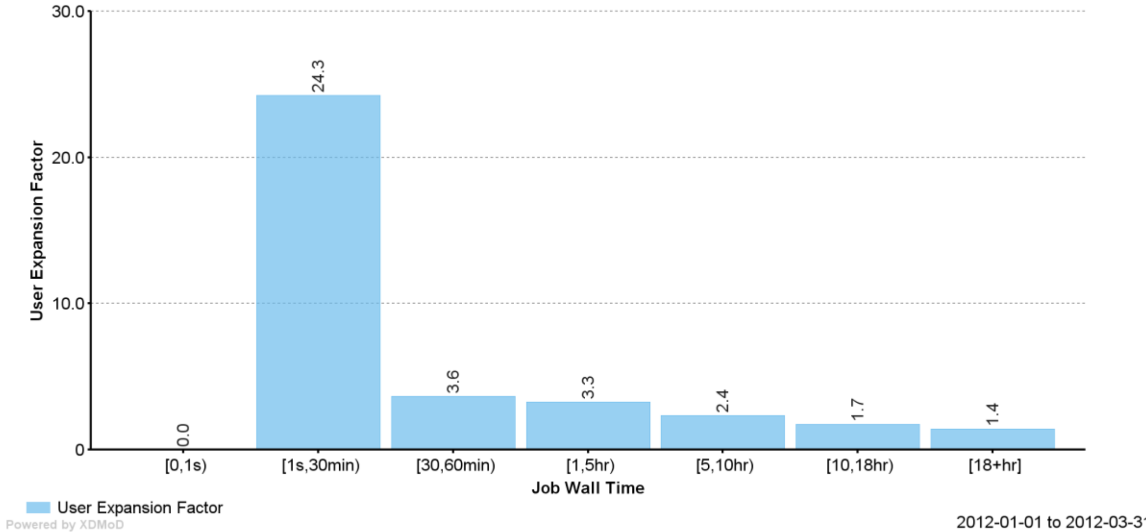

Total NUs Charged by Field of Science

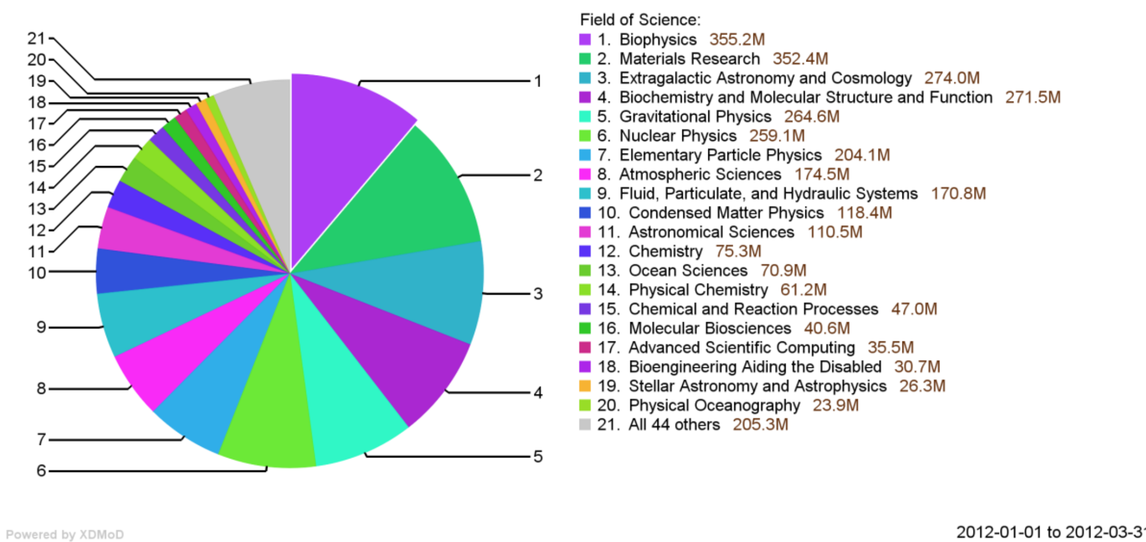

Total NUs Charged by Institution

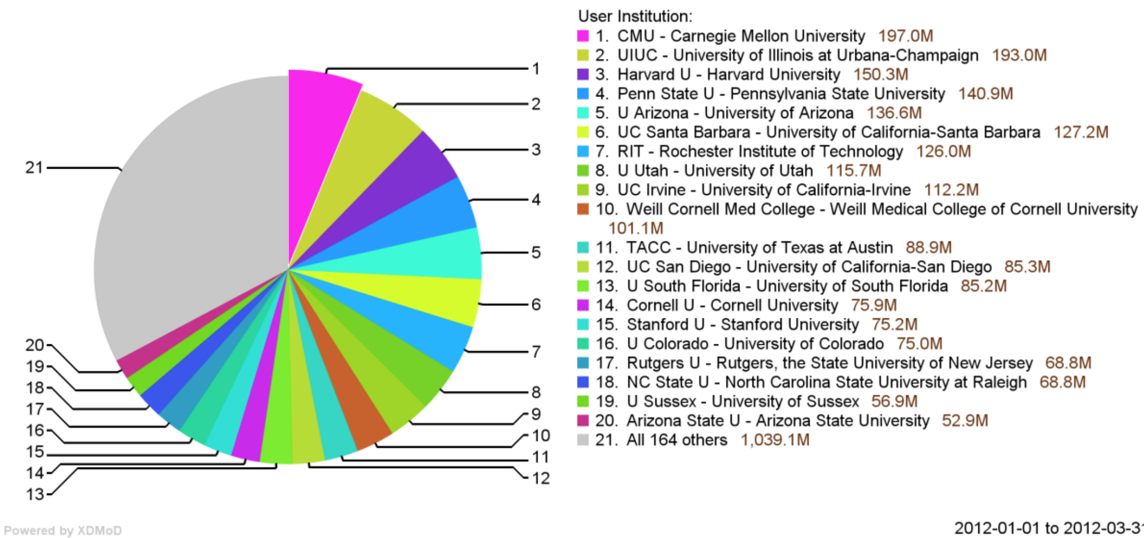

Total NUs Charged by Principal Investigator

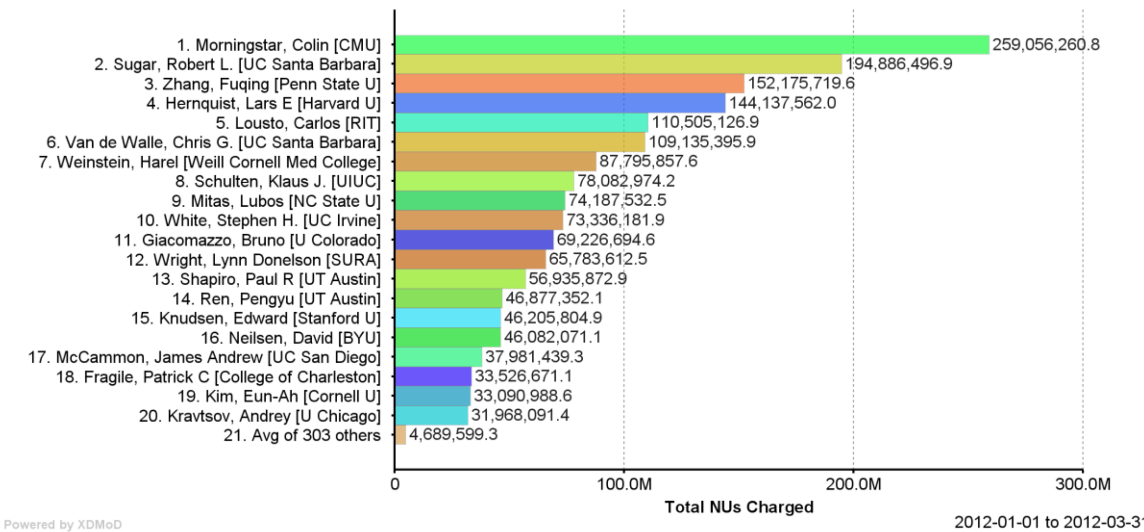

<u>Lonestar</u>

Total NUs Charged

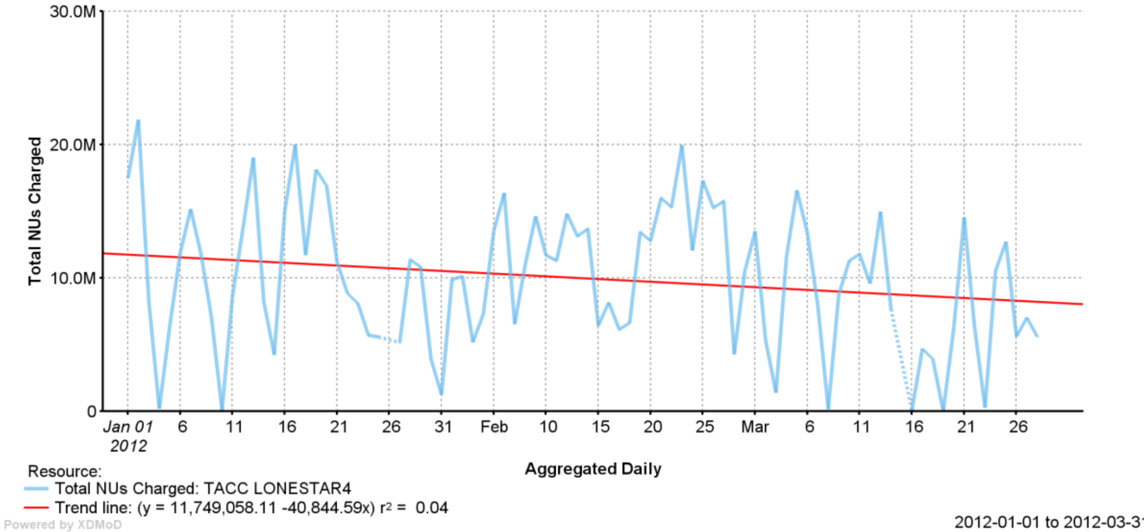

Total NUs Charged by Job Size

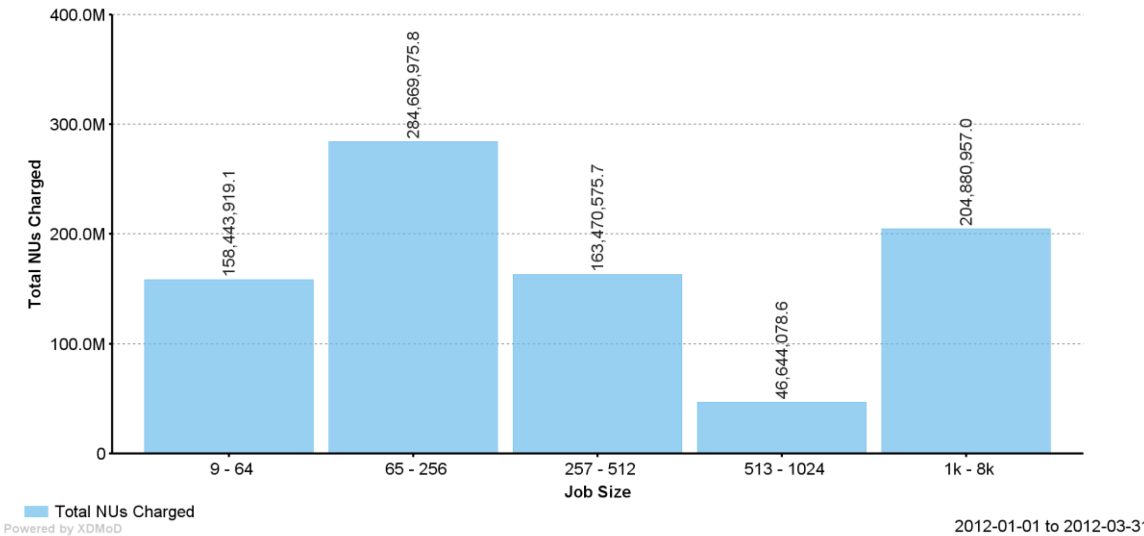

Avg Wall Hours Per Job by Job Size

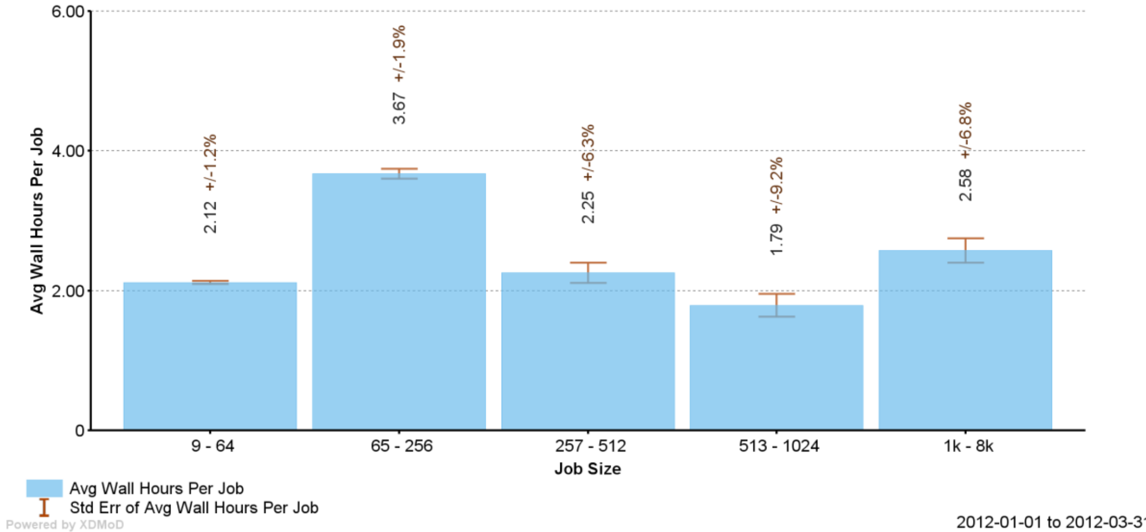

Avg Wait Hours Per Job by Job Size

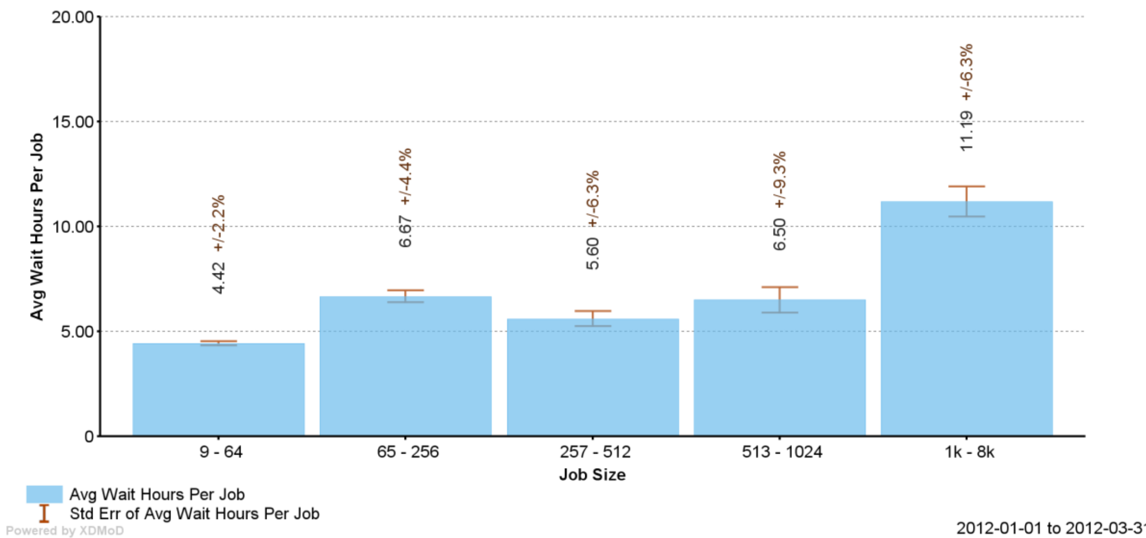

User Expansion Factor by Job Size

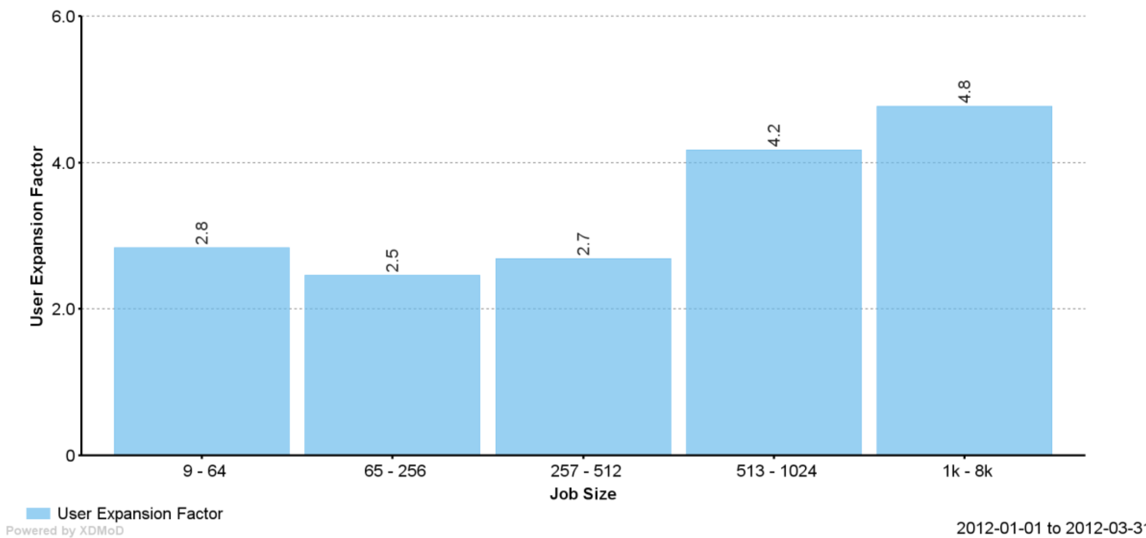

Total NUs Charged by Job Wall Time

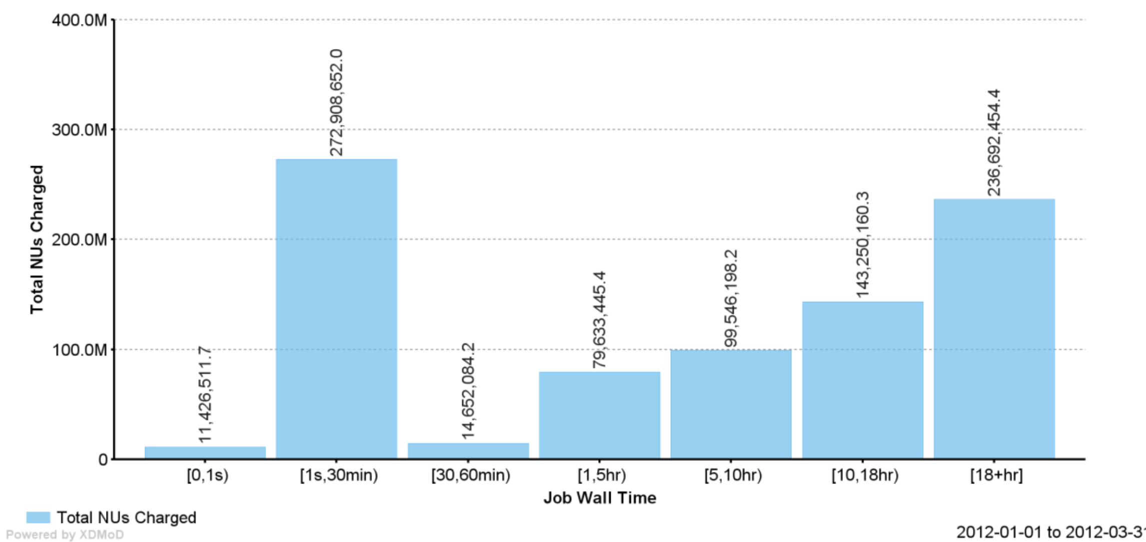

User Expansion Factor by Job Wall Time

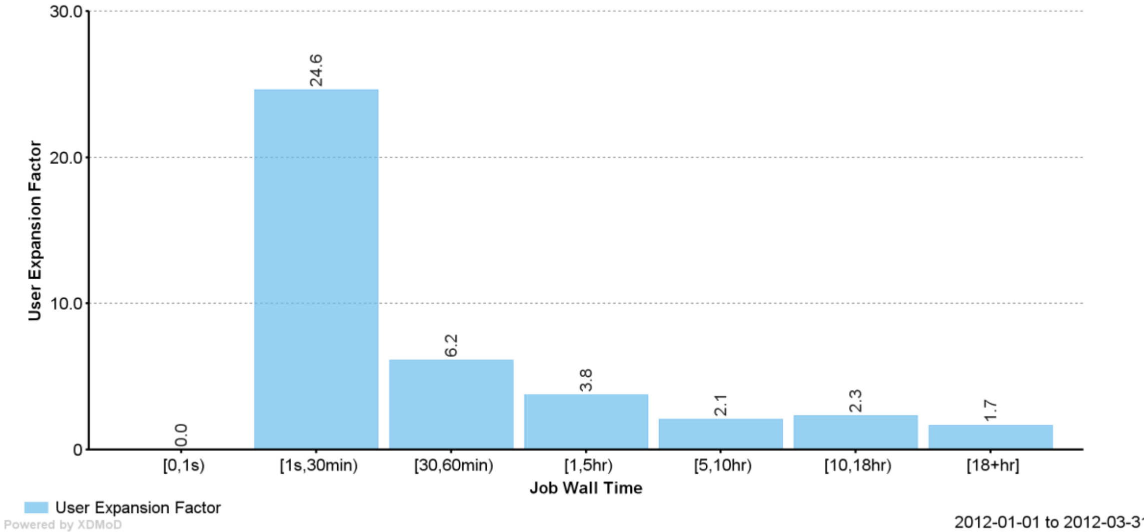

Total NUs Charged by Field of Science

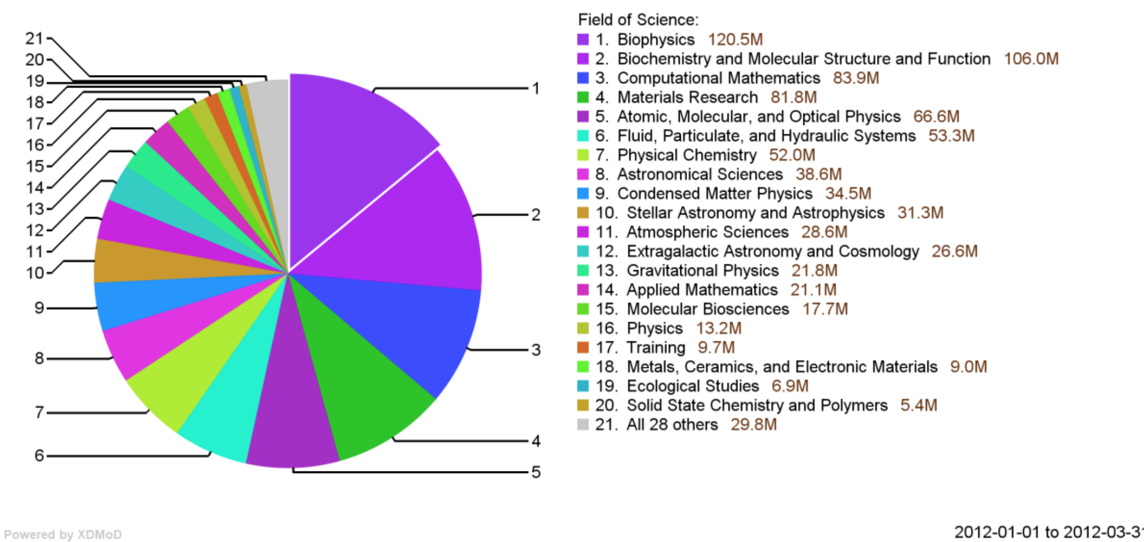

Total NUs Charged by Institution

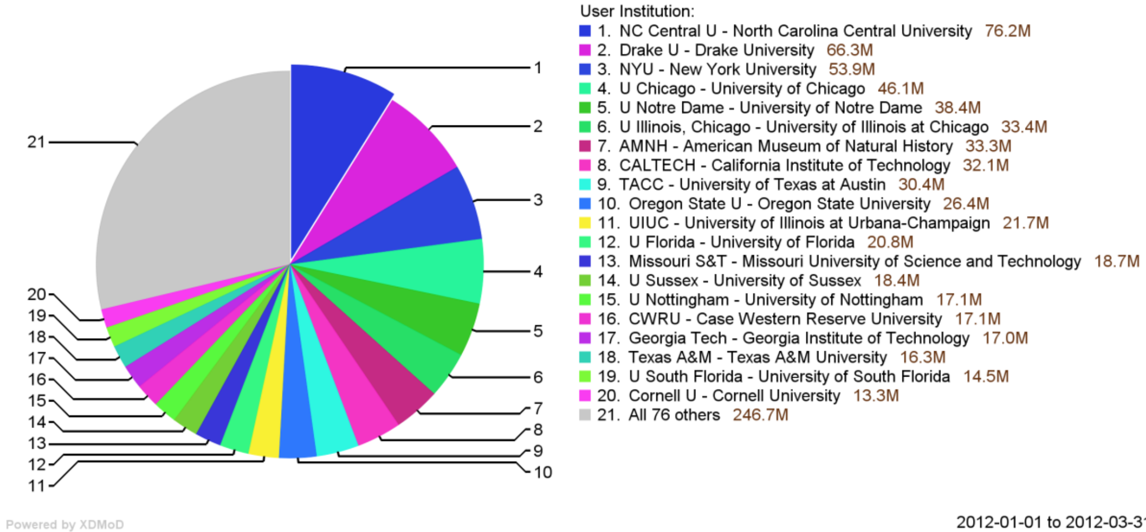

Total NUs Charged by Principal Investigator

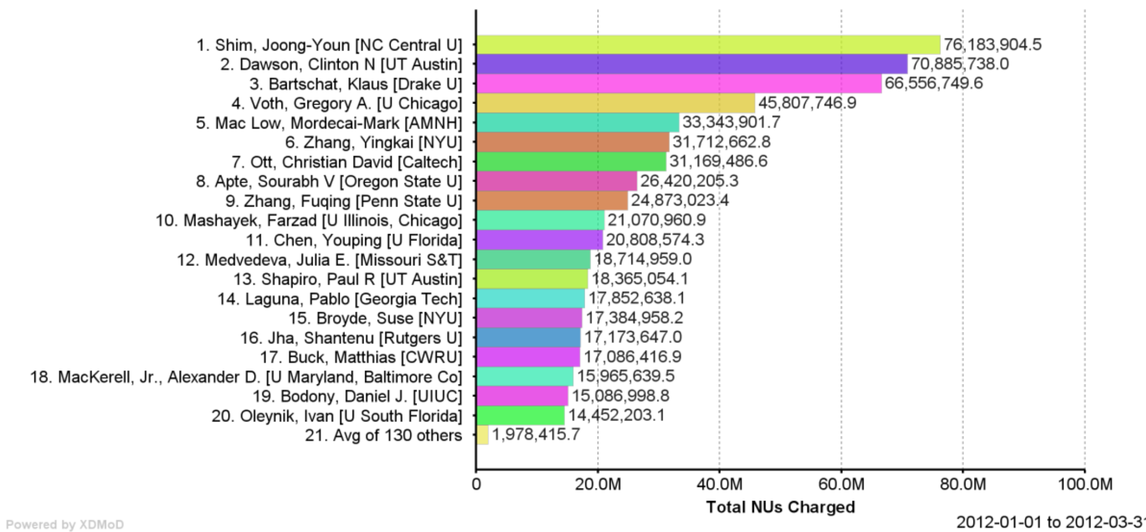

<u>Longhorn</u>

<u>Total NUs Charged</u>

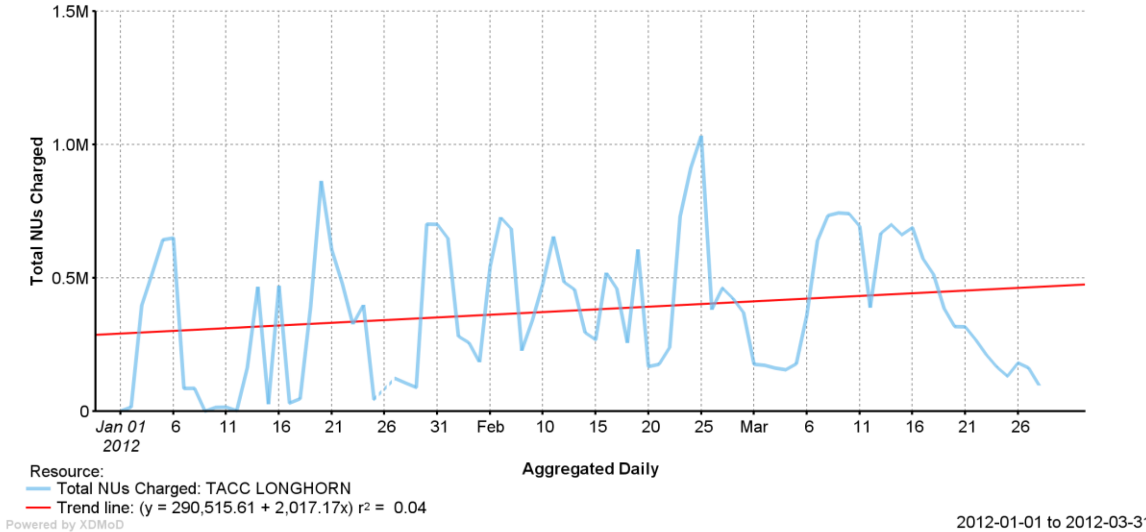

Total NUs Charged by Job Size

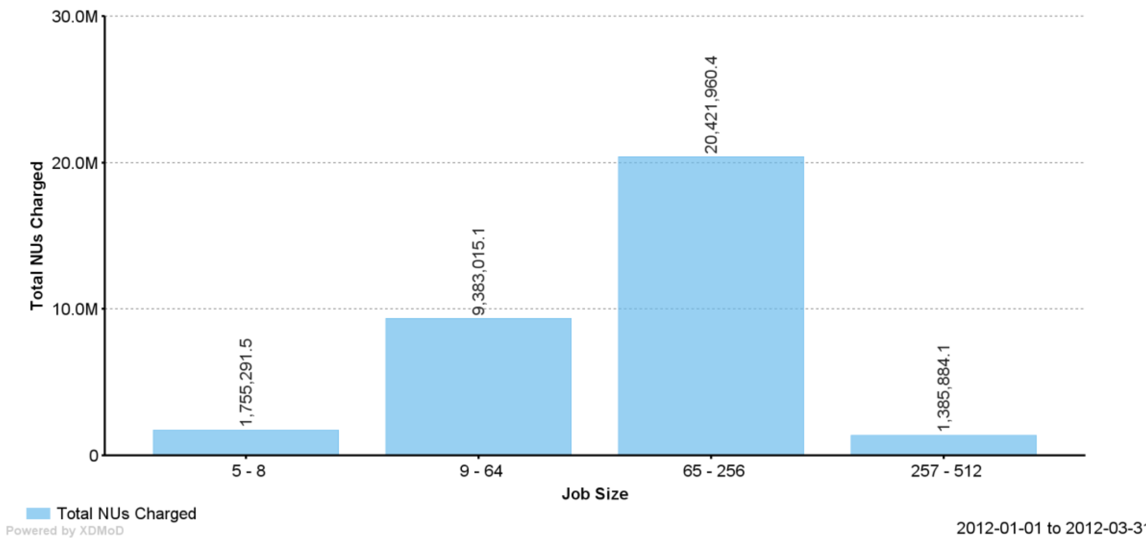

Avg Wall Hours Per Job by Job Size

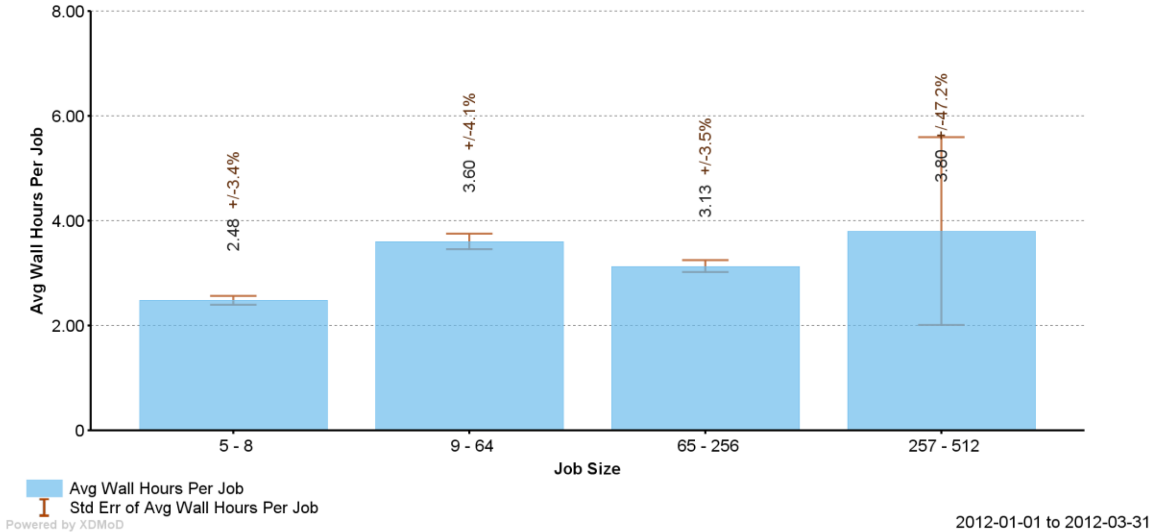

Avg Wait Hours Per Job by Job Size

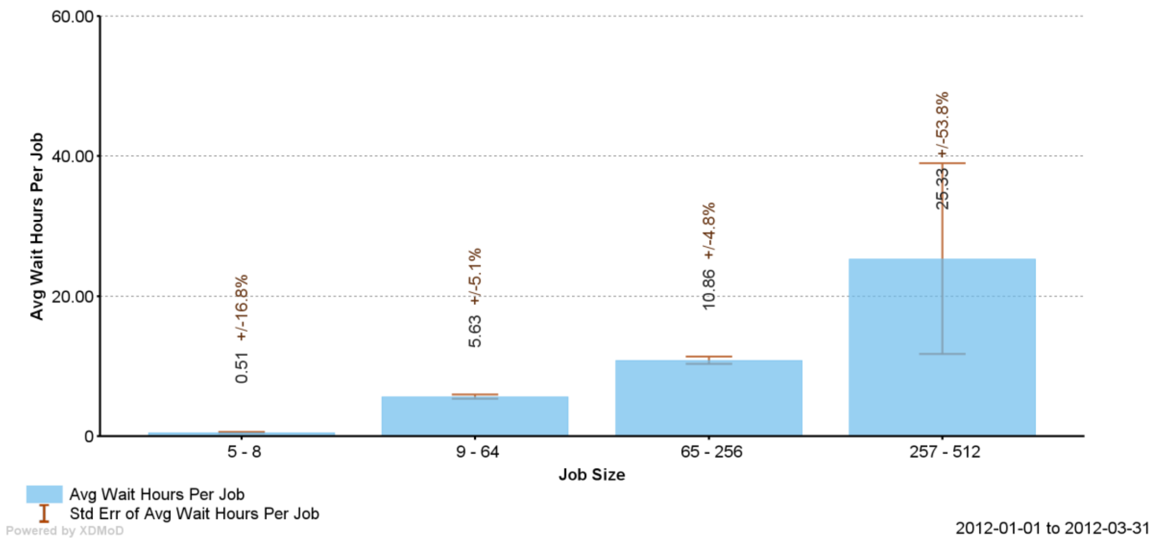

User Expansion Factor by Job Size

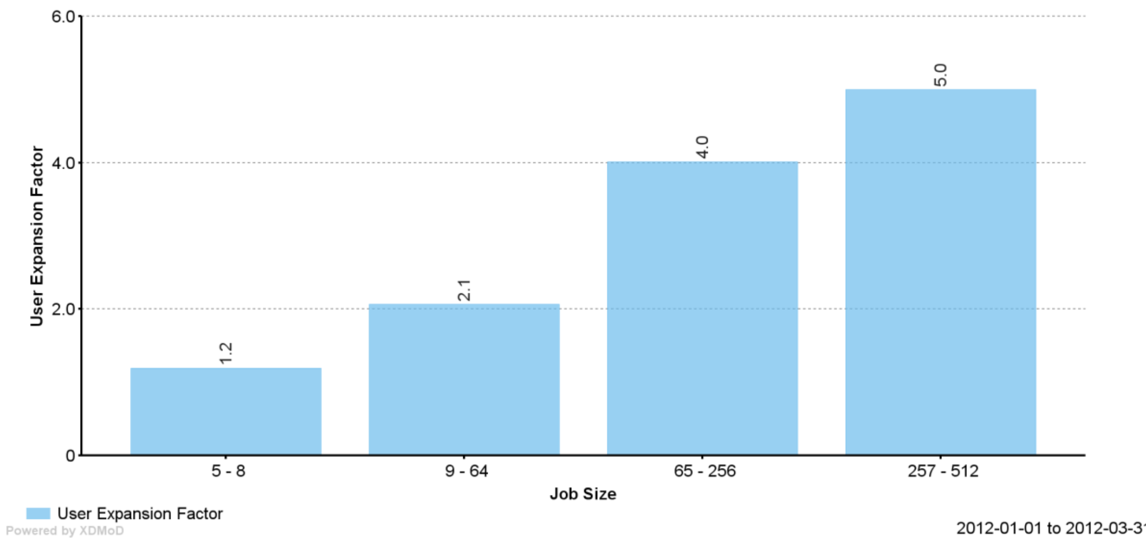

Total NUs Charged by Job Wall Time

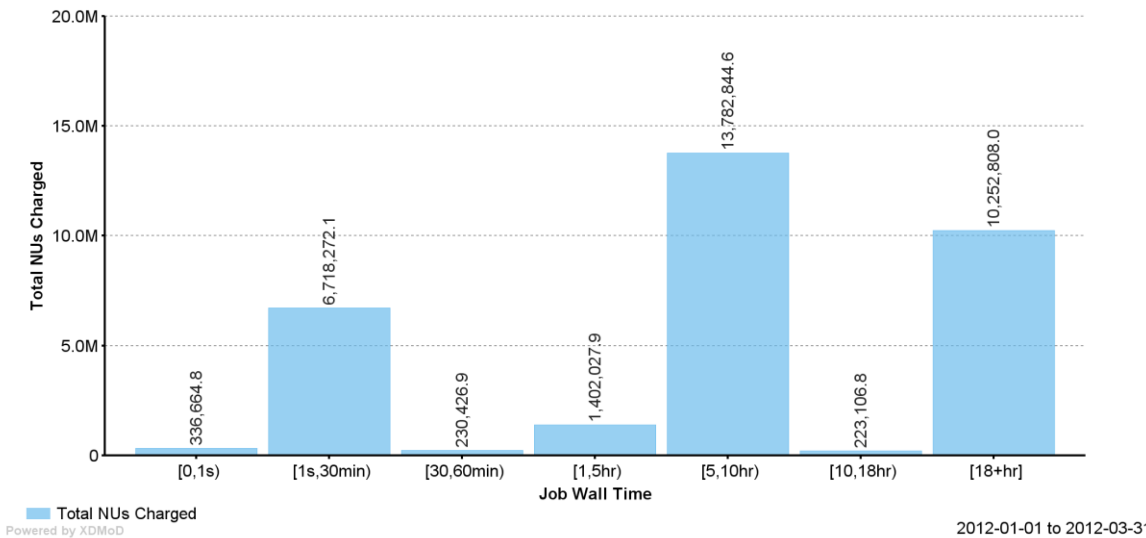

User Expansion Factor by Job Wall Time

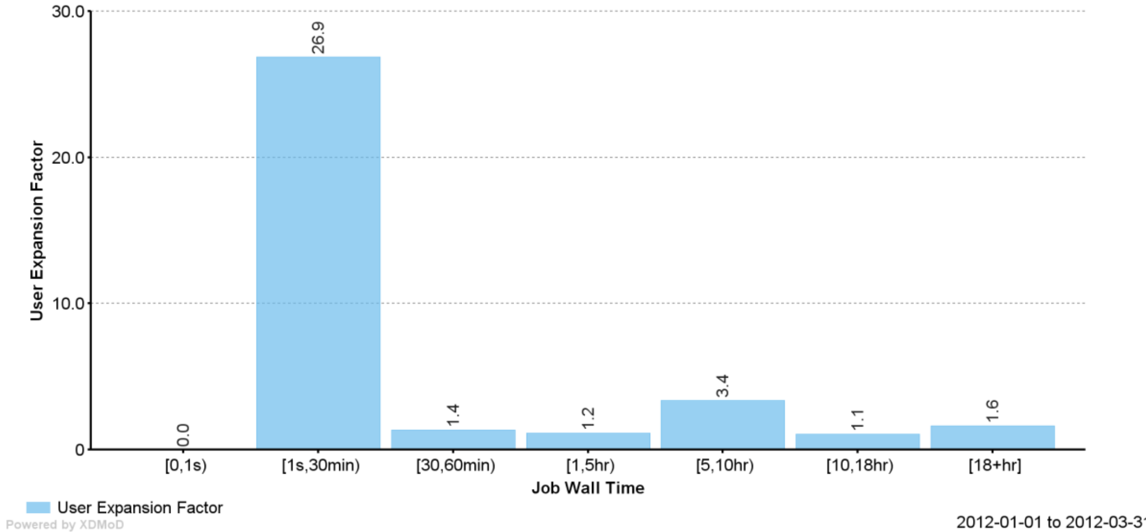

Total NUs Charged by Field of Science

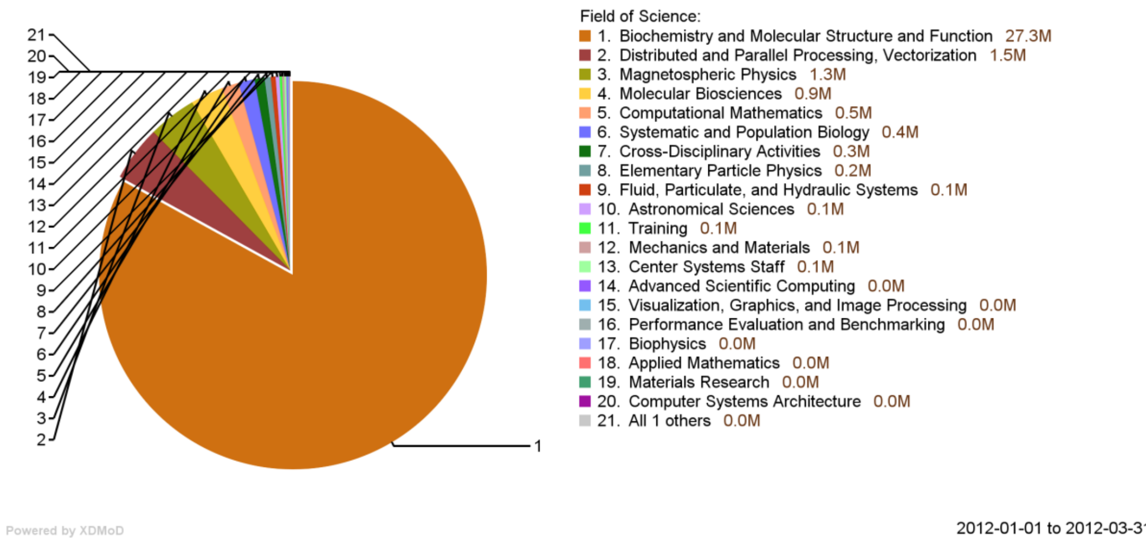

Total NUs Charged by Institution

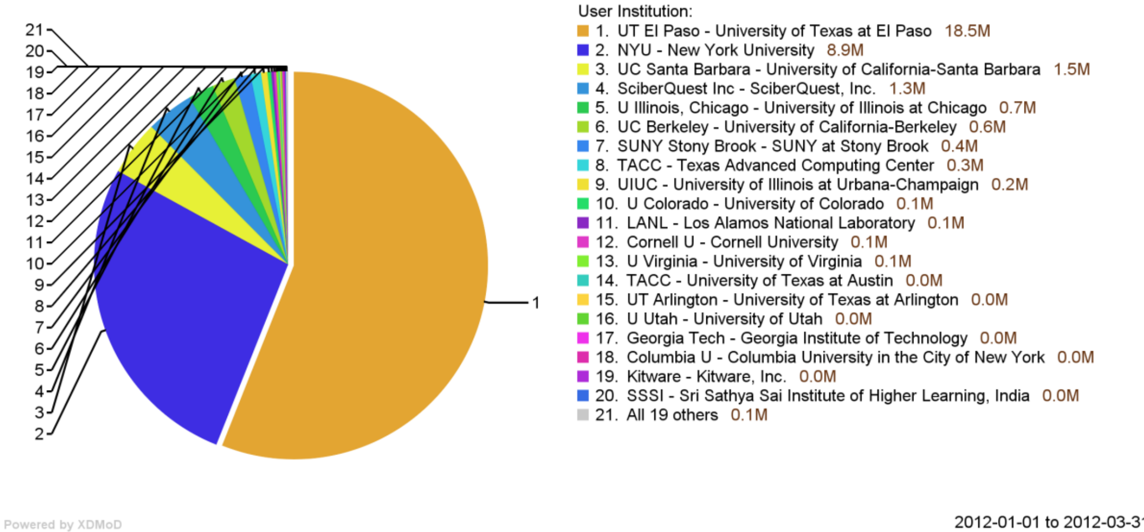

Total NUs Charged by Principal Investigator

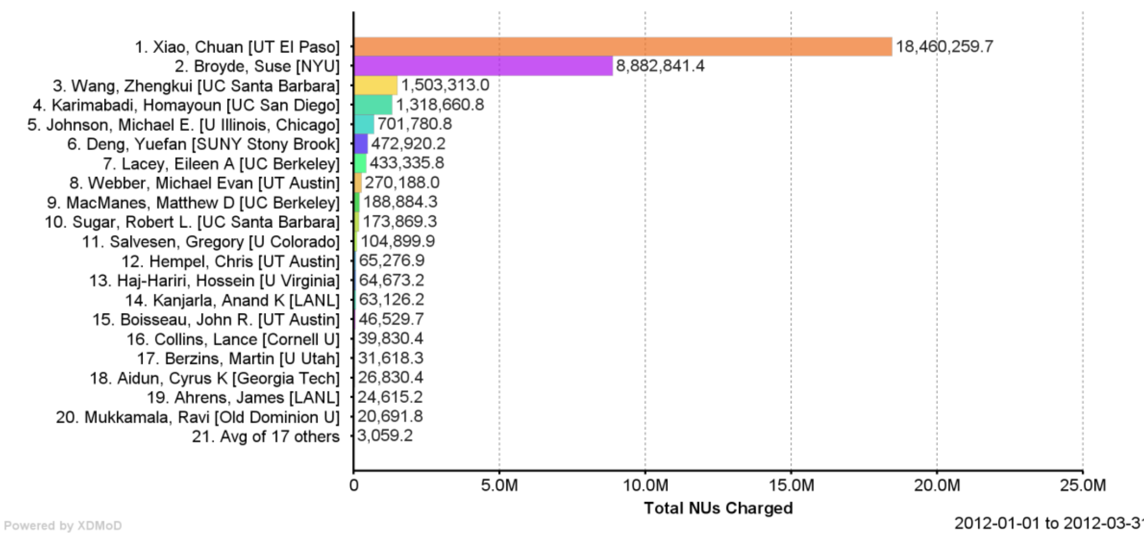

Spur

Total NUs Charged

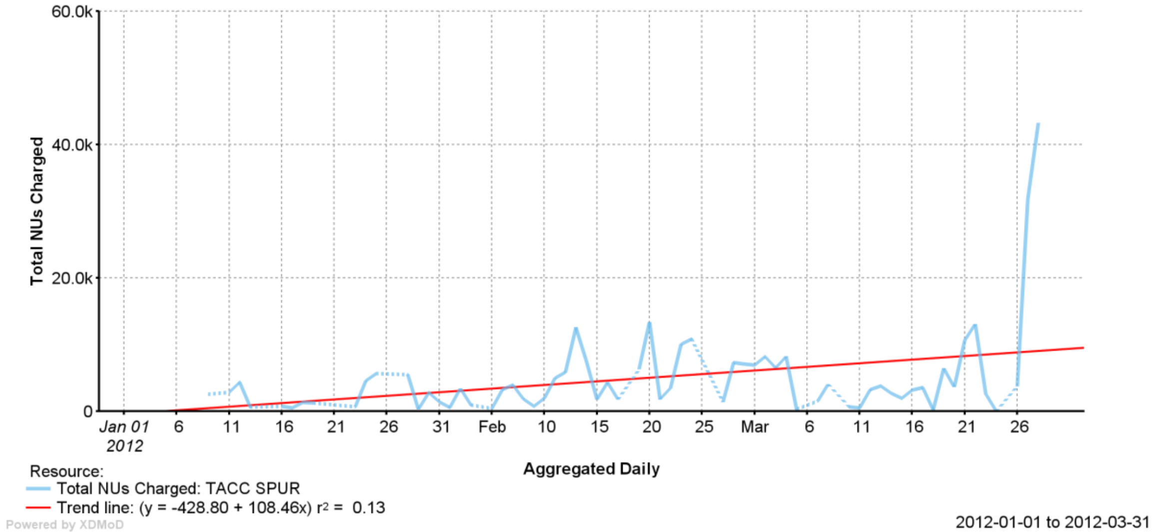

Total NUs Charged by Job Size

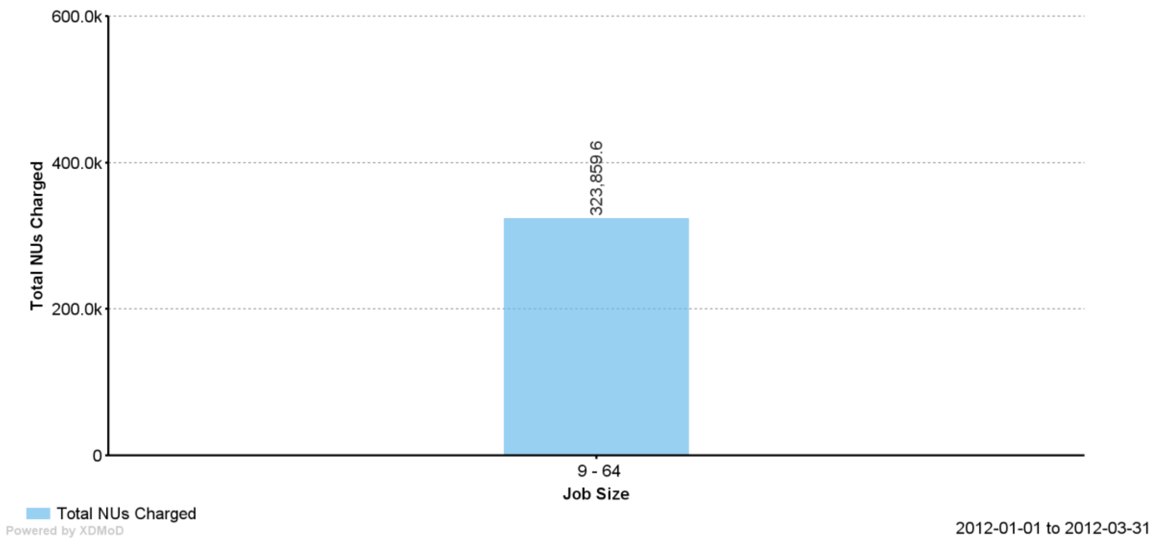

Avg Wall Hours Per Job by Job Size

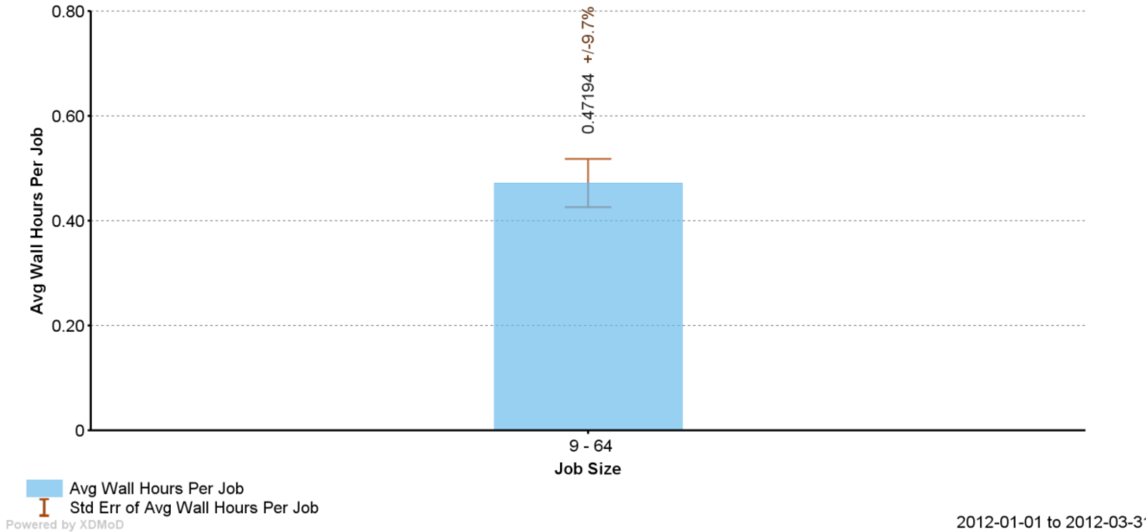

<u>Avg Wait Hours Per Job by Job Size</u>

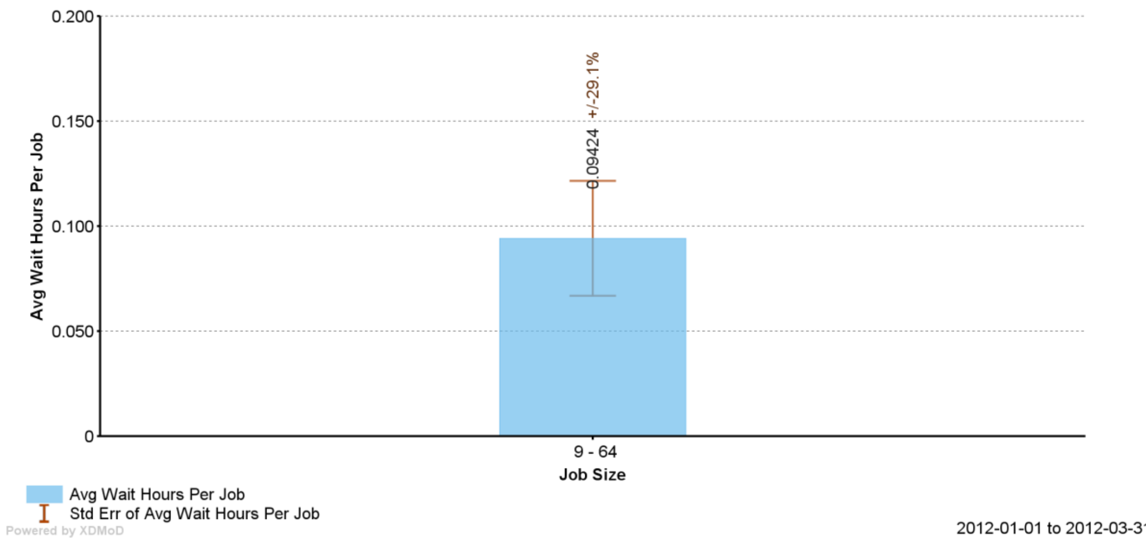

<u>User Expansion Factor by Job Size</u>

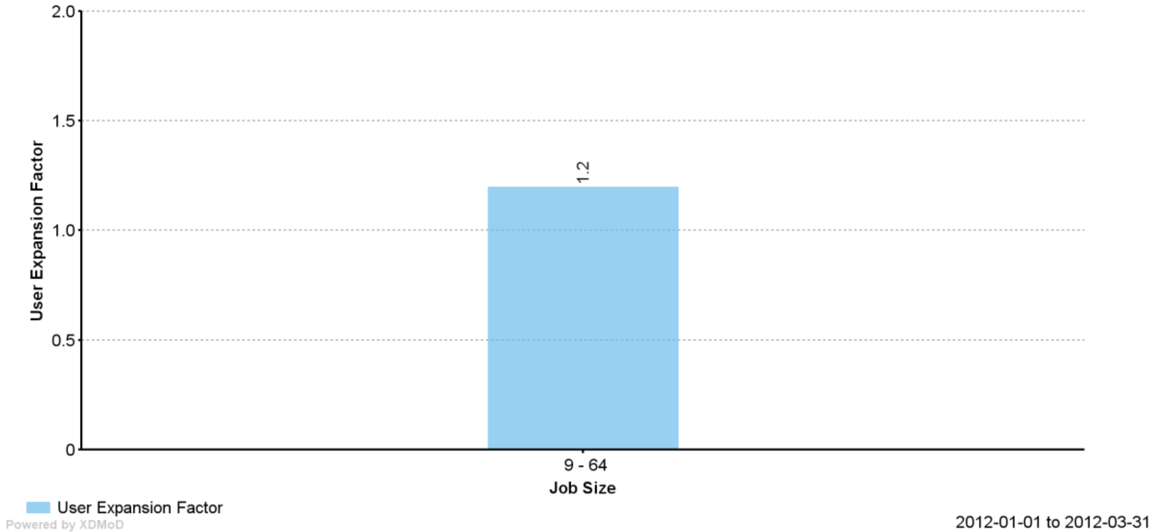

Total NUs Charged by Job Wall Time

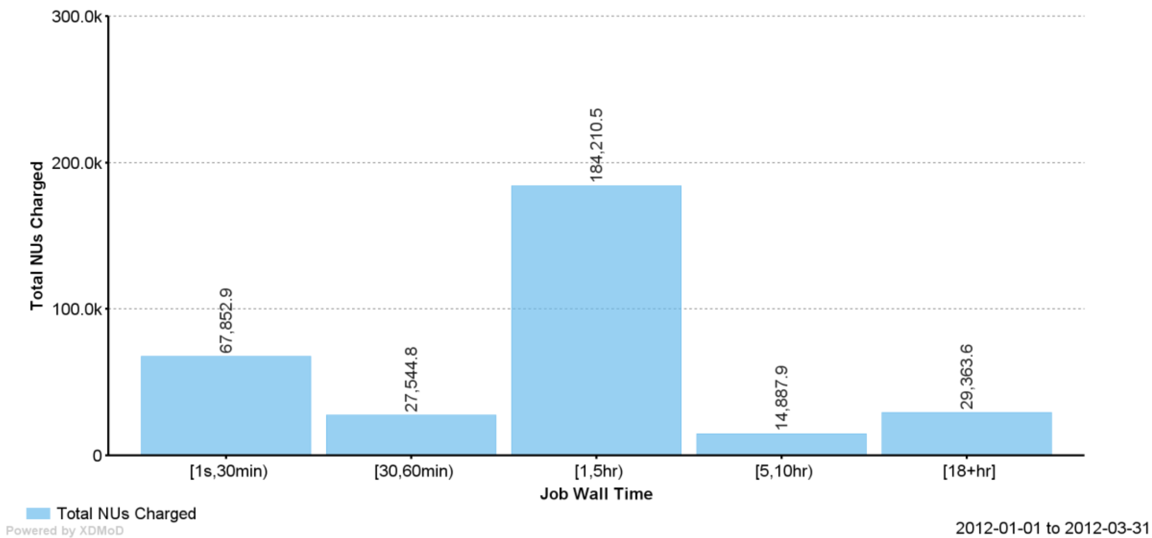

User Expansion Factor by Job Wall Time

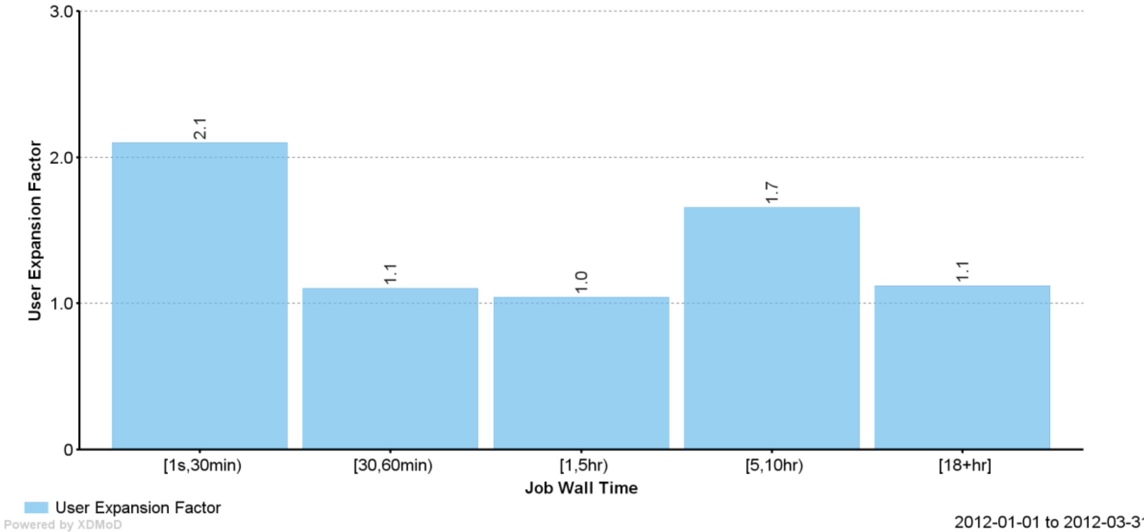

Total NUs Charged by Field of Science

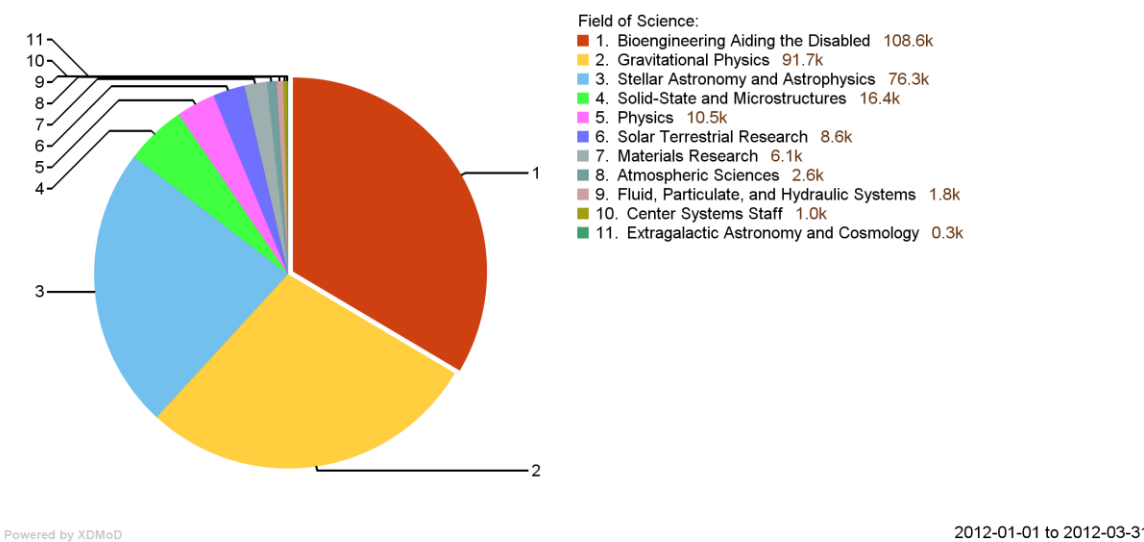

Total NUs Charged by Institution

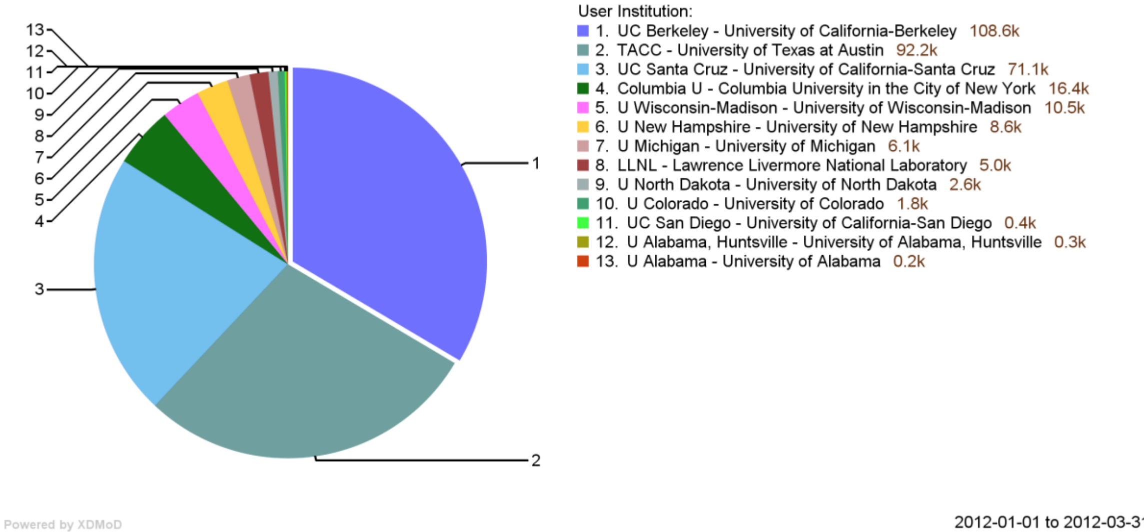

Total NUs Charged by Principal Investigator

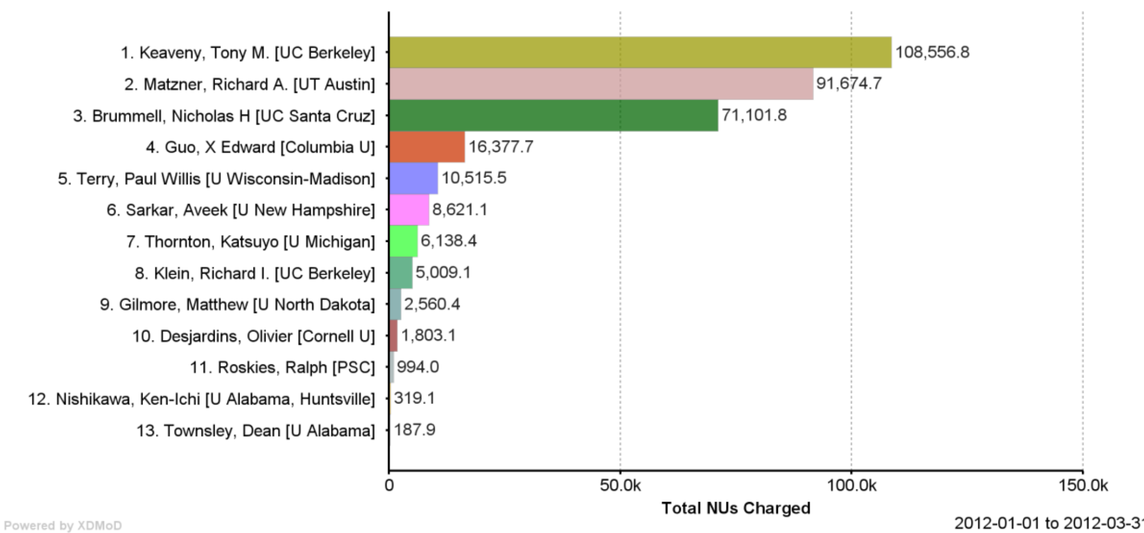

21.9.2 Standard User Assistance Metrics

TACC staff members continue to provide trouble ticket support via the XSEDE ticket system and the TACC Consulting System. 379 tickets, submitted via the XSEDE ticket system, were handled by TACC staff during the report period with 356 being closed. Both trouble ticket systems are monitored 7x24x365 and approximately 25 TACC staff members are engaged in this front-line support activity. The following table indicates the number of tickets opened, closed, and a breakdown of the ticket category.

Issue Category	Number of tickets opened	Number of tickets closed
Jobs/Batch Queues	100	93
Software/Applications	75	66
Login/Access Issues	113	112

System Issues	7	7
Account Issues	18	18
Filesystem Issues	10	8
Other	56	52

TACC ticket resolution times by category from the XSEDE ticket system.

Time to Resolution	account issues	file systems	grid software	jobs/batch queues	login/access issues	mss/data issues	network issues	software/apps	system issues	other
0-1 hr	1			1	8	1		1		
1-24 hr	6			11	40	3		10	4	6
1-7 d	6			19	21	5		12		4
1-2 wk	1	3	1	21	11	3	2	12	2	4
> 2 wk	1	3	1	39	8	7	1	27	1	2
Still Open		1		6	2		1	8		1

XSEDE users also may submit requests for assistance via the TACC User Portal. During the reporting period 117 tickets were submitted via through the TUP; 95 have been resolved, 18 are pending user response, and 4 are in progress.

21.9.3 SP-specific Metrics

Allocation usage in section 1.9.1 reflects utilization of TACC resources by the XSEDE user community. There are allocation pools on TACC resources for the non-XSEDE community; the following table indicates the breakdown of available allocation and usage during the reporting period for both communities. Allocation and usage information is reported in system units (SUs) with an SU being a core hour.

TG/UT Quarter Usage

System	SUs Available	TG SUs Delivered	TG Usage (%)	UT SUs Delivered	UT Usage (%)	Total SUs Delivered
Ranger	134,611,200	104,228,507	89.74	11,917,876	10.26	116,146,383
Lonestar	50,829,320	12,869,647	38	20,832,280	62	33,701,927
Longhorn	4,454,912	821,075	46	973,122	54	1,794,197
Spur	273,600	11,603	59	7,933	41	19,536

The following table contains uptime statistics for the reporting period for TACC compute, visualization, and storage resources.

TACC Resource Uptime Statistics

	Lonestar					Spur				
	Uptime					Uptime				
	PM		Outage			PM		Outage		
Month	#	Hrs	#	Hrs	%Up	#	Hrs	#	Hrs	%Up
2012-1	1	9.50	0	0.00	98.72	1	59.00	0	0.00	92.07

2012-2	0	0.00	0	0.00	100	1	11.50	0	0.00	98.40
2012-3	0	0.00	0	0.00	100	0	0.00	0	0.00	100
	Ranger					Ranch				
	Uptime					Uptime				
	PM		Outage			PM		Outage		
Month	#	Hrs	#	Hrs	%Up	#	Hrs	#	Hrs	%Up
2012-1	1	59.00	0	0.00	92.07	1	5.25	2	13.50	97.48
2012-2	1	11.50	0	0.00	98.4	0	0.00	0	0.00	100
2012-3	0	0.00	0	0.00	100	0	0.00	3	27.50	96.3
	Longhorn									
	Uptime									
	PM		Outage							
Month	#	Hrs	#	Hrs	%Up					
2012-1	0	0.00	2	23.50	96.84					
2012-2	1	9.50	0	0.00	98.68					
2012-3	0	0.00	0	0.00	100					

A XSEDE Project Milestones Update

Content for this appendix is pending finalizing the XSEDE Architecture that is being reworked due to the merging of the XSEDE and XROADS proposals.

B XSEDE Schedule with Progress Update

WBS	Task Name	Start	Finish	% Complete
1	**XSEDE**	**1/3/11**	**9/23/16**	**18%**
1.1	**Project Office**	**1/3/11**	**9/23/16**	**5%**
1.1.1	**Project Management and Reporting**	**7/1/11**	**8/25/16**	**15%**
1.1.1.1	**Management**	**7/1/11**	**6/30/16**	**15%**
1.1.1.1.1	Ongoing: Senior Management Team - Oversight and management of XSEDE	7/1/11	6/30/16	15%
1.1.1.1.2	Ongoing: Business and finance office support - subaward management	7/1/11	6/30/16	15%
1.1.1.1.3	Ongoing: Annual Planning, Budgeting, Change Control resulting from SEMP Spiral Design Process	7/1/11	6/30/16	15%
1.1.1.2	**Project Reporting**	**10/3/11**	**8/25/16**	**11%**
1.1.1.2.1	Prepare 2011 Q3 Quarterly Report	10/3/11	10/28/11	100%
1.1.1.2.2	Milestone: 2011 Q3 Quarterly Report completed	10/28/11	10/28/11	100%
1.1.1.2.3	Prepare 2011 Q4 Quarterly Report	1/2/12	1/27/12	100%
1.1.1.2.4	Milestone: 2011 Q4 Quarterly Report completed	1/27/12	1/27/12	100%
1.1.1.2.5	Prepare 2012 Q1 Quarterly Report	4/2/12	4/27/12	100%
1.1.1.2.6	Milestone: 2012 Q1 Quarterly Report completed	4/27/12	4/27/12	100%
1.1.1.2.7	Prepare Y1 Annual Report	7/2/12	8/24/12	0%
1.1.1.2.8	Milestone: Y1 Annual Report completed	8/27/12	9/21/12	0%
1.1.1.2.9	Prepare 2012 Q3 Quarterly Report	10/1/12	10/26/12	0%
1.1.1.2.10	Milestone: 2012 Q3 Quarterly Report completed	10/26/12	10/26/12	0%
1.1.1.2.11	Prepare 2012 Q4 Quarterly Report	1/1/13	1/28/13	0%
1.1.1.2.12	Milestone: 2012 Q4 Quarterly Report completed	1/28/13	1/28/13	0%
1.1.1.2.13	Prepare 2013 Q1 Quarterly Report	4/1/13	4/26/13	0%
1.1.1.2.14	Milestone: 2013 Q1 Quarterly Report completed	4/26/13	4/26/13	0%
1.1.1.2.15	Prepare Y2 Annual Report	7/1/13	8/23/13	0%
1.1.1.2.16	Milestone: Y2 Annual Report completed	11/1/13	11/1/13	0%
1.1.1.2.17	Prepare 2013 Q3 Quarterly Report	10/1/13	10/29/13	0%
1.1.1.2.18	Milestone: 2013 Q3 Quarterly Report completed	10/29/13	10/29/13	0%
1.1.1.2.19	Prepare 2013 Q4 Quarterly Report	1/1/14	1/28/14	0%
1.1.1.2.20	Milestone: 2013 Q4 Quarterly Report completed	1/28/14	1/28/14	0%
1.1.1.2.21	Prepare 2014 Q1 Quarterly Report	4/1/14	4/28/14	0%
1.1.1.2.22	Milestone: 2014 Q1 Quarterly Report completed	4/28/14	4/28/14	0%
1.1.1.2.23	Prepare Y3 Annual Report	7/1/14	8/25/14	0%
1.1.1.2.24	Milestone: Y3 Annual Report completed	8/25/14	8/25/14	0%
1.1.1.2.25	Prepare 2014 Q3 Quarterly Report	10/1/14	10/28/14	0%
1.1.1.2.26	Milestone: 2014 Q3 Quarterly Report completed	10/28/14	10/28/14	0%
1.1.1.2.27	Prepare 2014 Q4 Quarterly Report	1/1/15	1/28/15	0%
1.1.1.2.28	Milestone: 2014 Q4 Quarterly Report completed	1/28/15	1/28/15	0%
1.1.1.2.29	Prepare 2015 Q1 Quarterly Report	4/1/15	4/28/15	0%
1.1.1.2.30	Milestone: 2015 Q1 Quarterly Report completed	4/28/15	4/28/15	0%
1.1.1.2.31	Prepare Y4 Annual Report	7/1/15	8/25/15	0%
1.1.1.2.32	Milestone: Y4 Annual Report completed	8/25/15	8/25/15	0%
1.1.1.2.33	Prepare 2015 Q3 Quarterly Report	10/1/15	10/28/15	0%
1.1.1.2.34	Milestone: 2015 Q3 Quarterly Report completed	10/28/15	10/28/15	0%

Page 1

WBS	Task Name	Start	Finish	% Complete
1.1.1.2.35	Prepare 2015 Q4 Quarterly Report	1/1/16	1/28/16	0%
1.1.1.2.36	Milestone: 2015 Q4 Quarterly Report completed	1/28/16	1/28/16	0%
1.1.1.2.37	Prepare 2016 Q1 Quarterly Report	4/1/16	4/28/16	0%
1.1.1.2.38	Milestone: 2016 Q1 Quarterly Report completed	4/28/16	4/28/16	0%
1.1.1.2.39	Prepare Y5 Annual Report	7/1/16	8/25/16	0%
1.1.1.2.40	Milestone: Y5 Annual Report completed	8/25/16	8/25/16	0%
1.1.1.2.41	Prepare Final Report	7/1/16	8/25/16	0%
1.1.1.2.42	Milestone: Final Report completed	8/25/16	8/25/16	0%
1.1.1.3	**Project Management - Risk Management**	**7/1/11**	**4/1/16**	**67%**
1.1.1.3.1	Identify risks with Level 3 WBS Managers	7/1/11	7/28/11	100%
1.1.1.3.2	Review risk register with Level 2 WBS Managers	7/29/11	8/11/11	100%
1.1.1.3.3	Quarterly review of risk register - 2011 Q3	10/3/11	10/3/11	100%
1.1.1.3.4	Quarterly review of risk register - 2011 Q4	1/2/12	1/2/12	100%
1.1.1.3.5	Quarterly review of risk register - 2012 Q1	4/2/12	4/2/12	100%
1.1.1.3.6	Quarterly review of risk register - 2012 Q2	7/2/12	7/2/12	0%
1.1.1.3.7	Quarterly review of risk register - 2012 Q3	10/1/12	10/1/12	0%
1.1.1.3.8	Quarterly review of risk register - 2012 Q4	1/1/13	1/1/13	0%
1.1.1.3.9	Quarterly review of risk register - 2013 Q1	4/1/13	4/1/13	0%
1.1.1.3.10	Quarterly review of risk register - 2013 Q2	7/1/13	7/1/13	0%
1.1.1.3.11	Quarterly review of risk register - 2013 Q3	10/1/13	10/1/13	0%
1.1.1.3.12	Quarterly review of risk register - 2013 Q4	1/1/14	1/1/14	0%
1.1.1.3.13	Quarterly review of risk register - 2014 Q1	4/1/14	4/1/14	0%
1.1.1.3.14	Quarterly review of risk register - 2014 Q2	7/1/14	7/1/14	0%
1.1.1.3.15	Quarterly review of risk register - 2014 Q3	10/1/14	10/1/14	0%
1.1.1.3.16	Quarterly review of risk register - 2014 Q4	1/1/15	1/1/15	0%
1.1.1.3.17	Quarterly review of risk register - 2015 Q1	4/1/15	4/1/15	0%
1.1.1.3.18	Quarterly review of risk register - 2015 Q2	7/1/15	7/1/15	0%
1.1.1.3.19	Quarterly review of risk register - 2015 Q3	10/1/15	10/1/15	0%
1.1.1.3.20	Quarterly review of risk register - 2015 Q4	1/1/16	1/1/16	0%
1.1.1.3.21	Quarterly review of risk register - 2016 Q1	4/1/16	4/1/16	0%
1.1.2	**Systems Engineering**	**7/1/11**	**6/23/16**	**0%**
1.1.2.1	Execute Systems Engineering Management Plan	7/1/11	6/23/16	0%
1.1.3	**Systems Architecture**	**1/3/11**	**3/30/16**	**2%**
1.1.3.1	Deploy Grid Middleware Infrastructure	3/1/11	3/28/11	0%
1.1.3.2	Milestone: Grid Middleware Infrastructure deployed	3/28/11	3/28/11	0%
1.1.3.3	Deploy Data Management software	4/1/11	3/29/12	0%
1.1.3.4	Milestone: Data Management software deployed	3/29/12	3/29/12	0%
1.1.3.5	Deploy Account Management software	3/1/11	3/28/11	0%
1.1.3.6	Milestone: Account Management software deployed	3/28/11	3/28/11	0%
1.1.3.7	Deploy Information Services Infrastructure	3/1/11	3/28/11	0%
1.1.3.8	Milestone: Information Services Infrastructure deployed	3/28/11	3/28/11	0%
1.1.3.9	Deploy Common User Environment	3/1/11	3/28/11	0%

WBS	Task Name	Start	Finish	% Complete
1.1.3.10	Milestone: Common User Environment deployed	3/28/11	3/28/11	0%
1.1.3.11	Deploy System of Systems Test Environment	3/1/11	3/28/11	0%
1.1.3.12	Milestone: System of Systems Test Environment deployed	3/28/11	3/28/11	0%
1.1.3.13	**Spiral 1.0**	**1/3/11**	**12/21/11**	**0%**
1.1.3.13.1	Inc 1.0 Start	1/3/11	1/3/11	0%
1.1.3.13.2	Refine Engineering Plan	1/3/11	1/21/11	0%
1.1.3.13.3	Prod. Baseline 1.0 Approved	1/3/11	1/3/11	0%
1.1.3.13.4	UNICORE Stack (1.0)	1/3/11	3/4/11	0%
1.1.3.13.5	Genesis II Stack (1.0)	1/3/11	2/25/11	0%
1.1.3.13.6	Execution Mgmt. (1.0)	1/3/11	2/25/11	0%
1.1.3.13.7	Execution Mon. Sys. (1.0)	1/3/11	1/28/11	0%
1.1.3.13.8	Information Serv. (1.0)	1/3/11	1/28/11	0%
1.1.3.13.9	GFFS Initial	1/3/11	1/11/11	0%
1.1.3.13.10	XWFS Initial	1/3/11	1/21/11	0%
1.1.3.13.11	Data Management (1.0)	1/3/11	3/25/11	0%
1.1.3.13.12	G&VO Management (1.0)	1/3/11	3/25/11	0%
1.1.3.13.13	API&CLT (1.0)	1/3/11	3/25/11	0%
1.1.3.13.14	GUI Tools (1.0)	1/3/11	1/28/11	0%
1.1.3.13.15	Replica Management Sys (1.0)	1/3/11	3/25/11	0%
1.1.3.13.16	AAM System (1.0)	1/3/11	1/28/11	0%
1.1.3.13.17	Ticket System (1.0)	1/3/11	1/28/11	0%
1.1.3.13.18	Ops Ctr Spt System (1.0)	1/3/11	1/28/11	0%
1.1.3.13.19	XSEDE Gateway Framework (1.0)	1/3/11	3/25/11	0%
1.1.3.13.20	Campus Interoperability Framework (1.0)	1/3/11	3/25/11	0%
1.1.3.13.21	Investigate Identity Mgmt and Auth w InCommon	1/3/11	4/22/11	0%
1.1.3.13.22	Security Mgmt and Services(1.0)	1/3/11	6/17/11	0%
1.1.3.13.23	Create Security PSG	9/1/11	10/12/11	0%
1.1.3.13.24	Approval of Security PSG	10/12/11	10/12/11	0%
1.1.3.13.25	XSEDE Certificate Authority	9/1/11	12/21/11	0%
1.1.3.13.26	Unified System Log (1.0)	1/3/11	1/28/11	0%
1.1.3.13.27	CM & Deployment Mgmt System (1.0)	1/3/11	1/21/11	0%
1.1.3.13.28	OSG Bridge (1.0)	1/3/11	2/25/11	0%
1.1.3.13.29	XSEDE Portal	1/3/11	3/25/11	0%
1.1.3.13.30	System Test Planning (1.0)	5/23/11	6/10/11	0%
1.1.3.13.31	Usability Panel Planning (1.0)	5/23/11	6/3/11	0%
1.1.3.13.32	TRR for Inc 1.0	7/4/11	7/5/11	0%
1.1.3.13.33	Inc 1.0 Integration	6/20/11	7/1/11	0%
1.1.3.13.34	Inc 1.0 System Test	7/6/11	8/2/11	0%
1.1.3.13.35	Inc 1.0 Spiral I&T Complete	8/2/11	8/2/11	0%
1.1.3.13.36	Inc 1.0 Deployment Planning	7/4/11	7/15/11	0%
1.1.3.13.37	Training Material Dev. (1.0)	3/28/11	5/20/11	0%
1.1.3.13.38	Inc 1.0 Deployment	8/3/11	8/16/11	0%

WBS	Task Name	Start	Finish	% Complete
1.1.3.13.39	Inc 1.0 Deployment Complete	8/16/11	8/16/11	0%
1.1.3.13.40	XSEDE IOC	8/16/11	8/16/11	0%
1.1.3.14	**Spiral 2.0**	**5/3/11**	**10/18/11**	**0%**
1.1.3.14.1	Prod. Baseline 2.0 Approved	5/3/11	5/3/11	0%
1.1.3.14.2	Syslog Analysis Sys (2.0)	5/3/11	5/3/11	0%
1.1.3.14.3	AAM System (2.0)	5/3/11	6/27/11	0%
1.1.3.14.4	Execution Monitoring Sys (2.0)	5/3/11	6/27/11	0%
1.1.3.14.5	RQOSA Mgmt Sys (2.0)	5/3/11	7/25/11	0%
1.1.3.14.6	Security Mgmt and Services (2.0)	5/3/11	8/22/11	0%
1.1.3.14.7	API&CLT (2.0)	5/3/11	6/27/11	0%
1.1.3.14.8	GUI Tools (2.0)	5/3/11	5/30/11	0%
1.1.3.14.9	CM & Deploy Mgmt (2.0)	5/3/11	5/30/11	0%
1.1.3.14.10	Campus Interop Fwrk (2.0)	5/3/11	7/25/11	0%
1.1.3.14.11	Ticket System (2.0)	5/3/11	6/27/11	0%
1.1.3.14.12	XSEDE Gateway Framework (2.0)	5/3/11	7/25/11	0%
1.1.3.14.13	Replica Mgmt (2.0)	5/3/11	7/25/11	0%
1.1.3.14.14	File System (2.0)	5/3/11	8/22/11	0%
1.1.3.14.15	Ops Ctr Spt Sys (2.0)	5/3/11	5/3/11	0%
1.1.3.14.16	Execution Mgmt (2.0)	5/3/11	6/27/11	0%
1.1.3.14.17	Info Services (2.0)	5/3/11	8/22/11	0%
1.1.3.14.18	Foreign Grid Policy Adapters (2.0)	5/3/11	6/13/11	0%
1.1.3.14.19	Bandwidth Provisioning (2.0)	5/3/11	7/25/11	0%
1.1.3.14.20	XSEDE Portal (2.0)	5/3/11	6/27/11	0%
1.1.3.14.21	G&VO Mgmt (2.0)	5/3/11	7/25/11	0%
1.1.3.14.22	Collab. Tools (2.0)	5/3/11	5/30/11	0%
1.1.3.14.23	Data Management (2.0)	5/3/11	7/25/11	0%
1.1.3.14.24	Inc 2.0 Integration	8/23/11	9/5/11	0%
1.1.3.14.25	Inc 2.0 System Test	9/7/11	9/20/11	0%
1.1.3.14.26	Training Material Dev. (2.0)	5/3/11	6/27/11	0%
1.1.3.14.27	System Test Planning (2.0)	8/15/11	9/2/11	0%
1.1.3.14.28	Usability Panel Planning (2.0)	8/22/11	8/26/11	0%
1.1.3.14.29	TRR for Inc 2.0	9/6/11	9/6/11	0%
1.1.3.14.30	Inc 2.0 Deployment Planning	9/21/11	10/4/11	0%
1.1.3.14.31	Inc 2.0 Deployment	10/5/11	10/18/11	0%
1.1.3.14.32	Inc 2.0 Deployment Complete	10/18/11	10/18/11	0%
1.1.3.14.33	Inc 2.0 Spiral I&T Complete	9/20/11	9/20/11	0%
1.1.3.15	**Spiral 3.0**	**10/18/11**	**4/4/12**	**0%**
1.1.3.15.1	Prod. Baseline 3.0 Approved	10/18/11	10/18/11	0%
1.1.3.15.2	Execution Mgmt (3.0)	10/18/11	12/12/11	0%
1.1.3.15.3	G&VO Mgmt (3.0)	10/18/11	1/9/12	0%
1.1.3.15.4	Data Managementt (3.0)	10/18/11	12/12/11	0%
1.1.3.15.5	Info Services (3.0)	10/18/11	2/6/12	0%

Page 4

WBS	Task Name	Start	Finish	% Complete
1.1.3.15.6	Replica Mgmt Sys (3.0)	10/18/11	10/18/11	0%
1.1.3.15.7	Bandwidth Provisioning (3.0)	10/18/11	10/18/11	0%
1.1.3.15.8	RQOSA Mgmt (3.0)	10/18/11	10/18/11	0%
1.1.3.15.9	AAM Sys (3.0)	10/18/11	12/12/11	0%
1.1.3.15.10	Execution Mon Sys (3.0)	10/18/11	11/14/11	0%
1.1.3.15.11	Collab Tools (3.0)	10/18/11	11/14/11	0%
1.1.3.15.12	API&CLT (3.0)	10/18/11	11/14/11	0%
1.1.3.15.13	GUI Tools (3.0)	10/18/11	10/31/11	0%
1.1.3.15.14	File System (3.0)	10/18/11	1/9/12	0%
1.1.3.15.15	Ops Ctr Support Sys (3.0)	10/18/11	10/18/11	0%
1.1.3.15.16	Syslog Analysis Sys (3.0)	10/18/11	10/31/11	0%
1.1.3.15.17	XSEDE Portal (3.0)	10/18/11	11/14/11	0%
1.1.3.15.18	System Test Planning (3.0)	2/1/12	2/21/12	0%
1.1.3.15.19	Usability Panel Planning (3.0)	2/7/12	2/13/12	0%
1.1.3.15.20	TRR for Inc 3.0	2/22/12	2/22/12	0%
1.1.3.15.21	Security Mgmt and Services (3.0)	10/18/11	1/9/12	0%
1.1.3.15.22	Inc 3.0 Integration	2/7/12	2/20/12	0%
1.1.3.15.23	Inc 3.0 System Test	2/23/12	3/7/12	0%
1.1.3.15.24	Inc 3.0 Spiral I&T Complete	3/7/12	3/7/12	0%
1.1.3.15.25	Training Material Dev. (3.0)	11/7/11	12/30/11	0%
1.1.3.15.26	Inc 3.0 Deployment Planning	3/8/12	3/21/12	0%
1.1.3.15.27	Inc 3.0 Deployment	3/22/12	4/4/12	0%
1.1.3.15.28	Inc 3.0 Deployment Complete	4/4/12	4/4/12	0%
1.1.3.15.29	XSEDE FOC	4/4/12	4/4/12	0%
1.1.3.16	Ongoing: Incremental improvements continue via SEMP Spiral Design Process	4/5/12	3/30/16	0%
1.1.3.17	**Public Facing XSEDE Architecture Document**	**12/5/11**	**3/5/12**	**100%**
1.1.3.17.1	Agreement of contents and level of detail	12/5/11	12/8/11	100%
1.1.3.17.2	Establish time frame to produce public facing architecture document	12/9/11	1/5/12	100%
1.1.3.17.3	Outline for initial level 1 & level 2-decomposition documentation	1/6/12	1/12/12	100%
1.1.3.17.4	First draft of public facing document	1/13/12	2/1/12	100%
1.1.3.17.5	Revise, comment add content to document as necessary	2/2/12	2/9/12	100%
1.1.3.17.6	Architects review first draft with Bachman	2/10/12	2/10/12	100%
1.1.3.17.7	Identify remaining steps to complete first draft	2/13/12	2/16/12	100%
1.1.3.17.8	A&D team including liaisons from SD&I, Security, Campus Bridging and management review first draft	2/17/12	2/27/12	100%
1.1.3.17.9	Architects address comments/revisions and request endorsement from A&D team and liaisons	2/28/12	3/1/12	100%
1.1.3.17.10	First version of the XSEDE Architecture Document (Level 1 & 2 Decomposition) made publicly available at Quarterly Meeting	3/2/12	3/5/12	100%
1.1.3.18	**Campus Bridging - Architectural Response to Stakeholder Requirements**	**1/19/12**	**5/3/12**	**61%**
1.1.3.18.1	Preliminary background work	1/19/12	2/16/12	100%
1.1.3.18.2	Documentation and review of use cases and requirements matrix completed	2/17/12	2/23/12	100%
1.1.3.18.3	Architectural response at a Level 3 Decomposition prepared by the architects	2/24/12	3/22/12	50%
1.1.3.18.4	Stakeholder review of Architectural response	3/23/12	4/19/12	50%

Page 5

WBS	Task Name	Start	Finish	% Complete
1.1.3.18.5	Stakeholder approved Architectural response for Campus Bridging incorporated into next release of Public Facing XSEDE Architecture Document	4/20/12	5/3/12	0%
1.1.3.19	**Gateways - Architectural Response to Stakeholder Requirements**	**5/4/12**	**8/9/12**	**0%**
1.1.3.19.1	Documentation and review of use cases and requirements matrix completed	5/4/12	5/31/12	0%
1.1.3.19.2	Architectural response at a Level 3 Decomposition prepared by the architects	6/1/12	6/28/12	0%
1.1.3.19.3	Stakeholder review of Architectural response	6/29/12	6/29/12	0%
1.1.3.19.4	Stakeholder review of Architectural response (cont.)	7/2/12	7/26/12	0%
1.1.3.19.5	Stakeholder approved Architectural response for Gateways incorporated into next release of Public Facing XSEDE Architecture Document	7/27/12	8/9/12	0%
1.1.3.20	**Security - Architectural Response to Stakeholder Requirements**	**8/10/12**	**11/15/12**	**0%**
1.1.3.20.1	Documentation and review of use cases and requirements matrix completed	8/10/12	9/6/12	0%
1.1.3.20.2	Architectural response at a Level 3 Decomposition prepared by the architects	9/7/12	10/4/12	0%
1.1.3.20.3	Stakeholder review of Architectural response	10/5/12	11/1/12	0%
1.1.3.20.4	Stakeholder approved Architectural response for Security incorporated into next release of Public Facing XSEDE Architecture Document	11/2/12	11/15/12	0%
1.1.3.21	**Integration with Other Architectures - Architectural Response to Stakeholder Requirements**	**11/16/12**	**3/7/13**	**0%**
1.1.3.21.1	Documentation and review of use cases and requirements matrix completed	11/16/12	12/13/12	0%
1.1.3.21.2	Architectural response at a Level 3 Decomposition prepared by the architects	12/14/12	1/25/13	0%
1.1.3.21.3	Stakeholder review of Architectural response	1/28/13	2/21/13	0%
1.1.3.21.4	Stakeholder approved Architectural response for Integration with Other Architectures incorporated into next release of Public Facing XSEDE Architecture Document	2/22/13	3/7/13	0%
1.1.3.22	**Replication and Fault Tolerance - Architectural Response to Stakeholder Requirements**	**3/8/13**	**6/13/13**	**0%**
1.1.3.22.1	Documentation and review of use cases and requirements matrix completed	3/8/13	4/4/13	0%
1.1.3.22.2	Architectural response at a Level 3 Decomposition prepared by the architects	4/5/13	5/2/13	0%
1.1.3.22.3	Stakeholder review of Architectural response	5/3/13	5/30/13	0%
1.1.3.22.4	Stakeholder approved Architectural response for Replication and Fault Tolerance incorporated into next release of Public Facing XSEDE Architecture Document	5/31/13	6/13/13	0%
1.1.3.23	**Interprocess Communication Mechanisms (e.g., REST, WSI) - Architectural Response to Stakeholder Requirements**	**6/14/13**	**9/19/13**	**0%**
1.1.3.23.1	Documentation and review of use cases and requirements matrix completed	6/14/13	6/28/13	0%
1.1.3.23.2	Documentation and review of use cases and requirements matrix completed (cont.)	7/1/13	7/11/13	0%
1.1.3.23.3	Architectural response at a Level 3 Decomposition prepared by the architects	7/12/13	8/8/13	0%
1.1.3.23.4	Stakeholder review of Architectural response	8/9/13	9/5/13	0%
1.1.3.23.5	Stakeholder approved Architectural response for Interprocess Communication Mechanisms incorporated into next release of Public Facing XSEDE Architecture Document	9/6/13	9/19/13	0%
1.1.4	**External Relations**	**1/3/11**	**4/19/16**	**0%**
1.1.4.1	**Generate Publications: Highlights (Science, EOT, Digital Resources)**	**3/31/11**	**3/25/16**	**0%**
1.1.4.1.1	SciHi subcommittee and chair from XSEDE ER established	3/31/11	3/31/11	0%
1.1.4.1.2	Story assignments established	4/1/11	4/1/11	0%
1.1.4.1.3	Story choices approved by XSEDE leadership	4/15/11	4/15/11	0%
1.1.4.1.4	About 15 science highlights stories completed incl tech review	6/15/11	6/15/11	0%
1.1.4.1.5	Graphic designer selected	6/15/11	6/15/11	0%
1.1.4.1.6	Copy editing and internal ER review complete	7/1/11	7/1/11	0%
1.1.4.1.7	Overall design and test story mock up complete and reviewed	8/1/11	8/1/11	0%
1.1.4.1.8	Final design complete	9/1/11	9/1/11	0%

WBS	Task Name	Start	Finish	% Complete
1.1.4.1.9	Completed book delivered to printer	9/7/11	9/7/11	0%
1.1.4.1.10	Milestone: Science Highlights published	10/14/11	10/14/11	0%
1.1.4.1.11	Ongoing: Repeat previous 10 tasks annually	4/2/12	3/25/16	0%
1.1.4.2	**Create XSEDE website and translate relevant TG website content**	**1/3/11**	**4/18/11**	**0%**
1.1.4.2.1	XSEDE WebsiXSEDE website committee established	1/3/11	1/3/11	0%
1.1.4.2.2	Website requirements document complete	1/17/11	1/17/11	0%
1.1.4.2.3	Content requirements document complete	1/17/11	1/17/11	0%
1.1.4.2.4	First rev of design reviewed by website committee	2/15/11	2/15/11	0%
1.1.4.2.5	Rev of website reviewed by XSEDE leadership	3/1/11	3/1/11	0%
1.1.4.2.6	Short form usability test completed	3/15/11	3/15/11	0%
1.1.4.2.7	Final version of website reviewed by website committee	3/22/11	3/22/11	0%
1.1.4.2.8	Content approved	4/1/11	4/1/11	0%
1.1.4.2.9	Final build complete	4/7/11	4/7/11	0%
1.1.4.2.10	Content ported and built	4/15/11	4/15/11	0%
1.1.4.2.11	Initial version of XSEDE website launched	4/18/11	4/18/11	0%
1.1.4.2.12	Milestone: XSEDE website completed	4/18/11	4/18/11	0%
1.1.4.3	**Generate Annual Conference Proceedings**	**2/1/11**	**4/19/16**	**0%**
1.1.4.3.1	Proceedings chair selected and incorporated into event planning	2/1/11	2/1/11	0%
1.1.4.3.2	Submission site open (This isn't an ER function)	2/1/11	2/1/11	0%
1.1.4.3.3	Approval from ACM or whatever publisher received	3/1/11	3/1/11	0%
1.1.4.3.4	Templates and copyright permission forms out to authors	4/15/11	4/15/11	0%
1.1.4.3.5	All papers and copyright permission forms received from authors and OKed by track chairs	6/15/11	6/15/11	0%
1.1.4.3.6	All papers sent to production house for reproduction to USB key	7/7/11	7/7/11	0%
1.1.4.3.7	Milestone: Annual Conference Proceedings Complete	8/1/11	8/1/11	0%
1.1.4.3.8	Ongoing: Repeat previous 7 tasks annually	2/1/12	4/19/16	0%
1.1.4.4	Milestone: Ongoing - Generate press releases & website content	4/4/11	3/25/16	0%
1.1.5	**Industry Relations**	**7/1/11**	**9/23/16**	**0%**
1.1.5.1	**Workforce Development**	**7/1/11**	**6/30/16**	**0%**
1.1.5.1.1	Increase industry partners' awareness of all XSEDE SP's training opportunites	7/1/11	6/30/16	0%
1.1.5.1.2	Elicit industry partners' input to enhacne training programs for workforce development	7/1/11	6/30/16	0%
1.1.5.2	**Small and Medium Enterprise Program**	**10/3/11**	**9/23/16**	**0%**
1.1.5.2.1	Co-develop marketing materials to attract Small and Medium Enterprises	10/3/11	10/28/11	0%
1.1.5.2.2	Market XSEDE SPs to SME companies	10/31/11	9/23/16	0%
1.1.5.3	**Software Development**	**10/3/11**	**6/5/15**	**0%**
1.1.5.3.1	Solicit industry to propose software co-development projects	10/3/11	6/8/12	0%
1.1.5.3.2	Select and execute the software development project	6/11/12	6/5/15	0%
1.1.6	**SD&I**	**9/1/11**	**1/6/12**	**100%**
1.1.6.1	PDR	9/1/11	9/1/11	100%
1.1.6.2	**Increment Planning**	**9/1/11**	**9/6/11**	**100%**
1.1.6.2.1	develop incrment plan	9/1/11	9/6/11	100%
1.1.6.3	**IRR**	**9/7/11**	**9/7/11**	**100%**
1.1.6.3.1	conduct IRR	9/7/11	9/7/11	100%

WBS	Task Name	Start	Finish	% Complete
1.1.6.3.2	IRR complete and passed	9/7/11	9/7/11	100%
1.1.6.4	**CI Detailed Design**	**9/8/11**	**9/8/11**	**100%**
1.1.6.4.1	**develop detail design**	**9/8/11**	**9/8/11**	**100%**
1.1.6.4.1.1	GFFS	9/8/11	9/8/11	100%
1.1.6.4.1.2	Execution Management	9/8/11	9/8/11	100%
1.1.6.4.1.3	XUAS Data	9/8/11	9/8/11	100%
1.1.6.5	**CDR**	**9/9/11**	**9/15/11**	**100%**
1.1.6.5.1	conduct CDR	9/9/11	9/15/11	100%
1.1.6.5.2	CDR complete and passed	9/15/11	9/15/11	100%
1.1.6.6	**CI Development**	**9/16/11**	**9/16/11**	**100%**
1.1.6.6.1	**develop CI**	**9/16/11**	**9/16/11**	**100%**
1.1.6.6.1.1	GFFS	9/16/11	9/16/11	100%
1.1.6.6.1.2	Execution Management	9/16/11	9/16/11	100%
1.1.6.6.1.3	XUAS Data	9/16/11	9/16/11	100%
1.1.6.7	**CI TRR**	**9/19/11**	**10/7/11**	**100%**
1.1.6.7.1	**conduct CI TRR**	**9/19/11**	**10/7/11**	**100%**
1.1.6.7.1.1	GFFS	9/19/11	10/7/11	100%
1.1.6.7.1.2	Execution Management	9/19/11	10/7/11	100%
1.1.6.7.1.3	XUAS Data	9/19/11	10/7/11	100%
1.1.6.7.2	CI TRR complete and passed	10/7/11	10/7/11	100%
1.1.6.8	**CI Tests**	**10/10/11**	**12/15/11**	**100%**
1.1.6.8.1	**conduct Ci tests**	**10/10/11**	**12/15/11**	**100%**
1.1.6.8.1.1	GFFS	10/10/11	12/15/11	100%
1.1.6.8.1.2	Execution Management	10/10/11	12/15/11	100%
1.1.6.8.1.3	XUAS Data	10/10/11	12/1/11	100%
1.1.6.8.2	CI Tests complete and passed	12/15/11	12/15/11	100%
1.1.6.9	**STRR**	**12/16/11**	**12/16/11**	**100%**
1.1.6.9.1	conduct STRR	12/16/11	12/16/11	100%
1.1.6.9.2	STRR complete and passed	12/16/11	12/16/11	100%
1.1.6.10	**System Integration Test**	**12/19/11**	**12/28/11**	**100%**
1.1.6.10.1	conduct system test	12/19/11	12/28/11	100%
1.1.6.10.2	System test complete and passed	12/28/11	12/28/11	100%
1.1.6.11	**ORR**	**12/29/11**	**12/29/11**	**100%**
1.1.6.11.1	conduct ORR	12/29/11	12/29/11	100%
1.1.6.11.2	ORR complete and passed	12/29/11	12/29/11	100%
1.1.6.12	**Increment De-brief**	**12/30/11**	**1/6/12**	**100%**
1.1.6.12.1	increment reflection workshop	12/30/11	12/30/11	100%
1.1.6.12.2	reflection and practice report	1/2/12	1/6/12	100%
1.1.6.12.3	de-brief complete	1/6/12	1/6/12	100%
1.2	**Operations**	**4/1/11**	**7/1/16**	**5%**
1.2.1	**Cybersecurity**	**5/12/11**	**6/30/16**	**1%**
1.2.1.1	Setup coordination of XSEDE incident response	5/12/11	7/1/11	100%

Page 8

WBS	Task Name	Start	Finish	% Complete
1.2.1.2	Setup and deploy XSEDE Certificate Authority	7/1/11	6/28/12	0%
1.2.1.3	Maintain Certificate Authority	6/29/12	6/30/16	0%
1.2.1.4	Develop and deploy Security Awareness program	7/1/11	6/28/12	0%
1.2.1.5	Maintain Security Awareness program	6/29/12	6/30/16	0%
1.2.1.6	Develop and deploy two factor authentication service	7/1/11	6/28/12	0%
1.2.1.7	Maintain two factor authentication service	6/29/12	6/30/16	0%
1.2.1.8	Integrate and deploy InCommon Federated Authentication service	7/1/11	6/28/12	0%
1.2.1.9	Maintain InCommon Federated Authentication service	6/29/12	6/30/16	0%
1.2.1.10	Ongoing: Develop and implement improvements based on SEMP Spiral Design Process	7/1/11	6/30/16	0%
1.2.2	**Data Services**	**7/1/11**	**7/1/16**	**0%**
1.2.2.1	**XSEDE Wide File System**	**7/4/11**	**7/1/16**	**0%**
1.2.2.1.1	Evaluate/Select global parallel file system(s)	7/4/11	3/30/12	0%
1.2.2.1.2	Deploy/extend global parallel file system(s)	4/2/12	6/29/12	0%
1.2.2.1.3	Maintain/extend global parallel file system(s)	7/2/12	7/1/16	0%
1.2.2.2	Evaluate and design archival replication framework	7/1/11	6/28/12	0%
1.2.3	**XSEDENet Networking**	**4/1/11**	**6/30/16**	**6%**
1.2.3.1	**Transition current TG network to XSEDE network**	**4/1/11**	**7/29/11**	**100%**
1.2.3.1.1	Execute NLR contracts	4/1/11	7/1/11	100%
1.2.3.1.2	Connect XSEDE sites to NLR	6/27/11	7/29/11	100%
1.2.3.2	**xsede.org nameservice**	**5/9/11**	**7/29/11**	**100%**
1.2.3.2.1	delegate subdomains to XSEDE partners: NICS, TACC, PSC	5/9/11	7/29/11	100%
1.2.3.2.2	establish DNS entries for non-subdomain names (network.xsede.org, etc.)	5/9/11	7/29/11	100%
1.2.3.3	Maintain and Monitor XSEDEnet	7/29/11	6/30/16	0%
1.2.3.4	Preserve the PSC speed page	5/9/11	7/29/11	100%
1.2.3.5	Tune end-to-end performance for user applications	7/1/11	6/30/16	0%
1.2.3.6	Ongoing: Develop and implement improvements based on SEMP Spiral Design Process	7/1/11	6/30/16	0%
1.2.3.7	Setup perfSONAR	7/1/11	7/13/12	0%
1.2.3.8	Provide Operations tools to monitor the network	8/1/11	3/30/12	0%
1.2.3.9	Replace network.teragrid.org functions	8/1/11	3/30/12	0%
1.2.4	**Software Testing and Deployment**	**7/4/11**	**7/1/16**	**0%**
1.2.4.1	**Test XSEDE software**	**7/4/11**	**3/30/12**	**0%**
1.2.4.1.1	Test Grid Middleware - EMS (Unicore and Genesis II client)	7/4/11	3/30/12	0%
1.2.4.1.2	Test Grid Middleware - GFSS (Genesis II)	7/4/11	3/30/12	0%
1.2.4.1.3	Test Data Movement - Globus Online	7/4/11	3/30/12	0%
1.2.4.2	**Deploy XSEDE software**	**4/2/12**	**6/29/12**	**0%**
1.2.4.2.1	Deploy Grid Middleware - Unicore	4/2/12	6/29/12	0%
1.2.4.2.2	Deploy Grid Middleware - Genesis II (campus bridging)	4/2/12	6/29/12	0%
1.2.4.2.3	Deploy Grid Middleware - Globus Online	4/2/12	4/2/12	0%
1.2.4.3	Support testing of XSEDE configuration items from SW Eng process (ongoing)	7/1/11	6/30/16	0%

WBS	Task Name	Start	Finish	% Complete
1.2.5	**Accounting and Account Management**	**5/9/11**	**6/30/16**	**3%**
1.2.5.1	Maintain accounting and account management databases	7/1/11	6/30/16	0%
1.2.5.2	Investigate and improve accounting/account management processes	7/1/11	6/30/16	0%
1.2.5.3	Provide reports	7/1/11	6/30/16	0%
1.2.5.4	Optimize and upgrade the XDCDB system	7/1/11	3/29/12	0%
1.2.5.5	Ongoing: Develop and implement improvements based on SEMP Spiral Design Process	7/1/11	6/30/16	0%
1.2.5.6	**Make appropriate name changes to {item}.xsede.org**	**5/9/11**	**7/29/11**	**100%**
1.2.5.6.1	domain name changes	5/9/11	7/29/11	100%
1.2.5.6.2	web site changes	5/9/11	7/29/11	100%
1.2.5.6.3	documentation changes	5/9/11	7/29/11	100%
1.2.5.7	Complete the 2nd phase of vetted/unvetted account creation effort	5/9/11	3/30/12	0%
1.2.6	**System Operational Support**	**5/9/11**	**7/1/16**	**17%**
1.2.6.1	Setup XSEDE Operations Center	6/1/11	6/28/11	100%
1.2.6.2	Operate XSEDE Operations Center	7/1/11	6/30/16	0%
1.2.6.3	Deploy centralized XSEDE cyberinfrasture servers/services	5/30/11	8/31/12	0%
1.2.6.4	Maintain centralized XSEDE cyberinfrastructure servers/services	9/3/12	7/1/16	0%
1.2.6.5	**Upgrade XDCDB hardware and split database and AMIE parts**	**5/9/11**	**8/12/11**	**100%**
1.2.6.5.1	upgrade hardware at SDSC and operating system at PSC	5/9/11	8/12/11	100%
1.2.6.5.2	migrate AMIE to stand alone VM server (PSC and SDSC)	5/9/11	6/17/11	100%
1.2.6.6	Evaluate XDCDB hardware at PSC and determine if refresh needed	5/9/11	6/29/12	100%
1.2.6.7	Evaluate current TG services and classify them into XSEDE HA tiers	5/9/11	6/29/12	0%
1.2.6.8	Plan and schedule a semi-annual XDCDB failover test (SDSC to PSC)	5/9/11	6/29/12	100%
1.2.6.9	TeraGrid ticket system transition	7/1/11	8/11/11	100%
1.2.6.10	Evaluate/Deploy XSEDE Centralized monitoring software	11/1/11	6/29/12	0%
1.2.6.11	Participate in Ticket System evaluation in conjunction with User Support	7/1/11	2/2/12	0%
1.2.6.12	Transition Ticket System: Legacy to New	2/3/12	6/29/12	0%
1.2.6.13	**Backup XOC setup at IU**	**11/1/11**	**7/2/12**	**0%**
1.2.6.13.1	Documentation/Training for XOC Backup	11/1/11	7/2/12	0%
1.2.6.13.2	XOC Failover Test	11/1/11	7/2/12	0%
1.2.6.14	**Operations Annual Report**	**6/1/12**	**6/28/12**	**0%**
1.2.6.14.1	Prepare operational metrics annual report/internal Operational Assessment	6/1/12	6/28/12	0%
1.2.6.14.2	Milestone: Operational metrics annual report completed	6/28/12	6/28/12	0%
1.3	**User Services**	**3/1/11**	**7/1/16**	**12%**
1.3.1	**Training**	**3/1/11**	**6/30/16**	**12%**
1.3.1.1	Milestone: Develop guidelines for online and in-person training materials across XSEDE Sites	3/1/11	3/28/11	0%
1.3.1.2	**Milestone: Develop 10 Training Modules**	**4/1/11**	**3/30/12**	**38%**
1.3.1.2.1	Develop and Post 2 new online training tutorials	4/1/11	7/1/11	100%
1.3.1.2.2	Develop and Post 2 new online training tutorials	7/5/11	9/30/11	50%
1.3.1.2.3	Develop and Post 3 new online training tutorials	10/3/11	12/30/11	0%
1.3.1.2.4	Develop and Post 3 new online training tutorials	1/4/12	3/30/12	0%

WBS	Task Name	Start	Finish	% Complete
1.3.1.3	**Milestone: Conduct 50 Training sessions**	**4/1/11**	**3/30/12**	**69%**
1.3.1.3.1	Conduct first 10 in-person or webcast training sessions	4/1/11	7/1/11	100%
1.3.1.3.2	Conduct first 15 in-person or webcast training sessions	7/5/11	9/30/11	100%
1.3.1.3.3	Conduct first 10 in-person or webcast training sessions	10/3/11	12/30/11	75%
1.3.1.3.4	Conduct first 15 in-person or webcast training sessions	1/4/12	3/30/12	0%
1.3.1.4	Ongoing: Repeat 1.3.1.3 and 1.3.1.3 annually	4/2/12	6/30/16	0%
1.3.1.5	Milestone: Complete Federation of existing online training materials with XSEDE Repositories	4/4/11	3/30/12	40%
1.3.1.6	**Milestone: Complete 2 targeted community workshops**	**7/1/11**	**6/30/12**	**0%**
1.3.1.6.1	Conduct first targeted community workshop	7/1/11	12/31/11	0%
1.3.1.6.2	Conduct second targeted community wirkshop	1/2/12	6/30/12	0%
1.3.1.7	Ongoing: Repeat 1.3.1.6 annually	7/2/12	6/30/16	0%
1.3.2	**User Information Resources**	**4/1/11**	**6/30/16**	**36%**
1.3.2.1	**Milestone: Release Production User Portal & Web Site**	**4/1/11**	**7/1/11**	**100%**
1.3.2.1.1	Transition existing production portal capabilities and web site	4/1/11	7/1/11	100%
1.3.2.1.2	Define user information architecture	4/1/11	6/23/11	100%
1.3.2.2	**Milestone: Maintain Production User News System**	**4/1/11**	**7/1/11**	**100%**
1.3.2.2.1	Transition existing user news system to production	4/1/11	7/1/11	100%
1.3.2.3	**Milestone: Release Allocation & User Guide for New and Transitioning Users**	**4/1/11**	**10/11/11**	**100%**
1.3.2.3.1	Document instructions for new users coming to XSEDE	4/1/11	7/1/11	100%
1.3.2.3.2	Document instructions for existing users	7/4/11	9/5/11	100%
1.3.2.3.3	Document allocation policies for resources	9/6/11	10/11/11	100%
1.3.2.4	**Milestone: Release new user guide with user comment capabilities**	**4/1/11**	**7/1/11**	**100%**
1.3.2.4.1	Create user guide template for HPC, Viz, Storage, etc.	4/1/11	6/23/11	100%
1.3.2.4.2	Create user guide examples for each resource type	4/1/11	6/23/11	100%
1.3.2.4.3	Ensure all user guides have been transitioned to template	4/1/11	7/1/11	100%
1.3.2.4.4	Publish all user guides across web presence	4/1/11	7/1/11	100%
1.3.2.5	**Milestone: Production mobile user portal**	**4/1/11**	**7/1/11**	**59%**
1.3.2.5.1	Transition existing mobile framework	4/1/11	4/28/11	100%
1.3.2.5.2	Evaluate requirements for mobile features	4/29/11	6/23/11	50%
1.3.2.5.3	Create schedule for releasing future mobile features	4/29/11	7/1/11	50%
1.3.2.6	**Milestone: Release updated user news**	**7/1/11**	**9/22/11**	**100%**
1.3.2.6.1	Define requirements for user news system	7/1/11	7/28/11	100%
1.3.2.6.2	Evaluate existing and alternative technologies	7/29/11	8/11/11	100%
1.3.2.6.3	Release new user news system (if appropriate)	8/12/11	9/22/11	100%
1.3.2.7	**Milestone: Create new social media presence for XD**	**10/3/11**	**12/23/11**	**0%**
1.3.2.7.1	Define requirements for social media	10/3/11	10/28/11	0%
1.3.2.7.2	Evaluate requirements that come out of User Engagement	10/31/11	11/25/11	0%
1.3.2.7.3	Create twitter and facebook presence for users of XD	11/28/11	12/23/11	0%
1.3.2.7.4	Display twitter feeds on user portal	11/28/11	12/23/11	0%
1.3.2.8	**Milestone: Release collaborative capabilities to user portal**	**10/3/11**	**2/17/12**	**0%**
1.3.2.8.1	Define requirements based on User Engagement feedback	10/3/11	11/25/11	0%

WBS	Task Name	Start	Finish	% Complete
1.3.2.8.2	Enable users to be able to share calendars, chat, files, etc. with collaborators via user portal	11/28/11	2/17/12	0%
1.3.2.9	**Milestone: Release integrated training system**	**1/2/12**	**3/23/12**	**100%**
1.3.2.9.1	Define requirements for integrated training system based on TEOS and User Requiremenst	1/2/12	1/27/12	100%
1.3.2.9.2	Enable sites to post training courses on user portal	1/30/12	3/23/12	100%
1.3.2.9.3	Enable sites to add online training resources to user portal	1/30/12	3/23/12	100%
1.3.2.9.4	Enable users to register for training online via user portal	1/30/12	3/23/12	100%
1.3.2.9.5	Create one stop shop for user training on user portal with calendar, SMS notification, online registering	1/30/12	3/23/12	100%
1.3.2.10	Ongoing: Develop and implement improvements based on SEMP Spiral Design Process	4/4/11	6/30/16	0%
1.3.3	**User Engagement**	**4/1/11**	**7/1/16**	**4%**
1.3.3.1	**Annual User Surveys**	**1/2/12**	**6/30/16**	**0%**
1.3.3.1.1	Develop and Implement Y1 Annual User Survey	1/2/12	3/30/12	0%
1.3.3.1.2	Milestone: Y1 Annual User Survey report	6/30/12	6/30/12	0%
1.3.3.1.3	Develop and Implement Y2 Annual User Survey	1/1/13	4/1/13	0%
1.3.3.1.4	Milestone: Y2 Annual User Survey report	6/30/13	6/30/13	0%
1.3.3.1.5	Develop and Implement Y3 Annual User Survey	1/1/14	3/31/14	0%
1.3.3.1.6	Milestone: Y3 Annual User Survey report	6/30/14	6/30/14	0%
1.3.3.1.7	Develop and Implement Y4 Annual User Survey	1/1/15	3/31/15	0%
1.3.3.1.8	Milestone: Y4 Annual User Survey report	6/30/15	6/30/15	0%
1.3.3.1.9	Develop and Implement Y5 Annual User Survey	1/1/16	3/31/16	0%
1.3.3.1.10	Milestone: Y5 Annual User Survey report	6/30/16	6/30/16	0%
1.3.3.2	**XSEDE CRM**	**1/2/12**	**7/2/12**	**0%**
1.3.3.2.1	Design and Implemnt initial CRM system	1/2/12	6/29/12	0%
1.3.3.2.1	Milestone: Initial CRM system deployed	7/2/12	7/2/12	0%
1.3.3.3	**Data Mining**	**10/3/11**	**7/1/16**	**5%**
1.3.3.3.1	Q1Y1 Data mining operations	10/3/11	12/30/11	100%
1.3.3.3.2	Milestone: Q1Y1 Data mining report	12/31/11	12/31/11	100%
1.3.3.3.3	Q2Y1 Data mining operations	1/2/12	3/31/12	0%
1.3.3.3.4	Milestone: Q2Y1 Data mining report	4/2/12	4/2/12	0%
1.3.3.3.5	Q3Y1 Data mining operations	4/3/12	7/2/12	0%
1.3.3.3.6	Milestone: Q3Y1 Data mining report	7/2/12	7/2/12	0%
1.3.3.3.7	Q4Y1 Data mining operations	7/3/12	10/1/12	0%
1.3.3.3.8	Milestone: Q4Y1 Data mining report	10/2/12	10/2/12	0%
1.3.3.3.9	Q1Y2 Data mining operations	10/2/12	12/31/12	0%
1.3.3.3.10	Milestone: Q1Y2 Data mining report	1/1/13	1/1/13	0%
1.3.3.3.11	Q2Y2 Data mining operations	1/2/13	4/2/13	0%
1.3.3.3.12	Milestone: Q2Y2 Data mining report	4/2/13	4/2/13	0%
1.3.3.3.13	Q3Y2 Data mining operations	4/1/13	6/30/13	0%
1.3.3.3.14	Milestone: Q3Y2 Data mining report	7/1/13	7/1/13	0%
1.3.3.3.15	Q4Y2 Data mining operations	7/1/13	9/30/13	0%
1.3.3.3.16	Milestone: Q4Y2 Data mining report	10/1/13	10/1/13	0%
1.3.3.3.17	Q1Y3 Data mining operations	10/1/13	12/31/13	0%

WBS	Task Name	Start	Finish	% Complete
1.3.3.3.18	Milestone: Q1Y3 Data mining report	1/1/14	1/1/14	0%
1.3.3.3.19	Q2Y3 Data mining operations	1/1/14	3/31/14	0%
1.3.3.3.20	Milestone: Q2Y3 Data mining report	4/1/14	4/1/14	0%
1.3.3.3.21	Q3Y3 Data mining operations	4/1/14	6/30/14	0%
1.3.3.3.22	Milestone: Q3Y3 Data mining report	7/1/14	7/1/14	0%
1.3.3.3.23	Q4Y3 Data mining operations	7/1/14	9/30/14	0%
1.3.3.3.24	Milestone: Q4Y3 Data mining report	10/1/14	10/1/14	0%
1.3.3.3.25	Q1Y4 Data mining operations	10/1/14	12/31/14	0%
1.3.3.3.26	Milestone: Q1Y4 Data mining report	1/1/15	1/1/15	0%
1.3.3.3.27	Q2Y4 Data mining operations	1/1/15	3/31/15	0%
1.3.3.3.28	Milestone: Q2Y4 Data mining report	4/1/15	4/1/15	0%
1.3.3.3.29	Q3Y4 Data mining operations	4/1/15	6/30/15	0%
1.3.3.3.30	Milestone: Q3Y4 Data mining report	7/1/15	7/1/15	0%
1.3.3.3.31	Q4Y4 Data mining operations	7/1/15	9/30/15	0%
1.3.3.3.32	Milestone: Q4Y4 Data mining report	10/1/15	10/1/15	0%
1.3.3.3.33	Q1Y5 Data mining operations	10/1/15	12/31/15	0%
1.3.3.3.34	Milestone: Q1Y5 Data mining report	1/1/16	1/1/16	0%
1.3.3.3.35	Q2Y5 Data mining operations	1/1/16	3/31/16	0%
1.3.3.3.36	Milestone: Q2Y5 Data mining report	4/1/16	4/1/16	0%
1.3.3.3.37	Q3Y5 Data mining operations	4/1/16	6/30/16	0%
1.3.3.3.38	Milestone: Q3Y5 Data mining report	7/1/16	7/1/16	0%
1.3.3.4	**Develop Focus Group Topics and conduct focus groups**	**9/30/11**	**6/30/16**	**7%**
1.3.3.4.1	Milestone: Q1Y1 Focus Group Report	9/30/11	9/30/11	100%
1.3.3.4.2	Milestone: Q2Y1 Focus Group Report	12/31/11	12/31/11	100%
1.3.3.4.3	Milestone: Q3Y1 Focus Group Report	3/31/12	3/31/12	0%
1.3.3.4.4	Milestone: Q4Y1 Focus Group Report	6/30/12	6/30/12	0%
1.3.3.4.5	Milestone: Q1Y2 Focus Group Report	9/30/12	9/30/12	0%
1.3.3.4.6	Milestone: Q2Y2 Focus Group Report	12/31/12	12/31/12	0%
1.3.3.4.7	Milestone: Q3Y2 Focus Group Report	3/31/13	3/31/13	0%
1.3.3.4.8	Milestone: Q4Y2 Focus Group Report	6/30/13	6/30/13	0%
1.3.3.4.9	Milestone: Q1Y3 Focus Group Report	9/30/13	9/30/13	0%
1.3.3.4.10	Milestone: Q2Y3 Focus Group Report	12/31/13	12/31/13	0%
1.3.3.4.11	Milestone: Q3Y3 Focus Group Report	3/31/14	3/31/14	0%
1.3.3.4.12	Milestone: Q4Y3 Focus Group Report	6/30/14	6/30/14	0%
1.3.3.4.13	Milestone: Q1Y4 Focus Group Report	9/30/14	9/30/14	0%
1.3.3.4.14	Milestone: Q2Y4 Focus Group Report	12/31/14	12/31/14	0%
1.3.3.4.15	Milestone: Q3Y4 Focus Group Report	3/31/15	3/31/15	0%
1.3.3.4.16	Milestone: Q4Y4 Focus Group Report	6/30/15	6/30/15	0%
1.3.3.4.17	Milestone: Q1Y5 Focus Group Report	9/30/15	9/30/15	0%
1.3.3.4.18	Milestone: Q2Y5 Focus Group Report	12/31/15	12/31/15	0%
1.3.3.4.19	Milestone: Q3Y5 Focus Group Report	3/31/16	3/31/16	0%
1.3.3.4.20	Milestone: Q4Y5 Focus Group Report	6/30/16	6/30/16	0%

WBS	Task Name	Start	Finish	% Complete
1.3.3.5	**Develop and Conduct BoF Sessions**	**9/30/11**	**12/31/15**	**13%**
1.3.3.5.1	Milestone: Conduct Y1 XSEDE BoF and Report	9/30/11	9/30/11	100%
1.3.3.5.2	Milestone: Conduct Y1 SC BoF and Report	12/31/11	12/31/11	100%
1.3.3.5.3	Milestone: Conduct Y2 XSEDE BoF and Report	9/30/12	9/30/12	0%
1.3.3.5.4	Milestone: Conduct Y2 SC BoF and Report	12/31/12	12/31/12	0%
1.3.3.5.5	Milestone: Conduct Y3 XSEDE BoF and Report	9/30/13	9/30/13	0%
1.3.3.5.6	Milestone: Conduct Y3 SC BoF and Report	12/31/13	12/31/13	0%
1.3.3.5.7	Milestone: Conduct Y4 XSEDE BoF and Report	9/30/14	9/30/14	0%
1.3.3.5.8	Milestone: Conduct Y4 SC BoF and Report	12/31/14	12/31/14	0%
1.3.3.5.9	Milestone: Conduct Y5 XSEDE BoF and Report	9/30/15	9/30/15	0%
1.3.3.5.10	Milestone: Conduct Y5 SC BoF and Report	12/31/15	12/31/15	0%
1.3.3.6	**Conduct Usability Panels and Testing (as needed)**	**7/1/11**	**6/30/16**	0%
1.3.3.7	**User Engagement General Operations**	**7/1/11**	**6/30/16**	0%
1.3.3.8	**XSEDE Ticket System**	**4/1/11**	**6/30/12**	**34%**
1.3.3.8.1	Transition existing TG ticket system to production in XSEDE	4/1/11	7/7/11	100%
1.3.3.8.2	Milestone: Release Production Ticketing System	7/7/11	7/7/11	0%
1.3.3.8.3	**Deploy new/improved XSEDE Ticket System**	**7/1/11**	**6/30/12**	**20%**
1.3.3.8.3.1	Define requirements for ticket system	7/1/11	8/14/11	100%
1.3.3.8.3.2	Evaluate candidate ticket systems	8/15/11	9/30/11	100%
1.3.3.8.3.3	Milestone: Select new/improved ticket system	10/1/11	10/1/11	100%
1.3.3.8.3.4	Develop and deploy new/improved ticket system	10/3/11	3/31/12	0%
1.3.3.8.3.5	Milestone: Release new/improved ticket system	3/31/12	3/31/12	0%
1.3.3.8.3.6	Integrate new/improved ticket system with CRM	4/2/12	6/30/12	0%
1.3.3.8.3.7	Milestone: Release Integrated CRM interface	6/30/12	6/30/12	0%
1.3.3.8.3.8	Integrate new/improved ticket system with XSEDE User Portal	4/2/12	6/30/12	0%
1.3.3.8.3.9	Milestone: Release Integrated XUP interface	6/30/12	6/30/12	0%
1.3.3.9	**Consulting Policies**	**10/3/11**	**7/2/12**	**0%**
1.3.3.9.1	Create consulting policies, procedures, and support guide	10/3/11	6/29/12	0%
1.3.3.9.2	Milestone: Deploy consulting policies, procedures, and support guide	7/2/12	7/2/12	0%
1.3.4	**Allocations**	**4/1/11**	**6/1/16**	**3%**
1.3.4.1	Allocations policy in place	4/1/11	7/1/11	100%
1.3.4.2	**Host Year 1 Quarterly Allocations Meetings**	**9/1/11**	**6/1/12**	**50%**
1.3.4.2.1	Host Quarterly Allocations Meeting	9/1/11	9/1/11	100%
1.3.4.2.2	Host Quarterly Allocations Meeting	12/1/11	12/1/11	100%
1.3.4.2.3	Host Quarterly Allocations Meeting	3/1/12	3/1/12	0%
1.3.4.2.4	Host Quarterly Allocations Meeting	6/1/12	6/1/12	0%
1.3.4.3	Ongoing: repeat 1.3.4.2 annually	9/3/12	6/1/16	0%
1.3.4.4	**Conduct Year 1 How to Write a Successful Proposal Webcast**	**9/1/11**	**6/1/12**	**50%**
1.3.4.4.1	Conduct How to Write a Successful Proposal Webcast	9/1/11	9/1/11	100%
1.3.4.4.2	Conduct How to Write a Successful Proposal Webcast	12/1/11	12/1/11	100%
1.3.4.4.3	Conduct How to Write a Successful Proposal Webcast	3/1/12	3/1/12	0%

WBS	Task Name	Start	Finish	% Complete
1.3.4.4.4	Conduct How to Write a Successful Proposal Webcast	6/1/12	6/1/12	0%
1.3.4.5	Ongoing: repeat 1.3.4.4 annually	9/3/12	6/1/16	0%
1.4	**Extended Collaborative Support Service-Projects**	**7/1/11**	**7/22/16**	**43%**
1.4.1	Create/ test proj. mgmt. framework for ESCC work plans/reporting	7/1/11	10/28/11	100%
1.4.2	Add at least 1 external FTE to fill an identified skills gap	7/1/11	6/30/16	100%
1.4.3	**Extended Support for Research Team (ESRT)**	**7/1/11**	**7/22/16**	**20%**
1.4.3.1	**Establish ESRT group**	**7/1/11**	**7/21/11**	**100%**
1.4.3.1.1	Set up ESRT staff and management teams and communications	7/1/11	7/21/11	100%
1.4.3.2	**Milestone: Support 20 ESRT Projects Annually**	**7/1/11**	**7/22/16**	**19%**
1.4.3.2.1	Work w/TG AUS to transition ASTA Projs. To ESRT mgmt.	7/1/11	10/28/11	100%
1.4.3.2.2	Milestone: All TG ASTA proj. managed as XD ESRT Projs.	10/28/11	10/28/11	100%
1.4.3.2.3	**Work with XD CMS to Generate 20 new XD ESRT projs. Annually**	**7/1/11**	**7/1/16**	**20%**
1.4.3.2.3.1	Year 1 Generate 20 new XD ESRT projects	7/1/11	6/29/12	100%
1.4.3.2.3.2	Year 2 Generate 20 new XD ESRT projects	7/2/12	6/28/13	0%
1.4.3.2.3.3	Year 3 Generate 20 new XD ESRT projects	7/1/13	6/30/14	0%
1.4.3.2.3.4	Year 4 Generate 20 new XD ESRT projects	7/1/14	6/30/15	0%
1.4.3.2.3.5	Year 5 Generate 20 new XD ESRT projects	7/1/15	7/1/16	0%
1.4.3.2.4	**20 ESRT work plans documented and actively managed annually**	**10/31/11**	**7/22/16**	**13%**
1.4.3.2.4.1	Year 1 - 20 work plans documented	10/31/11	6/29/12	100%
1.4.3.2.4.2	milestone: Y1 prepare and complete Final Reports	7/2/12	7/27/12	0%
1.4.3.2.4.3	Year 2 - 20 work plans documented	7/2/12	6/28/13	0%
1.4.3.2.4.4	milestone: Y2 prepare and complete Final Reports	7/1/13	7/26/13	0%
1.4.3.2.4.5	Year 3 - 20 work plans documented	7/1/13	6/30/14	0%
1.4.3.2.4.6	milestone: Y3 prepare and complete Final Reports	7/1/14	7/25/14	0%
1.4.3.2.4.7	Year 4 - 20 work plans documented	7/1/14	6/30/15	0%
1.4.3.2.4.8	milestone: Y4 prepare and complete Final Reports	7/1/15	7/24/15	0%
1.4.3.2.4.9	Year 5 - 20 work plans documented	7/1/15	6/30/16	0%
1.4.3.2.4.10	milestone: Y5 prepare and complete Final Reports	7/1/16	7/22/16	0%
1.4.4	**Novel & Innovative Projects**	**7/1/11**	**6/30/16**	**30%**
1.4.4.1	**Establish NIP group**	**7/1/11**	**9/30/11**	**100%**
1.4.4.1.1	Set up NIP staff and management teams and communications	7/1/11	9/30/11	100%
1.4.4.2	**Milestone: Generate 20 new XSEDE+ ECS NIPs Annually**	**7/1/11**	**6/30/16**	**20%**
1.4.4.2.1	Year 1-work w/XD CMS,TEOS,TIS to generate 20 XD NIP projects	7/1/11	6/29/12	100%
1.4.4.2.2	Year 2-work w/XD CMS,TEOS,TIS to generate 20 XD NIP projects	7/2/12	6/28/13	0%
1.4.4.2.3	Year 3-work w/XD CMS,TEOS,TIS to generate 20 XD NIP projects	7/1/13	6/30/14	0%
1.4.4.2.4	Year 4-work w/XD CMS,TEOS,TIS to generate 20 XD NIP projects	7/1/14	6/30/15	0%
1.4.4.2.5	Year 5-work w/XD CMS,TEOS,TIS to generate 20 XD NIP projects	7/1/15	6/30/16	0%
1.4.4.3	Milestone: Create ECSS Project work plans	7/1/11	12/30/11	100%
1.5	**Extended Collaborative Support Service - Communities**	**7/1/11**	**7/22/16**	**23%**

WBS	Task Name	Start	Finish	% Complete
1.5.1	Create/ test proj. mgmt. framework for ESCC work plans/reporting	7/1/11	10/28/11	100%
1.5.2	Add at least 1 external FTE to fill an identified skills gap	7/1/11	6/30/16	25%
1.5.3	**Extended Support for Community Codes (ESCC)**	**7/1/11**	**7/22/16**	**20%**
1.5.3.1	**Establish ESCC group**	**7/1/11**	**7/21/11**	**100%**
1.5.3.1.1	Set up ESCC staff and management teams and communications	7/1/11	7/21/11	100%
1.5.3.2	Transition TG ASP proj. to ESCC management	7/1/11	10/28/11	100%
1.5.3.3	Milestone:Active TG ASP proj. managed as XD ESCC proj.	10/28/11	10/28/11	100%
1.5.3.4	**Milestone: Support 10 ESCC Projects Annually**	**7/1/11**	**7/22/16**	**17%**
1.5.3.4.1	Year 1-Work w/TG CMS,TEOS,TIS to generate 10 XD ESCC proj.	7/1/11	6/29/12	100%
1.5.3.4.2	Year 2-Work w/TG CMS,TEOS,TIS to generate 10 XD ESCC proj.	7/2/12	6/28/13	0%
1.5.3.4.3	Year 3-Work w/TG CMS,TEOS,TIS to generate 10 XD ESCC proj.	7/1/13	6/30/14	0%
1.5.3.4.4	Year 4-Work w/TG CMS,TEOS,TIS to generate 10 XD ESCC proj.	7/1/14	6/30/15	0%
1.5.3.4.5	Year 5-Work w/TG CMS,TEOS,TIS to generate 10 XD ESCC proj.	7/1/15	7/1/16	0%
1.5.3.4.6	**Milestone: Create 10 ESCC work plans Annually**	**10/31/11**	**7/22/16**	**13%**
1.5.3.4.6.1	Y1 - 10 work plans documented	10/31/11	6/29/12	100%
1.5.3.4.6.2	milestone: Y1 prepare and complete Final Reports	7/2/12	7/27/12	0%
1.5.3.4.6.3	Y2 - 10 work plans documented	7/2/12	6/28/13	0%
1.5.3.4.6.4	milestone: Y2 prepare and complete Final Reports	7/1/13	7/26/13	0%
1.5.3.4.6.5	Y3 - 10 work plans documented	7/1/13	6/30/14	0%
1.5.3.4.6.6	milestone: Y3 prepare and complete Final Reports	7/1/14	7/25/14	0%
1.5.3.4.6.7	Y4 - 10 work plans documented	7/1/14	6/30/15	0%
1.5.3.4.6.8	milestone: Y4 prepare and complete Final Reports	7/1/15	7/24/15	0%
1.5.3.4.6.9	Y5 - 10 work plans documented	7/1/15	6/30/16	0%
1.5.3.4.6.10	milestone: Y5 prepare and complete Final Reports	7/1/16	7/22/16	0%
1.5.4	**Extended Collaborative Support Service - Science Gateways**	**7/1/11**	**7/22/16**	**24%**
1.5.4.1	**Establish ESSGW group**	**7/1/11**	**7/29/11**	**100%**
1.5.4.1.1	Set up ESSGW staff and management teams and communications	7/1/11	7/29/11	100%
1.5.4.2	**Milestone: Support 10 Science Gateways Annually**	**7/1/11**	**7/22/16**	**13%**
1.5.4.2.1	Year 1-Work w/TG CMS,TEOS,TIS to generate 10 XD ESSGW proj.	7/1/11	6/29/12	100%
1.5.4.2.2	Year 2-Work w/TG CMS,TEOS,TIS to generate 10 XD ESSGW proj.	7/2/12	6/28/13	0%
1.5.4.2.3	Year 3-Work w/TG CMS,TEOS,TIS to generate 10 XD ESSGW proj.	7/1/13	6/30/14	0%
1.5.4.2.4	Year 4-Work w/TG CMS,TEOS,TIS to generate 10 XD ESSGW proj.	7/1/14	6/30/15	0%
1.5.4.2.5	Year 5-Work w/TG CMS,TEOS,TIS to generate 10 XD ESSGW proj.	7/1/15	7/1/16	0%
1.5.4.2.6	**Milestone: Create ESSGW work plans**	**10/31/11**	**7/22/16**	**7%**
1.5.4.2.6.1	**At least 10 ESSGW work plans documented&actively managed annually**	**10/31/11**	**7/22/16**	**7%**
1.5.4.2.6.1.1	Y1 - 10 work plans documented	10/31/11	6/29/12	50%
1.5.4.2.6.1.2	milestone: Y1 prepare and complete Final Reports	7/2/12	7/27/12	0%
1.5.4.2.6.1.3	Y2 - 10 work plans documented	7/2/12	6/28/13	0%
1.5.4.2.6.1.4	milestone: Y2 prepare and complete Final Reports	7/1/13	7/26/13	0%
1.5.4.2.6.1.5	Y3 - 10 work plans documented	7/1/13	6/30/14	0%
1.5.4.2.6.1.6	milestone: Y3 prepare and complete Final Reports	7/1/14	7/25/14	0%
1.5.4.2.6.1.7	Y4 - 10 work plans documented	7/1/14	6/30/15	0%

WBS	Task Name	Start	Finish	% Complete
1.5.4.2.6.1.8	milestone: Y4 prepare and complete Final Reports	7/1/15	7/24/15	0%
1.5.4.2.6.1.9	Y5 - 10 work plans documented	7/1/15	6/30/16	0%
1.5.4.2.6.1.10	milestone: Y5 prepare and complete Final Reports	7/1/16	7/22/16	0%
1.5.4.3	Gateway Outreach: Constantly reach out to new potential gateways independently and in collaboration with NIP	7/1/11	6/30/16	30%
1.5.4.4	XSEDE Requirements: Work with Gateway community in analyzing the XSEDE architecture requirements	7/1/11	6/30/16	30%
1.5.4.5	XSEDE Architecture Test Cases: Provide Test Cases to SD&I teams in nature of tests to evaluate scalaibility, eligibility and usage patterns of middleware services.	7/1/11	6/30/16	30%
1.5.5	**Extended Support for Training Education and Outreach (ESTEO)**	**7/1/11**	**6/30/16**	**24%**
1.5.5.1	**Establish ESTEO group**	**7/1/11**	**7/21/11**	**100%**
1.5.5.1.1	Set up ESTEO staff and management teams and communications	7/1/11	7/21/11	100%
1.5.5.2	**Milestone: Contribute content for TEO modules**	**7/1/11**	**10/28/11**	**100%**
1.5.5.2.1	Work w/TG AUS to transition ASEOT projs. to XD ESTEO mgmt.	7/1/11	10/28/11	100%
1.5.5.2.2	Milestone: All TG ASEOT projs. Managed as XD ESTEO projs.	10/28/11	10/28/11	100%
1.5.5.3	**Milestone: 50 ESTEO projects supported Annually**	**7/1/11**	**6/30/16**	**18%**
1.5.5.3.1	Year 1-work w/XD CMS, TEOS, TIS to generate 50 XD ESTEO projects	7/1/11	6/29/12	90%
1.5.5.3.2	Year 2- work w/XD CMS, TEOS, TIS to generate 50 XD ESTEO projects	7/1/12	6/28/13	0%
1.5.5.3.3	Year 3 - work w/XD CMS, TEOS, TIS to generate 50 XD ESTEO projects	7/1/13	6/30/14	0%
1.5.5.3.4	Year 4 - work w/XD CMS, TEOS, TIS to generate 50 XD ESTEO projects	7/1/14	6/30/15	0%
1.5.5.3.5	Year 5 - work w/XD CMS, TEOS, TIS to generate 50 XD ESTEO projects	7/1/15	6/30/16	0%
1.6	**Education and Outreach**	**4/4/11**	**6/30/16**	**62%**
1.6.1	**Education**	**7/1/11**	**6/30/12**	**88%**
1.6.1.1	Milestone: 2 HPC Graduate level summer schools annually	7/1/11	6/30/12	100%
1.6.1.2	Milestone: 5 summer workshops annually	7/1/11	6/30/12	100%
1.6.1.3	Milestone: Add certificate programs at specific universities	7/1/11	6/30/12	75%
1.6.1.4	Milestone: Provide online educational services	7/1/11	6/30/12	75%
1.6.2	**Outreach**	**6/1/11**	**6/30/16**	**46%**
1.6.2.1	**Underrepresented Engagement**	**7/1/11**	**6/30/12**	**67%**
1.6.2.1.1	Milestone: 10 campus visits (SURA)	7/1/11	6/30/12	75%
1.6.2.1.2	Milestone: Engage 40 underrepresented individuals (Rice)	7/1/11	6/30/12	75%
1.6.2.1.3	Milestone: Create Faculty Council with 20 minority faculty (Rice)	7/1/11	6/30/12	50%
1.1.1.4	**Speakers' Bureau**	**7/1/11**	**6/30/12**	**75%**
1.1.1.4.1	Milestone: 10 National/Regional presentations	7/1/11	6/30/12	75%
1.6.2.3	**Student Engagement**	**7/1/11**	**6/30/12**	**75%**
1.6.2.3.1	Milestone: Recruit 20 students for training/mentoring/internship	7/1/11	6/30/12	75%
1.6.2.4	**Campus Champions**	**7/1/11**	**6/30/12**	**75%**
1.6.2.4.1	Milestone: Increase membership in Campus Champions program	7/1/11	6/30/12	75%
1.6.2.5	**XSEDE Annual Conference**	**6/1/11**	**7/1/12**	**75%**
1.6.2.5.1	XSEDE 12	6/1/11	7/1/12	75%
1.6.2.6	Ongoing: Repeat previous 5 tasks annually	7/2/12	6/30/16	0%
1.6.3	**Evaluation**	**4/4/11**	**3/25/16**	**58%**
1.6.3.1	Ongoing: External Evaluator Quarterly Reports	4/4/11	3/25/16	75%
1.6.3.2	Ongoing: Semi-Annual consultation with TEOS Advisory Group	4/4/11	3/25/16	50%

WBS	Task Name	Start	Finish	% Complete
1.6.3.3	Ongoing: Annual collection and analysis of community needs and requirements	4/4/11	3/25/16	50%
1.6.4	**Infratructure**	**4/4/11**	**3/25/16**	**75%**
1.6.4.1	Ongoing: Curation of TEOS information on public web and XSEDE portal	4/4/11	3/25/16	75%
1.6.5	**Campus Bridging**	**4/4/11**	**3/25/16**	**75%**
1.6.5.1	Ongoing: Share information with campuses interested in campus bridging	4/4/11	3/25/16	75%

C XSEDE Risk Register

Home > Risks Register > XSEDE Risks Summary

Risks Summary

Risk Summary Stats	
Total Number of Risks:	91
Number of non-retired Risks:	76
High Impact Risks:	34
High Probability Risks:	6
High Risk-Level:	2
Key Risks:	2
Current Risks:	70
(stats count non-retired risks)	

All Risks (Probability × Impact)

1	3	3
8	16	22
9	11	18

Non-retired Risks (Probability × Impact)

1	3	2
8	15	16
5	10	16

Retired Risks (Probability × Impact)

		1
	1	6
4	1	2

Non-retired Risks

76 risks

Risk Id	Risk	Risk Level	Probability	Impact	Status	Subproject	Monitor Date	Owner
273	Grid Software Scaling	High	High	High	Monitor	1.1.3 Architecture & Design	Apr 01, 2011	Andrew Grimshaw
272	Grid Software System Integration	High	High	High	Monitor	1.1.3 Architecture & Design	Apr 01, 2011	Andrew Grimshaw
328	Campus Bridging expectation management	High	High	Medium	Monitor	1.6.5 Campus Bridging	Oct 01, 2012	Craig Stewart
278	Implementation delays and inconsistencies	High	High	Medium	Monitor	1.2.4 Software Testing & Deployment	Jul 01, 2011	Troy Baer
280	Local root compromise	High	Medium	High	Monitor	1.2.1 Security	Jul 01, 2011	Randal Butler
287	Loss of Project PI or co-PI	High	Medium	High	Monitor	1 XSEDE	Jan 01, 2011	John Towns
282	Malicious Social Engineering	High	Medium	High	Monitor	1.2.1 Security	Jul 01, 2011	Randal Butler
258	Mismatch between research teams' needs, XD resources, and AUSS staff availability.	High	Medium	High	Inactive	1.4 Extended Collaborative Support - Projects	Apr 01, 2011	Sergiu Sanielevici
312	Network Disruption Disables Critical Services	High	Medium	High	Monitor	1.2.3 XSEDEnet	Jul 01, 2011	Linda Winkler
331	Program Plan	High	Medium	High	Monitor	1.1.1 Project Management and Reporting	Feb 16, 2012	Timothy Cockerill

279	Remote root compromise	High	Medium	High	Monitor	1.2.1 Security	Jul 01, 2011	Randal Butler
320	Security Compromise of XSEDE Infrastructure	High	Medium	High	Monitor	1.2.1 Security	Jul 01, 2011	Christopher Jordan
250	Suitable project management framework and process is not available.	High	Medium	High	Monitor	1.4.1 Extended Collaborative Research Teams Support	Jul 01, 2011	Mark Fahey
274	Usage of deployed software and services	High	High	Medium	Monitor	1.2.4 Software Testing & Deployment	Jul 01, 2011	Troy Baer
242	XRAC Meeting Costs	High	Medium	High	Monitor	1.3.4 Allocations	Jul 01, 2011	Ken Hackworth
Risk Id	Risk	Risk Level	Probability	Impact	Status	Subproject	Monitor Date	Owner
292	XROADS Opportunity: Integrating services receive other support	High	Medium	High	Inactive	1.1 Project Office	Jul 01, 2011	John Towns
301	XROADS Opportunity: Leverage new/novel resources	High	Medium	High	Inactive	1.1 Project Office	Jul 01, 2011	John Towns
296	XROADS: Capacity of XD reources less than demand	High	Medium	High	Inactive	1.1 Project Office	Jul 01, 2011	John Towns
304	XROADS: Dynamic network performance/costs may be bad	High	Medium	High	Inactive	1.2.3 XSEDEnet	Jul 01, 2011	Linda Winkler
291	XROADS: Insufficient capacity for high-quality help desk support	High	Medium	High	Inactive	1.3 User Services	Jul 01, 2011	Robert Brook
298	XSEDE archival storage gap	High	Medium	High	Monitor	1.2.2 Data Services	Jul 01, 2011	Christopher Jordan
252	Campus Infrastructure	Medium	Medium	Medium	Monitor	1.6.5 Campus Bridging	Jul 01, 2011	Craig Stewart
318	Communication Breakdown	Medium	Medium	Medium	Monitor	1.2.6 Systems Operational Support	Jul 01, 2011	Stephen McNally
325	Delay in program implementation	Medium	Medium	Medium	Monitor	1.6.1 Education	Sep 19, 2011	Steven Gordon
264	Failure of Security Systems and Procedures	Medium	Medium	Medium	Monitor	1.2.1 Security	Jul 01, 2011	Randal Butler
308	GPFS availability	Medium	Medium	Medium	Monitor	1.2.2 Data Services	Jun 01, 2011	Christopher Jordan
271	Insufficient Service Provider Integration into XSEDE activities	Medium	Medium	Medium	Monitor	1.1 Project Office	Jan 01, 2011	John Towns
251	Mentoring Program	Medium	Medium	Medium	Monitor	1.6.2 Outreach	Jul 01, 2011	Laura McGinnis
277	Noncompliant service provider	Medium	Medium	Medium	Monitor	1.1 Project Office	Apr 01, 2011	Timothy Cockerill
253	Remote User Support	Medium	Medium	Medium	Monitor	1.6.5 Campus Bridging	Jul 01, 2011	Craig Stewart
Risk Id	Risk	Risk Level	Probability	Impact	Status	Subproject	Monitor Date	Owner

322	Suitable project management framework and process is not available	Medium	Medium	Medium	Monitor	1.5.2 Extended Collaborative Science Gateways Support	Jul 01, 2011	Suresh Marru
255	Suitable project management framework and process is not available.	Medium	Medium	Medium	Monitor	1.5.1 Extended Support for Community Codes	Jul 01, 2011	John Cazes
257	Suitable project management framework and process is not available.	Medium	Medium	Medium	Monitor	1.5.3 Extended Collaborative EOT Support	Jul 01, 2011	Galen Arnold
246	Training - Sync delivery	Medium	Medium	Medium	Monitor	1.3.1 Training	Jul 01, 2011	Dan Stanzione
281	User Account Compromise	Medium	Medium	Medium	Monitor	1.2.1 Security	Jul 01, 2011	Randal Butler
305	XROADS: Federated identity management does not catch on	Medium	Medium	Medium	Inactive	1.1 Project Office	Jul 01, 2011	John Towns
266	Architecture Obsolescence & Software Risks	Medium	Low	High	Monitor	1.1.2 Systems and Software Engineering	Apr 01, 2011	Janet Brown
288	Differing architectural views hinder deployment and operation of XSEDE	Medium	Low	High	Monitor	1.1.2 Systems and Software Engineering	Mar 01, 2011	Janet Brown
260	Failure of XSEDE Operational Infrastructure	Medium	Low	High	Monitor	1.2.6 Systems Operational Support	Jul 01, 2011	Stephen McNally
295	Funds are not provided for initial network "gap" costs	Medium	Low	High	Monitor	1.2.3 XSEDEnet	Jul 01, 2011	Linda Winkler
276	Inadequate communication with SPs	Medium	Low	High	Monitor	1.1 Project Office	Jan 01, 2011	Timothy Cockerill
254	Mismatch between research teams' needs, XD resources, and ECSS staff availability.	Medium	Low	High	Monitor	1.5.1 Extended Support for Community Codes	Jul 01, 2011	John Cazes
256	Mismatch between research teams' needs, XD resources, and ECSS staff availability.	Medium	Low	High	Monitor	1.5.3 Extended Collaborative EOT Support	Jul 01, 2011	Galen Arnold
248	Mismatch maintained between research teams' needs, XD resources, and ECSS staff availability.	Medium	Low	High	Monitor	1.4.1 Extended Collaborative Research Teams Support	Jul 01, 2011	Mark Fahey
311	NLR Services Are Inadequate	Medium	Low	High	Monitor	1.2.3 XSEDEnet	Jul 01, 2011	Linda Winkler
Risk Id	Risk	Risk Level	Probability	Impact	Status	Subproject	Monitor Date	Owner
283	Network Security and Operation	Medium	Low	High	Monitor	1.2.1 Security	Jul 01, 2011	Randal Butler
316	Prohibitive Cost for Required Software	Medium	Low	High	Monitor	1.2.4 Software Testing & Deployment	Jul 01, 2011	Troy Baer

286	Reduced overall budget for the project	Medium	Low	High	Monitor	1 XSEDE	Sep 15, 2010	John Towns
259	Service Provider Non-Participation in SOC	Medium	Low	High	Monitor	1.2.1 Security	Jul 01, 2011	Randal Butler
265	Technical Obsolesence	Medium	Low	High	Monitor	1.1.3 Architecture & Design	Apr 01, 2011	Andrew Grimshaw
317	TeraGrid Domain Name Hardcoded into SP Software	Medium	High	Low	Monitor	1.2.6 Systems Operational Support	Jul 01, 2011	Stephen McNally
314	XSEDE DNS Service Availability	Medium	Low	High	Monitor	1.2.3 XSEDEnet	Jul 01, 2011	Linda Winkler
315	XSEDE RS Service Availability	Medium	Low	High	Monitor	1.2.3 XSEDEnet	Jul 01, 2011	Linda Winkler
275	Deviation from project management procedures	Low	Medium	Low	Monitor	1.1.1 Project Management and Reporting	Jan 01, 2011	Timothy Cockerill
309	Genesis II FUSE/RNS scalability and performance	Low	Low	Medium	Monitor	1.2.2 Data Services	Jul 31, 2011	Christopher Jordan
330	Lack of Testing Resources	Low	Medium	Low	Monitor	1.2.4 Software Testing & Deployment	Jan 10, 2012	Troy Baer
270	Loss of Senior Technical Personnel	Low	Low	Medium	Inactive	1 XSEDE	Apr 01, 2011	John Towns
321	Mismatch between non-traditional users' needs, XSEDE services, and NIP staff allocation.	Low	Medium	Low	Monitor	1.4.2 Novel & Innovative Projects	Jul 01, 2011	Sergiu Sanielevici
326	NLR Business Focus Changes	Low	Low	Medium	Monitor	1.2.3 XSEDEnet	Sep 01, 2011	Linda Winkler
327	Namespace collisions between XSEDE software and local software	Low	Medium	Low	Monitor	1.2.4 Software Testing & Deployment	Oct 03, 2011	Troy Baer
Risk Id	Risk	Risk Level	Probability	Impact	Status	Subproject	Monitor Date	Owner
319	Overloading XSEDE Operations Center Staff	Low	Medium	Low	Monitor	1.2.6 Systems Operational Support	Jul 01, 2011	Stephen McNally
239	Potential equipment failure can disrupt database services.	Low	Low	Medium	Monitor	1.2.5 Accounting/Account Mgmt	Jul 01, 2011	Steven Quinn
241	Prohibitive Operating Costs for Hardware or Software	Low	Low	Medium	Monitor	1.2.6 Systems Operational Support	Jul 01, 2011	Stephen McNally
268	Software Partner Failure	Low	Low	Medium	Monitor	1.1.3 Architecture & Design	Apr 01, 2011	Andrew Grimshaw
249	Student Internships	Low	Medium	Low	Monitor	1.6.2 Outreach	Jul 01, 2011	Laura McGinnis
245	Training - AUSS support	Low	Low	Medium	Monitor	1.3.1 Training	Jul 01, 2011	Dan Stanzione
261	UNICORE project is cancelled	Low	Low	Medium	Monitor	1.1.3 Architecture & Design	Jan 01, 2011	Andrew Grimshaw

244	XD Architecture not fully implemented at an XD Service Provider	Low	Low	Medium	Inactive	1.1 Project Office	Apr 01, 2011	John Towns
262	XD Service Provider has insufficient resources to implement XSEDE Architecture	Low	Low	Medium	Monitor	1.2.4 Software Testing & Deployment	Jul 01, 2011	Troy Baer
307	XROADS: Integrating services fail to meet standards in service level agreements	Low	Medium	Low	Inactive	1.1 Project Office	Jul 01, 2011	John Towns
306	XROADS: Plan for long term service scaling may not meet short term needs	Low	Medium	Low	Inactive	1.1 Project Office	Jul 01, 2011	John Towns
329	Delay in ticketing system interface release	Low	Low	Low	Monitor	1.3.2 User Information & Interfaces	Jan 01, 2012	Maytal Dahan
247	Education Workshops	Low	Low	Low	Monitor	1.6.1 Education	May 01, 2011	Steven Gordon
269	Lack of Qualified System Administration Personnel	Low	Low	Low	Monitor	1.1.2 Systems and Software Engineering	Mar 01, 2011	Janet Brown
263	Poor cooperation by non-XSEDE TG staff with transition on AUS projects	Low	Low	Low	Inactive	1.4 Extended Collaborative Support - Projects	Jul 31, 2010	Sergiu Sanielevici
Risk Id	Risk	Risk Level	Probability	Impact	Status	Subproject	Monitor Date	Owner
267	Software Complexity	Low	Low	Low	Monitor	1.2.4 Software Testing & Deployment	Jul 01, 2011	Troy Baer

Retired risks

Retired Risks

D XSEDE Change Control Report

Nothing to report this quarter.

E Metrics

To demonstrate its success and help focus management attention on areas in need of improvement, XSEDE monitors a wide range of metrics in support of different aspects of "success" for the program. The metrics presented in the quarterly reports provide a view into XSEDE's user community, including its success at expanding that community, the projects and allocations through which XSEDE manages access to resources, and the use of the resources by those projects (§E.1). In addition, XSEDE has identified metrics describing the program's success at delivering centralized services to this community, including operations, user support, advanced user support, and education and outreach activities (§E.2). Together, these metrics provide perspectives on how XSEDE works to ensure that the XSEDE-associated services and resources deliver science impact for the science and engineering research community.

E.1 XSEDE Resource and Service Usage Metrics

Table 6 highlights a few key XSEDE measures that summarize the user community, the projects and allocations, and resource utilization for the quarter. Expanded information and five-year historical trends are shown in three corresponding subsections.

Notably for Q1 2012, XSEDE saw a large jump in open user accounts and active individuals, with more than 1,063 individuals getting new XSEDE accounts during the quarter—surpassing by almost 200 users the quarterly record of 882 set in Q1 2011. XSEDE also showed increased use by EPSCoR state institutions and MSIs, as well as notable international engagement with Central and South America—the quarter's users represented three institutions from each of Argentina, Brazil, and Colombia, and two institutions from Chile and from Mexico. Gateways continued to represent a major part of the XSEDE community, with 1,039 users submitting jobs via science gateways, a decline due in part to the end of XSEDE service by NCSA's Ember. More details are in §E.1.1.

Table 6. Quarterly activity summary

User Community	**Q3 2011**	**Q4 2011**	**Q1 2012**
Open user accounts	6,056	5,829	6,313
Active individuals	1,965	1,819	2,165
Gateway users	1,158	1,389	1,039
New user accounts	471	545	1,063
Active fields of science	28	26	26
Active institutions	337	316	340
Projects and Allocations			
NUs available at XRAC	13.398B	16.780B	17.377B
NUs requested at XRAC	33.164B	28.012B	36.959B
NUs recommended by XRAC	21.928B	16.268B	20.562B
NUs awarded at XRAC	13.973B	16.268B	18.230B
Open projects	1,588	1,560	1,545
Active projects	971	917	979
Active gateways	21	16	18
New projects	187	188	216
Closed projects	253	254	219
Resources and Usage			
Resources open (all types)	32	23	23
Total peak petaflops	2.55	2.50	2.92
Resources reporting use	19	15	14
Jobs reported	1.62M	1.13M	1.04M
NUs delivered	14.84B	15.21B	14.76B
Avg wtd run time (hrs)	22.6	22.5	23.3
Avg wtd wait time (hrs)	27.7	26.8	23.2
Avg wtd slow down	3.1	3.5	3.3

Project and allocation activity remained strong, with XSEDE resources requested at 167% of what was available. The XRAC recommendations, however, fit within the resources available. During the quarter, open projects decreased to

1,545, but an increased number (979) made use of the resources. More details are in §E.1.2.

XSEDE computing resources represented 2.9 Pflops (peak) at the end of the quarter, due to the arrival of the Gordon system at SDSC. The central accounting system showed 14 resources reporting activity, and together they delivered 14.76 billion NUs of computing. This represents an decrease of approximately 3% over the previous quarter, but March showed an increase as usage from Gordon began to arrive. At the same time, XSEDE users experienced shorter wait times, on average, according to several metrics. More details are in §E.1.3.

E.1.1 User community metrics

Figure 26 shows the five-year trend in the XSEDE user community, including open user accounts, total active XSEDE users, active individual accounts, active gateway users, and the number of new accounts during the quarter. The numbers of individual users, both open and active accounts, rebounded significantly and set new record levels after two quarters of decline. The 1,063 new user accounts created in Q1 surpassed the previous record set in Q1 2011. Gateway users declined in part due to the end of XSEDE service by NCSA's Ember system. Figure 27 shows the activity on XSEDE resources according to field of science, including the relative fraction of PIs, open accounts, active users, allocations, and NUs used according to discipline. Disciplines with more than 2% of NUs used are listed in the figure. PIs and users are counted more than once if they are associated with projects in different fields of science. The Q1 data show that the percentages of PIs and accounts associated with the "other" disciplines represent nearly 30% of all PIs and user accounts, and more than 20% of active users. Collectively the "other" fields of science represent about 6% of total quarterly usage, holding steady from the previous quarter.

Table 7 and Table 8 highlight aspects of the broader impact of XSEDE. The former shows that graduate students, post-doctoral researchers, and undergraduates make up 65% of the XSEDE user base. The latter table shows XSEDE's reach into targeted institutional communities, including a 50% increase in usage at MSIs. Institutions with Campus Champions represent a large portion of XSEDE's usage (this table shows all users at Campus Champion institutions, not just those on the champion's project). The table also shows XSEDE's reach into EPSCoR states, the MSI community, and internationally. Also of note for Q1, the international institutions included strong showing from Central and South America with three institutions from each of Argentina, Brazil, and Colombia, and two institutions from Chile and from Mexico.

Table 7. End of quarter XSEDE open user accounts by type, excluding XSEDE staff.

Category	Q3 2011	Q4 2011	Q1 2012
Graduate Student	2,368	2,275	2,466
Faculty	1,336	1,299	1,293
Postdoctorate	1,008	1,002	1,059
Undergraduate Student	505	416	587
University Research Staff (excluding postdocs)	506	501	509
High school	5	5	5
Others	328	331	394

Table 8. Active institutions, overall and in selected categories. Notes: "Total" reflects institutions not in any specially designated category. Institutions may be in more than one category.

Category	Institutions	Users	NUs	% NUs
Campus Champion site	59	776	5,779,628,966	39%
EPSCoR state	64	287	2,383,944,729	16%
MSI	13	33	147,210,705	1%
International	52	83	503,638,223	3%

Total	340	2,165	14,759,096,345	100%

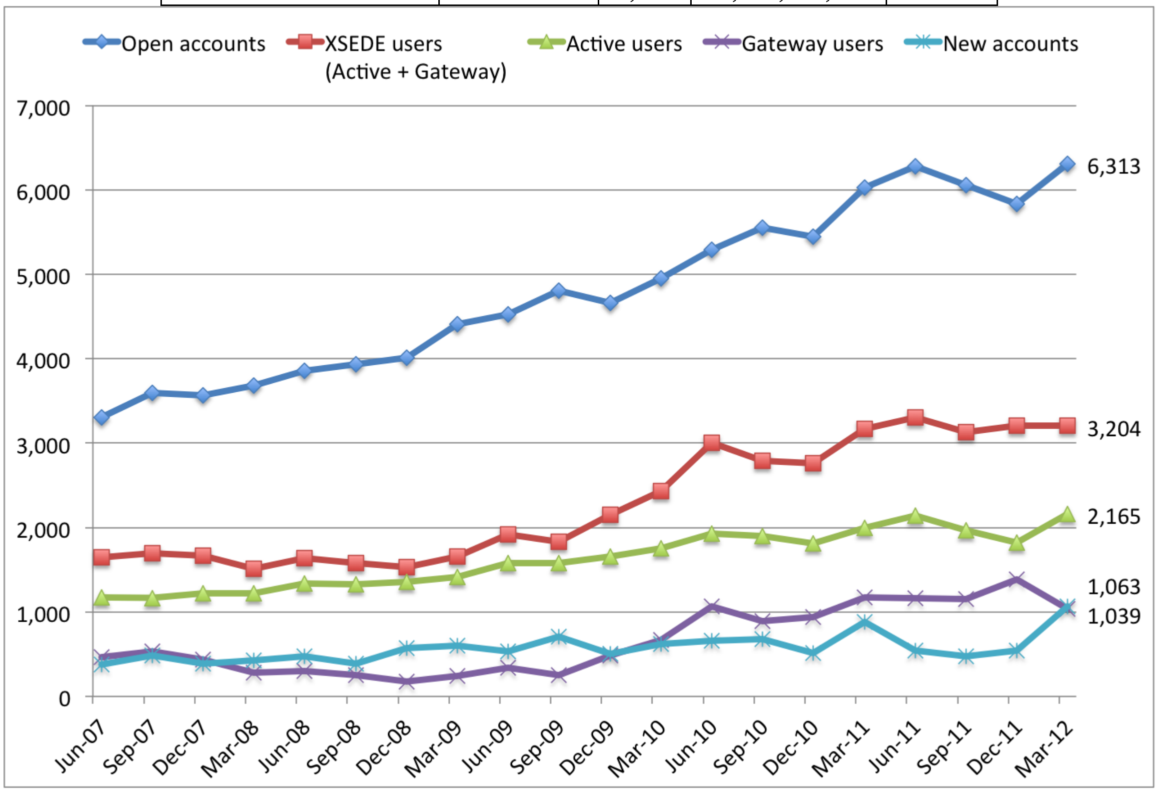

Figure 26. XSEDE user census, excluding XSEDE staff.

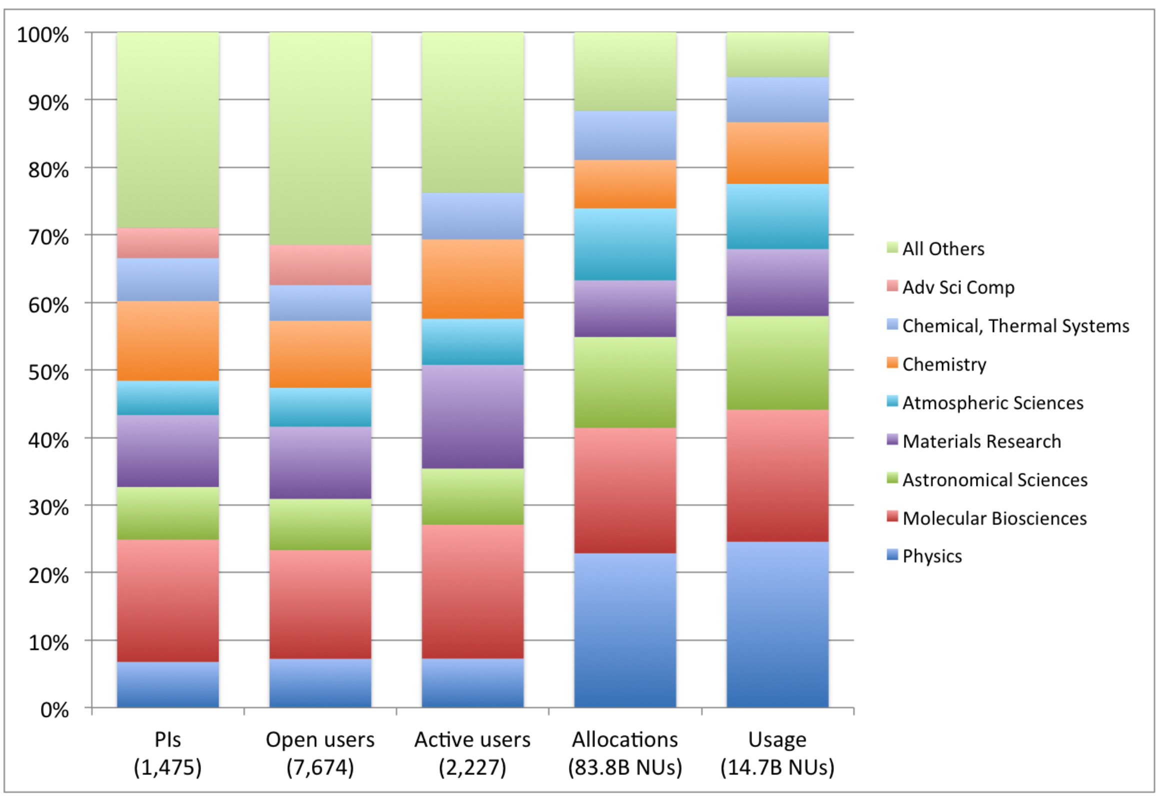

Figure 27. Quarterly XSEDE user, allocation, and usage summary by field of science, in order by usage, excluding staff projects. Note: PIs, users may appear under more than one field of science.

E.1.2 Project and allocation metrics

Figure 28 shows the historical trends for requests and awards at XSEDE quarterly allocation meetings. The figure shows the continued growth in demand against the relative plateau in available resources; NUs requested were 213% of NUs available, while the XRAC recommended awarding almost 18% more NUs than were available to the 123 requests for resources.

Table 9 presents a summary of overall project activity. Notably, 96% of XRAC requests received an award, and 36% were new awards.

Table 10 and Table 11 show projects and activity in key project categories—Campus Champion, Staff, and Science Gateways—and by allocation board type. (Science Gateways may appear under any board.) Table 10 also shows new and closed projects for the quarter, while Table 11 shows the number of open and active user accounts with each type of project. Table 12 shows detailed information about allocations activity for the various request types available for the different classes of projects. Notably, XSEDE had 123 Research (XRAC) requests, of which 118 (96%) received awards, including 43 new projects. There were also 168 Startup requests, of which 149 (87%) received awards; 12 Education requests with 11 awards; and 27 Campus Champion requests with 26 awards.

As a special class of projects, science gateway activity is detailed in Figure 29 showing continued high levels of usage and users from these projects.

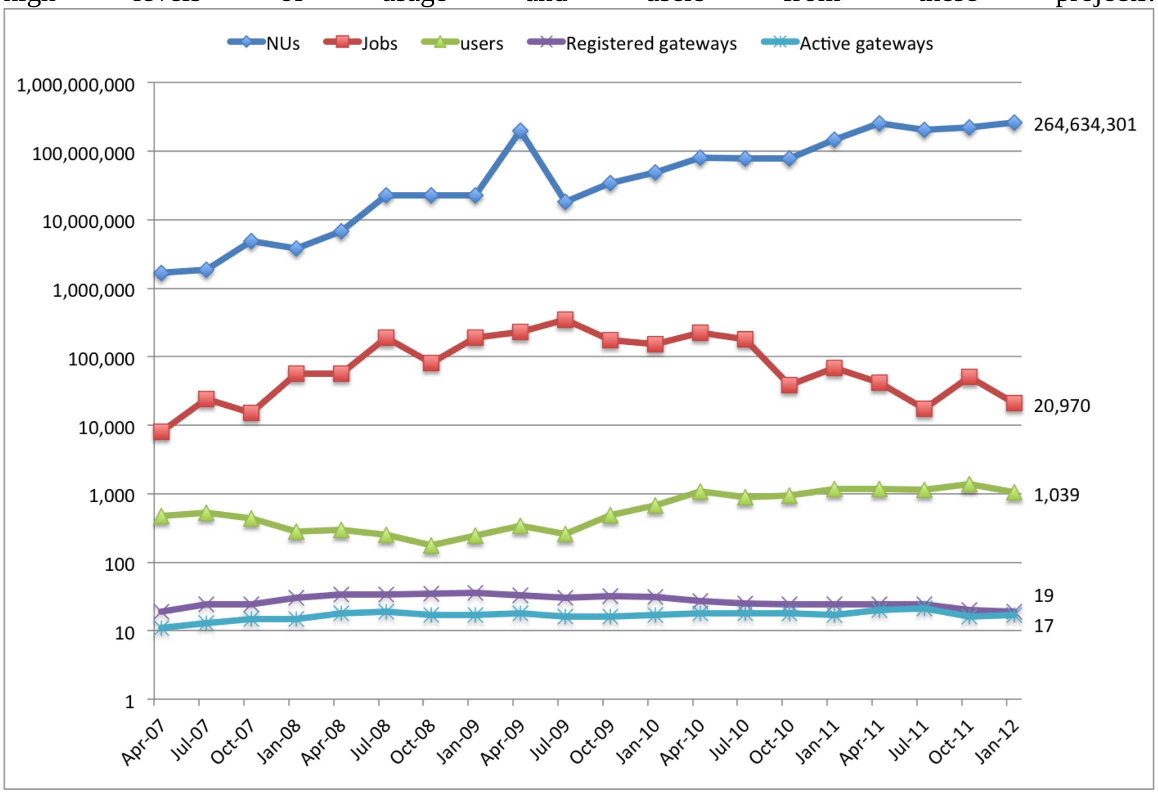

Figure 29. Quarterly gateway usage (NUs), jobs submitted, users (reported by AUSS), registered gateways, and active gateways.

Table 13 shows gateway activity supported by specific XSEDE resources.

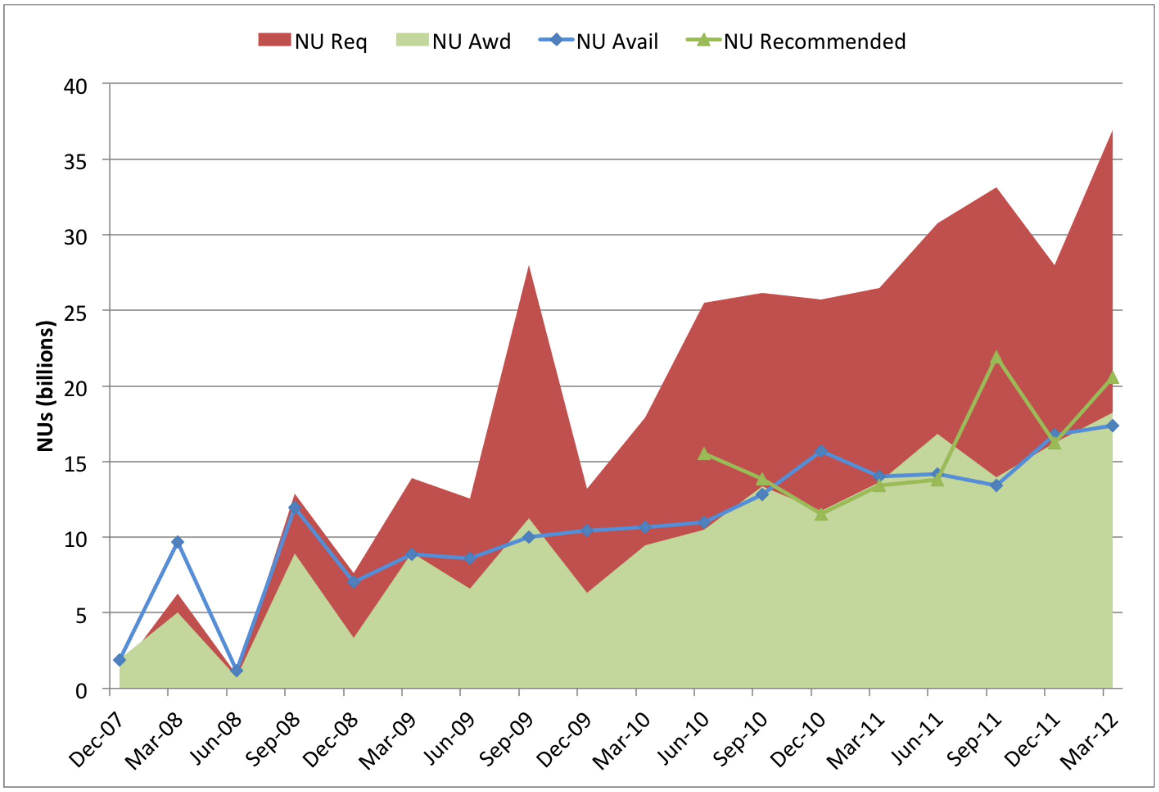

Figure 28. Allocation meeting history, showing NUs requested, awarded, available, and recommended. Note that December 2007 and June 2008 were the last two MRAC meetings, i.e. only had “medium” requests.

Table 9. Project summary metrics

Project metric	Q3 2011	Q4 2011	Q1 2012
XRAC requests	157	138	123
XRAC request success	91%	88%	96%
XRAC new awards	37%	29%	36%
Startup requests	137	129	
Startup request success	86%	89%	
Projects open	1,588	1,560	1,545
Projects new	187	188	216
Projects active	971	917	979
Projects closed	252	254	219
Resource diversity (wtd)	1.5 (2.3)	1.4 (2.1)	1.4 (1.9)
SP diversity (wtd)	1.4 (1.8)	1.3 (1.7)	1.3 (1.6)

Table 10. Project activity in designated categories.

Type	Open	New	Closed	Active	NUs	% NUs
Campus Champion	92	16	13	33	25,704,233	0.2%
Science Gateway	19			18	527,891,908	3.6%
TG Staff Project	25		1	12	39,156,689	0.3%
Others	1,409	200	205	916	14,205,500,207	96.0%
TOTALS	1,545	216	219	979	14,798,253,037	100.0%

Table 11 Project activity by allocation board type.

Board	Open projects	Open users	Active projects	Active users	NUs
XRAC/TRAC	594	4,301	513	1,477	14,291,718,328
Startup	799	2,005	390	495	375,683,021
Discretionary	2	25	3	7	44,147,827
Staff	27	497	12	84	39,156,689
Educational	62	1,109	38	329	29,211,429
Campus Champions	61	375	23	56	18,335,741
Totals	1,545	8,312	979	2,448	14,798,253,035

Table 12. Allocations activity in POPS, excluding staff and discretionary projects.

	Research				Startup			
	# Req	SU Req	# Awd	SUs Awd	# Req	SU Req	# Awd	SUs Awd
New	48	198,995,720	44	54,778,458	149	22,272,682	133	15,601,409
Prog. Report	4	16,313,212	4	14,402,325	n/a	n/a	n/a	n/a
Renewal	74	592,391,322	72	327,640,399	19	2,960,117	16	1,935,087
Advance	31	25,997,280	28	9,213,840	n/a	n/a	n/a	n/a
Justification	2	3,500,002	1	1,710,000	0	0	0	0
Supplemental	23	42,483,955	19	19,127,441	13	3,154,001	10	683,002
Transfer	97	32,189,200	83	21,948,391	40	2,944,273	37	1,527,278
Extension	83	n/a	67	n/a	42	n/a	37	n/a

	Education				Campus Champions			
	# Req	SU Req	# Awd	SUs Awd	# Req	SU Req	# Awd	SUs Awd
New	9	1,696,014	9	1,706,013	18	12,181,120	17	10,592,017
Prog. Report	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
Renewal	3	945,019	2	40,001	9	6,959,131	9	5,288,017
Advance	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
Justification	0	0	0	0	0	0	0	0
Supplemental	2	100,000	1	20,000	4	350,000	4	350,000
Transfer	3	69,284	2	56,000	1	2,600,000	1	190,000
Extension	1	n/a	1	n/a	0	n/a	0	n/a

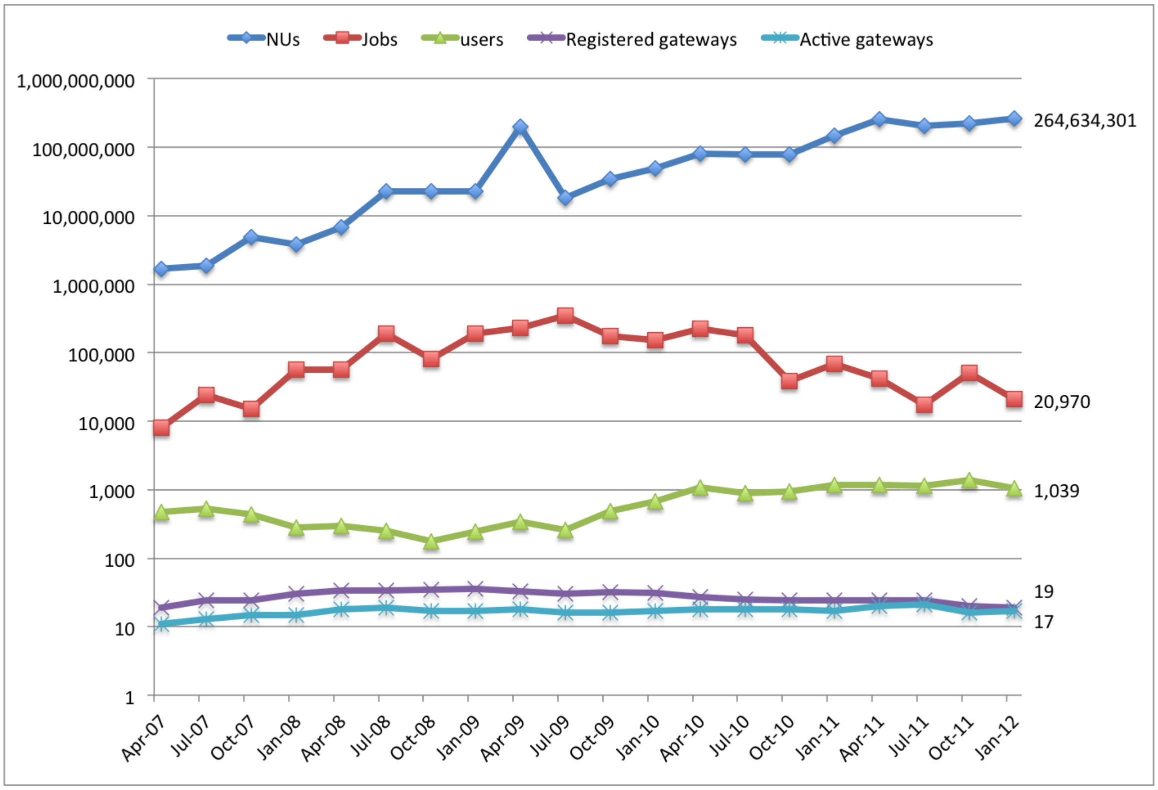

Figure 29. Quarterly gateway usage (NUs), jobs submitted, users (reported by AUSS), registered gateways, and active gateways.

Table 13. Gateway activity by resource.

Resource	Gateways	Jobs	Nus
SDSC Trestles	5	11,904	173,481,488
NICS Kraken	4	4,778	64,915,408
TACC Ranger	10	1,296	21,601,507
Purdue Steele	3	2,199	2,710,837
TACC Lonestar4	4	604	1,859,293
PSC Blacklight	2	78	64,194
NCSA Forge	1	111	1,575

E.1.3 Resource and usage metrics

Figure 30 shows the total NUs delivered by XSEDE computing systems, as reported to the central accounting system over the past five years. In Q1, the systems delivered 14.8 billion NUs, an decrease of about 3% from Q4 2011, but still represented 7% more NUs than the year ago quarter. The chart shows a rise as the first usage from SDSC's Gordon system began arriving in March. Table 14 breaks out the resource activity according to different resource types.

Figure 31 presents a perspective of the capacity and capability use of XSEDE resources by project. The figure shows the cumulative percentage of projects and resource usage according to each project's largest reported job size (in cores). The point at which the proverbial 80/20 rule holds precisely is at 76/24; that is the 76% of projects whose largest jobs were between 512 and 1,024 cores consumed only 24% of the delivered NUs, while the remaining 24% of projects consumed the remaining 76% of delivered NUs.

Finally, Table 15 presents some summary metrics to reflect aggregate "usage satisfaction," including the average run time, wait time, response time (run + wait), and slow down (or expansion factor). These values are presented as unweighted averages, which show the impact of small jobs, and as averages weighted by each job's portion of the workload (in core-hours), which show responsiveness to the jobs responsible for most of the delivered NUs. Notably for Q1, while the "average" job is only 5 hours long, the average weighted job is just over 23 hours long. On a number of usage satisfaction metrics, XSEDE showed modest improvements, including average weighted wait time, average weighted response time, and average weighted slow down. The weighted average for slow down (3.3) eliminates the skew in the job slow down attributed to small jobs and shows a much more realistic average perceived slowdown for the work delivered.

XSEDE provides central monitoring of GRAM5 and GRAM job submission activity at XSEDE SP sites (Figure 32). GRAM has been deprecated in favor of GRAM5.

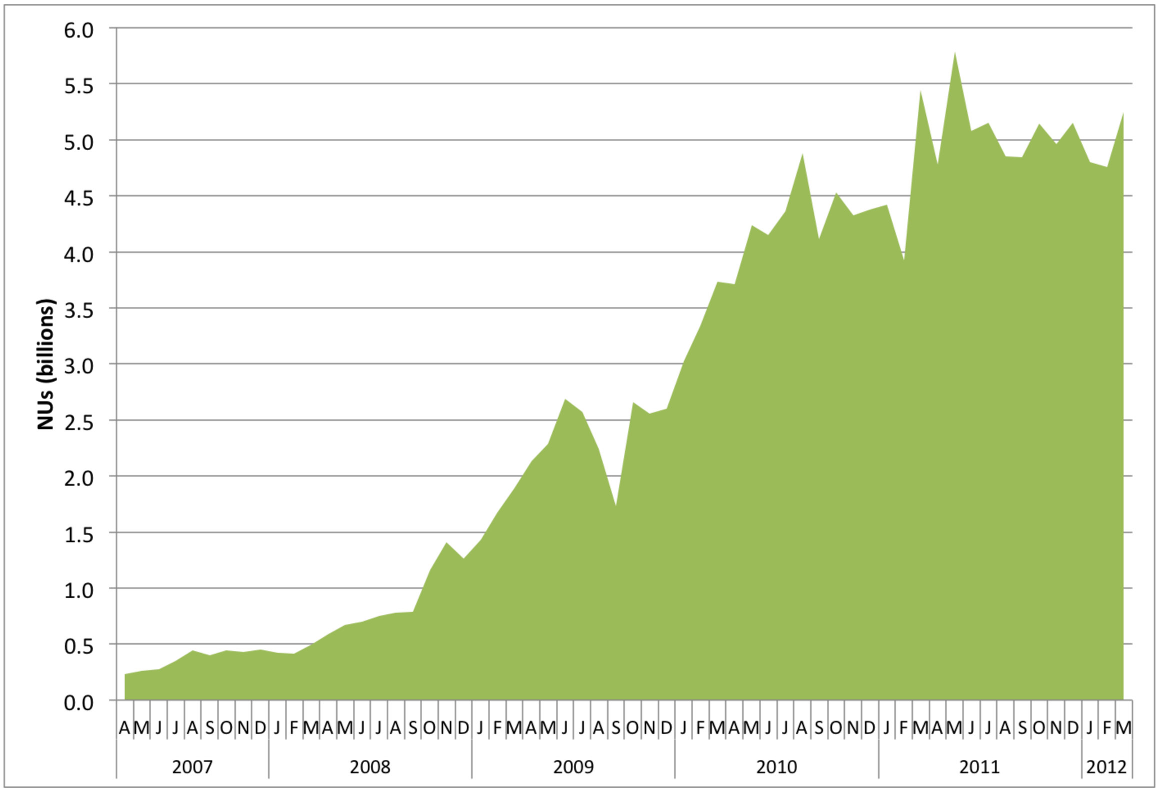

Figure 30. Total XSEDE resource usage in NUs.

Table 14. Resource activity, by type of resource, excluding staff projects.
Note: A user will be counted for each type of resource used.

Type	Resources	Jobs	Users	NUs
High-performance computing	7	894,006	1,773	14,300,137,069
Data-intensive computing	3	27,899	402	394,124,537
Visualization system	3	11,856	162	60,859,663
High-throughput computing	1	114,897	16	3,975,080
Total	14	1,048,658	2,353	14,759,096,349

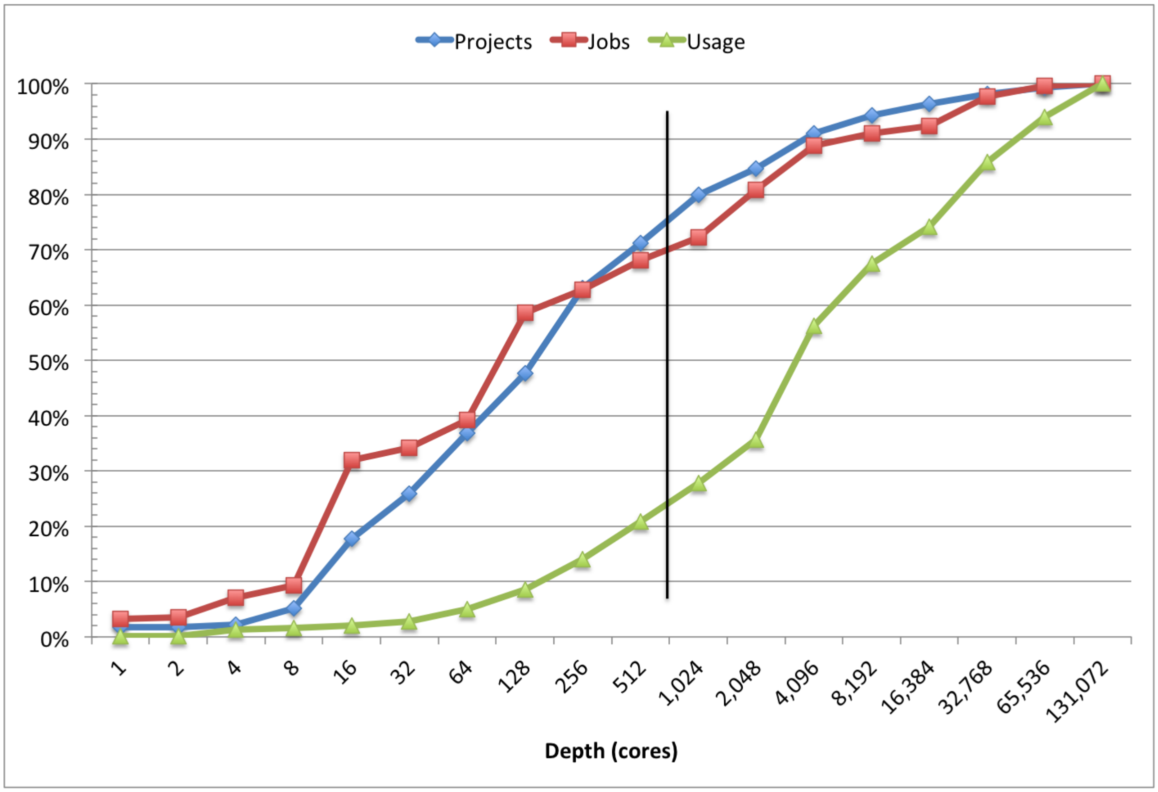

Figure 31. Cumulative distribution of projects, jobs, and usage according to project's maximum job size in cores (excludes staff projects). Vertical line (black) shows "joint ratio" of 76/24 at between 512 and 1,024 cores. I.e., 76% of projects use fewer than 1,024 cores and consume 24% of XSEDE NUs; the other 24% of projects have jobs larger than 1,024 cores and consume the other 76% of XSEDE NUs.

Table 15. Usage satisfaction metrics, for HPC and data-intensive computing resources only, excluding staff projects.

	Job attribute	Q3 2011 Average	Q4 2011 Average	Q1 2012 Average
Unweighted	Run time (hrs)	2.0	3.0	5.0
	Wait time (hrs)	4.3	3.1	6.1
	Response time (hrs)	7.2	7.0	12.1
	Slow down	699.4	252.3	334.4
Weighted	Wtd run time (hrs)	22.6	22.5	23.3
	Wtd wait time (hrs)	27.7	26.8	23.2
	Wtd response time (hrs)	50.3	49.4	46.5
	Wtd slow down	3.1	3.5	3.3

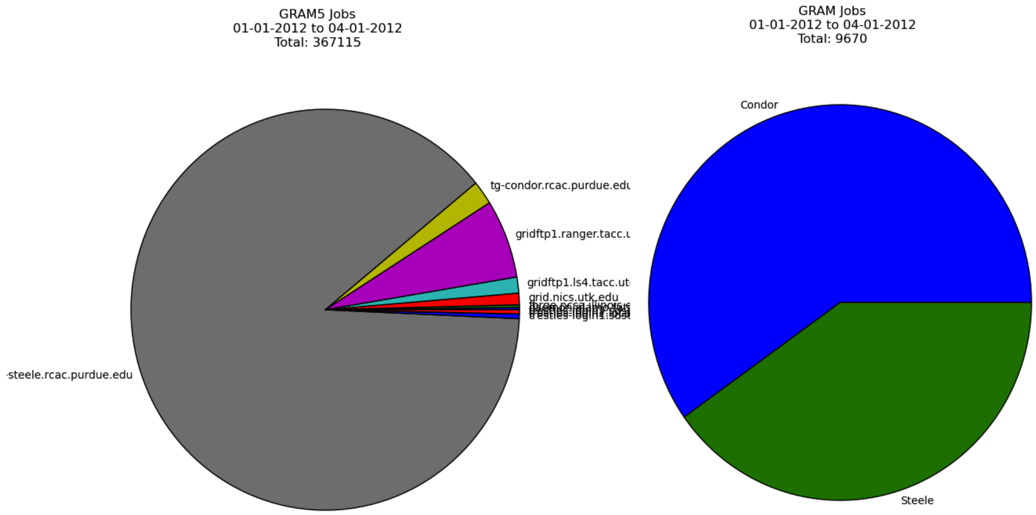

Figure 32. GRAM5 and GRAM jobs recorded per resource.

E.2 XSEDE Program Metrics

E.2.1 Project Office 1.1

1.1.4 External Relations

Table 16. XSEDE media hits.

Date	Source	Item	Notes
3/30/12	HPCWire	XSEDE Allocating Time to HPC Projects with Shared Memory	
3/29/12	GRIDcast	US-XSEDE ripple effect crosses the pond: UK-NGS Campus Champion Program takes flight!	
3/29/12	SupercomputingOnline	With Shared Memory in Pittsburgh, XSEDE Expands Horizons for Supercomputer Research	
3/27/12	NASDAQ GlobeWire	The Apache Software Foundation Announces Apache Rave as a Top-Level Project	Quote about XSEDE Science Gateways from Marlon Pierce, Indiana U.
3/22/12	Harvard blogspot	Free new-user training for XSEDE supercomputing network offered April 19	
3/22/12	HPCWire	Blue Waters and XSEDE Call for Presentations	
3/20/12	EurekAlert!	SDSC's 'big data' expertise aiding genomics research	
3/19/12	HASTAC	XSEDE Reaching out to Humanities, Arts and Social Science - eScience 2012 Workshop	
3/16/12	Georgia Tech PACE	New details about XSEDE12 Conference	
3/16/12	LiveScience	Movie Packs Virtual Birth of Universe into Mere Minutes	
3/16/12	Science magazine	Data Deluge Drives Demand	
3/15/12	LiveScience	Movie Packs Virtial Birth of Universe into Mere Minutes	
3/12/12	HPCWire	XSEDE Awards Supercomputing Time to UIUC I-CHASS Student and Professor	Also in: Primeur Live Mag.: http://primeurmagazine.com/weekly/AE-PR-04-12-87.html
3/6/12	ICHEC	European-US Summer School 2012 in Dublin, 24-28 June	
3/5/12	Newswise	SDSC's 'Gordon' Supercomputer: Ready for Researchers	Also on HPCWire: http://www.hpcwire.com/hpcwire/2012-03-

Date	Source	Item	Notes
			05/sdsc's_'gordon'_supercomputer_ready_for_researchers.html
3/5/12	YouTube	XSEDE facilitating scientific discovery	Phil Blood, PSC and XSEDE, delivers presentation at Univ. of Michigan
3/1/12	HiPACC	SDSC, HiPACC to Host Summer School on Astroinformatics	
3/1/12	HPCWire	Mobile Mayhem	
2/27/12	HPCWire	XSEDE's Extended Collaborative Support Conducts Symposium Series	
2/20/12	APPRO	Blog Video Series: SDSC Presentation on Gordon and Customer Interview	
2/15/12	Science Daily	Lava formations in Western U.S. linked to rip in giant slab of Earth	Study also published in Nature magazine; news also appeared in Real Clear Science, Science Newsline, YouTube, and numerous blogs
2/8/12	Inside HPC	Upgrading the hurricane forecast	
2/7/12	HPCWire	PRACE announces EU-U.S. summer school	
2/3/12	HPC in the Cloud	PRACE, XSEDE announce third EU-U.S. summer school on HPC challenges in computational sciences	Also in ISGTW and HPCWire
2/1/12	Inside HPC	Inaugural XSEDE12 conference issues call for participation	
1/31/12	HPCWire	PSC, Drexel light up cross-state optical link	
1/25/12	Inside HPC	Video: Gordon Supercomputer Wows TV Audience	FOX News video
1/24/12	Inside HPC	Nanomappers of the mind	
1/23/12	Wired	Intel Sees Exabucks in Supercomputing's Future	Maverick mentioned in caption and photo.
1/18/12	ISGTW	Drug delivery	
1/18/12	ISGTW	Space junk in 3D	(NOTE: Although this doesn't mention TG/XSEDE involvement, the visualization was based on a simulation run on a TG resource at NCSA.)
1/17/12	InsideHPC	Green revolution cooling	
1/17/12	HPCWire/HPC in the Cloud	Blue Waters, XSEDE speakers to headline GlobusWORLD 2012	ALSO IN: http://www.newswise.com/articles/blue-waters-xsede-speakers-to-headline-globusworld-2012
1/16/12	R&D magazine	Scientists predict an 'out-of-this-world' kind of ice	ALSO IN: http://www.physorg.com/news/2012-01-scientists-out-of-this-world-kind-ice.html
1/9/12	ECN magazine	High-performance computing tackles hepatitis	
1/6/12	Science Magazine	Ohm's Law survives to the atomic scale	
1/3/12	GIGAom	Why cloud could make crowdsourcing the norm for scientists	

E.2.2 Operations 1.2

1.2.1 Security

The XSEDE security team has identified the following metrics for tracking security incidents and response. Table 12 summarizes the metrics, and details on any incidents are provided in the main body of the report.

Table 17. XSEDE security metrics and incident response

	Q4 2011	Q1 2012
XSEDE-wide notice of vulnerability	0	0
Compromised user accounts	5	0
Other incident response	1	0
Critical rollout of vulnerability patches	0	0
Security enhancement rollouts	0	0

1.2.2 Data Services

XSEDE supports monitoring for two central data movement services: the gridFTP service connecting the XSEDE service providers and the Globus Online service for connecting XSEDE service providers as well as external sites.

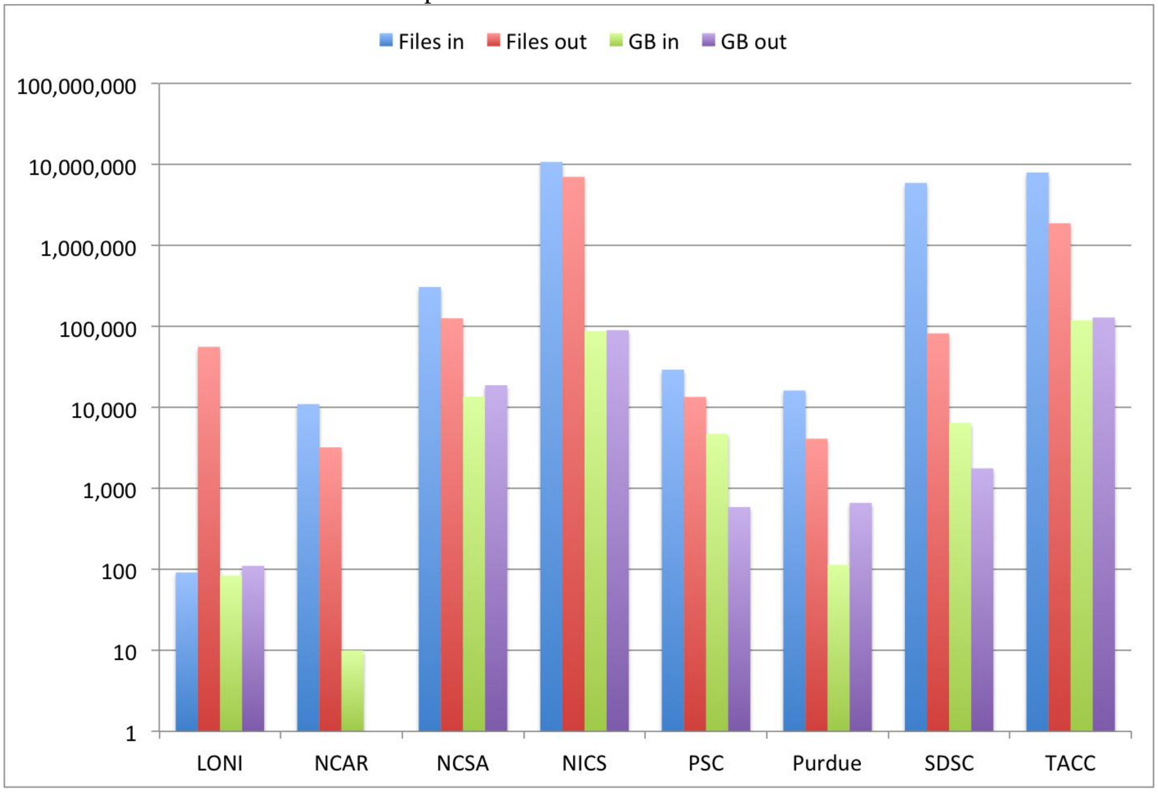

Figure 33. Globus Online activity into and out of XSEDE SP end points

Table 18 shows quarterly summary metrics for the past several quarters, while Figure 33 and Figure 34 show Globus Online and GridFTP activity, respectively, by SP site.

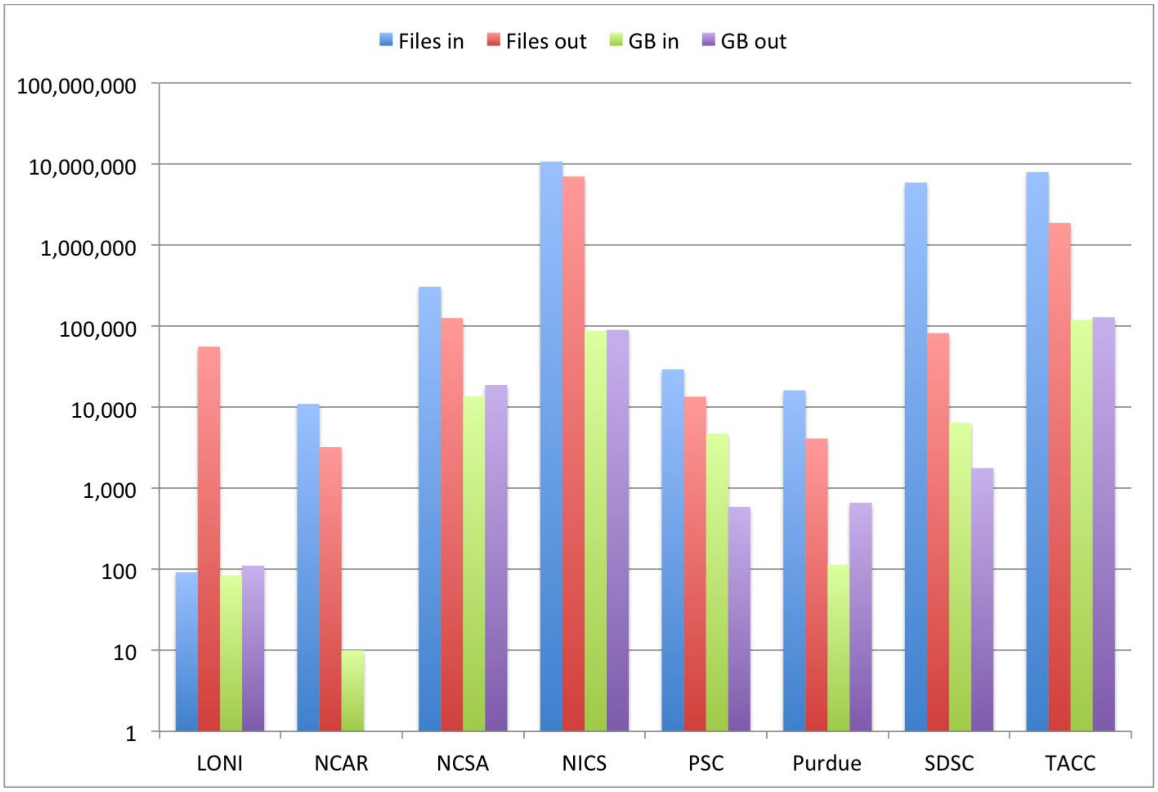

Figure 33. Globus Online activity into and out of XSEDE SP end points

Table 18. Globus Online activity to and from XSEDE endpoints, excluding GO XSEDE speed page user.

		Q3 2011	Q4 2011	Q1 2012
To/from XSEDE endpoint	Files to XSEDE	12.7 million	8 million	27.6 million
	TB to XSEDE	206	216	232
	Files from XSEDE	13.5 million	21.8 million	17.8 million
	TB from XSEDE	186	180	242
	Faults detected	732,000	573,000	584,000
	Users	116	124	149
To/from XSEDE via Globus Connect	Files to XSEDE	2.2 million	1.5 million	6.3 million
	TB to XSEDE	41	19	45
	Files from XSEDE	1.7 million	1.4 million	0.6 million
	TB from XSEDE	18	20	25
	Faults detected	422,000	102,000	242,000
	Users	71	74	87

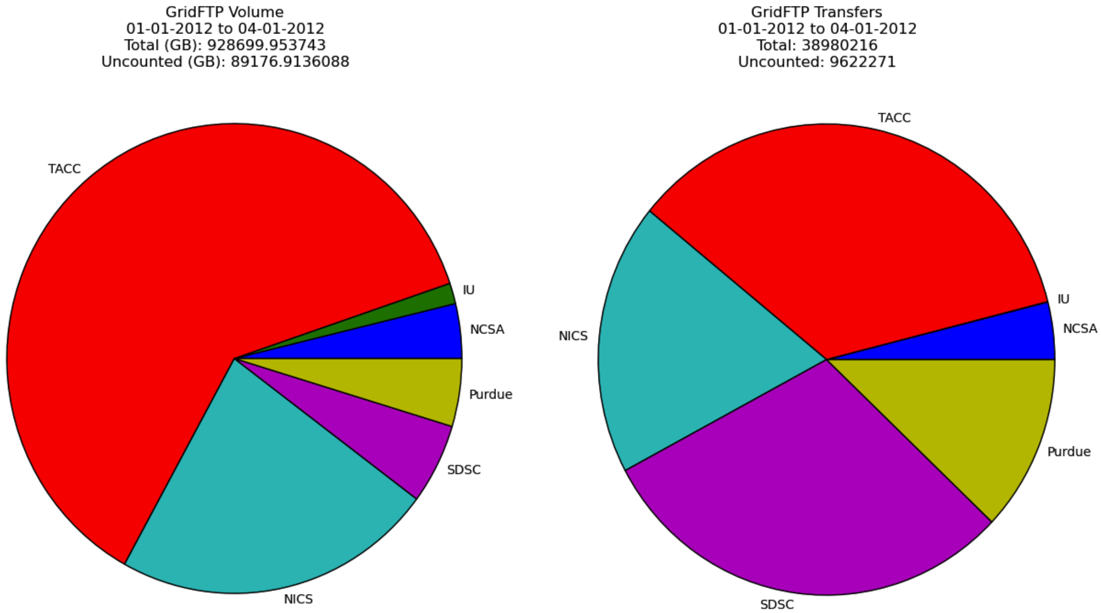

Figure 34. GridFTP volume and file transfers, per SP site.

1.2.3 XSEDEnet

Traffic utilization of the Chicago-Denver XSEDEnet link is shown in two figures below. Figure 35 shows the peak bandwidth across the link for the period. Figure 36 shows link utilization as a percentage. Traffic utilization across all XSEDEnet links is shown in Figure 37.

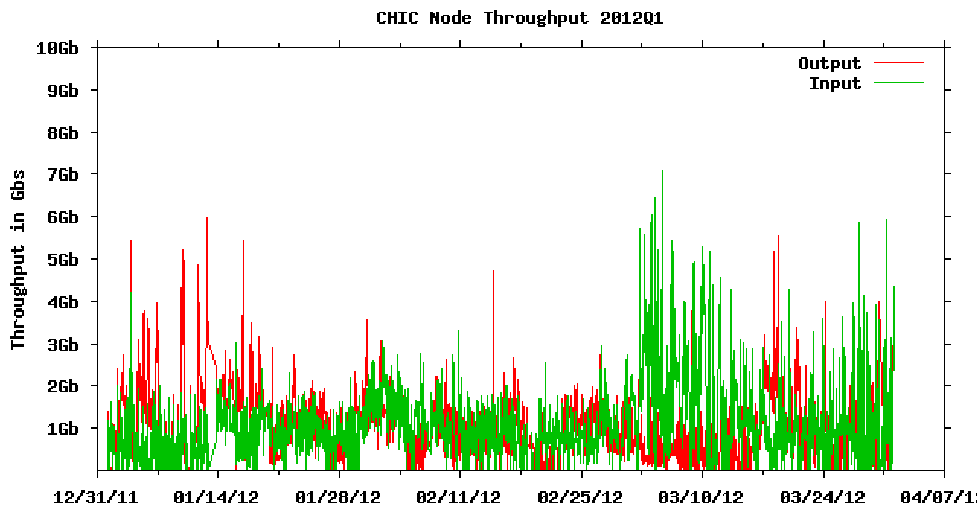

Figure 35. XSEDEnet Chicago-Denver peak bandwidth

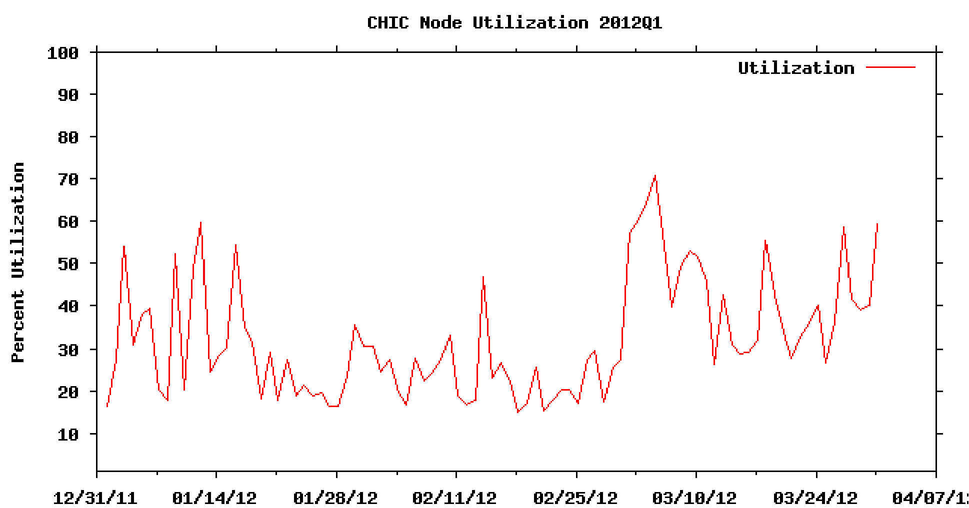

Figure 36. XSEDEnet Chicago-Denver utilization (as a percentage)

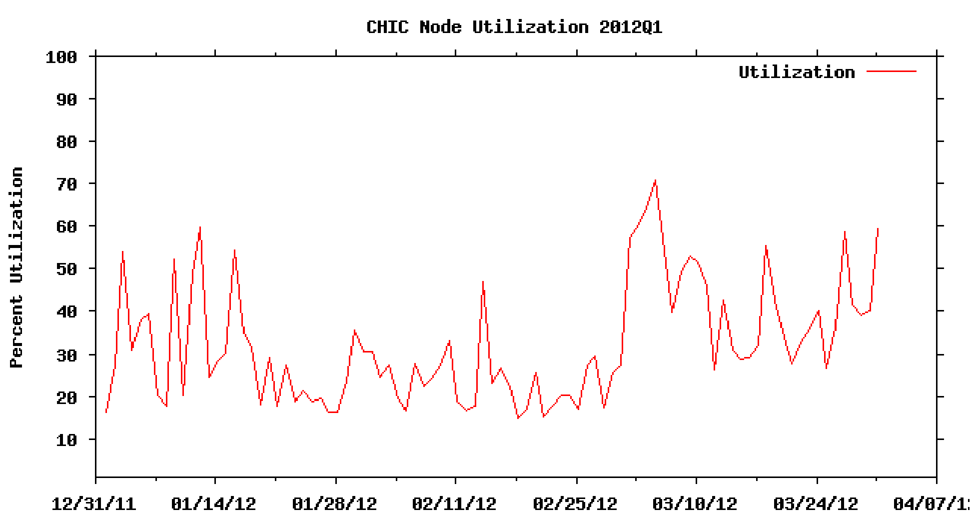

Figure 37. XSEDEnet utilization across all links

1.2.4 Software Testing and Deployment

XSEDE is evaluating what, if any, software testing and deployment metrics may be provided in the future.

1.2.5 Accounting and Account Management

The Accounting and Account Management group administers and operates the software for the XSEDE allocations system (POPS), the accounting system, and user account management. XSEDE reduced the account creation time by more than 50% by training a Help Desk team member to provide additional support in this area.

Table 19. Average time to process allocation requests and account creation requests, in days. (Quarterly XRAC requests omitted)

ALLOCATION REQUESTS		Q3 2011	Q4 2011	Q1 2012
Research	Advance	15	12	6
	Transfer	5	4	3
	Supplement	23	27	25
	Justification	37	54	15
Startup, Education, Campus Champions, Discretionary	New	13	13	9
	Renewal	11	11	10
	Transfer	5	4	3
	Supplement	10	11	4
Account creation requests		3.2	4.6	1.8

1.2.6 Systems Operational Support

The Systems Operational Support group encompasses the XSEDE Operations Center (XOC), which includes front-line user support and the ticket system, and the system administration of all XSEDE centralized services.

Table 20. XSEDE Operations Center ticket system metrics.

	Q3 2011	Q4 2011	Q1 2012
Total tickets opened	2,800	2,140	2,651
Tickets opened – email	1,788	1,318	2,381
Tickets opened – portal	812	684	254
Tickets opened – phone	200	138	16
Total tickets closed	2,965	2,111	2,335
Tickets, response in 2 bus. days	1,985 (71%)	1,550 (72%)	2,263 (85%)
Tickets closed within 2 bus. days	925 (33%)	739 (35%)	1,044 (39%)

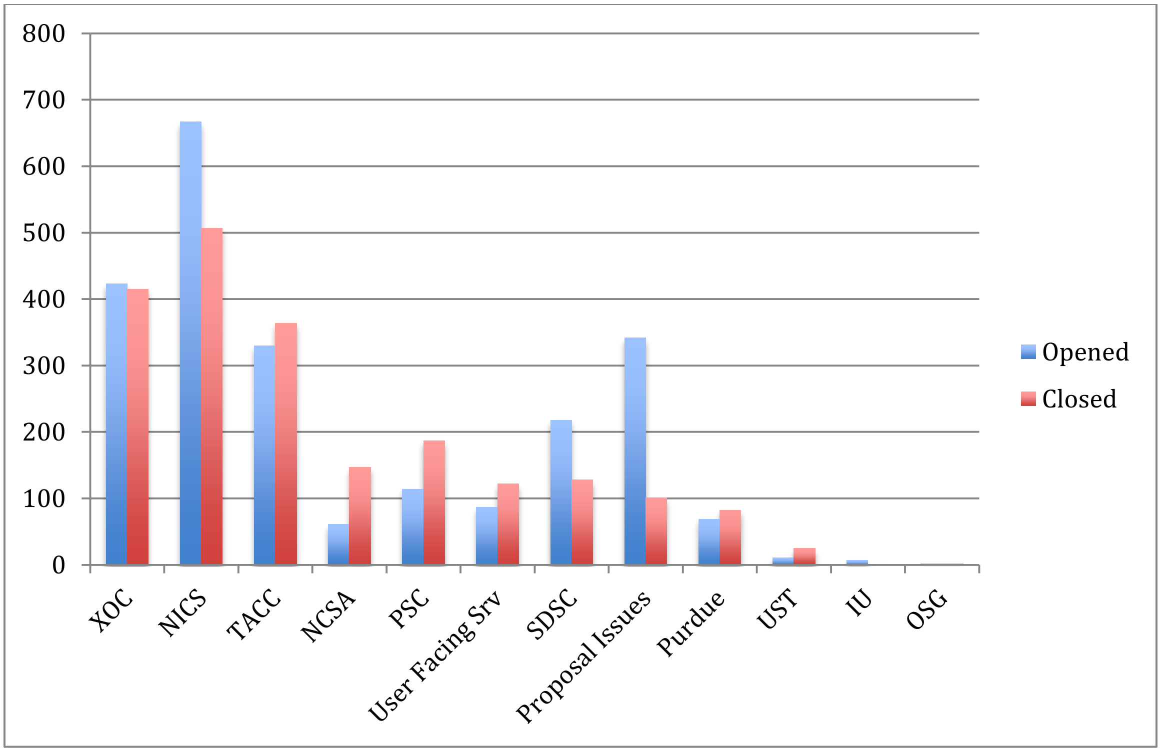

Figure 38. Ticket breakdown (opened/closed) for each major resolution center.

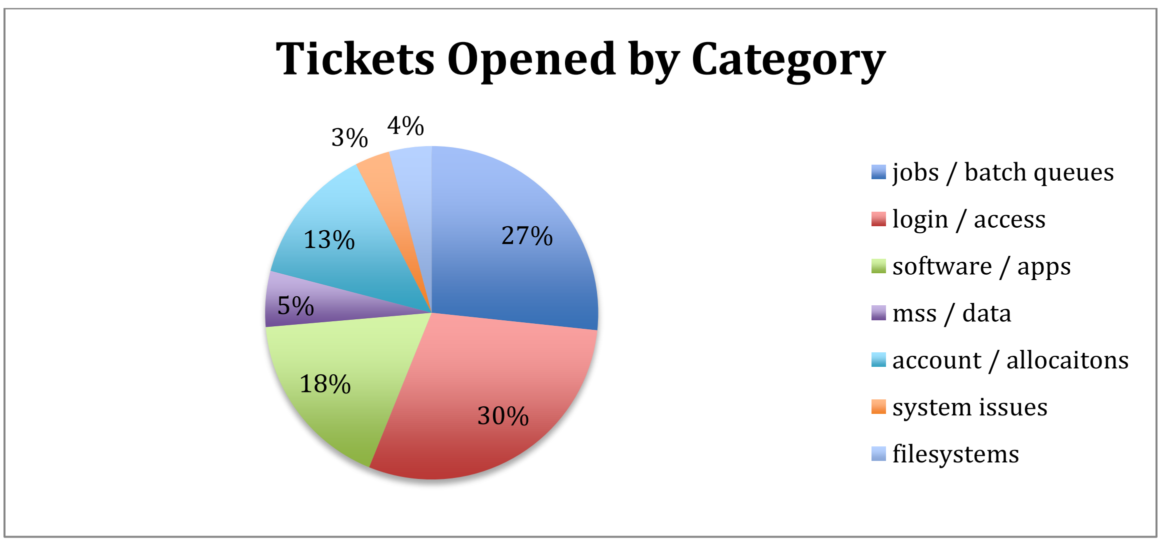

Figure 39. Tickets in seven primary problem categories. This pie chart represents a significant portion of tickets but does not represent the entire range. Tickets largely fit in the seven displayed categories; other categories are not significant enough to visually represent.

Table 21. XSEDE centralized service uptime and outages. Services not shown did not have any outages during the reporting period.

Service	Percentage of Uptime (Number of hours)	Nature of the Outage
RDR	**90%** (24 hours)	Unplanned
Information Services	**97.4%** (56 hours)	Unplanned
Karnak	**97.4%** (57.5 hours)	Unplanned
IIS Metrics	**97.4%** (56 hours)	Unplanned
XSEDE User Portal	**99.97%** (45 minutes)	Unplanned
TG Wiki	**97.4%** (56 hours)	Unplanned
Globus Listener	**97.4%** (56 hours)	Unplanned
Source Repository	**97.4%** (56 hours)	Unplanned
Build and Test	**97.4%** (56 hours)	Unplanned
Software Distribution	**97.4%** (56 hours)	Unplanned
Bugzilla	**97.4%** (56 hours)	Unplanned
User Profile Service	**97.4%** (56 hours)	Unplanned
Kerberos Backup	**99.99%** (5 mins)	Unplanned
XDCDB Backup	**99.99%** (30 mins)	Planned
AMIE Backup	**99.99%** (30 mins)	Planned
RDR	**99.77%** (5 hours)	Planned
Speedpage	**99.91%** (2 hours)	Planned
Portal/WWW backup	**99.6%** (8.5 hours)	Planned
Certificate Authority (not in production)	**99.6%** (8.5 hours)	Planned
Kerberos Backup	**99.91%** (2 hours)	Planned

Table 22. Inca-monitored XSEDE central services, Inca-detected uptimes, and outages.

Service	Definition of outage	Uptime (Details of outages)
Inca	Inca status pages are unavailable or not able to fetch data from the database (i.e., test details page fails to load). Tests every 5 mins.	**100%** (One partial outage for 16 hours – test results not stored)
Information Services	Information Web pages are unavailable. Tests every 15 mins.	**99.996%** (Two outages for 45 minutes of downtime)
Karnak	Karnak front page fails to load. Tests every 30 mins.	**99.93%** (Two outages for a total of 1.5 hours of downtime)
MyProxy	MyProxy server does not respond to credential query check. Tests every hour.	**100%** (No outages detected)
User Portal	Portal front home page fails to load correctly. Tests every 30 mins.	**99.97%** (One outage for a total of 30 minutes of downtime)
XDCDB	Connection to database refused or slow (using check_postgres.pl script). Tests every 5 mins.	**100%** (No outages detected)

E.2.3 User Services 1.3

The User Information and Interfaces group provides XSEDE users with central information and services via the XSEDE User Portal (XUP), web site, XUP mobile, and knowledgebase. Table 23 shows increasing activity on most user information interfaces, as well as increases in the numbers of logged-in users accessing these interfaces. Table 24 shows the most popular XUP applications, by visits in Q1 2012. After being integrated into XUP last quarter, User News received the most hits this quarter, but Allocations/Usage had the most users.

Table 23. XSEDE web site, user portal and XUP Mobile activity. (Note: “Users” indicates logged-in users.)

UII Activity	Q3 2011	Q4 2011	Q1 2012
Web hits	1,114,358	1,924,648	2,667,120
Web visitors	11,693	24,208	38,502
XUP hits	751,526	1,164,994	1,340,704
XUP visitors	6,290	10,346	14,913
XUP users	2,694	3,024	3,919
XUP Mobile hits	886	1,702	1,307
XUP Mobile users	41	48	37
KB docs retrieved	17,139	64,679	68,619
Total KB docs	348	405	478
New KB docs	24	25	73

Table 24. XUP and Web site application visits. “Users” indicates logged-in users.

	Q3 2011		Q4 2011		Q1 2012	
Application	**Visits**	**Users**	**Visits**	**Users**	**Visits**	**Users**
User News			4,312	25	62,784	526
Allocations/Usage	27,589	2,686	69,290	2,991	49,261	3,803
Resource Listing	64,688	652	100,544	864	36,244	1,104
File Manager	49,537 (transfers)	81 users (7.3 TB)	60,009 (transfers)	60 (3.9 TB)	31,952 (transfers)	79 (2.5 TB)
GSI-SSH	11,438	919	34,132	1,083	24,714	1,343
Training Registration			12,648	345	19,535	828
Science Gateways Listing	10,803	177	51,684	180	16,338	229
System Accounts	10,949	1,515	13,457	1,609	12,897	1,841
Knowledge Base	4,882	285	17,454	261	11,436	377
System Monitor	5,861	1,047	13,350	1,050	8,213	1,257
Software Search	5,249	150	14,812	176	8,204	265
Add User Form	3,483	533	8,406	611	5,751	626
Consulting Form	2,520	434	6,206	557	4,659	748
My Jobs	2,465	834	5,640	908	3,500	960
Ticketing System	1,503	559	3,708	609	2,366	682
Online Training Listing	1,111	175	2,720	215	2,093	347
SU Calculator (web)	1,001	141	4,376	180	2,048	243
User Profile	870	398			1,954	833
Karnak Queue Prediction	751	215	1,770	212	1,108	222
Feedback form			1,174	57	791	53
DN Listing	498	371	1,092	399	554	371

E.2.4 Extended Collaborative Support Service 1.4, 1.5

The Extended Collaborative Support Service pairs members of the XSEDE user community with expert staff members in projects lasting weeks to a year in order to solve challenging science and engineering problems through the application of cyberinfrastructure. Table 25 shows project and staffing metrics and, notably, a 30-project increase for Q1 2012. Table 26 shows activity metrics for Extended Support for Training, Education, and Outreach.

Table 25. Extended Collaborative Support project and staffing activity

		Q3 2011	Q4 2011	Q1 2012
Project requests	XRAC	22	16	22
	Supplemental/Startups	5	21	41
	ECSS In-house project			1
Projects initiated	Research Team	6	18	19
	Community Codes	2	9	8
	Science Gateways		4	6
	Unassigned			7
Projects cancelled/no-go	XRAC		1	4
	Supplemental/Startups		9	20
Projects active	Research Team (XRAC)		7	27
	Research Team (S/S)		4	17
	Research (TG)		29	12
	Subtotal	**37**	**40**	**56**
	Community Codes (XRAC)		3	6
	Community Codes (S/S)		3	9
	Community Codes (TG)		2	2
	In-house			1
	Subtotal	**2**	**8**	**18**
	Science Gateways (XRAC)			4
	Science Gateways (S/S)			7
	Science Gateways (TG)	2	10	4
	Subtotal	**2**	**10**	**15**
	Total Projects Active		**58**	**89**
Work plans			6	14
Projects completed		7	11	16
Novel, Innovative Projects (NIP)	User groups engaged	27	15	20
	NIP-led ECS planning efforts	8	8	9
	NIP ECS requests (prospective user groups)	5	14	23
	NIP ECS projects active	5	5	7
	NIP outreach events	2	9	9
ECSS staffing (FTE)	Research Team	13.62	13.72	13.28
	Community Codes	6.98	6.73	7.11
	Science Gateways	4.83	4.83	4.92
	NIP	4.7	5.30	5.79
	TEO	4.07	4.02	4.42
	Total	**34.20**	**34.60**	**35.52**

Table 26. Extended Support for Training, Education and Outreach 1.5.3

	Q3 2011		**Q4 2011**		**Q1 2012**	
Description	**#**	**Staff**	**#**	**Staff**	**#**	**Staff**
Requests for service	1*		1	1	4	4
User meetings and BOFs	5	15	7	9	2	2
Mentoring	2	4	3	3	3	2
Talks and presentations	6	8	7	8	12	9
Tutorials	33	58	3	3	21	25
Online tutorials and webinars	2	5	1	3	4	6
Online tutorial reviews	8	10	5	4	6	6

F XSEDE Publications Listing

F.1 XSEDE Staff Publications

F.1.1 Project Office 1.1

F.1.2 Operations 1.2

F.1.3 User Services 1.3

F.1.4 Extended Collaborative Support Service 1.4, 1.5

1. Arora R., Bangalore P., Mernik M. "A FRAMEWORK FOR SEMI-AUTOMATIC EXPLICIT PARALLELIZATION" *Authored book chapter*
2. Arora R., Bangalore P., Mernik M. "Formal and Practical Aspects of Domain-Specific Languages: Recent Developments" *Authored book chapter*
3. Jeong B., Navratil P., Gaither K., Abram G., Johnson G. "Configurable Data Prefetching Scheme for Interactive Visualization of Large-Scale Volume Data" *Proceedings of Visualization and Data Analysis (VDA) 2012*
4. Marker B., Terrel A., Poulson J., Batory D., van de Geijn R. "Mechanizing the Expert Dense Linear Algebra Developer" *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*
5. Navratil P. "General-Purpose GPU Programming: Hardware Considerations" *Guest lecture for Dr. Calvin Lin's (UT at Austin) parallel programming class (CS 380P)*
6. Navratil P. "Texas A&M GPU Workshop" *Invited talk at Texas A&M University*
7. Navratil P. "Visualization Clusters: from Tiled Displays to Remote Visualization" *Invited talk at Microsoft Research*
8. Navratil P., Fussell D., Lin C., Childs H. "Dynamic Scheduling for Large-Scale Distributed-Memory Ray Tracing" *Proceedings of Eurographics Symposium on Parallel Graphics and Visualization (EGPGV) 2012*
9. Navratil P., Johnson G., Westing B. "Visualization Clusters: from Tiled Displays to Remote Visualization" *IEEE Cluster 2011*
10. Terrel A., Mandli K. "ManyClaw: Slicing and dicing Riemann solvers for next generation highly parallel architectures" *TACC-Intel Highly Parallel Computing Symposium*

F.1.5 Education and Outreach 1.6

F.2 Publications from XSEDE Users

The following publications were gathered from Research submissions to the March 2012 XSEDE Resource Allocations Committee (XRAC) meeting. Renewal submissions are required to provide a file specifically to identify publications resulting from the work conducted in the prior year. The publications are organized by the proposal with which they were associated. This quarter, users identified 638 publications and conference papers that were published, in press, accepted, submitted, or in preparation.

ASC070031

1. G. GHIGLIOTTI, A. RAHIMIAN, G. BIROS, AND C. MISBAH, Vesicle migration and spatial organization driven by flow line curvature, Physical Review Letters, (2011). In press.
2. B. KAOUI, N. TAHIRI, T. BIBEN, H. EZ-ZAHRAOUY, A. BENYOUSSEF, G. BIROS, AND C. MISBAH, What dictates red blood cell shapes and dynamics in the microvasculature Physical Review E, (2011), pp. 1–11. in print.
3. L. MOON, D. LONG, S. JOSHI, V. TRIPATH, B. XIAO, AND G. BIROS, Parallel algorithms for clustering and nearest neighbor search problems in high dimensions, in 2011 ACM/IEEE conference on Supercomputing, Poster Session, Piscataway, NJ, USA, 2011, IEEE Press.
4. Massively parallel algorithms for nearest neighbor searches, SIAM Journal on Scientific Computing, (2012). In preparation.
5. A. RAHIMIAN, T. LOYSEL, AND G. BIROS, Vesicle collisions in shear flow, (2011). Submitted for publication.

6. A. RAHIMIAN, S. K. VEERAPANENI, D. ZORIN, AND G. BIROS, Fully discrete Galerkin spectral boundary integral method for vesicle flows with viscosity contrast, SIAM Journal on Scientific Computing, (2012), pp. 1–28. Submitted for publication.

AST020001

7. Marcus, P. S., and Shetty, S., "Jupiter's Zonal Winds: Are They Bands of Homogenized Potential Vorticity and Do They Form a Monotonic Staircase?, *Philosophical Transactions of the Royal Society A*., **369**, 701-832, 2011.
8. "The Universal Aspect Ratio of Vortices in Rotating Stratified Flows: Theory and Simulation", Hassanzadeh, P., Marcus, P. S., and Le Gal P., *Journal of Fluid Mechanics*, submitted, 2012.
9. "The Universal Aspect Ratio of Vortices in Rotating Stratified Flows: Experiment and Observation", Aubert, O., Le Bars, M, Le Gal, P. and Marcus, P. S., *Journal of Fluid Mechanics*, submitted, 2012.
10. Marcus, P. S., Asay-Davis, X., Wong, M. H., and de Pater, I., "Jupiter's New Red Oval: Its Dynamics; Why it Changed Color; and Its Relationship to Jovian Climate Change", *Journal of Heat Transfer*, accepted, 2011.
11. Wong M. H., de Pater. I., Asay-Davis, X., Marcus. P. S., and Go, C. "Vertical structure of Jupiter's Oval BA before and after it reddened: What changed?, *Icarus*, **215**, 211-225, 201.
12. Martin S. H., de Pater, I., and Marcus, P. S., "Neptune's Zonal Winds from Near-IR Keck Adaptive Optics Imaging in August 2001", *Astrophysics and Space Sciences*, accepted for publication, 2011 (available online, but not yet published in print).
13. Asay-Davis, X., Marcus, P. S. Wong, M. and de Pater, I., "Changes in Jupiter's Zonal Velocity Profile Between 1979 and 2008", *Icarus*, **211**, 1215-1232, 2011.
14. Jiang, C.-H. and Marcus, P. S., "Selection Rules for the Interactions between Inertial-Internal Waves", *Physics of Fluids*, in preparation, 2012.
15. Jiang, C.-H, Pei, S. and Marcus, P. S., "Critical Layers and Beams of Inertial-Internal Wave in Rotating Flows with Vertical Stratification and Shear, *Journal of Fluid Mechanics*, in preparation, 2012.

AST090005

16. Daruru, S., Gupta, G., Iliev, I. T., Xu, W., Navratil, P., Marin, N., & Ghosh, J. 2011, "Distributed, Scalable Clustering for Detecting Halos in Terascale Astronomy," in refereed proceedings of *KDCloud-10*, in press (http://ieeexplore.ieee.org/xpls/abs all.jsp?arnumber=5693293\&tag=1)
17. Obradovic, M., Kunz, M., Hindmarsh, M., & Iliev, I. T. 2011, "Particle motion in weak relativistic gravitational fields," *ArXiv e-prints*, (astro-ph/1106.5866) (http://adsabs.harvard.edu/abs/2011arXiv1106.5866O)
18. Riera, A., Raga, A. C., Mellema, G., Esquivel, A., & Vel´azquez P. F. 2011, "Physical conditions and excitation of FLIERs," in *Asymmetric Planetary Nebulae 5 Conference* (http://adsabs.harvard.edu/abs/2011apn5.confE.198R)
19. Arthur, S. J., Henney, W. J., Mellema, G., de Colle, F., & V´azquez-Semadeni, E. 2011, "Radiation-magnetohydrodynamic simulations of H II regions and their associated PDRs in turbulent molecular clouds," *MNRAS*, 414, 1747 (http://adsabs.harvard.edu/abs/2011MNRAS.414.1747A)
20. Arthur, S. J., Henney, W. J., Mellema, G., de Colle, F., & V´azquez-Semadeni, E. 2011, "Radiation-MHD Simulations of HII Region Expansion in Turbulent Molecular Clouds," in *Computational Star Formation*, eds. J. Alves, B. G. Elmegreen, J. M. Girart, & V. Trimble, Vol. 270 of *IAU Symposium*, 6 297 (http://adsabs.harvard.edu/abs/2011IAUS..270..297A)
21. Iliev, I. T., Moore, B., Gottl¨ober, S., Yepes, G., Hoffman, Y., & Mellema, G. 2011, "Reionization of the Local Group of galaxies," *MNRAS*, 413, 2093 (http://adsabs.harvard.edu/abs/2011MNRAS.413.2093I)
22. Paciga, G., Chang, T.-C., Gupta, Y., Nityanada, R., Odegova, J., Pen, U.-L., Peterson, J. B., Roy, J., & Sigurdson, K. 2011, "The GMRT Epoch of Reionization experiment: a new upper limit on the neutral hydrogen power spectrum at z ¼ 8.6," *MNRAS*, 413, 1174 (http://adsabs.harvard.edu/abs/2011MNRAS.413.1174P)

AST090105

23. 1. Log-normality in Global Models of Magnetized Accretion Disks, S.M.O'Neill, C.S.Reynolds, 2012, Astrophysical Journal, in preparation
24. Locality of Global Accretion Disk Turbulence, K.A.Sorathia, C.S.Reynolds, J.M.Stone, K.Beckwith, 2012, Astrophysical Journal, in preparation
25. Global Simulations of Accretion Disks I: Convergence and Comparisons with Local Models, K.A.Sorathia, C.S.Reynolds, J.M.Stone, K.Beckwith, 2012, Astrophysical Journal, submitted

26. Turbulent Transport in Global Models of Magnetized Accretion Disks, K.Sorathia, 2011, PhD thesis, University of Maryland College Park.
27. Low-frequency Oscillations in Global Simulations of Black Hole Accretion, S.M.O'Neill, C.S.Reynolds, M.C.Miller, K.A.Sorathia, 2011, Astrophysical Journal, 736, 107

AST090107

28. Bazot, M., Ireland, M.J., Huber, D., et al. (2011) *The radius and mass of the close solar twin 18Scorpii derived from asteroseismology and interferometry*. Astronomy & Astrophysics, 526:L4.
29. Brand˜ao, I.M., Do˘gan, G., Christensen-Dalsgaard, J., Cunha, M.S., Bedding, T.R., Metcalfe, T.S., Kjeldsen, H., Bruntt, H., & Arentoft, T. (2011) *Asteroseismic modelling of the solar-type subgiant star ßHydri*. Astronomy & Astrophysics, 527:A37.
30. Mathur, S., Handberg, R., Campante, T.L., et al. (2011) *Solar-like Oscillations in KIC 11395018 andKIC 11234888 from 8 Months of Kepler Data*. Astrophysical Journal, 733:95.
31. Creevey, O., Do˘gan, G., et al. (2011) *Fundamental properties of five Kepler stars using global asteroseismic quantities and ground-based observations*. Astronomy & Astrophysics, accepted (arXiv:1111.4615).
32. Borucki, W.J., Kock, D.G., Batalha, N., et al. (2011) *Kepler-22b: A 2.4 Earth-radius Planet in theHabitable Zone of a Sun-like Star*. Astrophysical Journal, accepted (arXiv:1112.1640).
33. Howell, S.B., Rowe, J.F., Bryson, S.T., et al. (2011) *Kepler-21b: A* 1.6 R_ *Planet Transiting the Bright Oscillating F Subgiant Star HD 179070*. Astrophysical Journal, accepted (arXiv:1112.2165).
34. Mathur, S., Metcalfe, T.S.,Woitaszek, M., et al. (2011) *A uniform asteroseismic analysis of 22 solartype stars from Kepler observations*. Astrophysical Journal, submitted.
35. Ceillier, T., Ballot, J., Garcia, R.A., Davies, G.R. Mathur, S., Metcalfe, T.S., & Salabert, D. (2011) *HD 49933: A laboratory for magnetic activity cycles*. Proceedings of the 2nd CoRoT symposium, submitted.
36. Metcalfe, T.S., et al. (2012) *Asteroseismology of the solar analogs 16 Cyg A & B from Kepler observations*. Astrophysical Journal Letters, submitted.
37. Doğan, G., Metcalfe, T.S., Woitaszek, M., et al. (2012) *A uniform asteroseismic analysis of bright solar-type stars from ground-based observations*. Astrophysical Journal, in preparation.
38. Brandão, I.M., et al. (2012) *An in-depth asteroseismic study of the Kepler target KIC 10273246*. Astronomy & Astrophysics, in preparation.
39. Doğan, G., et al. (2012) *Characterizing two solar-type Kepler subgiants with asteroseismology:*
40. *KIC 11395018 and KIC 10920273*. Astronomy & Astrophysics, in preparation.
41. Deheuvels, S., et al. (2012) *Asteroseismology using Mixed Modes and Avoided Crossings*. Astrophysical Journal, in preparation.
42. Silva Aguirre, V., et al. (2012) *Detection and properties of convective cores from Kepler data*. Astrophysical Journal, in preparation.

AST100014

43. "Neutron Star Crust and Molecular Dynamics Simulation", C. J. Horowitz, J. Hughto, A. Schneider, D. K. Berry, to appear as a chapter in the book "Neutron Star Crust", edited by C. A. Bertulani and J. Piekarewicz, available as arXiv:1109.5095.
44. "Direct MD simulation of liquid-solid phase equilibria for two-component plasmas", A. S. Schneider, J. Hughto, C. J. Horowitz, D. K. Berry, submitted to Phys. Rev. E, available as arXiv:1108.3101.
45. "Durability of neutron star crust", Andrey Chugunov and C. J. Horowitz, submitted to Contributions to Plasma Physics, 2011.
46. "Multi-messenger observations of neutron rich matter", C. J. Horowitz, Int J. Mod. Phys. E 20, 2077 (2011).
47. "Diffusion in Coulomb Crystals", J. Hughto, A. S. Schneider, C. J. Horowitz, D. K. Berry, Phys Rev E 84, 016401 (2011).
48. "Neutron rich matter, neutron stars and their crusts", C. J. Horowitz, J. Phys.: Conf. Ser. 312, 042003 (2011).

AST110016

49. Global baryon statistics of the Universe (Vogelsberger et al. 2011)
50. Hydrodynamics of galaxy formation (Sijacki et al. 2011)
51. Characteristics of galaxies and haloes (Keres et al. 2011)
52. Properties of gas disks (Torrey et al. 2011)

ATM050012

53. Houston, A. L., and R. B. Wilhelmson, 2011: The maintenance of deep convection in a low-shear environment and the impact of airmass boundaries, *Monthly Weather Review* (accepted).
54. Houston, A. L., and R. B. Wilhelmson, 2011: The impact of airmass boundaries on the propagation of deep convection: A modeling-based study in a high-CAPE, low-shear environment, *Monthly Weather Review* (accepted)
55. Houston, A. L., and R. B. Wilhelmson, 2011: The dependence of storm longevity on the pattern of deep convection initiation in a low-shear environment, *Monthly WeatherReview* (accepted).
56. Dorier, M., G. Antoniu, F. Cappello, M. Snir, and L. Orf, 2011: Damaris: Leveraging Multicore Parallelism to Mask I/O Jitter. Submitted for formal review to the International Conference for High Performance Computing, Networking, Storage and Analysis.
57. Jewett, B. F., R. M. Rauber and G. M. McFarquhar, and M. Pitcel, 2012: The role of stability and shear on mesoscale gravity wave evolution. In preparation for submission to *Monthly Weather Review.*

ATM050014

58. Lei, H., 2011: Simulations and Projections of Major Air Pollutants over the United States and Uncertainty Analyses. Ph.D. Thesis, University of Illinois.
59. Markus, M., D. J. Wuebbles, X.-L. Liang, K. Hayhoe, and D. A. R. Kristovich, 2012: Diagnostic analysis of future climate scenarios applied to urban flooding in the Chicago metropolitan area. Climatic Change, in press.
60. Olsen, S., D. Wuebbles, et al., 2012: Comparison of Global 3-D Aviation Emissions Datasets. Submitted to journal.
61. Patten, K., and D. Wuebbles, 2010: Atmospheric lifetimes and ozone depletion potentials of trans-1-chloro-3,3,3-trifluoropropylene and trans-1,2-dichloroethylene in a three-dimensional model. Atmos. Chem. Phys., 10, 10867-10874.
62. Patten, K. O., V. G. Khamaganov, V. L. Orkin, S. L. Baughcum, and D. J. Wuebbles, 2011: OH reaction rate constant, IR absorption spectrum, ozone depletion potentials and global warming potentials of 2-bromo-3,3,3-trifluoropropene, J. Geophys. Res. 116, D24307, 13 pp., doi:10.1029/2011JD016518.
63. Stoner, A. M., 2011: Development of a Quantile-Based Approach to Statistically Downscale Global Climate Models, Ph.D. Thesis, University of Illinois.
64. Wuebbles, D. J., K. O. Patten, D. Wang, D. Youn, M. Martínez-Avilés, and J. S. Francisco, 2011: Three-dimensional model evaluation of the ozone depletion potentials for n-propylbromide, trichloroethylene and perchloroethylene. Atmos. Chem. Phys., 11, 2371-2380.
65. Youn, D., K. O. Patten, D. J. Wuebbles, H. Lee, and C.-W. So, 2010: Potential impact of iodinated replacement compounds CF_3I and CH_3I on atmospheric ozone: A threedimensional modeling study. Atmos. Chem. Phys., 10, 10129-10144.
66. Zhang, X., X. Shao, K. Hayhoe and D. J. Wuebbles, 2011: Testing the structural stability of temporally dependent functional observations and application to climate projections. *Electronic Journal of Statistics*, 5, 1765-1796.

ATM70007

67. P.L. Pritchett and F.S. Mozer, \Rippling Mode in the Subsolar Magnetopause Current Layer and its In°uence on Three-Dimensional Magnetic Reconnection," J. Geophys.Res., 116, A04215, doi:10.1029/2010JA016190, 2011.
68. P.L. Pritchett and F.V. Coroniti, \Plasma Sheet Disruption by Interchange-Generated Flow Intrusions," Geophys. Res. Lett., 38, L10102, doi:10.1029/2011GL047527, 2011.
69. J. Egedal, A. Le, P.L. Pritchett, and W. Daughton, \Electron Dynamics in Two-Dimensional Asymmetric Anti-Parallel Reconnection," Phys. Plasmas, 18, 102901, 2011.
70. F.S. Mozer, D. Sundkvist, J.P. McFadden, P.L. Pritchett, and I. Roth, \Satellite Observations of Plasma Physics Near the Magnetic Field Reconnection X Line," J. Geophys. Res., 116, A12224, doi:10.1029/2011JA017109, 2011.
71. P.L. Pritchett, F.S. Mozer, and M. Wilber, \Intense Perpendicular Electric Fields Associated with Three-Dimensional Magnetic Reconnection at the Subsolar Magnetopause," J. Geophys. Res., submitted, 2011.

ATM110006

72. J.W. DePalma, B.R. Bzdek, D.J. Doren, M.V. Johnston, “Structure and Energetics of Nanometer Size Clusters of Sulfuric Acid with Ammonia and Dimethylamine”, Journal of Physical Chemistry A (2012), in press, doi:10.1021/jp210127w.

CHE040014

73. Peng Liu, Michael J. Krische, and Kendall N. Houk: "Mechanism and Origins of Regio- and Enantioselectivities in RhI-Catalyzed Hydrogenative Couplings of 1,3-Diynes and Activated Carbonyl Partners: Intervention of a Cumulene Intermediate," *Chem. Eur. J*., 17, 4021-4029 (2011).
74. Elizabeth H. Krenske, K. N. Houk, Andrew B. Holmes, and John Thompson: "Entropy versus Tether Strain Effects on Rates of Intramolecular 1,3-dipolar Cycloadditions of N-alkenylnitrones," *Tet. Lett*., 52, 2181-2184 (2011).
75. Peng Liu and K. N. Houk: "Theoretical Studes of Regioselectivity of Ni- and Rh-Catalyzed C-C Bond Forming Reactions with Unsymmetrical Alkynes," *Inorg. Chim. Acta*, 369, 2-14 (2011).
76. Nihan Celebi-Olcum, Ben W. Boal, Alexander D. Huters, Neil K. Garg, and K. N. Houk: "Why Do Some Fischer Insolizations Fail?" *J. Am. Chem. Soc*., 133, 5752-5755 (2011).
77. Kyle W. Quasdorf, Aurora Antoft-Finch, Peng Liu, Amanda L. Silberstein, Anna Komaromi, Tom Blackburn, Stephen D. Ramgren, K. N. Houk, Victor Snieckus, and Neil K. Garg: "Suzuki-Miyaura Cross-Coupling of Aryl Carbamates and Sulfamates: Experimental and Computational Studies," *J. Am. Chem. Soc*., 133, 6352-6363 (2011).
78. Peng Liu, John Montgomery, and K. N. Houk: "Ligand Steric Contours to Understand the Effects of N-Heterocyclic Carbene Ligands on the Reversal of Regioselectivity in Ni-Catalyzed Reductive Couplings of Alkynes and Aldehydes," *J. Am. Chem. Soc*., 133, 6956-6959 (2011).
79. Hao Wang, K. N. Houk, Damian A. Allen, and Michael E. Jung: "Computational Elucidation of the Origins of Reactivity and Selectivity in Non-Aldol Aldol Rearrangements of Cyclic Epoxides," *Org. Lett*., 13, 3238-3241 (2011).
80. Yu Lan and K. N. Houk: "Mechanism of the Palladium-Catalyzed Addition of Arylboronic Acids to Enones: A Computational Study," *J. Org. Chem*., 76, 4905-4909 (2011).
81. Elizabeth H. Krenske, Russell C. Petter, Zhendong Zhu, and K. N. Houk: "Transition States and Energetics of Nucleophilic Additions of Thiols to Subtituted alpha,beta-Unsaturated Ketones: Substituent Effects Involve Enone Stabilization, Product Branching, and Solvation," *J. Org. Chem.*, 76, 5074-5081 (2011).
82. Joann M. Um, Daniel A. DiRocco, Elizabeth L. Noey, Tomislav Rovis, and K. N. Houk: "Quantum Mechanical Investigation of the Effect of Catalyst Fluorination in the Intermolecular Asymmetric Stetter Reaction," *J. Am. Chem. Soc*., 133, 11249-11254 (2011).
83. Elizabeth L. Noey, Xiang Wang, and K. N. Houk: "Selective Gold(I)-Catalyzed Formation of Tetracyclic Indolines: A Single Transition Structure and Bifurcations Lead to Multiple Products," *J. Org. Chem.*, 76, 3477-3483 (2011).
84. Nihan Çelebi-Ölçüm, Yu-hong Lam, Kenneth B. Ling, Andrew D. Smith, and K. N. Houk: "Pericyclic cascade with chirality transfer: Reaction pathway and origin of enantioselectivity of the hetero-Claisen approach to oxindoles," *Angew. Chem., Int. Ed*., 50, 11478-11482 (2011) 2
85. Lai Xu, Charles E. Doubleday, and K. N. Houk: "Dynamics of Carbene Cycloadditions," *J. Am. Chem. Soc.* 133, 17848-17854. (2011).
86. Ryan J. Marcheschi, Han Li, Kechun Zhang, Elizabeth L. Noey, Seonah Kim, Asha Chaubey, K. N. Houk, and James C. Liao: "A Synthetic Recursive "+1" Pathway for Carbon Chain Elongation," *ACS Chem. Biol.* Submitted.
87. Elizabeth L. Noey, Yingdong Luo, Liming Zhang, and K. N. Houk: "Mechanism of Gold(I)- Catalyzed Rearrangements of Acetylenic Amine-N-Oxides: Computational Investigations Lead to a New Mechanism Confirmed by Experiment," *J. Am. Chem. Soc.* In Press.
88. Daniel A. DiRocco, Elizabeth L. Noey, Tomislav Rovis, and K. N. Houk: "Catalytic Asymmetric Intermolecular Stetter Reactions of Enolizable Aldehydes with Nitrostyrenes: Computational Study Provides Insight into the Success of a New Catalyst," *Angew. Chem. Int. Ed.* In Press.
89. Blanton N. Martin, Hsiang-Yu Chen, Yang Yang, and K. N. Houk: "Influence of Side Chains on Intermolecular Interactions in Cyclopentadithiophene- and Silacyclopentadithiophene-based Polymers," *Adv. Funct. Mater.* Submitted.

CHE090047

90. "O + C2H4 Potential Energy Surface: Initial Excited States, Biradicals, and Lowest-lying Singlet at the Multireference Level", A.C. West, J.D. Lynch, B. Sellner, H. Lischka, W.L.Hase, and T.L. Windus, *Theor. Chem. Acc.*, Accepted.
91. "Fragmentation and reactivity in collisions of protonated diglycine with chemically modified perfluorinated alkylthiolate-self-assembled monolayer surfaces", G. L. Barnes, K. Young, L. Yang, and W. L. Hase, *J. Chem. Phys.*, **2011**, *134*, 094106.
92. "Photodynamics and Time-Resolved Fluorescence of Azobenzene in Solution: A Mixed Quantum-Classical Simulation", T. Cusati, G. Granucci, and M. Persico, *J. Am. Chem.Soc.* **2011**, *113*, 5109-5123.

CHE100005

93. Lim, D.-H., Negreira, A.S., and Wilcox, J., "DFT Studies on the interaction of defective graphene-supported Fe and Al nanoparticles", ***J.** Phys. Chem. C 115 (18*), 8961 (2011).
94. Lim, D.-H. and Wilcox, J., "DFT-based study on oxygen adsorption on defective graphene-supported Pt nanoparticles", ***J.** Phys. Chem. C*, 115 (46), 22742 (2011). (Partial Support from XSEDE)

CHE100033

95. Y. Dai and E. Blaisten-Barojas, "Polaron Pairs and Bipolarons in Oligopyrrole Dications," submitted (2012).
96. R. Massaro and E. Blaisten-Barojas,"Theoretical investigation of the photophysics of methyl salicylate isomers," J. Chem. Phys. 135, 164306 (2011).
97. R. Massaro and E. Blaisten-Barojas, "Density Functional Theory Study of Dipicolinic Acid Isomers and Crystalline Polytypes." Computational and Theoretical Chemistry 977, 148–156 (2011).
98. J. Lyver and E. Blaisten-Barojas, "Lattice Thermal Conductivity in SiC Nanotubes, Nanowires and Nanofilaments: a Molecular Dynamics Study," J. of Computational and Theoretical Nanoscience 8, 529-534 (2011).
99. Y. Dai, S. Chowdhury, E. Blaisten-Barojas, "First-principles study of the structure, energetics and electron affinity of negatively charged oligopyrroles," Int. J. Quantum Chem. 111, 2295-2305 (2011).

CHE100039

100. S. Fonseca dos Santos, N. Balakrishnan, R. C. Forrey, and P. C. Stancil, Vibrationvibration and vibration-rotation energy transfer in H_2-H_2 collisions, J. Chem. Phys. (in preparation).

CHE100123

101. Experimental and Computational Evidence for Gold Vinylidenes: Generation from Terminal Alkynes via a Bifurcation Pathway and Facile C-H Insertions, Longwu Ye, Yanzhao Wang, Donald H. Aue,* and Liming Zhang,* *J. Amer. Chem. Soc*. **2012**, *134*, 31-34.
102. Regioselective Reductions of □□□-Disubstituted Enones Catalyzed by Nonracemically Ligated Copper Hydride, Tetrahedron, **2012**, full paper, in press, Karl R. Voigtritter, Nicholas A. Isley, Ralph Moser, Donald H. Aue,* and Bruce H. Lipshutz,*.
103. Carbocations. Wiley Interdisciplinary Reviews: Computational Molecular Science **2011**, *1*, 487-508, Donald H. Aue.
104. Synthesis & Redox Properties of Quinone-modified Analogs of Coenzyme Q_{10}. Did Nature Choose Wisely?, Bruce H. Lipshutz,* Danielle M. Nihan, Donald H. Aue,* manuscript completed, but awaiting biological testing.
105. Calculation of Electron Affinities, Ionization Potentials, and Reduction and Oxidation Potentials of Large Polycyclic Aromatic Hydrocarbons, Leon D. Betowski* and Donald H. Aue,* *J. Phys. Chem*. to be submitted.

CHE110011

106. Manna, K.; Xu, S.; Sadow, A. D. "A Highly Enantioselective Zirconium Catalyst for Intramolecular Alkene Hydroamination: Activity Below Room Temperature and a New Mechanism" *Angew. Chem Int. Ed.*, **2011**, *50*, 1865-1868. http://dx.doi.org/10.1002/anie.201006163
107. Manna, K.; Kruse, M. L.; Sadow, A. D. "Concerted C-N/C-H Bond Formation in Highly Enantioselective Yttrium(III)-Catalyzed Hydroamination." *ACS Catalysis*, **2011**, *1*, 1637-1642. http://dx.doi.org/10.1021/cs200511z
108. Dunne, J. F.; Neal, S. R.; Engelkemier, J.; Ellern, A.; Sadow, A. D. "Tris(oxazolinyl)boratomagnesium-catalyzed Cross-dehydrocoupling of Organosilanes with Amines, Hydrazine, and Ammonia" *J. Am. Chem. Soc.* **2011**, *133*, 16782-16785. http://dx.doi.org/10.1021/ja207641b

CHE120020

109. B. C. Tlach, A. L. Tomlinson, A. Bhuwalka and M. Jeffries-EL "Tuning the Optical and Electronic Properties of 4,8-Disubstituted Benzobisazoles via Alkyne Substitution," *Journal of Organic Chemistry*, ASAP, DOI: 10.1021/jo201078w.
110. J. S. Klimavicz, J. F. Mike, A. Bhuwalka, A. L. Tomlinson, and M. Jeffries-EL "Synthesis of Benzobisazoles Based on $D-\pi-A-\pi-D$ Organic Chromophores with Variable Optical and Electronic Properites," accepted by the *J. Pure Appl Chem*. **In Progress**
111. A. Tomlinson, A. Ryno, and D. Lehman, "Benchmarking TDDFT Methods for Benzobiszaole Polymers," will be submitted to *Journal of Physical Chemistry* in progress.

112. B. C. Tlach, A. L. Tomlinson, and M. Jeffries-EL "Improving the Optical and Electronic Properties of Benzobisoxazoles via Aryl Substitution," in progress.

CHE120021

113. Wei-Li Li, Constantin Romanescu, Timur R. Galeev, Zachary A. Piazza, Alexander I. Boldyrev, and Lai-Sheng Wang. "Transition-Metal-Centered Nine-Membered Boron Rings: $M©B_9$ and $M©B_9^-$ (M=Rh, Ir)". *J. Am. Chem. Soc.* Published on Web Dec. 13, 2011. *In press.*
114. Zachary A. Piazza, Wei-Li Li, Constantin Romanescu, Alina P. Sergeeva, Lai-Sheng Wang, and Alexander I. Boldyrev. "Anionic Boron Clusters Continue to be Planar at 21: A Photoelectron Spectroscopy and Ab Initio study of the B_{21}^- Cluster". *J. Chem. Phys. In Progress.*

CHE120025

115. Temelso, B.; Archer, K. A.; Shields, G. C. Benchmark Structures and Binding Energies of Small Water Clusters with Anharmonicity Corrections. ***J. Phys. Chem. A*** **2011**, *115*, 12034.
116. Temelso, B.; Shields, G. C. The Role of Anharmonicity in Hydrogen-Bonded Systems: The Case of Water Clusters. ***J. Chem. TheoryComput.*** **2011**, *7*, 2804.
117. Husar, D.; Temelso, B.; Ashworth, A.; Kirschner, K. N.; Shields, G. C. AQuantum Chemical Study of HSO_4^- Hydrate Formation. ***J. Phys. Chem.A*** **2012**, In preparation.
118. Temelso, B.; Morrell, T. E.; Shields, R. M.; Allodi, M. A.; Wood, E. K.; Kirschner, K. N.; Castonguay, T. C.; Archer, K. A.; Shields, G. C. Quantum Mechanical Investigation of the Binary Nucleation of Sulfuric Acid and Water the Atmosphere. ***J. Phys. Chem. A*** **2012**, Submitted.
119. Temelso, B.; Phan, T. N.; Shields, R. M.; Kirschner, K. N.; Shields, G. C. The Hydration of Sulfuric Acid Dimers and Its Atmospheric Implications. ***J.Phys. Chem. A*** **2012**, In Preparation.

CTS030025

120. Yazdani, A., and Bagchi, P. 2011. Phase diagram and breathing dynamics of a single red blood cell and a biconcave capsule in dilute shear flow. *Physical ReviewE*, 84, 026314.
121. Zhao, M., and Bagchi, P. 2011. Dynamics of microcapsules in oscillating shear flow. *Physics of Fluids*, 23, 111901.
122. Yazdani, A., Kalluri, R.M., and Bagchi, P., 2011, Tank-treading and tumbling frequencies of capsules and red blood cells. *Physical Review E* 83, 046305.
123. Bagchi, P., & Kalluri, R.M.. 2011. Dynamic Rheology of a Dilute Suspension of Elastic Capsules: Effect of Capsule Tank-Treading, Swinging, and Tumbling. *Journal of Fluid Mechanics*, 669, 498.

CTS050054

124. Samarth S. Raut, Anirban Jana, Ender A. Finol, Effects Of Shape Versus Material Model Variations On AAA Wall Mechanics, *Sixth M.I.T. Conference on Computational Fluid and Solid Mechanics,*
125. *Boston, MA,* June 2011 (Peer-reviewed abstract and presentation). • S. S. Raut, P. Liu, A. Jana, E.A. Finol, Aortic Wall Mechanics: A Geometry-driven Problem, Summer Bioengineering Conference, Farmington, PA, June 2011 (Peer-reviewed abstract and poster).
126. Santanu Chandra, Samarth S. Raut, Anirban Jana, Robert Biederman, Mark Doyle, Satish Muluk, Ender A. Finol, Efficient Fluid-Structure Interaction Modeling of Abdominal Aortic Aneurysms with Patient Specific Inflow Boundary Conditions, *Journal of Biomechanical Engineering* (manuscript in preparation).
127. S. S. Raut, J. Shum, S. Chandra, A. Jana, P. Liu, K. Lee, E. S. Di Martino, T. C. Doehring, and E. A. Finol, AAAVASC: A novel Integrated Approach for Image Based Modeling toward Individualized AAA Rupture Risk Assessment, *Proceedings of the 2011 Biomedical Engineering Society Annual Fall Meeting,* Hartford, CT, October 12-15.
128. Peng Liu, Samarth S. Raut, Anirban Jana, Todd Doehring, Ender A. Finol, A Volume Meshing Strategy for Patient Specific Abdominal Aortic Aneurysms, *Annals of Biomedical Engineering* (manuscript in preparation).
129. Judy Shum, Risk Assessment of Abdominal Aortic Aneurysms by Geometry Quantification Measures, PhD Thesis, Carnegie Mellon University, Pittsburgh, PA, 2011.
130. Washington, C.B., Shum, J., Muluk, S.C., and Finol, E.A., 2011, "The association of wall mechanics and morphology: a case study of abdominal aortic aneurysm growth," *Journal of Biomechanical Engineering,* Vol. 133, No. 10, 104501.
131. Gail M. Siewcorek and Ender A. Finol, Computational Modeling of Distal Protection Filters, *Journal of Endovascular Therapy,* Vol. 17, No. 6, pp. 777-788, 2010.

CTS070066

132. Cruz Perez B., Toro J., Lucena J., Bonilla R. & Leonardi S. (2011) Turbulent flow in a square channel with transverse square bars and V-shaped turbulators on the walls, ASME technical paper (under review).

CTS100024

133. Baraldi, A. & Ferrante, A. (2012) "DNS of fully-resolved droplet-laden isotropic turbulence: a mass-conserving volume of fluid method". ***23rd International Congress of Theoretical and Applied Mechanics, ICTAM***, *Submitted in January 2012 Bejing, China, August, 2012*
134. Baraldi, A. & Ferrante, A. (2012) "A VoF method for DNS of droplet-laden incompressible turbulence"
135. ***7th International Conference on Computational Fluid Dynamics,*** *Submitted in December 2011 Kohala Coast, Hawaii, July 9-13, 2012*
136. Dodd M., Webster K. & Ferrante A. (2012) "Direct numerical simulation of particle dispersion in a spatially developing turbulent boundary layer". ***7th International Conference on Computational Fluid Dynamics,*** *Submitted in December 2011Kohala Coast, Hawaii, July 9-13, 2012*
137. Baraldi, A. & Ferrante, A. (2011) "DNS of droplet-laden incompressible turbulence: surface tension in a VoF method" ***64th American Physical Society Meeting, Division of Fluid Dynamics****Baltimore, MD, November 20-22, 2011.*
138. Baraldi, A. & Ferrante, A. (2011) "A mass-conserving volume of fluid method: volume tracking in incompressible isotropic turbulence". ***Journal of Computational Physics***, *Submitted in October 2011*.
139. Ferrante A. & Baraldi A. (2011) "A mass-conserving volume of fluid method: volume tracking in incompressible isotropic turbulence". ***20th AIAA Computational Fluid Dynamics Conference,*** *2011-3829, pp. 1-16. Honolulu, Hawaii, June 27-30, 2011*
140. Lucci F., Ferrante A. & Elghobashi S. (2011) "Is Stokes number an appropriate indicator for turbulence modulation by particles of Taylor-length-scale size?" ***Physics of Fluids***, *Vol. 23, 025101, pp. 1-7*

CTS100055

141. Predictive Structure-Reactivity Models for Rapid Screening of Pt-based Multimetallic Electrocatalysts for the Oxygen Reduction Reaction H. Xin, A. Holewinski, S. Linic, *ACSCatalysis*, **2**, 12, 2012. (Selected as cover article of issue.)
142. Development of physically transparent models for predicting the local chemical reactivity of metal surfaces with perturbed electronic properties, H. Xin, T. Cleve, S. Linic, To be submitted.
143. First-principles Investigation of the Environment-dependent Alkali Promotion Mechanism for Oxidation Reactions on the Ag(111) Surface, H. Xin, S. Linic, In Preparation

CTS110027

144. E. Knudsen, Shashank, and H. Pitsch. \An LES Study Of Spray And Pollutant Relationships In A Liquid Fueled Aircraft Combustor." In preparation for submission to Combustion and Flame.
145. E. Knudsen, Shashank, and H. Pitsch. \LES Of A Model Aircraft Combustor: Spray Model Assumptions And Pollutant Formation." 64th APS Division of Fluid Dynamics Meeting. Baltimore, Maryland, 2011.
146. E. Knudsen. \Understanding And Reducing Reactive LES Model Sensitivities." General Electric's Symposium On Large Eddy Simulations Of Turbulent Reacting Flows For Gas Turbine Combustor Design. Niskayuna, New York, 2011.
147. Shashank, E. Knudsen, H. Pitsch. \Direct numerical simulation study of spray combustion in liquid-fuel jets.", 13th international conference on numerical combustion, Corfu, Greece, 27-29 April 2011
148. V. Mittal, H. Pitsch. \A Flamelet Model For Premixed Combustion Under Variable Pressure Conditions.", Submitted to Proceedings of Combustion Symposium, 2012.

CTS110029

149. E. Knudsen, Shashank, and H. Pitsch. \An LES Study Of Spray And Pollutant Relationships In A Liquid Fueled Aircraft Combustor." In preparation for submission to Combustion and Flame.
150. E. Knudsen, Shashank, and H. Pitsch. \LES Of A Model Aircraft Combustor: Spray Model Assumptions And Pollutant Formation." 64th APS Division of Fluid Dynamics Meeting. Baltimore, Maryland, 2011.
151. E. Knudsen. \Understanding And Reducing Reactive LES Model Sensitivities." General Electric's Symposium On Large Eddy Simulations Of Turbulent Reacting Flows For Gas Turbine Combustor Design. Niskayuna, New York, 2011.
152. Shashank, E. Knudsen, H. Pitsch. \Direct numerical simulation study of spray combustion in liquid-fuel jets.", 13th international conference on numerical combustion, Corfu, Greece, 27-29 April 2011

153. V. Mittal, H. Pitsch. \A Flamelet Model For Premixed Combustion Under Variable Pressure Conditions.", Submitted to Proceedings of Combustion Symposium, 2012.
154. Donzis D.A. (2011) “Asymptotic states in turbulent mixing: the role of Pecl´et number in scalar fluxes, dissipation, spectra and intermittency,” Turbulence, Mixing and Beyond (invited lecture), International Centre for Theoretical Physics, Triest, Italy.

CTS120015

155. S. L. Yilmaz, M. B. Nik, M. R. H. Sheikhi, P. Strakey, and P. Givi, “An Irregularly Portioned Lagrangian Monte Carlo Method for Turbulent Flow Simulation,” Journal of Scientific Computing, 47:1, 109-125, (2011).
156. S. L. Yilmaz, N. Ansari, P. H. Pisciuneri, M. B. Nik, C. C. Otis, P. Givi, “Advances in Filtered Density Function Modeling and Simulation,” 47th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, San Diego, CA, AIAA Paper 2011-5918, 1-11, July 31, (2011).
157. P. H. Pisciuneri, S. L. Yilmaz, P. A. Strakey and P. Givi, “An Irregularly Portioned FDF Solver,” 64th Annual Meeting of the Division of Fluid Dynamics of the American Physical Society, Baltimore, MD, Nov. 20-22, (2011).
158. S. L. Yilmaz, P.H. Pisciuneri, M. B. Nik, P. Givi. “Modern, Scalable, Reliable Modeling of Turbulent Combustion,” Proceedings of the 2011 TeraGrid Conference: Extreme Digital Discovery, Salt Lake City, Utah, Jul. 18-21, American Computing Machinery Proceedings, 11, 1-2, (2011).

DBS110003

159. Clark, J. H., Dyer, C., Lavie, A., and Smith, N. A. (2011). Better hypothesis testing for statistical machine translation: Controlling for optimizer instability. In Proceedings of the Annual Meeting of the Association for Computational Linguistics, companion volume, Portland, OR.
160. Cohen, S. B. (2011). Computational Learning of Probabilistic Grammars in the Unsupervised Setting. PhD thesis, Carnegie Mellon University, Pittsburgh, PA.
161. Cohen, S. B., Das, D., and Smith, N. A. (2011). Unsupervised structure prediction with non-parallel multilingual guidance. In Proceedings of the Conference on Empirical Methods in Natural Language Processing, Edinburgh, UK.
162. Das, D. and Smith, N. A. (2011). Semi-supervised frame-semantic parsing for unknown predicates. In Proceedings of the Annual Meeting of the Association for Computational Linguistics, Portland, OR.
163. Das, D. and Smith, N. A. (2012). Semi-supervised construction of compact lexicons using sparse gaussian and entropic _elds. In review.
164. Dyer, C., Clark, J. H., Lavie, A., and Smith, N. A. (2011). Unsupervised word alignment with arbitrary features. In Proceedings of the Annual Meeting of the Association for Computational Linguistics, Portland, OR.
165. Yogatama, D., Heilman, M., O'Connor, B., Dyer, C., Routledge, B. R., and Smith, N. A. (2011a). Predicting a scientic community's response to an article. In Proceedings of the Conference on Empirical Methods in Natural Language Processing, Edinburgh, UK.
166. Yogatama, D., Heilman, M., O'Connor, B., Dyer, C., Routledge, B. R., and Smith, N. A. (2011b). Predicting responses and discovering social factors in scienti_c literature. Technical Report CMU-LTI-11-015, Carnegie Mellon University, Pittsburgh, PA.
167. Yogatama, D., Sim, Y., and Smith, N. A. (2012). A probabilistic model for canonicalizing named entity mentions. In review.

DMR090071

168. *Calcium coated B80 fullerene: A study on various coating configurations of B80* Marco Olguin, Tunna Baruah, Rajendra R. Zope Chemical Physics Letters, 514,66 (2011).
169. *Snub boron nanostructures: Chiral fullerenes, nanotubes and planar sheet* Rajendra R. Zope and T. Baruah Chemical Physics Letters, 501, 193 (2011).

DMR100033

170. Y.U. Wang, D.Q. Tan, “Computational Study of Filler Microstructure and Effective Property Relations in Dielectric Composites,” *J. Appl. Phys.*, **109**, 104102-1-8, 2011.
171. Y.U. Wang, D.Q. Tan, J. Krahn, “Computational Study of Dielectric Composites with Core-Shell Filler Particles,” *J. Appl. Phys.*, **110**, 044103-1-6, 2011.
172. Y.M. Jin, “Competitive Roles of Elastic and Magnetic Interactions in Twin Boundary Behaviors of Magnetic Shape Memory Alloys,” *Appl. Phys. Lett.*, **99**, 062507-1-3, 2011.

173. Y.M. Jin, H.D. Chopra, "Altering Magnetostrictive Strain Pathways via Morphology of Spontaneously Aligned Domains," *Phys. Rev. B.*, **84**, 140401(R)-1-4, 2011.
174. J.E. Zhou, T.L. Cheng, Y.U. Wang, "Correlated Nucleation and Self-Accommodating Kinetic Pathway of Ferroelectric Phase Transformation," *J. Appl. Phys.*, 2012 (accepted).
175. T.L. Cheng, Y.U. Wang, "Spontaneous Formation of Stable Capillary Bridges for Firming Compact Colloidal Microstructures in Phase Separating Liquids: A Computational Study," *Langmuir*, 2012 (accepted).
176. B. Wang, Y.M. Jin, "Magnetization and Magnetostriction of Terfenol-D near Spin Reorientation Boundary," *J. Appl. Phys.*, 2012 (submitted to *J. Appl. Phys.*)
177. T.L. Cheng, Y.U. Wang, "Diffuse Interface Field Approach to Modeling Sterically-Shielded Shape-Anisotropic Particles at Curved Fluid Interfaces: The Role of Laplace Pressure," (in preparation to be submitted to *J. Colloid Interface Sci.*)
178. W.F. Rao, J.E. Zhou, Y.U. Wang, "Poling-Induced Single-Crystal-Like Piezoelectric Anisotropy in Untextured Ferroelectric Polycrystalline Ceramics: Phase Field Modeling," (in preparation to be submitted to *Appl. Phys. Lett.*)
179. T.L. Cheng, I. Kuyanov, Y.U. Wang, "Charge Compensation and Aging Mechanisms in Iron-Doped Ferroelectric Titanate Perovskites," (ongoing, to be submitted to *Phys. Rev. B*)
180. L.D. Geng, I. Kuyanov, Y.U. Wang, "Nanostructures and Novel Properties of Ferroelectric Domain Walls," (ongoing, to be submitted to *Appl. Phys. Lett.*)

DMR110037

181. Predrag S. Krstic, Jean Paul Allain, Alain Allouched, Jacek Jakowski , Jonny Dadras , Chase N. Taylor, *"Dynamics of deuterium retention and sputtering of Li–C–O surfaces",* Fusion Engineering and Design, (2011), doi:10.1016/j.fusengdes.2011.07.009
182. A. Allouche and P.S. Krstic, "Atomic hydrogen adsorption on lithium-doped graphite surfaces", Carbon 50, 510 (2012), online- doi:10.1016/j.carbon.2011.09.006
183. Taylor, C. N. , El-Atwani, O., Heim, B. , Allain, J. P. , Skinner, C. H. , Roquemore, A. L. , Kugel, H., Krstic, Predrag S , Surface Chemistry and Physics of D-Retention in Lithited Graphite, J. Nuclear Mat., (in press, 2011), doi:10.1016/j.jnucmat.2010.09.049
184. J. Jakowski, S. Irle, K. Morokuma, "*Time –dependent quantum dynamical simulations of C2 condensation under extreme conditions*", Phys. Chem. Chem. Phys., (in press) DOI:10.1039/C1CP22035G.
185. Jacek Jakowski, Stephan Irle, Bobby Sumpter and Keiji Morokuma , *"Modeling charge transfer in fullerene collisions via real time electron dynamics simulations"* (in preparation)
186. R C Ehemann, J Dadras, P R C Kent, J Jakowski, and P S Krstić, "Sensoring of hydrogen by grapheme", Nano Research Letters (submitted)
187. Jingsong Huang, Ariana Beste, Jarod Younker, Alvaro Vazquez-Mayagoitia, Edurado Cruz-Silva, Miguel Fuentes-Cabrera, Jacek Jakowski, Alejandro Lopez-Bezanilla, Vincent Meunier, Bobby G. Sumpter , *"Advancing Understanding and Design of Functional Materials through Theoretical and Computational Chemical Physics"* in J. Leszczynski, M. K. Shukla, Eds. "Practical Aspects of Computational Chemistry II: An Overview of the Last Two Decades and Current Trends", Springer-European Academy of Sciences (2011). ISBN 978-94-007-0922-5. (in press)
188. Jacek Jakowski, Stephan Irle, Keiji Morokuma, "Quantum Chemistry: Propagation of electronic structure in GPU" in "GPU Computing Gems. Emerald Edition", Eds. Wen-mei Hwu, ch 5. pp. 59-74, Morgan-Kauffman (2011) , ISBN:9780123859631
189. J. Jakowski, Steven Stuart, Dulwa Nugawela, Keiji Morokuma, Stephan Irle, *"From fullerenes to nano-devices. Modeling reactions between carbon-nanostructures from quantum chemical molecular dynamics simulations",* TeraGrid 2011, Salt Lake City
190. Jacek Jakowski, Steve Stuart, Dulma Nugawela, Predrag Krstic, Jonny Dadras, Sophya Garashchuk, Stephan Irle, "Building Community of practice through collaborative research: Modeling of nano-scale carbon and metalized carbon materials", TeraGrid 2011, Salt Lake City
191. Jacek Jakowski, "Time dependent quantum dynamical simulations of C2 condensation under extreme conditions", 7th Congress of International Society for Theoretical Chemical Physics, Sept.2-8,2011, Waseda University, Tokyo, Japan
192. Hu-Jun Qian, Stephan Irle, "QM/MD simulations of a Nanotube Welding process on Graphene Hole" (in preparation)
193. S. J. Stuart, D. E. Nugawela, and J. Jakowski, "Multiscale Approach to Prediction of Fullerene Fusion Products", *Phys. Rev. Lett.* (to be submitted), Jan. 2012.

DMS060014

194. Maurini, C., Bourdin, B., Gauthier, G., and Lazarus, V. (2011). Crack patterns in directional drying: numerical simulations using a variational approach and experiments. Submitted.

195. Bourdin, B., Larsen, C., and Richardson, C. A time-discrete model for dynamic fracture based on crack regularization. International Journal of Fracture, 168, 133143, 2011.

EAR100012

196. Wang, J., E. Mehnert and Y.F., Lin, 2011. Protecting Drinking Water by Coupling Two Models Associated with Geologic Carbon Sequestration in Deep Saline Aquifers, MODFLOW and More 2011: Integrated Hydrologic Modeling - Conference Proceedings, p. 363-365 (abstract, presentation by Y.F. Lin)
197. Mehnert, E., J. Damico, S. Frailey, H. Leetaru, R. Okwen, Y.-F. Lin, J. Wang, N. Adams, O. Garcia-Cabrejo, B. Storsved and A. Valocchi, 2011. Basin-scale Modeling of CO_2 Sequestration in the Mount Simon Sandstone of the Illinois Basin—Status Report, *Midwest Carbon SequestrationScience Conference*, November 2011, Champaign, IL

EAR100021

198. Cande, S., and D. Stegman, Indian and African plate motions driven by the push force of the Reunion plume head, Nature, 475 (7354), 47{52, 2011.
199. Capitanio, F. A., C. Faccenna, S. Zlotnik, and D. Stegman, Subduction dynamics and the origin of Andean orogeny and Bolivian Orocline, Nature, 480 (7375), 83{86, 2011.
200. Druken, K. A., D. R. Stegman, and C. Kincaid, Plume-slab interaction: Insight from numerical models, Phys. Earth Planet. Int., (in preparation) 2012.
201. Liu, L., and D. R. Stegman, Segmentation of the Farallon slab, Earth and Planetary Science Letters, 311, 1{10, doi:10.1016/j.epsl.2011.09.027, 2011.
202. Liu, L., and D. R. Stegman, Onset of Columbia River ood basalt controlled by propagating rupture of the Farallon slab, Nature, (in press) 2012.

EAR120003

203. Wang, J, Zhou, B, Shi, XT, Liu, XS, and Guo, XE, “Bone Mineral Density (BMD) of Individual Trabeculae in Human Trabecular Bone”, ASBMR 2011 Annual Meeting, September 16-20, 2011, San Diego, CA.
204. 2. Shi, XT, Zhou, B, Adams, MF, and Guo, XE, “Boundary Conditions in Nonlinear μCT Image Based Finite Element Modeling Do Not Affect Predictions in Yield Strength of Human Trabecular Bone”, ORS 2012 Annual Meeting, February 4-7, 2012, San Francisco, CA.
205. 3. Wang, J, Zhou, B, Shi, XT, Liu, XS, and Guo, XE, “Bone Mineral Density (BMD) of Individual Trabecula in Human Trabecular Bone”, ORS 2012 Annual Meeting, February 4-7, 2012, San Francisco, CA.

IRI120005

206. Edwards, R., J. New, and L. Parker, “Predicting Future Hourly Residential Electrical Consumption: A Machine Learning Case Study”. Journal article currently in review by the editor of “Energy and Buildings”, December 2011

MCA01S007

207. A. Alexandru, T. Draper, I. Horvath, and T. Streuer, The Analysis of Space-Time Structure in QCD Vacuum II: Dynamics of Polarization and Absolute X-Distribution Annals Phys. 326, 1941 (2011), [arXiv:1009.4451].
208. A. Alexandru and U. Wenger, QCD at non-zero density and canonical partition functions with Wilson fermions, Phys. Rev. D83, 034502 (2011), [arXiv:1009.2197].
209. J.M.M. Hall, F.X. Lee, D.B. Leinweber, K.F. Liu, N. Mathur, R.D. Young, and J.B. Zhang, Chiral extrapolation beyond the power-counting regime Phys. Rev. D 84, 114011 (2011), [arXiv:1101.4411]. UK/11-02
210. A. Li, Andrei Alexandru, and Keh-Fei Liu, Critical point of $N_f = 3$ QCD from lattice simulations in the canonical ensemble Phys. Rev. D84 071503 (2011), [arXiv:1103.3045].
211. A. Alexandru, C. Pelissier, B. Gamari, Multi-mass solvers for lattice QCD on GPUs submitted for publication, [arXiv:1103.5103].
212. A. Alexandru, M. Lujan, C. Pelissier, B. Gamari, and F. X. Lee Efficient implementation of the overlap operator on multi-GPUs submitted for publication, [arXiv:1106.4964]
213. UK/11-05 K.F. Liu Charge-dependent Azimuthal Correlations in Relativistic Heavy Ion Collisions and Electromagnetic Effects Accepted for publication in Phys. Rev. C, [arXiv:1109.4883],
214. UK/11-07 H.W. Lin and K.F. Liu Comment on ‘Controversy concerning the definition of quark and gluon angular momentum’Accepted for publication in Phys. Rev. D, [arXiv:1111.0678]

MCA01S018

215. C. Gajewski, A. Dagcan, B. Roux & C. Deutsch. Biogenesis of the pore architecture of a voltage-gated potassium channel. Proc Natl Acad Sci U S A 108, 3240-3245, (2011). -PMC ID: 3044366
216. H. I. Ingolfsson, Y. Li, V. V. Vostrikov, H. Gu, J. F. Hinton, R. E. Koeppe, 2nd, B. Roux & O. S. Andersen. Gramicidin A backbone and side chain dynamics evaluated by molecular dynamics simulations and nuclear magnetic resonance experiments. I: molecular dynamics simulations. The Journal of Physical Chemistry B 115, 7417-7426, (2011). -PMC ID: 3107394
217. W. Jiang, D. J. Hardy, J. C. Phillips, A. D. Mackerell, Jr., K. Schulten & B. Roux. High-performance scalable molecular dynamics simulations of a polarizable force field based on classical Drude oscillators in NAMD. The Journal of Physical Chemistry Letters 2, 87-92, (2011). -PMC ID: 3092300
218. A. Y. Lau & B. Roux. Structural biology: One domain, multiple conformations. Nature Chemical Biology 7, 130-131, (2011).
219. A. Y. Lau & B. Roux. The hidden energetics of ligand binding and activation in a glutamate receptor. Nature Structural & Molecular Biology 18, 283-287, (2011). -PMC ID: 3075596
220. A. C. Pan, L. G. Cuello, E. Perozo & B. Roux. Thermodynamic coupling between activation and inactivation gating in potassium channels revealed by free energy molecular dynamics simulations. The Journal of General Physiology 138, 571-580, (2011). -PMC ID: 3226968
221. B. Roux. Computational electrophysiology: the molecular dynamics of ion channel permeation and selectivity in atomistic detail. Biophysical Journal 101, 755-756, (2011). -PMC ID: 3175060
222. B. Roux, S. Berneche, B. Egwolf, B. Lev, S. Y. Noskov, C. N. Rowley & H. Yu. Ion selectivity in channels and transporters. The Journal of General Physiology 137, 415-426, (2011). -PMC ID: 3082929
223. W. Sandtner, B. Egwolf, F. Khalili-Araghi, J. E. Sanchez-Rodriguez, B. Roux, F. Bezanilla & M. Holmgren. Ouabain binding site in a functioning Na+/K+ ATPase. The Journal of Biological Chemistry 286, 38177-38183, (2011). -PMC ID: 3207469
224. E. Vargas, F. Bezanilla & B. Roux. In search of a consensus model of the resting state of a voltage-sensing domain. Neuron 72, 713-720, (2011).
225. V. V. Vostrikov, H. Gu, H. I. Ingolfsson, J. F. Hinton, O. S. Andersen, B. Roux & R. E. Koeppe, 2nd. Gramicidin A backbone and side chain dynamics evaluated by molecular dynamics simulations and nuclear magnetic resonance experiments. II: nuclear magnetic resonance experiments. The journal of Physical Chemistry B 115, 7427-7432, (2011). -PMC ID: 3114435
226. S. Yang & B. Roux. EROS: Better than SAXS! Structure 19, 3-4, (2011).
227. H. Yu, I. M. Ratheal, P. Artigas & B. Roux. Protonation of key acidic residues is critical for the K-selectivity of the Na/K pump. Nature Structural & Molecular Biology 18, 1159-1163, (2011). -PMC ID: 3190665
228. J. Gumbart, F. Khalili-Araghi, M. Sotomayor & B. Roux. Constant electric field simulations of the membrane potential illustrated with simple systems. Biochimica Biophysica Acta 1818, 294-302, (2012).
229. K. I. Lee, S. Jo, H. Rui, B. Egwolf, B. Roux, R. W. Pastor & W. Im. Web interface for brownian dynamics simulation of ion transport and its applications to beta-barrel pores. Journal of Computational Chemistry 33, 331-339, (2012). -PMC ID: 3240732

MCA01S026

230. Li, X., Kolorenc, J., and Mitas, L. Atomic fermi gas in the unitary limit by quantum Monte Carlo methods: Effects of the interaction range. Phys. Rev. A 84, 023615 (2011).
231. S. Bour, X. Li, H.-W. Hammer, D. Lee, U.-G. Meissner, L. Mitas, Precision benchmark calculations for four particles at unitarity. Phys. Rev. A 83, 063619 (2011).
232. Kolorenc, J. and Mitas, L., Applications of quantum Monte Carlo in condensed systems, Rep. Prog. Phys. 74, 026502 (2011) (review paper).
233. Ambrosetti, A., Silvestrelli, P., Toigo, F., Mitas, L., and Pederiva, F. Variational Monte Carlo for spin-orbit interacting systems. Phys. Rev. B, accepted, arXiv 1103.6209.
234. Ambrosetti, A., Silvestrelli, P., Toigo, F., Mitas, L., and Pederiva, F., Quantum Monte Carlo in the presence of spin-orbit interactions. In Advances in Quantum Monte Carlo, Tanaka, S. and Rothstein, S.,editors, American Chemical Society, (peer reviewed), accepted.
235. Hu, S., Rasch, K., and Mitas, L. Many-body nodal hypersurface and domain averages for correlated wave functions. In Advances in Quantum Monte Carlo, Tanaka, S. and Rothstein, S., American Chemical Society, (peer reviewed), accepted.
236. Horvathova, L., Dubecky, M., Mitas, L., and Stich, I., Spin multiplicity and symmetry breaking in vanadium-benzene complexes. Phys. Rev. Lett., submitted, (2011).
237. Rasch, K. M. and Mitas, L. Impact of electron density on the fixed-node errors in quantum Monte Carlo. Chem. Phys. Lett., submitted, arXiv: 1107.3520.

238. L. Horvathova, M. Dubecky, L. Mitas, and I. Stich, QMC study of energetics and structure of neutral and cationic vanadiumbenzene half-sandwich, submitted, (2011).
239. R. Derian, S. Guo, S. Hu, M. Zhu and L. Mitas, Dynamical spins in diffusion Monte Carlo method: spin-orbit operators, in preparation.

MCA03S007

240. 1. Z. Yu and S. Fan, "Angular constraint on light-trapping absorption enhancement in solar cells", Applied Physics Letters, vol. 98, art. No. 011106 (2011).
241. 2. Z. Ruan and S. Fan, "Design of subwavelength superscattering nanospheres", Applied Physics Letters, vol. 98, art. No. 043101 (2011). (Featured on cover)
242. 3. V. Liu, Y. Jiao, D. A. B. Miller and S. Fan, "Design of a compact photonic crystal based wavelength division multiplexer", Optics Letters, vol. 36, pp. 591-593(2011).
243. 4. H. Wu, L. Hu, T. Carney, Z. Ruan, D. Kong, Z. Yu, Y. Yao, J. Cha. J. Zhu, S. Fan and Y. Cui, "Low reflectivity and high flexibility ITO Nanofiber Transparent Electrodes", Journal of the American Chemical Society (Communications), vol. 133, pp. 27-29 (2011).
244. 5. Y. Huo, S. Sandhu, J. Pan, N. Stuhrmann, M. L. Povinelli, J. M. Kahn, J. S. Harris, M. M. Fejer, and S. Fan, "Experimental demonstration of two methods for controlling the group delay in a system with photonic-crystal resonators coupled to a waveguide", Optics Letters, vol. 36, pp. 1482-1484 (2011).
245. 6. Z. Yu and S. Fan, "Extraordinarily high spectral sensitivity in refractive index sensors using multiple optical modes", Optics Express, vol. 19, pp. 10029-10040 (2011).
246. 7. V. Intarprasonk, Z. Yu and S. Fan, "Image transfer with sub-wavelength resolution to metal-dielectric interface", Journal of Optical Society of America B vol. 28, pp. 1335-1338 (2011).
247. 8. A. Raman and S. Fan, "Perturbation theory for plasmonic modulation and sensing", Physical Review B, vol. 83, art. No. 205131 (2011).
248. 9. P. B. Catrysse and S. Fan, "Transverse electromagnetic modes in apertures containing a meta-material with extreme anisotropy", Physical Review Letters vol. 106, art. No. 223902 (2011).
249. 10. V. Intaraprasonk and S. Fan, "Non-volatile bistable all-optical switch from mechanical buckling", Applied Physics Letters, vol. 98, art. No. 241104 (2011).
250. 11. V. Liu and S. Fan, "Efficient computation of equi-frequency surfaces and density of states in photonic crystals using Dirichlet-to-Neumann maps", Journal of the Optical Society of America B. vol. 28, pp. 1837-1843 (2011).
251. 12. X. Huang and S. Fan, "Complete all-optical fiber isolator via stimulated Brillouin scattering", Journal of Lightwave Technology, vol. 29, pp. 2267-2275 (2011).
252. 13. K. Fang, Z. Yu and S. Fan, "Microscopic theory of photonic one-way edge mode", Physical Review B, vol. 84, art. No. 075477 (2011).
253. 14. A. K. Hafeli, E. Rephaeli, S. Fan, D. G. Cahill and T. E. Tiwald, "Temperature dependence of surface phonon polaritons from a quartz grating", Journal of Applied Physics, vol. 110, art. No. 043517 (2011).
254. 15. Z. Yu and S. Fan, "Optical isolation: a non-magnetic approach", Nature Photonics (News and Views), vol. 5, pp. 517-519 (2011).
255. 16. A. Raman, Z. Yu and S. Fan, "Dielectric nanostructures for broad-band light trapping in organic solar cells", Optics Express, vol. 19, pp. 19015-19026 (2011).
256. 17. K. Fang, Z. Yu, V. Liu and S. Fan, "Ultracompact non-reciprocal optical isolator based on guided resonance in a magneto-optical photonic crystal slab", Optics Letters, vol. 36, pp. 4254-4256 (2011).
257. 18. X. Yu, S. Sandhu, S. Beiker, R. Sassoon and S. Fan, "Wireless energy transfer with the presence of metallic planes", Applied Physics Letters, vol. 99, art. No.214202 (2011).
258. 19. T. Shi, S. Fan, and C. P. Sun, "Two-photon transport in a waveguide coupled to a cavity with a two-level system", Physical Review A, vol. 84, art. No. 063803 (2011).
259. 20. Z. Yu, A. Raman and S. Fan, "Nanophotonic light trapping theory for solar cells", Applied Physics A, vol. 105, pp. 329-339 (2011, Invited paper).
260. 21. C. Otey and S. Fan, "Exact microscopic theory of electromagnetic heat transfer between a dielectric sphere and a plate", Physical Review B vol. 84, art. No. 245431 (2011).
261. 22. H. Wen, M. Terrel, S. Fan, and M. J. F. Digonnet, "Sensing with slow light in fiber bragg gratings", IEEE Sensors Journal (accepted).
262. 23. E. Rephaeli, S. E. Kocabas and S. Fan, "Few photon transport in a waveguide coupled to a pair of collocated two-level atoms", Physical Review A (accepted).
263. 24. M. A. Terrel, M. Digonnet and S. Fan, "Resonant fiber optic gyroscope using an air-core fiber", Journal of Lightwave Technology (accepted).
264. 25. S. Fan, R. Baets, A. Petrov, Z. Yu, J. D. Joannopuolos, W. Freude, A. Melloni, M. Popovic, M. Vanwolleghem, D. Jalas, M. Eich, M. Krausse, H. Renner, E.Brinkmeyer, C. R. Doerr, "Comment on

‘Nonreciprocal light propagation in a silicon photonic circuit’, by Feng et al Science 333, 729, 2011”, Science (accepted)

MCA08X014

265. Chen, W., et al., Molecular dynamics simulations of forced unbending of integrin αVβ3. PLoS Comput Biol, 2011. **7**(2): p. e1001086.
266. Huang, Q., et al., Conformational Transition of Glycoprotein Ibalpha Mutants in Flow Molecular Dynamics Simulation. Cell Mol Bioeng, 2011. **4**(3): p. 495-504.
267. Xiang, X., et al., Structural basis and kinetics of force-induced conformational changes of an alphaA domain-containing integrin. PLoS One, 2011. **6**(11): p. e27946.

MCA08X013

268. C. Dawson, E.J. Kubatko, J.J. Westerink, C. Trahan, C. Mirabito, C. Michoski, N. Panda, Discontinuous Galerkin methods for modeling hurricane storm surge, Advances in Water Resources, doi:10.1016/j.advwatres.2010.11.004, 34, pp. 1165-1176, 2011.
269. S. Tanaka, S. Bunya, J.J. Westerink, C. Dawson, and R.A. Luettich, Scalability of an unstructured mesh continuous Galerkin based hurricane storm surge model, Journal of Scientific Computing, 46, pp. 329-358, 2011.
270. J.C. Dietrich, J.J. Westerink, A.B. Kennedy, J.M. Smith, R.E. Jensen, M. Zijlema, L.H. Holthuijsen, C. Dawson, R.A. Luettich, Jr., M.D. Powell, V.J. Cardone, A.T. Cox, G.W. Stone, H. Pourtaheri, M.E. Hope, S. Tanaka, L.G. Westerink, H.J. Westerink, Z. Co-bell, Hurricane Gustav (2008) Waves and Storm Surge: Hindcast, Synoptic Analysis and Validation in Southern Louisiana, Monthly Weather Review, 139, pp. 2488-2522, 2011.
271. J.C. Dietrich, S. Tanaka, J.J. Westerink, C.N. Dawson, R.A. Luettich Jr., M. Zijlema, L.H. Holthuijsen, J.M. Smith, L.G.Westerink, H.J.Westerink, Performance of the Unstructured-Mesh, SWAN+ADCIRC Model in Computing Hurricane Waves and Surge, Journal of Scienti_c Computing, to appear.
272. S. Mattis, C. Dawson, C.E. Kees and M.W. Farthing, Numerical modeling of drag for flow through vegetated domains and porous structures, Advances in Water Resources, to appear.

MCA08X004

273. [P1] J. Luitjens and M. Berzins., Parallel regridding algorithms for block-structured adaptive mesh refinement. Concurrency and Computation: Practice and Experience, Vol. 23, No. 13, John Wiley & Sons, Ltd., pp. 1522–1537. 2011.
274. [P2] M. Berzins, Q.Meng, J.Schmidt, J. Sutherland., DAG-Based Software Frameworks for PDEs, In Proceedings of HPSS 2011, Europar Bordeaux, Springer, (to appear), August, 2011.
275. [P3] Q. Meng, M. Berzins, J. Schmidt., Using Hybrid Parallelism to improve memory use in Uintah, In Proceedings of the Teragrid 2011 Conference, Salt Lake City, Utah, ACM, July, 2011.
276. [P4] L.T. Tran, M. Berzins., IMPICE Method for Compressible Flow Problems in Uintah, International Journal For Numerical Methods In Fluids, Note: Published online 20 July, 2011.
277. [P5] J. van Rij, T. Harman, T. Ameel., Model Development for Fluid Structure Interaction in the Slip Regime, in 8th International Conference on Nanochannels, Microchannels and Minichannels., 2011, pp. 1055-1064.
278. [P6] J. van Rij, T. Harman, and T. Ameel., Slip Flow Fluid-Structure Interaction, International Journal of Heat and Mass Transfer, in press 2012.
279. [P7] J. Peterson and C.Wight., An Eulerian-Lagrangian Computational Model for Deflagration to Detonation Transition, (accepted by Combustion and Flame Fall 2011).

MCA08X003

280. 1. Olsen BN, Schlesinger PH, Ory DS, Baker NA. Side-chain oxysterols: From cells to membranes to molecules, *Biochimica et Biophysica Acta*, 1818 (2), 330-6, 2012.
281. 2. Olsen BN, Schlesinger PH, Ory DS, Baker NA. 25-Hydroxycholesterol Increases the Availability of Cholesterol in Phospholipid Membranes, *Biophysical Journal*, 100 (4), 948-956, 2011.
282. 4. Lee SJ, Schlesinger PH, Wickline SA, Lanza GM, and Baker NA. Interaction of melittin peptides with perfluorocarbon nanoemulsion particles. *J. Phys. Chem. B* (2011) 115 (51), 15271-15279
283. 5. Lee SJ, Schlesinger PH, Wickline SA, Lanza GM, and Baker NA. Simulation of fusion-mediated nanoemulsion interactions with model lipid bilayers. Submitted to *ACS Nano*.

MCA08X002

284. 7. Wang L.; Berne B. J.; Friesner R. A.; Ligand binding to protein-binding pockets with wet and dry regions. Proc. Nat. Acad. Sci. USA. 108, 4, 1326-1330 (2011).
285. 8. Abel, R., Salam, N. K., Shelley, J., Farid, R., Friesner, R. A. and Sherman, W. (2011), Contribution of Explicit Solvent Effects to the Binding Affinity of Small-Molecule Inhibitors in Blood Coagulation Factor Serine Proteases. ChemMedChem, 6: 1049–1066.
286. 9. Wang L, Friesner RA, Berne BJ (2011) Replica exchange with solute scaling: A more e_cient version of replica exchange with solute tempering (rest2). J. Phys. Chem. B 115:9431-9438. 10. Wang L, Berne BJ, Friesner RA (2011) On achieving high accuracy and reliability in the calculation of relative protein ligand binding affinities. Proc. Nat. Acad. Sci. USA. 2011 (In Press).

MCA06N007

287. Cerove_cki, I., L. D. Talley, and M. R. Mazlo_, 2011: A comparison of southern ocean air-sea buoyancy ux from an ocean state estimate with _ve other products. Journal of Climate, doi:10.1175/2011JCLI3858.1.
288. Firing, Y. L., T. K. Chereskin, and M. Mazlo_, 2011: Vertical structure and transport of the Antarctic Circumpolar Current in Drake Passage from direct velocity observations. Journal of Geophysical Research, 116.
289. Griesel, A., M. Mazlo_, and S. T. Gille, 2011: Mean Dynamic Topography in the Southern Ocean: Evaluating Antarctic Circumpolar Current Transport. Journal of Geophysical Research, in press.
290. Mazlo_, M., 2011: On the sensitivity of the Drake Passage transport to air-sea momentum ux. Journal of Climate, In press.
291. Roquet, F., C. Wunsch, and G. Madec, 2011: On the patterns of wind-power input to the ocean circulation. Journal of Physical Oceanography, doi:10.1175/JPO-D-11-024.1.
292. Sallee, J. and S. Rintoul, 2011: Parameterization of eddy-induced subduction in the Southern Ocean surface-layer. Ocean Modelling, 39 (1-2), 146{153.
293. Woloszyn, M., M. Mazlo_, and T. Ito, 2011: Testing an eddy-permitting model of the southern ocean carbon cycle against observations. Ocean Modelling, 39 (1-2), 170{182}.

MCA06N038

294. Mukhopadhyay, S. (2011) *Ignition and Early Flame Development in Stratified-Charge Mixtures*, Ph.D. Dissertation, Purdue University, West Lafayette, IN.
295. 2. Reddy, H. (2011) *Numerical Studies of Turbulence Effects on Developing Flames in Lean Methane/Air Mixtures*, Ph.D. Dissertation, Purdue University, West Lafayette, IN.
296. Mukhopadhyay, S. and Abraham, J. (2011) Influence of Compositional Stratification on Autoignition in n-Heptane/Air Mixtures, *Combustion and Flame*, 158(6):1064-1075.
297. Reddy, H. and Abraham, J. (2011) A Numerical Study of Vortex Interactions with Flames Developing from Ignition Kernels in Lean Methane/Air Mixtures, *Combustion and Flame*, 158(3):401-415.
298. Owston, R. and Abraham, J. (2011) Exploratory Studies of Modeling Approaches for Hydrogen Triple Flames, *International Journal of Hydrogen Energy*, 36(14):8570-8582.
299. Sayeed, M., Magi, V., and Abraham, J. (2011) On the Efficiency of a High-Performance Parallel Solver for Turbulent Reacting Flow Simulations, *Numerical Methods in Heat Transfer, Part B*, 59(3):169-189.
300. Reddy, H. and Abraham, J. (2010) Ignition kernel development studies relevant to lean-burn natural gas engines, Fuel, 89(11):3262-3271.
301. Mukhopadhyay, S. and Abraham, J., Direct Numerical Simulations of n- Heptane/Air Autoignition in Stratified Mixtures, *Combustion Science and Technology*, submitted January 2012.
302. Mukhopadyay, S. and Abraham, J., Influence of Heat Release and Turbulence on Scalar Dissipation Rate in Autoigniting n-Heptane/Air Mixtures, *Combustion and Flame*, revision submitted December 2011.
303. Mukhopadhyay, S. and Abraham, J., Evaluation of an Unsteady Flamelet Progress Variable Model For Autoignition and Flame Development in Compositionally Stratified Mixtures, *Physics of Fluids*, revision submitted January 2012.
304. Reddy, H. and Abraham, J., Influence of Turbulence-Kernel Interactions on Flame Development in Lean Methane/Air Mixtures under Engine Conditions, *Combustion and Flame*, revision submitted December 2011.
305. Reddy, H. and Abraham, J., Evaluation of the Flame-Surface-Density Model for Turbulent Flames Developing from an Ignition Kernel in Lean Methane/Air Mixtures under Engine Conditions, *Physics of Fluids*, submitted October 2011.
306. Reddy, H. and Abraham, J., Transient Evolution of Strain and Curvature in Flames Developing from Kernels, *Proceedings of the Combustion Institute*, submitted December 2011.

MCA07S005

307. K. Yang, W. Setyawan, S. Wang, M. Buongiorno-Nardelli, and S. Curtarolo, Where are the missing topological-insulators? The search with high-throughput descriptors, in review, Nature (2012).
308. S. Curtarolo, W. Setyawan, S. Wang, J. Xue, K. Yang, R. H. Taylor, G. L. W. Hart, S. Sanvito, M. Buongiorno-Nardelli, N. Mingo, and O. Levy, AFLOWLIB.ORG: a distributed materials properties repository from high-throughput ab initio calculations, submitted to Comp. Mat. Sci. (2012).
309. S. Curtarolo, W. Setyawan, G. L. W. Hart, M. Jahnatek, R. V. Chepulskii, R. H. Taylor, S. Wang, J. Xue, K. Yang, O. Levy, M. Mehl, H. T. Stokes, D. O. Demchenko, and D. Morgan, AFLOW: an automatic framework for high-throughput materials discovery, submitted to Comp. Mat. Sci. (2012).
310. O. Levy, J. Xue, S. Wang, G. L. W. Hart, and S. Curtarolo, Uncovering technetium binary ordered structures from first principles, in press, Phys. Rev. B (2012).
311. S. Wang, Z. Wang, W. Setyawan, N. Mingo, and S. Curtarolo, Assessing the thermoelectric properties of sintered compounds with high-throughput ab-initio calculations, Phys. Rev. X 1, 021012 (2011). Note. This paper was featured in Nature Materials “Research Highlights” for Volume 11(1). Research Highlights: F. Pulizzi, Thermopowerful guidelines, Nature Materials 11, 8 (2012).
312. M. Jahnatek, O. Levy, G. L. W. Hart, L. J. Nelson, R. V. Chepulskii, J. Xue, and S. Curtarolo, Ordered phases in ruthenium binary alloys from high-throughput first principles calculations, Phys. Rev. B 84, 214110 (2011).
313. R. V. Chepulskii and S. Curtarolo, Revealing low-temperature atomic ordering in bulk Co-Pt with the high-throughput ab initio method, Appl. Phys. Lett. 99, 261902 (2011).
314. W. Setyawan, R. M. Gaum´e, S. Lam, R. S. Feigelson, and S. Curtarolo, High-Throughput Combinatorial Database of Electronic Band Structures for Inorganic Scintillator Materials, ACS Comb. Sci. 13(4), 382-390 (2011).
315. O. Levy, M. Jahnatek, R. V. Chepulskii, G. L. W. Hart, and S. Curtarolo, Ordered Structures in Rhenium Binary Alloys from First-Principles Calculations, J. Am. Chem. Soc. 133(1), 158-163 (2011).
316. R. V. Chepulskii and S. Curtarolo, Ab Initio Insights on the Shapes of Platinum Nanocatalysts, ACS-NANO 5(1), 247-254 (2011).
317. R. H. Taylor, S. Curtarolo, and G. L. W. Hart, Guiding the experimental discovery of magnesium alloys, Phys. Rev. B 84, 084101 (2011).

MCA07S017

318. Improved stochastic estimation of quark propagation with Laplacian Heaviside smearing in lattice QCD, C. Morningstar, J. Bulava, J. Foley, K.J. Juge, D. Lenkner, M. Peardon, C.H. Wong, Phys. Rev. D 83, 114505 (2011).
319. Excited-state hadron masses using the stochastic LapH method, J. Bulava, J. Foley, Y.C. Jhang, K.J. Juge, D. Lenkner, C. Morningstar, C.H. Wong, PoS(Lattice 2011):131 (2011), proceedings of the XXIX International Symposium on Lattice Field Theory, July 10-16 2011, Squaw Valley, Lake Tahoe, CA, USA, (arXiv:1111.0845 [hep-lat])
320. Progress on Excited Hadrons in Lattice QCD, J. Bulava, PoS(Lattice 2011):021 (2011), proceedings of the XXIX International Symposium on Lattice Field Theory, July 10-16 2011, Squaw Valley, Lake Tahoe, CA, USA, (arXiv:1112.0212 [hep-lat])
321. Group-theoretical construction of finite-momentum and multi-particle operators for lattice hadron spectroscopy, J. Foley et al., PoS(Lattice 2011)120, submitted to the proceedings of Lattice 2011.
322. A Novel algorithm for computing quark propagators for lattice hadron spectroscopy, J. Bulava, J. Foley, Y.C. Jhang, D. Lenkner, C. Morningstar, C.H. Wong, K.J. Juge, M. Peardon, proceedings of the 2011 Meeting of the APS Division of Particles and Fields (DPF-2011), Brown University, Providence, RI, Aug. 9-13, 2011, (arXiv:1109.6614 [hep-lat])
323. Excited-State Hadron Masses from Lattice QCD, C. Morningstar, J. Bulava, J. Foley, Y.C. Jhang, K.J. Juge, D. Lenkner, C.H. Wong, proceedings of the 19th Particles and Nuclei International Conference (PANIC 2011), Massachusetts Institute of Technology, Cambridge, MA, USA, July 24-29, 2011, (arXiv:1109.0308 [hep-lat])
324. Progress Report on Computing Excited-State Hadron Masses in Lattice QCD, C. Morningstar, A. Bell, J. Bulava, J. Foley, K.J. Juge, D. Lenkner, C.H. Wong, AIP Conf. Proc. 1388, 34 (2011), proceedings of Baryons 2010: International Conference on the Structure of Baryons, Osaka, Japan, Dec. 7-11, 2010, (arXiv:1103.2783 [hep-lat])
325. Excited-State Hadrons using the Stochastic LapH Method, C. Morningstar, A. Bell, J. Bulava, J. Foley, K.J. Juge, D. Lenkner, C.H. Wong, AIP Conf. Proc. 1354, 97 (2011), proceedings of the T(r)opical QCD Workshop, Cairns, Australia, Oct. 2010, (arXiv:1101.5398 [hep-lat])

326. Nucleon, Δ and Ω excited states in Nf = 2 + 1 lattice QCD, J. Bulava, R.G. Edwards, E. Engelson, B. Joo, H-W. Lin, C. Morningstar, D.G. Richards, S.J. Wallace, Phys. Rev. D82, 014507 (2010).

MCA08X032

327. L. PETRIDIS, J. XU, M. CROWLEY, J.C. SMITH and X. CHENG. Atomistic Simulationof Lignocellulosic Biomass and Associated Cellulosomal Protein Complexes Biofuels: ACS. (2011)
328. B. COLLIGNON, R. SCHULZ, J.C. SMITH and J.BAUDRY. Task-Parallel MPI Implementation of Autodock4 for Docking of Very Large Databases of Compounds using High Performance Super-Computers. Journal of Computational Chemistry. 32: 1202–1209 (2011). ORNL press release
329. Y. MIAO and J. BAUDRY Active-Site Hydration and Water Diffusion in Cytochrome P450cam: A Highly Dynamic Process. Biophysical Journal 101(6):1493-1503 (2011)
330. M. ZAHRAN, T. BEREZNIAK, P. IMHOF, J.C. SMITH. Role of Magnesium Ions in DNA Recognition by the EcoRV Restriction Endonuclease. FEBS Letters 585(17):2739-2743 (2011).
331. B. MOSTOFIAN, J.C. SMITH, X. CHENG. The Solvation Structures of Cellulose Microfibrils in Ionic Liquids. Interdisciplinary Sciences - Computational Life Sciences, in press
332. H.B. GUO, A. JOHS, J.M. PARKS, J.C. SMITH. Mercury Detoxification by Bacteria: Simulations of Transcription Activation and Mercury-Carbon Bond Cleavage, in Modeling of Molecular Properties, Wiley-VCH, in Press.
333. M. BISWAS, K. VOLTZ, J.C. SMITH, J. LANGOWSKI . Role of Histone Tails in Structural Stability of the Nucleosome. PLoS Computational Biolology 7(12): e1002279. doi:10.1371/journal.pcbi.1002279 (2011).
334. B.J. MINTZ, J.M. PARKS. Benchmark Interaction Energies for biologically relevant noncovalent complexes containing divalent sulfur. Journal of Physical Chemistry A, doi 10.1021/jp209536e

MCA08X034

335. Differential cross sections for electron impact excitation of the n = 2 states of helium at intermediate energies (80, 100 and 120 eV) measured across the complete angular scattering range (0° -180°) R. Ward, D. Cubric, N. Bowring, G.C. King, F.H. Read, D.V. Fursa, I. Bray, O. Zatsarinny, and K. Bartschat, J. Phys. B 44 (2011) 045209
336. Benchmark calculations of atomic data for plasma and lighting applications K. Bartschat and O. Zatsarinny, Plasma Sources Science and Technology 20 (2011) 024012
337. Electron impact excitation of the (4p55s) states in krypton: High-resolution electron scattering experiments and B-spline R-matrix calculations, M. Allan, O. Zatsarinny, and K. Bartschat, J. Phys. B 44 (2011) 065201
338. Breakup of the aligned H□2 molecule by xuv laser pulses: A time-dependent treatment in prolate spheroidal coordinates, X. Guan, K. Bartschat, and B.I. Schneider, Phys. Rev. A 83 (2011) 043403
339. High-resolution experiments and B-spline R-matrix calculations for elastic electron scattering from krypton O. Zatsarinny, K. Bartschat, and M. Allan, Phys. Rev. A 83 (2011) 032713
340. Unexpected e_ects in spin-polarized electron impact excitation of the (3d104s5s)3S1 state in zinc L. Pravica, J.F. Williams, D. Cvejanovi_c, S. Samarin, K. Bartschat, O. Zatsarinny, A.D. Stau_er, and R. Srivastava, Phys. Rev. A 83 (2011) 040701
341. Electron collision cross sections for iodine O. Zatsarinny, K. Bartschat, G. Garcia, F. Blanco, L.R. Hargreaves, D.B. Jones, R. Murrie, J.R. Brunton, M.J. Brunger, M. Hoshino, and S.J. Buckman, Phys. Rev. A 83 (2011) 042702
342. Nonperturbative treatment of ionization and excitation of helium by electron impact O. Zatsarinny and K. Bartschat, Phys. Rev. Lett. 107 (2011) 023203
343. Alignment effects in two-photon double ionization of H2 in femtosecond xuv laser pulses X. Guan, K. Bartschat, and B.I. Schneider, Phys. Rev. A 84 (2011) 033403 Multi-photon ionization of H+2 in laser pulses X. Guan, E. Secor, K. Bartschat, and B.I. Schneider, Phys. Rev. A 84 (2011) 033420 16
344. Time-resolved photoemission by attosecond streaking: extraction of time information S. Nagele, R. Pazourek, J. Feist, K. Doblho_-Dier, C. Lemell, K. T}ok_esi, J. Burgd□orfer, J. Phys. B 44, 081001 (2011) (Fast Track Communication).
345. Universal features in sequential and nonsequential two-photon double ionization of helium R. Pazourek, J. Feist, S. Nagele, E. Persson, B. I. Schneider, L. A. Collins, J. Burgd□orfer, Phys. Rev. A 83, 053418 (2011).
346. Attosecond Two-Photon Interferometry for Doubly Excited States of Helium J. Feist, S. Nagele, C. Ticknor, B. I. Schneider, L. A. Collins, and J. Burgd□orfer, Phys. Rev. Lett. 107, 093005 (2011) Probing scattering phase shifts by attosecond streaking
347. R. Pazourek, S. Nagele, K. Doblho_-Dier, J. Feist, C. Lemell, K. T}ok_esi, J. Burgd□orfer, submitted to Journal of Physics: Conference Series (2011), arXiv:1111.4172.

348. Attosecond streaking of correlated two-electron transitions in helium R. Pazourek, S. Nagele, J. Feist, and J. Burgd□orfer, submitted to Phys. Rev. Lett. (2011), arXiv: 1112.3273.
349. Time shifts in photoemission from a fully-correlated two-electron model system S. Nagele, R. Pazourek, J. Feist, and J. Burgd□orfer, submitted to Phys. Rev. A. (2011) Refereed Book Contributions
350. Multi-photon single and double ionization of complex atoms by ultrashort intense laser pulses K. Bartschat, X. Guan, C.J. Noble, B.I. Schneider, and O. Zatsarinny in: Quantum Dynamic Imaging, A.D. Bandrauk and M. Ivanov (Eds.), Springer (2011)
351. Recent Advances in Computational Methods for the Solution of the Time-Dependent Schrdinger Equation for the Interaction of Short, Intense Radiation with One and Two-Electron Systems: Application to He and H+2
352. B. I. Schneider, J. Feist, S. Nagele, R. Pazourek, S. X. Hu, L. A. Collins, J. Burgd□orfer in: Quantum Dynamic Imaging, A.D. Bandrauk and M. Ivanov (Eds.), Springer (2011) Invited Talks at International Conferences
353. K. Bartschat, Benchmark calculations of atomic data for plasma and discharge applications 16th Gaseous Electronics Meeting, January 31□February 3, 2010, in Batemans Bay, New South Wales (Australia)
354. J. Feist, Ab initio calculations of correlated electron dynamics in ultrashort pulses March Meeting of the American Physical Society, March 15□19, 2010, Portland (Oregon, USA)
355. J. Burgdorfer, Probing and controlling electronic dynamics on the attosecond scale 41th Meeting of the Division for Atomic, Molecular and Optical Physics (DAMOP) of the American Physical Society (APS), May 25 - 29, 2010, Houston (Texas, USA)
356. J. Burgdorfer, Probing few-body dynamics in atoms by photons and neutrons International Focus Workshop Few Body Dynamics in AMPS, Max Planck Institute for the Physics of Complex Systems, June 27□July 1, 2010, Dresden (Germany)
357. J. Burgdorfer, Interactions of ions and atoms with surfaces, clusters, and capillaries Plenary Talk 10th European Conference on Atoms, Molecules and Photons, July 4□9, 2010, Salamanca (Spain)
358. K. Bartschat, Benchmark calculations of atomic data for plasma and lighting applications 20th European Conference on Atomic and Molecular Processes in Ionized Gases, July 14□17, 2010, in Novi Sad (Serbia)
359. O. Zatsarinny, Cross section calculations for low energy electron-atom collisions Seventh International Conference on Atomic and Molecular Data and Their Applications, September 21□24, 2010, Vilnius (Lithuania)
360. J. Burgd□orfer, Attosecond physics Symposium in Honor of Siegfried Grossmann, Centre for Applied Mathematics and Theoretical Physics, Dec. 9 -11, 2010, Maribor (Slovenia)
361. K. Bartschat, Benchmark calculations of atomic data for plasma applications 4th International Conference on Plasma Nanotechnology & Science, March 10 - 12, 2011, in Takayama (Japan)
362. J. Burgd□orfer, Attosecond Physics: time-dependent electronic dynamics in atoms, molecules, and solids Focus Session on Ultrafast Dynamics, APS March Meeting 2011, March 21□25, 2011, Dallas (Texas, USA)
363. J. Burgd□orfer, Ultrashort pulses: quantum dynamics in the time-domain Lecture at Graduate School on New Coherent Sources, University of Hamburg, May 10, 2011, Hamburg (Germany)
364. J. Burgd□orfer, Time-resolved photoemission by attosecond streaking: extraction of the Wigner-Smith time delay Workshop on Attosecond Science-Exploring and Controlling Matter on its Natural Time Scale, KAVLI Institute for Theoretical Physics China (KITPC), May 18, 2011, Beijing (China)
365. J. Burgd□orfer, Linear and non-linear Schr□odinger equation: where is chaos? 8th International Summerschool Let's face chaos through nonlinear dynamics June 27 - 30, 2011, Maribor (Slovenia)
366. J. Burgd□orfer (talk given by S. Nagele), Time-resolved photoemission by attosecond streaking 12th International Conference on Multi-Photon Processes, July 3 - 6, 2011, Sapporo (Japan)

MCA93SO13

367. Markwick, P.R.L., L.C.T. Pierce, D.B. Goodin, J.A. McCammon. Adaptive accelerated molecular dynamics (Ad-AMD) reveals the molecular plasticity of P450cam. J. Phys. Chem. Letters 2, 158–164 (2011).
368. Wang, Y., P.R.L. Markwick, C.A.F. de Oliveira, J.A. McCammon. Enhanced lipid diffusion and mixing in accelerated molecular dynamics. J. Chem. Theory Comp. 7, 3199-3207 (2011).
369. Wang, Y., C. Harrison, K. Schulten, J.A. McCammon. Implementation of Accelerated Molecular Dynamics in NAMD. Comp. Sci. Discov. 4, article 015002 (10pp) (2011).
370. Lawrenz, M., R. Baron, Y, Wang, J.A. McCammon. Effects of Biomolecular Flexibility on Alchemical Calculations of Absolute Binding Free Energies. J. Chem. Theory Comp. 7, 2224–2232 (2011).

371. Baron, R., C. Binda, J.A. McCammon, A. Mattevi. Molecular Mimicry and Ligand Recognition in Binding and Catalysis by the Histone Demethylase LSD1-CoREST Complex. Structure 19, 212-220 (2011).
372. Grant, B.J., D. Gheorghe, W. Zheng, M. Alonsoc, G. Huber, M. Dlugosz, J.A. McCammon, R.A. Cross. Electrostatically Biased Binding of Kinesin to Microtubules. PLoS Biology 9, article e1001207 (2011).

MCA94P014 (could not paste pubs list from PDF)

373. H. Gao, B. Zhang, S. G. Johnson, and G. Barbastathis, “Design of thin–film photonic metamaterial Lüneburg lens using analytical approach,” *Optics Express,* vol. 20, pp. 1617–1628, January 2012.
374. A. W. Rodriguez, O. Ilic, P. Bermel, I. Celanovic, J. D. Joannopoulos, M. Soljačić, and S. G. Johnson, “Frequency- selective near- field radiative heat transfer between photonic crystal slabs: A computational approach for arbitrary geometries and materials,” *Physical Review Letters,* vol. 107, p. 114302, September 2011.
375. D. S. Deng, J.- C. Nave, X. Liang, S. G. Johnson, and Y. Fink, “Exploration of in- fiber nanostructures from capillary instability,” *Optics Express,* vol. 17, pp. 16273–16290, August 2011.
376. H. Hashemi, A. Oskooi, J. D. Joannopoulos, and S. G. Johnson, “General scaling limitations of ground- plane and isolated- object cloaks,” *Physical Review A,* vol. 84, p. 023815, August 2011.
377. M. Levin and S. G. Johnson, “Is the electrostatic force between a point charge and a neutral metallic object always attractive?,” *American Journal of Physics,* vol. 79, pp. 843– 849, August 2011.
378. X. Liang, D. S. Deng, J.- C. Nave, and S. G. Johnson, “Linear stability analysis of capillary instabilities for concentric cylindrical shells,” *Journal of Fluid Mechanics,* vol. 683, pp. 235–262, August 2011.
379. M. T. H. Reid, J. White, and S. G. Johnson, “Computation of Casimir interactions between arbitrary three- dimensional objects with arbitrary material properties,” *Physical Review A,* vol. 84, p. 010503(R), July 2011.
380. X. Sheng, S. G. Johnson, J. Michel, and L. C. Kimerling, “Optimization- based design of surface textures for thin- film Si solar cells,” *Optics Express,* vol. 19, pp. A841–A850, June 2011.□A. Kurs, J. D. Joannopoulous, M. Soljačić, and S. G. Johnson, “Abrupt coupling between strongly dissimilar waveguides with 100% transmission,” *Optics Express,* vol. 19, pp. 13714–13721, June 2011.
381. S. G. Johnson, “Numerical methods for computing Casimir interactions,” in *Casimir Physics* (D. Dalvit, P. Milonni, D. Roberts, and F. da Rosa, eds.), vol. 834 of *Lecture Notes in Physics,* ch. 6, pp. 175–218, Berlin: Springer, June 2011.
382. A. P. McCauley, F. S. S. Rosa, A. W. Rodriguez, J. D. Joannopoulos, D. A. R. Dalvit, and S. G. Johnson, “Structural anisotropy and orientation- induced Casimir repulsion in fluids,” *Physical Review A,* vol. 83, p. 052503, May 2011.
383. W. Rodriguez, D. Woolf, P.- C. Hui, E. Iwase, A. P. McCauley, F. Capasso, M. Loncar, , and S. G. Johnson, “Designing evanescent optical interactions to control the expression of Casimir forces in optomechanical structures,” *Applied Physics Letters,* vol. 98, p. 194105, May 2011. Invited paper in May 23, 2011, issue of the *Virtual Journal of Nanoscale Science & Technology.*
384. Oskooi and S. G. Johnson, “Distinguishing correct from incorrect PML proposals and a corrected unsplit PML for anisotropic, dispersive media,” *Journal of Computational Physics,* vol. 230, pp. 2369–2377, April 2011.
385. J. Varela, A. W. Rodriguez, A. P. McCauley, and S. G. Johnson, “Casimir micro- sphere diclusters and three- body effects in fluids,” *Physical Review A,* vol. 83, p. 042516, April 2011.
386. K. Pan, A. P. McCauley, A. W. Rodriguez, M. T. H. Reid, J. K. White, and S. G. Johnson, “Calculation of nonzero- temperature Casimir forces in the time domain,” *Physical Review A,* vol. 83, p. 040503(R), April 2011.
387. A. Kurs, M. Kesler, and S. G. Johnson, “Optimized design of a low- resistance electrical conductor for the multimegahertz range,” *Applied Physics Letters,* vol. 98, p. 172504, April 2011.
388. D. Ramirez, A. W. Rodriguez, H. Hashemi, J. Joannopoulos, M. Soljačić, and S. G. Johnson, “Degenerate four- wave mixing in triply- resonant Kerr cavities,” *Physical Review A,* vol. 83, p. 033834, March 2011.
389. A. W. Rodriguez, F. Capasso, and S. G. Johnson, “The Casimir effect in microstructured geometries,” *Nature Photonics,* vol. 5, pp. 211–221, March 2011. Invited review.
390. K. Y. K. Lee, M. K. Nyein, D. F. Moore, J. D. Joannopoulos, S. Socrate, T. Imholt, R. Radovitzky, and S. G. Johnson, “Blast- induced electromagnetic fields in the brain from bone piezoelectricity,” *NeuroImage,* vol. 54, pp. S30–S36, January 2011. Invited paper, published online June 2010.
391. P. Bermel, M. Ghebrebrhan, M. Harradon, Y. X. Yeng, I. Celanovic, J. D. Joannopoulos, and M. Soljačić, “Tailoring photonic metamaterial resonances for thermal radiation,” *Nanoscale Research Letters,* vol. 6, p. 549, October 2011.
392. S.- L. Chua, C. A. Caccamise, D. J. Phillips, J. D. Joannopoulos, M. Soljačić, H. O. Everitt, and J. Bravo- Abad, “Spatio- temporal theory of lasing action in optically- pumped rotationally excited molecular gases,” *Optics Express,* vol. 19, pp. 7513–7529, April 2011.

393. M. Ghebrebrhan, P. Bermel, Y. X. Yeng, I. Celanovic, M. Soljačić, and J. D. Joannopoulos, "Tailoring thermal emission via *q* matching of photonic crystal resonances," *Physical Review A,* vol. 83, p. 033810, March 2011.
394. S. Danto, Z. Ruff, Z. Wang, J. D. Joannopoulos, and Y. Fink, "Ovonic memory switching in multimaterial fibers," *Advanced Functional Materials,* vol. 21, pp. 1095–1101, March 2011.
395. D. Chester, P. Bermel, J. D. Joannopoulos, M. Soljačić, and I. Celanovic, "Design and global optimization of high- efficiency solar thermal systems with tungsten cermets," *Optics Express,* vol. 19, pp. A245–A257, March 2011.
396. Mohammad Araghchini, Yi Xiang Yeng, Natalia Jovanovic, Peter Bermel, Leslie A. Kolodziejski, Marin Soljacic, Ivan Celanovic, John D. Joannopoulos, "Fabrication of two- dimensional tungsten photonic crystals for high- temperature applications", *Journal of Vacuum Science and Technology* **B 29,** 061402 (2011).
397. Wenjun Qiu, Zheng Wang, and Marin Soljacic, "Broadband circulators based on directional coupling of one- way waveguides," *Invited Paper in a Focus Issue of Optics Express* Vol.19, No.22, p. 22248, (2011).
398. Song- Liang Chua, Yidong Chong, A. Douglas Stone, Marin Soljacic, and Jorge Bravo- Abad, "Low- threshold lasing action in photonic- crystal slabs enabled by Fano resonances," *Optics Express* Vol.19, No.2, p. 1539, (2011).

MCA94P018

399. J. Bellovary, M. Volonteri, F. Governato, S. Shen, T. Quinn, and J. Wadsley. The First Massive Black Hole Seeds and Their Hosts. Ap. J., 742:13, November 2011.
400. J. M. Bellovary, F. Governato, T. Quinn, and M. Volonteri. The Formation and Evolution of Massive Black Holes in Cosmological Simulations. In American Astronomical Society Meeting Abstracts #217, volume 43 of Bulletin of the American Astronomical Society, page #128.01, January 2011.
401. C. B. Brook, F. Governato, R. Roˇskar, G. Stinson, A. M. Brooks, J. Wadsley, T. Quinn, B. K. Gibson, O. Snaith, K. Pilkington, E. House, and A. Pontzen. Hierarchical formation of bulgeless galaxies: why outflows have low angular momentum. MNRAS, 415:1051–1060, August 2011.
402. A. M. Brooks, A. R. Solomon, F. Governato, J. McCleary, L. A. MacArthur, C. B. A. Brook, P. Jonsson, T. R. Quinn, and J. Wadsley. Interpreting the Evolution of the Size-Luminosity Relation for Disk Galaxies from Redshift 1 to the Present. Ap. J., 728:51–+, February 2011.
403. C. Christensen, T. Quinn, and F. Governato. Molecular Hydrogen in Cosmological Simulations of Galaxies. In American Astronomical Society Meeting Abstracts #217, volume 43 of Bulletin of the American Astronomical Society, page 216.02, January 2011.
404. C. R. Christensen. The Connection Between the Interstellar Media and Star Formation in Simulations of Evolving Galaxies with a Focus on Molecular Hydrogen. PhD thesis, University of Washington, 2011.
405. N. A. Kaib, T. Quinn, and R. Brasser. Decreasing Computing Time with Symplectic Correctors in Adaptive Timestepping Routines. A. J., 141:3, January 2011.
406. N. A. Kaib, R. Roˇskar, and T. Quinn. Sedna and the Oort Cloud around a migrating Sun. Icarus, 215:491–507, October 2011.
407. S. R. Loebman, R. Roˇskar, V. P. Debattista, ˇZ. Ivezi´c, T. R. Quinn, and J. Wadsley. The Genesis of the Milky Way's Thick Disk Via Stellar Migration. Ap. J., 737:8, August 2011.
408. F. Munshi, C. Christensen, F. Governato, T. Quinn, and J. Wadsley. Testing the Effectiveness of the H-alpha/UV flux Ratio in Determining Star Formation Rates in Dwarf Galaxies. In American Astronomical Society Meeting Abstracts #217, volume 43 of Bulletin of the American Astronomical Society, page #328.05, January 2011.
409. R. Roˇskar, V. P. Debattista, S. R. Loebman, ˇZ. Ivezi´c, and T. R. Quinn. Implications of Radial Migration for Stellar Population Studies. ArXiv e-prints, January 2011.
410. R. Roˇskar, V. P. Debattista, T. R. Quinn, and J. Wadsley. Radial Migration in Disk Galaxies I: Transient Spiral Structure and Dynamics. ArXiv e-prints, October 2011.4
411. M. Valluri, V. P. Debattista, T. R. Quinn, R. Roˇskar, and J. Wadsley. Probing the shape and history of the Milky Way halo with orbital spectral analysis. MNRAS, 419:1951–1969, January 2012.

MCA94P030

412. J. Noffsinger and M.L. Cohen, "Electron-phonon coupling and superconductivity in double-walled carbon nanotubes," Phys. Rev B 83, 165420 (2011).
413. T. Koretsune, S. Saito, and M.L. Cohen, "One-dimensional alkali-doped C60 chains encapsulated in BN nanotubes," Phys. Rev. B 83, 193406 (2011).
414. Cohen and B.D. Malone, "Wave-function engineering: Other phases of Si for photovoltaic applications," J. Appl. Phys. 109, 102402 (2011).

415. K. Kim, Z. Lee, B.D. Malone, K.T. Chan, B. Aleman, W. Regan, W. Gannett, M.F. Crommie, M.L. Cohen, and A. Zettl, "Multiply folded graphene," Phys. Rev. B 83, 245433 (2011).
416. H. Oh, H.J. Choi, S.G. Louie, and M.L. Cohen, "Fermi surfaces and quantum oscillations in the underdoped high-Tc superconductors YBa2Cu3O6.5 and YBa2Cu4O8," Phys. Rev. B 84, 014518 (2011).
417. T. Bazhirov, J. Noffsinger, and M.L. Cohen, "Electron-phonon coupling in bcc lithium," Phys. Rev. B 84, 125122 (2011).
418. D.A. Siegel, C.-H. Park, C. Hwang, J. Deslippe, A.V. Federov, S.G. Louie, and A. Lanzara, "Many-body Interactions in Quasi-freestanding Graphene," Proc. Natl. Acad. Sci. U.S.A. 108 11365 (2011).
419. L.Z. Tan, C.-H. Park, and S.G. Louie, "New Dirac Fermions in Periodically Modulated Bilayer Graphene," Nano Lett. 11 2596 (2011).
420. C. Tao, L. Jiao, O.V. Yazyev, Y.-C. Chen, J. Feng, X. Zhang, R.B. Capaz, J.M. Tour, A. Zettl, S.G. Louie, H. Dai, and M.F. Crommie, "Spatially Resolving Edge States of Chiral Graphene Nanoribbons," Nature Physics 7, 616 (2011).
421. H.C. Hsueh, G.Y. Guo, and S.G. Louie, "Excitonic Effects in the Optical Properties of an SiC Sheet and Nanotubes," Phys. Rev. B 84, 85404 (2011).
422. G.P. Zhang, D.A. Strubbe, S.G. Louie, and T.F. George, "First-principles Prediction of Optical Second-order Harmonic Generation in the Endohedral N@C60 Compound," Phys. Rev. A 84, 023837 (2011).
423. S. Blackburn, M. Cote, S.G. Louie, and M.L. Cohen, "Enhanced electron-phonon coupling near the lattice instability of superconducting NbC1-xNx from densityfunctional calculations," Phys. Rev. B 84, 104506 (2011).
424. C. Hwang, C.-H. Park, D.A. Siegel, A.V. Fedorov, S.G. Louie, and A. Lanzara, "Direct Measurement of Quantum Phases in Graphene via Photoemission Spectroscopy," Phys. Rev. B 84, 125422 (2011).
425. O.V. Yazyev, R.B. Capaz, S.G. Louie, "Theory of Magnetic Edge States in Chiral Graphene Nanoribbons," Phys. Rev. B 84, 115406 (2011).
426. G. Samsonidze, M. Jain, J. Deslippe, M.L. Cohen, and S.G. Louie, "Simple approximate physical orbitals for *GW* quasiparticle calculations, " Phys. Rev. Lett 107, 186404 (2011).
427. G. Samsonidze, M.L. Cohen, and S.G. Louie, "Compensation-doped silicon for photovoltaic applications," Phys. Rev. B 84, 195201 (2011).
428. K.T. Chan, H.K. Lee, and M.L. Cohen, "Possibility of transforming the electronic structure of one species of graphene adatoms into that of another by application of gate voltage: First-principles calculations," Phys. Rev. B 84, 165419 (2011).
429. M.L. Cohen, "Conceptual progress for explaining and predicting semiconductor properties," J. Mater. Res. 26, 2815 (2011).
430. M. Jain, J.R. Chelikowsky, and S.G. Louie, "Reliability of Hybrid Functionals in Predicting Band Gaps," Phys. Rev. Lett. 107, 216806 (2011).

MCA95C014

431. Greene, P.T., Eldredge, J.D., Zhong, X. & Kim, J., 2011 Numerical Study of Hypersonic Flow Over an Isolated Roughness with a High-Order Cut-Cell Method, AIAA Paper 2011-3249, 2011.
432. Lagha, M., Kim, J., Eldredge, J.D. & Zhong, X., 2011 Near-wall Dynamics of Compressible Boundary Layers, Phys. Fluids, 23.
433. Lagha, M., Kim, J., Eldredge, J.D. & Zhong, X., 2011 A Numerical Study of Compressible Boundary Layers, Phys. Fluids, 23.

MCA98N017

434. Guvench, O., S.S. Mallajosyula, E.P. Raman, E. Hatcher, K. Vanommeslaeghe, T.J. Foster, F.W. Jamison, 2nd, and A.D. Mackerell, Jr., *CHARMM additive all‐atom force field for carbohydrate derivatives and its utility in polysaccharide and carbohydrate‐ protein modeling.* Journal of Chemical Theory and Computation, 2011. **7:** p. 3162‐ 3180.
435. Abel, S., F.Y. Dupradeau, E.P. Raman, A.D. MacKerell, Jr., and M. Marchi, *Molecular simulations of dodecyl‐beta‐maltoside micelles in water: influence of the headgroup conformation and force field parameters.* The journal of physical chemistry. B, 2011. **115:** p. 487‐ 99.
436. Hatcher, E., E. Sawen, G. Widmalm, and A.D. MacKerell, Jr., *Conformational properties of methyl beta‐maltoside and methyl alpha‐ and beta‐cellobioside disaccharides.* The journal of physical chemistry. B, 2011. **115:** p. 597‐ 608.
437. Mallajosyula, S.S. and A.D. MacKerell, Jr., *Influence of solvent and intramolecular hydrogen bonding on the conformational properties of o‐linked glycopeptides.* The journal of physical chemistry. B, 2011. **115:** p. 11215‐ 29.

438. Jo, S., K.C. Song, H. Desaire, A.D. MacKerell, Jr., and W. Im, *Glycan Reader: automated sugar identification and simulation preparation for carbohydrates and glycoproteins.* Journal of Computational Chemistry, 2011. **32:** p. 3135- 41.
439. Denning, E.J., U.D. Priyakumar, L. Nilsson, and A.D. Mackerell, Jr., *Impact of 2'- hydroxyl sampling on the conformational properties of RNA: Update of the CHARMM all- atom additive force field for RNA.* Journal of Computational Chemistry, 2011. **32:** p. 1929- 43.
440. Pastor, R.W. and A.D. Mackerell, Jr., *Development of the CHARMM Force Field for Lipids.* The journal of physical chemistry letters, 2011. **2:** p. 1526- 1532.
441. Baker, C.M., V.M. Anisimov, and A.D. MacKerell, Jr., *Development of CHARMM polarizable force field for nucleic acid bases based on the classical Drude oscillator model.* The journal of physical chemistry. B, 2011. **115:** p. 580- 96.
442. Jiang, W., D.J. Hardy, J.C. Phillips, A.D. Mackerell, Jr., K. Schulten, and B. Roux, *High- performance scalable molecular dynamics simulations of a polarizable force field based on classical Drude oscillators in NAMD.* The journal of physical chemistry letters, 2011. **2:** p. 87- 92.
443. Shim, J., A. Coop, and A.D. MacKerell, Jr., *Consensus 3D model of mu- opioid receptor ligand efficacy based on a quantitative Conformationally Sampled Pharmacophore.* The journal of physical chemistry. B, 2011. **115:** p. 7487- 96.
444. Shim, J. and J.A.D. MacKerell, *Computational ligand- based rational design: role of conformational sampling and force fields in model development.* MedChemComm, 2011. **2:** p. 356- 370.
445. Boston, S.R., R. Deshmukh, S. Strome, U.D. Priyakumar, A.D. MacKerell, Jr., and P. Shapiro, *Characterization of ERK docking domain inhibitors that induce apoptosis by targeting Rsk- 1 and caspase- 9.* BMC cancer, 2011. **11:** p. 7.
446. Yap, J.L., S. Worlikar, A.D. MacKerell, Jr., P. Shapiro, and S. Fletcher, *Small- molecule inhibitors of the ERK signaling pathway: Towards novel anticancer therapeutics.* ChemMedChem, 2011. **6:** p. 38- 48.
447. Velvadapu, V., T. Paul, B. Wagh, D. Klepacki, O. Guvench, A. Mackerell, Jr., and R.B. Andrade, *Desmethyl Macrolide Analogues to Address Antibiotic Resistance: Total Synthesis and Biological Evaluation of 4,8,10- Tridesmethyl Telithromycin.* ACS medicinal chemistry letters, 2011. **2:** p. 68- 72.
448. Kolhatkar, V., L. Diao, C. Acharya, A.D. Mackerell, Jr., and J.E. Polli, *Identification of novel nonsteroidal compounds as substrates or inhibitors of hASBT.* Journal of Pharmaceutical Sciences, 2012. **101:** p. 116- 26.
449. Xu, Z., Y.W. Chen, A. Battu, P. Wilder, D. Weber, W. Yu, A.D. Mackerell, L.M. Chen, K.X. Chai, M.D. Johnson, and C.Y. Lin, *Targeting zymogen activation to control the matriptase- prostasin proteolytic cascade.* Journal of Medicinal Chemistry, 2011. **54:** p. 7567- 78.
450. Underwood, K.F., D.R. D'Souza, M.T. Mochin, A.D. Pierce, S. Kommineni, M. Choe, J. Bennett, A. Gnatt, B. Habtemariam, A.D. Mackerell, Jr., and A. Passaniti, *Regulation of RUNX2 transcription factor- DNA interactions and cell proliferation by vitamin D3 (cholecalciferol) prohormone activity.* Journal of bone and mineral research : the official journal of the American Society for Bone and Mineral Research, 2011.
451. Oashi, T., A.L. Ringer, E.P. Raman, and A.D. Mackerell, *Automated Selection of Compounds with Physicochemical Properties To Maximize Bioavailability and Druglikeness.* J Chem Inf Model, 2011. **51:** p. 148- 158.
452. Raman, E.P., W. Yu, O. Guvench, and A.D. Mackerell, *Reproducing crystal binding modes of ligand functional groups using Site- Identification by Ligand Competitive Saturation (SILCS) simulations.* Journal of chemical information and modeling, 2011. **51:** p. 877- 96.
453. Darian, E., O. Guvench, B. Yu, C.K. Qu, and A.D. MacKerell, Jr., *Structural mechanism associated with domain opening in gain- of- function mutations in SHP2 phosphatase.* Proteins, 2011. **79:** p. 1573- 88.
454. Busby, B., T. Oashi, C.D. Willis, M.A. Ackermann, A. Kontrogianni- Konstantopoulos, A.D. Mackerell, Jr., and R.J. Bloch, *Electrostatic interactions mediate binding of obscurin to small ankyrin 1: biochemical and molecular modeling studies.* Journal of Molecular Biology, 2011. **408:** p. 321- 34.
455. Ringer, A.L. and A.D. Mackerell, *Calculation of the Vibrational Stark Effect Using a First- Principles QM/MM Approach.* The journal of physical chemistry letters, 2011. **2011:** p. 553- 556.
456. Denning, E.J. and A.D. Mackerell, Jr., *Impact of arsenic/phosphorus substitution on the intrinsic conformational properties of the phosphodiester backbone of DNA investigated using ab initio quantum mechanical calculations.* Journal of the American Chemical Society, 2011. **133:** p. 5770- 2.
457. Barr, D., T. Oashi, K. Burkhard, S. Lucius, R. Samadani, J. Zhang, P. Shapiro, A.D. MacKerell, and A. van der Vaart, *Importance of domain closure for the autoactivation of ERK2.* Biochemistry, 2011. **50:** p. 8038- 48.
458. Fribourg, M., J.L. Moreno, T. Holloway, D. Provasi, L. Baki, R. Mahajan, G. Park, S.K. Adney, C. Hatcher, J.M. Eltit, J.D. Ruta, L. Albizu, Z. Li, A. Umali, J. Shim, A. Fabiato, A.D. MacKerell, Jr., V. Brezina, S.C. Sealfon, M. Filizola, J. Gonzalez- Maeso, and D.E. Logothetis, *Decoding the signaling of a GPCR heteromeric complex reveals a unifying mechanism of action of antipsychotic drugs.* Cell, 2011. **147:** p. 1011- 23.

MCA98N020

459. Wise, J., Turk, M., Norman, M. L., & Abel, T. 2012a "The Birth of a Galaxy: Primordial Metal Enrichment and Stellar Populations", Ap. J., 745, 50 (simulation was begun on SDSC Trestles and completed on NASA Pleiades)
460. ________________. 2012b. "The Birth of a Galaxy: II. Effect of Primordial IMF and Other Physics", Ap. J., submitted
461. ________________. 2012c. "The Birth of a Galaxy: III. Galaxy Evolution in a Rare Peak", in preparation
462. Norman, M. L., Reynolds, D., So, G., & Harkness, R. 2012. "Direct Numerical Simulation of Cosmic Reionization in Large Volumes: I. Numerical Methods and Tests", Ap. J. Supp., submitted
463. Norman, M. L., So, G., Harkness, R., & Reynolds, D. 2012a. "Direct Numerical Simulation of Cosmic Reionization in Large Volumes: II. Suppression of Star Formation in 108-9 Ms Protogalaxies due to Radiative and Supernova Feedback", in preparation
464. ________________. 2012b. "Direct Numerical Simulation of Cosmic Reionization in Large Volumes: III. Late Stages of Reionization from the INCITE Simulation", in preparation
465. So, G., Norman, M. L., Reynolds, D., & Harkness, R. 2012. "Direct Numerical Simulation of Cosmic Reionization in Large Volumes: IV. Are Clumping Factors a Good Predictor of When Reionization Completes?", in preparation

MCA99S007

466. 1. Rojas, AV, Liwo, A, Scheraga, HA. A Study of the alpha-Helical Intermediate Preceding the Aggregation of the Amino-Terminal Fragment of the beta Amyloid Peptide (A beta(1-28)). JOURNAL OF PHYSICAL CHEMISTRY B, 12978-12983, 115, 2011 DOI 10.1021/jp2050993
467. 2. Makowski, M, Liwo, A, Sobolewski, E, Scheraga, HA. Simple Physics-Based Analytical Formulas for the Potentials of Mean Force of the Interaction of Amino-Acid Side Chains in Water. V. Like-Charged Side Chains. JOURNAL OF PHYSICAL CHEMISTRY B, 6119-6129, 115, 2011 DOI 10.1021/jp111258p
468. 3. Makowski, M, Liwo, A, E, Scheraga, HA. Simple Physics-Based Analytical Formulas for the Potentials of Mean Force of the Interaction of Amino-Acid Side Chains in Water. VI. Oppositely Charged Side Chains. JOURNAL OF PHYSICAL CHEMISTRY B, 6130-6137, 115, 2011 DOI: 10.1021/jp111259e
469. 4. He, Y, Liwo, A, Weinstein, H, Scheraga, HA. PDZ Binding to the BAR Domain of PICK1is Elucidated by Coarse-grained Molecular Dynamics. JOURNAL OF MOLECULAR BIOLOGY, 298-314, 405, 2011 DOI 10.1016/j.jmb.2010.10.051
470. 5. Liwo, A, He, Y, Scheraga, HA. Coarse-grained force field: general folding theory. PHYSICAL CHEMISTRY CHEMICAL PHYSICS, 16890-16901, 13, 2011. DOI 10.1039/c1cp20752k
471. 6. Vila, JA, Arnautova, YA, Vorobjev, Y, Scheraga, HA. Assessing the fractions of tautomeric forms of the imidazole ring of histidine in proteins as a function of pH. PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES OF THE UNITED STATES OF AMERICA, 5602-5607, 108, 2011.

MCB040002

472. Willenbring, D, Liu, LT, Mowrey, D, Xu, Y, Tang, P, Isoflurane Alters the Structure and Dynamics of GLIC, *Biophysical J*, 101, 1905-1912, (2011)
473. Wang, HJ, Kleinhammes, A, Tang, P, Xu, Y and Wu, Y, Temperature dependence of lysozyme hydration and the role of elastic energy, Phys. Rev. E 83, 031924 (2011)
474. Cui, T, Mowrey, D, Bondarenko, V, Tillman, T, Ma, D, Landrum, E, Perez-Aguilar, JM, He, J, Wang, W, Saven, JG, Eckenhoff, R, Tang, P, Xu, Y, NMR structure and dynamics of water-solubilized transmembrane domains of nicotinic acetylcholine receptor, *BBA- Biomembranes*, in press. Published online December 2, 2011. doi:10.1016/j.bbamem.2011.11.021 (2011)

MCB060006

475. 1. Molecular Dynamics Analysis of Conserved Hydrophobic and Hydrophilic Bond Interaction Networks in ErbB Family Kinases, A. Shih, S. E. Telesco, S. H. Choi, M. A. Lemmon, R. Radhakrishnan, Biochemical Journal, 2011, 436(2), 241-251. Pubmed ID: 21426301.
476. 2. Investigating Molecular Mechanisms of Activation and Mutation of the HER2 Receptor Tyrosine Kinase through Computational Modeling and Simulation, S. E. Telesco, A. Shih, Y. Liu, R. Radhakrishnan, Cancer Research Journal, 2011, **4**(4)**,** 1-35.
477. 3. Analysis of Somatic Mutations in Cancer: Molecular Mechanisms of Activation in the ErbB family of Receptor Tyrosine Kinases, A. J. Shih, S. E. Telesco, R. Radhakrishnan, Cancers, 2011, 3(1), 1195-1231; doi:10.3390/cancers3011195. Pubmed ID: 21701703
478. 4. Dynamic factors controlling carrier anchoring on vascular cells, T. N. Swaminathan, J. Liu, B. Uma, P. S. Ayyaswamy, R. Radhakrishnan, D. M. Eckmann, IUBMB Life, 2011, 63(8):640-647. Pubmed ID: 21721099

479. 5. A Multiscale Modeling Approach to Investigate Molecular Mechanisms of Pseudokinase Activation and Drug Resistance in the HER3/ErbB3 Receptor Tyrosine Kinase Signaling Network, S. E. Telesco, A. J. Shih, F. Jia, R. Radhakrishnan, Molecular Biosystems (RSC Journal), 2011, 7 (6), 2066 - 2080. DOI: 10.1039/c0mb00345j. Pubmed ID: 21509365.
480. 6. Dynamic factors controlling targeting nanocarriers to vascular endothelium, D. M. Eckmann, R. Radhakrishnan, V. Muzykantov, R. Radhakrishnan, D. M. Eckmann, Current Drug Metabolism, 2010, in press.
481. 7. Multiscale Modeling of Functionalized Nanocarriers in Targeted Drug Delivery, J. Liu, R. P. Bradley, D. M. Eckmann, P. S. Ayyaswamy, R. Radhakrishnan, Current Nanoscience, 2011, 7(5): 727-735.
482. 8. Multivalent binding of nanocarrier to endothelial cells under shear flow, J. Liu, N. J. Agrawal, A. Calderon, P. S. Ayyaswamy, D. M. Eckmann, R. Radhakrishnan, Biophysical J., 2011, 101(2):319-326. Pubmed ID: 21767483
483. 9. Systems Biology and Physical Biology of Clathrin-Mediated Endocytosis: An Integrative Experimental and Theoretical Perspective, V. Ramanan, N. J. Agrawal, J. Liu, S. Engles, R. Toy, R. Radhakrishnan, Integrative Biology (RSC Journal), 2011, 3(8), 803-815. DOI: 10.1039/c1ib00036e. Pubmed ID: 21792431.
484. 10. Nanoparticle Brownian motion and hydrodynamic interactions in the presence of flow fields, B. Uma, T. N. Swaminathan, R. Radhakrishnan, D. M. Eckmann, P. S. Ayyaswamy, Physics of Fluids, 2011, 23, 073602. DOI: 10.1063/1.3611026.
485. 11. Generalized Langevin dynamics of a nanoparticle using a finite element approach: Thermostating with correlated noise, B. Uma, T. N. Swaminathan, P. S. Ayyaswamy, D. M. Eckmann, R. Radhakrishnan, J. Chem. Phys., 2011, 135, 114104. DOI: 10.1063/1.3635776; Erratum, in press, 2012.
486. Modeling of a nanoparticle motion in a Newtonian fluid: A comparison between fluctuating hydrodynamics and generalized Langevin methods, B.Uma, R.Radhakrishnan, D.M.Eckmann and P.S.Ayyaswamy, Proceedings of the ASME 3rd Micro/Nanoscale Heat and Mass Transfer International Conference, Paper No. MNHMT2012-75019, 2012, in press.
487. 2. Modeling of Binding Free Energy of Targeted Nanocarriers to Cell Surfaces, J. Liu, D. M. Eckmann, P. S. Ayyaswamy, R. Radhakrishnan, The Seventh Interdisciplinary Transport Phenomena Conference, Proceedings of ITP2011, Interdisciplinary Transport Phenomena VII: Fluid, Thermal, Biological, Materials and Space Sciences, 2011, Dresden, Germany, pp6-6 to 6-11.
488. 1. Linking oncogenic signaling to molecular structure, Multiscale Cancer Modeling, J. Purvis, A. Shih, Y. Liu, R. Radhakrishnan, Eds. T.S. Deisboeck, G. Stamatakos, Chapman & Hall-CRC Mathematical and Computational Biology Series, 2011, pp31-43.
489. 2. Computational methods related to reaction chemistry, A. J. Shih*, S. T. Telesco*, Y. Liu, R. Venkatramani, R. Radhakrishnan, Comprehensive Biomaterials, eds. P. Ducheyne, K.E. Healy, D.W. Hutmacher, D.W. Grainger, C.J. Kirkpatrick, Elsevier London, 2011, vol. 3, pp155-169. *These authors contributed equally.
490. 3. Investigating Molecular Mechanisms of Specificity in Regulation of the HER2 Receptor Tyrosine Kinase through Molecular Modeling and Simulation, S. E. Telesco, A. Shih, Y. Liu, R. Radhakrishnan, HER2 and Cancer: Mechanism, Testing and Targeted Therapy, Nova Science Publishers, New York, 2011, in press.
491. 4. Top-down Mesoscale Models and Free Energy Calculations of Multivalent Protein-Protein and Protein-Membrane Interactions in Nanocarrier Adhesion and Receptor Trafficking, J. Liu, N. J. Agrawal, D. M. Eckmann, P. S. Ayyaswamy, R. Radhakrishnan, Innovations in Biomolecular Modeling, Royal Society of Chemistry Publishing, 2011, in press.
492. 1. A molecular docking study of the binding of ruthenium complex compounds to PIM1, GSK3, and CDK2 protein kinases, Y. Liu, N. J. Agrawal, R. Radhakrishnan, submitted to J. Molecular Modeling, 2011.
493. 2. Mesoscale Modeling and Simulations of Spatial Partitioning of Curvature Inducing Proteins under the Influence of Mean Curvature Fields in Bilayer Membranes, J. Liu, R. Tourdot, V. Ramanan, N. J. Agrawal, R. Radhakrishnan, submitted to Molecular Physics, 2011
494. 3. Modeling of Binding Free Energy of Targeted Nanocarriers to Cell Surfaces, J. Liu, D. M. Eckmann, P. S. Ayyaswamy, R. Radhakrishnan, submitted to International J. Transport Phenomena, 2011
495. 4. Fluctuating hydrodynamics approach for the simulation of nanoparticle Brownian motion in a Newtonian fluid, B.Uma, R.Radhakrishnan, D.M.Eckmann and P.S.Ayyaswamy, Proceedings of the 21st National and 10th ISHMT-ASME Heat and Mass Transfer conference, 2011.
496. 5. A hybrid formalism combining fluctuating hydrodynamics and generalized Langevin dynamics for the simulation of nanoparticle thermal motion in an incompressible fluid medium, B. Uma, D. M. Eckmann, P. S. Ayyaswamy, R. Radhakrishnan, submitted to Molecular Physics, 2011
497. 6. Temporal multiscale approach for nanocarrier motion with simultaneous adhesion and hydrodynamic interactions in targeted drug delivery, R. Radhakrishnan, B. Uma, J. Liu, P. S. Ayyaswamy, D. M. Eckmann, submitted to J. Computational Physics, 2011.

498. 1. Energy landscapes in tyrosine phosphorylation and mechanisms of phosphoryl transfer in ErbB1 kinase, Y. Liu, R. Radhakrishnan, to be submitted.
499. 2. Mechanism of erlotinib binding to EGFR through crystallography and molecular dynamics simulations, Y. Liu, J. Park, M. A. Lemmon, R. Radhakrishnan, to be submitted.
500. 3. Bridging atomic resolution models with signal dynamics in Multiscale systems biology, S. E. Telesco, R. Radhakrishnan, to be submitted.
501. 4. A Multiscale Modeling Approach to Investigate Molecular Mechanisms of Activation in the HER4/ErbB4 Receptor Tyrosine Kinase Signaling Network, S. E. Telesco, R. Vadigepalli, R. Radhakrishnan, to be submitted.

MCB070038

502. 1. Watanabe S, Resch M, Lilyestrom W, Clark N, Hansen JC, Peterson C, Luger K. *Structural characterization of H3K56Q nucleosomes and nucleosomal arrays.* (2010 Epub: 2010-01-25) Biochim Biophys Acta 1799:480-6
503. 2. Merdanovic M, Mamant N, Meltzer M, Poepsel S, Auckenthaler A, Melgaard R, Hauske P, Nagel-Steger L, Clarke AR, Kaiser M, Huber R, Ehrmann M. *Determinants of structural and functional plasticity of a widely conserved protease chaperone complex.* (2010) Nat Struct Mol Biol. 17(7):837-843
504. 3. Tapley TL, Franzmann TM, Chakraborty S, Jakob U, Bardwell JC. *Protein refolding bypH-triggered chaperone binding and release.* (2010-01-07 Epub: 2010-01-07) Proc Natl Acad Sci U S A :
505. 4. Lee YS, Lee S, Demeler B, Molineux IJ, Johnson KA, Yin YW. *Each monomer of thedimeric accessory protein for human mitochondrial DNA polymerase has a distinct rolein conferring processivity*. (2010-01-08 Epub: 2009-10-26) J Biol Chem 285:1490-9
506. 5. Hoopes JT, Liu X, Xu X, Demeler B, Folta-Stogniew E, Li C, Ha Y. *Structuralcharacterization of the E2 domain of APL-1, a Caenorhabditis elegans homolog of human amyloid precursor protein, and its heparin binding site.* (2010-01-15 Epub: 2009-11-10) J Biol Chem 285:2165-73
507. 6. Cölfen H, Laue TM, Wohlleben W, Schilling K, Karabudak E, Langhorst BW, Brookes E, Dubbs B, Zollars D, Rocco M, Demeler B. *The Open AUC Project.* (2010-02-01 Epub:2009-03-19) Eur Biophys J 39:347-59
508. 7. Brookes E, Cao W, Demeler B. *A two-dimensional spectrum analysis for sedimentation velocity experiments of mixtures with heterogeneity in molecular weight and shape.*(2010-02-01 Epub: 2009-02-27) Eur Biophys J 39:405-14
509. 8. Brookes E, Demeler B, Rosano C, Rocco M. *The implementation of SOMO (Solution MOdeller) in the UltraScan analytical ultracentrifugation data analysis suite: enhanced capabilities allow the reliable hydrodynamic modeling of virtually any kind ofbiomacromolecule.* (2010-02-01 Epub: 2009-02-21) Eur Biophys J 39:423-35
510. 9. Voss JE, Scally SW, Taylor NL, Atkinson SC, Griffin MD, Hutton CA, Parker MW, Alderton MR, Gerrard JA, Dobson RC, Dogovski C, Perugini MA. *Substrate-mediatedstabilization of a tetrameric drug target reveals Achilles heel in anthrax.* (2010-02-19 Epub: 2009-11-30) J Biol Chem 285:5188-95
511. 10. Demeler B. *Methods for the design and analysis of sedimentation velocity andsedimentation equilibrium experiments with proteins.* (2010-04-01 Epub: 0000-00-00) Curr Protoc Protein Sci Chapter 7:Unit 7.13
512. 11. Chen J, Feige MJ, Franzmann TM, Bepperling A, *Buchner J. Regions outside the alphacrystallindomain of the small heat shock protein Hsp26 are required for its dimerization.*(2010-04-23 Epub: 2010-02-18) J Mol Biol 398:122-31
513. 12. Demeler B, Brookes E, Wang R, Schirf V, Kim CA. *Characterization of Reversible Associations by Sedimentation Velocity with UltraScan.* (2010-05-18 Epub: 2010-05-18) Macromol Biosci.
514. 13. Ghosh RP, Nikitina T, Horowitz-Scherer RA, Gierasch LM, Uversky VN, Hite K, Hansen JC, Woodcock CL. *Unique physical properties and interactions of the domainsof methylated DNA binding protein 2.* (2010-05-25 Epub: 0000-00-00) Biochemistry 49:4395-410
515. 14. Xiao F, Demeler B, Guo P. *Assembly mechanism of the sixty-subunit nanoparticles viainteraction of RNA with the reengineered protein connector of phi29 DNA-packaging motor.* (2010-06-22 Epub: 0000-00-00) ACS Nano 4:3293-301
516. 15. Brookes E, Demeler B, Rocco M. *Developments in the US-SOMO bead modeling suite:new features in the direct residue-to-bead method, improved grid routines, and influenceof accessible surface area screening.* (2010-07-07 Epub: 0000-00-00) Macromol Biosci 10:746-53
517. 16. Seetharaman SV, Winkler DD, Taylor AB, Cao X, Whitson LJ, Doucette PA, Valentine JS, Schirf V, Demeler B, Carroll MC, Culotta VC, Hart PJ. *Disrupted zinc-binding sites in structures of pathogenic SOD1 variants D124V and H80R.* (2010-07-13 Epub: 0000-00-00) Biochemistry 49:5714-25
518. 17. Nardozzi J, Wenta N, Yasuhara N, Vinkemeier U, Cingolani G. *Molecular Basis for theRecognition of Phosphorylated STAT1 by Importin alpha5.* (2010-07-17 Epub: 2010-07-17) J Mol Biol :
519. 18. Gilmore JM, Bieber Urbauer RJ, Minakhin L, Akoyev V, Zolkiewski M, Severinov K,

520. Urbauer JL. *Determinants of affinity and activity of the anti-sigma factor AsiA.* (2010-07-27 Epub: 0000-00-00) Biochemistry 49:6143-54
521. 19. Rosano C, Rocco M. *Solution properties of full-length integrin alpha(IIb)beta3 refinedmodels suggest environment-dependent induction of alternative bent /extended restingstates.* (2010-08-01 Epub: 2010-06-24) FEBS J 277:3190-202
522. 20. Wang R, Taylor AB, Leal BZ, Chadwell LV, Ilangovan U, Robinson AK, Schirf V, Hart PJ, Lafer EM, Demeler B, Hinck AP, McEwen DG, Kim CA. *Polycomb group targeting through different binding partners of RING1B C-terminal domain.* (2010-08-11 Epub:0000-00-00) Structure 18:966-75
523. 21. Comoletti D, Miller MT, Jeffries CM, Wilson J, Demeler B, Taylor P, Trewhella J, Nakagawa T. *The macromolecular architecture of extracellular domain of alphaNRXN1:domain organization, flexibility, and insights into trans-synaptic disposition.* (2010-08-11Epub: 0000-00-00) Structure 18:1044-53 22. Zaytsev DV, Xie F, Mukherjee M, Bludin A, Demeler B, Breece RM, Tierney DL, Ogawa MY. *Nanometer to millimeter scale peptide-porphyrin materials.* (2010-10-11 Epub: 0000-00-00) Biomacromolecules 11:2602-9
524. 23. Szerlong HJ, Prenni JE, Nyborg JK, Hansen JC. *Activator-dependent p300 acetylation of chromatin in vitro: enhancement of transcription by disruption of repressive nucleosomenucleosomeinteractions.* (2010-10-15 Epub: 2010-08-18) J Biol Chem 285:31954-64
525. 24. Zhuo Y, Ilangovan U, Schirf V, Demeler B, Sousa R, Hinck AP, Lafer EM. *Dynamicinteractions between clathrin and locally structured elements in a disordered proteinmediate clathrin lattice assembly.* (2010-11-26 Epub: 2010-09-25) J Mol Biol 404:274-90
526. 25. Panza G, Luers L, Stöhr J, Nagel-Steger L, Weiss J, Riesner D, Willbold D, Birkmann E. *Molecular interactions between prions as seeds and recombinant prion proteins assubstrates resemble the biological interspecies barrier in vitro.* (2010-12-09 Epub: 2010-12-09) PLoS One 5:e14283
527. 26. Lin AP, Demeler B, Minard KI, Anderson SL, Schirf V, Galaleldeen A, McAlister-Henn L. *Construction and Analyses of Tetrameric Forms of Yeast NAD(+)-Specific Isocitrate Dehydrogenase.* (2010-12-21 Epub: 2010-12-21) Biochemistry :
528. 27. Li J, Richter K, Buchner J. *Mixed Hsp90-cochaperone complexes are important for theprogression of the reaction cycle.* (2011-01-01 Epub: 2010-12-19) Nat Struct Mol Biol 18:61-6
529. 28. Rana MS, Riggs AF. *Indefinite noncooperative self-association of chicken deoxy hemoglobin D.* (2011-01-04 Epub: 2011-01-04) Proteins :
530. 29. Hochdorffer K, Ma?rz-Berberich J, Nagel-Steger L, Epple M, Meyer-Zaika W, Horn AH, Sticht H, Sinha S, Bitan G, Schrader T. *Rational Design of beta-Sheet Ligands AgainstA(42)-Induced Toxicity.* (2011-03-07 Epub: 2011-03-07) J Am Chem Soc :30. Gu L, Eisman EB, Dutta S, Franzmann TM, Walter S, Gerwick WH, Skiniotis G, Sherman DH. *Tandem acyl carrier proteins in the curacin biosynthetic pathway promote consecutive multienzyme reactions with a synergistic effect.* (2011-03-14 Epub: 2011-02-23) Angew Chem Int Ed Engl 50:2795-8
531. 31. Muthurajan UM, McBryant SJ, Lu X, Hansen JC, Luger K. *The linker region of macroH2A promotes self-association of nucleosomal arrays.* (2011-04-30 Epub: 2011-04-30) J Biol Chem :
532. 32. Yang C, van der Woerd MJ, Muthurajan UM, Hansen JC, Luger K. *Biophysical analysisand small-angle X-ray scattering-derived structures of MeCP2-nucleosome complexes.*(2011-05-01 Epub: 2011-01-29) Nucleic Acids Res 39:4122-35
533. 33. Franzmann TM, Czekalla A, Walter SG. *Regulatory circuits of the AAA+ disaggregaseHsp104.* (2011-05-20 Epub: 2011-03-23) J Biol Chem 286:17992-8001
534. 34. Laue T, Demeler B. *A postreductionist framework for protein biochemistry.* (2011-06-01 Epub: 0000-00-00) Nat Chem Biol 7:331-4
535. 35. Lee S, Xue Y, Hu J, Wang Y, Liu X, Demeler B, Ha Y. *The E2 domains of APP andAPLP1 share a conserved mode of dimerization.* (2011-06-21 Epub: 2011-05-26) Biochemistry 50:5453-64
536. 36. Sathiyamoorthy K, Mills E, Franzmann TM, Rosenshine I, Saper MA. *The crystal structure of Escherichia coli group 4 capsule protein GfcC reveals a domain organization resembling that of Wza.* (2011-06-21 Epub: 2011-05-27) Biochemistry 50:5465-76
537. 37. Childers BM, Cao X, Weber GG, Demeler B, Hart PJ, Klose KE. *N-terminal residues of the Vibrio cholerae virulence regulatory protein ToxT involved in dimerization and modulation by fatty acids.* (2011-08-12 Epub: 2011-06-14) J Biol Chem 286:28644-55
538. 38. Catanese DJ Jr, Matthews KS. *Disconnected Interacting Protein 1 binds with highaffinity to pre-tRNA and ADAT.* (2011-10-28 Epub: 2011-09-24) Biochem Biophys Res Commun 414:506-11
539. 39. Borries Demeler and Raminderjeet Singh and Pierce, Marlon and Emre H. Brookes and Suresh Marru and Bruce Dubbs. "UltraScan Gateway Enhancements: in Collaboration with TeraGrid Advanced User Support." ACM TeraGrid 2011.
540. 40. Demeler, B., Singh, R., Pierce, M., Brookes, E. H., Marru, S., and Dubbs, B. 2011. *UltraScan gateway enhancements: in collaboration with TeraGrid advanced user support.* In Proceedings of the 2011 Teragrid

Conference: Extreme Digital Discovery (Salt Lake City, Utah, July 18 - 21, 2011). TG '11. ACM, New York, NY, 1-8. DOI=http://doi.acm.org/10.1145/2016741.2016778

541. 41. Brookes, E. H. and Demeler, B. 2010. *Performance optimization of large non-negatively constrained least squares problems with an application in biophysics.* In Proceedings of the 2010 Teragrid Conference (Pittsburgh, Pennsylvania, August 02 - 05, 2010). TG '10.
542. ACM, New York, NY, 1-9. DOI= http://doi.acm.org/10.1145/1838574.1838579 **B. Submitted publications:** 1. Rocco, M, Brookes, E.H., Byron, O. *US-SOMO: Methods for the Construction andHydration of Macromolecular Hydrodynamic Models.* G.C.K. Roberts (ed.), Encyclopedia of Biophysics, DOI 10.1007/978-3-642-16712-6, Springer-Verlag Berlin Heidelberg 2012

MCB070118

543. Yanzi Zhou and Yingkai Zhang, Serine Protease Acylation Proceeds with a Subtle Re-orientation of the Histidine Ring at the Tetrahedral Intermediate, *ChemComm,* 47, 1577–1579, **2011**.
544. Ruibo Wu, Zhenyu Lu, Zexing Cao and Yingkai Zhang, Zinc Chelation with Hydroxamate in Histone Deacetylases Modulated by Water Access to the Linker Binding Channel *J. Am. Chem. Soc.*, 133, 6110-6113, **2011**.
545. Yanzi Zhou, Syed M. Moin, Sinisa Urban and Yingkai Zhang, An Internal Water-Retention Site in the Rhomboid Intramembrane Protease GlpG Ensures Catalytic Efficiency, *Structure*, in revision.

MCB080013

546. Shim JY, Rudd J, Ding TT. Distinct second extracellular loop structures of the brain cannabinoid CB(1) receptor: Implication in ligand binding and receptor function. Proteins. 2011 Feb;79(2):581-597. (*TeraGrid Allocations support acknowledged*)
547. Shim JY, Bertalovitz AC, Kendall DA. Identification of essential cannabinoid-binding domains: structural insights into early dynamic events in receptor activation. J. Biol. Chem. 2011 Sep 23;286(38):33422-33435. (*TeraGrid Allocations support acknowledged*)

MCB080071

548. A. D. Friesen and D. V. Matyushov, Local polarity excess at the interface of water with a nonpolar solute, Chem. Phys. Lett. 511, 256–261 (2011).
549. 2 D. V. Matyushov and A. Y. Morozov, Electrostatics of the proteinwater interface and the dynamical transition in proteins, Phys. Rev. E 84, 011908 (2011).
550. 3 A. D. Friesen and D. V. Matyushov, Non-Gaussian statistics of electrostatic fluctuations of hydration shells, J. Chem. Phys. 135, 104501 (2011).
551. 4 D. R. Martin, A. D. Friesen, and D. V. Matyushov, Electric field inside a “Rossky cavity” in uniformly polarized water, J. Chem. Phys. 135, 084514 (2011).
552. 5 D. V. Matyushov, Nanosecond Stokes Shift Dynamics, Dynamical Transition, and Gigantic Reorganization Energy of Hydrated Heme Proteins, J. Phys. Chem. B 115, 10715–10724 (2011).
553. 6 D. V. Matyushov, Dielectric response of protein solutions, J. Chem. Phys. , submitted, arXiv:1108.2713 (2012).
554. 7 D. R. Martin, S. B. Ozkan, and D. V. Matyushov, Dissipative electro-elastic network model of protein electrostatics, Biophys. J. , submitted, arXiv:1111.4726 (2011).
555. 8 D. R. Martin and D. V. Matyushov, Non-Gaussian Statistics and Nanosecond Dynamics of Electrostatic Fluctuations Affecting Optical transitions in a Green Fluorescent Protein, Biophys. J. , in preparation (2011).
556. 9 D. V. Matyushov, On the theory of dielectric response of protein solutions, J. Phys. Chem. B , submitted (2012).

MCB090003

557. S. Gopal, S. Gul, Y. Cheng, S. M. Law, **M. Feig**: PRIMO: A transferable coarsegrained force field for the simulation of proteins. *under review*
558. L. Fang, S. Mukherjee, H. J. Cho, C. Chan, **M. Feig**: Interaction of palmitate with the human kinase IRE1 from computer simulations. *under review*
559. V. Mirjalili, N. Liu, K. Vadivel, **M. Feig**: Protein structure refinement with long molecular dynamics simulations and an improved force field. *under review*
560. M. Sayadi, **M. Feig**: Computational study of Ser16 and Thr17 phosphorylated phospholamban in implicit membrane. *under review*

561. R. B. Best, X. Zhu, J. Shim, P. Lopes, J. Mittal, **M. Feig**, A. D. MacKerell Jr: Optimization of the additive CHARMM all-atom protein force field targeting improved sampling of the backbone Φ, Ψ and sidechain X_1 and X_2 dihedral angles. *to be submitted*
562. S. M. Law, A. Panahi, **M. Feig**: A Path-Based Reaction Coordinate for Biased Sampling of Nucleic Acid Translocation, *under review*
563. Y.-M. Cheng, S. M. Gopal, S. M. Law, **M. Feig**: Molecular dynamics trajectory compression with a coarse-grained model. IEEE/ACM Transactions on Computational Biology and Bioinformatics (2011) *in press*
564. S. M. Law, **M. Feig**: Base-flipping mechanism in DNA mismatch recognition by MutS. *Biophysical Journal* (2011) 101, 2223-2231
565. A. Jaskierny, A. Panahi, **M. Feig**: Effect of flanking residues on the conformational sampling of the internal fusion peptide from Ebola virus. *Proteins* (2011) 79, 1109-1117
566. S. Seibold, B. Singh, C. Zhang, M. Kireeva, C. Domecq, A. Bouchard, A. Nazione, M. Feig, R. Cukier, B. Coulombe, M. Kashlev, M. Hampsey, Z. Burton: Conformational coupling, bridge helix dynamics and active site dehydration in catalysis by RNA polymerase. *Biochmica et Biophysica Acta – Gene Regulatory Mechanisms* (2010) 1799, 575-587
567. H. Cho, S. Mukherjee, P. Palasuberniam, C. Nezich, **M. Feig**, C. Chan: Molecular mechanism by which palmitate inhibits PKR autophosphorylation. *Biochemistry (2011)* 50, 1110-1119

MCB090108

568. **Judith M. LaLonde**, Young Do Kwon, David M. Jones, Alexander W. Sun, Joel R. Courter, Takahiro Soeta, Toyohara Kobayashi, Amy M. Princiotto, Xueling Wu, Arne Schön, Ernesto Freire, Peter D. Kwong, John R. Mascola, Joseph Sodroski, Navid Madani, Amos B. Smith, III "Dual Hot Spot Structure-Based Design, Synthesis and Characterization of Small-Molecule HIV-1 Entry Inhibitor", Submitted, Jan. 2012.

MCB090131

569. Zhu, S., & Elcock, A. H. "A combined computational and experimental study of the effects of amino acid mutations on the binding thermodynamics of a small molecule inhibitor to the p38α MAP kinase." (to be submitted to *J. Am. Chem. Soc.*)
570. Thomas, A. S., Mao, S., & Elcock, A. H. "Molecular dynamics simulations of the bacterial chaperone trigger factor show that it is capable of extraordinary conformational flexibility." (to be submitted to *Structure*)
571. Thomas, A. S., & Elcock, A. H. "Molecular dynamics simulations of a complete atomic model of an oligonucleosome array"
572. Thomas, A. S., Zhu, S., Andrews, C. A., & Elcock, A. H. "Solvent behavior in a fully atomic model of the bacterial cytoplasm"

MCB100049

573. Roxbury, D.; Manohar, S.; Jagota, A. *J. Phys. Chem. C* **2010**, *114*, 31
574. Roxbury, D.; Jagota, A.; Mittal, J., *J. Am. Chem. Soc.* **2011,** *133* (34), 13545-13550.
575. Roxbury, D.; Mittal, J.; Jagota, A., *Nano Lett.* (Under Review).

MCB100065

576. 1) Nagarian A. and Woolf TB (2011) Coarse grained simulations of transitions in the E2 to E1 conformations for Ca ATPase (SERCA) show Entropy:Enthalpy compensation.
577. 2) Nagarian A. and Woolf TB (2011) Studies on the coarse-grained conformational changes of the E1P to E2P states of SERCA: The importance of water.
578. 3) Perilla J. and Woolf TB (2011) Towards a better order parameter: Improved sampling of transitions in DIMS
579. 1) Nagarian A. and Woolf TB (2012) Comparing all-atom and coarse-grained simulations for optimal sampling in SERCA (2012 pending)
580. 2) Denning L., Perilla J., Beckstein O., and Woolf TB (2012) Holo AdK transitions can be inferred from apo AdK: Simulations of transition events in DIMS (2012 pending).

MCB100146

581. Mehrbod M, Mofrad MRK. On the Significance of Microtubule Flexural Behavior in Cytoskeletal Mechanics. PLoS One. 2011. 6(10):e25627.
582. Jamali Y, Azimi, M, Mofrad MRK. (2010) A Sub- Cellular Viscoelastic Model for Cell Population Mechanics, PLoS One. 5(8):e12097

583. Kolahi KS, Mofrad MR. Mechanotransduction: a major regulator of homeostasis and development. Wiley Interdiscip Rev Syst Biol Med. 2010 Nov- Dec;2(6):625- 39. Review.
584. Golji J, Lam J, Mofrad MRK. Vinculin Activation Is Necessary for Complete Talin Binding. Biophysical Journal Volume 100((2):332- 40
585. Golji J, Mofrad MRK. A Molecular Dynamics Investigation of Vinculin Activation. Biophysical Journal Volume 99(3) August 2010
586. Peter SJ, Mofrad MRK. Computational Modeling of Axonal Microtubule Bundles under Tension. Biophysical Journal. In press.
587. Mehrbod M, Mofrad MR, Localized Lipid Packing of Transmembrane Domain Impedes Integrin Clustering, Submitted to PLoS Computational Biology, 2012.
588. Trisno S, Mehrbod M, Mofrad MR, Talin Binding Significantly Reduces the Affinity Between Integrin Transmembrane Helices; In preparation
589. Hao H, Mehrbod M, Mofrad MR, New Binding Sites for FAK on Integrin Cytoplasmic Domain; In Preparation.
590. Azimi M, Jamali Y, Mofrad MRK. Accounting for Diffusion in Agent Based Models of Reaction- Diffusion Systems with Application to Cytoskeletal Diffusion. PLoS One. 2011;6(9):e25306.
591. Moussavi- Baygi RM, Jamali Y, Karimi R, Mofrad MRK. Brownian Dynamics Simulation of Nucleocytoplasmic Transport: A Coarse- Grained Model for the Functional State of the Nuclear Pore Complex. PLoS Computational Biology. June 2011, Volume 7, Issue 6, e1002049
592. Jamali T, Jamali Y, Mehrbod M, Mofrad MRK. Nuclear Pore Complex: Biochemistry and Biophysics of Nucleocytoplasmic Transport in Health and Disease. International Review of Cell and Molecular Biology, 2011; Volume 287: 233- 286.
593. Moussavi- Baygi RM, Jamali Y, Karimi R, Mofrad MRK. Biophysical Coarse- Grained Modeling Provides Insights into Transport through the Nuclear Pore Complex. Biophysical Journal,2011 100(6):1410- 1419.
594. Zhao C, Moussavi- Baygi R, Mofrad M. Exploring the Interaction between the Nuclear Pore Protein Tpr and the Exportin Crm1: A Computational Method of Finding Protein- Protein Binding Sites. Submitted to Biophysical Journal. 2011.
595. Azimi, M ; Mofrad, M. R., First and Second Order Chemical Kinetics in Dilute and Crowded Agent Based Models, 2012. In preparation
596. Moussavi- Baygi; Karimi; Mofrad, Do FG- repeat domains form a gel inside the NPC?, 2012. In preparation.

MCB110056

597. J. Hsin, A. Gopinathan, and K. C. Huang. Nucleotide-dependent conformations of FtsZ dimmers and force generation observed through molecular dynamics simulations, 2012. Submitted.
598. J. G. Cueva, J. Hsin, K. C. Huang, and M. B. Goodman. Alpha tubulin acetylation regulates proto_lament number in native microtubules, 2012. Submitted.
599. J. Hsin, J. Strumpfer, E. H. Lee, and K. Schulten. The molecular origin of the hierarchical elasticity of titin: simulation, experiment and theory. Annu. Rev. Biophys., 40:187-203, 2011.
600. W. Chen, J. Lou, J. Hsin, K. Schulten, S. C. Harvey, and C. Zhu. Molecular dynamics simulations
601. of forced unbending of integrin _V _3. PLoS Comput. Biol., 7(2):e1001086, 2011.
602. J. Hsin and K. Schulten. Improved resolution of tertiary structure elasticity in muscle protein. Biophys. J., 100:L22-L24, 2011.
603. Y. Liu, J. Hsin, H. Kim, P. R. Selvin, and K. Schulten. Extension of a three-helix bundle domain of myosin VI and key role of calmodulins. Biophys. J., 100:2964{2973, 2011.
604. J. Hsin, L. LaPointe, A. Kazy, C. Chipot, A. Senes, and K. Schulten. Oligomerization state of photosynthetic core complexes is correlated with the dimerization a_nity of a transmembrane helix. J. Am. Chem. Soc., 133:14071{14081, 2011.

MCB120047

605. 3. Zerbetto, M, Buck, M., Meirovitch, E & Polimeno, A (2011) Combining experiment and computation for the characterization of ps-ns time-scale protein motions: An integration of main-chain and global dynamics to interpret NMR relaxation in the Rho GTPase binding domain of plexin B1 J. Phys. Chem. B. 115, 376-88.
606. 4. Mehdi Bagheri Hamaneh, Liqun Zhang, Matthias Buck, (2011) A direct coupling between global and internal motions in a single domain protein? A molecular dynamics investigation of extreme case scenarios Biophysical Journal, 101, 196-204.
607. 5. Zhang, L., Bouguet-Bonnet, S. & Buck , M. (2010) Dynamic and conformational coupling in protein-GTPase interactions from experiments, sequence analysis and molecular dynamics simulations in Methods in Molecular Biology: Allostery Methods & Protocols (Humana/Springer Press), 2012, Volume 796, Part 3, 235-259, DOI:10.1007/978-1-61779-334-9-13.

608. 1. Zhang, L., Sodt, A., Venable, R.M., Pastor R.W. Buck, M. Prediction and persistency of transmembrane helix dimer structures in lipid bilayers using implicit and explicit solvent and lipid molecular dynamics simulations with enhanced sampling and the new CHARMM36 lipid force-field. in preparation for Biophysical J.

MSS080019

609. Cheng-Hsiang Lin, Zhi Liang, and Hai-Lung Tsai, "Femtosecond and Nanosecond Dual-laser Optical Emission Spectroscopy of Gaseous Mixtures" Optics Letters (submitted).
610. Cheng-Hsiang Lin, Zhi Liang, and Hai-Lung Tsai, "Time-resolved observation of femtosecond and nanosecond dual-laser optical manipulation of radical gaseous mixtures" Applied Physics Letters (submitted).
611. Zhi Liang and Hai-Lung Tsai, "Effect of Thin Film Confined between Two Dissimilar Solids on Interfacial Thermal Resistance," Journal of Physics: Condensed Matter 23, 495303 (2011).
612. Zhi Liang and Hai-Lung Tsai, "Thermal Conductivity of Interfacial Layers in Nanofluids," Physical Review E 83, 041602 (2011).
613. Zhi Liang and Hai-Lung Tsai, "Effect of Molecular Film Thickness on Thermal Conduction across Solid–Film Interfaces," Physical Review E 83, 061603 (2011).

PHY100035

614. D. Vrinceanu, R. Onofrio, and H. R. Sadeghpour, ApJ, doi:10.1088/0004-637X/745/1/1 (2012).
615. M. Guglielmino, V. Penna, and B. Capogrosso-Sansone, Phys. Rev. A 84, 031603(R) (2011).

PHY110010

616. 1. "Nucleon structure from 2+1f dynamical DWF lattice QCD at nearly physical pion mass" S. Ohta [for the RBC and for the UKQCD Collaborations]. arXiv:1111.5269 [hep-lat]. Proceedings for Lattice 2011, Squaw Valley, CA, USA, July 2011. To be published in Proceedings of Science (Lattice 2011), 168 (2012).
617. 2. "Nucleon structure from 2+1-flavor dynamical DWF lattice QCD at nearly physical pion mass" M. Lin and S. Ohta [for the RBC and UKQCD Collaborations]. arXiv:1112.5489 [hep-lat]. Proceedings for Erice School "From Quarks and Gluons to Hadrons and Nuclei," September 16-24, 2011, Erice, Sicily. To be published in Progress in Particle and Nuclear Physics (2012).

PHY110015

618. 1. Higgs Boson Mass Predictions in SUGRA Uni_cation and Recent LHC-7 Results," S. Akula, B. Altunkaynak, D. Feldman, P. Nath and G. Peim, Submitted Phys. Rev. D; arXiv:1112.3645 [hep-ph].
619. 2. Naturalness, Supersymmetry and Implications for LHC and Dark Matter," S. Akula, M. Liu, P. Nath and G. Peim Submitted Phys. Lett. B; arXiv:1111.4589 [hep-ph].
620. 3. Excess Observed in CDF β⊙8 → $\mu+\mu$ - and SUSY at the LHC," S. Akula, D. Feldman, P. Nath and G. Peim Phys. Rev. D 84, 115011 (2011); arXiv:1107.3535 [hep-ph].
621. 4. An Explanation of the CDF Dijet Anomaly within a U(1)X Stueckelberg Extension," Z. Liu, P. Nath and G. Peim Phys. Lett. B 701, 601 (2011); arXiv:1105.4371 [hep-ph].
622. 5. \New Constraints on Dark Matter from CMS and ATLAS Data," S. Akula, D. Feldman, Z. Liu, P. Nath and G. Peim Mod. Phys. Lett. A 26, 1521 (2011); arXiv:1103.5061 [hep-ph].
623. 6. \Interpreting the First CMS and ATLAS SUSY Results," S. Akula, N. Chen, D. Feldman, M. Liu, Z. Liu, P. Nath and G. Peim Phys. Lett. B 699, 377 (2011); arXiv:1103.1197 [hep-ph].
624. 7. \Predictive Signatures of Supersymmetry: Measuring the Dark Matter Mass and Gluino Mass with Early LHC data," D. Feldman, K. Freese, P. Nath, B. D. Nelson and G. Peim Phys. Rev. D 84, 015007 (2011); arXiv:1102.2548 [hep-ph].
625. 8. \Low Mass Gluino within the Sparticle Landscape, Implications for Dark Matter, and Early Discovery Prospects at LHC-7"N. Chen, D. Feldman, Z. Liu, P. Nath and G. Peim, Phys. Rev. D 83, 035005 (2011); arXiv:1011.1246 [hep-ph]
626. 9. \Higgsino Dark Matter Model Consistent with Galactic Cosmic Ray Data and Possibility of Discovery at LHC-7" N. Chen, D. Feldman, Z. Liu, P. Nath and G. Peim Phys. Rev. D 83, 023506 (2011); arXiv:1010.0939 [hep-ph].
627. 10. \SUSY Discovery Potential and Benchmarks for Early Runs at ps = 7 TeV at the LHC" B. Altunkaynak, M. Holmes, P. Nath, B. D. Nelson and G. Peim Phys. Rev. D 82, 115001 (2010); arXiv:1008.3423 [hep-ph].

PHY110016

628. Mason, J., Perez, J.C., Boldyrev, S. & Cattaneo, F. Numerical simulations of strong incompressible magnetohydrodynamic turbulence. Submitted to Phys. of Plasmas.

629. Perez, J.C., Mason, J., Boldyrev, S. & Cattaneo, F. High-resolution numerical study of energy spectra in strong MHD turbulence. In preparation.
630. Zhdankin, V. et al. Magnetic discontinuities in magnetohydrodynamic turbulence and the solar wind. Submitted to Phys. Rev. Lett.

PHY110018

631. \Probing Novel Scalar and Tensor Interactions from (Ultra)Cold Neutrons to the LHC", T. Bhattacharya, V. Cirigliano, S. D. Cohen, A. Filipuzzi, M. Gonzalez-Alonso, M. L. Graesser, R. Gupta, H.-W. Lin, submitted to PRD (arXiv:1110.6448 [hep-ph])

PHY120016

632. Bode, T., Haas, R., Bogdanovi´c, T., Laguna, P., and Shoemaker, D., “Mergers of Supermassive Black Holes in Astrophysical Environments,” Astrophys. J. 744, 45 (2012)
633. O’Shaughnessy, R., Vaishnav, B., Healy, J. and Shoemaker, D., “Efficient asymptotic frame selection for binary black hole spacetimes using asymptotic radiation,” accepted for publication by Physical Review D (2012)
634. Bode, T., Laguna, P., and Matzner, R., “Super-Extremal Spinning Black Holes via Accretion,” Physical Review D, vol. 84, 064044 (2011)
635. Loffler, F., Faber, J., Bentivegna, E., Bode, T., Diener, P., et. al ,“The Einstein Toolkit: A Community Computational Infrastructure for Relativistic Astrophysics,” arXiv 1111.3344
636. Healy, J., Bode, T., Haas, R., Pazos, E., Laguna, P., Shoemaker, D., and Yunes, N., “Late Inspiral and Merger of Binary Black Holes in Scalar-Tensor Theories of Gravity,” submitted to Physical Review Letters (2011)
637. O’Shaughnessy, R., Healy, J., London, L., Meeks, Z., and Shoemaker, D., “Is J enough? Comparison of gravitational waves emitted along the total angular momentum direction with other preferred orientations,” submitted to Physical Review D (2012)
638. R. Haas, R. V. Shcherbakov, T. Bode, P. Laguna, “Tidal disruptions of white dwarfs from ultra-close encounters with intermediate mass spinning black holes,” under review Astrophysical Journal (2012).

G XSEDE Citation Report

Content for this appendix is pending the publication and subsequent citation of a research paper on the XSEDE project. XSEDE allocated users will be requested to cite this paper in their publications, thus enabling a more comprehensive analysis of the science impact of the XSEDE project.

H XSEDE EOT Event Details

The following are events conducted during the quarter.

During the first quarter, 19 training events were recorded by the XSEDE User Portal, with 831 total registered trainees, 580 by webcast, 251 in person. The details of these courses are in the table below.

Type	Title	Location	Date(s)	Hours	Number of Participants	Number of Under-represented people	Method	Funding Sources
Training	Introduction to PETSc	TACC			29		Webcast	XSEDE
Training	Data Analysis on Ranger and Longhorn	Cornell			29		Live	XSEDE
Training	Introduction to Scientific Visualization on Longhorn	Old Dominion University			20		Live	XSEDE
Training	Linux/Unix Basics	TACC			66		Webcast	XSEDE
Training	XSEDE New User Training	PSC			31		Webcast	XSEDE
Training	Data Management Planning and Execution	TACC			71		Live/Webcast	XSEDE
Training	Introduction to Parallel Computing on Lonestar and Ranger	TACC			32		Live/Webcast	XSEDE
Training	Advanced PETSC	TACC			32		Live	XSEDE
Training	Introduction to Scientific Visualization on Longhorn	TACC			6		Live	XSEDE
Training	Fortran 90/95/2003 Programming	TACC			77		Webcast	XSEDE
Training	NICS/RDAV Spring Training: Getting Started at UTK's NICS and RDAV Centers	NICS			31		Live	XSEDE
Training	C Programming Basics	TACC			86		Webcast	XSEDE
Training	NICS/RDAV Spring Training: Basics of Programming HPC	NICS			35		Live	XSEDE

	systems							
Training	NICS/RDAV Spring Training: Novice and Intermediate MPI programming	NICS			34		Live	XSEDE
Training	NICS/RDAV Spring Training: Advanced programming topics	NICS			31		Live	XSEDE
Training	C++ Programming Basics	TACC			97		Webcast	XSEDE
Training	TACC Training: Writing a Data Management Plan : A Guide for the Perplexed	TACC			46		Webcast	XSEDE
Training	Introduction to Scientific Visualization on Longhorn	Florida A&M			26		Live	XSEDE
Training	PSC at Lehigh HPC Symposium	Lehigh			35		Live	XSEDE
Presentations	Meetings with NC Carolina Institutions and central administration	Raleigh-Durham; Asheville	Jan. 17-18	14	16		Live	XSEDE; RENCI
Workshop	Identifying a Dissertation Topic: Fay Cobb Payton	Online	Jan. 15	1	19	14	Online	ELA
Workshop	Introduction to XSEDE User Portal: Maytal Dahan	Online	Jan. 19	1	25	23	Online	XSEDE
Webinar	Intro to XSEDE Portal for Minority Research Community Monthly Call		Jan. 13	1	8	6	Webcast	XSEDE
Workshop	Intro to Scientific Visualization@ODU	ODU	Jan. 20	8	12	1	In-Person	TACC
Campus Visit	North Carolina Central University	NCCU	Jan. 23	4	10	6	In-Person	XSEDE
Campus Visit	Southern University & A&M College	Southern	Jan. 31	6	16	10	In-Person	XSEDE
Workshop	Conferencing: Publishing and Presenting Your Work: Yolanda Rankin	Online	Feb. 9	1	16	14	Online	ELA
Works	Parallel Programming with MPI, part I: Dan	Online	Feb. 9	1	20	15	Online	XSEDE

hop	Stanzione							
Advisory Board	Cyber-ShARE Project	El Paso, TX	Feb. 16-17	16	20	15	Live	CREST Program
Workshop	Parallel Programming with MPI, part II: Dan Stanzione	Online	Feb. 23	1	10	6	Online	XSEDE
Tutorial	XSEDE for Physicists	Boston Convention Center	Feb 29	4	4		Live	XSEDE
Worshop	Cluster on demand - Coalition for Academic Scientific Computation	Arlington, VA	Feb. 29	0.75	60	20	Live	XSEDE
Conference	Emerging Researchers National Conference	Atlanta	Feb.23-25	16	750	650	In-Person	XSEDE
Campus Visit	Clark Atlanta University – XSEDE Presentation	Atlanta	Feb. 24	2	8	4	In-Person	XSEDE
Webinar	Intro to Computational Chemistry Tools for the Classroom		Feb. 29	1.5	8	7	Webcast	XSEDE
Presentations	Meeting with faculty and administrators at Clark Atlanta University	Atlanta, GA	March 19	6	25		Live	XSEDE
Workshop	What is the XSEDE Campus Champions program; Intro to HPC; How to submit jobs on an Xsede machine: Roger Moye	Online	March 21	1	16	13	Online	XSEDE
Campus Visit	Bowie State University – XSEDE Seminar	Bowie	March 14	2	14	8	In-Person	XSEDE
Campus Visit	Clark Atlanta University – XSEDE Education Programs	Atlanta	March 19	4	19	14	In-Person	XSEDE
Campus Visit	Florida A & M University – XSEDE Seminar	FAMU	March 29	2	27	24	In-Person	XSEDE
Workshop	Intro to Scientific Visualization @FAMU	FAMU	March 30	8	20	19	In-Person	TACC