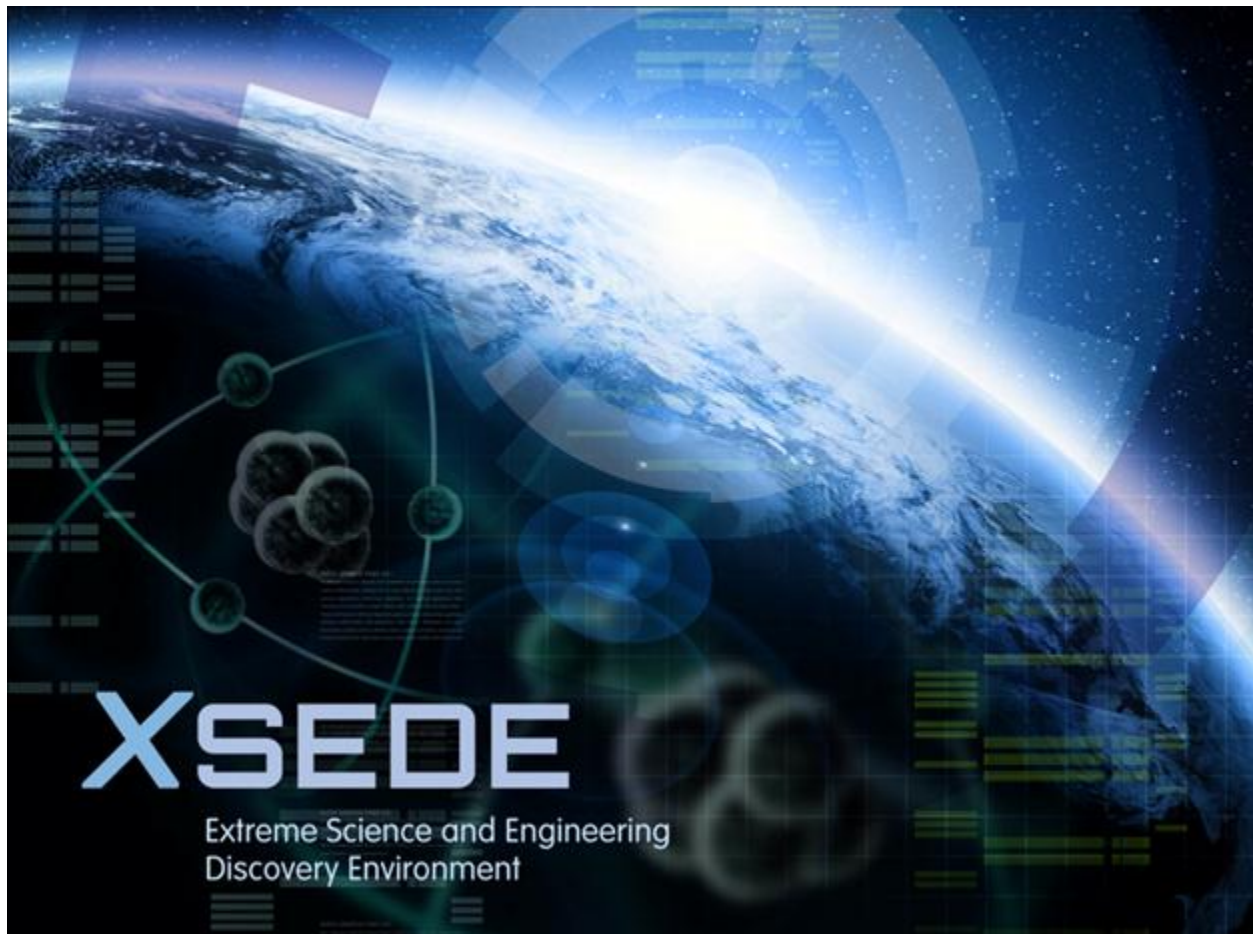


XSEDE: eXtreme Science and Engineering Discovery Environment

Program Plan for Program Year 2

14 May 2012



XSEDE Senior Management Team

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
Thom Cheatham (Utah)	Chair of the User Advisory Committee
Carol Song (PU)	Interim Chair of the XD Service Providers Forum

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)

XSEDE Year 2 Program Plan

Contents

1	Overview	2
1.1	Project Context.....	2
1.2	Program Year 2 Highlights	4
2	The Science Impact of XSEDE.....	9
2.1	Program Year 1 Science Accomplishments.....	10
2.2	Sampling of Allocated Projects for Program Year 2	18
3	XSEDE Project Office 1.1.....	22
3.1	Overview	22
3.2	Project Management and Reporting 1.1.1.....	23
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.....	26
3.6	Industry Relations 1.1.5	27
3.7	Software Development and Integration 1.1.6	28
4	XSEDE Operations 1.2.....	31
4.1	Overview	31
4.2	Cybersecurity 1.2.1	31
4.3	Data Services 1.2.2.....	32
4.4	XSEDEnet 1.2.3.....	33
4.5	Software Testing and Deployment 1.2.4.....	34
4.6	Accounting and Account Management 1.2.5.....	34
4.7	Systems Operational Support 1.2.6.....	35
5	User Services 1.3.....	37
5.1	Overview	37
5.2	Training 1.3.1	37
5.3	User Information & Interfaces 1.3.2	38
5.4	User Engagement 1.3.3	39
5.5	Allocations 1.3.4	40
6	Extended Collaborative Support Service – Projects 1.4	42
6.1	Overview	42
6.2	Extended Support for Research Teams 1.4.1	42
6.3	Novel and Innovative Projects 1.4.2	43

7	Extended Collaborative Support Service – Communities 1.5	45
7.1	Overview	45
7.2	Extended Support for Community Codes 1.5.1	46
7.3	Extended Support for Science Gateways 1.5.2	46
7.4	Extended Support for Training, Education and Outreach 1.5.3	47
8	Education and Outreach 1.6.....	49
8.1	Overview	49
8.2	Education 1.6.1	49
8.3	Outreach 1.6.2	51
8.4	Community Requirements and External Evaluation 1.6.3	52
8.5	E&O Infrastructure 1.6.4	53
8.6	Campus Bridging 1.6.5	53

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. This Program Plan documents our vision, objectives and plans for Program Year 2 (July 1, 2012 through June 30, 2013).

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 provides the expertise and tools 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 engages 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. During the first three quarters of Program Year 1, 33 new education projects, 332 new start-up projects, and 2,079 new users were added to allocations of XSEDE resources. The number of open individual accounts climbed to 6,313 at the end of this period; adding the 1,039 gateway users in Q1 2012, the XSEDE community numbered 7,352 users at the end of Q1 2012. Over the three quarters, 3,227 non-gateway individuals charged jobs, and at least 1,389 users (peak from Q4 2011) submitted jobs via science gateways - that is, at least 30% of XSEDE users submitting jobs worked through gateways.

Further details can be found in Appendix E of the XSEDE Quarterly Reports.

1.1.2 XSEDE's Integrated, Distributed Environment

XSEDE is taking on the difficult-but-necessary ongoing task of defining and documenting a complete architectural design for its distributed systems architecture. Given how the XD proposal competition ultimately resulted in the XSEDE award that also required merging in XROADS proposal ideas including architectural elements, 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 publicly available 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 Program Year 2 Highlights

Building on the considerable work of initiating the project in Program Year 1, XSEDE has nearly every planned activity area fully ramped up and executing effectively. Some of the unexpected challenges of initiating various areas have delayed some work, but the vast majority of planned activities for Program Year 1 will complete on time or ahead of schedule. With the gathering momentum of the project we plan to complete delayed tasks and push forward taking us from what has largely been a period of transition into a period of delivery of new capabilities to the community.

XSEDE's fundamental goal is to increase the productivity of scientists and engineers. A very interesting set of anticipated scientific and engineering advances are on the horizon as identified by working with teams with newly awarded allocations ranging from more traditional communities we serve such as astronomy and astrophysics, high-energy physics, computational fluid dynamics, and chemistry to communities who have arrived more recently or are just emerging in the cyberinfrastructure landscape such as genomics and humanities. In §2 of the Program Plan we recap some of the most significant work we have supported to date in Program Year 1 and provide a preview of some of the exciting work planned by scientists and engineers during Program Year 2.

XSEDE will aid these projects in a variety of ways, but most particularly by some of our Extended Collaborative Support Service (ECSS) staff members who will be paired with these and other projects for an extended period to work together to solve their challenging science and engineering problems through the application of cyberinfrastructure. These staff are primarily from the Extended Support of Research Teams (ESRT, §6.2) group and our original proposal indicated we would engage at least 20 new ESRT projects each year with documented work plans. We believe that we will exceed this number in Program Year 2, as we have done in Program Year 1.

In addition, the Novel and Innovative Projects (NIP, §6.3) team—which provides proactive, sustained efforts to jump-start XSEDE projects by non-traditional (to HPC/CI) users—has a Program Year 2 goal of 20 new projects in non-traditional areas. In Program Year 1, we completed the first ECSS flexible contract hiring process to bring in an expert in Digital Humanities. In Program Year 2, 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 of communities.

Extended Support for Community Codes (ESCC, §7.2) projects are focused on helping users with community codes and tools on XSEDE systems. Over the past year, most ESCC projects were initiated as a result of requests for assistance during the resource allocation process similar to ESRT projects. In Program Year 2, ESCC will work with the Service Providers (SPs) to determine which community codes are used most extensively and catalog and advertise them. These community codes will be prime candidates for internal projects. The goal continues to be initiating and completing at least 10 ESCC projects per year.

The primary mission of the Extended Support for Science Gateways (ESSGW, §7.3) Group is to existing and new scientific communities who would like to use XSEDE resources through gateways. As with most areas in ECSS, new requests for extended support come in continually and thus the exact priorities and plans have a dynamic element to them. Current Program Year 2

top priorities are in the areas of: contributing gateway requirements to the Architecture & Design team, documenting portions of the XSEDE architecture particularly relevant to gateways, and re-establishing the TeraGrid gateways effort in providing attribute-based authentication to automatically submit gateway end user attribute information with each job submitted to XSEDE resources.

Extended Support for Training, Education, and Outreach (ESTEO, §7.4) staff provide the person-power, expertise and technical content for activities described in the User Services Training area (§5.2) and in the Education and Outreach areas (§8). A new ESTEO focus in Program Year 2 will be on intensive internal training seminars to identify and train those who will become experts in newly requested areas and to address the specifics of use of several new resources anticipated to be deployed during Program Year 2.

XSEDE User Services (§5) activities will continue to grow in Program Year 2, in three fundamental ways: 1) expanding and completing the coverage of diverse resource types, 2) offering improved guidance for both beginning and advanced users on how to select among XSEDE's diverse resource types for the most appropriate capabilities, and 3) more features for using XSEDE effectively delivered through the XSEDE User Portal (XUP) directly. Allocations policies and procedures will be developed for storage systems, and new guidance will be provided for users of GPU-accelerated and MIC-coprocessor systems. Program Year 2 will also see the creation of more advanced documentation, building on the strong set of basic documentation developed in Program Year 1. There will be more advanced topics covered in documentation, and this will apply to new training classes as well. The XUP will continue to integrate new documentation, add new capabilities for managing files and data, and explore new opportunities including social media and other new methods for user engagement and collaboration.

The Education and Outreach (E&O, §8) team had a very productive first year pursuing the plans outlined in the proposal for XSEDE. All of the activities begun in Program Year 1 are continuing in Program Year 2 with modifications based on the lessons learned from the first year. The E&O team plans for Program Year 2 have also benefitted from the feedback from the external evaluations, the comments and recommendations from the XSEDE Advisory Board, and the recommendations of the TEOS Advisory Committee. In addition, the E&O team conducted a new round of collecting requirements from the community that have helped to refine the plans in Program Year 2.

E&O will build upon and expand external collaborations to leverage activities with other groups pursuing complementary efforts. The E&O team will expand upon its collaboration with PRACE (Partnership for Advanced Computing in Europe¹) to include other continents in the HPC Summer School in 2013 and beyond and will pursue more in-depth collaborations with the Blue Waters project as it evolves from deployment to operations. The Education team will continue to push forward with the revision of computational science competencies, assistance in creating formal computational science programs, developing education courses and workshops for students and faculty, and creating a virtual computational science education community resource.

The newly initiated XSEDE Scholars Program will raise awareness of the opportunities available to all minority and underrepresented students to pursue computational science and engineering studies and careers, together with training and mentoring to prepare them for advanced studies and careers in computational science and engineering. Likewise the broader and wildly successful Student Engagement program will review and refine processes established in Program Year 1, moving toward supporting 20 students in a full-year program as originally intended. Based on the

¹ <http://www.prace-ri.eu/>

popularity and expected success of the initial program, XSEDE will apply for an NSF REU site award, to expand the program to more students, sites, and projects.

The Campus Champion program, which supports campus representatives as a local source of knowledge about digital services, opportunities and resources, has achieved unprecedented growth through the transition from TeraGrid to XSEDE. The Campus Champions program is gaining momentum and having a significant impact on the member campuses and extending their impact to benefit researchers, educators and students on other campuses. The focus for Program Year 2 is to provide depth of support for the Champions to help ensure they and their campus personnel benefit from all that XSEDE has to offer the community. The Campus Champions Fellows Program conducted in collaboration with ECSS will directly contribute to enhancing the community of practice.

The response from the community for campus bridging has been even greater than we imagined when the planning for this effort began. During Program Year 2, XSEDE has committed an additional 1.25 FTE to enhance the effort to respond to community requirements. The Campus Bridging pilot project is allowing XSEDE to work out organizational understandings about governance and group-to-group handoffs within XSEDE that will be helpful in the long run. Initial Pilot Program will be conducted during Program Year 2 using the Execution Management Service (EMS) and the Global Federated File System (GFFS), which will provide transparent file-system level access to a collection of XSEDE data resources to machines outside of XSEDE. In addition, the Campus Bridging team will develop software packages and installers that will allow the implementation of campus systems that resemble XSEDE resources. The team will build tools that will enable clusters on campus to be built similar to XSEDE clusters and increase the interoperability among clusters.

Behind the scenes, XSEDE will continue to provide an excellent level of operational support for the services it operates and with some specific and notable goals for Program Year 2 (XSEDE Operations, §4). All of these are intended to improve the level of support provided and instantiate additional lower level capabilities to make the lives of the researchers and engineers we support easier. In Program Year 2, the operational security aspects of prevention, monitoring, and response will continue as will coordination of the XSEDE-wide security infrastructure across the SPs. In addition, we will be documenting and updating security policies and designing XSEDE's risk-based approach to cybersecurity based on our Risk, Threat, and Analysis Assessment conducted during Program Year 1. We will implement a two-factor authentication service for XSEDE which will federate with the NICS and Blue Waters two-factor authentication systems. This will be done in concert with the goal of having this new set of certificate services in place before the end of Program Year 2. In parallel with these efforts we are involved in on-going discussions with the Open Science Grid with respect to certificate services and how XSEDE and OSG might collaborate.

Based on the outcomes of an extensive evaluation of options for an XSEDE-Wide File System (XWFS) — a single high-performance, distributed file system encompassing storage elements at the leading service provider sites — we will initiate a beta deployment of a file system involving multiple SPs and allowing for collaboration with the Campus Bridging group. Similarly, the outcomes of investigations from Program Year 1 on options for unified interfaces to archive systems will be utilized to determine whether additional efforts on the existing Data Replication Service (DRS) are needed or whether to extend the existing infrastructure to all archive resources at XSEDE SPs.

The XSEDEnet Operations team will be considering options for ongoing connectivity providers given the limited initial contract with National Lambda Rail (NLR) for services starting in Program Year 1 and running through September 2012. The team continues to monitor utilization

and other metrics to help drive that decision process. In addition, the XSEDEnet team will document and maintain accurate network resource information on the XSEDE User Portal. The XSEDEnet team will continue to provide proactive expert consulting and tuning support through a special tuning team to improve the end-to-end performance of network-based applications. In addition, the team will work with developers to investigate mechanisms for integrating network metrics into applications. This is designed to understand how to use metrics in order to allow end-users and applications to make effective use of network resources.

The Software Testing and Deployment team will continue to work with SDI on software progressing through the XSEDE software engineering process. In addition, some new goals for Program Year 2 have been identified based on experiences during the first year of the project. These include assignments specifically on service provider coordination and campus bridging coordination. Systems Operational Support (SysOps) will continue to operate the 24x7 XSEDE Operations Center for initial user support and trouble ticket routing for the Program. SysOps will continue to provide support for existing central services, installation and implementation of new services, global authentication implementation for all users/staff, and semi-annual failover testing for all primary services.

During Program Year 2, one of the primary goals of the XSEDE Accounting and Account Management (A&AM) group will be to continue supporting timely processing of allocation request submissions for the various XSEDE allocations boards and prompt turnaround time for subsequent account requests for these awards. Work will also proceed on providing enhanced user feedback on the XSEDE User Portal (XUP) with respect to the status of request submissions, user account requests, and system utilization. Planning is also underway to make significant updates and improvements to these basic services, beyond these incremental improvements. Beginning in the third quarter of Program Year 1, a requirements-gathering process was initiated, directed at identifying areas where these dated systems and technologies can be brought up-to-date. This requirements-gathering process will continue into early Program Year 2, and will result in a set of requirements directed at updating all of these systems.

Program Year 1 saw an incredible amount of activity behind the scenes as the project organized itself, set up processes to manage, coordinate and communicate, and established the three key advisory functions to the project—the XSEDE Advisory Board (XAB), the User Advisory Committee (UAC), the Service Providers Forum (SPF), and the TEOS Advisory Committee. The Training and Education & Outreach (TEOS) Advisory Committee (with members also serving on the XAB) is providing in-depth review and recommendations to enhance E&O activities. As we enter Program Year 2, many of these activities will become routine. Still, several additional elements will come into place building on the work of Program Year 1. Completing the deployment and adopting the use of advanced project management software will further support these processes. Fully employing the recently established User Requirements Evaluation and Prioritization (UREP) working group will support our requirements prioritization which will affect almost every aspect of the project.

Architecture and Design (A&D, §3.4) team has laid out a new plan and is defining effective processes for moving forward; we will see a significant level of architectural work manifested as design documents. The work has been clearly defined and prioritized with input from stakeholders; there are great expectations from this are in Program Year 2 including providing detailed architectural specifications for the areas of Science Gateways, Computing, Big Data and Connecting Instrumentation.

Similarly, the Software Development and Integration (SDI, §3.7) team is beginning to realize the benefits of considerable effort in Program Year 1 in establishing a set of processes and practices to support their work. As the role of SDI has become better understood to include evolving the

current environment to delivering essential technical enhancements, they have organized themselves to be much more efficient and agile in integration of new capabilities and preparing them for operational use. While our agile process does not allow for long term forecasting, features to be focused on early in Program Year 2 in the areas of data services, resource management, security services and integration services are emerging as priorities. To facilitate this, further engineering enhancements will be put in place to provide automation in our processes to manage activities and we will consolidate and document these processes and support mechanisms for use by others within XSEDE.

While the Industry Relations (§3.6) area has largely been uncoordinated effort in Program Year 1, we also have great hopes for a reinvigorated effort following an interaction with the XAB on our ideas in this space. Principal among them, we plan to run an open competition for industry to propose an innovative development project collaboratively with XSEDE to produce solutions useful to industry in enabling them to harness the capabilities of XSEDE—a program about which the XAB was particularly enthusiastic.

Finally, we will continue to develop our resources and capabilities in making known the existence of XSEDE and its successes (External Relations, §3.5). We will finalize our Strategic Communications Plan and act on a renewed commitment to creating strong working relationships with the media. As the project has stabilized we will create a basic introductory slideshow which will be available to XSEDE staff for presentations and distributing at meetings, conferences, and trade shows.

2 The Science Impact of XSEDE

XSEDE's fundamental goal is to increase the productivity of scientists and engineers by providing them with new capabilities, facilitating scientific discovery and enabling transformational science and engineering. For example, making codes run faster and more easily allows researchers to get more science done in a fixed amount of time. Likewise, by lowering the barrier to entry to use of a wide array of digital services, we enable additional research in current communities we support² and new communities of researchers³ who have not harnessed these services to date. Such productivity increases are sometimes the difference between a feasible project and an infeasible one.

XSEDE strives to provide users and developers the perception of a single environment rather than a set of different resources with different administrative domains. That system can include not only local resources connecting to XSEDE, but also any additional resources that the user frequently accesses (e.g. campus systems, DOE systems and collaborators' computers, databases, or instruments). Compute and data resources should be accessible from anywhere in a uniform fashion with no need to manually copy things around or launch a shell from system to system. Productivity for users utilizing multiple sites is also greatly enhanced by such XSEDE features as: single sign-on capability extended to support campus credential-based authentication, submission of a single allocation request with a single review committee for all XSEDE services and resources, a single name space for files, end-to-end tuning to enable the rapid transfer of data from one site to another, and a coordinated system of help desk support, general user support, and extended collaborative services that XSEDE offers. XSEDE also makes it easier for researchers to migrate jobs from heavily used systems to those with more availability. The lowering of usability barriers facilitates new communities' incorporation of higher-performing computational resources into their regular work environment and will unleash new developments in science that these communities are only beginning to articulate.

XSEDE has a very vigorous and, compared to TeraGrid, a much more coordinated effort in Training, Education and Outreach. These also have considerable scientific impact, which is addressed elsewhere in this report (§5.2 and §8).

Scientific advancement on single projects often requires a variety of resources and services which, in turn, requires the availability of comprehensive cyberinfrastructure composed of heterogeneous digital resources. Thus, powerful tightly coupled systems, large shared memory systems, massive loosely-coupled systems linked by efficient workflows enabled by fast networks must all come together in an easy-to-use environment. The actual compute resources that users exploit are operated by independent service providers (SPs). However, many functions that were traditionally offered by the individual SPs are now coordinated XSEDE services, such as the help desk, general user services, extended collaborative support services, and coordinated training services.

XSEDE supports a huge variety of users working in a very large number of disciplines. Since the beginning of the project on July 1, 2011 through the first Quarter of calendar 2012, XSEDE has supported 1,978 projects involving 8,695 distinct users in 32 NSF divisions. The distribution among disciplines is shown in the chart below. Some of these began as TeraGrid projects but all have been smoothly transitioned to XSEDE. XSEDE has reported many examples of important scientific achievements, especially in our quarterly reports. Some selected highlights are reported

² See for example the highlights of Di Matteo and Croft, Karimabadi, and Amaro and Bush below.

³ See for example the highlights of Kuhn and Schatz below.

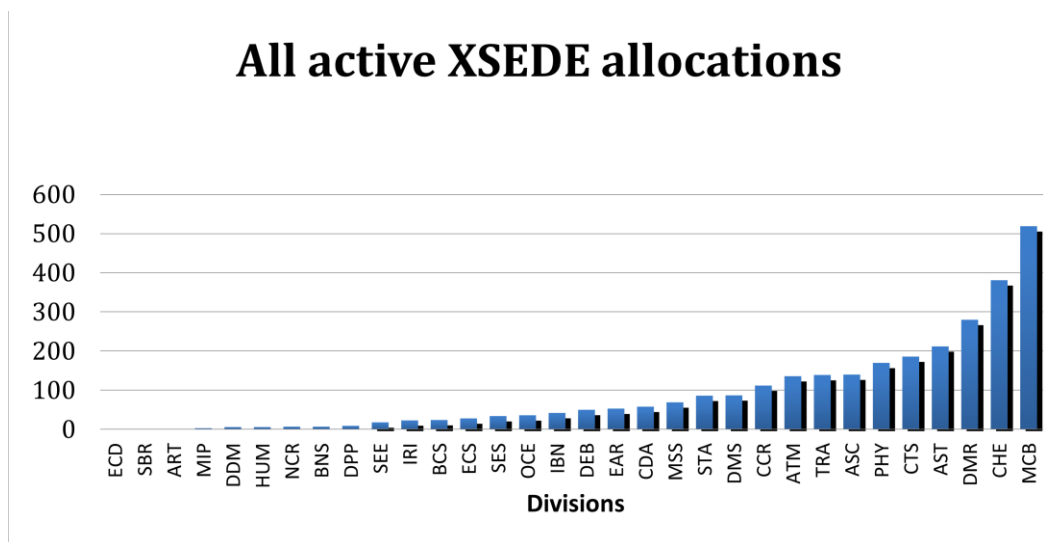


Figure 1: XSEDE allocations by NSF division active since July 1, 2011

below in §2.1. §2.2 gives a sampling of projects awarded resources on XSEDE platforms that may yield significant scientific insight in the coming year.

A novel and potentially revolutionary aspect of XSEDE is our proactive effort in identifying subfields that are not traditional XSEDE users that may be enabled to make transformative breakthroughs with resources brought about by XSEDE’s scale and availability of resources. Driven by the Novel and Innovative Project (NIP, §6.3) sub-activity of XSEDE’s Extended Collaborative Support Services (ECSS, §6 and §7), XSEDE has already identified and reached out to several such areas. By overcoming entry barriers and making these groups aware of new opportunities, XSEDE has already empowered non-traditional areas to reach beyond what they thought they could achieve. Some examples (genomics, digital humanities) are presented in §2.2.

XSEDE has also greatly increased interaction with the Open Science Grid (OSG), the nation’s premier high-throughput computing infrastructure. Soon, users will be able to move rather seamlessly between the XSEDE and OSG infrastructures.

2.1 Program Year 1 Science Accomplishments

Below we include a dozen Science Highlights, most of which have been included in XSEDE’s past three quarterly reports. Some of them represent work begun under TeraGrid.

2.1.1 *Astrophysics: Cold Gas Flows and the First Quasars (Tiziana Di Matteo & Rupert Croft, Carnegie Mellon University)*

With the largest smooth-particle hydrodynamics cosmological simulation to date and a new approach to visualizing the results, astrophysicists solved a puzzle about how some of the first black holes in the universe became supermassive in such a short time. Recent astronomical observations, such as the Sloan Digital Sky Survey, have discovered quasars associated with supermassive black holes in the first billion years after the big bang. While the existence of black holes isn’t surprising, supermassive ones at this early stage of the universe challenge the cold dark matter model. To resolve this puzzle, in 2010, Tiziana Di Matteo, Rupert Croft, and their collaborators mounted a very large-scale simulation with Kraken at NICS. They called their simulation MassiveBlack. Using all of Kraken’s nearly 100,000 compute cores, the researchers simulated a cube of 0.75 gigaparsecs per side (about 2.5 billion light years). Within this volume, MassiveBlack included 65.5 billion particles to represent matter as the universe evolved from 10

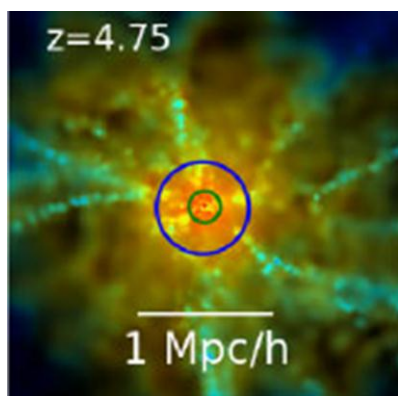


Figure 2: Snapshot from the MassiveBlack simulation at redshift 4.75, approximately 1.3 billion years after the big bang, showing a quasar associated with a supermassive black hole. Gas distribution is color coded by temperature (blue through red). Cold streams of gas (green) penetrate the dark matter “halo” (blue circle) of the black hole (green circle) at the galaxy center.

million years after the big bang through early structure formation and emergence of the first galaxies and quasars to around 1.3 billion years later. To visualize their results, researchers turned to Blacklight at PSC, which made it possible to hold a snapshot of the entire simulation volume—three terabytes—in memory at one time. With help from this visualization, the researchers identified a physical phenomenon that goes far toward explaining the existence of supermassive black holes so early in the universe. They showed that, in regions of high density, gas comes straight into the center of the black holes, and in these places black holes grow very rapidly. This cold gas flow phenomenon had been gaining acceptance as a phenomenon involved in galaxy formation but only at much lower redshifts. This was the first simulation to see this phenomenon at high redshift. XSEDE consultant Anirban Jana helped scale the Kraken code and supervised the data transfer from NICS to PSC.

2.1.2 Molecular Dynamics: UC Irvine Study Points to New Approach to Influenza's Antiviral Resistance (Rommie Amaro, University of California, Irvine)

Researchers from the University of California, Irvine, with assistance from the San Diego Supercomputer Center at UC San Diego, have found a new approach to the creation of customized therapies for virulent flu strains that resist current antiviral drugs. The findings, published online in *Nature Communications*, could aid development of new drugs that exploit so-called flu protein “pockets.” Using powerful computer simulations on SDSC’s Trestles system, UCI’s Rommie Amaro and Robin Bush, together with SDSC’s Ross Walker, created a method to predict how pocket structures on the surface of influenza proteins promoting viral replication can be identified as these proteins evolve, allowing for possible pharmaceutical exploitation. “Our results can influence the development of new drugs taking advantage of this unique feature,” said Amaro, an assistant professor of pharmaceutical sciences and computer science at UCI. The search for effective flu drugs has always been hampered by the influenza virus itself, which mutates from strain to strain, making it difficult to target with a specific pharmaceutical approach. The most common clinical flu treatments are broad-based and only partially effective. They work by interrupting the action of an enzyme in the virus called neuraminidase, which plays a critical role in viral

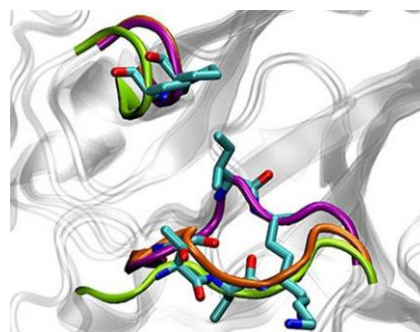


Figure 3: The 150- and 430-loop structures are shown for 09N1 crystal structure (purple), 09N1 second most dominant molecular dynamics (MD) cluster representative structure (green backbone) and VN04N1 crystal structure (orange), indicating that the pandemic N1 adopts an open 150-loop conformation. Gly147, Ile149, Lys150 and Pro431 are shown in stick representation. This simulation was conducted on SDSC’s Trestles supercomputer. (Credit: R. Amaro/UCI, Ross Walker, UCSD/SDSC.)

replication. In 2006, scientists discovered that avian influenza neuraminidase (N1) exhibited a distinctive, pocket-shaped feature in the area pinpointed by clinically used drugs. They named it the 150-cavity. Amaro and Bush, associate professor of ecology and evolutionary biology, conducted research using resources at SDSC and NICS to learn the conditions under which the pockets form. They created molecular simulations of flu proteins to predict how these dynamic structures move and change, as well as and where and when the 150-cavity pockets will appear on the protein surface. This sequence analysis method could be used on evolving flu strains, providing vital information for drug design, Amaro said. “Having additional antivirals in our treatment arsenal would be advantageous and potentially critical if a highly virulent strain, for example, H5N1, evolved to undergo rapid transmission among humans or if the already highly transmissible H1N1 pandemic virus was to develop resistance to existing antiviral drugs,” she added. Walker, an assistant research professor who runs the Walker Molecular Dynamics Lab at SDSC, developed a customized version of the AMBER software, a widely used package of molecular simulation codes, to run these simulations on Trestles under the NSF’s TeraGrid Advanced User Support program. The research involved about 2.5 million CPU hours on Trestles. “We initially used the Athena supercomputer at NICS, which provided us with all the initial comparison data before Trestles came online earlier this year,” said Walker, who is also an adjunct assistant professor in UC San Diego’s Department of Chemistry and Biochemistry. “We had Trestles all ready to go as soon as the first H1N1 protein structure was available, and using the earlier work we did on Athena, we were able to put Trestles immediately to work to conduct simulations of the structure as part of this research.” Robert Swift and Lane Votapka of UCI and Wilfred Li of UC San Diego also contributed to the study, which received support from the NIH and the NSF.

2.1.3 *Biophysics: Large-scale molecular dynamics simulations of anesthetic effects on ion channels (Pei Tang, University of Pittsburgh School of Medicine)*

Despite their frequent use in medicine, how general anesthetics function at the molecular level isn’t well understood. Although there’s been some success in identifying discrete anesthetic binding sites, there’s little understanding about how anesthetics work after they bind to the biomolecules with which they interact. Recent work by Pei Tang has shown how general anesthetics, by binding to molecules called ion channels, can affect the flow of ions that transmit signals in the brain and lead to the loss of consciousness. The general anesthetic isoflurane can bind to a particular ion channel, the *Gloeobacter violaceus* Ligand-gated Ion Channel (GLIC). Little is known about the atomistic changes that this binding produces, which ultimately results in the patient’s unconsciousness. By means of molecular dynamics simulations, Tang and members of her group were able to study the position and motion of GLIC atoms, in the presence and absence of isoflurane. Their results show that isoflurane bound to several locations within GLIC and disrupted the relative position of the five subunits that compose GLIC. They also found that the helical structure that lines the pore of the ion channel changed in ways that contributed to closing off the ion flow. Isoflurane binding also induced strong anticorrelated motions among the five subunits. Knowledge of these structural and dynamical changes induced by isoflurane binding advances understanding of the underlying mechanism of how the anesthetic inhibits GLIC and possibly other structurally similar ion channels and offers a direction toward development of anesthetics more precisely targeted by dosage and with fewer side effects. XSEDE resources Blacklight, Kraken, and Ranger allowed Tang to perform these long-range simulations. The work was published in *Biophysics Journal* in 2011.

2.1.4 *Atmospheric Sciences: Whatever the Weather (Ming Xue, University of Oklahoma)*

Since 2009, scientists from the Center for Analysis and Prediction of Storms (CAPS) at the University of Oklahoma have spent several weeks of spring—the meteorological peak of severe weather in the continental United States—using supercomputers at the NICS to improve real-time

weather forecasts. This past year, led by CAPS director Ming Xue, the team ran—at 4-kilometer grid spacing—real-time weather forecasts on the Cray XT4 known as Athena. The CAPS team also used the Cray XT5 Kraken for another daily forecast at 1-kilometer grid spacing in near real time. Both the 4- and 1-kilometer forecasts covered the entire continental United States. This year’s real-time forecasts dedicated roughly half of Athena’s 18,048 computer processor cores solely to this project. The CAPS team used four numerical weather prediction models in their 2011 forecasting system. Using multiple models, the team is able to produce ensemble forecasts that properly represent the uncertainty in the atmospheric initial conditions and, in the prediction model, consequently producing reliable probability forecasts of severe weather. The 2011 CAPS Spring Forecast Experiment project was the most efficient and productive yet, due in part to help from staff at NICS. A profiling tool, as well as load balance optimization, created by NICS staff members, allowed the research team to monitor the status of their runs in real time, permitting early detection of problematic runs and optimal core usage on Athena. They were able to successfully simulate the outbreak of tornadoes around Joplin, MO, in May 2011. “At 4-kilometer resolution, you cannot resolve the tornado, but you can forecast supercell storms that generate tornadoes,” explained Fanyou Kong, a senior research scientist at CAPS. “This includes some diagnostic parameters, which can show you the location of the highest probability of tornado generation.” A Ph.D. student in the CAPS group conducted some very high-resolution (25-meter) simulations where tornado generation and touchdown were visible, but Kong explained that, even with today’s most powerful supercomputers, real-time forecast of tornadoes at such high resolution is not yet a possibility. The team hopes that, as computational resources grow in power in the future, their real-time weather forecasts will be ready and able to precisely pinpoint these violent storms.

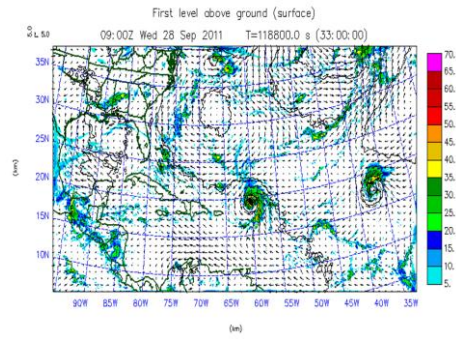


Figure 4: A still shot of an animated movie showing real-time forecasts for the Atlantic hurricane season. Dark swirls of green represent building storms. Credit: CAPS, University of Oklahoma

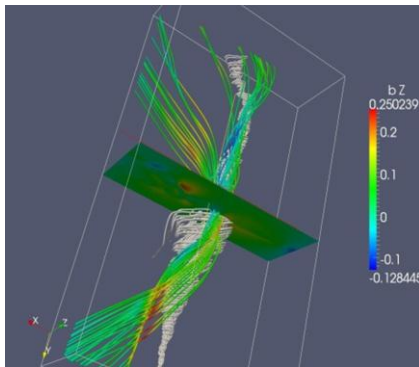


Figure 5: This 3D graphic from visualization on Blacklight shows magnetic-field lines (intensity coded by color, blue through red, negative to positive) and associated tornado-like streamlines (white) of a large flux rope formed due to tearing instability in thin electron layers.

2.1.5 Geophysics: Role of Electron Physics in Development of Turbulent Magnetic Reconnection in Collisionless Plasmas (Homa Karimabadi, University of California San Diego)

Homa Karimabadi works with a team of physicists who have used petascale supercomputing to carry out, for the first time, realistic 3D simulations of magnetic reconnection, the physical process involved in the auroras, when ionized particles blasted into space by the sun crash into Earth’s magnetic field. Some of these solar wind events are perilous to Earth-based electronics, having knocked out satellites and occasionally causing power blackouts. With very large-scale 3D simulations using Kraken at NICS (also Roadrunner at Los Alamos), Karimabadi and his collaborators have been able to characterize—with much greater realism than was previously possible—how turbulence within sheets of electrons generates helical magnetic structures called “flux ropes,” which physicists believe play a large role in magnetic reconnection. In contrast to previous theories, they found that the majority of

flux ropes are produced by secondary instabilities within the electron layers. Through this process, Earth's magnetic field lines break and reconnect in ways that allow solar particles to penetrate Earth's magnetosphere, which under most circumstances serves as a shield against solar plasma. XSEDE consultant Mahidhar Tatineni from SDSC helped to substantially improve necessary IO performance. Karimabadi is using PSC's Blacklight to visualize their simulations, one run of which can generate more than 200 terabytes of data. Blacklight's shared-memory architecture was critical, says Karimabadi, for the researchers being able to solve the physics of flux rope formation. XSEDE consultants R. Reddy and Joel Welling helped materially in the data transfer and visualization on Blacklight respectively. Karimabadi and colleagues reported their findings in *Nature Physics* (April 2011), and they are writing up recent results for *Physical Review Letters* and *Geophysical Research Letters*. Their findings are important for NASA's upcoming Magnetosphere Multiscale Mission to observe and measure magnetic reconnection.

2.1.6 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 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 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 is 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 TACC to conduct their research — the Dell Linux cluster Lonestar, the Sun Constellation Linux cluster Ranger, and the Sun visualization cluster Spur.

2.1.7 *Biochemistry: Insights from Spinning Samples (Borries Demeler, The University of Texas Health Science Center)*

Analytical ultracentrifugation (AUC) experiments spin samples at very high speeds to study how large molecules such as proteins, DNA, and RNA act in solution. As they rotate over time, these materials undergo sedimentation and diffusion processes, revealing aspects of the molecules' behavior under physiological conditions in solution. With support from the NSF, Borries Demeler, associate professor of biochemistry at The University of Texas Health Sciences Center, created the UltraScan software package and the UltraScan LIMS portal to make AUC analysis more effective. The tools allow researchers to analyze their experimental data over the web using advanced computing methods and systems. UltraScan also lets researchers address entirely new classes of research questions and widens the application of the AUC method. Demeler applied his methodology to a collaboration with researchers in Germany, characterizing fluorescent nanoparticles made out of cadmium telluride crystals for use in solar panels. Using Ultrascan, he was able to measure the hydrodynamic properties of the nanoparticles, observe their individual absorption spectra, and correlate absorbance properties with particle size, helping to guide experimental efforts to optimize the material. Demeler's simulations use anywhere between 40 and 14,000 processors simultaneously and speed up the analytic processing by as much as 10,000 times. In 2010-2011, Demeler used 3.5 million computing hours on the Ranger and Lonestar supercomputers at TACC to perform simulations for the open science community. Through an Advanced Support for TeraGrid Applications (ASTA) grant from the NSF, staff at Indiana University helped Demeler develop a web-based gateway where researchers log in, access their data, and submit jobs as if they were running a very simple web application. The portal has made Ultrascan accessible to thousands of lab scientists around the world.

2.1.8 *Stellar Astronomy and Astrophysics: From Star Formation To Supernovae: Astrochemistry Of Turbulent Gmc Models And Studies Of Type Ia Supernovae (Robert Fisher, University of Massachusetts Dartmouth)*

With master's student David Falta and UMass Dartmouth physics colleague Gaurav Khanna, Robert Fisher used NCSA's recently retired Abe and LONI's QueenBee to explore the expected gravitational wave signature from Type Ia supernovae, running first-principles 3D simulations of the gravitationally confined detonation scenario for Type Ia supernovae. Their team is the first to successfully predict the gravitational wave signature for Type Ia supernovae. They not only predicted the expected gravitational wave signature but also determined that it would be detectable for nearby extragalactic events by the next-generation follow-on to the European Space Agency's Laser Interferometry Space Antenna (LISA) or the Big Bang Observer (BBO). If detectable, these gravitational waves will offer a window into the very first second of the Type Ia explosion, information that is not obtainable by studying visible light. They discovered the wave signature by exploring the consequences of the Type Ia supernova explosion instead of just simulating it. By focusing on the consequences, they noticed the explosion doesn't originate from the star's center; it goes off asymmetrically. Recent optical measurements by Keiichi Maeda of the University of Tokyo and colleagues—published in *Nature*—have independently confirmed that Type Ia supernovae are asymmetric. Albert Einstein's Theory of Relativity predicts that, if you take a large amount of mass that is asymmetric, it will give off

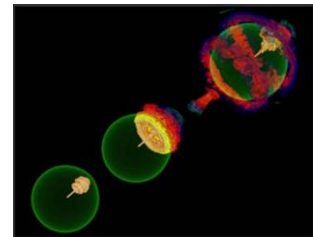


Figure 6: The simulation is initiated near the endpoint of the white dwarf's lifespan, when a single nuclear flame bubble has been ignited near its center. The buoyancy of the bubble rapidly pushes it up to the surface of the white dwarf. Within one second of physical time, the bubble emerges from the surface of the white dwarf. An intense jet sets off a detonation on the opposite end of the star about one second later.

another form of radiation that is not seen by visible light but is actually a distortion in space and time. Type 1a supernovae are a leading candidate for the source of the waves we might detect, and knowing what they might look like makes it much easier to try to find them. The team's results were published in 2011 in *Physical Review Letters*.

2.1.9 Materials Science: Materials for Next-Generation Solar Cells (Richard Hennig, Cornell University)

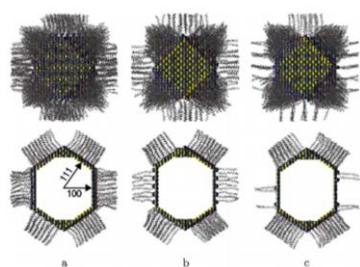


Figure 7: Oleic acid ligands (gray) are shown capping a cuboctahedral lead selenide (PbSe) nanocrystal, illustrating the change in effective shape as the ligand coverage on the {100} facets decreases. The upper panel shows the nanocrystal viewed along the [001] direction, while the lower panel shows a central nanocrystal slice (the interior core atoms are not shown), perpendicular to the [110] direction.

Over the last decade, researchers have created useful nanoparticles for all sorts of industrial applications, however a comprehensive understanding of the design principles behind nanotechnology continues to evolve. One of the most promising areas of nanotechnology innovation lies in new energy applications, where nanoparticles can sometimes perform nano-processes with far greater efficiency than larger molecules. Using high-performance computing systems at the NCSA, TACC, LONI, and the Computation Center for Nanotechnology Innovation at Rensselaer Polytechnic Institute, Richard Hennig and his research team showed that the concentration and location of small molecules (ligands) on the surface of lead-selenium nanoparticles cause the particles to form different shapes, some with markedly improved energy potentials. Lead-selenium and other lead salts are a common and well-studied system used in photovoltaic cells. The results of the study were published in *ACS Nano* in February 2012.

When the individual nanocrystals line up into periodic superstructures with long-range order, they maintain the ideal band gap and electronic properties to produce electricity from the sun. The nanocrystals can form a range of shapes and assemble into different superstructures that are more or less efficient. Hennig's study focused on what controls the shape of the nanocrystals and what controls their assembly. By altering the concentration of ligands present when the nanocrystals form, Hennig and his team produced a range of shapes from octahedrons to cubes with cut-off corners. These results, determined by computer simulations, were later confirmed by laboratory experiments. After World War I, chemists learned how organic polymers derived from petrochemicals could be transformed through industrial processes to make a broad range of plastic products. Today, materials scientists like Hennig are working on a similar problem (albeit at a scale hundreds of times smaller) with the potential to provide an equal or greater benefit to society.

2.1.10 Organic and Macromolecular Chemistry: Exploring New Bifurcation Mechanisms, Radical Ion Chemistry and Copper Chemistry (Donald Aue, University of California, Santa Barbara)

For several years, Donald Aue and his colleagues, Bruce Lipshutz and Liming Zhang, at the University of California, Santa Barbara, relied on NCSA's Abe cluster for their computational chemistry needs. When Abe retired in 2011, the team switched their work to NCSA's Ember and PSC's Blacklight systems. Aue and his team constructed energy diagrams for reaction mechanisms that explained experimental results in the areas of oxidations, epoxide ring openings, and organometallic reactions using the Gaussian and Molpro software. Results of the team's research have been published in the *Journal of the American Chemical Society* and *Tetrahedron*. Other aspects of the team's work, which is funded by NSF and NIH, also was published in 2011 in *Wiley Interdisciplinary Reviews: Computational Molecular Science*. The work published in the

Journal of the American Chemical Society focuses on a novel mechanism for formation of a five-membered ring product **B** and six-membered ring product **C**. Density functional theory (DFT) calculations indicate that this reaction occurs by a “bifurcation” after the initial transition state **TS1**, falling down from the ridge between **TS1** and **TS1a** to form **B** or **C**. In such reactions, product ratios are not predictable from normal transition-state theory so other methods, like trajectory calculations, are used to predict the products. Control of product ratios is a key to the synthesis of new materials, and this new way of predicting the ratios constitutes a dramatic shift from normal methods for mechanisms involving bifurcation. The team’s work published in *Tetrahedron* focused on the ability to predictably change the mode of reactivity via a change in ligand, a goal highly sought in synthetic chemistry. To shed more light on the factors affecting the selectivity of reagents attacking unsaturated ketones, DFT calculations were carried out directly on the transition states for the 1,2- and 1,4-addition to *E*- and *Z*-isomers, as the team believed them to be good models for ethyl derivatives (for which a reversal in selectivity had been observed experimentally). The trends in energy differences for the various additions are consistent with experimental data with the SEGPHOS ligand, but not with other bulky ligands. The team concluded this suggests that these shifts in selectivity likely result from subtle and complex interactions between the substrate, solvent, and bulky ligands, rather than trends inherent in the substrate alone. The team successfully developed a method that uses the SEGPHOS ligand to shift the normal 1,4-addition mode to a 1,2-mode. The 1,2-mode is more useful in some situations in the synthesis of complex molecules.

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

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 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

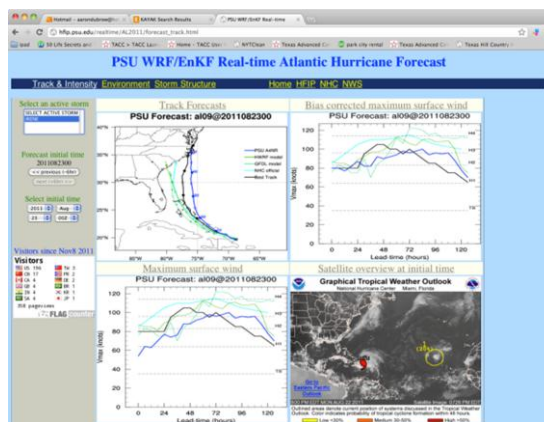


Figure 8: Zhang's [website](http://hfip.psu.edu/realtime/AL2011/forecast_track.html) shows his group's real-time Atlantic Hurricane Forecast (created using Ranger) plotted alongside comparable model forecasts, satellite images of the storm, and other pertinent information..

prediction system is one of a handful the NHC is assessing to become part of the operational forecasting system used in emergency situations.

2.1.12 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)

As climate science advances, climate models are producing larger and larger datasets. Setting the data in motion is the best way to decipher what's to come. "In one sense, all our analysis comes down to the ability to visualize it," said Benjamin Cash, a research scientist at the Center for Ocean-Land-Atmosphere Studies (COLA), one of six international organizations involved in an ambitious climate science undertaking dubbed "Project Athena." The project sought to explore the impact higher resolution has on a variety of climate features. Researchers from COLA used the Cray XT4 "Athena" at NICS to perform climate experiments with two suites of simulation—the Integrated Forecast System (IFS) and the Non-hydrostatic ICosahedral Atmospheric Model (NICAM). Project Athena generated more than a petabyte of data.

Nautilus (an SGI UltraViolet system) at NICS offered the team four terabytes of global memory to analyze their large data set. High-resolution simulations of the Asian monsoon area showed degraded representations of rainfall patterns and magnitude for both IFS and NICAM—an unexpected result that will require more investigation. The monsoon trough is an area of low atmospheric pressure that often causes heavy rainfall when it meets a landmass. "It's known that coarser resolution models cannot even resolve the monsoon trough because it's actually narrower than the model's grid. There's some indication that higher resolution picks up the monsoon trough, but it can't sustain it." Changes in future rainfall were measured using IFS, with sea surface temperatures from a future climate scenario used as input. This simulation was found to be strongly dependent on resolution (15 kilometers), suggesting that high resolution will play a key role in any future climate change simulations. Another area of exploration was the difference in how the two models represent convective precipitation and which method provides a better simulation. IFS employs convective parameterization, meaning that the model adds up statistical representations of variables like humidity and temperature to simulate rainfall. NICAM uses direct representation of convection, a more realistic but computationally intensive endeavor. While preliminary observations suggest that parameterization of convection might never match some observed features, a high-resolution configuration of NICAM still showed weaker performance than IFS in some simulations, including those involving rainfall. Particularly important for future configurations of the IFS model were simulations of blocking over the eastern Atlantic Ocean and western Europe. Climate simulations chronically underestimate the observed occurrence of blocking. Although the precise mechanism by which blocks form and disperse isn't understood, these high-pressure areas in the upper atmosphere have serious consequences on local weather patterns as they can situate in an area for days and even weeks, bringing blistering waves of heat, massive rainfall, and vicious cold snaps. At a resolution of 15 kilometers, IFS was able to significantly improve representation of upper air circulation (atmospheric eddies), which vastly improves the representation of blocking. "A resource like Nautilus will only become more necessary in the future," said Cash. "As climate models reach higher resolutions, the size of the data will grow—you have to provide more than just CPU hours if your customers are going to make the most out of it."

2.2 Sampling of Allocated Projects for Program Year 2

Below is a sampling of the wide variety of projects that will be carried out on XSEDE resources in the coming year. We begin with those from fields not commonly associated with high-performance computing.

Humanities: *Kuhn, USC*. One issue that plagues cinema and new media scholars is the inability to index, tag, and search vast media archives. This includes films that have been digitized, as well as emergent born-digital media. The team will consolidate multiple multimedia databases into one common resource. They will also experiment with automatic metadata extraction of the multimedia resources and with image and video analytics of the media assets. They will also try to provide real-time, on-demand ability to do sophisticated queries, which require the use of computationally intensive algorithms to determine which frames in which assets to return. NIP staff helped Kuhn's team to map their research requirements onto XSEDE resources, including first-ever experimentation with a persistent database on the new Gordon system, and to define an ECSS project, currently under way, to implement this solution. A University of Illinois at Urbana-Champaign team consisting of collaborators from the library (*Harriett Green, Kirk Hess, and Richard Hislop*) and the Institute for Computing in Humanities, Arts and Social Science (I-CHASS) will build a scalable system for library collection analysis and enhanced recommender system for library catalogs and digital libraries that retrieves richer search results from a library collection search based on network analysis of subject relevancy, circulation data of items, and usage data for items that share interrelated subjects.

Genomics: *Michael Schatz, Cold Spring Harbor*. The goal is to assemble four strains of wheat, including one at 16 billion base pairs. The cereal plants, including wheat, corn, and rice, are some of the most significant crops in the world. Wheat, corn, and rice alone provide 60 percent of the world's food intake every day, and just 15 species make up 90 percent of the world's food intake. The sequencing technology now exists to sequence their large genomes at low cost, but it remains a substantial computational challenge to assemble the individual sequencing reads back into their complete genome sequences. One of the simpler species of wheat, *Aegilops tauschii*, has a genome size of ~5 GB, making it >50% larger than the human genome. Christopher Mason, Weill Cornell Medical College, intends to use a different assembler (Trinity) for de novo Transcriptome Assembly for the Non-Human Primate Reference Transcriptome (NHPRT) Resource. Several of the data sets consist of several billion sequence reads, which sum to approximately 6.3TB of raw input data. (By way of explanation, Trinity, along with other assemblers, requires significant shared memory to create and manipulate large de Bruijn graphs using thread-level parallelism.) NIP staff helped the Schatz and Mason teams, among many others, to understand how to use PSC's Blacklight for their projects by identifying the best codes to run and the best ways to run them. They helped define an ECSS Community project, currently under way, to optimize, strengthen, and document these capabilities for the benefit of the next generation genomics community.

The Clean Energy Project (CEP): *Aspuru-Guzik, Harvard*. In collaboration with the IBM World Community Grid (WCG), this group launched CEP to find new organic semiconductors for the next generation of solar cells. This second phase of CEP is screening approximately 5 million small and medium organic molecules (mostly aromatic oligomers) and populating a public database. The results will be employed to determine candidate compounds for experimental synthesis. The WCG is a public grid-computing effort that allows computer-time donors around the world to provide spare CPU cycles for scientific computation. The massive resources coordinated by XSEDE are an ideal complement to WCG computations, allowing the team to perform calculations at a higher level of theory (e.g., coupled cluster or active-space perfect pairing methods) for the most interesting and promising compounds as well as larger molecular systems.

Bioengineering *Keaveny, UC Berkeley*. This project addresses the micro-mechanics and failure mechanisms of osteoporotic hip fractures, as well as spine fractures with particular emphasis on bending loads, typically associated with clinical spine wedge fractures. Another sub-project addresses bone quality issues, again in the overall context of osteoporosis, and focuses on micro-

scale effects such as effect of bone tissue ductility that can uniquely be analyzed with a combination of medical imaging, biomechanical experimentation, and high-performance computation. This work should provide unique insight into the etiology of osteoporotic hip and spine fractures, lead to improvements in non-invasive clinical assessment of fracture risk, direct the development of new drug therapies to target regions within the bone where those therapies will be most effective, and eventually perhaps may help identify which patients will benefit most from such therapies.

Atmospheric Chemistry: *Paesani, UCSD*. This project involves molecular-level simulations of physicochemical processes related to the formation, growth, and fates of atmospheric particles (aerosols). Recent field studies have emphasized the importance of aerosols in influencing the chemical composition of the atmosphere and in altering the Earth's radiative balance, both of which have major implications for climate, ecosystems, and public health.

Quantum Chemistry: *Beran, UC Riverside*. This project focuses on predicting organic crystal structures. How organic molecules pack in the solid state critically affects properties of pharmaceuticals, organic electronics, and many other systems. The ability to predict which packing "polymorphs" will form and to rationally engineer new molecular crystals would have tremendously value. This group has developed a new, affordable, and accurate fragment-based, hybrid quantum mechanical/classical mechanical model for predicting molecular crystal structures.

Molecular Biosciences: *Beckham, NREL*. This project investigates the catalytic function of cellulose degrading enzymes using a QM/MM computational methodology. The research will enhance the understanding of biological routes to convert cellulosic raw material to fermentable sugars and provide strategies for rational design approaches to enhance industrial biofuel production.

Engineering: *Uzun, Florida State U*. This group studies the active microjet-based control of hot supersonic impinging jets, which are an important problem for short take-off and vertical landing (STOVL) military aircraft. In this study, the team investigates pulsed actuators, which generate unsteady microjets with high mean and unsteady momentum for flow and noise control, with a view toward optimal design of impinging jet noise suppression devices. Complementary physical experiments conducted at Florida State University will provide data both on pulsed actuators and uncontrolled/controlled hot supersonic impinging jets for the validation of simulation results.

Particle physics: *Sugar, UCSB*. The *MILC* collaboration is generating gauge configuration ensembles with up, down, strange, and charm quarks using the Highly Improved Staggered Quark (HISQ) action. By significantly reducing lattice artifacts, simulating lighter up and down quark masses, and including charm quarks in the simulations, the HISQ ensembles hold the promise of enabling major improvements in the calculation of a variety of physical quantities of importance in high energy physics, including the leptonic decay constants of the π , K , D , D_s , B and B_s mesons, the semileptonic form factors of D and B mesons, and the mixing of neutral B and B_s mesons with their antiparticles.

Astrophysics: *McKinney, Stanford*. This project proposes to test Einstein's general relativity theory by comparing observations of accreting black holes against radiative transfer results from fully three-dimensional magnetohydrodynamical simulations using different black hole metrics. The team led by PI *Athanassiadou, Arizona State U.*, has undertaken a comprehensive computational project to analyze collisions and mergers of two white dwarfs as a pathway for producing Type Ia supernovae, tracking both the dynamics and the nucleosynthesis emerging from such a scenario. The team led by PI *Blandford, Stanford*, is studying core-collapse supernovae by performing direct numerical astrophysical simulations to determine the important yet unresolved effects of magnetic fields and general and special relativity.

Geophysics: *Pouquet, UCAR*. This group is exploring the role of helicity in turbulent geophysical and astrophysical fluids. Helicity (the correlation between the velocity field and its curl, which measures departures from complete mirror symmetry) affects the energy cascade toward small scales. The team led by PI *Hartzell (USGS)* is conducting an interdisciplinary study of the source process and the resulting landslides of the 2008 magnitude 7.9 earthquake in Wenchuan, China. The research objective is to obtain an accurate description of the time evolution of the earthquake rupture from the available data and using a 3D earth model. Under the leadership of PI *Wright, SURA*, The Super-Regional Modeling Testbed uses XSEDE in the evaluation and readiness testing of marine predictions for key environmental and societal decision-making, such as preparing for, and recovery from, storm surge flooding or seasonal dead zones impacting water quality. This NOAA IOOS-funded test bed is developing processes for assessment of models, with standard metrics that will facilitate transitioning to operations, scenario planning, event reconstruction, observing system design, and further model development.

3 XSEDE Project Office 1.1

3.1 Overview

The XSEDE Project Office encompasses several functions that span the project in important ways. It also is the area through which the project as a whole is coordinated. Program Year 1 saw an incredible amount of activity in this area in the project organizing itself; setting up processes for management, coordination, and communication; and establishing the three key advisory functions to the project—the XSEDE Advisory Board (XAB), the User Advisory Committee (UAC), and the Service Providers Forum (SPF). This effort has begun to bear fruit for the project in facilitating our planning and reporting processes and establishing the key functions for guiding the project as it moves forward.

As we enter Program Year 2, some of the Project Office activities will become routine. Still, several additional elements will come into place, building on the work of Program Year 1. Completing the deployment and adopting the use of advanced project management software will further support these processes. Fully employing the recently established User Requirements Evaluation and Prioritization (UREP) working group will support our requirements prioritization project, which will affect almost every aspect of the project in planning of activities.

With the Architecture and Design (A&D, §3.4) overcoming the multiple challenges of organizing itself, laying out a new plan and defining effective processes for moving forward—something needed following the integration of the original competing proposals—we will see a significant level of architectural work manifested as design documents to be updated quarterly as we detail the XSEDE distributed systems environment. The work has been clearly defined and prioritized with input from all stakeholders, and we have great expectations from this area in Program Year 2, including providing detailed architectural specifications for the areas of Science Gateways, Computing, Big Data, and Connecting Instrumentation.

The activities will be realized through the efforts of the Software Development and Integration (SDI, §3.7) team, which likewise is beginning to realize the benefits of considerable effort in Program Year 1 in establishing a set of processes and practices to support their work. The role of SDI has become better understood in practice to be more than just implementation of the architectural definition, but also to include evolving the currently deployed and operational software and services to deliver essential technical enhancements required to maintain the quality and functionality of the environment. The SDI team members have organized themselves to be much more efficient and agile in integration of new capabilities and preparing them for operational use. Further enhancements to planning processes will include support for the UREP to assign and revise priorities and additional mechanisms to obtain input from stakeholders. While our agile process does not allow for long-term forecasting, features to be focused on early in Program Year 2 in the areas of data services, resource management, security services, and integration services are emerging as priorities. To facilitate this, further engineering enhancements will be put in place to provide automation in our processes to manage activities, we will establish baseline engineering performance metrics to better understand time and effort spent satisfying particular requirements, and we will consolidate and document these processes and support mechanisms for use by others within XSEDE.

We also have great hopes for a reinvigorated effort in our Industry Relations area (§3.6). While this has largely been an uncoordinated effort in Program Year 1, interaction with the XAB on our ideas in this space have helped to drive us toward a strong program in this area. Among other activities, we plan to run an open competition for industry to propose an innovative development project collaboratively with XSEDE to produce solutions useful to industry in enabling them to

harness the capabilities of XSEDE—a program about which the XAB was particularly enthusiastic.

Finally, we will continue to develop our resources and capabilities in making known the existence of XSEDE and its successes. We will finalize our Strategic Communications Plan and act on a renewed commitment to creating strong working relationships with the media. As the project has stabilized, we will create a basic introductory slideshow that will be available to XSEDE staff for presentations and distribution at meetings, conferences, and trade shows.

3.2 Project Management and Reporting 1.1.1

The project management (PM) team members have extensive experience applying project management principles to large, complex, distributed projects. They have developed and will continue to maintain a Project Execution Plan (PEP) that establishes the means to execute, monitor, and administer XSEDE. The PEP will be updated each year to reflect the current project baselines against which the project will be measured and serve as the primary reference document for the project. The PM team selects and adapts appropriate project management practices and tools to accommodate evolving XSEDE project needs and provides the discipline of planning, organizing, and managing resources to bring about the successful completion of project goals and objectives. The PM team provides communication and integrated reporting to all project stakeholders: NSF, XSEDE partner institutions, Service Providers, and the researchers and engineers XSEDE supports. The PM team also manages the baseline change control process, maintains a risk management plan and register, creates and updates a resource loaded project schedule, and manages other sub-projects within the overall XSEDE project.

In Program Year 1 the PM team acquired the Sciforma high-end project management software to provide the team with a more effective tool for managing the unique characteristics of XSEDE as a highly distributed complex project by a distributed team. The acceptance testing and training for this software will complete as we enter Program Year 2. In Program Year 2 we will fully implement this tool for XSEDE-wide use, especially in ECSS, where many sub-projects need to be managed.

Each year the project will determine the schedule, budget, and scope to be accomplished, and this information will be documented at Level 3 of the Work Breakdown Structure (WBS), providing the project baseline from which to measure project accomplishments. Changes to these baselines (cost, schedule, and scope) are managed through a change management process. The project change request (PCR) form was developed during Program Year 1 and will be used to document any mid-year baseline change requests in Program Year 2. Depending on the level and complexity of the change requested, the appropriate level of XSEDE management will review and approve the request. The Project Director may convene a Change Control Board (CCB) to discuss significantly complex change requests.

Risk management is incorporated into the project at all WBS levels. The XSEDE risk tool is used for risk identification and monitoring. The PM team provides expertise and guidance to the WBS managers so they can effectively manage risk for their work scope. As risks are added, changed, or retired, they are captured and updated regularly in a risk register. A risk discussion takes place at each quarterly review as part of the Level 3 WBS managers' presentations.

The PM team, in conjunction with the Level 3 WBS managers, has identified the Program Year 2 activities and deliverables. The duration and interconnectivity of these activities is captured in a schedule baseline subject to the change control process and measureable milestones developed for tracking and reporting progress. Progress against the resource loaded schedule (RLS) will be monitored regularly, and the schedule will be revised as required throughout the year using the change control process.

The project team provides NSF with regular updates via teleconference, as well as formal quarterly and annual reports. The PM team leads the effort to capture the relevant information and produce the required project reports.

3.3 Systems and Software Engineering 1.1.2

Requirements management during Program Year 1 focused on merging and cleaning up XSEDE and XROADS requirements, putting the requirements into the XSEDE DOORS requirements database, and starting work to design a system architecture to satisfy them. During Program Year 2 we will continue to work on implementing the functionality prescribed by our baseline requirements but will also actively seek new requirements to guide future development of XSEDE.

New requirements will come from many sources: from XSEDE staff members; direct requests from users to XSEDE PIs and leaders; requirements derived from the original baseline requirements; from necessary bug-fixes; and, most importantly, from the XSEDE groups that engage in regular communications with users, including User Services (and in particular User Engagement, which conducts regular feedback solicitation activities), E&O, ECSS, and the User Advisory Committee. Systems and Software Engineering (SSE) staff will work with all of these groups to understand and document new needs and capabilities in each of their areas, research and analyze these needs and capabilities and develop them into sufficiently detailed requirements that can be evaluated by the User Requirements Evaluation and Prioritization (UREP) working group.

Precisely defining a requirement is a difficult task for staff and users alike. An effective approach is to ask the staff/users to provide a detailed description, written in normal language, of the functionality or capability they want to achieve. This description, or “use case,” is then used to tease out the requirements associated with the desired functionality. To facilitate this, we will develop a template and guidelines to help users get started with this process. Similarly, SSE will also develop more accessible documentation to help both staff and users understand the requirements process and how they can engage with it and use it to address emerging user needs and new technologies.

In addition to establishing the XSEDE DOORS requirements database and establishing the processes for maintaining its contents, SSE will develop a safe, “read only” access mechanism for the DOORS web client so that a wide variety of XSEDE staff may review any and all requirements information (definition, traceability, activity, etc.) whenever they have a need to do so. Reports from the DOORS requirements database will be generated on a quarterly basis that will provide information such as number and definition of new requirements, identify requirements that have been satisfied, and review traceability between requirements and work being done to satisfy them. This will also enable us to identify requirements that have “fallen through the cracks” or where sufficient progress is not being made. In Program Year 1, we initiated the UREP working group and a regular schedule for electronic meetings will be established early in Program Year 2. These e-meetings will approve new requirements, review changes to existing requirements, and prioritize activities that will determine implementation priorities. During Program Year 1, the XSEDE Advisory Board, the User Advisory Committee, and the SP Forum were all formed. SSE will leverage the expertise of these advisory groups by inviting members to participate in the UREP decision process that will guide the development of XSEDE.

SSE coordinates and assures execution of many processes to ensure consistency in how XSEDE operates. In conjunction with User Engagement, SSE will conduct periodic Usability Panels, at the discretion of the project PI, to assess the effectiveness of what has been implemented by XSEDE. SSE will also revamp the System Engineering Management Plan (SEMP) to reflect the

merger of the XSEDE and XROADS teams and the significant organizational and resource changes that were the result. This document will be simplified to make it both more understandable and reflective of the relationships and responsibilities of the XSEDE groups (SSE, A&D, SDI, Operations) that are responsible for designing, implementing, and operating XSEDE.

3.4 Architecture and Design 1.1.3

The Architecture and Design (A&D) team faced particular challenges starting out the project due to the redefinition of the overall project plan prior to award. Architecture was the area of most significant change with respect to original proposals, and a redefinition of the architecture was created under a very tight time schedule. This incurred what was termed “technical debt”—the eventual consequences of hasty definition of the architecture. The starting days of the project required that we begin to pay that technical debt.

A considerable amount of Program Year 1 was spent in the social engineering project of establishing an A&D team that incorporates participants with fundamentally different philosophical views of architecture yet allows us to develop a single architecture that benefits from the strengths of all views. The architectural description in the Revision Narrative from January gave a reasonably good high-level description, but many details had not been fleshed out. Several areas were left without detail being specified, and there were several areas of ambiguity/disagreement that needed to be sorted out.

After a couple of missteps the A&D team has gelled and is gaining momentum. The team has produced the first architecture definition document and made it publically available⁴. The public-facing XSEDE architecture document includes the Level 1 and Level 2 decomposition of the architecture as well as a Level 3 decomposition for each key area of the architecture (see below). A&D will produce the first successor to this document providing the Level 3 decomposition addressing campus bridging issues before the close of Program Year 1.

Going forward into Program Year 2, the A&D team will continue to document the XSEDE architecture and incrementally release a new version of the public-facing document on a quarterly basis, focusing primarily on additional Level 3 decompositions. These Level 3 decompositions will cover a range of things that XSEDE users (an identifiable stakeholder group) want to be able to *do* with XSEDE resources, each of which can be defined in terms of a set of use cases and scenarios. The current list and order of areas reflects feedback from the XAB. The ordered prioritization of areas is reviewed and updated as necessary on a quarterly basis. The current list of areas and the lead stakeholders in prioritized order 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

⁴ <https://www.xsede.org/documents/10157/281380/XSEDE+Architecture+2012-03-03+Level-1-and-2+ver+1.pdf>

The Campus Bridging Level 3 decomposition will be completed by the end of Program Year 1, but may require some revisions early in Program Year 2 based on feedback from XSEDE leadership and stakeholders. Level 3 decompositions for Science Gateways, Computing, Big Data, and Connecting Instrumentation are also targeted in Program Year 2.

Science Gateways are web portals focused on enabling scientific research in domains on specific topics or sets of topics. Gateways can provide access to relevant events, publications, data, and tools and in some cases also seamless execution of simulations and analyses. *Computing* focuses on capacity, capability and high-throughput models supported by XSEDE service providers. *Big Data* was called out by the XAB as an important area to address early on. It includes data movement, global file systems, storage, archive, replication, sharing, analysis, and tools (e.g., Hadoop). *Connecting Instrumentation* focuses on services and interfaces for collecting data directly from scientific instruments, such as telescopes and remote sensors, and analyzing that data.

The A&D team has adopted the following methodology to create Level 3 decomposition documentation for each area. It is important to note that the process is driven by stakeholder requirements in the form of use cases and quality attribute scenarios that describe the desired behavior and performance. A Level 3 decomposition is not complete until stakeholders agree that it meets their requirements.

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

In order to maintain a holistic and consistent architecture, cross-cutting evaluations are made with support from stakeholders from the XSEDE Security, SDI, SSE, and SYS-OPS teams focusing on security, integration and interoperability, replication and fault tolerance, and all the Level 2 service layer components.

3.5 External Relations 1.1.4

XSEDE communications and external relations are accomplished by a core XSEDE team of science writers, public relations professionals, a webmaster, and designer, and are coordinated by a communications lead, all of whom work for the co-PI institutions. The XSEDE External Relations team built, refined, and provided content for the XSEDE website during Program Year 1. The team also promoted the impact of research and educational efforts enabled by XSEDE through a comprehensive and complementary set of venues and media, including: spinoff websites for XSEDE-related events; informational handouts; booths and displays; media releases; postings on the XSEDE website, Facebook, Twitter, and LinkedIn; two monthly newsletters —

one for XSEDE staff and one for public promotion; an annual highlights book; and support for the annual Supercomputing conference and annual XSEDE conference.

Program Year 1 also included communications support for efforts of other XSEDE teams, primarily Extended Collaborative Support Services and Education, Outreach, and Training, including print materials for use at conferences and the creation of web pages, application forms, and announcements for several workshops, the Student Engagement Program, XSEDE Scholars Program, and Campus Champions Fellowship. The ER team plans to further expand this collaboration to assist additional XSEDE teams and provide more substantive support during the planning stages of projects, helping to ensure that communications efforts are effective and strategic.

New initiatives in Program Year 2 include finalization of a Strategic Communications Plan to be used as a guide for establishing and maintaining effective communications support. Also new in Program Year 2 is a renewed commitment to creating strong working relationships with media who have an interest in the work and successes of XSEDE. Specifically, the ER team plans to aggressively push at least eight large-scale XSEDE stories to national media and smaller-scale stories or story ideas to science and technology bloggers.

Program Year 2 also will include a refinement of XSEDE basic working documents and creation of a basic introductory slideshow, all of which will be available to XSEDE staff for use in making presentations and distributing at meetings, conferences, and trade shows. Finally, the ER team will make a more concerted effort at using social media more effectively and embracing appropriate new promotion venues as they emerge and gain traffic.

Outcomes will be assessed by the level of satisfaction of XSEDE teams with which ER collaborates and by monitoring the number of media pickups of XSEDE stories, the frequency of XSEDE mentions in the media, and traffic trends on the XSEDE website and social media sites.

3.6 Industry Relations 1.1.5

Unfortunately, Program Year 1 saw little progress on XSEDE's larger goals and plans in this area, primarily due to the loss of a key staff member to another project just before the start of XSEDE. Activities supporting industrial partners' users at four SPs (see below) were conducted, but these activities were not integrated fully into the larger XSEDE plan. During our XSEDE Advisory Board meeting on April 13, 2012, we presented the thoughts originally proposed to NSF and asked the XAB members whether this was something strategically important to push forward or if it would be more beneficial to the program to cancel this activity and redirect efforts elsewhere. The XAB members were emphatic that this was an important effort. They further provided some suggestions on how to better craft an integrated XSEDE industry relations program. Here we present the current state of those thoughts, but the XAB also committed to a follow-up discussion regarding this topic. That discussion will help to better define the plans and we will be able to detail the program out at that point.

The four industrial programs engaged in XSEDE–NCSA's Private Sector Program; PSC's Corporate Affiliates Program; TACC's Science & Technology Affiliates for Research Program; and the NICS Industrial Partnerships Program—bring industry to the table to use the centers' advanced resources and services to drive scientific achievement, increase national competitiveness, and improve the economy. This team of centers is superbly qualified to promote and share the benefits of the XSEDE digital technologies and services with industry. The XSEDE partnership will engage industry by leveraging these powerful industrial programs.

In Program Year 2, XSEDE will continue to leverage extensive training classes for industry partners and will develop additional tailored training activities to help industry address its primary challenge in using advanced digital resources: the shortage of trained talent. Industry

partners will be asked to provide input on workforce requirements so that training classes may be further tailored to their needs. The XSEDE service providers will use their current industrial partnership programs as a springboard to reach these firms by coordinating marketing and outreach materials and programs.

In addition, XSEDE will run an open competition for industry to propose an innovative development project that could either be at an algorithmic level, a completely new code written from scratch, or modifying, rewriting, or porting an existing application. XSEDE will commit to providing staff—likely from the Extended Collaborative Support Service staff—in conjunction with a commitment of staff effort from the proposers from industry. In addition, XSEDE will provide access to appropriate computing and data resources to support the projects. The intent is for the outcome of this activity to produce something useful to industry in enabling them to harness the capabilities of XSEDE. The results must also be made publically available. The XAB was very enthusiastic about this and we will be having a special conference call with them to further refine these ideas and plan to release a call in the first quarter of Program Year 2. One initial concept suggested by an industry person on the XAB would be developing a scalable algorithm for *temporal* decomposition as opposed to strictly spatial decomposition available today. This would greatly enhance the scalability of applications and is an area of great need for a multitude of industries. Most importantly, this is a project that companies will participate in openly without conflicts of a proprietary nature.

3.7 Software Development and Integration 1.1.6

The original mission of Software Development and Integration (SDI) was to implement the XSEDE architecture. However, during the first quarter of work we discovered that this original mission was too narrow. XSEDE is not a green field effort because it inherited a TeraGrid production cyberinfrastructure system. SDI's task is more than implementing a new architecture, but also includes evolving the currently deployed and operational software and services to incorporate new architecture features and to deliver a myriad of routine but essential technical enhancements required to maintain the quality and functionality of any large and complex production software system. The proposed Program Year 2 plan for SDI reflects this necessary shift in focus and includes the following:

- an open **planning task** drawing on inputs from multiple project and stakeholder sources to identify, document, and prioritize activities (§3.7.1),
- an **implementation task** proceeding by implementing as many high-priority planned activities as possible within quarterly “time boxes” (§3.7.2), and
- a practice **improvement task** consolidating, documenting, and automating a growing collection of XSEDE engineering processes and best practices (§3.7.3).

3.7.1 *Open, Continuous Planning*

In February 2012 the SDI team introduced the basic component of an engineering planning system, currently implemented via the XSEDE wiki, to support ongoing efforts to identify, plan, and prioritize work. The planning system manages a queue of work (called “activities”) through a lifecycle of stages (proposed, planning, planned), with highest-priority planned activities scheduled as resources become available. One aspect of “openness” is that all information about the work (objectives, requirements addressed, software affected, cost and schedule effort, etc.) is visible. Another is that XSEDE's stakeholders (initially limited to XSEDE partners) can engage with any and all aspects of the planning process, including, ultimately, prioritization of work. Within two weeks of initiating planning, SDI received over 90 proposed system enhancements in areas as diverse as: security, interoperability, user documentation, testing, performance and reliability. This reflects a substantial backlog of improvement work that XSEDE has inherited from TeraGrid.

Several enhancements to the planning capability will be initiated in Program Year 2, including:

- Mechanisms to support the UREP, including the ability to assign and revise priorities; and to identify technical dependencies and assess the impact of new work, new requirements, and new priorities.
- Mechanisms to obtain input from stakeholders, beginning with service providers, including the ability to propose baseline enhancements or new capabilities (a “wish list”), and to provide input about the impact and priority of planned activities.

3.7.2 Continuous Development and Integration

Beginning April 2012, SDI initiated quarterly increments. Activities that have reached the planned state are selected for launch in priority order. Once launched, activities move through another lifecycle (pending, launched, active, complete). SDI launched 16 activities in April 2012, all of which are scheduled to complete before the end of the Program Year 1. Information about these activities and their current state is managed on the staff wiki.

The combination of continuous planning and quarterly increments reflects the basic principles of agile software development. One advantage of an agile approach is that it allows for course corrections in work and quick response to changing stakeholder needs. One consequence of an agile approach is a shorter planning horizon—it is possible to talk about work currently under way and about work likely to be scheduled in the next sprint, but forecasting beyond the next sprint is speculative.

The following activities are likely candidates for launch in the first quarter of Program Year 2:

- Data Services: XSEDE Wide File System (XWFS) deployment; archival replication services; data integrity services
- Resource Management Services: Rich client interfaces; improved prediction services (Karnak); improved Science Gateway integration;
- Security Services: Improved interactive single sign-on (SSO); InCommon integration; improved certificate authority (CA) certificate installers; improvements to MyProxy;
- Integration Services: Improved interoperation (REST and WS* Interoperation); improved system information publishing, discovery, and donitoring; enhanced central service registration; streamlined SP integration.

3.7.3 Engineering Improvements

The planning activity and development and integration practices established during Program Year 1 have had a significant and very positive effect on the efficacy of SDI. While the three initial activities accepted by NSF in the revised narrative took multiple quarters to deliver (primarily due to the lack of foundational TeraGrid engineering processes), the second group of activities initiated during the fourth quarter of Project Year 1 are expected to take 12 weeks. The planning activity (§3.7.1) and development and integration practices (§3.7.2) reflect a range of engineering processes introduced and piloted by SDI in: technical reviews (especially security), technical documentation, software integration testing, configuration management, bug tracking, and coordination of SDI with Operations on acceptance testing. SDI has also introduced automated status tracking and reporting for all planning, development, and integration activities. For example, XSEDE can track which requirements are being/have been addressed by past and ongoing work, which components are/have been affected, and how much time and effort is being invested in particular requirements and components.

Engineering improvements proposed for Program Year 2 will target:

- Substantial automation management processes, including automated status reporting on all launched activities (milestone completion rate, deliverable completion rate, effort expenditure rate).
- Baseline engineering performance metrics, including: investment profile (e.g., time and effort spent satisfying particular requirements, integrating specific technologies, enhancing specific components); per component measures of defect and help desk rates; and quality and usability of end-user and engineering documentation.
- Consolidation and documentation of SDI processes and support mechanisms for use by others within XSEDE, with emphasis on: design review practices; document review practices; technology review practices; synchronized quarterly releases; and a project-wide pipeline to funnel prototypes and technology evaluation through SDI as the preferred “front door” to operational deployment.

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 of unprecedented scale that incorporates a wide range of digital capabilities to support the national scientific and engineering research effort. Significant progress was made in Program Year 1 in operating and maintaining the existing cyberinfrastructure, accomplishing 84% of the stated tasks from the XSEDE project schedule, coordinating and participating with the Software Development and Integration (SD&I) team on the items under development in the XSEDE Software Engineering process, and participating in developing the XSEDE Requirements, Software and Services Baseline document, the Service Provider Definition document, and the Technical Security Baseline document. Incomplete Program Year 1 tasks will be extended into Program Year 2; the delays can be attributed to inaccurate effort estimates, staffing issues, and unaccounted for cross-group communication and interactions (mostly between Operations and SD&I).

XSEDE Operations 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 Account Management
- 1.2.6 Systems Operational Support

Operations will maintain the current level of service and meet the goals of Program Year 2 at the existing staffing level with the small addition of 0.75 FTE to Accounting and Account Management to address some software development issues. Program Year 2 activities for each Operations group are described in the following sections.

4.2 Cybersecurity 1.2.1

The role of the XSEDE Security Operations (Sec-Ops) team is to safeguard the services provided by XSEDE. Sec-Ops coordinates cybersecurity across all SPs, including risk, threat and vulnerability assessments, design and implementation of a risk-based approach to cybersecurity, documenting and enforcing security policies, development of security training materials, scanning for vulnerabilities, monitoring and analysis, preventative measures including security patch management, responding to XSEDE incidents, and participating in both SD&I and Operations configuration item security reviews.

In Program Year 1, Sec-Ops successfully managed the cybersecurity for XSEDE without any major incidents or downtime caused by security incidents. Specific achievements include: 1) development and handing off training materials to E&O and Training, 2) deploying a new practice for XSEDE-wide security log analysis, 3) performing a Risk, Threat and Vulnerability Assessment, 4) design and deployment of an XSEDE vulnerability scanning service, and 5) documenting XSEDE X.509 Certificate requirements and options for moving forward with an upgraded set of services that would serve both the traditional supercomputing access needs of XSEDE as well as the needs of the new Campus Bridging project.

In Program Year 2, the operational aspects of prevention, monitoring, and response will continue, as will coordination of the XSEDE-wide security infrastructure across the SPs. In addition, Sec-Ops will pursue the following:

Engagement with A&D and SD&I: Expand the engagement with the A&D and SD&I teams to collaboratively define XSEDE's security architecture and formalize software security reviews.

Engagement with Campus Bridging: Engage with the Campus Bridging effort, first by documenting the operational security requirements for Campus Bridging activities and then developing a Campus Bridging Security Best Practices Guide.

Risk, Threat, and Analysis Assessment: Transition the Risk, Threat, and Analysis Assessment activities to documenting and updating security policies and designing XSEDE's risk-based approach to cybersecurity.

Two-factor Authentication: Implement a two-factor authentication service for XSEDE in two phases. Phase One is the federation of the NICS and Blue Waters two-factor authentication systems. Phase Two is the development of an XSEDE-wide two-factor authentication system.

XSEDE Certificate Authority: In Program Year 2 we will further refine the XSEDE Certificate Authority requirements and present various implementation options to Operations and SD&I for further evaluation and discussion. Once a strategy has been decided upon, we will move forward with a deployment plan with the goal of having this new set of certificate services in place before the end of Program Year 2. Further, we are involved in ongoing discussions with the Open Science Grid with respect to certificate services and how XSEDE and OSG might collaborate. An initial meeting between OSG and XSEDE on this topic is scheduled for June 2012, and that is expected to help identify where the two projects could align.

The XSEDE Security Officer (XSO) oversees these efforts and provides a single point of contact for both internal and external security. The XSO serves as the responsible party for operational computer security for XSEDE, the advancements described in our defined security thrusts, and coordination with other XSEDE teams. The XSEDE Deputy Security Officer oversees and leads the XSEDE-wide Incident Response Team.

4.3 Data Services 1.2.2

The Data Services team tests, supports, and improves data transfer services, file systems, and archive services for use by the XSEDE user community. The team maintains existing services and works with SD&I to test and deploy new services working within the systems engineering process. In addition, some specific larger goals have been identified for the project in the data services area.

4.3.1 Global File Systems

An important effort in Program Year 1 has been the XSEDE-Wide File System (XWFS) evaluation, involving extensive definition and evaluation of requirements for a single high-performance, distributed file system encompassing storage elements at the leading SP sites and client mount points on the diverse resource types across XSEDE service providers. The next stage of the XWFS effort will encompass beta deployment of a file system involving multiple SPs and allowing for collaboration with the Campus Bridging group. This will include both native and NFS mounts to support the diverse platforms and networks involved in bridging both XSEDE and campus resources, and will use effort in both Operations and SD&I.

4.3.2 Data Transfer Performance Monitoring and Metrics

Two critical non-technical components of providing data transfer services are setting realistic user expectations and provisioning services to meet user demand. In pursuit of these goals and in coordination with the XSEDEnet team, efforts will be undertaken to improve the reach of performance monitoring for data transfers, based on the existing "Speedpage" framework developed at PSC and XSEDE perfSONAR capabilities, to include new data services such as the

Global Federated File System, and to improve the quality and consistency of metrics for data transfers to allow for better tracking of data flows between XSEDE SPs and across the “border” of XSEDE to the campus level. This will allow the service infrastructure to better adapt to the needs of the community.

4.3.3 Archive Replication Framework

An archival replication framework for XSEDE exists in the form of the Data Replication Service (DRS), but many other options exist and not all archive resources are currently included in the DRS. The results of investigations performed in Program Year 1 on options for unified interfaces to archive systems will be used, in cooperation with the SD&I team, to determine whether additional efforts in this area are needed or whether to extend the existing infrastructure to all archive resources at XSEDE SPs, and to implement the plan resulting from the systems engineering process.

4.4 XSEDEnet 1.2.3

The XSEDEnet Operations team will support and improve networking services for the XSEDE user community. The team will maintain existing network connectivity/services and work with SD&I to test and integrate new networking technologies and services into the XSEDE environment.

In Program Year 2 the team will focus on tasks in these areas:

XSEDE Backbone: The team continues to monitor, review, and update the XSEDE backbone network topology to ensure it meets the needs of the XSEDE community. Monitoring tools play an essential role in operational monitoring and capacity planning. A transition is under way in the process for collection of operational statistics due to the replacement of the NLR Network Operations Center. If a reasonable alternative cannot be identified soon, a development project may be required in Program Year 2. Ongoing troubleshooting/problem resolution related to network infrastructure is an important activity of the team.

Network Resources: The XSEDEnet team will document and maintain accurate network resource information on the XSEDE User Portal. Engineers will document what is possible based on the different types of network connections. This ongoing task will be tailored to the information required by the XSEDE community to effectively use XSEDE resources across the wide-area network. The focus for Program Year 2 will be on documenting capabilities needed to support user requirements.

Performance Monitoring: During Program Year 1, the team deployed perfSONAR hosts at XSEDE SP sites. The team will build on that foundation to make performance-monitoring information available to the XSEDE user community. Activities will include documenting existing performance-monitoring infrastructure and determining metrics for managed services, collecting metrics, and reporting on metrics in a user-friendly manner.

End-to-end Performance: The XSEDEnet team will provide expert consulting and tuning support through a special operations subgroup to improve the end-to-end performance of network-based applications. The subgroup will work with XSEDE users to identify and resolve end-to-end performance issues. Solutions may include better documentation on how to tune an end-host, recommendation of specific tools or software, recommendation for configuration changes in the network path to use, or the integration of specific technologies into the XSEDE environment.

Integrating Networking Metrics and Technologies into Applications: The team will work with developers to investigate mechanisms for integrating network metrics into applications. This

is designed to understand how to use metrics in order to allow end-users and applications to make effective use of network resources.

4.5 Software Testing and Deployment 1.2.4

The XSEDE Software Testing and Deployment (ST&D) team tests, deploys, and maintains production versions of software needed to coordinate and operate the XSEDE cyberinfrastructure. This team works closely with the SD&I team to bring items going through the XSEDE software

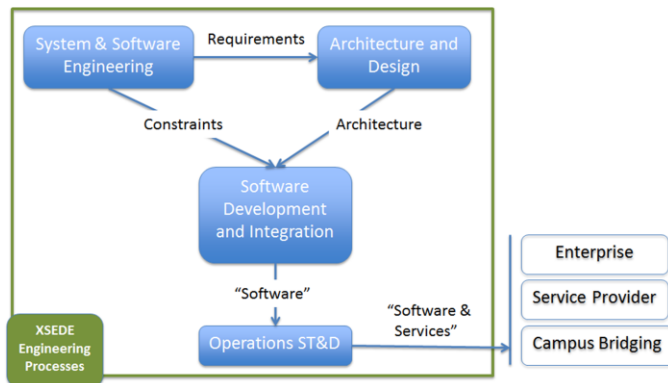


Figure 9: XSEDE Software Engineering process overview.

Bridging. In addition, some new goals for Program Year 2 have been identified based on experiences during Year 1 of the project. These include service provider coordination and campus bridging coordination.

Service Provider Coordination: This activity provides coordination for testing and deployment of XSEDE software to XSEDE Service Providers. This effort also leads the process of integrating new SPs into XSEDE and coordinates the process of deploying new XSEDE software components at the SP level.

Campus Bridging Coordination: This activity provides coordination for testing and deployment of XSEDE software to Campus Bridging partners. This effort also coordinates the process of deploying new XSEDE software components intended for use with science gateways and Campus Bridging projects.

4.6 Accounting and Account Management 1.2.5

During Program Year 2, one of the primary goals of the XSEDE Accounting and Account Management (A&AM) group will be to continue supporting timely processing of proposal submissions for the various XSEDE allocations boards (Research, Startups, Campus Champions, Educational, etc.), and prompt turnaround time for subsequent account requests for these awards. Significant improvements in these areas were made during Program Year 1, and these efforts toward improvement will continue.

Work will proceed on providing enhanced user feedback on the XSEDE User Portal (XUP) with respect to the status of proposal submissions, user account requests, and system utilization – these efforts will be coordinated with the XSEDE Allocations and User Information and Interfaces groups.

Activities will continue on providing enhanced reporting to XSEDE administrative staff, the NSF, and the XSEDE user community. These efforts will be directed toward providing detailed statistics on the various A&AM related metrics, including: numbers of new XSEDE awards across the various awarding boards; requested, recommended, and awarded amounts of resources

at the regular XRAC panels; volume of new user account requests processed; turnaround time on allocation and account requests; resource utilization, etc.

In addition to the ongoing efforts at maintaining and enhancing the basic A&AM services mentioned above, planning is under way to make significant updates and improvements to these basic services. As an example, the Partnerships Online Proposal System (POPS)—the system responsible for handling all XSEDE resource requests and supporting the processing of these requests (initial triage, reviewer assignment, management of board meetings, etc.), and the entry of reviews—is now 12 years old. While the system currently meets most of the requirements for these services, the technology is dated and there are many new requirements that are beyond the scope of the original implementation. In addition, the XSEDE Central Database (XDCDB—the centralized database containing all the definitive information on awards, allocations, resources, users and usage) and the Account Management Information Exchange (AMIE) system are both based on development efforts and technologies from 10 or more years ago and are likely to begin facing issues of scale going forward.

Beginning in the third quarter of Program Year 1, a requirements-gathering process was initiated, directed at identifying areas where these dated systems and technologies can be brought up to date. This process will continue into early Program Year 2 and will result in a set of requirements directed at updating all of these systems. This work will begin in early Program Year 2 and will likely continue into Program Year 3. Ultimately, these efforts should result in a much more streamlined and efficient XSEDE A&AM infrastructure.

Additional XSEDE initiatives that the A&AM group will be involved in during Program Year 2 include:

- Exploring ways to assist Campus Bridging, including possible representation in POPS and the XDCDB
- Continuing to assist in the integration of XSEDE with the Open Science Grid (OSG)
- Additional support for Campus Champions initiatives, including support for varied resource limits in CC proposals, alternate fast-track submission mechanisms, etc.
- Integrating new Service Provider sites into the XSEDE A&AM infrastructure

4.7 Systems Operational Support 1.2.6

The Systems Operational Support team encompasses the core services for XSEDE: 24x7 XSEDE Operations Center and system administration for all centralized services. These services include various web, database, virtual, and account management servers. This group monitors and maintains the underlying servers that host the core of XSEDE's user-facing services.

XSEDE Operations Center: The XSEDE Operations Center (XOC) continues to operate at an extremely efficient rate. Frontline user support and initial ticket triage are essential for user satisfaction. The XOC fielded approximately 7,600 tickets as of March 2012 and was able to resolve approximately 19% of them prior to assigning the remainder to other groups. The XOC, located at NCSA, will continue to maintain the frontline user support role, as the expertise and knowledge base were established there. Future plans for the XOC include real-time system monitoring, XSEDE ticket system transition and staff training, and establishment of a backup operations center at Indiana University, including annual XOC failover testing.

System Administration: Central server system administration continues to be a fundamental function within XSEDE. The system administrators are committed to providing reliable services with an uptime of 95% or better. They have performed well in the transition from TeraGrid and in rolling out new services for XSEDE users and staff. Future plans for the system administrators include continued support for existing central services, installation and implementation of new

services, global authentication implementation for all users/staff, and semi-annual failover testing for all primary services.

Software Administration and Development: SDSC's Inca administrators continue to improve Inca interfaces and overall uptime. As new resources are added to XSEDE, Inca administrators work with the SPs to properly integrate the Inca monitoring suite on those resources. Future plans for Inca include continued support and development, new SP integration, and central Inca failover deployment and testing.

Developers and administrators at Cornell are focused on providing a real-time log-monitoring utility that can accurately predict failures based on correlations in the logs. Future work includes SP testing and integration, application acceptance, and application rollout for all XSEDE resources.

5 User Services 1.3

5.1 Overview

XSEDE User Services activities will continue to grow in Program Year 2 in three fundamental ways: 1) expanding and completing the coverage (in allocations, documentation, training, and user support) of diverse resource types (not just HPC, but HTC, storage, visualization, software, etc.), 2) offering improved guidance for both beginning and advanced users on how to select among XSEDE's diverse resource types for the most appropriate capabilities, and 3) more features for using XSEDE effectively, delivered through the XSEDE User Portal (XUP) directly. Each of the four areas of User Services will continue to evolve through internal team efforts and through increased integration with the Service Providers. The XUP will increasingly become an integration point for every User Services offering and increasing user capabilities.

In Program Year 2, allocations policies and procedures will be developed for storage systems and new guidance will be provided for users of GPU-accelerated and MIC-coprocessor systems. New training materials for these and other resource types will be created, including for the HTC capabilities of the Open Science Grid. New documentation will help users understand the capabilities offered by the increasingly rich set of XSEDE systems and make informed decisions on what to request based on how these capabilities map to their science. There will be revised information on understanding "service units," performance, etc. to account for the new types of machines and to guide allocations requests for these systems.

Documentation will remain a fundamental means of supporting a growing and increasingly diverse—in expertise as well as resource needs—user community, and the Program Year 1 efforts to provide consistent documentation across HPC systems will be complemented by new documentation consistently structured for all other resource types. Program Year 2 also will see the creation of more advanced documentation, building on the strong set of basic documentation developed in Program Year 1. There will be more advanced topics covered in documentation, and this will apply to new training classes as well.

The XSEDE User Portal (XUP) has become the sole means for requesting allocations and for training class signups and the primary means for many other user services: reading documentation, monitoring and managing allocations and usage, updating XSEDE user profile information, etc. The XUP will continue to integrate new documentation, add new capabilities for managing files and data, and explore new opportunities, including social media and other new methods for user engagement and collaboration. Community views and personalization also will be explored, and significant improvements to the mobile interface will be implemented. The XUP is the leading web-based interface for advanced, distributed CI in the world, and its capabilities will be extended in the browser while increasing the subset available through smartphones. The goal is not only for every user to use it but also that it be useful on a daily basis to all current and potential users.

5.2 Training 1.3.1

The Training Coordinator will continue to oversee a comprehensive training program designed to increase the expertise and effectiveness of new and current users and to prepare future users. This training program will include both interactive, in-person training classes that also are broadcast in real time, and online training materials that can be accessed and used anytime, anywhere.

Program Year 2 will bring some shifts in focus, with efforts freed up from tasks completed in Program Year 1. In addition to delivering training content (which was successful, exceeding the proposed level of 50 events per year), much effort was spent, as proposed, putting in place a specific area of the user portal for training that provides centralization and long-term tracking of

training efforts. In addition to providing longitudinal data on training courses and student experiences, this portal has provided a consistent user experience and “one-stop shopping” for XSEDE training. Another significant effort in Program Year 1 was porting existing content into the XSEDE online formats (36 modules are now available online through XSEDE).

Efforts shifted from these tasks will now be spent in several new areas, as promised in the original proposal. Focus will turn to developing additional content driven by user requirements and identified gaps in the available offerings. Program Year 2 will see the advent of several significant new systems in XSEDE, all of which will feature, to some degree, a heterogeneous architecture. Preparing new content for these systems and the architectural trends they embody will be a significant focus. In addition, the Extended Collaborative Support Services has identified a continued need for staff training. These offerings will be made available to the user community, as well.

Now that a substantial amount of content exists, we will focus on developing and delivering training certifications, paths through the training courses that define competencies when completed. This work will continue in close coordination with the Education and Outreach area and in collaboration with other HPC groups, including Blue Waters. The effort also will work closely with Education and Outreach staff to develop shared materials repositories and continue ongoing efforts to provide an ontology and readily searchable index of training and education materials.

The training section of the user portal will continue to see new development. Its success has led to numerous feature requests and expanded its use into other areas, including the collection of survey information, registration of groups in workshops and for education accounts on systems, support for hybrid in-person/webcast courses, etc.

As the portal database of longitudinal information on training activities grows, and as we continue to collect user requirements, we will work with the E&O external evaluator to assess the impact of the training program and individual courses and use this information to adjust course content, assignment of instructors, new course development, etc.

The bulk of the training efforts in Program Year 2 will continue as in Program Year 1: the delivery of high-quality, timely training to as many users as possible, via in-person, web broadcast, and asynchronous online formats. This delivery will continue to expand into discipline-specific areas as begun in Program Year 1 with biology (next-generation sequencing). Future domain-specific training classes will be determined based on user input via User Services efforts, allocations and software usage data, User Advisory Committee recommendations, and other means of evaluating interests and needs.

5.3 User Information & Interfaces 1.3.2

The User Information and Interfaces team is in charge of the development of the XSEDE User Portal and XSEDE mobile interfaces, the preparation of technical information presented through these interfaces and through the main XSEDE website (and supports the overall website implementation), and the delivery mechanisms for user content such as user news (web, XUP, RSS, etc.). In Program Year 2 the team’s goal is to move forward all three of these services to make XSEDE a more constructive and collaborative environment for users.

A major goal is to create a more productive environment for users by increasing communication and collaboration among users and XSEDE staff and by highlighting scientific accomplishments. This includes creating communities via online forums, sharing of files and data, sharing publication information, and helping users become aware of services, events, and upcoming workshops, conferences, and training offerings. The integration of social media services also will facilitate communication and collaboration. Social media has revolutionized how users

communicate with each other and how they make effective use of the services available to them. The challenge is to integrate social media to help advance scientific research. Examples of this include the ability to share social media information among users, easily share and tweet information and receive news, announcements, and notifications via a social media presence.

Another goal is the continued improvement of user documentation. Users will be able to provide immediate feedback on all pages across the site, and improvements will be made to templates and site navigation. Work also will continue on increasing the scope and depth of XSEDE documentation.

The XUP also will continue to expand its services and features. This will be based on supporting new software features that come out of the XSEDE architecture and improvements to existing features such as resource monitoring that further integrate XSEDE services and new Service Providers to offer a complete and cohesive view for users. Another service that we have started to fully integrate into the User Portal in Program Plan Year 1 and will continue to improve is the POPS allocations system. Currently, the allocations effort is still a relatively separate effort as far as users and staff are concerned. Program Year 2 milestones will further integrate POPS into XUP. Interfaces and capabilities will be expanded and improved to give the user a more cohesive experience when it comes to applying for an allocation in XSEDE.

Lastly, a large effort will be focused on improving the mobile interfaces for XSEDE. Users are constantly connected via their mobile devices and rely on them heavily for interfacing with the world. Effort will be put into expanding the XSEDE mobile services, based on user feedback and requirements, to develop new features that expand the mobile presence of XSEDE and assist users in being productive on the go. Furthermore, the team will continue its ongoing support of additional communities and XSEDE activities such as Campus Champions, student fellows, Campus Bridging, and service provider forums through the maintenance and development of features such as calendars, applications for services offered, and more, based on requirements.

5.4 User Engagement 1.3.3

XSEDE User Engagement is organized as two working teams: Feedback and Consulting. The Feedback team focuses on gathering, recording, and mining information obtained from interactions with the XSEDE user community to distill the needs of stakeholders. The Consulting team coordinates seamless support of the user community across XSEDE. Requirements derived from consulting and feedback efforts are mapped to efforts within XSEDE to guide improvements to the resources and services offered.

5.4.1 Feedback

The feedback team will focus on gathering, recording, and mining information obtained from interactions with the HPC community through surveys, forums, focus groups, usability panels, interviews, and email submissions to feedback@xsede.org. An annual user survey will be conducted in the third quarter of each project year to gauge user satisfaction with XSEDE services and to aid in understanding the current and future needs of the users. The feedback team will also coordinate smaller, targeted surveys to explore issues of interest to elements of XSEDE, as needed. Online user forums will provide a 24x7 source for community-provided support that will be monitored by XSEDE User Engagement during business hours. Birds-of-a-feather forums will be held annually at the Supercomputing and XSEDE conferences to provide opportunities for face-to-face contact and open discussion between XSEDE staff and users, and focus groups will be held quarterly to allow the feedback team to explore specific topics with a targeted group of users. Usability panels will be held, as needed, at the request of the XSEDE Project Office, to assess the usability of proposed changes and additions to XSEDE services, in conjunction with the activities of the XSEDE Technology Insertion Service. One-on-one interviews will be used,

as necessary, to supplement other feedback efforts and to follow up with specific users to obtain a more in-depth understanding of their needs, concerns, and recommendations. The feedback team will mine the data gathered through these activities and from submissions to feedback@xse.de.org to produce reports that will be posted on the XSEDE staff wiki and focused notifications that will be sent directly to relevant XSEDE staff. Finally, the User Engagement staff, and the entire User Services team, will be directly involved in evaluating and implementing many of the suggestions from the User Advisory Committee (UAC), as well as some of the guidance from the XSEDE Advisory Board.

5.4.2 *Consulting*

The consulting team will monitor consulting activities across XSEDE and deploy new policies, procedures, and software to ensure that seamless, high-quality support is delivered to users in a timely manner. In addition to directly handling XSEDE tickets that are not clearly associated with the activities of a specific service provider site or XSEDE team, the consulting team will monitor the progress of all XSEDE tickets and quantify the provided quality of service by collecting and analyzing metric data. Members of the consulting team or the Consulting Coordinator who oversees the team will address any service problems related to the XSEDE consulting process. The XSEDE ticket system will be modified, as necessary, to enable accurate tracking of both help and feedback tickets and to facilitate automation of monitoring processes. The consulting team will devise new policies and procedures built around the XSEDE ticket system to improve the efficiency and efficacy of the consulting process, and members of the team will train consulting staff across XSEDE in their use to ensure uniform adoption and application. A customer relations management (CRM) system will be developed (mostly by integrating existing XSEDE tools, with some additional software choices and development) and deployed to provide comprehensive tracking of all XSEDE interactions with users, including training, allocation requests, user consulting calls and tickets, extended collaborative support projects, workshops and conferences, and any other relevant data to ensure that support staff across XSEDE have quick, easy access to any information necessary to provide comprehensive, personalized support. The consulting team will mine ticket data and metric data to produce appropriate feedback and recommendations for XSEDE and the service providers, facilitating the continued improvement of XSEDE services as the project progresses.

5.5 Allocations 1.3.4

The XSEDE allocations process will continue to provide the simplest, most flexible and most equitable allocations procedure possible. The team will continue to focus on the three main objectives from the original plan:

1. minimize the barriers faced in applying for resources
2. provide users with access quickly whenever possible
3. maintain an appropriate level of review for scarce, valuable resources

XSEDE allocations staff will continue to foster excellent working relationships with other XSEDE groups, mainly the accounting and XUP teams. Allocations staff will work with these groups to revamp the entire allocations front-end and back-end process of allocations and accounting.

The allocations team expects to usher in several new compute resources. With this comes the need to re-think both how to measure and how to allocate a “resource unit” currently known as a “service unit.” The unique architectures of the new compute resources will require allocations staff to work closely with service providers to explain to the user community the definition of the new “resource unit.” For example, we will be looking to define the service unit for Keeneland,

the GPU cluster coming online at NICS, in terms of GPU hours, as opposed to CPU core hours, as it has traditionally been defined.

The XSEDE allocations process will continue to provide the current levels of allocations: Small (Startup, Education, Development, Campus Champions, and Research) as well as Large (XRAC). The team is also working to put in place Standard allocations, significant but fixed amounts of resource or service allocations specified for each resource or service. Standard allocations are expected to reduce the effort by PIs to prepare requests, shorten wait times on awards, and lessen demands on the XRAC review panel.

The allocations team will continue to conduct quarterly webinars (coordinated through User Services Training) focused on the allocations process. Held prior to each proposal submission window, these webinars are designed to assist users in the preparation of successful proposals. Also, in concert with the User Engagement team, Allocations will continue to maintain a dialog with Campus Champions, providing updates on allocations processes and procedures and identifying issues Campus Champions report on behalf of their users.

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 ECSS 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 from 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 dataflow systems. ECSS staff members also participate in reviewing adaptive proposals associated with XRAC meetings.

We divided ECSS efforts in two areas: Projects, headed by Ralph Roskies; and Communities, headed by Nancy Wilkins-Diehr. These groups have very close interactions, with common Project Management support. All told, ECSS consists of 37 FTEs, spread over ~80 people at about one dozen sites.

The “ECSS-Projects” area consists of ESRT (Extended Support for Research Teams) and NIP (Novel and Innovative Projects). 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. There will be at least 20 new ESRT projects in Program Plan Year 2, each with documented work plans. 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. Our goal for Program Plan Year 2 is 20 new projects in non-traditional areas.

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 up to 4 flexible hires for a year at a time. In Q3 of Program Plan Year 1, we identified such an area (Digital Humanities) and made our first flexible hire.

In addition to staff members’ participation in the regular training activities of User Services, we recognized in Program Plan Year 1 that we need to increase initiatives in the area of staff training. One way of addressing that was our monthly symposium, at which we had two web-based presentations either on lessons learned from ongoing projects, or an introduction to new capabilities that are being planned for XSEDE that ECSS consultants should be aware of. Also, together with the Blue Waters team we are planning a workshop on Extreme Scaling, to be held July 15-16, just before the XSEDE12 conference in Chicago.

Finally, we will put more effort in Program Plan Year 2 into performing staff evaluations.

6.2 Extended Support for Research Teams 1.4.1

An Extended Support for Research Teams (ESRT) project is a collaborative effort between an XSEDE user group and one or more ECSS staff members, the goal of which is to enhance the research group’s capability to transform knowledge using XD resources and related technologies. Typical ESRT projects will last for six months to one year and might include the optimization and scaling of application codes to use 10,000–way parallelism 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.

A request for ESRT support will be made by the principal investigator of the research team via the resource allocation process. For requests recommended by reviewers, an ESRT lead

consultant is assigned to work with the PI to discuss the potential project and develop a work plan if appropriate (assuming staff resources are available.) The work plan will be up to one year in length, and will have milestones for each quarter. The work plan is submitted to the ESRT project manager for approval. It will include a staff assignment from the pool of available advanced support experts who have the necessary skills. The ESRT lead consultant, working with the ESRT manager and ECSS project management, will be responsible for project tracking and reporting. This team is also responsible for making changes to the work plan as appropriate and requesting additional resources or assistance from XSEDE management as needed.

The XSEDE proposal calls for at least 20 new ESRT projects each year with documented work plans. We believe that we will exceed this number in Program Plan Year 2, as we have done in Program Plan Year 1. All ESRT projects will have quarterly reports as well as a final report. These are maintained on the XSEDE wiki and are available to all XSEDE staff. Useful lessons learned will be publicly disseminated and fed to the Training arm of User Services. Managing the assignment of ESRT personnel to projects has been a challenge because the spreadsheet tools we had to identify consultants with the availability and expertise to undertake new projects, and to track the progress of projects assigned to particular consultants, were very crude. These inefficiencies left managers less time to actually manage projects. In Program Plan Year 2, we are hoping that the new project management software will enable us to carry this out more efficiently.

ESRT staff members will be required to participate in, and occasionally give presentations to, the monthly ECSS symposium as part of their ongoing education. They will occasionally review adaptive XRAC proposals.

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 (minority serving institutions) and EPSCoR (Experimental Program to Stimulate Competitive Research⁵) institutions, and SBIR (Small Business Innovation Research⁶) recipients. Bringing these communities to XSEDE has already begun 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. The implementation of campus bridging processes and technologies will be particularly important for these communities.

In Program Plan Year 1, our main strategy for developing such projects has been to leverage “bridgehead” researchers, institutions, and programs with an awareness of the benefits of HPC/CI, by asking them to introduce us to other members of their community whose projects are likely to benefit as well. For example, we have generated more than 20 new genome assembly startup grants by initially working with a handful of existing XSEDE users who blogged about their positive experiences, and by forming a partnership with the influential code developers at the Broad Institute (see www.broadinstitute.org). In computational economics, an interview with a previous user of TeraGrid led to an invitation to present a tutorial at a prestigious summer school at the University of Chicago -- an opportunity to introduce HPC/CI to top graduate students, which will be repeated in 2012. In the areas of digital humanities and arts, we began by

⁵ <http://www.nsf.gov/od/oia/programs/epscor/about.jsp>

⁶ <http://www.nsf.gov/eng/iip/sbir/>

assisting a few existing projects that needed some help, and built on their success to establish relationships with influential programs such as I-CHASS, which are now attracting additional projects to XSEDE.

We plan to leverage these relationships and to build new ones following this model. An important enabler will be the workshop we are organizing at the 2012 IEEE eScience conference in October (see <http://www.ci.uchicago.edu/escience2012/workshops.php>), in collaboration with leaders in digital humanities, computational economics and social science, citizen science, GIS, genomics, and public health. Its goal is to discuss examples of successful projects, as well as barriers and practical approaches to overcoming them. This will lead to an improved understanding of actions that should be taken by various stakeholders in order to enable a wide spectrum of practitioners to use HPC/CI resources as part of their work and dataflows, and to establish an informal network of people and communities interested in this outcome. In doing so, we are confident that we will create at least 20 new projects in non-traditional areas in Program Plan Year 2.

Our main challenge is to balance our ability to attract novel projects with our ability to help them succeed in becoming efficient users of XSEDE resources. This requires matching the expertise of ECSS staff to the characteristic requirements of NIP, which often differ from those of traditional HPC applications. In Program Plan Year 1, we completed the first ECSS flexible contract hiring process to bring in an expert in Digital Humanities. 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. In Program Plan Year 2, we will continue to monitor the need to rebalance the ECSS staff expertise pool and, if necessary, recruit additional contract positions to meet the support requirements we discover as we interact with new communities.

As part of our team members' continuing education and adaptation of their skills to the evolution of XSEDE technologies and community requirements, we will actively promote and require their participation in, and occasional presentation to, the monthly ECSS symposium, and their participation in the reviewer pool for adaptive XRAC proposals.

7 Extended Collaborative Support Service – Communities 1.5

7.1 Overview

Extended Collaborative Support for Communities focuses efforts on collaborative projects that benefit larger user communities rather than individual research groups. The Communities area has three components—Extended Support for Community Codes (ESCC); Extended Support for Science Gateways (ESSGW); and Extended Support for Training, Education and Outreach (ESTEO). 16.5 FTEs (nearly half of the total ECSS staff) spread across 63 individuals are devoted to the Communities area. While this is a large number, ECSS benefits from involvement of staff in additional programs such as PetaApps, Blue Waters, and their own grants. This leverage of expertise allows us to retain top-notch, well-qualified staff members. Also, ESTEO activities tend to be shorter-term and allow for the involvement of larger numbers of staff for shorter time periods.

An ESCC project (§7.2) is a collaboration whose goal is to port or optimize a community code for XSEDE-allocated compute systems. In Program Year 2, the team will first conduct an assessment of community codes installed throughout XSEDE and make sure these are catalogued and advertised. ESCC projects can be both requested by users and initiated by staff. The User Advisory Committee and the XD Technology Insertion Service will provide guidance as staff evaluate the many NSF-funded software projects for suitability within XSEDE.

The ESSGW area (§7.3) helps project teams develop and deploy science gateways that use XSEDE resources. Science gateways are community-designed interfaces (often web-based) that provide a user-friendly interface to many services. Depending on the community, these might include data collections, user work spaces, collaboration spaces, and computation and visualization capabilities. The ESSGW teams are experienced in the integration of high-end resources into gateways and help with user-requested projects. In addition, gateway staff work on infrastructure tasks that benefit all gateways. The team is looking at application programming interfaces (APIs) for interfacing with the XSEDE infrastructure, and is also working closely with the Architecture & Design and System Development and Integration teams to assess readiness of the XSEDE infrastructure for gateway use.

The ESTEO area (§7.4) provides all of the staffing for training, education, and outreach activities across XSEDE. In Program Year 2, staff members will continue to create and review online tutorials; travel to conferences, campuses and user sites to deliver training classes and presentations about XSEDE resources and services; give remote webinars; host and mentor students; and much more. As the need for diverse expertise (from petascale computing to visualization of large datasets to high-throughput computing) expands, staff training and development is all the more important. The majority of staff members in the ECSS program contribute to ESTEO activities.

ECSS has many incoming demands. User requests for extended support come in continually through startup requests. Every quarter, larger numbers of requests are received through the XRAC allocations process. Staff assist in review of adaptive proposals (user requests for time that are not large enough to be considered by the full XRAC committee) and routinely participate in reviews and delivery of training materials as requested by the User Services Training and E&O arms of XSEDE. ECSS management works to continually monitor the load on staff and assess the team's ability to tackle proactive work in ECSS. These types of activities are described below.

7.2 Extended Support for Community Codes 1.5.1

Extended Support for Community Codes (ESCC) efforts are aimed at deploying, hardening, and optimizing software systems necessary for extensive research communities to create new knowledge using XSEDE resources and related technologies. ESCC projects are focused on helping users with community codes and tools on XSEDE systems. Priority is given to NSF-funded software initiatives such as PetaApps, Software Infrastructure for Sustained Innovation (SI2), Software Development for Cyberinfrastructure (SDCI), and Strategic Technologies for Cyberinfrastructure (STCI).

ESCC projects may be proposed in several ways. During Program Year 1, most ESCC projects were initiated by requests for assistance during the resource allocation process similar to ESRT projects. Projects that request assistance with community codes are allocated to ESCC rather than ESRT. Also, a couple of projects were initiated internally based on a recognized need for widespread support. Although ESCC will continue to consider projects proposed via these methods, we will also support projects suggested by the newly formed User Advisory Committee (UAC) and the XD Technology Insertion Services (TIS) team.

In addition, ESCC will work with the SPs to determine which community codes are used most extensively and to catalog and advertise them. These community codes will be prime candidates for internal projects. Furthermore, an effort will be made to provide users with sample scripts, benchmark results, and documentation for the most popular community codes. For these codes, users will have access to repeatable test cases, documentation tailored for specific XSEDE systems, sample batch scripts that demonstrate the proper way to run the code, and benchmark results to provide an expectation of performance. Staff activities will need to be balanced between activities that are not initiated by user requests and those that are.

ESCC will continue to use the project management format developed during Program Year 1 in which each project has a work plan describing its goals. ESCC will have two types of work plans: a normal work plan (similar to an ESRT work plan) and an initiated work plan. An initiated work plan will have an ESCC staff member as the PI to coordinate the project and provide updates to multiple XSEDE research projects dependent upon the effort. A list of supported research groups will be included in the work plan.

The goal is to continue to support at least 10 new ESCC projects per year. Each project will provide quarterly reports as well as a final report to the ECSS management team. Lessons learned are shared through the monthly ECSS Symposium.

7.3 Extended Support for Science Gateways 1.5.2

A science gateway is a community-developed set of tools, applications, and data that are integrated via a portal or a suite of applications, usually in a graphical user interface, that is further customized to meet the needs of a specific community. Gateways enable entire communities of users associated with a common discipline to use national resources through a common interface that is configured for optimal use. Researchers who use gateways can focus on their scientific goals and less on assembling the cyberinfrastructure they require. Gateways can also foster collaborations and the exchange of ideas among researchers.

The primary mission of the Extended Support for Science Gateways (ESSGW) Group is to provide Extended Collaborative Support to existing and new scientific communities that would like to use XSEDE resources through gateways. Support is requested through the allocations process and lasts from months to up to one year. ESSGW will continue to assign staff members as these requests are approved by the allocations committee during Program Year 2.

Often the ESSGW members working on these projects can benefit from lessons learned from other gateway teams. The group holds weekly science gateway telecons. Half are devoted to internal activities, but alternate weeks feature gateway speakers both from within and outside XSEDE and are advertised via XSEDE news. Mailing list discussions, a gateway track at the XSEDE conference, and the Gateway Computing Environments conference at SC all help build community. Proceedings from these events are published.

The plans below are listed in priority order.

- Priority activities in Program Year 1 included contributing gateway requirements to the Architecture & Design team (§3.4). This very important work continues in Program Year 2 as ESSGW continues to engage the gateway developer community in documenting the use cases and working with Architecture & Design and System Development and Integration (§3.7) teams in creating XSEDE use cases and filling in a system requirements matrix. ESSGW will also develop tests to help determine whether architectural elements will perform at the scale needed by gateways. Reliability at scale is very important for gateway developers.
- ESSGW will help document portions of the XSEDE architecture particularly relevant to gateways, including job submission and certificate management. The group will proactively work with end gateway developer communities and to analyze the programmatic interfaces (APIs) required to efficiently develop gateways and seamlessly maintain them.
- The ESSGW effort will also resurrect the TeraGrid gateways effort in providing attribute-based authentication to automatically submit gateway end-user attribute information with each job submitted to XSEDE resources. This effort will require coordination with the Service Provider Forum and Operations groups to capture this data via the job logs sent to the XSEDE Central Database, and will allow programmatic counting of gateway end-users that will enhance our analysis for improvements to the ESSGW effort.

7.4 Extended Support for Training, Education and Outreach 1.5.3

ESTEO staff provide the person-power, expertise, and technical content for activities described in the User Services Training area (§5.2) and in the Education and Outreach areas (§8). These areas coordinate training, education, and outreach activities across all XSEDE partners and Service Providers to reach out most effectively to a diverse community of users and potential users, to ensure broad coverage of pertinent topics, to engage the best and brightest experts, and to minimize duplication of efforts. Activities are selected based on user requirements analysis and feedback from advisory groups. E&O is also significantly involved in STEM workforce development and curricula development. ESTEO assists by providing and reviewing technical content. ESTEO consists of 4.4 FTEs spread across 45 individuals. This breadth is possible because of the short-term nature of most training activities when compared to the longer-term involvement in project support. Because ESTEO involves many ECSS staff, lessons learned through the projects are being passed on to others.

In Program Year 1, ESTEO team members participated in dozens of events, including “train-the-trainers” events, onsite classes requested by Campus Champions, conferences, summer schools (including those in collaboration with PRACE), and extreme scalability workshops (in collaboration with Blue Waters). The staff also creates and reviews online documentation and training modules. This just-in-time training is increasingly popular with the user community when both time and travel budgets are limited.

Program Year 1 saw an unprecedented expansion of the types of expertise needed to address user requests for extended support. A new ESTEO focus in Program Year 2 will be on intensive internal training seminars to identify and train those who will become experts in newly requested

areas such as genomics, bioinformatics, high-throughput computing (Condor and Open Science Grid), data analysis, and digital humanities. We envision combining this training with planned time for hands-on work that will result in staff members gaining in-depth expertise at a level that will allow them to serve as project experts.

Several new systems with specialized architecture will be coming online in Program Year 2,, prompting new training materials. The GPU system Keeneland, the experimental grid FutureGrid, the high-throughput Open Science Grid, and the Stampede system with Intel's Many Integrated Core (MIC) technology will all be new or relatively newly allocated XSEDE systems in Program Year 2.

A second major focus in Program Year 2 will be supporting the work of Jim Demmel at UC Berkeley. His work on graduate education complements XSEDE efforts in the undergraduate space, while benefiting XSEDE efforts to assist other institutions as they develop undergraduate and graduate curricular materials. ESTEO staff will test and document examples for his computer science class on XSEDE resources.

Also new for XSEDE is the Campus Champions Fellows program, which seeks to extend ECSS expertise to the community. Program development occurred in Program Year 1, culminating in a program launch in April 2012. Operation of the program will begin in earnest in Program Year 2. Campus Champions are non-XSEDE-funded local university representatives who are knowledgeable about XSEDE resources and who engage large numbers researchers, educators, and students among their campus community. Fellows are paid a stipend to facilitate their engagement in an intensive one-year collaboration with an ECSS staff member on user projects where ECSS support has been requested. The ultimate goal is to strengthen the user support fabric that is being woven throughout the country via the Champions initiative. Eight fellows will be selected per year. The program structure was developed over many months through an in-depth collaboration with ECSS management and the Campus Champions leadership team. The program will be evaluated by the E&O external evaluation team. At this writing, the application period is closing and initial applicants are being reviewed.

8 Education and Outreach 1.6

8.1 Overview

The Education and Outreach (E&O) team had a very productive first year. All of the activities begun in Program Year 1 are continuing in Program Year 2 with modifications based on the lessons learned from the first year. The E&O team plans for Program Year 2 also have benefitted from the feedback from the external evaluations, the comments and recommendations from the XSEDE Advisory Board, and the recommendations of the TEOS Advisory Committee. In addition, the E&O team conducted a new round of collecting requirements from the community which have helped to refine the plans in Program Year 2.

The E&O team will seek advice from the advisory groups, seek community requirements, and gain a deeper understanding from the external evaluations of the E&O programs and offerings. E&O will benefit tremendously from the expertise, knowledge and collaborative efforts of all XSEDE teams. The commitment of ECSS/ESTEO staff is essential to the success of Education and Outreach, and the demands are expected to increase from about 4.5 FTE in Program Year 1 to the full 6.0 FTE complement prescribed in the original proposal during Program Year 2.

E&O will build upon and expand external collaborations to leverage activities with other groups pursuing complementary efforts. The E&O team will expand upon its collaboration with PRACE to include other continents in the HPC Summer School in 2013 and beyond. The E&O, User Services and ECSS teams are building a new collaboration with the HPC Wales team, an effort that has been in the formative stages for nearly two years. In addition, E&O will pursue more in-depth collaborations with the Blue Waters project as it evolves from deployment to operations.

Planning for XSEDE13 has begun and a hotel contract is in place, at significant savings to XSEDE. The committee has been formed with Nancy Wilkins-Diehr as the Conference Chair and the XSEDE12 deputies as the XSEDE13 committee chairs, which will significantly advance the planning. The conference will be held July 22-25, 2013, in San Diego.

Since XSEDE began, 33 new education projects, 332 new start-up projects were allocated. XSEDE is successfully engaging new communities.

8.2 Education 1.6.1

The education team will build on the activities started in Program Year 1 and add some additional tasks to further disseminate computational science education materials and help to build computational science programs. The major tasks can be broken into four parts: revision of computational science competencies; assistance in creating formal computational science programs; education courses and workshops for students and faculty; and creating a virtual computational science education community resource.

8.2.1 *Revision of Computational Science Competencies*

During Program Year 1, revisions were made to both the undergraduate and graduate computational science competencies (<http://www.rscs.org/competencies>) based on the feedback from faculty and professionals. The undergraduate competencies were updated to reflect the infusion of multi-core and many-core computers and the requisite need to understand parallel computing earlier in the curriculum. The core graduate competencies were refined and their general format accepted, but examples of domain-specific requirements must still be developed. The competencies will include recommended and optional components. E&O expects that, as we assist institutions with the implementation of specific programs in computational science, the institutions will adapt and add to the competencies to meet their needs. We will use those

examples as opportunities to document and update the competencies, as well as to define graduate competencies for major domain areas.

8.2.2 Assistance in Creating Formal Computational Science Programs

During Program Year 1, we engaged 14 different institutions in discussions concerning the adoption of formal computational science programs. We currently are working with 10 institutions on the creation of programs. The other four have indicated persistent interest but have deferred activity until fall 2012 because of changes in personnel or administration. The institutions that have been active this year are Clark Atlanta University, Kean University, Montgomery Community College, Norfolk State University, North Carolina State University, North Carolina Asheville, Southern University, Stockton College, The Ohio State University (OSU), and University of North Carolina. In Program Year 2, Clemson University, University of Mary Washington, University of Arkansas, and North Carolina Central University may become active. We plan to work with all of these institutions, helping them to move toward formal approval of academic programs.

8.2.3 Education Courses and Workshops for Students and Faculty

There are several different activities that have been initiated with respect to education courses and workshops during Program Year 1. These activities will be continued and several will evolve in form in Program Year 2, as discussed below.

During Program Year 1, we visited a number of campuses to introduce the pedagogy of computational science, computational modeling tools, and potential XSEDE services. We also scheduled workshops of various lengths to provide faculty with a greater depth of experience in using computational science tools and educational materials. These have ranged from invited talks to single-day or multi-day workshops under the sponsorship or co-sponsorship of XSEDE. We plan to continue these activities in Program Year 2.

In addition to the short-term training courses being provided by the XSEDE Training and User Services, and ECSS/ESTEO, we are working with the faculty at the University of California, Berkeley to pilot several more extensive parallel computing courses. This includes a summer bootcamp on parallel computing and two regular computer science courses that can be extended to serve a larger audience. In Program Year 1 the bootcamp served more than 250 people, including almost 70 from industry. We expect expanded participation through this and other online class activities in Program Year 2.

Another activity under way in Program Year 1 has been the investigation of modeling and simulation instructional materials that could be incorporated into pre-service science and math teacher preparation courses. Working with the College of Education at The Ohio State University (OSU), we will identify three or four modeling exercises spanning several science and mathematics topics that can be introduced into the methods courses at OSU.

8.2.4 Create a Virtual Computational Science Resource

During Program Year 1, we initiated several activities aimed at facilitating the creation of a virtual community of educators in computational science fields. We provided support to the Journal of Computational Science Education (<http://jocse.org/>), publishing articles on a variety of faculty and student experiences with computational science education topics. We also began to integrate previous work on a computational science education roadmap, the computational science competencies, and existing metadata and science library ontologies that will help guide the creation of a central index of computational science training and education materials.

8.3 Outreach 1.6.2

Outreach programs include engagement of underrepresented communities, a speaker's bureau, student engagement, and Campus Champions, with improvements based on lessons learned during Program Year 1. The programs help identify strengths, weaknesses, and opportunities among XSEDE's programs and activities, to maximize resources and services within XSEDE to engage and retain underrepresented communities in becoming competitive for XSEDE and STEM workforce opportunities.

8.3.1 Under-represented Engagement

The Southeastern Universities Research Association (SURA) will work with faculty and researchers at minority-serving institutions (MSIs) to deepen their involvement in XSEDE and their use of HPC and computational science in their classrooms and research activities. The team will conduct at least 10 visits per year, with the focus on improving sustained engagement of underrepresented faculty, researchers, and campuses. SURA will encourage, mentor, and support the Minority Research Community, a peer-group organization established to provide opportunities for researchers (especially from MSIs and underrepresented communities) to discuss issues and approaches of relevance to their work. SURA will leverage other XSEDE programs (Campus Champions, ECSS/ESTEO, ECSS/NIP, etc.) to distribute the work of engaging new communities more broadly, especially in bringing faculty and researchers up to speed more quickly by utilizing the improved tools available to all XSEDE users.

The XSEDE Scholars Program, led by Rice University, will raise awareness of the opportunities available to all minority and underrepresented students to pursue computational science and engineering studies and careers, as well as provide participating students with skills to develop highly parallel scientific codes, and provide mentoring to sustain the motivation of the students to pursue advanced studies and careers in computational science and engineering. The XSEDE Scholars Program is recruiting the second cohort of 35 underrepresented students in XSEDE training and mentoring, which will have their kickoff meeting at the XSEDE12 conference in July 2012. During Program Year 2, in addition to providing students with awareness of opportunities in the field and mentoring opportunities, the XSEDE Scholars Program will provide a series of webinars to teach students parallel programming skills, culminating in an evaluation process to certify they have mastered the requisite skills, thereby enhancing their skills and resumes. The Campus Champions program will work with XSEDE Scholars to identify potential student candidates for inclusion in the Student Campus Champions program. Students who have been immersed in XSEDE through the Scholars Program will be strong advocates for, and supporters of, XSEDE on their campuses, as well as becoming viable candidates for internships and research opportunities.

Rice University's Minority Faculty Council is charged with serving as a resource for the XSEDE partnership, providing advice for broadening participation within XSEDE and the HPC community in general. The Council will serve in an advising capacity with "bite" and as mentors for students, with the goal of raising retention among minority students. The Council will begin Program Year 2 with a face-to-face meeting at XSEDE12 to start detailed planning for advising XSEDE.

8.3.2 Speakers Bureau

Program Year 2 of the Speakers Bureau will participate in at least 10 national/regional professional events, with the goal of raising awareness and recruiting a larger and more diverse STEM community to utilize XSEDE's services and resources to support their research and education needs. Materials will be improved to be more effective tools for communicating XSEDE's availability, usability, and potential benefits to the national science and engineering

community. Processes and programs designed for prioritizing, executing, and following up on outreach opportunities will continue to be refined to be more efficient and effective. Using feedback from the community, especially the E&O community requirements survey, the Speakers Bureau will participate in events that the community has identified as highly relevant targets for extending the message of XSEDE. The Speakers Bureau will work with ECSS/ESTEO to conduct outreach and with XSEDE External Relations to improve information dissemination.

8.3.3 *Student Engagement*

Based on the overwhelming enthusiasm for its initial rollout, the Student Engagement program will review and refine processes established in Program Year 1, including application, execution, and evaluation mechanisms. This will include evaluating issues related to supporting 20 students in a full-year program as originally intended. Based on the popularity and expected success of the initial program, XSEDE will apply for an NSF REU site award, to expand the program to more students, sites, and projects.

8.3.4 *Campus Champions*

The Campus Champion program has achieved unprecedented growth through the transition from TeraGrid to XSEDE. The focus for Program Year 2 is to provide depth of support for the Champions to help ensure they and their campus personnel benefit from all that XSEDE has to offer the community. The Campus Champions program is gaining momentum and having a significant impact on the member campuses and in extending their impact to benefit researchers, educators, and students on other campuses. The Champions will build strength through their increasing interactions and mutual support mechanisms as they build an ever more vibrant community of practice. The Campus Champions Fellows Program, conducted in collaboration with ECSS, will directly contribute to enhancing the community of practice. Two working groups have been formed from among a number of Champions to address the group's self-identified most pressing needs: 1) how to promote XSEDE on their campuses through more effective outreach, and 2) how to engage campus users more effectively. The efforts include collaboration with all XSEDE teams, to develop appropriate training roadmaps and materials for Champions and their users, and mentoring and supporting students and faculty involved in utilizing and benefitting from the broad range of XSEDE resources and services.

8.4 Community Requirements and External Evaluation 1.6.3

It is critical that the E&O activities be responsive to the needs of the community and that the programs and activities meet or exceed their goals in serving the community. To accomplish this, the E&O team will formally collect community requirements annually, conduct rigorous external evaluations, and respond to the advice of the XSEDE Advisory Board, the XSEDE User Advisory Committee, and the TEOS Advisory Committee.

While community feedback is collected on an ongoing basis, E&O will conduct an annual community requirements gathering process to assess evolving community needs. The process will include surveys, interviews, and focus groups. The community requirements will be gathered during the spring to influence the planning for Program Year 3. E&O will conduct quarterly meetings with the TEOS Advisory Committee to seek their ongoing advice and recommendations for refining, improving, and altering E&O programs and activities. The TEOS Advisory Committee will ensure that at least one person is also on the XSEDE Advisory Board and the XSEDE User Advisory Committee, to ensure cross-communications among the various XSEDE advisory groups.

E&O will work with the external evaluation team to pursue formative and summative evaluations to help E&O improve its services. In addition to the continuous evaluation activities conducted in

Program Year 1, the evaluation will include the following activities for Program Year 2 in each area of E&O:

Training evaluation: The external evaluators will implement the training evaluation suite by conducting interviews with XSEDE's largest users of resources and surveying training activities.

Education evaluation: During Program Year 2, evaluation of education will involve more follow-up on workshops to identify best practices of institutional HPC concentration implementation.

Outreach evaluation: The external evaluators will add a new Program Year 2 component to conduct interviews with the Campus Champions Fellows and their ECSS supervisors to assess this new program's impact. Also new in Program Year 2 will be evaluations of the newly formed Campus Champions working groups.

Campus Bridging evaluation: Site visits for the GFFS pilot program will occur in Program Year 2. Evaluators will follow up with pilot sites one year after implementation to assess the impact of the pilots on the campus and the extent to which the infrastructure enhances the productivity of XSEDE users.

A E&O Evaluation Matrix has been developed (and documented in the XSEDE quarterly reports) to provide additional details on the evaluation process. The evaluation matrix is regularly revised and updated to reflect lessons learned throughout the year.

8.5 E&O Infrastructure 1.6.4

The E&O Infrastructure group serves the internal needs of the E&O team, and provides a conduit to the XSEDE External Relations team and to the public for information and opportunities arising from the E&O team. The Infrastructure group will consult with E&O leads to keep website content fresh, informative, and easy for the community to navigate. In addition, the infrastructure team will 1) be pro-active in posting information to the blog, Facebook pages, and the XSEDE public website, 2) place an emphasis on disseminating community impact stories, and 3) collect internal requirements from the E&O team to address ongoing portal needs for the various E&O programs.

8.6 Campus Bridging 1.6.5

The response from the community for Campus Bridging has been even greater than we imagined when the planning for this effort began. During Program Year 2, XSEDE has committed an additional 1.25 FTE at IU and Cornell to enhance the effort to respond to community requirements. The Campus Bridging pilot project is allowing XSEDE to work out organizational understandings about governance and group-to-group handoffs within XSEDE that will be helpful in the long run.

8.6.1 Campus Bridging Pilot Program for Campus Bridging

The Campus Bridging team will conduct pilot programs using a combination of the Execution Management Service (EMS) and the Global Federated File System (GFFS). A dry run of the pilot program will be conducted at Indiana University, with work continuing at four selected pilot sites. The pilot implementations are being done in collaboration with the A&D, Operations, and SD&I teams. The Campus Bridging team provides communication between pilot sites and XSEDE and coordinates training at the pilot sites. By the beginning of Program Year 2, the Campus Bridging team will have worked with half of the pilot sites and hosted a panel session on the experiences to date at XSEDE12. The pilot program is expected to be completed in Program Year 2. The Campus Bridging team is collecting requirements on an ongoing basis. XSEDE Operations will roll out EMS and GFFS to the 12 institutions that submitted proposals to participate in the pilot

project but were not selected for participation. These institutions will be considered “early adopters” of the EMS and GFFS software.

8.6.2 Campus Bridging Use Cases

The Campus Bridging team will work with the A&D team to transform the use cases into a Level 3 decomposition document that will be made public in Program Year 2. A paper titled, “What is Campus Bridging?” is being prepared for submission to the XSEDE12 conference to provide an explanation of Campus Bridging and how it will benefit users. This will help clarify Campus Bridging and the related set of challenges that exist at the interface between XSEDE and campus-based cyberinfrastructure.

8.6.3 Software packaging

The Campus Bridging team will develop software packages and installers that will allow the implementation of campus systems that resemble XSEDE resources. One FTE at Cornell will be funded with the responsibility of creating software packages using ROCKS as the basic distribution method. The team will build tools to enable clusters on campus to be built similar to XSEDE clusters and increase the interoperability among clusters. Finally, the team will track the Open Science Grid’s work using RPM-based distribution of their software, utilizing FutureGrid systems as a test bed for implementation.

8.6.4 Training, Documentation and Communications

One FTE at Indiana University will dedicate time to the creation of system documentation based on XSEDE resources and training materials for Campus Bridging. To facilitate internal knowledge and communication on the project, the Campus Bridging team will produce quarterly updates on the progress of the project for the “Inside XSEDE” internal newsletter. The Campus Bridging team will produce features for the XSEDE external newsletter at the same pace, providing information that may be presented through outreach efforts. Two of the biggest steps in communication during Program Year 2 will be a “What is Campus Bridging?” paper at XSEDE12, and the public-facing use cases and Level 3 decompositions of the use cases. These documents will provide clarity on XSEDE’s Campus Bridging initiatives during the next four years.