

XSEDE Value Added, Cost Avoidance, and Return on Investment

Craig A. Stewart
IU Pervasive Technology Institute
& IU School of Informatics and
Computing
2709 E. Tenth Street
Bloomington, IN 47408-2671
stewart@iu.edu

Ralph Roskies
Pittsburgh Supercomputing Center
300 S Craig St.
Pittsburgh, PA 15213
roskies@psc.edu

Richard Knepper
IU Pervasive Technology Institute
2709 E. Tenth Street
Bloomington, IN 47408-2671
rknepper@iu.edu

Richard L. Moore
San Diego Supercomputer Center
10100 Hopkins Dr.
La Jolla, CA 92093
rlm@sdsc.edu

Justin Whitt
National Institute for Computational
Sciences
PO Box 2008, BLDG 5100
Oak Ridge, TN 37831-6173
justin-whitt@tennessee.edu

Timothy M. Cockerill
Texas Advanced Computing Center
University of Texas at Austin
Road A, Austin, TX 78758
cockerill@tacc.utexas.edu

ABSTRACT

It is difficult for large research facilities to quantify a return on the investments that fund their operations. This is because there can be a time lag of years or decades between an innovation or discovery and the realization of its value through practical application. This report presents a three-part methodology that attempts to assess the value of federal investment in XSEDE: 1) a qualitative examination of the areas where XSEDE adds value to the activities of the open research community, 2) a “thought model” examining the cost avoidance realized by the National Science Foundation (NSF) through the centralization and coordination XSEDE provides, and 3) an assessment of the value XSEDE provides to Service Providers in the XD ecosystem. XSEDE adds significantly to the US research community because it functions as a unified interface to the XD ecosystem and because of its scale. A partly quantitative, partly qualitative analysis suggests the Return on Investment of NSF spending on XSEDE is greater than 1.0, indicating that the aggregate value received by the nation from XSEDE is greater than the cost of direct federal investment in XSEDE.

Categories and Subject Descriptors

C.5 [Computer System Implementation]; J. [Computer Applications]; K.6 [Management of Computing and Information Systems].

General Terms

Management, Measurement, Design, Economics, Reliability

Keywords

XSEDE; Return on Investment; ROI; Value added, cost avoidance; advanced cyberinfrastructure

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

XSEDE '15, July 26 - 30, 2015, St. Louis, MO, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3720-5/15/07 ... \$15.00

DOI: <http://dx.doi.org/10.1145.2792745.2792768>

1. INTRODUCTION

Does the current structure of national cyberinfrastructure (CI) funded by the National Science Foundation, with XSEDE and multiple CI service providers, make sense? Is this structure fiscally efficient? And how did we end up with this structure anyway? A quick historical answer to the last question helps create a foundation for addressing the first two questions.

The National Science Foundation (NSF) has implemented a variety of approaches to the organization and delivery of what we now call advanced cyberinfrastructure services [1]. Starting in the 1980s the NSF funded operations of five distinct supercomputer centers. In 1997, NSF established the PACI (Partnerships for Advanced Computational Infrastructure) program. Under PACI it funded two large collaborative groups, one led by the National Center for Supercomputer Applications (NCSA) and the other led by the San Diego Supercomputer Center (SDSC). In 2001 the NSF refocused its support for national CI facilities based on a grid approach with four sites included in the resulting grid facility: NCSA, University of Chicago/Argonne, SDSC, and the California Institute of Technology. This was the start of TeraGrid. The NSF later funded a large system at Pittsburgh, then added it into TeraGrid, creating a heterogeneous grid environment and began using the term “cyberinfrastructure” (CI) to describe the national advanced information technology resources it funded in support of open science and engineering research [2,3]. TeraGrid was later expanded by incorporation of resources at the University of Texas, Purdue University, Indiana University, and the University of Tennessee. From 2004 to 2011 TeraGrid was the dominant organizing structure for NSF-funded advanced CI. During most of its operational phases TeraGrid consisted of multiple independent awards to a group of centers then called Resource Providers (who delivered the computational, storage, and visualization facilities within TeraGrid) and one award for a coordinating function (the Grid Integration Group) led by the University of Chicago [4]. Resource Providers were not obligated to adhere to standards set by the Grid Integration Group. Decisions made on a TeraGrid-wide basis thus required unanimous agreement of as many as seven Principal Investigators. The impact of “veto by a minority”

was that fewer services were offered consistently across TeraGrid than otherwise might have been.

In 2008 the NSF introduced a new structure under the newly named XD (eXtreme Digital) solicitation [5]. Two grant awards were executed under this solicitation for implementation of XSEDE — the eXtreme Science and Engineering Discovery Environment — with NCSA as the lead organization and John Towns as Principal Investigator (PI) [6-8]. A related grant for audit of XSEDE activities was awarded to SUNY Buffalo with Thomas Furlani as PI [9]. The XSEDE award initially included 16 funded partner institutions.

In 2011 XSEDE became the central coordinating entity organizing the delivery and support of NSF-funded advanced CI services for the nation via the XD program. The NSF charged XSEDE with setting standards and developing formal mechanisms for integration of the computational, storage, visualization, and consulting services allocated via XSEDE. The entities that deliver such services have been called Service Providers since 2011. This led to creation of the Service Provider (SP) Forum, which defines a Service Provider [10] as:

Service Providers are independently funded projects or organizations that provide cyberinfrastructure (CI) services to the science and engineering community.

The SP Forum includes but is not limited to entities that provide advanced CI services and has an organizational structure clearly defined in a charter document [10]. SP Forum members have a role in XSEDE governance. The SP Forum Charter and a companion document describe how SPs request membership in the XSEDE Federation [11], a subset of the SP Forum. This document defines three Levels of XSEDE SPs. Level 1 SPs are most tightly linked to XSEDE standards and services. SPs funded by the NSF through the High Performance Supercomputer Acquisition solicitations are required to integrate with the XSEDE architecture. The most recent such solicitation [12] states:

The resource acquired will become integrated into the XD program. Within this context, describe which elements of the proposed resource will be integrated into XD via XSEDE and what steps such integration will require. This description should be based on the current XSEDE architecture.

The SPs funded through the NSF High Performance Supercomputer Acquisition solicitations are generally categorized as Level 1 SPs – those that integrate and comply most closely with XSEDE-specified architecture and standards. XSEDE and the SP Forum also define Level 2 and Level 3 XSEDE SP characteristics. Level 2 SPs may have some resources allocated via the XSEDE allocations process and adhere to some (but generally not all) of the XSEDE standards. Level 3 SPs don't necessarily offer resources via the XSEDE allocations process and have minimal requirements, but do at least publish via XSEDE information about their service offerings. Level 2 SPs may for example be funded via NSF Major Research Infrastructure awards to implement a system that provides resources to a local community and the national research community [13]. Level 3 SPs typically do not provide XSEDE-allocated CI resources but publicize services and the nature of the communities they serve via XSEDE information dissemination mechanisms. The SP Forum Charter also allows for SPs to be members but have no affiliation with XSEDE.

The XSEDE governance model provides a clear decision-making hierarchy, including formal roles for the XSEDE PI and Co-PIs and the SPs that are members of the XSEDE Federation. An in-

depth description of XSEDE and its governance model is available at [10].

Over the course of roughly 30 years the NSF-funded resources delivering what we now term advanced CI have moved from a model of multiple independent centers serving the national research community to a model of a single central organizing, supporting entity (XSEDE) and multiple independent awards to SPs, which integrate with XSEDE standards. The mission of NSF-funded advanced computing facilities has also changed. In the 1980s the supercomputer center program was intended to act as a “time machine” to provide a small group of researchers with access to extreme-scale computing capabilities that would not be available for several years in the future [14]. Today, the mission for NSF advanced CI facilities has expanded to serve several thousands of researchers and research students, enabling cutting-edge research in many areas of study, accelerating research and helping to develop a technically savvy and diverse 21st-century workforce.

This is the historical answer to the question of how we arrived at the current structure of XSEDE and multiple SPs. Part of the structure of NSF-funded advanced CI services was specified by the NSF. Other aspects, such as the current structure of the SP Forum, were organized by leaders within the SP community and XSEDE leadership. This suite of services is funded by multiple awards from the NSF, along with facilities contributions by awardees and in some cases additional funding from other federal or state agencies. In this paper we use the term *XD ecosystem* to refer to XSEDE itself and the SP members of the XSEDE Federation designated as XSEDE Level 1, 2, or 3 SPs. The XD ecosystem provides computational, data analysis, visualization, and consulting and programming services that support many researchers with a wide variety of needs and interests.

The NSF budget for the XD ecosystem in recent years has averaged more than \$50M per year, including XSEDE-related activities, new system acquisition, and system operation and maintenance. XSEDE-related awards include more than \$26M annually for three awards: XSEDE- \$23.6M [6], the XSEDE evaluation service - \$1.5M [7], and the Technology Insertion Service - \$1.6M [9]. Solicitations for investments in major systems (the “High Performance Computing System Acquisition” solicitations) are typically issued annually by the NSF and over the past few years have ranged from \$30M [15] to \$12M [16]. Total annual awards average at least \$15M. Annual operations and maintenance typically comprise at least 20% of each system's acquisition cost, and each system is typically operated for four years. The assumption that in any year there are systems in operation resulting from four annual cycles of acquisitions implies that annual expenditures in operation and maintenance of major systems operated with NSF funding at SPs is at least an annual \$12M. In other words, NSF's annual investment in the XD ecosystem is more than \$50M a year. (This figure is consistent with the NSF FY2014 budget request summary available at [16].) This is a significant investment.

The question with which we started this report — “Does the current structure of national cyberinfrastructure (CI) funded by the National Science Foundation, with XSEDE and multiple CI service providers, make sense?” — is a lot to grapple with. The following questions are more tractable: “Are the XSEDE organization and services a cost-effective way to deliver CI services to the US open research community? Is the money invested in XSEDE well spent?” Answering these two questions well is our goal in this report.

We attempt to analyze the value and financial efficiency of the organization of XSEDE and SPs, using these four business concepts [17]:

- Value Added, defined as “an activity that increases the worth of the product or services to the customer...”
- Committed fixed costs, defined as “a cost related to either the long-term investment in plant and equipment of a business or the organizational personnel who are deemed essential....”
- Cost avoidance, defined as “the practice of finding acceptable alternatives ... and/or not spending money for unnecessary goods or services.” This is measured as the cost difference between doing something one way vs. a hypothetical other way.
- Return On Investment (ROI): “a ratio that relates income generated ... to the resources (or asset base) used to produce that income,” calculated typically as “income or some other measure of return on investment.” Values greater than 1.0 indicate that return is greater than investment.

Objectively estimating value of a national investment in research or its support is difficult. In science and engineering the identification of a return on an investment may take decades. For example, Peter Higgs proposed what was later named the Higgs Boson in 1964, but was not awarded a Nobel Prize for this major scientific advance until 2013.

Analyzing return on investment in cyberinfrastructure is particularly difficult because advances enabled by use of advanced CI may come only after years or decades, and may often result from many contributing factors, of which advanced CI is just one. An SC14 conference panel discussed ROI for supercomputing and outlined types of value and ROI [18], but no panelist presented actual calculations. Work by Apon *et al.* suggests that institutional investment in high performance computing yields increased federal grant funding and publications [19,20]. Apon’s analysis is correlative and intriguing, but it does not determine whether the overall return is greater than the investment cost. Two prior reports [21,22] contain dozens of examples of advanced CI adding value to health, science, and quality of life, and benefiting industry and business. One presumes large businesses don’t invest in supercomputers without a positive ROI. However, the details of such analyses are of a competitive nature and not publicly released. In these two reports combined there are few quantitative financial statements. Among them is this from Alcoa in the 1980s: “[T]o develop a new can design in the laboratory costs about \$100,000 compared to about \$2,000 with computer modeling.” ROI for a given project may seem high, but such analyses do not offer comprehensive assessments of the overall ROI in the facilities that enabled the discoveries.

Figures 1 and 2 below are schematic representations of two approaches to the organization of federally funded national CI. Figure 1 shows individual users interacting independently with multiple national centers. In this model, it may not be clear to a given end user which national CI center provides the most appropriate resources. National centers may even compete for users in ways that reflect a mix of the best interests of the users and the national centers, rather than just the best interests of the users. Figure 2 shows XSEDE serving as a single interface between the national user community and multiple centers. In this paper we analyze the value added that derives from XSEDE as a single interface for the national research community and estimate cost avoidance from XSEDE providing essential support services for that community vs. multiple independent national centers. For the latter we compare committed fixed costs for XSEDE essential

core functions and models of two or four independent national CI centers. We also calculate an overall ROI for essential core XSEDE functions, based in part on SP valuation of XSEDE services.

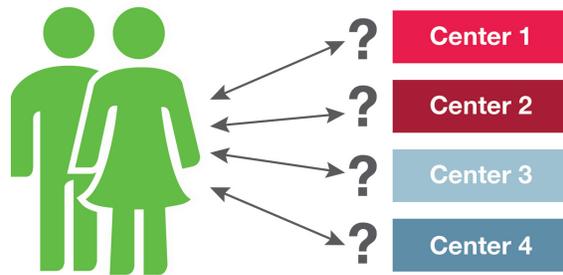


Figure 1. User view of a hypothetical model of four independent national CI centers.

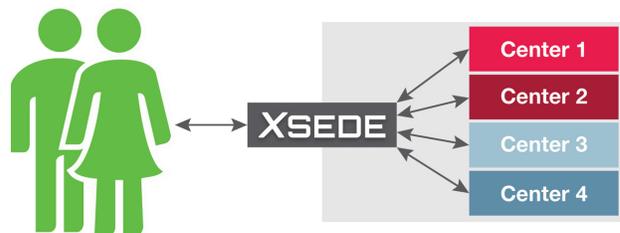


Figure 2. User view of current NSF-funded XD ecosystem

2. XSEDE VALUE ADDED

As a national, unified, virtual entity supporting open science and engineering R&E, XSEDE creates value because it is a single entity. Other benefits derive from the project scale. We group the values XSEDE adds into one of two categories, depending on our view of the primary source of that value:

- * Value resulting primarily from one entity serving the national research community, usable as a common national resource, and providing unified standards with which other projects can align.
- * Value that results primarily from the scale of XSEDE.

2.1 Value derived from XSEDE centralization

Multiple benefits derive from XSEDE’s being one entity serving the national user community. The list of such benefits below moves from values largely internal to XD ecosystem organization to those benefiting the US user community, then benefits to US researchers beyond the XSEDE community:

Security effectiveness. One unified security team within XSEDE has proven highly effective in responding to security incidents as compared to TeraGrid. Under TeraGrid, between 2004 and 2011, two major security incidents propagated from one center to multiple centers. Under XSEDE, while there have been isolated breaches at individual SPs, no incidents have propagated from one SP to another. This is largely thanks to XSEDE’s coordinated approach to security. The potential cost of a security breach spanning multiple computational and data storage systems that are among the most powerful in the US could be very large.

Disaster resilience. Essential core services are configured by XSEDE to provide resilience even in the face of a regional disaster, e.g. they are operated in one US geographic region with backups in a different region. This includes authentication and

user interfaces. Failover to secondary services has been successful in practice.

Uniformity and unification of computing environment improves user productivity, including allocations and authentication. The XD ecosystem has one interface and one central supporting organization: XSEDE. A single web interface provides access to support and services for the entire XD ecosystem, including a unified authentication service. XSEDE help desk support is coordinated: Documentation at all sites has a consistent look and feel. A unified set of tools for data movement means consistency in tools for moving data from SP to SP, or from outside XSEDE to within XSEDE and back. The configuration and use of SPs may vary, but users experience more uniformity among systems overall than during TeraGrid because of the NSF requirement that Level 1 SPs (typically the largest systems) align with the essential aspects of XSEDE standards.

XSEDE has a single process for matching researcher needs to CI resources in the XD ecosystem and a single allocation process. It provides expert help to users in selecting the right resource from national XD resources. Each quarter the XSEDE Resource Allocation Committee (XRAC) evaluates all allocation requests for XD ecosystem resources, selects the most meritorious, and awards allocations worth millions of dollars. Its operation is consistent with NSF peer-review principles, ensuring NSF investments in XD resources are put to the best possible use to support the US research community.

Software optimization and support. XSEDE provides high-quality software infrastructure for national service. Its engineering process hardens and vets software used as part of the XD ecosystem, ensuring high-quality functionality and security of such essential software.

XSEDE also provides significant benefit in support of widely used community codes. In a way analogous to the allocation process, XSEDE prioritizes and coordinates software support efforts on a national scale, coordinating dozens of experts to optimize software needed to support the most important and innovative research in the national open research community. XSEDE provides continuity over time of science and engineering domain support as computing architectures change. This support is enabled by application experts well versed in particular applications and who participate as staff in XSEDE for periods of years, optimizing important codes for multiple architectures. Improving code substantially is better than buying more hardware and shortens researcher time to new insights. Further, XSEDE code optimizations are generally released back to the US research community as open source. Thus investments in widely used community codes improve the effectiveness of the entire national CI within and beyond XSEDE.

National cyberinfrastructure software leadership. XSEDE is an entity around which the national CI ecosystem can align. It creates a set of interfaces that enables bridging between XSEDE and US college and university campuses. This is achieved largely by the architecture and set of standards XSEDE created and enables the national research community to better leverage other investments in campus CI, including federally funded systems located on and serving a single campus, state, or region.

2.2 Value derived from XSEDE's scale

A number of benefits derive largely from the fact that XSEDE is a large organization with more than 200 individuals funded annually in part or in whole to carry out XSEDE functions. These benefits derive mainly from the scale of XSEDE, and are listed

here in order from values largely internal to the XD ecosystem to those benefiting the national user community:

Broader participation among institutions. One of the side effects of XSEDE's scale is it includes many partners, increasing the breadth of participation in delivering national CI services. Some 15% of XSEDE's budget goes to the 14 institutions outside the five institutions represented by the XSEDE PI & Co-PIs.

International coordination. XSEDE serves as a coordinating point for national and international entities engaging and collaborating with NSF advanced CI programs, including, for example, PRACE [23] and Compute Canada [24]. These discussions include work to develop international standards in CI implementation.

Disciplinary breadth. XSEDE is able to employ a breadth of expertise across many diverse disciplines. In some cases, it is possible to integrate need across the nation and have a person dedicated full time to a domain that would be a small portion of the use of any one center, e.g. digital humanities. Because of XSEDE's distributed staffing model it is also possible to obtain a fraction of an expert's time so as to allow coverage of domains composed of diverse sub-domains that two or four independent centers could not cover. XSEDE's Novel and Innovative Projects (NIP) group provides a single team supporting new and emerging needs in the national research community [25]. In the case of NIP and software optimization, one collaborating and coordinating team, as opposed to two or more competing teams, provides more effective use of talent, deeper support, and greater availability, and enables innovation integrated at the national level.

Confidence in continuity. Under the former centers models, it was never certain how many centers would be funded and for how long. The question for users was often: "Which one will disappear next?" SPs may indeed lose funding. And while it is natural to expect XSEDE to change over time and to eventually be replaced by a successor organization, there is a strong sense that XSEDE will not suddenly disappear. This helps staff feel confident in their employment (a problem at centers in the past). It also enables other scientific projects to depend on XSEDE and the XD ecosystem overall. This continuity is less likely with multiple, competing national centers and shorter durations of grant awards.

An effective CI resource for major national science projects. Several science projects depend on XSEDE as a resource. XSEDE Principal Investigator John Towns has written more than 40 letters committing XSEDE assistance to proposals for funding of science programs that seek to depend on XSEDE for CI resources. TAIR (The Arabidopsis Information Resource) [26] is among those that use XD resources in addition to or instead of operating their own CI. NCGAS, the National Center for Genome Analysis Support [27] is an interesting model, aiding XSEDE and SPs with managing and maintaining genome analysis software, while assisting researchers in its use. IU operates the large-memory cluster Mason in support of NCGAS as an XSEDE Level 2 SP. Mason is allocated through XSEDE. Together XSEDE and NCGAS provide the CI resources and detailed software expertise to serve US biologists doing genome research. NCGAS could not do this without XSEDE partnership, and NCGAS provides deep expertise in genome analysis complementing XSEDE services.

Science gateways. Science gateways are "a community-specific set of tools, applications, and data collections that are integrated together via a portal or a suite of applications" that can "support a variety of capabilities including workflows, visualization as well as resource discovery and job execution services" [28]. Science gateways enable sophisticated and intensive simulations and data analyses without requiring computational science expertise. The

consistency of interfaces to XSEDE resources is particularly important in supporting distributed workflows orchestrated by science gateways and utilizing multiple XSEDE SP resources. XSEDE supports and facilitates development of science gateways by providing one national CI environment serving as a comprehensive CI foundation on which they operate. The importance of science gateways in enabling access to large-scale CI resources is demonstrated by a significant first for XSEDE and predecessors. As of March 31, 2014, the number of users accessing XSEDE-supported resources via science gateways surpassed the number using command-line interfaces.

Training, education, and outreach programs of national scope. There is a well-known shortage of people leaving undergraduate and graduate programs with strong backgrounds in computational science. XSEDE attracts attention and XSEDE training reaches a nationwide audience, adding to the national supply of academics and professionals entering the workforce with good computational skills. The XSEDE portal shows a unified view of all training programs available via XSEDE. XSEDE reaches the nation with users in all 50 states, the District of Columbia, and US territories. XSEDE staff are located in 18 states and the District of Columbia. Campus Champions, IT professionals or faculty volunteers who are funded by their home institutions, convey information and answer questions on XSEDE. There are 247 Campus Champions at 189 institutions in 49 states [29]. The annual XSEDE conference [30] brings together a national community of scientists and computational experts to share expertise and collaborate. XSEDE gives all these efforts national impact.

XSEDE as a technology adoption accelerant. XSEDE's scale, particularly as regards training, education, and outreach, aids the adoption of XSEDE and XSEDE-supported software. XSEDE's approach is consistent with the current social-science-based understandings of technology adoption [31] that suggest adoption is driven by performance expectancy (perceived value), effort expectancy (perceived ease of use), social influence, and facilitating conditions (including knowledge of a technology and the belief that end users will find it accessible). Performance expectancy, or perceived value, is likely not dramatically different for XSEDE as compared to multiple independent centers. Effort expectancy, or perceived ease of use, is crucial in adopting CI and central in XSEDE's approach to training, outreach, education, and overall usability, providing users with a more favorable impression of usability. Social influence and facilitating conditions play an important role in encouraging adoption. Because there is just one XSEDE, those who ought to be making use of NSF-funded CI may more readily do so because they hear a consistent message about it from their colleagues.

3. COST AVOIDANCE

Is XSEDE a fiscally efficient way to use taxpayer money to enhance US global competitiveness and enable discoveries that improve quality of life and enhance understanding and appreciation of our universe? XSEDE's budget can be seen as having two major components: essential core supporting and organizing functions (~\$16M), and consulting and programming activities that contribute to the "value added" of XSEDE (~\$7M). We can calculate some measures of financial effectiveness in terms of the essential core functions that any national CI center must provide using the accounting concept of cost avoidance.

We created a "thought model" comparing the cost of one national integrating core service (XSEDE) with the committed fixed-cost

minimum for a hypothetical two-center and four-center model. First, we calculated committed fixed costs of the essential core functions of a national CI center. For example, there has to be at least one person to manage user accounts. It takes a minimum of five people (assuming a maximum 40-hour week) to provide 7 x 24 operational support. In multiple discussions at XSEDE meetings and in conversations with CI experts, we arrived at staffing levels for essential core functions for any national cyberinfrastructure service center. We used the \$200,000 per-year cost of a "canonical staff member" used in setting XSEDE budgets (covering salary, benefits, travel, and facilities and administration costs) to approximate the cost of a high-performance computing expert including salary (less than half the total cost), benefits, facilities & administration costs associated with grant awards, travel, and computer equipment. We used a lower figure of \$125,000 for some functions such as 7 x 24 operational support. **Table 1** on the next page shows the minimum fixed costs per national CI center. **Table 2** shows XSEDE's actual annual investments in these functions. We assumed that staffing for these functions scales with the number of centers except parallel computing support. We estimated that 24 FTEs (Full Time Equivalents) are needed for parallel computing support for the US research community regardless of number of centers.

We calculated cost avoidance figures for XSEDE, comparing XSEDE expenditures for essential core functions with the thought-model costs for essential functions of a national model of two or four national centers. Cost avoidance was calculated as

$$\text{Cost avoidance} = n \times (\text{minimum fixed cost of a national CI center other than parallel computing support}) + \text{cost of 24 parallel computing support experts} - \text{actual XSEDE expenditures on essential core functions (which includes 24 parallel computing support experts)}$$

where $n = 2$ or 4 . This equation is lengthy in order to emphasize that the cost avoidance calculations are independent of the number of individuals it takes to provide parallel computing support for the national research community, as long as our assumption is correct that the number of such individuals scales with the size of the research community rather than the number of centers. These values for n echo the past in some ways as the PACI program of the 1990s involved two major grant awardees, and four centers operated during most of the NSF supercomputer center program of the 1980s and early 1990s. As of the writing of this analysis there were four SPs operating large advanced CI systems as part of the XD ecosystem.

We calculated \$18,418,250 per year as the cost avoidance attributable to XSEDE as compared to four independent national advanced CI centers, and \$3,428,250 per year as the cost avoidance attributable to XSEDE as compared to two independent centers. In delivering the essential core functions of a national advanced CI center, XSEDE seems more cost effective than two or four independent national CI centers for delivering essential services. This analysis considers only the functions XSEDE is charged and funded to deliver through the NSF grant awards and does not consider management and operations budgets of individual SPs. While this is simply one approach to comparing financial efficiencies, to a first order this analysis suggests XSEDE is financially efficient in delivering services to the national research community.

Activity	FTEs	FTE Costs	Non-FTE Costs
Allocations	2.0	\$400,000	\$100,000
Account Management	2.0	\$400,000	
Authentication Services	3.0	\$600,000	
24 x 7 Operations	5.0	\$625,000	
Ticket Support	1.5	\$187,500	
User Survey	0.5	\$62,500	
Leadership	3.0	\$600,000	
Project Managers and Financial Functions	2.0	\$400,000	\$30,000
User Information Services	3.0	\$600,000	\$20,000
Training	2.0	\$400,000	
Education	3.0	\$600,000	
Outreach	1.0	\$200,000	
Network	1.0	\$200,000	\$250,000
Central Services (web portal, etc.)	5.0	\$1,000,000	\$20,000
Systems Engineering and Deployment	5.0	\$1,000,000	
Parallel Computing Support	24.0	\$4,800,000	
ANNUAL TOTALS	63.0	\$12,245,000	

Activity	FTEs	FTE Costs	Non-FTE Costs
Allocations	2.25	\$450,000	\$100,000
Account Management	2.50	\$500,000	
Authentication Services	4.30	\$860,000	
24 x 7 Operations	5.0	\$625,000	
Ticket Support	2.0	\$250,500	
User Survey	1.25	\$156,250	
Leadership	5.0	\$1,000,000	
Project Managers and Financial Functions	1.50	\$30,000	\$30,000
User Information Services	5.0	\$1,200,000	\$20,000
Training	6.5	\$1,300,000	
Education	4.5	\$900,000	
Outreach	3.0	\$600,000	
Network	3.5	\$700,000	\$250,000
Central Services (web portal, etc.)	5.0	\$1,000,000	\$20,000
Systems Engineering and Deployment	7.0	\$1,400,000	
Parallel Computing Support	24.0	\$4,800,000	
ANNUAL TOTALS	82.3	\$16,191,750	

4. VALUE TO SERVICE PROVIDERS

XSEDE delivers services directly to two sorts of entities: SPs and end users. The cost avoidance analysis did not include consideration of SP operations and management budgets. However, we assessed XSEDE’s value to SPs by asking Level 1 and Level 2 SPs with resources allocated through XSEDE to estimate how many FTEs (Full Time Equivalents) it would take to offer the same service level they judge it to receive from XSEDE, rounded to the nearest 0.25 FTE. We allowed negative values (something XSEDE does creates work for an SP and is a net negative in terms of the operational activities of an SP) and zero (either the SP judges the service to be without value, or the SP could offer it at zero incremental cost). We requested these values per organization (e.g. the Texas Advanced Computing Center) not per resource (e.g. Stampede vs. Ranch).

As of January 2015 four organizations offered Level 1 XSEDE SP resources: the National Institute for Computational Sciences, Pittsburgh Supercomputing Center, San Diego Supercomputer Center, and Texas Advanced Computing Center. Based on their responses the aggregate value of XSEDE to these SPs, calculated using the \$200,000-per-year cost of a canonical XSEDE FTE, is \$10,546,250 per year. Level 2 SPs are less tightly tied to XSEDE, and in several cases represent services and systems allocated through XSEDE but not funded directly

through the NSF XD program. Among Level 2 SPs are institutions with Major Research Infrastructure awards that offer allocations on those systems through the XSEDE allocation process. Three Level 2 SPs responded to our request for information (not all current Level 2 SPs). Based on their responses the total value of XSEDE services to them, calculated using the \$200,000-per-year cost of a canonical XSEDE FTE, is \$1,111,250. Level 2 SP comments also give anecdotal evidence of XSEDE value:

- “There seems to be incrementally more administrative work needed to support XSEDE than if we didn't. But those negatives get offset by values that aren't quantified.”
- “We don’t have the resources to explore software improvements, so see XSEDE’s efforts as helping make sure we have efficient code to deploy.”
- “We would not be able to offer our resource as a national resource without XSEDE.”

The total benefit of XSEDE to seven SPs that responded to the survey (four Level 1 SPs, three Level 2 SPs) came to \$11,657,500 per year. This is a good bit lower than the \$18,418,250 we calculated as a cost avoidance value for XSEDE. There are two reasons for this. First, the cost avoidance measurement takes into consideration benefits to end users and SPs; the \$11M figure reflects only value to SPs. Also, there

were differences between SPs and the cost basis we used for the minimal fixed cost of a national CI center. For example, the single biggest difference between the SP-assessed value and the source of cost avoidance in the thought model is in the value assigned to XSEDE 7 x 24 operations. This was viewed as much less important by SPs than the value indicated in calculations based on the thought model would suggest. XSEDE leadership is considering this and other SP survey data in future planning.

5. ROI – GREATER THAN 1.0 OR NOT?

We do not have enough information to calculate ROI on NSF XSEDE expenditures overall. It is impossible to calculate a financial return on value-added aspects of XSEDE from the results measureable in the few years of XSEDE's existence because so many benefits will be realized in years to come as innovations enabled by XSEDE result in new consumer products and contribute to improved national quality of life. Qualitatively we can tally XSEDE value added with analyses and anecdotes as in Section 2.0. The \$7.6M difference between investment in essential core functions and XSEDE's annual budget (using PY3 figures) represents XSEDE's direct investment in value-added activities. One singular accomplishment might well cost justify the entire XSEDE enterprise. For example, with help from the XSEDE NIP program, Mao Ye (U. Illinois College of Business) and collaborators used XSEDE resources to show that automated high-speed trading can distort stock markets [32]. They found 20% of NASDAQ trades are done automatically in "odd lots" not previously reportable to the moment-by-moment ticker tape. This find results in a change in NASDAQ and New York Stock Exchange rules. Now all trades are reportable and visible each moment, reducing chances of a stock market crash created by automated trading. The value of averting one such crash could cost justify XSEDE's total budget.

As to the budget for essential core services, we can make some qualitative assessment on whether ROI is greater than 1.0. From data on the perceived value of XSEDE essential core services we can calculate a partial ROI and break down the value of these as value delivered to SPs and to end users. The total SP-assessed value of XSEDE was \$11,657,500. XSEDE is by design focused on service to individual users and had well more than 10,000 as of the end of PY3 – 30 June 2014. The difference between total investment in essential core services and the value of those services to SPs as assessed by SPs is \$4,604,250. For ROI on XSEDE essential core services to fail to exceed 1.0, the average value of these services to end users would have to be below \$460 per year. It seems most unlikely that the real value to end users is lower than that, which implies that the ROI on XSEDE core services is in fact over 1.0.

6. DISCUSSION AND CONCLUSIONS

XSEDE's value proposition is that XSEDE, as a single coordinating entity, enables services that would not accrue with multiple, independent national CI centers, and delivers them more cost-effectively than would those centers.

The cost-avoidance thought model suggests in the essential core functions of operating an advanced CI center, XSEDE achieves significant cost avoidance as compared to a model of multiple independent national CI centers. Quantitative and qualitative arguments together suggest the return on investment in XSEDE for the nation is greater than 1.0 (i.e. the value it delivers to the nation exceeds the cost of federal investment).

To answer our first question, "Are XSEDE organization and services a cost-effective way to deliver CI services to the US

open research community?" we say, "Based on the best available data, XSEDE seems more cost effective than an alternate model of multiple independent national cyberinfrastructure service centers." To our second question, "Is the money invested in XSEDE well spent?" we can say, "Based on best data, yes, as ROI is greater than 1.0."

As far as we know, this is the first published ROI analysis for an entire CI enterprise. This analysis has imperfections but we hope it is a useful analysis of XSEDE and a starting point for further analysis of the value of US federal investment in CI. We plan to continue this analysis again next year, including analyses of SP budgets and taking into account user estimates of the value to them of XSEDE and the XD ecosystem.

7. ACKNOWLEDGMENTS & DISCLAIMER

This report was based on research funded in part by the National Science Foundation through grant number ACI-1053575, which funds XSEDE, and in part by the Indiana University Pervasive Technology Institute, founded with funding from IU and the Lilly Endowment, Inc. Richard Moore's effort here was funded by the National Science Foundation grant number OCI-0910847. Moore suggested the phrase "XD ecosystem." Any opinions expressed here are those of the authors and do not necessarily reflect the views of the National Science Foundation, IU, the Lilly Endowment, or XSEDE leadership as a whole. We thank XSEDE leadership and staff and NSF staff for thoughtful discussion and assistance in making this report possible. Editorial and graphics help from Jan Holloway (IU), Jeremy Fischer (IU), and Mathew Stelmaszek (TACC) is gratefully acknowledged. We thank anonymous reviewers whose comments helped us improve the quality of this report.

Disclaimer: All authors receive funding through XSEDE award subcontracts to the University of Illinois Urbana Champaign, and for high performance computing systems operated as part of the XD ecosystem. All authors thus have a vested interest in the overall XD ecosystem being successful and being viewed as a success. Each author is also inclined to believe that their home institution would get significant additional funding if XSEDE were eliminated, so there is not that much vested interest in the outcome, positive or negative, of the analysis of XSEDE itself.

8. REFERENCES

- [1] National Science Foundation. 2006. Cyberinfrastructure: From Supercomputing to the TeraGrid. http://www.nsf.gov/news/special_reports/cyber/fromsctotg.jsp
- [2] Stewart, C.A., R.D. Knepper, M.R. Link, M.A. Pierce, E.A. Wernert, N. Wilkins-Diehr. 2014. Cyberinfrastructure, Science Gateways, Campus Bridging, and Cloud Computing. In: Encyclopedia of Information Science and Technology. Hershey, PA. <http://hdl.handle.net/2022/18608>
- [3] Stewart, C.A., S. Simms, B. Plale, M.R. Link, D. Hancock and G.C. Fox. What is Cyberinfrastructure? In: Proceedings of SIGUCCS 2010. Norfolk, VA, 24-27 Oct 2010. <http://dx.doi.org/10.1145/1878335.1878347>
- [4] Catlett, C., W.E. Allcock, P. Andrews, R. Aydt, R. Bair, N. Balac, B. Banister, T. Barker, M. Bartelt, P. Beckman, F. Berman, G. Bertoline, A. Blatecky, J. Boisseau, J. Bottum, S. Brunett, J. Bunn, M. Butler, D. Carver, J. Cobb, T. Cockerill, P.F. Couvares, M. Dahan, D. Diehl, T. Dunning, I. Foster, K. Gaither, D. Gannon, S. Goasguen, M. Grobe, D. Hart, M. Heinzl, C. Hempel, W. Huntoon, J. Insley, C. Jordan, I. Judson, A. Kamrath, N. Karonis, C. Kesselman, P. Kovatch, L. Lane,

- S.L. Lathrop, D. Lifka, L. Liming, M. Livny, R. Loft, D. Marcusiu, J. Marsteller, S. Martin, D.S. McCaulay, J. McGee, L. McGinnis, M.A. McRobbie, P. Messina, R. Moore, J.P. Navarro, J. Nichols, M.E. Papka, R. Pennington, G. Pike, J. Pool, R. Reddy, D. Reed, T. Rimovsky, E. Roberts, R. Roskies, S. Sanielevici, J.R. Scott, A. Shankar, M. Sheddon, M. Showerman, D. Simmel, A. Singer, D. Skow, S. Smallen, W. Smith, C. Song, R. Stevens, C.A. Stewart, R.B. Stock, N. Stone, J. Towns, T. Urban, M. Vildibill, E. Walker, V. Welch, N. Wilkins-Diehr, R. Williams, L. Winkler, L. Zhao and A. Zimmerman. 2008. TeraGrid: Analysis of Organization, System Architecture, and Middleware Enabling New Types of Applications. L. Grandinetti, ed. *Advances in Parallel Computing*. Amsterdam, The Netherlands. <http://hdl.handle.net/2022/14524>
- [5] National Science Foundation. 2008. TeraGrid Phase III: eXtreme Digital Resources for Science and Engineering (XD) Program Solicitation NSF 08-571. <http://www.nsf.gov/pubs/2008/nsf08571/nsf08571.htm>
- [6] National Science Foundation. 2014. XSEDE: eXtreme Science and Engineering Discovery Environment (Award Abstract). http://www.nsf.gov/awardsearch/showAward?AWD_ID=1053575
- [7] National Science Foundation. 2014. Evaluating and Enhancing the eXtreme Digital (XD) Cyberinfrastructure for Maximum Usability and Science Impact (Award Abstract). http://www.nsf.gov/awardsearch/showAward?AWD_ID=0946505
- [8] Towns, J., T. Cockerill, M. Dahan, I. Foster, K. Gauthier, A. Grimshaw, V. Hazlewood, S. Lathrop, D. Lifka, G.D. Peterson, R. Roskies, J.R. Scott and N. Wilkins-Diehr. 2014. XSEDE: Accelerating Scientific Discovery. In: *Computing in Science & Engineering* 16(5): 62-74. doi:10.1109/MCSE.2014.80
- [9] Furlani, T., M. Green, G. von Laszewski, V. Chaudhary and M. Jones. 2010. Technology Audit and Insertion Service for TeraGrid (Award Abstract). http://www.nsf.gov/awardsearch/showAward?AWD_ID=1025159
- [10] XSEDE. Service Provider Forum: Charter, Membership, and Governance, V12, 20 July 2014. <http://hdl.handle.net/2142/49980>
- [11] XSEDE. Requesting Membership in the XSEDE Federation, V2, 10 July 2014. <http://hdl.handle.net/2142/49981>
- [12] National Science Foundation. 2014. High Performance Computing System Acquisition: Continuing the Building of a More Inclusive Computing Environment for Science and Engineering Program Solicitation NSF 14-536. <http://www.nsf.gov/pubs/2014/nsf14536/nsf14536.htm>
- [13] National Science Foundation. 2014. Major Research Instrumentation Program: (MRI) Instrument Acquisition or Development Program Solicitation NSF 15-504. <http://www.nsf.gov/pubs/2015/nsf15504/nsf15504.htm>
- [14] Jones, A. Keynote Presentation. TeraGrid '07 (Madison, WI, 4 June 2007). <http://www.union.wisc.edu/teragrid07/>
- [15] National Science Foundation. 2015. High Performance Computing System Acquisition: Enhancing the Petascale Computing Environment for Science and Engineering Program Solicitation 11-511. <http://www.nsf.gov/pubs/2011/nsf11511/nsf11511.htm>
- [16] National Science Foundation. 2014. Directorate for Computer and Information Science and Engineering (CISE): CISE Funding. http://www.nsf.gov/about/budget/fy2014/pdf/18_fy2014.pdf
- [17] Kinney, M.F. and C.A. Raiborn. 2011. *Cost Accounting*. South-Western, Mason, OH, 2011. ISBN: 978-1-111-97172.
- [18] Newby, G., A. Apon, N. Berente, R. Eigenmann, S. Fratkin, D. Lifka and C.A. Stewart. 2014. Return on Investment from Academic Supercomputing: Panel presented at: SC14 (Supercomputing 2014), New Orleans, LA, November 21, 2014. <http://hdl.handle.net/2022/19242>
- [19] Apon, A., S. Ahalt, V. Dantuluri, C. Gurdgiev, M. Limayem, L. Ngo and M. Stealey. 2010. High Performance Computing Instrumentation and Research Productivity in U.S. Universities. *Journal of Information Technology Impact*, 10(2), 87-98. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1679248
- [20] Apon, A.W., L.B. Ngo, M.E. Payne and P.W. Wilson. 2014. Assessing the effect of high performance computing capabilities on academic research output. *Empirical Economics* 48: 283-312. DOI: 10.1007/s00181-014-0833-7
- [21] Joseph, E.C., C. Dekate and S. Conway. Real-World Examples of Supercomputers Used For Economic and Societal Benefits: A Prelude to What the Exascale Era Can Provide (Special Study). International Data Corporation (IDC). 2014. <http://case.org/wp-content/uploads/2014/07/IDCReportRealWorldExamplesOfBenefitsOfSupercomputers.pdf>
- [22] Tabor Griffin Communications. High-Performance Computing Contributions to Society. Tabor Griffin Communications, 1998. <http://www.tgc.com/hpcbook/toc1.html>
- [23] PRACE. PRACE (Partnership for Advanced Computing in Europe). Home page. <http://www.prace-ri.eu>
- [24] Compute Canada. Compute Canada. Home Page. <https://www.computeCanada.ca>
- [25] XSEDE. Extended Collaborative Support (ECSS). <http://www.xsede.org/ecss>
- [26] The Arabidopsis Information Resource (TAIR). Home Page. <http://www.arabidopsis.org>
- [27] NCGAS. National Center for Genome Analysis Support (NCGAS). Home Page. <http://ncgas.org>
- [28] Wilkins-Diehr, N. 2007. Special Issue: Science Gateways — Common Community Interfaces to Grid Resources: Editorials. *Concurr. Comput. : Pract. Exper.*, 19(6), 743-749. <http://dl.acm.org/citation.cfm?id=1232528>
- [29] XSEDE. XSEDE Campus Champions. <http://www.xsede.org/campus-champions>
- [30] XSEDE. XSEDE Conferences. <https://conferences.xsede.org>
- [31] Venkatesh, V., Morris, M.G., Davis, F.D., Davis, G.B. 2003. User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly*, 27(3), 425-478. [http://csdl-techreports.googlecode.com/svn/trunk/techreports/2005/05-06/doc/Venkatesh2003.pdf](http://csdl.techreports.googlecode.com/svn/trunk/techreports/2005/05-06/doc/Venkatesh2003.pdf)
- [32] O'Hara, M., C. Yao and M. Ye. What's Not There: Odd Lots and Market Data. 2014. *The Journal of Finance*, 69(5), 2199-2236. <http://onlinelibrary.wiley.com/doi/10.1111/jofi.12185/abstract>