2018 IEEE/ACM 11th International Conference on Utility and Cloud Computing (UCC)

Return on Investment for three cyberinfrastructure facilities: a local campus supercomputer; the NSF-funded Jetstream cloud system; and XSEDE (the eXtreme Science and Engineering Discovery Environment)

Craig A. Stewart[1*], David Y. Hancock[1*], Julie Wernert[1], Matthew R. Link[1],
Nancy Wilkins-Diehr[2], Therese Miller[1], Kelly Gaither[3], and Winona Snapp-Childs[1]

[1]Pervasive Technology Institute, Indiana University, Bloomington, IN 47408
[2]San Diego Supercomputing Center, University of California, San Diego, La Jolla, CA 92093
[3]Texas Advanced Computing Center, University of Texas at Austin, Austin, TX 78758

***Abstract*—The economics of high performance computing are rapidly changing. Commercial cloud offerings, private research clouds, and pressure on the budgets of institutions of higher education and federally-funded research organizations are all contributing factors. As such, it has become a necessity that all expenses and investments be analyzed and considered carefully. In this paper we will analyze the return on investment (ROI) for three different kinds of cyberinfrastructure resources: the eXtreme Science and Engineering Discovery Environment (XSEDE); the NSF-funded Jetstream cloud system; and the Indiana University (IU) Big Red II supercomputer, funded exclusively by IU for use of the IU community and collaborators. We determined the ROI for these three resources by assigning financial values to services by either comparison with commercially available services, or by surveys of value of these resources to their users. In all three cases, the ROI for these very different types of cyberinfrastructure resources was well greater than 1 – meaning that investors are getting more than $1 in returned value for every $1 invested. While there are many ways to measure the value and impact of investment in cyberinfrastructure resources, we are able to quantify the short-term ROI and show that it is a net positive for campuses and the federal government respectively.**

***Keywords*-cost benefit analysis; high performance computing; scientific computing; supercomputing;**

I. INTRODUCTION

In most, perhaps all, countries in the world, there is more worthwhile research that could be done than there is funding to support it. Pressure on the budgets of institutions of higher education and research organizations funded by national research organizations has led to the questioning of many research expenses. “Does this investment matter to the mission of our organization?” and, “If so, is the current level of investment essential, or can it be reduced without damage to our core mission?” are the canonical questions asked about any and all expenses. This line of questioning is particularly challenging where impact is indirect, hard to quantify, or quantifiable only through lagging indicators (such as research papers or patent filings). Investments in research — and particularly investment in infrastructure that supports research — can be hard to quantify and is rarely signified by anything other than lagging indicators.

Because of escalating prices of local computing resources and the claims made about the economies available through cloud computing, the cost savings of different choices for providing computing facilities to research activities. Questions often center on the overall return on investment in computing facilities, and are routinely concerned with the economic choices of “local vs. cloud.” Of course, there is more to this topic than just the financial aspect. Effectiveness, efficiency, and support of discovery and innovation all matter. These matters and finances are all relevant to the issue of utility and cloud computing, particularly as regards the following specific topics:

- Economic and business models of clouds and services
- High performance computing (HPC) and the cloud
- Cloud computing middleware, stacks, tools, delivery networks and services at all layers (XaaS)
- Virtualization, containerization, composition, orchestration and other enablers
- Designs and deployment models for clouds: private, public, hybrid, federated, aggregated

We view cloud and utility computing as specific forms of computing and storage resources within the general term “cyberinfrastructure” as used in the United States (US) — consisting of “. . . computing systems, data storage systems, advanced instruments and data repositories, visualization environments, and people, all linked by high speed networks to make possible scholarly innovation and discoveries not otherwise possible” [1]. This is similar to the term e-Science, as used in Europe. In this paper we will analyze the Return on Investment (ROI) for three very different kinds of cyberinfrastructure resources:

- A “local to campus” supercomputer — Indiana Uni-

---

*Denotes equally contributing authors.

978-1-5386-5504-7/18/$31.00 ©2018 IEEE
DOI 10.1109/UCC.2018.00031

versity's (IU) 1 PetaFLOPS Cray supercomputer called Big Red II [2]. Big Red II is funded exclusively by IU for use of the IU community and collaborators working with members of the IU community.
- The NSF-funded Jetstream cloud system — A national resource, intended in many ways as a national utility for federally supported research, particularly for scientists who need interactive and cloud services rather than batch-style HPC resources [3], [4], [5], [6].
- The eXtreme Science and Engineering Discovery Environment (XSEDE) [7]. XSEDE is a federally-funded cyberinfrastructure support organization.

This paper expands upon and continues two threads of earlier research in ROI of cyberinfrastructure: research on the ROI of federal funding for XSEDE [8], and ROI on university investment in local cyberinfrastructure resources [9]. As such, it should help advance understanding of options relative to the economics of cloud computing and computing in support of research as a utility. This paper also considers factors related to enabling use of computing resources and supporting research activities, and if those resources can be considered a commodity.

II. RETURN ON INVESTMENT AS APPLIED TO CLOUD, UTILITY, AND HIGH PERFORMANCE COMPUTING

There are many ways to define the value of an investment in something. Here, we will focus primarily on return on investment, defined by the financial community as, "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 [10]. There are other financial concepts relevant in a discussion of the value of cyberinfrastructure, including all definitions from [10]:

- 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 versus a hypothetical other way.
- Value added, defined as, "An activity that increases the worth of the product or services to the customer. . . "
- Opportunity cost, defined as, "The benefit forgone because one course of action is chosen over another."

The ROI in cyberinfrastructure is not as simple to calculate as buying and selling a stock. Investments in cyberinfrastructure often precede the creation of *returns* by years or sometimes decades. Even restricting oneself to financial returns, there are many possible approaches to measuring ROI in cyberinfrastructure.

Among recent published papers, two are particularly notable. Hyperion Research estimated the ROI for high performance computing (HPC) systems by surveying across a number of industries and academics, asking about the return on investment of HPC for particular successful projects [11]. That survey found an ROI of up to $673 in revenue per dollar of HPC invested. As they are careful to point out, this result is based on a collection of *success stories*, so this figure should be regarded as an upper bound on what ROI in HPC might be. It is not an unbiased estimate in that it is not based on the cost of HPC investments by an organization and the returns on those investments over time; it is a sampling of self-reported, specific, best case project results.

The National Center for Supercomputer Applications at the University of Illinois Urbana-Champaign took a different approach in a report analyzing the financial impact of the Blue Waters supercomputer project [12], [13]. The primary measures of impact considered in this report were overall regional impact, number of jobs created, and grant income to the University of Illinois Urbana-Champaign (UIUC). The largest part of this ROI analysis was done with a methodology based on "regional impact multipliers." This is a well-known and widely used technique based on the commercial software package IMPLAN [14] but it is somewhat unsatisfying as a way to measure the economic impact of HPC since it is a "dollars in, dollars out" model — it provides the same results for an impact in a road construction project as for an investment in HPC. On the other hand, the analysis of impact on grants awarded to UIUC addresses perhaps one of the most important ways to assess ROI for an institution of higher education.

In Table I, we identify a number of types of economic and other impacts that could be measured as benefits and outcomes of investment in cyberinfrastructure, with references to previously published studies of ROI and impact of cyberinfrastructure investments. In the case of a campus or national computational resource, one way to think about a measure of the value is that the "return" is the value of what a similar resource would cost on the commercial market. This is our approach in general: measure full costs as investment, measure the financial value of a return in some reasonable way to calculate ROI. We are careful in all cases to calculate a conservative estimate of the return, so that the calculated ROI values are reasonable estimates of the actual ROI values. We recognize that we are, to a certain extent, mixing the concepts of ROI and cost-avoidance by treating "avoided costs" as a form of return. This seems the only practicable way to apply the concept of ROI to investments in cyberinfrastructure in public sector research.

In addition to the economic measures of investments in cyberinfrastructure, there is the very general term *impact* that is often asked about in the form of "what is the impact of investment in . . . ?" Impact can take many forms, perhaps the weakest of which is the number of publications derived from the use of a particular cyberinfrastructure resource. Publication counts are often viewed as a weak form of impact because, given the proliferation of journals and

Table I
POSSIBLE KINDS OF BENEFIT

Area of Benefit	Measure of Benefit	Ways to Measure Benefit or Examples of Prior Studies Using This Methodology
	Economic Benefits	
Benefit to CI facilities user	Financial value of time saved	Cost of the time that would have been spent by end users in absence of CI resource
Training	Value of training materials used directly by end users	Comparison of value of training materials with market processes for commercially created training materials
Training	Value of re-use of training materials created by organizations operating CI facility	
Training	Value of a well-trained employee ready to enter STEM workforce	
Value of hardware resources	Value of investment in CI facilities that would have had to have been made by a research organization without use of some other CI facility	[9]
Value in terms of personnel resources	Value of support and consulting from CI resource provider	[9]
Grant income to CI facilities users	Monetary income	[12], [13]
Patents and commercialized products	Monetary income	[15]
Economic impact	Regional economic impact as measured by economic models (IMPLAN)	[15]
Economic impact	Jobs created directly as a result of investment	[16], [17], [18]
	Impact and other Benefits	
Research innovations	Number of papers produced, number of citations, impact of resulting innovations and products	[19], [20], [21]
Grants awarded	Number of grants awarded, number an impact of papers produced resulting from such grant awards	[22], [12]
Prestigious awards	Number of Nobel Prizes	

conferences, practical impact is no longer assured to be correlated with the existence of publications. One of the strongest forms of impact is, “How many Nobel prizes have been awarded to people who used a particular piece of cyberinfrastructure?” However, this is an indicator that may lag years to decades. Nevertheless, such examples do exist — XSEDE has enabled research results that have contributed directly to three different Nobel Prizes: the 2013 Nobel Prize in Chemistry “for the development of multiscale models for complex chemical systems” [23]; the 2013 Nobel prize in Physics for the discovery of the Higgs Boson [24] and the 2017 Nobel Prize in Physics “for decisive contributions to the LIGO detector and the observation of gravitational waves” [25].

III. METHODS AND RESULTS

The three resources we analyze and the means by which we estimate ROI vary. For this reason, we present the methods for estimating ROI and the results of those analyses by resource in three sections below, and then the results are discussed together in the following Discussion section.

A. *ROI for a campus-based CI resource: Big Red II*

For analysis of ROI of campus-based CI, we examined the acquisition and operations costs of an IU resource — the Big Red II supercomputer [2]. In 2012, Big Red II was announced as the first petascale supercomputer exclusively funded by and operated for a US university. It was put into production in 2013 and the analysis provided is based on usage through calendar year (CY) 2017. For the “return” of this system, we used the value of the cost to purchase the same number of hours from Amazon Web Services (AWS) as were actually utilized on Big Red II during the given period.

The investment costs are detailed beginning with Equation 1. This comparison is conservative in that actual charges for AWS would be higher. The high speed network is a significant contributor to the cost of Big Red II — it is very important to the performance of the system due to the usage involving large-scale parallel jobs with high communication requirements and tightly-coupled data movement. Moving these calculations to AWS would require more charged-to-the-customer processor time than the time used on Big Red II. These data do not include interactive use for building applications and testing workflows on interactive portions of the system which are included in the locally provided service, but would need to be provisioned and built in the commercial cloud, a convenience that would incur cost over time.

Processor varieties refresh faster in the commercial cloud and those advances result in costs that are passed on to the consumer. By using today’s rates with less-expensive instance types, we feel the comparison is an adequate compromise. That is a contributing factor to the challenge of comparison to AWS — resources that precisely match what is provided in Big Red II are not available, and calculating cost over time involves using a litany of frequently changing matrices of AWS costs. Therefore, we calculate ROI using

AWS rates from February 2018. This means that the "value" in terms of calculated use of AWS is conservative, and lower than what the time would have cost, because the current cost of AWS instances is lower now than during prior years of use of Big Red II, a system that has been in production for five years.

Calculating data storage and transfer costs (which can sometimes be waived as an academic institution) is a separate exercise in itself. The acquisition cost used is for the computational system only, with minimal internal storage, and does not include external storage system cost for system or site-wide file systems. For facilities and operations, we include the cost of system administration (to represent the cost for maintaining system software) and operation cost for power based on monthly measured load at the switchboard.

As illustrated in Equation 1 total investment TI_{local} is the sum of the acquisition cost AC, system maintenance cost MC, total system administration salary TS, power cost TPC, and co-location (space) cost TCC.

$$TI_{local} = AC + MC + TS + TPC + TCC \tag{1}$$

Total system administration salary TS shown in Equation 2 is the sum of the yearly salary YS and yearly salary YS times the fringe benefit rate FBR to account for total compensation over n years where the yearly salary YS and the fringe benefit rate FBR vary, and typically increase, each year n. Year one in our case was a partial year with system deployment beginning in April.

$$TS = (YS_1 + YS_1 \times FBR_1) \times 0.75 + \sum_{n=2}^{5} YS_n + YS_n \times FBR_n \tag{2}$$

Total power cost TPC as shown in Equation 3 is the product of average consumed power P_{avg}, 379 kW in the case of Big Red II, average power usage efficiency PUE_{avg}, and the average per kilowatt hour cost kWh_{avg} over a period of 4.75 years.

$$TPC = P_{avg} \times PUE_{avg} \times kWh_{avg} \times 24 \times 365.25 \times 4.75 \tag{3}$$

The cost for space, total co-location cost TCC shown in Equation 4 is a simple product of the per rack co-location cost RCC, the number of racks NR, twelve in this case, and the number of years n, with our partial year calculated as a full year for a total of 5 years. The rack co-location cost RCC has not varied significantly in our case, so the most recent rate was used.

$$TCC = RCC \times NR \times n \tag{4}$$

The "value" of AWS is a straightforward calculation as illustrated in Equation 5 where the total value TV_{AWS} is the core hours consumed CH multiplied by the reserved instance hourly cost IC. The calculations can also be done using consumed node hours and dedicated host rates but the results for total cost are similar so core hours were used since they are more commonly reported.

$$TV_{AWS} = CH \times IC \tag{5}$$

The AWS instance types used for this analysis are c4.large (two vCPUs based on hyper-threaded Intel Xeon 2666v3 cores and 3.75GB RAM) and p2.xlarge (four vCPUs of Intel Xeon v4 with 61 GB RAM and an NVIDIA K80). Pricing varies widely depending on how one would structure such a large agreement with Amazon — neither on-demand or spot pricing seem appropriate for such analysis. Spot pricing may work for research groups with extensive check-pointing, no deadlines, or short running jobs, but does not reflect the sort of resource that researchers require in order to conduct their mission effectively. On Big Red II, jobs can be as long as two weeks, often without check-pointing for serial workloads. Among instance types, we chose to calculate value based on reserved instance costs, a significant discount over on-demand rates, but more expensive than spot instances that can be preempted with two minutes notice.

From mid-2013 through 2017, Big Red II delivered 637,874,648 core hours to the IU community. The acquisition cost (AC) of the system was \$7.5M; including system administration, power, cooling, space, and extended system maintenance, the total investment TI_{local} for that period was \$10,132,097. If one were to procure the same number of comparable instance hours within AWS (c4.large), the total value TV_{AWS} would be \$24,877,111–\$37,634,604, assuming reserved up-front payments in 3-year or 1-year terms, respectively. In other words, a lower bound on the ROI of investment in local resources for Big Red II ranges from 2.5 to 3.7 (TV_{AWS} / TI_{local}) as highlighted in Table II.

Table II
ROI FOR BIG RED II 2013–2017

Value Proposition	Cost/Value
Total Value AWS (TV_{AWS})	\$24,877,111–\$37,634,604
Total Local Investment (TI_{local})	\$10,132,097
ROI	**2.5–3.7**

Our analysis is also conservative in that it does not include the cost of comparable AWS GPU instance types. Big Red II contains 676 GPUs — usage of those GPUs is not included in the core hours delivered above. Considering the available GPU hours during that period would add another \$10,244,635 to the value, in terms of "AWS-equivalents," moving the range of ROI estimates up to a range of 3.1 to 4.2. Put another way, after approximately 172M core hours of use, the total investment in the local system TI_{local} is equivalent to the total value on AWS TV_{AWS} at the 1-year reserved instance rate, or 260M core hours at the 3-year rate, which equates to 1-2 years of normal usage of Big Red II.

B. ROI for an NSF-funded cloud resource: the Jetstream cloud system

Jetstream is a cloud computing platform funded by the National Science Foundation, specifically for researchers, through the 2013 round of the ongoing High Performance Computing System /Acquisition solicitation [26]. Of the resources discussed here, it is the one closest to being a "utility-type" cloud computing system similar to commercially available Infrastructure as a Service (IaaS) environments. It is based on source projects from the OpenStack cloud computing software suite and standard Dell hardware suitable for hosting cloud systems with production hardware at IU and the Texas Advanced Computing Center (TACC). However, Jetstream operates with the Atmosphere interface created by the University of Arizona (UA) and the CyVerse project (formerly known as iPlant) and in this regard presents an interface specifically designed for use by research scientists who may or may not be computational experts [27], [28].

We now have seven full quarters of utilization of the resource and adoption has climbed steadily over that period, with over 118M vCPU hours in this case. The past two quarters of usage are the best measure of projecting future usage as the number of average active instances (virtual machines in this case) has climbed significantly toward what appears to be a steadier state. Usage has risen steadily from an average number of active virtual machines (VMs) of 342 during our first quarter of production to over 1100 for each of the last two quarters. For calendar year 2018, to date the average monthly running instance totals have a standard deviation of 6%.

Projecting the two most recent quarters of usage for an entire year is, we believe, the most reasonable approximation to consider a cost comparison. Doing a projection for a longer period of time has more variables and we expect utilization to continue to increase, perhaps with some additional costs. At the time of calculation, within the two production regions (IU and TACC), the average number of vCPUs/instance is 18.5 and 16.9, with memory consumption per instance (in gigabytes) at 48.4 and 44.6, respectively. The comparable instance size for these larger memory and often computational workloads would be a c4.4xlarge within AWS, which provides 16 vCPUs and 30 GB of RAM. A jump to a larger instance size for more memory would skew the data too much. Jetstream uses the same Intel Haswell processors as this instance type but a higher bin part, 2680v3 vs 2666v3. Since Jetstream also uses Ethernet instead of InfiniBand for networking, it's the most closely matched comparison described in this paper. The on-demand cost for a c4.4xlarge instance type at the time of analysis was $0.796 per hour in the US East region and drops to $0.504 per hour if a 1-year commitment is made.

The Jetstream system as acquired and installed at IU, TACC, and UA had an initial system acquisition cost to the National Science Foundation of $5.5M within an 18-month acquisition award that totaled $6,576,101 with an allowable 20% annual budget for operations and maintenance (O&M), $1,315,220 per year in this case. The award is intended to support at least 5 years of operations and the O&M budget not only includes operations of the system, but also aspects of user support and management of the program, costs that do not disappear if utilizing a commercial environment. To calculate ROI conservatively, we still include all O&M funding and prorated system acquisition and implementation cost for this one year period, not just system administration and facilities costs as was the case for Big Red II. Specifically in Equation 6 below, ROI_J equals the average number of running virtual machine instances NVM_{avg} times the running hours over a year at the AWS on-demand instance cost *IC* divided by the sum of the prorated system acquisition and implementation cost *AIC* (the value of the implementation award spread over 5 years) and the yearly operations and maintenance cost *OM*.

$$ROI_J = \frac{NVM_{avg} \times 24 \times 365.25 \times IC}{AIC \times 0.2 + OM} \tag{6}$$

As summarized in Table III below, if average monthly usage for the two most recent quarters was the yearly average (NVM_{avg}=1152), those cycles would have a cost of $8,038,352 at the on-demand rate listed above, the relevant basis for comparison for Jetstream as it operates as an on-demand cloud allowing for extended execution times. This results in a calculated ROI value of 3.1 over a 1-year period; by prepaying annually for those instances the ROI would be reduced to 1.9. The Jetstream acquisition and implementation award also included storage costs (approximately 1 PB of aggregate usable block and object storage), an additional cost in the commercial cloud that increases the Jetstream ROI but not considered in TV_{AWS}.

Table III
1–YEAR ESTIMATED ROI FOR JETSTREAM

Value Proposition	Cost/Value
Total Value AWS (*1152 c4.4xlarge instances*)	$5,089,609–$8,038,352
Total NSF Investment (*AIC x 0.2 + OM*)	$2,630,440
ROI	**1.9–3.1**

C. ROI for XSEDE

XSEDE [7] is a different sort of cyberinfrastructure resource than Big Red II or Jetstream in that XSEDE itself does not provide cycles or storage. Rather, XSEDE supports the use of computational resources funded by the NSF, such as the Stampede2 supercomputer at the Texas Advanced Computing Center, Bridges supercomputer at the Pittsburgh Supercomputing Center, the Comet supercomputer at the San Diego Supercomputer Center, the Wrangler storage system, led by TACC with the IU Pervasive Technology

Institute (PTI) as a partner and host of hardware; and Jetstream, led by IU PTI with several partners and hardware located at TACC and UA. XSEDE supports the use of these resources with consulting, programming, online help, educational materials, and by managing the process of allocating these resources. The institutions that provide resources supported and allocated by XSEDE are generally funded by the NSF and referred to as "Service Providers (SPs)." We published the first comprehensive analysis of ROI for any cyberinfrastructure facility in 2015 when we did our first evaluations of the ROI on XSEDE.

Like Jetstream, the costs accrue to the NSF and the benefits accrue to a variety of different constituencies such as end users, principal investigators, campus CI administrators, service providers, and the NSF. A description of the benefits and the communities to which they accrue follows.

We are most concerned here with ROI on federal investments in cyberinfrastructure, so we take the cost of XSEDE to be the NSF budget for XSEDE, and the value of XSEDE to be the aggregate value of all of the above sorts of return to the US and its populace generally. Since our initial paper on ROI for XSEDE in 2015, we have been collecting more thorough data on the value of XSEDE. We have organized our data collection by the program year (PY) of XSEDE. The time periods for which we have calculated ROI are PY4 (07/01/2014 to 06/30/2015), PY5 (07/01/2015 to 08/31/2016), and PY6 (09/01/2016 to 08/31/2017; in some reports this is also described as PY1 of XSEDE2).

We measured the value of XSEDE to various constituencies and via assorted mechanisms through several different methods described below in Table IV. Much of the data collected for the ROI analysis was gathered through surveys, asking recipients of XSEDE services to estimate the benefit they received from XSEDE in terms of the question, "How much time would it have taken you to do what you did without the help of XSEDE?" All surveys were administered with support from the IU Center for Survey Research (CSR) to ensure that respondents knew that their responses would be confidential and managed by a well-regarded survey research center. End-users were asked to consider the value they received from XSEDE in terms of the amount of their own additional time it would have taken to do their research or educational activities without XSEDE support. The end-user assessment was conducted via the XSEDE User Portal (XUP), using an anonymous link to the online survey instrument, where any user could participate upon logging in to the XUP. Participation was completely optional and recruitment methods were passive, meaning that users were not broadly contacted and solicited to participate in the survey, nor were they precluded from accessing resources via the XUP if they opted not to participate.

Service Providers were similarly surveyed, again via the IU Center for Survey Research, to assure anonymity. Each SP lead was sent a direct email with a worksheet, asking the lead to factor the value derived for their institution from a broad range of XSEDE computational, training, consulting and extended collaborative support, and systems integration resources. SP leads were also asked to indicate the approximate number of FTEs and dollars that would be required for the institution to provide (or replace) the resources and services presently operated or provided by XSEDE. This assessment also asked SP leads to consider whether or not and, if so, at what level, services would (or could) be replaced if they were not accessible through the XSEDE. Responses to the SP assessment were aggregated, redacted, and anonymized by IU CSR staff before being shared with the XSEDE ROI lead to ensure confidentiality and that no bias was inadvertently introduced into the analysis process. Analysis includes a comparison of the FTEs and dollars required for individuals to acquire, or institutions to provide (or replace), the resources and services presently operated or provided by XSEDE to the actual FTEs and dollars allocated through the XSEDE.

The value of XSEDE's training services, through in-person training, webinars, and online videos, was assessed by measuring actual participation and use of these types of training activities and then calculating the value of these services according to current commercially available services. The value of XSEDE across these different user constituencies and types of value provided is summarized in Table V. A report detailing the data collected as part of the analysis of ROI for XSEDE is available at [31].

IV. DISCUSSION

In this paper we have used a somewhat modified version of the definition of return on investment than typical, as discussed in Section II, because a "cash in, cash back" direct relationship does not exist for things like infrastructure used for academic research and development. We have measured costs in a traditional way, and measured "return" as the value of investments in cyberinfrastructure where that value is measured conservatively and based on the value of similar services had they been bought on the open market. As noted earlier, we have mixed the concepts of ROI and cost avoidance, but believe we have used the ROI concept responsibly and in a way that clearly conveys the value of investments in cyberinfrastructure.

Below we discuss the results from the three very different cyberinfrastructure entities in terms of ROI as a financial benefit and in terms of other types of value added and impact. After that, we discuss the implications of our findings on more general matters related to the economy of computational resources on the Internet and the idea of cloud computing resources as utilities.

A. Big Red II

In terms of financial Return on Investment, the value of IU's investment in a Cray supercomputer is highly positive:

Table IV
CONSTITUENCIES AND METHODS USED TO ESTIMATE "RETURN" FROM XSEDE SERVICES

Constituency	Service	Means used to estimate value
Service providers	Services provided by XSEDE to SPs	Surveyed SPs to determine how many person-months SP representatives believe XSEDE services cost them; calculate value based on actual average cost per relevant group of XSEDE staff
Campus CI resource administrators	Online Tools	Not yet measured
Principal Investigators with ECSS	Consulting and programming	Surveyed PIs to determine the number of person-months that they estimate they saved as a result of use of ECSS services; costed at the average per month cost of XSEDE ECSS staff
End users	Overall value of XSEDE	Surveyed end users as to the number of hours of their own time that use of XSEDE had saved them; calculated Return based on an average estimated value of $50 per hour across all user types (beginning students to Nobel laureates)
End users	In-person training	Total hours of training delivered (session length x participants) x $11.29 per hour (cost for training from Fred Pryor on use of spreadsheet software [29])
End users	Online training	Total hours consumed x $4.62 per hour (cost for on-demand basic web development class from Udemy [30])
Educators	Documentation as a source for curricular materials	Not yet measured

from just over 2.5 times to 4.2 times the value of what could have been obtained if computer resources had been purchased from Amazon Web Services. The cycles delivered to the local IU community are actually more 'valuable' to the local community than the resource we are using to quantify the financial value of the local resource, because the high speed interconnect of Big Red II means that workloads can be run locally that might be impractical to have run in even relatively similar amounts of elapsed time on AWS. Additionally, Big Red II will remain in service for at least 12-18 months with minimal additional operational costs, further increasing returns.

The delivery of HPC resources to the IU community as a "free good" available to all members of the research community goes back more than 60 years, to the creation of the first IU Research Computing Service and its first appointed Director, Marshall Wrubel. Since the 1990s, IU leadership has believed that local HPC and supercomputing resources constitute a strategic advantage for the university in terms of recruiting researchers and students, retaining faculty, and accelerating innovation. This was codified in the first Indiana University Information Technology Strategic Plan [32]. In fact, IU's current IT strategic plan states that "IU should pursue strategies that approximate a philosophy of abundance, within reason, towards unmetered availability of basic IT services, support, and infrastructure for creative activity, storage, computation, communication, and other activities fundamental to the work of the university via any appropriate sourcing strategy." A highlight of this is open access to the Big Red II supercomputer, along with storage resources of significant scale.

This strategy works well for researchers who need supercomputing resources and is enabled at IU by acquiring large, centralized computing resources that are shared amongst the institution without usage fees. Compared to what a single department or user could acquire and operate individually, the resources are abundant and unmetered, even if still finite.

Furthermore, the use of Big Red II as a local resource sharpens focus on the research, scholarly, and artistic communities of IU. It represents a significant investment visible to the entire university community, acquired with financial support from budgets under the direct control of the IU President, with a commitment that Big Red II would support a diversity of disciplines and sub-disciplines.

As of the writing of this report, Big Red II has been used by a total of 134 departments and 214 academic disciplines within the two core (research) campuses of IU – IU Bloomington, and Indiana University – Purdue University Indianapolis. This is particularly important for IU as a liberal arts school, as the local ownership of HPC resources means that any researcher, scholar, or artist from any discipline can use the resource. This is reflected in the diversity of departments that use the system, from traditional disciplines like physics and astronomy to literature scholars in many disciplines, and even artists and students in the Department of Interior Design and Apparel Merchandising.

Over the past 10 years, we have correlated IU grant data and IU users who are PIs, Co-PIs, or Senior Personnel on grants awarded to the university. Using these data we are able to determine additional information such as department, school, campus, group or lab with which these associated researchers are affiliated. This insight affords us the ability to reach out to those researchers to better understand how they are utilizing our CI resources and assist them with more efficient, and most times, faster use of these systems enabling them to reduce their time to science or discovery.

While we are in these consultations, we often have the opportunity to discuss how the use of CI has changed their ability to be successful in obtaining current and future opportunities. For example, in FY17, IU brought in $504M in grants and contracts – 74% of that total was associated with those PIs, Co-PIs or Senior Personnel as users on our systems. In FY18, IU brought in \$608M and our personnel percentage increased to 86%. Correlating these data over time has allowed us to have targeted and successful interactions with our researchers in successful grant awards. For the 5-year operational period of Big Red II, researchers on that system have been awarded \$444M, with an approximate F&A return to IU of \$142M. Pro-rated over 5 years, the yearly F&A contribution would be \$28.4M. With the cost of BRII factored over 5 years, this accounts for a positive variance of \$26.4M/year – paying for the yearly cost nearly 14 times over.

B. Jetstream

The ROI analysis for Jetstream can be used as a proxy for any organization considering investment in a private cloud. What we have found on the financial front is significant benefit to the NSF in investment in a private cloud as compared to the cost of purchasing resources from roughly equivalent cloud providers, with ROI ranging from 1.9 to 3.1. While the ROI estimates for Jetstream are in the early stages, given increasing usage levels and similar levels of "returns", continued growth in ROI is likely.

Our work with Jetstream provides particular insight on matters important to utility and cloud computing such as cloud computing middleware, stacks, tools, delivery networks and services, and virtualization, containerization, composition, and orchestration. A critical part of the success of Jetstream has been that it was designed to meet a particular need in cloud computing. Jetstream is intended, as specified by the NSF [26], to serve communities needing flexible and user-friendly environments. We thus implemented Jetstream with a default interface being the Atmosphere interface developed by the University of Arizona, which is very intuitive to use and which presents users with a catalog of pre-built VM images to use for a variety of purposes. At present there are more than 15 such pre-built images that researchers can use ranging from VMs with "bare" Linux operating system variants installed, to particular applications, to gateways into entire workflow systems (such as the bioinformatics workflow system Galaxy [33]).

One of the critical success factors in deploying Jetstream and attracting users from communities that have not previously made much use of NSF-funded cyberinfrastructure is the fact that it was designed and implemented with those communities in mind. And indeed, the system is attracting new researchers as intended and as called for in NSF solicitation 14-536. More than 80% of Jetstream users have never executed a batch job on any other XSEDE resource; some of those users are having, through Jetstream, their first use of NSF-funded cyberinfrastructure. Additionally, the experience in operating and supporting users within such an environment should provide "returns" in the form of a skilled workforce both for administrators and users that can now apply their knowledge to not only Jetstream but other private or public cloud environments (see [3], [4] for more information).

At the same time as providing tools to empower researchers who are not themselves computational science experts, Jetstream is designed to allow considerable freedom to those who have deep computational skills. Users may interact directly with the OpenStack API or the Atmosphere API, and as a result many researchers have made extensive use of current cloud orchestration services without significant help from XSEDE or Jetstream support staff. This autonomy to implement advanced tools has been particularly valuable to enabling very large-scale analyses with Jetstream. Another aspect of the effective use of the Jetstream system has to do with system management policies.

- Restrictions on levels of over-subscription. We take a different approach to over-subscription of vCPUs to physical CPUs than that taken by some commercial cloud providers. We observe a maximum of two vCPUs per physical core to provide a cloud computing experience that feels both very responsive to user needs and that provides a consistent sense of performance as experienced by the user (one of the common compliments we hear about Jetstream as compared to commercial cloud environments is that the user experience feels more consistent day to day and month to month).
- Percent usage vs capacity for bursting. Commercial cloud providers are said to run their cloud resources at 50% or less utilization, with a of very large cloud resources distributed essentially worldwide. This level of utilization is the critical factor that enables cloud resources to meet bursts in demand for resources. Based on community needs, we have chosen to operate at somewhat higher levels of average utilization aiming on average for about 75% of the maximum possible utilization. This, plus the smaller overall scale of Jetstream, means that it has less bursting capacity than commercially operated clouds provide.
- Retirement of unused VMs. VMs that are left in stopped or suspended state for two weeks are moved to disk storage and then deleted if left untouched for 6 months. Users receive email notification before VMs are deleted.

All in all, we believe that we have put in place a set of system implementation and usage policies that strike a good balance between relatively high utilization of a taxpayer-funded national cloud computing resource within the context of enabling researchers (whose time is also paid for at least partly through public funding) to do their work effectively.

C. XSEDE

Our analysis of ROI of XSEDE has expanded considerably since our first attempt in 2015 [8]. Thus far, and working with still incomplete data, we are able to document

Table V
VALUE OF XSEDE TO VARIOUS XSEDE USER COMMUNITIES

Value Proposition	PY4	PY5	PY6
XSEDE to SPs			
Level 1 SPs	$15,958,215	$13,395,398	$7,338,831
Level 2 SPs	$1,372,766	$1,376,413	$1,908,322
Level 3 SPs	$1,012,758	$1,241,642	$5,977,500
Sub-Totals	$18,343,739	$16,013,453	$15,224,653
ECSS to Principal Investigators	$1,700,000	$7,153,333	$7,358,333
XSEDE to End Users	–	$3,889,800	Survey results still being analysed
Training to Recipients			
In Person	–	–	$837,140.43
Live Online	–	–	Not yet assessed
Recorded Online	–	–	$247,584
Sub-Totals	–	–	$1,084,724.43
Re-use of Training Materials	–	–	Survey still underway
Total value of XSEDE quantified to date	$20,043,739	$27,056,586	$24,752,435
Total XSEDE Budget	$23,562,931	$23,067,000	$18,442,569
ROI	**Greater than 1***	**1.17**	**1.34**

*With qualitative argument about value of particular projects supported this PY

**This table includes minor corrections to data presented in [8] (minor enough that they do not impact the conclusions of that report)

a ROI of 1.3 on NSF investment in XSEDE for PY6. This ROI value is higher than the ROI recorded for the prior program year, which was in turn higher than the ROI recorded for the prior year. This is primarily the result of inclusion of more types of quantifiable “return” data related to the value of XSEDE.

In our original 2015 analysis of ROI for XSEDE, we focused on the ROI as perceived by Service Providers (SPs). We now have three program years of data for value of XSEDE as perceived by SPs. Overall, the aggregate value of XSEDE as perceived by SPs has remained relatively stable but decreased slightly over the past three years. This overall trend is driven by two interesting trends within the SP valuation of XSEDE a decrease in the aggregate value of XSEDE to the so-called Level 1 SPs the largest NSF SPs that are (as dictated by NSF policy) 90% allocated through XSEDE. The other trend is the aggregate increase in value of services delivered to the so-called Level 3 SPs.

We believe that the value of XSEDE as assessed by Level 1 SPs has decreased for a number of factors. First, in its earlier years XSEDE delivered services that were novel, valuable, and not generally or easily available outside of XSEDE. For example, in PY4, XSEDE had just recently worked with Globus to make Globus Transfer available and properly secured as a means to move data into and out of XSEDE, and among resources within XSEDE. This was then new and important. Today, Globus Online [34] is a well-established and widely used subscription service; no one gives XSEDE any significant credit anymore for the fact that this still works. Furthermore, XSEDE no longer does any development of new tools or applications from scratch, and the overall budget for services that target the very largest SPs has been significantly reduced. So, benefit as perceived by the largest XSEDE SPs has gone down at the same time as the investment in services most relevant to these large SPs has gone down.

An interesting trend is the increase in value of XSEDE as perceived by Level 3 SPs. According to XSEDE web sites, “Level 3 Service Providers are progressively more loosely coupled with XSEDE” [35] and, in fact, generally do not offer any resources that are allocated through the XSEDE-run allocations process. Instead, Level 3 providers benefit from their affiliation with XSEDE through information exchange, shared expertise, and the opportunity to publish information about the services they provide via XSEDE web pages and documentation. There has been a strong rise in the value of XSEDE as perceived by Level 3 SPs driven by increases in the assessment of the value of XSEDE per Level 3 SP and by an increase in the number of such SPs.

As a whole, particularly to students and adult learners, XSEDE has delivered tremendous value to the nation. One striking thing about XSEDE is the significant value of training resources. Training resources, using value estimated conservatively from relatively inexpensive commercial offerings for simple computing programs, shows a value of more than $1M in PY6 alone delivered to the US research and education community.

The overall value of Extended Collaborative Support Services offered by XSEDE is terrifically high as shown in Table V. XSEDE’s Extended Collaborative Support Services (ECSS) program [36] provides in depth expert consulting and programming services, delivered in a collaborative setting. An XSEDE staff member is assigned to spend on to several person-months working on a project with a Principal Investigator and their team. A typical project involves creating or adapting a parallel computing program to run on one or more XSEDE-supported cyberinfrastructure platforms. The XSEDE ECSS program is modeled on previous programs that pair experts in computational science with domain experts for in depth collaborations designed to achieve transformative changes in the way research is conducted. Centers introduced programs like the Strategic Applications Collaborations (SAC) at the San Diego Supercomputer Center in the 1990s. The NSF-funded TeraGrid program included Advanced Support for TeraGrid Applications (ASTA), but it was ECSS that really moved the field forward, both through increased funding levels and a more structured approach to collaborations.

Collaboratively developed workplans set expectations for both research teams and ECSS staff. Interviews with PIs after the collaborations uncover valuable feedback that helps the program adapt and improve. PIs are also asked how many person months it would have taken them to achieve the same results without the support of ECSS. It is this return on investment which is reported in this paper - \$1.7M in PY4, \$7.1M in PY5 and \$7.3M in PY6. Often principal investigators say they could not have conducted the work at all without help and direction from ECSS staff. For these responses we use a figure of 24 months. Each ECSS project consists of an investment of 25% (3 months) of a staff member's time over the course of a year.

D. The Internet Economy and Cyberinfrastructure as a Utility

1) Economic and business models of cloud and HPC services: What do the results of our analyses say regarding the economy of advanced cyberinfrastructure and the Internet, and the evolution of cloud computing as a commodity? And what is a commodity? *The Economist* [37] explains that "commodities are vital components of commerce that are standardized and hence easy to exchange for goods of the same type, have a fairly uniform price around the world (excluding transport costs and taxes) and help make other products." Entities that are products such as electricity, Internet access, and computational cycles, become commodities by being first very specialized products, then becoming widely available, and then at some point standardization becomes sufficient to meet the vast majority of consumer needs and pricing becomes a more critical factor driving consumers toward the product provided as a utility.

In considering the examples of Big Red II and Jetstream, we see that the capabilities that each provides are such that they have not yet reached the level of being a commodity. First, in the classic "lease or buy" finance question, it is still cheaper to buy the source of the product (a computational system than lease (purchase cycles as needed from a cloud provider) given sufficient utilization and duration.

Second, differentiation among computational systems is not yet such that computing cycles can be treated as a utility, even amongst commercial cloud providers, particularly if one is leveraging provider-specific application programming interfaces and tools. Our own experience with Jetstream shows that a relatively modest amount of modification of a resource based on the OpenStack cloud software has made it tremendously valuable to the US research community. Our experience serving the IU community with Big Red II has also supported the value of crafting a computational resource to local needs. There are additional quantitative analyses that support the conclusion that ownership of resources (and consequent with that customization to local needs) is valuable. The studies by Apon et al. show everything from more rapid innovation to increases in publication rates and grant income for institutions that invest in local HPC resources [22], [19].

2) Opportunity costs and patterns of investment: There are also considerations in terms of the financial models of institutions that use advanced cyberinfrastructure. IU, the NSF, and many other institutions use a model based on periodic cash investments in advanced cyberinfrastructure resources, between which expenses are often limited to personnel (support and operations) and maintenance. Because the personnel and maintenance costs tend to be far lower than acquisition costs, this approach allows an institution to accumulate one-time monies and make an investment without committing to any particular future and high operational costs. At the scale of the resources purchased by IU, for example, this is preferable to purchasing resources from a commercial cloud provider.

IU has indeed upgraded some number of its HPC systems incrementally over the years. Our experience on average is that incremental hardware upgrades to a system are more expensive financially, take more human effort, and disrupt user activities in a way that makes such upgrades less beneficial overall than the our usual approach of "buy, setup, and run (upgrading software as needed by keeping hardware static)." The issues in particular include the following:

- Disruption to user activities during upgrades and personnel effort. The system setup and configuration are sufficiently complex that an approach of incremental upgrades to hardware are difficult to effect without disruption to the activities of system users. It also takes significantly more staff effort to setup a system once and then put in place a subsequent system upgrade later than to set up a system once (with software updates as needed, of course).
- The effects of hardware heterogeneity. No matter how hard one tries one can never really manage to make a system bought with components months apart be homogeneous. Such heterogeneity can be as modest in its effects or may be sufficiently noticeable to users that they have to think about heterogeneity in systems (where none might otherwise exist) when submitting and running jobs on an everyday basis. Managing vendor maintenance and support on a system containing components of varying ages can also be difficult. When

components leave vendor maintenance and/or support, great complexity is added to system management and maintenance.

- Quality of vendor offerings. While we have no empirical basis for comparisons, our strong impression is that we get better deals from vendors on acquisitions and on maintenance with fewer, larger deals than we would get with more, smaller deals.
- Growth of usage is so fast trying to pace upgrades to expanded usage is unlikely to be a net gain financially. If the adoption and growth of new systems was slow, one might be able to conserve funds by incrementally adding hardware to a system to keep pace with growth in demand. Our experience is that usage expands to nearly the available capacity within a matter of a few months, so that for IU at least more gradual expansions of hardware would be unlikely to stretch out over a long enough period of time to be financially beneficial.

From the standpoint of an individual university or college, the above should not be taken to take away from the value of regular investments in cyberinfrastructure. While IUs pattern of purchases in general is to “buy as large as possible” and also make regular expansions of overall capacity by purchasing systems regularly. Currently, PTI operates one cloud system, one supercomputer, and two parallel computing clusters 9standard Linux clusters with 10 Gbps interconnects). On average and in aggregate, then, we have added capacity to locally purchased systems every two to three years. Similarly, the National Science Foundation tends to make investments in its overall portfolio of advanced cyberinfrastructure resources every year or two. Indeed, Apon et al. [22], [19] have demonstrated that regular investment in HPC systems results in higher numbers of publications and greater success in obtaining grant awards. The key question is how that investment is implemented over time. We note that IU’s colleagues to the northeast at the University of Notre Dame practice with great success a pattern of incremental upgrades to their parallel computing clusters. Purdue University and Indiana University the other two large research universities in the State of Indiana tend to make repeated investments in new systems but then leave each system alone once it is purchased.

We also note that fiscal management processes can dictate options. One of the factors that enables the Office of the Vice President for Information Technology and the IU Pervasive Technology Institute to practice the pattern of investment that we have come to prefer is that, since 1997, we have had an equipment replacement fund for our advanced cyberinfrastructure. Each year, money from the base budget for the Office of the Vice President for Information Technology goes into this fund, and any accumulated monies not used in one fiscal year roll over automatically and without exception to the next year. That lets us accumulate funds for larger deals and wait for opportune times in business and technology cycles to make purchases. If we lived in a fiscal environment of “use it or lose it” each fiscal year, our acquisition practices would be very different.

As regards the use of IU’s Big Red II and other systems funded by IU and made available to the IU community, there is another very subtle aspect to the question of how effectively researchers and learners make use of HPC systems. This has to do with how carefully researchers and students think out analyses and simulations in advance. The actual utilization of such systems is chronically high generally such systems are used at more than 80% of the total capacity available, and when new resources are added, usage of those resources climbs to this high level very quickly (six months to a year). But, an open question is what would happen if IU invested, say, 2/3 as much as it had over recent years. If researchers and students were careful and clever, would that level of investment have resulted in as much new knowledge and innovation as the actual amount of investment made by IU? If so, could the other 1/3 of the funding have been spent in some other way that would have resulted in greater benefit to the university as a whole? This is a question of opportunity cost and is a very difficult question. We know of no good way to approach answering it quantitatively right now. However, the capacity of the local system and fair-shared access to its use provides a modest, but real, check on use. That we routinely get requests for more resources than we can provide from large numbers of our local users suggests that we are not under-investing.

The “buy and run” approach we take at IU also has financial management benefits to IU beyond the favorable ROI of investments in local systems. If IU were to move this workload to an unmetered cloud environment, we would lose the existing social controls on expenditure levels, and convert a periodic cash expense into an ongoing operational cost. This would create significant financial management problems for the university, and would most likely lead, under present levels of usage at least, to not just increased expenses to the university but also to institutional controls on computing resource needs in ways that would decrease overall resource availability. This would fundamentally change the philosophy of the provisioning of computational cycles at IU (and, we believe, at other institutions that employ this approach of periodic purchases). IU’s investment in local resources does not imply a failure to use commercial and public cloud resources. They are indeed used in educational settings and when short term bursts of capacity are needed.

There may be two circumstances in which the economics of cloud computing make it a better choice for some institutions than the “periodic purchase” approach of IU. IU and the NSF operate fairly large cyberinfrastructure resources often with purchase prices in the $10M and above range. For institutions with the financial necessity to make smaller investments in computational resources, it may be wiser to purchase cloud resources on the fly and provide local support and customization of those resources for use by a particular set of users. At the other end of the scale, entities that use resources at a large enough scale and with a pattern of use that enables them to use spot pricing may benefit right now from use of cloud resources rather than "owning your own." Globus [34] is one example of this based in the US.

Utilizing spot instances does not guarantee benefit versus local ownership. Prior work by Bauerdick et al. [38] leveraged spot instances with workflow-specific optimization within a domain for the CMS experiment using the Fermilab HEPCloud Facility which shows that closing the ROI gap is possible to some degree. But after using over 15 million hours (approximately a month of time on a modern supercomputer) they concluded, "The steady-state cost [of AWS] came to 1.4±12% cents per core-hour, which is not much larger than the estimated 0.9±25% cents per core-hour for the Fermilab data center." Even in the best case, a difference of 0.11 cents per core hour will accumulate over the life of a project to the tune of $700K using the core hours from our first example and in the worst case, a 0.89 cents per core hour difference would result in a $5.7M increase in required investment. That is large in the context of an entire CI facility, but could be considered reasonable (or not very large) for a specific project or lab.

3) Progress towards utility services: XSEDE may represent an interesting step toward the commoditization of computing cycles and advanced cyberinfrastructure resources supporting research in the US. XSEDE supports a set of diverse resources, but in ways that provide consistency of user experiences across systems. For example, authentication and file movement work the same way for XSEDE-supported systems, and there are many other operational similarities. Centralized allocations, consulting, and programming provided by XSEDE staff make it easier to move from one resource to another. Indeed, in handling requests for allocations on XSEDE-supported systems, XSEDE is able to treat computational resources in part as if they were commodities. Since there is more demand for XSEDE resources than available time, requests usually exceed available resources by a factor of three or more; approved allocations are often at reduced levels or provided on systems other than those requested. XSEDE staff are often able to provide resources to support such a person's request by first providing an allocation on a similar supercomputer, and then helping that person adapt to a system they did not intend to use. In summary, the resources supported by XSEDE are not yet a commodity, but the impact of XSEDE is to make them seem more like a commodity than a standalone product.

V. CONCLUSIONS

First, the financial analysis of IUs Big Red II supercomputer and the NSF-funded, IU-led Jetstream cloud system suggests that for both there is still financial efficiency in operating resources locally, where expertise and facilities exist, rather than buying resources from commercial cloud providers. In addition to that, customization of these facilities to particular needs and policies of particular constituencies seems to positively affect the impact of those resources for their intended uses. So at least at the scales of the examples we have considered here, neither HPC nor cloud computing is yet to the point of being a "commodity."

The analysis of the cost effectiveness of XSEDE is perhaps the most interesting in terms of understanding the position of the US research community, which appears to be on the road toward managing advanced cyberinfrastructure resources as a commodity. XSEDE creates one national review and prioritization process for more than a half-dozen advanced cyberinfrastructure resources funded by the US government, and then sorts and adapts the most highly prioritized requests to the resources available. These advanced cyberinfrastructure systems are not yet commodities, but the impact of XSEDE is to make them more like commodities, more easily accessible, and used by a community of US researchers that grows larger each year, and in so doing aids high impact research [21].

Overall, we note that the assessment of return on investment in cyberinfrastructure is an area now with many people doing very interesting research. While there are inherent challenges in assessing the value of investments in advanced cyberinfrastructure, this area of research is both important and evolving rapidly over time. We plan future work in ROI of investments in cyberinfrastructure that will continue to expand the financial aspects of the analyses presented here and add more analysis of the value of the research enabled by cyberinfrastructure supporting research.

ACKNOWLEDGEMENT

This work has been supported in part by the Indiana University Pervasive Technology Institute, created with the help of a grant by the Lilly Endowment, Inc. Other grants and awards that have contributed to this report include: NSF Award 1445604 for Jetstream, and NSF Awards 1053575 and 1548562 for XSEDE. Any opinions expressed here represent those of the authors and the authors alone, and may not necessarily be shared by the agencies that have funded this research.

We thank all of the many partner institutions involved in XSEDE and Jetstream for their contributions to this work (and to making the services of XSEDE and Jetstream possible). We thank anonymous reviewers for thoughtful suggestions on earlier versions of this report.

REFERENCES

[1] C. A. Stewart, R. Knepper, M. R. Link, M. Pierce, E. Wernert, and N. Wilkins-Diehr, "Cyberinfrastructure, Cloud Computing, Science Gateways, Visualization, and Cyberinfrastructure Ease of Use," in *Encyclopedia of Information Science and*

Technology (M. Khosrow-Pour, ed.), Hershey, PA: IGI Global, fourth ed., 2018.

[2] Indiana University, "Big Red II at Indiana University," 2013.

[3] C. A. Stewart, T. M. Cockerill, I. Foster, D. Hancock, N. Merchant, E. Skidmore, D. Stanzione, J. Taylor, S. Tuecke, G. Turner, M. Vaughn, and N. Gaffney, "Jetstream - A self-provisioned, scalable science and engineering cloud environment," in *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*, (St. Louis, MO), pp. 29:1–29:8, 2015.

[4] C. A. Stewart, D. Y. Hancock, M. Vaughn, J. Fischer, L. Liming, N. Merchant, T. Miller, J. M. Lowe, D. Stanzione, J. Taylor, and E. Skidmore, "Jetstream - Performance, Early Experiences, and Early Results," in *Proceedings of the XSEDE16 Conference*, (St. Louis, MO), 2016.

[5] Indiana University, "Getting Started with Jetstream," 2017.

[6] D. Y. Hancock, C. A. Stewart, M. Vaughn, J. Fischer, J. M. Lowe, G. Turner, T. L. Swetnam, T. K. Chafin, E. Afgan, M. E. Pierce, and W. Snapp-Childs, "Jetstream-Early operations performance, adoption, and impacts," *Concurrency and Computation: Practice and Experience*, p. e4683, 9 2018.

[7] J. Towns, T. Cockerill, M. Dahan, I. Foster, K. Gaither, A. Grimshaw, V. Hazelwood, S. Lathrop, D. Lifka, R. Roskies, J. R. Scott, and N. Wilkins-Diehr, "XSEDE: Accelerating Scientific Discovery," *Comput. Sci. Eng.*, vol. 16, no. October, p. 62, 2014.

[8] C. A. Stewart, R. Roskies, R. Knepper, R. L. Moore, J. Whitt, and T. M. Cockerill, "XSEDE Value Added, Cost Avoidance, and Return on Investment," in *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*, XSEDE '15, (New York, NY, USA), pp. 23:1–23:8, ACM, 2015.

[9] A. S. Thota, B. Fulton, L. M. W. Weakley, R. Henschel, D. Y. Hancock, M. Allen, J. Tillotson, M. Link, and C. A. Stewart, "A PetaFLOPS Supercomputer As a Campus Resource: Innovation, Impact, and Models for Locally-Owned High Performance Computing at Research Colleges and Universities," in *Proceedings of the 2016 ACM on SIGUCCS Annual Conference*, SIGUCCS '16, (New York, NY, USA), pp. 61–68, ACM, 2016.

[10] M. R. Kinney and C. A. Raiborn, *Cost Accounting - Foundations and Evolutionse*. 9th ed., 2012.

[11] IDC and HPC User Forum, "IDC Economic Models Linking HPC and ROI," 2018.

[12] S. Dall'erba and Z. Chen, "The impact of Blue Waters on the economy of Illinois," tech. rep., 2017.

[13] National Center for Supercomputing Applications, "NCSAs Blue Waters project provides $1.08 billion direct return to Illinois economy," 2017.

[14] IMPLAN, "Economic Impact Analysis for Planning," 2018.

[15] Hyperion Research, "High Performance Computing (HPC) Research," 2017.

[16] C. A. Stewart, B. Plale, V. Welch, M. R. Link, T. Miller, E. A. Wernert, M. J. Boyles, B. Fulton, D. Y. Hancock, R. Henschel, S. A. Michael, M. Pierce, R. J. Ping, T. Gniady, G. C. Fox, and G. Miksik, "Pervasive Technology Institute Annual Report: Research Innovations and Advanced Cyberinfrastructure Services in Support of IU Strategic Goals During FY 2015," tech. rep., Indiana University, Bloomington, 2015.

[17] C. A. Stewart, B. Plale, V. Welch, G. C. Fox, M. R. Link, T. Miller, E. A. Wernert, M. J. Boyles, B. Fulton, D. Y. Hancock, R. Henschel, S. A. Michael, M. Pierce, R. J. Ping, T. Gniady, G. Miksik, and W. Snapp-Childs, "Pervasive Technology Institute Annual Report: Research Innovations and Advanced Cyberinfrastructure Services in Support of IU Strategic Goals During FY 2016," tech. rep., Indiana University, Bloomington, 8 2016.

[18] C. Stewart, B. Plale, V. Welch, M. Pierce, G. C. Fox, T. G. Doak, D. Y. Hancock, R. Henschel, M. R. Link, T. Miller, E. Wernert, M. J. Boyles, B. Fulton, L. M. Weakley, R. Ping, T. Gniady, and W. Snapp-Childs, "Pervasive Technology Institute Annual Report: Research Innovations and Advanced Cyberinfrastructure Services in Support of IU Strategic Goals During FY 2017," tech. rep., Indiana University, Bloomington, 7 2017.

[19] A. Apon, S. Ahalt, V. Dantuluri, C. Gurdgiev, M. Limayem, L. Ngo, and M. Stealey, "High Performance Computing Instrumentation and Research Productivity in U.S. Universities," *Journal of Information Technology Impact*, vol. 10, pp. 87–98, 9 2010.

[20] R. L. DeLeon, T. R. Furlani, S. M. Gallo, J. P. White, M. D. Jones, A. Patra, M. Innus, T. Yearke, J. T. Palmer, J. M. Sperhac, R. Rathsam, N. Simakov, G. von Laszewski, and F. Wang, "TAS View of XSEDE Users and Usage," in *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*, XSEDE '15, (New York, NY, USA), pp. 21:1–21:8, ACM, 2015.

[21] G. von Laszewski, F. Wang, G. C. Fox, D. L. Hart, T. R. Furlani, R. L. DeLeon, and S. M. Gallo, "Peer Comparison of XSEDE and NCAR Publication Data," in *2015 IEEE International Conference on Cluster Computing*, pp. 531–532, 9 2015.

[22] A. Apon, C. Gurdgiev, S. Ahalt, M. Limayem, and M. Stealey, "High Performance Computing Instrumentation and Research Productivity in U.S. Universities," *Journal of Information Technology Impact*, vol. 10, no. 2, pp. 87–98, 2010.

[23] nobelprize.org, "The Nobel Prize in Chemistry 2013," 2013.

[24] nobelprize.org, "The Nobel Prize in Physics 2013," 2013.

[25] nobelprize.org, "The Nobel Prize in Physics 2017," 2017.

[26] National Science Foundation, "NSF Solicitation 14-536," 2014.

[27] N. Merchant, E. Lyons, S. Goff, M. Vaughn, D. Ware, D. Micklos, and P. Antin, "The iPlant Collaborative: Cyberinfrastructure for Enabling Data to Discovery for the Life Sciences," *PLOS Biology*, vol. 14, p. e1002342, 1 2016.

[28] S. Goff, R. Jorgensen, D. Stanzione, G. Andrews, V. Chandler, S. Ram, and L. Stein, "PSCIC Full Proposal: The iPlant Collaborative: A Cyberinfrastructure-Centered Community for a New Plant Biology," 2008.

[29] Pryor Learning Solutions, "Microsoft® Excel® 2007/2010 Basics," 2018.

[30] Udemy, "Online Courses - Learn Anything, On Your Schedule," 2018.

[31] C. A. Stewart, J. A. Wernert, N. Wilkins-Diehr, and K. Gaither, "XSEDE Return on Investment Data and Analysis, 2015-2017," tech. rep., Indiana University, Bloomington Indiana, 2018.

[32] Indiana University, "Action 5: Philosophy of Abundance : Empowering People : Indiana University," 2009.

[33] E. Afgan, D. Baker, B. Batut, M. vandenBeek, D. Bouvier, M. Čech, J. Chilton, D. Clements, N. Coraor, B. A. Grüning, A. Guerler, J. Hillman-Jackson, S. Hiltemann, V. Jalili, H. Rasche, N. Soranzo, J. Goecks, J. Taylor, A. Nekrutenko, and D. Blankenberg, "The Galaxy platform for accessible, reproducible and collaborative biomedical analyses: 2018 update," *Nucleic Acids Research*, vol. 46, pp. W537–W544, 7 2018.

[34] I. Foster, "Globus Online: Accelerating and democratizing science through cloud-based services," *IEEE Internet Computing*, no. May/June, pp. 70–73, 2011.

[35] XSEDE, "XSEDE Ecosystem," 2018.

[36] N. Wilkins-Diehr, S. Sanielevici, J. Alameda, J. Cazes, L. Crosby, M. Pierce, and R. Roskies, "An Overview of the XSEDE Extended Collaborative Support Program," in *High Performance Computer Applications 6th International Conference*, (Mexico City, Mexico), pp. 3–13, Springer International Publishing, 2016.

[37] HT, "What makes something a commodity? - The Economist explains," *The Economist*, 1 2017.

[38] S. Timm and G. Cooper, "Experience in using commercial clouds in CMS," tech. rep., Fermi National Accelerator Laboratory, Batavia, IL, 2017.