

# Assessment of financial returns on investments in cyberinfrastructure facilities: A survey of current methods

Craig A. Stewart  
stewart@iu.edu  
orcid.org/0000-0003-2423-9019  
Pervasive Technology Institute  
Indiana University, Bloomington

David Y. Hancock  
dyhancoc@iu.edu  
orcid.org/0000-0001-8082-8980  
Pervasive Technology Institute  
Indiana University, Bloomington

Julie Wernert  
jwernert@iu.edu  
orcid.org/0000-0002-5705-9527  
Pervasive Technology Institute  
Indiana University, Bloomington

Thomas Furlani  
thomas.furlani@roswellpark.org  
orcid.org/0000-0002-4683-0814  
Roswell Park Comprehensive Cancer  
Center

David Lifka  
lifka@cornell.edu  
orcid.org/0000-0003-0069-6530  
Cornell University

Alan Sill  
alan.sill@ttu.edu  
orcid.org/0000-0003-2527-764X  
High Performance Computing Center  
Texas Tech University

Nicholas Berente  
nberente@nd.edu  
orcid.org/0000-0002-1403-4696  
Mendoza College of Business  
University of Notre Dame

Donald F. McMullen  
mcmullendf@gmail.com  
orcid.org/0000-0001-9996-1888  
Texas A&M University

Thomas Cheatham  
tec3@utah.edu  
orcid.org/0000-0003-0298-3904  
University of Utah

Amy Apon  
aapon@clemsun.edu  
orcid.org/0000-0001-5617-5334  
Computer Science Division  
Clemson University

Ron Payne  
rpayne@illinois.edu  
orcid.org/0000-0001-7834-0951  
University of Illinois,  
Urbana-Champaign

Shawn D. Slavin  
slavin@iu.edu  
orcid.org/0000-0001-6602-4310  
Pervasive Technology Institute  
Indiana University, Bloomington

## ABSTRACT

In recent years, considerable attention has been given to assessing the value of investments in cyberinfrastructure (CI). This paper includes a survey of current methods for the assessment of financial returns on investment (ROI) in CI. Applying the financial concept of ROI proves challenging with regard to a service that, in most academic environments, does not generate a “sold amount” such as one would find in the buying and selling of stocks. The paper concludes with a discussion of future research directions and challenges in the assessment of financial ROI in CI. This work is intended less as a definitive guide than as a starting point for further exploration in the assessment of CI’s value for scientific research.

## CCS CONCEPTS

• **General and reference** → **Metrics**; • **Social and professional topics** → **Government technology policy**; • **Computer systems organization** → **Architectures**.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).  
PEARC '19, July 28-August 1, 2019, Chicago, IL, USA  
© 2019 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-7227-5/19/07.  
<https://doi.org/10.1145/3332186.3332228>

## KEYWORDS

ROI; Total Cost of Ownership; TCO; cost avoidance

### ACM Reference Format:


Craig A. Stewart, David Y. Hancock, Julie Wernert, Thomas Furlani, David Lifka, Alan Sill, Nicholas Berente, Donald F. McMullen, Thomas Cheatham, Amy Apon, Ron Payne, and Shawn D. Slavin. 2019. Assessment of financial returns on investments in cyberinfrastructure facilities: A survey of current methods. In *Practice and Experience in Advanced Research Computing (PEARC '19)*, July 28-August 1, 2019, Chicago, IL, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3332186.3332228>

## 1 INTRODUCTION

Across the world, increasingly tight national government budgets prompt increasingly stringent requirements for scientific research receiving federal funding. Similar concerns persist in private sector research and development organizations, which demand profitability and viability. For higher education institutions in the United States, this issue is particularly acute due to a combination of increased competition for federal research grants coupled with other financial pressures driven by, among other things, projected decreases in enrollment [14, 16]. When the acquisition of new, locally operated cyberinfrastructure (CI) resources is proposed, or when annual budgets are set, *what do we get for this money?* is a particularly challenging and useful question (with CI defined as “... computing systems, data storage systems, advanced instruments and data repositories, visualization environments, and people, all linked together by software and high performance networks to

improve research productivity and enable breakthroughs not otherwise possible” [37]). This question is challenging for a number of reasons, including the difficulty of properly assessing the many hidden costs associated with investments in CI. Further, unlike other infrastructure investments, the relationship between investment in CI and returns on that investment is both indirect and far from immediate. The value of the question posed above is that it leads all involved in CI system acquisition and deployment to think carefully about these decisions. Taking these matters seriously gives stakeholders confidence that investments have been made prudently.

The basic question *what do we get for this money?* is generally posed with intuitive, rather than precise notions of what the benefits of an investment might be. Many peer-reviewed papers, panels, and technical reports use or describe a variety of approaches to assessing the value of, and returns on investment in, CI (see Table 1). This paper offers a survey of current methods used to measure financial returns based on direct outcomes of investment in CI. Considerations regarding non-financial impacts resulting from investment in science infrastructure are addressed in detail in a companion paper [33]. This is consistent with a logic model-based view of program evaluation [12]. Logic models, in general, offer a way to formalize an understanding of what one plans to do and what one intends as results. Figure 1 below is adapted from [12].



**Figure 1: Logic Model of Organizational Processes**

This logic model leads us to create a clear separation between *what we are doing* and *how we are doing it*. The usual objectives for investment in CI include enabling outcomes from innovation, creativity, discovery, engineering, analyses, and research, along with education and training integral to 21st century workforce development. Indeed, while discoveries and high-profile awards such as Nobel Prizes may have a financial values associated with them, for many research institutions, the primary products of research are not measured first and foremost in financial terms.

Table 1 (based on a similar table by Stewart [36]) enumerates several beneficial impacts of investment in CI, and includes references to previously published studies of ROI and impact of CI investments in higher education and research organizations, generally. These include both financial and non-financial analyses.

Financial ROI analysis offers a way to evaluate and optimize procedures that enable the outcomes enumerated above. Research institutions intend for their investments in CI to be used effectively and efficiently. Considering investments in CI in strictly financial terms, there may be a variety of inputs and outcomes. An institution of higher education might be the primary investor in its own CI, but federal funding agencies may also contribute. An institution, its researchers, and students may benefit from financial ROIs in CI. Some benefits might accrue to individuals and entities outside the institutions making the investments: workers in a high-tech community organized around a major research university (who

benefit financially from higher paying jobs) and local and state governments (that benefit from a more affluent tax base) are among potential secondary recipients of financial benefits from research institutions’ investments in CI.

We begin this paper by describing ways of defining and analyzing outcomes (investments and benefits) in financial terms. We then address some of the challenges and implications of calculating ROIs in academic environments, within which actual purchase of services is not the norm. This report will be of value to anyone interested in facilitating research processes via investment in CI, from practitioners to Chief Financial Officers, Chief Information Officers, Chief Research Officers, and Chief Executive Officers.

## 2 FINANCIAL DEFINITIONS

To be precise about what we mean by financial return on investment, a textbook definition explains ROI as “A ratio that relates income generated . . . to the resources (or asset base) used to produce that income” [24]. There is an alternate definition, which uses the formula (income - cost)/cost. For the formula we are using here, a value greater than 1.0 for ROI or a proxy of ROI indicates that the financial value of the return is greater than the cost of the investment. The concept of ROI was invented in cost accounting to assess profit (or loss) on products created or purchased and then sold, such as investments in the stock market. In many kinds of commercial organizations, ROI is a straightforward calculation based on the revenue increase or cost savings associated with an investment. In the case of CI as a facilitator of accomplishment in research organizations, there is no simple *money returned on money invested* relationship. Furthermore, outcomes and impacts associated with a particular investment may be difficult to trace to the investment itself and may be far removed in time [32]. With regard to CI in a research institution, an intuitive sense of *what we’re getting out of this investment* might persist, but the situation’s particularities do not allow for calculating ROI in the manner described in cost accounting textbooks. Below are several financial terms relevant to analysis of CI that are useful when considering financial returns on investment:

- Total cost of ownership (TCO): The total cost of owning something, from acquisition to disposal.
- Value for money: “... assesses the extent to which the program has obtained the maximum benefit from the outputs and outcomes it has produced within the resources available to it” [13].
- Cost avoidance: This is measured as the cost difference between, or ratio of, doing something one way versus another other way [24].
- Value added: “an activity that increases the worth of the product or services to the customer . . .” This can be measured in financial terms, but is often measured in other terms.
- Impact: generally discussed in terms of “the impact of investment in . . .” Impact can take many forms, but they are typically not financial.
- Cost center: “a department within an organization that does not directly add to profit but still costs the organization money to operate. Cost centers only contribute to a company’s profitability indirectly” [23].

**Table 1: Possible Kinds of returns on investments in cyberinfrastructure**

Area of Benefit	Measure of Benefit	Ways to Measure Benefit	Citations
<i>Financial Returns</i>			
Benefit to end user of CI facilities in research	Financial value of time saved	Cost of the time that would have been spent by end user doing research without use of CI resources	
CI system resources	Value of investment in other CI facilities that would have been made without use of a particular facility	Actual costs, cost avoidance	[18, 40]
Personnel resources	Value of support and consulting from CI resource provider	Evaluate allocated usage as a fraction of total costs for providing support and consulting	[40]
Training	Value of training materials created by organizations operating CI facilities	Perceived value, equivalent cost of commercial training, value of CI skills held by employee entering job market	[25]
Grant income	Monetary income	Measure income attributable to use of the resource	[19, 27]
Products & patents	Monetary income	Allocate part of income attributable to use of the resource	[8]
Economic impact	Regional economic impact as measured by economic models (IMPLAN)	Indirect financial benefits, jobs, & tax income attributable to existence of resource	[5, 27, 34]
<i>Non-financial impacts and benefits</i>			
Research innovations	Number of papers, citations, impact of resulting innovations	Impact Factors, patents awarded, license fees generated that are attributable to use of the resource	[15, 20, 41]
Grants awarded	Number of grants awarded, number and impact of papers produced resulting from such grant awards	Grants awarded to users of the resource that can be attributed to usage of the resource	[15, 19]

- Profit center: “a division of a company that directly adds or is expected to add to the bottom-line profitability of the entire organization” [31].

Organizational subunits that manage and support CI within research organizations are often viewed as *cost centers*. A cost center’s expenses are readily apparent; their benefits are generally less direct. This is as opposed to, say, a university registrar’s office, which is a profit center, as students enrolling in classes provide a large portion of a higher education institution’s income.

Depending on institutional policies, ROI may be defined at various levels of an organization. A researcher, department, or college may invest in CI directly, or the costs for such investments may be provided in whole or in part at the institutional level. Some CI groups are sufficiently autonomous within a higher education institution or other research organization that it makes sense to calculate a measure of returns on investment based on financial

exchanges within the larger organization. This can be viewed as a proxy ROI of the CI group *within* the institution, not factoring in research productivity of users, grants received, etc. It is, however, generally not possible to measure ROI for investments in CI within research organizations in a sense that aligns precisely with the textbook definition of ROI, because such CI groups tend not to function in any sort of economy that involves buying and selling. As a result, we cannot answer the intuitive question *what do we get for this money?* in ways that correspond to textbook definitions of ROI. Still, the term ROI resonates with stakeholders in ways that make use of proxies for ROI valuable. One must be careful to define ROI clearly, whether using its textbook definition or a proxy. Even the authors of this paper do not entirely agree on definitions and usage for the terms listed above. The most widely used approaches attempt to create a ratio assessing the total cost of ownership (the denominator of a calculation) and the total value of the return (the

numerator), and refer to this as a financial ROI or as a proxy for financial ROI. But, as an example, Texas Tech University (TTU) takes a different approach, referring to and calculating what they call a cost recovery ratio (CRR). TTU defines this as fees collected divided by costs to provide services. But at TTU, the benefits of the CI center to the university as a whole outweighs the importance of the CI center's financial costs within the TTU ecosystem.

Other approaches to financial aspects of ROI are described below. In general, there are no commonly accepted standards within the CI community for financial ROI and ROI-like ratios; so long as methodologies and terms are clearly defined, it will be possible to compare the results of different analyses. The following section describes specific approaches that have been used heretofore in the literature to measure TCO, ROI, and proxies for ROI.

### 3 TCO CALCULATIONS

As noted earlier, the challenge of measuring the total cost of any CI facility lies in costs that are hidden and difficult to estimate.

Many commercial cloud vendors provide online estimators for simple metrics such as raw bulk cost per core-hour of computation, and raw cost per terabyte for storage and retrieval of data; more complicated considerations, and additional costs, exist. Locally managed resources (clusters, supercomputers, cloud systems) have clear acquisition and maintenance costs, but others (e.g., cybersecurity) are more difficult to quantify.

To be meaningful in academic contexts, TCO analyses must be sophisticated enough to assess factors such as the time-cost of adopting new technologies to both CI providers and users, and the end-to-end complexity of managing CI workflows. This type of analysis depends upon accurate measurement of TCO, including hidden costs. Relevant costs include:

- Acquisition cost (computational and storage) and depreciation for capital assets.
- Training associated with acquisition (cost to the organization of adopting new technology).
- Maintenance cost.
- Facilities cost (floorspace, on-going facility maintenance).
- Electricity cost (including facility overhead).
- Networking cost (data transfer and connection fees).
- System administration staff (for core system operations as well as for design and implementation of new solutions).
- Software costs (including the time cost of adopting new software, and the time and academic progress lost if software breaks or goes out of date).
- Security costs.
- Monitoring and billing infrastructures.

Some of these costs exist regardless of the service delivery model, though they may appear in different ways. User support cost is one such type of cost. Electrical power may not be billed directly for cloud services, but it is included in the overall service billing. In locally installed systems, a university computing or high performance computing (HPC) center may not have to pay for electricity out of its own budget, but electrical power remains part of the system's cost to the university. Egress fees may not exist for certain on-premises delivery methods, but may be a major factor in cloud or external co-location costs.

In the case of local facilities, estimating the costs of security, policy compliance, and risk presents a significant challenge, but these challenges are also present in the commercial cloud. Major commercial cloud cost components include the cost of an appropriate set of instance sizes, times the actual usage, plus storage costs, as well as the cost of extracting data from a cloud resource. Hidden costs for commercial clouds might include contract and billing management, and the effort required to move or optimize usage to secure the greatest discount. For now, the best course of action is to be clear about what is being calculated. If we enunciate all assumptions used in comparing costs for each environment, then we can understand what was done, and how alternative approaches might be applied.

## 4 RETURNS AND PROXIES FOR AND ROI

Once we have calculated TCO (the denominator), we come to the problem of what to put in the numerator in order to calculate ROI.

### 4.1 Value for money as proxy for ROI

Value for money (VFM), defined earlier, is often the best proxy for ROI available for research and development organizations, particularly in higher education. With this approach, VFM as a proxy for ROI is measured as shown below:

$$\text{ROI proxy} = \frac{\text{cost of running a workflow some other way}}{\text{cost actually experienced}} \quad (1)$$

A value of greater than one in equation 1 indicates that what was done would have cost more if done by an alternative method. XSEDE (the NSF-funded eXtreme Science and Engineering Discovery Environment [10]) has been a leader in ROI analysis for some time [38]. XSEDE's ROI analyses are based on this sort of logic, as in the analysis in [36], which compares investment in local CI resources to use of commercial clouds.

The financial efficacy of investment in local CI resources versus the purchase of services from a commercial cloud provider is widely discussed, at present. Extending the approach described above, ROI for local resources vs. use of commercial cloud services is measured in this way:

$$\frac{\text{cost of purchasing resources from a cloud provider}}{\text{TCO for local system}} \quad (2)$$

These calculations do not account for potential application performance differences between platforms. A local resource may be faster or slower than a cloud resource depending on the specific application. But this approach provides a starting point for considering options, and was the approach taken in assessing ROI on investments in a supercomputer and a university-operated cloud system (the NSF-funded Jetstream system [35]) in [36].

### 4.2 ROI on specific workflows

The works cited above generally show that investing in a local facility is more cost-effective than using commercial clouds; it is also possible to assess the financial efficiency of using cloud resources for specific workflows. Bauerdick et al. [17] performed an experiment comparing the use of a commercial cloud with workflow-specific

optimization to the use of a local facility. Specifically, they compared AWS (Amazon Web Services) spot instances for analysis of CMS experiment data [9] versus the Fermilab HEPCloud [6]. The comparison showed that it is possible to get costs of commercial cloud services close to the actual costs of local facilities. But, after using over 15 million hours (approximately a month of time on a modern super-computer) they concluded, “The steady-state cost [of AWS] came to  $1.4 \pm 12\%$  cents per core-hour, which is not much larger than the estimated  $0.9 \pm 25\%$  cents per core-hour for the Fermilab data center” [17]. Such small differences do not sound like much, but they can accumulate over a project lifetime to a cost difference of hundreds of thousands to millions of dollars. Bauerdick’s work also used an automated algorithm to bid opportunistically on spot-market pricing. Most research CI facilities provide custom, dedicated, “available when needed” resources, not spot-market resources. When considering financial efficiency, the research needs and style of use of research workflows must be taken into account. This study is applicable to low-resource-need, high-throughput computing workflows, and not to MPI workflows more typical for HPC use cases.

Hyperion Research [7] also analyzed financial ROI for specific workflows. Hyperion surveyed research and development organizations, asking about costs and incomes associated with particularly successful uses of HPC. This report is useful for highlighting the value of CI in many areas of research and engineering, but the ROI values it reports (\$600 or more per \$1 invested) should be considered an upper bound on ROI for investments in advanced CI.

### 4.3 ROI for services provided by people

The approach described in equation 1 above has also been used to estimate the financial efficacy of resources that provide consulting services. The eXtreme Science and Engineering Discovery Environment (XSEDE) provides a number of services to the US national research community, including one called Extended Collaborative Consulting Services (ECCS). Researchers receiving help via ECSS are assigned an XSEDE staff allocation to support their research activities (at no cost to the researcher). Staff time is funded by XSEDE’s NSF grant. To assess the value of this service, XSEDE asked researchers receiving such services how many person months it would have taken them to do the work they did with ECSS help, in the absence of such help. A financial ROI can be calculated by taking the ratio of an estimated salary for a postdoctoral fellow working for that estimated number of months, divided by the actual staff cost to XSEDE to provide such help [36, 38]. Similar approaches to comparisons between training provided in person versus online can be calculated by tracking usage of training programs (i.e., hours of training watched), and comparing the cost of purchasing such services from commercial providers to the actual cost of delivering such resources via XSEDE or local training programs. These approaches are by no means flawless, but provide a reasonable way of assessing the value of services provided by humans. Financial ROI for these XSEDE services is well above 1.0 [39].

The XSEDE Campus Champions program [11] and, independently, the Advanced Cyberinfrastructure and Research and Education Facilitator (ACI-REF) [1] project are leaders in a systematic approach to the organization and delivery of services from people.

ACI-REF started as an experiment funded by the NSF for two research computing and data facilitators at six different institutions, and evolved into the Campus Research Computing Consortium (CaRCC) [3]. CaRCC offers consulting with expert facilitators who understand research computing and can connect researchers with the information and computing systems essential to their work. Surveys done by the Center for High Performance Computing at the University of Utah demonstrate that facilitators provide assistance with research computing, data services, and facilitation critical to research groups around the campus. Similar qualitative studies undertaken by outside consultants for Indiana University (IU) have provided comparable results [21].

The cost of provisioning local staff expertise versus hiring outside consultants is still at issue. Calculating the equivalent cost of external services versus the human effort of local experts and research facilitators (such as those at the University of Utah) proves challenging in that, in many cases, the services offered by university staff simply have no commercially available equivalent. Quantifying the importance of people who facilitate research remains an area for further research.

### 4.4 Other approaches to ROI estimation

Prior sections focus on the most common approach to ROI assessment: some proxy for a measure of return divided by an actual total cost of owning a system or running a workflow. Another significant and well-grounded approach addresses ROI for CI investments in a regional sense, that is, beyond the benefits to the institution making the investment. IMPLAN [5] estimates the number of jobs created in a given area (e.g., region or state) as the result of investments in that area, such as investments in advanced CI facilities. The NCSA Blue Waters project published a report using regional multipliers to demonstrate its economic value to the State of Illinois and the Midwest [27]. Multiplier-based methods offer useful explanations of the economic value of investment in CI (and research in general) to the general public, including the taxpayers who are ultimately the source of investments in government-funded research.

## 5 OPEN XDMOD AS AN ANALYSIS TOOL

Open XDMoD (Metrics on Demand) is a tool that aids analyses related to financial ROI on investments in CI, and offers a number of other capabilities for optimizing use of HPC systems. Open XDMoD was developed at the University at Buffalo, with contributions from the IU Pervasive Technology Institute (PTI). Development is now led by the Roswell Park Comprehensive Cancer Center. Open XDMoD has the following key capabilities relative to assessment of financial aspects of CI [28]:

- Providing operational staff with the ability to monitor, diagnose, and tune system performance, and to measure the performance of all applications running on their system(s)
- Providing software developers with the ability to obtain detailed analyses of application performance, helping optimize code performance
- Providing metrics analyzing grants and contracts to a research institution, enabling comparisons of success by researchers who use advanced CI facilities relative to those who do not

When software application performance and workflow efficiency on a CI system is improved, more computing can be done per dollar of investment in hardware. Put differently, improving software and workflow efficiency can have the same impact as buying more hardware. Taking two versions of software and workflows (one of which runs faster than the other), one can calculate a proxy ROI on software improvements by taking the same sort of *cost avoidance* approached identified earlier.

$$\text{hardware costs avoided} = \left[ \begin{array}{c} \text{cost of hardware} \\ \text{that would have been} \\ \text{required to run workloads} \\ \text{with unimproved software} \end{array} \right] - \left[ \begin{array}{c} \text{cost of hardware} \\ \text{that was used} \\ \text{to run workloads} \\ \text{with improved software} \end{array} \right] \quad (3)$$

And as with equation 1, a ratio of VFM in the form of hardware costs avoided to cost, in this case the cost of software improvements, can be used as a proxy for ROI:

$$\text{ROI proxy} = \frac{\text{hardware costs avoided}}{\text{cost of the software improvements}} \quad (4)$$

In the case of open source software, the benefits that accrue across a community can be dramatic. For example, Henschel, et al. [22] optimized the widely used Trinity RNA sequence assembly software. Their optimizations decreased the time required to complete an analysis by roughly a factor of 8. That is, the investment in code optimization reduces the necessary hardware to 1/8th of what was formerly required.

Another type of financial ROI involves actual cash income to research institutions. Competing for grant dollars constitutes an important part of the economics of most research organizations, prompting the question *how much does investment in CI contribute to success in competing for grants and contracts?* The Open XDMoD Value Analytics (VA) [30] allows analysis of grant income to an organization as it relates to use (or not) of CI services at that institution. That is, one can discover how many grants, and how many grant award dollars, were received by an institution overall, and of those funds, how much went to users of advanced CI systems. This module links income in the form of grant and contract awards to use of CI systems. When local policy allows, the OpenXDMoD VA module can draw financial data directly from an institution's financial systems, or it can use publicly available data from the NSF and NIH.

Figure 2 shows total grant income to IU per fiscal year from 2013 to 2018 for users of IU's advanced CI systems. This income is divided into the two main categories recognized by funding agencies: direct costs, and facilities and administration (F&A) costs. Direct costs are those budgeted to perform a particular project. F&A monies are negotiated between US research institutions and the federal government, and reflect the cost of the institution's investment in research infrastructure generally (buildings, labs, instruments, etc.). Part of this amount reflects the minimal investment required to perform certain tasks, and part represents institutional decisions to invest in particular areas in order to enhance grant competitiveness. There is, of course, no guarantee that an institution will 'make back' its investment any more than a professional sports team's investment in a star player guarantees it will win. One can look at the costs of facilities investments and measure them against grant wins in the form of F&A income. In the case of IU, the total F&A income associated with extramural grants and contracts where the

PI and/or the PI's research team are active HPC users is much larger than the total cost of operating the systems. This creates a basis for the belief that investments in CI may be sound from the narrow view of just the university's finances.




Figure 2: Grant Income Associated with IU's HPC Users

## 6 CI FACILITIES WITH A (PROXY) ROI OF ~1

Core facilities within Cornell University, by university policy, operate within an internal-to-Cornell economy of buying and selling services. If one accepts this as analogous to any other open market, then one can calculate a proxy of ROI for such core facilities. The Center for Advanced Computing (CAC) has operated within Cornell as a core facility since 2007 [26]. It receives nominal base funding, but is expected to generate enough income from within the university economy that it operates in a nearly cost-neutral manner. That is, CAC has a *money in, money out* relationship within Cornell's internal economy, and by university policy, the cost recovered / cost expended ratio, which can be considered a form of proxy for ROI, is required to be approximately 1.

CAC has, like many other CI centers within universities, a primary mission of enabling the success of its research community. Cornell's core facilities must provide high-value, leading-edge computing resources with professional support; otherwise, researchers will seek better value from external providers, putting the core facility at risk. CAC must recover 75-80% of its operating costs for staff and infrastructure. Power, space, and cooling are covered by F&A income. CAC maintains Cornell's cloud and storage infrastructure, as well as 26 private faculty and department-owned computing clusters [4]. CAC also competes for grants that will benefit the Cornell research community, including XSEDE, Jetstream, and Aristotle (which CAC leads) [2].

CAC has met its annual cost recovery goals for the past 12 years while gradually increasing the size of its staff based on demand. Each year, CAC submits its core facility rates to Financial Affairs for approval; core facilities are not permitted to make a profit. The Vice Provost for Research provides a subsidy to ensure competitive rates. CAC's operational model provides clarity to the university in value returned for money. Its budget is mostly determined by what other members of the community are willing to pay for, plus what CAC is able to secure in external funding. The Minnesota Supercomputing Institute employs a similar model [18], and with CAC, the two function as important examples, showing that advanced computing delivery and support organizations can operate on the basis of

selling services within strict fiscal boundaries that require raising the cost of operations from the user community.

## 7 DISCUSSION

There is much work yet to be done in assessing the financial cost and value of CI investments. For example, assessing the cost of cybersecurity, policy compliance, and risk remains challenging. Such efforts are often provided organization-wide, and pro-rating costs may be difficult. The real cost of security for local on-premises CI systems is never zero, since local cybersecurity staff must be present in order to secure the institution's information technology (IT) environment. Security and policy compliance in the cloud are also not free. In fact, policy compliance proves particularly challenging in cloud environments, as evidenced by fines assessed for improperly uploading protected health information to commercial cloud environments [29]. Another open question involves determining how to assign a cost to risk, though examples from the fields of risk management and insurance will likely be useful in this regard. Measuring the cost of change is similarly problematic, as the cost of change may involve data egress charges, the cost of retraining users, and the cost to users of temporarily decreased productivity.

A significant challenge in these sorts of analyses involves the difficulty of creating counterfactuals, or, reliable inferences about what would happen under circumstances different from those that actually occur. Such a counterfactual could involve the following: *if a particular researcher had used a cloud computing system instead of a local system, it would have cost so-and-so many dollars.* This can be a reasonable basis for comparisons as regards computational cycles. Using computational cycles requires researcher effort, so researchers are generally careful about their use of computational time as a side effect of being careful about use of their own time. Use of storage resources likely runs in the opposite direction, as careful use of storage resources may take more of researchers' time than less careful use. Thus, researchers may not manage resources as efficiently when there is no direct charge for them. A comparison of what actual use of local storage resources would cost in the cloud may then conflate costs and behavioral issues. In a circumstance involving real and obvious charges for storage, a researcher might be more vigilant about storage use than otherwise.

Even when attending strictly to financial concerns, matters other than cost effectiveness are relevant; expenditure categories or the way an expense is experienced may be equally as important. For instance, commercial cloud resources are billed as operational costs (often monthly); some institutions may find it impossible to take on such a burden. The only practicable approach to funding CI resources at some higher education institutions is to accumulate cash so that the purchase of a local CI resource appears in budgetary terms as a one-time capital expense. In this case, *how you can buy it* becomes more important than the amount of resources per dollar.

As mentioned in Stewart et al. [36], statements extolling the virtues of commercial cloud computing over those of local CI resources are sometimes too naive or general to be helpful. The results of most analyses to date show that when a local CI facility can be kept very busy most of the time, local facilities are more cost effective than commercial cloud facilities. However, factors other than cost effectiveness might provide a compelling basis for using one

type of CI resource over another. As Apon et al. point out [15], and as indicated in almost any advertisement for cloud computing resources, commercial cloud resources may expand the scope of an analysis that can be completed within a certain amount of wall clock time; in some cases, that speed may be more important than financial effectiveness.

## 8 CONCLUSIONS

This paper has discussed tools useful in understanding the financial value of investments in CI. We focused on use of a ratio of (Value for Money) : (Total Cost of Ownership) as a proxy for ROI as defined in cost accounting textbooks. ROI-like ratios as described here can be valuable in informing acquisition decisions and strategies for enabling innovation and research. One open problem is the issue of defining a *space* within which returns on investment are measured. Cornell University and the Minnesota Supercomputing Institute offer examples in which the returns on investment established within the internal marketplace of the university are a primary focus. This might be considered an endogenous approach to returns on investment. Indiana University and Texas Tech University, among many other potential examples, show more emphasis on returns as related to financial returns that come from outside the university. This might be considered an exogenous approach to returns on investment. How to integrate and balance these measures is something that has not yet been considered conclusively so far as this group of authors know. Additionally, further research is necessary on the impact of individual CI facilities and aggregate national investment in CI on national economies and nations' global economic competitiveness.

The approaches and example data cited herein suggest a direct tie between CI investments and financial income, challenging the traditional view that groups managing, delivering, and supporting CI facilities should be considered *cost centers*. Indeed, some publications suggest that CI facilities may be so directly tied to grant award income that they are more accurately considered *profit centers*. *Monetize everything* may be an intellectually unsatisfying approach to measuring the benefits of investments in CI, and the monetization of CI services may, in some contexts, distort our understanding of their value. Some benefits, like the development of novel capabilities for creative expression and increased understanding of the world around us, defy monetization. Indeed, the logic model of expected outcomes resulting from investment in CI makes clear that cost effectiveness is generally a measure of how effectively such investments are made, which is a different matter than how effectively such investments have facilitated desired outcomes. However, as higher education institutions face increasing fiscal pressure, it may prove useful to be able to assign as full a financial value as possible to CI investments, which makes it possible to use purely financial terms during related fiscal discussions. What remains most important is the successful, fiscally responsible facilitation of academic innovation and discovery.

## REFERENCES

- [1] [n. d.]. Advanced CyberInfrastructure - Research and Education Facilitators. <https://aciref.org/>
- [2] [n. d.]. Aristotle Cloud Federation: Federated Cloud Model. <http://federatedcloud.org>
- [3] [n. d.]. Campus Research Computing Consortium. <https://carcc.org>

- [4] [n. d.]. Cornell University Center for Advanced Computing. <https://www.cac.cornell.edu/>
- [5] [n. d.]. Economic Impact Analysis for Planning. <http://www.implan.com/>
- [6] [n. d.]. Fermilab HEPCloud. <https://hepcloud.fnal.gov/>
- [7] [n. d.]. High Performance Computing (HPC) Research. <https://hyperionresearch.com/>
- [8] [n. d.]. ROI with HPC. <https://www.hpcuserforum.com/ROI/>
- [9] [n. d.]. The CMS Experiment. <https://cms.cern/>
- [10] [n. d.]. XSEDE. <https://www.xsede.org/>
- [11] [n. d.]. XSEDE Campus Champions. <https://www.xsede.org/web/site/community-engagement/campus-champions>
- [12] 2004. W.K. Kellogg Foundation Logic Model Development Guide.
- [13] 2007. *Sourcebook for Evaluating Global and Regional Partnership Programs: Indicative Principles and Standards*. Washington, D.C. <http://www.oecd.org/development/evaluation/dcdndep/37981082.pdf>
- [14] 2018. Colleges With the Greatest Percentage Decreases in Full-Time Undergraduates. *Chronicle of Higher Education* 65, 15 (12 2018). <https://www.chronicle.com/article/Colleges-With-the-Greatest/245285>
- [15] Amy Apon, Stanley Ahalt, Vijay Dantuluri, Constantin Gurdgiev, Moez Limayem, Linh Ngo, and Michael Stealey. 2010. High Performance Computing Instrumentation and Research Productivity in U.S. Universities. *Journal of Information Technology Impact* 10, 2 (9 2010), 87–98. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1679248](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1679248)
- [16] Jill Barshay. 2018. College students predicted to fall by more than 15% after the year 2025. *The Hechinger Report* (9 2018). <https://hechingerreport.org/college-students-predicted-to-fall-by-more-than-15-after-the-year-2025/>
- [17] L. Bauerdick, B. Bockelman, D. Dykstra, S. Fuess, G. Garzoglio, O. Gutsche, B. Holzman, D. Hufnagel, Hanjin Kim, R. Kennedy, D.I. Mason, P. Spentzouris, S. Timm, A. Tiradani, and E. Vaandering. 2017. Experience in using commercial clouds in CMS. *Journal of Physics Conference Series* 898, 5 (2017). <https://doi.org/10.1088/1742-6596/898/5/052019>
- [18] Evan F. Bollig and James C. Wilgenbusch. 2018. From Bare Metal to Virtual. In *PEARC '18 Proceedings of the Practice and Experience on Advanced Research Computing*. ACM, New York, NY, USA, 13:1–13:8. <https://doi.org/10.1145/3219104.3219164>
- [19] Sandy Dall'erba and Zhangliang Chen. 2017. *The Impact of Blue Waters on the Economy of Illinois*. Technical Report. 4 pages. [http://www.ncsa.illinois.edu/assets/pdf/about/bw\\_economic\\_impact.pdf](http://www.ncsa.illinois.edu/assets/pdf/about/bw_economic_impact.pdf)
- [20] Robert L DeLeon, Thomas R Furlani, Steven M Gallo, Joseph P White, Matthew D Jones, Abani Patra, Martins Innus, Thomas Yearke, Jeffrey T Palmer, Jeanette M Sperhac, Ryan Rathsam, Nikolay Simakov, Gregor von Laszewski, and Fungang Wang. 2015. TAS View of XSEDE Users and Usage. In *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure (XSEDE '15)*. ACM, New York, NY, USA, 21:1–21:8. <https://doi.org/10.1145/2792745.2792766>
- [21] Lizanne DeStefano and Lorna Rivera. 2015. *Cyberinfrastructure Value Assessment Report*. Technical Report. <http://hdl.handle.net/2022/20568>
- [22] Robert Henschel, Phillip M. Nista, Matthias Lieber, Brian J. Haas, Le-Shin Wu, and Richard D. LeDuc. 2012. Trinity RNA-Seq assembler performance optimization. (2012), 1. <https://doi.org/10.1145/2335755.2335842>
- [23] Will Kenton. 2018. Investopedia: Cost Center. <https://www.investopedia.com/terms/c/cost-center.asp>
- [24] Michael R. Kinney and Cecily A. Raiborn. 2012. *Cost Accounting - Foundations and Evolution* (9th ed.).
- [25] D.L. Kirkpatrick. 1998. *Evaluating Training Programs: The Four Levels* (2nd ed.). BerrettKoehler Publishers, San Francisco, CA.
- [26] D. Lifka, R. Reynolds, Susan Mehringer, and P. Redfern. [n. d.]. Overview of the Cornell University Center for Advanced Computing Sustainable Funding Model. In *NSF Workshop on Sustainable Funding and Business Models for High Performance Computing Centers*. Ithaca, NY.
- [27] National Center for Supercomputing Applications. 2017. NCSA's Blue Waters project provides \$1.08 billion direct return to Illinois' economy. [http://www.ncsa.illinois.edu/news/story/ncsas\\_blue\\_waters\\_project\\_provides\\_1.08\\_billion\\_direct\\_return\\_to\\_illinois\\_e](http://www.ncsa.illinois.edu/news/story/ncsas_blue_waters_project_provides_1.08_billion_direct_return_to_illinois_e)
- [28] Jeffrey T. Palmer, M. Gallo, Steven, Thomas R. Furlani, Matthew D. Jones, Robert L. DeLeon, Joseph P. White, Nikolay Simakov, Abani K. Patra, Jeanette Sperhac, Thomas Yearke, Ryan Rathsam, Martins Innus, Cynthia Cornelius, James Browne, William L. Barth, and T. Evans, Richard. 2015. Open XDMoD: A Tool for the Comprehensive Management of High Performance Computing Resources. *Computing in Science and Engineering* 17, 4 (2015), 52–62.
- [29] HHS Office of Civil Rights. 2016. Widespread HIPAA vulnerabilities result in \$2.7 million settlement with Oregon Health & Science University [Press release]. <https://wayback.archive-it.org/3926/20170128024121/https://www.hhs.gov/about/news/2016/07/18/widespread-hipaa-vulnerabilities-result-in-settlement-with-oregon-health-science-university.html>
- [30] Olga Scrivener, Gangadeep Singh, Sara E. Bouchard, Scott C. Hutcheson, Ben Fulton, Matthew R. Link, and Katy Borner. 2018. XD Metrics on Demand Value Analytics: Visualizing the Impact of Internal Information Technology Investments on External Funding, Publications, and Collaboration Networks. *Frontiers in Research Metrics and Analytics* (2018). <https://doi.org/10.3389/frma.2017.00010>
- [31] Troy Segal. 2019. Investopedia: Profit Center. <https://www.investopedia.com/terms/p/profitcentre.asp>
- [32] P.E. Stephan. 2012. *How Economics Shapes Science, Vol. 1*. Harvard University Press, Cambridge, MA.
- [33] Craig Stewart, David Hancock, Amy Apon, Alan Sill, D.F. McMullen, Thomas Furlani, Julie Wernert, Nicholas Berente, Ron Payne, David Lifka, and Thomas Cheatham. 2019. Assessment of return on investment in cyberinfrastructure facilities: A survey of current methods for assessing non-financial impact. In *Workshop Proceedings: Humans in the Loop: Enabling and Facilitating Research on Cloud Computing*. PEARC19. Chicago, IL.
- [34] Craig Stewart, Beth Pale, Von Welch, Marlon Pierce, Geoffrey C. Fox, Thomas G. Doak, David Y. Hancock, Robert Henschel, Matthew R. Link, Therese Miller, Eric Wernert, Michael J. Boyles, Ben Fulton, Le Mai Weakley, Robert Ping, Tassie Gniady, and Winona Snapp-Childs. 2017. *Pervasive Technology Institute Annual Report: Research Innovations and Advanced Cyberinfrastructure Services in Support of IU Strategic Goals During FY 2017*. Technical Report. Indiana University, Bloomington. <https://scholarworks.iu.edu/dspace/handle/2022/21809>
- [35] Craig 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. 2015. 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, 29:1–29:8. <https://doi.org/10.1145/2792745.2792774>
- [36] Craig A. Stewart, David Y. Hancock, Julie Wernert, Matthew R. Link, Nancy Wilkins-Diehr, Therese Miller, Kelly Gaither, and Winona Snapp-Childs. 2018. 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). In *2018 IEEE/ACM 11th International Conference on Utility and Cloud Computing (UCC)*. IEEE, 223–236. <https://doi.org/10.1109/UCC.2018.00031>
- [37] Craig A. Stewart, R. Knepper, M.R. Link, M. Pierce, E.A. Wernert, and N. Wilkins-Diehr. 2017. Cyberinfrastructure, Cloud Computing, Science Gateways, Visualization, and Cyberinfrastructure Ease of Use. In *Encyclopedia of Information Science and Technology, Fourth Edition* (fourth ed.), M. Khosrow-Pour (Ed.). IGI Global, Hershey, PA, 1063–1074.
- [38] Craig A Stewart, Ralph Roskies, Richard Knepper, Richard L Moore, Justin Whitt, and Timothy M Cockerill. 2015. XSEDE Value Added, Cost Avoidance, and Return on Investment. In *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure (XSEDE '15)*. ACM, New York, NY, USA, 23:1–23:8. <https://doi.org/10.1145/2792745.2792768>
- [39] Craig A. Stewart, Julie A. Wernert, Nancy Wilkins-ÂrpDiehr, and Kelly Gaither. 2018. *XSEDE Return on Investment Data and Analysis, 2015-2017*. Technical Report. Indiana University, Bloomington Indiana. <http://hdl.handle.net/2022/22491>
- [40] Abhinav S Thota, Ben Fulton, Le Mai Weakley Weakley, Robert Henschel, David Y Hancock, Matt Allen, Jenett Tillotson, Matt Link, and Craig A Stewart. 2016. 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)*. ACM, New York, NY, USA, 61–68. <https://doi.org/10.1145/2974927.2974956>
- [41] G von Laszewski, F Wang, G C Fox, D L Hart, T R Furlani, R L DeLeon, and S M Gallo. 2015. Peer Comparison of XSEDE and NCAR Publication Data. In *2015 IEEE International Conference on Cluster Computing*. 531–532. <https://doi.org/10.1109/CLUSTER.2015.98>

## ACKNOWLEDGMENTS

This report was supported by the Indiana University Pervasive Technology Institute and by NSF Awards 1445604, 1053575, 1548562, 1362134, 1405767, 1243436, 1148996 and 0946726. Thanks to our many collaborators, particularly R. Eigenmann who suggested this line of research. Thanks to Harmony Jankowski for editing, Lauren Huber for graphics, and Monica Shannon for logistical support.

CAS and DYH contributed equally to this report. CAS bears responsibility for any and all mistakes.