

Re-centering the Research Computing Enterprise

The insatiable institutional demands for computing cycles, network bandwidth, and storage clearly demonstrate that IT is a mission-critical function in nearly all areas of higher education. Nowhere is this clearer than in the challenges facing central data centers. Not too long ago, the important issue for the central data center was physical size and floor space. Although the development of smaller, compact systems alleviated this problem by condensing more computers into less space, this only shifted the challenge. The important issue now for the central data center is providing reliable power and cooling for these dense computing systems. The gains in floor space consolidation are being trumped by growth in demand, heat, and power. IT-enabled research is the leading exemplar of this challenge.

As IT leaders struggle to meet relentlessly increasing demands, schools and departments often argue, either implicitly or explicitly, that managing many smaller data centers is better than spending millions of dollars on a central data center. Yet the reverse is actually the case: *The centralization of data centers is needed now more than ever.* It is time to consolidate, not distribute, resources—especially for research computing. To put it as starkly as possible, I argue that the devolved model of research computing, popular over the last decade, is no longer efficient and is grievously wasteful of both human

and financial resources. That is not to say that such resources are not well used—they clearly are, in most cases. Rather, I assert that these resources could be used much more efficiently and effectively to provide a better quality of services supporting an even larger amount of excellent research.

The Cyberinfrastructure Imperative

The National Science Foundation coined the term *cyberinfrastructure* to describe the IT environment necessary to conduct today's scientific research: a combination of high-performance computing systems, massive data storage, visualization systems, and advanced instrumentation, all linked by a high-speed network. Research in nearly all areas of science and scholarly inquiry requires advanced cyberinfrastructure, and an institution's ability to attract and retain research faculty now depends, in part, on its ability to provision and support cyberinfrastructure. How, then, can an institution fund these needs in times of severe budget pressures? The cost of cyberinfrastructure and the growing demand for cyberinfrastructure are clearly outpacing the ability of most institutions to pay for it.

The optimal approach is to prioritize research needs, coordinate grant submissions, and use indirect cost recoveries to assist with funding the growing cyberinfrastructure needs in our institutions. This calls for a coordinated, central plan for supporting IT in research—coordination

that seems all too often to be lacking in higher education. In a presentation based on a forthcoming ECAR research study, Harvey Blustain indicated that the responsibility for most elements of research infrastructure is typically distributed between central and local IT.¹

The Problem Heats Up

Many CIOs worldwide face problems with their data centers. Facilities that are only a few years old were probably designed to support 30–40 watts per square foot, a traditional measurement of power and cooling, and building costs have been roughly \$400 per square foot. However, most new data center designs call for 100 watts per square foot, and there are blade computers being designed that may require up to 500 watts per square foot. In the future, a modest 5,000-square-foot data center could cost \$30 million or more. How can our financially strapped institutions prepare for such possibilities? The best answer is to use a coordinated approach, and nowhere is that more important than with research computing.

Collaboration Is the Key

CIOs should encourage a critical review of their approach to research computing. CIOs need to ask, "Can a coordinated and leveraged approach to centralized research computing provide more and better research computing for all stakeholders?"

The very fact that individual departments see the need to have their own mini-data centers is at least in part indicative of the current underfunded state of central data centers and quality staff. Blustain also noted that central

The devolved model of research computing, popular over the last decade, is no longer efficient and is grievously wasteful of both human and financial resources.

IT budgets for research computing are relatively modest. Those institutions that are considered “research essential” in his study must have state-of-the-art data centers as well as the necessary resources to maintain them. Life-cycle funding needs to be allocated in anticipation of replacement and enhancement of these critical facilities. Through a coordinated, centralized approach, the resources that departments use to run their own systems can be reallocated and pooled to achieve economies of scale and to serve populations incrementally larger than the combined departmental headcount.

structure. The same coordinated planning process should be applied. Central data centers are more cost-effective than satellite centers, and although satellite centers can support disaster planning, they need to be a part of the institution’s coordinated central IT plan and management.

Data Centers for the Future

How are leading-edge institutions designing their data centers with an eye to flexibility and growth? How can institutions look beyond their immediate IT requirements and position themselves to participate in a national community of shared

Many competent electrical and mechanical engineers are not familiar with the complex requirements for the fault-tolerant, never-fail electrical and cooling infrastructure needed for these facilities. They may not have the expertise to accurately assess the cooling and electrical capacity necessary for the next generation of systems, which are completely different from the typical issues on which these engineers usually work. Though doing so will be expensive in up-front costs, CIOs must engage design engineers who understand the present and future issues as they plan for modern data centers.

Ultimately, data center issues form an ideal catalyst for discussions about collaboration and strategic planning in IT within an institution and represent an opportunity to share best practices for data center design among institutions. Although distributing research computing may seem to result in problems that are smaller and easier to solve, the results of that approach will not serve our institutions or the advancement of scientific research in the long run. There are clear economies of scale in provisioning many high-quality institutional services, and cyberinfrastructure for research computing will necessitate leveraged funding models. Institutions that learn how to steer research computing investments through a centralized approach will be better equipped to participate and lead in the future world of cyberinfrastructure. This is no time for “business as usual” in higher education; it is time for collaborative thinking to get the most from institutional investments in IT.

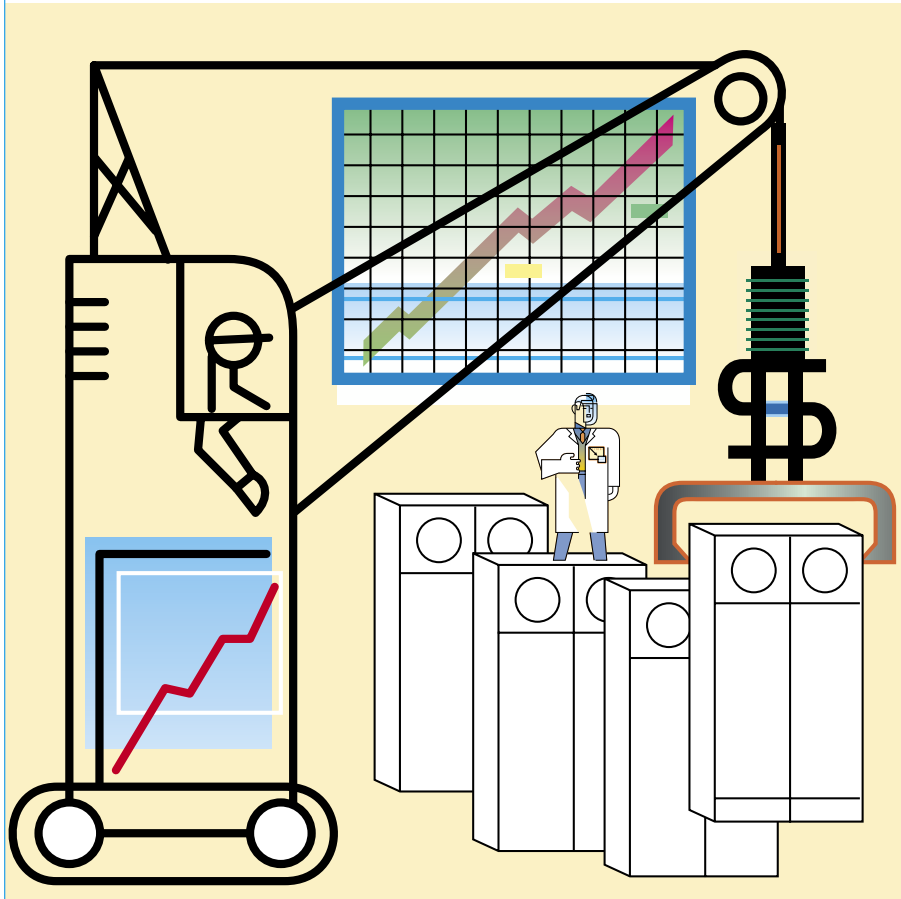


Illustration by Hank Osuna, © 2006

Moreover, centralized maintenance and support services can be delivered at lower overall cost and higher quality of service to the institution as economies of scale are realized and fewer staff are needed.

A similar argument is applicable where uncoordinated development of ERP systems, data warehouses, and learning management and other systems also strains central IT data centers—though their growth seems more manageable relative to the trajectory of research cyberinfra-

IT resources? What planning is being done to meet the mounting challenges of security and the necessity of business continuity and redundancy in the event of disaster? Although these decisions must be made by the senior IT leader, collaboration between the CIO and the top researchers in the institution is essential in planning for cyberinfrastructure.

Data center design has become a highly specialized art that necessitates the services of data center design experts.

Note

1. Harvey Blustain, “The Engagement of Information Technology and Research,” presentation, EDUCAUSE Center for Applied Research (ECAR) Symposium 2005, Carefree, Arizona, December 8–9, 2005, <<http://www.educause.edu/LibraryDetailPage/666?ID=ECR0507>>. Blustain’s presentation was based on preliminary findings from the ECAR research study *IT Engagement in Research*, to be published later this year.

Michael A. McRobbie is Interim Provost and Vice President for Academic Affairs at Indiana University, Bloomington. He was Vice President for IT and CIO from 1997 to 2005 and Vice President for Research from 2003 to 2005. He is a Professor of Computer Science and a Professor of Philosophy.

