


CYBER

I
N
F
R
A
S
T
R
U
C
T
U
R
E




In Tune for the FUTURE

By James R. Bottum, James F. Davis, Peter M. Siegel,
Brad Wheeler, and Diana G. Oblinger

Rapid advances in the speed, power, and ubiquity of computers, computing networks, and related technologies continuously redefine what is possible today. Aiding significantly in that redefinition is *cyberinfrastructure* (CI), also known as *e-research*, *e-science*, and *e-infrastructure*. Cyberinfrastructure connects institutions, researchers, educators, and students with high-performance computing, remote sensors, large data sets, middleware, and sophisticated applications such as visualization tools and virtual environments. Allowing the sharing not only of tools and data but also of expertise, cyberinfrastructure merges technology, data, and human resources into a seamless whole.

James R. Bottum is Vice Provost for Computing & IT and CIO at Clemson University. James F. Davis is Associate Vice Chancellor and CIO at UCLA. Peter M. Siegel is CIO and Vice Provost for Information & Educational Technology at the University of California, Davis. Brad Wheeler is Vice President for Information Technology and CIO at Indiana University. Diana G. Oblinger is President and CEO of EDUCAUSE.

Cyberinfrastructure allows researchers to tackle this mountain of information and to answer questions that could hardly even be asked a decade ago.



The Evolution of Cyberinfrastructure

The idea of cyberinfrastructure and the word itself moved more widely into use after the 2003 publication of the report by the National Science Foundation (NSF) Blue-Ribbon Advisory Panel on Cyberinfrastructure. *Revolutionizing Science and Engineering through Cyberinfrastructure* stated: "The term *infrastructure* has been used since the 1920s to refer collectively to the roads, power grids, telephone systems, bridges, rail lines, and similar public works that are required for an industrial economy to function. Although good infrastructure is often taken for granted and noticed only when it stops functioning, it is among the most complex and expensive thing[s] that society creates. The newer term *cyberinfrastructure* refers to infrastructure based upon distributed computer, information and communication technology. If *infrastructure* is required for an *industrial* economy, then we could say that *cyberinfrastructure* is required for a *knowledge* economy."¹

Since that time, the concept of cyberinfrastructure has expanded beyond the report's focus on science and engineering to include areas such as economics, social sciences, and the arts and humanities.² The evolution has been driven by several overlapping forces:

- **Research challenges.** A wide range of disciplines need more sophisticated, cyberinfrastructure-enabled research approaches.
- **Institutional competitiveness.** Cyberinfrastructure is emerging as a competitive element among institutions, with colleges and universities that possess effective cyberinfrastructure outcompeting others for external research funding as well as for highly sought-after faculty and students.
- **Education and learning.** Cyberinfrastructure can enhance learning by allowing students to learn-by-doing rather than learn-by-listening.

Likewise, although the initial definition of cyberinfrastructure emphasized high-performance computing, the term now encompasses additional tools and applications. Beginning in 2005, for example, NSF created multidisciplinary teams and categorized the major components of cyberinfrastructure into four overlapping and complementary areas:

- High Performance Computing
- Data, Data Analysis, and Visualization
- Cyber Services and Virtual Organizations
- Learning and Workforce Development³

Most recently, in late 2007, NSF announced the Cyber-Enabled Discovery and Innovation (CDI) initiative, focusing on transformative research via *computational thinking*. CDI research will address three thematic areas:

- **From Data to Knowledge:** enhancing human cognition and generating new knowledge from a wealth of heterogeneous digital data
- **Understanding Complexity in Natural, Built, and Social Systems:** deriving fundamental insights on systems comprising multiple interacting elements
- **Building Virtual Organizations:** enhancing discovery and innovation by bringing people and resources together across institutional, geographical, and cultural boundaries⁴

No doubt, the evolution of cyberinfrastructure will continue.

More than the Internet

Cyberinfrastructure differs from traditional web and broadband access in its focus and its magnitude. Consumer broadband allows us to watch movies online, quickly download music files, or use visually engaging and media-rich online learning resources. In contrast, the high-performance computing and networking resources of cyberinfrastructure enable researchers, for ex-

ample, to create fully interactive, three-dimensional models of severe weather systems and to make those models available to other researchers across the country—instantly, at any time. CI resources give scientists and educators opportunities to create and collaborate in entirely new contexts—to *experience* processes and results even if the technologies and data sets are thousands of miles away.

Cyberinfrastructure permits a new kind of scholarly inquiry and education, empowering communities to innovate and to revolutionize what they do, how they do it, and who participates. Data is being collected, archived, and analyzed on a scale previously unimaginable. Cyberinfrastructure allows researchers to tackle this mountain of information and to answer questions that could hardly even be asked a decade ago.

NSF, major research institutions, and colleges and universities of all types have already made significant investments in this infrastructure. Researchers in science, engineering, medicine, the social sciences, and the arts and humanities today are thus able to share resources and data across significant distances at near-instantaneous speeds, permitting an unprecedented level of collaboration. Visualization and virtualization technologies allow them to model and interact with research findings—to "see" and "touch" their objects of focus—in a way that generates fresh insight, that speeds discovery, and that creates new opportunities.

By providing coordinated, dynamic resource sharing and aggregation, cyberinfrastructure can enable institutions, researchers, and/or educators to

- advance the ability of the scholarly community to interact, collaborate, and conduct research;

- support teaching and learning in new interactive and experiential ways;
- build and access capability that cannot be singly supported by individual researchers or even institutions;
- tap the capacity of midrange resources in support of high-end users and researchers doing IT-enabled research;
- create a researcher pipeline that grows and builds experience in IT-enabled research by escalating capability with sophistication of solution;
- address the CI requirements across various disciplines even though the potential for resources and funding also varies;
- identify unused capacity that can be repurposed for use based on agreed-upon policies;
- create a national network of grid-accessible resources by building collaborating grids; and
- generate major new social and economic benefits, such as improved medicines and medical care and greater understanding of and ability to manage changes in the environment, the economy, and social structures.

These benefits extend from research to teaching and learning—and not just at the graduate level. Instant access to remote instrumentation and analysis tools allows undergraduates to gain hands-on experience in research methods and projects and to contribute to research activities in a variety of disciplines. Data-visualization applications and graphically rich virtual environments enable humanities and social sciences students to engage in practical problem-solving in their communities or to experience historical contexts in a direct, personal way.

An excellent example of how cyberinfrastructure can

advance not only research but also teaching and learning at both the undergraduate and the graduate levels is nanoHUB.org (<http://www.nanohub.org>), a gateway for researchers, faculty, and students in nanotechnology.⁵ NanoHUB provides users with fingertip access to more than eighty simulation tools for research and education. In addition to being able to launch jobs that are executed on the state-of-the-art computational facilities of Open Science Grid and TeraGrid, users can interactively visualize and analyze the results—all via an ordinary web browser. The nanoHUB middleware hides the complexity of grid computing, handling authentication, authorization, file transfer, and visualization and letting the researcher focus on research. This approach also helps educators bring these tools to the classroom, letting them bypass the difficulties of grid computing and focus instead on teaching and learning.

With over 58,000 users in early 2008, the site serves undergraduates, graduate students, faculty, and industry researchers. In the twelve months ending April 2008, more than 270,000 simulations have been run, and web hits exceed 35 million. Instructors used nanoHUB in 40 classes during the 2007–8 academic year, and to date there are 270 citations of its tools and resources in the research literature.

Undergraduate students use nanoHUB to learn about nanotechnology, watching lectures or visualizations, completing problem sets, running experiments, and participating in the community. The site allows them to use the same applications that researchers use, introducing them to the tools and techniques of the discipline as early as possible. Graduate students access research papers, lectures, and other educational resources. The tools provided via nanoHUB.org allow them to advance their research and careers. For example, Saumitra Mehrotra, a University of Cincinnati master's student, tried 26 simula-

tion tools in 10 months, spent 52 hours with 134 items of tutorial and seminar content, and did 2,855 simulations of a nanowire model using 8,242 CPU hours, resulting in a paper for a 2007 IEEE workshop. Along the way, Mehrotra improved the tools for other users, whether staff, students, or faculty.

As with many other CI activities, nanoHUB has changed processes as well. So that students and researchers can gauge the value of their work to the community, nanoHUB has built in a tracking mechanism. The usage statistics provided are analogous to those provided in paper citation indices. And the underlying structure of nanoHUB is being used to create other communities; globalHUB.org, pharmaHUB.org, and thermalHUB.org are three recent ones.

From Computing to Collaboration

Cyberinfrastructure is about more than the technology; it involves creating a culture of collaboration, both within and across disciplines. As research has grown increasingly computational and data-driven, collaboration has become essential. Even in business, competitiveness hinges on the aggregation, analysis, and application of data across industries. Massive amounts of data and information are utilized to conduct research and development, validate models, and assess or predict risks. More and more, experts from across states, institutions, and countries must come together to tackle highly complex, large-scale problems such as the environmental impacts of pollution or the identification and development of new energy sources. And their efforts increasingly depend on being able to work together, regardless of location, via high-performance networks and computational resources. These trends leave little doubt that the research enterprise in the twenty-first century will be information-based and will emphasize teams of researchers who can tackle large problems holistically.

Cyberinfrastructure is about more than the technology; it involves creating a culture of collaboration, both within and across disciplines.

The scale, cost, and flexibility of the cyberinfrastructure needed by any one researcher argue against “go it alone” approaches.

In this environment, the scale, cost, and flexibility of the cyberinfrastructure needed by any one researcher argue against “go it alone” approaches. The problem is that in spite of its holistic nature, cyberinfrastructure often still develops “at the edge,” where each academic department or faculty member with a grant develops precisely the IT capabilities needed to support his or her research. This leads to the chemistry department, for example, running its own computing clusters in makeshift machine rooms, with system administration often done by graduate assistants possessing varying levels of IT competence. Faculty can configure the computing, storage, or visualization environment to their precise needs without worrying about disturbing the needs of other faculty.

The core task of CI alignment thus involves enabling effective CI services at the edge while aggregating some services for leverage. A better institutional model is to provision cyberinfrastructure as a leveraged service in which a campus IT group aggregates funding to provide large-scale common systems. The services can be made available to any researcher, faculty, or graduate student, with or without a charge-back mechanism to users. Common systems are open to all professional staff and are life-cycle funded. From this perspective, one may envision cyberinfrastructure as a collection of independently owned and administered (faculty, center, institutional) resources joined together by a shared grid of hardware, software, networking, and support services. This model addresses competing dimensions (e.g., individual researcher vs. research team, specialization vs. scale, grant funding vs. institutional investment) by focusing on an infrastructure that provides coordinated, dynamic resource-sharing that is dependable, consistent, and administered according to policies on which all parties have agreed.⁶

With a focus on joining independently

owned and administered resources and at the same time aggregating like needs, cyberinfrastructure balances a rich spectrum of capability with increased efficiency across a base of standard resources. Vital IT support staff who have specialized expertise with particular types of applications and equipment are sustained as groups distributed among campuses and labs. The management of short resource life-cycles is spread throughout, the probability of state-of-the-art resources being available at any given time is increased, and the risk that facilities will become out-of-date is lessened. Individual researchers or research teams gain access to the best resources for their needs while overall resource utilization is maximized. In addition, a given capability can be made available to a particular research team and a wider range of researchers regardless of location. Through cyberinfrastructure, an institution can coalesce expertise, development capability, facilities, and tools into a coordinated resource that provides capability far beyond what any one research group or even institution could singly produce, while offering a research engine to a much larger base of users.


Moving Forward

To successfully align CI resources, institutions must operate within a fabric of trust. Cyberinfrastructure requires building collaborative, cooperative, and responsible trust relationships between IT providers and users and between research units, faculty members, and their respective institutions. Shared, grid-based research environments must have a governance and decision framework with strong faculty involvement—a framework that fits the institutional culture and generates buy-in to meet users’ needs.

Decisions must be made, individually at each campus and collectively across institutions and government agencies, on where and how to invest in cyberinfrastructure:

- At what level should CI services be provided?
- What is the appropriate campus role and investment in cyberinfrastructure?
- What is the appropriate role at the research group level? in the multi-institutional research communities?
- How can an institution create the right incentives for collaborative behavior?
- In what ways should an institution support its researchers and students in the context of very large data management?
- What is the role of CI planning beyond the research arena?
- How can federal and state attention to the investment needs for cyberinfrastructure at the campus level be increased?

Cyberinfrastructure holds significant potential for higher education, but these and many other questions remain. To help the higher education community answer these questions, EDUCAUSE will be placing a sustained focus on cyberinfrastructure. This issue of *EDUCAUSE Review* is just the beginning of an ongoing exploration of the topic of cyberinfrastructure and its potential. We begin with the following articles, which address cyberinfrastructure as it relates to research and education in science and engineering (Francine Berman) and the liberal arts (David Green and Michael Roy), along with an article reporting on ECAR survey respondents’ opinions regarding the importance of CI technologies to research and to teaching and learning now and in the near future (Mark C. Sheehan). The online version of this issue of the magazine also offers a short background on the EDUCAUSE Campus Cyberinfrastructure (CCI) Working Group and its activities in this area.



The way cyberinfrastructure works or fails can be compared to an electronic signal that is broadly transmitted across a network. The overall signal is made of many waves, which in the case of cyberinfrastructure are determined by a highly distributed spectrum of hardware, software, and support resources, encompassing the conflicting needs of individual vs. team research, autonomy vs. capacity, specific discipline vs. standardization, ownership vs. sharing, specialized vs. scale, grant funding vs. institutional investment, and sustainability vs. short life-cycles. If the individual waves align, there is resonance and a strong, overall signal at the point CI capabilities are needed, but if they do not align, they cancel out each other's effects, and the resonance of capabilities is lost. It is thus imperative for those of us in higher education to tune resources, policy, and investments to find the right balance and achieve the greatest impact for cyberinfrastructure. e

Notes

1. *Revolutionizing Science and Engineering through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure*, January 2003, p. 5, <<http://www.nsf.gov/od/oci/reports/CH1.pdf>>.
2. For example, see the American Council of Learned Societies Commission on Cyberinfrastructure for the Humanities and Social Sciences, *Our Cultural Commonwealth* (2006), <http://www.acls.org/uploadedfiles/publications/programs/our_cultural_commonwealth.pdf>; and "Cyberinfrastructure and the Liberal Arts," special issue of *Academic Commons*, December 2007, <<http://www.academiccommons.org/commons/announcement/table-of-contents>>.
3. National Science Foundation, Cyberinfrastructure Council, *Cyberinfrastructure Vision for 21st Century Discovery* (Washington, D.C.: National Science Foundation, March 2007), <<http://www.nsf.gov/pubs/2007/nsf0728/nsf0728.pdf>>.
4. See the CDI page on the NSF website: <http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503163>.
5. This description of nanoHub is drawn from George B. Adams III, "NanoHUB.org: Future Cyberinfrastructure Serving over 58,000 Users Today," presentation at the CNI Task Force Meeting, Minneapolis, April 8, 2008. See also Carie Windham, "The nanoHUB: Community and Collaboration," *EDUCAUSE Review*, vol. 42, no. 6 (November/December 2007), pp. 144-45, <<http://www.educause.edu/er/erm07/erm07612.asp>>.
6. Brad Wheeler, "Research Technologies: Edge, Leverage, and Trust," in "The Organization of the Organization: CIOs' Views on the Role of Central IT," *EDUCAUSE Review*, vol. 42, no. 6 (November/December 2007), p. 44, <<http://www.educause.edu/er/erm07/erm0761.asp>>.