

University of Utah Campus Cyberinfrastructure Plan
Information Technology (IT) Governance Research Portfolio
March 28, 2013

Document context

The University of Utah Information Technology Research Portfolio, currently chaired by Prof. Thomas Cheatham, is a component of the newly implemented Information Technology (IT) governance structure of the University of Utah. The portfolio has replaced the earlier Campus Cyberinfrastructure (CI) Council that was established in 2007 in response to a report of a campus *ad hoc* committee and operated under the leadership of Prof. Martin Berzins. The Research Portfolio brings together faculty members and cognizant administrators to determine the strategy and project priorities of the University's cyberinfrastructure initiatives. The Research Portfolio then conveys these priorities to the associated campus IT service unit for implementation.

The Research Portfolio has developed this document to outline the high-level strategy and goals of the Cyberinfrastructure at the University of Utah. This document should receive review and revision on an annual basis or more frequently as needed. This document is itself a revision of previous CI strategic plans developed under the Campus CI Council. As needed, the Research Portfolio will coordinate its strategy recommendations and budget and project prioritization recommendations with those of the other IT governance portfolios (Teaching & Learning, Infrastructure, and University Support Services) with the oversight of the campus Operational IT Council (OITC).

The overall cyberinfrastructure strategy of the University of Utah is to provide a rich, responsive set of research and training focused services and capabilities to the faculty and other professionals and students comprising the campus community. This strategy should closely align with national objectives of the National Science Foundation (NSF), the National Institutes of Health (NIH), the Department of Energy (DOE), and other relevant agencies.

Introduction

In recent years, increasing recognition has been devoted to the need for U.S. research universities to develop a balance of coordinated and centralized information technology (IT) capabilities for computationally based and highly collaborative research. At the national level, CI planning and investments have been driven by the growing complexity, the importance of distributed collaboration, and the increasing IT demands in the research disciplines within the physical and social sciences, engineering, and medicine. The National Science Foundation (NSF) recently has expounded an updated vision for the CI model in its Cyberinfrastructure Framework for 21st Century Science and Engineering (CIF21).¹ This vision specifically highlights the need for developing an integrated and scalable national cyberinfrastructure that supports all areas of research.

¹ <http://www.nsf.gov/od/oci/cif21/CIF21Vision2012current.pdf>

In addition to the aforementioned efforts at NSF, other major federal R&D funding agencies including the National Institutes of Health (NIH) and the Office of Science within the Department of Energy (DoE) have placed increasing emphasis on identifying CI needs within their supported disciplines and are making initial investments. Part of the strategic vision for NIH specifically states:

By 2025, we will have left behind a world of very expensive, personally held knowledge in which people with trained intuitions based on years of education and practice could produce acceptable results in either health care or research. A worldwide Internet based cyber-infrastructure of knowledge provided in real time and mediated by expert systems exploring massive databases will be useful tools for healthcare and research. The health sciences will likely share some of the infrastructure with other scientific disciplines as part of the new world of e-science, both within the U.S. and internationally.²

The EDUCAUSE working group on Campus Cyberinfrastructure advanced the following definition of CI, and we feel this remains a good working definition for our purposes.

Cyberinfrastructure consists of computing systems, data storage systems, data repositories and advanced instruments, visualization environments, and people, all linked together by software and advanced networks to improve scholarly productivity and enable breakthroughs not otherwise possible.

At the University of Utah, existing research service centers such as the Center for High Performance Computing (CHPC) and the CCTS Bioinformatics Core and faculty-led institutes such as the Scientific Computing and Imaging (SCI) Institute, the Institute for Clean and Secure Energy (ICSE), the Center for Extreme Data Management Analysis and Visualization (CEDMAV), and the Brain Institute already provide a number of specific CI resources and services to campus researchers. Numerous programs already exist or are in the advanced planning stage to leverage these resources and research service centers. Students in these programs make use of the CI resources, and in many cases, participate in the development and deployment of these resources by interacting in short courses and internships. Examples include:

- Computational Science and Engineering (<http://www.ces.utah.edu/>)
- Data Center Engineering Certificate (new interdisciplinary engineering program among Electrical and Mechanical Engineering and the School of Computing)
- Big Data Professional Masters (a new program to be introduced in the School of Computing)
- Geographic Information Science (<http://geog.utah.edu/graduate/msgisinfo.php>)
- School of Business Information Systems (<http://msis.business.utah.edu/>)
- Biomedical Informatics (<http://medicine.utah.edu/bmi/>)
- Medical Physics Program (<http://medicine.utah.edu/radoncol/medPhysics/index.htm>)

A principal charge of the Research Portfolio is to identify opportunities for greater coordination and improvement in existing CI investments and to develop plans for meeting emerging CI needs through a combination of campus and external funding.

² <http://www.nlm.nih.gov/pubs/plan/lrp06/report/healthcaresystem.html>

Strategic Plan

Mission

The mission of the University's collective effort in Cyberinfrastructure is to develop, provide, and communicate an evolving, customizable set of integrated, broadly defined, high-performance, cost-effective, and sustainable research IT capabilities and services to the University community and its collaborators.

Vision

As many research disciplines either advance or develop increasing computational requirements for data acquisition and analysis, modeling, and simulation, University faculty and other researchers will be able to maximize their daily effectiveness and overall competitiveness by seamlessly utilizing campus and national CI capabilities and services in the course of their research. To support this vision, we will aim to expand, generalize, and clearly define the portfolio of CI service offerings at the University of Utah to meet the emerging CI requirements of an increasingly broad set of University faculty and researchers.

Core Enablers

1. We will facilitate continuous faculty, graduate program and center/institute engagement in establishing the overall CI strategy and the development and customization of new services. This will be accomplished through:
 - a. Research Portfolio leadership, direction, and oversight within the broader IT Governance process,
 - b. Faculty outreach – particularly for the evolution of existing services and the development of new capabilities, and
 - c. Development of regular processes for assessing CI service delivery and faculty satisfaction.
 - d. Dissemination of the CI strategy and services through the Research Portfolio web site, the CHPC web site and newsletter, and the newly expanded capabilities of the University Information Technology (UIT) Strategic Communication group.
2. We will ensure the continuing evolution of computational capabilities to support most high-end computational research efforts on campus.
 - a. We will establish targets for relative system performance and impact metrics *viz a viz* those of national centers.
 - b. We will seek to provide computational platforms meeting the distinct needs of standard users requiring high job throughput (*capacity* clusters) as well as those running highly parallel jobs (*capability* clusters).
 - c. We will continue to push the envelope through use of specialized resources, such as with GPUs or accelerators (*heterogeneous* clusters), e.g. the nVIDIA Center for Excellence in SCI and generally available GPU resources in CHPC, and also through novel HPC hardware and network testbeds (e.g. SCI, Flux Research Group).
 - d. We will seek out and respond in a coordinated fashion to national funding opportunities for CI development as they emerge through NSF CISE ACI, NSF EPSCoR, NIH CTSA, and similar programs.

3. We will establish and enhance collaborative relationships with national HPC and CI centers (e.g., NSF XSEDE, Open Science Grid, NIH CTSA) to facilitate the smooth transition of high-end users and their most demanding applications from campus systems and to facilitate technology transfer back to campus CI efforts. This includes strengthening connectivity on campus, to regional partners, and to the national CI infrastructure.
4. We will enhance CI collaboration with the other research universities and undergraduate institutions within the state of Utah through the following:
 - a. Shared services among the Utah System of Higher Education (USHE) research campuses for distributed computation and storage,
 - b. Shared use of the University's new off-campus Downtown Data Center (with nearly 100 racks and 1.2 MW of electric power for HPC and CI use),
 - c. Expansion of the Research@UEN optical network beyond the core set of universities (University of Utah, Utah State, and BYU),
 - d. Collaboration with both active Utah EPSCoR projects - iUTAH and CI-WATER - on the provisioning of CI services supporting the research initiatives and
 - e. Computational science curricula and CI technology transfer.
5. We will adapt and develop where necessary clearly defined and sustainable research data storage models and provide an extensible framework for data management and curation that is adaptable by distinct research disciplines.
 - a. We will engage the campus libraries in this effort and define their long-term roles clearly.
6. We will extend to campus researchers new network services that support reliable, high-performance connectivity among research groups and clusters on campus, the CI capabilities located in the new off-campus data center, and University collaborators both nationally and globally.
7. We will address common workflow and middleware requirements of campus research projects and coordinate these with national initiatives
8. As recommended in many campus CI committee summary reports over the past decade, we will continue to consider the creation of a broad computational science institute on campus for research, education and sustainability in campus CI efforts by bringing together broad expertise and experience.
9. We will work to increase baseline CI knowledge and technology adoption among students, research staff, and faculty.
10. We will continue to encourage the adoption of new core Internet technologies, i.e. IPv6, in expanding areas of the University infrastructure via specific requirements and focused projects.

List of prioritized, near-term actions

Specific goals continued in 2013 include:

- Foster collaboration among campus stakeholders to enhance identity management on campus, to plug IT vulnerabilities, and facilitate integration among central IT and research applications.
- Increase statewide visibility and to build the case for expanded state support for CI, including a presence at the SC13 high performance computing conference and exhibition in Denver this November and advance planning for SC16, which will return to Salt Lake City in three years.

- Expand regional collaboration among research universities in large-scale CI activities expanding on existing network-based collaboration with research universities and their affiliates in Utah, Idaho, and Montana.
- Seek to build critical mass in knowledge and informational resources for faculty and students.
- Increase the University of Utah's IPv6 presence by:
 - Making key computational user interactive nodes and certain campus research facing services available via IPv6 in 2013.
 - Expediting the current University Infrastructure Portfolio project of IPv6 Readiness whose scope is to create standards for all new projects, packages and equipment for the support of IPv6.
 - Implementing security policies and purchases so that the University's IPv6 security posture is equivalent to or better than the existing IPv4 stance.
 - Supporting the University's architectural review of all new IT projects, packages, and equipment for IPv6 compatibility.
- Continue the collaboration with Internet2 DYNES effort (supported by an NSF MRI award) to implement dynamic circuit capabilities.
- Implement the Science/Performance DMZ on separate discrete hardware to allow for a completely separate ingress/egress to the campus.
- Develop a response to the FY13 NSF CC-NIE solicitation that incorporates the University's plans for a Science DMZ and the strong innovation-based network research program under the leadership of Profs. Kobus Van der Merwe and Rob Ricci in the School of Computing.
- Expand the existing perfSONAR active network monitoring system
- Support Big Data management, storage, and analytics implementations such as the Sloan Digital Sky Survey 4 (SDSS-4) site at the University under the leadership of Prof. Adam Bolton of Physics and Astronomy.
- Collaborate with the Marriott (academic campus) and Eccles (medical campus) Libraries and the Office of the Vice President of Research regarding data management and curation initiatives.
- Expand and enhance the Identity and Access Management mechanisms and coordinate them tightly with national initiatives by
 - Supporting the expanded use of the InCommon trust federation.
 - Supporting the expansion of the internal use of the single sign-on infrastructure.
- Expand metropolitan optical network connectivity with additional GENI specific wavelengths and additional campus 100-Gbps wavelengths.
- Implement 100-Gbps connectivity from the campus network to the Downtown Data Center and the Internet2 Network node in Salt Lake City.
- Develop plans to address needs for increased computational cycles on campus to support research and plans for a sustainable funding model for campus research computing.
- Work with the NSF XSEDE collaboration to support new University users and transition of codes to the national HPC environment via the XSEDE Campus Champions program.
- Develop policies and recommendations for addressing growth in research computing and Cyberinfrastructure, including addressing space needs, allocations and plans for growth in both the University of Utah Downtown Data Center and in campus buildings.