CYBERINFRASTRUCTURE VISION
FOR 21ST CENTURY DISCOVERY

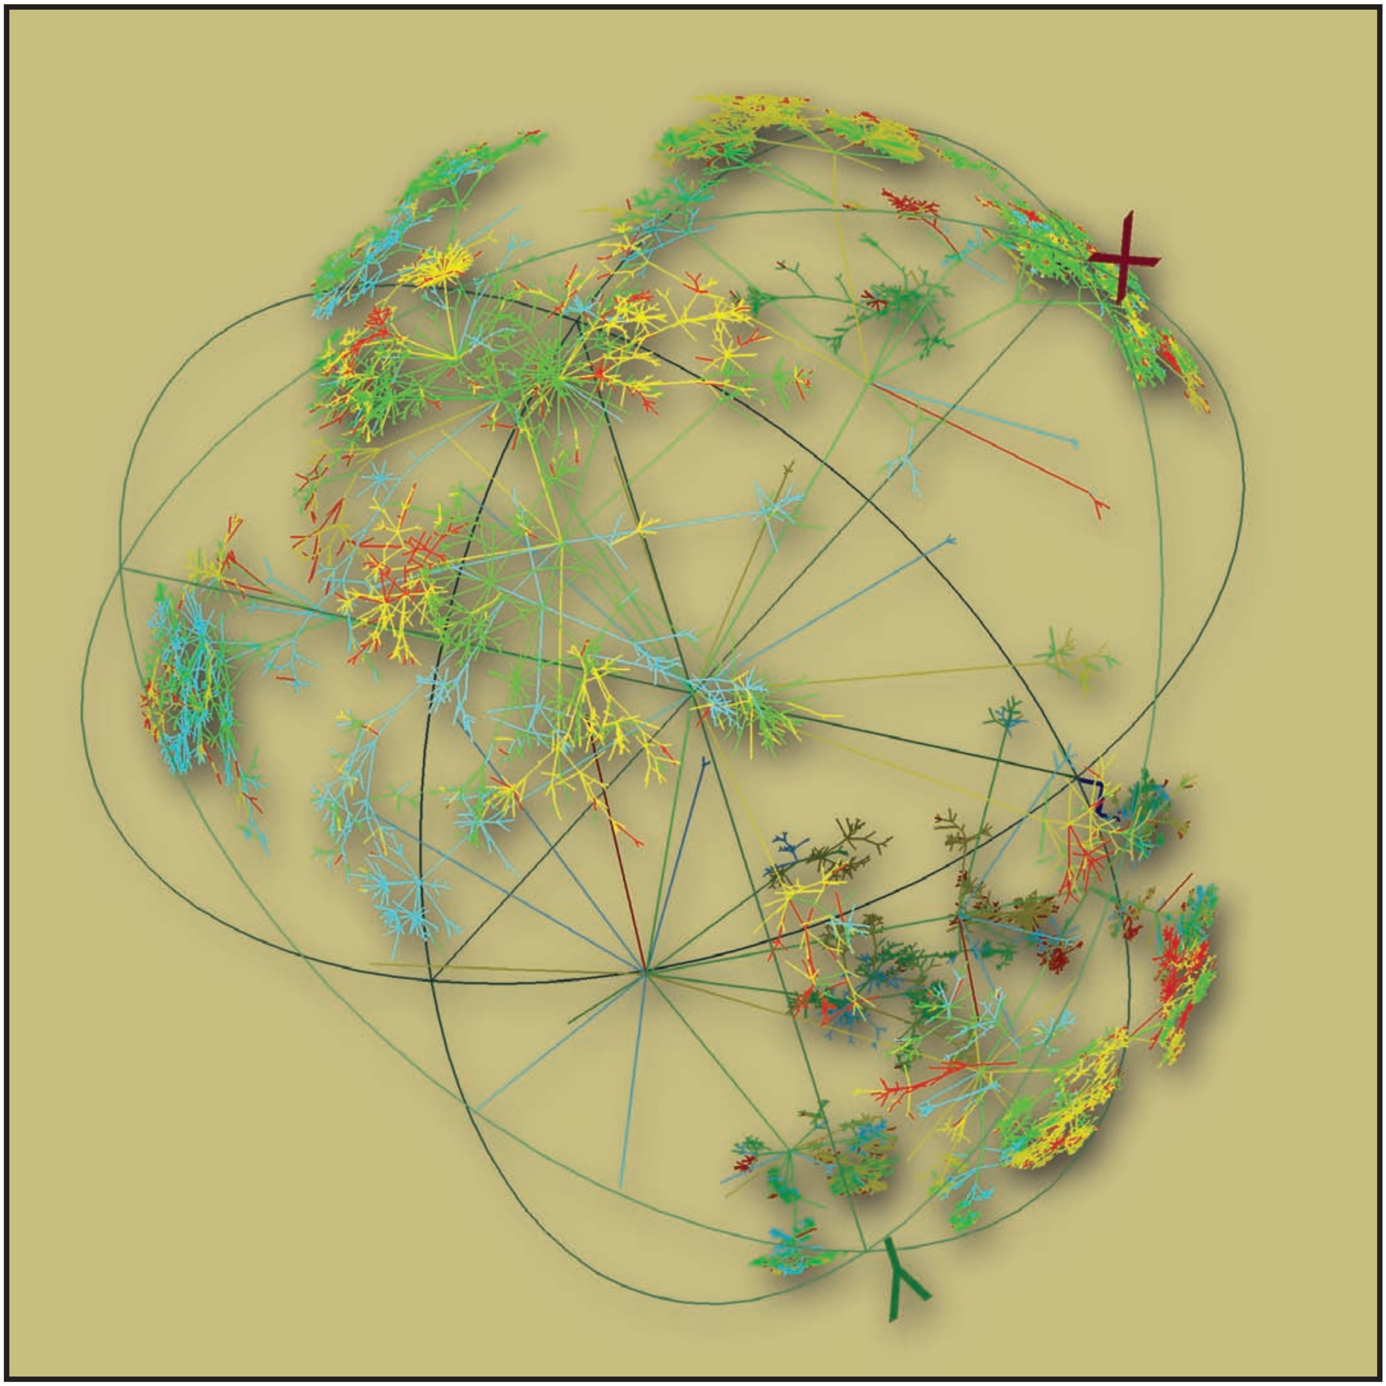

National Science Foundation
Cyberinfrastructure Council
March 2007

ABOUT THE COVER

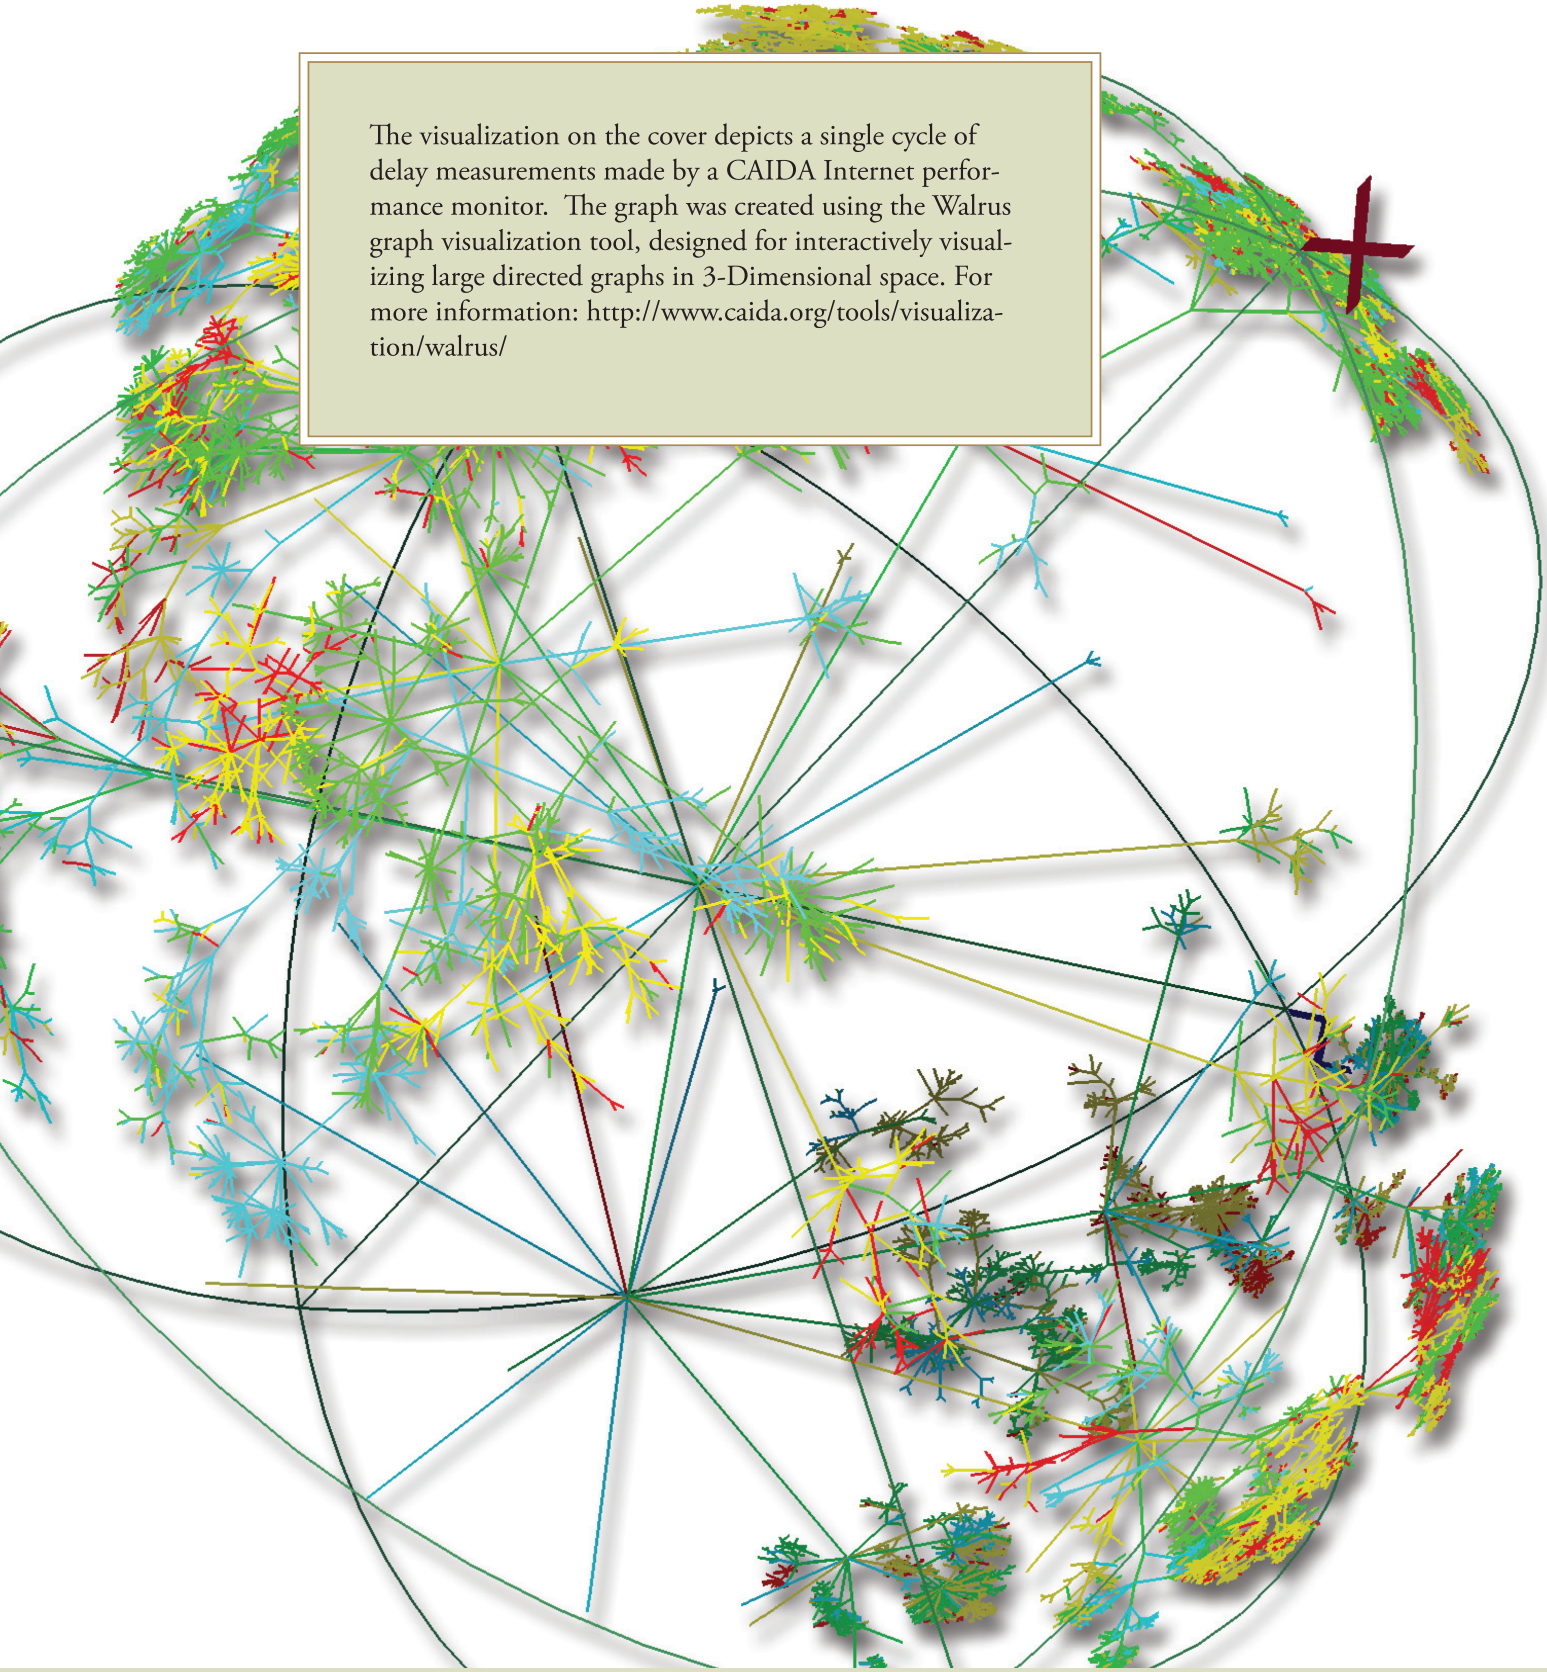

The visualization on the cover depicts a single cycle of delay measurements made by a CAIDA Internet performance monitor. The graph was created using the Walrus graph visualization tool, designed for interactively visualizing large directed graphs in 3-Dimensional space. For more information: http://www.caida.org/tools/visualization/walrus/

LETTER FROM THE DIRECTOR

Dear Colleague:

I am pleased to present NSF's Cyberinfrastructure Vision for 21st Century Discovery. This document, developed in consultation with the wider science, engineering, and education communities, lays out an evolving vision that will help to guide the Foundation's future investments in cyberinfrastructure.

At the heart of the cyberinfrastructure vision is the development of a cultural community that supports peer-to-peer collaboration and new modes of education based upon broad and open access to leadership computing; data and information resources; online instruments and observatories; and visualization and collaboration services. Cyberinfrastructure enables distributed knowledge communities that collaborate and communicate across disciplines, distances and cultures. These research and education communities extend beyond traditional brick-and-mortar facilities, becoming virtual organizations that transcend geographic and institutional boundaries. This vision is new, exciting and bold.

Dr. Arden L. Bement, Jr. , Director of National Science Foundation

Realizing the cyberinfrastructure vision described in this document will require the broad participation and collaboration of individuals from all fields and institutions, and across the entire spectrum of education. It will require leveraging resources through multiple and diverse partnerships among academia, industry and government. An important challenge is to develop the leadership to move the vision forward in anticipation of a comprehensive cyberinfrastructure that will strengthen innovation, economic growth and education.

Sincerely,

Arden L. Bement, Jr.
Director

PREFACE

The National Science Foundation's Cyberinfrastructure Council (CIC)[1], based on extensive input from the research community, has developed a comprehensive vision to guide the Foundation's future investments in cyberinfrastructure (CI). In 2005, four multi-disciplinary, cross-foundational teams were created and charged with drafting a vision for cyberinfrastructure in four overlapping and complementary areas: 1) High Performance Computing, 2) Data, Data Analysis, and Visualization, 3) Cyber Services and Virtual Organizations, and 4) Learning and Workforce Development. Draft versions of the document were posted on the NSF website and public comments were solicited from the community. These drafts were also reviewed for comment by the National Science Board. The National Science Foundation thanks all of those who provided feedback on the Cyberinfrastructure Vision for 21st Century Discovery document. Your comments were carefully reviewed and considered during preparation of this version of the document, which is intended to be a living document, and will be updated periodically.

ACKNOWLEDGEMENTS

We acknowledge the following NSF personnel who served on the strategic planning teams and whose efforts made this document possible. We especially acknowledge Deborah Crawford, who served as acting director for OCI from July 2005 to June 2006, and whose leadership was instrumental in the formulation of this document.

High Performance Computing (HPC) CI Team: Deborah Crawford (Chair), Leland Jameson, Margaret Leinen (CIC Representative), José Muñoz, Stephen Meacham, Michael Plesniak

Data CI Team: Cheryl Eavey, James French, Christopher Greer, David Lightfoot (CIC Representative), Elizabeth Lyons, Fillia Makedon, Daniel Newlon, Nigel Sharp, Sylvia Spengler (Chair)

Virtual Organizations (VO) CI Team: Thomas Baerwald, Elizabeth Blood, Charles Boudin, Arthur Goldstein, Joy Pauschke (Co-Chair), Randal Ruchti, Bonnie Thompson, Kevin Thompson (Co-Chair), Michael Turner (CIC Representative)

Learning and Workforce Development (LWD) CI Team: James Collins (CIC Representative), Janice Cuny, Semahat Demir, Lloyd Douglas, Debasish Dutta (Chair), Miriam Heller, Sally O'Connor, Michael Smith, Harold Stolberg, Lee Zia

---

[1] Complete list of acronyms can be found in Appendix A.

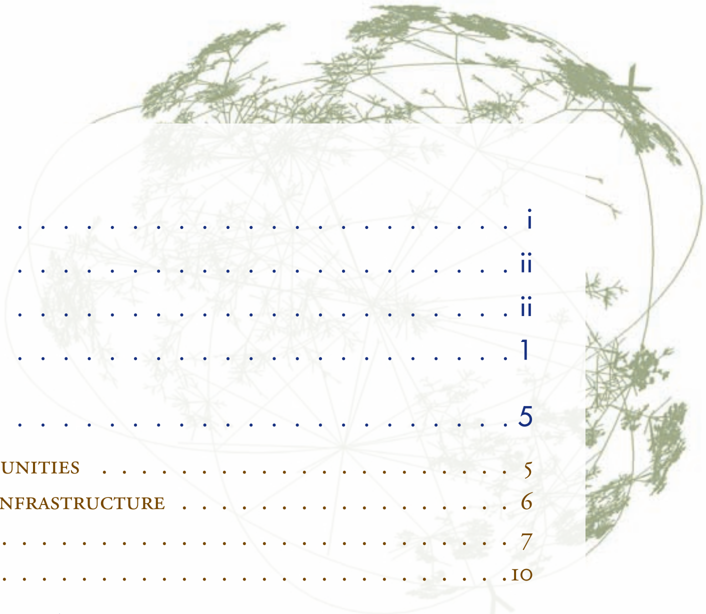

TABLE OF CONTENTS

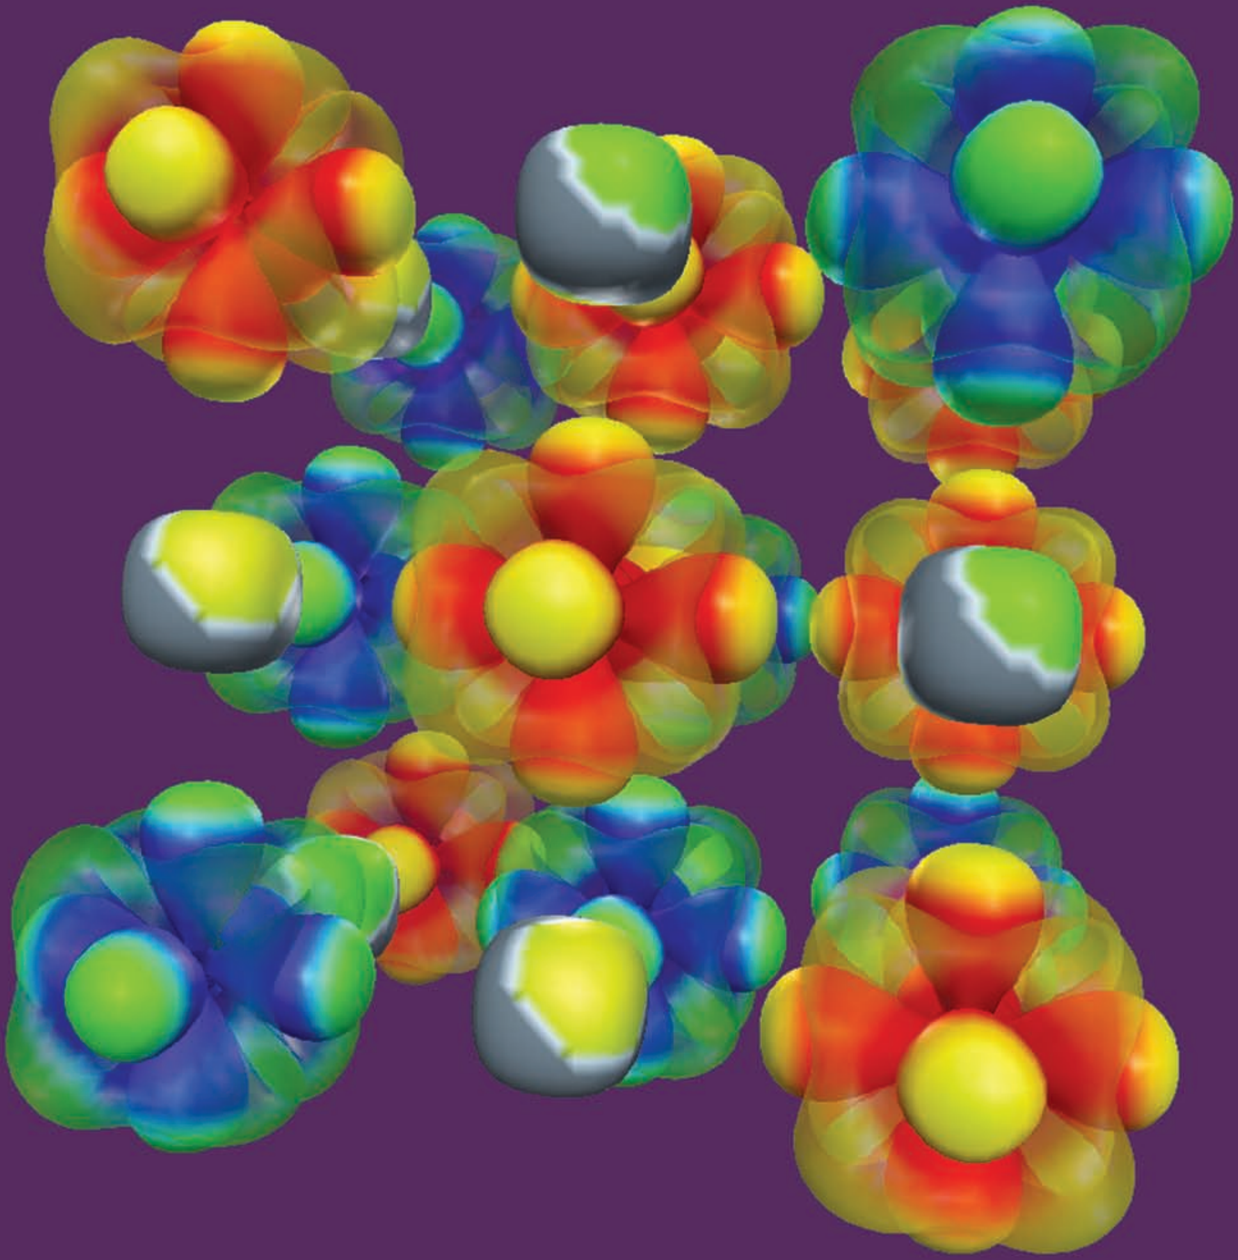

EXECUTIVE SUMMARY

NSF's Cyberinfrastructure Vision for 21st Century Discovery is presented in a set of interrelated chapters that describe the various challenges and opportunities in the complementary areas that make up cyberinfrastructure: computing systems, data, information resources, networking, digitally enabled-sensors, instruments, virtual organizations, and observatories, along with an interoperable suite of software services and tools. This technology is complemented by the interdisciplinary teams of professionals that are responsible for its development, deployment and its use in transformative approaches to scientific and engineering discovery and learning. The vision also includes attention to the educational and workforce initiatives necessary for both the creation and effective use of cyberinfrastructure.

The five chapters of this document set out NSF's cyberinfrastructure vision. The first, ***A Call for Action***, presents NSF's vision and commitment to a cyberinfrastructure initiative. NSF will play a leadership role in the development and support of a comprehensive cyberinfrastructure essential to 21st century advances in science and engineering research and education. The vision focuses on a time frame of 2006-2010. The mission is for cyberinfrastructure to be human-centered, world-class, supportive of broadened participation in science and engineering, sustainable, and stable but extensible. The guiding principles are that investments will be science-driven, recognize the uniqueness of NSF's role, provide for inclusive strategic planning, enable U.S. leadership in science and engineering, promote partnerships and integration with investments made by others in all sectors, both national and international, and rely on strong merit review and on-going assessment, and a collaborative governance culture. This chapter goes on to review a set of more specific goals and strategies for NSF's cyberinfrastructure initiative along with brief descriptions of the strategy to achieve those goals.

High Performance Computing (HPC) in support of modeling, simulation, and extraction of knowledge from huge data collections is increasingly essential to a broad range of scientific and engineering disciplines, often multi-disciplinary (e.g. physics, biology, medicine, chemistry, cosmology, computer science, mathematics), as well as multi-scalar in dimensions of space (e.g., nanometers to light-years) time (e.g., picoseconds[1] to billions of years), and complexity. A vision for petascale[2] science and engineering for the academic community, enabled by high performance computing, is presented along with a series of principles that would be used to guide NSF science-driven HPC investments. This would result in a sustained petascale capable system deployed in the FY 2010 timeframe. The plan presented addresses HPC acquisition and deployment and various aspects of HPC software and tools, in addition to the necessary scalable applications that would execute on these HPC assets.

An effective computing environment designed to meet the computational needs of a range of science and engineering applications will include a variety of computing systems with complementary performance capabilities. NSF will invest in leadership class environments in the 0.5-10 petascale performance range. Strong partnerships involving other federal agencies, universities, industry and state government are also critical to success. NSF will also promote resource sharing between and among academic institutions to optimize the accessibility and use of HPC assets deployed and supported at the campus level. Supporting software services include the provision of intelligent development and problem-solving environments and tools. These tools are designed to provide improvements in ease of use, reusability of modules, and portable performance.

---

[1] A picosecond is 10^{-12} second

[2] A petascale is 10^{15} operations per second with comparable storage and networking capacity

The image shows computed charge density for iron oxide (FeO) within the local density approximation, with spherical ions subtracted. The colors represent the spin density, showing the antiferromagnetic ordering.

Data, Data Analysis, and Visualization are vital for progress in the increasingly data-intensive realm of science and engineering research and education. Any cogent plan addressing cyberinfrastructure must address the phenomenal growth of data in all its various dimensions. Scientists and engineers are producing, accessing, analyzing, integrating, storing and retrieving massive amounts of data daily. Further, this is a trend that is expected to see significant growth in the very near future as advances in sensors and sensor networks, high-throughput technologies and instrumentation, automated data acquisition, computational modeling and simulation, and other methods and technologies materialize. The anticipated growth in both the production and repurposing of digital data raises complex issues not only of scale and heterogeneity, but also of stewardship, curation and long-term access.

Responding to the challenges and opportunities of a data-intensive world, NSF will pursue a vision in which science and engineering digital data are routinely deposited in well-documented form, are regularly and easily consulted and analyzed by specialist and non-specialist alike, are openly accessible while suitably protected, and are reliably preserved. To realize this vision, NSF's goals for 2006-2010 are twofold: to catalyze the development of a system of science and engineering data collections that is open, extensible, and evolvable; and to support development of a new generation of tools and services for data discovery, integration, visualization, analysis and preservation. The resulting national digital data framework will be an integral component in the national cyberinfrastructure framework. It will consist of a range of data collections and managing organizations, networked together in a flexible technical architecture using standard, open protocols and interfaces, and designed to contribute to the emerging global information commons. It will be simultaneously local, regional, national and global in nature, and will evolve as science and engineering research and education needs change and as new science and engineering opportunities arise.

Virtual Organizations for Distributed Communities, built upon cyberinfrastructure, enable science and engineering communities to pursue their research and learning goals with dramatically relaxed constraints of time and distance. A virtual organization is created by a group of individuals whose members and resources may be dispersed geographically and/or temporally, yet who function as a coherent unit through the use of end-to-end cyberinfrastructure systems. These CI systems provide shared access to centralized or distributed

Researchers create cyberenvironments—secure, easy-to-use interfaces to instruments, data, computing systems, networks, applications, analysis and visualization tools, and services.

resources and services, often in real-time. Such virtual organizations supporting distributed communities go by numerous names: collaboratory, co-laboratory, grid community, science gateway, science portal, and others. As such environments become more and more functionally complete they offer new organizations for discovery and learning and bold new opportunities for broadened participation in science and engineering.

Creating and sustaining effective virtual organizations, especially those spanning many traditional organizations, is a complex technical and social challenge. It requires an open technological framework consisting of, for example, applications, tools, middleware, remote access to experimental facilities, instruments and sensors, as well as monitoring and post-analysis capabilities. An operational framework from campus level to international scale is required, as well as a need for partnerships between the various cyberinfrastructure stakeholders. Overall effectiveness also depends upon the appropriate social, governance, legal, economic and incentive structures. Formative and longitudinal evaluation is also necessary both to inform iterative design as well as to develop understanding of the impact of virtual organizations on enhancing the effectiveness of discovery and learning.

Learning and Workforce Development opportunities and requirements recognize that the ubiquitous and interconnected nature of cyberinfrastructure will change not only how we teach but also how we learn. The future will see increasingly open access to online educational resources including courseware, knowledge repositories, laboratories, and collaboration tools. Collaboratories or science gateways (instances of virtual organizations) created by research communities will also offer participation in authentic inquiry-based learning. These new modes and opportunities to learn and to teach, covering K-12, post-secondary, the workforce and the general public, come with their own set of opportunities and challenges. New assessment techniques will have to be developed and understood; undergraduate curricula must be reinvented to fully exploit the capabilities made possible by cyberinfrastructure; and the education of the professionals that are being relied upon to support, develop and deploy future generations of cyberinfrastructure must be addressed. In addition, cyberinfrastructure will have an impact on how business will be conducted and members of the workforce must have the capability to fully exploit the benefits afforded by these new technologies.

Cyberinfrastructure-enhanced discovery and learning is especially exciting because of the opportunities it affords for broadened participation and wider diversity along individual, geographical and institutional dimensions. To fully realize these opportunities NSF will identify and address the barriers to utilization of cyberinfrastructure tools, services, and resources; promote the training of faculty, educators, students, researchers and the public; and encourage programs that will explore and exploit cyberinfrastructure, including taking advantage of the international connectivity it provides - particularly important as we prepare a globally engaged workforce.