



Research Directions in the I/O Technology and Infrastructure

Almadena Chtchelkanova

Program Director

**National Science Foundation (NSF)
Directorate for Computer and Information
Science and Engineering (CISE)**

achtchel@nsf.gov

Outline

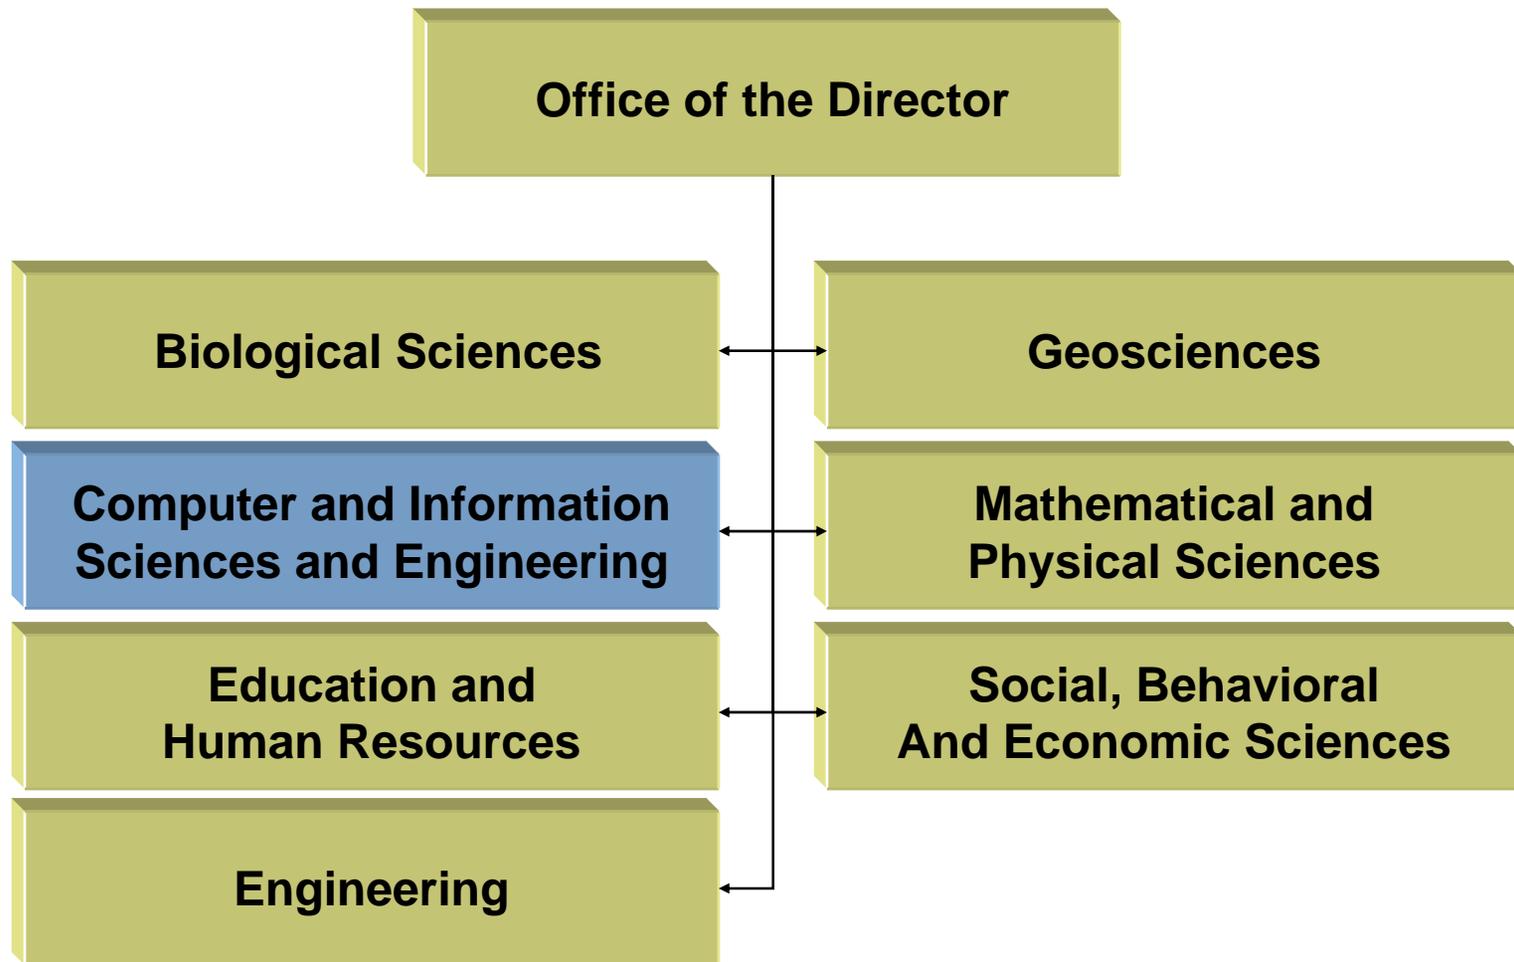
- **Introduction to NSF**
 - NSF Mission, organization & strategic objectives
 - CISE Organization
 - Divisions, Clusters, Programs
 - FY 2006 activities; FY 2007 plans
- **Highlighted Programs**
 - Cyber Trust
 - Science of Design
 - GENI Initiative
 - CRI, MRI, I/UCRC
- **NSF Review Process**



Outline (cont.)

- **HEC RTF**
- **HEC IWG**
- **HEC URA**
- **ST-HEC 2004**
- **HEC IWG FS I/O Workshop**
- **HECURA 2006**

How NSF is organized



National Science Foundation

- **Basic scientific research & research fundamental to the engineering process;**
- **Programs to strengthen scientific and engineering research potential;**
- **Science and engineering education programs at all levels and in all fields of science and engineering; and,**
- **A knowledge base for science and engineering appropriate for development of national and international policy**

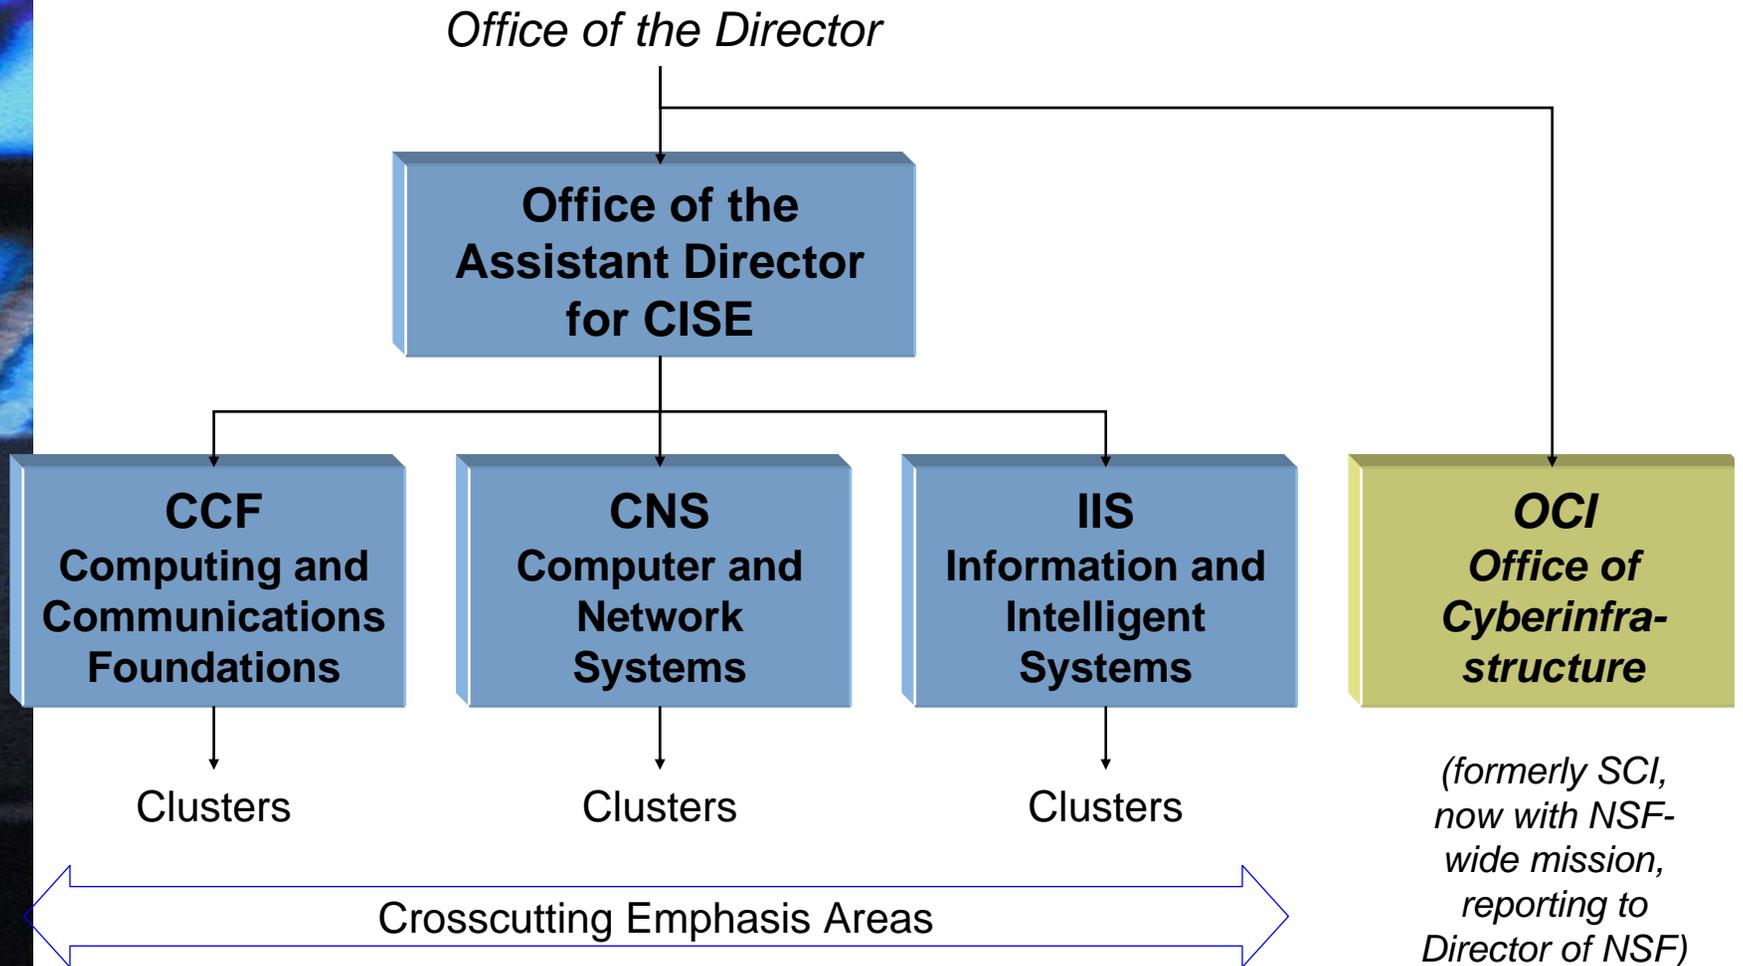
NSF Strategic Mission

- **People:**
to develop a diverse, internationally competitive and globally-engaged workforce of scientists, engineers, and well-prepared citizens
- **Ideas:**
to provide a deep and broad fundamental science and engineering knowledge base
- **Tools:**
to provide widely accessible, state-of-the-art science and engineering infrastructure
- **Organization Excellence:**
to develop an agile, innovative organization that fulfills its mission through leadership in state-of the-art business practices

CISE Mission

- **CISE has three goals:**
 - **to enable the United States to remain competitive in computing, communications, and information science and engineering;**
 - **to promote understanding of the principles and uses of advanced computing, communications, and information systems in service to society; and**
 - **to contribute to universal, transparent, and affordable participation in an information-based society.**

How CISE is organized



Office of CyberInfrastructure (OCI)

- **New Office of CyberInfrastructure (OCI) created in the Office of the Director**
 - **Division of Shared CyberInfrastructure (SCI) in CISE no longer exists**
- **NSF's plans for CyberInfrastructure may be found in: NSF's Cyberinfrastructure Vision for 21st Century Discovery**
 - **Consists of four chapters covering the period (2006-2010):**
 - **Strategic plan for High Performance Computing**
 - **Strategic plan for Data, Analysis and Visualization**
 - **Strategic plan for Collaboratories, Observatories and Virtual Organizations**
 - **Strategic plan for Education and the Workforce**

CISE Strategic Objectives

- **Push the Frontiers of Computer Science**
 - **Cyber Trust (cybersecurity)**
 - **Science of Design**
 - **Emerging models of computation**
 - **Theory**
- **Advanced Applications**
- **Research for cyberinfrastructure**
- **Broaden participation**
 - **Education & Workforce Preparation**
- **Improve organizational effectiveness**

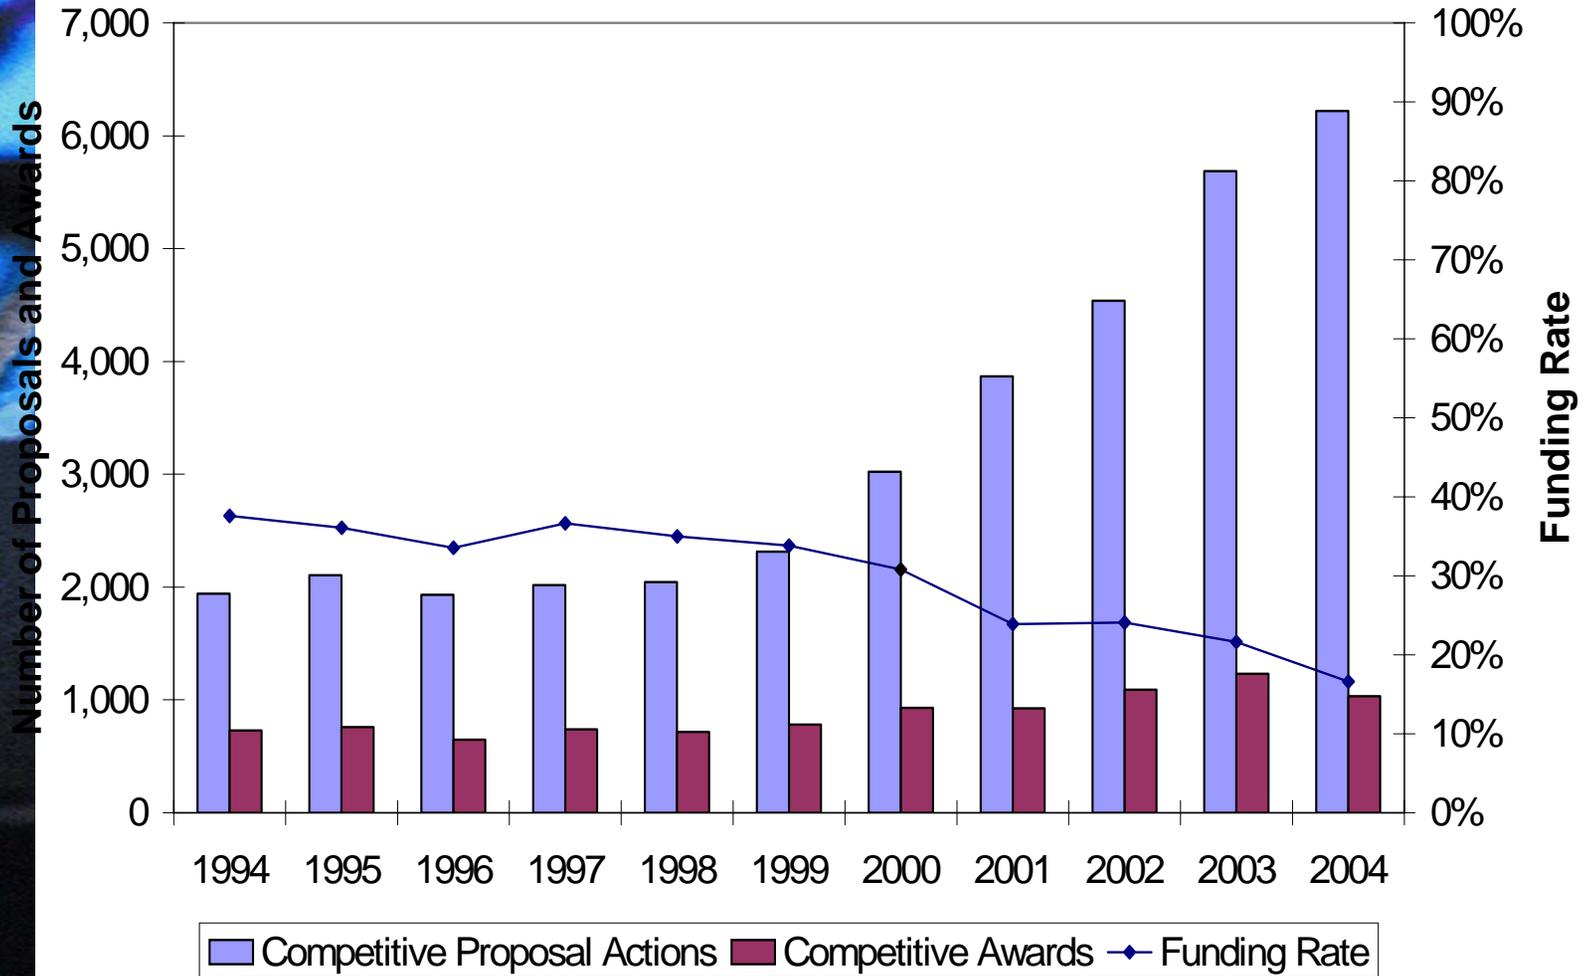
CISE (and related) Budget FY05 Actual (\$M)

| <i>CISE - Divisions</i> | <i>FY 2005</i> |
|---|-----------------|
| CCF | \$91.29 |
| CNS | \$132.17 |
| IIS | \$92.31 |
| ITR (not a division; cross-CISE) | \$174.43 |
| <i>CISE Total Research</i> | <i>\$490.20</i> |
| OCI | \$123.40 |
| <i>Major Research Equipment and Facilities Construction (MREFC)</i> | <i>\$165.60</i> |



**CISE provides 86%
of all Federal support for
computer science research**

Funding Rate for Competitive Awards in CISE

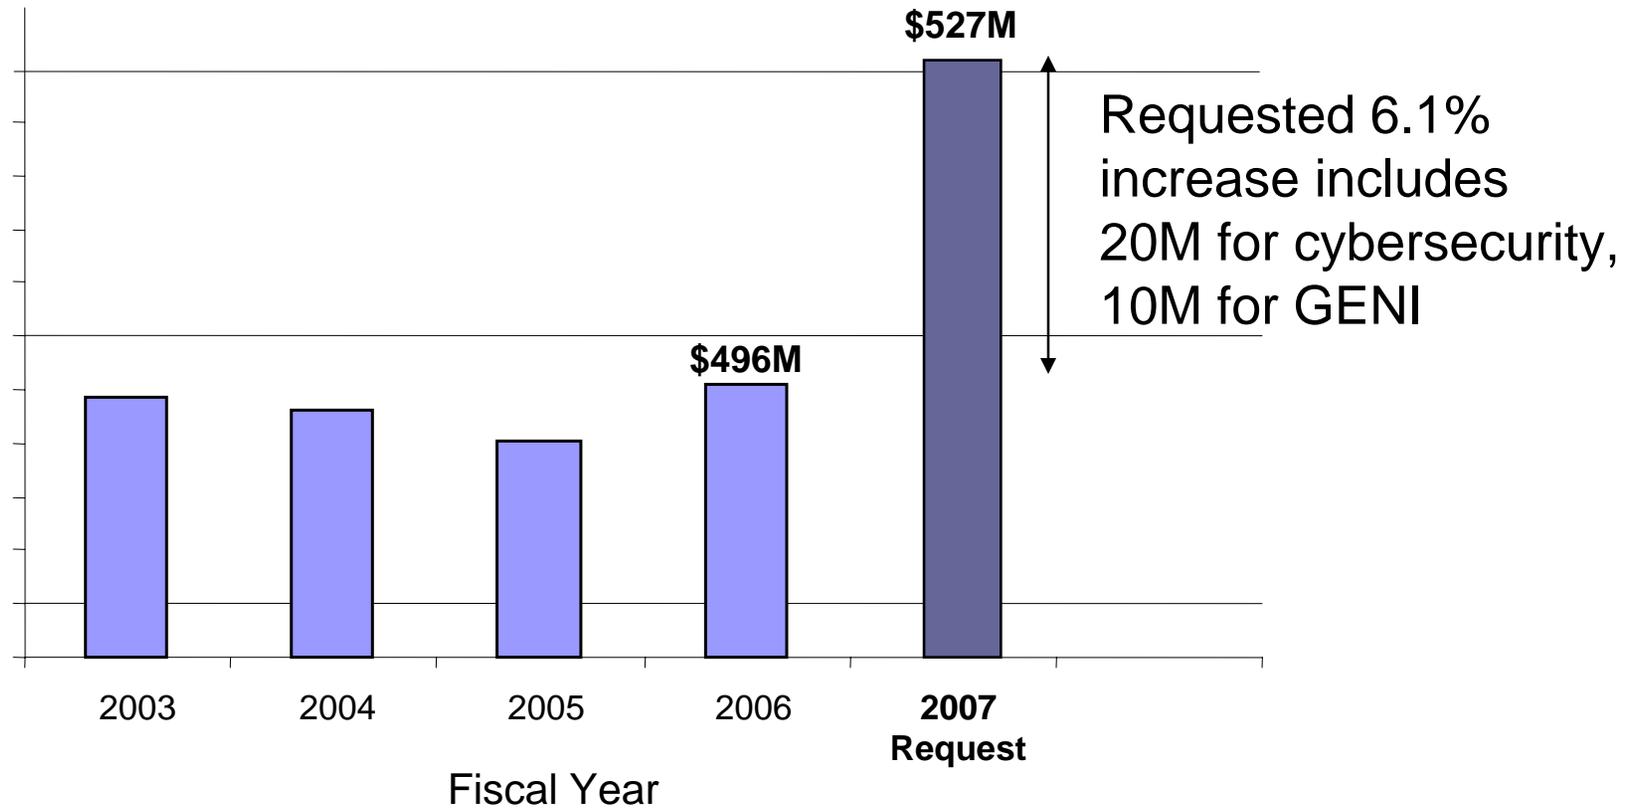


CISE Proposal/Award Statistics

| FY | Proposals | Awards | Funding Rate | CGIs | Supplements |
|-------------|------------------|---------------|---------------------|--------------|--------------------|
| 2005 | 4,962 | 1,086 | 23% | 1,398 | 581 |
| 2004 | 6,266 | 1,017 | 16% | 1,297 | 400 |
| 2003 | 5,346 | 1,174 | 22% | 1,023 | 354 |
| 2002 | 4,314 | 1,038 | 24% | 918 | 308 |
| 2001 | 3,579 | 885 | 25% | 768 | 231 |
| 2000 | 2,853 | 903 | 32% | 547 | 210 |
| 1999 | 2,209 | 746 | 34% | 493 | 301 |
| 1998 | 1,885 | 667 | 35% | 476 | 211 |
| 1997 | 1,894 | 684 | 36% | 527 | 219 |
| 1996 | 1,760 | 601 | 34% | 610 | 183 |
| 1995 | 1,941 | 708 | 36% | 631 | 215 |

**ADJUSTED*

CISE Budget: 2003-2007



Funding Outlook

- **NSF funds available to support computing have nearly doubled in the past five years**
- **However, proposals have almost tripled**
- **From less than one per year per CS faculty member to more than one per year**
- **Greatly expanded scope of research**

Computing and Communication Foundations Division (CCF)

- **Theoretical Foundations**
 - Computer science theory; numerical computing; computational algebra and geometry; signal processing and communication
- **Foundations of Computing Processes and Artifacts**
 - Software engineering; software tools for HPC; programming languages; compilers; computer architecture; graphics and visualization
- **Emerging Models and Technologies for Computation**
 - Computational biology; quantum computing; nano-scale computing; biologically inspired computing

Computer and Network Systems Division (CNS)

- **Computer Systems**
 - Distributed systems; embedded and hybrid systems; next-generation software; parallel systems
- **Network Systems**
 - Networking research broadly defined plus focus areas
- **Computing Research Infrastructure**
 - Equipment and infrastructure to advance computing research
- **Education and Workforce**
 - IT workforce; special projects; cross-directorate activities (e.g., REU sites, IGERT, ADVANCE)

Information and Intelligent Systems Division (IIS)

- **Systems in Context**
 - Human computer interaction; educational technology; robotics; computer-supported cooperative work; digital government
- **Data, Inference & Understanding**
 - Databases; artificial intelligence; text, image, speech, and video analysis; information retrieval; knowledge systems
- **Science & Engineering Informatics**
 - Bioinformatics; geoinformatics; cognitive neuroscience; ...

CISE Cross-Cutting Emphasis Areas

- **Characteristics**
 - cut across clusters and divisions (and directorates)
 - address scientific or national priority
- **FY 2005 Emphasis Areas**
 - Cyber Trust
 - Science of Design
 - Information Integration
 - Broadening Participation in Computing
 - Computational Science/High End Computing
 - Dynamic Data Driven Application Systems
- **FY 2006 Emphasis Areas**
 - Cyber Trust: February 6, 2006
 - Science of Design: January 6, 2006
 - High End Computing University Research Activity February 3, 2006
 - Broadening Participation: April 5, 2006

Examining Networked Systems

- **Part of President's 2006 Budget Request to Congress (written in December 2004)**
- **Partially justifies FY06 funding for Networking Technology and Systems in CNS**
- **Notes that network models are from 1970-1990**
- **Calls for architectures based on technology advances and new requirements**
- **Method is systematic redesign of current network systems**
- **Full text available at <http://www.nsf.gov/about/budget/fy2006/>**

CyberTrust

- **Vision: A society in which**
 - **Computing systems operate securely and reliably**
 - **Computing systems protect sensitive information**
 - **Systems are developed and operated by a well-trained and diverse workforce**
- **Research on foundations, network security, systems software, and information systems**
- **Integrated education and workforce activities**

Science of Design

- ***About Computing: computers, computation, information, communication***
- ***Not about: buildings, bridges, airplane wings, traditional engineering design, nano, biotech, ..***
- ***However: desirable to import design research from other fields***
- ***How is “software” different from other materials from which artifacts are designed?***
- ***How is design of (distributed, embedded, heterogeneous,...) systems different or the same as design of other artifacts?***

Global Environment for Networking Investigations (GENI) Initiative

- **The GENI Initiative envisions the creation of new networking and distributed system architectures that, for example:**
 - **Build in security and robustness;**
 - **Enable the vision of pervasive computing and bridge the gap between the physical and virtual worlds by including mobile, wireless and sensor networks;**
 - **Enable control and management of other critical infrastructures;**
 - **Include ease of operation and usability; and**
 - **Enable new classes of societal-level services and applications.**

GENI Components

GENI Research Program

- **Support research, design, and development of new networking and distributed systems**
- **Build on many years of knowledge and experience**
- **Encouraging researchers and designers to:**
 - **reexamine all networking assumptions**
 - **reinvent where needed**
 - **design for intended capabilities**
 - **deploy and validate architectures**
 - **build new services and applications**
 - **encourage users to participate in experimentation**
 - **take a system-wide approach to the synthesis of new architectures**

GENI Components

(continued)

The GENI Facility will enable:

- Shared use through slicing and virtualization in time and space domains (i.e., where "slice" denotes the subset of resources bound to a particular experiment);
- Access to physical facilities through programmable platforms (e.g., via customized protocol stacks);
- Large-scale user participation by "user opt-in" and IP tunnels;
- Protection and collaboration among researchers by controlled isolation and connection among slices;
- A broad range of investigations using new classes of platforms and networks, a variety of access circuits and technologies, and global control and management software; and
- Interconnection of independent facilities via federated design.

GENI Outreach

- **CISE has supported numerous community workshops in support of GENI**
- **CISE is supporting on-going planning efforts, including needs assessment and requirements for the GENI Facility.**
- **CISE will hold town meetings and continue to support future workshops to broaden community participation.**
- **CISE will work with industry, other US agencies, and international groups to broaden participation in GENI beyond NSF and the US government.**

CISE Computing Research Infrastructure (CRI) annual

- **Infrastructure Acquisition.** These awards have budgets up to \$2,000,000.
- **Community Resource Development.** These awards have budgets from \$300,000 to \$2,000,000: medium from \$300,000 to \$800,000 and large over \$800,000. Development projects create a resource for an entire CISE research community, such as a testbed for evaluating research results or a large data resource that contains problems a community is trying to solve (e.g., annotated speech data).
- **Planning.** These awards facilitate the preparation of a proposal for a medium or large infrastructure acquisition grant. They have budgets up to \$50,000 for one institution or up to \$100,000 if more than one institution is involved.



NSF-wide crosscutting programs Major Research Instrumentation Program (MRI)

FY2006 budget \$90M

The Major Research Instrumentation Program (MRI) is designed to increase access to scientific and engineering equipment for research and research training in our Nation's organizations of higher education, research museums and non-profit research organizations. This program seeks to improve the quality and expand the scope of research and research training in science and engineering, and to foster the integration of research and education by providing instrumentation for research-intensive learning environments. The MRI program encourages the development and acquisition of research instrumentation for shared inter- and/or intra-organizational use and in concert with private sector partners.

NSF-wide crosscutting programs

Industry University Cooperative Research Program (I/UCRC)

- **Partnering Industries and Universities to Innovate.**
- **I/UCRCs stimulate highly leveraged industry/university cooperation by focusing on fundamental research recommended by Industrial Advisory Boards.**
- **I/UCRC develops long-term partnerships among industry, academic institutions, and government.**
- **The centers are catalyzed by a small investment from the National Science Foundation (NSF) and are primarily supported by center members, with NSF taking a supporting role in their development and evolution.**



Directorate for Computer &
Information Science & Engineering

NSF Merit Review Criteria

- **What is the intellectual merit and quality of the proposed activity?**
- **What are the broader impacts of the proposed activity?**

What is the intellectual merit of the proposed activity?

Potential Considerations:

- How important is the proposed activity to advancing knowledge and understanding within its own field or across different fields?
- How well qualified is the proposer (individual or team) to conduct the project? (If appropriate, reviewers will comment on the quality of prior work)
- How creative and original are the concepts?
- How well conceived and organized is the proposed activity?
- Is there sufficient access to resources?

What are the broader impacts of the proposed activity?

Potential Considerations:

- How well does the activity advance discovery and understanding while promoting teaching, training and learning?
- How well does the activity broaden the participation of underrepresented groups (e.g., gender, ethnicity, disability, geographic, etc.)?
- To what extent will it enhance the infrastructure for research and education, such as facilities, instrumentation, networks and partnerships?
- Will the results be disseminated broadly to enhance scientific and technological understanding?
- What may be the benefits of the proposed activity to society?

How to Help NSF

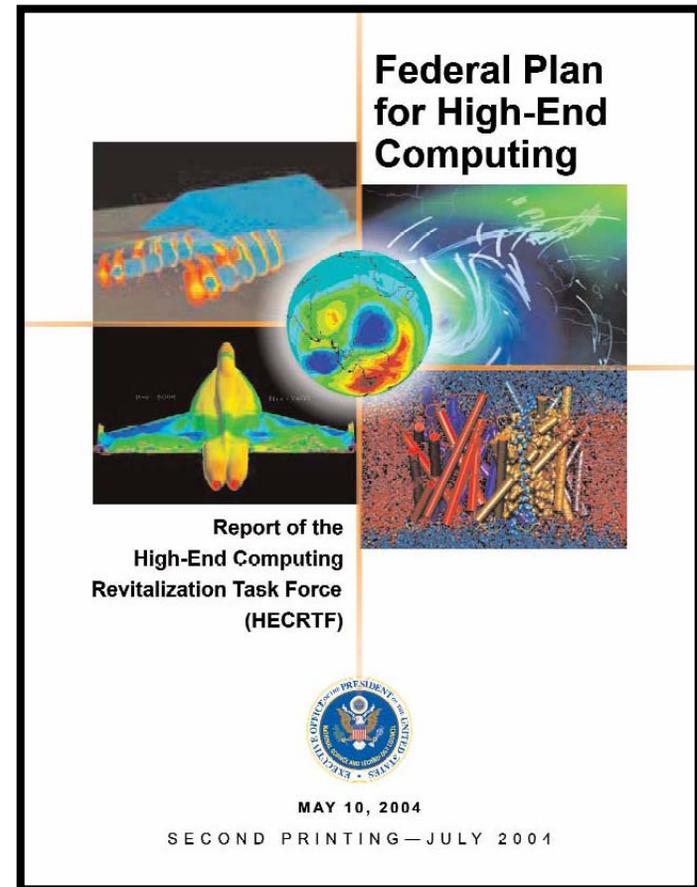
- **Send us your best ideas**
- **Participate in the process**
- **Keep us informed of your accomplishments**
- **Work within your institutions to support collaborative, interdisciplinary research**
- **Call our attention to things that need improvement**
- **Suggest transition strategies from basic research to prototyping and production**
- **Serve as a program officer (“rotator”)**

HECRTF HEC IWG HECURA

High-End Computing Revitalization Task Force (HECRTF)

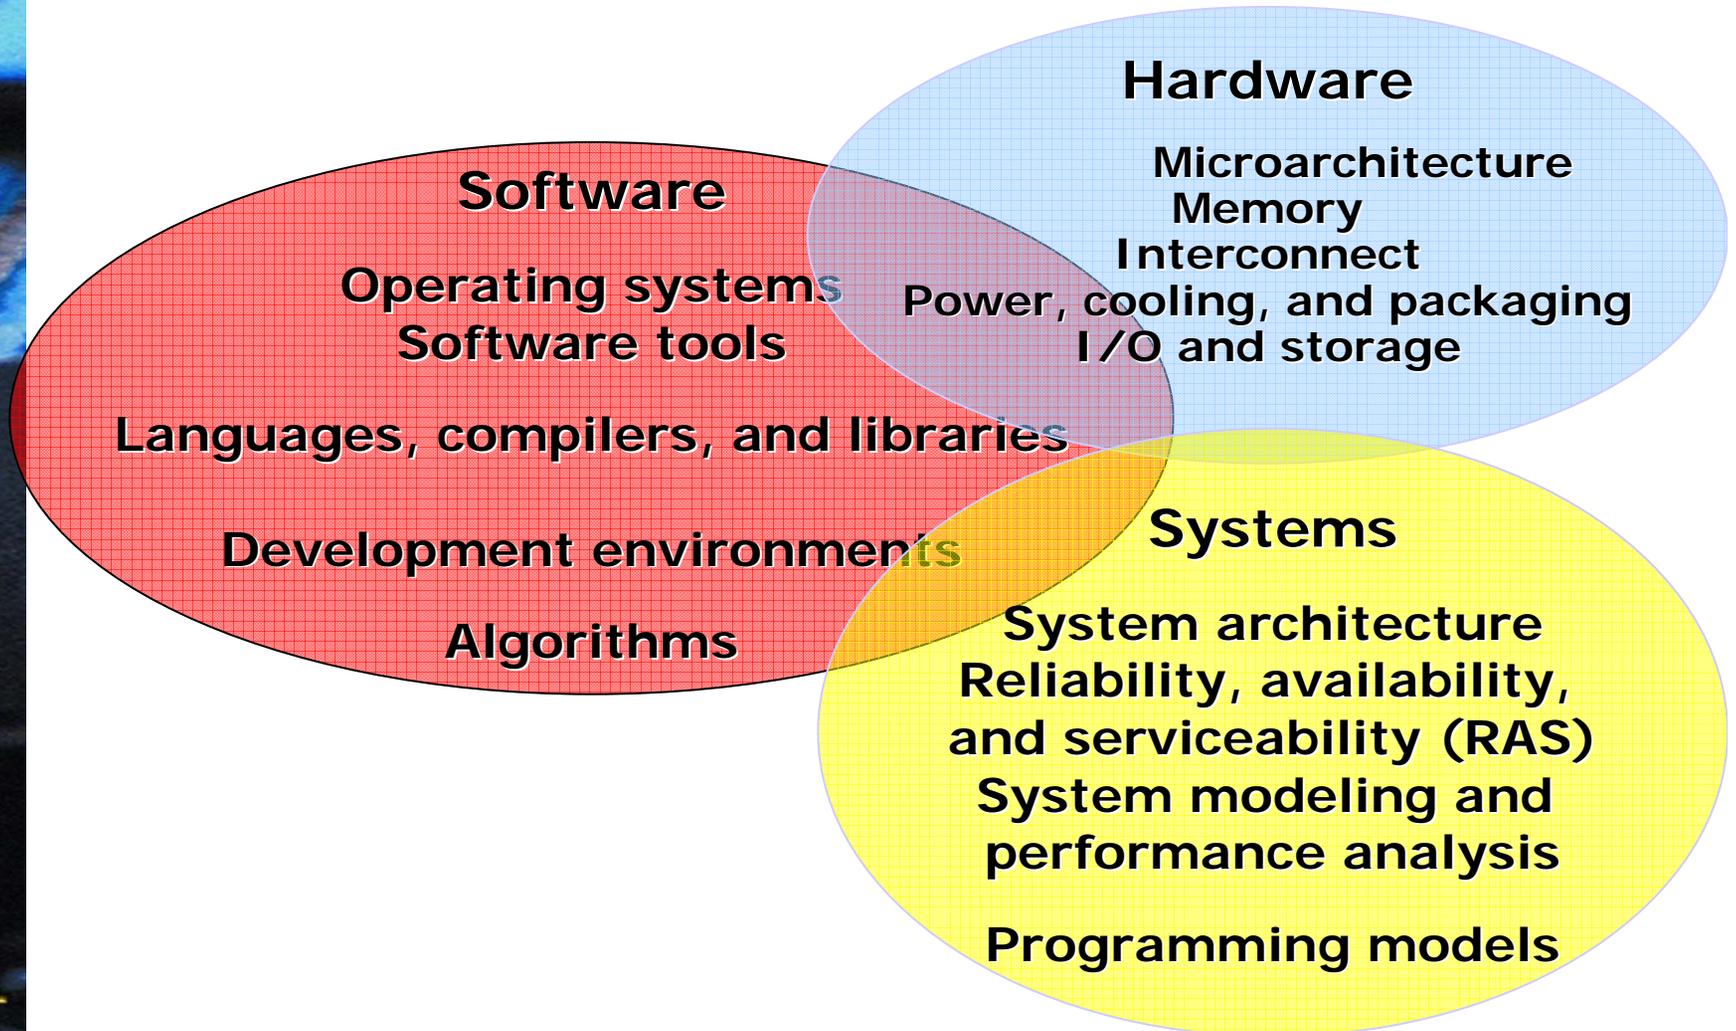
March 2003

NSF NIH DOE/SC
DOE/NNSA
DARPA NSA NASA
AHRQ NIST OSD
NOAA EPA AFRL
ARO DHS FAA
FDA NARA ONR



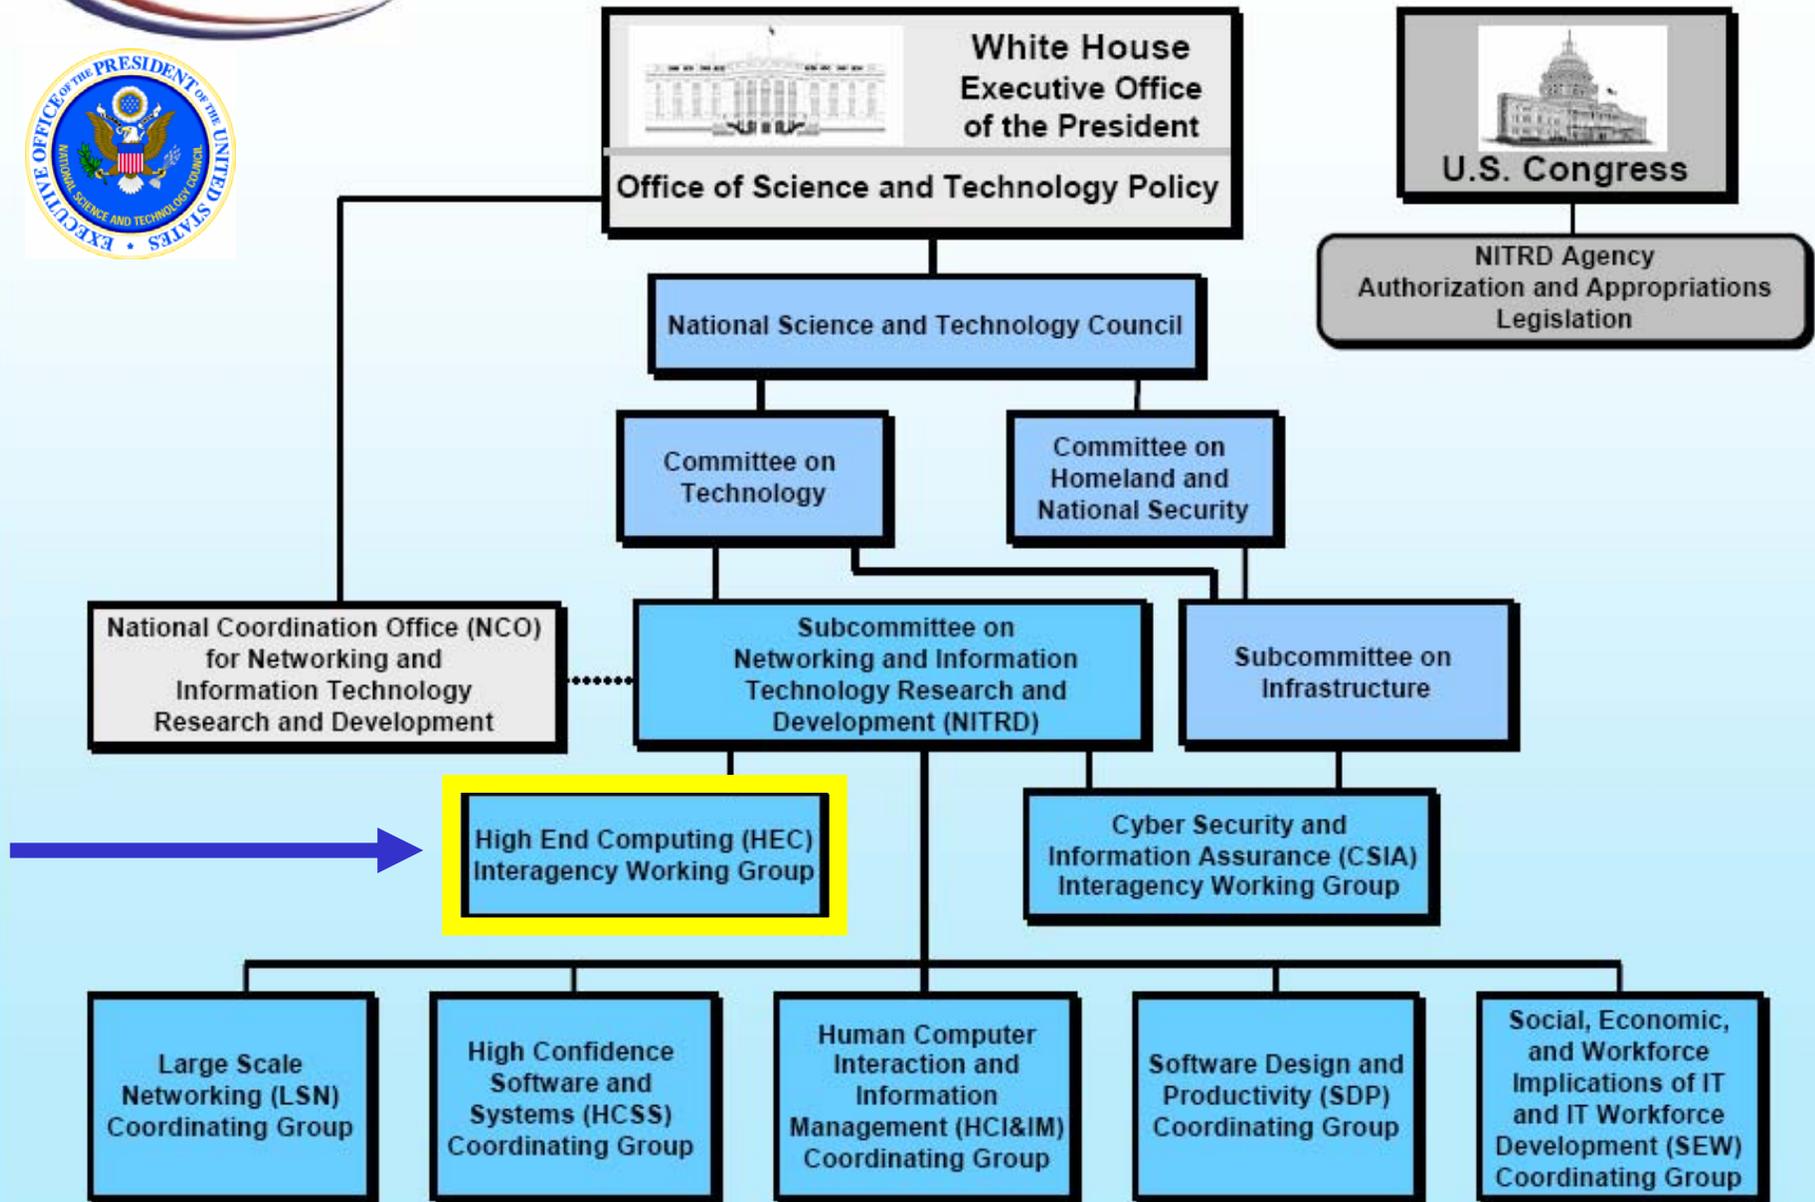


Potential High-End Computing Research Areas





NITRD Program Coordination





Interagency Collaborations in High-End Computing

- **High-End Computing University Research Activity (HECURA)**
- **Leadership Systems (DOE/OS, NASA)**
- **Performance/Productivity Measurement (DARPA, DOE/ OS, DOE NNSA, NASA, NSA, NSF, and others)**
- **DARPA High Productivity Computing Systems Program (DOE/OS, DOE NNSA, NSA, NASA, NSF, and others)**



High-End Computing University Research Activity (HECURA)

- FY 2004 Research Areas (22 awards)

- Runtime operating systems, languages, compilers, libraries
- Projects executed out of DOE/SC and NSF
- Funding provided by DOE/SC, NSF, DARPA, NSA

- FY 2006 Research Areas

- I/O, file and storage systems
- NSF solicitation deadline February 3, 2006
- Funding provided by NSF, DARPA, other agencies



NSF/DARPA ST-HEC 2004

- **100 proposals submitted in July 2004**
 - 82 projects submitted by 57 US academic institutions and non-profit organizations
 - Includes no-cost national lab and industrial lab collaborators
- **Nine projects were awarded**
 - Tools and libraries for high-end computing
 - Resource management
 - Reliability of high-end systems

DOE NNSA/Office of Science Recent and Current I/O R&D Focus

- Cluster File Systems that scale for aligned data operations and Symmetric NFS
 - Lustre, Panasas, PVFS2 (NNSA, Office of Science, OSC, and friends)
- User space file systems (ANL, NNSA)
- Beginning engineering and early R&D for other scalability dimensions
 - Metadata (Lustre, Panasas, PVFS2, NNSA, ANL, UCSC, SDSC)
 - Security (Lustre, Panasas, NNSA, UCSC, U of Minn, CMU, DOD/NSA)
- NFSv4 security, WAN, Performance (NNSA, U of Mich, and industry, NSF)
- pNFS (NNSA, U of Mich, CMU, and industry)
- Middleware to address usage, unaligned I/O, and efficient connection to High Level I/O libraries (ANL, Northwestern, NNSA)
- High Level I/O libraries (ANL, U of Chicago, NCSA, NNSA)
- Early utilization of intelligence near the disk drive (PNL, U of Minn/NNSA)
- Relaxation of Posix Semantics becoming more common (NNSA, ANL, PNL)
- Enterprise Sharing of Global Parallel File System – access from many clusters, (NNSA, UCSC, U of Minn)
- Autonomics (UCSC, NNSA, ANL, CMU)
- New storage devices for hierarchy – mems, mram (UCSC, NNSA, NSF, INSIC)
- Archive (UCSC, U of Minn, NNSA)
- Standards work - POSIX, NFSv4, pNFS, OBSD (NNSA, Office of Science)



HEC IWG File Systems and I/O Workshop

Dallas, TX August 15-18, 2005

- **Agencies/Sites**
 - Government HEC Sites
- **Top research institutions in this area**
 - Universities with major emphasis or center of excellence in the file systems and I/O area
- **Select research funding industry leads**
 - **ONLY** companies that fund file systems and I/O research, particularly at the top research institutions, derived from university sponsor lists



Areas of Research

- **Metadata**
 - Evolution - **Scalability, Extensibility, Archival considerations, ACL's**
 - Revolution - **Scalability, Extensibility, Name Spaces, Archival considerations, Hybrid devices**
- **Measurement and Understanding**
 - Evolution - **Understanding layering contribution, End to end benchmarking and tracing, Visualization**
 - Revolution - **End to end modeling and simulation, VM as tool for large scale simulation**
- **QOS**
 - Evolution - **Determinism with multiple applications and priorities**
 - Revolution - **Adaptive, End to end QOS**
- **Security**
 - Evolution - **Usability, Long term key management, Distributed authentication for file systems, Dealing with security overhead**
 - Revolution - **Novel security as related to file systems and I/O, Novel encryption at rest over time, Key Management, ACL's, End to End encryption API**



Areas of Research (cont.)

- **Next generation I/O architectures**
 - Evolutionary - POSIX, Archive considerations, Access aware interfaces, HEC considerations, Small/unaligned I/O, Mixed large and small I/O, Collaborative caching, Impedance matching
 - Revolutionary - Redistribution of intelligence and what abstractions we need, Adaptive/reconfigurable stack (application specific perhaps), User space component considerations, File systems that are semantically aware of the data. Novel devices/hybrid devices exploitation
- **File System related communications and protocols**
 - Evolution - Exploitation of RDMA/one sided etc., OBSD (transports, security, extensions, applications), NFSv4 (extensions and applications), pNFS (proof of concept, extensions, applications)
 - Revolution - Server to server communication
- **Management and RAS**
 - Evolution - Reliability and availability at scale end to end and its overhead, Management scaling, Continuous versioning, Power management
 - Revolution - Autonomics (adaptive/self healing/predictive), VM as a RAS enabler, Novel devices as enabler
- **Archive (as it relates to I/O and File Systems)**
 - Content addressable, Deep archive on disk, Object archives/parallel archives, Scheduling movement/ILM
- **Assisting research**
 - Testbeds, Clearing houses (providing traces, reliability info, etc.), Support growth of I/O students



HECURA - 2006 FOCUS

- I/O, file and storage systems design for efficient, high throughput data storage, retrieval and management in the HEC environment.
- hardware and software tools for design, simulation, benchmarking, performance measurement and tuning of file and storage systems.



HECURA - 2006 SCOPE

- File Systems Research
- Quality of Service
- I/O middleware
- Archives/Backups as extensions to file systems
- Novel storage devices for the IO stack
- I/O Architectures
- Management, reliability, and availability at scale
- Future File Systems related protocols
- Hardware and software tools for design, simulation of I/O, file and storage systems.
- Efficient benchmarking, tracing, performance measurement and tuning tools of I/O, file and storage systems

HECURA Status

- 62 proposals
- 8 collaborative projects
- 44 single PI projects

- 2 panels March 27-28 and March 30-31, 2006
- 9 reviewers out of 17 participated in both panels
- Panelists represented Academia, industry, DARPA, NSA, NASA, NRO, DOE/OS and DOE/NNSA
- Access to proposals was given to NIH HEC PMs



HECURA Planning Group

- **Candy Culhane (NSA) - Chair**
- **Almadena Chtchelkanova, Jose Munoz (NSF)**
- **Robert Graybill, William Harrod (DARPA)**
- **John Grosh (DOD)**
- **Fred Johnson (DOE/SC)**
- **Robert Meisner, Thuc Huang (DOE/NNSA)**



Acknowledgments

Gary Grider (LANL)

HECURA planning group

Michael Foster (NSF)

HECURA proposal reviewers



Agency NITRD Budgets by PCA

FY 2006 Budget Requests (dollars in millions)

| Agency | HEC I&A | HEC R&D | HCI&IM | LSN | HCSS | SDP | SEW | Totals |
|------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|-------------|
| NSF | 201.8 | 105.0 | 168.5 | 94.5 | 76.0 | 65.4 | 92.1 | 803 |
| NIH | 135.1 | 67.3 | 171.0 | 76.6 | 12.3 | 26.4 | 12.0 | 501 |
| DOE/SC | 105.7 | 82.0 | | 36.2 | | | 3.5 | 227 |
| DARPA | | 81.0 | 74.4 | 20.8 | | | | 176 |
| NSA | | 36.9 | | 1.5 | 62.2 | | | 101 |
| NASA | 34.0 | | 14.5 | 13.0 | 12.8 | | | 74 |
| AHRQ | | | 38.0 | 30.0 | | | | 68 |
| NIST | 5.4 | 0.6 | 8.6 | 4.6 | 18.0 | 4.8 | | 42 |
| OSD | | | | | 2.5 | 20.0 | | 22 |
| NOAA | 13.7 | 1.8 | 0.5 | 2.8 | | | | 20 |
| EPA | 3.3 | | 3.0 | | | | | 6 |
| Subtotals | 499.0 | 374.6 | 478.5 | 280.0 | 183.8 | 118.1 | 107.6 | 2041 |
| DOE /NSA | 33.1 | 30.5 | | 14.3 | | 31.6 | 4.4 | 114 |
| Totals | 532.1 | 405.1 | 478.5 | 294.3 | 183.8 | 149.7 | 112.0 | 2155 |

Other participating agencies: AFRL, ARO, DHS, FAA, FDA, NARA, ONR