Software Infrastructure for Sustained Innovation $(SI)^2$

Rajiv Ramnath and Dan Katz

Program Directors
Software Cluster
Division of Advanced Cyberinfrastructure
Directorate for Computer and Information Science and Engineering

rramnath@nsf.gov
dkatz@nsf.gov

Version: 2/15/16 20:49

Outline

- The NSF CI Vision and ACI
- Software CI Programs
- Software Programs – By The Numbers
- Example Projects
- Future Directions
- Other Programs of Relevance
- Credits

NSF CI VISION AND ACI


NSF
IRMINGER SEA
PIONEER ARRAY
REGIONAL SCALE NODES
ENDURANCE ARRAY
ARGENTINE BASIN
SOUTHERN OCEAN

NSF embraces an expansive view of cyberinfrastructure motivated by research priorities and the scientific process

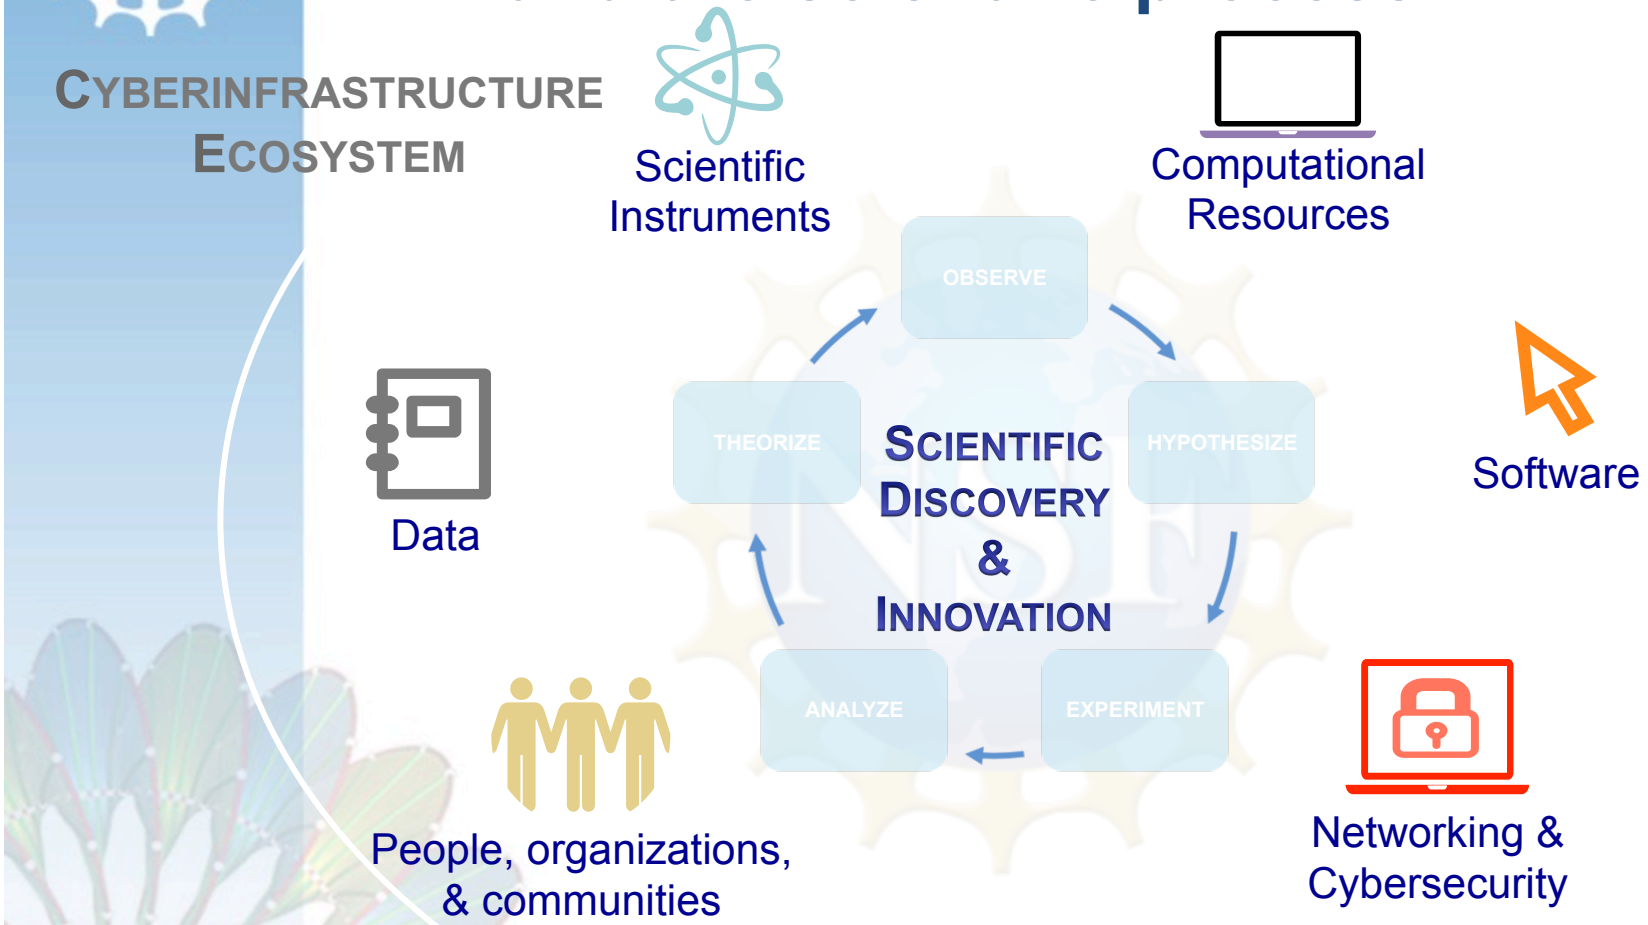

NSF's vision for cyberinfrastructure is informed by community input, development, and experience

Initial Vision (2007-2010)

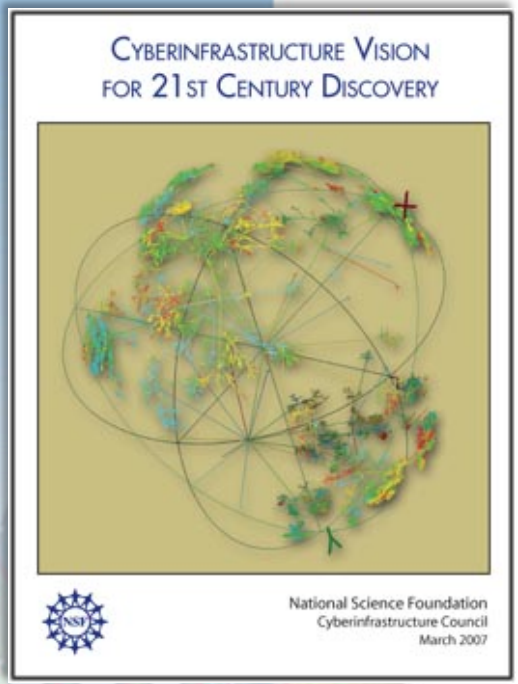

NSF-Wide Task Force Reports (2009-2011)

National Academies Study (On going)

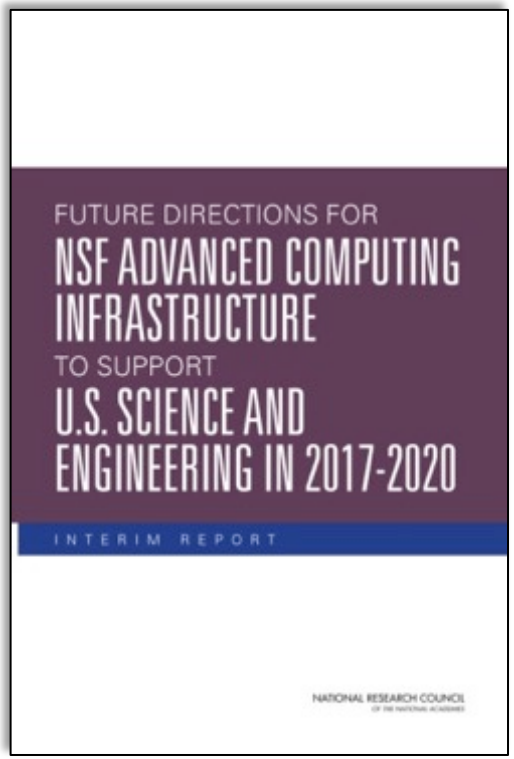

Interim Report, Oct 2014
Final Report expected Fall 2015

CI Challenge: User-Centric Viewpoint

Revolution in the scientific workflow: many interfaces to shared services

Large Facilities

Software

Shared Data/Software Gateway Resources

Researcher

Collaboration Networks

Data

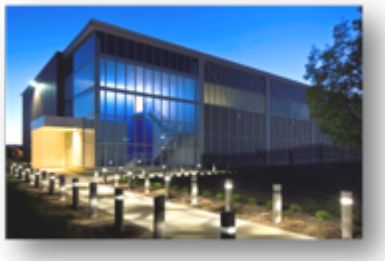

National Computing Resources

Cloud Services

Identities?
Resources?
Persistence?

ACI: Operational View

Supporting advance CI to accelerate discovery and innovation

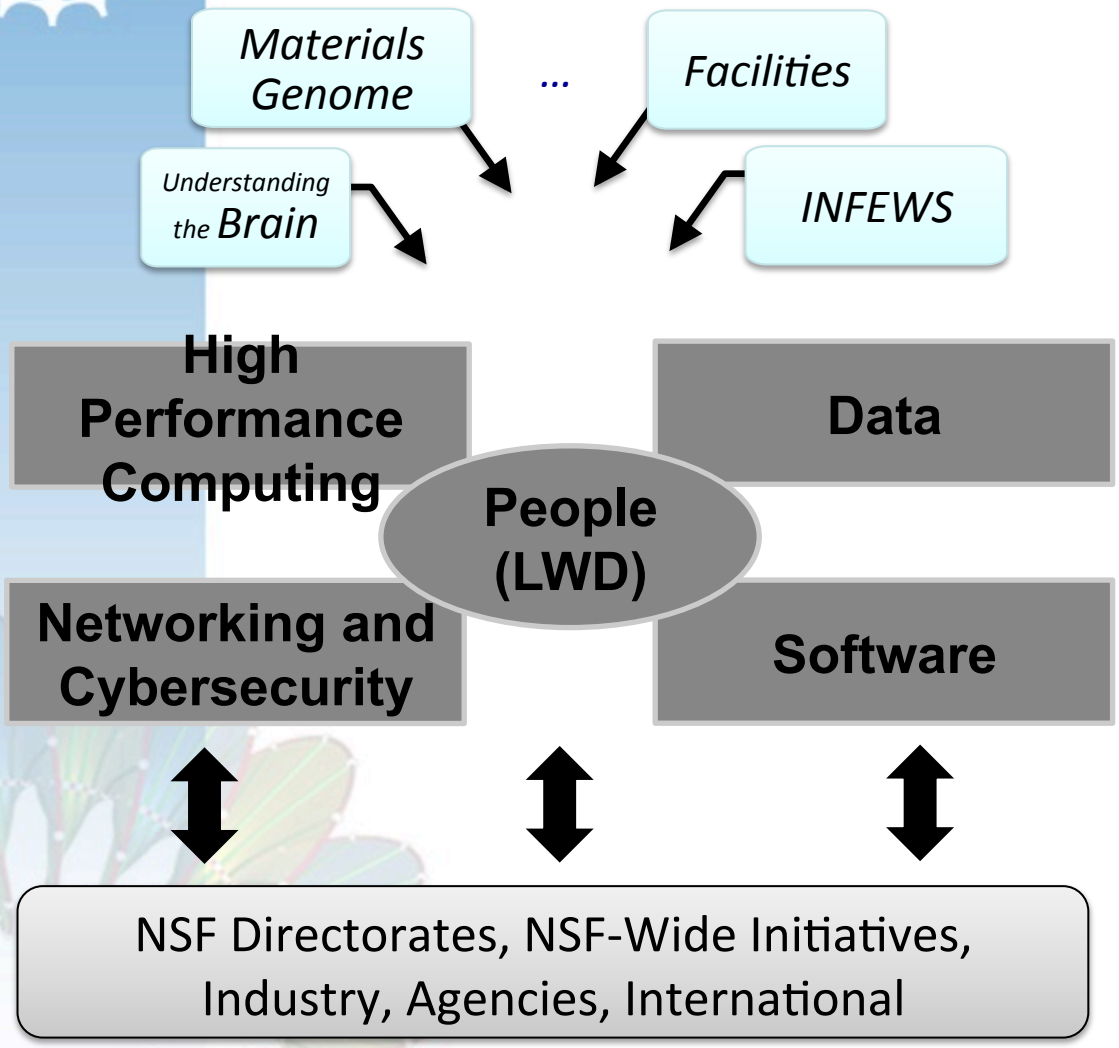

Science Drivers

Constant exchange with NSF Directorates, Divisions and Programs

ACI investments

Convergent investments in technologies and communities to maximize impact

Leadership, Coordination, Partnership

OVERVIEW OF SOFTWARE CI PROGRAMS

Software

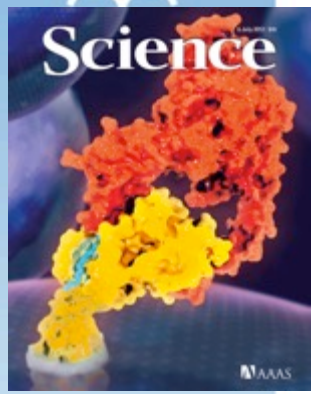

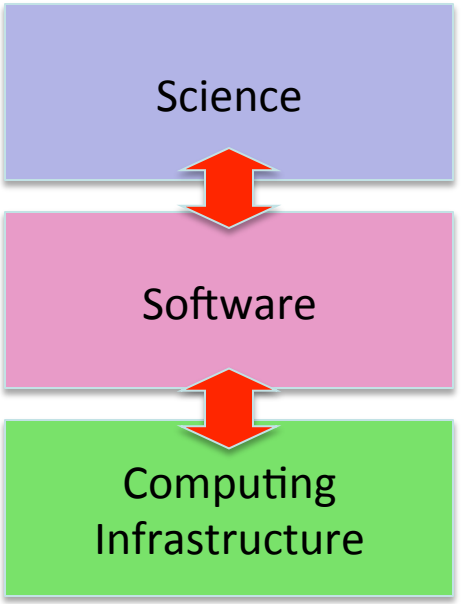

- Software (including services) essential for the bulk of science
  - About half the papers in recent issues of Science were software-intensive
  - Research becoming dependent upon advances in software
  - Wide range of software types: system, applications, modeling, gateways, analysis, algorithms, middleware, libraries
  - Significant software-intensive projects across NSF: e.g. NEON, OOI, NEES, NCN, iPlant, etc
- Software is not a one-time effort, it must be sustained
  - Development, production, and **maintenance** are people intensive
  - Software life-times are long vs hardware
  - Software has under-appreciated value

Many Software Programs - Cross-Cutting and Within Directorates

- Software Infrastructure for Sustained Innovation (S2I2)
- Computational and Data-Enabled Science and Engineering (CDS&E)
- Designing Materials to Revolutionize and Engineer our Future (DMREF)
- Software Venture (Reuse) Fund
- Exploiting Parallelism and Scalability (XPS)
- Advances in Biological Informatics (ABI)
- Metadata for Long-standing Large-Scale Social Science Surveys (META-SSS)
- Geoinformatics Program in the Division of Earth Sciences (EAR)
- Polar Cyberinfrastructure Program in Polar Programs
- Critical Techniques, Technologies and Methodologies for Advancing Foundations and Applications of Big Data Sciences and Engineering (BIGDATA)
- Critical Resilient Infrastructure Systems and Processes (CRISP)
- Integrative Graduate Education and Research Traineeship Program-CIF21 Track (IGERT)

ACI Software Program Strategy

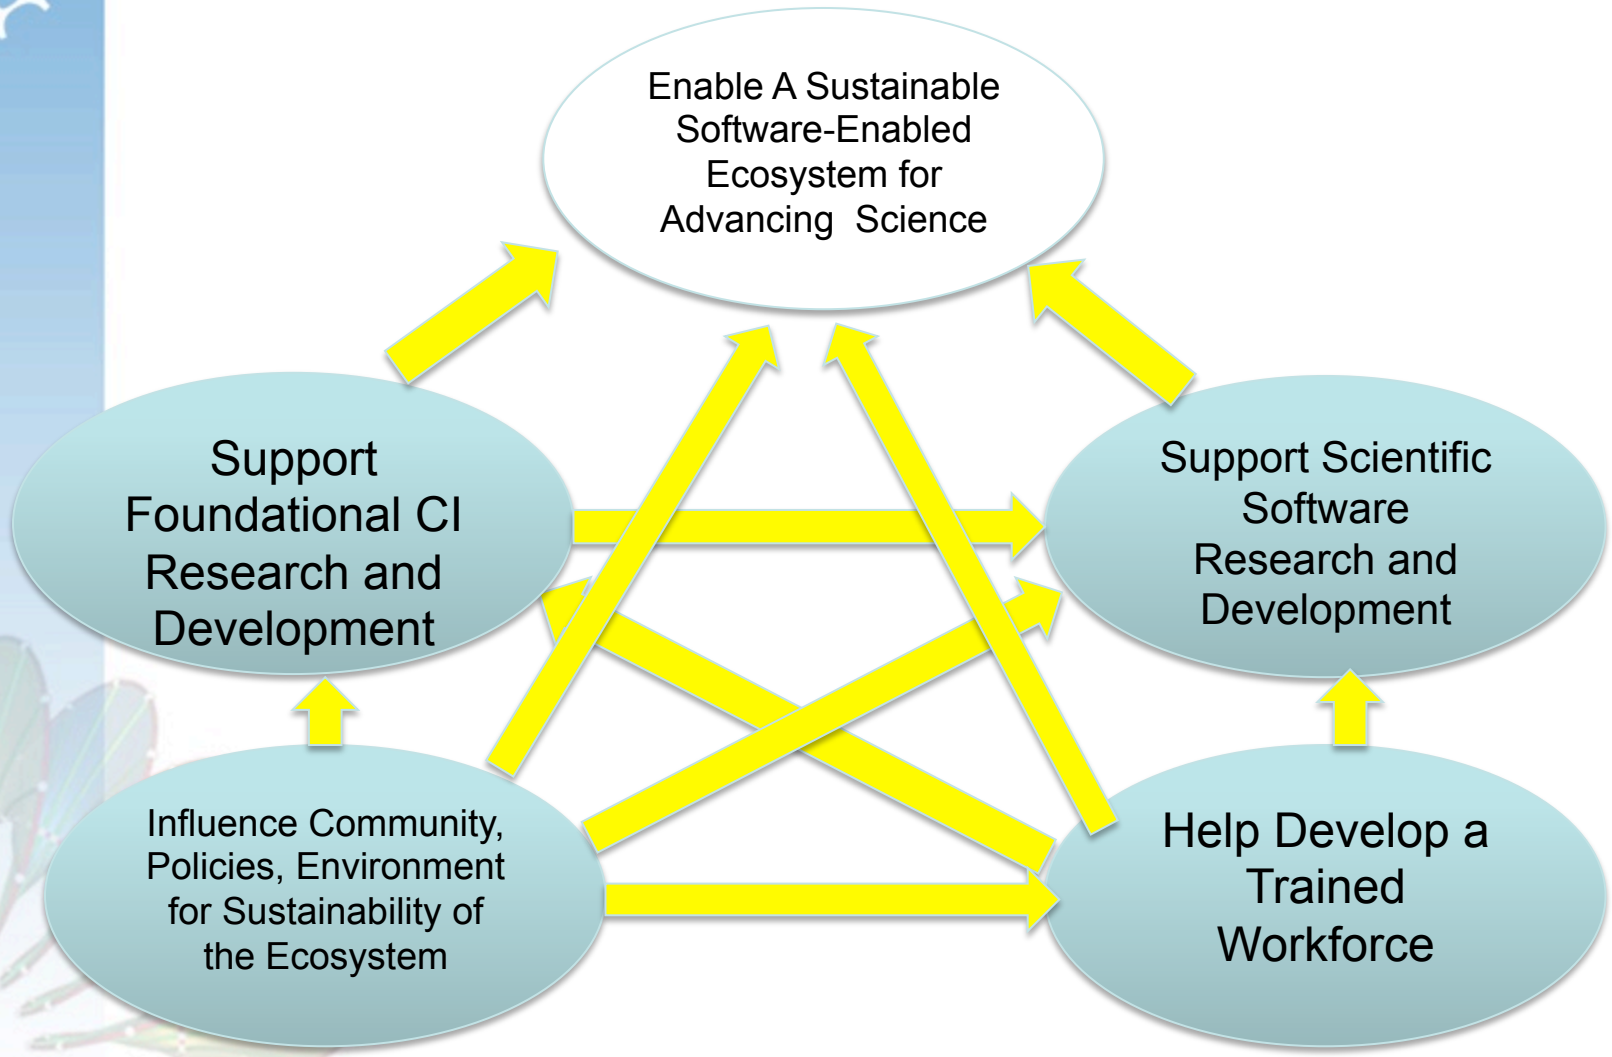

ACI Flagship - Software Infrastructure for Sustained Innovation (SI2)

Elements: $500K/3 years

Frameworks: $1M/ year 3-5 years

Institutes: $3-$5m/ year 5-10 years

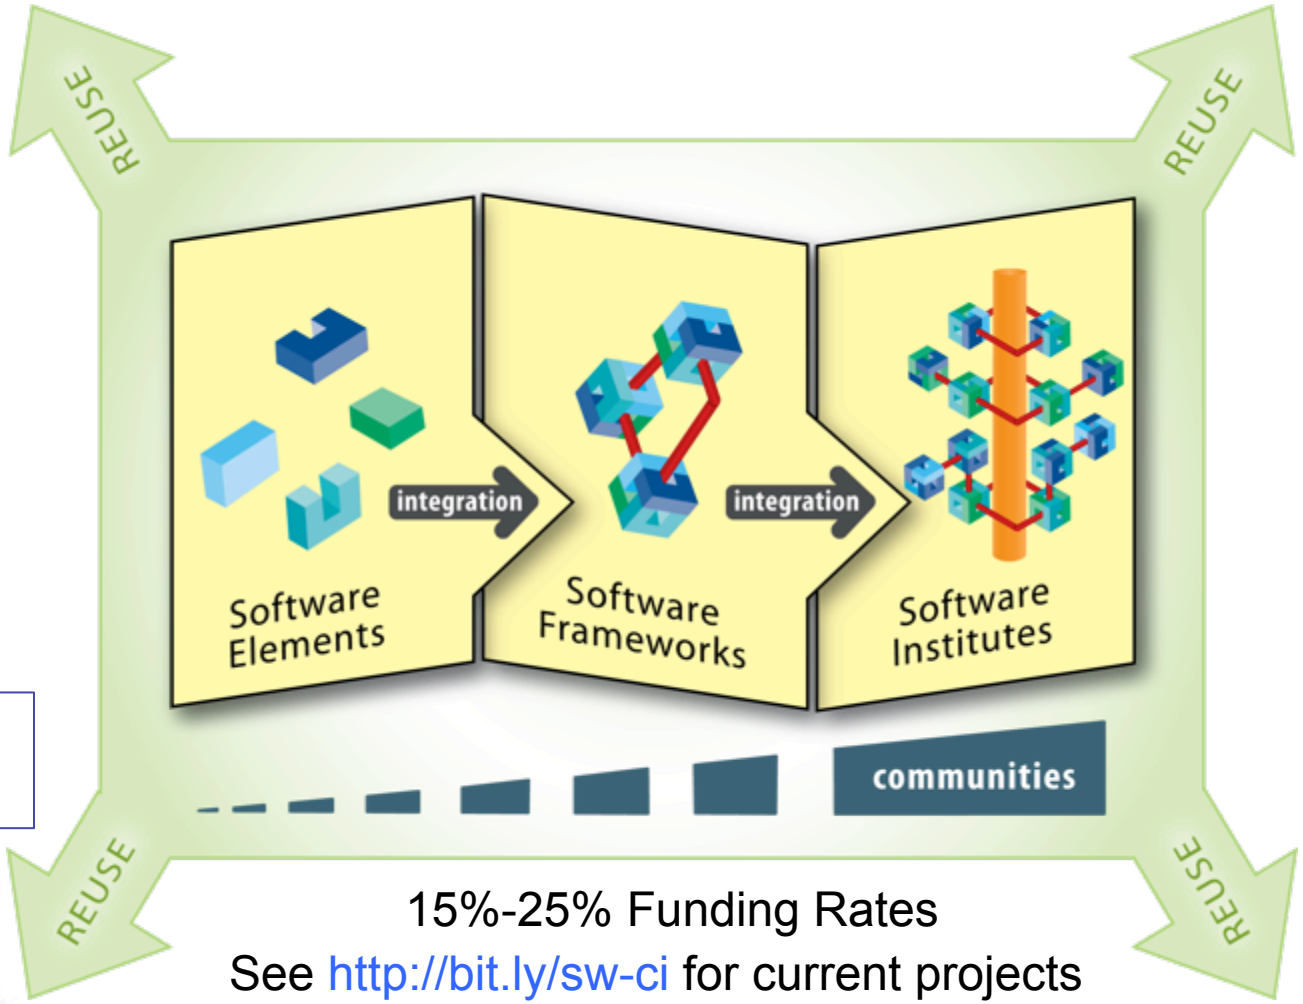

15%-25% Funding Rates

See http://bit.ly/sw-ci for current projects

Also: EAGERS, RAPIDs and Workshops to target areas of opportunity

Current SI2 Priorities - Towards a National Cyberinfrastructure Ecosystem

- Multidisciplinary and omni-disciplinary software
- That builds on other ongoing NSF-supported programs.
- Techniques, tools and processes for rapid integration of software that reduces cost of custom solutions and custom integrations
- Embedded innovation and research on the development, effectiveness, usability, adoption, and organizational aspects of the software and the project.
- Serious considerations of security, trustworthiness and reproducibility.
- Comprehensive, innovative approaches to sustainability (e.g. SAAS, incorporation into university offerings, commercialization)
- Science-inspired education and LWD
- Comprehensive metrics (ideally of impact)

Broader Software Challenges

- **Funding models**. NSF supports projects for up to 5 years; lifetime of software projects can span 20+ years
- **International funding.** Software collaborations span countries, but funding agencies don't
- **Career paths for software-focused researchers**. University structure and academic culture rewards publications; what about researchers whose main products are software? Stop the Google Bus…
- **Incentives, including credit**. How should software be cited? How are all software contributions recognized?
- **Training**. What software engineering practices work in science software?
- **Interdisciplinary work**. Many software contributors work in both computer science and another science or engineering area,or multiple areas, but doing so doesn't fit our siloed system and culture and is often discouraged.
- **Portability**. How to deal with changing hardware, middleware, and languages?
- **Dissemination.** Documenting available software and examples of how it has been used, strengths, weaknesses, and the experience of other users.

Example - Models for Sustainable Funding

- Community supported:
  - Open source licensing (of various types)
  - Suported by volunteer efforts by community
  - “Club” i.e. consortium fees
- Institutes
- Research Organizations (e.g. Universities):
  - Direct support of research infrastructure (like any other infrastructure)
  - Usage paid for via indirect costs on projects
  - Incorporation into curriculum (and paid for by tuition $$)
- NSF/Funding Agencies:
  - Long term funding from NSF (not for CI in NSF model)
  - CI line item on research projects
    - E.g. Software “credits” to projects (like a computation allocation to HPC/XSEDE) coupled with long-term funding of foundation projects
    - E.g. Direct budget line item
- Commercialization
  - License fees, royalties
- Hybrids of the above

SOFTWARE PROGRAMS - BY THE NUMBERS

SI2 Award Statistics

Award Year	Type	Funding	Percentage
2012	SSE	19/124	15%
2012	SSI	13/90	14%
2012	S2I2 (concept)	13/32	40%
2013	SSE	14/60	23%
2013	SSI	11/50	22%
2014	SSE	24/78	28%
2014	S2I2 (concept)	1/3	33%
2015	SSE	16/58	28%
2015	SSI	15/60	25%
2016	SSI	10/65 (expected)	12%
2016	SSE	10/? (projected)	?
2016	S2I2 (Implementation)	CMRSI, SGSI (1 each planned)	

Additional Software Proposal Stats – XPS, DMREF, CRISP

- CDS&E
  - FY13 - 2 of 8 (25%) ACI proposals funded, 8 projects co-funded
  - FY14 - 2 of 10 (20%) ACI projects funded, 12 projects co-funded,
  - FY15 - 16 projects co-funded
  - FY16 – in process
- XPS
  - FY13 – 2 projects co-funded
  - FY14 – 3 projects co-funded
  - FY15 – 2 projects co-funded
  - FY16 – projected - 2-3 projects
- DMREF
  - FY14 – 0 DMREF projects & 4 SI2 projects co-funded
  - FY15 – 2 DMREF projects & 3 SI2 projects co-funded
  - FY16 – TBD
- CRISP
  - FY15 – 4 projects co-funded
  - FY16 – ~5 projects projected

Additional Software Proposal Stats – Venture, EAGER, RAPID, REU, RIDIR

- Software Venture (Reuse) Fund
  - FY12 – 22 projects co-funded
  - FY13 – 14 projects co-funded
  - FY14 – 28 projects co-funded
  - FY15 – 12 projects co-funded
- EAGER - 3
  - SI2: Software Discovery Index
  - SI2: Crosswalk table across existing metadata schema
  - LWD: Remote sensing + cloud curriculm
  - CDS&E:
    - Improving exascale scientific simulations using in situ machine learning
- CAREER – 1:
  - Transition of research infrastructure from grant-based to long term sustainability
- Ebola RAPIDs - 3
  - Trial software + BW allocation
    - Can allocations be better estimated?
  - Showcase NSF responsiveness to immediate needs
- REUs – 4 (including one interdisciplinary VOSS)
- Workshops (or travel support) – 5
  - CECAM, CCCGrid, SPD, Scientific Software Days, STAR Metrics

EXAMPLE PROJECTS

SSE: A Next-Generation Open-Source Computational Fluid Dynamic Code for Polydisperse Multiphase Flows in Science and Engineering

Alberto Passalacqua, Iowa State University

This project research aims to develop innovative open-source computational tools for the numerical simulation of turbulent multiphase flows based on OpenFOAM® and ported to power Processing Units (GPU) Applications, include nanoparticle formation in chemical reactors for the production of cancer drugs, bubble columns used in the oil and energy industry for the production chemicals, fluidized-bed reactors, with application to bio-mass combustion, hydraulic fracturing flows used to extract natural gas, and dispersion of particles in the atmosphere, environmental, biological and geological flows. The community includes scientists and engineers whose work involves multiphase flows, together with a broader community of students and researchers in the pharmaceutical and environmental science areas. Documentation will be provided of the developed theory and of the software, and by tutorials and demonstration modules that illustrate the application of the software to real-world cases, in particular, in the chemical industry. The developed codes will be publicly released under the GNU General Public License 3, together with their documentation and test cases.

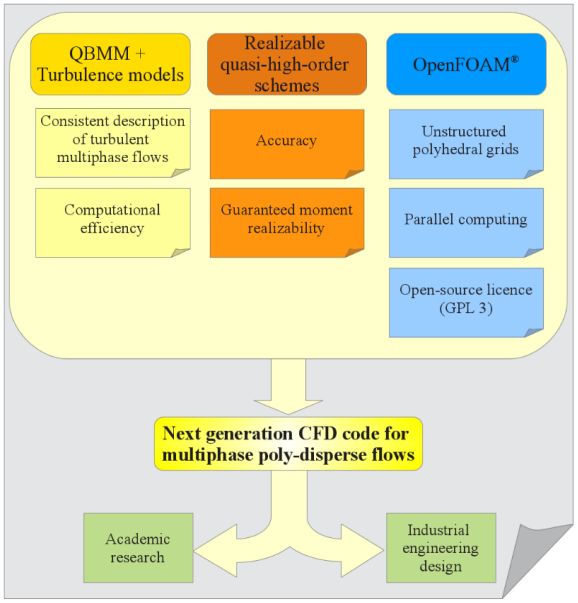

Fig. 4: Schematic representation of the project objectives and of the target audience.

SSE: Adding Research Accounts to the ASSISTments' Platform: Helping Researchers Do Randomized Controlled Studies with Thousands of Students

Neil T. Heffernan, Worcester Polytechnic Institute

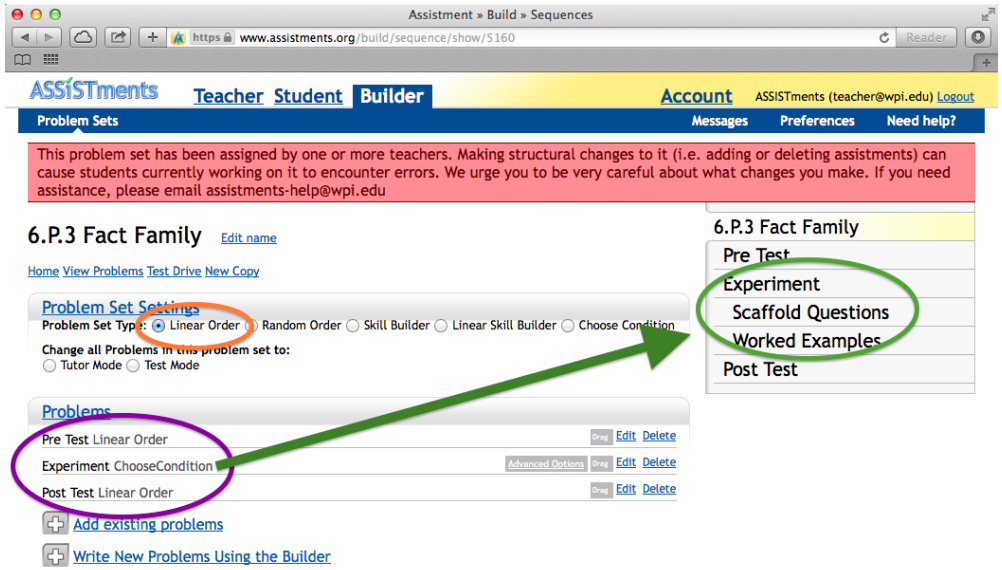

Figure 1: Shows that a researcher can create a problem set. In this case the problem set is of type "Linear Order" (see orange oval) where students will do three sections in a row (see purple oval): first they will do a section called "Pre Test," followed by one called "Experiment," and then finally a section called "Post Test" (the names for sections are created by the experimenter and have no meaning except to help the experimenter to remember his design). The experimental section is of type "ChooseCondition," which means it will randomly pick one of the objects contained within it. The conditions in the section labeled as "Experiment" are called "Scaffold Questions" and "Worked Examples" (see the green oval).

ASSISTments is a free, platform for randomized controlled student-focused trials (RCTs) to help increase the quality, speed, and reliability of K-12 education while not compromising student learning. This project will add Researcher accounts to ASSISTments. Researchers will create their own experiments with IRB approval for release to teachers, and get anonymized data. Its long-term goal is to have a community of hundreds of scientists that use this tool to do their studies.

SI2-SSI: Distributed Workflow Management Research and Software in Support of Science

Ewa Deelman, USC Information Sciences Institute
Miron Livny University of Wisconsin

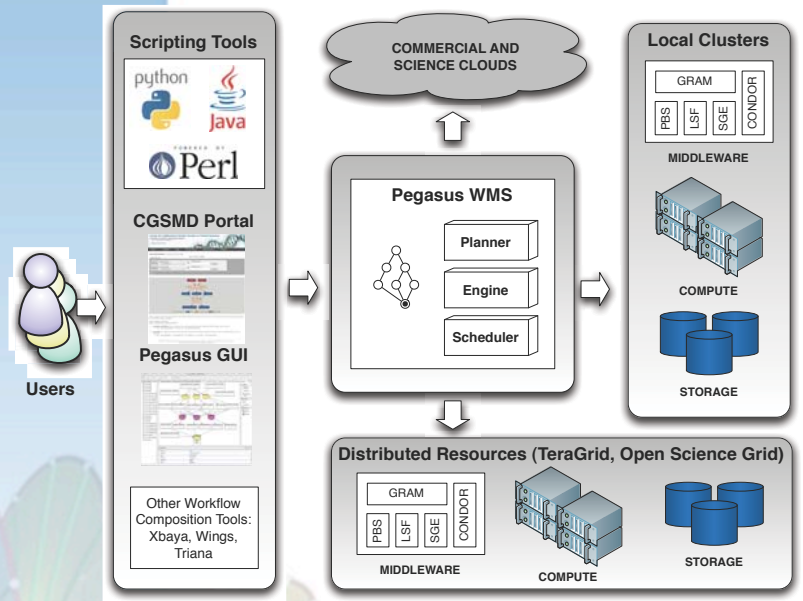

CI tools are needed that translate a high-level description of the computation into a detailed execution plan and reliably execute that plan. Pegasus provides an easy to use, effective, and dependable workflow management system. Pegasus allows users to declara- tively describe their workflow definition, then makes a plan that maps this description onto the available execution resources and executes the plans. This approach is scalable, reliable, and supports applications running on campus resources, clouds, and national cyberinfrastructure. Built on broadly accepted abstractions and a proven computer science framework, Pegasus has been serving scientists from a broad range of disciplines: astronomy, bioinformatics, earthquake science, gravitational wave physics, limnology, and others. In order to sustain scientific methods that rely on Pegasus, this software needs to be sustained.

Instrumental in the data analysis in the recent LIGO announcement!

Science Gateways Software Institute (SGSI)

A science gateway (portal or hub):

- "Synergistic focal points" for communities of scientists
- Provide scalable access to digital resources, facilities, instruments, collaboration
- Integrate assemble complex CI and make them available and *accessible* to users

A gateway institute: Hub for a software ecosystem for developing science gateways

Chemistry and Materials Research Institute (CMRSI)

- A focal point of a sustainable software ecosystem containing reliable, interoperable, verified, and accessible software tools
- Catalyze the application of computation and associated data-centric methods across chemical and materials research.
- Science goal: the development of a quantitative and predictive understanding of materials and chemistry, such as:
  - The computational design of chemicals and materials for specific functions starting from atoms, molecules, or other fundamental building blocks;
  - The prediction of new synthesis pathways;
  - Advancing understanding of how catalysts work;
  - Advancing fundamental understanding of systems far from equilibrium with application to biological systems and the synthesis of soft materials;
  - Enabling meaningful simulation of polymeric and other materials across scales of length and time leading to insights for synthesis and performance;
  - Advancing understanding of quantum dynamics of complex chemical and condensed phase systems; and
  - Understanding macroscopic materials or chemical properties from their atomic or molecular origins.
  - Activities in support of the Materials Genome Initiative.

SOME THOUGHTS ON FUTURE DIRECTIONS

In 2026

- Example societal grand challenges:
  - Disaster scenarios imposed by climate change
  - Sustainable provisioning of food, water and energy
  - Economic shifts imposed by a decentralized world order
  - Education for all
- Research needs:
  - Distributed, dynamic, multi-disciplinary collaborations
  - Lots of research “in the wild” in situ with the object of research
  - Data-driven, integrative over multiple scales, from “dual-use” sources
  - A wide spectrum of stakeholders and participants – from scientists to citizens to industry organizations.
- Research processes will need to be dynamic:
  - Range from explorations to repeatable workflows and back.
  - Research methods – ranging from quantitative to qualitative - will need to be drawn from across disciplines and then integrated.
  - Participants will engage and disengage depending upon the stage of the research.
  - Datasets, instruments and computation will be brought in and utilized as and when needed.

Dynamic CI Processes on Interoperable, Configurable CI

- Toolboxes of composable computational research methods
- Workflows that adapt to data and humans-in-the-loop.
- Use dynamically configurable systems, software and networks (software defined everything, software injection, parameterized components)
- Need security, access, reproducibility and trustworthiness techniques for dynamic situations (human back in the loop?).
- HCI that evolves
- Business models for "on-demand negotiation”
- Learning aimed towards “integrative synthesis” rather than disciplinary depth.

Research Directions

Systems Integration:

- New integration techniques - auto-generation of integration code from interface specifications
- Studies of software engineering methods for s/w integration – DevOps, continuous deployment
- Studies of integrative methods for data science
- Empirical studies on software reuse in science
- Analytical models for understanding/evaluating performance, scalability, security during integration
- Service-based integration of data analytics and HPC system architectures

HCI:

- Search based composition of services
- Human-computer interfaces and interaction design and evaluation during integration - e.g. when “surprise" is a given

SBE:

- Ethnographic studies on how scientists actually work
- Economic and social aspects of reuse
- Economic and social aspects of integration
- Science of team science in dynamic situations

Education:

- Learning theories for "just-in-time" application (e.g. how is novice vs. expert learning different)

Domain science:

- End-to-end composition of models across scales (neuron->cognition, chip->data center)

AN SI2 SHOUT-OUT TO OTHER PROGRAMS

Innovations in Cyberinfrastructure Learning and Workforce Development

Advanced Cyberinfrastructure Division (ACI)
Computer and Information Science & Engineering (CISE)

Sushil K Prasad, sprasad@nsf.gov





National Science Foundation
WHERE DISCOVERIES BEGIN

LWD: The Career Pipeline

- Goal: Build robust careers paths in Cyber-Infrastructure (CI) and Computational and Data-enabled Science and Engineering (CDSE)
- Techniques: Leverage existing programs for early-stage researchers. Develop new programs in areas of need/challenge

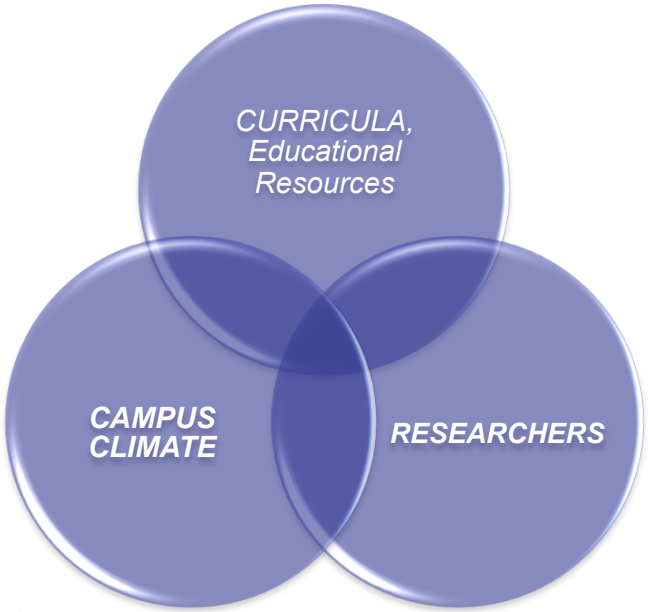

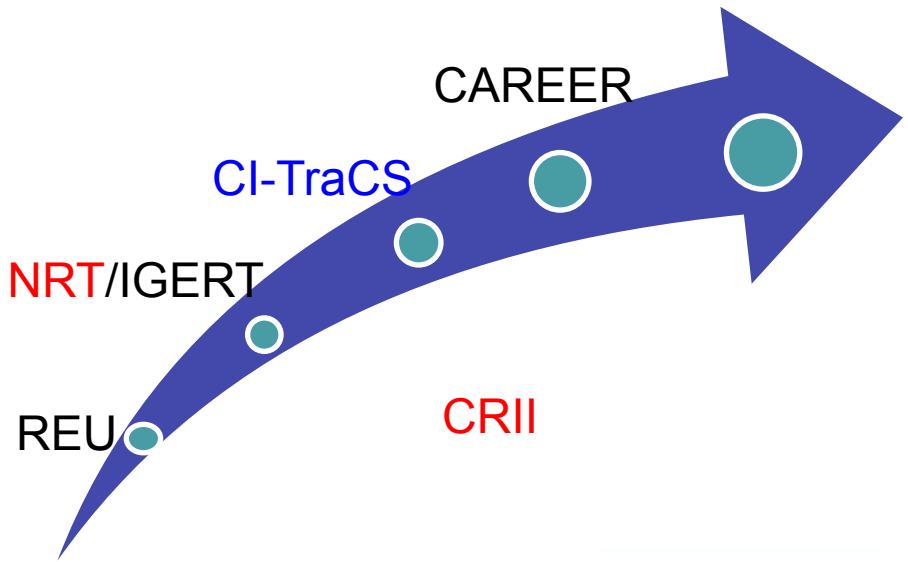

Other Opportunities within ACI

- Fellowships (e.g., Blue Waters)
- CI Engineers
- EAGERs, Workshops – Standards
  - Curriculum Development: Computational Science into domain courses
  - Easy parallel/HPC program development
    - Parallel programming models/paradigms – standards
    - Domain-specific languages/programming environments - standards
  - HPC + Data Analytics: models/paradigms/environments
  - Career paths for non-tenure track CI Scientists and Engineers
- Student travel grants
- REU supplements to existing awards
- *Discuss with ACI Program Officers*

Cybersecurity Innovation for Cyberinfrastructure (CICI) NSF 16-533 - ANIKOLIC@nsf.gov

Activities that impact the security of science, engineering and education environments

Target community is operational cyberinfrastructure

- $7M available. Estimated 7 – 9 awards.
- 2 Areas (due April 19th):
  - Secure and Resilient Architecture - $1M awards
  - Regional Cybersecurity Collaboration - $500K awards
- Sample Topic Areas of Interest:
  - Secure interoperability of cloud/campus
  - Security Metrics
  - DDoS Defense
  - Shared assessments and security design reviews

Secure and Trustworthy Cyberspace (SaTC)

- Cross Directorate Program
- Aims to support fundamental scientific advances and technologies to protect cyber-systems from malicious behavior, while preserving privacy and promoting usability.
- Develop the foundations for engineering systems inherently resistant to malicious cyber disruption
- Cybersecurity is a *multi-dimensional problem*, involving both the strength of security technologies and variability of human behavior.
- Encourage and incentivize socially responsible and safe behavior by individuals and organizations

SaTC: Transition to Practice (TTP)

- Supports later stage activities in the research and development lifecycle such as prototyping and experimental deployment
- Review Criteria:
  - Impact on deployed environment
  - Value in terms of needed capability and potential impact across the NSF community
  - Feasibility, utility, and interoperability in operation
  - Project plan including goals, milestones, demonstration and evaluation
  - Tangible metrics to evaluate effectiveness of capabilities developed
- Due: Sept 2016 (Medium) and Nov 2016 (Small)
- Funding: Small up to $500K; Medium up to $1.2M

Center for Trustworthy Scientific Cyberinfrastructure (CTSC)

- **Mission**: Establish a coherent cybersecurity ecosystem for NSF computational science and engineering, while allowing projects to focus on their science endeavors.
- **Trustedci.org -** webinars, project documents, best practices, online free training
- FREE software assurance code reviews!
- FREE security architecture design reviews!

CENTER FOR APPLIED CYBERSECURITY RESEARCH
INDIANA UNIVERSITY
Pervasive Technology Institute

NSF

NCSA

PITTSBURGH SUPERCOMPUTING CENTER

UNIVERSITY of WISCONSIN MILWAUKEE

THE UNIVERSITY of WISCONSIN MADISON

XD Network of Computational Resources and Services - REIGENMA@nsf.gov

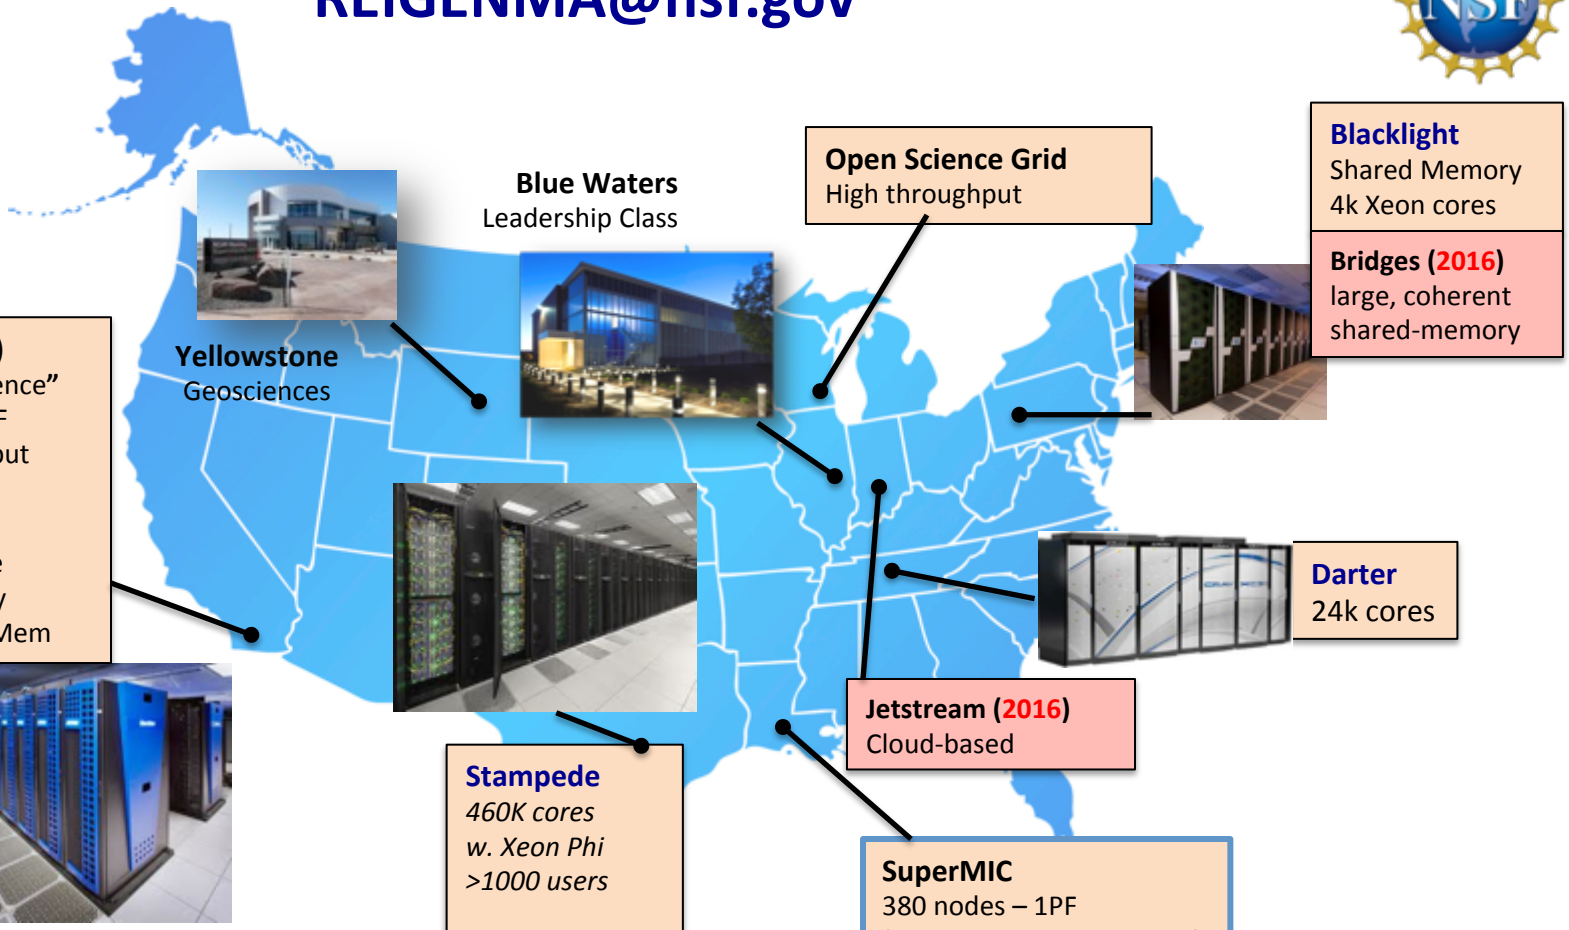

Blue Waters
Leadership Class

Open Science Grid
High throughput

Blacklight
Shared Memory
4k Xeon cores

Bridges (2016)
large, coherent
shared-memory

Yellowstone
Geosciences

Comet (new)
"Long Tail Science"
47k cores/2 PF
High throughput

Gordon
Data intensive
64 TB memory
300 TB Flash Mem

Darter
24k cores

Jetstream (2016)
Cloud-based

Stampede
460K cores
w. Xeon Phi
>1000 users

Wrangler (new)
Date Analytics

Maverick
Visualization
Data Analytics

SuperMIC
380 nodes – 1PF
(Ivy bridge, Xeon Phi, GPU)

Coordination though XSEDE

- Resource Allocation
- Advanced User Support
- Education and Outreach
- Common Software

XSEDE: Supporting collaborative expertise, services, education and outreach in an overall ecosystem - REIGENMA@nsf.gov

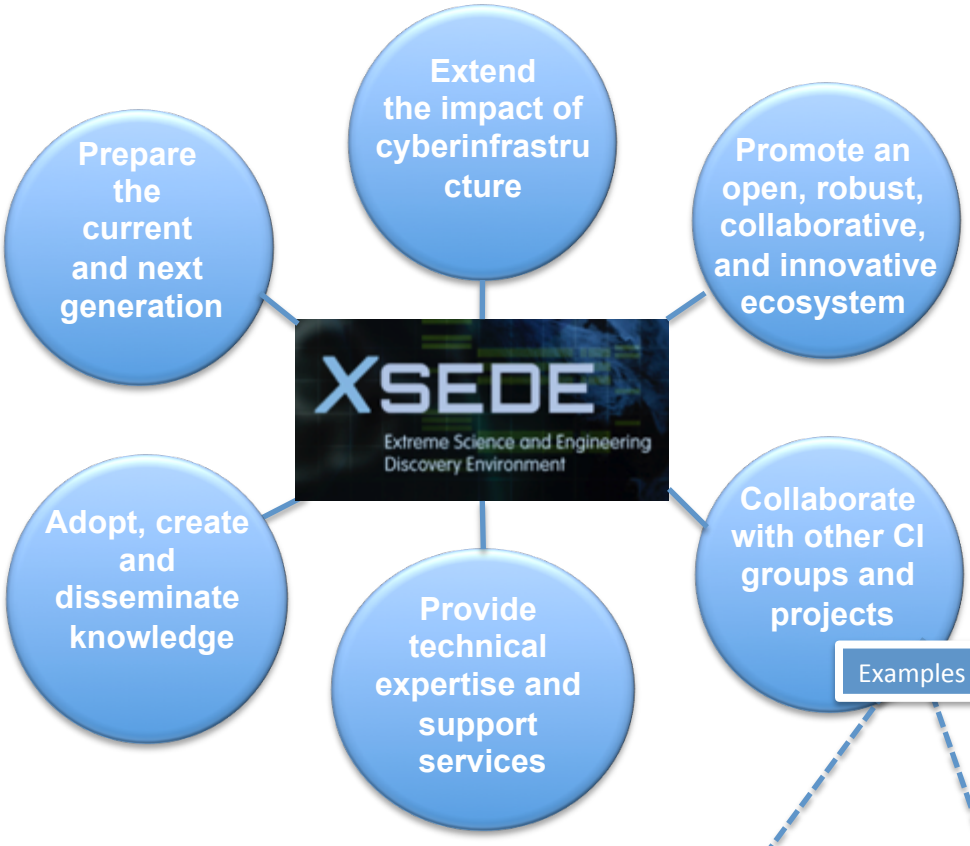

XDMoD
METRICS ON DEMAND

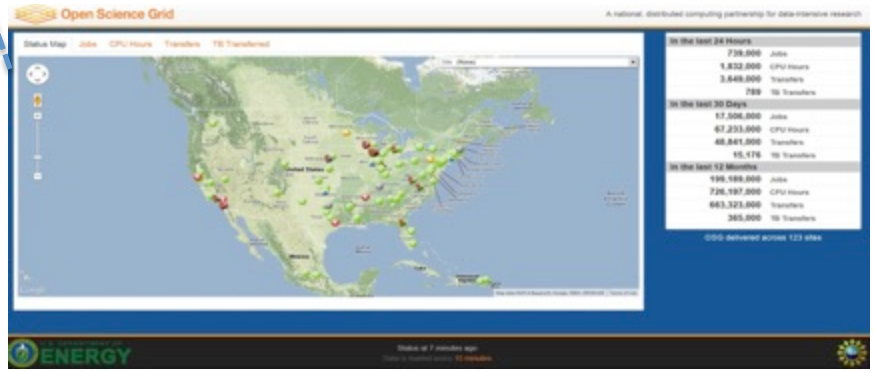

Petascale Computing Resource Allocations (PRAC) - edwalker@nsf.gov

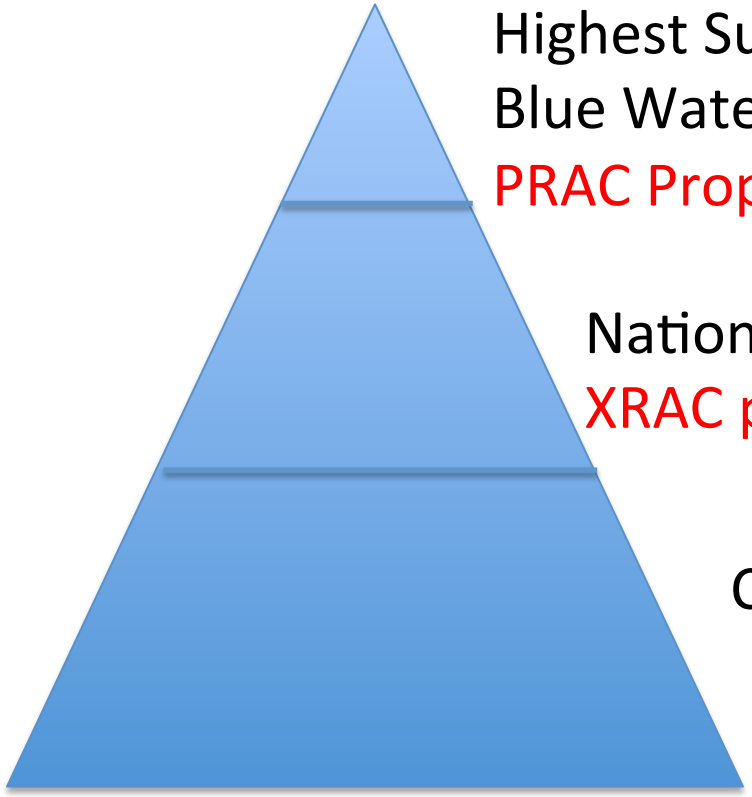

Highest Sustained Performance:
Blue Waters machine
PRAC Proposals to NSF <= focus

National HPC systems : XSEDE Virtual Org.
XRAC proposals to xsede.org

Campus machines

FY 16 Solicitation (NSF 14-518)

- Request computing allocations on Blue Waters
- You must demonstrate that
  - You have a breakthrough science or engineering research problem that requires petascale computing capabilities
  - Your code requires and can effectively exploit the capabilities offered by Blue Waters.
- Proposals from or including junior researchers are encouraged
- Next PRAC deadline is April 4, 2016

Networking Programs in CISE/ACI – kthompso@nsf.gov

- Networking as a fundamental layer and underpinning of CI
- CC* (Campus Cyberinfrastructure – Data, Networking, and Innovation)
  - Campus networking upgrade (re-design to scienceDMZ at campus border and 10/100Gbps) and innovation program
- IRNC – International R&E Network Connections
  - Scientific discovery as a global collaborative endeavor
  - Provide network connections linking U.S. research with peer networks in other parts of the world
  - Supports all R&E US data flows (not just NSF-funded)
  - Stimulate the deployment and operational understanding of emerging network technology and best practices
  - 10Gbps and 100Gbps connections

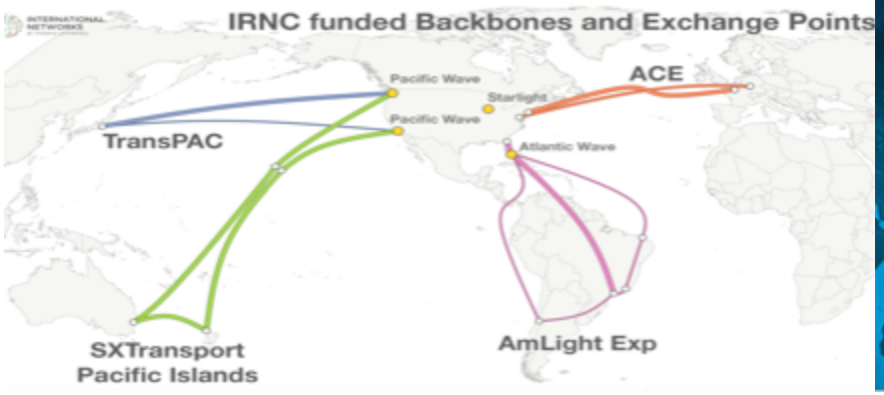

Data Infrastructure Building Blocks (DIBBS)

NSF Solicitation 16-530
Webinar -- March 3, 2016

Amy Walton, Program Director
Advanced Cyberinfrastructure
National Science Foundation

Questions: DIBBsQueries@nsf.gov

The DIBBs Solicitation

- Seeks proposals that develop robust, scalable, well-designed cyberinfrastructure (the 'building blocks') contributing to future discovery and innovation across disciplines
  - Guided by science and engineering research priorities
  - Built upon recognized community data collections
  - Result in clear, tangible cyberinfrastructure products
  - Implemented through collaborations between cyberinfrastructure experts and specific science and engineering research communities
- Focus areas:
  - Early Implementation Awards: up to 6 awards, each up to $4M total for up to 5 years
  - Pilot Demonstrations: up to 5 awards, each up to $500K total for up to 3 years
- Webinar March 16th

Critical Resilient *Interdependent* Infrastructure Systems and Processes (CRISP)

Research goals:

1. Improved resilience, performance, readiness, innovation
2. Multidisciplinary, multilevel data-driven models of interdependent infrastructures, for analysis, prediction, real-time control and adaptation in response to system and policy changes, while addressing organizational, social, psychological, legal, economic, technical obstacles
3. Validation, verification of software integration, software engineering processes

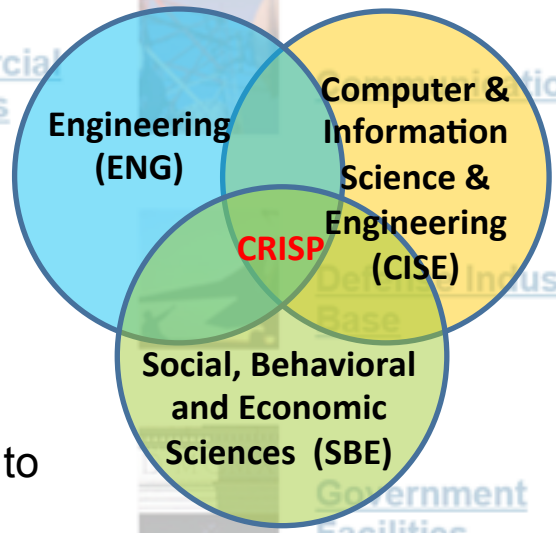

At least one *engineer, computer, information or computational scientist, and a social, economic or behavioral scientist as a PI*

Type 1 Projects: Theory, modeling, data collection, metrics, problem case studies.
Type 2 Projects: Empirical analysis and studies, process design and simulation

Anticipated Funding: $26.5M

- *Type 1:* 2 years, $500K
- *Type 2:* 3-4 years, $1M-$2.5

Solicitation URL: http://www.nsf.gov/pubs/2016/nsf16519/nsf16519.htm

Full Proposal Due Date: March 9, 2015 (4:59 pm proposer's local time)

ACI Contact: rramnath@nsf.gov

Proposals of ACI Interest in CRISP

- With significant computing relevance, in addition to the ENG and SBE relevance,
- Lead to a path where the physical infrastructure becomes a long-lasting cyber infrastructure (in a sense similar to a facility)
- Require modelling and simulation at scale to need our current science CI, like XSEDE, and create and sustain software for this
- Develop large scale data assets,
- Are not so focused on disaster scenarios; rather propose innovative new uses of the physical infrastructure.
- When computing is the major component of the proposed work, better integrate SBE components.

ACI EPSCoR investments cover all program areas

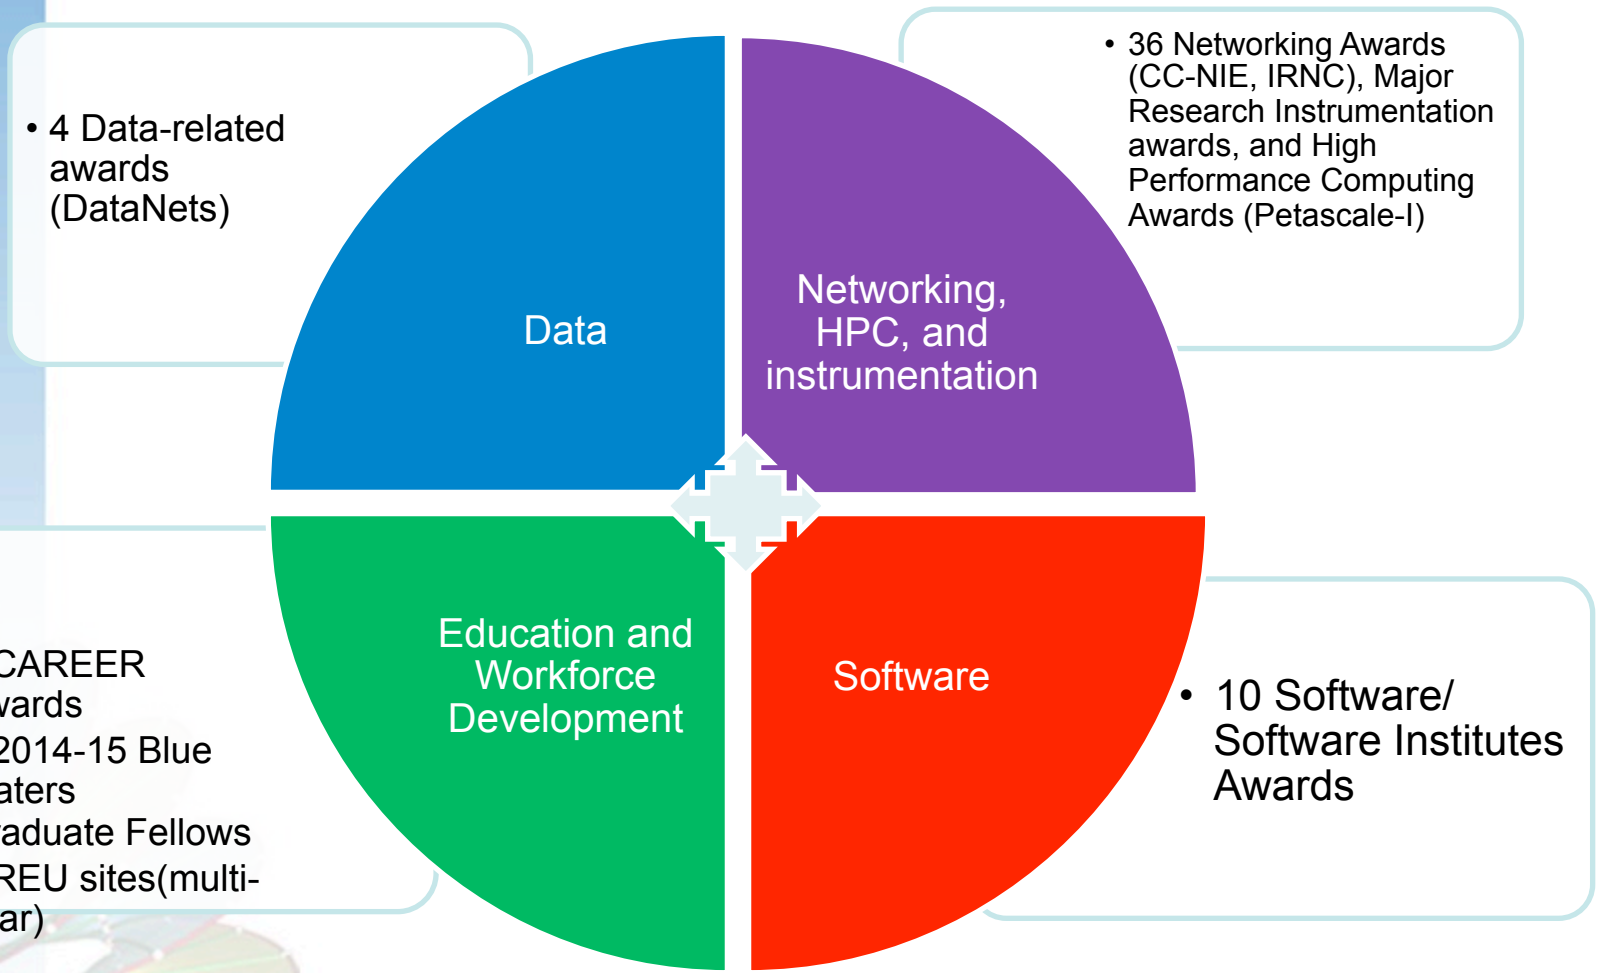

CREDITS

2016 PI Meeting Organizers – Thanks!

- Frank Timmes, Arizona State University
- Stan Ahalt, RENCI
- Matthew Turk, University of Illinois
- Shaowen Wang, University of Illinois
- Karl Gustafson, RENCI
- Ray Idaszak, RENCI
- Chris Lenhardt, RENCI
- Richard Brower, Boston University

2017 PI Meeting Organizers – Welcome!

- Ganesh Gopalakrishnan, University of Utah
- Jonathan Hauenstein, University of Notre Dame
- Matt Knepley, University of Chicago
- Yung-Hsiang Lu, Purdue University
- Amarda Shehu, George Mason University
- Matthew Turk, University of Illinois
- Nancy Wilkins-Diehr, San Diego Supercomputing Center
- Kyle Niemeyer, Oregon State University

Daniel S. Katz – Bon Voyage!

Snow – NO THANKS!

LET THE GAMES BEGIN!