



High Performance Computing System Acquisition: Building a More Inclusive Computing Environment for Science and Engineering (NSF 13-528)

Barry I. Schneider

Irene M. Qualters

&

Bob Chadduck

Program Directors
Division of Advanced
Cyberinfrastructure



Preliminaries

- **OCI now part of CISE Directorate**

- CISE:Division of Advanced Cyberinfrastructure (ACI)
- OCI website will transition
- Details in progress
- Small diversion - A word about open access- The John Holdren memo- Please see <http://www.whitehouse.gov/blog/2013/02/22/expanding-public-access-results-federally-funded-research>





Preliminaries

- **Revised NSF Proposal & Award Policies & Procedures Guide (NSF 13- 1) – for proposals submitted after Jan 14, 2013**

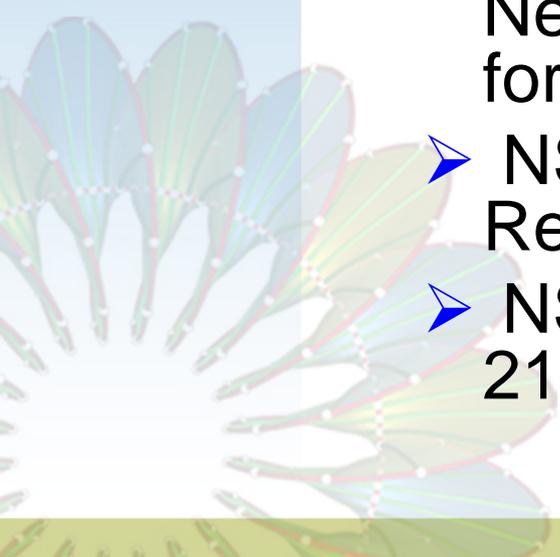
- *While the two merit review criteria remain unchanged (Intellectual Merit and Broader Impacts), **guidance has been provided to clarify and improve the function of the criteria.** Changes will affect the project summary and project description sections of proposals. Annual and final reports also will be affected.*
- *GPG changes also include: Biosketches; Indirects; Budget (senior personnel at zero dollars); Facilities, Equipment and other Resources*
- **Reminder that cost-sharing is prohibited.**
Facilities, Equipment and other Resources

*.....Proposers should include an aggregated description of the internal and external resources (both physical and personnel) that the organization and its collaborators will provide to the project, should it be funded. Such information must be provided in this section, in lieu of other parts of the proposal (e.g., budget justification, project description). **The description should be narrative in nature and must not include any quantifiable financial information.***



Preliminaries

- Other ACI solicitations in the same time frame as this Webinar
 - NSF 13-530 Campus Cyberinfrastructure - Network Infrastructure and Engineering Program (Apr 3, 2013)
 - NSF 13-525 Software Infrastructure for Sustained Innovation - SSE & SSI (Mar 19, 2013)
 - NSF 13-520 Research Coordination Networks (Mar 4 for SEES, June 14, 2013 for UB)
 - NSF 12-516 Industry/University Cooperative Research Centers Program (Mar 5, 2013)
 - NSF 12-609 Computing Education for the 21st Century (Mar 13, 2013)





ACI contribution to NSF's Advanced Computational Infrastructure

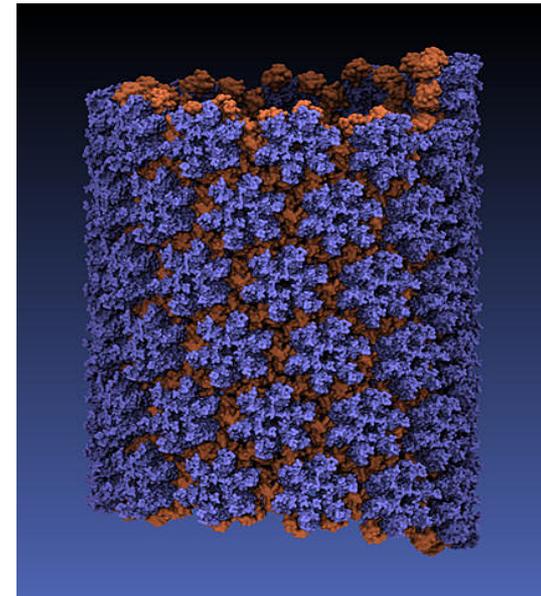
Vision: Sustained and comprehensive portfolio of advanced computing infrastructure, programs and other resources to facilitate cutting-edge foundational research in Computational and Data Enabled Science and Engineering (CDS&E) and its applications to all disciplines.

Model spread of H1N1 MD modeling of HIV Capsid

Picture will be added after credit is obtained

Type I Supernova

Picture will be added after credit is obtained



Credit: Courtesy of Klaus Schulten, University of Illinois at Urbana-Champaign Theoretical and Computational Biophysics Group/Beckman Institute; Angela Gronenborn and Peijun Zhang, University of Pittsburgh School of Medicine Center for HIV Protein Interactions/Department of Structural Biology.



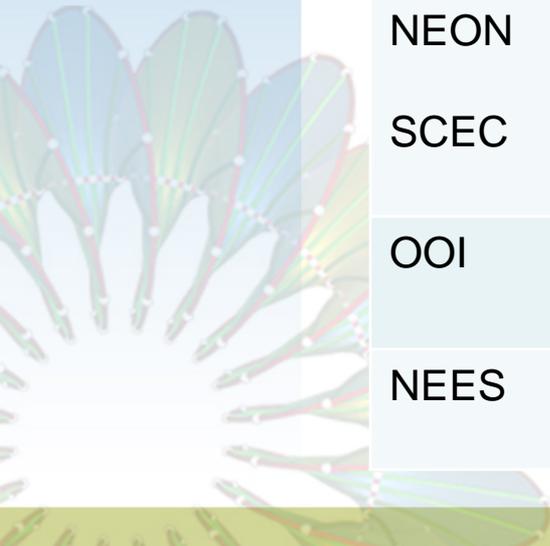
The Big Picture: NSF Long-Range Goals for a Comprehensive Cyberinfrastructure

- 1) Enable computational and data driven science research and education across a broad set of disciplines and at many levels- *Leverage and invest in collaborative flexible “fabrics” dynamically connecting scientific communities with computational resources and services at all scales (campus, regional, national, international)*
- 2) Support a cyberinfrastructure that is capable of managing a distributed, heterogeneous set of digital resources to support 1)
- 3) Provide the flexibility to accommodate computational and data systems as well as instruments that will change in time - *Anticipate and invest in diverse and innovative national scale shared resources, outreach and education complementing campus and other national investments*
- 4) **Goals:**
 - ✓ **Maintain the HPC environment for the power users - Deep**
 - ✓ **Bring in new communities requiring advanced digital services to enable their productivity – Wide**
- 5) Provide capabilities for campus, regional and national resources to interact transparently – **Even wider and broader**
- 6) Collaborate easily with other CI's here and abroad

NSF Cyberinfrastructure



System	URL
XSEDE	http://www.xsede.org
Blue Waters	http://www.ncsa.illinois.edu/BlueWaters
Open Science Grid	http://www.opensciencegrid.edu
FutureGrid	http://www.futuregrid.org
iPlant	http://www.iplantcollaborative.org
NanoHub	http://www.nanohub.org
NEON	http://www.neoninc.org
SCEC	http://www.scec.org
OOI	http://www.oceanobservatories.org
NEES	http://www.nees.org



Recent NSF Investments in Computational CI

Tier I

- Blue Waters – **Sustained PF** (peak >13 PF)
 - Targeting the most challenging computational research
 - Contains a well balanced mix of CPU and GPU processors, a very fast interconnect and large storage
 - 10's of projects annually
- Yellowstone – peak of 1.5 PF –
 - Targetting GEO community applications –
 - Resources managed by NCAR

Tier II

- eXtreme Digital Program (XD)
 - Targetting the widest diversity of computational and data enabled science – 1000s projects annually
 - resources managed by XSEDE project
 - Stampede – 2 PF multicore base moving to 15PF by 2015 using novel processors
 - List of all resources – see www.xsede.org for details



Current Allocatable XD Compute Resources

- Kraken @ NICS
 - 1.2 PF Cray XT5
- Stampede @ TACC
 - 2 PF Sandy Bridge CPU
 - Intel MiC added to 10-15 PF by 2015
- Gordon @ SDSC
 - 341 TF Appro Distributed SMP cluster
- Lonestar (4) @ TACC
 - 302 TF Dell Cluster
- Trestles @ SDSC
 - 100TF Appro Cluster
- Blacklight @ PSC
 - 36 TF SGI UV (2 x 16TB shared memory SMP)
- Keeneland @ NICS
 - 615 TF NVIDIA GPU Cluster with Sandy Bridge CPU
- Open Science Grid @ Wisconsin
 - High Throughput
- Steele/Condor Pool @ Purdue
 - 67 TF Dell Cluster
 - High Throughput
- FutureGrid @ Indiana
 - Experimental test bed for HPC applications, cloud computing and virtual machines

<https://www.xsede.org/web/xup/resource-monitor>



Current Allocatable XD Visualization and Data Resources

- Visualization

- Nautilus @ UTK
 - 8.2 TF SGI/NVIDIA SMP
 - 960 TB disk
- Longhorn @ TACC
 - 20.7 TF Dell/NVIDIA cluster
 - 18.7 TB disk

https://www.xsede.org/web/xup/resource-monitor#advanced_vis_systems

- Storage

- Data Capacitor @ Indiana
 - 535 TB Lustre WAN filesystem
- Data Replication Service
 - 2.8 PB iRODS distributed storage
- HPSS @ NICS
 - 17PB tape& disk
- MSS @ NCSA
 - 10 PB tape
- SuperCell @ PSC
 - 4PB disk
- Ranch @ TACC
 - 60 PB tape & disk
- NCAR GLADE

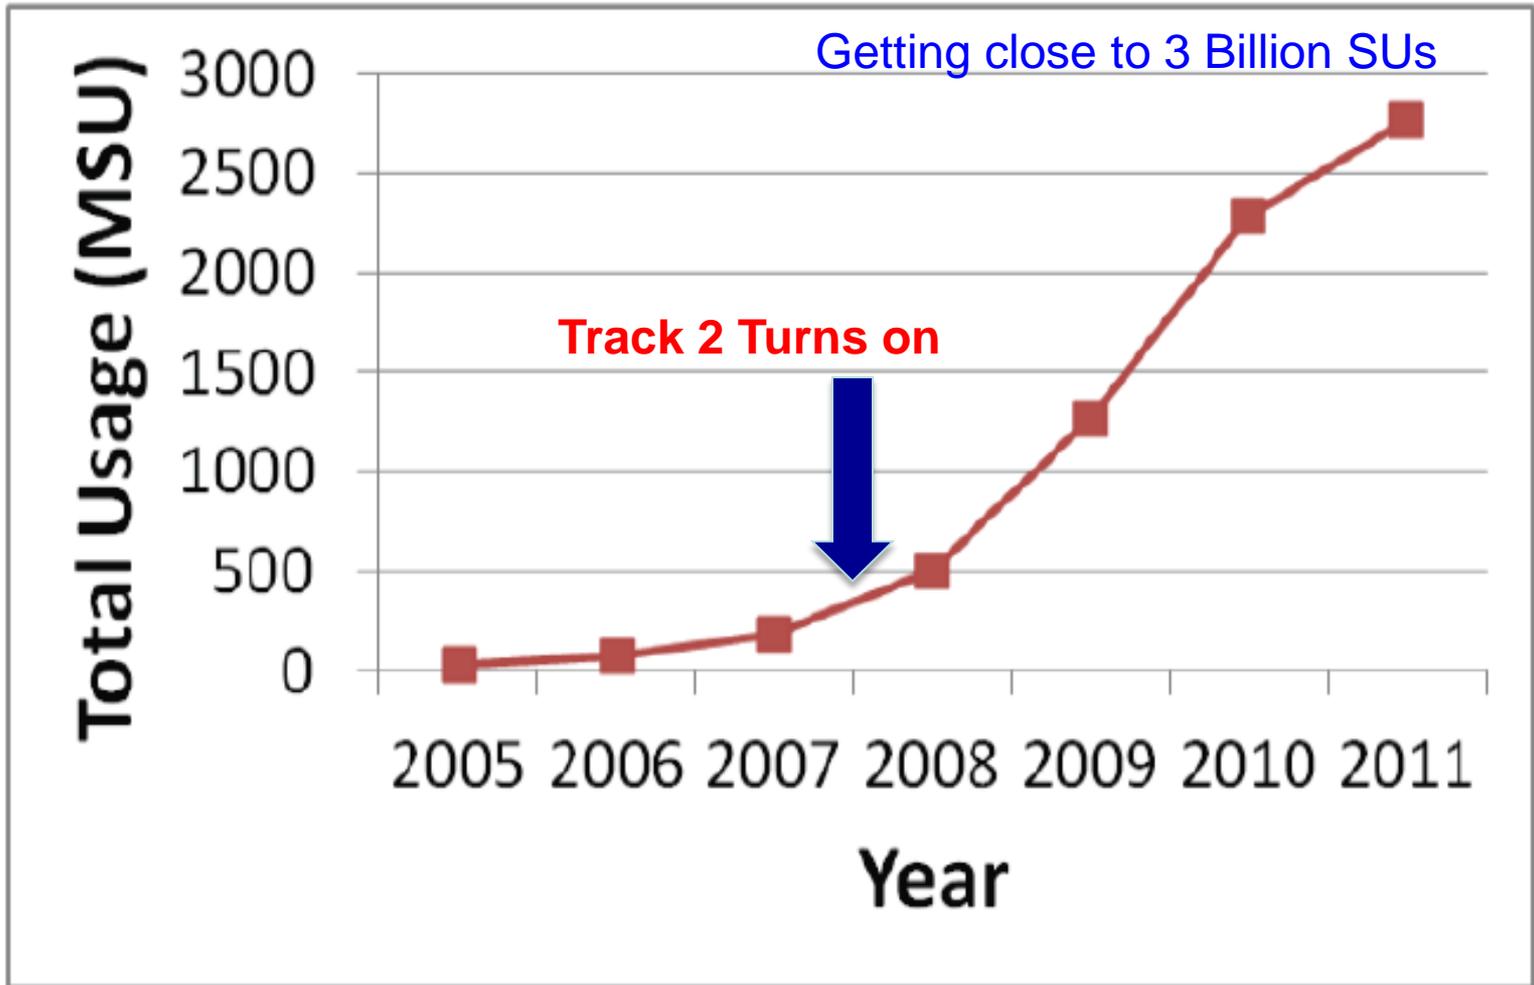
Read Only

Generally allocable

https://www.xsede.org/web/xup/resource-monitor#storage_systems



Total Usage of Service Units





Solicitation: Context, Motivation and Details

- Aligned with ACI Strategic Plan
 - Position and support the entire spectrum of NSF-funded communities - **Diversification**
 - Promote a more comprehensive and balanced portfolio
 - Include capabilities suitable for addressing emerging computationally and/or data intense scientific and engineering research topics, and workflows
 - Support multidisciplinary computational AND data-enabled science and engineering
 - Enlarge the horizon to include research communities that are not optimally served by current XD or Blue Waters and who would benefit from advanced computational capabilities at the national level





Solicitation: Context, Motivation and Details

- **Complementing Current Investments**
 - Facilitate transitions from local to national CI via the use of virtual machines
 - Support novel approaches to the utilization of cloud resources within the current XD environment
 - Expand the range of data-intensive and/or computationally-challenging S&E applications that can be tackled with current XD resources
 - Provide reliable high throughput resources not available in current XD investments
 - Provide a useful interactive environment for users needing to develop and debug codes using hundreds of cores or for scientific workflows/gateways requiring highly responsive computation





Solicitation: Details

- Total funds available; up to \$30M
- Submission date is April 15, 2013
- Project description is limited to 30 pages
- An organization may submit only one proposal but may be a sub-awardee on other proposals responding to this solicitation
- Anticipate two to three awards; the actual outcome will depend on the proposals submitted
 - ✓ **Novel ideas are encouraged and there is flexibility in the final funding profile**
- The computational resource awards will be capped at \$12,000,000 each and the data resource award at \$6,000,000
 - ✓ Proposals above the caps will be RWR
- Combined awards at \$18,000,000 are within scope
- Project duration up to 4 years
- Acquisition and deployment of final system by end of FY 2015 – **but**
 - ✓ Expect a high degree of stability and usability by January, 2015
 - ✓ Earlier deployment is encouraged, where possible
 - ✓ An appropriate use of these funds not prohibited after full deployment
- Expansion of a current resource is allowed
- Operating and maintenance costs provided in a separate budget action up to 20% of initial cost
- Proposal should address innovative capability
- Convincing evidence presented that the resource will perform as described





Solicitation: Storage-specific Details

- Acquisition and support of a large scale instrument that will have the capability of storing, sustainably accessing, analyzing, disseminating, securing and migrating data across the NSF cyberinfrastructure.
- Data may come from computation, instruments/sensors, or other sources but needs to be available to the scientific community independent of geographical location.
- **Not** designed to address long term archival storage issues





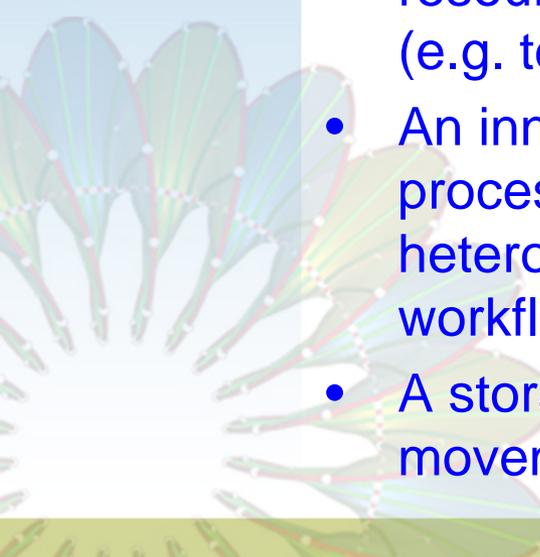
Considerations for Competitive Proposals

- Critically: Complement existing XD capabilities
- Provide computational resources attuned to less traditional computational science communities
- Provide resources for innovative and reliable dynamic workflows
- Enable transitions from local to national environments via the use of virtual machines
- Consider cost efficient cloud computing 1) for new modes of computationally intensive research or 2) to expand current resources as required by changing demands
- Provide resources for emerging data intensive and/or computationally-challenging science and engineering
- Expand high-throughput computing capabilities
- Provide capabilities for researchers needing 100's of cores or for scientific workflows/gateways requiring highly responsive and interactive environments
- Efficiently provide a high degree of stability and usability by January, 2015 or earlier



Potential New capabilities

- A novel data-intensive, high-performance computing capability suitable for new scientific communities
- An innovative, power efficient, highly usable, high-performance computing capability with sustained, high throughput performance
- An innovative capability that expands XD boundaries, for example, by the introduction of domain specific capabilities, high throughput capabilities, time-sharing, efficient use of virtualization and/or clouds
- Dynamic interactive research workflows across XD resources or between other cyberinfrastructure resources (e.g. telescopes, sequencers) and XD resources.
- An innovative high performance approach to "end to end" processing and sustainable throughput of ultra-large, heterogeneous data collections for complex, real time workflows
- A storage resource designed to enable rapid access and movement of data across the NSF CI





Solicitation Specific Review Criteria

- **Resource Specification**
- **Performance and Innovation for Science and Engineering Applications**
- **Reliability and Usability**
- **Implementation, Project Management, and Risk Mitigation**
- **Quality of the Physical Infrastructure**
- **Plan for Operations, including User Support and Training**
- **Broadening Participation**

