

# TeraGrid Overview, Integration and Structure

Rob Pennington

Deputy Director NCSA

March 4, 2008

[robp@ncsa.uiuc.edu](mailto:robp@ncsa.uiuc.edu)



TeraGrid™

# The NSF TeraGrid

- TeraGrid Background and History
  - Supercomputing Centers, PACI, CORE, DTF, ETF, TeraGrid
- NSF Cyberinfrastructure Strategy
  - CI Vision document
- Current Organization, Operation and Status
  - Governance:
  - Grid Infrastructure Group (GIG), Resource Providers (RPs), and Science Gateways
  - Working groups and RATs
  - Funding, organizational reflection
  - Requirements and commitments to participate
  - Meetings and reporting
  - Public presence
  - Advisory committees, review committees and oversight
- Future of TeraGrid Planning



# Background Reports on Cyberinfrastructure

**ESTABLISHING  
A PETASCALE  
COLLABORATORY  
FOR THE GEOSCIENCES**

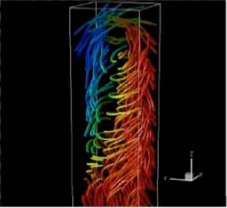


Geoscience Education  
and Cyberinf...

**Building a  
Cyberinfrastructure  
for the Biological  
Sciences (CIBIO)**

2005 and Beyond:  
A Roadmap for Consolidation and Exponentiation

**Materials Research  
Cyberscience  
enabled by  
Cyberinfrastructure**



Authors:  
Mark A. Novotny, Mississippi State University  
David Ceperley, University of Illinois, Urbana-Champaign  
Chakram S. Jayanthi, University of Louisville  
Richard M. Martin, University of Illinois, Urbana-Champaign

November 7, 2004

The Draft Report of the  
American Council of Learned Societies'  
Commission on Cyberinfrastructure for  
Humanities and Social Sciences  
(for public comment)

Contributing Members:

Paul Coates Professor & Professor of Economics University of Michigan	Paul B. Kaufman Professor, Swarthmore College
Mark France Assistant Professor and Chair Art History, Northwestern University	James Madison John Sorenson Bruce Professor English, University of Virginia
Walter Griesel Director, Center for Spatially Integrated Social Research	Ray Rosenzweig Mark and Barbara Fried Professor of History & New York
Robert Professor, Geography University of California, Santa Barbara	Walter D. Rouse Director, Center for History & New Media George Mason University
Margaret Stebbins Assistant Professor, School of Information University of Michigan	John Unsworth (Chair) Dean and Professor Graduate School of Library and Information Science
Charles Stacey Vice President and Chief Information Officer Rice University	Robert Zuckerman Professor, School of Religion Director, Technological Research Collection University of Southern California

Editor:  
Mark W. Williams  
Assistant Chair for Publications and Communications  
Graduate School of Library and Information Science  
University of Illinois, Urbana-Champaign

Executive Advisory to the Commission:

Don Adams Assistant Director for Publications and Communications Library of Congress	Deanna Marston Assistant Librarian for Library Services Library of Congress
Scott Brink Senior NSF/NSB Advisor National Science Foundation Washington, DC	Alvin Neff Director of Programs Council on Library and Information Research Washington, DC
David Foray Director for National Information	Steve Whalley, Vice President American Council of Learned Societies

**Computing, in  
technology**


**GRAND CHALLENGES:  
Science, Engineering, and  
Societal Advances  
Requiring Networking and  
Information Technology  
Research and Development**

Interagency Working Group on  
Information Technology  
Research and Development

Second Printing—MARCH 2004

**Computation As a  
Tool for Discovery  
in Physics**

A Report to the  
National Science Foundation  
by the Steering Committee on  
Computational Physics



**Simulation Based Engineering Science**

A Report on a Workshop  
Held Under the Auspices of the  
National Science Foundation  
April 15-16, 2004

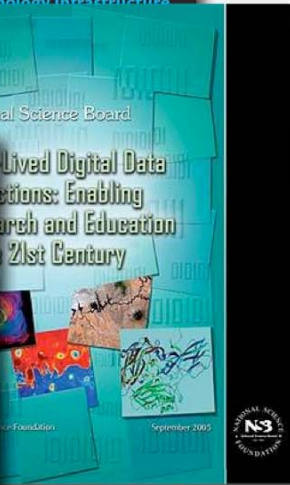


**Trends in  
Information Technology Infrastructure**

**Lived Digital Data  
Environments: Enabling  
Research and Education  
in the 21st Century**

National Science Board

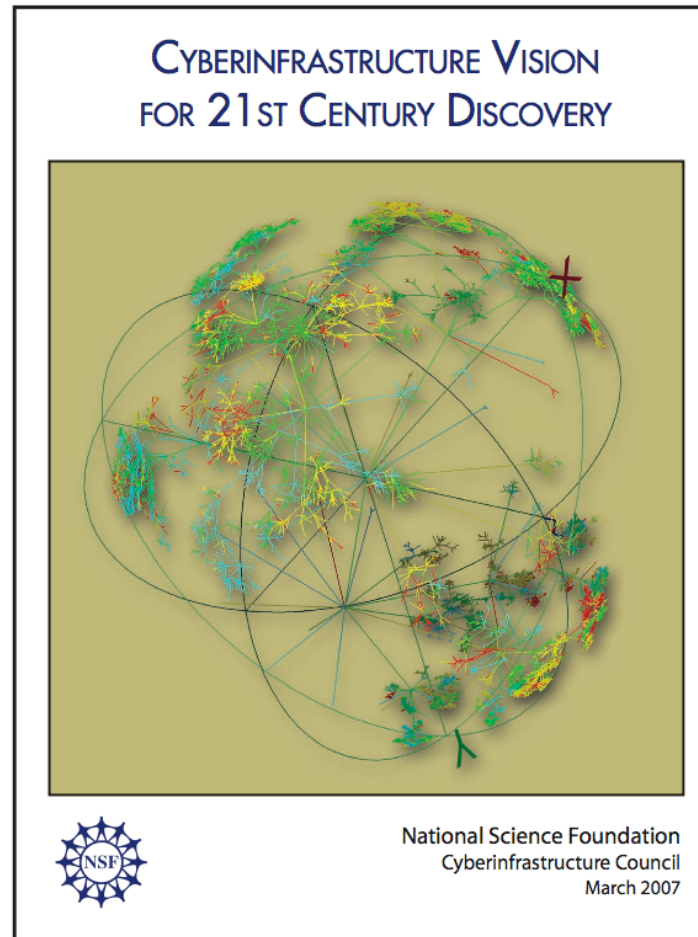
September 2005



# NSF Cyberinfrastructure Strategy

High Performance  
Computing

Data, Data Analysis  
and Visualization



Virtual Organizations  
For Distributed  
Communities

Learning and  
Workforce  
Development



# Vision of TeraGrid

- **Vision**

- Large-scale operational cyberinfrastructure
  - Provides very high-end computational capabilities for leading-edge research in the national open science community
- Reach out to new users and communities
  - Broaden the impact of cyberinfrastructure in research, education, and society
- Collaboration of over a dozen organizations and NSF awards
  - Working together to provide collective services that cannot be provided by individual institutions

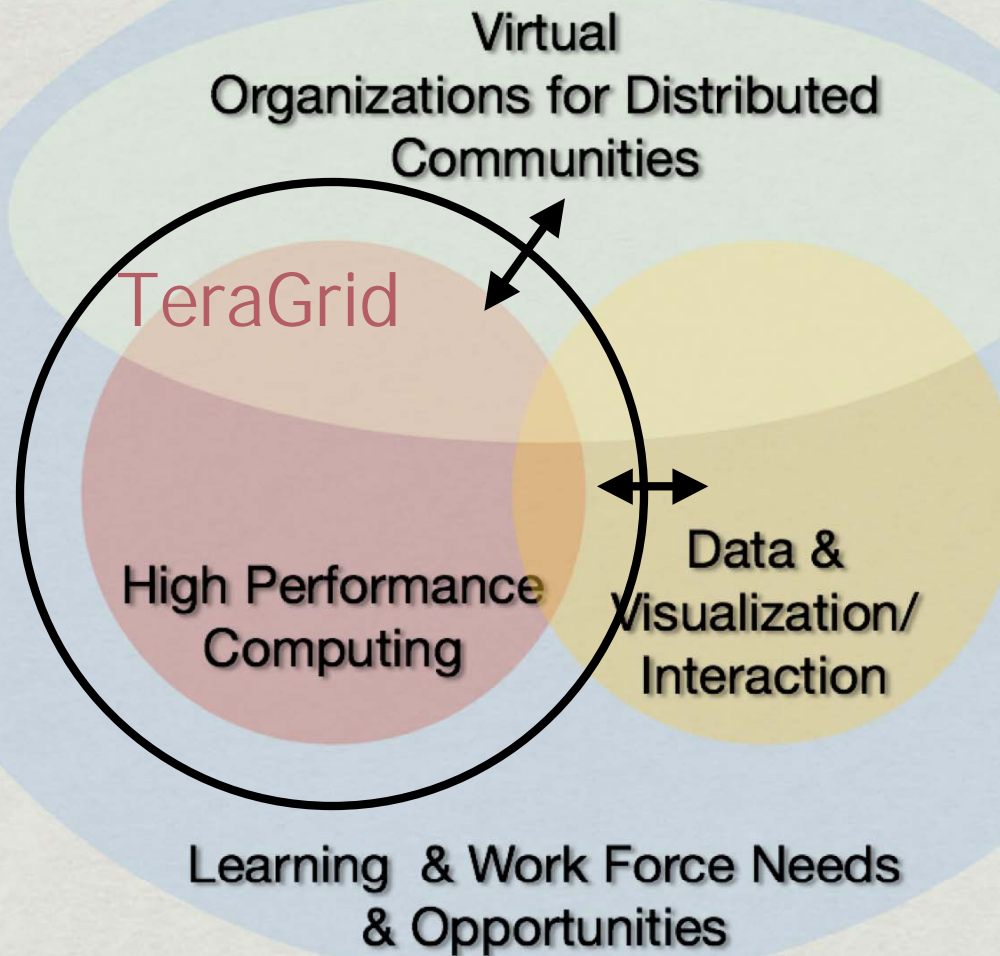
- **High level project strategic objectives**

- Enabling science that could not be done without TeraGrid
- Broadening the user base
- Simplify use of the resources
- Improve operations
- Enabling connections to external resources




# TeraGrid in Context

## CI Vision for 21st Century Discovery



# TeraGrid Timeline



- **Distributed Terascale Facility (DTF)**

- Design/Deploy
- UC/ANL, Caltech, NCSA, SDSC
- Intel 64bit - Homogeneous

- **Extended Terascale Facility (ETF)**

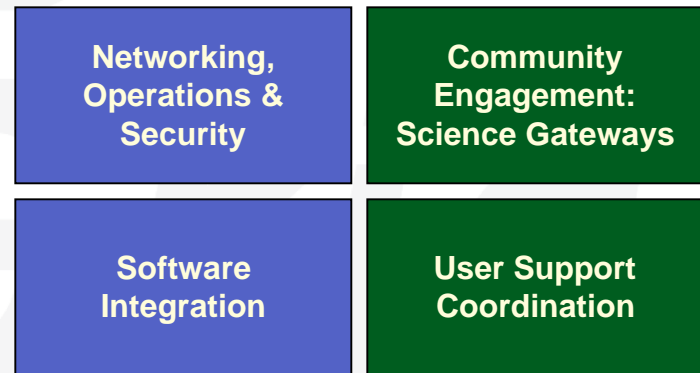
- Add Heterogeneity with PSC
- Add IBM SMP, HP

- **Terascale Expansion Program (TEP)**

- Add TACC, PU, IU, ORNL
- Add IA32

- **FY05-09: Operations & Science Outreach, HPC Operations**


- FY05-6 Add XT3, SGI SMP, BG/L
- FY08 Add NICS
- FY09 Add ?
- 8 RP Sites => 11 RP sites
- Grid Infrastructure Group



# System Resource Allocations Process

- **National peer-review process**
  - Each site provide forecast for resources that will be available
    - Funding for operation of resources by NSF
  - Experienced peers allocate computational and data resources
    - Medium scale and large scale allocations
  - Panel may recommend allocation of advanced direct support services
- **Process managed by TeraGrid**
  - GIG and RP Participants in reviews
  - Separate award to manage shared responsibilities
    - TACC: Meeting coordination
    - SDSC: TG Central DB
    - NCSA: POPS, TG Allocations group
- **Currently awarding >1B Normalized Units of resources**
  - Cost of operation paid by NSF

# User Community is Growing

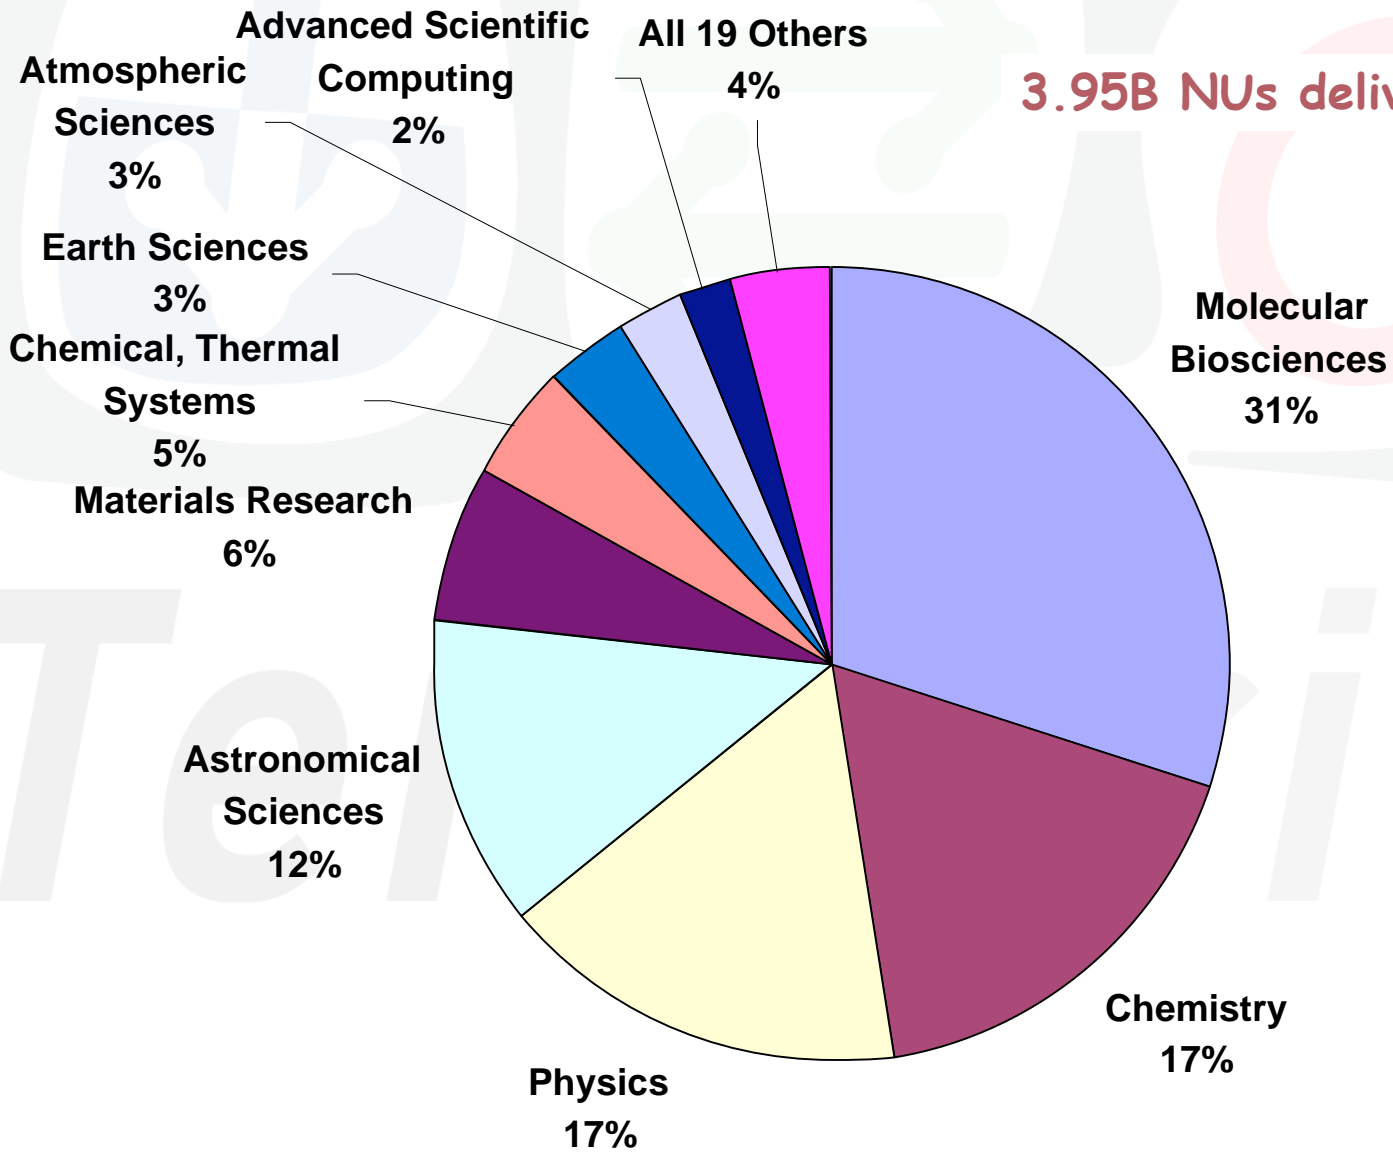


Source: TeraGrid Central Database



# CY2007 Usage by Discipline

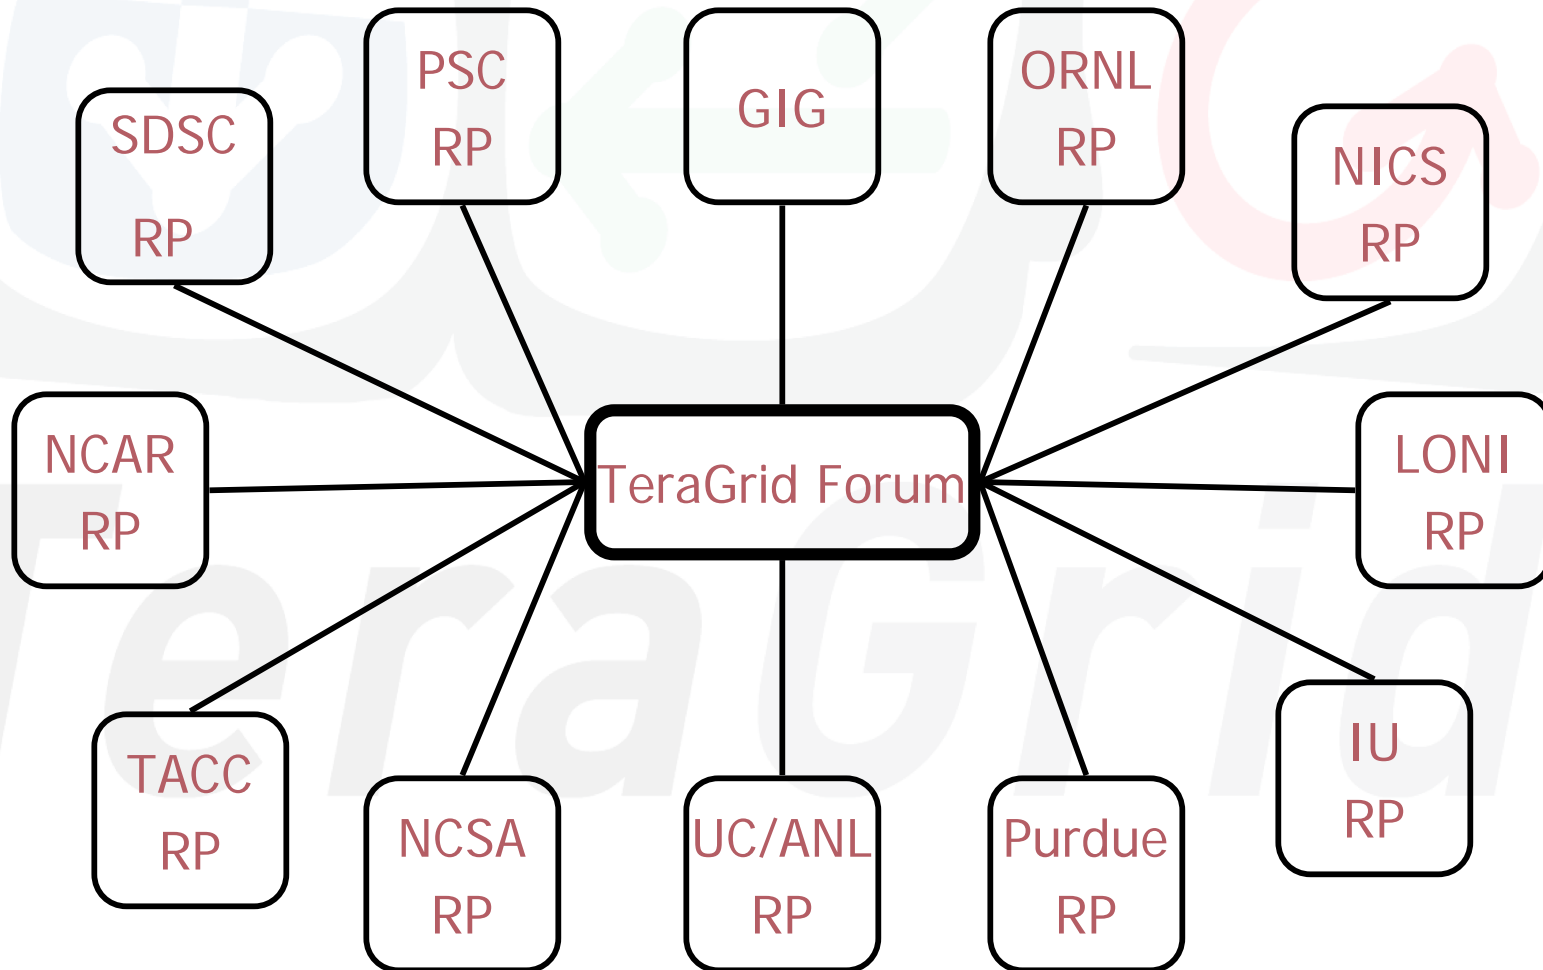
3.95B NUs delivered in CY2007



# How is TeraGrid Organized?

- Not a classic hierarchical organizational structure
  - Evolved from many years of collaborative arrangements between the centers.
  - Continuing to evolve
- Set up more like a large cooperative research group
  - Multiple grants funding the activities
- Made up of 11 Resource Providers (RP's) and the Grid Infrastructure Group (GIG)
  - Specific responsibilities for each member
- Centrally coordinated through the TeraGrid forum

# TeraGrid Organizational Chart



Grid Infrastructure Group, Oak Ridge National Lab, National Institute for Computational Sciences, Louisiana Optical Network Initiative, Indiana University, Purdue University, University of Chicago/Argonne National Lab, National for Supercomputing Applications, Texas Advanced Computing Center, National Center for Atmospheric Research, San Diego Supercomputer Center, Pittsburgh Supercomputing Center



TeraGrid™

# TeraGrid Forum

- Made up of the PI's from each RP and the GIG
  - Grid Infrastructure Group, Oak Ridge National Lab, National Institute for Computational Sciences, Louisiana Optical Network Initiative, Indiana University, Purdue University, University of Chicago/Argonne National Lab, National for Supercomputing Applications, Texas Advanced Computing Center, National Center for Atmospheric Research, San Diego Supercomputer Center, Pittsburgh Supercomputing Center
- Led by the RP forum chair, who is responsible for coordinating the group
  - Interim chairs: John Towns (NCSA), Jay Boisseau (TACC)
- Forum is responsible for the strategic decision making that affects the collaboration

# Members of the TeraGrid Forum

- GIG: “Grid Infrastructure Group”
- Manages local Area Directors (AD’s) who direct project activities across multiple RP’s
- Responsible for coordinating and maintaining TeraGrid central services
- Helps facilitate any joint RP activities that could benefit the entire collaboration
- RPs are Resource Providers
- Large scale systems, data sets, visualization systems
- Resources are primarily NSF funded
  - Sites house other facilities and activities
- Access for use is primarily under the control of a national peer review board



# GIG Management

- Director – Dane Skow
  - Matthew Heinzl
- Area Directors
  - EOT (4.5) Scott Lathrop
  - Gateways (10.5) Nancy Wilkins-Diehr
  - Data, Vis, Scheduling (5.5) Kelly Gaither
  - Software Integration (5) Lee Liming/  
J.P. Navarro
  - Network, Ops, Security (10) Tony Rimovsky
  - User Services (7.5) Sergiu Sanielevici

M. Heinzl, SAB Meeting, January 14-15, 2008



# TeraGrid Coordination

- Weekly Area Directors meetings
- Weekly Architecture calls
- Weekly Gateway Coordination Calls.
- Biweekly TG forum leadership calls
- Education, Outreach and Training (EOT) activities
- TG Quarterly meeting
  - Usually aligned with the resource allocations meeting
- Area Summits
- Working groups and RAT's
  - Software
  - Scheduling
  - User Services
  - Etc...

M. Heinzl, SAB Meeting, January 14-15, 2008



# Area Director Objectives

- Each Area Director (AD) identifies a set of high level objectives to accomplish before the end of the project
- Decompose high level objectives into detailed objective breakouts
  - Tie each objective back to at least one strategic goal
  - Break out what could be done in 2008 vs. 2009
  - Rank objectives by importance
- Look for cross area connections between objectives
- Review objectives with the project

# Statements of Work

- A call for SOW's made during the presentation of 2008 and 2009 objectives review
- Each AD works with the RP's to clarify work that needs to be done to accomplish their objectives
- Each GIG sub award submits a statement of work to address the work required
- ADs review the SOW and match priorities to resources

# Tying the Process Together

- The budget is formed based on the items in the SOW
- Connecting the pieces
  - The budget will tie directly back to the SOW's from the RP's
  - Each SOW item address one or more of the AD objectives
  - Each AD objectives ties back to one or more strategic goal
  - The strategic goals come from user input
- Once a final budget is completed, all the pieces are combined into a detailed program and project plan
- This is a evolving process, and changes and improvements are implemented where they make sense

# User Support Coordination by GIG

- **Mission:**
  - Manage joint work by user support staff at all RP sites in order to provide excellent assistance to all TeraGrid users.
- **Current Staff:**
  - 2.1 FTE in GIG
  - Working with ~15 FTE of user consultants in the 11 RPs.
- **Complementary to TG User Facing Presence**
  - [www.teragrid.org/userinfo](http://www.teragrid.org/userinfo)
  - [www.teragrid.org/kb](http://www.teragrid.org/kb)
  - [portal.teragrid.org](http://portal.teragrid.org)

# GIG User Support Coordination Tasks

- **Lead the User Services Working Group**
  - Forum for discussions and concerted action by user consultants from all RP sites
  - Resolve cross-site and “difficult” tickets referred by TG help desk
  - Weekly teleconferences, email list for assigning/solving problems
- **User Champion Coordination**
  - Assign and coordinate contact teams for each xRAC user group
  - Foster and report transformative science achievements; extract exemplary scenarios of TG resource use
  - Manage cross-site user support
  - Coordinate TG-wide aspects of integration of Track 2 systems

# GIG User Support Coordination Tasks


- **Advanced Support for Teragrid Applications (ASTA)**
  - With help from the xRAC reviewers, foster mappings of transformative science to exemplary use of TG resource combinations, where a focused insertion of TG staff expertise for up to 1 year is likely to result in important progress
  - Manage ASTA project execution
  - 35 projects supported to date: 25 graduated, 10 active.
- **Coordinate User Requirements Gathering**
  - Conduct and interpret annual user surveys
  - Extract user needs from contact team interactions
  - This is the basis for 2008 and 2009 planning

Src: S. Sanielevici, SAB Meeting, January 14-15, 2008



# TeraGrid User Community in 2006

Use Modality	Community Size (est. number of projects)
Batch Computing on Individual Resources	850
Exploratory and Application Porting	650
Workflow, Ensemble, and Parameter Sweep	250
Science Gateway Access	500
Remote Interactive Steering and Visualization	35
Tightly-Coupled Distributed Computation	10



# User Needs Survey

- **Anonymous survey via online service**
  - survey period was Sep 4 – Oct 19, 2007
  - allowed respondents to self-identify if they wanted potential follow-up from TG staff
    - assured identities not revealed in any reports
- **Focused on:**
  - understanding usage patterns
  - importance users place on various capabilities and services
- **Developed by TG GIG Office**
  - consultation with Ann Zimmerman

# TG 2007 User Needs Survey Highlights

- Full set of responses from 203 of 800 active users in Sep-Oct 2007
- Use modes:
  - ~70% indicate they practice “traditional supercomputing”
    - Use machines independently
  - ~60% use multiple resources
    - ~35% indicate cross site run support is important to them
    - ~13% use “workflow” middleware linking job execution on multiple machines
  - ~25% have a need for interactive supercomputing
    - Dynamic visualization, data analysis

# TG 2007 User Needs Survey Highlights

- **Access methods:**

- 73% log in with *ssh* with the local passwords
- 38% *ssh* into one TeraGrid machine and then use certificate based authentication to access other resources
- 13% log in to the TeraGrid User Portal and access all other resources from that session
- 12% use a *MyProxy* client on their local workstation

- **Programming paradigms**

- 67% require a MPI parallel programming environment
- 18% need a SMP parallel programming environment
- 29% need to run serial code.
- 5% require distributed workflows

# TG 2007 User Needs Survey Highlights

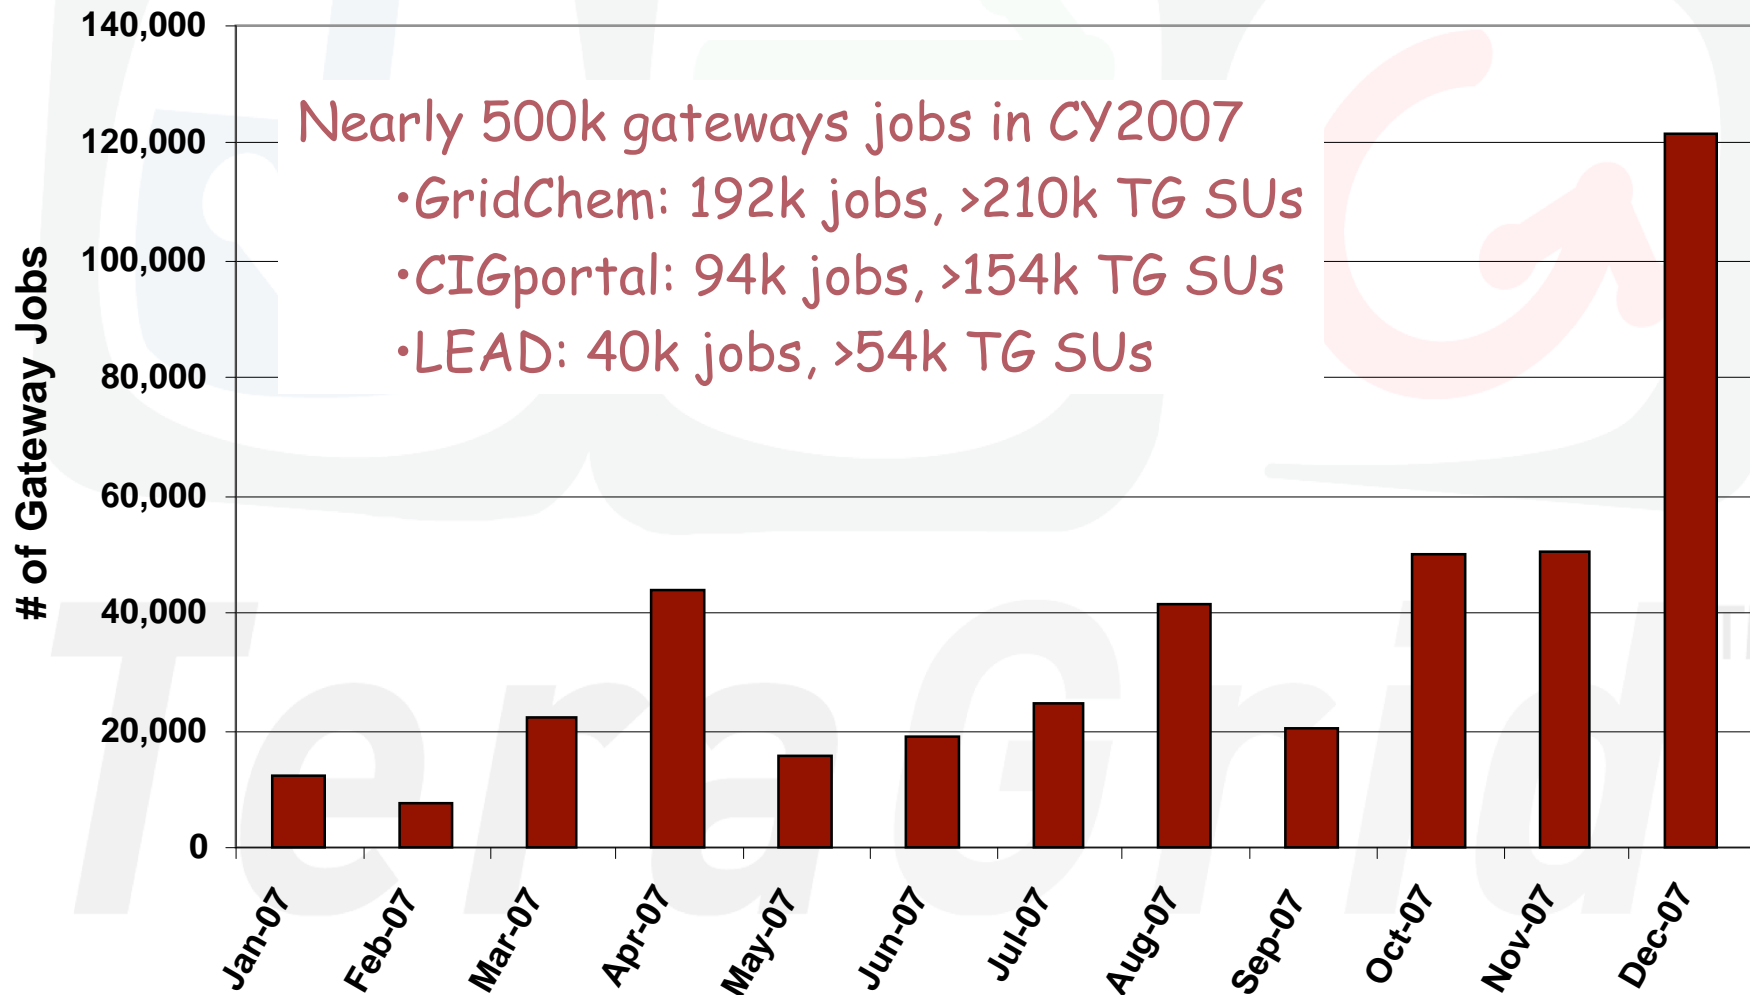
- **Job submission:**

- 70% submit jobs locally to each machine they use
- 24% require a scheduling system that runs and manages a set of independent jobs on one or more machines
  - 14% consider this unnecessary for at least two years
- 21% require a system that routes to the best machine to meet their needs
  - 15% who don't foresee needing this capability for at least 2 years
- 18% require a system that runs and manages a set of jobs *simultaneously* on several machines, and the capability to accept reservations on one or more machines
  - 20% do not foresee needing simultaneous runs
  - 21% do not foresee needing reservations
- 22% require the capability to have a job start on demand (no waiting period)
  - 22% do not foresee needing this capability for at least 2 years
- 11% require a scheduling system that can route jobs to local or campus machines as well as to TeraGrid resources
  - 20% do not foresee needing this capability for at least 2 years
- 8% require a scheduling system that can route to resources on the TeraGrid as well as those on its peer Grids
  - 23% of respondents do not think this applicable for at least 2 years

J. Towns, SAB Meeting, January 14-15, 2008



# Gateways Usage is Growing



J. Towns, SAB Meeting, January 14-15, 2008



# Science Gateway Timeline

- **October, 2004** “TeraGrid Science Gateway” term originates
  - We will help them build gateway portals that leverage TeraGrid capabilities and provide web-based interfaces to community tools. Typical services provided will include access to the following:
    - **Data:** metadata catalogs for the community data resources, the user’s experiments, and remote files, with access via browsable directories, query interfaces, or indexes.
    - **Analysis:** hyperlinked visualization and other data analysis and grid-enabled desktop tools.
    - **Applications:** applications encapsulated as web services and given a user interface in the portal. The portal manages back-end job management and, based on the user’s authorization capabilities, the level of resources applied to the user’s request.
    - **Collaboration:** newsgroups, shared data spaces, “publication” mechanisms.
    - **Workflow:** tools that enable the user to compose TeraGrid and application services to create new applications to be “published” for others to use.

# TeraGrid RATs (Requirements Analysis Teams)

- Science Gateway Requirements Analysis Team (RAT)

- Origin of the RAT

- Needs of the communities



- October 2005 most gateways begin charging in earnest due to funding delays

- Identification of common needs across the gateways
- Production use of TG resources
- Development of process and policy within TG for scalable gateway svcs

- Results on teragrid forum website

Src: N. Wilkens-Diehr, SAB Meeting, January 14-15, 2008



# 2006 – Implementing Common Gateway Requirements

## •Web Services

- GT4 deployment, identification of remaining capabilities
- Information services, WebMDS

## •Auditing

- Need to retrieve job usage info on production resources
- GRAM audit deployed in test mode in September, inclusion in CTSSv4

## •Community Accounts

- Policy finalized, security approaches being tested by RPs
- Attribute-based authentication testing

## •Allocations

- Changes in allocation procedures, the mechanisms used to evaluate science impact, and models for identity management, authentication and authorization that are more tuned to virtual organizations.

## •Scheduling

- Metascheduling RAT
- On-demand via SPRUCE framework

## •Outreach

- Talks, Schools/workshops (NVO, GISolve), major project demonstrations (LEAD)
- SURA, HASTAC, GEON, CI-Channel, SC, Grace Hopper, MSI-CI2, Lariat, Science Workflows and On Demand Computing for Geosciences Workshop

## •Primer

- Living document in wiki, provides up-to-date overview and instructions for new gateway developers (“how to make your portal a TeraGrid science gateway”)

# 2007 – Gateways move into production

- **Web Services**
  - Development of common services
    - QBETS “where can I run soonest” service
- **Auditing**
  - Provide capability to regularly report number of gateway users
- **Community Accounts**
  - Finalize community account implementation policy
  - Provide web interface to account details for TG security staff
- **Allocations**
  - Collaboration with xRAC reviewers to develop instructions for gateway Pis
- **Scheduling**
  - Metascheduling working group
  - Urgent computing workshop
- **Outreach**
  - “Build a Gateway” tutorial at TG07
  - Gateways featured in student competition at TG07
  - Cross directorate presentation at NSF
  - LEAD collegiate forecast competition, April
  - GISolve used in 2 classes
  - NVO announcement of production TG capabilities
  - EOT supplement evaluation of gateway use in education
- **Primer**
  - TG documentation staff identified to move the Primer into fully functional documentation
- **Addressing issues that prevent current gateways from using TG in production**

# Gateway Objectives for 2008 and 2009

- TeraGrid integration will be straightforward for new and existing gateway developers
- There will be a set of easy to discover general services provided by and for Gateways
- The targeted support program will be well-organized
- We will be able to routinely count end gateway users, who will total 25% of total TeraGrid users
- There will be a funded cross-directorate gateway program at the NSF

# Review of the Integration Process

- The requirements are determined from the needs of the user
  - Prioritized and analyzed based capabilities and GIG funding
- RPs propose to carry out the work
  - Commitments to supporting the baseline integration efforts
  - Separately funded
- The budget is formed, based on the items in the SOW
  - The budget will tie directly back to the SOW's from the RP's
  - Each SOW item address one or more of the AD objectives
  - Each AD objectives ties back to one or more strategic goal
  - The strategic goals come from user input
- Once a final GIG budget is completed, all the pieces are combined into a detailed program and project plan
- Follow up and tracking for the efforts

# Planning For the Future

- **TeraGrid:**
  - Continue to focus on providing leadership in HPC for high-end computational research.
  - Work more closely with other CI projects to help develop a comprehensive national CI
    - Also integrate with CI of all scales to provide easier access to entire spectrum of resource needed by diverse research programs
  - Expand the number of users able to utilize high-end resources
    - Working with campus and science facilities to provide pathways for users to more seamlessly scale up into TG (tools, procedures, etc.)
    - Expanding support for users of Science Gateways that simplify usage of TG2's high-end resources.
- **NSF planning for follow on to TeraGrid program:**
  - Future of TeraGrid planning: <http://teragridfuture.org/>



# Acknowledgements

Funding for TeraGrid is provided by the US National Science Foundation.

TeraGrid is the result of the efforts of a large number of people at multiple institutions since 2001.

This presentation is based on information and slides from Matt Heinzl, Dennis Gannon, John Towns, <sup>TM</sup> Sergiu Sanielevici, Nancy Wilkins-Diehr, and many others.

