TeraGrid: Service-Oriented Science Infrastructure

Charlie Catlett
Director, TeraGrid
www.teragrid.org
Senior Fellow, Computation Institute
The University of Chicago and Argonne National Laboratory

December 2006

TeraGrid is supported by the National Science Foundation *Office of Cyberinfrastructure*

"NSF Cyberinfrastructure Vision for 21st Century Discovery"

From Draft 7.1 CI Plan at www.nsf.gov/oci/

4. Education and Workforce

3. Collaboratories, observatories, virtual organizations

"sophisticated" science application software

1. Distributed, scalable up to petaFLOPS HPC

includes networking, middleware, systems software

2. Data, data analysis, visualization

includes data to and from instruments

- *provide sustainable and evolving CI that is secure, efficient, reliable, accessible, usable, and interoperable*
- *provide access to world-class tools and services*

Source: Dan Atkins, NSF Office of Cyberinfrastructure

TeraGrid A National Production Facility

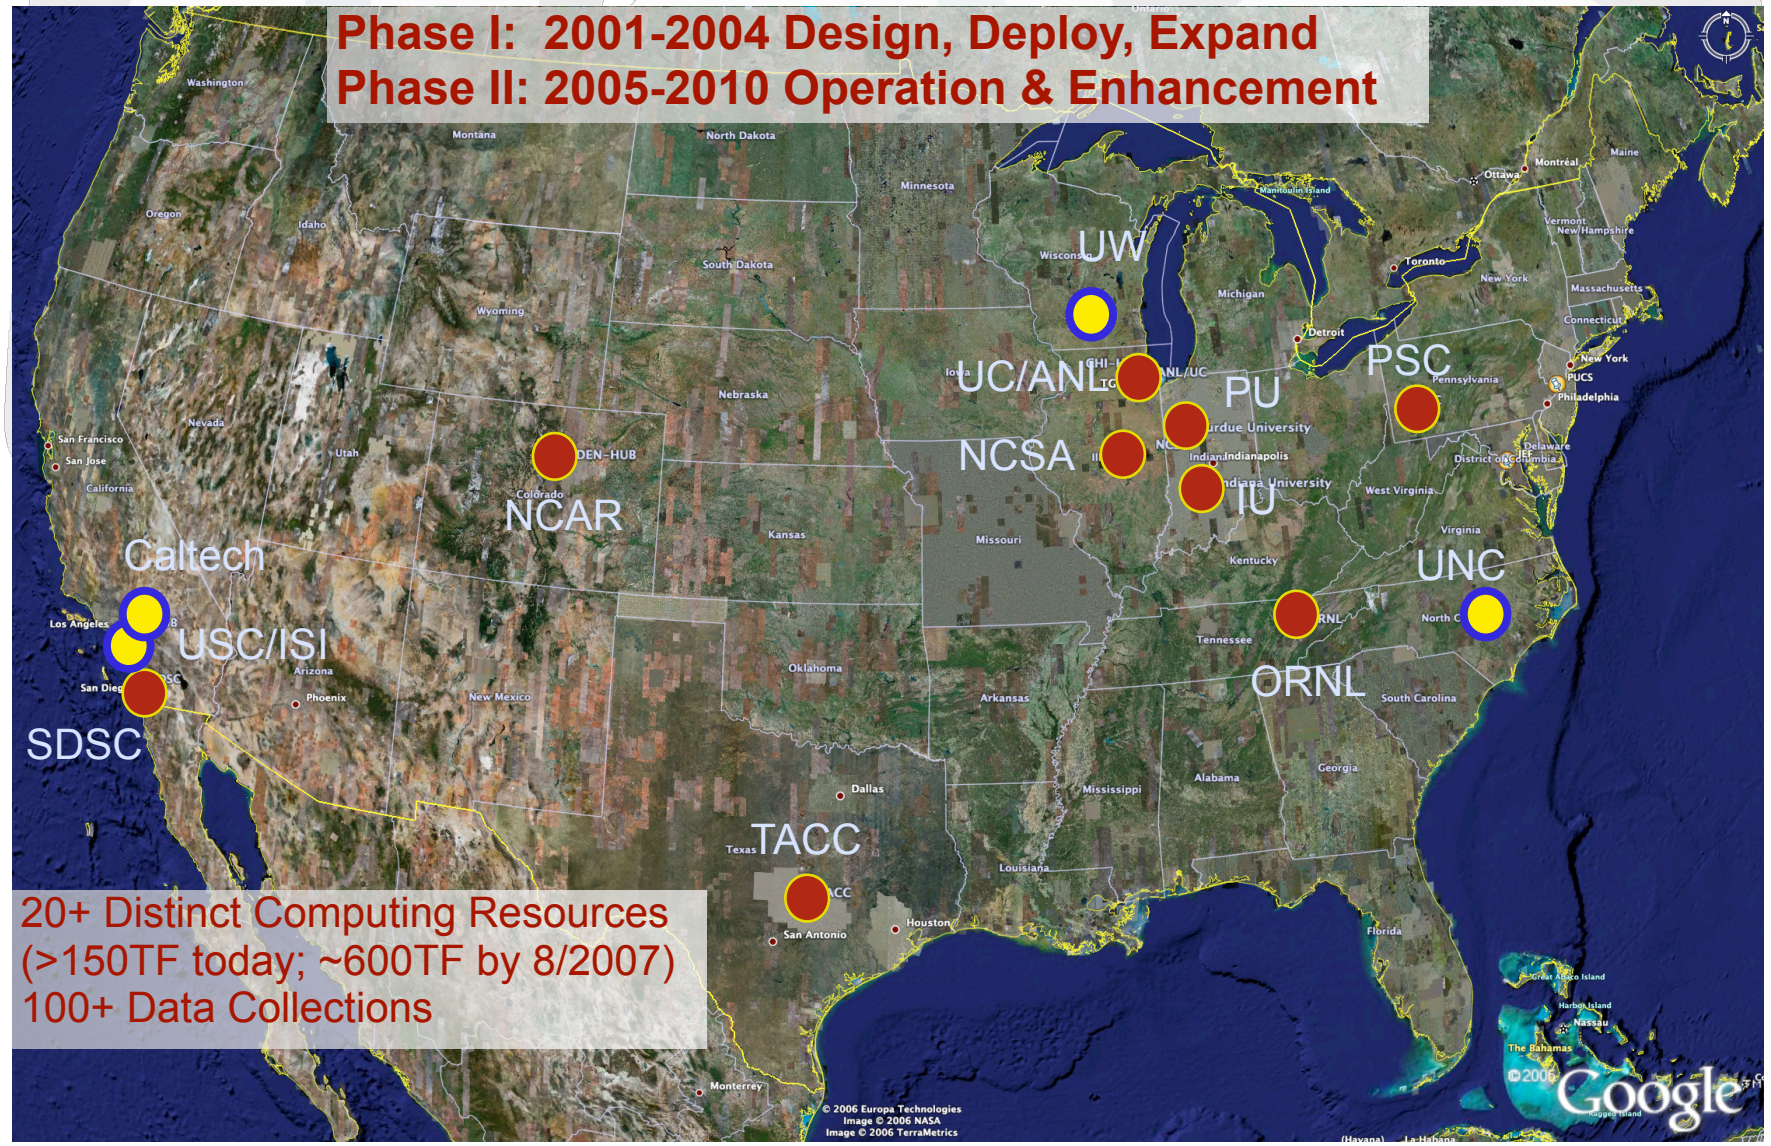

Charlie Catlett (cec@uchicago.edu)

December 2006

TeraGrid Projects by Institution

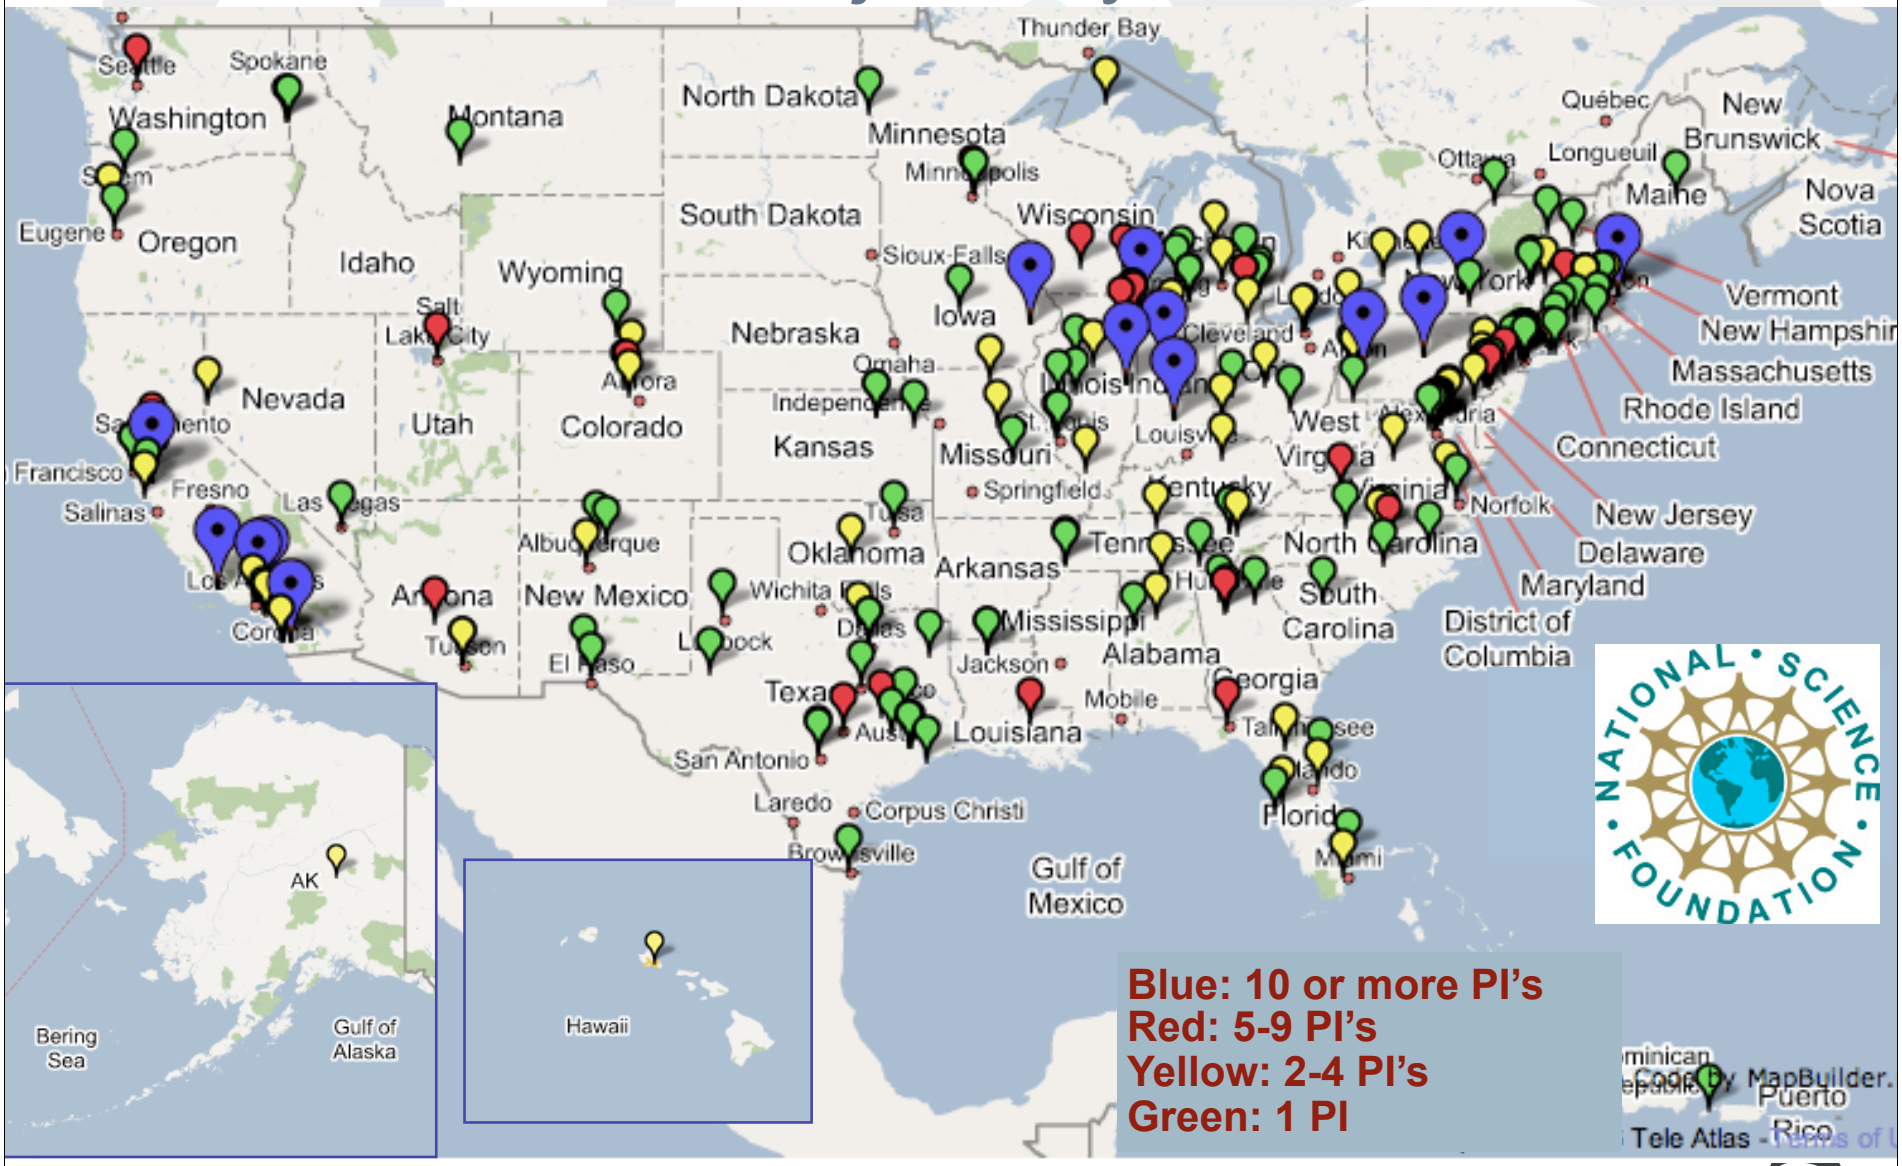

1000 projects, 3200 users

Charlie Catlett (cec@uchicago.edu)

December 2006

TeraGrid

Who Uses TeraGrid?

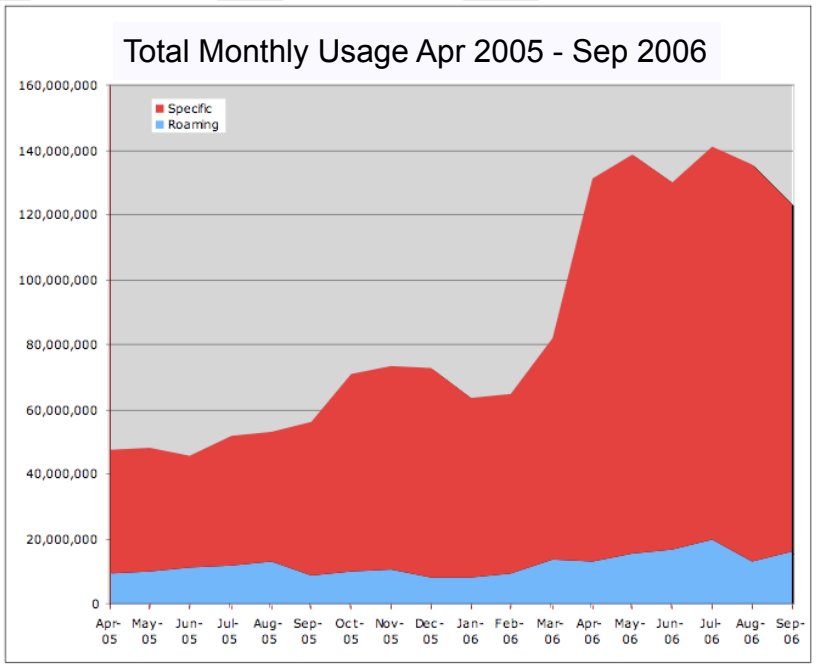

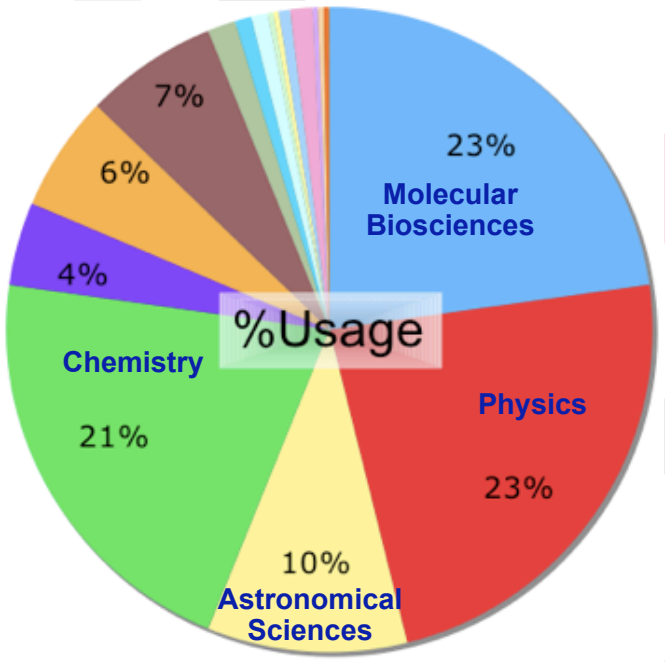

- Molecular Biosciences
- Physics
- Astronomical Sciences
- Chemistry
- Materials Research
- Chemical, Thermal Systems
- Atmospheric Sciences
- Advanced Scientific Computing
- Earth Sciences
- Biological and Critical Systems
- Ocean Sciences
- Cross-Disciplinary Activities
- Computer and Computation Research
- Integrative Biology and Neuroscience
- Mechanical and Structural Systems
- Mathematical Sciences
- Electrical and Communication Systems, Design and Manufacturing Systems, Environmental Biology

1000 projects, 3200 users

Charlie Catlett (cec@uchicago.edu)

TeraGrid™

TeraGrid Resources

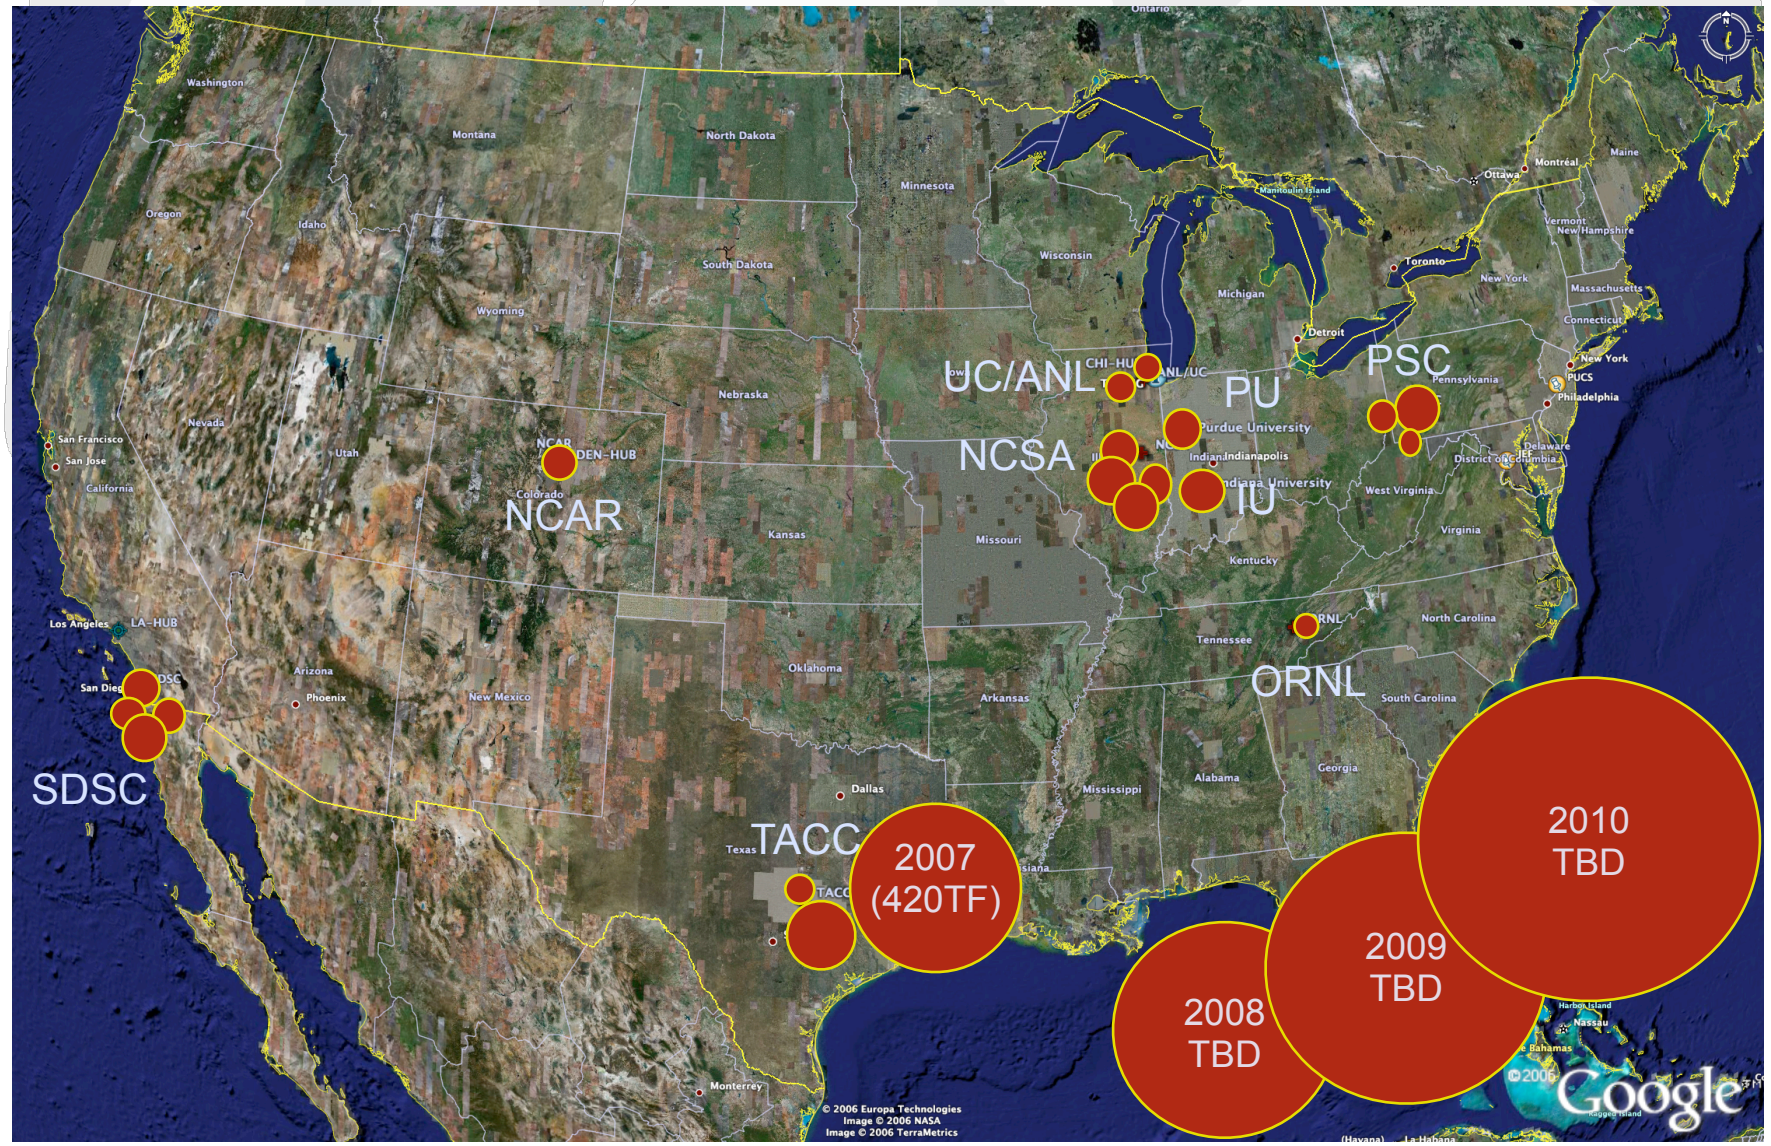

Charlie Catlett (cec@uchicago.edu)

December 2006

Initial TeraGrid Integration: Implementation-based

- Coordinated TeraGrid Software and Services (CTSS)
  - Provide software for heterogeneous systems, leverage specific implementations to achieve interoperation.
  - Evolving understanding of “minimum” required software set for users
  - Defines a single “compute environment/execution” service, includes software-based verification/validation
    - Globus Toolkit, Condor, Inca, internal accounting, allocations and account management tools
- Motivation
  - Attract traditional (client-server) user community
  - Reproducible environment, human-out-of-the-loop.
  - No other feasible approaches!

Challenges

- Scale
  - What works for 4 sites and identical machines is difficult to scale to 10+ sites and 20+ machines with many architectures
- Sociology
  - Requires high-level of buy-in from autonomous sites
    - (to run software or adopt conventions not invented here...)
- Interoperation (e.g. with other Grids)
  - Requires adoption of common software stack
    - ("everything would be interoperable if everyone just ran my software...")
      - (see Sociology)

TeraGrid 2006-2007: Services-based

- Core services: define a "TeraGrid Resource"
  - Authentication & Authorization Capability
  - Information Service
  - Auditing/Accounting/Usage Reporting Capability
  - Verification & Validation Service
- Provides a foundation for value-added services.
  - Each Resource also runs one or more added services, or "kits"
    - Enables a given resource to have a significantly smaller set of components than the previous "full" CTSS
    - Allows subsets of resources to offer advanced capabilities, exploiting architectures or common software choices.
    - Allows portals (science gateways) to customize service offerings
  - Core and individual kits can evolve incrementally, in parallel
  - Largely a re-organization and re-casting of CTSS, i.e. does not require radical new approach, software development, etc.

Example Value-Added Service Kits

- Job Execution
- Application Development
- Science Gateway Hosting
- (Web) Service Hosting
  - dynamic service deployment
- Data Movement
- Data Management
- Science Workflow Support
- Visualization

FLASHBACK: August 2000

A Proposal

- Rethink "Centers," Abandon "Metacenters"
  - "Railroads" vs. "transportation companies"
  - "Computer Center" vs "Application Server"
    - Defined by service & intellectual content
- Focus on Embedded and Discrete Services
  - WWW depends on "free" access to storage resources
    - Value is in content (information), not capacity
  - Use this model for processing resources
    - Value in content (applications) and unique service (including "high performance")
    - Generic processing "free" (e.g. via Condor)
    - Value+ services –e.g. performance, applications
  - Encourage "portals" building via exposing API's

TeraGrid Science Gateways Initiative: Service-Oriented Approach

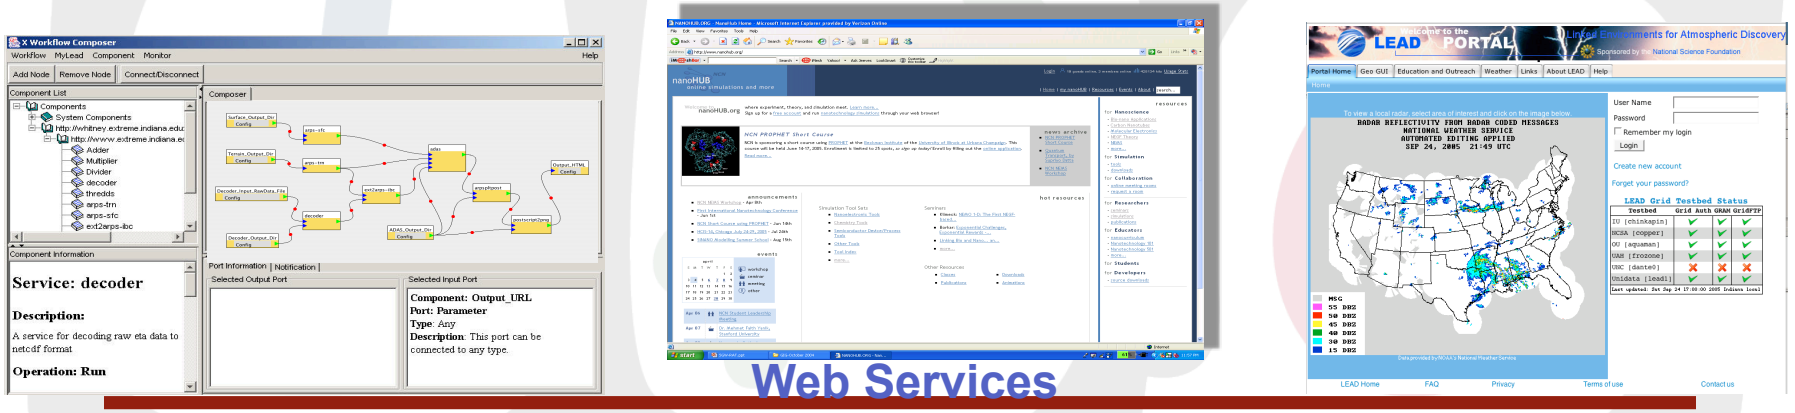

- **Common Web Portal or application interfaces (database access, computation, workflow, etc), exploit standards (primarily web services)**
- **"Back-End" use of grid services such as computation, information management, visualization, etc.**
- **Standard approaches so that science gateways may readily access resources in any cooperating Grid without technical modification.**

Charlie Catlett (cec@uchicago.edu)

TeraGrid™

Science Gateway Partners

Science Gateways

Below is a complete list of current science gateways, to see a detailed project description please click on the name science gateway.

Title	Field of Science	Portal Homepage
Biology and Biomedicine Science Gateway	Molecular Biosciences	Visit Portal
Computational Chemistry Grid	Chemistry	Visit Portal
Computational Science and Engineering Online	Chemistry	Visit Portal
GEON(GEOsciences Network)	Earth Sciences	Visit Portal
GIScience Gateway	Geography and Regional Science	Visit Portal
Grid Analysis Environment	Physics	N/A
Linked Environments for Atmospheric Discovery	Atmospheric Sciences	Visit Portal
National Biomedical Computation Resource	Integrative Biology and Neuroscience	Visit Portal
National Virtual Observatory	Astronomical Sciences	Visit Portal
Network for Computational Nanotechnology and nanoHUB	Emerging Technologies Initiation	Visit Portal
Network for Earthquake Engineering Simulation	Earthquake Hazard Mitigation	Visit Portal
Neutron Science Instrument Gateway	Physics	Visit Portal
Open Life Sciences Gateway	Molecular Biosciences	Visit Portal
Open Science Grid	Advanced Scientific Computing	N/A
SCEC Earthworks Project	Earthquake Hazard Mitigation	Visit Portal
Special PRiority and Urgent Computing Environment	Advanced Scientific Computing	Visit Portal
TeraGrid Visualization Gateway	Visualization, Graphics, and Image Processing	Visit Portal
The Earth System Grid	Global Atmospheric Research	Visit Portal
The Telescience Project	Neuroscience Biology	Visit Portal
Virtual Laboratory for Earth and Planetary Materials	Materials Research	Visit Portal

For more information on the science gateways effort please visit the Science Gateways program page.

Science Gateways access via the TeraGrid User Portal (portal.teragrid.org).
Additional gateways are currently working with TeraGrid to join this list of active gateways.
For more information please contact **Nancy Wilkins-Diehr (wilkinsn@sdsc.edu)**

Charlie Catlett (cec@uchicago.edu)

TeraGrid

TeraGrid User Portal

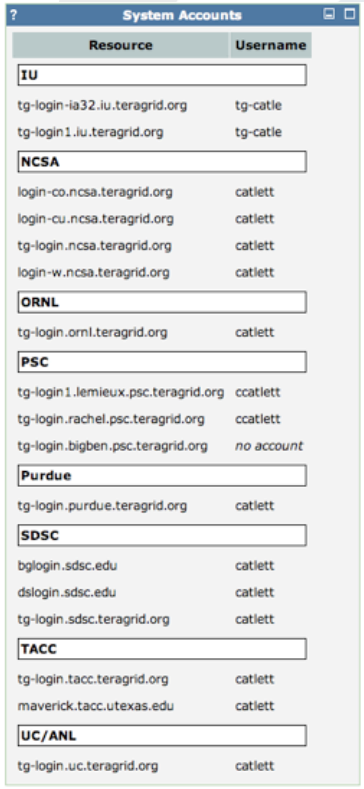

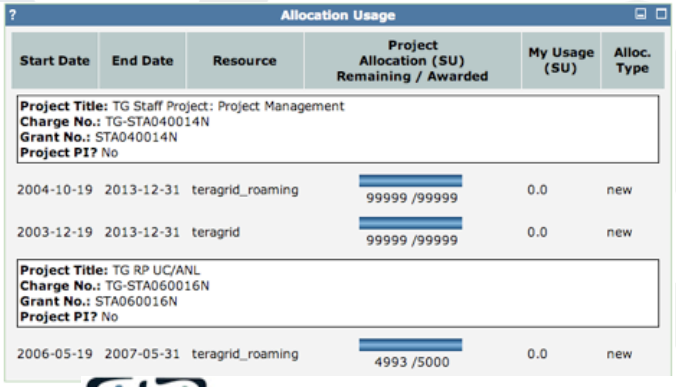

Tracking usage for my allocations

Managing Credentials

Current State of all Resources

TeraGrid User Portal

Logout

Welcome, Charles E Catlett

Home | My TeraGrid | Resources | Documentation | Consulting | Allocations

Systems Monitor | Science Gateways | Data Collections

TeraGrid Systems Monitor

Refresh

High Performance Computing Systems

Name	Institution	System	CPUs	Peak TFlops	Memory TBytes	Disk TBytes	Load	Jobs* R	Q	O
Lonestar	TACC	Dell PowerEdge Linux Cluster	5200	55.00	10.40	94.90		186	64	0
Big Red	IU	IBM e1350	2048	20.40	8.20	266.00		1	126	1892
Tungsten	NCSA	Dell Xeon IA-32 Linux Cluster	2560	16.38	3.75	109.00		76	1149	34
DataStar p655	SDSC	IBM Power4+ p655	2176	14.30	5.75	115.00		34	48	33
TeraGrid Cluster	NCSA	IBM Itanium2 Cluster	1744	10.23	4.47	60.00		57	430	70
Bigben	PSC	Cray XT3	2090	10.00	2.02	48.80		1	32	59
Lear	Purdue	Dell EM64T Linux Cluster	1024	6.60	2.00	28.00		218	385	4
Cobalt	NCSA	SGI Altix	1024	6.55	3.00	100.00		33	262	12
Lemieux	PSC	HP Alpha Cluster	3000	6.00	2.93	78.13		46	2	38
Blue Gene	SDSC	IBM Blue Gene	2048	5.70	0.50	19.50		3	5	7
TeraGrid Cluster	SDSC	IBM Itanium2 Cluster	524	3.10	1.02	48.80		6	246	68
Copper	NCSA	IBM Power4 p690	384	2.00	1.44	30.00		88	108	2
DataStar p690	SDSC	IBM Power4+ p690	192	1.30	0.88	115.00		18	3	12
TeraGrid Cluster	UC/ANL	IBM Itanium2 Cluster	124	0.61	0.24	4.00		7	10	0
NSTG	ORNL	IBM IA-32 Cluster	56	0.34	0.07	2.14		0	0	0
Rachel	PSC	HP Alpha SMP	128	0.31	0.50	6.00		20	75	1
		Total:	**24322**	**158.82**	**47.17**	**1125.27**		**794**	**2945**	**2232**

Advanced Visualization Systems

Name	Institution	System	CPUs	Peak TFlops	Memory TBytes	Disk TBytes	Graphics HW
TeraGrid Cluster	UC/ANL	Intel Xeon Cluster	192	0.61	0.38	4.00	nVIDIA GeFORCE 6600GT AGP graphics cards
Maverick	TACC	Sun E25K	128	0.27	0.50	0.56	16 nVIDIA QuadroFX 3000G graphics cards
		Total:	**320**	**0.88**	**0.88**	**4.56**	

TeraGrid User Portal -
Eric Roberts, Texas Advanced Computing Center
(ericrobe@tacc.utexas.edu)

TeraGrid

Real-Time Usage Mashup

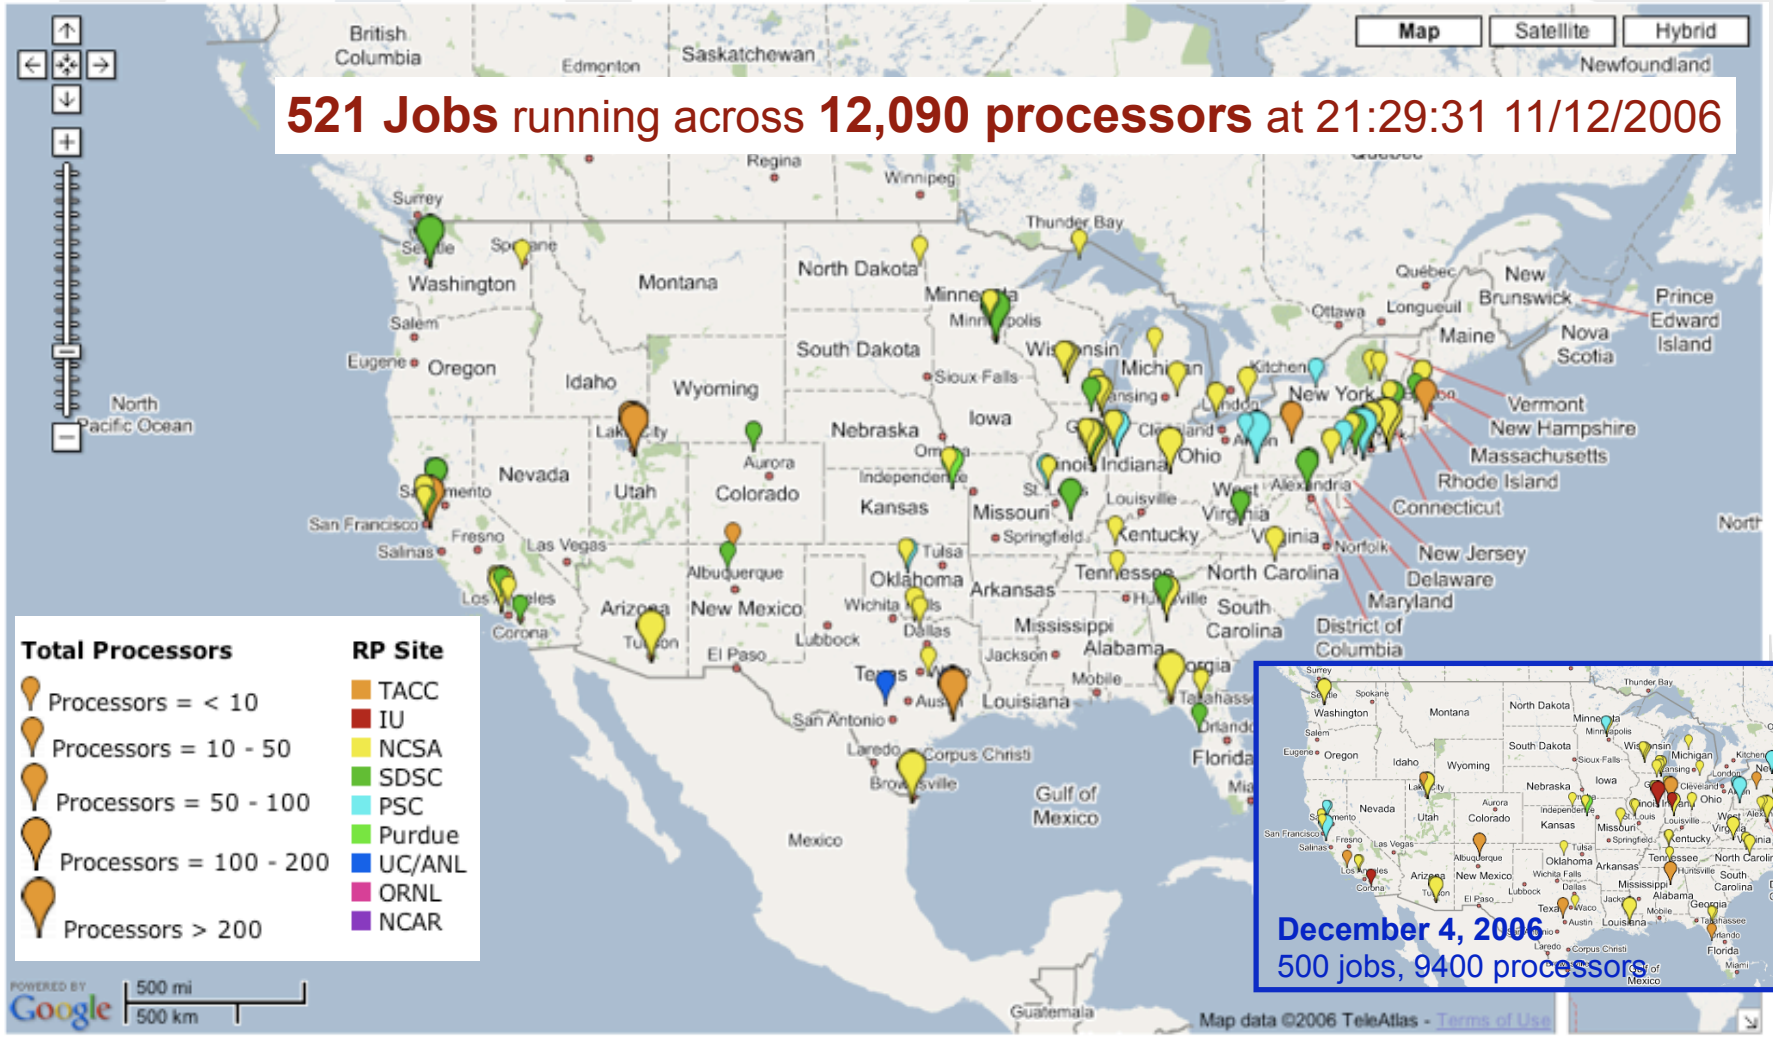

The TeraGrid job map displays the current running jobs across TeraGrid. Each pin location denotes the location of the job owner, the color of the pin denotes the RP site of the job(s), the size of the pin denotes the total number of processors for the jobs. By clicking on the pin you can see the users job informormation - RP site, total number of jobs running, total number of processors - in addition to the user's location, department, and institution.

Continuing to Improve User Tools

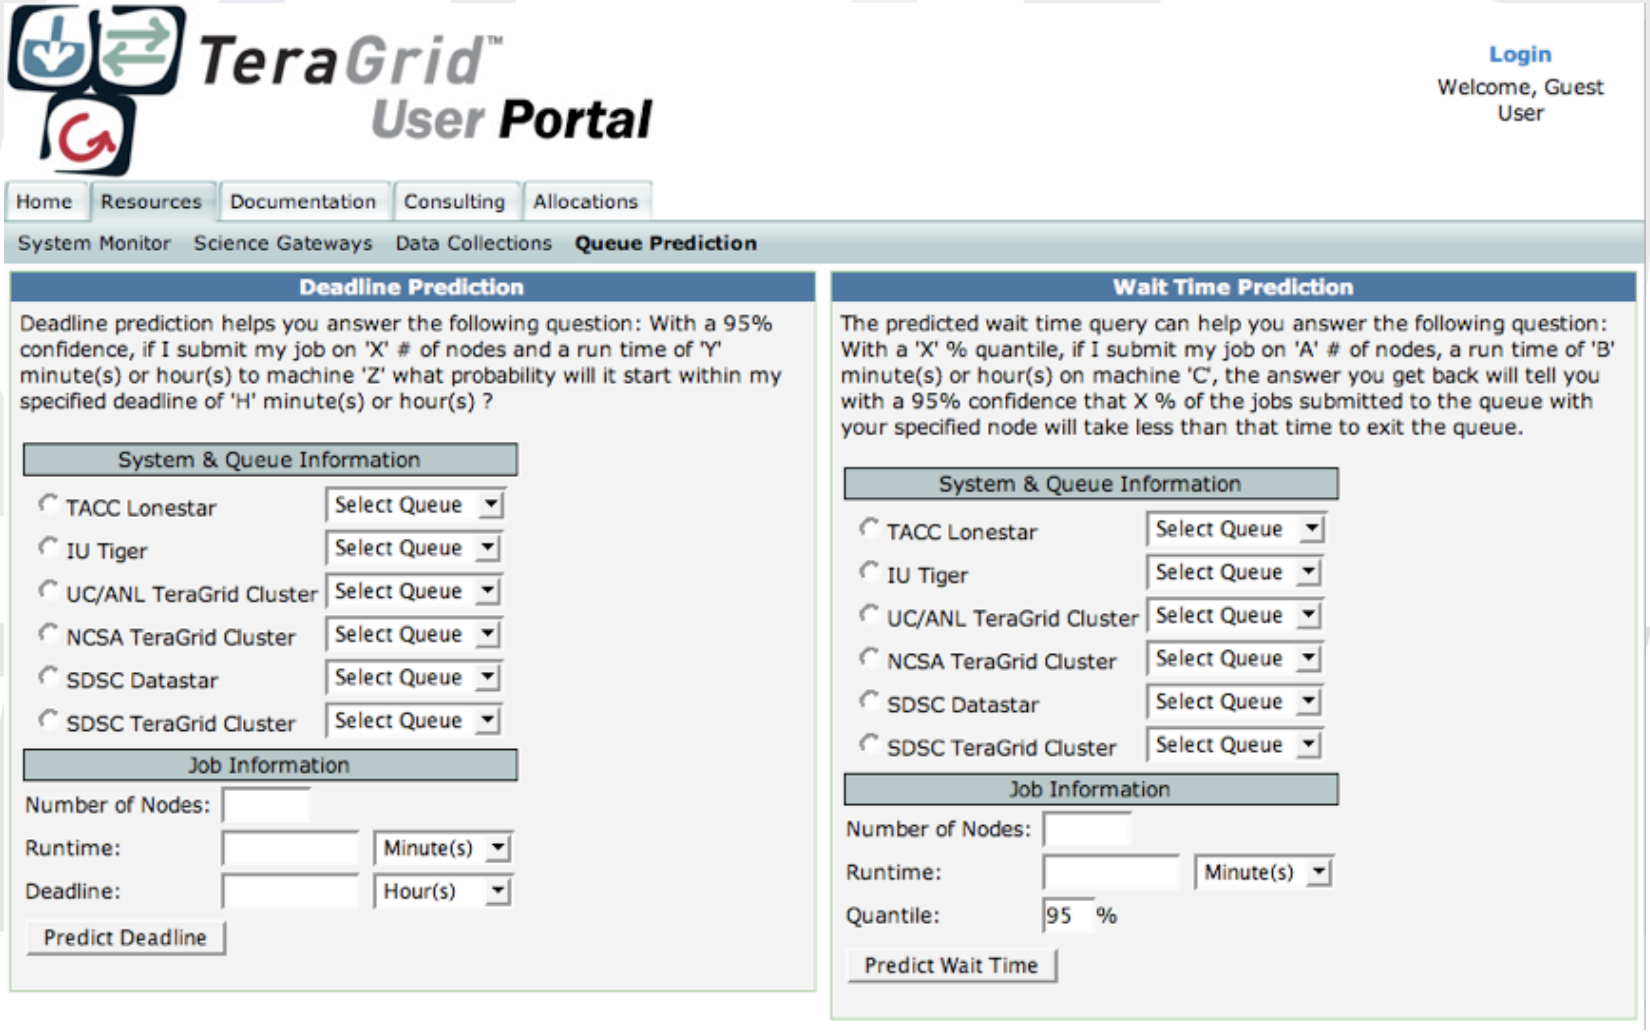

Alpha-test

Deadline prediction - Network Weather Service "Batch Queue Prediction" (BQP) - Rich Wolski (rich@cs.ucsb.edu)
Wait time prediction - Warren Smith (wsmith@tacc.utexas.edu)

Evolving TeraGrid

- Service-Oriented Architecture
  - Core services on all resources
  - Kits for value-added services
- Implications and Opportunities
  - Kits do not necessarily come from TeraGrid team
  - A key kit will be web services support
    - Our aim is to enable the community to develop new services using TeraGrid resources
- How to Get Started
  - Apply for a Development allocation (DAC)
    - A “roaming” account provides you with a login on nearly all TeraGrid resources.

Charlie Catlett (cec@uchicago.edu)