The TIGRE Project:
an Example of a State-Wide Grid

A HiPCAT Project

TIGRE

Texas Internet Grid for Research and Education

Alan Sill, TTU
for the TIGRE
Development Team

Open Science Grid All-Hands Meeting Chapel Hill, NC
March 4, 2008

Background and History: HiPCAT, LEARN and TIGRE

SuperComputing 1998: Representatives from five Texas institutions met and agreed to cooperate and exchange notes on a variety of topics of interest in computing.

High Performance Computing Across Texas (HiPCAT) was born as a result.

This organization has since grown to ten universities in Texas and is considering the application of another.

TIGRE is a grid software project of HiPCAT begun two years ago by the original five research universities.

Another organization and project, LEARN, was funded by the same legislative bill to do optical networking.

FUNDING: TEXAS ENTERPRISE FUND

TEF is a tool to "bring jobs to Texas." Funding normally used to attract business interests.

Intended to create quality jobs and leverage private investment to strengthen economic future of the State. (This mandate differs from that of many grid projects.)

Accompanied by a formal Economic Development Agreement with State DIR.

TIGRE

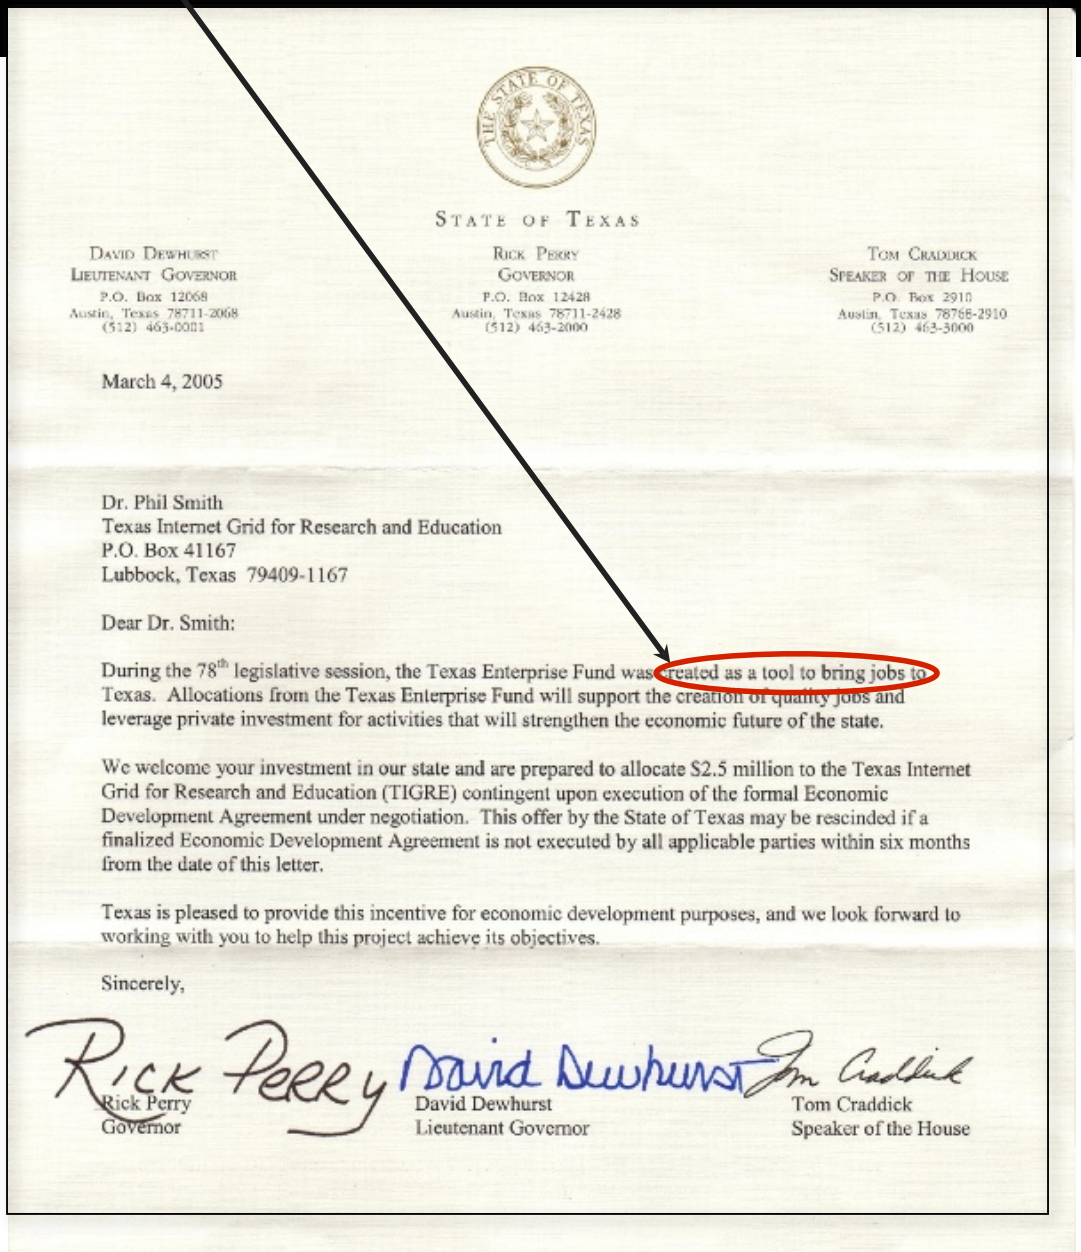

STATE OF TEXAS

DAVID DEWHURST
LIEUTENANT GOVERNOR
P.O. Box 12068
Austin, Texas 78711-2068
(512) 463-0001

RICK PERRY
GOVERNOR
P.O. Box 12428
Austin, Texas 78711-2428
(512) 463-2000

TOM CRADDICK
SPEAKER OF THE HOUSE
P.O. Box 2910
Austin, Texas 78768-2910
(512) 463-3000

March 4, 2005

Dr. Phil Smith
Texas Internet Grid for Research and Education
P.O. Box 41167
Lubbock, Texas 79409-1167

Dear Dr. Smith:

During the 78th legislative session, the Texas Enterprise Fund was created as a tool to bring jobs to Texas. Allocations from the Texas Enterprise Fund will support the creation of quality jobs and leverage private investment for activities that will strengthen the economic future of the state.

We welcome your investment in our state and are prepared to allocate $2.5 million to the Texas Internet Grid for Research and Education (TIGRE) contingent upon execution of the formal Economic Development Agreement under negotiation. This offer by the State of Texas may be rescinded if a finalized Economic Development Agreement is not executed by all applicable parties within six months from the date of this letter.

Texas is pleased to provide this incentive for economic development purposes, and we look forward to working with you to help this project achieve its objectives.

Sincerely,

Rick Perry
Governor

David Dewhurst
Lieutenant Governor

Tom Craddick
Speaker of the House

Project Organization

Targeted application areas were selected by the steering committee, which also has responsibility for the overall direction of the project. These initial areas are:

- Biosciences and Medicine
- Air Quality Modeling
- Energy Exploration

Implementation is carried out by a *development team* consisting of two people at each of the five primary TIGRE institutions.

Technical development organized into *activities* to meet project *milestones*, with work targeted to accomplish these milestones in collaboration across all TIGRE institutions.

Original Goals of the TIGRE Project

- Provide a grid infrastructure that enables integration of computing systems, storage systems, databases, visualization labs and displays, even instruments and sensors across Texas.
- Facilitate new academic - government - private industry research partnerships by dramatically enhancing both computational capabilities and research infrastructure.
- Address research areas of interest to the State of Texas in which substantial increase of computing power, data access, and collaboration are necessary.
- Demonstrate new, preferred, enhanced or increased computing and storage handling capabilities offered by a statewide grid infrastructure.

Details of TIGRE Software Stacks

- Based on the Virtual Data Toolkit (VDT), working in close cooperation with VDT team members, Globus and the Open Science Grid (OSG).
- Uses a simplified VDT set including GSI-OpenSSH, omitting much monitoring and accounting in favor of lightweight status reporting.
- TIGRE implementation based on Web Services (GRAM4) only; pre-WS available only upon request (no requests). ***Translation: a modern grid***
- Client and server software stacks separately available.
- Goal was "one page" installation instructions that can be implemented quickly by newcomers. (Achieved!)
- Authentication via X.509 (new TACC CA is now accredited by IGTF); authorization local, mostly via grid-mapfiles. (TTU uses OSG tools.)
- Installed on systems at all five primary TIGRE institutions; also running at other locations throughout the state.

Project Progress Timeline

Stated goal was to achieve a "quick build" toward working status:

YEAR 1 (Begun Dec. 2005)

Q1:
Project plan ✔
Web site ✔
Certificate Authority ✔
Minimum testbed requirements ✔
Select 3 driving applications ✔

Q2:
Alpha quality user portal ✔

Q3:
Define server software stack ✔
Distribution Mechanism ✔
Simple demo of 1 TIGRE app ✔

Q4:
Alpha client software stack and installation method distributed ✔

YEAR 2

Q1:
Alpha customer management services system ✔
Demonstrate applications in three areas ✔

Q2:
Project-wide global grid scheduler deployed ✔

Q3:
Stable software status (only bug fixes after this) ✔
Required services for TIGRE specified ✔

Q4 (Completed Dec. 2007):
Complete hardening of software ✔
Complete documentation ✔
Finalized procedures and policies to join TIGRE
(suggested drafts completed ✔; need action here.)
Demonstrate TIGRE at SC ✔

*-- Development was completed **on time** and **within budget, according to original plan!***

TIGRE User Portal

http://tigreportal.hipcat.net/

Parallel Computing Resources

Name	Institution	Department	System	CPUs	Peak GFlops	Memory GBytes	Disk GBytes	Status	Load	Jobs
Ada	Rice University	Computer and Information Technology Institute	Cray XD1	632	2939	1320	33736	↑		0R-154Q-0O
Alamo	University of Texas Health Science Center at San Antonio	Department of Biochemistry	Rocks i386 Linux Cluster	19	0	8.7	1126	↑		0R-0Q-0O
Cosmos	Texas A&M University	Texas A&M Supercomputing Facility	SGI Altix	128	666	256	4096	↑		29R-15Q-0O
Eldorado	University of Houston	Advanced Computing Research Laboratory	Eldorado Itanium2 Cluster	126	67	4	2232	↑		0R-0Q-1O
Jacinto	University of Texas Health Science Center at San Antonio	Department of Biochemistry	Microway Linux Cluster	44	66	85.1	8433.3	↑		0R-0Q-1O
Laredo	University of Texas Health Science Center at San Antonio	Department of Biochemistry	Dual Athlon Cluster	32	0	31.6	9509	↑		0R-0Q-1O
Lonestar	The University of Texas at Austin	Texas Advanced Computing Center	Dell PowerEdge Linux Cluster	5200	55000	10400	94900	↑		59R-0Q-191O
Minigar	Texas Tech University	High Performance Computing Center	Dell Linux Cluster	32	230	64	70	↑		8R-0Q-88O
RTC	Rice University	Computer and Information Technology Institute	HP Itanium II Linux Cluster	290	1044	596	7000	↑		0R-16Q-0O
TTU-Antaeus	Texas Tech University	High Performance Computing Center	Dell Xeon Cluster	192	2300	96	6000	↑		10R-0Q-17O
			Total:	**6695**	**62312**	**12861.4**	**167102.3**			

High Throughput Computing Resources

Name	Institution	Department	System	Active PCs	Active CPUs	Memory GBytes	Disk GBytes	Resource Details	Jobs
glb-test	Texas Tech University	High Performance Computing Center	Condor	**0** / 0	**68** / 68	14	1301	🔍	🔍
Rodeo	The University of Texas at Austin	Texas Advanced Computing Center	Condor	**26** / 26	**26** / 26	17	1114	🔍	🔍
			Total:	**26 / 26**	**94 / 94**	**31**	**2415**		

INSTALLATION INSTRUCTIONS ON TIGRE PORTAL

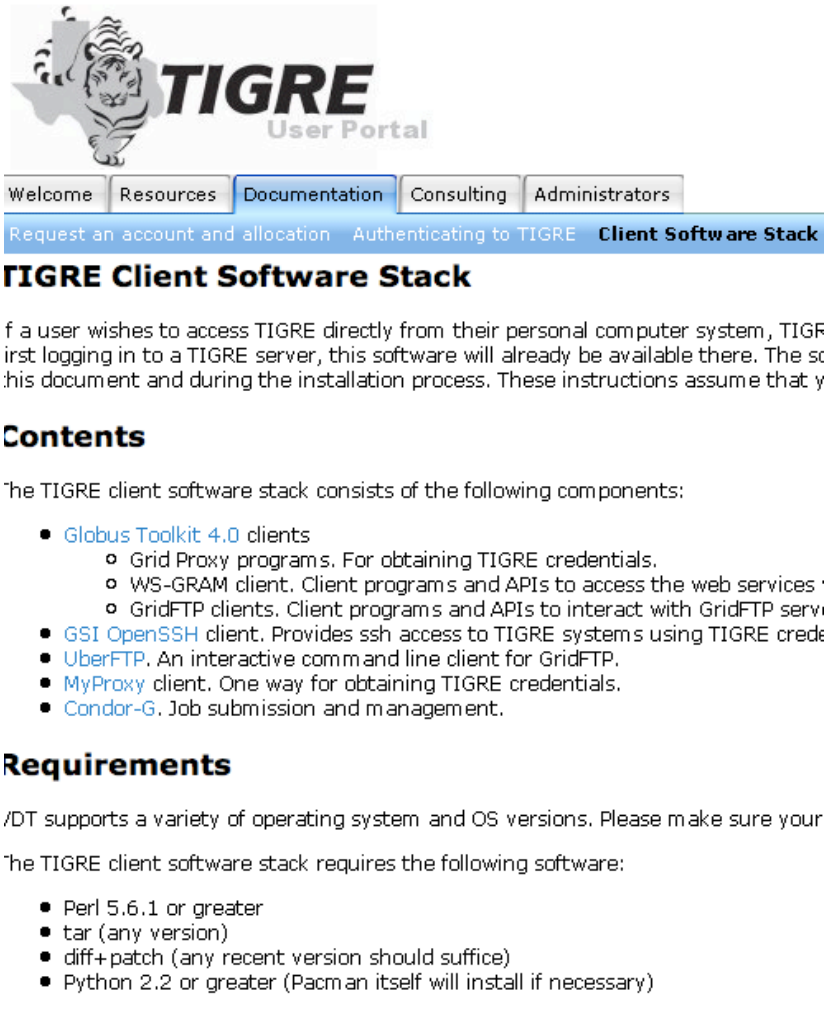

TIGRE User Portal

Welcome | Resources | Documentation | Consulting | Administrators

Request an account and allocation | Authenticating to TIGRE | **Client Software Stack**

TIGRE Client Software Stack

f a user wishes to access TIGRE directly from their personal computer system, TIGR
irst logging in to a TIGRE server, this software will already be available there. The sc
his document and during the installation process. These instructions assume that y

Contents

he TIGRE client software stack consists of the following components:

- Globus Toolkit 4.0 clients
  - Grid Proxy programs. For obtaining TIGRE credentials.
  - WS-GRAM client. Client programs and APIs to access the web services
  - GridFTP clients. Client programs and APIs to interact with GridFTP serv
- GSI OpenSSH client. Provides ssh access to TIGRE systems using TIGRE crede
- UberFTP. An interactive command line client for GridFTP.
- MyProxy client. One way for obtaining TIGRE credentials.
- Condor-G. Job submission and management.

Requirements

/DT supports a variety of operating system and OS versions. Please make sure your

he TIGRE client software stack requires the following software:

- Perl 5.6.1 or greater
- tar (any version)
- diff+patch (any recent version should suffice)
- Python 2.2 or greater (Pacman itself will install if necessary)

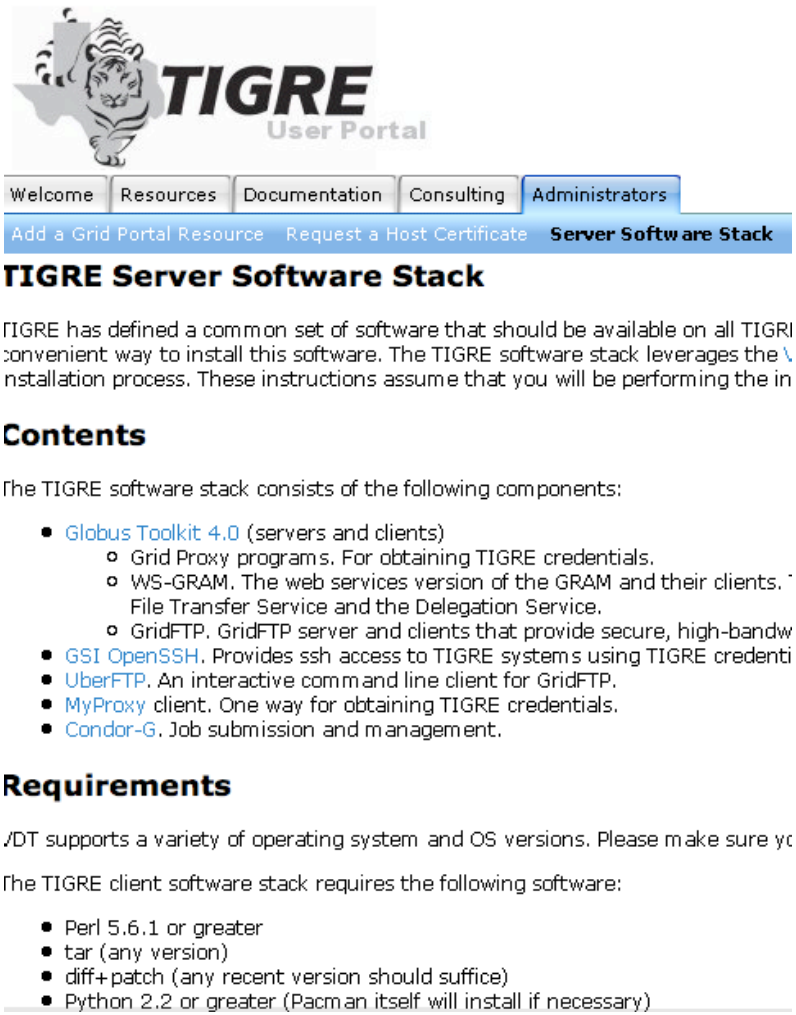

TIGRE User Portal

Welcome | Resources | Documentation | Consulting | Administrators

Add a Grid Portal Resource | Request a Host Certificate | **Server Software Stack**

TIGRE Server Software Stack

TIGRE has defined a common set of software that should be available on all TIGRI
convenient way to install this software. The TIGRE software stack leverages the \
nstallation process. These instructions assume that you will be performing the in

Contents

The TIGRE software stack consists of the following components:

- Globus Toolkit 4.0 (servers and clients)
  - Grid Proxy programs. For obtaining TIGRE credentials.
  - WS-GRAM. The web services version of the GRAM and their clients.
    File Transfer Service and the Delegation Service.
  - GridFTP. GridFTP server and clients that provide secure, high-bandw
- GSI OpenSSH. Provides ssh access to TIGRE systems using TIGRE credenti
- UberFTP. An interactive command line client for GridFTP.
- MyProxy client. One way for obtaining TIGRE credentials.
- Condor-G. Job submission and management.

Requirements

/DT supports a variety of operating system and OS versions. Please make sure yc

The TIGRE client software stack requires the following software:

- Perl 5.6.1 or greater
- tar (any version)
- diff+patch (any recent version should suffice)
- Python 2.2 or greater (Pacman itself will install if necessary)

Authentication and Authorization

Every distributed or grid computing project in existence seems to develop its own approach to authentication and authorization.

We dealt with this in TIGRE by adopting a system to provide strong credentials based on X.509 grid certificates, and leave mapping of these credentials at individual institutions to systems based on their own policies. (TTU uses GUMS/PRIMA)

This worked, but it is probably time for us to adopt common policies and procedures, encourage adoption of common software, and follow new developments in this area.

The system we did use is compatible with OSG and Teragrid approaches and can be extended further with modest efforts.

Project-Wide Metascheduler Studies

We learned early on in the project that *scientific usage models employ workflows*, and need tools that extend beyond writing isolated executables to organizing such workflows. Also, a project-wide metascheduler was called for in the milestones.

We pursued implementation of two different project-wide job workflow metascheduling options and carried out a head-to-head comparison:

1. Grid Resource Management System (GRMS)
2. GridWay (soon to be built into the Globus stack)

No clear winner emerged, but both were made to work and were used in Grid2007 demos. A choice between these two options would require more user experience.

Applications Developed During This Effort:

The application examples pursued in TIGRE were the result of *selection of areas by the steering committee* and not directly from pre-existing self-organized collaborations.

In each case we were able to make progress by putting developers *directly into contact with researchers* to find the optimum solution, and then try to form collaborations where appropriate to pursue common needs.

It seems that this works best when the engagement with the researchers is mediated where necessary by developers making contact to learn and implement best practices in a given field.

We also had several examples of cross-institutional collaboration.

Targeted Applications Summary

Analysis of Macromolecular Assemblies using UltraScan

- UT Health Science Center San Antonio (Borries Demeler, Emre Brooks)
- http://ultrascan.uthscsa.edu
- Composition analysis
  - Analyzes ultracentrifugation data
  - Data are the concentration profiles of diffusing and sedimenting solutes as a function of radius and time
  - Derive composition from this data
- Questions to answer:
  - How many components?
  - What are their molecular weights?
  - What are their shapes?
  - Are there interactions between components?

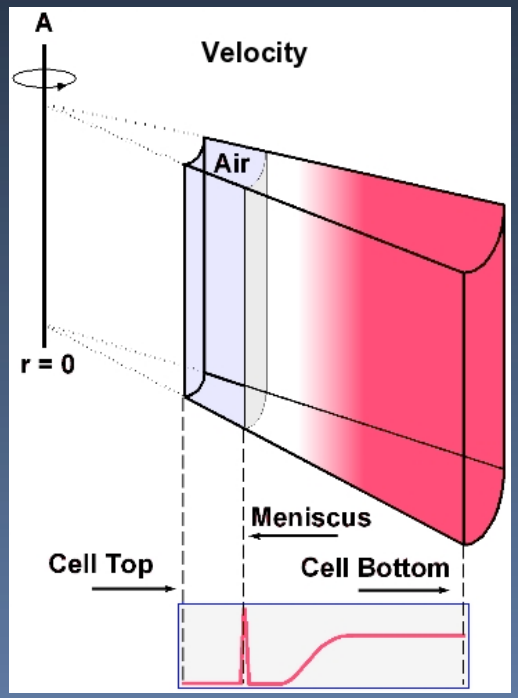

Application Example: Biosciences and Medicine

Initial demo was "UltraScan" : an analysis tool for reduction of data from ultracentrifuge biomolecular optical spectra.

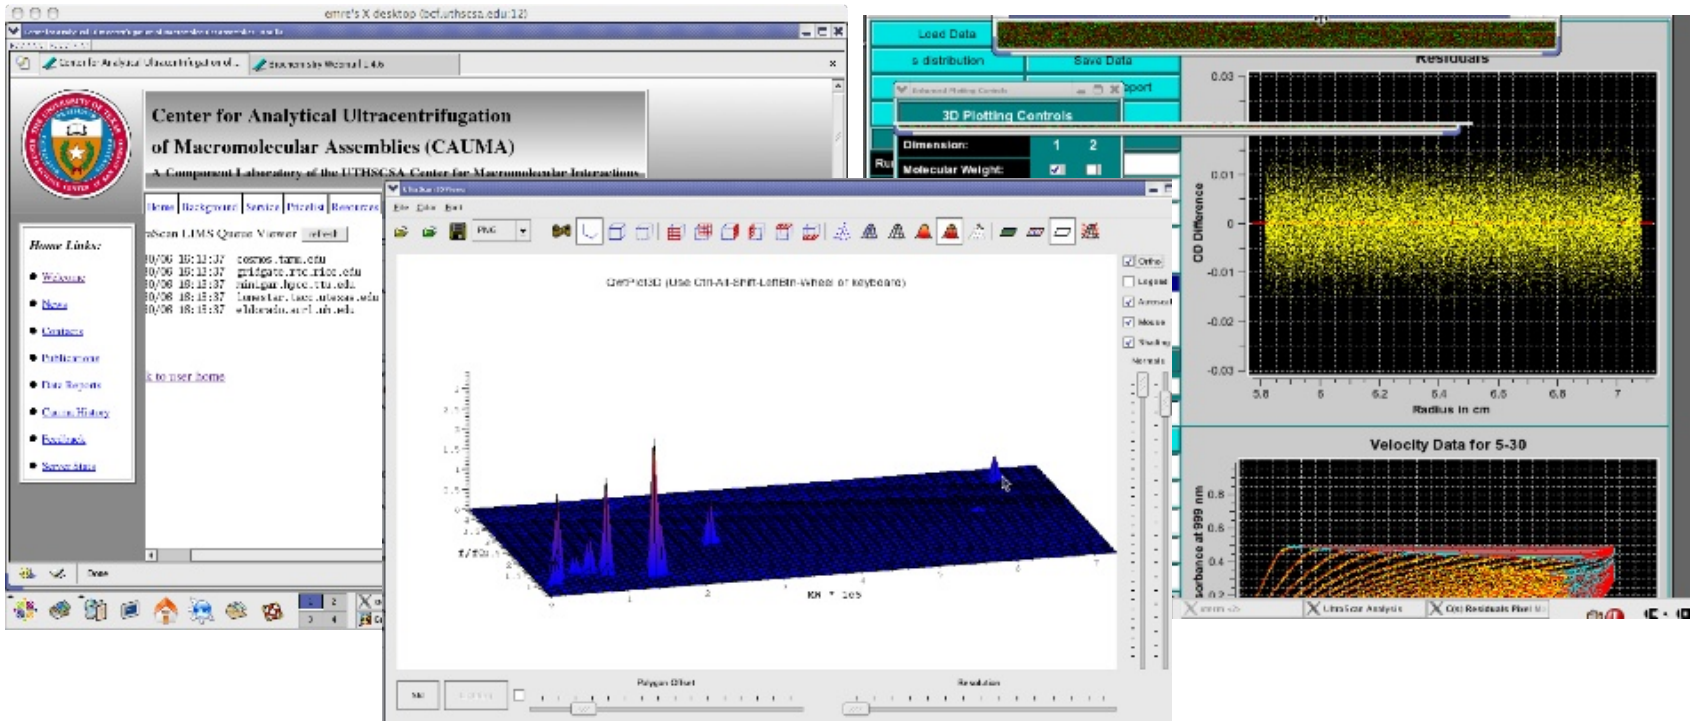

(Prof. Borries Demeler and postdoc Emre Brookes, UTHSCSA: "*We can do science we never did before!*")

Atmospheric Modeling

Collaboration with Texas Mesonet and related researchers

http://mesonet.tamu.edu

Collaborators: UH, TAMU, TTU, Rice

Application goal: Robust, fault-tolerant processing of WRF and MM5 air quality and atmospheric modeling data on a 24x7 basis.

1. Initial engagement: Movement of data to and from grid-enabled execution locations.
2. Medium-term goal: Regular execution of grid jobs on data.
3. Long-term target: ***Fault-tolerant job flow and execution*** on continuous weather data stream for air quality and atmospheric modeling codes.

TIGRE outcomes and goals:

1. Multi-institutional collaborations.
2. Disaster recovery and model delivery reliability.
3. Demonstration of service robustness.

Atmospheric Modeling Application: Texas Mesonet Surface Wind Tracking

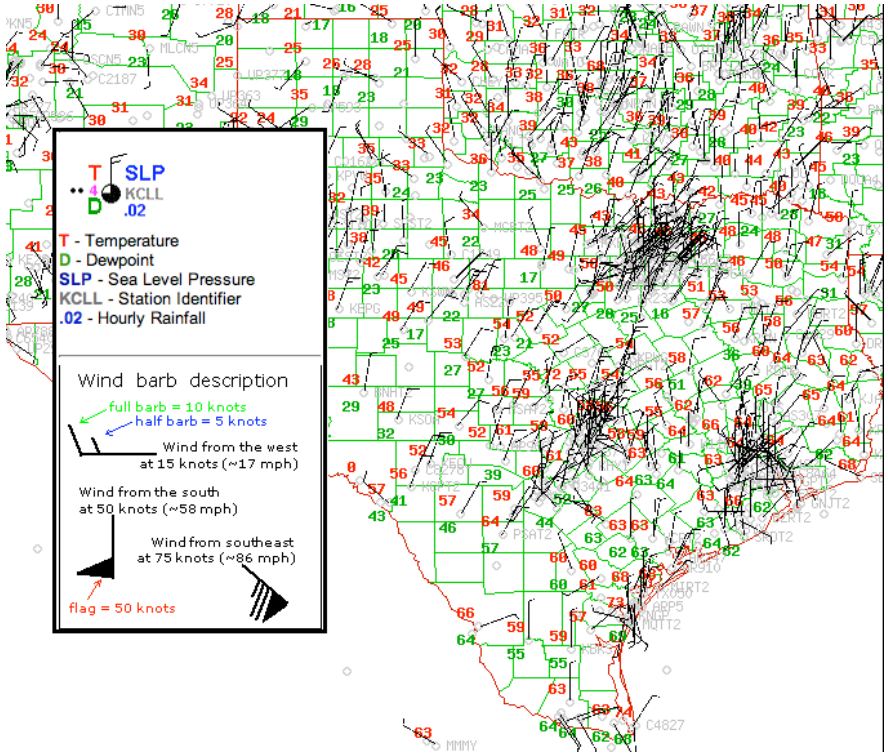

The Texas Mesonet provides Mesoscale monitoring for:

- Citizens of Texas
- Researchers
- Emergency Planners/Mgrs
- Public Utilities

TIGRE metascheduler and data services (UCAR Local Data Manager):

- Produce and migrate data
- Schedule jobs
- Inspect data flow
- Analyze results at multiple locations via grid methods

http://mesonet.tamu.edu

http://unidata.ucar.edu/software/ldm

TIGRE

Atmospheric Data Flow

Subscriptions at each site to and from repository via UCAR Local Data Manager (LDM).

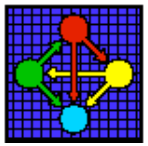

http://www.unidata.ucar.edu/software/ldm/

Provide data from "nearest" location to jobs via gridftp-enabled interfaces.

> (Any grid job can always find a local and alternate locations for the input data, and can write to multiple redundant locations)

Researchers can inspect data files through GridFTP directly on resource or through grid-enabled data GUI.

Present LDM/GridFTP deployment (10/2007):

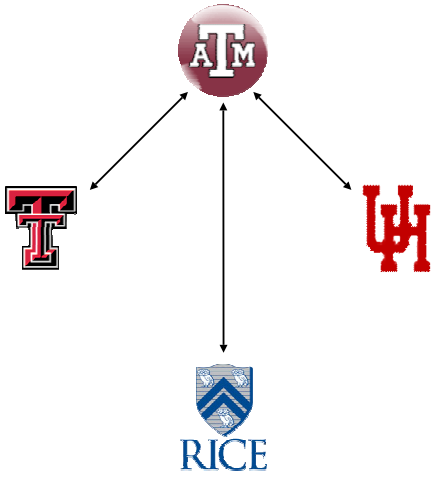

(Other data flows also possible)

Energy Exploration

Ensemble Kalman Filter (EnKF): http://enkf.nersc.edu

Collaborators: TTU, TAMU, UH and UT Austin

Application goal:

1. Characterize petroleum reservoirs for optimal development and management.
2. Apply over 50 geological models, make use of 50-60 years of experimental data in time steps ranging from 3 months to 1 year.
3. Support high-strength industry data security requirements.

TIGRE outcomes and goals:

1. Multi-institutional collaborations.
2. Demonstrate energy exploration via common software paradigm.

EnKF Workflow

(From Grid2007 demo)

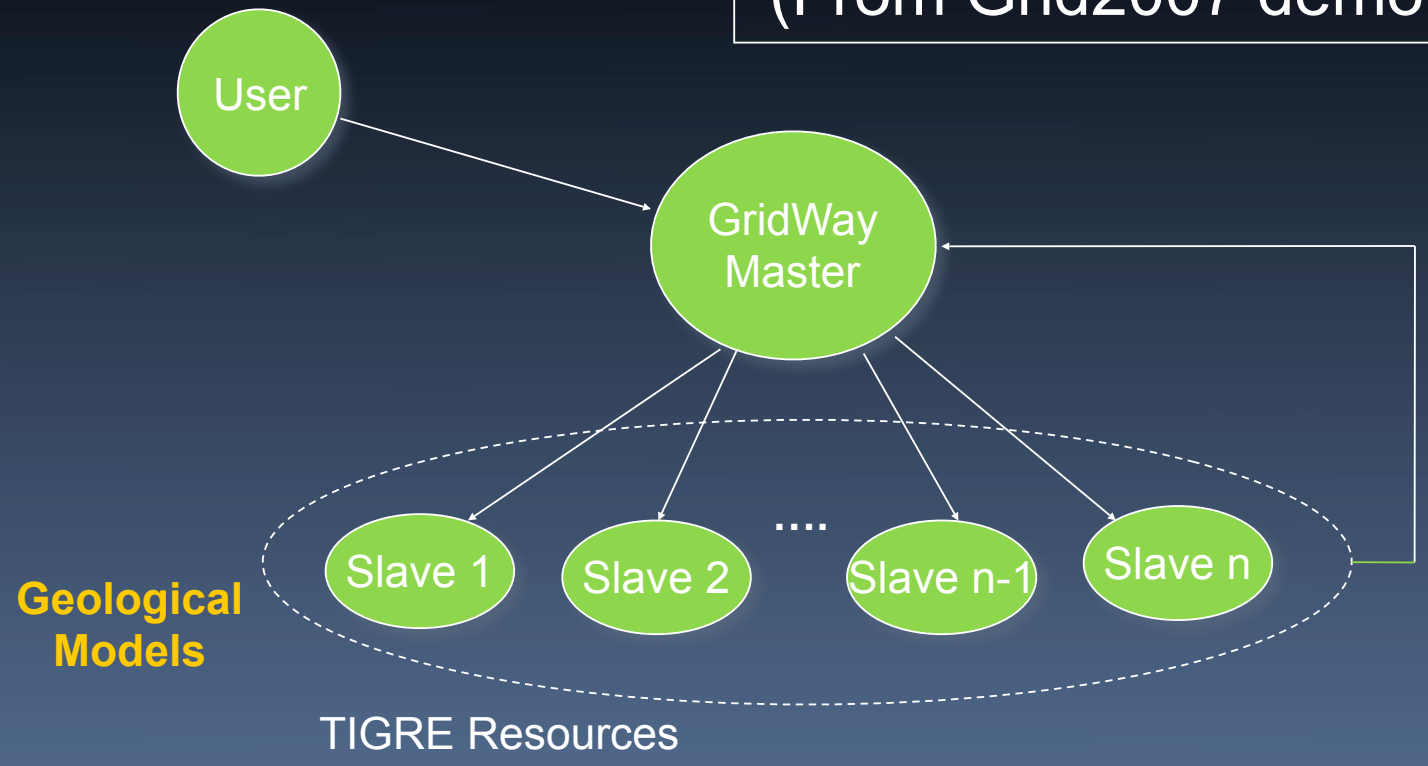

Degree of parallelism

Subject to availability of ECLIPSE licenses and CPUs

SUPERCOMPUTING 2007 DEMOS

Demonstrations in each of the above areas were carried out over a period that extended across quarterly milestones and culminated in on-floor live demonstrations at the SC2007 conference in Reno, Nevada, and at Grid2007 in Austin, TX.

Of the three different TIGRE application areas targeted, two remain in production and continue to consume cycles on the demonstration resources put together for TIGRE. The third, atmospheric modeling, is in a highly ready state and could easily be put into production upon demand from researchers.

Other application areas can easily be added with the TIGRE software. New systems can also be added with minimal effort. SURAgrid has also adopted the TIGRE stack on a pilot basis.

Future Directions and Options

Near-term actions we can now pursue:

1. Open opportunities for participation in TIGRE to all HiPCAT institutions.
2. Open investigations into pursuit of extension of TIGRE software to connect to cycle-scavenging (campus) grids.
3. Pursue specialty hardware and development directions for TIGRE.

Long-term directions/possibilities:

1. Large-scale cycle-scavenging grid.
2. Development front-end for web-services portal and job submission to Teragrid, OSG, SURAgrid and TIGRE.
3. Training, CI engagement, outreach, workshops, etc.

TIGRE's Role in Regional Computing

The TIGRE project has a state-oriented partner for high energy physics computing, THEGrid (Texas High Energy Grid) for research in high energy physics, nuclear physics, astrophysics and radiation therapy modeling (also a HiPCAT project).

THEGrid partners have agreed in principle to share access to local computational resources through TIGRE software.

TIGRE software is OSG-compatible.

In addition, Teragrid has agreed to make 5% of TACC's new system available to Texas researchers by the usual allocation application methods on a competitive basis. We are considering the best methods to pursue this on a science-field-specific basis.

Other Regional Projects

Additional projects are also in place or under development to aggregate resources across a number of institutions on a regional and/or state-wide basis.

SURAgrid (Southeastern US Research Association) and TIGRE (Texas Internet Grid for Research and Education) are examples of science-oriented aggregation around regional interests.

Our participation in these projects has been heavily influenced by, and has benefited from our participation in OSG for CMS Tier-3 computing. We are also investigating ways to make use of our membership in the above projects to enhance and extend our computing capabilities.

SURAgrid:

Lowering the barriers for deploying and utilizing local, regional and national CI

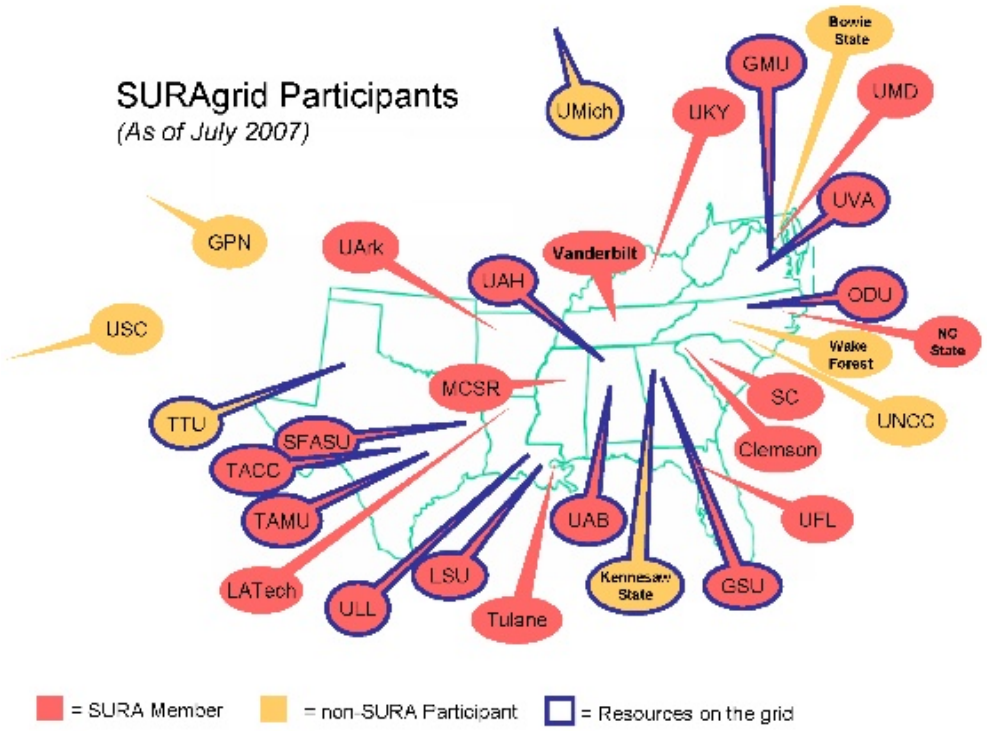

- 30 participating institutions
- Shared accessible grid computing environment
- Access to group negotiated discounted HPC systems
- Enabling CI supported research & education
- On-Ramp to National CI projects

See: www.sura.org/suragrid for more information.

Open Science Grid in 2008:

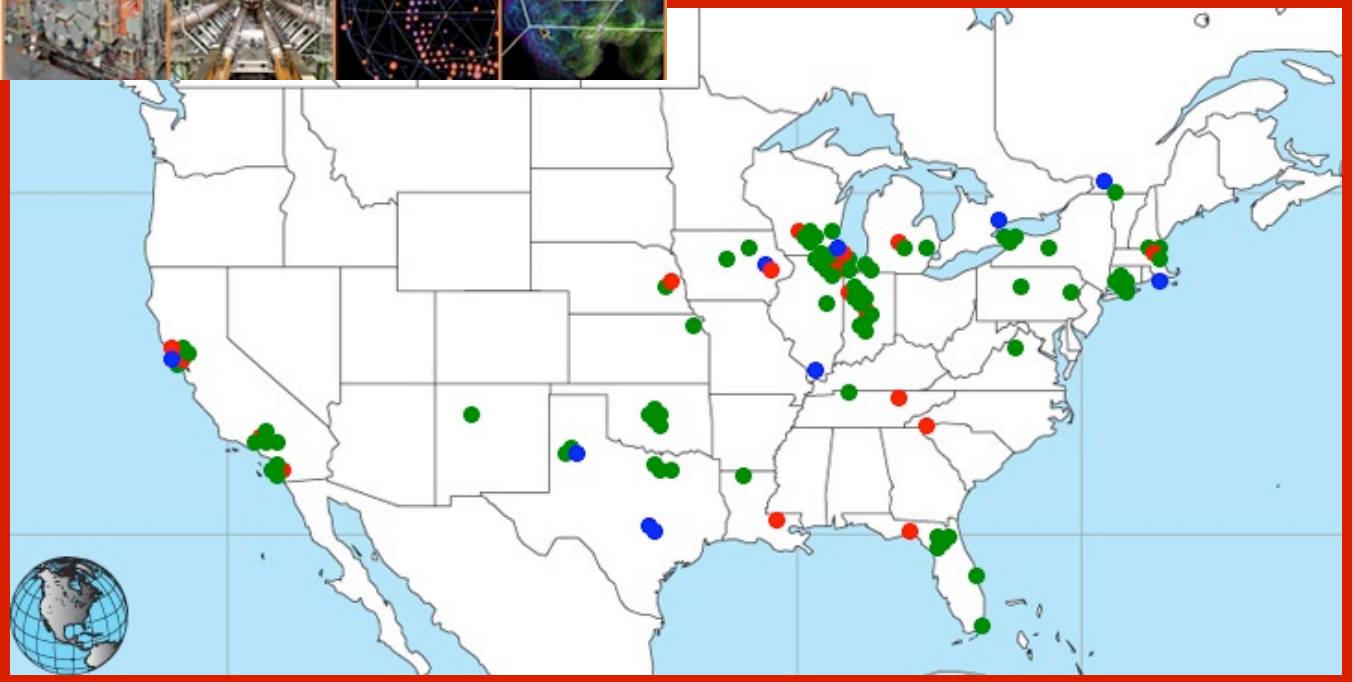

30 Virtual Organizations, hundreds of members, >10,000 cpus.
Total aggregate capability across all VOs ~80 TF

TeraGrid High Performance Computing

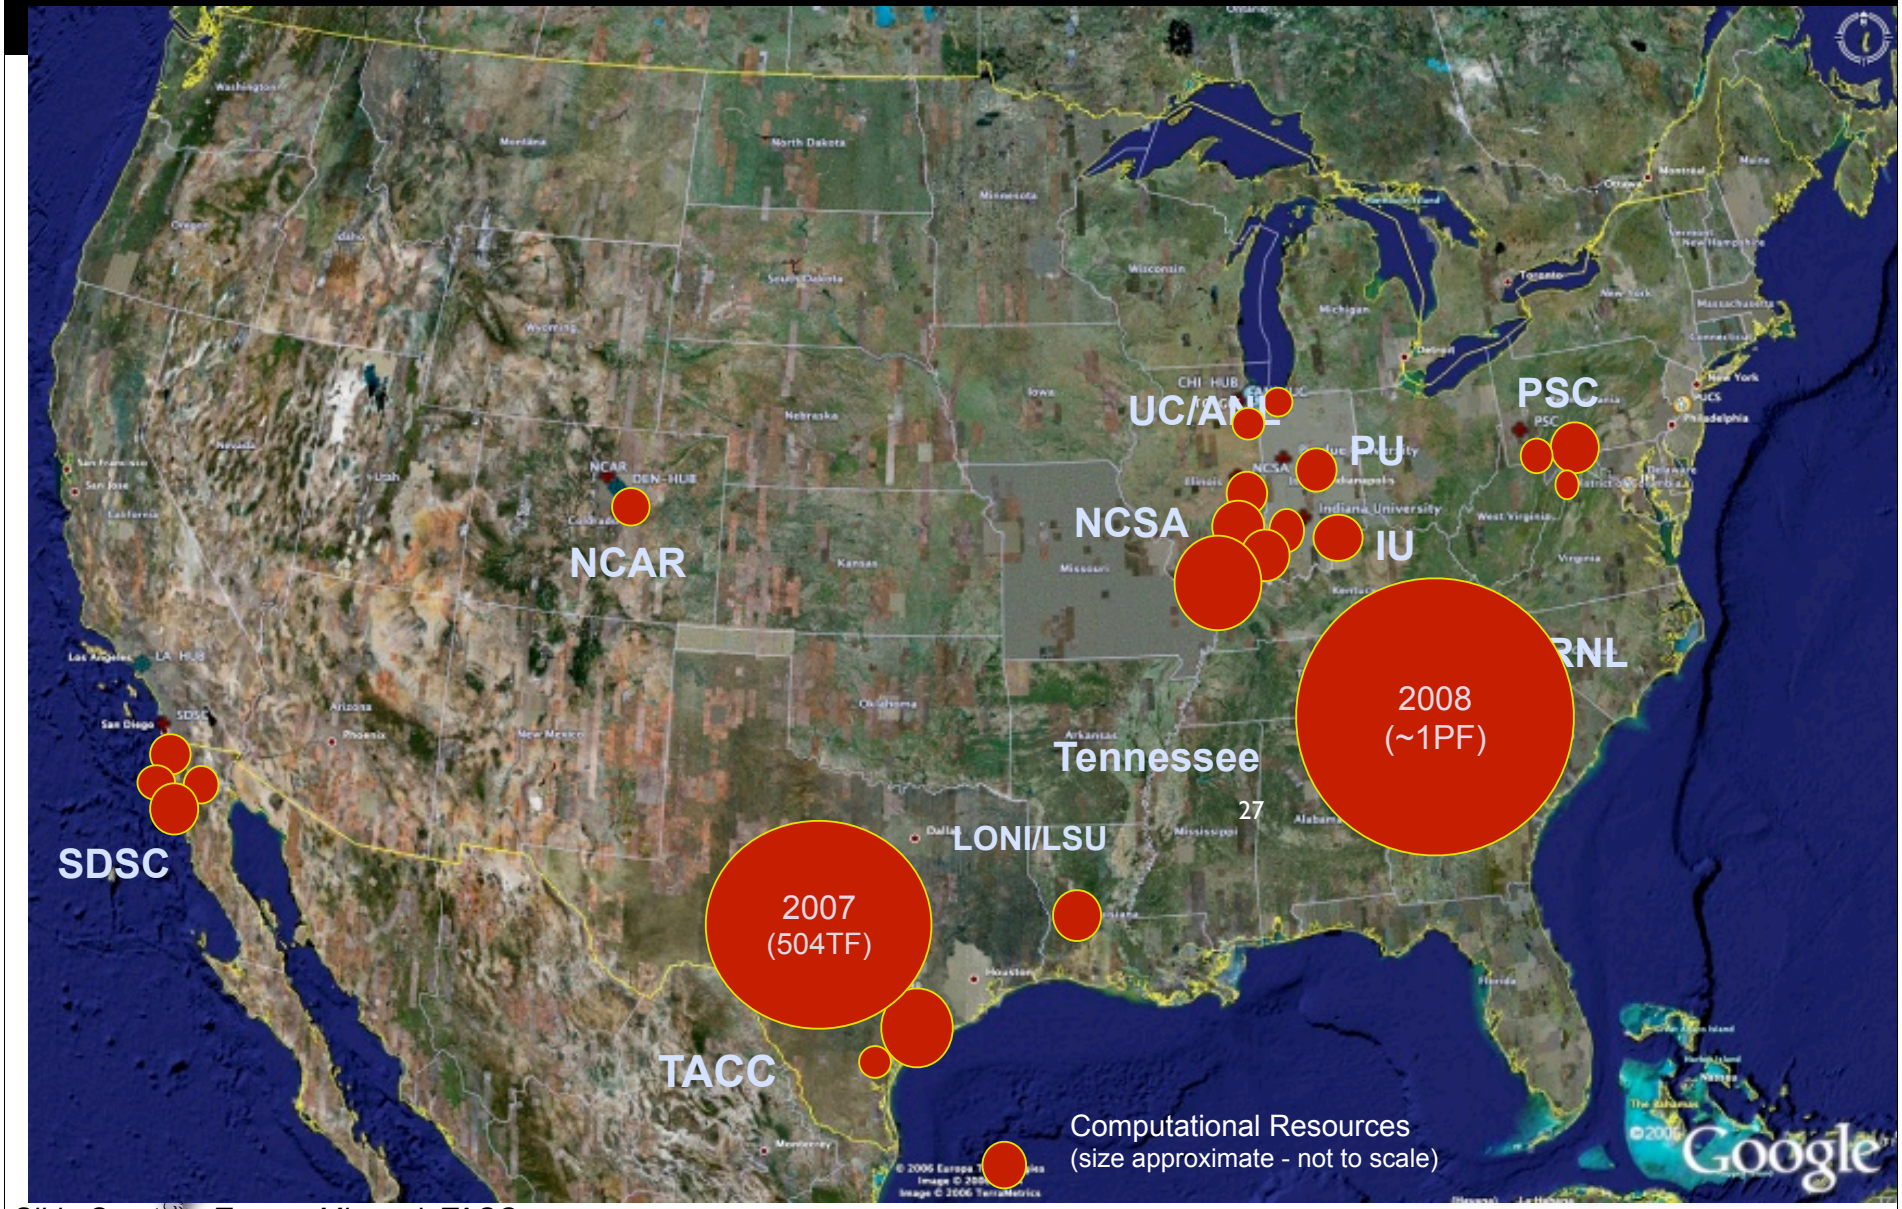

Slide Courtesy Tommy Minyard, TACC

Summary and Conclusions

- TIGRE set out to involve staff and researchers at five major HiPCAT universities directly in development and deployment of modern grid software. We succeeded in this effort.
- Demos were conducted in *all targeted application areas,* and are now complete. The resulting simple, easily maintained TIGRE software stack serves as an entry point for researchers and their institutions into high performance computing, and also provides tools compatible with use of other large-scale projects such as LCG, OSG and Teragrid.
- TIGRE has a great potential future in *education and outreach*, e.g. bringing high performance computing to institutions, people and areas not yet familiar with its use. My personal opinion is that this is one of the best future uses of TIGRE.