TIGRE Status as of Feb. 2008

A HiPCAT Project

TIGRE

Texas Internet Grid for Research and Education

Alan Sill, TTU
for the TIGRE
Development Team

HiPCAT Meeting UT Arlington Feb. 29, 2008

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.

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**		

TIGRE

HiPCAT Meeting
Feb. 29, 2008

Texas Internet Grid for Research and Education

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.

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
this 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 has had 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.

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 Teragrid and IGTF 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.

Targeted Applications Summary

Application Example: Biosciences and Medicine

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

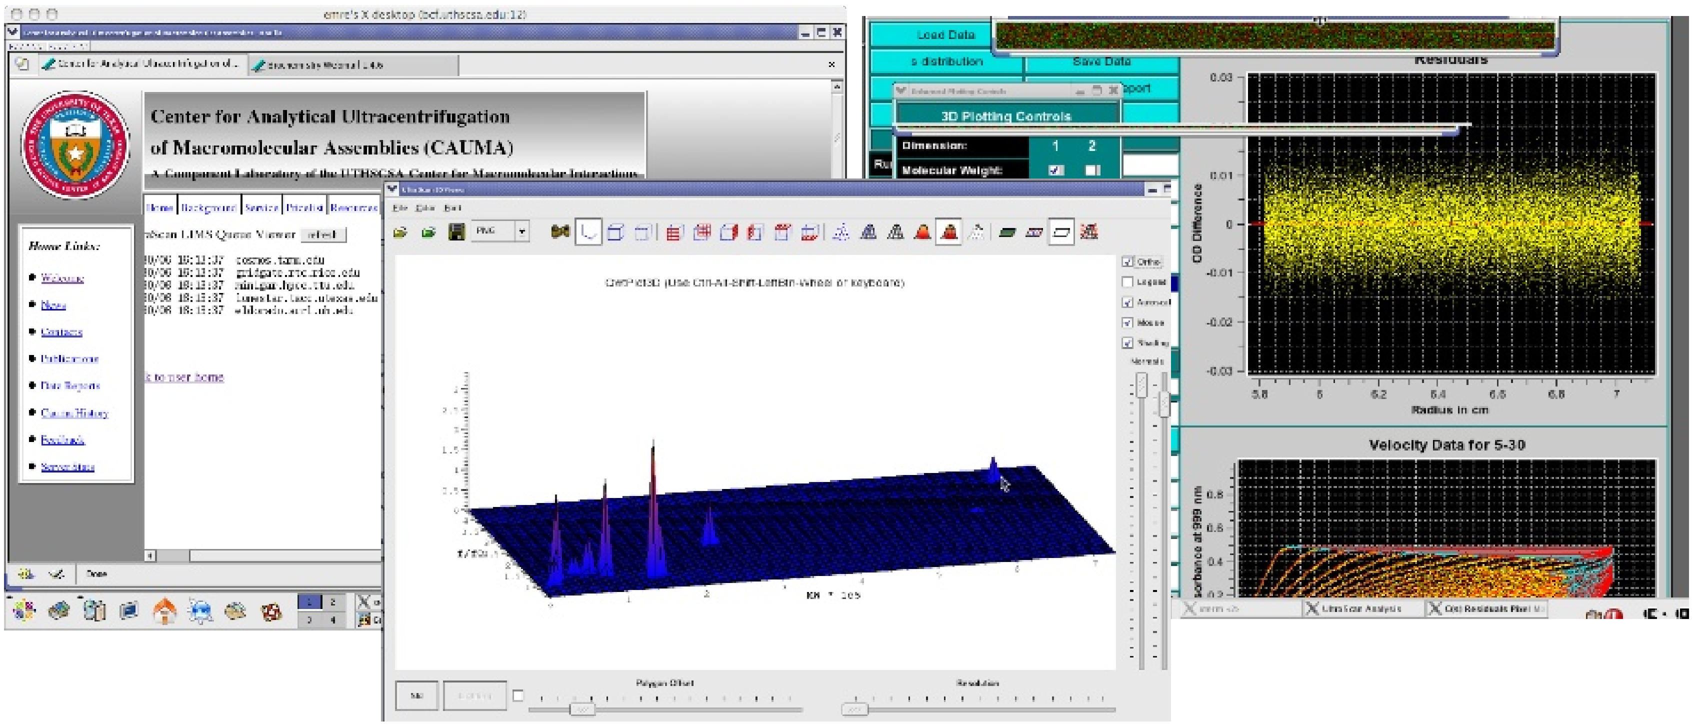

(Prof. Borries Demeler, UTHSCSA, will talk more about this application and its results in the next talk.)

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 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):

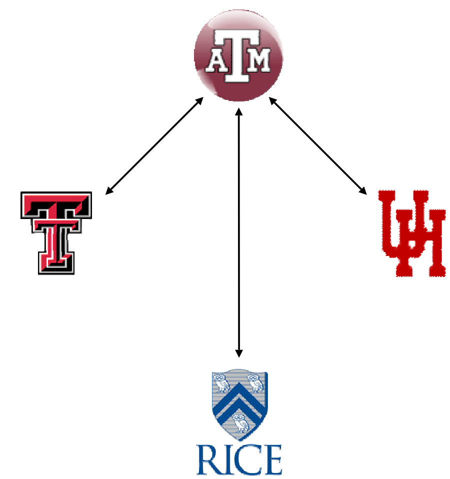

(Other data flows also possible)

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)

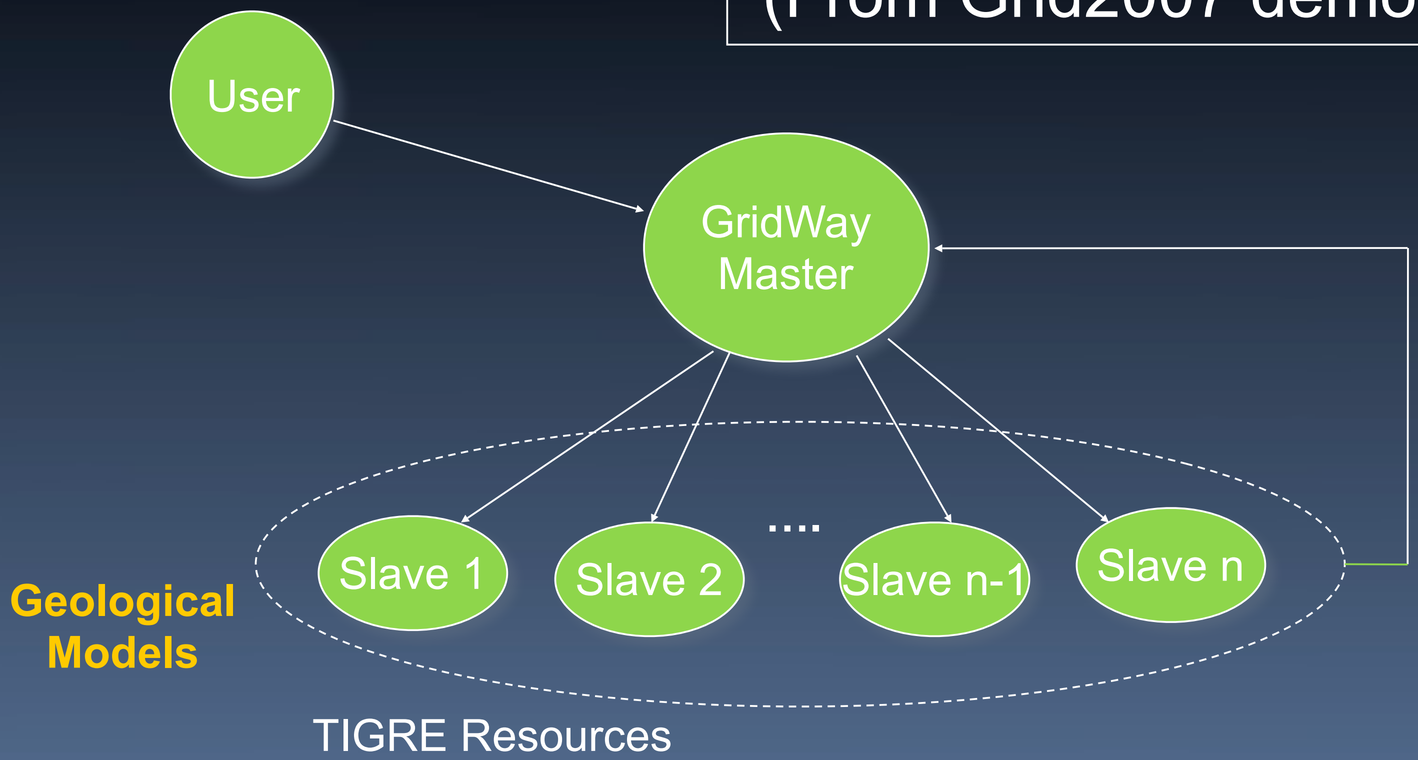

Degree of parallelism

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.

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.

Teragrid Compatibility

Teragrid uses the *same* authentication and authorization system (in its fundamentals) as is used in TIGRE. In fact, the certificates issued by TACC for TIGRE can work if authorized on any VDT-based or other compliant software stack. IGTF accreditation of the new TACC CA will only increase the future compatibility of these credentials.

The basic X.509-authenticated file transfer capabilities, gsissh login and other major features of TIGRE are common to all globus-based grid projects and can work without alteration (requiring only local allocation and authorization) with lonestar, ranger, as client tools for any such resource. TIGRE is essentially *compatible* with Teragrid.

Project Shortcomings

As a development project, TIGRE did have certain shortcomings that hold it back from full production status:

1. The monitoring system used is very light-weight, and not well coupled to either of the prototype metaschedulers.
2. The level of resources provided by project participants is uneven, and not of sufficient scale to serve many users.
3. The number of participating institutions was small, limited to only slightly more than the five developing institutions.

None of the above limitations is fundamental, though, and motivation to solve them can easily be found by growing the number of participants and applications.

Action Items

The following action items were put forward during the past few weeks at TIGRE steering committee meetings:

1. Put out a call (e.g. at the this HiPCAT meeting) for TIGRE project participation. Minimal "ante" would be about 0.5 TFlops per site.

2. If successful, decide how to formulate a TIGRE allocations board and what its functions will be. (For example, review of applications for usage, discussion of methods of enforcing usage limits, etc.) This board should be given some guidance from the rest of TIGRE in terms of policy, for example, whether to limit TIGRE allocations to Texas users and their collaborators, etc.

3. Clarify the role of cycle-scavenging serial systems in the above. Pursue research in this area. Possibilities include Condor, BOINC, home-grown software, etc. Demonstrate that TIGRE software can be used to submit to such cycle-scavenging grids. (In progress.)

4. Decide whether to submit a TIGRE development allocation request to TACC for development, application porting, etc. (not for large-scale running). Large-scale science users should be pointed to the existing allocations request process.

5. Delineate what you get when you install the TIGRE stack on a local system. (Done.) Answer: basic grid-enabled tools for job submission, file transfer and monitoring.

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.

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* 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.