TEXAS TECH UNIVERSITY
Information Technology Division™

High Performance Computing Center

TIGRE Software Stack: Status and Goals as of 3/2007

Alan Sill

TIGRE Senior Scientist, TTU

SURAgrid

SURAgrid All-Hands Meeting
Wash. DC Mar. 14-16, 2007

Background and History

- 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.

- TIGRE was proposed as one of HiPCAT's first projects. (Rice, Texas Tech. U. of Houston, Texas A&M, and UT/TACC.)

- Combined by Legislature into one funding bill with LEARN: (Lone Star Education and Research Network)

TIGRE

$2.5M

Grid Software Development

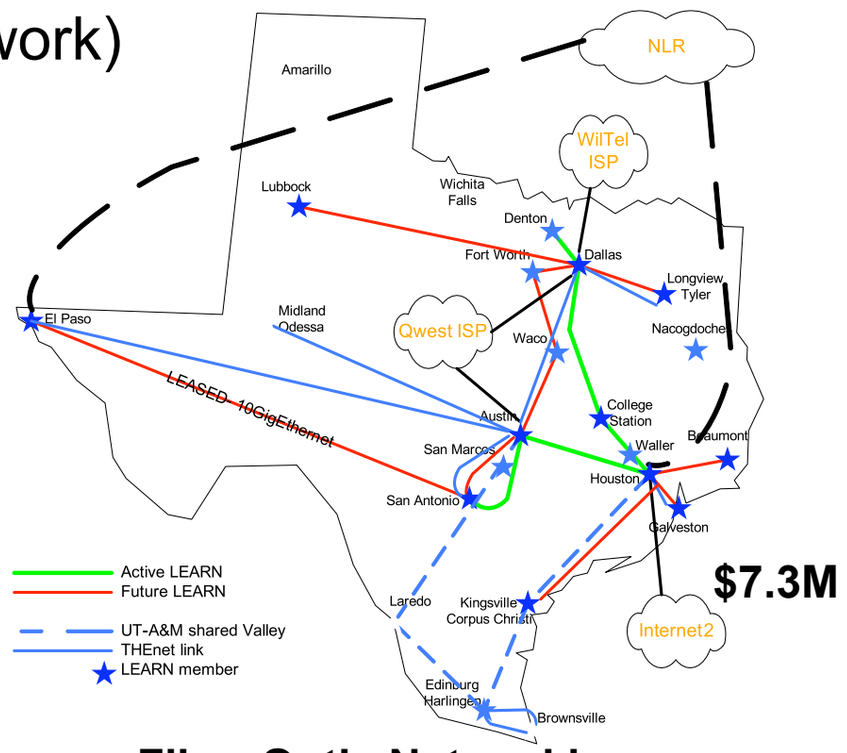

$7.3M

Fiber Optic Networking

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 manifold 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 to Date

Goal is to achieve "quick build" toward working status:

YEAR 1

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 deployed & demonstrated ✔
Applications in three application areas *(in progress)*

Q2:
Project-wide global grid scheduler deployed *(in progress)*

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

Q4:
Complete hardening of software
Complete documentation
Finalized procedures and policies to join TIGRE
Demonstrate TIGRE at SC

... So Far Making Excellent Progress!

Server and Client Software Stack Contents

- Contents of the client software stack:
  - Globus 4.x pre-web services and web services clients
  - GSI OpenSSH client
  - UberFTP
  - MyProxy client
  - Condor-G
- Contents of the cluster/compute server stack:
  - Globus 4.x GRAM4 (web services) server
  - GSI-OpenSSH server (note: as of this version works with PRIMA)
  - GPIR monitoring (added manually after the VDT pieces)
  - The client software stack
- Storage is not yet resolved in TIGRE, but initial investigations using SRM/DRM, SRB/iRODS, etc. are in progress.

Details of TIGRE Software

- Both server and client software stacks available.
- Based on the VDT, working in close cooperation with OSG.
- Uses a simplified set including GSI-OpenSSH, omitting much monitoring in favor of GPIR for lightweight status reporting.
- Installed on systems at all five primary TIGRE institutions and also several other locations throughout the state.
- Goal is “one page” installation instructions that can be implemented quickly by newcomers.
- Authentication via X.509 (IGTF + TACC CA recently accredited); authorization local mostly via grid-mapfiles. (TTU uses GUMS/PRIMA.)

TIGRE User Portal http://tigreportal.hipcat.net

Grid Information Browser

Refresh

Parallel Computing Resources

Name	Institution	Department	System	CPUs	Peak GFlops	Memory GBytes	Disk GBytes	Status	Load	Jobs
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	↑		40R-38Q-0O
Eldorado	University of Houston	Texas Learning and Computation Center	Eldorado Itanium2 Cluster	124	472	256	2232	↑		0R-0Q-0O
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	↑		70R-0Q-168O
Minigar	Texas Tech University	High Performance Computing Center	Dell Linux Cluster	32	230	64	70	↑		0R-0Q-0O
RTC	Rice University	Computer and Information Technology Institute	HP Itanium II Linux Cluster	290	1044	596	7000	↑		31R-0Q-0O
Weland	Texas Tech University	High Performance Computing Center	AMD Athlon MP 2000+	64	69	64	780	↑		0R-0Q-0O
			Total:	**5933**	**57547**	**11761.4**	**128146.3**			

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	**60** / 60	**60** / 60	49	1494	🔍	🔍
			Total:	**60 / 60**	**128 / 128**	**63**	**2795**		

Jobs Key: **#R** - Number of Jobs Running, **#Q** - Number of Jobs Queued, **#O** - Number of Jobs in an Other State

How to Get and Install It

- Download and install from instructions on the TIGRE portal:

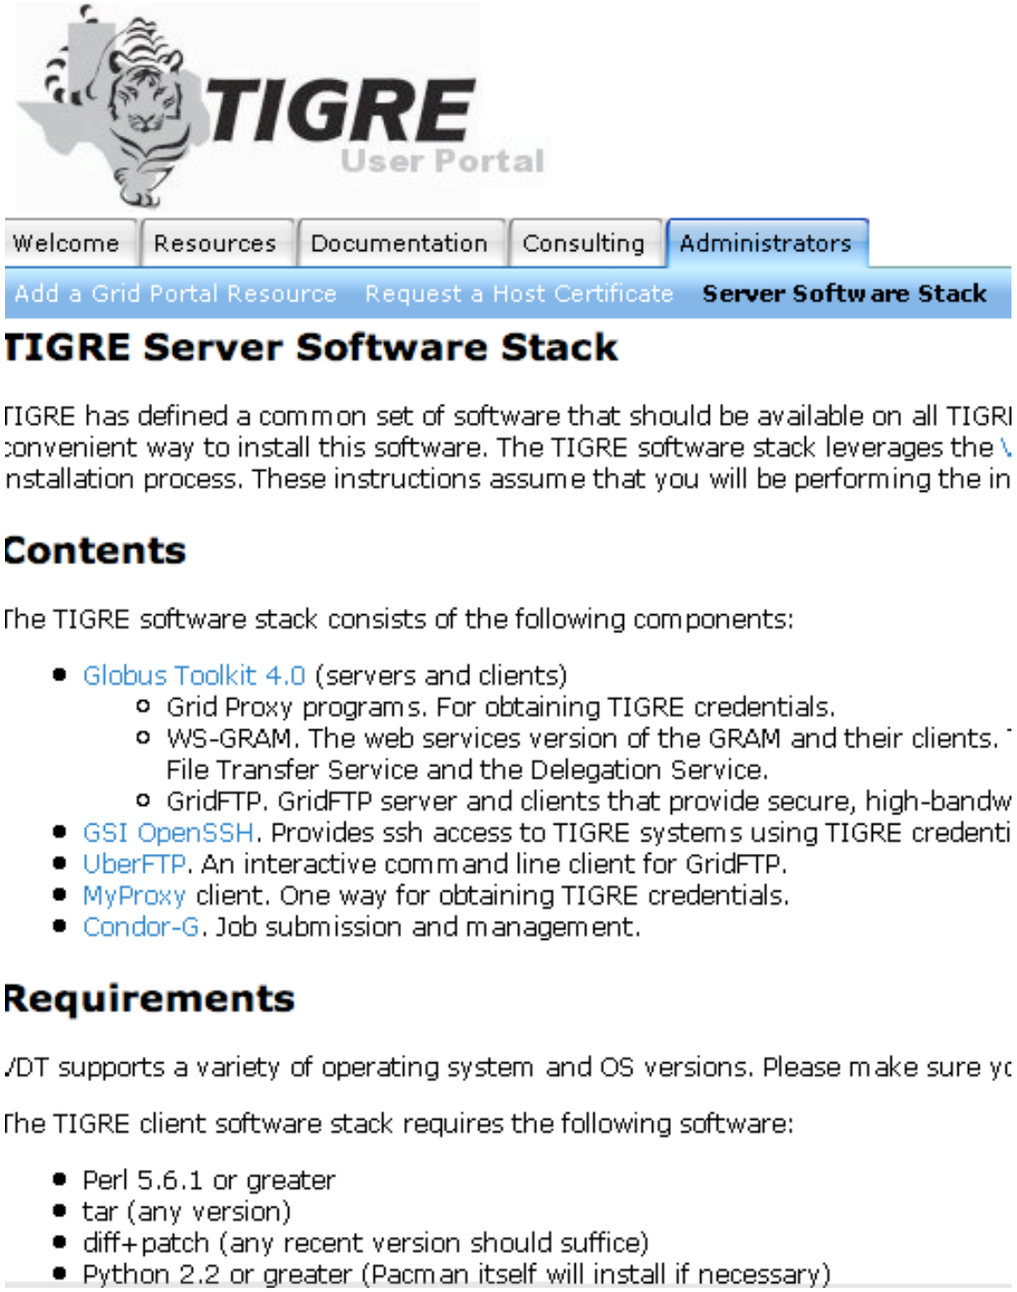

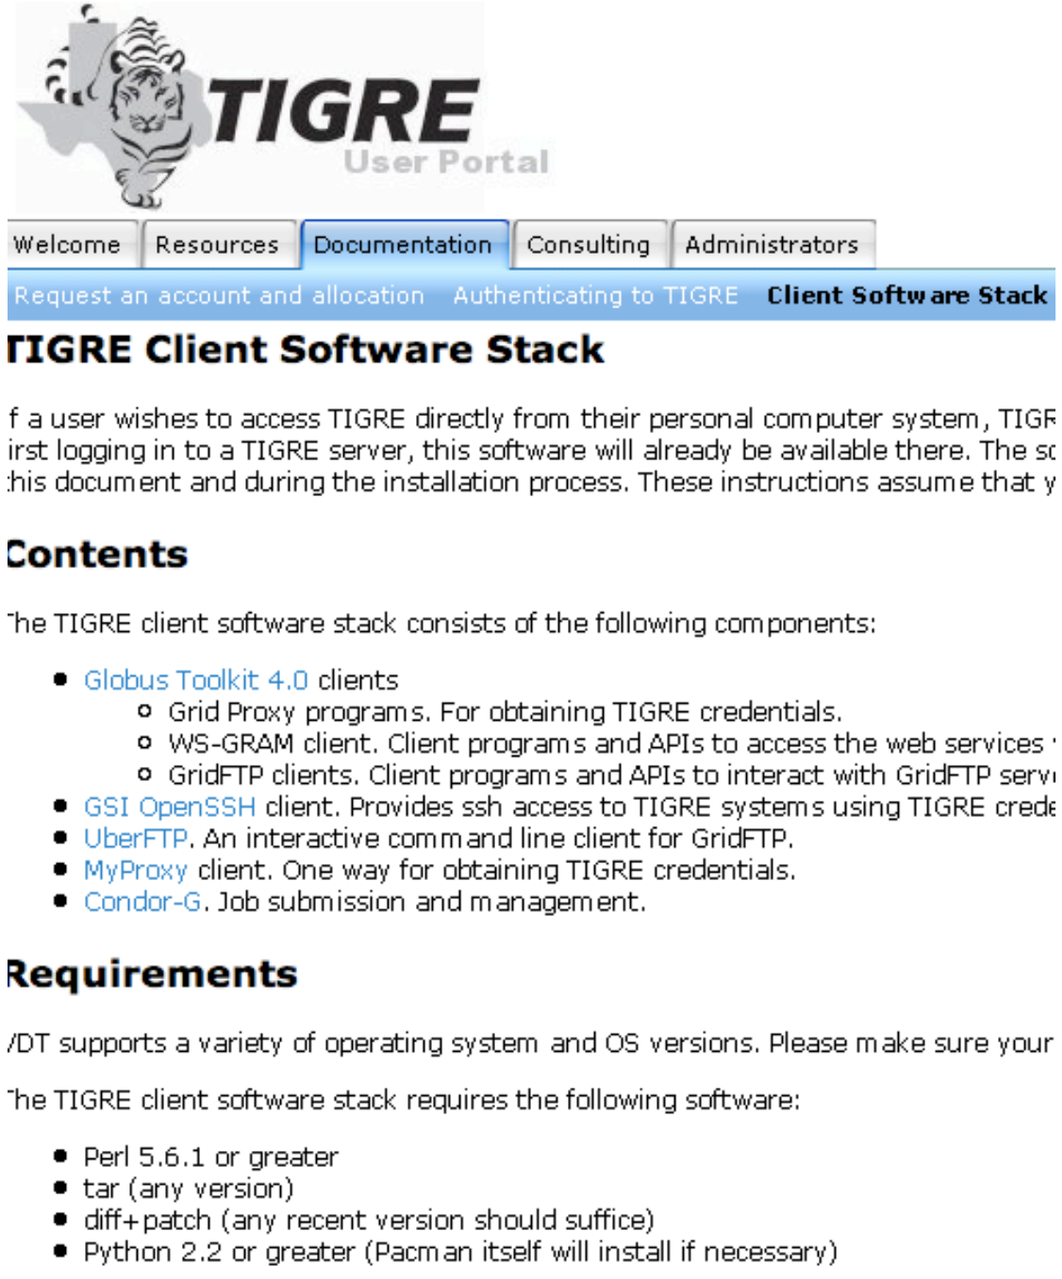

Details of Platforms and Capabilities

- Installation of the Gatekeeper (gram2) or Container (gram4) features and of many of the server and client functions are only supported on Linux platforms.
- Unofficially many client functions work also on Windows and Mac OS X. (Separate installation by hand may be needed.)
- Linux system can serve as the front end for grid submissions to a non-Linux batch scheduler (PBS, Condor, LSF, SGE). Might be the best way to support non-Linux execution.
- TIGRE only uses GT4 Web Services (gram4) for grid jobs, but the stack supports pre-WS (gram2) as an option.
- All non-job-oriented services use GSI for AuthN/AuthZ, which is compatible with the SURA approach.

Conclusions

- TIGRE's goals to involve support staff and researchers at major Texas universities in production-quality grid software are similar to those of the SURA effort.
- We have put together a simple software stack based on the VDT that can serve as an entry point for researchers and institutions into other large-scale national, international, regional and local projects. This is a simpler subset than used by the OSG, with some additional components for our needs.
- Demos in targeted application areas using this software stack are complete or in progress.
- SURA has similar goals and could use the TIGRE stack either as-is or by interacting with us regarding needed changes.
- TIGRE is open to this type of interaction.

TIGRE

Texas Internet Grid for Research and Education

RICE

HiPCAT
High Performance Computing Across Texas

A HiPCAT Project