

PEGrid: The Path from Research Project to Grid Computing In The Real World

Alan Sill, Ph.D

Senior Scientist, High Performance Computing Center

Texas Tech University

For the PEGrid Project



TEXAS TECH UNIVERSITY

Information Technology Division™

High Performance Computing Center

Goals

- ▶ The goals of this talk are NOT
 - ▶ To present particular narrow science topic or area.
 - ▶ To convince you of the value of our particular widget development.
 - ▶ To lay out the funding case for a particular subset or approach for software development.
- ▶ They are instead to give a practical example of how innovation makes the transition from funded research project into practical use in the real world, and get you to think about how that works.

Purpose

The purpose of all of this is to get you to understand, or at least think more deeply about, the methods by which research makes it out of the lab and acquires real value in ways that are NOT anticipated by funding agencies, even when they make such practical development their goal.

I want you to carry away an understanding of your interdependency as funding agencies with partnerships at universities, with individual researchers, with industries that you see only indirectly, and with the software development process as a whole.

Only then can you properly fund such projects.

Methods

- ▶ I will take examples with which I am familiar, which span the range from science computing with large scale distributed resources through development of the Open Science Grid and its associated tool set (also used in Teragrid and other projects),
- ▶ Combine this with the story of a state-funded project for economic development applied to grid computing,
- ▶ Describe how this led to practical demonstrations of computing technology in three different areas of economic interest to the State of Texas, and
- ▶ End with industry involvement today in a project called PEGrid.



THE TIGRE PROJECT: AN EXAMPLE OF A STATE-WIDE GRID



A HiPCAT Project



TIGRE

Texas Internet Grid for
Research and Education

5-University
development team
spanning 5 different
University systems



Project ran Dec. 2005 through Dec. 2007
(Funded by authorization from the Texas Legislature)

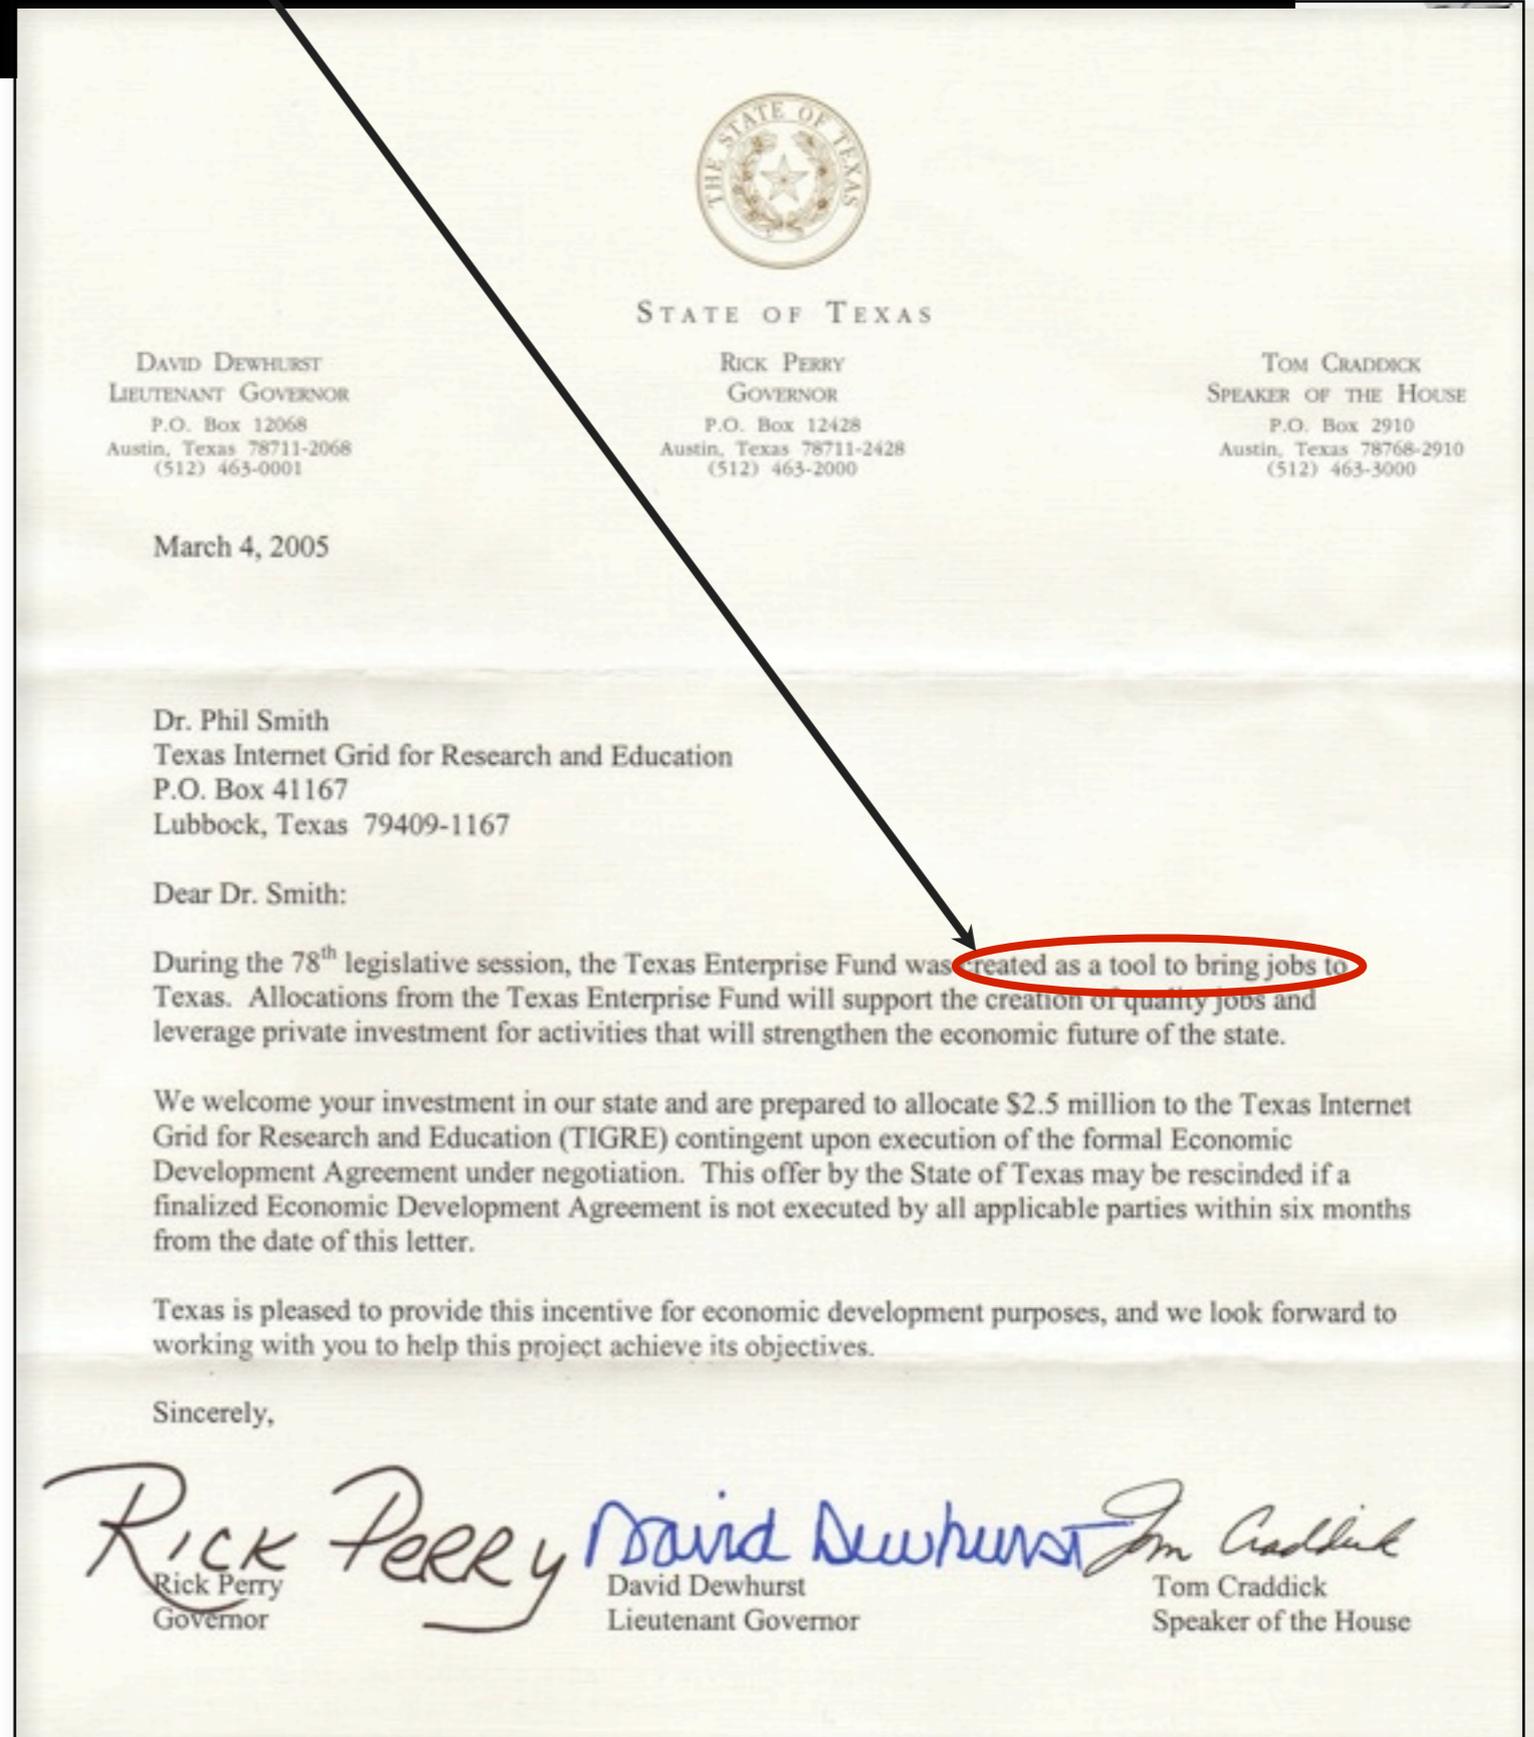
FUNDING: TEXAS ENTERPRISE FUND



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.



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 ORGANIZATION



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

- Biosciences and Medicine
- Air Quality Modeling
- Energy Exploration

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

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



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!



Where did TIGRE come from?

(It's a question worth asking...)

Antecedents

- ▶ TIGRE was the result of the desire of the group called “High Performance Computing Across Texas” (HiPCAT)

which is a self-assembled group of high performance computing centers and their directors in the state of Texas

to create a project that would demonstrate their ability to share and aggregate resources across many centers for mutual benefit and economic practicality.

Their goal was to show what is possible by combining resources.

- ▶ TIGRE itself was inspired by shared projects such as OSG, High Energy Physics distributed computing, NCAR, etc.

Participants and Partners

- ▶ TIGRE hired developers to implement their grid, but could not simultaneously achieve their given goals and create new grid software out of whole cloth.
- ▶ To achieve their goals, they evaluated and extended existing software stacks, condensing some to smaller form.
- ▶ In particular, components from the Virtual Data Toolkit used by Open Science Grid, GSI-enabled software for login from NCSA, and elements from UCAR and Berkeley were used, along with some custom development for portal and workflow control.
- ▶ We couldn't do this ourselves, and borrowed heavily from other projects.

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). At the time, most modern grid.
- Client and server software stacks made separately available.
- Goal was “one page” installation instructions that can be implemented quickly by newcomers. (This was achieved and was a project strength.)
- Authentication via X.509 (new TACC CA is now accredited by IGTF); authorization local, mostly via grid-mapfiles.
- Installed on systems at all five primary TIGRE institutions; also running at other locations throughout the state.



OTHER COMPONENTS



- Berkeley Storage Manager (BeStMan) for storage at TTU.
- GridWay workflow engine used in Petroleum Engineering demo.
- National Center for Atmospheric Research / UCAR Local Data Manager (LDM) for atmospheric data routing and distribution.
- Custom development by Texas Advanced Computing Center for portal software used in Biosciences and Medicine demo.
- License management tools from industrial partner (Schlumberger Information Systems).
- Monitoring system borrowed from Teragrid.

These systems were assembled as needed to produce demos needed for milestone progress. Custom development was not encouraged beyond that needed to achieve an operational state.



Results

- ▶ TIGRE achieved all of its milestones, and was able to deliver each of the required demonstrations on time and within budget.
- ▶ Working examples in Biosciences/Medicine, Energy, and Atmospheric Modeling were developed, delivered and used.
 - ▶ Biosciences app produced publishable results, led to funding for the involved researcher, who then used some of the funding to purchase clusters and connect them with TIGRE. Still runs components of the TIGRE stack 2 years after project completion.
 - ▶ Petroleum Engineering demo led to deep involvement with industry that continues to the present day.
 - ▶ Atmospheric modeling demo successful but not adopted by the user community; led to closer involvement in HPC by some researchers.

- ▶ Remember, the goal was to demonstrate the use of grid computing in contexts that would be of economic value to industries pursued in the State of Texas.
- ▶ One such industry is deeply involved in reservoir modeling (for oil, gas, water, etc.) and comprises a major user community for high performance computing.
- ▶ We had to learn how to approach each industry or area separately, and listen to the needs of researchers and proponents in each.
- ▶ In the area of reservoir modeling, our demonstration was in the simultaneous use of licenses across many institutions by researchers without compromising the integrity of the data or of access to the license server.
- ▶ It turned out that security was of far more interest to this user community than access to computing resources.

Other Outcomes

- ▶ Greatly improved the knowledge and technical skill on the part of staff and researchers at participating universities in methods of grid computing.
- ▶ Software stack developed by TIGRE fed back to upstream developers of tools, bug and security fixes in these components, and integration of new components into the Virtual Data Toolkit. *(Extremely positive relationships with many developers achieved!)*
- ▶ Compact version of the VDT developed by TIGRE now in use by other projects, such as SURAGrid.
- ▶ Led to close cooperation with researchers across campuses and in fields that were not familiar with grid computing.

How Did This Lead To PEGrid?

TIGRE DEMO: ENERGY EXPLORATION



Ensemble Kalman Filter (EnKF): <http://enkf.neresc.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:

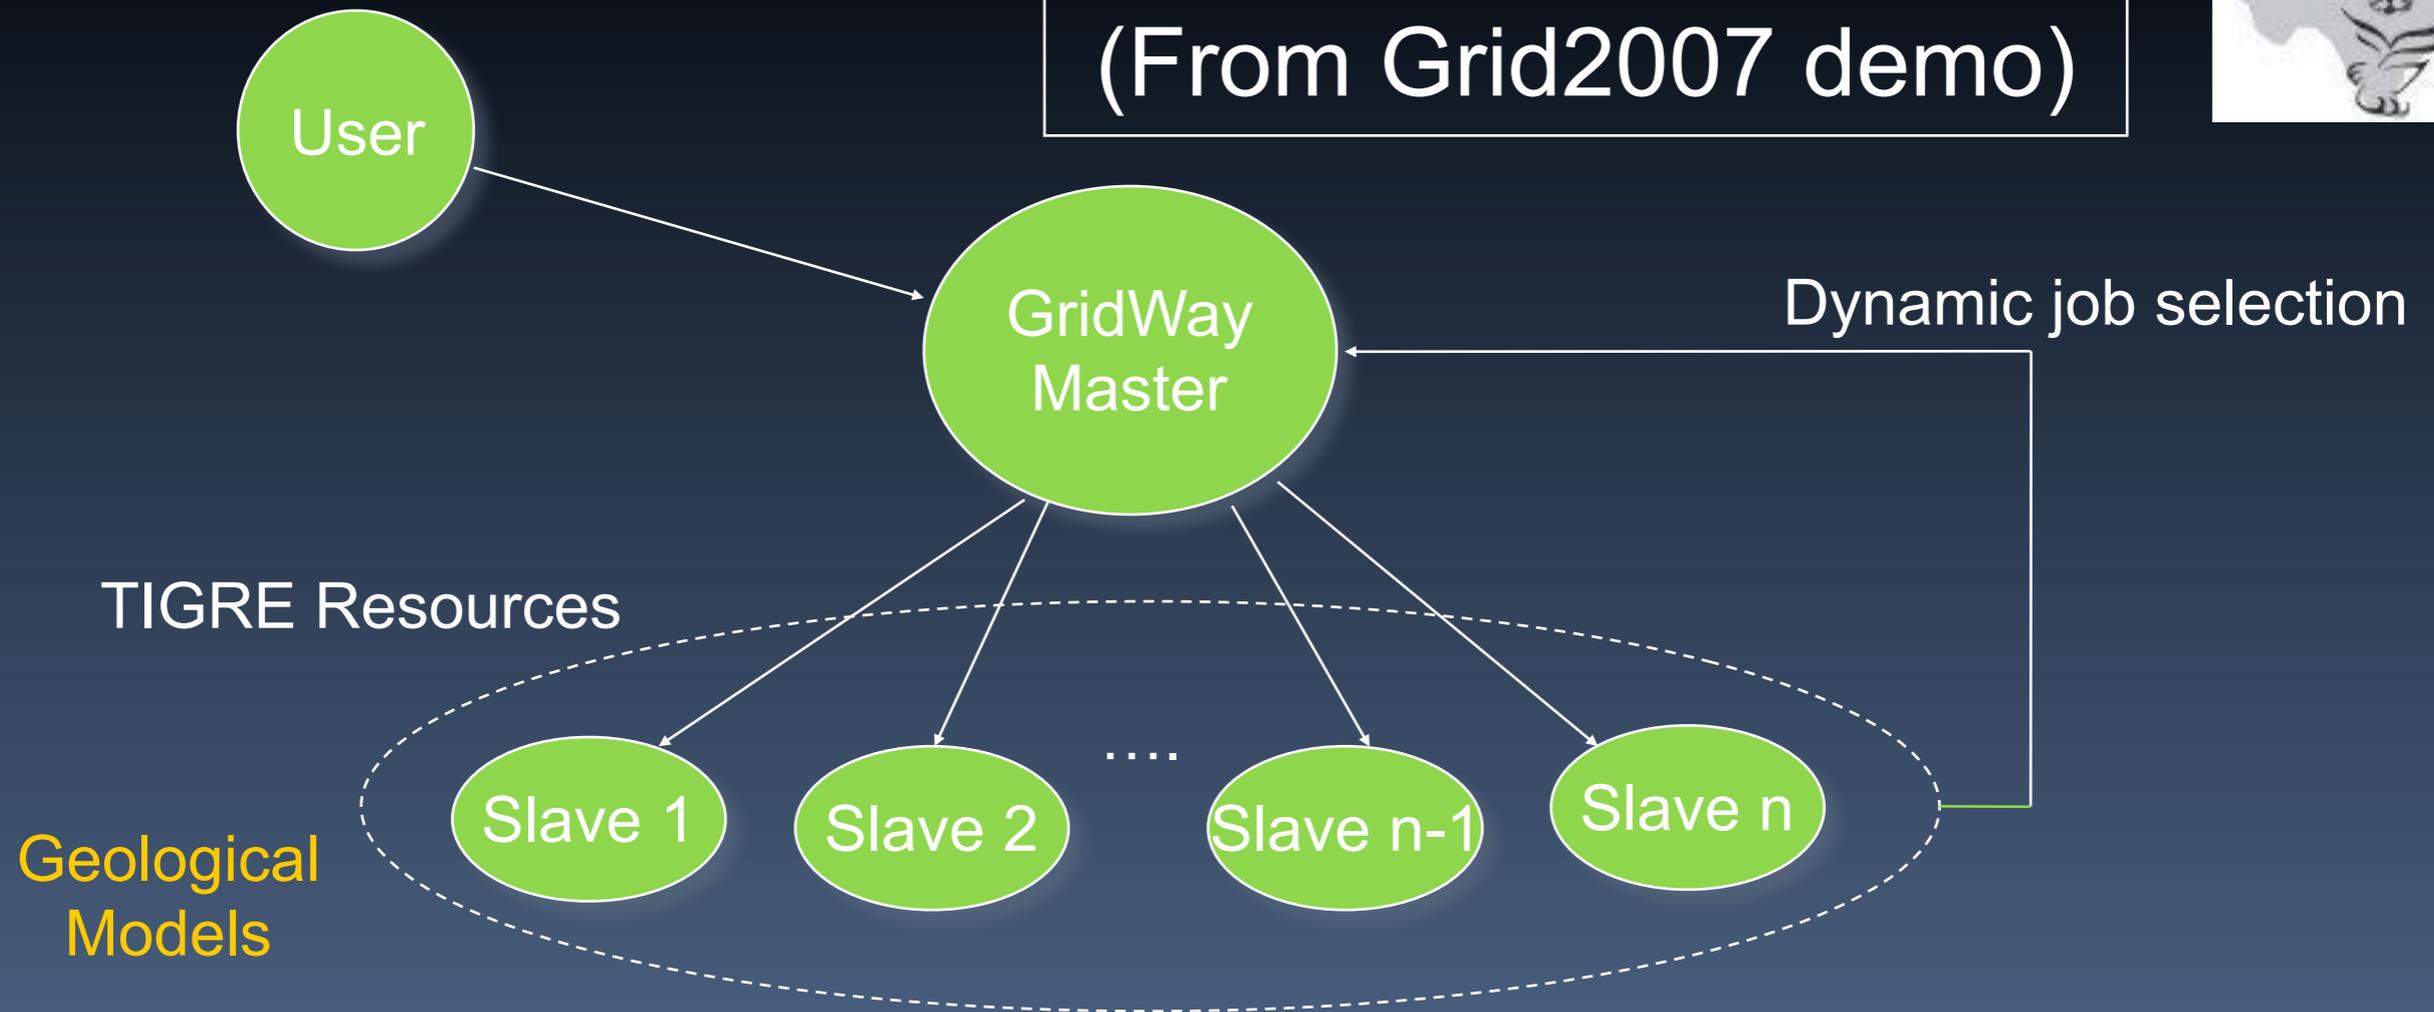
1. Multi-institutional collaborations.
2. Demonstrate energy exploration via common software paradigm.
3. Internships and collaborative research opportunities
4. Workforce development.



EnKF Workflow



(From Grid2007 demo)



Degree of parallelism subject to real-time availability of ECLIPSE licenses and Grid CPUs.

Master/slave processing steps managed dynamically to maximize licensed vs. non-licensed component use and throughput; system routing transparent to user.

Petroleum Engineering Grid (PEGrid): Outcome of TIGRE Work With Industry



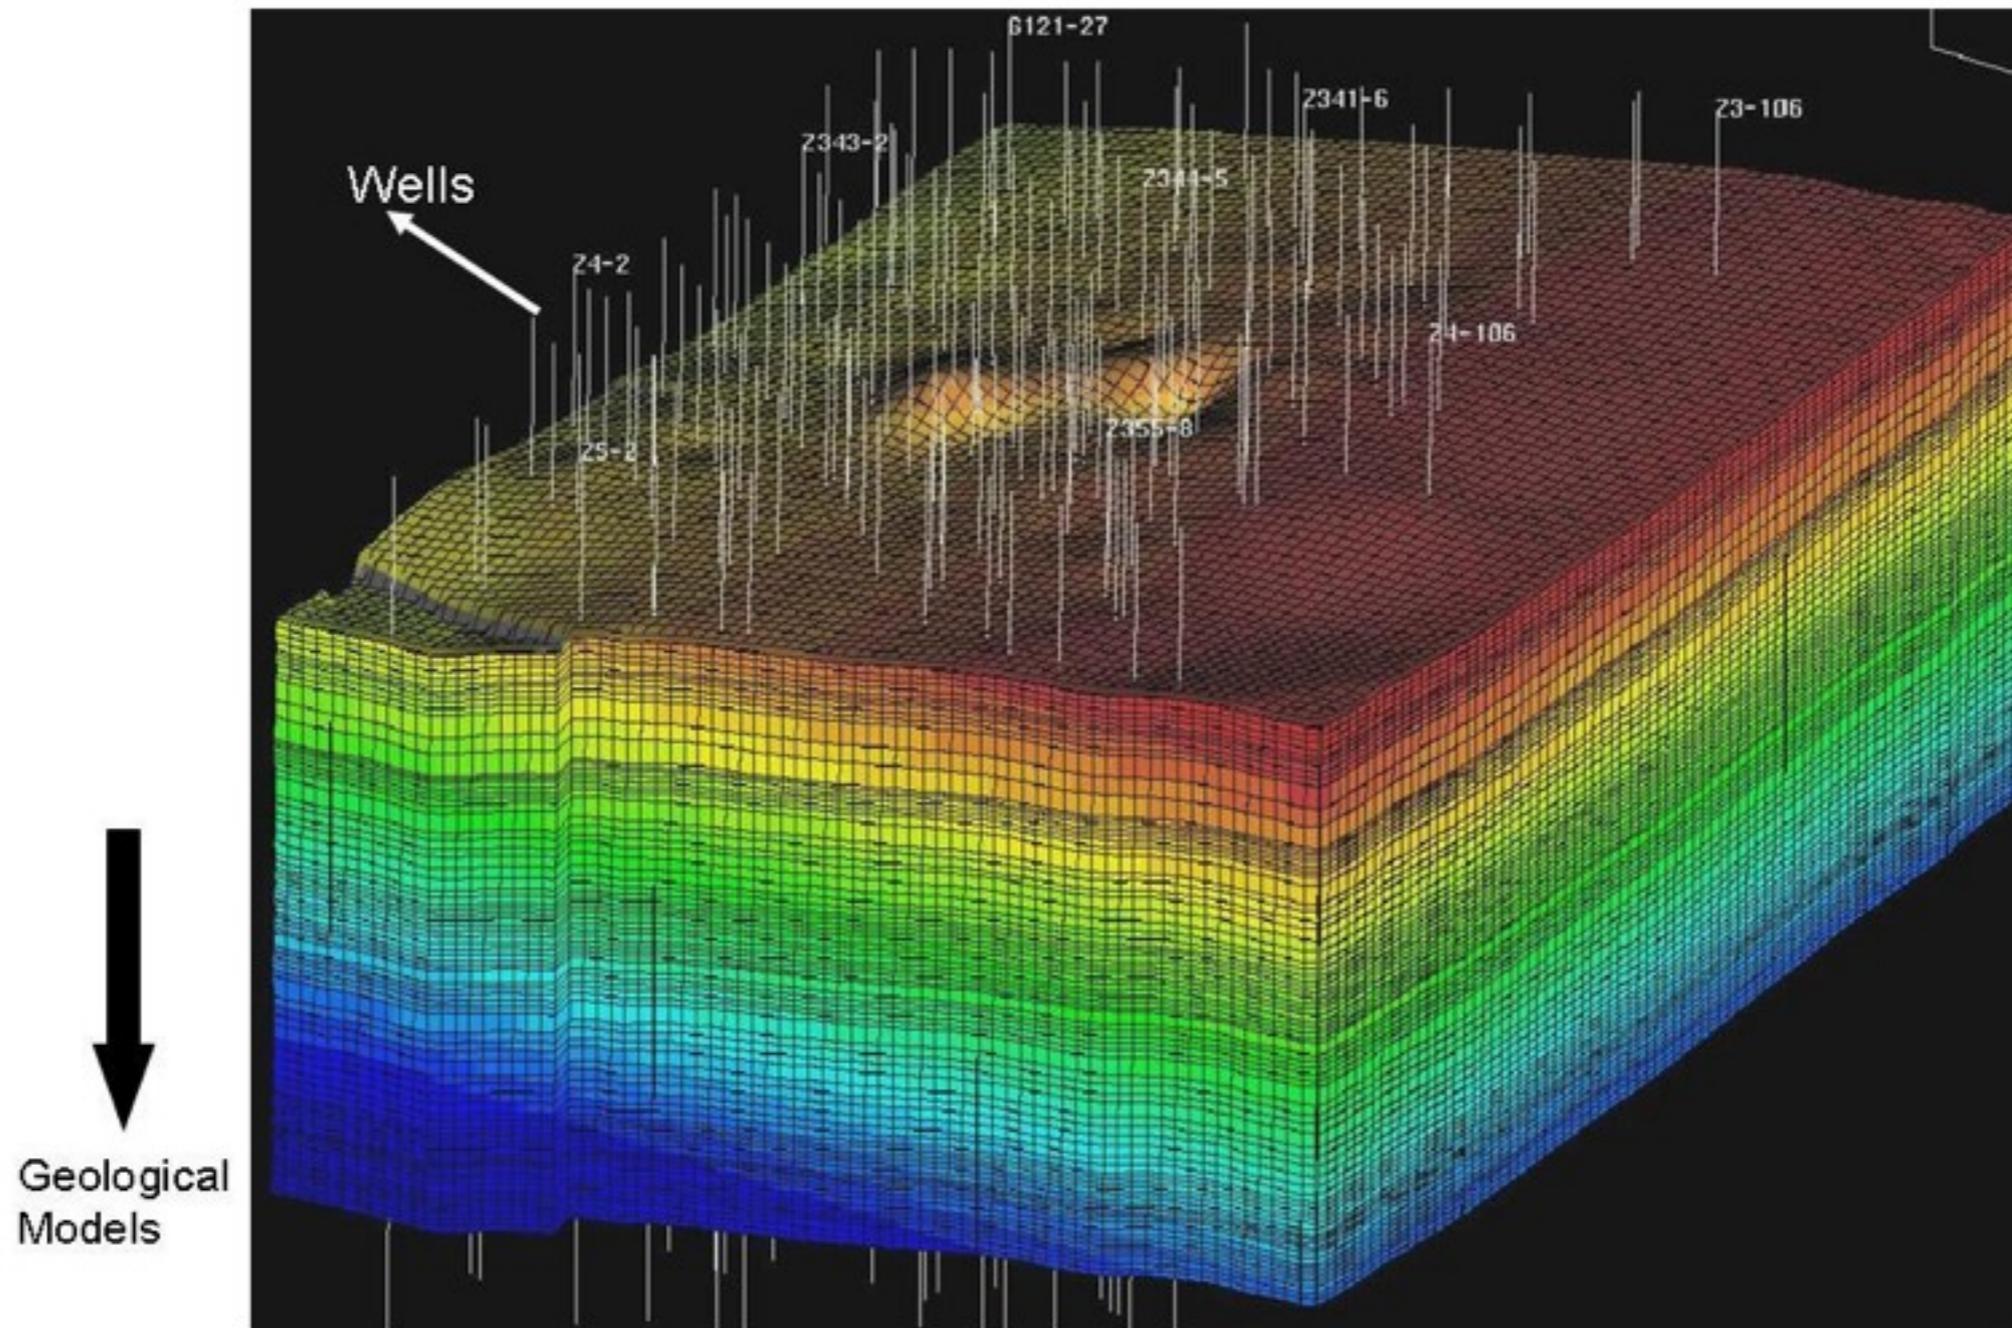
- Students and researchers before PEGrid generally do not have access to best practices in reservoir modeling and software tools.
- Industry as a result often struggles with difficulty identifying and hiring trained people with experience in high performance computing and the use of these methods.

The goal of PEGrid is to provide learning environments for complete workflows in reservoir modeling, exposing students and researchers to best practices in science and software methods.

- Comprehensive understanding of complex multi-scale models spanning many types and scales of physical systems is required. Such multi-scale modeling is computationally challenging and therefore HPC centers form the backbone of this effort. None of these capabilities exist in current industry standard software architectures.



Reservoir Modeling: Complex Multi-Scale Physical Processes



Many types of science from micron to many-km scales:
comprehensive approach is necessary

PEGrid Software Partners



\$596,445 Web portal software grant, internships



\$2.3 Million optimization software grant, internships

Participating Industry Members



Taking on the world's toughest energy challenges.™



\$90,000 software grant & Collaborative projects for metaschedulers



\$42.5 Million software grant, \$75,000 training support, internships and collaborations

Academic Collaborators



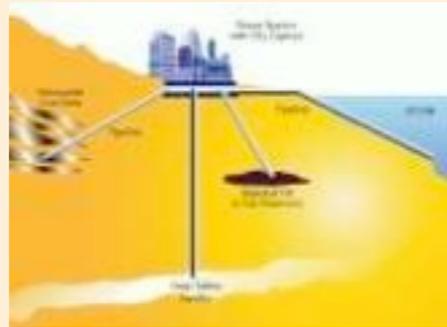
(Amounts are for 3-year period beginning in Fall 2009.)

Example: Schlumberger CO2 Storage Model Workflow

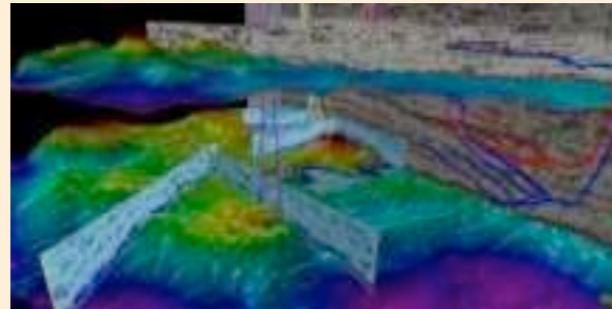
Pre-Operation Phase

~ 1-2 year

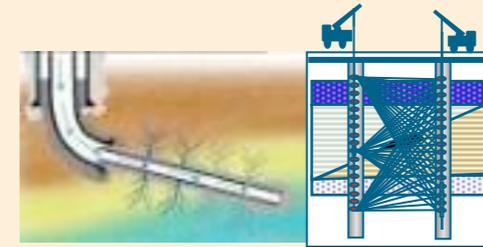
Site Selection



Site Characterization (SCP)



Field Design



Operation Phase

~ 10-50 years

Site Construction



Site Preparation



Injection



Monitoring (M&V)

• Operation

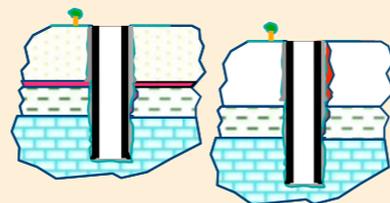


• Verification

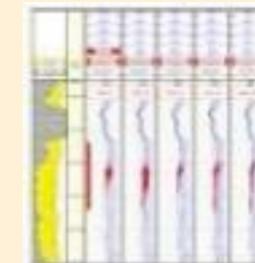
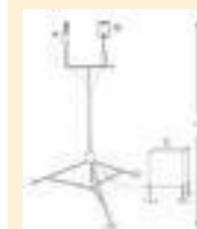
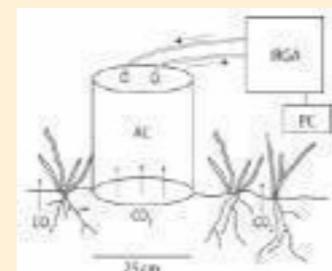
Post-Injection Phase

~ 100+ years

Site Retirement Programme (SRP)



• Environmental



Performance & Risk Management System (PRSM)
Communication and Public Acceptance

RISK		LEVEL	
Category	Value	Level	Color
Low	1-2	Low	Green
Medium	3-4	Medium	Yellow
High	5-6	High	Red
Critical	7-8	Critical	Black

Integrated workflow modeling is the path to new directions in energy research

Technical Challenges: PEGrid



- Operate science with high data security while providing controlled access to researchers
 - Data controls, strong authentication where needed
 - Execution controls for submission and retrieval of jobs
 - Full access for software manager
- Provide large-scale scalable infrastructure
 - Full conformance to appropriate grid standards
 - WS-Trust, WS-Security, Secure Conversation
 - X.509 infrastructure for large-scale submission of jobs
 - Web services with signed applications
 - Interface with .NET and web services portals (EnginFrame etc.)
- Create environment that is optimized for use in both commercial and educational settings.



Data, Access and Application Security



While creating the environment for full application access, we need controls in place for:

- Who has access to the applications
- Who can control and move job execution
- Specific controls on who can use and access data



We achieve this on a routine basis with existing grid controls by using:

- X.509 personal credentials for all grid access
- Host and user controls with individual account mapping
- Grid authentication and authorization for all job submission and data movement

In addition, we can impose any or all of the following:

- WS-Trust Secure Conversation, Transport Layer Security, Secure Message
- Separate account mapping for all PEGrid users
- X.509 secured storage and/or signed applications



High Speed Data Transfer Required



Features needed:

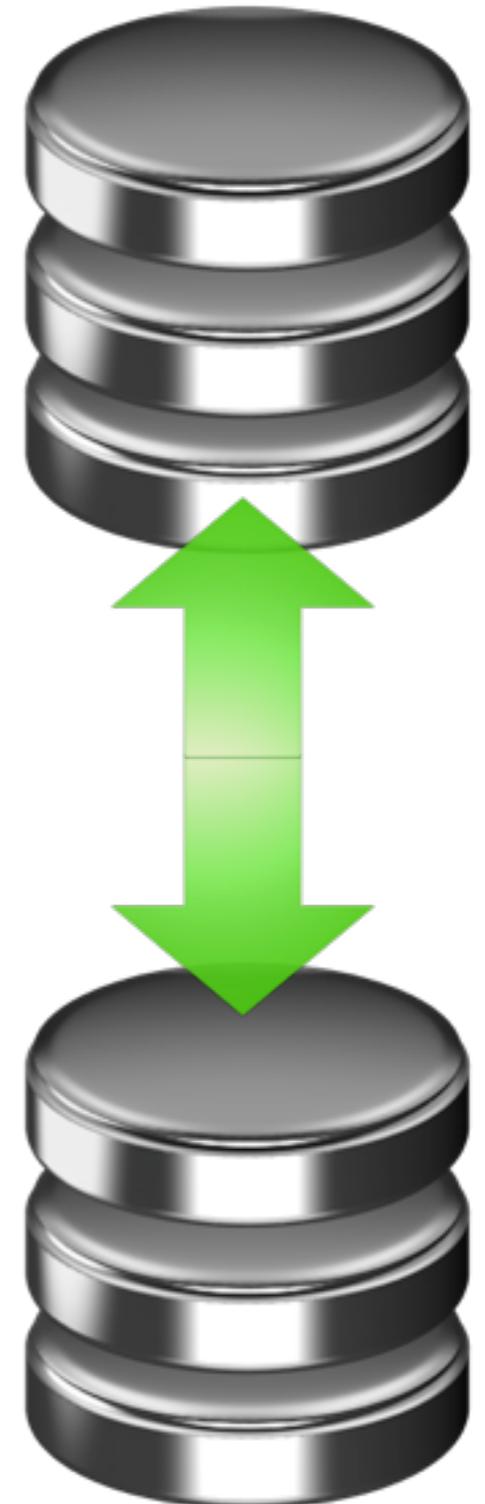
- Complete end-to-end security
- “Just in time” delivery as well as pre-staged data
- Specific controls on who can use and access data

Need security **from the outset!**

- Integrate X.509 credentials into portal access system
- Same host and user controls as used in job submission
- Separate account mapping for all PEGrid users

Also need very high performance data movement

- Systems used in collaborating grids already handle petabytes
- Must work across all centers and systems
- Multi-stream GridFTP with SRM (Storage Resource Manager) fits the bill



The projects I have described here were NOT Federally funded!

- ▶ TIGRE was a state project: \$2.5M to demonstrate grid computing.
- ▶ PEGrid is entirely an industry effort in partnership with academia, based on software donations and contributions by stakeholders of hardware, personnel etc.

They were, however, ENABLED by Federally funded middleware

- ▶ TIGRE, SURAgid and other cooperative regional projects cannot subsist or survive without appropriate middleware development.
- ▶ PEGrid would have been impossible on its own from scratch, but resulted in more than \$45M in industry-supplied software and funding.

What Is The Takeaway Lesson?

PEGrid is “two projects removed” in hierarchy from Federally-funded middleware software development, but would have been impossible to achieve without the VDT, International Grid Trust Federation, Globus, GridWay, and related computing projects.

Caveats and Warnings

- ▶ Outcomes like this can only be achieved through deep involvement by support staff with researchers.
- ▶ Demonstrations often outpace, but sometimes drive software development for infrastructure (web services, data and access security, load balancing and cloud computing, etc.)
- ▶ You didn't even know we were pursuing this topic, and we did not know it would lead to fruition in this area (although of course we were trying to do this all along)!
- ▶ If the infrastructure tools change out from under us during implementation, all progress could be lost. (Example: Globus changing web services interfaces.)

Final Lessons

- ▶ The role of intermediate groups such as individual university IT departments and high performance support groups for computing, such as state-wide efforts like HiPCAT, cannot be overestimated. Developments such as this would be impossible without cooperative engagement *for their own purposes* by such groups. Social efforts to promote and encourage such involvement should and *must* be undertaken.
- ▶ Development for its own sake is not always a sufficient goal; *operation* of the resulting systems must also be encouraged.
- ▶ Goal-oriented development with milestones can be a good thing.
- ▶ Involvement by university-based IT to support research is crucial.

Acknowledgments

Members of the TIGRE development and steering teams:

Francis Dang, Rion Dooley, Mark Huang, Keith Jackson, Taesung Kim, Ping Luo, Margaret Murray, Warren Smith, Kiran Thyagaraja, Ravi Vadapalli, Alan Sill Jay Boisseau, Lennart Johnsson, Charles Koelbel, Lee Panetta, Philip Smith

Petroleum Engineering Departments of TTU, UT Austin and Texas A&M
including Shameem Siddiqui and Lloyd Heinze (TTU), Kamy Sepehrnoori, Matthew Balhoff (UT Austin), Akhil Datta-Gupta, and Tom Blasingame (TAMU)

The Virtual Data Toolkit and Condor Teams

Miron Livny, Tim Freeman, Alain Roy, et al.

Many people in the Open Science Grid

especially Ruth Pordes, Robert Gardner, Doug Olson, Robert Quick, and others

International Grid Trust Federation

Scott Rea, Michael Helm, David Groep, James Marsteller and many others

Open Grid Services Architecture Work Group/Open Grid Forum

Andrew Grimshaw, Hiro Kishimoto, Duane Merrill, Andreas Savva

Southeastern Universities Research Association

Linda Akli, Gary Crane, James Jokl, John-Paul Robinson, Mary Fran Yafchak, etc.

Thank you.