Final Project Report

Texas Internet Grid for Research and Education (TIGRE)

Rice University,

Texas A & M University,

Texas Tech University,

University of Houston, and

The University of Texas at Austin

Dec 14, 2007

1. Introduction

The Texas Internet Grid for Research and Education (TIGRE) project was a two-year development effort conducted to provide a grid computing infrastructure that enables integration of computing systems, storage systems and databases, visualization laboratories and displays, and even instruments and sensors across Texas. The overall goal of this project was to facilitate new modalities for academic-government-private research partnerships, fostering collaborations and partnerships among universities and industries by dramatically enhancing both computational capabilities and collaborative research infrastructure. The TIGRE project was targeted to address research areas of interest to the State of Texas where manifold increase of computing power, data access, and collaboration would be necessary.

The TIGRE project consisted of institutional participants comprising Principal Investigators (PIs) and TIGRE-funded personnel at each of five universities: Rice University, Texas A&M University, Texas Tech University, the University of Houston, and the University of Texas at Austin. Collectively these will be referred to in this document as the "TIGRE Institutions".

The scientific goal of TIGRE was to provide an operational grid software that enables leading-edge applications in at least three at least three areas of importance to the state of Texas. As a starting point, one example application was selected each in the areas of bio-sciences and medicine, energy, and environmental modeling. TIGRE brought researchers together at multiple institutions into coordinated grid application development communities towards this objective, with the goal of demonstrating new, preferred, enhanced or increased computing and storage handling capabilities offered by a statewide grid infrastructure.

TIGRE has succeeded in achieving these goals, and has implemented both the software and organizational infrastructure called for in its original formulation. This report provides a summary of the objectives, scope, organization, milestones, and achievements of the TIGRE project and lays out a path toward future production status.

2. Objectives

The high-level objective for this project was to deploy a collaborative grid infrastructure and set up working grid services architecture (GSA). TIGRE was intended to promote collaboration among Texas research institutions and companies, thereby allowing them to compete nationally and internationally in science and engineering research and for large multi-institution programs and awards. An explicit goal at the outset was to support extensibility of these grid services to additional institutions in Texas as the project proceeded.

The lower-level objectives of this project were to:

1. Provide client- and server-side software with capabilities for resource sharing and coordination that enhances what any single institution can contribute.
2. Create and distribute a software stack that primarily supports the three driving applications of interest to the State of Texas and is easily extensible to a variety of applications.
3. Provide the service infrastructure for TIGRE such as services for authentication and authorization, scheduling work, customer relations management, and a web portal.
4. Provide the operational infrastructure for TIGRE such as a defined organization, policies and procedures, and documentation.

3. Scope

The goal of TIGRE was to deploy an extensible architecture for grid computing in Texas. As funded, the project included support for development activities at the five founding TIGRE institutions, but not for production operations. While all development objectives for the TIGRE project were met, it was anticipated from the outset that significant further effort would be required to create support conditions under which a large number of Texas institutions can adopt the resulting architecture. The purpose of the TIGRE development project was therefore to prepare production-quality software and infrastructure, and to demonstrate that such expansion is practical.

The TIGRE project also did not include explicit funding to provide a large amount of dedicated computational resources. As an initial step, the TIGRE members provided access to a limited number of computational and storage resources on a volunteer basis. These resources were provisioned with copies of the TIGRE software for users to access so that the project could evaluate software, services, and operational approaches. A goal of the project was to ensure that the software, including grid middleware and client tools, support center tools, and other developed resources were capable of handling large operational loads when deployed. If and when, as a further extension of this project, the user base and/or number of participating institutions are to grow, it will be necessary to explore ways to enhance the available grid resources substantially in the future.

The final section of this report lays out the preparations and framework for such an expansion of the TIGRE development project into production status.

4. Project Management

4.1 Organization and Leadership

The TIGRE development project was organized into two primary components: a project Steering Committee (SC), and a project Development Team (DT).

The SC was a group consisting of the Principal Investigator (PI) or his or her representative from each TIGRE institution. This committee provided strategic direction and high-level leadership for the project by making decisions regarding major project directions. These decisions were made through unanimous agreement among the five institutional participants. Functions of the project explicitly under the control of the SC included establishing procedures regarding usage allocations, policies and usage priorities, selection of the initial project application areas, and working interactively with researchers and educators to stimulate, formulate, and guide applications development for TIGRE.

The DT consisted of TIGRE-funded staff at each of the institutions organized into a body that constructed and operated TIGRE during its development stages. This consisted of a range of technical tasks related to selecting, testing, and deploying software, defining and carrying out operational policies, providing documentation, providing user support and working with the pilot application deployment. The DT operated by identifying tasks that needed to be performed and creating activities to perform them. The primary organizational paradigm used by the DT was to group its effort around such activity areas, and to work closely with the SC in meeting the quarterly milestones of the project.

4.2 Roles & Responsibilities

The SC was responsible for making high-level decisions for providing status reports quarterly to DIR. The DT was responsible for carrying out technical activities, satisfying milestones and deliverables and reporting progress to the Steering Committee.

To assist in satisfying milestones and deliverables, the DT formed activity sub-groups that were responsible for performing specific pieces of work. Each activity had a lead and one or more members. The lead for the activity was responsible for reporting on the status of the activity at each weekly DT meeting and ensuring that the activity was completed successfully. A set of internal web pages were maintained as appropriate for each activity or sub-activity.

4.3 Tracking and Control

Current and upcoming milestones were reviewed on a regular basis among the members of the DT through telephone conferences and in-person meetings. The DT typically met weekly, and in-person "face to face" meetings were held approximately quarterly throughout the project. The lead and his or her team on each activity reported the progress and/or problems to the DT as a whole. The completed milestones within each quarter were reported to the SC, and consequently by the SC to the DIR. The SC met via telephone and exchanged e-mail as required to review the progress of the project and to provide feedback to the DT. Internal TIGRE web pages were created and updated as needed to reflect the progress and upcoming activities for common reference by the DT and SC.

The DT held face-to-face meetings as required when there were discussions needed or tasks to perform that were best accomplished in person. Potential users were involved frequently in these meetings to gather their feedback as well as to identify new applications that could gainfully exploit the TIGRE grid services architecture.

4.4 Status Reporting

The SC appointed Prof. Philip W. Smith of Texas Tech University to disseminate the quarterly status reports to the DIR with consensus from all the TIGRE institutions. These quarterly reports were submitted regularly as the project proceeded forward.

The main internal mechanism for reporting status was the project web site, http://hipcat.net/Projects/tigre/ (public page) and http://hipcat.net/groups/tigre/ (site open only to project participants). These web sites provided information about milestones and deliverables for the project, activities as they were ongoing and as completed, and serve as a repository for meeting minutes. A secondary mechanism for status reporting and discussion was through the use of TIGRE internal email lists. Archives of these lists are available for reference at the UH list server page: http://mail.tlc2.uh.edu/mailman/listinfo/ .

Further pubic-oriented pages were developed and maintained at http://tigreportal.hipcat.net/ to provide status reporting on resources available to the project, installation instructions for the client and server software stacks, and tutorials and instruction for users. This site, supported by a wiki to allow maintenance of the included information by developers, became the primary method of communication with the public and with users about the project.

4.5 Project Assessment and Control

We used the following procedures to control and document the project to ensure adherence to project goals and a high standard of performance:

- Project documents such as the project plan, requirements documents, an architecture description and a reference software implementation were written and maintained on the web sites described above. These documents were version controlled for subsequent refinements in the course of the progress of the project.

- We expected from the outset to adopt or modify software or create new software in an iterative manner to satisfy our deliverables and meet user needs. Such progressive refinements were introduced in a controlled way through specific version numbers as appropriate. Such changes were conducted with the explicit goal of establishing initial criteria and then gathering user feedback that led to further refinements in the software stack and associated instructions.
- Further change control, quality assurance, and risk assessment steps were carried out in a controlled manner according to the project plan.

5. Statement of Work

This section presents the milestones and deliverables for the project along with the work breakdown structure, resources, and schedules to satisfy the milestones and deliverables. All milestones were achieved within a reasonable period close to or ahead of their planned interval, and the overall development project was completed on time as designed.

5.1 Milestones

Table I shows the quarterly milestones of the TIGRE contract with the State of Texas. All milestones were achieved, as recorded in this table.

5.2 Deliverables

TABLE II shows TIGRE project deliverables along with their timeline. The indicated deadlines were derived from the TIGRE project milestones (see TABLE I) described in the previous section. Further, as described in a previous letter to DIR, we had selected a start date for the TIGRE project as December 1, 2005. This led to the due dates specified in TABLE II. All deliverables were achieved as specified in conformance with this plan.

5.3 Work Breakdown Structure

A Work Breakdown Structure (WBS) was used to categorize and track project work components, and revised by the project as necessary. As the primary organization of the technical work of the project was by milestones, deliverables, and project activities, the WBS goals were adjusted to be reflective of these project activities. The current final status of the WBS is below.

1. **Project management**. The creation of an initial project plan document, one of the deliverables during the first quarter of the project will be carried out. Periodic iteration of the project plan throughout the duration of the project with emphasis on the vision of immediate, near and future goals of TIGRE helps in laying stronger foundation in the project management task.

   1.1. Create initial project plan. This element provides deliverable Y1Q1.1 **(Achieved.)**

   1.2. Revise project plan at the end of Y1Q3. (No revision needed.)

   1.3. Revise project plan at the end of Y2Q1. (Slight delay in demonstration deliverables.)

   1.4. Revise project plan at the end of Y2Q3. (Project caught up, no revision needed.)

2. **Software**. This module consists of identifying initial software stack for at least three driving applications of interest to the State of Texas. The chosen software stack will be adopted or modified, new software tools will be written for transparent deployment of these applications using GSA. Important steps involved in this process include the following:

2.1. **Gather initial requirements.** These requirements will be used to evaluate software packaging and distribution technologies as well as determine which software should be included in the initial TIGRE software stack. **(Achieved.)**

2.2. **Select packaging/distribution technology.** The packaging technology will allow TIGRE software components to be bundled into easily installable packages. The distribution technology will allow TIGRE users and administrators to easily download and install TIGRE software packages. This task will provide deliverable Y1Q3.2. **(Achieved.)**

2.3. **Create initial software stack.** This task defines, packages, tests, and documents the initial software stack. This task will simultaneously provide deliverable Y1Q3.1 and Y1Q4.1. **(Achieved.)**

2.4. **Gather requirements for second version of software stack. (Achieved through ongoing activity.)**

2.5. **Create second software stack.** This task defines, packages, tests, and documents the second software stack. This element provides deliverable Y2Q3.1. **(Achieved.)**

2.6. **Gather requirements for final version of software stack. (Achieved.)**

2.7. **Create final software stack.** This task defines, packages, tests, and documents the final software stack. This task provides deliverable Y2Q4.1 and Y2Q4.2. **(Achieved . Note that software will be kept up to date through production use as needed for software security and functionality updates.)**

3. **Services**. This module provides an interface for interoperability of the GSA and applications. The services include user level certificates for authentication and authorization, resource allocation and registration, application deployment, and work load management through scheduling services, and a customer services management system for monitoring the quality of service and bug fixing.

3.1. **Provide a web site.** The public web site will be used to provide information to users, potential users and partners, and the state. This element provides deliverable Y1Q1.2. **(Achieved as described above.)**

3.2. **Certificate Authority.** This task provides the procedures and software to issue credentials to TIGRE users and resources, after their identity has been proven. **(Achieved.)**

3.2.1. **Provide the initial CA.** The initial CA will be suitable for use in the early phases of TIGRE. This element provides deliverable Y1Q1.3 **(Achieved.)**

3.2.2. **Provide the final CA.** This task provides an improved CA based on lessons learned when using the initial CA. **(Achieved. Final CA is now accredited by the International Grid Trust Federation, and is in the latter stages of deployment.)**

3.3. **Registration Authority.** This element provides the procedures and software needed so that individuals who wish to be part of TIGRE can prove their identity so that they can then have credentials issued to them. **(Achieved by use of TIGRE Developer Team members as Registration Authority representatives to the project.)**

3.3.1. **Provide the initial RA.** This task provides an initial RA that is suitable for the early phases of TIGRE. **(Achieved as described above.)**

3.3.2. **Provide the final RA.** This element provides an improved RA based on lessons learned when using the initial RA and that will scale to when TIGRE has 30+ member sites. **(Deferred to production project.)**

3.4. Provide a web portal. The TIGRE web portal is targeted toward TIGRE users and will allow them to obtain information about and use TIGRE resources. This task provides deliverable Y1Q2.1. **(Achieved.)**

3.5. Customer management system. A customer management system is used to monitor and analyze customer behavior. In business, this data can be used to improve profits, but in the case of TIGRE, this data will be used to improve the software and services provided to users. **(Achieved.)**

3.5.1. Provide the initial customer management system. This element provides deliverable Y2Q1.1 **(Achieved.)**

3.5.2. Provide the final customer management system. (Achieved with no changes needed to the above.)

3.6. Provide a global grid scheduler. A global grid scheduler accepts jobs from users and then sends those jobs to TIGRE systems for execution while trying to optimize properties such as completion time, throughput, and cost. This element provides deliverable Y2Q2.1. **(Achieved. Two such global schedulers were identified and tested. Either is capable of production use.)**

4. Operations. This module offers operations to interface basic services of TIGRE that includes providing policies and procedures for joining TIGRE, utilization of TIGRE resources, monitoring services, policies and procedures for establishing site support.

4.1. Deployment. This element contains the effort to deploy the TIGRE software stack onto TIGRE resources.

4.1.1. Identify minimum testbed requirements. This element provides deliverable Y1Q1.4. **(Achieved.)**

4.1.2. Deploy the initial software stack. (Achieved.)

4.1.3. Deploy the second software stack. (Achieved.)

4.1.4. Deploy the final software stack. (Achieved.)

4.2. Monitoring and testing. This element identifies requirements and deploys a software system to monitor and test TIGRE resources.

4.2.1. Deploy initial monitoring and testing system. (Achieved.)

4.2.2. Deploy final monitoring and testing system. (Achieved with no changes needed.)

4.3. User support. This task performs various user support activities including working with users to resolve problems and creating documentation.

4.3.1. Author user guide. (Achieved.)

4.3.2. Author application porting guide. (Revised; achieved through example projects.)

4.4. Site support. This element defines and documents the procedures for a site to be part of TIGRE.

4.4.1. Document the minimum software and services needed to join TIGRE. This element provides deliverable Y2Q3.2. **(Achieved; see attached.)**

4.4.2. Document how to join TIGRE. The documents that result from this element will describe what a new site needs to do to join TIGRE. This element partially provides deliverable Y2Q4.3. **(Achieved; see attached.)**

4.4.3. **Document TIGRE operational policies and procedures.** This element partially provides deliverable Y2Q4.3. **(Achieved; see attached.)**

5. **Application porting**. Many potential TIGRE users will have legacy serial or parallel codes. Porting such codes to a distributed environment, such as TIGRE, can be difficult. This element will provide assistance to users to port their applications to TIGRE. The primary effort will be on working with selected pilot applications.

   5.1. **Select initial three applications.** This element provides deliverable Y1Q1.5. **(Achieved.)**

   5.2. **Port first application.** For each application, we identify the interfaces and functionality the application will provide, design how the application will operate in TIGRE, identify the interfaces and functionality the application needs from TIGRE, enhance TIGRE to provide these interfaces and functionality, implement the application design, measure and tune performance, and finally document successes, failures, and any additional functionality that TIGRE and the application could provide to make the application more successful. **(Achieved.)**

   5.3. **Port second application.** The second application will be ported to TIGRE using the same procedures as used to port the first application. **(Achieved.)**

   5.4. **Port third application.** The third application will be ported to TIGRE using the same procedures as used to port the first application. **(Achieved.)**

6. **Demonstrations.** The TIGRE project will perform a number of demonstrations. The purposes of the demonstrations are to demonstrate that TIGRE is operational to the state and to demonstrate to potential users and member sites the benefits of joining TIGRE.

   6.1. **Provide initial demonstration.** For each demonstration we define the demonstration, prepare it, provide the demonstration, and then document how the successes and failures of the demonstration. This element provides deliverable Y1Q3.3. **(Achieved through UltraScan application in Biosciences and Medicine area.)**

   6.2. **Provide second demonstration.** This task provides deliverable Y2Q1.2. **(Achieved through EnKF application in Energy Exploration area.)**

   6.3. **Provide third demonstration.** This element provides deliverable Y2Q3.3 and Y2Q3.4. **(Achieved through data transfer and processing applications in Atmospheric Modeling area.)**

   6.4. **Provide final demonstration.** The final demonstration of the TIGRE project will be performed at the SC conference. This task provides deliverable Y2Q4.4. **(Achieved; demonstrations in all application areas at Supercomputing 2007 conference in Reno, NV, November 2007.)**

6. Schedule

A Gantt chart that describes the schedule for working on and completing the elements of our work breakdown structure for TIGRE is shown in Figure 1.

7. Resources

The financial resources available to TIGRE through direct State funding were utilized almost entirely in the form of staff support. The participating institutions typically provided two full-time equivalent personnel through project funding. In some cases, TIGRE institutions ramped up their staff and activity

in order to make project progress before the official project start date, but most institutions synchronized their overall project activity to match fairly close to the project time line. Each institution was responsible for its own staffing plan to meet project goals.

Project resources for computing and data transfer in the course of the development project were furnished on a volunteer basis by the participating institutions. The current list of resources connected through TIGRE software is attached as Figure 2. Although not all of these sites provide full access to all of their internal resources to TIGRE, it should be noted that the list includes several computing resources that were added to TIGRE by institutions that were beyond the initial list, including some from an institution (UT Health Sciences Center, San Antonio) that was not one of the TIGRE funded institutional participants.

8. Budget and financial reporting

Out of the $9.8 million Texas Enterprise Fund award for LEARN and TIGRE, $500,000 was allocated for each of the five TIGRE institutions, amounting to a total of $2.5 million, to deploy grid computing infrastructure that enables the integration of resources at Texas institutions. Financial reports will be submitted by each of the TIGRE institutions separately.

9. Summary

TIGRE has achieved its goals in terms of developing, testing, deploying and delivering a grid software infrastructure for the State of Texas that is fully state-of-the-art, reliable, tested and well matched to project requirements in its targeted application areas. The resulting software stack is available in both client and server-oriented formats, and can be downloaded from the TIGRE portal web site.

Our experience with this software stack has shown that it is easy to use, easily deployed and suited for a broad range of grid-related purposes. The software includes elements that allow obtaining and using grid credentials for secure access to computing and storage resources, tools that allow users to deploy a broad range of applications in a modern web-services context, tools to allow central submission of jobs through submission services across a range of resources, and tools for monitoring and data transfer in association with these jobs. The range of applications was chosen by the Steering Committee to demonstrate TIGRE capabilities in three different areas of considerable interest to the State of Texas: Energy Exploration, Atmospheric Modeling, and Biosciences/Medicine. Each of the demonstrations in these areas succeeded in showing the utility of grid computing for carrying out these applications across a broad range of resources.

The TIGRE stack has attracted a great deal of interest from the Globus software team, which sent managers and representatives twice to meet with our group, from the Open Science Grid and Virtual Data Toolkit development teams, and from other grid projects, including SURAGrid. In particular, SURAGrid has decided to adopt a variant of our client and server software packages as the basis for its own similar efforts.

Several institutions beyond those of the initial five TIGRE participants have expressed interest in joining TIGRE as it leaves the development process and enters its production phase. In particular, we have received strong interest from UT El Paso and the UT Health Sciences Center in San Antonio, the latter of which has already begun contributing access to TIGRE for some of its resources and has become a steady TIGRE user in the Biosciences and Medicine area.

We look forward to many applications of TIGRE and to further collaborations as it enters its production stage.

APPENDIX

A. Definitions

Alpha quality	Software that is not feature complete and may be known to have bugs is considered alpha quality. Alpha quality software is occasionally provided to users that have an urgent need for some of the features in the software.
Beta quality	Software that is feature complete, ready for regular users, thought to be mostly bug free, and in final stages of testing is considered to be beta quality. Beta quality software is provided to users to help with the final testing of the software.
Grid	A network that coordinates distributed compute and data resources using standard, open, general-purpose protocols and interfaces to deliver nontrivial quality of service.
Middleware	Software that interfaces two separate and often already existing application programs. It is neither a part of the operating system nor application programs.

B. Glossary

CA	Certificate Authority
DIR	Department of Information Resources
DT	Developers Team
GSA	Grid Services Architecture
HiPCAT	High Performance Computing Across Texas
LEARN	Lonestar Education and Research Network
RA	Registration Authority
SC	Steering Committee
TIGRE	Texas Internet Grid for Research and Education

TABLE I: Quarterly Milestones of the TIGRE contract with the State of Texas

Quarter	Milestone
Y1Q1	Present a project plan and scope of work in a mutually acceptable format to Department of Information Resources (DIR). Establish a common web site, shared Certificate of Authority (CA) service and minimum testbed requirements within the first quarter of the first year. This will enable TIGRE institutions to immediately perform outreach activities by publishing lessons learned and experience papers to the community and the public at large. Select three driving applications of importance to the State that will be used to derive grid software stack requirements for construction, as well as provide a set of applications for testing, verification and assessment of success.
Y1Q2	A prototype alpha quality portal service will be deployed within the second quarter of year 1 to facilitate requirement gathering of application and site specific resource submission characteristics amongst the varied scientific and educational communities that will be using the portal.
Y1Q3	Within the third quarter of year 1, a site service software stack will have been defined. A distribution mechanism, possibly using packaging technologies like ***pacman***, will be agreed upon and customized for TIGRE. Demonstrate, on a small scale, basic functionalities of TIGRE capabilities using at least one of the driving applications.
Y1Q4	A prototype alpha-quality client tools package will be distributed to friendly users by the end of year 1. This client tools package will allow friendly users to start using site resources earlier on to identify missing components, and harden tools which will eventually be released to the wider TIGRE user community.
Y2Q1	An alpha-quality customer management services system will be deployed in the first quarter of year 2 to enable testing and site specific customization. Demonstrate across resources from the majority of the TIGRE institutions capabilities using all three driving applications.
Y2Q2	In the middle of year 2, a global grid scheduler service will be deployed. This grid scheduler will enable sites to advertise resources, and TIGRE user jobs to be routed to the most appropriate site for execution.
Y2Q3	In the third quarter of year 2, the alpha-quality packages will be upgraded with bug fixes and new feature requirements and a cool down period will be enforced. During this period only critical changes to the infrastructure will be done, allowing a stable environment to be migrated into production at the end of year 2. Also the minimum set of services needed to use and/or join TIGRE production grids will be defined. Prepare and execute a major demonstration of TIGRE software capabilities using the three driving applications and evaluate ease of use and success.
Y2Q4	During the final quarter, the software hardening will be completed and the installation and support documentation will be prepared. Procedures and policies for joining TIGRE and sharing resources will be developed and documented, and posted to the TIGRE web sites. Demonstrate TIGRE at the national supercomputing conference.

TABLE II: Quarterly Deliverables of the TIGRE contract (from TABLE I)

ID	Date	Metric	Description
Y1Q1.1	02/28/2006	Project plan document to DIR	Project Plan. Create and present a project plan and scope of work to the Department of Information Resources (DIR).
Y1Q1.2	02/28/2006	Public TIGRE web site accessible	Web site. Establish a common web site to publish lessons learned and experience papers.
Y1Q1.3	02/28/2006	CA service available	Certificate Authority. Establish a shared Certificate Authority.
Y1Q1.4	02/28/2006	Requirements document available	Testbed requirements. Define the minimum requirements.
Y1Q1.5	02/28/2006	Three applications identified	Driving applications. Select three driving applications of importance to the State to derive selection of software stack and for testing, verification, and assessment.
Y1Q2.1	05/30/2006	Portal deployed and usable	Portal. Deploy a prototype alpha-quality portal service.
Y1Q3.1	08/31/2006	Software stack identified	Software stack. Deliver alpha quality version of middleware service software stack to be deployed at each site.
Y1Q3.2	08/31/2006	Distribution mechanism identified	Distribution mechanism. Deploy a software distribution mechanism for TIGRE and perform any needed customization.
Y1Q3.3	08/31/2006	Demonstration performed	Small-scale demonstration. Perform a demonstration, on a small scale, of basic TIGRE capabilities using at least one driving application.
Y1Q4.1	11/30/2006	Client tools distributed	Client tools. Prototype alpha-quality client tools distributed to users.
Y2Q1.1	02/28/2007	Customer management system available for use	Customer management. Deploy an alpha-quality customer management system.
Y2Q1.2	02/28/2007	Demonstration performed	Demonstration. Demonstrate using TIGRE capabilities to execute all three driving applications on resources from the majority of the TIGRE institutions.
Y2Q2.1	05/30/2007	Grid scheduler available	Grid scheduler. Deploy a global grid scheduler that enables sites to advertise resources and routes TIGRE jobs to the most appropriate sites.
Y2Q3.1	08/31/2007	Client and server packages available.	Stable packages. Provide stable client and server packages.
Y2Q3.2	08/31/2007	Document describing the minimum services	Minimum services. Define the minimum software services needed to use/join TIGRE.
Y2Q3.3	08/31/2007	Demonstration performed	Major demonstration. Provide a major demonstration of TIGRE capabilities using the three driving applications
Y2Q3.4	08/31/2007	Document describing ease of use and success	Evaluation. Evaluate the ease of TIGRE using the major demonstration and the success of the demonstration.

TABLE II: Quarterly Deliverables of the TIGRE contract … (Continued)

ID	Date	Metric	Description
Y2Q4.1	11/30/2007	Final software packages available	Final grid middleware packages. Harden the software packages to create release-quality stable versions of these packages.
Y2Q4.2	11/30/2007	Documentation available	Documentation. Release production-quality installation and support documentation.
Y2Q4.3	11/30/2007	Documents prepared and posted to TIGRE web site	Procedures and policies. Prepare procedures and policies for joining TIGRE and sharing resources.
Y2Q4.4	11/30/2007	Demonstration performed at SC	SC demonstration. Demonstrate TIGRE at the national supercomputing conference.

ID	Task Name
1	**Project Management**
2	Initial project plan
3	Y1Q3 project plan revision
4	Y2Q1 project plan revision
5	Y2Q3 project plan revision
6	**Software**
7	Gather initial requirements
8	Select packagin/distribution tech
9	Create initial software stack
10	Gather 2nd set of requirements
11	Create second software stack
12	Gather final requirements
13	Create final software stack
14	**Services**
15	Provide a web site
16	**Certificate Authority**
17	Provide initial CA
18	Provide final CA
19	**Registration Authority**
20	Provide initial RA
21	Provide final RA
22	Provide web portal
23	**Customer management system**
24	Provide initial CMS
25	Provide final CMS
26	Provide global grid scheduler
27	**Operations**
28	**Deployment**
29	Identify minimum testbed reqs
30	Deploy initial software stack
31	Deploy second software stack
32	Deploy final software stack
33	**Monitoring and testing**
34	Deploy initial system
35	Deploy final system
36	**User support**
37	Author user guide
38	Author application porting guide
39	**Site support**
40	Author minimum software and services
41	Author how to join TIGRE
42	Author operational policies and procedures
43	**Application porting**
44	Select initial three applications
45	Port first application
46	Port second application
47	Port third application
48	**Demonstrations**
49	Provide initial demonstration
50	Provide second demonstration
51	Provide third demonstration
52	Provide final demonstartion

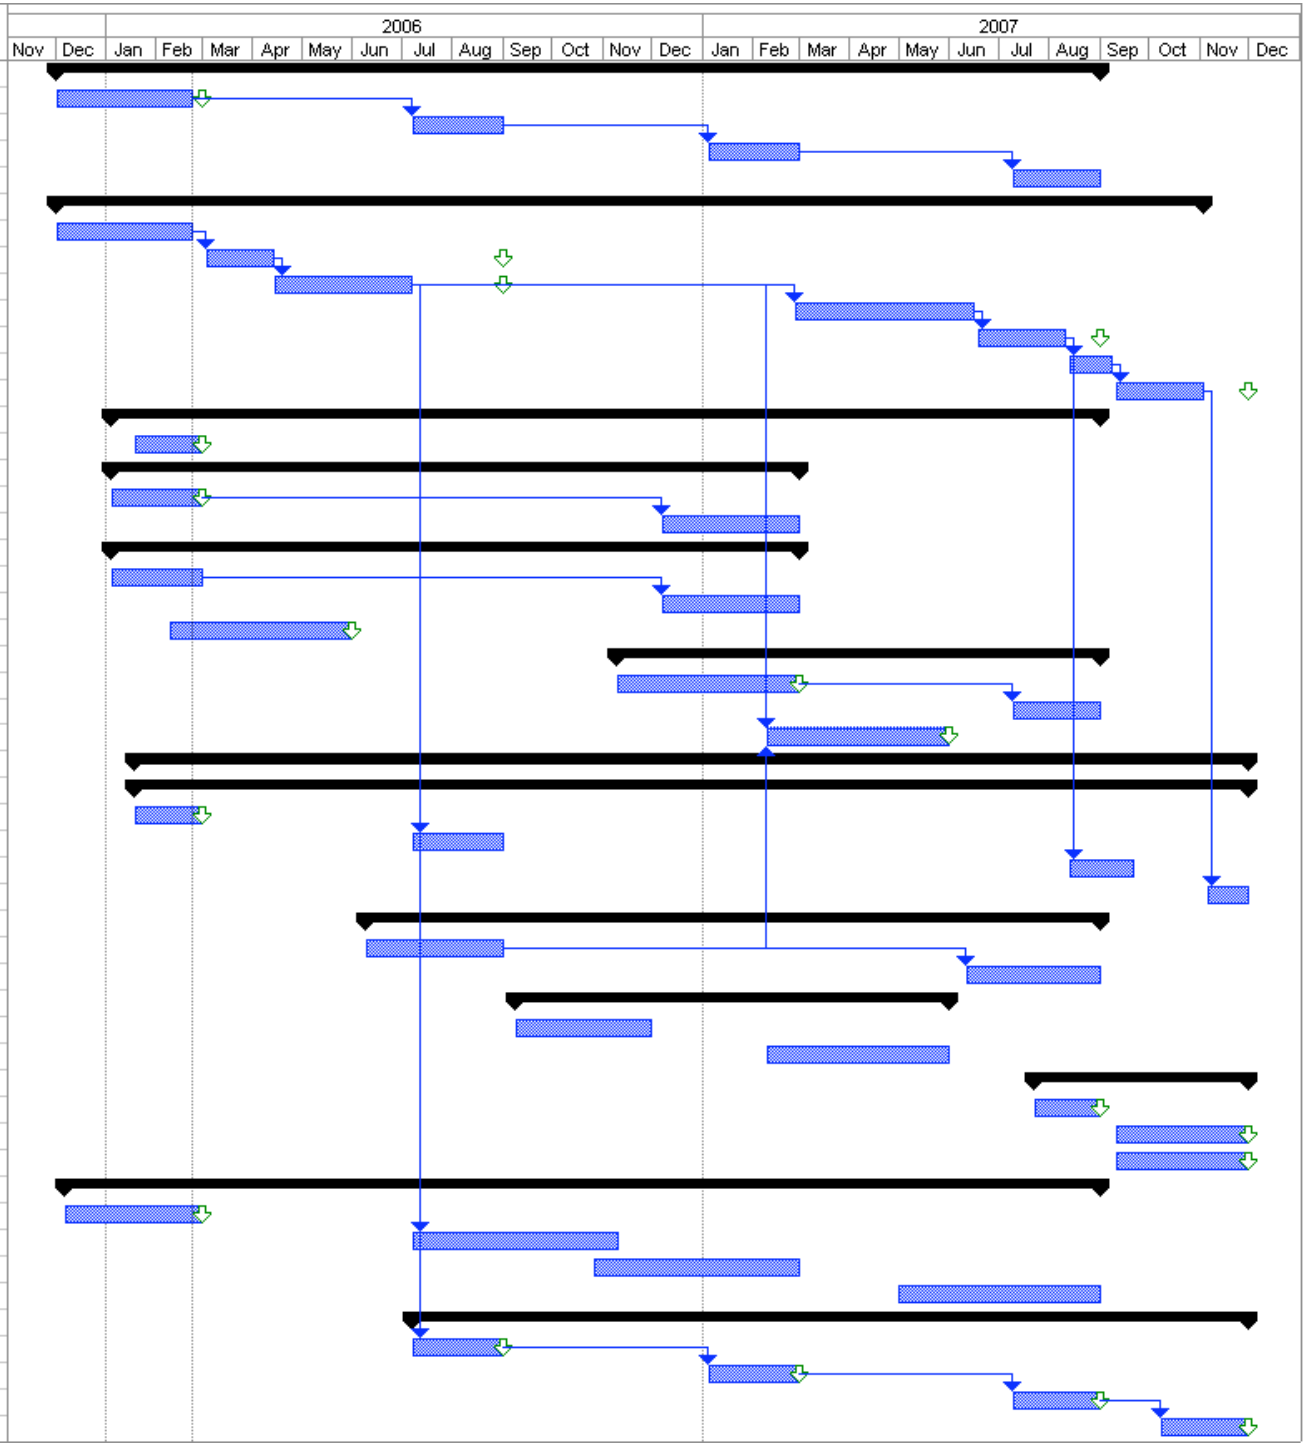

FIGURE 1: The project schedule