

TeraGrid: Analysis of Organization, System Architecture, and Middleware Enabling New Types of Applications

Charlie CATLETT^a, William E. ALLCOCK^a, Phil ANDREWS^b, Ruth AYDT^f,
Ray BAIR^a, Natasha BALAC^b, Bryan BANISTER^b, Trish BARKER^f,
Mark BARTELT^k, Pete BECKMAN^a, Francine BERMAN^b, Gary BERTOLINE^c,
Alan BLATECKY^j, Jay BOISSEAU^d, Jim BOTTUM^l, Sharon BRUNETT^k,
Julian BUNN^k, Michelle BUTLER^f, David CARVER^d, John COBB^c,
Tim COCKERILL^f, Peter F. COUVARES^o, Maytal DAHAN^d, Diana DIEHL^b,
Thom DUNNING^f, Ian FOSTER^a, Kelly GAITHER^d, Dennis GANNON^g,
Sebastien GOASGUEN^l, Michael GROBE^q, Dave HART^b, Matt HEINZEL^a,
Chris HEMPEL^d, Wendy HUNTOON^h, Joseph INSLEY^a, Christopher JORDAN^b,
Ivan JUDSON^a, Anke KAMRATH^b, Nicholas KARONIS^{m,a}, Carl KESSELMANⁿ,
Patricia KOVATCH^b, Lex LANE^f, Scott LATHROP^a, Michael LEVINE^h,
David LIFKA^s, Lee LIMING^a, Miron LIVNY^o, Rich LOFTⁱ, Doru MARCUSIU^f,
Jim MARSTELLER^h, Stuart MARTIN^a, Scott MCCAULAY^g, John MCGEE^j,
Laura MCGINNIS^h, Michael MCROBBIE^g, Paul MESSINA^{a,k}, Reagan MOORE^b,
Richard MOORE^b, J.P. NAVARRO^a, Jeff NICHOLS^c, Michael E. PAPKA^a,
Rob PENNINGTON^f, Greg PIKE^c, Jim POOL^k, Raghu REDDY^h, Dan REED^t,
Tony RIMOVSKY^f, Eric ROBERTS^d, Ralph ROSKIES^h, Sergiu SANIELEVICI^h,
J. Ray SCOTT^h, Anurag SHANKAR^g, Mark SHEDDON^b, Mike SHOWERMAN^f,
Derek SIMMEL^h, Abe SINGER^b, Dane SKOW^a, Shava SMALLEN^b,
Warren SMITH^d, Carol SONG^c, Rick STEVENS^a, Craig STEWART^g,
Robert B. STOCK^h, Nathan STONE^h, John TOWNS^f, Tomislav URBAN^d,
Mike VILDIBILL^{b,r}, Edward WALKER^d, Von WELCH^f, Nancy WILKINS-DIEHR^b,
Roy WILLIAMS^k, Linda WINKLER^a, Lan ZHAO^c and Ann ZIMMERMAN^p

^aUniversity of Chicago/Argonne National Laboratory; ^bSan Diego
Supercomputer Center; ^cPurdue University; ^dTexas Advanced Computing Center;
^eOak Ridge National Laboratory; ^fNational Center for Supercomputing Applications;
^gIndiana University; ^hPittsburgh Supercomputing Center; ⁱNational Center for
Atmospheric Research; ^jUniversity of North Carolina; ^kCalifornia Institute of
Technology; ^lClemson University; ^mNorthern Illinois University; ⁿUniversity of
Southern California Information Sciences Institute; ^oUniversity of Wisconsin;
^pUniversity of Michigan; ^qIndiana University – Purdue University Indianapolis;
^rSun Microsystems; ^sCornell University; ^tMicrosoft

Abstract. TeraGrid is a national-scale computational science facility supported through a partnership among thirteen institutions, with funding from the US National Science Foundation [1]. Initially created through a Major Research Equipment Facilities Construction (MREFC [2]) award in 2001, the TeraGrid facility began providing production computing, storage, visualization, and data collections services to the national science, engineering, and education community in January 2004. In August 2005 NSF funded a five-year program to operate, enhance, and

expand the capacity and capabilities of the TeraGrid facility to meet the growing needs of the science and engineering community through 2010. This paper describes TeraGrid in terms of the structures, architecture, technologies, and services that are used to provide national-scale, open cyberinfrastructure. The focus of the paper is specifically on the technology approach and use of middleware for the purposes of discussing the impact of such approaches on scientific use of computational infrastructure. While there are many individual science success stories, we do not focus on these in this paper. Similarly, there are many software tools and systems deployed in TeraGrid but our coverage is of the basic system middleware and is not meant to be exhaustive of all technology efforts within TeraGrid. We look in particular at growth and events during 2006 as the user population expanded dramatically and reached an initial “tipping point” with respect to adoption of new “grid” capabilities and usage modalities.

Keywords. Grids, distributed computing, computational science, infrastructure, high-performance computing

Introduction

The TeraGrid¹ facility is an integrated portfolio of more than twenty high-performance computational (HPC) systems, several specialized visualization resources and storage archives, and a dedicated continental-scale interconnection network. *Policy and planning* integration allows the national user community to request access through a single national review process and use the resources of the facility with a single allocation. *Operational and user support* integration enables the user community to interact with many distinct resources and HPC centers through a common service, training, and support organization – masking the complexity of a distributed organization. *Software and services* integration creates a user environment and standard service interfaces that lower barriers to porting applications, enable users to readily exploit the many TeraGrid resources to optimize their workload, and is catalyzing a new generation of scientific discovery through distributed computing modalities.

The TeraGrid mission is to advance science through three integrated initiatives:

- **Deep: Enable Terascale/Petascale Science:** TeraGrid will enable scientists to pursue scientific discovery through an integrated set of Terascale resources and services.
- **Wide: Empower Communities:** TeraGrid will make Terascale resources and services broadly available through partnerships with community-driven service providers.
- **Open: Provide an Extensible Foundation for Cyberinfrastructure:** TeraGrid will provide, and use where provided by others, a set of foundational services and resources to support nation-wide cyberinfrastructure, using open standards, policy, and processes.

The user community that relies on this national facility has dramatically expanded, from under 1,000 users in October 2005 to over 4,000 users at the close of 2006. Nearly 2000 of these are new users with development allocations to explore TeraGrid, port their codes, and incorporate HPC services into their science (§1.1).

¹ The “TeraGrid” project name, chosen in 2001, now more appropriately describes the individual resources, however the aggregate capacity of TeraGrid computing resources was over 700 Teraflops by early 2008 and will exceed one Petaflops by the end of 2008.

TeraGrid resources are also growing exponentially. In early 2006 the largest capability computing resources within TeraGrid were 10–15 Teraflops and ~2,000 processors. By the end of 2007 the largest resource, an NSF-funded system at TACC, will be over 500 Teraflops and 60,000 processor cores, and similar scale systems are planned for 2008, 2009, and 2010 [3]. Storage systems are also growing significantly. Thus, the TeraGrid team is beginning a multi-year challenge to work with the user community to provide training, porting, and optimizing support in order to fully exploit this fundamentally new scale of capability moving into the Petascale regime, while continuing to support a growing user community accessing dozens of resources nation-wide.

The multi-level integration of the TeraGrid facility is also enabling new usage modalities – and corresponding new user communities – that harness HPC, storage, visualization and data resources through advanced software applications and services, often through web portals (§1.2). The introduction of TeraGrid-wide distributed computing building blocks (§2) such as information services, remote job submission, single sign-on, parallel file transfer, and workflow support – in part based on emerging web services technologies – has catalyzed a set of discipline-specific, community-provided “Science Gateways” (§3.2). Gateways interact with TeraGrid resources through these services and related policies to serve communities of 100’s to 1000’s of scientists and educators.

Enabling Petascale science, supporting the increasing number of new users, and the growth in adoption of new usage modalities and science gateways all require a coordinated approach to building and sustaining a workforce that can fully realize the promise of cyberinfrastructure. Our user support and operations teams leverage the expertise across the TeraGrid partner institutions (§3) and our education, outreach, and training work is focused on a comprehensive set of programs – “HPC University” (§3.3).

The TeraGrid facility and organizational model consists of a set of independent, cooperating resource providers (RPs) working together with a Grid Infrastructure Group (GIG), which facilitates coordination, software and service integration, operations, management, and planning [4]. The GIG is a distributed team with staff located at multiple TeraGrid RP sites as well as other partner institutions. TeraGrid governance borrows from concepts developed in other types of organizations such as open source software projects and standards bodies, which harness the efforts and creativity of many independent participants. Policy and key decisions regarding all aspects of the TeraGrid facility are developed and approved through a forum comprised of representatives from each of the RPs and from the GIG. Results of these decisions are recorded in a persistent document series with a record of consensus among representatives.

In this paper we provide an overview and analysis of TeraGrid in four major sections. First we examine TeraGrid resources and usage as of late 2006, with an analysis of usage modalities and growth. We then turn to three aspects of TeraGrid integration: software and services, user and operational support, and policy and planning.

1. Resources, Usage, and User Community Analysis

TeraGrid resources include computational, storage, and visualization systems as well as specialized data collections and information management capabilities. TeraGrid RP sites are interconnected via a dedicated, optical, wide-area network with individual site connections ranging from 10 Gb/s to 30 Gb/s (Fig. 1). Each computational resource




Figure 1. TeraGrid Partner Sites.

Table 1. TeraGrid Resource Growth from 2004 to 2006

| Resources | 2004 | 2005 | 2006 |
|--|-------|-------|-------|
| HPC Resources | 8 | 16 | 23 |
| Storage Resources (supporting allocations) | 3 (0) | 6 (0) | 9 (3) |
| Data Collections | 40 | 90 | 101 |
| Science Gateways | 0 | 11 | 20 |

provides, in addition to traditional local user environments and services, a set of coordinated TeraGrid software and services that enable TeraGrid-wide capabilities such as single sign-on, common allocations and accounting, or advanced features such as workflow. We begin with an overview of the portfolio of resources, the policies and structures that grant access to users, and an analysis of the user community and their use of the TeraGrid facility. Subsequent sections will detail the software and services, user and community support and engagement, and overall TeraGrid organization.

1.1. Overview of TeraGrid Resources and Usage

Many metrics can be used to examine the adoption of a given set of resources and services and their impact on the science and engineering research and education. In Table 1 we summarize four key growth metrics: resource portfolio, allocation awards granted, resource usage, and the size and nature of the user base. In this section we examine these four metrics.

1.1.1. Resource Portfolio

During 2006 the number of TeraGrid HPC computational resources increased from 16 to 23, expanding aggregate computational capacity by a factor of 2.5, from roughly 120

to nearly 300 Teraflops. TeraGrid's largest computational resources as of early 2007 range from 50–100 Teraflops, with the smallest under 5 Teraflops. The overall computational portfolio includes nearly every type of system architecture, microprocessor, and operating system available for high-performance computing.

TeraGrid HPC systems are owned and operated, in most cases, by the local resource provider institution. Some of the systems were purchased using NSF awards while others were purchased with local institutional or other funding sources. The resources are provided to the national community, through the TeraGrid facility, based on cooperative agreements between the resource providers and NSF. These agreements specify funding levels for resource and user support as well as the percentage of the resource that will be made available through the national allocations process (described below in §1.1.2).

Five TeraGrid RP sites provide storage archives supporting long-term data management. Together these archives currently hold approximately 10 Petabytes of user data, up 50% since 2005.

Users access resources through available wide area Internet access including commercial as well as Internet2 and National Lambda Rail national backbone networks. Interconnection of TeraGrid resources themselves employs a dedicated nationwide optical backbone network with hubs in Los Angeles, Denver, and Chicago. RPs are responsible for maintaining connectivity to other TeraGrid resources to support high-performance data transfer and resource interaction (e.g. workflow), typically through one or more 10 Gb/s connections to a TeraGrid network hub.

A high-performance wide-area parallel filesystem supports tight coupling between TeraGrid resources at three sites using IBM's GPFS-WAN [5] and harnessing TeraGrid's dedicated optical network. Over 500 Terabytes of storage at SDSC can be mounted for remote file I/O from SDSC, ANL, and NCSA. In 2007 several additional RPs will begin offering this service and multiple TeraGrid RP sites are experimenting with alternative wide area distributed filesystems such as Lustre [6]. As of early 2007, filesystems at Indiana University, based on Lustre, can be mounted from TeraGrid systems at NCSA, PSC, and ORNL.

Increasingly, users are also interested in making data collections available in standardized fashion to enable their use in grid applications, workflow, etc. There are over 100 data collections available through TeraGrid, and in 2006 it became clear that a common set of information would be needed in order to ensure that users could readily find and access data collections. We developed a set of minimum requirements for TeraGrid data collections that includes a general description and information about data provenance, access mechanisms, and necessary metadata. Based on these requirements [7] we maintain a data collections directory within the TeraGrid User Portal.

1.1.2. Peer-Reviewed Access to TeraGrid Resources and Services

Users access TeraGrid based on allocation awards made through a national peer-review process. The resource allocation committee (RAC) is a rotating team of several dozen computational scientists, from a variety of disciplines, serving 2- to 3-year terms. RAC members are nominated by TeraGrid resource provider sites, with input from program officers at the National Science Foundation.

Eligibility for TeraGrid use is limited to researchers (post-doctoral included) or educators at U.S. academic or non-profit research institutions. A qualified advisor may

apply for an allocation for his or her class, but high school, undergraduate, and graduate students may not be principal investigators (PIs).

Review criteria are related to the computational requests and appropriate resource use, rather than a review of the underlying scientific theory or research objectives. As proposals are typically associated with a funded research program, the scientific peer review has already taken place and is not repeated by the TeraGrid RAC.

The allocations process has traditionally been focused primarily on computational resources, but was expanded to include major storage requests in 2006. In 2007 the process will be expanded further to facilitate input from the RAC regarding the allocation of dedicated support staff to assist specific projects (see §3.1). This process will involve requests for assistance in units of FTE-months. The RAC will rank these requests and these rankings will be used as input to the user support staffing allocations decisions made by management at TeraGrid RP sites.

Computational Allocation Awards

Computational allocations are measured in Service Units, or “SUs.” TeraGrid SUs translate to CPU hours on a given resource based on a “normalized unit” (NU) which is the equivalent of one hour of CPU time on a Cray X-MP. The relative performance rating for a given resource is taken by scaling performance on the HPL [8] benchmark (a standard component of determining the Top 500 rankings [9]). A rough conversion for modern processors is 1 SU = 20 NU.

Proposals are grouped into three categories based on the number of SUs requested. Large requests (currently defined as >500 k SUs) are reviewed semi-annually and medium requests (30–500 k SUs) are reviewed quarterly. Allocations are granted for periods of 12 months, with extensions to 18 months upon request. Small requests (<30k SUs) are reviewed on an ongoing basis by an internal TeraGrid review team using the same qualifying criteria. These requests are generally new projects exploring how TeraGrid resources may help the project’s scientific goals, and often involving benchmarking and code porting or for classroom instruction.

In addition to tracking allocations by overall size of award, as shown in Table 2, there are two types of allocations granted for TeraGrid computational resources:

- ***Specific*** allocations are tied to a particular TeraGrid resource, at a particular site.
- ***Roaming*** allocations are usable on any TeraGrid compute resource.

Users may request either of these types of allocations or a combination of the two. Many projects also have multiple *specific* type allocations for different machines. To illustrate the overall mix of *specific* versus *roaming* allocations for large and medium sized awards, we examine the 88 medium allocations in detail in Table 3. All development awards are *roaming* allocations.

Most TeraGrid allocation awards involve a principal investigator (PI) and a small group of collaborators and/or students. However, some awards are used to support larger communities of dozens or even hundreds of users through science gateways (§3.2). We refer to these as *community allocations*.

Storage Resource Allocation Awards

Traditionally, access to storage archives has not been regulated for TeraGrid users – any project has been permitted to store as much data in tape archives as required for their project without special arrangement. Due to sustained, exponential growth in stor-

Table 2. TeraGrid Allocations, Usage, and User Community Growth from 2005 to 2006

| Allocations | 2005 | 2006 | % Growth |
|-----------------------------------|-----------|-------------|-------------|
| Large proposals awarded (new) | 62 (13) | 88 (22) | 42% (69%) |
| Medium proposals awarded (new) | 70 (50) | 160 (92) | 129% (84%) |
| Dev. proposals awarded (new) | 123 (115) | 229 (209) | 86% (82%) |
| Active TeraGrid PIs | 361 | 1,019 | 182% |
| Usage | | | |
| NUs Requested | 1.3 B | 2.96 B | 128% |
| NUs Awarded | 844 M | 1.92 B | 127% |
| NUs Available (max) | 881 M | 2.23 B | 153% |
| NUs Delivered | 565 M | 1.28 B | 127% |
| NUs used by TG Staff | 10.4 M | 10.1 M | -3% |
| Jobs run | 594,756 | 1,686,686 | 184% |
| Users (Total) | | | |
| Users with accounts during 2006 | 1,712 | 4,190 | 145% |
| Users charging jobs during 2006 | 876 | 1,731 | 98% |
| Users with accounts on 31-Dec | 1,468 | 3,126 | 113% |
| User Institutions (charging jobs) | 151 | 265 | 75% |
| US states (charging jobs) | 37 | 47 | 27% |
| Users by Allocation Size | | | |
| Large Users (# charging jobs) | 509 (238) | 1,152 (496) | 126% (108%) |
| Medium Users (# charging jobs) | 542 (248) | 1,087 (423) | 101% (71%) |
| Dev. Users (# charging jobs) | 661 (365) | 1,948 (783) | 195% (115%) |

Table 3. Breakdown of allocations for the 88 large (>500 k SU) projects

| | No roaming | Roaming |
|------------------------|------------|---------|
| Multiple resources | 53 | 4 |
| Single resources | 26 | 0 |
| No specific allocation | 0 | 5 |

age for those projects dealing with very large data sets and the associated costs of managing this data, the TeraGrid team evaluated storage costs, trends, and user storage requirements in 2006 [10]. Based on this analysis, the TeraGrid peer-review allocations process now requires proposals from projects that anticipate needs for long-term tape storage above a threshold defined as a function of computational allocation. Requests for tape storage independent of a compute allocation can also be requested. Similarly, projects requiring long-term dedicated space on the TeraGrid wide area parallel file system (GPFS-WAN) obtain this space through the peer-review process.

1.2. Analysis of Technology Usage Patterns of the TeraGrid User Community

During the allocation proposal process each PI designates a discipline area, selecting from a menu of NSF science divisions. Usage by discipline is summarized in Fig. 3. Note that while there are 10 disciplines that account for 94% of TeraGrid usage, there are 20 disciplines that collectively consume the remaining 6%. At the same time these 20 disciplines account for nearly 30% of the overall user population.

As shown in Table 2, overall TeraGrid HPC computational resource usage grew by over a factor of two (127%) during 2006, and the number of jobs executed grew by nearly a factor of 3 (184%). The large increase in jobs can be attributed in part to in-

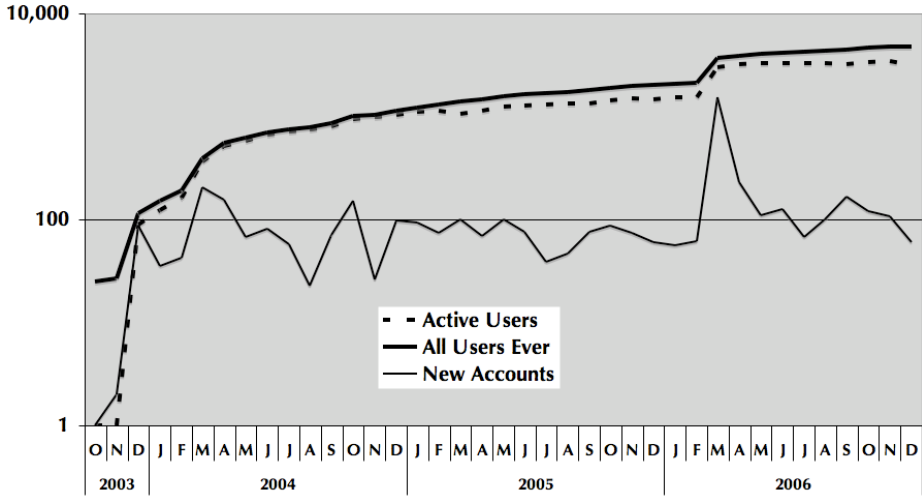


Figure 2. The TeraGrid user population grew by 2800 users during 2006, including a one-time addition of order 1500 users when NCSA and SDSC Core resources were integrated in April 2006.

creased usage of Condor to support large numbers of single-processor jobs. However, even excluding the Condor usage there was a 35% increase in the number of jobs executed in 2006.

We look at several measures, each of which produce different (but correct!) answers to the question “how many users are there?” As shown in Table 2 and Fig. 2, we differentiate between the following three groups of “users:”

- Current:** The group of people with user accounts.
- Active:** The subset of current users who have run one or more jobs (active) in a particular reporting period.
- Cumulative:** The group of people who have, or at one time have had, user accounts.

Differences between these measures are due to effects such as turnover in the user community and cleanup effects as allocations and machines are retired. We use snapshot numbers to track user population growth to minimize these effects on our estimate of current user population.

The growing user community reflects greater geographic distribution as well as growth in numbers during 2006, with users from 114 new institutions, active users in almost every US state and eight users from Puerto Rico.

Beyond the overall measure of the user community, however, effective planning and resource allocation requires an understanding of how the users are utilizing the facility. Traditional measures of system usage (jobs executed, bytes moved, etc.) are useful for examining individual systems, but they provide only a limited, indirect picture of the use of resources in coordinated use modalities as are becoming more and more common in distributed facilities such as TeraGrid.

In 2006 we began to collect additional data such as software use and usage of particular distributed services and we are in the process of expanding the number of “markers” we track. These “markers,” combined with the experience of our user sup-




Figure 3. Publications and Usage by Science Discipline. Figure courtesy Dave Hart, SDSC.

port staff through their direct interactions with users, begin to reveal patterns of use. Shown in Fig. 4, these markers include:

- Remote Job Submission and Workflow: Remote job submission logs. GRAM [11] usage shows the trend in remote job submission (jobs and users) to five frequently used resources.
- Condor: Initial Purdue Condor flock use in 2006 was heavy and grew significantly, with a population of several dozen TeraGrid users.

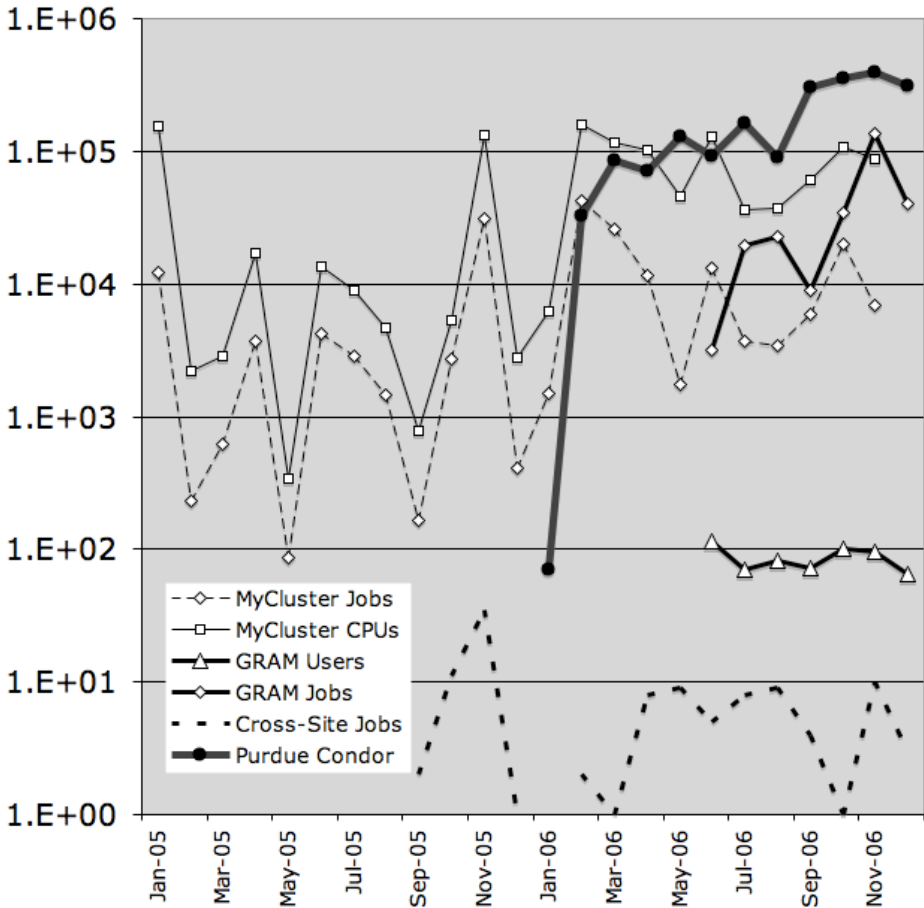


Figure 4. Initial metrics for CY2005-6 indicate growth in a sampling of distributed use modalities including MyCluster, GRAM remote job submission, and cross-site application runs. Many other modes of use are not shown in this graph, such as traditional job submission (see Table 4).

- **Parameter Search:** Use of the MyCluster [12] parameter search and ensemble simulation tool. MyCluster users increased from 24 to 36 this year.
- **Co-Scheduling:** Cross-site reservations made to our manual co-scheduling service. 15 projects ran 109 cross-site jobs. In terms of duration, 50% ran for 2–8 hours and 11% for more than 32 hours. With respect to size of jobs, 68% used 64–256 processors 15% of the jobs used more than 1024 processors.

Based on the markers above and on discussions among user services, science gateway, and technical staff members, we estimate of the number of active users (those who executed at least one job in 2006) whose typical interaction with the TeraGrid facility is best described in one of five broad categories as shown in Table 4. We note, however, that this is a very simplified overview because many teams use multiple of these modalities. For example, even those traditional HPC teams who predominantly compute in batch mode on a particular system will also use roaming allocations for opportunistic throughput enhancement.

Table 4. Estimated number of users and percent of user population for each modality of use

| Use Modality | % of Active Users |
|---|-------------------|
| Batch Computing on Individual Resources | 45% |
| Exploratory and Application Porting | 35% |
| Workflows, Ensemble and Parameter Sweep | 10% |
| Science Gateway access | 5% |
| Remote Interactive Steering and Visualization and Tightly-Coupled Distributed Computation | <5% |

In 2006 there were 1,702 active users (see Table 2), roughly half of which (919) were using large or medium allocations that are primarily – but by no means exclusively – of the *specific* type. The remaining 783 users were using development allocations that are of the *roaming* type. We find that roughly half of these development users have submitted jobs to multiple TeraGrid resources, thus we characterize them as “exploratory and application porting” users. Below we describe the six different usage modalities.

1.2.1. Batch Computing on Individual Resources

The largest group of consumers of TeraGrid resources is the community of experienced HPC users who run very large jobs through batch queues. In some cases, a small research group will have a single *specific* allocation on a platform particularly suited to their work. In other cases groups will have multiple *specific* allocations to make use of a number of TeraGrid resources. These are typically large teams who are early adopters of HPC systems and adept at porting and optimizing their codes. Collaborators on these large teams may divide the simulation workload into subprojects to be executed by one designated co-PI on one designated platform. Many of these teams are beginning to explore multi-site usage scenarios through *roaming* allocations in addition to their *specific* allocations, and several of these teams use only *roaming* allocations.

1.2.2. Exploratory and Application Porting

This category represents users who have received development awards and are porting and benchmarking codes. As shown in Table 2, development award users in 2006 accounted for 47% of all TeraGrid users and 46% of those users who submitted at least one job. Roughly half of these users are interacting with multiple resources in an exploratory way, including porting and benchmarking applications to different TeraGrid resources.

While these users do not consume nearly the amount that the medium or large allocation user communities consume, they are the largest population of TeraGrid users – nearly twice the size of the medium allocation community and 70% larger than the large allocation community. Development awards are *roaming* allocations to enable users to experiment with a variety of TeraGrid systems. The base level of familiarity afforded by TeraGrid’s coordinated software and services on all TeraGrid platforms makes it easier for new users to do so without having to learn the nuances of each of the dozens of individual systems. This is particularly useful in support of benchmarking codes on different TeraGrid machines, which is the traditional objective of development awards and is a required component of proposals for medium and large allocation awards.

1.2.3. Workflows, Ensemble and Parameter Sweep

Workflows are computations and data analysis tasks that are composed of a sequence of related but distinct jobs. This might include one or more data preprocessing steps, e.g. data assimilation and cleaning, followed by a series of computational steps in which one or more steps may depend upon a preceding step, followed by post analysis. Ensembles and Parameter sweeps are actually a category of workflow in which a large number of identical tasks are run logically “in parallel” followed by an ensemble analysis step. Even a job control script with more than one subtask is a type of workflow.

However, TeraGrid is ideal for much more general cases of workflows in which different tasks are executed on different resources. A workflow control process manages the orchestration of the workflow. Typical examples of these workflows use tools such as Condor’s *DAGman* [13], *Kepler* [14], *Taverna* [15] and *BPEL* [16] (the standard web service workflow language.) It is not easy to distinguish jobs that are managed by one of these workflow engines, because the individual tasks look identical to other submitted jobs. However, most use either Condor or GRAM for the job submission and thus the corresponding “markers” shown in Fig. 4 are indicative of the growth of this user community. For example, we find that monthly use of remote job submission has grown from several thousand jobs in mid-2006 to several tens of thousand jobs per month by late 2006. The number of users remotely submitting jobs in this same period of time was roughly constant at 90–100, however many of these “users” are actually community allocations associated with science gateways, which use workflow tools to orchestrate complex job scenarios on behalf of the tens to hundreds of gateway users.

1.2.4. Science Gateway Usage

A typical science gateway user is interacting with a web portal, invoking applications specific to the science community supported by that gateway. Gateways provide specific application services to their user communities, executing those applications on TeraGrid platforms on behalf of those users. For security reasons, special authorization and authentication provisions apply to community user accounts, which are usually constrained to execute a fixed set of commands on the TeraGrid systems. There currently are about 25 active community allocations. For both this type of usage and the visualization usage discussed below, we are seeing demand for a different type of allocation: small but persistent awards to gain periodic access to specialized resources. As of early 2007 we are receiving reports from gateway providers with hundreds of new active users accessing TeraGrid through their community accounts, indicating rapid growth in this sector of the user community.

1.2.5. Remote Interactive Steering and Visualization and Tightly Coupled Distributed Applications

Over the past year, we have seen a growth in the need for remote interactive visualization that can only be accomplished with specialized graphics hardware. Some of these users have been traditional HPC users in the past and now have need for more sophisticated visualization capabilities; and some of these users have observational or experimental data that they are analyzing using TeraGrid resources. Users typically stage their data on the platform of interest, log into the machine, prepare the remote interac-

tive session and launch the job giving them scheduled remote interactive access to the resources. This can also be done in a variety of ways including use of the TeraGrid Visualization Gateway [17] as well as via workflows using tools such as “Portals Direct I/O” (PDIO) [18].

Another small pioneering group of users employ multiple TeraGrid resources in tightly coupled applications, comprised by two general cases: functionally decomposed simulations and data-decomposed simulations. In the first case a simulation consists of large subtasks, each of which is best suited to a specific HPC architecture. For example, model building and definition of initial conditions is done on a large-memory SMP (site A) followed by a simulation on a very large MPP (site B), and by data analysis back on the SMP (site A). In the second case, a simulation is distributed in order to harness aggregate computing or memory capacity of multiple platforms. In these cases, MPICH-G2 [19] is used to distribute a single parallel job across several Clusters, necessitating advanced capabilities such as co-scheduling.

2. Creating Sustainable Cyberinfrastructure: Software and Services Architecture

The dramatic growth of the TeraGrid user community, dominated by development allocations that are typically new HPC users, reinforces our strategy to provide a coordinated software environment across TeraGrid resources. This coordinated environment enables users to roam between machines with a single allocation, a single sign-on, and access to a common core software environment of compilers, libraries, and tools. In parallel, the science gateways initiative – with ten partner gateways in 2005 and over 20 by late 2006 – relies on a common set of interfaces, specifications, and policies that allow the gateway developers to interact in a consistent way with the heterogeneous set of TeraGrid resources.

2.1. Software Services and Capabilities

Initially the Coordinated TeraGrid Software and Services (CTSS) involved a large set of software components including the Globus Toolkit, Condor, and other capabilities. In 2006 we added capabilities to support key science use cases and to improve the robustness and scalability of existing capabilities. These enhancements were done in the context of moving toward a *Service Oriented Architecture* that relies on emerging Web Services technologies. This architecture shift also involves modularization of our software to reduce packaging, deployment, and support costs while enabling us to more rapidly respond to the software requirements of our user community. Finally, this approach introduces the potential for our users to combine TeraGrid services with other components of national cyberinfrastructure – such as campus authentication and authorization services through Internet2’s Shibboleth [20] framework.

CTSS, the mechanism by which we deliver a set of common software capabilities to TeraGrid users, includes both local software packages and remote service interfaces, providing TeraGrid users with a common set of expectations for the software and services that will be present on any TeraGrid resource. We deployed the third version of CTSS in mid-2006 and are presently deploying the fourth version, CTSS 4, scheduled to complete in mid-2007.

2.1.1. New Software Capabilities in CTSS 3

The third version of CTSS was the first new version since completion of TeraGrid construction in 2004, and it represents a significant improvement over CTSS 2 with the addition of capabilities requested by our users and science gateway partners as well as process changes designed to streamline the deployment and support work of GIG and RP staff.

The key requirements that drove CTSS 3 design were as follows.

- Provide a suite of service interfaces to enable the science gateway operational model.
- Improve cross-site file transfer performance for mainstream users.
- Provide software tools to support the popular parameter sweep usage scenario.
- Provide software tools to support advanced, multi-site MPI applications.

In response to the first requirement, CTSS 3 included new service interfaces from the Globus Toolkit 4.0 (GT4) [21]. These include the WS GRAM service (for remote job submission and management), the RFT service (for managing file transfers), and the MDS4 Index Service (the basis for a TeraGrid-wide information service). These service interfaces provide high-quality mechanisms for science gateways to use TeraGrid resources via the popular Web services programming model. Usage data shows that 236,000 jobs were submitted to TeraGrid systems via the WS GRAM interface after the CTSS 3 production date in June (Fig. 4).

In response to the second requirement, CTSS 3 includes two important data movement capabilities. The first, striped GridFTP server, allows resource providers to construct multi-node data movement services that can make full use of TeraGrid's dedicated national 10–30 Gbps network for large file transfers. Usage data collected in 2006 shows that these GridFTP servers moved, on average, between 0.6 and 1.6 TB of data into and out of TeraGrid systems every day. To make this capability more accessible to new users, CTSS 3 includes a second capability: *tgcp* (TeraGrid copy). *Tgcp* presents a familiar Unix secure copy (*scp*) syntax to perform serial, striped, and reliable (RFT) GridFTP file transfers. *Tgcp* also uses knowledge about TeraGrid's GridFTP server configurations to automatically apply tuning parameters that optimize file transfer performance and to select appropriate service endpoints. With this capability, users with no little or no knowledge of the specific systems or network configurations (or GridFTP) can benefit from tuned, high-performance file transfer services.

In response to the third requirement, CTSS 3 included MyCluster [22], developed at TACC. MyCluster works in conjunction with Condor to simplify the execution of large-scale parameter sweep applications. CPUs on multiple TeraGrid systems are allocated using standard job submission interfaces and then used to execute parallel tasks. Figure 4 shows the growth in MyCluster usage during 2006.

The fourth requirement arose from the fact that a small number of early adopter science teams have found that they can now solve previously intractable problems by executing their simulation codes across multiple TeraGrid resources, optimizing their algorithms to compensate for wide-area communication latency. CTSS 3 includes an updated, more robust, and better-documented deployment of MPICH-G2, a tool that supports both local and wide-area execution of applications built using the standard MPI programming model. Several user teams have used this deployment of MPICH-

G2 for award-winning simulation runs that are on the leading edge of scientific simulation capabilities.

2.1.2. Use of TeraGrid's Integrated Software

Understanding usage and usage trends is essential to identify and prioritize improvements. Most TeraGrid service interfaces (e.g., ssh, GridFTP) record usage information in local log files, though the kinds of information tracked vary from service to service. This log data can be collected from multiple systems to produce usage reports. However, several services (from the Globus Toolkit in particular) are instrumented to report usage to a central *listener* service, where data is stored in a database.

In addition to the software and capability use data we described earlier, we note the following adoption indicators:

- **Single sign-on.** In the second half of 2006 there were 597 unique users who used TeraGrid's MyProxy server, which enables access through a single login (credential) to all TeraGrid systems.
- **Remote Job Submission.** TeraGrid GRAM interfaces (see Fig. 4) were used by at least 50 individual and community accounts across TeraGrid, supporting at least 526,000 job submissions across TeraGrid sites from June through December. The GRAM interface supports science gateway integration (e.g., NanoHUB, BioPortal), large-scale parameter studies (e.g., MyCluster users), and large-scale science workflows (e.g., SCEC, GADU).
- **Striped GridFTP.** TeraGrid GridFTP servers moved between 0.6 and 1.6 TB of data per day on average during 2006 into and out of TeraGrid systems. Data movement illustrates the use of TeraGrid systems as elements of the end-to-end scientific workflow.

2.1.3. New Capabilities Developed for CTSS 4

The primary driver for CTSS 4 is to reduce deployment and support effort required by RP sites while increasing the control an RP site has regarding the services offered through TeraGrid. By improving the capability delivery process we also enable improved stability and reduced cost for RPs. By modularizing CTSS we allow for more agile deployment of new capabilities – such as upgrading or adding capabilities without changing the entire CTSS deployment.

We will complete the deploy CTSS 4 in mid-2007, introducing significant improvement in our ability to deliver new capabilities to the user community by decentralizing the integration process that had been used for CTSS 1–3. Our new process allows other software providers (and RPs) to add capabilities to CTSS independently. CTSS 4 deployment and enhancement will employ a formal, open change management process for proposing and deploying new CTSS capabilities.

We are deploying a number of new capabilities with CTSS 4, many of which are driven by the requirements of science gateways and the TeraGrid user portal. These include:

- Queue prediction using the Batch Queue Predictor service from the Network Weather Services team [23] – This will allow users to compare the expected queuing delay of a given job on various TeraGrid systems.

- Science Gateway Audit Interface – This allows science gateway providers to track the use of their community allocations based on usage by local (gateway) users.
- Integrated Information Service – This enables automatic service registration, service discovery interfaces, and automated mechanisms for updating documentation.
- CTSS 4 Science Workflow Support kit – This provides documentation and tools to support this style of scientific operation, including mechanisms to make it easier to add applications into workflow models.

2.2. Integrated Software Environment Management

2.2.1. CTSS 4 Design and Costs

Deployment of CTSS involves cooperative effort between GIG and RP staff. The GIG produces software packages for RPs that include pre-built software for individual TeraGrid systems with customized deployment, configuration, and testing instructions. RPs deploy these packages on their HPC resources and manage operational issues. GIG coordinates operational issues, including enhancing and operating Inca (our validation and verification suite [24]) and coordinating service outage review and response. The Inca system provides monitoring of CTSS and was upgraded in 2006, adding features to better identify, analyze, and troubleshoot user-level Grid failures, thereby improving TeraGrid stability. The GIG also coordinates interactions between RPs and software vendors to ensure that problems and fixes identified by resource or software providers were made available to the entire TeraGrid community.

Overall, RPs estimate that maintaining CTSS on their systems requires an additional 0.25 to 1.75 FTEs beyond what they would otherwise spend on software maintenance. The large variation is driven by two dominant factors: (1) how much of CTSS would be deployed if the RP were not a TeraGrid partner, and (2) how able the RP is to take advantage of the GIG assistance. The first factor is largely determined by the degree to which the RP's resources are used by other cyberinfrastructure initiatives. The second factor is determined by the uniqueness of the RP's resources and practices. There is a wide range of diversity on both points.

2.2.2. CTSS Design and Structure

The first step to moving to a modular CTSS 4 and a service oriented architecture was to restructure the CTSS 3 software into a series of capability kits, each focusing on a specific set of related user capabilities. Examples include: remote job submission, remote login, and science workflow support. A single "TeraGrid Core Integration kit" was designed to provide the management capabilities that integrate any TeraGrid resource with the rest of the TeraGrid facility. (These include common security mechanisms, capability registration mechanisms, verification & validation mechanisms, and capability deployment and instrumentation mechanisms.) We also documented the CTSS design and delivery process so that anyone in the TeraGrid community can now define and deliver new CTSS capabilities.

Beginning with CTSS 4, CTSS capabilities will be managed in a distributed fashion by teams of experts in particular capability areas drawn from the RPs and GIG, and potentially external software partners. New capabilities can be deployed on independ-

ent schedules. The Software Working Group and GIG operations staff will continue to coordinate new capability deployment schedules throughout the TeraGrid community.

2.2.3. Software Build and Test

The NMI [25] Software Build-Test mechanisms play a key role in our software management process. A significant subset of CTSS 4 capabilities are prepared using these build-test mechanisms, which will be deployed on all TeraGrid resources. Our partnership with the Virtual Data Toolkit (VDT) team at the University of Wisconsin ensures that software prepared for CTSS and software prepared for VDT (used by other Grid projects such as OSG and EGEE) use the same versions, patches, and builds for common platforms.

2.3. Authentication, Authorization, and Accounting

TeraGrid supports two authorization methods: a) users are individually registered with TeraGrid and associated with a particular project, or b) users register with a Gateway that acts as a proxy to invoke TeraGrid services through a community account. Current work is focused on introducing more broadly the virtual organization approach of (b) while retaining the accountability benefits of (a). We have already put in place a TeraGrid Kerberos [26] Realm to support the User Portal authentication and are leveraging this to support single sign-on functions across the TeraGrid.

U.S. campuses, home for most of our users, are creating robust and interoperable Identity Management systems. Most notable among these is the *inCommon* Shibboleth federation. Working with Internet2 partners, we are deploying a testbed in 2007 for using campus credentials to authenticate to the TeraGrid. An evaluation will be held in June 2007 on whether we should proceed to an initial production deployment currently targeted for early 2008. We are working with the community to establish a set of guidelines and agreements that can readily be re-used by new participants.

To support a national allocations system that allows users to have either specific or “roaming” allocations, as outlined earlier, we employ a distributed resource accounting and accounts management system. Originally developed at NCSA, this system uses a central user and usage database and a standard messaging system to exchange accounting information as well as requests such as to create accounts or map user credentials to accounts. Each TeraGrid resource reports usage to the central database, which tracks balances on allocations and supports reporting and queries from resource providers as well as from the users via the user portal.

Load in our distributed accounting system grew by more than a factor of three in 2006. This was fueled by growth factors outlined in §1.1: seven additional HPC systems, 2500 new users, more than double the usage. Changes in use modalities also have a non-linear effect on the system. For example, parameter sweep studies supported by Condor flocks or MyCluster software can produce thousands of usage records where on a traditional supercomputer they would produce a single record.

We are exploring strategies to influence overall TeraGrid workload to exploit lower utilization on some RP systems. Given the large variety of system sizes and types within TeraGrid, we find that there is opportunity- and need- to optimize load across the system in order to provide improved service for end users. Our distributed accounting system and policies allow for non-uniform billing where an RP would set a charging rate to either promote increased – or decreased – usage based on load, queue

length, etc. A simple strategy we are exploring in 2007 is a lower rate for *roaming*, or opportunistic resource use, on lower utilized systems.

We are also investigating “on-demand” computing services that become more practical with a large set of resources than for a single center. We have developed a system called Special Priority and Urgent Computing Environment (SPRUCE [27]) that supports priority “tokens” that users can use to flag a job as “urgent,” with several levels of priority. Resource providers determine local policy for responding to urgent computing requests. For example, one resource provider may elect to suspend all running jobs in order to immediately run a high-priority urgent job while another may allow the job to go to the “front of the queue” for “next-run” status. Initial implementation of SPRUCE at the University of Chicago and Argonne National Laboratory is exploring various policies including the notion of offering a cluster as a “pre-emptible” service for a lower “price” because the user expects that his or her job may be suspended in the event that an urgent job arrives.

3. Engaging the Community

TeraGrid’s user community is not only growing in size but is also diversifying with new programs such as Science Gateways. Thus our user and community engagement strategies attempt to map to the needs of various types of users. Below we describe traditional user support strategies, early-adopter support strategies, and our Science Gateway program that effectively uses a “wholesale-retail” model. In such a model, TeraGrid support is focused on gateway service providers who in turn are engaging entire communities of users.

3.1. User Support Strategies

User support within a computing center is well understood, but differs significantly from user support in a distributed infrastructure. Beyond the complexity of a distributed system from a problem diagnosis and resolution perspective are the multi-organizational dynamics of coordinating the efforts of staff. A key focus of our coordination strategy in this area has been to develop a rich set of interactions and interconnections among support staff at multiple centers, harnessing the unique strengths each participating group brings to bear.

Our user support approach involves four integrated programs that are coordinated by the GIG and staffed from both GIG (3 FTE) and the RPs (20 FTE):

- ***Proactive*** User Training and Support
- ***Persistent*** Online Tools and Resources: Website, Knowledgebase, User Portal
- ***Responsive*** TeraGrid Operations Center (TOC) Helpdesk
- ***Advanced*** TeraGrid Applications Support (ASTA).

In addition to the 23 FTE dedicated to user support, the full user services team effectively involves over 80 staff members (many of whom are not funded by TeraGrid) from throughout the project who monitor the user services and trouble ticket distribution mailing list and participate in problem resolution.

Our user support strategy is driven by several user support requirement patterns. Table 5 summarizes the top ten issues that largely characterize the work of the user services team. The table also helps to illustrate how the team addresses various types of

Table 5. Top Ten User Support Issues/Questions

| Locally Handled | Coordinated | Centrally Handled |
|-----------------------------|----------------------|-------------------|
| Job turnaround (wait) | Job submission | Account balances |
| Job failures | Data transfers (WAN) | Access/outages |
| Code porting & optimization | | Logins/passwords |
| Third-party packages | | |
| File system problems | | |

issues. Some items are handled directly by the central helpdesk at the TeraGrid Operations Center (TOC), others are assigned to the relevant RP site support team, and some require cooperative diagnosis among different RP sites.

3.1.1. User Training and Support and Helpdesk

To ensure that new users are able to smoothly begin working, we assign a member of the user support team to each new PI with either a large or medium allocation award. This “ombudsman” will contact the user and will provide ongoing personal support. As was discussed in §1.1, we have seen a dramatic increase in new users through development awards, such that these users now make up roughly half of the user community. For these users, we provide startup instruction materials and TeraGrid help desk contact information with their “welcome packet” of materials. Much of our work in online tools and resources, including the user portal, is aimed at ensuring that we have a robust support structure in place for this most rapidly growing portion of the community.

3.1.2. Online Tools and Resources: Website, Knowledgebase, User Portal

TeraGrid online resources include the TeraGrid main website, TeraGrid User Portal, and a TeraGrid Community Wiki. The main website is the primary access point today for documentation, knowledgebase, and for accessing the allocations proposal system. A content management system allows staff from the RP sites to update site-specific information regarding resources and services. The TeraGrid wiki is the primary collaboration site for project planning, internal policy development, internal and public reports, and other activities. The TeraGrid makes use of a Knowledge Base (based on Knowledge Base Technology from Indiana University) to provide a convenient interface for users to search and find solutions to technical problems.

The common issues experienced by the support team as shown in Table 5 also drive the development of specific user tools and the evolution of the user portal. For example, several of the “top ten” issues were among the first to be addressed in the initial launch of the user portal. Users can manage logins and passwords, check allocation account balances, and use prediction tools and resource queue status information to optimize job turnaround.

The TeraGrid User Portal, launched in May 2006, provides a single point of entry for TeraGrid users to all TeraGrid processes, including cross-referenced access to the website, Wiki, and other online resources. The initial version included basic tools for users to manage allocations: allocation usage monitoring, system account directory, resource system monitoring, and user documentation. During its first 6 months of operation, more than 20% of TeraGrid users had authenticated and used the User Portal. Among the top 20 most active user portal users, six were from large allocation projects,

demonstrating that even experienced users are finding value in this gateway to TeraGrid.

We expect continued adoption of the User Portal in 2007 as we expand the available tools such as batch queue and data bandwidth prediction interfaces, and as we provide richer allocation management tools. We will also introduce the concept of science domain views to provide portal users customized default interfaces depending on the domain of science views they select.

3.1.3. Advanced Support for TeraGrid Applications (ASTA)

The Advanced Support for TeraGrid Applications (ASTA) Program associates application consultants with specific user teams, on average involving 25% of a consultant's time for 6–12 months. Each project involves a detailed scope of work that generally focuses on application software development necessary to maximize the effectiveness of the use of TeraGrid resources and capabilities toward the user's scientific goals. ASTA projects are typically attempting to harness advanced, often new, TeraGrid capabilities for which there is not yet sufficient operational experience to optimize our general support and documentation. Thus, ASTA projects help to “debug” these advanced features, including assisting us in determining their utility.

Due to the growth in demand for the ASTA program we have developed a policy to introduce peer review, whereby ASTA support can be requested through the national allocations process using the same mechanisms that grant TeraGrid allocation awards.

We have planned a new ASTA selection model that will employ the xRAC committee to review requests. Among our strategic 2007 ASTA goals are to promote:

- a) Scaling to Petascale and $\sim 10^4$ cores,
- b) Complex workflow development embedding Petascale applications assisted via (a) into the entirety of TG infrastructure, and
- c) Diversity of disciplines and modalities of resource usage.

3.2. TeraGrid Science Gateways Initiative

The Science Gateways program promotes and supports the use of HPC resources through community-designed interfaces, recognizing that many of today's scientists routinely use desktop computing applications and web browsers to conduct their work, including utilizing remote HPC resources. The gateway program began with eight specific prototypes spanning seven disciplines, each of which had external funding to build a community-specific infrastructure and an existing user community. This included seven web portals and a community Grid (Open Science Grid).

The gateway model to access resources is available to any academic developer, and this opportunity has been highlighted at the TeraGrid website as well as widely advertised through workshops and outreach events. TeraGrid support staff work directly with developers who are providing capabilities for their communities in the same fashion that our user services programs (see §3.1) assist individual HPC users.

The gateway architecture defines programming interfaces and web services that developers can use to bring resources of the TeraGrid to a community of users within the environment – generally a web portal – with which they are already familiar. This effectively adds HPC resources to the scientific, and education portfolio of these communities without introducing a steep learning curve. The adoption of this model has been rapidly growing, increasing the importance of our providing the necessary func-

tionality and documentation to enable developers to incorporate TeraGrid resources into their community infrastructure. The gateway program is supported by GIG-funded staff who focus on specific technologies required by multiple gateways, such as job audit and on-demand computing support. Initially (in 2005–6) we dedicated GIG staff members to the eight original gateways in order to focus on developing and building a scalable set of processes and policies to start the program. During 2007, the GIG-funded staff will transition from focusing on the initial gateway prototypes to forming an integration team that can assist new gateways, selected through a peer-review process similar to that used with the ASTA program. Already during 2006 this team has expanded their support to over 20 gateway partners.

Gateway work began in 2005 with a survey of 10 projects from a variety of disciplines and with a variety of access models, and the results of this study drove our initial set of priorities for gateway work as well as requirements for GIG software integration work (e.g. the requirement for web services, see §2.1). While gateway developer needs are often distinct from scientists using HPC systems at the command line, we found a significant core of common service and capability requirements shared by *Deep* HPC users and *Wide* gateway users and developers. By working with some pathfinders, we have been able to deploy needed capabilities, while educating developers on changes they need to make to operate in a shared, production environment. As we develop methods and processes for gateway integration, we implement these and document them via an online primer, so that integration is easier both for subsequent gateways and for subsequent resource providers.

3.3. Training, Outreach and Community Engagement

TeraGrid training, education, and public outreach programs leverage the efforts of RP sites in a coordinated fashion to allow for common planning, optimal event scheduling, and sharing of expertise and materials. In 2006 alone we supported over 100 training, education, and public outreach events, reaching thousands of educators and students in hundreds of secondary, undergraduate, and graduate institutions. Evaluation of these programs is extensive in terms of both internal and external (a separate NSF award to the University of Michigan) methods ranging from focus groups to surveys and interviews.

We organize our education, outreach, and training programs with a comprehensive approach we refer to as “HPC University,” where integrated scheduling and event information allow a user or prospective user to develop a personalized training plan drawing from tutorials, workshops, and other events across TeraGrid as well as in the broader international and US community. This concept involves:

- A regular series of training sessions conducted by TeraGrid RP and GIG staff to address a variety of topics from introductory to advanced topics (new user startup, user portal, parallel programming, data management, data analysis, visualization, grid computing, etc.).
- Coordinated summer institutes and workshops at RP sites to introduce users to TeraGrid resources. Training will also be provided to TeraGrid staff to ensure they are fully up-to-date on the latest tools, technologies, and methods.
- Curriculum development, with the support from the SC07-09 [28] Education Programs, working with undergraduate faculty and high school teachers on

integrating computational science, scientific computing and grid computing resources, tools and methods into the curriculum.

- Student internships at Resource Provider and GIG sites.

4. Implementation: Organization, Structure, and Governance

TeraGrid is a facility involving resources and services that are provided by multiple, autonomous, institutions. Resource providers (RPs) are independently funded, by the National Science Foundation and other sources, to deliver sophisticated portfolios of HPC resources, support, and related services. The Grid Infrastructure Group (GIG), funded through a separate grant to the University of Chicago, is a distributed team charged with a variety of integrative functions. GIG leadership and management roles are filled with individuals selected from across the project based on expertise and merit. Similarly, the majority of GIG staffing is drawn from partner institutions via subawards to partner sites, each with a detailed statement of work identifying the tasks and responsibilities of the individuals at that site who are funded as part of the GIG.

The distributed GIG function provides for several key TeraGrid-wide capabilities and services, predominantly through subawards to experts at RP sites, including:

- A TeraGrid operations center and helpdesk,
- Common services such as authentication and authorization services, information services, a user portal, and various internal and external websites,
- Integration of new capabilities (software, policy, interfaces) to address user requirements,
- A nation-wide, dedicated optical network backbone and hubs to interconnect the RP sites, each of whom is responsible for maintaining a connection to a hub on the backbone network (hubs are located in Los Angeles, Denver, and Chicago).
- Common processes such as accounting, authorization, and allocations peer review.

In addition, the GIG provides coordination for a larger set of activities that involve both GIG and RP staff, including:

- User support functions and programs such as ASTA,
- Education, Outreach, and Training initiatives,
- Overall software and service operation and coordination,
- Planning and prioritization of strategies and architecture,
- Organizational structures for coordination, policy, and governance.

4.1.1. Technical Strategies and Planning

Because the GIG function of TeraGrid is inherently distributed – at both the management and staffing levels – it provides an excellent platform for coordination and facilitation for project-wide planning. As with most institutions and projects we have a set of “working groups” that are focused on technical or service areas, involve staff from each participating institution, and are effective at ongoing coordination of services. However, working groups can readily become technical “stovepipes” and because they are responsible for operational services they tend to minimize change in order to opti-

mize for stability. As priorities and requirements change, however, it is necessary to explore changes and new approaches as well as to bring together individuals from different technical areas.

We created a structure called a “Requirement Analysis Team” (RAT) that addresses these two issues. A RAT is created for a finite period of time – generally 8–10 weeks – in order to explore a particular opportunity or challenge that cuts across multiple technical and policy realms. A RAT operates based on a written charter and the responsibility of the RAT is to create a set of recommendations. There have been nearly 20 RATs in the past two years of the TeraGrid project, each of which has produced a detailed set of analysis and recommendations that inform and in most cases define TeraGrid strategy and policy.

4.1.2. Governance: The TeraGrid Forum

The decision-making process in the TeraGrid project includes both local and collective functions. As each resource provider, and the GIG, are independent entities there is no single, top-down, authority that dictates policy and technology. Cooperative decision-making is accomplished through the TeraGrid Forum, a body including one representative from each participating institution. In March 2006 the Forum “ratified” a consensus-based democracy decision-making process where major policies are recorded, along with consensus records, in a persistent document series.

5. Conclusions

TeraGrid began as a cooperative project of four institutions and 6 resources in late 2001 and has grown to over a dozen institutions with over 20 major HPC resources. Since completing construction in 2004 TeraGrid’s user community has grown from several hundred to several thousand users, including many users who are new to the HPC environment.

As with many “grid” projects, the bulk of early use has been traditional, rather than distributed, usage modalities. However, during 2006 the TeraGrid project saw dramatic adoption of a number of new use modalities including workflow, remote job submission, and parameter sweep. During the same period of time a new generation of HPC user emerged comprised of two types of new users, each with unique expectations. The first are researchers whose base expectation is that they can access any of the TeraGrid systems rather than being tied to a particular computer at a particular computing center. The second are users who access HPC services through web portals and via web services technologies.

We believe that one of the keys to continued success at providing for the needs of these different, and growing, user communities is the coordinated provision of standard services on HPC systems, recognizing that they are not stand-alone capabilities to which users must adapt their work. Rather, they are rather building blocks users wish to incorporate into their own science environments. Critical to this approach is the need to actively engage the scientific user community in order to understand their requirements, deploying services aimed at meeting the scientific objectives of the user community.

Acknowledgments

The TeraGrid project has been supported through a variety of funding and in-kind contributions in addition to multiple grants from the National Science Foundation. State support has come from the states of California, Illinois, Indiana, Pennsylvania, and Texas. Institutional support has come from Carnegie Melon University, Indiana University, Purdue University, University of California-San Diego, University of Chicago, University of Illinois at Urbana-Champaign, University of Pittsburgh, the University of North Carolina, California Institute of Technology, and the University of Texas. Corporate support has come from Cray, Dell, IBM, Lilly Endowment, Qwest Communications, and Sun Microsystems. Several hundred staff members from partner institutions contribute to the TeraGrid facility.

References

- [1] National Science Foundation: www.nsf.gov.
- [2] "Guidelines for Planning and Managing the Major Research Equipment and Facilities Construction (MREFC) Account," www.nsf.gov/bfa.
- [3] The NSF High Performance Computing System Acquisition: Towards a Petascale Computing Environment for Science and Engineering (NSF 05-625) has proposal deadlines of 2/06, 11/06, 11/07, and 11/08 for acquisition of HPC systems to be integrated as part of TeraGrid, with anticipated \$30M available annually. The "Track 2" solicitation is available at www.nsf.gov/pubs/2005/nsf05625/nsf05625.htm.
- [4] Catlett, C., Beckman, P., Skow, D., and Foster, I. "Creating and Operating National-Scale Cyberinfrastructure Services," *CTWatch Quarterly*, Volume 2, Number 2, May 2006. www.ctwatch.org/quarterly/articles/2006/05/creating-and-operating-national-scale-cyberinfrastructure-services.
- [5] Phil Andrews, Bryan Banister, Patricia Kovatch, Chris Jordan, Roger Haskin, "Scaling a Global File System to the Greatest Possible Extent, Performance, Capacity, and Number of Users," Proceedings of the 22nd IEEE/13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST'05), 2005.
- [6] Hedges, R., Loewe, B., McLarty, T., and Morrone, C. 2005. Parallel File System Testing for the Lunatic Fringe: The Care and Feeding of Restless I/O Power Users. In Proceedings of the 22nd IEEE/13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST'05) – Volume 00 (April 11–14, 2005). MSST. IEEE Computer Society, Washington, DC, 3-17. DOI= [dx.doi.org/10.1109/MSST.2005.22](https://doi.org/10.1109/MSST.2005.22)
- [7] Data Collections requirement analysis team report, December 2006, Kelly Gaither et al. www.teragridforum.org/mediawiki.
- [8] High-Performance Linpack benchmark (HPL) is a software package that solves a (random) dense linear system in double precision (64 bits) arithmetic on distributed-memory computers. It can thus be regarded as a portable as well as freely available implementation of the High Performance Computing Linpack Benchmark. From Dongarra et al., University of Tennessee, www.netlib.org/benchmark/hpl.
- [9] The Top500 project, initiated in 1993, is a collaboration between the University of Tennessee, the University of Mannheim, and the U.S. Department of Energy Office of Science National Energy Research Scientific Computing Center (NERSC) at the Lawrence Berkeley National Laboratory (LLNL). See www.top500.org.
- [10] "Allocations and Usage Requirements Analysis Team" (AURA) report, October 2006, Richard Moore et al. www.teragridforum.org/mediawiki.
- [11] Globus Resource Allocation Manager (GRAM), www.globus.org/toolkit.
- [12] MyCluster: www.teragrid.org/userinfo/jobs/gridshell.php.
- [13] Directed Acyclic Graph Manager (DAGman): www.cs.wisc.edu/condor/dagman.
- [14] Kepler is an open source, cross-platform workflow software project: kepler-project.org.
- [15] Taverna development is led by Tom Oinn at the European Bioinformatics Institute with the main development team at the University of Manchester and funded through the Open Middleware Infrastructure Institute UK (OMII-UK) myGrid node: taverna.sourceforge.net.
- [16] Business Process Execution Language (BPEL): en.wikipedia.org/wiki/Business_Process_Execution_Language.

- [17] Binns, J., DiCarlo, J., Insley, J. A., Leggett, T., Lueninghoener, C., Navarro, J., and Papka, M. E. 2007. Enabling community access to TeraGrid visualization resources: Research Articles. *Concurr. Comput.: Pract. Exper.* 19, 6 (Apr. 2007), 783-794. DOI= [dx.doi.org/10.1002/cpe.v19:6](https://doi.org/10.1002/cpe.v19:6).
- [18] Stone, N. T. B., Balog, D., Gill, B., Johanson, B., Marsteller, J., Nowoczynski, P., Porter, D., Reddy, R., Scott, J.R., Simmel, D., Sommerfield, J., Vargo, K., and Vizino, C. "PDIO: High-Performance Remote File I/O for Portals-Enabled Compute Nodes," Proceedings of the 2006 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'06/ISBN#:1-932415-89-0/CSREA), Editor: Hamid R. Arabnia, pp. 925-930, Las Vegas, USA, 2006.
- [19] MPICH-G2, Nick Karonis, Northern Illinois University: www.niu.edu/mpi.
- [20] Internet2 Shibboleth Framework: shibboleth.internet2.edu.
- [21] WS GRAM, RFT, and MDS4: www.globus.org/toolkit.
- [22] MyCluster: www.teragrid.org/userinfo/jobs/gridshell.php.
- [23] Network Weather Service project, Rich Wolski et al., University of California at Santa Barbara: nws.cs.ucsb.edu.
- [24] S. Smallen, K. Ericson, J. Hayes, C. Olschanowsky, "User-level Grid Monitoring with Inca 2", to appear in the Proceedings of the HPDC 2007 Workshop on Grid Monitoring, June 2007.
- [25] NSF Middleware Initiative: archive.nsf-middleware.org.
- [26] Kerberos: web.mit.edu/Kerberos.
- [27] SPRUCE, Beckman et al., Argonne National Laboratory: spruce.teragrid.org.
- [28] SCxy, The International Conference for High Performance Computing, Networking, Storage, and Analysis: www.supercomp.org.