Science Gateways
Requirement Analysis Team (RAT) report

Sebastien Goasguen, John Cobb, Dennis Gannon, Jeff Gardner,
Eric Roberts, Von Welch, Nancy Wilkins-Diehr, Roy Williams.

April 2005

Abstract

Over the course of two months the Science Gateway Requirements Analysis Team interviewed the 10 original TeraGrid gateways as identified in the GIG proposal. These gateways are at the core of the TeraGrid Wide effort and will bring a new set of users to the TeraGrid through a community approach instead of an individual per user basis.

The RAT efforts were focused on getting a first order understanding of the communities involved, the technologies being used and the mode of operations while unveiling the challenges ahead for the TeraGrid team at large in our effort to enable this paradigm shift.

A list of recommendations (see section 3) to the current working groups has been developed as well as some thoughts about future efforts in web services and portal technology. As the GIG gets started and the Science Gateways staff up, the gateway working group should build on this report to refine the analysis and move towards implementation. Several actions can be already started –group accounts, security policies – while others like web services activities and OSG interaction call for some organizational structure and careful planning. A primer template is also suggested in section 5.

Overall from all the interviews a great sense of interest could be felt. Interviewees seemed eager to get started and use the TeraGrid. Now it is up to us to make it happen.

1 Introduction

Science Gateways signal a paradigm shift from traditional high performance computing use by enabling entire communities of users associated with a common scientific goal to use the national resources through a common interface. Science gateways are enabled by a community allocation whose goal is to delegate account management, accounting, certificates management and user support to the gateway developers.

In the GIG proposal, two main streams are identified TeraGrid Deep and TeraGrid Wide. While Deep addresses fundamental software and infrastructure needs to tie the resource partners (RP) together and provide enhanced experience to users, TeraGrid

Wide addresses the science Gateway effort and its goal to reach out to large community of users who would benefit from the national resources provided by TeraGrid.
In order to prepare for the 2005 TeraGrid All Hands Meeting (AHM), a Requirement Analysis Team (RAT) was formed to analyze the science gateways listed in the GIG proposal and determine commonalities between them as well as specificities. The SGW RAT aimed at laying the ground work to jump start the science gateways activity and linking them to the TeraGrid resources. This document represents the results of the RAT's work.

The initial set of ten science gateways represents a fairly complete spectrum of science community requirements. These can be classified into three categories:

- Type 1 - A gateway that is packaged as a web portal with users in front and TeraGrid services in back.
- Type 2 - Grid-bridging Gateways – often communities run their own Grid devoted to their science. In these cases the Science gateway is a mechanism to extend the reach of the community Grid so it may use the resources of the TeraGrid.
- Type 3 - A gateway that involves application programs running on users' machines (i.e. workstations and desktops), accessing services in TeraGrid (and elsewhere).

Some of these gateways are about exposing specific set of community codes so that anonymous scientists can run them. Others are making what we can call a "Metaportal", meaning a community portal that can be used in general to bring new services and applications to the community. A common trait of all three types of gateways is that they will interact with the TeraGrid through the various service interfaces provided by TeraGrid. While the gateways may be instantiated on TeraGrid resources, it is expected that many will be instantiated on community resources and be administered by the community itself.

The document is organized as follows: Section 2 presents the analysis and findings of the RAT. These findings are organized by technical topic and aim at creating a classification of the SGW. The specific topics of concern include: scheduling, security, accounts and accounting, web services and portal middleware. Section 3 is a list of proposed action items for each TeraGrid working group that the SGW RAT has identified as necessary to make the SGW effort a success. Due to recent discussions within the TeraGrid, section 4 gives some thoughts about web services. In section 5, we address the need for a SGW primer that would be used by future SGW who would want to use the TeraGrid resources. Conclusions of the RAT findings are provided in section 6. The Appendix gives a brief summary of each science gateway interview. Complete interviews can be found at http://wg.TeraGrid.org/Gateways.

2 Requirement Analysis

2.1 Accounts and Accounting

Requirements summary:

- Most Gateways need to support a variety of accounts (individual, group, student, high-end researcher, developer). Different types of accounts can have different capabilities. For example a developer can log on to a machine, compile and run any code. A group account may be restricted to running installed binaries through a portal.
- OSG provides allocations for a virtual organization (VO), which manages accounts via the Grid User Management System (GUMS)
- Most Gateways could most benefit from a scheme that allows portal registration and tracks and limits access to resources per portal user.
    - Almost all Gateways mention that they have plans to do this.
    - Almost none have this implemented.
    - Most Gateways have different classes of users – from students to higher end researchers
    - Some Gateways plan to map multiple portal users to a single account and therefore a single DN. Need to determine if functionality of these accounts must be limited to prevent security problems?
    - NVO interview notes refers to other grids that have accounting working, but I couldn't find out which projects this was
- Need to set up developer accounts with defined allocations for all Gateways now.
- About half the Gateways plan to use TeraGrid certificates, about half have other certificates that they honor. Need to have clear policies in place for certificate acceptance,
- OSG is offering their resources to TG users as well. Need to think about how this will work, will they become a TG resource provider?
- There was no common methodology for accounting and auditing across the SGWs.

2.2 Security

Gateway summaries:

- ***LEAD****:* Community allocation. Want to do own fine-grained access control through capabilities issued by community to web services running on TG resources. Certificates from NCSA. Firewalls have been an issue in the past, defining open port range.
- ***nanoHUB****:* Group accounts with restricted privileges. Expressed desire to have dynamic/shadow accounts so each user is sandboxed from other nanoHUB users.
- ***Neutron Science***: Both community users w/community allocation and individual users with own allocation. Use DOE Grids certificates, maybe their own.
- ***Flood Modeling****:* Community users running under a community allocation (one

certificate for hydrology gateway user). Users registered by portal and known by email address. May adjust security in reaction to amount of usage.

- ***Homeland Security****:* Pretty basic security model in terms of accounts, user on portal has account on relevant TG systems. Weekly or bi-weekly testing of codes may require automation? On-demand assumption of large number of resources, who makes the policy decision on when that should happen and how does it get enacted? Application codes are under export control.
- ***NVO****:* Anonymous, quickly registered, or well-known (have TG account) users. Community allocation. Stronger authentication allows for more privileges. NVO or HotGrid CA-issued certificates.
- ***HEP***: Basic portal model – Standard and Portal users have own certificates and accounts. Will use Clarens on resource.
- **OSG**: Users have DOE Grids certificates, they also honor European LCG certificates. Use VOMS. Hierarchical allocations through VOs, sub-VOs. This is really a case of needing to do federation.
- **Evolutionary Biology**: Using community account. Local CA (based on SimpleCA) for issuing user certificates. Restricted set of applications to be run through portal.
- ***NMBR:*** Basic portal model – Standard and Portal users with own accounts.

Requirements summary:

- **Community Allocations/Accounts:** A number of portals have expressed the requirement, or desire at least, to have allocations covering subsets of their community, explicitly so that these members do not have to obtain individual resource allocations. Gateways expressing this requirement were: LEAD, nanoHUB, Neutron Source (SNS), Pathogen Data (Stevens), NVO. OSG expressed a more complicated picture involving role-based mapping to accounts.
- **Anonymous users:** nanoHUB, flood modeling and NVO expressed the desire to support anonymous users (or at least users who are basically not vetted to any degree and known only by email address).
- **Dynamic Accounts:** nanoHUB has requirement that community users be sandboxed – they call this shadow account.
- **Providing Best Effort Support:** None of the gateways interviews provide a point of contact or support that is more than best effort (basically 9-5, M-F). TeraGrid cannot expect any guarantee of being able to reach points of contact regarding a gateway, particularly off-hours.
- **Certificate Authorities:** Two portals (OSG and SNS) have users with DOE Grids CA-issued certificates. OSG also has users with certificates from the European CAs involved in the LCG project. Several projects plan on (or have) establishing their own local CAs: Evolutionary Biology (Reed), SNS, LEAD, NVO (HotGrid CA).
- **Restricted Accounts:** For all the SGWs interviewed that plan on supporting community allocations, the users would need only limited privileges on TG resources. E.g. only running a prescribed set of applications on TG compute resources. For community allocations this would appear to be a highly preferred

mechanism. An implied need here is a method for managing the restricted accounts (e.g. installing new applications, updating existing ones, debugging). If developers or administrators in the community take on this function

- **Accounting and Auditing:** There was no common methodology for accounting and auditing across the SGWs.
- **Firewalls:** The issue of firewalls was raised in the LEAD interview. A set of services and ports which SGW developers should be able to count on to be available to the TeraGrid should be developed. The GT Firewall Requirements document [http://www.globus.org/security/firewalls/] would be a good start.
- **Export-controlled codes:** The homeland security community has a need to have codes with legally mandated confidentiality concerns.

2.3 Scheduling:

Gateway summaries:

- ***LEAD**:* Implication is that a central TG job submission mechanism would benefit user (with the ability to submit jobs to "the TeraGrid"), especially since users will be running "ensembles" of jobs. Users will also define workflows that will need to make preemptive, on-demand use of the TeraGrid. However, the workflows will be adaptive, meaning that they will want the ability to migrate running applications from one resource to another.
- ***nanoHUB**:* Requirement for scheduling of jobs across TG sites to provide fastest turnaround. Mentions Condor-G as a possible interest. Job size ranges from single-PE to parallel.
- ***Neutron Science**:* Scheduling not explicitly mentioned. They do mention "no MPI" and "roaming computation," implying the desire to interact with a central TeraGrid job submission interface.
- ***Flood Modeling**:* Users will specify workflow models to allow automated visualization, implying the need for orchestrating the execution of tasks and data movement.
- ***Homeland Security**:* Biggest requirement here is preemptive scheduling: knock users off resources when the need demands. A further implication of this might be co-scheduling.
- ***NVO**:* Implication is that a central TG cross-site job management interface might be useful, although they may use the same resource discovery tools as HEP, negating for TG to provide one.
- ***HEP**:* They mention "large numbers of CPUs" but it is not clear if they actually have parallel applications. They will use Clarens and MonALISA systems for resource discovery. These packages could be tuned to submit jobs to TG resources individually, circumventing any TG cross-site job management interface.
- ***OSG**:* "On-demand allocation and configuration of virtual clusters with negotiated service level agreements." As far as scheduling this kind of thing goes, probably trivial.

- ***Evolutionary Biology****:* mention of using Globus and local schedulers such as PBS/LSF.
- ***NMBR****:* No scheduling requirements mentioned in interview or gateway description on repo.

Requirements summary:

Homeland sec portal is all about getting instant nodes. Other portals talk about web services, presumably meaning either (a) somewhere are non-scheduled "instant" nodes, or (b) asynchronous services (i.e. methods are submit, monitor, notify, retrieve)

2.4 Web services

"Web services" refer to the technology that is becoming standard in industry for managing enterprise-wide application and services. These "WS-*" standard are being defined by Microsoft, IBM, Sun, HP, Oracle and the standard bodies like W3C, Oasis and GGF. It is also the core set of technology that forms the foundation of most new Grid middleware standards including the Open Grid Service Architecture.

Gateway summaries:

- ***LEAD***: lots of Web service applications on backend. There will be several persistent services including MyLead, a user metadata catalog, notification services and application factories which launch application workflows onto TeraGrid. Additional web-services are needed for resource reservation and allocation.
- ***nanoHUB*** : no web service requirements.
- ***Neutron Science***: no web services requirement yet defined.
- ***Flood Modeling***: Will most likely need Web Service interfaces to all steps in workflow (data transformations, job submissions, data movement)
- ***Homeland Security***: no web service requirements defined at this point.
- ***NVO***: building Web services for exposure of large astronomical datasets and to allow "grid-enabling" of a pop. Data analysis environment called IRAF. Job monitoring via Web services. When job finished, products available via URLs.
- ***HEP*****:** concept is to expose the computing backend available as well as analysis environments such as ROOT, Java analysis studio, CAVES, PhySh.
- ***OSG***: set of VO-specific services on which VO developed user portal and applications. Low level services. GAS – grid access service.
- ***NMBR***: WS interface to scheduler. WS interface to monitoring. Programmatic web service interfaces needed from TG and need to be provided to their end users. Standardization of Web services.

Requirements summary:

5 of 9 (Flood Modeling, LEAD, OSG, NVO, NMBR) state integration of portal with backend Web services. Many of these services will build on those available from Globus. Others are being developed within the group of developers for that community. A major issue that will arise is web-service interoperability. Globus GT4 is based on the Web Service Resource Framework proposed standards and many, if not most of the others are based on simple WS-I based models. However, it is likely that these will be able to make use of the GT4 services. Another issue that will arise is the use of web service security. Current implementations of WS-security are slow, but performance is expected to improve.

The most important set of web service requirements are standardized interfaces to resource scheduling and account management that will be provided by TeraGrid. A second, but still important issue is where these web services will be hosted. Some will be hosted on community resources but others, such as the resource and account management will need to be hosted by TeraGrid.

2.5 Portal middleware

Gateway summaries:

- ***LEAD**:* Based on OGCE codebase. Leverage TG User Portal as much as possible. Portal to most likely run on one of LEAD testbed grid machines. Single-signon w/ MyProxy. Portal should log usage to keep audit trail. GridPort. URL: http://lead.ou.edu/
- ***nanoHUB**:* In-VIGO middleware. Some users will have TG User Portal accounts. Portal action is through SSL. Users will have Portal account. Portal users mapped to TG accounts as needed. Time multiplex accounts through portal and TG. Everything in portal logged. GUI run on virtual machine. Output piped through portal to browser. Can host VM's on portal server; define, save and deploy VM's on portal server for later use by all users. Saving data in portal environment. Import/export data through browser. URL: http://www.nanoHUB.org/
- ***Neutron Science**:* Tomcat + Apache, pojo framework, custom code developed – designed to be flexible to incorporate legacy applications. Build portal on many-to-many mapping. Portal as simple username/password auth. Mech, portal carries creds transparently. Worried about 24/7 ops support. Will portal be mirrored? URL: http://home.ornl.org/ , http://tg-web1.ornl.TeraGrid.org:8080/jetspeed/SNS-Portal/index.jsp.
- ***Flood Modeling**:* GridPort, OGCE for portal and middleware. No portal framework yet but will follow TACC portal technology trends (GridSphere). Portal should track user level data. May only need to know their email address. URL: none.
- ***Homeland Security**:* Gateway should support policy, scheduling, list of resources, certs, data movement, pushing data. URL: none
- ***NVO**:* Clarens, Apache. Large cluster that runs Clarens web server + disk for

storing large datasets. Web form written in astronomer-friendly lang. Params to specify job. User authentication via X.509 cert either in browser or via proxy server. MODES: (a) no auth. (b) fill in form and get quick auth. ("hotgrid") (c) has TG account already.

- ***HEP***: Concept is exposing computing backend available, exposing analysis environments such as ROOT, Java analysis studio, CAVES, PhySh. Clarens middleware. Users will need portal account. Would like Clarens software to be part of CTSS. URL: https://tg-tst-h.caltech.TeraGrid.org:8443/clarens/ and https://plato.cacr.caltech.edu:8443/clarens/
- ***OSG****:* Portal interface to grid services. portal built upon a set of VO-specific services. Middleware: OGCE, Clarens, homegrown, vdt. Usage information collected via monitoring tools. Looking to develop real accounting infrastructure. URL: http://www.opensciencegrid.org/
- ***NMBR****:* Framework called The Seed which is a web interface for annotation and analysis of genomes. Gateway will provide access to The Seed interface. The Gateway will be a combination of open source middleware, job submission via GRAM, WS interface to scheduler, WS interface to monitoring. Needs accounting, tracking and logging. URL: http://theseed.uchicago.edu/FIG/index.cgi

Requirements Summary:

About half already have some type of portal with existing URL. Most will require some type of registration through the portal. Most require some form of single signon in order to use grid proxy certificates for interaction. 5 of 9 (Flood Modeling, LEAD, OSG, NVO, NMBRC) state integration of portal with backend Web services. 4 of 9 (NMBRC, nanoHUB, SNS, HEP) state straight portal integration with existing science middleware. 3 of 9 (OSG, LEAD, Flood Modeling) state usage of OGCE. 3 of 9 (OSG, HEP, NVO) state usage of Clarens.

2.6 User models

Three user types have been identified:

- **Standard:** User has TG account and own certificate which they manage.
- **Portal:** User has TG account, but only uses TG through portal. They authenticate with username and password.
- **Community:** User does not have own TG account, uses community allocation. They authenticate with username and password, and present community certificate to TG resources.

Security models:

- **Standard:** All users have valid TG login and cert. May also be group structure or shadow accounts. (Homeland sec, NANO, Stevens).

- **Portal:** Users log into to portal machine with user/pass. Portal uses community or local cert to authenticate to TeraGrid. Different users get different capability. May be group permissions -- LEAD calls these shared experiments (LEAD, Hydro, SNS).
- **Community:** Users have community cert and pass it to portal, which authenticates and then tries to use TeraGrid on behalf of user. (NVO, HEP, OSG)

2.7 Classifications

So far, three types of Gateways can be identified: Portal based, Applications based and Grid interoperability focused.

From our TeraGrid leader himself (email excerpt) :

"*It's extremely tempting to use "gateway" and "portal" as synonyms but this is an instance where the difference is exposed. Of our Science Gateway partners, most are in fact portals - not all. OSG falls outside of that envelope and the life sciences gateway that Rick Stevens is doing may also.*

Here is how I see the gateways:

***Type 1:** - A gateway that is packaged as a web portal with users in front and TeraGrid services in back. Type 1 gateways include:*

nanoHUB
LEAD (see also type 2)
User Portal
NVO
Instruments (ORNL)
Flood Modeling

***Type 2:** - A gateway that involves a set of bridging activities/services between two Grids, for which we are not yet sure exactly what the best approach is, but is not a portal.*

OSG
LEAD (if we decide also to link their Grid with TeraGrid, but I don't understand their Grid well enough to say more or to make this claim strongly)

***Type 3**: - A gateway that involves application programs running on users' machines (i.e. workstations and desktops), accessing services in TeraGrid (and elsewhere)*

Life Sciences (Stevens)

Summary:

Some I'm not yet sure about. The UNC efforts are likely to involve a portal, but may involve other things. I assume that the CMS portal Roy talked about is type 1 as well. Then there is the homeland security work Pete has initiated with LANL - and I'm not sure what that will focus on. It is quite possible it won't be a portal at all, but rather some codes and data that are stored "inside" TeraGrid and can be fired up at a moment's notice. In that case "gateway" isn't even the best way to describe it."

3 Action requested of working groups

3.1 Account management working group

- Creation of individual accounts for Gateway developers, need policy decision on how many CPU hours to allocate to each Gateway
- Creation of group accounts for the Gateways where needed. This would be for "anonymous" runs through the portal where the developers track individual usage at the portal level. The group account can be associated with the Gateway PI.
    - Can a PI have both his/her own account and be the human associated with a group account, e.g. mapped to two distinct usernames?
    - Is any special security mechanism required (e.g. mapping multiple users to a single certificate) required to support portal-level tracking of the anonymous account?
    - Can anonymous be restricted as to what commands they are allowed to run?

3.2 Security working group

- Examine OSG interactions in greater detail.
- Develop policy regarding anonymous portal users. Fine-grained requirements and methodologies need to be developed.
- Consider dynamic accounts in long-term plans.
- Propose procedures to address how TeraGrid will deal with incidents involving a portal when a contact cannot be reached? (How is this different from PIs who cannot be reached?)
- Work with Gateways who are interested in establishing their own CAs to see if TeraGrid could possibly fill this need
- A set of requirements for SGW auditing and accounting needs to be developed (for example a job on a group account needs to be traceable to an individual portal user). Guidance to the SGW developers would be beneficial.

- Develop this set of ports which portal developers can expect to be available.

- Develop procedures or alternatives to address requests such as TeraGrid hosting of webservers, cgi-bin code running as root, passwordless ssh. One idea is that users make their own web page, but call on a TeraGrid security server.

- Develop general set of security definitions and services that can be offered, based on initial Gateway requests.

- Need clear process for non-TeraGrid certificate acceptance.

- Develop policy for export-controlled codes on TG resources.

3.3 Portals working group

Provide information on group account usage including details on all jobs run at all sites.

- Be aware of: Action item for Security WG: A set of requirements for SGW auditing and accounting (tracing jobs on group accounts to individual portal users) needs to be developed. Guidance to the SGW developers would be beneficial. (not sure that the portals-wg will need to worry about group accounts??)

A longer term recommendation is to study portal technologies other than OGCE/GridPort that have been identified during the survey. A Portal-RAT may be re-created with a new charter to investigate the following portal/gateway framework.

- The Clarens framework is server-side application environment provides an easy to use framework for developing Grid services, with mechanisms for PKI-based authentication, authorization based on access control lists in the context of distributed virtual organizations. Services are automatically published for discovery using a global service registry, and can be used from a variety of clients, including web browsers, scripts and full-scale applications.

- In-VIGO is the framework used by nanoHUB. This NSF NMI project provides a complete framework for a virtualized application hosting environment and portal front-end. It has been already tested with large numbers of users in the nanoHUB environment.

3.4 Software working group

- Compare VDT and CTSS overlap for OSG interaction. (Globus,Condor-G, Chimera, Pegasus, Monalisa)

- Install UML for nanoHUB
- Install Clarens
- Install OGCE/gridport

We recommend that a special OSG RAT be formed in conjunction with the security and gateway working group to investigate in deeper details OSG interoperability. Some TeraGrid members are already part of the OSG interoperability activity and could lead this RAT.

3.5 User Services working group

- Provide list of installed applications
- Improve Community Software Area request and approval process
- Provide instructions and policy information for staging data collections (see also data-wg).
- The SGW RAT or gateways working group will work with the documentation group to create a Science Gateway web area on the TeraGrid web site.

3.6 Scheduling (software-working group / sched-RAT)

A central TeraGrid job management middleware will probably benefit at least 4 of the 7 gateways. A further 2, HEP and NVO, may bring their own resource discovery/management implementation with which any TG job management middleware will have to interact. At least 2 of the 4 (LEAD and Flood Modeling) will also allow their users to specify workflow dependencies. Furthermore, LEAD will require the ability to migrate running jobs from one resource to another. Homeland Security will require job prioritization, and pre-emption/migration of jobs and presumably also co-scheduling of resources. For at least one (Reed) current TG schedulers seem sufficient. The remaining gateways do not address scheduling.

Who owns this scheduling capabilities and will deliver on the milestones that are critical for SGW operations, software-wg?

3.7 Data working group

- Need access to datasets through web services.
- Need Metadata catalog using MyLEAD
- OGSA-Day
- Need instructions and policy information for staging data collections

Few specific survey answers have allowed us to make a good assessment of Datasets needs. However several gateways (LEAD, NVO and NMBR) will definitely benefit from the work of the Datasets collections RAT and we expect new gateways to solely focused on datasets access.

4 Thoughts on Web Services

From a service oriented architecture point of view, we can envision a TG resource running web services. For type 1 and 3 gateways this allows a portal, or a locally-running application (on user's desktop), to invoke a TG-hosted service. In both cases we have similar issues to deal with in terms of certificates/authorization (perhaps the application has an embedded certificate, which we treat as an "anonymous user" (with associated restrictions, etc.). For type 2 gateways, it's a matter of compatibility. It may be the case that a community grid, such as OSG, has a job scheduling service that is able to launch jobs at TG resources, but the launching is coming from an OSG job scheduler service rather than a web portal or a client Hence in this case, web service interoperability is an issue.

Four of the ten initial gateways need back end web services to handle various tasks. These include services that are provided by Globus and TeraGrid itself:

- Resource allocation – a web service that can be used provide information about resources and schedule their use.
- Account information services – what is the status of a user's account. This is likely to be provided by the TeraGrid User portal services.

Other services that may be used by many gateways, but not necessarily hosted by TeraGrid, include

- User-level metadata catalogs that provide searchable indexes of files, applications, experiment logs.
- File management services that are capable of staging files from one location to another or managing file replication.
- Application services – these are web services which are capable of launching a specific application on a TeraGrid resource given configuration input parameters and input file urls.
- Workflow execution services which may manage, on behalf of a user, the execution of several sequenced applications on TeraGrid.
- Logging and Notification services that keep track of the execution of a user's workflows.
- Authorization services that provide information to a portal about the specific levels of authorization a user has to invoke a specific application service on TeraGrid.

Moreover, Web services are increasingly being used by communities to expose both the community data resources and specialized software components. The service is available from filling in a web form, and also programmable, therefore providing an attractive uptake path by the community.

A major objective of TeraGrid should be objective is to make such services scalable -- in the sense that the user is not thwarted by changing interfaces as jobs increase by a factor of 10 and 100 and 10,000, and the big jobs are executed with good parallel speedup.
Small jobs can run in a web-server directly, and medium-sized jobs farmed out to multitasking worker nodes. Large jobs run in a batch queue, and the client has a session ID to monitor the job and retrieve results. Very large jobs can be farmed out to other clusters.

In bulk servicing of web-service requests, the challenges of scale arise from several factors: the number of users, the large size of a single request, and from a large synchronized collection of requests. Each of these will require different job-management strategy. A possible architecture is a combination of "multitasking" nodes along with nodes controlled by a batch queue system.

In many cases these services may be provided as tool-kits to the gateway developers by GIG staff or other gateway developers. Finding a way for gateway developers to share this work would be a significant contribution.

This RAT recommends the creation of a web services RAT to investigate avenues in which TeraGrid resources and services may be accessed through standard web services.

5 A Primer for Science Gateways

We believe that there are prospects for more science gateways in addition to the ones currently deployed or in development. We would like to encourage new gateways. There have already been discussions in this vein concerning BIRN and Geodise. Other possibilities also exist. We suggest the creation of a short primer for Science Gateways as a touchstone document for a process to guide the suggestion, adoption, and implementation of additional science gateways. The purpose of this primer will be to provide a context and starting point for those discussions.

We suggest that this be tasked to the gateways working group, when created and the cognizant Area Director. As a “primer” the target length should be around 15 pages much of the content for such a document can be copied from this report. Below is an initial outline of content that might be included.

1. Introduction

   The TeraGrid provides cyber-infrastructure for science within the NSF community. It consists of a geographically distributed computing infrastructure containing high-performance computing, large scale and high performance storage, advanced software services all connected via a dedicated high speed IP network. The TeraGrid is currently being used as a platform for high performance computing (both leadership class and distributed computing) in the tradition of the NSF supercomputer centers.

However, the TeraGrid is also supporting a use model called a “Science Gateway” whose idea is to provide services that augment this traditional model. Such services include community based portals, application hosting, simplified credential management, simplified job submission, and simplified interfaces. The notion is that Science Gateway access to TeraGrid resources can augment TeraGrid services by adding one or more of the following:

- Subject based focus for specific projects or communities
- Simplified user environment for non-traditional high-performance computing users
- Better integration and partnership with other community, science, or computing efforts external to the TeraGrid

Many TeraGrid Science Gateways are providing these augmented services as part of a portal effort

The hope is that access via a science gateway can make TeraGrid resources available to several thousand science users, many of whom are not traditional high performance users.

2. Science Gateway in Context
    a. Science Gateway (SGW) Definition(s)
    b. Science Gateway user modes (see sec. 2.6 and “types” in introduction)
    c. Distinction between SGW and other TeraGrid user modes
3. Components of a Science Gateway
    a. User Model
    b. Gateway targeted community
    c. Gateway Services
    d. Integration with TeraGrid external resources (data collections, services, …)
    e. Organizational and administrative structure
4. TeraGrid services and policies available for Science Gateways
    a. Portal middleware tools (user portal and other portal tools)
    b. Account Management (user models, community accounts,)
    c. Security environment (security models)
    d. Web Services
    e. Scheduling services (and meta-scheduling)
    f. Community accounts and allocations
    g. Community Software Areas
    h. All traditional TeraGrid services and resources
    i. Ability to propose additional services and how that would interact with TeraGrid operations
5. Responsibilities and Requirements for Science Gateways
    a. Interaction with and compatibility with TeraGrid communities
    b. Control procedures
        i. Community user identification and tracking (map TeraGrid usage to Portal user)

   ii. Use monitoring and reporting
   iii. Security and trust
   iv. Appropriate use
6. How to get started
   a. Existing resources
      i. Publication references
      ii. Web areas with more details
      iii. Online tutorials
      iv. Upcoming presentations and tutorials
   b. Who to contact for initial discussions
   c. How to propose a new Gateway
   d. How to integrate with TeraGrid Gateways efforts.
   e. How to obtain a resource allocation
7. Conclusion:
   a. We have resources
   b. We are interested in additional collaborations
   c. We have a ser of capabilities (resources, services, outreach) that might assist your efforts.
8. Appendices: List of current TeraGrid Science Gateways
9. Appendix: Primer lifecycle: revision log, update schedule, etc.

Appendix- A (Interviews summary)

A-1 LEAD

User community and science: This Gateway addresses the need of the atmospheric science community especially students and researchers interested in modeling and prediction of severe storms. Close to forty scientists are involved in LEAD and on the order of one hundred additional scientists from outside LEAD will be affected by this gateway, students will also interact with the gateway to run small jobs.
Accounts/Accounting: The initial group is about 30 atmospheric scientists and computer scientists that are developing the LEAD technologies. The second group consists of several hundred scientific colleagues (atmospheric science researchers) who will use the LEAD portal because it provides them with access to LEAD data and simulation capabilities. Third group is students.
Portal server fetches proxy certificate and set of capability tokens. User interactions are secure web service requests, service verifies authorization and executes request. Users will have certs local to the portal, map to single TG username. Does not currently keep track of accounting.
Applications: All applications are community codes developed by the weather research community at UCAR/NCAR/Unidata. Twelve to fifteen different applications (WORF,ADAS…)
Middleware: Portal based in the OGCE portals code base. It allows the user to login and manage their experiment information space and to run applications services that run on TG and the LEAD Grid testbed.
User model and security: A user defines couple workflows in the portal, depending on its privileges the user will have access to specific datasets, resources and parameter ranges for the simulation. Advanced users will have the ability to launch workflows on large amounts of resources on demand. The portal fetches the user proxy certificate and a set of capability tokens that allow the user to access different services. Limited set of users will have direct TG cert while students and casual users will have a certificate local to the portal that will give them very limited capability.
Maturity: In demo stage, operational by June/summer 05

A-2 nanoHUB

User community and science: The nanoHUB serves the computational nanotechnology community but also the nanotech community at large, it is seen as a user facility by the National Nanotechnology Initiative (NNI), it is the place for the nanotechnology community faculty, researcher, students, public and industry) to access knowledge in nanoelectronics, nanobio and nano electromechanical systems. Knowledge is seminars, publications, video streams, educational modules and on-line web based simulation.
Accounts/Accounting: 1100 simulation users, expecting to grow that to up to 100 computational nano-scientists, 1k – 5k users interested in educational simulations in next 2 years. Some users will be regular users – own usernames, log in, use community allocation. Others will be anonymous with restricted privileges, certs assigned to group accounts assigned at runtime to real nanoHUB users. Group accounts have restricted

privileges, don't run arbitrary code. Keeps logs for accounting. Existing community allocation to give access to TeraGrid to its power users who may decide to run applications such as NEMO3D on their own. NCN will also use its community allocation to tie the nanoHUB to TeraGrid and use TeraGrid as a compute backend to bring HPC application to the classroom.

Applications: A wide range of applications are being supported, from serial apps to parallel jobs and Matlab scripts. Some applications are developed by the Network for Computational Nanotechnology (NCN) but there are also community codes such as GAMESS, Abinit etc… Some applications have licensing restriction and will not be brokered to TeraGrid.

Middleware: Current middleware is PUNCH (Purdue University Computing Hubs) but it is currently being retired and being replaced by In-VIGO (http://invigo.acis.ufl.edu) that will provide an interface to Condor-G. Middleware effort is supported by NMI integration and deployment grant.

User model and security: Some users will be regular TG users who will get accounts through the community allocation, other users will be anonymous their jobs will run on TG through group accounts with restricted privileges. Portal accounts and TG group accounts will be time multiplexed and provide sandboxing of users for file access.

Maturity: Operational since 2001, http://www.nanoHUB.org, applications are being added regularly, it serves over thousand users yearly.

A-3 Neutron Science

User community and science: This gateway is primarily intended for the neutron science community but also for material science, biology, physics it does cover a wide range of science.The goal is to provide a platform for users to take advantage of compute resources, large datasets and web access to instruments interface. Seventeen intstruments will be coming on-line at the Spallation Neutron Source (SNS), these instruments will be made available to scientists all over the world (1-2000 users). The main challenge is to give access to the data generated by the instrument and make that data available to the scientists as well as provide him with post-processing capability.

Accounts/Accounting: Users shouldn't care whether certs are TG or other, though there might be problems issuing DOE certs. Not currently tracking accounting, but will need to think about it. User credential cross-mapping between ETF credentials and neutron science facility credentials;

Applications: A wide range of applications will be available mostly serial for now. C and fortran programs but also java and Matlab. Some applications run on Windows and need to made available through the portal. Netsolve is also being looked at carefully as a new service that could enhance the user capabilities (parallelization..).

Middleware: It is currently a custom code developed at SNS based on a tomcat/apache infrastructure with the pojo framework. The architecture is designed to be flexible to incorporate legacy applications and commercial packages.

User model and security: It is a difficult problem with multiple facets, SNS users are people who come to run an experiment, they may not be ORNL users and may not be eligible for a DOE certificate. SNS may decide to deploy it's own CA. There is a

possibility to generate a certificate based on a one time password used to access the instruments.
Maturity: In demo stage currently, 40 friendly users. The plan is to be operational by October 2005.

A-4 Flood Modeling

User community and science: This gateway is intended for the hydrology and hydrography community and may extend to the Urban Flood Modeling Community. Primary goal is to model flood based on rainfall data (real time through NEXRAD data or archived data), elevation data (SQL format). Three modes of operation are envisioned all three happening through a portal. First mode is historical flood analysis to compare model and actual event, second is simulation of possible flood and third is emergency response.
Accounts/Accounting: Anonymous user model initially. If takes off might have to change. Will not have their own certificates. All runs under one allocation, application will have certificate (hydrology gateway user). Portal needs to track user level data. Users will have to register. We may not know about them then their e-mail address.
Applications: Several codes run in serial mode: Map2Map handles transformation data, HMS computes metrics during and after a storm. HECRAZ describes flow of water unidirectionaly. As of now the applications involved in this gateway are serial but there is a need for parallelization to increase the resolution and run larger problem sizes, it is envisioned that the problem will scale well on large SMP machines.
Middleware: There is no portal framework as of yet but it will most likely use OGCE and follow TACC technology. Applications will run on lone-star first and maybe Maverick through a discretionary allocations, there might be a need for a community allocation but it will depend on the size and extent of the actual gateway.
User model and security: Users will be anonymous, only their email address will be known. They will register through the portal and run applications through the portal allocation. Users will not have their own certificates. However this model may have to change if the usage increases.
Maturity: The gateway is not already operational and no URL are available, It is unknown when it will be online a first guess is fourth quarter 2005. First priority is to work on the applications, parallelize the hydrology codes and access LIDAR data.

A-5 Homeland Security

User community and science: Gateway to access massive amounts of resources at a moment's notice to simulate emergency scenarios (storms, gaseous clouds etc…)
Accounts/Accounting: Traditional user model, individual users log on, may need mechanism to simultaneously launch jobs across multiple sites. TG certs.
Applications: Based on a code from LANL that simulates avian flu, other applications simulate spread of gaseous clouds.
Middleware: No specific middleware or portal technology are being mentioned, there will be a need for a simplified way to move data across sites. TG wide policy to marshal large amounts of resources simultaneously. There may be a need in the future for a gateway that facilitates the workflow: sites notification, priorities bump, sites monitoring and pushing input decks to code.

User model and security: Standard user model and security, users will have TeraGrid certificates. Users will need ability to launch jobs across multiple sites.
Maturity: Planning stage, this gateway is a proof of concept for on-demand access to large computing resources in order to deal with emergency response. Jobs will launched by hands within 6 months.

A-6 NVO (National Virtual Observatory)

User community and science: The NVO gateway is intended for the astronomy community, its primary goal is to enable astronomical image mosaicking. The client inputs a survey and a sky region, images are fetched from a NVO service, reprojected, backgrounded and coadded. The gateway will also work on building web services for exposure of large astronomical datasets and to allow "grid-enabling" of a popular data analysis environment called IRAF.
Accounts/Accounting: Three classes of users - the NVO research team, science collaborators and students. Study interaction of authentication and access mechanisms with TeraGrid resource policies and potential abuse scenarios. Separate the safe application structures from the unsafe.
Applications: The survey has not identified specific applications for the NVO gateway, the most important aspects of this gateway seems to be the ability to fetched data from the NVO image service. Data could be at FermiLab, NASA Goddard, NRAO, Soccorro.
Middleware: Currently uses the Clarens portal and needs to run apache server on TeraGrid machines.
User model and security: Users may have their own certificates issued from NVO CA or HotGrid CA. Some users may have their own TeraGrid certificate. They submit a set of parameters to specify a job using a web form. The philosophy is that a stronger authentication mechanism allows for more privileges when running jobs.
Maturity: Gateway is not currently in production, URL is in a intermittent state but this gateway does have an accounting mechanism.

A-7 HEP (High Energy Physics)

User community and science: Initially the gateway will target users in the astronomy and particle physics communities, and later a wider scientific clientele. The main concept of the gateway is "Grid Based Analysis Environment (GAE)" which makes the grid computing backends available to analysis environments. TeraGrid resources are sought for event reconstruction and analysis from large datasets. Users are looking for the benefit of large distributed computing infrastructure without the burden of learning unfamiliar software,system and middleware.
Accounts/Accounting: Initial services will include authentication and support for "guest" clients, remote shell functions, file catalogue browsing, file upload/download, and simple task execution. Users would need to get an account and a certificate,.Gateway does not keep track of accounting in a user-specific way, although there are resource monitoring and accounting tools.
Applications: HEP physicists are familiar with ROOT, Java Analysis Studio, CAVES, PhySh etc… and would like the power of distributed computing simply by adding 2 or 3

new commands to the command set that they already know. Software needed are part of VDT.
Middleware: Clarens (http://clarens.sourceforge.net)is the entry point or gateway into the grid for end users who wish to greatly extend the capabilities of their existing analysis tools by using them with grid resources.
User model and security: No group allocations. Individuals will get accounts as part of the large HEP grant (Newman/Litvin). When they present their certificate, Clarens authenticates against the gridmap file. Clarens portal security will need to be analyzed to determine if Clarens portals can be installed on TeraGrid hosts and automatically started on reboot.
Maturity: This gateway middleware infrastructure is already operational on other grids (OSG partners, Grid3 and the LHC computing grid). Still in development on Caltech TeraGrid machines (https://tg-tst-h.caltech.TeraGrid.org:8443/clarens/)

A-8 OSG (Open Science Grid)

User community and science: OSG stands out in the pool of TeraGrid science gateways as being a grid infrastructure in itself. OSG develops principles and architectures using common software to offer services on which Virtual Organizations (VOs) develop their own user portal and applications. Therefore the user community is VO dependent. Supporting OSG grid architecture is a matter of exposing TeraGrid resources to OSG VOs and maybe vice versa.
Accounts/Accounting: DOE Science Grid CAs, honor European LCG certs, D0 and CMF have their own certs. Accounting through monitoring, 3 different monitoring system, not complete, will start accounting project within CMS. No allocation per user, not run out of resources, allocation for a VO then VO will manage allocation for multiple groups within the organization.
Applications: Applications are VO dependent, CMS applications are serial codes. Similar case as the HEP gateway, HEP/CMS is a OSG VO.
Middleware: Clarens portal is used by VO to publish resources and monitor sites status. VDT is the main middleware infrastructure of OSG. TeraGrid will need to work on VDT interoperability in conjunction with OSG interoperability activity.
User model and security: Group accounts are used with a roll based uid mapping at the sites based on the VO, can be modified by account manager.
Maturity: All Grid3 sites should convert to OSG by the first of June 2005, interoperability effort just started.

A-9 Evolutionary Biology

User community and science: The gateway users will be biologists and research in medicine. Three efforts/communities will be leveraged: First, National Evolutionary Synthesis Center (NESC) that is developing federated data models and hosting visitors to develop grid tools, vis, collaboration infrastructure, comparison software. Second, exploratory genetics (data federation and correlation tools for broad access) and finally, Bioinformatics portal (building a bioinformatics portal that will combine access to standard databases and computational analysis tools).

Accounts/Accounting: Three categories of users: biomedical and biology research team, external collaborators, and students. Community allocation, users log on to portal, map to single account on back end. Local CA in early stages. Plans for accounting, nothing operational yet.
Applications: Biological applications mostly serial with multi-threaded capability at times. There is a possibility for parallel version of these applications. As of now they run on 32 bit architecture. Using PISE to generate cgi script and corresponding html for the application interface. An XML description of the application generates the html and the logic for the actual command sent to the globus gatekeeper.
Middleware: Will use OGCE for portal software, globus and myproxy and will gain access to local scheduler such as PBS and LSF. The tools will not be real web services but will be represented as portlets in OGCE, they will use globus gatekeeper and gridftp on the backend. The portal will need to be hosted on TeraGrid machines.
User model and security: Currently has mechanism to give users local certificates, this gateway will most likely benefit from the portal account management effort as this is a very similar setup. Application portlet created by PISE which generates HTML and CGI scripts to interface with the applications.
Maturity: These three gateways are just getting started therefore the gateways are in planning stages. Even though there are three different projects the technology and the staff will be the same.

A-10 Pathogen Data (NMBR)

User community and science: The gateway is intended for the biologist and microbiologists community. Goal is to expand to meet the needs of a variety of bio communities and bio resources – BIRN, NBCR, BISTI. It will provide a set of tools that enable biologists to predict and annotate the function of genes.
Accounts/Accounting: Three categories of users: the NMPDR research team, science collaborators, and students. This gateway as plans for accounting, tracking, logging in portal. They will use a community allocation and will be be happy to use TeraGrid certificates even though it is undetermined right now what they will need precisely.
Applications: Based on http://theseed.uchicago.edu/FIG/index.cgi , the data available through seed is precomputed but there is a need to dynamically generate this data upon request. The applications involved will require database hosting, storing queries and pre-computed data.
Middleware: The middleware used will be homegrown and built from a set of open source middleware components. Job submission will be done via GRAM and there will be a need for web services interface to scheduler and monitoring software.
User model and security: Four types of gateway access will be used: standard web portal model, two more programmatic models (Users write high level scripts, Users write program that compile and run locally), final user model being through web services interaction. Web services being offered by TeraGrid.
Maturity: Gateway is not operational yet but there is a model in the “seed” project. Summer will be critical with a large effort to have a prototype by August and reach production phase in Fall 2005.