Laura F. McGinnis,
Pittsburgh Supercomputing Center
Editor
1-Feb-2006

Grids on Campus: Proceedings from the Workshop

Status of This Memo

This memo provides information to the Grid community. It does not define any standards or technical recommendations. Distribution is unlimited.

Copyright Notice

Abstract

This is the proceedings from the Grids on Campus Workshop, organized by the Production Grid Services Research Group and held at Harvard University in conjunction with GGF15. This document includes the presentations that were accepted by the program committee.

Contents

Abstract 1
1. Forward 2
2. Organizers 2
3. Program Committee 2
4. Speakers and Their Abstracts 2
4.1 Jayanta Sircar, Harvard University, CrimsonGrid 2
4.2 Sridhara Dasu, University of Wisconsin-Madison, GLOW 3
4.3 Glenn Wasson, University of Virgina, UVaCG 3
4.4 Jill Gemmill, University of Alabama-Birmingham, UABGrid 4
4.5 David Wallom, Oxford University, UK e-Science Campus Grids 8
4.6 Arvind Gopu, Indiana University, Hydra 8
4.7 Scott McCaulay, Indiana University, Campus Grids at Indiana 9
4.8 Preston Smith, Purdue University, Campus Grids at Purdue 12
4.9 Joel Snow, Langston University, DOSAR 14
4.10 Valeria Bartsch, Fermi Lab, SAMGrid 14
4.11 Ognjen Prnjat, Greek Research & Technology Network, EGEE 15
4.12 Alan Sill, Texas Tech, TIGRE, THEGrid 15
5. Roundtable Discussions 16
5.1 Users 16
5.2 Applications 17
5.3 Infrastructure 17
5.4 Management 17
6. Security Considerations 17
7. Author/Editor Information 18
8. Intellectual Property Statement 18
9. Full Copyright Notice 18
Appendix: Presentations from the Workshop 19

1. Forward

It is clear that grids are becoming more pervasive as environments for corporate, scientific, and academic work. As grids mature, so must the communities that provide and support them. The Grid on Campus workshop provided an opportunity for members of campus grids to share their experiences and build a support network to improve and enhance their work and the effectiveness of their grid implementations.

Presentations consisted of case studies of grids that are currently providing production- or near-production quality resources to their users. In particular, speakers were asked to address four common areas of grid implementation: Users, Applications, Infrastructure, and Management.

Because the primary goals of this workshop are to build community and identify areas of common interest, papers were not solicited. Instead, selected speakers were asked to prepare a 30 minute presentation about their grid. Slides were collected from the speakers at the workshop and are included as an appendix to this GGF informational document. As more immediate follow-up, a series of roundtable discussions were also held at GGF15, each focusing on one of the key areas. Summaries of these discussions are also included in this document.

2. Organizers

This workshop was Jointly organized by GGF Production Grid Services Research Group (PGS-RG) and Harvard University Division of Engineering and Applied Sciences.

3. Program Committee

- Ian Foster, Argonne National Laboratory and The University of Chicago, GGF External Advisory Committee
- Geoffrey Fox, University of Indiana, AD, Community Affairs
- Laura McGinnis, Pittsburgh Supercomputing Center, Co-chair PGS-RG
- Jayanta Sircar, Harvard University
- Judith Utley, Old Dominion University, Co-chair PGS-RG
- David Wallom, Bristol University, Co-chair PGS-RG

4. Speakers and Their Abstracts

4.1 Jayanta Sircar, Harvard University, CrimsonGrid

Collaborative and interdisciplinary research appears to be the hallmark of current and future directions of leading-edge science research, and such scholarship cuts across departmental and campus administrative domains and policies. However, few campuses today have seamless interoperability and the ability to leverage all available resources--from desktops to other labs, departments, and units on and off campus. In addition, unless research universities quickly begin to leverage the new opportunities in grid architecture, the benefits from large national grid systems may remain out of reach to the individual faculty or student.

The CrimsonGrid Initiative at Harvard, started in April, 2004, aims at building a "test bed" for a next-generation campus "technology infrastructure". A major goal of the Harvard project is to test a template for the transformation of traditional "stovepipe or client/server" based university campuses, to a more seamless technology eco-system, using a switched or "grid" framework. The early use and success of the CrimsonGrid among interdisciplinary and collaborative researchers suggests new possibilities for campus technology environments. An additional key aspect of the Harvard initiative is the development of a grid-technology benchmarking sandbox. The sandbox resources supports a learning and research environment for grid technologies and is available as a resource to faculty, students, and IT staff at the Harvard campus and in the near future, we hope will be available to University sites outside of Harvard.

Collaborators: Joy Sircar, Aaron Culich, Robert Parrott, Lars Kellogg-Stedman, Thaddeus Sze, John Fisher: Information Technology, Harvard Engineering & Applied Sciences; Steve Sakata, Chris McMahon: IBM Universities & Research

4.2 Sridhara Dasu, University of Wisconsin-Madison, GLOW

The Grid Laboratory of Wisconsin

The Grid Laboratory of Wisconsin (GLOW) is a campus grid of the University of Wisconsin-Madison that operates using Condor technologies. GLOW hardware comprises about a thousand Xeon CPUs and about a hundred terabytes of RAID storage. These computing elements are distributed amongst six laboratory sites within the UW-Madison campus, and are connected by 1-10 Gbps campus LAN. The GLOW community includes astro-physicists (Icecube), biologists (genomics), computer scientists, engineers (chemical and materials), medical physicists, and particle physicists (CMS & ATLAS). The GLOW proposal was jointly put together by these groups and was funded by NSF MRI program. The planning of the facilities and operations review are conducted in a technical board. Standardized computing elements are deployed in the domain science group laboratories, but are centrally managed. The local groups retain highest priority for the usage of machines located in their labs. However, any resources unused by them are available to all other GLOW members. Jobs from each of the collaborating groups, and other guests from the world-wide grid flock to all GLOW sites. Domain scientists from all these fields have benefited significantly from this arrangement, by harnessing large amount of resources in short amount of time, opportunistically. The successful use of GLOW for their research has encouraged them to make substantial additions, from other funding sources, to original GLOW infrastructure. In this talk we will describe GLOW, its usage, and its benefits to the UW-Madison research community.

4.3 Glenn Wasson, University of Virgina, UVaCG

As part of an NSF NMI-funded project for the development of WSRF on the Microsoft .NET Platform (WSRF.NET), we have deployed a near-production quality Grid at the University of Virginia based on Globus Toolkit v4 and our own WSRF.NET. Our focus on this deployment effort (and in this talk) is to re-use as much existing campus infrastructure as possible (only deploy *NEW* infrastructure if no existing campus infrastructure already exists).

Because there is not enough time in the talk to get into details of all of the aspects of our campus grid identified in the call for presentations, our presentation will focus on the following topics:

1) User Access: We are convinced that a portal environment is necessary, and further it should be built by the community (we see no utility in creating our own). We selected OGCE/uPortal because UVa ITC is using uPortal in a separate effort (http://myuva.virginia.edu). We are first developing a "Grid-specific" portal; once users are acclimated, we plan to merge the two portals as necessary. How best to provide documentation/user services is still unresolved.

2) Applications: We have built a generic infrastructure that supports all applications. That said, we are particularly focusing right now on matlab and computer architecture simulations. Other UVa campus researchers have expressed an interest in application-specific portals for high-energy physics, chemistry, chemical engineering, and biomedical engineering.

3) Infrastructure:
   - Our software is the NMI stack (particularly GT4), plus our own WSRF.NET.
   - Authentication: UVa Institutional PKI is being utilized; PubCookie-MyProxy integration supports "transparent access" to the Grid, including the TeraGrid
   - Authorization: We wrote an LDAP-based service is used that is compliant with the GGF OGSA SAML Authorization Service.
   - Resources: CS-owned resources, School of Engineering resources, and campus-wide resources. Note: we recently stood up a windows cluster (Windows Compute Cluster Edition, in pre-Beta) and plan to discuss its integration in our Campus Grid via WSRF.NET.
4) Management: This is one of the more challenging aspects that we believe has the most open issues. We are coordinating management with the UVa Central Computing Department (Jim Jokl from ITC), but there remain a lot of tricky policy issues that are addressed now in only an ad hoc basis. We are particularly interested in seeing how others approach this.

In addition to discussing the overall design and implementation of our campus grid, we would like to make a presentation at this workshop so that we can expose the community to our .NET-based Grid software (and its interoperability with Globus) and show that it is stable and used in a near-production environment. We feel the strength of the presentation would be the discussion of novel uses of Windows in the grid environment, plus our novel integration with the campus IT infrastructure.

Collaborators: Marty Humphrey, Department of Computer Science, University of Virginia
Glenn Wasson, Department of Computer Science, University of Virginia
Jim Jokl, Information Technology and Communication (ITC), University of Virginia

4.4 Jill Gemmill, University of Alabama-Birmingham, UABGrid

Uabgrid - A Campus-Wide Distributed Computational Infrastructure

Users:
UABgrid is a collaboration between academic and administrative IT units at the University of Alabama at Birmingham, including office of the Vice President for

Information Technology, Computer and Information Sciences, Mechanical Engineering, Biostatistics, Microbiology, and Structural Biology. Additional computational scientists include mathematicians, physicists, and geologists. These users currently run jobs directly on one or more clusters and also a supercomputer managed by the Alabama Supercomputer Authority. We view UABgrid as an opportunity for our institution to maximize use of its investments in computational resources through shared access, and to minimize the administrative effort required in doing so.

At our institution, High Performance Computing for traditional application areas such as surface modeling and engineering simulations has recently coincided with the rapidly growing field of Bioinformatics. As one of the top research institutions in NIH funding (ranked 20th) and 4th in the SouthEast (behind Johns Hopkins, Duke, and UNC Chapel Hill) UAB's research administration is highly attuned to supporting biomedical research requirements and has recently begun a concerted effort to build up the computational resources available for campus use. We view grid computing as a means of aggregating centrally funded and department-owned computational cycles for the benefit of all participants.

Numbers: approximately 25 computational/engineering scientists; 35 bioinformaticists and a community of over 100 faculty/staff who use the BLAST application. Approximately 100 PhD/grad students are working on some aspect of these applications and up to 150 students/semester use for coursework. This amounts to 2-5% of faculty (depending on whether you include BLAST users) and 1.5% of students. [Note: there is a larger amount of computation occurring in department owned computer labs; it is feasible to consider porting these applications to the campus grid at some future point in time].

Training:
Almost all of these users are trained in their specific application field, and have learned how to use one or another cluster or parallel supercomputer by logging in and submitting a job file. The staff from the central IT Academic Computing department have been working with Globus and condor software through the NMI TestBed program and have become quite familiar with it; these staff people have a combined 30 years of Unix system experience, scientific application programming and support, and experience developing web applications. One computer science faculty member focuses on grid computing as his research area and has taught a graduate level course in grid computing once per year for each of the past two years. Together, the IT and CIS grid people constitute the core of trained personnel, and about 8 graduate students work half time under their supervision.

As we have only recently completed our mastery of the grid components and portal implementation, we have written very little documentation to date. Our goal for this coming academic year is to move into operational mode, which will of necessity include some training and documentation.

User Access:
We have developed a customized portal, based on OGCE. Its front door is located at http://uabgrid.uab.edu/. Details involved in creating, using, and managing certificates are hidden from end users. The key features of UABgridCA are its use of authoritative identity management, its ability to hide certificate management from end-users, and its function as gateway from username and password-based to digital certificate-based

identity. The UABgridCA has been cross-certified with University of Virginia's NMI Testbed Bridge CA, a stand in bridge CA for the planned integration with Higher Education Bridge Certification Authority (HEBCA) [1]. The Testbed Bridge CA can be used for inter-institutional collaboration and resource sharing on the SURAGrid [2]. A detailed description of our implementation can be found in [3, 4]. A key feature is that users access the grid via the campus single sign on, called the 'BlazerID'. The UAB 'BlazerID', named for the university mascot, serves as the user's identity for many network services on campus and is integral to UAB's central, authoritative LDAP-based directory and authentication infrastructure.

Currently, our portal provides access to the one-time registration, web-enabled grid-login, and basic workspace (proxy manager with GRAM job launch fields, GridFTP, and Ganglia view. Most of the students use a command line interface and manage their own private keys, as this is part of their learning process. We are planning for computational students (*e.g.,* math, physics, chemistry) to be able to use this web interface. For applications in common use, like BLAST, we are building an application-specific interface for ease of use; our first such endeavor was GridBlast. The user needs only to browse to their input file, set a few parameters (mostly from drop-down list) and go.

We expect students working on grid technology to use command line interface, and we expect most users of the grid to employ the web interface. We believe a portal or other web interface is a necessity for this class of users.

Applications actually run:

- BLAST (with and without MPI libraries)
- PovRAY
- Genomics matching program from Ga State
- G-BLAST – A GT4.0 based Grid Service for BLAST

This year we will be adding at least one surface modeling application and also a structural biology application. We may also be looking at running a distributed version of a commercial statistics application.

Job scheduling:
Since we are still in early development, current job scheduling is available at the specific resource. The Collaborative Computing Lab (CCL) in the Department of Computer and Information Sciences is involved in developing a resource broker and metascheduler to determine which machine to selecting for submitting a job.

Resource requirements:
Overall, the biggest problem identified has been long queue waits for access to individual clusters (i.e., not enough nodes). The performance of the bioinformatics applications tends to be dependent on the staging of large data repositories involved. The engineering simulations run on clusters tend to be storage bound (in terms of what is available within the cluster).

Infrastructure:
Globus Toolkit Versions 2.X, 3.X and 4.X Condor (NMI release) Pubcookie (will be replaced by Shibboleth this year) MyProxy OGCE PHPki LDAP campus authentication service & directory

Software needed: C, C++, Fortran, Java compilers; BLAS and LAPACK libraries, MPI and Pthread libraries, BLAST Program, R Package.

Resources: heterogeneous (although there is a pair of homogeneous clusters). We have not yet addressed synchronizing software resources; however, we are interested in the UMICH model that introduces a grid "head node" so that the grid stack is installed in front of a cluster rather than on every node. This approach also provides a clean administrative border between grid admin and cluster admin.

Specialized components: we are building a general-purpose, highly configurable collaboration environment where the grid portal interface is just one of several applications sharing a common system environment (including identity and user attributes) by using Shibboleth in a VO Service Provider model. This implementation is called myVOC, developed under NSF grant [] and that has been demonstrated at [].

Account Management:
UAB's central IT organization provides a BlazerID Central service as part of the campus architecture, a single management point for password resets, email account creation or deactivation, and user-authoritative attribute management. There is also a centralized Help Desk called AskIT that has a trouble ticket system. We will leverage AskIT and their trouble ticket system. Therefore, UABgrid does not need to determine who is a member of the UAB community; in addition, attributes such as "STUDENT", currently enrolled course numbers, department affiliation are readily available to us using Shibboleth. It is possible for application-specific attributes or VO-specific role information to be made available, either by adding new attributes to the central LDAP directory or through our VO Service.

Authorization is currently managed by contents of the grid mapfile. Each system administrator can make use of that information to implement their respective usage policies. For the systems available campus-wide, we temporarily map the BlazerID provided in the grid map file to one of 20 accounts that have been pre-established on the system. In the Department of Computer and Information Sciences, individual accounts are created for users based on the BlazerID and the corresponding entries are added to the grid map file manually.

Our overall architecture does application-specific account provisioning and it is worth noting that it would be possible to apply this same approach to automatically provision a permanent account for a particular BlazerID. There are multiple usage policies in place: systems managed by ITAC are available to anyone with a valid BlazerID. Systems managed by the Enabling Technologies Lab give first priority to engineering faculty, then other faculty, then students. Systems available through the Department of Computer and Information Sciences are made available to any UAB faculty and staff whenever the cluster is available and students in the department are given the highest priority. We consider the ability to support multiple policies to be a plus in our campus culture.

Usage Tracking:
We are just beginning to think about this, but we believe it may be central to the success or failure of a campus grid. We will definitely track both users and jobs. We will be accountable to owners of the participating systems; to be determined is whether that information would be public, available to respective owners, or available to all participating sysadmins. There has been a user committee advising the Enabling

Technologies Lab on usage policies; that process will be expanded to include campus-owned clusters and our job will be to implement those policies. Especially for "frequently run jobs" where we would be developing a customized web interface we will monitor the job itself in order to optimize performance. This type of job monitoring will also be necessary at a later phase when we are attempting to run a single job on multiple resources.

Collaborators: Jill Gemmill and Purushotham Bangalore

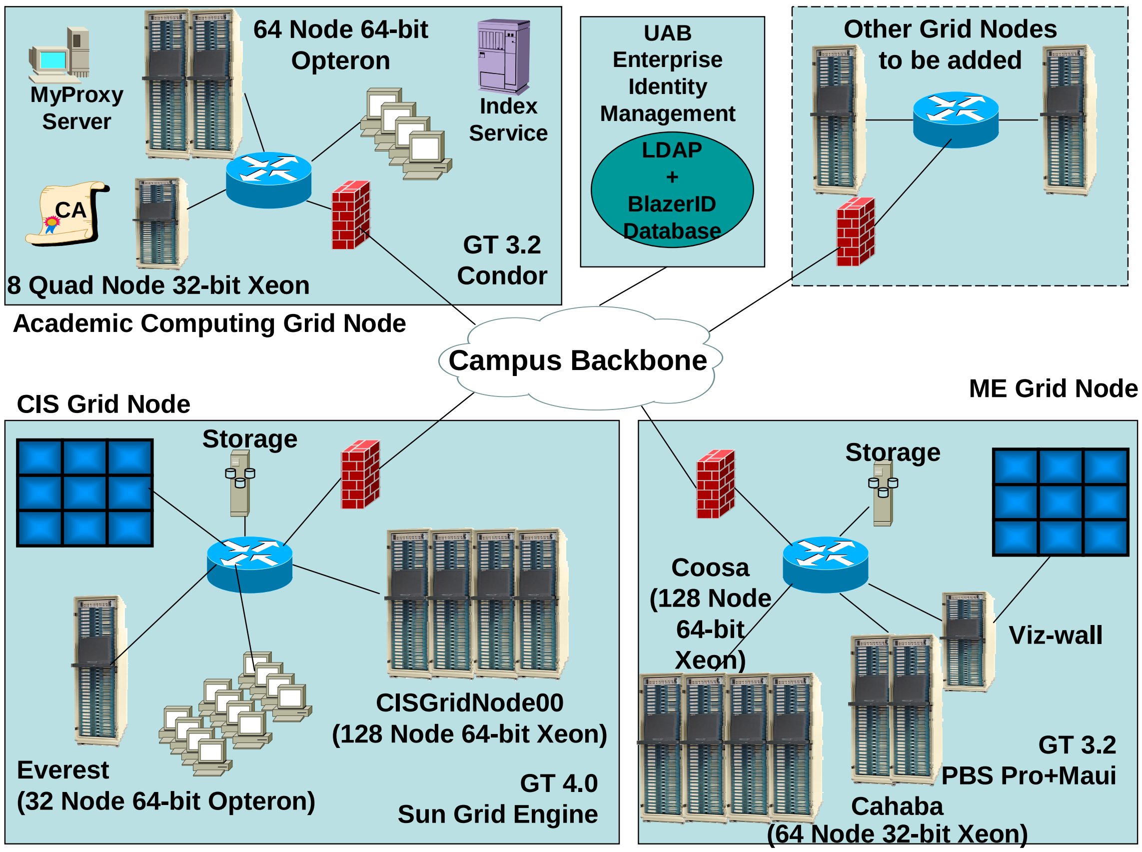

4.5 David Wallom, Oxford University, UK e-Science Campus Grids

4.6 Arvind Gopu, Indiana University, Hydra

Hydra Cluster at Indiana University

The Hydra cluster consists of a Condor pool of computers, running MS Windows spread across the Bloomington campus of Indiana University, that currently lets Bioinformatics researchers to run specific computational Biology software in parallel. The availability of computing cycles is on an opportunistic basis -- Hydra users can run jobs if a machine is untouched (by students in the lab where the machine is located at) for a specified period of time. The cluster is accessible through a web portal that uses Kerberos authentication.

The primary components of the Hydra cluster include:

- A Condor server - To take care of scheduling, match-making, so forth.
- Simple Message Brokering Library (SMBL) - A library used to enable parallel computing on sporadically available desktop systems by introducing a server to keep track of the processing nodes and route messages between them.
- Process and Port Manager (PPM) -- A program that manages resources (processes, ports, so forth) consumed by multiple parallel sessions on the server machine.
- A PHP based web-portal used for authentication and job submission

Future Direction:
Among others, one of the current objectives, that the Hydra project team is working on, is to bring the cluster onto the Teragrid. Efforts are being made to use Globus certificates for authentication apart from Kerberos. The project team is also looking into the possibility of deploying virtual machines/networks in an effort to increase the range of applications that Hydra users can use (given that the worker machines are MS Windows based). A JSR-168 compliant portal is another thing that the project team is working on.

Answers to Questions given in "Suggested Topics for Discussion in the Four Key Areas":

USERS: As mentioned above, currently the user base is restricted to Bioinformatics researchers though that is expected to change in the near future. Researchers in Bioinformatics include experts in the field of Biology and/or Computer Science. Since Hydra is accessible through a web portal -- home grown, requires authentication -- the learning curve is not steep.

APPLICATIONS: Currently the Hydra cluster lets users to run Blast, Meme and FastDNAml in parallel on multiple CPUs at the same time. The project team is looking into various possibilites including deployment of virtual machines/networks to increase the range of applications that can run on the Hydra's grid. Jobs are scheduled by a Condor server on an opportunistic basis -- on a computer that is not being used by a student at a given time.

INFRASTRUCTURE: As mentioned earlier, the Hydra cluser uses Condor, SMBL and PPM along with the software that the end-user runs (currently Blast, Meme and FastDNAml. Account management is done on a per-user basis at this point of time. Users who have accounts authenticate to a Kerberos server. Usage tracking is done via Condor logs. The network connectivity ranges from 10 Mbps to 100 Mbps speeds while the server has Gigabit network connectivity.

MANAGEMENT: Currently, the Condor server is managed by a group of experts from the Unix Systems Support Group (USSG) at Indiana University. The Windows based worker machines are managed by the Student Technology Cluster group at Indiana University. The restriction on software that runs on the Hydra cluster arises mostly out of security concerns – running arbitrary code on Windows machines -- and Windows compatibility issues. Once again, the project team is trying to address this issue (Virtual machines, so forth).

4.7 Scott McCaulay, Indiana University, Campus Grids at Indiana

Campus Grids at Indiana University

Note: Indiana University and Purdue University propose coordinated talks, starting with Purdue University and followed by Indiana University. This will allow each university to discuss their own campus grids, and leave a few minutes at the end of the second talk to discuss statewide grid efforts. A description of IU's campus grid activities follows below, followed by statewide grid efforts.

Indiana University Grid efforts

Population

IU campus grid efforts support local and national users, as follows:

- Computational biology users. We have deployed parallel versions of important bioinformatics applications on a Condor flock that encompasses our student computing labs, and deliver this as a production resource for users of these particular applications. This facility serves a relatively small number of users who use a tremendous amount of computer resources.
- Local researchers. We maintain a grid of high performance computing applications that includes facilities at IU Bloomington and the Indiana University – Purdue University Indianapolis campus. Local researchers from both campuses run across this intercampus grid, providing resilience in case of disaster and also better service for users on each campus. Overall our high performance computer grids serve more than 1,000 local users who use these facilities routinely.
- National research community. Indiana University is part of the TeraGrid and participates in the Open Science Grid. We serve users from the national user community who are not part of IU through these grids and their local instantiations at IU. We have generated accounts on our grid facilities for well over a thousand users.

We provide user education and training through a number of vehicles. A key issue is simply interesting people in using campus grids. We host periodic symposia (www.i-light.org for more information) which focus on researchers explaining how their own research programs have benefited from use of grid computing. In addition, we provide regular training sessions, online help, and a small and very popular booklet entitled "IU's advanced information technology resources: the least every researcher needs to know."

Applications

- Our Condor pool runs three very popular bioinformatics applications: BLAST, fastDNAml, and MEME
- Our campus research grid runs a very large number of applications – literally hundreds. The applications are concentrated in life sciences, chemistry, and physics
- National users of our campus grids tend to focus on chemistry and physics applications

Infrastructure

- Software Stack
    - Our resources in general are heterogeneous, and we use multiple software distribution and packaging tools to manage a very diverse suite of software. Synchronization of software in compliance with multiple grid research efforts (especially the TeraGrid and the Open Science Grid) requires significant human effort, even though we use a variety of tools – in some case specific to particular projects – within a given pool of resources.

- Our Condor pool runs Condor on windows, with an IU-developed parallel communications library called SMBL – Simple Message Brokering Library – which manages a subset of the MPI standard communications between Condor workers in a way that enables MPI applications to tolerate a constantly shifting pool of compute resources.
- Our campus research grid generally includes the TeraGrid CTSS (Common TeraGrid Software Stack), based on Globus.
- Our facilities involved in the Open Science Grid run OSG-compliant installations of Condor on Linux

Account management

- Any member of the university research community is eligible to receive accounts on our systems, and there are no usage limits or allocations. (Jobs are prioritized on many resources via a fair-share scheme). Users of our campus grids include faculty, staff, graduate students, and a small number of undergraduates who have faculty sponsors supervising their research projects. Local users are managed via an IU-created account management system.
- We also support users within the TeraGrid and Open Science Grid communities, with accounts managed by the facilities pertinent to each.

Usage Tracking

- We track usage, and use it for local management and national reporting. We do not reveal any detailed user information in an 'identifiable' way, but we do use usage information as a way to plan our support activities. Users that use extremely large amounts of CPU time are invited to receive consulting help with code optimization. We do report usage on a department-by-department basis as a matter of public record. This helps the academic units of the university understand the value they receive from computing grids on campus.
- We also collect and report usage via project-specific mechanisms for the TeraGrid and the Open Science Grid, which are reported back to the relevant contact points within these projects.
- We track usage on a system-by-system basis, or an an architecture-by-architecture basis, which allows us to manage system heterogeneity (except within the Condor pool, where CPU heterogeneity is simply ignored at present and all we track is CPU hours). We track CPU usage, disk usage, and archival tape usage.

Network Connectivity

- We have local campus networks that are primarily 100 baseT ethernet, and it is via this network that the Condor pool is connected. At each campus (Bloomington and Indianapolis) we have large Force10 Gigabit ethernet switches, connected to each other and national networks. The Force10 switches function essential as a backplane for each campus grid. The Indianapolis and Bloomington campus grids are connected to each other by multiple 10GigE connections, and are connected to the TeraGrid network via a 20 Gbit connection. We have two 1 GigE connections to Abilene/Internet2. As a result we have relatively few network bottlenecks.

Security

- We user kerberos authentication within our research grid generally. For our Condor pool we use ads authentication.

Management

- There are significant technical issues in terms of managing our campus grids, especially as regards security. These simply require a considerable amount of expertise and effort to manage. For example, one of the principles of operation of the TeraGrid is that there be no firewalls between TeraGrid resources and the TeraGrid network. For this reason, we have dual network connections (one to the TeraGrid, one to the campus network) for each system connected to the TeraGrid, and IU firewall policy is enforced between such TeraGrid-connected resources and other campus resources.
- Administrative issues: While our campus grids are significant in size and widespread, most of the resources are controlled by the central campus IT organization (University Information Technology Services - UITS). We have collaborative relationships with other resource providers and participants in the grid which enable effective participation of such resources in our campus grids. In some cases, such as the widely used Flybase, the IU Department of Biology runs servers connected to Abilene and the general national and international research community, while UITS operates a mirror server which is directly connected to the TeraGrid.
- Political issues. We have had relatively few 'political' issues to deal with, in large part thanks to a perception that UITS, which operates the majority of the campus grid efforts, is diligent in trying (and succeeding) to meet the needs of local researchers.

Other issues

- Indiana University operates a production local grid and participates in two national grid projects. IU's unusual status of having two large research campuses, connected by university-owned optical fiber, provides very interesting opportunities for grids that span multiple campuses. For example, we have a storage grid that includes distributed data movers and mirrored data silos on IU's two core campuses, providing disaster resilience and reliable protection of data in case of natural disaster.
- Indiana University and Purdue University operate a statewide network called the Ilight network. This network, jointly owned by Purdue and Indiana universities, connects the Purdue campus in West Lafayette, the Indiana University Purdue University joint campus in Indianapolis, and the Indiana University campus in Bloomington. This statewide network provides extremely important intercampus grid possibilities, as well as disaster resilience capabilities for both universities.
- Indiana University (alone and in collaboration with Purdue University) are implementing production and experimental campus grids that are delivering important research services today, as well as providing research and development platforms for the grids of tomorrow.

4.8 Preston Smith, Purdue University, Campus Grids at Purdue

Purdue Service Oriented Campus Grid Architecture

Introduction: In this presentation, we will discuss in-depth the infrastructure deployed at Purdue University to support three types of user communities with varying requirements - namely local users from the Purdue University community, grid users both from the TeraGrid and the Open Science Grid communities, and thirdly the science gateways such as the nanoHUB and the Purdue Terrestrial Observatory (PTO). As outlined in the first ECAR study, research and education at Purdue University is supported by Information Technology at Purdue (ITaP) - the central IT organization. In particular, the cyberinfrastructure required for all discovery activities on-campus is supported through the Rosen Center for Advanced Computing (RCAC). However, in addition to providing computing services to campus users, RCAC also functions in the role of resource provider (RP) to the nation through TeraGrid (TG) and to the world through its involvement in Open Science Grid (OSG) as a Tier-2 site serving the Compact Muon Solenoid (CMS) experiment being conducted at CERN. We will show the system administration challenges involved in supporting all these users and projects and outline the plan that RCAC has to deploy a service oriented on demand architecture to build community grids.

Computational Services: The computing infrastructure is focused on the development and operation of community clusters. This notion is built on the foundation that individual investigators on local campuses can secure sufficient funding for the purchase of their own computational resources. By centralizing and coordinating the purchase of individual clusters, each university can build a larger compute resource. The incentive to individual researchers to contribute towards community clusters is the potential access to additional compute power than what was bought by an individual PI. It promotes the use of opportunistic cycles through Condor and work on new allocation process for preemptable jobs as well as on-demand computing. External grid projects link to these clusters through Globus gatekeepers and some appropriate job managers. Data oriented services: One important aspect of the central infrastructure maintained by RCAC is that data oriented projects are systematically being tied to the Storage Resource Broker (SRB) for easy data access and management through the TG network. Furthermore, RCAC is also developing data portals unifying SRB portlets, GIS clients such as ArcGIS and climate tools such as IDV from Unidata. Metadata is extracted or defined as closely to existing standards as possible and published into an SRB MCAT server. The datasets are served from a central spinning storage system and backed-up on tape storage. PTO which includes data from Laboratory for Applications of Remote Sensing (LARS), National Weather Service, and climate modeling data from the CCSM simulator is connected to the TG and data are accessible through the PTO portal. Both, static and streaming datasets are made available while access to relational datasets is also being enabled.

Community based services: The nanoHUB, which is the service delivery vehicle of the NSF Network for Computational Nanotechnology (NCN) is a key example of the Purdue campus grid capability. The nanoHUB relies on middleware technology such as virtual machines, virtual networks, Condor, and Globus. It builds on top of the community cluster setup to allow for on demand provisioning of compute resources including virtual cluster. We will show how through the use of virtual machine we can provision community specific grids on demand on top of a common physical infrastructure. Account management is left to the community and accounting responsibilities is partially delegated. By doing so the resource provider only focuses on dealing with the Virtual Organization (VO) and by virtualizing the resource it isolates the execution environment and can dynamically allocate resources as VO need them. To reach such a goal RCAC has started to develop a framework called Narwhal, that is a system whose goal is to provide network adapted resources with heterogeneous access layers---specifically to allow for graphical and text-oriented applications to be accessed through either a web portal or WSDL/SOAP interface. Our current implementation exists as a Mambo/PHP module that implements a session protocol that supports interactive applications using VNC. These applications run on a pool of Xen Virtual Machines (VM) in order to insulate the runtime environment from the physical domain. The session protocol allows the system to do intelligent load balancing and delegation of jobs to VMs while closely monitoring and recording statistics about resource consumption. An administrative interface allows an administrator to rapidly deploy new applications and computational resources. A similar interface also allows users to dynamically share interactive sessions with other users.

Educational services: One unique aspect of the work on campus grids that RCAC is pioneering relates to its use for educational purposes. In order to have maximum impact on learning, it is critical for cyberinfrastructure scientists to lower the barrier of entry into grid environments such as the TeraGrid. The vGrid – or Virtual Grid - project is focused on introducing grid environments to learners without over-exposing them to the complexities associated with grid use. It must, however, be pointed out that the same infrastructure is highly suitable for researchers in various disciplines to focus on science with minimal startup times and learning curves. There are two parts to the vGrid environment: the backend that provides computational resources and a portable CD that allows the access to this backend. The CD is basically a Knoppix CD with the installation of the Globus and Condor/Condor-G software. The vGrid backend consists of a Globus gatekeeper, a condor pool, an EJBCA (Enterprise Java Beans Certificate Authority), and a DNS (Domain Name Service) server. All these functionalities are served by the Xen virtual machines for security reasons. The virtual machines, except for the DNS server, are on a non-public subnet. Upon booting the CD, researchers are ready to learn essential parts of grid computing, such as getting a certificate, running basic Globus commands, and submitting jobs via Condor to the vGrid.

Conclusion: We believe that university campuses are leading CyberInfrastructure deployment across the nation and we will show how Purdue University's Rosen Center for Advanced Computing is addressing the issue by defining a model for CI and service oriented architecture that tailors to individual users as well as communities both local and national.

Collaborators: S. Goasguen, R. Kennell, W. Lin, S. Clark, T. Stef-Praun, C. Baumbauer, L. Zhao, T. Park, P. Smith, K. Madhavan, Rosen Center for Advanced Computing, West Lafayette, In 47906

4.9 Joel Snow, Langston University, DOSAR

We have a campus-based community grid organization, DOSAR (a Distributed Organization of Scientific and Academic Research), which is registered in OSG and has several institutions working on D0, a Fermilab Tevatron experiment, and both the LHC experiments. Several institutions in the organization play important roles in LHC experiments. For example. Univ. of Texas at Arlington and Univ. of Oklahoma are a joint ATLAS Tier 2 site and have been working with Grid 3 and OSG through the past several years. In addition, Univ. of Sao Paolo in Brazil has been working closely with CMS. Thus, I would like to request an opportunity to present our activities in DOSAR at this workshop. In addition to the overall talk on DOSAR, I think we might also want to present one or two specific campus activities at the meeting if your agenda slots allow.

4.10 Valeria Bartsch, Fermi Lab, SAMGrid

Deployment and Operation of Samgrid on Many Sites Including Our Local Campus Grid Fermigrid

SAMGrid is a GRID framework designed to be performant for experiments with large (petabyte-sized) datasets and widely distributed production and analysis facilities. Its focus is on high energy physics experiments and it is deployed at the CDF, DZero and Minos experiment. Therefore the user community addressed by this GRID system is quite homogenous though for every experiment special adaptations were made. The GRID system can be roughly divided into two sections, the data handling system SAM and the GRID extension SAMGrid. SAM provides a set of services for data transfer, data storage and process bookkeeping on distributed systems. SAMGrid enhances SAM incorporating standard Grid tools and protocols and developing a Grid/Fabric interface for DZero and CDF.The data handling system SAM has been successfully deployed for DZero and for CDF. SAMGrid has been deployed and used for DZero. For the future of SAMGrid it is vitale to interface to other GRID systems. There is a program on the way to integrate SAMGrid with otherGRID systems. A test bed for LCG is developed.

FermiGrid serves as a campus grid at Fermilab providing common grid services including a site wide globus gateway, a virtual organization membership service (VOMS), a grid user management system (GUMS) and a site authorization service (SAZ). It is a meta-facility composed of a number of existing resources, many of which dedicated to the use of a particular stakeholder. FermiGrid provides a open science grid (OSG) interface to enable opportunistic use of Fermilab compute elements and therefore optimizes the use of resources at Fermilab. DZero is one of the VOs participating in FermiGrid as well as using SAMGrid for job/data handling and environment preparation. Therefore SAMGrid will be interfacing to FermiGrid.

4.11 Ognjen Prnjat, Greek Research & Technology Network, EGEE

EGEE – European Grid Computing

The presentation will focus on the EGEE Grid. Since the South-east Europe Federation which I am managing is one of the 12 federations within EGEE, there are no particularities of this federation Grid which differ much from EGEE. So, I will focus on EGEE Grid covering:

- Infrastructure overview
- Middleware brief overview
- Focus on Grid operations:
    - Management hierarchy
    - M/w upgrades
    - Monitoring
    - Operational an user helpdesk and support
    - security brief overview
- Applications and VOs, Grid usage
- Some guidelines for green-field regions such as SEE how to join larger production Grid
- some political issues in this context

4.12 Alan Sill, Texas Tech, TIGRE, THEGrid

I will be attending GGF15 representing our university's involvement in the TIGRE and THEGrid projects (http://www.hipcat.net) and would be willing to speak about TTU's campus grid and involvement in other local, regional, national and international grid projects. I am the project director for THEGrid, the Texas High Energy Grid being set up to serve the needs of the nuclear physics, astrophysics, astronomy and high energy particle physics communities within Texas, and Senior Scientist at TTU for the TIGRE project (Texas Internet Grid for Research and Education), which is a production grid being set up to serve the more general education and research community.

I have been involved in building successful production grid and large-scale distributed computing projects for the Open Science Grid and for the CDF experiment (Collider Detector at Fermilab), the latter of which produced a successful collaborative computing environment of approximately 5000 GHz PIII-equivalent processing power spanning 12 sites in 6 different countries. I have also served internationally as a consultant to the national Center for High-performance Computing (NCHC) in Taiwan, and to other institutions.

Within Texas, the TIGRE project aims to serve the needs of 5 major university sites spanning three different university systems, and will build on the successful experience we have had at TTU in running a campus Avaki-based grid for the past 3 years, extending this to Globus-based methods in collaboration with our 4 other university collaborating institutions (UT Austin, University of Houston, Texas A&M, and Rice). The TIGRE project is a production grid, not an R&D project, funded out of the Texas Enterprise fund, which is normally reserved for projects with high potential for job creation. We will also extend TIGRE to serve the needs of other institutions within the state on a voluntary, self-supported basis.

I would be willing to speak about some or all of the above, including our operational experience in installing and running these production operating grids and our plans for the future if asked by the GGF workshop organizers.

5. Roundtable Discussions

5.1 Users

- Strategies for Recruiting Users
  - Open, easy access for students, esp. graduate students
  - Identify successful users and enlist their help to recruit colleagues
- Training
- Portals
  - Critical if user is already working in a browser/GUI environment
  - Need to be based on how user is already working
  - Should not be imposed, esp. unnecessarily
- Communities Are Important
  - Provide environment for networking and mutual support
  - Build critical mass for standards and policies
  - Can aid in recruiting more participants
- Barriers to Successful User Experiences
  - Libraries and common packages that aren't "gridified"
    Recommendation: Pull technical students from scientific disciplines to help work on this

  - Account management and access are (perceived to be) complicated
  - Concerns about the quality and reliability of the services and systems provided

5.2 Applications

- Issues to Consider:
  - MPI versus Workflow
  - Intensely versus trivially parallel
  - Low-hanging versus high visibility
- Campus communities can support each other if they can find each other
  Recommendation: Set up website to serve as point of contact, source for collaborative tools, etc.
- Licensing: From theory to practice
  - Need more cooperation from vendors
  - Current licensing discussions in GGF may not be including sufficient input on academic and research licensing
  - Too much attention to enterprise/commercial grid licensing?
- Base Systems
  - Are there applications of value that can run easily on multiple platforms (e.g. WinX, OSx, Linux)?
  - Are there configuration management, virtual machines that provide environments that users and applications need without major cost incursions?

5.3 Infrastructure

- Potential Problem areas:
  - AAA (Authentication/Authorization/Accounting) boundaries are complex and possibly vague.
  - Operations Center/Help Desk
  - Systems Operations, esp. human factors
  - Users and applications are more complex in grid environments
    Recommendation: PGS-RG should look at interoperability issues
- *Question*: Should "standards" be proactive or reactive?

5.4 Management

- Management decisions should be informed by the people/stakeholders involved:
  - Users
  - Funding agencies/organizations
  - Technical supporters
- Business models, best practices for grid technology
  - "Best" varies by application and population using a given grid
  - "Best" or "common"? Which would be useful?
- Grids need to fit into campus strategies

6. Security Considerations

Security for individual campus grid installations may be addressed in the presentations, but is not an explicit requirement for this document or its contents.

7. Author/Editor Information

Contact information for authors: Relevant contact information for the presenters is available in the appendix.

Editor Contact Information:
Laura F. McGinnis
Pittsburgh Supercomputing Center
4400 Fifth Avenue
Pittsburgh, PA 15213
USA
lfm@psc.edu

8. Intellectual Property Statement

9. Full Copyright Notice

Appendix: Presentations from the Workshop