Purdue Service Oriented Campus Grid Architecture

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

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 pre-emptable 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.