Grid Architecture for Scientitic Communities

Sebastien Goasguen
Purdue University

Abstract. Sharing of resources among virtual organizations (VO) has been termed the grid problem. At the heart of this problem is development and acceptance of a protocol to share, discover and compose services. However significant challenges arise in monitoring, accounting and securing any grid infrastucture. VOs of any size should be able to build their own cyberinfrastructure (CI) by discovering and then composing services to build higher level capabilities that they need. These community's CI could be dynamically expandable, persistent and migratable while using resources across adminstrative domains. In this paper we will present several VO based grid architecture and focuse on the nanotechnology community CI called the nanoHUB. The nanoHUB uses virtualization technologies to isolate the infrastructure from the local administrative domain. Virtual machine migration and virtual networking techniques allow the infrastructure to be dynamic and adapt to the underlying physical system. Moreover by using virtual machine the users are provided with a sand box that serve as a development platform for their software an as a entry point to physical grids. Finally, we will cover the security issues in VO access to multiple grid infrastructure and present our work towards using shibboleth witin the nanoHUB.

Please use the following format when citing this chapter:

Goasguen, S., 2007, in IFIP International Federation for Information Processing, Volume 239, Grid-Based Problem Solving Environments, eds. Gaffney, P. W., Pool, J.C.T., (Boston: Springer), pp. 397.

Q&A – Sebastien Goasguen

Questioner: Bill Applebe

Is the nanoHUB installed elsewhere, or is there a plan to distribute it more widely or make it compatible with other grid hubs?

Sebastien Goasguen

Currently the nanoHUB infrastructure is centralized and operated from Purdue. Remote resources on TeraGrid and Open Science Grid are used but the virtual infrastructure is based at Purdue. We have started work towards a much more distributed infrastructure using some virtual networking technologies. Resources at the University of Florida are being used to test migration of virtual machines between Purdue and UF. We also plan to package the middleware to make it available to other communities.

Questioner: Mladen Vouk

What are the failover (fault-tolerance) provisions in your system?

VMWare seems to have nicer failover properties (automatic migration to backup servers). Is there are reason you chose to use XEN?

NC State has a system that is in many ways similar to what you have. Ours is called Virtual Computing Laboratory and uses RDP as the remote desktop access protocol (and within web window display). Is there a reason you chose VNC? In my experience VNC sometimes has difficulties with graphics.

Sebastien Goasguen

The In-VIGO middleware has been deployed at Purdue and a reference implementation is used in Production. This implementation has been enhanced with monitoring tools such that user sessions, file system, connection to back-ends, VNC servers are being checked regularly and restarted if they fail. This system also allows us to stress test the middleware and simulate high user load.

We currently do not have automatic migration in our infrastructure. Migration is being used in our virtual cluster to adapt to the load and run time characteristics of the applications. When the project started we purchased two licenses of the Vmware GSX server. We ended up transitioning to the open source Xen system which is a patch to the Linux kernel. It offers less overhead and through experience has shown to be more reliable. Now that Vmware has released a free version of their player and their server we are starting to re-introduce some Vmware technologies, especially the player which allows us to give a "nanohub appliance" to our users.

So far VNC has seemed able to meet our needs. We are aware that there maybe some networking issues when applications have heavy graphics.

Questioner: Ron Boisvert

There are many simulation applications available to users of the nanoHUB. Presumably, these are contributed by the research community. On average, how much work is required to take one of these apps and make it presentable/reliable for use in the nanoHUB, and who does this work?

Sebastien Goasguen

Anyone can contribute an application to the nanoHUB. A project is created in our subversion repository and the contributor can upload his or her applications. The Network for Computational Nanotechnology has developed its own graphical user interface builder called Rappture, which can be used easily to create a nice GUI. An undergraduate can create the interface in a couple of hours. More complex applications can take weeks, but mostly due to a review process that has several iterations before an application is made publicly available. However the strength of the middleware is that any application can be made available, so technically it takes as much time as to get the source and compile the code.

Questioner: Xiaoge Wang

Will and how does this middleware affect the way that scientists work?

Sebastien Goasguen

Users of the applications can access them through their browser without installing or compiling any software components. More advanced users who develop applications can benefit from a standard development environment available through the workspaces. If developers wish to adopt Rappture to create their applications interface they will need to learn the Rappture API that they will use in their code. If we consider the access to remote resources outside the nanoHUB, the scientists need only know how to use condor and pbs. They actually don't have to learn too many new things from their standard mode of operation.