FutureGrid

Quarterly Report – Q1 2009 – 1 October to 31 December 2009

Geoffrey Fox, Indiana University (PI)
Kate Keahey, University of Chicago (co-PI)
Jose Fortes, University of Florida (co-PI)
Andrew Grimshaw, University of Virginia (co-PI)
Warren Smith, University of Texas (co-PI)

Introduction

This quarterly report provides a summarization of the following bi-weekly reports submitted to NSF each Monday:

October 11, 2009
October 26, 2009
November 9. 2009
November 23, 2009
December 7, 2009
December 21, 2009

FutureGrid started on October 1 2009 and had its first All Hands Meeting October 2-3 2009 in Indianapolis. All partners were represented and key organizational issues were discussed. There is a project web site http://www.futuregrid.org/ built using Drupal allowing authorized participants to update information.

There are seven (7) committees around which will project activities will be organized. These are:

- System Administration & Network Management Committee (lead David Hancock IU)
- Software Committee (lead Gregor von Laszewski IU)
- User Requirements Committee (lead Andrew Grimshaw UVa)
- Performance Committee (lead Shava Smallen SDSC/UCSD)
- Training, Education and Outreach Services Committee (lead Renato Figueiredo UF)
- User Support Committee (lead Jonathan Bolte IU)
- Operations and Change Management Committee (including Change Control Board (lead Craig Stewart and Gary Miksik IU)

Science Outcomes

No science outcomes during this quarter, as the project's priorities are the ordering, receiving, and installation of all hardware and networking, and the creation of all requisite committees and workgroups

Systems Administration & Network Management Committee

Compute and Storage Systems

- IBM clusters have been ordered and we anticipate system delivery the week of January 4th for IU, UC/ANL, and UF

- The Cray XT5m for IU has been ordered and we anticipate system delivery the week of January 11, 2010. Delivery was intentionally delayed from December 15th to after January 7th so that power upgrades that are required to support the Cray could be completed at the IU data center.

- The Dell system for TACC is still in the ordering process. Delivery is still on schedule for January 2010.

- The current test IBM iDataplex at IU is being prepared for delivery to UCSD/SDSC in parallel to benchmarking work, along with xCAT and Moab integration. We anticipate sending the system to UCSD/SDSC on February 1, 2010.

- The DDN 6200 storage system for UC/ANL has been ordered and delivery is scheduled for December 18, 2009.

- The Sun X4540 storage system has been delivered at UCSD/SDSC and acceptance testing should begin on December 21, 2009.

Networks

- Juniper EX 8208 for the core router has arrived and IU and is being configured before installation in Chicago.

- Spirent Networking Impairment Device has arrived at IU and is being configured before installation in Chicago.

- The co-location contract for networking central networking equipment in Chicago, at Starlight, has been negotiated. Additional power has been ordered, and will be available on January 1, 2010 to host the central networking hardware.

- The Central networking contract with NLR has been negotiated and a purchase order is in progress.

- Site networking order for UCSD/SDSC, UC/ANL, and UFL are in progress and awaiting NLR local termination locations.

- We are working with Linda Winkler at UC/ANL in her role as Networking Working Group chair to peer FutureGrid networking with the TeraGrid and provide shared access to TACC

Software Committee

Summary of Recent Events

We have refined the software architecture plan, including the collection of more use cases. Several meetings with individual group leaders including our experts from system administration and scheduling, operations, performance, portals, and social tools were held to improve upon the collection of the requirements driving the software development and priorities. Additionally, a test license for Moab was obtained, the scheduling system, which is a major component of the FG software architecture. A 3 phase configuration plan for the use of Moab to support dynamic provisioning has been developed and is refined as part of the software engineering process

Accounts for FG on Grid5000 were created while using Pegasus to deploy hardware images on Grid5000 and identify how to interact with xCAT-like tools, and schedulers like Torque, Moab, etc.

Architecture Definition

The following architecture diagram integrates the many subcomponents to be developed and needed as part of FutureGrid:

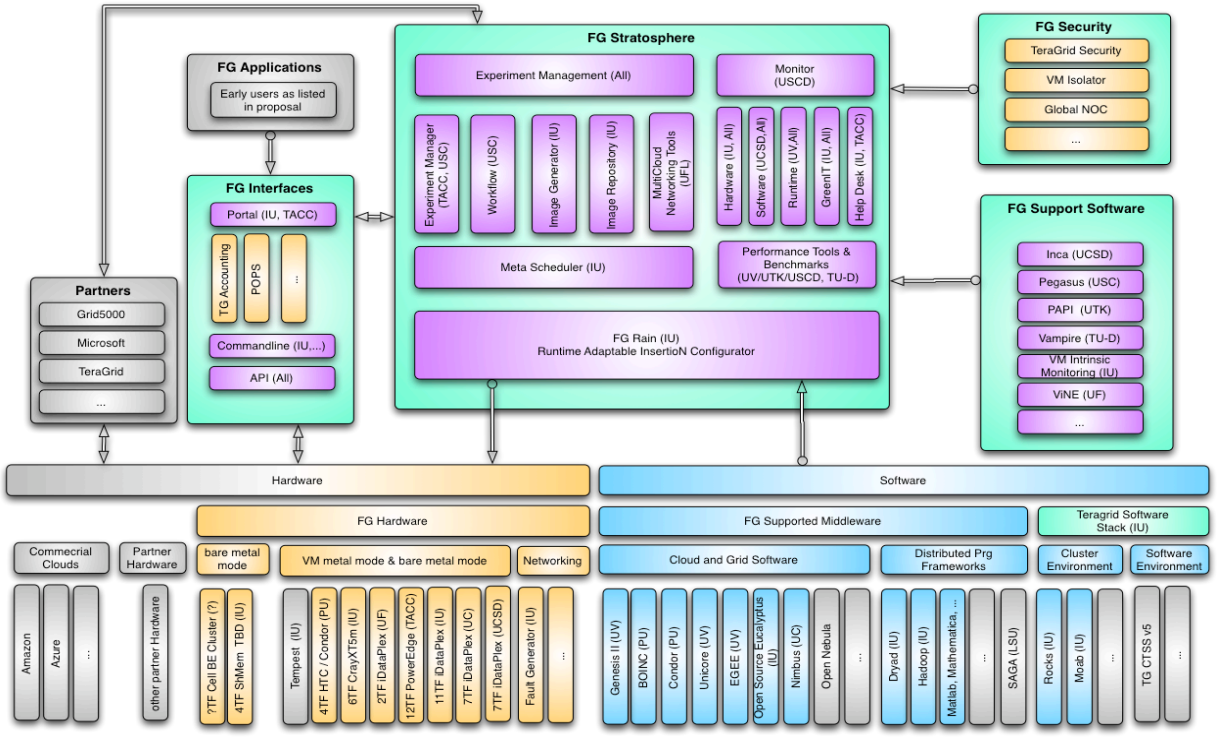

Software Task Management

The management of such a large activity requires the establishment of an infrastructure that is suitable to coordinate the many tasks and groups. The software committee has evaluated a number of tools and JIRA has been chosen to facilitate this task. The system is in the process of being deployed.

Requirements and Use Cases

Use cases are being gathered that address many of the components and software needs motivating the architectural diagram. These use cases are motivators and provide guidance for the implementation of the software needed. Several meetings with individual group leaders including our experts from system administration and scheduling, operations, performance, portals, and social tools were held to improve upon the collection of the requirements driving the software development and priorities. Each committee has been asked to produce a document on what use cases motivate software developments and deployment a part of the FG architecture while addressing their particular goals. All of this data will be included in the creation of a FutureGrid software architecture plan,

Subgroups

We have started the creation of subgroups to work towards the implementation of software and features of FG. These include at present a portal group, a scheduler group, a software operations group, a performance group, a knowledge base group. Other activities are enhanced through individual contacts to group members and other committee chairs.

Simulated FG Environment (based on VMs)

We have started to evaluate the creation of a simulated FG environment so we can set up a scheduler testbed before we actually have hardware available. Based on virtual machines, this will also allow for the design and development of Pegasus workflows to simulate some use cases

Experimental ("Test") Harness

We started the development of the experimental test harness using TACC's Lonestar system to assemble the base software stack (until FG hardware arrives).

Moab Scheduling Software

A test license for Moab was obtained and a 3-phase configuration plan for its use in supporting dynamic provisioning was developed.

User Requirements Committee and User Advisory Board

Summary

The size of the User Advisory Board was determined to be six to nine members. Several names have been accumulated as candidates, one of which is Shantenu Jha. The balance between middle-ware and cloud researchers/developers, end-user scientists and infrastructure leaders needs to be decided.

The requirements for a typical middleware developer using FutureGrid have been drafted.

A new early adopter Registry was added to the FutureGrid website to solicit information from potential early adopters on the nature of their use of FutureGrid.

Performance Analysis Committee

Summary

The goals of the Performance Analysis committee are to 1) evaluate the FutureGrid infrastructure as it is developed and deployed and 2) help users to leverage the FutureGrid infrastructure for their experiments.

UC San Diego is leading the effort and there are four other paid participants from IU, USC/ISI, UTK (starting early next year), and TU-D. Since October, the group has been meeting weekly to design an architecture document to accomplish the above goals. There are four main components:

1) A set of benchmarks that can be used to establish and analyze the baseline performance of FutureGrid and a framework that can automate periodic execution in order to detect performance problems

2) Instrumentation of the FutureGrid software to collect performance measurements in order to validate its usability and detect performance problems

3) Performance tools that will be deployed within the FutureGrid environment to help users analyze the performance of their software and/or applications

4) A subcomponent of the FutureGrid Portal for users and staff to view and analyze performance tools and data. These components will leverage partner software such as Inca, Pegasus, Vampir, and PAPI.

The group also participated in SC'09 poster and talk activities and is exploring two potential collaborations with the DOE Magellan project from LBL and ANL and the IPM team at NERSC. Work is also underway to get a server setup for an initial Inca monitoring deployment and other early performance activities.

Training, Education, and Outreach Committee

Summary

The TEOS committee has been iterating on the design of virtual appliances and tutorials for FutureGrid, and has created a presence on the FutureGrid Web site to disseminate tutorials and provide information to attract early adopters. The tutorials to date provide instructions on how to create virtual clusters with Condor, MPI, and Hadoop. Higher-level tutorials can be built on top of these virtual infrastructures. One example is a tutorial on MPI by an early adopter (Amy Apon, U. Arkansas) that uses a virtual appliance infrastructure. Currently the tutorials are based on the Grid appliance system from U. Florida and rely on user-hosted or cloud-provided resources to create virtual clusters. Given the use of appliances and standard virtualization technologies, these appliances will also seamlessly deploy on FutureGrid resources once available.

The committee is refining the definition of TEOS-centric requirements for the FutureGrid Web site and software infrastructure, including the ability to define and manage user groups for private virtual clusters. This capability will enable dynamic creation of virtual clusters for education/training activities on FutureGrid by instructors.

Summary of EOT Events for Q1

Type of Event	Title	Location	Presenter	Type of Audience	Month
Workshops and Tutorials	Introduction to the Grid Appliance	FutureGrid web site	University of Florida	On-line tutorial	Nov-Dec 2009
	Creating Grid Appliance clusters	FutureGrid web site	University of Florida	On-line tutorial	Nov-Dec 2009
	Building a Ubuntu-based Grid Appliance on cloud or local resources	FutureGrid web site	University of Florida	On-line tutorial	Nov-Dec 2009
	Deploying Grid Appliances using Nimbus	FutureGrid web site	University of Florida	On-line tutorial	Nov-Dec 2009
	Virtual MPI clusters with the Grid Appliance and MPICH2	FutureGrid web site	University of Florida	On-line tutorial	Nov-Dec 2009
	Introduction to Hadoop using the Grid Appliance	FutureGrid web site	University of Florida	On-line tutorial	Nov-Dec 2009
Presentations and Talks	FutureGrid	Open Grid Forum 27, Banff, Canada	Andrew Grimshaw, University of Virginia	Advanced technical	Oct 2009
	FutureGrid Overview	CCA-09 Cloud Computing and its Applications Workshop, Chicago, IL	Geoffrey Fox, Indiana University	Technical	Oct-2009
	FutureGrid and Green Aware Computing	Keynote at 3rd International Conference on Networks & System Security, Gold Coast, Australia	Gregor von Laszewski, Indiana University	Technical	Oct 2009
	FutureGrid and Green Aware Computing	Northern Illinois University, DeKalb, IL	Gregor von Laszewski, Indiana University	Technical	Oct 2009
	FutureGrid: An Experimental High-Performance Grid Testbed	NCSA, Champaign-Urbana, IL	Craig Stewart, Indiana University	Technical	Oct 2009
	FutureGrid Overview	SC09 Conference, Indiana University booth, Portland, OR	Geoffrey Fox, Indiana University	Mixed technical	Nov 2009
	FutureGrid Overview	SC09 Conference, AIST booth, Portland, OR	Geoffrey Fox, Indiana University	Mixed technical	Nov 2009
	FutureGrid Overview	SC09 Conference, Indiana University booth, Portland, OR	Geoffrey Fox, Gregor von Laszewski, Marlon Pierce, Judy Qiu, Indiana University	Mixed technical	Nov 2009
	FutureGrid Cloud Technologies and Bioinformatics Applications	Keynote at 1st International Conference, CloudCom 2009, Jiaotong University, Beijing, China	Geoffrey Fox, Indiana University	Mixed technical	Dec 2009
	Cloud Technologies and GeoScience Applications, including FutureGrid	International Symposium on Geo-Computation and Analysis (ISGA) 2009, Laboratory for Information Engineering in Surveying, Mapping and Remote Sensing (LIESMARS), Wuhan University, China	Geoffrey Fox, Indiana University	Technical	Dec 2009
	FutureGrid and Applications	Innovation Center, Indiana University	Geoffrey Fox, Indiana University	Mixed technical	Dec 2009

User Support Committee

Summary

Activities of the User Support Committee during the first 3 months have focused on organization and infrastructure. Actual support incidents have been few and all have been internal: accounts, website issues, etc. Organizationally, the committee has defined a workflow with network operations to manage issues and reporting. A tiered support model is in place. Because most support issues have been internal at this point, 1st tier activity has not been activated. This will change as the systems and users come on board.

Core support is being provided by several mature support units: IU Support Center, IU Research Technology, and the Global NOC. Staff will continue to use tracking systems currently in use. A ticket transfer system will be implemented to manage tickets across systems and a dashboard is planned to view aggregate data relative to support issues.

A central content repository is ready to manage information needed by users and support providers. Plans are underway for full integration of access of this content in the website. There will be further exploration regarding how this information can be integrated into other support channels. In addition, we have proposed a resource catalog to manage details about software and hardware available using web services.

Network Operations

- Produced a workflow for issues implementing a tier approach ensuring access to higher tiers when appropriate
- Identified standard reports for internal and external audiences
- Reached agreement on use of ticket transfer process
- Building contact list for the NOC

FutureGrid Knowledge Base

- All infrastructure completed to manage internal and external content
- Shared FGKB web service specifications with the software group
- Defined requirements for tool to insert FGKB docs dynamically into the FG website
- New FG menu document created to support menu navigation as distinguished from search
- Created internal documentation describing use of FG subversion repository at Sourceforge
- Framework for content contribution for hardware, software, and general use is in place and has been published

Operations & Change Management Committee

Summary

- All committee charters finalized and posted on FutureGrid website
- All subaward documents finalized and submitted to NSF for requisite review and approval
- FutureGrid website created to support initial project ramp up
- Accounting structure for tracking project expenditures, including extensive IU cost-sharing, finalized
- Process for reimbursing non-employee travel for FutureGrid partners finalized
- Process for transferring assets to partner institutions, including all requisite insurance coverage, finalized
- Includes the Change Control Board (CCB), separate charter document attached below

Expenditures Report

Expenditures Report	Q1		Q2		Q3		TOTAL	
	NSF $	Match $	NSF $	Match $	NSF $	Match $	NSF $	Match $
Total Salaries and Wages	15,150	722					15,150	722
Total Fringe Benefits	3,489	51					3,489	51
Total Salaries, Wages and Fringe Benefits	18,639	773					18,639	773
Total Equipment	0	266,981					0	266,981
Total Travel (Domestic)	5,745	56					5,745	56
Total Participant Support	10,804	0					10,804	0
Total Other Direct Costs	0	0					0	0
Subawards	0	0					0	0
Other (Networking, etc.)	0	32,288					0	32,288
Total Direct Costs	35,188	300,098					35,188	300,098
Total Indirect Costs	13,167	0					13,167	0
Total Direct and Indirect	**48,355**	**300,098**					48,355	300,098