Annual Report for Period: 10/2009 - 09/2010
Submitted on: 07/30/2010
Principal Investigator: Fox, Geoffrey C.
Award ID: 0910812
Organization: Indiana University
Submitted By:
Fox, Geoffrey - Principal Investigator
Title:
FutureGrid: An Experimental, High-Performance Grid Test-bed

Project Participants

Senior Personnel

Name: Fox, Geoffrey
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Fortes, Jose
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Keahey, Katarzyna
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Smith, Warren
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: von Laszewski, Gregor
Worked for more than 160 Hours: Yes
Contribution to Project:
Software Architect

Name: Stewart, Craig
Worked for more than 160 Hours: Yes
Contribution to Project:
FutureGrid Executive Director, funded 20% by NSF

Name: Hancock, Dave
Worked for more than 160 Hours: Yes
Contribution to Project:
Chair of the FutureGrid Systems and Networks Committee, funded 25% by NSF and 75% by Indiana University

Name: Smallen, Shava
Worked for more than 160 Hours: Yes
Contribution to Project:
Site Lead for the University of California, San Diego, hair of the Performance Analysis Committee, and oversight for Inca performance monitoring software

Name: Figueiredo, Renato
Worked for more than 160 Hours: Yes
Contribution to Project:
Chair of the Training, Education, and Outreach Committee

Name: Deelman, Ewa

Worked for more than 160 Hours: Yes

Contribution to Project:

Site Lead for the University of Southern California, Information Sciences Institute, and oversight for Pegasus workslow software

Name: Grimshaw, Andrew

Worked for more than 160 Hours: Yes

Contribution to Project:

Andrew Grimshaw replaced Rich Wolski as co-PI on the FutureGrid project during the award process. He is also Chair of the User Advisory Board

Name: Papadopoulos, Phil

Worked for more than 160 Hours: No

Contribution to Project:

co-PI with Shava Smallen at UCSD. Oversees San Diego Supercomputer Center efforts for FutureGrid.

Post-doc

Graduate Student

Undergraduate Student

Technician, Programmer

Name: Maini, Siddharth

Worked for more than 160 Hours: Yes

Contribution to Project:

User Portal developer

Name: Wang, Fugang

Worked for more than 160 Hours: Yes

Contribution to Project:

Software developer, funded 100% NSF.

Name: Tillotson, Jenett

Worked for more than 160 Hours: Yes

Contribution to Project:

System engineer for IBM iDataPlex support and FutureGrid software installations

Name: Wilson, Lucas

Worked for more than 160 Hours: Yes

Contribution to Project:

Software developer on the experiment harness

Name: Voeckler, Jens

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Mehta, Gaurang

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Rynge, Mats

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Matsunaga, Andrea
Worked for more than 160 Hours: Yes
Contribution to Project:
Technical support for FutureGrid demos

Name: Hoover, Paul
Worked for more than 160 Hours: No
Contribution to Project:
UCSD developer

Other Participant

Name: Miksik, Gary
Worked for more than 160 Hours: Yes
Contribution to Project:
FutureGrid Project Manager, funded 50% by NSF and 50% by Indiana University.

Name: Knepper, Richard
Worked for more than 160 Hours: Yes
Contribution to Project:
FutureGrid Site Lead, funded 100% by Indiana University

Name: Pierce, Marlon
Worked for more than 160 Hours: Yes
Contribution to Project:
FutureGrid Gateway Architect, responsible for the design and implementation of the User Portal.

Name: Bolte, Jonathan
Worked for more than 160 Hours: Yes
Contribution to Project:
Chair of the User Support Committee.

Name: Rinkovsky, Joe
Worked for more than 160 Hours: Yes
Contribution to Project:
System administrator for various FutureGrid clusters, including the IBM iDataPlex at Indiana. Funded 100% by Indiana University.

Name: Johnson, Tom
Worked for more than 160 Hours: Yes
Contribution to Project:
Network engineeer overseeing all network efforts connecting Indiana, Chicago, Florida, San Diego, USC, Florida, Texas, and Virginia

Name: Henschel, Robert
Worked for more than 160 Hours: Yes
Contribution to Project:
Benchmarking management for FutureGrid clusters, including Cray XT5m and IBM iDataPlex at Indiana.

Name: Li, Huian
Worked for more than 160 Hours: Yes
Contribution to Project:
Benchmarking support for FutureGrid clusters, including Cray XT5m and IBM iDataPlex at Indiana.

Name: Leggett, Ti
Worked for more than 160 Hours: Yes
Contribution to Project:
Oversight for the IBM iDataPlex at U Chicago and all networking

Name: Bolze, Raphael
Worked for more than 160 Hours: Yes
Contribution to Project:
Project support

Name: Tsugawa, Mauricio
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Collins, Matt
Worked for more than 160 Hours: Yes
Contribution to Project:

Research Experience for Undergraduates

Organizational Partners

GWT-TUD GmbH

GWT-TUD GmbH will support VampirTrace on the FutureGrid platforms and will work with IU to make a best effort to port VampirTrace functionality to other, novel systems added to the testbed. GWT-TUD GmbH will support use of VampirServer on a central system located at IU - most likely the IBM iDataPlex ? for the use of any researcher using FutureGrid.

Other Collaborators or Contacts

See attached document

Activities and Findings

Research and Education Activities: (See PDF version submitted by PI at the end of the report)

IBM iDataPlexes at IU, UC, UCSD & UF, Cray XT5m at IU, Dell cluster at TACC, DDN storage & SunFire storage, Spirent network impairment devide all installed, accepted and available for use. All network connections in place and functional.

Supported early users with approximately 20 applicants and 3 completed projects in benchmarking of Windows v Linux and MapReduce versus Commercial clouds. Supported virtual school with 200 students across the country.

See attached document for more details

Findings:

Demonstrated xCAT and MOAB can deploy dynamically general images in a time of O(5 minutes). Designed a novel software architecture to support general simulation environments based on this concept. Addressed image management, security and user interface supporting general experiments based on Hypervisor, Linux or Windows environments.

See Activities Attached File for more details

Training and Development:

FutureGrid early users included faculty, staff and PhD graduate students who were able to perform research on new approaches to

Bioinformatics Cyberinfrastructure and look at MapReduce technologies comparing Hadoop, Twister and Dryad. The FutureGrid software development probed new approaches to building simulation environments increasing skills of involved staff and students. Over 200 students were supported by FutureGrid in the Big Data tutorial enabling them to learn about MapReduce.

Outreach Activities:

See Activities Attached File, page 15, for table of all outreach activities.

Journal Publications

Thilina Gunarathne, Tak-Lon Wu, Judy Qiu, and Geoffrey Fox, "Cloud Computing Paradigms for Pleasingly Parallel Biomedical Applications", Proceedings of Emerging Computational Methods for the Life Sciences Workshop of ACM, p. , vol. , (2010). Published,

Stewart C.A., Link M., Simms S., Hancock D.Y., Plale, B., Fox G.C., "What is Cyberinfrastructure?", SIGUCCS, p. , vol. , (2010). Accepted,

Marshall, P., Keahey K., Freeman, T, "Elastic Site: Using Clouds to Elastically Extend Site Resources", IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2010), Melbourne, Australia. May 2010, p. , vol. , (2010). Published,

Books or Other One-time Publications

Web/Internet Site

URL(s):

http://www.futuregrid.org

Description:

This is the initial version of the FutureGrid project web site.

Other Specific Products

Contributions

Contributions within Discipline:

Developed and prototyped novel simulation environment spanning clouds, grid and parallel computing. This included enhancements of the Nimbus cloud technology including work on Sky Computing on FutureGrid and Grid5000, Cumulus: a Storage Cloud for Science, Clouds and HTC: a Match Made in Heaven.

Completed early user projects contributed to understanding MapReduce and comparing performance on a variety of platforms. We also developed educational software appliances based on virtual machine technology. Ongoing projects are probing a broad range of middleware and application science computational issues.

See Activities Attached File, Science Highlights

Contributions to Other Disciplines:

The early user completed projects contributed to Bioinformatics and ongoing early user applications cover Life Science and computational fluid dynamics. The CFD applications include work of Cummins -- a major US truck engine company.

Contributions to Human Resource Development:

The work on educational appliances supports laboratories for computing classes. The Big Data tutorial supported by FutureGrid July 26-30 directly trained over 200 students on Cyberinfrastructure and data handling

Contributions to Resources for Research and Education:

Contributions Beyond Science and Engineering:

Conference Proceedings

Special Requirements

Special reporting requirements: None
Change in Objectives or Scope: None
Animal, Human Subjects, Biohazards: None

Categories for which nothing is reported:

Any Book
Any Product
Contributions: To Any Resources for Research and Education
Contributions: To Any Beyond Science and Engineering
Any Conference

FutureGrid

Annual Report: Program Year 1

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)
Craig Stewart, Indiana University (Executive Investigator)

Introduction

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, which allows authorized participants to update information.

There are seven (7) committees around which project activities are 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 (leads: Craig Stewart and Gary Miksik, IU)

The organization of this report is as follows:

- Science Highlights
- Systems Administration and Network Management
- Software
- User Requirements Committee and User Advisory Board
- Performance Analysis
- Training, Education, and Outreach
- User Support
- Operations and Change Management

Science Highlights
Big Data for Science Workshop
Investigate performance overheads of clouds in parallel and distributed environments
Prototype applications (BLAST) across multiple FutureGrid clusters and Grid'5000
Cloud Computing Paradigms for Pleasingly Parallel Biomedical Applications
Cumulus: a Storage Cloud for Science
Clouds and HTC: a Match Made in Heaven?
Industry (Columbus IN) running CFD codes to study combustion strategies to maximize energy efficiency
Investigate metascheduling approaches on Cray and iDataPlex
Deploy Genesis II and Unicore end points on Cray and iDataPlex clusters
Develop new Nimbus cloud capabilities
Compare Amazon, Azure with FutureGrid hardware running Linux, Linux on Xen or Windows for data intensive applications
Test ScaleMP software shared memory for genome assembly
Develop Genetic algorithms on Hadoop for optimization
Attach power monitoring equipment to iDataPlex nodes to study power use versus use characteristics
Support evaluation needed by XD TIS and TAS services
Investigate performance of Kepler workflow engine
Study scalability of SAGA in difference latency scenarios
Test and evaluate new algorithms for phylogenetics/systematics research in CIPRES portal

Big Data for Science Workshop

FutureGrid supported the hands on component of the Big Data for Science Workshop http://salsahpc.indiana.edu/tutorial/index.html which was offered July 26-30 to 300 students across the country as part of VSCSE, the Virtual School of Computational Science and Engineering. Ten sites were involved: Arkansas High Performance Computing Center, University of Arkansas, Fayetteville; Electronic Visualization Laboratory, University of Illinois at Chicago; Pervasive Technology Institute at Indiana University, Bloomington; Institute for Digital Research and Education, University of California, Los Angeles; Michigan State University, East Lansing; Pennsylvania State University, University Park; University of Iowa, Iowa City; University of Minnesota Supercomputing Institute, Minneapolis; University of Notre Dame, Notre Dame, Indiana; University of Texas at El Paso. FutureGrid supported the students experimenting with Twister and Hadoop MapReduce technologies following tutorials on their theory and use. The workshop was organized by Judy Qiu and FutureGrid PI Geoffrey Fox and built on earlier use of FutureGrid on performance measurements for MapReduce and comparison with Azure and Amazon.

SPEC MPI2007

A collection of performance benchmarks were executed on the IBM iDataPlex cluster (the UCSD/SDSC FutureGrid machine) using two different operating systems. Windows HPC Server 2008 (WinHPC) and Red Hat Enterprise Linux v5.4 (RHEL5) were compared using SPEC MPI2007 v1.1, the High Performance Computing Challenge (HPCC), and National Science Foundation (NSF) acceptance test benchmark suites. Overall, we found the performance of WinHPC and RHEL5 to be equivalent, but significant performance differences exist when analyzing specific applications. The focus was on presenting the results from the application benchmarks and including the results of the HPCC microbenchmark for completeness.

A paper detailing the results was submitted to the Standard Performance Evaluation Corporation (SPEC) Workshop to be held in Paderborn, Germany in October 2010.

SKY COMPUTING ON FUTUREGRID AND GRID'5000

"Sky computing" is an emerging computing model where resources from multiple cloud providers are leveraged to create large-scale distributed infrastructures. This demonstration will show how sky computing resources can be used as a platform for the execution of a bioinformatics application (BLAST). The application will be dynamically scaled out with new resources as need arises. This demonstration will also show how resources across two experimental projects – the FutureGrid experimental testbed in the United States and the Grid'5000, an infrastructure for large scale parallel and distributed computing research in France – can be combined and used to support large-scale distributed experiments. The demo will showcase not only the capabilities of the experimental platforms, but also their emerging collaboration. Finally, the demo will showcase several open source technologies. Specifically, our demo will use Nimbus for cloud management, offering virtual machine provisioning and contextualization services, ViNe to enable all-to-all communication among multiple clouds, and Hadoop for parallel fault-tolerant execution of BLAST.

The University of Florida team successfully demonstrated their ViNe and CloudBLAST technologies at the 10th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2010) during the 3rd IEEE International Scalable Computing Challenge, and at the Open Grid Forum 29 (OGF29). The OGF demonstration marked the collaboration between FutureGrid and the Grid'5000 (France) efforts. ViNe was a key component to enable the communication among resources, especially for the heavily protected Grid'5000 resources.

- 150-node virtual cluster across 5 sites, connected through ViNe (including 2 FutureGrid sites UF and SDSC) and running CloudBLAST, demonstrated a large scale BLAST run in less than 1 hour (the same job takes over 17 days sequentially)
- OGF29: 457-node virtual cluster across 3 FutureGrid sites (UF, UC, SDSC) and 2 G5k sites (Rennes, Sophia) demonstrated the dynamic extension of a Hadoop cluster. A modified Nimbus, with faster VM deployment, was deployed on G5k. As new VMs became available on G5k, the Hadoop worker pool was increased speeding up the CloudBLAST throughput.

The project developed a prototype infrastructure capable of provisioning virtual resources over multiple distributed clouds (a "sky" of clouds) and combining them into one virtual cluster. This creates a way of creating a common platform over diverse resources: during a live demonstration at OGF-29, this infrastructure was used to provision a virtual cluster of over a thousand cores that was then used to run BLAST. The resources were provisioned from six sites configured as Nimbus clouds: 3 FutureGrid sites (UCSD, UFL, and UC), and 3 Grid5000 sites (Lille, Rennes, Sophia) demonstrating how the two different experimental testbeds can interoperate. The project was developed using open source technologies: Nimbus (UC), ViNE (UFL), TakTuk, and Kastafior (Grid5000).

The project was a collaboration between Pierre Riteau (University of Rennes), M. Tsugawa, A. Matsunaga, J. Fortes (UFL) and , T. Freeman, D. LaBissoniere, K. Keahey (UC). The project is presented as a poster "Sky Computing on FutureGrid and Grid5000" at TeraGrid'10.

Cloud Computing Paradigms for Pleasingly Parallel Biomedical Applications

A team from the Indiana University SALSA group compared MapReduce on FutureGrid with an alternative implementation using the "master-worker" approach offered on the commercial cloud infrastructure service based virtual machine utility computing models of Amazon AWS and Microsoft Windows Azure. On FutureGrid they used the MapReduce based computing frameworks Apache Hadoop (deployed on raw hardware as well as on virtual machines) and Microsoft DryadLINQ. This involved running Windows Linux and Xen on the same FutureGrid system -- the UCSD/SDSC IBM iDataPlex hardware. They compared performance showing strong variations in cost between different EC2 machine choices and comparable performance between the utility computing (spawn off a set of jobs) and managed parallelism (MapReduce). The MapReduce approach offered the most user friendly approach.

This paper was published as Thilina Gunarathne, Tak-Lon Wu, Judy Qiu, and Geoffrey Fox, Cloud Computing Paradigms for Pleasingly Parallel Biomedical Applications March 21 2010. Proceedings of Emerging Computational Methods for the Life Sciences Workshop of ACM HPDC 2010 conference, Chicago, Illinois, June 20-25, 2010.

Cumulus: a Storage Cloud for Science

Cumulus is an open source implementation of a storage cloud. Its interface is compatible with the Amazon S3 REST API. It is packaged with the Nimbus toolkit and provides scalable and reliable access to scientific data. Cumulus is implemented as a pluggable system allowing users to connect to storage management system popular for scientific projects. Existing and planned backends include: POSIX, Hadoop file system, Sector, and BlobSeer. Cumulus also extends the S3 API to provide quota management implementing econometrics suitable to the scientific community. We performed initial performance studies of the Cumulus and compared it to other systems used in the scientific community (such as GridFTP) on UC FutureGrid resources.

The project is described in "Cumulus: a Storage Cloud for Science", J. Bresnahan, T. Freeman, D. LaBissoniere, K. Keahey, submitted as poster to SC10

Clouds and HTC: a Match Made in Heaven?

A challenge for the providers of on-demand clouds, such as Amazon EC2 is to make them cost-effective: in order to provide on-demand availability a provider needs to either significantly overprovision (and thus put up with low utilization, which is expensive), or reject a large proportion of requests in practice (and thus no longer be "on-demand"). This investigation proposes to use high-throughput computing (HTC) as backfill to an on-demand cloud and thus provide improved utilization while still providing on-demand capabilities. This summer, UC FutureGrid resources were used to explore the effect of this solution on user codes as well as utilization.

The project is described in "Clouds and HTC: a Match Made in Heaven?", P. Marshall, K. Keahey, T. Freeman, submitted as poster to SC10.

INDUSTRIAL PARTNERSHIPS

Cummins, Inc.

Dr. Steven E. Koonin, DOE Undersecretary for Science, gave a presentation to the National Science Foundation Advisory Committee for Cyberinfrastructure on 27 May 2010 during which he stated that "Within the US we do not do a good job of marshalling our public cyberinfrastructure resources as effectively as we should to meet national priorities in collaboration with the private sector" and "We need more examples of HPC [High Performance Computing] being used to the benefit of US industry in solving some of our national energy and manufacturing needs."

Cummins, Inc., a global manufacturing company with major facilities in Columbus IN, has been interested in collaborating with Indiana University to use FutureGrid in measuring the efficacy of new methods for designing diesel engine combustion solutions. Cummins is particularly interested in testing an automated optimization code called ModeFRONTIER to run computational fluid dynamics (CFD) codes to study combustion strategies to maximize energy efficiency, minimize pollutant emissions, etc. The work being done is experimental in that Cummins is investigating new automated methods for performing this design process. There is therefore an excellent opportunity to perform a performance analysis of 'time to solution' using Cummins' existing cluster facilities as well as FutureGrid systems.

The interesting thing to publish is not the engine design information, but rather how use of systems in FutureGrid compare in terms of performance with each other and how they may help accelerate the design process. Comparison of the two FutureGrid systems described below against each other and Cummins' current computing cluster"

- The *IBM iDataPlex* system is an IBM e1350 distributed shared memory cluster with 1024 processor cores and 3 TB total memory capacity. The compute nodes consist of 128 dx360 M2 iDataPlex servers, each with two quad-core Intel Xeon processors, 24 GB of memory, and a PCIe Mellanox ConnectX 4x DDR InfiniBand adapter for high bandwidth, low-latency MPI applications.
- The *Cray XT5m* is a distributed shared memory cluster with 672 processor cores and 1.3 TB total memory capacity. The compute blades consist of 21 XT5 blades, each with eight quad-core AMD Shanghai processors, 64 GB of memory, and an integrated Cray SeaStar adapter for high bandwidth, low-latency MPI applications.

Name	Architecture	TFLOPS	Total RAM (TB)	Local Disk (TB)
FutureGrid	IBM e1350	12	6	128
FutureGrid	Cray XT5m	6	1.3	5.5

The published work will compare performance on these two platforms with each other and analyze the use of automated optimization software in industry.

Systems Administration & Network Management

Compute and Storage Systems

A summary of the status of FutureGrid hardware as of 31July 2010 is shown below:

System type	# CPUs	# Cores	TFLOPS	Total RAM (GB)	Secondary Storage (TB)	Site	Status
Dynamically configurable systems							
IBM iDataPlex	256	1024	11	3072	339*	IU	Operational
Dell PowerEdge	192	768	8	1152	30	TACC	Being installed
IBM iDataPlex	168	672	7	2016	120	UC	Operational
IBM iDataPlex	168	672	7	2688	96	SDSC	Operational
Subtotal	*784*	*3136*	*33*	*8928*	*585*		
Systems not dynamically configurable							
Cray XT5m	168	672	6	1344	339*	IU	Operational
Shared memory system TBD	40	480	4	640	339*	IU	New System TBD
IBM iDataPlex	64	256	2	768	1	UF	Operational
High Throughput Cluster	192	384	4	192		PU	Not yet integrated
Subtotal	*464*	*1792*	*16*	*2944*	*1*		
Total	***1248***	***4928***	***49***	***11872***	***586***		

The IBM iDataPlex systems were ordered with a purchase order placed on 10/28/2009. Acceptance tests were completed successfully, in the view of IU staff and the PI, on 5/6/2010. NSF notified IU that an external review had confirmed the substance of the acceptance tests and endorsed IU's view that the terms of the purchase contract and that IBM's invoice for this system should be paid on 7/21/2010.

The IU Cray XT5m for IU was ordered on 10/30/2009 and delivered on 01/08/2010. It was declared to have passed its acceptance tests on 02/02/2010.

An IBM iDataplex owned by IU prior to the start of the NSF FutureGrid award was delivered to SDSC in February 2010, and performance tests there replicated results of performance tests done earlier at IU. Other, smaller pieces of equipment, particularly storage systems, were installed and passed acceptance tests, but, with costs under $200,000, did not require NSF review of acceptance tests.

The Dell PowerEdge cluster planned for TACC is being installed as of the writing of this report.

The FutureGrid network is fully operational, and is schematically depicted in Figure 1.

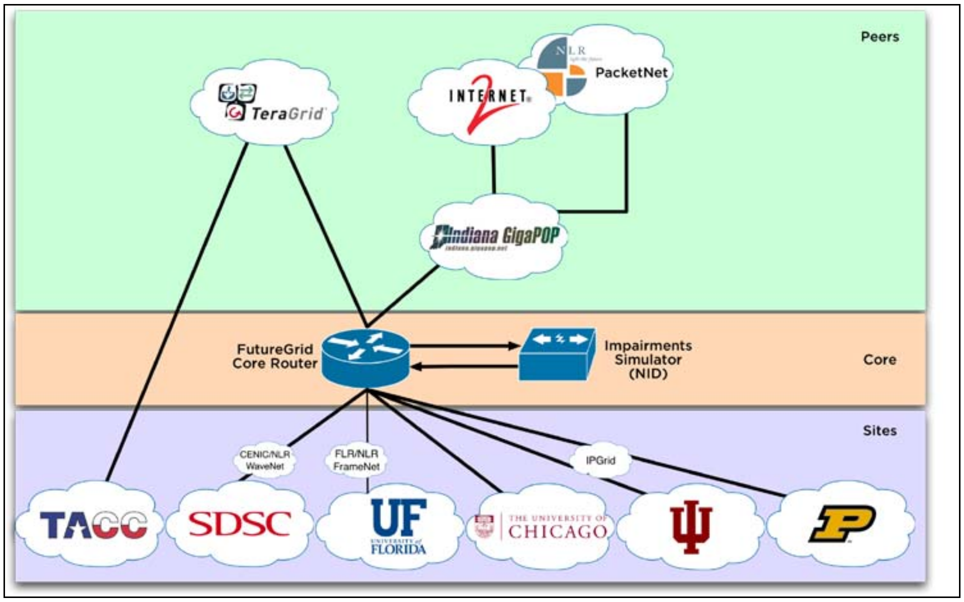

Figure 1. Schematic diagram of FutureGrid network.

Software

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. Moab is the scheduling system, which is a major component of the FutureGrid software architecture. A three-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 FutureGrid 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:

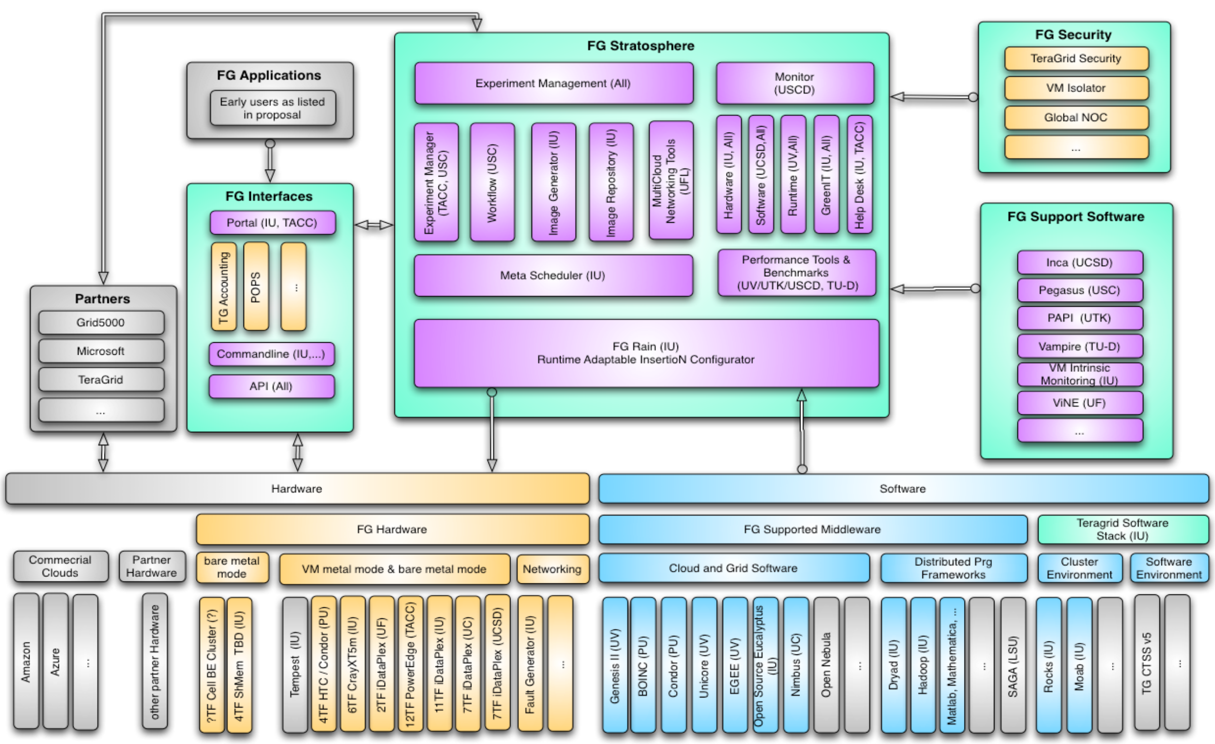

Figure 2. Schematic diagram of software architecture.

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 now in use.

Key information on software subsystems follows:

Performance: Activities to improve the performance monitoring software have continued. Noteworthy are the discussions of providing web space and support for users to publish their performance results (FutureGrid v. TeraGrid v. Azure v. Amazon). The HPCC performance tests are now available through Inca. This has also been shared with the external collaborators at University of Buffalo (see below). Installation of performance tools on test machine is complete. The next step is to provide packages to the hardware group for installation on FutureGrid machines (the IU applications group and Hardware group are currently designing the process).

Image Repository: We have set up a testing deployment for the image repository. We have also investigated and designed distributed data storage models for the images that are suitable for the FutureGrid network layout.

Image Creation: A FutureGrid image creation and validation system is under design and development. This system will let users pick their desired operating system, a set of preinstalled software specific to their development environment, and the amount of system resources they would like each virtual machine to acquire. A command line interface is targeted; however, the ultimate goal would be to incorporate such functionality within a portal.

Moab: The original setup of the queuing system used until now had some limitations. Thus, we set up a testbed as part of the software development activities to highlight dynamic provisioning controlled directly from Moab. This is a continuation of the work that we earlier did with a TORQUE/xCAT-based install. We believe the way this install is conducted will be a key issue towards dynamic provisioning of images via the Moab MSM on FutureGrid resources. Interestingly, it was reported that setting up a similar dynamic provisioning via TORQUE and xCAT instead of Moab and xCAT was easier to complete. We will be using Moab/xCAT for the general install on the resources, and use TORQUE/xCAT in cases we need to instantiate a virtual cluster within the staged images. A good example for a use case for this is the development of virtual Grid and cluster environments for the software development of FutureGrid.

Monitoring: For Inca, work has been done to plan out the Inca interfaces for the portal and improve documentation. In addition, the Inca graphing interfaces to display HPCC performance data are being set up. The group also discussed the setup of Netlogger (deployment milestone begins this month) in order to collect performance and load data from FutureGrid components. We pushed back HPCC performance test milestones to the end of July so we can include Tango from TACC.

Experiment Harness: We have evaluated the Qpid message broker to learn more about some advanced features, verify that it meets our needs, and plan how to use it in the Experiment Harness. In particular, we are investigating the dynamic creation of message queues, dynamic modification of who has access to message queues, message filtering, and X.509 authentication. The purpose of this is to be able to authenticate users to the harness and segregate experiment monitoring and management messages so that different user groups do not interfere with each other. We have also spent some time enhancing our experiment harness client and daemon implementations.

Dynamic Provisioning: A simple dynamic provisioning experiment has proven successful on the IU testbed (fg-gravel). This test includes provisioning virtual machines via xCAT using Torque to VirtualBox nodes.

Image Repository: More coding and testing on the Image Repository development has been conducted. Unix man page style documentation has been started and a deployment of the preliminary version at the fg-gravel cluster has been conducted. The application framework, CLI parsing, and back-end data persistence are done, and metadata put and query are functioning. More time is needed to complete the image file storage procedure as well as deal with the back end data file deployment issue in the planned ssh-based approach for the current phase. The Phase II image repository design document has been initiated.

Pegasus: A task to run a very simple job on all three sites (Xray, India, Sierra) was initiated and completed. The task assures the quality of the installed services and available documentation at their lowest level. An additional issue needs to be addressed to provide a user CA for Nimbus. We worked on updating the Netlogger installation on inca.futuregrid.org to update yet the newest release. The Netlogger broker is integrated into inca.futuregrid.org and comes with init.d scripts, and the API libraries for Python, Perl and Java were re-installed. MongoDB was installed. We started the activity to document and discuss a “FG Pegasus Software Stack” and added instruction on how to install Pegasus using a yum repository. To fulfill the requirements of FG new features will be added with the upcoming Pegasus 3.0 release. We have identified that special attention will be placed on dependencies with Globus and Condor.

USC provided preliminary documentation on the installation of Pegasus-WMS on FutureGrid VM images.

Eucalyptus: The long stalled activity of deploying Eucalyptus on india.futurgrid.org was restarted as the person responsible was on family leave and the system administrative group had no additional resources to work on it till now. Networking issues on sierra are still being worked on. The system was available at the end of June.

Nimbus: Several software updates have been conducted to Nimbus resulting in several releases. UC completed integration and testing for Nimbus 2.5 and released Nimbus cloud client version 15 (which includes the FutureGrid root CA certificates, simplifying the client setup for FutureGrid). Nimbus was also installed on the SDSC and UF clusters

Genesis 2: Installation of Genesis 2 on the IU Cray machine was started.

User Advisory Board

The first meeting of the FutureGrid User Advisory Board will be held on 2 August 2010 at the TeraGrid 2010 conference in Pittsburgh, PA. The User Advisory Board includes:

- Ann Chervenak ISI
- Andrew Grimshaw UVA -- Chair of UAB
- Shantenu Jha LSU
- Steven Newhouse EGI
- Ruth Pordes OSG
- Morris Riedel Unicore 6/FZJ
- John Towns NCSA/TIS
- Jon Weisman UMN
- Rich Wolski Eucalyptus/UCSB
- Nancy Wilkens-Diehr SDSC
- Frederic Desprez Inria

Performance Analysis

The Performance Analysis team has enhanced the initial Inca deployment at http://inca.futuregrid.org/ for detecting functionality and performance problems on FutureGrid. The main enhancement was deploying HPCC as part of our automated benchmarking work to detect performance problems. New graphing interfaces were deployed to display the data in historical graphs and a new Inca reporter was developed to execute HPCC, process, and report the statistics. The Performance Analysis Committee also decided to use Netlogger as an instrumentation API to passively collect performance and usage data from FutureGrid software components. This is a well-utilized tool and the ISI group members are also partnering with the Netlogger developer on another project, providing good access to consulting and support. An initial Netlogger deployment has been installed on http://inca.futuregrid.org/.

Training, Education, and Outreach

Table 1 enumerates the Training, Education, and Outreach (TEO) events completed during Q1 – Q3 of FutureGrid Program Year 1. A significant accomplishment to date is the creation of a slide deck for a general purpose tutorial called "FutureGrid 100 and 101." This slide deck contains material suitable for a half-day tutorial, and portions of it can be used for shorter talks. This set of slides grew out of early talks developed particularly by PI Fox, Craig Stewart, David Hancock, and Gregor von Lasziewski. The biggest educational effort to date is underway as this repot is being submitted – the use of these educational materials and FutureGrid as part of a Virtual School for Computational Science and Engineering class on "Big Data for Science" (http://www.vscse.org/summerschool/2010/bigdata.html). Approximately 200 students will use IU assistant professor Judy Qiu's Twister application running on FutureGrid as part of hands-on educational exercises associated with this class. Drs. Fox and Stewart gave the 'FutureGrid 100 & 101" talks as a 90-minute session as preparation for these exercises.

A key activity of the TEO group has been the preparation by the U. Florida group of its educational "Grid appliances" to the two cloud-provisioning technologies currently supported by FutureGrid – Eucalyptus and the Nimbus toolkit. Progress highlights include:

- Deployment of Grid appliance clusters via Nimbus – we have been able to deploy a version of the Hadoop Grid appliance that uses packages built upon Ubuntu 9.04 on FutureGrid's foxtrot. We have successfully deployed a small-scale 4-node virtual cluster connected through GroupVPN and running Hadoop through Nimbus command-line tools. Development is underway to create images based on Ubuntu 9.10.
- Deployment of Grid appliance via Eucalyptus – we have been able to deploy a version of the Condor Grid appliance that uses packages built upon Ubuntu 10.4 on FutureGrid's india. We have successfully deployed a small-scale 4-node virtual cluster connected through GroupVPN to the public Grid appliance pool maintained by University of Florida. Development is under way to create images for Hadoop and MPI educational appliances to be made available through the image repository on FutureGrid

Training, Education, and Outreach Events

Type of Event		Title	Location	Presenter	Type of Audience	Month
Workshops and Tutorials	Q1	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
	Q2	Performance Analysis Using the Vampir Toolchain	Innovation Center, Indiana University	Robert Henschel, Indiana University; Thomas William, ZIH, Dresden	Technical	April 2010
	Q4	Big Data for Science Workshop	Innovation Center, Indiana University	Geoffrey Fox, Judy Qiu, Indiana University, et al.	Mixed technical	July 2010
Presentations and Talks	Q1	FutureGrid	Open Grid Forum 27, Banff, Canada	Andrew Grimshaw, University of Virginia	Advanced technical	Oct 2009
	Q1	FutureGrid Overview	CCA-09 Cloud Computing and its Applications Workshop, Chicago, IL	Geoffrey Fox, Indiana University	Technical	Oct-2009
	Q1	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
	Q1	FutureGrid and Green Aware Computing	Northern Illinois University, DeKalb, IL	Gregor von Laszewski, Indiana University	Technical	Oct 2009
	Q1	FutureGrid: An Experimental High-Performance Grid Testbed	NCSA, Champaign-Urbana, IL	Craig Stewart, Indiana University	Technical	Oct 2009
	Q1	FutureGrid Overview	SC09 Conference, Indiana University booth, Portland, OR	Geoffrey Fox, Indiana University	Mixed technical	Nov 2009
	Q1	FutureGrid Overview	SC09 Conference, AIST booth, Portland, OR	Geoffrey Fox, Indiana University	Mixed technical	Nov 2009
	Q1	FutureGrid Overview	SC09 Conference, Indiana University booth, Portland, OR	Geoffrey Fox, Gregor von Laszewski, Marlon Pierce, Judy Qiu, Indiana University	Mixed technical	Nov 2009
	Q1	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
	Q1	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
	Q1	FutureGrid and Applications	Innovation Center, Indiana University	Geoffrey Fox, Indiana University	Mixed technical	Dec 2009
Presentations and Talks	Q2	Clouds and FutureGrid	All Hands Meeting, Minority Serving	Geoffrey Fox, Indiana University	Mixed technical	Jan 2010

			Institutions – Cyberinfrastructure Empowerment Coalition (MSI-CIEC), SDSC, San Diego, CA			
	Q2	Building Effective CyberGIS: FutureGrid	National Science Foundation TeraGrid Workshop on Cyber-GIS, Washington, DC	Marlon Pierce and Geoffrey Fox, Indiana University	Technical	Feb 2010
	Q2	FutureGrid: An Experimental High-Performance Grid Testbed	TeraGrid Quarterly Meeting, Tampa, FL	Craig Stewart, Indiana University	Technical	Mar 2010
	Q2	FutureGrid Introduction	All Hands Meeting, Open Science Grid, Fermilab, Batavia, IL	Gregor von Laszewski, Indiana University	Technical	Mar 2010
	Q3	FutureGrid Introduction	Department of Energy, MAGIC virtual meeting	Gregor von Laszewski, Indiana University	Mixed technical	April 2010
	Q3	Cloud Technologies and Data Intensive Applications	Instrumenting the Grid (INGRID) Workshop, Poznan, Poland	Geoffrey Fox, Indiana University	Technical	May 2010
	Q3	Cloud Technologies and Their Applications	Keynote at 5th International Workshop on Content Delivery Networks (CDN 2010) in The 10th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid) 2010, Melbourne, Victoria, Australia	Judy Qiu, Indiana University		May 2010
	Q3	Sky Computing: When Multiple Clouds Become One	The 10th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2010), Melbourne, Victoria, Australia	Jose Fortes, University of Florida	Technical	May 2010
	Q3	Cloud Cyberinfrastructure and Collaboration	The 2010 International Symposium on Collaborative Technologies and Systems (CTS 2010), Westin Lombard Yorktown Center, Chicago, IL	Geoffrey Fox, Indiana University	Technical	May 2010
	Q3	FutureGrid	Venus-C, Brussels, Belgium	Geoffrey Fox, Indiana University	Technical	May 2010
	Q3	Algorithms and Applications for Grids and Clouds	22nd ACM Symposium on Parallelism in Algorithms and Architectures, Santorini, Greece	Geoffrey Fox, Indiana University	Technical	June 2010
	Q3	FutureGrid: Supporting Next Generation Cyberinfrastructure	The 4th International Workshop on Virtualization Technologies in Distributed Computing (VTDC10) at HPDC, Chicago, IL	Geoffrey Fox, Indiana University	Technical	June 2010
	Q3	Cloud versus Cloud: How Will Cloud Computing Shape Our World?	Panel, HPDC, Chicago, IL	Geoffrey Fox, Indiana University	Mixed technical	June 2010
	Q3	Grids and Clouds for Cyberinfrastructure	Illinois Institute of Technology, Chicago, IL	Geoffrey Fox, Indiana University	Technical	June 2010
	Q3	Autonomic Computing Across Clouds	Grids Meet Autonomic Computing (GMAC), Washington, DC	Jose Fortes, University of Florida	Technical	June 2010
	Q3	Sky Computing on FutureGrid and Grid'5000	OGF-29, Chicago, IL	P. Riteau, et al., University of Florida	Technical	June 2010
	Q4	Cross-Cloud Computing	High Performance Computing GRIDS and Clouds (HPDC), Centraro, Italy	Jose Fortes, University of Florida	Technical	July 2010
	Q4	Virtually Networking the Clouds	The 9th IFIP Annual Mediterranean AD HOC Networking Workshop,	Jose Fortes, University of Florida	Technical	July 2010

			Juan-Les-Pins, France			
	Q4	FutureGrid Tutorial	TeraGrid 2010, Pittsburgh, PA	Scott McCaulay, Judy Qiu, Marlon Pierce, Rich Knepper, Indiana University	Mixed technical	Aug 2010
	Q4	FutureGrid BOF	TeraGrid 2010, Pittsburgh, PA	Geoffrey Fox, Indiana University, and FutureGrid Project Team	Mixed technical	Aug 2010
	Q4	Virtual Appliances for Training and Education on FutureGrid	TeraGrid 2010, Pittsburgh, PA	Renato Figueiredo et al, University of Florida	Mixed technical	Aug 2010
	Q4	Sky Computing on FutureGrid and Grid'5000	TeraGrid 2010, Pittsburgh, PA	P. Riteau, et al., University of Florida	Mixed technical	Aug 2010

Table 1. Outreach and training events completed between 1 October 2009 and 05 August 2010

User Support

Work continues on the revision of the futuregrid.org site, which will be an intermediate step on the way to the FutureGrid Portal. A new version of the FutureGrid web site and portal will go into production on July 30. Population of the FutureGrid Knowledge Base (KB) continues apace; 59 entries have been created to date. A critical component of the success of the KB as a "first source for help" is getting it well enough populated that users have a good chance of finding a helpful KB entry. Users are thus encouraged to use the KB before sending email or calling on the telephone for help. This aids the overall effectiveness of support and will allow the FutureGrid team to focus live, human help on problems that require human intervention for a solution.

Operations & Change Management

- All subaward purchase orders have been finalized (Chicago, Florida, USC, San Diego, Texas, and Virginia).
- A FutureGrid BOF abstract was submitted to TeraGrid 2010 and accepted. It is to be combined with those from other new TeraGrid sites for an expanded session.
- Early adopter allocation process has been refined to include a variant of the TeraGrid user responsibility form and information on important security aspects (country of citizenship of requestors).
- The Operations Committee has specifically examined and approved (with the endorsement of the UAB Chair) the following two notable requests for novel use of FutureGrid:
    - UCSD, with help from IU, put forward an internal proposal for SDSC user Catherine Olschanowsky to attach a power monitoring harness to one node of the UCSD Sierra cluster. This experiment will incur a cost for IBM to recertify the node and the proposal was approved this week, first by the Operations committee and then by PI Fox.
    - A proposal for collaboration with Cummins, Inc. to use FutureGrid was endorsed by the Operations Committee and approved by FutureGrid PI Fox.