Workshop Report:

High Performance Applications of Cloud and Grid Tools

Prepared for: US Department of Energy Advanced Scientific Computing Research Program

Prepared by: Alan Sill, Vice President of Standards, Open Grid Forum

May 4, 2011

Table of Contents

Workshop Summary

Original Objective

The workshop objective was to improve on the state of use of software tools designed for access, administration, job management, input/output standardization, performance optimization of job submission and execution, and other aspects of management of high performance computing resources by multiple users through incorporation of the latest advances in cloud and grid development tools. The workshop was announced to be focused specifically on tools in use or that could be used to produce and maintain production infrastructures for large-scale projects in scientific computing.

Attendance and content of the workshop

The event was held as scheduled April 14, 2011 at the Argonne National Laboratory, Building 240 Conference Center, near Chicago, Illinois. A total of 40 participants were registered of which nearly all attended, with several additional participants dropping in on the day of the event. The workshop followed immediately after the preceding Globus World (April 11-12) and Cloud Computing Applications 2011 conferences (April 13) at the same location. Co-location of this event with these two conferences allowed us to gain participants from the preceding events, although a significant fraction of the participants came in only or primarily to attend this workshop.

Eight speakers presented talks on talks spanning cloud and grid infrastructure software, security, authorization and authentication, scaling performance, on-demand networking and standards for large-scale deployment of both networking and middleware tools. Talks ranged from 20 to 30 minutes in length, with a short amount of time for discussion. Talks are available for download at the workshop web site, http://www.ogf.org/HPACG and the detailed agenda is given below. Although this was a rapid pace, the resulting concentration of topics allowed a high degree of concentration on each of the topics, and considerable information was exchanged between participants.

A working lunch followed, which was used to allow participants to engage in detailed discussion with one another about the topics covered in the talks in the morning session. The afternoon session was used to gather input from participants regarding the topics covered and to discuss plans for a larger workshop to be held in July to expand the topic range and to include a broader range of science agency project participants.

Part of the afternoon meeting was devoted to a conference call between the organizers and the Teragrid 2011 program committee to work out details of the plans for this second workshop.

Detailed agenda

Name	Organization	Topic	Start	End	Duration
Sill, Alan	Texas Tech University	Workshop Introduction	8:30 AM	8:40 AM	0:10
Sill, Alan	Texas Tech University	Recently Released OGF Standards; US Federal and Other Roadmap Efforts	8:40 AM	8:50 AM	0:10
Keahey, Kate	Argonne National Laboratory	Nimbus Cloud Computing for Science: Architecture, Status and Future Plans	8:50 AM	9:10 AM	0:20
Riteau, Pierre	IRISA/University of Rennes 1	Large Scale Sky Computing Applications with Nimbus	9:10 AM	9:30 AM	0:20
Timm, Steven	Fermilab	Security, AuthN/AuthZ and Scaling Performance Topics in FermiCloud	9:30 AM	10:00 AM	0:30
Break			*10:00 AM*	*10:20 AM*	*0:20*
Papaspyrou, Alexander	TU Dortmund University	Implementation of OGF's Open Cloud Computing Interface on top of libvirt	10:20 AM	10:40 AM	0:20
von Laszewski, Gregor	University of Indiana	Current Status and Plans for FutureGrid	10:40 AM	11:10 AM	0:30
Tannenbaum, Todd	University of Wisconsin-Madison	From 10K cores on EC2 to 60K cores on OSG - on the fly Condor deployment	11:10 AM	11:35 AM	0:25
Zurawski, Jason	Internet2	Scientific Networking with DYNES & LHCONE	11:35 AM	12:00 PM	0:25
Working Lunch	*All Participants*	*Roundtable: How Can Cloud and Grid Innovation Be Harnessed For Advanced High Performance Computing Needs?*	*12:00 PM*	*1:00 PM*	*1:00*
Afternoon session	Organizing Committee	July 18 Workshop on Science Agency Uses of Grids and Clouds (members of committee only)	1:00 PM	3:00 PM	2:00

Outcomes and evaluation against original goals

In this section, we evaluate the workshop results against the following original goals:

1. *Gather authors, maintainers, users and interested parties for production cloud and grid infrastructure tools. Survey the existing state of these tools with regard to use and implementation of existing cloud and grid standards.*

   Result: Achieved in the form of approximately three dozen participants from a broad range of current topical infrastructure projects. Participant feedback included comments that the concentrated nature of the workshop on topics directly related to high-performance infrastructure needs allowed concentrated exchange of information.

2. *Update participants on current US government and international efforts to promote cloud computing usage and associated efforts on related standards development, promotion and adoption.*

   Result: Achieved in the overall introductory talk given by the organizers to the workshop participants. We have seen higher participation in meetings related to the NIST, EGEE, SIENA and other ongoing efforts by participants since this workshop.

3. *Focus the talents of the community on options for cloud-to-grid and grid-to-cloud interoperability, high performance applications of existing tools and applicability of OGF and other emerging standards, and gather input from attendees regarding the current state of their projects with respect to these topics. Using the resulting input, explore options for future development that can be achieved on a short-term basis.*

   Result: Although cloud-to-grid and grid-to-cloud topics were discussed by only two or three of the speakers, considerable discussion did take place about existing and emerging standards, and one talk was presented by a member of the Open Cloud Computing Interface (OGF OCCI) team that was well received by the workshop participants.

   Considerable exchange of information took place between the OCCI, Nimbus, Condor and Fermi Cloud participants both within their talks on middleware infrastructure topics and in later lunchtime discussion, and with the Future Grid speaker and members of the audience on prospects for including these technologies in ongoing development-oriented test beds.

4. *Serve as a planning forum for a larger-scale workshop and forum aimed at evaluating, encouraging and promoting communication on this topic broadly across US-funded science agency computing projects and among their related international partners.*

   Result: This planning session was completed and the larger workshop tentatively planned. See further notes on this topic below.

Budget expense report

Expenses were budgeted for travel expenses for the workshop organizers and invited speakers, staff support, refreshments for participants, promotional materials and related ancillary expenses. The majority of these were expended as planned, but the portion of the expenses planned to cover travel for invited speakers was not required, as all speakers were able and willing to attend based on their available project resources.

Future workshop plans

One of the original goals for this workshop was to survey the field and determine the need for (and if and as appropriate, develop detailed plans for) a future event to cover science infrastructure, middleware and project needs for more organized use of clouds and grids in pursuit of science agency missions.

This second workshop is intended to be held of July 18, 2011 in Salt Lake City, Utah on the overlap day the Open Grid Forum 32 and Teragrid 2011 conferences (July 15-18 and July 18-21, respectively), with the title of "International Workshop on Science Agency Uses of Clouds and Grids" selected by participants, and an organizing committee formed to plan and screen talks and recruit participants. The focus of this next workshop will be on science infrastructures for large-scale or innovative projects that can serve a variety of science missions with similar or customized versions of grid and cloud middleware, as well as use of new technologies for cost savings, increased efficiency of delivery or other technical advances in delivery of computation, data storage and/or access, or network technology for science infrastructure projects.

An organizing committee has been selected and materials prepared to invite participants to this next workshop. Our intention is to present this report and the proposal for the next workshop to the LSN MAGIC meeting May 4, 2011 for agency discussion and consideration with the goal to identify interested US-based and international science agencies.

As with the workshop that is the subject of this report, co-location of the next event with the above two large-scale conferences will allow for convenient participation by prospective attendees of these other events, as well as direct attendance by participants who come in specifically for this event.

List of HPACG participants

Name	Company/Organization	Job Title
Behzad, Babak	University of Illinois at Urbana-Champaign	PhD Student
Binsteiner, Markus	University of Auckland	Software developer
Bui, Anh Huy	University of Illinois at Chicago	Phd Candidate
Chadwick, Keith *	Fermilab	Grid & Cloud Computing Dept Head
Cheng, Maggie	University of Missouri- Rolla	Professor
Clapp, George	Telcordia Technologies	Chief Scientist
Coghlan, Susan	Argonne National Lab	Division Director
Fraser, Dan	Open Science Grid	OSG Production Coordinator
Frey, Jaime *	Univ of Wisconsin-Madison	Researcher
Fuess, Stuart	Fermilab	Scientist
Hesselroth, Ted	Fermilab	Senior Application Specialist
Holzman, Burt	Fermilab	Computing Professional
Jin, Bo	UIUC	Dr.
Jones, Mike *	The University of Manchester	Research Computing Services Manager (Acting)
Keahey, Kate *	Argonne National Laboratory	Scientist
Kettimuthu, Rajkumar	Argonne National Laboratory	Principal Software Development Specialist
Lee, Jong	NCSA	Research Scientist
Lewis, Michael	Electronic Visualization Lab	Researcher, PhD Student
Liu, Yong	NCSA	Sr. Research Scientist
Maheshwari, Ketan	Argonne National Laboratory	postdoc
Mambelli, Marco	University of Chicago	OSG site coordinator
Martin, David *	Argonne National Laboratory	Manager, User Services and Outreach
Mukkamala, Ravi	Old Dominion University	Professor
Navarro, John-Paul *	UChicago/Argonne	TeraGrid Director for Software Integration
Nguyen, Phong	Fermilab	Software Developer

Papaspyrou, Alexander *	TU Dortmund University	Researcher, Computer Science
Perez, Pedro	Universidad de los Andes	Electronic Engineer
Quock, Deborah	Argonne National Laboratory	Computer Scientist
Replogle, Joel	OGF	Executive Director
Riteau, Pierre	IRISA/University of Rennes 1	PhD Student
Sill, Alan	Texas Tech University	Senior Scientist
Tannenbaum, Todd *	Univ of Wisconsin-Madison	Researcher
Teregowda, Pradeep	The Pennsylvania State University	Graduate Student
Thapa, Suchandra	Computation Institute / Univ. of Chicago	Software Engineer
Timm, Steven *	Fermilab	Comp. Svc Spec III
Tiradani, Anthony	Fermilab	Computing Specialist
Von Laszewski, Gregor	Indiana University	Assit. Dir. Cloud Comp.
Wegrzyn, Charles	Twisted Storage, Inc	Software Engineer
Yu, Liang	National Center for SuperComputing Applications, UIUC	Postdoc
Zurawski, Jason *	Internet2	Network Software Engineer III & Research Liaison

(* = volunteer to Science Agency Uses workshop organizing committee)

Contact Information

Further questions regarding the outcome of the HPACG workshop can be directed to the following persons:

Joel Replogle
Executive Director
(765) 228-9068
replogle@ogf.org

Melissa Osner
Program Manager
osner@ogf.org

http://www.ogf.org

Project Co-Organizers

The project co-organizers for this workshop wee Joel Replogle (OGF), Ian Foster (Argonne National Laboratory and University of Chicago), and Alan Sill (Texas Tech University and OGF Vice President of Standards).

Ian Foster
Director, Computation Institute
Professor, Department of Computer Science
Argonne National Laboratory
9700 Cass Ave., Building 240
Argonne, IL 60439
(630) 252-4619
foster@mcs.anl.gov
http://www.mcs.anl.gov

Alan Sill
Senior Scientist, High Performance Computing Center
Drane 162, Texas Tech University
Lubbock, TX 79409-1167
(806) 742-4350
Alan.Sill@ttu.edu
http://www.hpcc.ttu.edu