

Grids and Clouds Activities in KISTI

**Soonwook Hwang, Sunil Ahn, Taesang Huh, Seok Kyoo Kim, Geun Chul Park,
Sehoon Lee, Sangsu Ryu and Gianni M. Ricciardi**

Korea Institute of Science and Technology Information

52-11Eoeun-Dong Yuseong-Gu Daejeon, Korea

E-mail:

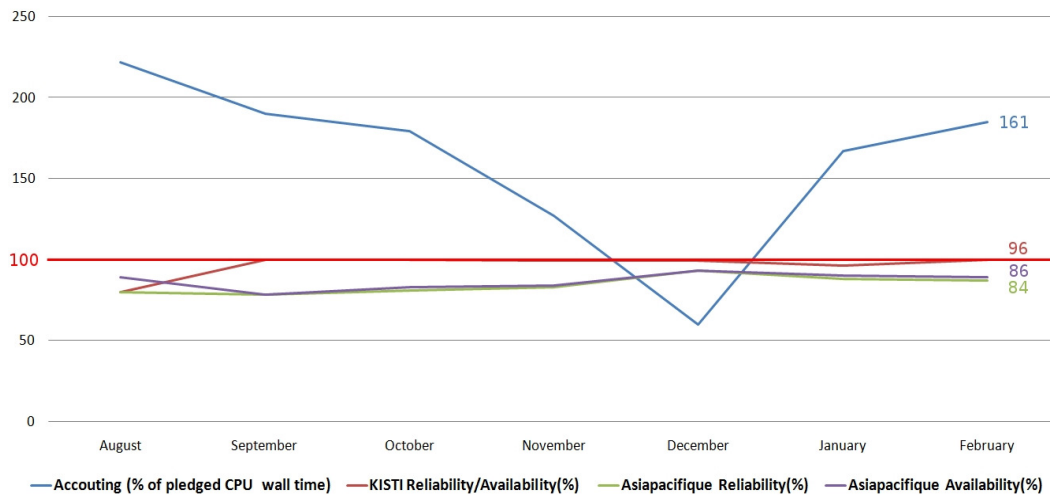
{hwang, sunil, tshuh, anemone, gcpark, sehooi, sangsuryu, gmr}@kisti.re.kr

KISTI was an official partner of EGEE, participating in many grid-related activities ranging from grid operation and maintenance to grid development activities. The EGEE project ended in April 2010 and some new Grid-related projects were launched. In a continuing effort to promote European Grid projects in Korea, KISTI has joined both the European Grid Initiative (EGI) and the European Middleware Initiative (EMI) as an official partner. Within EGI, KISTI is supposed to continue to provide a stable and reliable grid operation and production grid service as a member of Asia-Pacific Grid Initiative (APGI). KISTI has been part of the ALICE Computing Grid as a Tier2 since 2006, providing a stable and reliable node in the ALICE Grid. KISTI serves as one of the official EMI product teams, responsible for the maintenance and evolution of the AMGA software. Under the framework of EMI, in order to improve the usability of AMGA for user community, we have developed the AMGA GUI clients as an Eclipse RCP application released as part of the 2.1 version of AMGA in November 2010. We are now closely working with the other EMI product teams towards the release of EMI-1 scheduled for April 2011. KISTI has been keen on the seamless aggregation and use of heterogeneous distributed computing infrastructures (HDCI) for enabling large-scale science that would otherwise be impossible to address. To this end, we have developed some Ganga plug-ins for having access to different distributed infrastructures including Grids, Clouds, and Supercomputers as one of integral parts of our approach to HTC service.

*The International Symposium on Grids and Clouds and the Open Grid Forum
Academia Sinica, Taipei, Taiwan
March 19 - 25, 2011*

1. KISTI ALICE Tier2 Center

The last year of EGI-InSPIRE has been a good year for KISTI especially in terms of providing a reliable Grid infrastructure. KISTI signed a WLCG MoU agreement with CERN in 2007 to serve as an official Tier-2 for the ALICE experiment, pledging to provide a 150 KSI2K in CPU and 50 TBytes of disk storage dedicated to ALICE computing. Regarding CPU provisioning, for the past six months, KISTI has provided an average of 161% of what it was supposed to deliver in the WLCG MoU for 2010 (See the top blue line in Figure 1). In terms of disk storage delivery, we installed another 20 Tbytes of Disk last year so that we have at last fulfilled our pledge of the provision of 50 TBytes of disk storage dedicated to ALICE. In General, as far as grid operation and maintenance are concerned, KISTI has served as a stable and reliable site last year. For example, as can be seen in the red line in Figure 1, KISTI provided an average of 96% in service availability and reliability for the six-month period from August 2010 through February 2011 (See the red line in Figure 1), which we believe is quite good, compared to the average service reliability/availability for Asia Pacific region during the same period. It should be noted that the average service reliability and availability for Asia Pacific region during the same period is 84% and 86%, respectively (See the purple and green line in Figure 1)

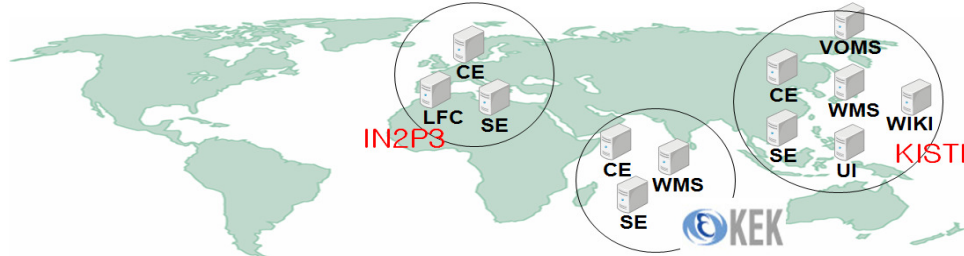


<Figure 1> KISTI Accounting/Reliability/Availability (From Aug. 2010 to Feb. 2011).

2. Production Grid Infrastructure: FKPPL VO

Another notable success achieved last year was the fact that we could observe the potential for expanding the FKPPL VO towards a France-Asia VO by having integrated some computing and disk resources of KEK in Japan last year. The FKPPL VO was established based on gLite middleware services across two computing centers, IN2P3 in France and KISTI in Korea at the end of the year of 2009, originally offering about 8,000 CPU cores and 3 TB disk storage to

scientists and researchers in the two countries. Last December, KEK joined the FKPPL VO, providing additional 1,600 CPUs and 27 TB of disk. With the KEK contribution, the FKPPL VO, as of now, consists of more than 10,000 CPUs and 30 TB of disk storage. Figure 2 shows grid services available in the three sites. We are now under discussion of integrating the IHEP of China as well by the end of this year.



<Figure 2> FKPPL VO Grid Infrastructure

For the long-standing operation and management of production grid infrastructures like FKPPL VO, it is crucial to develop applications fully exploiting the underlying infrastructures. Last year, we had a couple of applications successfully gridified, carrying out a large-scale deployment of the applications on top of the FKPPL VO. Among those applications are a Geant4 application and QCD (QuantumChromo Dynamcis) simulation from the displines of Bio-medical and theoretical Physics, respectivley. We ran a large-scale QCD simulation corresponding to approximately 75 CPU years on a single machine, in collaboration with the QCD research group of Sejong University in Korea. A preliminary result gained from the simulation was published in [1]. Table 1 shows how heavily the FKPPL VO was exploited last year with the deployment of QCD simulation and other applications. The FKPPL VO was ranked 16th among hundreds of existing VOs around the world in the list of the heaviest CPU consumption for the three-month period from September to November 2010.

	VO	Sep 2010	Oct 2010	Nov 2010	Total
1	atlas	196,073,904	274,240,936	242,671,204	712,986,044
2	cms	71,093,916	79,708,748	77,396,660	228,199,324
3	alice	61,897,556	78,409,636	64,368,804	204,675,996
4	lhcb	30,480,592	24,099,528	25,236,356	79,816,476
5	theophys	6,169,892	3,428,224	44,804,096	54,402,212
6	dzero	5,879,836	6,349,708	5,703,740	17,933,284
7	compchem	4,822,700	2,490,136	6,902,316	14,215,152
8	ilc	5,480,616	4,474,772	4,211,848	14,167,236
9	vo.cta.in2p3.fr	1,815,380	2,537,676	2,618,184	6,971,240
10	biomed	756,340	3,152,524	2,775,664	6,684,528
11	superbvo.org	4,676,060	699,836	0	5,375,896
12	auger	873,468	2,098,376	2,386,588	5,358,432
13	hone	750,552	3,118,336	1,268,064	5,136,952
14	pheno	1,312,896	1,708,028	1,845,340	4,866,264
15	icecube	1,287,804	1,897,840	1,177,468	4,363,112
16	fkppl.kisti.re.kr	918,352	2,022,848	1,209,720	4,150,920
17	see	1,044,224	1,033,832	2,008,676	4,086,732
18	cosmo	216,112	1,735,068	1,911,252	3,862,432
19	vo.lal.in2p3.fr	2,807,252	258,608	29,204	3,095,064
20	enmr.eu	980,180	807,260	1,011,816	2,799,256
	Total	411,611,532	506,320,224	502,216,812	1,420,148,568

<Table 1> Distribution of Normalized CPU times (HEPSPEC06) by VOs

3. Development of AMGA Metadata Catalog

We have developed and released the 2.1 version of AMGA [2] that supports both data federation and GUI client. The federation support gives AMGA users the ability to integrate distributed metadata. That is, it provides users with a single view to metadata that are actually dispersed geographically across multiple sites. It allows users to effectively handle partial failures with some user-defined policies. The AMGA Manager, an easy-to-use AMGA GUI toolkit, enables an interactive exploration on and searching for relevant metadata in a user-friendly manner, and hides the complexities of accessing Grid services. It allows users to easily manipulate metadata schema, entries, access control, user and group information through a GUI. It also supports some other easy-to-use features like automatic query composition, importing and exporting metadata into a spread sheet and filtered metadata searching.

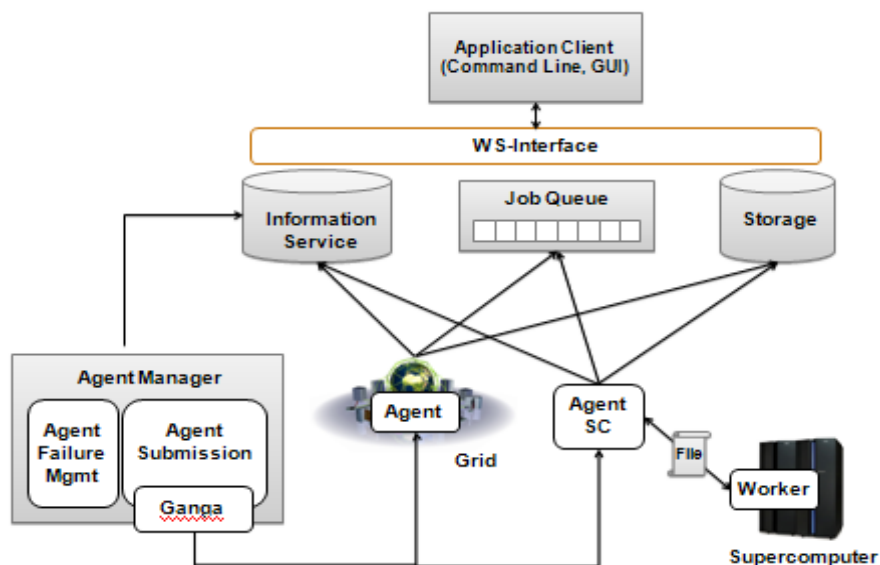
These AMGA new features were designed and implemented based on the requirements of diverse EGEE user communities collected over the past few years. The federation support is expected to greatly improve users' convenience and even address some security aspects of metadata due to the fact that less servers are required to keep metadata. In addition, it is possible to enhance performance through parallel metadata searching. The AMGA Manager helps users more easily and conveniently manipulate metadata so that it will boost productivity in managing large-sized metadata on the Grid. With these new features of easy-to-use GUI client and data federation, AMGA is expected to draw more attention to user communities in many different disciplines in EGI.

4. Development of HTC Service

We have developed a HTC problem solving environment for efficient large-scale distributed computing on heterogeneous infrastructures. It is intended to hide users from the complexity of using different resources and minimize the execution time of large-scale HTC problems. It uses a user-level job scheduling with a dynamic algorithm that automatically selects more responsive and effective resources based on the actual resource statistics collected during execution from agents launched on the resources. By providing fully automatic job submission and failure management, it reduces the burden of users from handling a large amount of jobs themselves. Internally, it provides a simple and uniform interface to different distributed infrastructures via Ganga, a tool for job submission and management on the Grid developed by CERN. Ganga shields the complexity of accessing different computing resources using its plugin modules that we have developed for access to distributed computing resources of different types like Amazon EC2 and the KISTI supercomputer. Figure 3 shows an overall system structure of the KISTI HTC service, where you could see grids and KISTI supercomputer being utilized as resources for the HTC service. We are currently working on the integration of Amazon EC2 resources as well.

The HTC Service focuses on how to effectively handle a large amount of jobs. As seen in Figure 3, a web service interface is provided for users to have access to the HTC service. Using

the web service interface, application and tool developers can easily develop high-level tools and applications tailored to the requirements of specific application domains.



<Figure 3> System Structure of HTC Service

5. User Support for Application Porting and Deployment

By the request of scientists in Chonnam National University, we conducted a large-scale deployment of in-silico docking on the Biomed VO to find potential drugs against SARS (Severe Acute Respiratory Syndrome). They came to us with a total of 1,100,000 chemical compounds (300,000 from Chambridge chemical database and 800,000 from ZINC database). We managed to mobilize up to 2,500 computers on the Biomed VO to screen the 1,100,000 compounds using Autodock, which took about three weeks to complete. Analysis was done by scientists in Chonnam National University and scientific results was published in [3].

Researchers in Sejong University requested some help to execute their two-color QCD simulation code with 1,647 different parameters, equivalent to about 75 CPU years. With access to up to 1,600 CPUs simultaneously on the FKPPL VO, we completed the whole calculation, which took about 3.5 months starting from the end of August until the 8th of December last year. Now it is under analysis by the QCD research group of Sejong University.

6. Outreach and Dissemination

Training is an integral part of Grid infrastructure projects like FKPPL VO. In order to promote the adoption of Grid technology and infrastructure, we organized Grid tutorials twice last year especially targeting the medical physics community in Korea. In February, we held a Geant4 and Grid tutorial 2010, where we had about 35 participants from major hospitals in Korea and about 20 users newly joined the FKPPL VO membership. In July, we organized the

2010 summer training course on Geant4, GATE and Grid computing where there were about 50 attendees from about 50 institutes in Korea (See Figure 4). For the two training events, we invited Geant4 and GATE experts from France and Japan. Students were given a chance to learn about grid computing including the basic concept of Grid computing, high-level grid tools, and porting of applications of their interest on the FKPPPL VO.



References

- [1] Soonwook Hwang, et. al, *Large-scale Deployment of two-color QCD simulations on the Grid using Ganga*, *Computational Physics Communications*, Vol. 182, No. 1, January 2011.
- [2] Sunil Ahn, et al, *Design of the Advanced Metadata Service System with AMGA for the Belle II Experiment*, *JKPS*, Vol. 57, No. 4, October 2010.
- [3] Thi Thanh Hanh, et al, *Virtual Screening identification of novel severe acute respiratory syndrome 3C-like protease inhibitors and in vitro confirmation*, BMCL-D-10-02200R1