OVERALL EXPERIENCE OF GRIDKA T1 OPERATIONS AND LHC EXPERIMENTS REPRESENTATION

A. Heiss, A. Petzold, M. Zvada
Karlsruhe Institute of Technology, Germany
{andreas.heiss,petzold,zvada}@kit.edu

The GridKa Computing Center at the Karlsruhe Institute of Technology is one of the biggest Tier-1 (T1) centers for the Worldwide LHC Computing Grid (WLCG) and one of the major resource providers in the EGI region NGI-DE. Since GridKa was established more than 10 years ago, GridKa staff has been closely cooperating with its different user communities also called virtual organizations (VOs) to ensure smooth operation and high availability of the resources and critical production services operated 24/7. This presentation will focus on overall experience how different experiments are represented in such a multi-VO computing infrastructure and what we are doing in order to keep GridKa running and serving T1 services successfully towards WLCG requirements.

1 The Grid and Steinbuch Center for Computing

Grid Computing is a key technology that provides scientists and engineers from research and industry worldwide distributed data and IT resources. Via high performance wide area networks grids integrate today's compute resources from desktops to cluster and visualization systems up to supercomputers. Also distributed data storage and archives from the terabyte to the multi-petabyte range as well as data of various types and disciplines. Transparent access from any location to such a distributed infrastructure allows to work on complex scientific and social problems and to collaborate in new interdisciplinary ways. Since many years the Steinbuch Center for Computing (SCC) department "Distributed Systems and Grid" (VSG from German term) contributes within national and international projects to the development and operations of this infrastructure, and works as a service and resource provider to the users [1].

Especially here in Europe, the European Grid Initiative (EGI) is mentioned, under whose umbrella national or international grid initiatives together form a common infrastructure [2]. The American equivalent to this is Open Science Grid (OSG), Asian countries and Australia are working together in the Asia Pacific Grid and Scandinavia has Nordic DataGrid Facility. None of these distributed infrastructures is complete itself rather than continuously evolving. However, despite the challenge of differences in the technologies, all resources to the user communities are available via standardized interfaces available worldwide.

1.1 VSG projects within SCC

The VSG department gives opportunity for diploma and doctoral thesis which are integral part of the grid and cloud computing evolution. In cooperation with the region Baden-Württemberg, VSG provides practical training of undergraduate students in the field of information technology. For example through participation in the winter semester 2011/12 the KIT first time offered lecture on "Distributed Systems: Grid and Cloud", and since 2003 has been hosting the annual Summer School for Grid and Cloud Computing [3].

In the context of the European Grid Infrastructure hierarchy, there are several Grid projects within the VSG department at SCC, for example:

- *European Grid-Initiative (EGI) / EGI-InSPIRE*: The European Commission in its 7th Framework Programme provided substantial support for a European Grid Initiative (EGI). Germany has responded to this initiative by the recently formed Gauss-Allianz e.V. consisting of 13 German supercomputing centers. The KIT is lead manager and was instrumental in the project planning of EGI-INSPIRE (Integrated Sustainable Pan-European Infrastructure for Researchers in Europe).

- *National Grid Initiative (Deutschland) Germany (NGI-DE)*: NGI-DE consist of about 50 grid initiatives, which are involved in EGI. Within the NGI-DE, VSG works in the areas of "Project Office and Sustainability", "Central Monitoring", "Helpdesk", "Security".
- *GGUS – Global Grid User Support System*: GGUS provides a central platform where users as administrators ask technical questions and address their issues accordingly to the Grid Site support unit. These support units are distributed teams of experts - administrators of local and global services through application supporters to grid middleware developers - connected to whom such requests of appropriate workflows are assigned. The GGUS platform is now available and used in several very large European IT projects.

2 The GridKa Tier-1 Center for LHC

GridKa – Grid Karlsruhe, project founded as part of the other SCC activities is driven by the need of thousands worldwide scientists use the computing and storage resources for data analysis of LHC experiments at CERN. There have been years involved into the grid development in addition to the construction of the LHC accelerator and the four detectors, ALICE, ATLAS, CMS and LHCb, where GridKa is now significantly important piece of the World Wide LHC Computing Grid (WLCG) [4]. Together with the German nuclear and elementary particle physicists SCC established Grid Computing Center Karlsruhe (GridKa) as the German contribution to the LHC Computing was developed specifically for high data throughput of the LHC and other high-energy physics experiments. WLCG ties together resources from the European Grid Initiative (EGI), Open Science Grid in the United States (OSG), and the Nordic DataGrid Facility in Scandinavia.

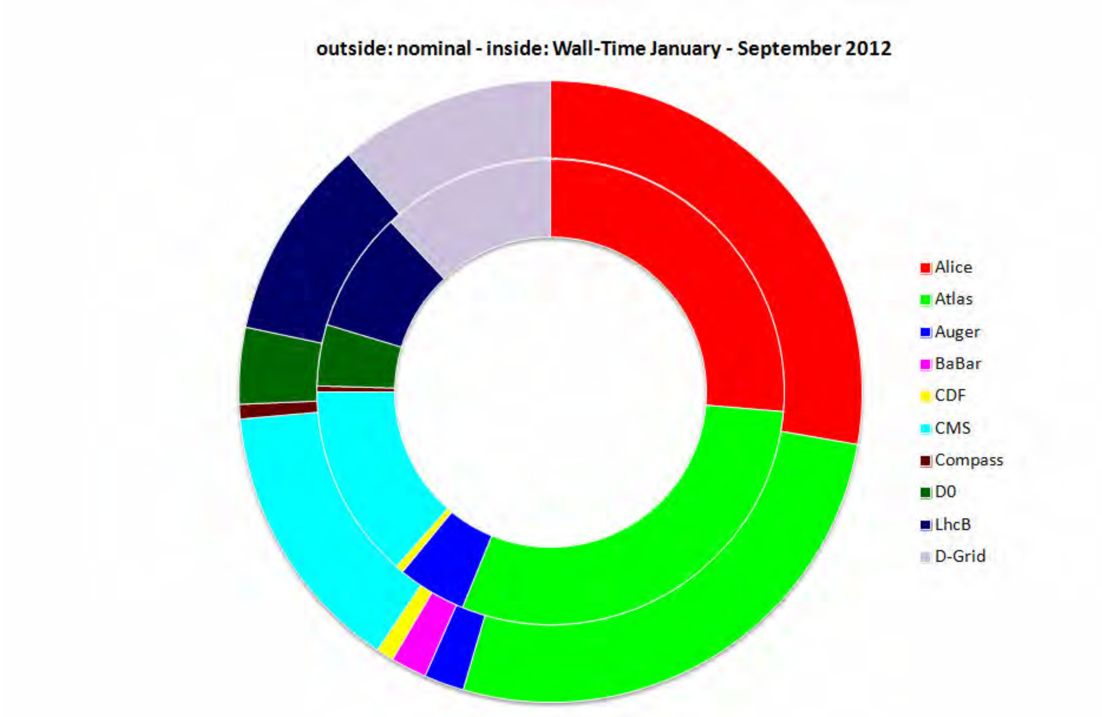

Fig. 1: GridKa computing resources share January-September 2012

In 2011, EGI sites provided the LHC VOs with more than 7 million kHS06hours [5] of CPU and more than 100 PB of online storage space. This is the largest single contribution of resources to

WLCG. The GridKa Tier-1 center is hosted by Steinbuch Centre for Computing at Karlsruhe Institute of Technology. GridKa was established in 2002 as a regional computing center for the LHC experiments which were still developing their computing models at that time and several other High Energy Physics experiments that were already taking data. Today, GridKa supports all four LHC experiment VOs and seven more VOs from High Energy Physics and Astroparticle Physics (Auger, BABAR, Belle, Belle2, CDF, Compass, DØ), and several other VOs from different fields of science. See Figure 1 for distribution of the resources from January till September 2012.

Among the 11 WLCG Tier-1 centers GridKa provides approximately 14% of the resources available to the experiments and it is the largest center supporting all four LHC VOs. In 2012, GridKa provides 130 kHS06 of CPU resources (more than 1200 compute nodes) split into two sub-clusters, 14PB of disk storage, and 17PB of tape storage to its users. For the regular archiving and reprocessing of the raw data of LHC, the experiment data flows via a 10 Gbit/s network connection imported from CERN while GridKa serves further connections to other Tier-1 and Tier-2 centers of the preprocessed data at Tier-1 level ready to be transferred. See Figure 2 with number running jobs per sub-cluster at GridKa over the year 2012.

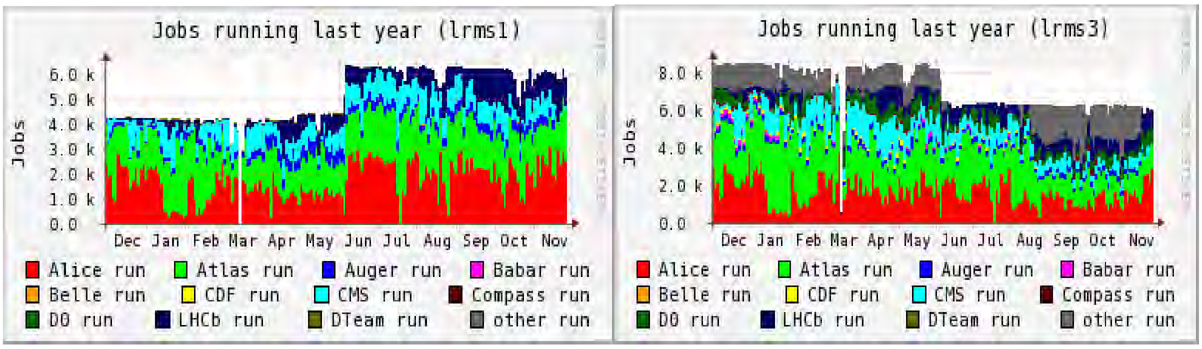

Fig. 2: Running jobs per VO in 2012 (~14k job slots in total)

3 Grid Core and Site Services

The amount of data from the LHC experiments with the latest algorithms and methods to (re-)processing is regularly ever-growing, so the computing power of GridKa increases yearly. For standardized measurement of computational power provided by all WLCG centers benchmark HEP-SPEC06 is used, where GridKa was deeply involved. These measurements helps in planning replacement an old and inefficient systems with new, more energy-efficient computers and thereby the overall computing power increases.

In addition to the resource management planning and accounting, GridKa provides high availability operations of grid services for the LHC and other user communities. It manages the Workload Management System (WMS) and compute element (CE), which forwards grid computing jobs to centers such GridKa.

Standardized Information Systems (Berkeley Database Information Index, BDII) deliver through Grid interfaces information about Grid computing centers, such as what virtual organizations are supported and whether free computing and storage resources are available. For data management there are file catalogs (LFC) and a file operated transfer service (FTS), which controls all data transfers between GridKa and its associated Tier 2 and Tier-1 centers. Different proxy services allow efficient access to database systems at SCC and at CERN. For community-specific services front-end VO-boxes are provided.

3.1 GridKA Certification Authority

The GridKa-CA [6] is present at KIT since 2002, as a member of the European Grid Policy Management Authority (EUGridPMA), an international organization, composed of over 40 countries.

Together with TAGPMA (North, Central and South America) and APGridPMA (Asia, Pacific, Australia) established a global network of the International Grid Trust Federation (IGTF).

Digital certificates contain a public key, additional information and the signature of the CA. The certificate body is accredited with signature of a public key belonging to a particular person. The certificates are used for authentication, encryption and decryption of sensitive data that is transmitted over the Internet and other networks. The GridKa-CA provides standard X.509 certificates to users, computers, and applications throughout Germany. The long list of organizations that use GridKa-CA certificates include both research institutions and universities, and industrial companies.

3.2 Continuous process optimization with ITIL

The coordination of global support needs well coordinated processes and workflows. The SCC has faced with a similar challenge. In merging the former Center of the University of Karlsruhe and the Institute for Scientific Computing of the Karlsruhe Research Center for SCC was clear that a data center in two locations requires special measures and coordination.

Therefore the decision was taken in the SCC within the rules of ITIL lifecycle stages "Service Strategy", "Service Design", "Service Transition" and "Service Operation" applying for continuous process optimization. An SCC internal ITIL Project was led by the department "ISM - IT Security and Service Management" initiative, which was also supported by the department of VSG. The focus in the grid environment was on the service process "Service Catalogue Management", the creation of service descriptions and the separation of grid services into so-called service modules and the service process "Incident Management", in particular towards the smooth 24/7 operation mode of GridKa [7].

To achieve improvements in terms of grid operation and support service in Germany "Incident Management" isn't enough, so keep maintenance of the German Tier-1 data center GridKa ITIL compliant, it requires in addition "Change Management" process. That is still work in progress though and on its way to the ITIL regulation fulfillment by SCC.

3.3 GridKa monitoring and 24/7 support

To monitor the function of the various grid services, VSG developed jointly with other research institutions monitoring system, which are largely based on the open source Nagios [8]. Nagios has been specially adapted to the needs of grid computing. However, as part of the modernization process and monitoring quality improvement, GridKa Tier-1 migrated all Nagios-based service to the Icinga [9], which is now used by the GridKa on-call engineer(s) (OCE) as a basic tool to solve problems or alarm local experts take an appropriate action.

At GridKa, we have two groups of OCEs: storage and grid services. Each group consist of 8-10 people and they are rotating on weekly basis. In addition, KIT has wide area network and infrastructure on-call service. Alarms are sent to mobile phones by Icinga plus GGUS alarm tickets are triggered by the incident thresholds configured within Icinga per service. Person on-call may or may not be an expert for the affected system (no experts expected on duty!). OCEs usually do their best try 'standard recipes' to fix the problem first and only if necessary they try to reach an expert.

Approximately 85% of problems could be solved without calling (additional) experts. Documentation is improved continuously, also people on-call gain experience by doing regular steps provided by 'recipes'. On-call services does not guarantee that problems can be fixed within few hours. The incident handling during on-call service hours is well illustrated on the Figure 3.

3.4 GGUS - Global Grid User Support System

Global Grid User Support [10], abbreviated GGUS, means globally distributed support system, which is managed and operated by a working group within the department of VSG. Scientists of international grid projects (e.g. EGI) or members of so-called virtual organizations such as the High Energy Physics (LHC experiments at CERN) use this system to report malfunctions related to the computing grid. It is also knowledge base for tips and tricks among site experts stored in the ticket database or place for providing important news.

Heart of GGUS is a complex web application that by providing web service interfaces allows the integration of other helpdesk systems to form in its principle knowledge federation. Currently 15 external systems are linked via these interfaces with GGUS. To date, over 100 expert groups are spread over all the continents of the world, which provide round the clock support for the continuously growing community of grid users. By now approximately 80,000 Grid problems solved by 1,500 grid experts worldwide.

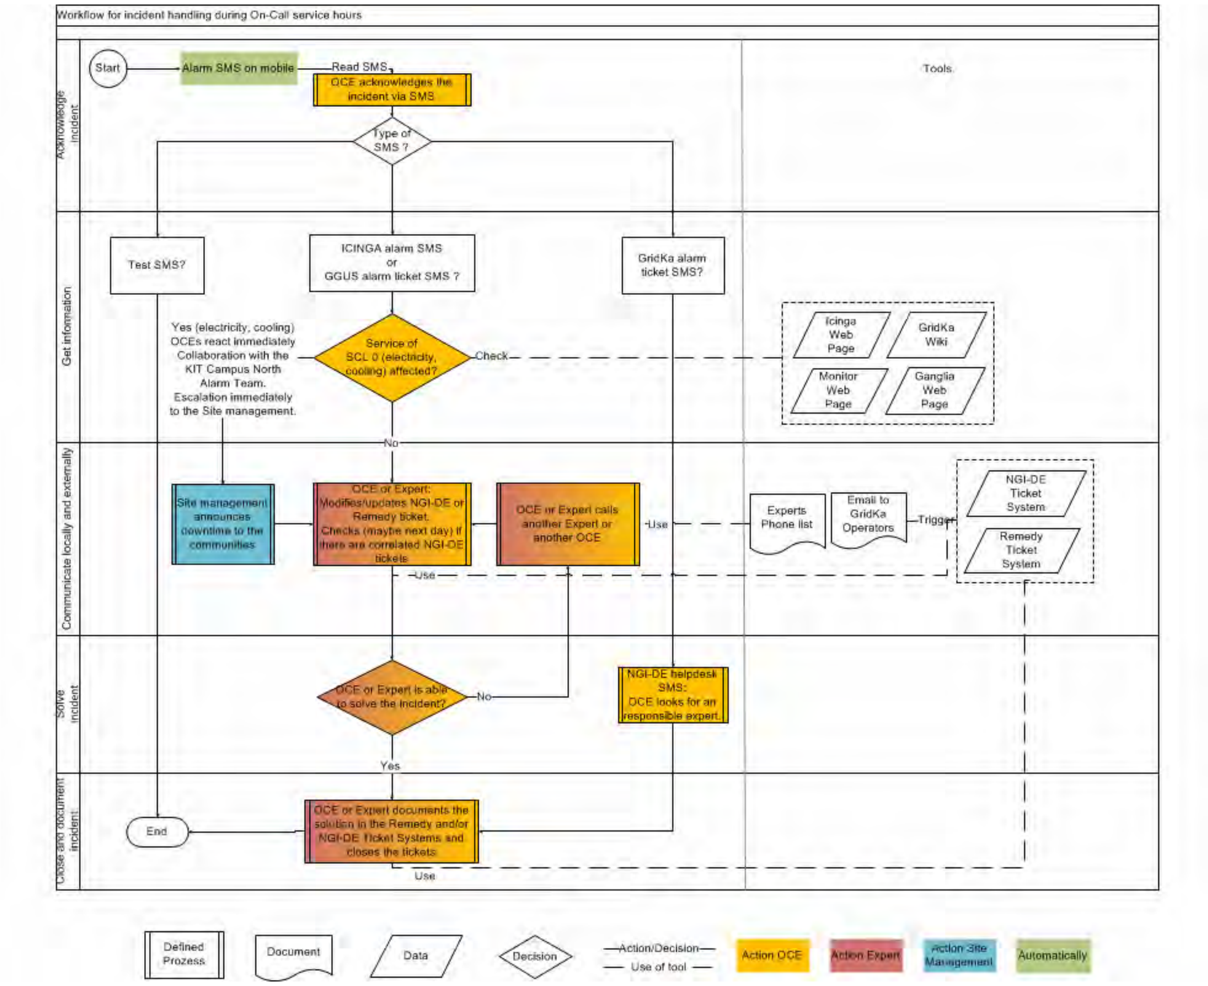

Fig. 3: Workflow for the incident handling during on-call service hours

4 LHC representation at GridKa

Several formalized support workflows, all based on ticket systems, have been set up within the LHC experiments, WLCG, and the major Tier centers. Experiment representation complements these workflows with experiment specific and proactive support.

In 2010, three positions of local VO representatives for ALICE, ATLAS, and CMS experiments were established. The representatives' duties are to represent both their respective LHC experiment at GridKa and GridKa within their respective LHC experiment [11].

4.1 Motivation and representatives' tasks

GridKa staff has been closely cooperating with its different user communities to ensure smooth operation and high availability of the resources and services. Even though all LHC VOs have used gLite-based services already for several years, there are major differences in the use of the components by the different experiment frameworks. These differences have a strong impact on the operations of the Tier-1 center.

The experiment representatives are fully integrated with GridKa staff and their experiments; within the latter, they mainly communicate with the computing experts, as the servers running at the Tier-1 are higher level ones. Both the experiments and GridKa benefit from the experiment representatives' inside knowledge of the details of the experiment specific computing and the on-site operations at the Tier-1.

The experiment contacts are responsible for several production critical services their respective experiment relies on at the Tier-1. These include site services like FTS (mainly used by ATLAS and CMS) and xrootd (used by ALICE) as well as other distributed experiment specific services.

The major task besides service administration is communication between the experiment and the site administrators. The experiment representatives serve as principal point of contact for the site administrators for all matters related to the VO, ranging from job behavior to storage system setup. At the same time the experiment contacts communicate with all parties within the experiment, e.g. Tier-0, Tier-1, and Tier-2 centers, users, and in particular experiment computing experts.

The close coordination among the experiment representatives has proved to be very beneficial for their daily work as well as intermediate and long term tasks. As the experiments share the Tier-1 infrastructure and many services, many issues affect several VOs. For instance, network problems are often discovered because of middleware service failures. The experiment representatives inform each other of these problems and together analyze the impact on the experiments and provide specific feedback to the local network administrators. In other cases, experiences of one VO with the effect of a VO workflow on the storage backend are shared among the on-site experiment contacts for the benefit of all VOs.

As part of the team of local administrators, the experiment representatives have privileges on the local infrastructure that would not have been granted to external representatives. Thus, they are enabled to quickly diagnose problems and assess the impact on the experiments' work much better than compared to a long chain of communication from externals to the local administrators and back. In addition, on-site representatives are perceived as internal members of the respective team by both the experiments and the site administrators. Thus, communication between VOs and sites has been improved on formal and informal levels.

The experiment representatives also contribute to organization of grid computing events as well as to the deployment and operation of services for the experiments. Also in this area, their work profits from their unique perspective which combines the view of the experiments and the Tier-1 center.

Experiment workshops like a collaboration-wide ALICE Tier-1/2 workshop or a face-to-face meeting of the grid administrators of computing centers of ATLAS DE cloud were organized by the experiment representatives at GridKa. In addition they make major contributions to the international GridKa summer school on grid and cloud computing.

Furthermore, experiment representatives are heavily involved in the deployment and operation of the glideInWMS [12] service for CMS and to the multi-VO meta-monitoring system HappyFace [13].

4.2 Feedback to on-site Representation Model

There has been positive feedback from the VOs represented at GridKa in several meetings and workshops, focusing on faster feedback and on better integration of GridKa into the experiments' computing groups. Service experts at GridKa highly appreciate to have expertise on the specific experiment workflows on-site and the easier and faster communication.

Contrary to formalized support via ticket systems, there are no obvious metrics for measuring the success of on-site experiment representation, as it uses a short and informal way of communication and as it provides proactive support. Its effectiveness is nevertheless reflected by the positive feedback mentioned above and the ongoing commitment by the experiments and the computing center to jointly co fund the positions.

5 Summary and Conclusion

In the long term future, we expect changes in the usage of the Tier-1 center both with respect to new user communities and with respect to new infrastructure as a service (IaaS) technology. If the current model of on-site VO representatives is still useful for computing centers that in the future will mainly work as IaaS providers depends largely on VO specific requirements which cannot be fulfilled with generic IaaS cloud resources. As long as VO specific knowledge is required to efficiently provide a service to the VO, it might be useful to have a VO representative integrated into the team of site administrators.

The GridKa Tier-1 center at its large profits very much from the on-site experiment representatives. Their work has resulted in improved availability and reliability of site services and in improved communication between the computing center and the experiments. In particular, the unique perspective offered by the integration in both the on-site team and the experiment enables the representatives to proactively address issues on most areas of site operations related to experiment work and even bring the new projects to participate on internally with significant outside GridKa visibility.

References

[1] SCC KIT News. Web site: http://www.scc.kit.edu/downloads/oko/SCC-News_2011_3_web.pdf
[2] European Grid Infrastructure. Web site: http://www.egi.edu
[3] GridKa School. Web site: http://gridka-school.scc.kit.edu
[4] Worldwide LHC Computing Grid. Web site: http://www.cern.ch/lcg
[5] HEPSPEC06. Web site: http://w3.hepix.org/benchmarks/doku.php
[6] GridKA-CA. Web site: http://www.gridka.de/ca
[7] ITIL and Grid services at GridKa: http://iopscience.iop.org/1742-6596/219/6/062018
[8] Nagios. Web site: http://www.nagios.org
[9] Icinga. Web site: http://www.icinga.org
[10] Global Grid User Support (GGUS). Web site: http://www.ggus.org/
[11] Experiment Representation at the WLCG Tier-1 Center GridKa. *PoS EGICF12-EMITC2*: http://pos.sissa.it/archive/conferences/162/032/EGICF12-EMITC2_032.pdf
[12] glideInWMS Project. Web site: http://www.uscms.org/SoftwareComputing/Grid/WMS/glideinWMS/doc.prd/index.html
[13] HappyFace Project. Web site: https://ekptrac.physik.uni-karlsruhe.de/trac/HappyFace