

Experiment Representation at the WLCG Tier-1 Center GridKa

Christopher Jung

Karlsruhe Institute of Technology

E-mail: christopher.jung@kit.edu

Andreas Petzold*

Karlsruhe Institute of Technology

E-mail: andreas.petzold@kit.edu

Marian Zvada

Karlsruhe Institute of Technology

E-mail: marian.zvada@kit.edu

The GridKa Computing Center at the Karlsruhe Institute of Technology is one of the biggest Tier-1 centers for the Worldwide LHC Computing Grid (WLCG) and one of the major resource providers in the National Grid Initiative of Germany (NGI-DE).

In 2010, three positions of local Virtual Organization (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. This presentation will focus on the representatives' technical and communication tasks and their experiences gained.

*EGI Community Forum 2012 / EMI Second Technical Conference,
26-30 March, 2012
Munich, Germany*

*Speaker.

1. The GridKa Tier-1 Center in WLCG and EGI

The WLCG is the distributed computing infrastructure for the LHC experiments ALICE, ATLAS, CMS and LHCb. It ties together resources from the European Grid Initiative (EGI), Open Science Grid in the United States, and the Nordic DataGrid Facility in Scandinavia. Each LHC experiment is represented by one VO, which bears the same name as its experiment.

In 2011, EGI sites provided the LHC VOs with more than 7 million kHS06hours [1] of CPU and more than 100PB 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.

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 130kHS06 of CPU resources, 14PB of disk storage, and 17PB of tape storage to its users.

2. Experiment 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, strengthening the collaboration of GridKa and the LHC experiments. It has been set up for the experiments which consume large resources and use complex services.

2.1 Motivation

Since GridKa was established more than 10 years ago, 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 File Transfer Service (FTS) is used by ATLAS, CMS, and LHCb but not by ALICE. LHCb and ATLAS both use local Oracle databases to store conditions data. While ATLAS utilizes FroNTier [2] to access the database, LHCb uses direct database access.

The ATLAS, CMS, and LHCb experiments use dCache as storage middleware, but have very different concepts for local data management. In particular, disk-only storage space and tape-backed storage space serve different use cases for each experiment. For instance, CMS utilizes dCache mainly as a hierarchical storage system, with the disk-only space used for intermediate storage.

On the other hand, ATLAS relies mainly on disk-only space for permanent storage with the tape-backed storage used only for archival of custodial data. ALICE is the only experiment to use xrootd as a storage middleware.

All VOs have built different multi-layer frameworks based on different middleware components for data management and distribution as well as for workload management. These different systems return high-level error messages which often do not allow the local operators to easily infer the problem with the underlying middleware components. In the past, these different view points have complicated communication with off-site experiment experts, especially since each VO considers different services to be critical.

Based on the different and complex setups for the LHC VOs and the resulting impact on Tier-1 operations, it was decided to establish positions for experiment representatives at the Tier-1 center, in order to achieve the best possible performance and service availability for each VO.

2.2 Representatives' Tasks

Three full-time positions for experiment representatives for ALICE, ATLAS, and CMS were established at the Tier-1 in 2010. 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 [3] service for CMS and to the multi-VO meta-monitoring system HappyFace [4].

2.3 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 cofund the positions.

2.4 Outlook for on-site VO Representation

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.

Whether or not establishing an on-site representative of a new VO at the Tier-1 warrants the investment, largely depends on the complexity of workflows used by the VO and if special dedicated services have to be run for the VO. The size of the VO is usually not a decisive factor. For instance, while LHCb is a large VO, their usage pattern of the Tier-1 resources does not require dedicated on-site support.

The current positions have been established at the request of the VOs, however the computing center or the VO can propose to establish a position for an on-site representative.

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.

Past experience with the grid middleware show that a "one size fits all" approach did not fit the requirements of the user communities from different fields of science. If these user communities will adopt the cloud services without special site requirements, thus rendering the dedicated VO support obsolete, remains to be seen.

3. Summary and Conclusions

The GridKa Tier-1 center 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.

It can be concluded that major computing centers and communities can profit from on-site representation of the user communities. This results in a better utilization of the site's potential and better mutual understanding between the user communities and the site administrators. The profit depends on the amount of resources consumed by the respective community and the complexity of its services.

The future of dedicated on-site support to VOs depends on the technologies that will be used to access the site resources. As long as complex community specific tools are used, on-site VO representatives can play an important part in helping to optimize the daily operations of the VO at the computing center.

References

- [1] <http://w3.hepiv.org/benchmarks/doku.php>
- [2] S. Kosyakov, J. Kowalkowski, D. Litvintsev, L. Lueking, M. Paterno, S. P. White, L. Autio and B. J. Blumenfeld *et al.*, *FroNTier: High Performance Database Access Using Standard Web Components in a Scalable Multi-Tier Architecture*, FERMILAB-CONF-04-367-CD.
- [3] <http://www.uscms.org/SoftwareComputing/Grid/WMS/glideinWMS/doc.prd/index.html>
- [4] <https://ekptrac.physik.uni-karlsruhe.de/trac/HappyFace>