SUPPORT FOR THE CMS EXPERIMENT AT THE TIER-1 CENTER IN GERMANY

J. Berger, C. Böser, T. Chwalek, M. Fischer, O. Oberst, G. Quast, N. Ratnikova[1], S. Röcker, S. Wayand, M. Zeise, M. Zvada

KIT - Karlsruhe Institute of Technology, KIT, Germany

[1] *also at ITEP - Institute of Theoretical and Experimental Physics, Russia*

Natalia.Ratnikova@kit.edu

The Compact Muon Solenoid (CMS) experiment is one of the two general purpose detectors operated at the Large Hadron Collider (LHC) at CERN. The CMS Collaboration includes 191 institute from 42 countries of the world. The Institut für Experimentelle Kernphysik (IEKP) at the Karlsruhe Institute of Technology (KIT) has a strong participation in the CMS programs in physics, statistics, detector study, software and computing. One of the major tasks is the support for the CMS Tier-1 Center at GridKa, KIT, including administration and operations of the experiment specific layer of the distributed computing infrastructure. We discuss the overall organization of support for the CMS activities at GridKa, the infrastructure and the tools used by the site administrators, and coordination with the GridKa and the CMS Central Computing Operations teams. We give an overview of related development projects conducted in this context. Finally, we draw some conclusions based on our experience of operating the Tier-1 for the CMS experiment during the active data taking period.

1. Introduction

The contribution of Germany to the Worldwide LHC Computing Grid (WLCG) [1] consists of a large Tier-1 centre, GridKa at the Karlsruhe Institute of Technology (KIT) and 7 Tier-2 centres at other national institutions or German universities. GridKa provides services to nine high energy physics experiments, among them the four LHC experiments ATLAS, Alice, CMS and LHCb. GridKa operates very closely also with many other Tier2 centres, mostly in eastern Europe. Thus, GridKa is at the heart of a very complex grid structure, with different computing models and workflows adopted by each of the supported Virtual Organisations (VOs). A clear separation line of responsibilities is defined at GridKa, including a formal service level agreement with each experiment. GridKa staff ensures the availability of the resources and is responsible for Grid middleware services. It is impossible for GridKa to provide experts for each of the experimental software set-ups and the specific workflows. Some operations require authorisation, which can only be granted to the members of the experiment. Therefore the responsibility for the experiment-specific tasks belongs to the group of experts within the experiment. Overall experience of GridKa T1 operations and LHC experiments representation are discussed in [2]. This paper highlights the support structures set up in Karlsruhe for the CMS experiment [3].

The Institut für Experimentelle Kernphysik (IEKP) [4] at the Karlsruhe Institute of Technology (KIT) is the fifth largest by membership from 192 institutes in the CMS Collaboration. IEKP members have strong participation in CMS physics data analysis, statistics, detector study, software, computing, and upgrade programs.

Our major task in computing is to provide smooth operation and efficient resource management at the CMS Tier-1 Center at GridKa. This includes configuration and administration of the experiment specific infrastructure at Tier-1, site monitoring, data management, and coordination activities.

Besides Tier-1 support, our institute members play leading roles in several other projects: enabling infrastructure and environment for the local CMS physicists to access and successfully use additional national resources at GridKa (NRG); set up infrastructure for CMS central workload management system based on GlideIn technology (GlideInWMS); development and validation of Physics Experiment Data Export (PhEDEx) project: CMS data consistency tool, validation tool suite, storage accounting tool; meta-monitoring tool HappyFace. The majority of these tasks are credited as IEKP contribution to CMS central computing services. On the other hand they help our group to build the expertise necessary for operating the complex CMS environment and tools.

In the following sections we present our experience in providing support for CMS Tier-1. Section 7 is devoted entirely to the development projects.

2. Resources

The CMS Offline computing system [5] consists of a large number of geographically distributed centers organised in an hierarchical structure of computing tiers. A single Tier-0 center at CERN accepts data from the CMS Online Data Acquisition System for archival storage, calibration and prompt reconstruction. Tier-0 distributes raw and processed data to a set of large Tier-1 centers for secondary archival storage, organized data processing and data serving to a more numerous set of smaller Tier-2 centers. Tier-2 centers provide resources for analysis, calibration and Monte Carlo simulation activities. The new data produced at Tier-2s are transferred to Tier-1s for custodial storage. Tier-3 centers provide interactive resources for local groups and additional best-effort computing capacity for the collaboration.

The CMS Computing model [6] specifies nominal Tier-1 Resource Requirements in terms of CPU, Disk and Mass Storage capacity, data rate from storage, WAN transfer capacity, CPU node I/O bandwidth. The actual Tier-1 resource contributions and expected service levels are agreed in the context of the WLCG MoU [7]. Accounting of the CMS computing resources is available via REBUS [8], the Resource, Balance and Usage website for the whole of WLCG, including topology information, resource pledges, and installed capacities. Resources pledged by all CMS Tier-1 sites and by the CMS Tier-1 at Gridka in the year 2012 are presented in table 1.

Table 1. CMS Tier-1 resource pledges in 2012

Resource	CMS T1 total	GridKa CMS T1 contribution
CPU power	145 kHS06	3’750.0 kSI2k, 1’456 job slots
Disk space	22 PB	1’950.0 TB
Tape storage	45 PB	5’000 TB
Local Support	217 FTE months of credited service work	31 FTE months of credited service work

In addition GridKa provides to CMS 175,0 TB local and 75,0 TB WAN storage, 10 Gbps national regional network and 10 Gbps OPN international connection speed.

3. Site Performance Metrics

Due to the importance of the Tier-1 services a high level of availability and operability of the site is expected. CMS computing has developed a so-called Site Readiness metric, which takes into account various criteria, such as the level of success of the test jobs sent to the site or the results of data transfer load-tests continuously running between the sites. Unavailability during scheduled down time periods is also properly accounted for. Figure 1 shows CMS Tier-1 Site Readiness status, with a

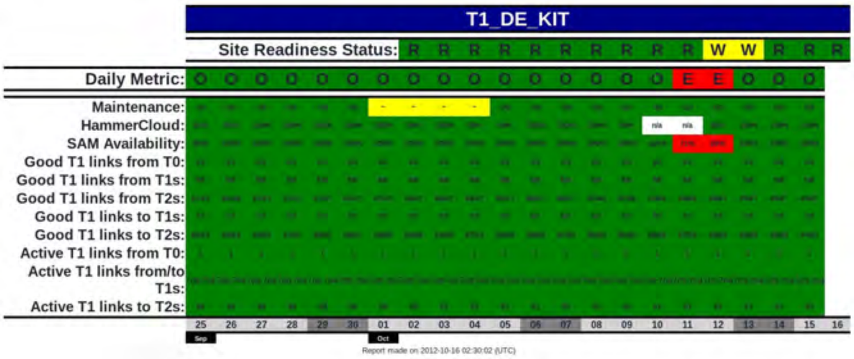

Figure 1: Site Readiness Status

short period of reduced availability below 90 percent which resulted in Warning status. The metric is updated and published daily. A cool-down time is automatically allowed for the site to recover after prolonged failures.

The failure of the automated test job is not necessary indicating the site unavailability. For example, on the multi-VO sites the system may no longer accept new jobs, if the virtual organisation has reached its fair-share limit. Carefully planned scheduling by the central operations helps to avoid this situation. Another example is the time when the system works under heavy load. Test jobs sent to the production system are competing for the resource with the real production jobs. The processing of the test jobs is slowed down, and the system appears to be unresponsive.

The monitoring of the site performance is a steadily ongoing effort [9]. CMS Computing community is constantly working to improve the monitoring and the overall efficiency of the CMS workflows.

Local site administrators and a worldwide distributed crew of central computing shifts are watching the system around the clock and apply monitoring and alarming procedures to the sites, central services, and operators.

4. Local Site Support Operations

Grid sites provide resources and basic Grid services and components, such as security, computing and storage elements, monitoring, accounting, workload management systems, information service and low-level file transfer service. The CMS Central Computing Operations team manages the large-scale distributed workflows over the Grid sites. These are two very different scopes. The ultimate goal of our local CMS support group is to match the resources and the workflows for smooth and efficient operation. This requires expertise in both areas.

Managing data [10] is by far the most labor-intensive task in Tier-1 support. The CMS distributed data transfer system [11] is the key ingredient which enables the optimal usage of all distributed resources. Data transfers from Tier-0 to Tier-1 sites must be done in a timely manner to avoid the overflow of the disk buffers at CERN. Simultaneously, the data are transferred in bursts to Tier- 2 level sites for analysis, and simulated Monte Carlo data produced at Tier-2 centers are moved to Tier-1 sites for archival. Additionally, data may be synchronized between different Tier-1 sites, and served to Tier-3 sites.

CMS uses PhEDEx tool [12] to initiate transfers and to keep track of data placement and transfer status.PhEDEx provides subscriptions and requests mechanisms to handle data operations. Every incoming data transfer or deletion request must be manually approved by the local data manager. Routine data operations include clean-up of obsolete data, maintaining consistency between the storage contents and the central data catalogues, monitoring and debugging transfer issues, attending meetings, providing prompt feedback on requests, managing PhEDEx software installation, configuration and upgrades, restarting PhEDEx agents, inspecting the log files, and providing support to the associated Tier-2 sites.

Another major task is support for data processing at the site. The job submissions to Tier-1 sites are managed by the CMS central operations. However a few actions for proper data handling are required from the site. For compact tape utilisation, sites are asked beforehand to create the tape families for the files that are to be produced and archived on tape. Site may be asked to pre-stage from tape input data required for re-processing or replicate files needed by many jobs. The output of the application jobs and the corresponding log files are first written to the local disk on the worker node. If necessary, the empowered on-site expert may login directly to the worker node, inspect the log files and troubleshoot any site specific problems. After job execution, the application output is staged out into the CMS namespace on the storage element. An additional merging step is applied to the small files before archival to tape in order to reduce the load on the mass storage system. Once the merging step is complete, the original unmerged output files can be removed. Sites are responsible for regular clean-up of those obsolete unmerged files. The CPU efficiency of the jobs is constantly monitored to identify any jobs with a low CPU to wall-clock-time ratio. This condition usually indicates a problem with data access, when jobs stay idle waiting for the required data. Monitoring of the total numbers of running and queued jobs per experiment helps to distinguish general infrastructure problems from the application specific ones.

Site support services and procedures are part of the Karlsruhe CMS group obligations to the Collaboration formalised in the Maintenance and Operations plans, and are credited with 31 months of service work per a calendar year.

To improve the quality of support and to ensure proper share of the expertise, we have introduced an expert rotation scheme. For the period of six to eight weeks a senior member of the group takes care of the routine operations, *i.e.* responding to service tickets, handling tape families and data transfers, troubleshooting, reporting, etc. Data consistency checking, software upgrades, site configuration, system tuning and optimisation, and other specific tasks are covered by the dedicated experts. The rest of the group, including junior members, provides support by taking local site monitoring shifts, and by participating in various development projects.

5. National Resources at GridKa - NRG

Tier-1 centers provide both regional and global services. The usage by the local community however should not interfere with the ability of the Tier-1 center to fulfil its obligations towards the whole CMS. The German CMS community (DCMS) is able to use a share of 1366 cores of the current KIT Tier 1 resources. CMS users which registered for the dcms group in the CMS VOMS server are able to create a VOMS proxy with a *cms:/cms/dcms* VOMS extension. These users are mapped at KIT to one of the dcms pool accounts which possess a common

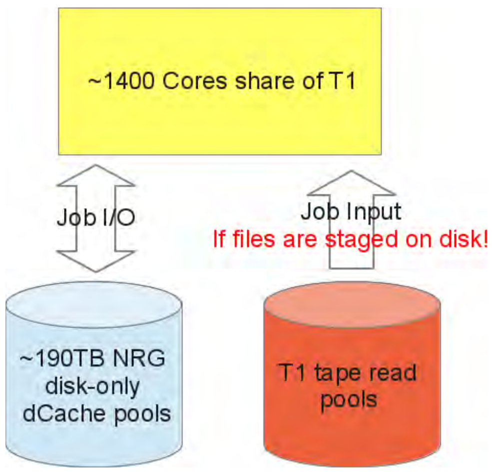

Figure 2: Sketch of the setup providing national resources at GridKa (NRG). Access to CMS T1 datasets is granted, if the required files are staged on disk

DCMS group fair share within the batch system accounting. To further limit the impact of the NRG users on the Tier-1 production operation, the maximum number of concurrently running DCMS pool account jobs is limited to 2000. Unused CPU resources within the NRG budget are available to the WLCG VOs.

Additionally to the computing resources DCMS users have read and write access to 190 TB of disk-only dCache storage via a specific storage element. This storage is foreseen as scratch space for users. Furthermore, access to CMS Tier-1 data is allowed as soon as the required official CMS datasets are staged on disk. A sketch of this setup is given in Fig. 2. Together the DCMS resources at KIT can be seen as a Tier-3 at the Tier-1, named "National Resources at GridKa" (NRG). As the whole setup is directly integrated within the T1 resources, no additional effort is needed to maintain the CMS software environment which the NRG users might need.

6. Coordination Activities

To insure an efficient use of Tier-1 resources and coherent and friendly computing environment for CMS users, a smooth cooperation of all involved parties is necessary.

The aim of the coordination efforts of our Tier-1 local support group is to maintain close contact both with CMS central computing operations and GridKa support teams, and also to stay alert to individual users' requests.

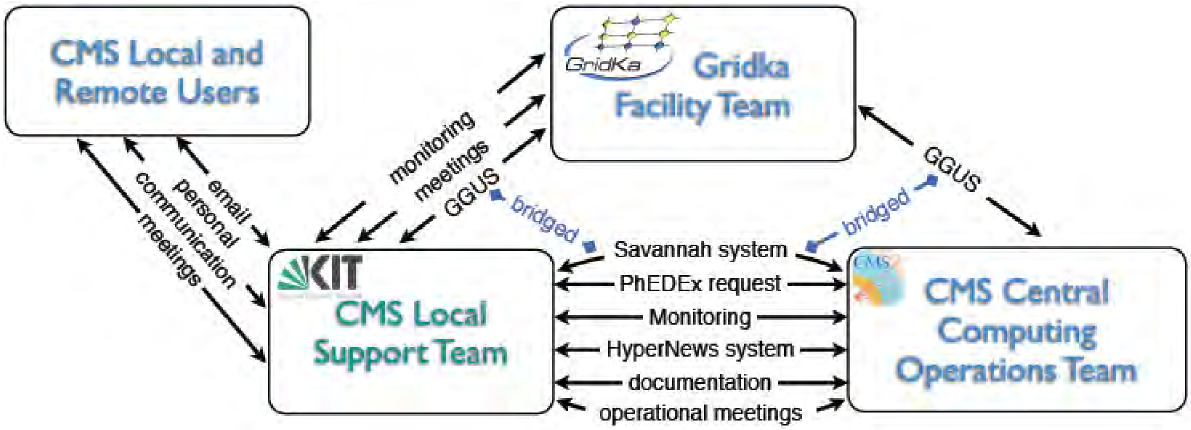

Figure 3: Involved parties and communication channels

Different channels have been established for the exchange of different types of information. Figure 3 illustrates communication channels between the involved parties. As one can see in the chart, CMS users do not communicate directly with the GridKa facility team. CMS Central operations may contact GridKa directly via a bridging mechanism between Savannah [13] issue tracking system used for internal CMS communications, and the WLCG problem tracking system GGUS [14]. According to CMS policies, bridging to GGUS is only used for well identified problems related to WLCG critical services at Tier-1 sites.

7. Development Projects

Senior experts in our group have extensive experience in the development of computer systems and continue to contribute to development of CMS computing infrastructure and tools. Playing a leading role in many CMS computing projects, they also involve students and young scientists in this work. The immediate benefits are the build-up of the group expertise, training, exchange of experiences, and a team-work spirit.

Development work is often conducted in collaboration with other groups, institutes and experiments. The results are regularly presented at various computing conferences and workshops, included in master or diploma and PhD theses, and published in scientific journals.

Most of this work is credited as CMS central computing services. This fits nicely with the CMS intention to extensively involve the youth in the technical aspects of scientific work. There is a general CMS requirement for every new member of the Collaboration to provide six months of service work to become CMS author.

In this section we give an overview of the most prominent development projects conducted at IEKP.

7.1. GlideinWMS

In order to facilitate an efficient management of the organised data reconstruction and Monte Carlo production, CMS uses the grid meta-scheduling system GlideinWMS [15]. The GlideinWMS workload management system is based on a virtual private Condor [16] pool. It is composed of several elements [17], and some of them can be multiplied for improved scalability.

GlideinWMS components and its flow of processes are shown in Figure 4. The operation logic [17], [18] of this system tries to maximise the amount of user jobs, while minimising the amount of wasted resources. It does this by keeping a steady pressure on the Grid pools; as long as there are jobs in the schedd queues that could potentially run on a Grid pool, a fixed number of pilot jobs is being kept in that pool queue. However, as soon as there are no more suitable jobs waiting in any of the schedds, no more pilot jobs are submitted. If any Glideins start after all the suitable user jobs have started, the Glidein itself will exit within a few minutes.

The pilot jobs are being submitted by the Glidein Factories, but it the job of the VO Frontends to decide how many pilot jobs to keep in each Grid pool. This number is calculated by matching the attributes of the user jobs, provided by the schedds, to the attributes of the Grid pools, provided by the Glidein Factories. If the number of matches is higher than the desired pressure, the Glidein Factories are told to keep the pressure. Else, the pressure is reduced as appropriate.

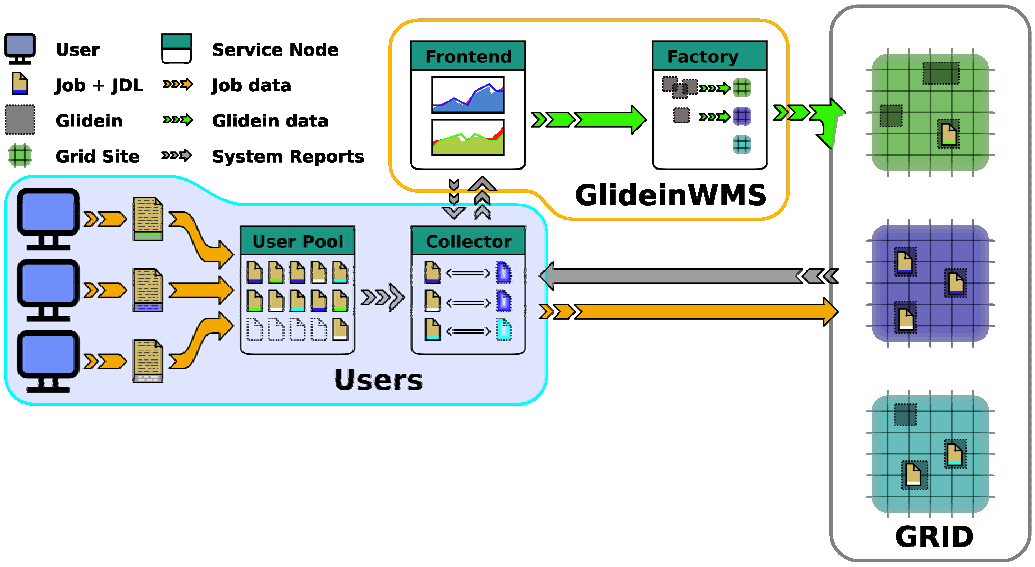

Figure 4: GlideinWMS components and processes flow

Once a pilot job starts on a Grid resource, it first validates it. This includes looking for the appropriate software libraries, ensuring sufficient disk space is available, and so on. Once these tests successfully complete, it configures the Condor daemons (i.e. the startd and supporting code) and starts them. All control is delegated to them, and the pilot job wrapper just waits for their termination to do the final cleanup.

Once the Condor daemons start, they behave like in a dedicated Condor pool. The startd registers back with the Condor central manager and waits to be matched. Once a suitable job is found, the schedd holding the job will contact the startd and the job starts running. The startd can run multiple consecutive user job, to keep the pilot wrapper overhead low for very short jobs. However, the pilot needs to end within the Grid back slot lease time, so the startd will stop accepting new jobs if they cannot complete in time.

The system has been designed to be highly scalable [19]. CMS commissioned this system into production at CERN/FNAL/UCSD as primary workload management system for the production and analysis jobs in HEP.

IEKP contributed in this work by providing expert support for the setup, configuration and tuning of the central components, the most complicated parts of the GlideinWMS deployment.

7.2. GlidInWMS access for users to NRG

The IEKP is currently in the last steps of establishing the usage of a private GlideinWMS system for IEKP physicists. Within this IEKP GlideinWMS factory the resources of the NRG become accessible from local IEKP login machines. The idea is to provide semi-interactive access as provided by GlideinWMS to the NRG Grid resources to IEKP users to help users to overcome the flaws of Grid computing, like the usual overhead in debugging, variety of tools, and the lower efficiency compared to the use of local resources.

The users will be able to submit the jobs locally via condor to the GlideinWMS factory, which handles all Grid related work. All Grid related problems will be intercepted by the GlideinWMS factory and are therefore directly in the hands of a Grid expert. The user thus can focus on the optimisation of his analysis code and will mostly face self-caused problems related to his executable or data.

7.3. PhEDEx Project

The PhEDEx (Physics Experiment Data Export) project provides the data placement and the file transfer system for the CMS experiment. Created in 2004, PhEDEx has by now become a mature software product. It currently runs on over hundred of CMS distributed computing sites, performing a variety of tasks such as managing data subscriptions and transfers between the sites, staging data from tape, data removal, storage consistency checking, bookkeeping and monitoring all related activities.

Originally implemented as a set of specialised tools, PhEDEx has been substantially refactored. It now adopts a concept of an open framework [20] providing a set of generalised technical solutions, which could also be used standalone. Examples are Core Agent, AgentLite, Namespace, Data Service, Website frameworks, and LifeCycele Agent.

7.4. PhEDEx Namespace Framework

One of our major areas of contribution is Namespace Framework, which provides a communication layer to various types of storage systems and is used for data consistency checking. Figure 5 demonstrates the use of Namespace Framework for the data consistency checks over distributed CMS sites. Communication between the operators and the sites happens via central PhEDEx database. The agents running at the sites check for test requests to be performed on the local data. The results are uploaded back to the database and published to the web site using PhEDEx Data Service.

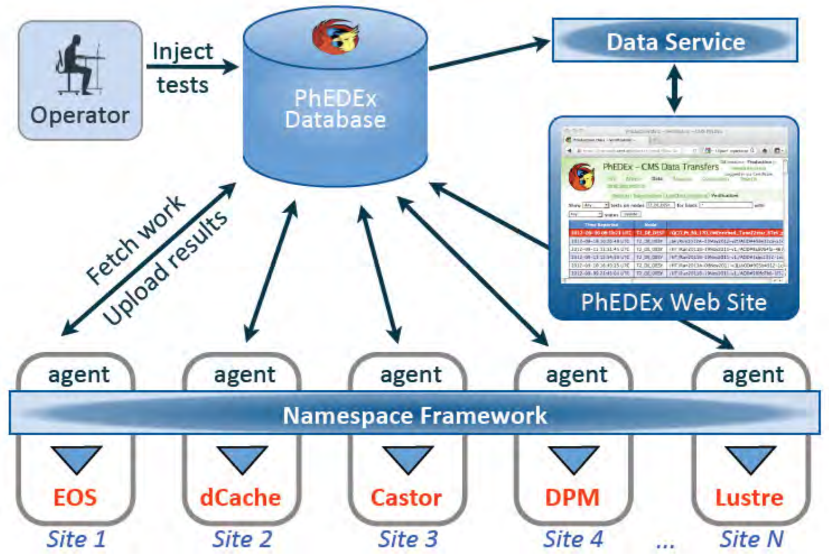

Figure 5: Namespace framework is used by PhEDEx agent for data consistency checking at distributed Grid sites

We have used our Tier-1 site for large scale testing aimed at improving the performance of the consistency checking agents. The time required for a full check of all CMS production data stored at the site have been reduced from several days to less than twelve hours. This allowed to use the system for regular data consistency checks throughout the CMS Tier-1 sites. It was shown that further three-fold improvement of timing can be achieved by using storage dumps instead of accessing the storage directly. These storage dumps are produced by the sites and used for detecting orphaned data not registered in the CMS central file catalogue.

This year, together with CMS data operations team, we have organised a CMS-wide campaign to add support for various storage technologies used at Tier-2 sites. By now automated consistency checks are running at over 50 % of CMS Tier-2s.

7.5. PhEDEx Validation

We keep track of new versions of the three PhEDEx components website, data service and agents, and provide information on the status of the releases, known bugs and bug fixes, and feature requests.

Besides this “bookkeeping” task we used the KIT Tier-1 site as a prototype for large scale testing for new PhEDEx agents releases. Basic test scenarios like testing the subscription and deletion of datasets on the KIT Tier-1 site were performed as well as advanced and specialised test scenarios, like deleting a dataset at Tier-0 and subscribe it to Tier-1, while the deletion is still ongoing.

The major project in the field of PhEDEx validation is the development of a validation tool suite based on the lifecycle agent to test the behaviour of new data service releases in a testbed using different roles (admin, data manager, site manager). Possible scenarios include testing that site admins cannot inject data anywhere, nor approve or disapprove requests and verifying that the data manager for T1 X can do so for T1 X, but not for T1 Y. The final goal is to run the test scenarios in a reliable and automated manner. Another tool is currently developed to test the email notifications of the PhEDEx system prior to the release of a new data service version. Whenever a transfer/deletion request is made via the PhEDEx system, the global PhEDEx admins and the responsible group admins, data managers and site managers of the respective groups and sites should receive an email so that they can further proceed and either approve or disapprove the request. In order to allow for a smooth workflow it is essential that this notification works.

7.6. Storage Accounting project

All major LHC experiments need to measure real storage usage at the Grid sites. This information is equally important for resource management, planning, and operations. To verify the consistency of central catalogs, experiments are asking sites to provide a full list of the files they have on storage, including size, checksum, and other file attributes. Such storage dumps, provided at regular intervals, give a realistic view of the storage resource usage by the experiments. This brought us to the idea of monitoring of storage use based on storage dumps, which resulted in storage accounting project [21]. IEKP contributes in all aspects of this project, including analysis of use cases and requirements, and concrete implementation. This work is conducted in close collaboration with WLCG, CERN IT Experiment Support group, ATLAS and LHCb experiments, CMS Monitoring Task Force [22], CMS PhEDEx development team, CMS central data operations, and a set of CMS pilot sites volunteered to participate in testing.

The Storage Accounting tool is re-using PhEDEx code base and components with some extentions. Particularly, Namespace framework has been extended to provide support for parsing of different types of storage dumps. It is envisioned to use PhEDEx Agents technology for driving the information provider on the local site. The API used to upload and retrieve the disk usage data to and from database is modeled based on PhEDEx Data Service. The data is stored in Oracle database at CERN. However it is using a separate DB instance to keep space usage information separately from PhEDEx data transfer details. Work with the remote sites is ongoing to deploy and test tools for producing the storage dumps. First version of Storage Accounting Data Service has been deployed on the production CMS web server at CERN.

7.7. HappyFace Project for Site Monitoring

The Tier-1 center at GridKa is monitored by a group consisting of local CMS members. The main tool for this job is the HappyFace framework. HappyFace is a meta-monitoring framework which gathers information from different monitoring sources, processes this information and provides an overview of all relevant information, which allows real-time site monitoring for both shifters and experts. HappyFace started as a development project at KIT and is now used by several other German grid sites, both ATLAS and CMS, to monitor their sites. It is also employed centrally by CMS to monitor the batch systems of all CMS Tier-1 and Tier-2 centers. More details about the architecture of the HappyFace framework and the current development can be found in a separate article in these proceedings [23].

8. Summary

The combination of Research Lab and Technical University at KIT allows one to combine Tier-1 operational services with computing development projects and thus provides an excellent environment for training of young scientists in the areas of high-performance and high- throughput computing.

Expertise within the CMS Tier-1 local support group at IEKP spans a wide range of Grid computing topics, including job submission, data management, monitoring, accounting, consistency checking and validation tools. Expert knowledge combined with coordination efforts and an expert rotation scheme help to provide sustained support of CMS-specific Tier-1 services during intensive data taking at LHC.

Members of our team provide strong contribution to CMS core computing development and support, and are open for collaboration with other sites and experiments.

References

[1] Knobloch J *et al.* 2005 ”LHC Computing Grid Technical Design Report” CERN-LHCC-2005-024.

[2] Heiss A, Petzold A, Zvada M, “Overall experience of GridKa T1 operations and LHC experiments representation”, ibid pp. 158-164.

[3] The CMS Collaboration 2008 ”The CMS experiment at the CERN LHC” JINST **3** S08004.

[4] Institut fu¨r Experimentelle Kernphysik, http://www-ekp.physik.uni-karlsruhe.de/

[5] CMS Collaboration 2005 “CMS: The computing project. Technical design report,” CERN-LHCC-2005-023.

[6] Bonacorsi D 2007 ”The CMS computing model” *Nucl. Phys.* B *(Proc. Suppl.)* **172** 53-56.

[7] WLCG Memorandum of Understanding, http://wlcg.web.cern.ch/collaboration/mou

[8] WLCG Resource, Balance and Usage website, http://wlcg.web.cern.ch/rebus

[9] Kreuzer P 2011 “CMS computing performance on the GRID during the second year of LHC collicions”, Proc. of NEC 2011, Varna.

[10] Kaselis R 2012 ”CMS Data Transfer operations after the first years of LHC collisions”, Proc. of CHEP 2012.

[11] N. Magini, N. Ratnikova et al 2011 ”Distributed data transfers in CMS,” J. Phys. Conf. Ser. **331**, 042036.

[12] Rehn J et al. 2006 “PhEDEx high-throughput data transfer management system” Proc. of CHEP06, Mumbai.

[13] Savannah - the LCG software development portal, https://savannah.cern.ch/

[14] T. Antoni et al “WLCG-specific special features in GGUS,” J. Phys. Conf. Ser. **219** (2010) 062032.

[15] GlideinWMS Web site, http://www.uscms.org/SoftwareComputing/Grid/WMS/glideinWMS/doc.html

[16] Condor Project: http://research.cs.wisc.edu/condor/

[17] Sfiligoi I 2008 “GlideinWMS - A generic pilot-based Workload Management System” Proc. of CHEP2007.

[18] M. Zvada M , Benjamin D, Sfiligoi I 2009 “CDF GlideinWMS usage in grid computing of High Energy Physics” Proc. of CHEP 2009, Prague.

[19] Bradley D, Sfiligoi I, et al 2009 ’Interoperability and scalability within glideinWMS’, Proc. of CHEP 2009.

[20] Sanchez-Hernandez A, Egeland R, Huang C-H, Ratnikova N, Magini N and Wildish T, 2012 ”From toolkit to framework - the past and future evolution of PhEDEx”, Proc. of CHEP 2012.

[21] Huang C-H, Lanciotti E, Magini N, Ratnikova N, Sanchez-Hernandez A, Serfon C, Wildish T, Zhang X “Data Storage Accounting and Verification at LHC experiments”, Proc. of CHEP 2012.

[22] Bauerdick L and Sciaba A 2012 ”Towards a global monitoring system for CMS computing”, Proc. of CHEP 2012.

[23] S. Röcker et al, ”Meta-monitoring with the HappyFace project”, ibid pp.200-203.