

Table of Contents

Support tools.....	1
Introduction and overview.....	1
Technology and tools.....	1
Ticketing tools and Request trackers.....	1
GGUS.....	2
Savannah.....	3
Trac.....	3
JIRA.....	4
Accounting tools.....	4
Administration tools.....	5
20111212 meeting.....	6

Support tools

Introduction and overview

This section covers all tools used by WLCG operations, in particular for what concerns problem reporting and tracking, resource accounting and administration information. Many of these tools were developed by other projects, as OSG and EGEE, and are all supported by OSG or EGI (some are open source solutions); however their development was heavily influenced by WLCG and are essential for the WLCG operations themselves.

Typically they have reached a high level of maturity, but WLCG should make sure that future developments do not diverge from the WLCG requirements.

Technology and tools

The tools covered in this section are:

- GGUS
- Savannah
- Trac
- JIRA
- GOCDB
- OIM
- EGI Operations Portal

Ticketing tools and Request trackers

The ticketing system of choice for WLCG is the Global Grid User Support (GGUS) system, accessible via email address Helpdesk@ggusNOSPAMPLEASE.org or the web entry point <https://ggus.org> (requires login or certificate-based authentication). GGUS' scope is wide. National Grid Initiatives (NGIs) i.e. Grid sites, Support Units (SUs) related to Virtual Organisations (VOs), middleware projects can accompany tickets towards their solution.

GGUS is flexible allowing direct site notification and automatic assignment to the relevant SU. This functionality bypasses the dispatching SU. TEAM tickets allow a number of knowledgeable VO members to co-own a ticket, hence remain up-to-date at all stages of the solution process. ALARM tickets allow a small number of Grid experts within the VO to raise ALARMS at the Tier0 or Tier1 sites for the appropriate problem areas according to the MoU.

GGUS maintains interfaces with the "CMS Computing" Savannah tracker, the GOCDB, the Operations' portal, the LCG VOMS service, the OSG Information Management (OIM) and the local ticketing systems in the Tier-0, several Tier-1 sites, the OSG and multiple NGIs.

Savannah is used by all 4 LHC experiments to track various experiment activities. Only the "CMS Computing" Savannah tracker is equipped with a GGUS bridge, to automatically convert a Savannah item into a GGUS ticket.

JIRA and TRAC are request trackers used by developers in areas like experiment applications, middleware, monitoring tools and others.

GGUS

Works well

- Most comments received agree GGUS overall works well.
- ATLAS' experience showed that sites understand better the importance of an incident report when this is done via TEAM or ALARM tickets. Hence, this functionality is very useful.
- The fact that GGUS is a **distributed** support system with **central** coordination is considered useful.
- WLCG collaboration with the GGUS developers for fixes and enhancements' planning is excellent.
- The use of GGUS tickets for further investigation at the WLCG daily meeting means its established adoption as The Reporting Tool of Choice for WLCG, which makes things simpler and clearer for follow-up.
- GGUS' integration with other tools like the operations' portal, GOCDDB, OIM for US sites makes up-to-date information available to users and supporters in an automated fashion.

Top issues

- Different priorities across collaborating projects (EGI, EMI, WLCG) require constant effort to make development plans.
- No fail-safe solution is available for GGUS so far, which means that if the web interface is not accessible one can't see ticket progress or open TEAM and ALARM tickets.
- The lack of a focused User Support Working Group, like the EGEE USAG that met once per month with a specific technical theme and brought development forward with consensus and commitment agreed by all partners (VOs, sites, mgnt) makes it very difficult to take implementation decisions that meet consensus by the collaborating communities.
- ATLAS would like GGUS to report what is going on at the site in terms of tickets and downtimes and actions. Programmatic access of GGUS infos should be stabilized.
- There are recurrent issues with ALARM tickets forwarding the ticket to the proper unit. Often this is human error, the operator calls the wrong service. Nevertheless, during monthly tests we regularly see holes in the automated processes as well.
- WLCG relies heavily on GGUS, the life of which, in fact, depends on EGI sustainability.
- GGUS workflow allows change of SU any time in the life of a ticket. CERN's ticketing system SNOW doesn't contain all the world-wide Grid Support structure, so two instances of the same ticket can get quite confused as they live different lives in the two systems.

To improve

- There is an important amount of information repetition in various meeting reports. This does not only concern information on tickets. It would be good to re-discuss the ways we use to disseminate information to the community.
- Today, when opening a GGUS ticket, this is classified as an incident by default. GGUS provides the option to create "Change requests". We should start using it and interface appropriately to external ticketing systems, see Savannah:120007 [↗](#)
- Some of the tools we need change names and locations creating a migration overhead with no added functionality (e.g. moving from Savannah to RT, changing the URL of the CIC portal etc).
- Put in place mechanisms to detect any break-up in automatic interfaces with external ticketing systems or other tools (VOMS, GOCDDB, OIM etc).
- Tickets transiting across ticketing systems, sometimes end up in states that prevent us from following up progress.
- GGUS email templates must be adjusted to take into consideration notifications sent by SMS with a limited length for the message text (frames of asteriscs hide the content). Will be done in November 2011, see Savannah:124169 [↗](#).

Savannah

Works well

- Savannah, as used by CMS works well as it has mapped all the site's admins/responsibles and the tickets are assigned to them efficiently. Also, it allows the creation of squads, not related to sites, but for services, or tools, which is helpful.
- For ATLAS, being able to move tickets from one unit to the other in Savannah (for example from Operations to development, from Data Management Development to workload management development) is appreciated.
- LHCb is also using Savannah heavily, there is one instance per project created. Savannah is also used for release management, i.e. grid deployment activities are being handled in dedicated instances.

Top issues

- Savannah lacks of GGUS functionality. CMS asked for the implementations of the Savannah-to-GGUS bridge. The Savannah-to-GGUS might be improved: we can still write comments on Savannah which were bridged. The Bridge creates normal GGUS tickets, which can be converted to Team and Alarm tickets from the GGUS portal only.
- The future of Savannah is unclear. If it became unsupported, GGUS would probably have to implement all Savannah features that are required by the experiments.

To improve

- It is not always obvious when a ticket has to be submitted to GGUS, to Savannah or to other tools.
- Reducing the number of official ticketing tools would reduce the need to bridge one to another.

Trac

Works well

- CMS has experience using Trac, a web-based project management tool, which has proven to be useful, at least for the projects development team.

Top issues

- Trac has proven to be a good tool, but the CERN deployment of it is limited to providing the tool "as is", as an add-on to SVN.
- The service has gone through periods of very low availability with slow response time for large projects, due to issues with the load balancing setup of the SVN servers and also the built-in Trac database backend. This has improved recently, thanks to the introduction of a new load-balancing scheme for the SVN/Trac server cluster. The support team has also taken measures to improve the underlying database back-end for Trac. (Migrated to MySQL.)
- CMS has expressed that when requests for improvements in the service have been made the response has been that changes will have to wait, due to a lack of resources. CMS is quite seriously considering moving to github, both for reliability and feature reasons.

To improve

- Response-time and stability of the service is important for the users of the service. Recent improvements and consolidation should be continued.

JIRA

Works well

- Known to be a very powerful issue tracking application. Those experiment groups who are using it are very happy.
- OSG uses JIRA for change management as well (see Savannah:120007#comment19[↗])

Top issues

- Some find JIRA to complex for simple bug-tracking.
- JIRA was earlier not provided as an official service at CERN, but some LHC experiment groups set up private JIRA installations. The BE/CO group at CERN in charge of LHC accelerator operations have also been using JIRA for a number of years, and are happy with this system as a powerful issue tracking tool.
- Following the IssueTrackingSurvey in 2011 and subsequent discussions in the IT Technical Users meeting and IT Service Review Meeting, it was decided to launch centrally supported issue tracking service based on JIRA. The central issue tracking service[↗] is fully operational, more information on: <http://cern.ch/its>[↗].

To improve

Accounting tools

Works well

- The APEL accounting infrastructure receives cpu accounting data from its own clients at three hundred EGI sites and from a number of other accounting systems (Gratia, SGAS, DGAS) to provide a single worldwide database of accounting data for the LHC experiments. Earlier this year it exceeded 10*9 jobs dating back to 2004.
- The Accounting Portal contains summaries by Site/Month/VO/UserDN/FQAN and users can visualise or download data for any combination of these, at any point in several hierarchical trees (Tier1/2, Country, NGI)
- Data on users is only available to authenticated users based on: the User (sees their own data); VO Manager (sees their VO); VO Member (sees VOMS roles/groups in their VO); site admin (sees users and FQAN running at their site).
- Nagios testing of whether sites publish and whether the data extracted at the site has been published centrally

Top issues

- Benchmarking: Normalisation of cpu data requires a reliable knowledge of the power of the cpus. The required benchmark is HEPSP06. The quality of the published data is not reliable. It is not certain that all sites actually run the benchmark on their clusters. Published values for the same nominal cpu vary greatly. Sites also make mistakes in averaging results across their cluster and forgetting to update when they add new cpus.
- While Nagios tests can compare the data collected locally by the parsers with what was published to APEL, there is no comparison with the total of all local batch systems. Thus if a site doesn't run the parser on a CE then no results are published for that CE. Completeness relies on the site noticing a discrepancy.
- Storage Accounting: under development in EMI for their supported storage systems. The new central infrastructure (mentioned below) will receive and store storage accounting usage records (StAR) and the portal will develop required visualisation of the data. Non-EMI storage solutions (Castor, xrootd, Bestman, ??) will also need to publish if WLCG is to have complete data on storage.

To improve

- The messaging infrastructure is under redevelopment. It will continue to support old clients through EMI-3 but the new client should be more reliable. Other systems (Gratia, etc) will publish by messaging. This is currently under test.
- More accounting systems will be supported but these generally won't affect WLCG sites where use is almost pervasive (an exception is GSI Damstadt, who run significant Alice work with no grid middleware)
- The Accounting Portal has a rudimentary URL query interface which returns XML. This could be developed into a full RESTful programmatic interface.
- The Accounting Portal should show data on more users and be able to search on a particular UserDN or FQAN (subject to the usual authorisation)

Administration tools

WLCG has adopted since many years some administration tools, the **Grid Configuration Database (GOCDB)** and the **EGI Operations Portal**, developed in the context of EGEE and now maintained by EGI. They provide several critical functionalities to WLCG, among which:

- an official information repository for all EGI sites and related information, including: contact emails, mailing lists and phone numbers, pointers to the local BDII and website, all registered services, all site administrators and their role;
- a system to publish service downtimes, their start and end time, their severity, a description and their impact, allowing to automatically send broadcasts to the VOs supported by the site;
- a catalogue of Virtual Organisations and their VO Identity Cards, containing static information such as a description of the purpose of the VO, a home page, the required resources, the endpoint of its VOMS server, a complete list of groups and roles, a free text description of its requirements and a list of contact people with their role;
- the ability to submit broadcasts targeted to site administrators, VO managers, VO members, NGI managers and official mailing lists, and to consult an archive of old broadcasts.

The GOCDB also provides a programmatic interface [☞](#) that is used by other software, including ATP (a SAM/Nagios component), Gridview, GSTAT, APEL, the Accounting Portal and some Experiment Dashboard applications, to extract information, in particular about registered sites and services and downtimes (for example, to feed downtime calendars).

The **OSG Information Management (OIM)** project [☞](#) is similar in scope to GOCDB: its goal is to gather necessary information from the OSG consortium to provide accurate information for administration and operation of the OSG. This includes topology and census data that can be used to answer the question "What is the physical structure of the OSG?".

OIM plays in OSG the same role of the GOCDB and the EGI Operations Portal in EGI, including a catalogue of VOs with contact and VOMS configuration details, a list of the Support Centers and a plethora of site and resource information. Finally, it allows to publish and retrieve information about service downtimes.

Works well

In general, GOCDB is perceived as a very useful, in particular for what concerns the downtime publication. The documentation is complete and up to date. There are well defined procedures to declare downtimes, extend them, deal with "at risk" downtimes that cause an outage and put nodes out of production.

Top issues

Among the main issues reported we can mention these:

- The information in GOCDB is not always up to date.
- GOCDB does not tell which VOs are supported by each service; this makes much more complicated for a VO to understand if it is affected by a downtime and requires to rely on external sources, like the information system, to find out what services are relevant for it. The real problem however is that no service provides VO-dependent downtime information, and it is not at all obvious that GOCDB is the best place where to insert VO-dependent information.
- Having more than one system (namely GOCDB and OIM) doing similar things but for different Grids (which are nevertheless part of WLCG) makes much more difficult to obtain information in a uniform and consistent way, to the point that can be partially achieved only via interfaces developed by monitoring projects and which have not yet achieved a satisfactory level of stability.
- Automatic downtime notifications are sometimes send more than once to the same person, if this person has more than one role. As a consequence they may be perceived as a nuisance and ignored.
- In OIM, there is no easy way to add new service technologies (Squid, Frontier, VM hosts, etc.) beyond what is supported now (CEs and SEs). This is considered to be its most important limitation.

To improve

- A feature that would be very useful for the LHC experiments and which does not necessarily fit into any of the existing tools is a place where to publish their "latest news", in a way similar to what the CERN IT does with the CERN IT Status Board.
- The Operations' portal sometimes doesnt send the site downtime broadcast for several hours and nobody notices it until WLCG users report it.
- At the 2012/01/23 f2f workshop at NIKHEF a need for good contact with the GOCDB developers was phrased. The answer to this is to address requests to gocdb-discuss@mailtalkNOSPAMPLEASE.ac.uk

20111212 meeting

Improvement Areas (10=Blocker , 5= Medium, 1=Low)	
Impact	Areas
5	GGUS fail-safe availability - whole stack (web front-end, remedy system and Oracle db)
5	Convince the entities funding GGUS of its sophisticated use by WLCG to ensure sustainability of our development priorities.
5	Ticketing systems and trackers' inter-connection (some entries are not mapped, users & supporters get incomplete stories)

-- MariaDimou - 08-Nov-2011

This topic: LCG > WLCGTEGOperationsSupportTools

Topic revision: r27 - 2012-09-24 - NilsHoeimyr



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback