

Report from WG2

Andrea Sciabà

- Support tools
 - Ticketing tools
 - Accounting tools
 - Request trackers
 - Administration tools
- Underlying services
 - Messaging services
 - Information services
- WLCG operations and procedures

- Overview
 - Tools mostly developed by other projects (OSG, EGEE, EGI...)
 - WLCG heavily influenced their development
 - Rather mature by now

- GGUS
- Savannah
- TRAC
- JIRA
- GOCDB
- OIM
- EGI operations portal

- **GGUS**
 - Used by all 4 experiments for incident reporting
- **Savannah**
 - Used by ATLAS, CMS, LHCb for internal investigation before bridging incidents to GGUS (CMS) or to other trackers (ATLAS) for development and/or release management (LHCb)
- **TRAC and JIRA**
 - Used by some experiments (as CMS) as development trackers but supporters make it available ‘as is’ so required improvements (e.g. on performance) are done on a best-effort basis

- Areas of improvement
 - GGUS
 - Some external interfaces periodically break
 - Ensure continuous availability
 - Savannah
 - Improve integration with other systems
 - TRAC / JIRA
 - Experiments would like them to be officially supported
- Areas of potential efficiency gains
 - GGUS: better reporting to avoid information repetition in multiple meetings
- Largest use of operational effort
- Missing areas
 - Savannah future uncertain

- Overview of technology and tools
 - APEL, Gratia, SGAS, DGAS
 - APEL receives CPU accounting data from its clients and the other accounting systems
 - Provides a single database of WLCG accounting data (~ 1 G jobs since 2004)
 - EGI Accounting Portal
 - Provides summaries by site/month/VO/user/FQAN and data can be plotted and downloaded
 - Authorisation to see data on users depends on role
 - SAM/Nagios used to check that sites publish data and if this is published centrally

- Areas of improvement
 - Benchmarking: published data not reliable
 - SAM tests for accounting data publication do not check the total of all batch systems, hence missing info may pass unnoticed
 - Storage accounting: development of a portal under way in EMI; non-EMI SEs will have to provide data in the correct format
 - Evolve Accounting Portal API in a full RESTful interface
- Areas of potential efficiency gains
 - Improved reliability from the redevelopment of the messaging infrastructure; messaging used also by Gratia, etc.
- Largest use of operational effort
 - Not reported
- Missing areas
 - Not reported

- Overview
 - GOCDB and EGI Operations portal provide several critical functionalities
 - Information repository for all EGI sites and VOs
 - Downtime publication
 - Broadcasts
 - GOCDB has a programmatic interface used to get info about registered sites, services and downtimes
 - OIM provides very similar functionality for OSG

- Areas of improvement
 - More updated info in GOCDB
 - Supported VOs
- Areas of potential efficiency gains
 - Seamless integration of GOCDB and OIM
 - Smarter and more reliable downtime notifications
 - Easier definition of new service types
- Largest use of operational effort
 - None identified
- Missing areas
 - A way to publish experiment news to a portal (similar to the CERN IT Status Board)

- Overview
 - Messaging system and the information system
 - Both developed by WLCG
 - Will have to include batch systems as well

- Active-MQ MSG system
- BDII
- GLUE
- LDAP



- Overview
 - Operated by EGI: two brokers at CERN, one at AUTH and one at SRCE
 - Two more broker services at CERN for testing and validation, one for ATLAS/DDM, one for IT-ES (each consisting of 2 prod and 1 test broker)
 - Used by several applications
 - APEL
 - SAM
 - Ganga/DIANE monitoring
 - LFC catalogue synchronisation (EMI prototype)
 - ATLAS/DDM tracer service (prototype)
 - FTS monitoring

- Areas of improvement
 - Security
 - scalability
- Areas of potential efficiency gains
 - Improve availability and reliability: now the service must be stopped during some interventions
- Largest use of operational effort
 - None identified
- Missing areas
 - None identified

- Overview
 - Covers several use cases
 - Service discovery
 - Installed software
 - Storage capacity and accounting
 - Batch system queue status
 - Configuration
 - Installed capacity
 - Fully distributed, hierarchical set of BDIIs, based on OpenLDAP
 - Implements GLUE schema
 - Information providers generate the service information

- Areas of improvement
 - *Stability*: service info is prone to disappear, bad because use cases shifted towards needing more stability
 - *Information validity*: info provider info very fragile, configuration very error prone
 - Better policies for resource publication
 - Lower latency for dynamic information
- Areas of potential efficiency gains
 - Better validation tools
 - Accurate storage information would make storage accounting a lot easier
 - Provide more powerful and user-friendly client tools
- Largest use of operational effort
 - Configuration and validation of information
 - Debugging IS problems for users and sites
- Missing areas
 - A continuous certification and auditing of the BDII information by WLCG

- Overview
 - Goals are:
 - Efficient communication
 - Quick resolution of issues according to agreed targets
 - Coordination and decision
 - Well defined procedures
 - Describes roles, bodies, communication channels and procedures
 - Lots of experience accumulated
 - Quality is good but still manpower intensive
 - No visible decrease of incidents

- Technology and tools (so to speak...)
 - Daily meeting
 - Tier-1 service coordination meeting
 - GDB
- Roles and bodies
 - Security, information, data management officers
 - Site administrators
 - Site security officers
 - Experiment contact persons

- Procedures and policies
 - Scheduling downtimes
 - Well defined rules to declare them
 - Problem handling
 - Little in terms of formal procedures, issues and incidents are handled and discussed in the daily meeting and the T1SCM
 - SIR for major incidents are an essential tool
 - GDB also useful to discuss issues at a general level

- Areas of improvement
 - Sometimes the strength of the link between an experiment and a site is not enough
 - The very need of site contacts can be seen as an issue...
 - Improve communication of the experiment requirements to the sites (e.g. via VO cards)
- Areas of potential efficiency gains
 - To have a real WLCG operations team: now experiments do most the computing operations
 - A better communication channel for the Tier-2's (now only the GDB)
- Largest use of operational effort
- Missing areas