


WLCG Operations and Tools TEG Report to the GDB

Jeff Templon and Maria Girone

9 November 2011



Areas of Work

Large spectrum of topics in charge

- Monitoring, (SAM, Nagios, etc.), metrics, and analysis of these data
- Support tools (APEL, GGUS, etc.)
- Underlying services (ActiveMQ, BDII, etc.)
- Application software management (e.g. cernvmfs)
- Operational requirements on middleware
- Configuration management
- Deployment management
- Middleware distribution management
- ... plus of course WLCG operations in the large

How we divided the work

- **WG1: Monitoring and Metrics**
 - Editors: Simone Campana and Pepe Flix
 - Contributors: Costin Grigoras , Andrea Sciaba, Alessandro Di Girolamo; (sites) Ian Collier , Xavier Espinal, Vera Hansper , Alexandre Lossent
 - WG1 Sub-Wiki: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTegOperationsWG1>
- **WG2: Support tools + Underlying Services+ WLCG Operations**
 - Editors: Andrea Sciaba, Maria Dimou (support tools), Lionel Cons (underlying services) and Stefan Roiser (WLCG Operations)
 - Contributors: Simone Campana, Andrea Sciaba, Alessandro Di Girolamo, Joel Closier, Pepe Flix, John Gordon, Alison Packer; (sites) Alexandre Lossent, Xavier Espinal
 - WG2 Sub-Wiki: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTegOperationsWG2>
- **WG3: Application Software Management**
 - Editor: Stefan Roiser
 - Contributions: Marco Cattaneo, Steve Traylen; (sites) Ian Collier, Alexandre Lossent, Alessandra Forti
- **WG4: Operational Requirements on Middleware**
 - Editors: Maarten Litmaath, Tiziana Ferrari
 - Contributors: Maria Dimou, Laurence Field; (sites) Alexandre Lossent, Jeff Templon, Vera Hansper, Anthony Tiradani, Paolo Veronesi
- **WG5: Middleware Configuration, Deployment and Distribution**
 - Editors: Oliver Keeble, Rob Quick
 - Contributors: Cristina Aiftimiei, Simone Campana, ; (sites) Ian Collier, Alexandre Lossent, Pablo Fernandez, Felix Lee
- Horizontal Representation from OSG (R.Quick), EGI (T. Ferrari), and EMI (C. Aifitimiei, L. Field)

Timeline

- Oct 24th: Kick-off meeting
 - Weekly phone conferences since (check points)
- Nov 28th F2F: Deliverable on assessment of current situation completed (5 weeks)
 - Presentation at GDB (14th Dec)
- 12th December: Workshop on Future Strategy
 - At CERN, prior to the GDB
- Jan 23rd F2F: Deliverable on medium to long term strategy
- Feb 7th: TEG reports (10 more weeks, including Xmas break)
 - Plan and needs for the next 2-5 years

Where we are headed

- “Provocative Question”: how can we reduce the manpower needed for “operations” to zero?
- Starting point: ask three related questions
 - What works well?
 - What are the three main problems?
 - Where could we save the most manpower?
- These questions were asked to each TEG member for each sub-area.

Current phase

- Develop “current status” document
Details the current state of affairs of “operations”, e.g.
 - What is being used?
 - Which of it works well?
 - Which of it does not work well, and why?
 - What is (unnecessarily) manpower intensive?
- Input to this phase mentioned a lot of tools by name, but the idea at this phase is not a contest between tools! It is an **ASSESSMENT**.



Examples of Work so far

Excerpt, raw input received

Support tools, underlying services

- ActiveMQ (underlying service)
 - works well (no problems ever observed)
- Accounting “top three problems”
 - There is no open access to control/monitoring information that would allow to know which sites are badly reporting accounting data.
 - The accuracy of the accounting information is unknown.
 - There are no requirements defined by WLCG for storage accounting.

Excerpt, assessment draft

Monitoring

The perspectives from experiments and sites are however rather distant: experiments monitor activities while sites know about availabilities of baseline services used by experiments.

- Not always the experiment can easily breakdown the failure modes (or at least, it is a time consuming operation for an expert) and indicate the right service the site should look at; this complicates the site operations and affects the level of site support for critical services.
- The experiment feels it is responsibility of the sites to monitor the correct functionality of the services they provide.
- Notifications for site downtimes proved to be very useful, but they are still limited and difficult to interpret in many cases (downtimes in federated sites, partial downtimes)
- Experiment computing systems have also downtimes, hardly exposed to sites. So the sites can hardly tell if a lack of activity is due to a lack of experiment activity, a problem in the experiment stack or a malfunctioning service at a site

In general, today this **gap is bridged by the intervention of experts** both in the experiment operations and at the site (experiment contact people). Sites with a strong contact with the experiment operations tend to have less of those problems, but **obviously the cost in manpower is not negligible.**

Issues (from MB)

- The membership of the group – e.g. missing stakeholders? Site representation is thin, only 5 of 10 listed on the wiki have responded.
- Other issues? It is a huge amount of work, we are thankful for the help of the subgroup editors, we will do our best to deliver within the stated time frame.

Input still welcome, but hurry ... Nov 28th is less than three weeks away!!!!

[More information on the wiki.](#)