


WLCG operations

A. Sciabà, M. Alandes, J. Flix, A. Forti

WLCG collaboration workshop
July 7-9 2014, Barcelona





Outline

- The WLCG operations group
- Experience before and during Run1
- Achievements and current work
- Issues
- Priorities beyond Run2

WLCG operations

- Evolved from a decade long experience in prototypes, service and data challenges
 - Result of the effort of many individuals
 - Use also procedures and tools provided by the federated Grid projects
 - GGUS, GOCDDB, OIM, EGI portals, etc.
 - Initial focus on delivering a stable service
 - Very successful but at the expense of a high manpower cost
- In 2011 WLCG decided that it was the right time for an internal review
 - The Technical Evolution Groups



Operations before and during Run1

- WLCG operations ran since 2008 as a loosely coordinated effort – effective but manpower intensive
- The 2012 TEG review identified several weak points

No effective communication with the Tier-2's

No central operations team

No central body where to take operative decisions

Fragile experiment software installation procedures

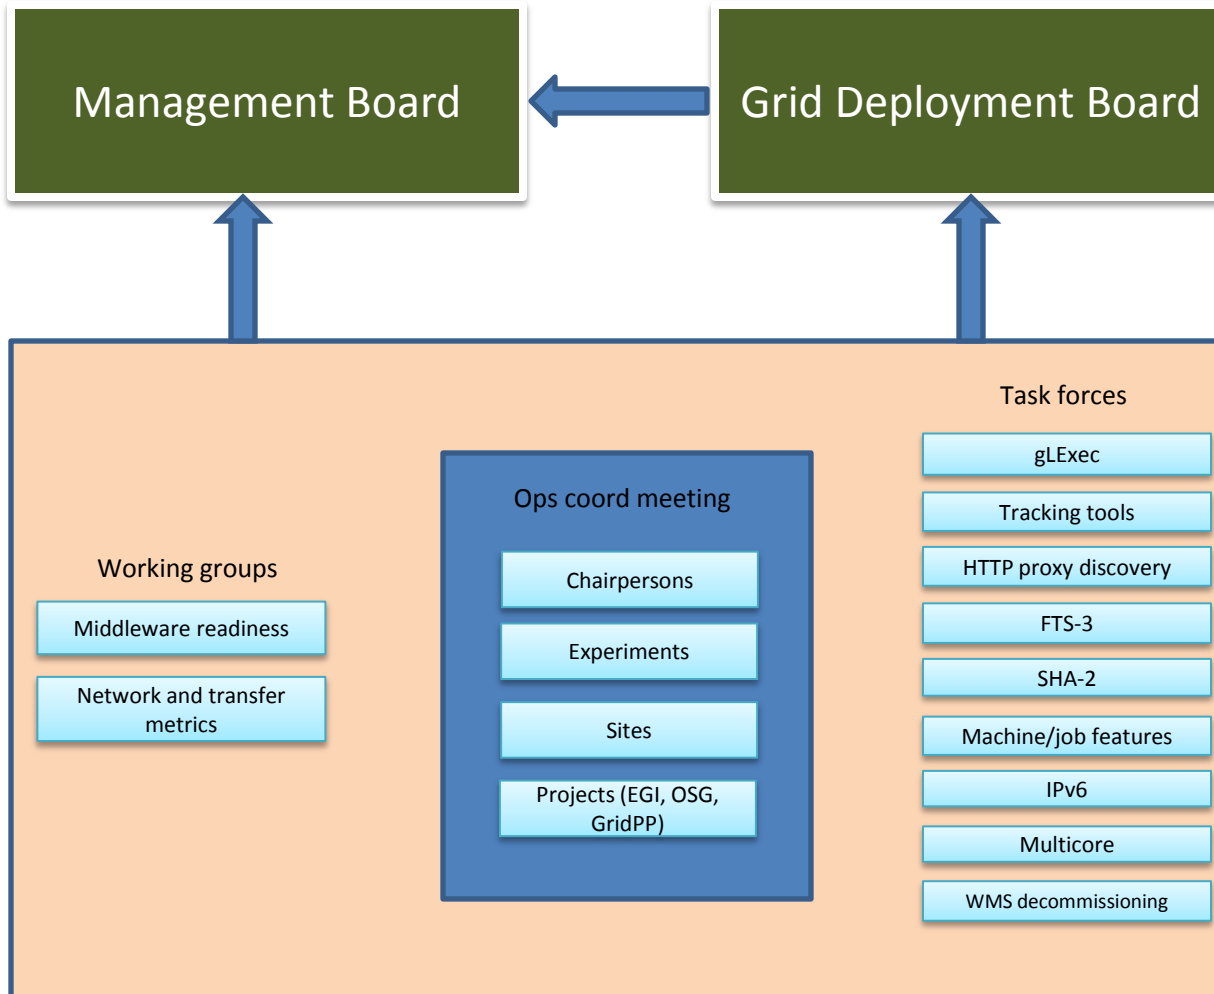
Poor documentation and logging for services

Insufficient middleware validation

The WLCG operations coordination working group

- Established in October 2012
- Core operations and deployment coordination team
 - Manages operational issues and service deployment in synergy with EGI and OSG
 - Discusses experiments plans and needs
 - Defines actions and work plans
 - Forms working groups and time-limited task forces on specific issues
 - Ensures communication among experiments and sites
- All stakeholders are represented
 - LHC experiments, Tier-0/1/2's, Grid projects
 - Largely based on voluntary effort from the entire WLCG community

Organisation



Operations during Run1

- Main goals
 - Implement the recommendations given by the TEG
 - Manage the daily operations
 - Solve specific problems requiring complex validation or deployment activities

TEG recommendations

Establish a core team for coordinating WLCG operations

Expand the scope of existing meetings to fully involve Tier-2 sites

Adopt CVMFS to distribute experiment software (and middleware)

Simplify the middleware stack and improve documentation and training

Improve middleware distribution and configuration mechanisms

Strengthen the participation of sites and experiments to the next phase of validation

The WLCG operations

T2 regional representatives

CVMFS used by all LHC
LFC and WMS

decommissioned

Rationalise repositories

Middleware readiness
validation working group

Completed

Work in progress

Achievements (1/2)

- Several other goals were accomplished
- The gLExec deployment is basically completed
 - Mandatory at CMS sites (affects their availability)
- CVMFS deployment at all sites and adopted by all experiments
 - Much more robust and scalable software distribution mechanism
- Migration to Scientific Linux 6 of all compute resources
 - Defined procedures for sites, a WLCG repository and a “HEP” metapackage for experiment software dependencies
 - Completed in the agreed timescale thanks to the considerable effort spent in keeping experiments and sites focused
 - As a bonus, validated the EMI-3 WNs for WLCG ahead of staged rollout


Achievements (2/2)

- Centralised squid monitoring for Frontier and CVMFS
 - Harmonised squid installations for different experiments and set up central monitoring pages
- WLCG-wide perfSONAR deployment
 - pS instances at almost all sites and central monitoring
 - the first important step towards network-aware applications
- Xrootd deployment
 - Enabled detailed monitoring of all xrootd traffic for ATLAS and CMS
 - Tested and distributed different plugins
- Validation of SHA-2 certificates
 - Validated all middleware and experiment software with SHA-2 certificates and proxies

Current activities

- Covering many high priority objectives for Run2
 - Enabling multicore resources
 - Commissioning of FTS-3
 - Enabling network-aware applications
 - Efficient job-compute node interaction in batch and cloud resources
 - Collaborative middleware validation using experiment applications
 - Testing and deployment of IPv6-enabled services
 - In collaboration with the HEPiX IPv6 working group
 - Many of these topics have dedicated talks in this workshop

Task force progress summary



Recent changes and new roles

- Recently reorganised the structure of the meetings and the reports
 - to make them more effective and useful by focusing on the most relevant issues and having more time for discussions
 - Task forces now maintain lists of specific goals and measure their progress
 - Eliminated the quarterly planning meeting, deemed redundant
- Created the role of the WLCG middleware officer
 - Tracks all relevant middleware issues
 - Decides which versions are the baseline
 - Works with the MW readiness WG to decide what to test and assess the results

Issues during LS1

- Site participation can be improved
 - Unfortunately many Tier-1 sites and Tier-2 regional representatives do not follow ops coordination meetings regularly
 - Some steps already taken, but still to be seen if effective
 - Hard also to get feedback
- Long tails in deployment campaigns
 - It's very hard to reach 100% completion
 - Tickets are a must but often not enough
 - Sometimes huge effort needed by the task force coordinators to keep things moving
- Uncertainty in middleware support
 - After the end of EMI, continuation of middleware support is a concern and we already face problems (e.g. ARGUS)
- Lack of risk analysis
 - Do we need it?
- Manpower
 - Not so severe, but contributions from sites are absolutely essential

Priorities for the next years

- Several task forces close to completion
 - SHA-2, gLExec, WMS, tracking tools
 - No new task forces on the horizon
- Future evolution will heavily impact operations
 - Changes in the experiment computing models
 - Changes in the infrastructure

Priorities for the experiments?

- Some clear commonalities
 - Funded CPU and storage resources will become increasingly scarce with respect to needs
 - Aggressive storage space management
 - Adapting to new types of resources beyond traditional tiers
 - Should the scope of WLCG operations change accordingly?
 - E.g. by providing VO-independent tools, procedures, recipes appropriate also for new types of resources?
 - Commissioning of new services, protocols, etc.
 - Here we can assume that WLCG operations will always play a role

Priorities from ATLAS

(from the ATLAS computing management)

- Reduce the overall effort from sites and experiments
 - The operations effort is still too heavy
 - Sites need very skilled people to run services which are still far from IT standards
 - Need to find a suitable model
- Fully enable the usage of cloud resources
 - If today 50% of Tier-2's moved to OpenStack, experiments would have a hard time
 - Need to address the operational side
 - Where the experiment becomes the sysadmin of the resources

Priorities from CMS

(from the CMS computing management)

- Continue to operate and support existing services and actively participate in their evolution during Run2
 - Central role in mediating between middleware developer teams and experiment requests
- Provide solutions for availability monitoring, accounting and resource utilisation in Openstack and other cloud resources
 - Covering also other resource provisioning interfaces like BOSCO
 - To contribute to enable opportunistic resources
- Usage and overload monitoring and service health for storage federations
 - Global storage usage accounting is a long standing request from the C-RSG
- Support in introducing new standards
 - Machine/job features, IPv6, multicore provisioning, etc.
- Operations of common services
 - Myproxy, gLExec, CVMFS, OS upgrades, perfSONAR, etc.

Priorities for WLCG

(input from WLCG management)

- Cloud computing
 - Not today, but will become an operations issue if sites start moving to clouds
- Understanding performance
 - Benchmarks, I/O metrics on storage federations, global monitoring
- Optimising the system for performance and cost
 - What can WLCG do to help experiments getting the most from the resources?
- “Virtual” sites
 - Will be possible to have sites that can run almost unattended and save on operational effort?

Conclusions

- WLCG Operations coordinate most activities of common interest for the experiments
- Site participation is mostly via task forces
 - But still to be improved
- Long term middleware support is already a problem
- Current activities are mostly targeted for Run2, need to look beyond
 - Good agreement in expectations between experiments and WLCG
 - Manpower, cloud resources, monitoring, service commissioning and operations are the main areas

Questions?