Operations Automation Team Kickoff Meeting

James Casey

CERN. 6-7 May 2008

(Slides from Maite @ WLCG Workshop)

EGEE-III Introduction

What is EGEE III?

- **A continuation of the EGEE program, building on its achievements and preparing the transition towards a sustainable infrastructure**
- **Goals:**
  - EGEE-III will work with EGI to transfer experience in operating Grid infrastructures, ensuring the development of a viable model. Based on the plans to be produced by EGI, EGEE-III will start implementing the changes for a transition to the EGI model
  - To maintain, enhance and simplify the use of the production quality computing infrastructure for an increasing range of researchers in diverse scientific fields

EGEE III Operations

Goals

- **The provision of a large-scale, production Grid infrastructure that interoperates at many levels, offering reliable services to a wide range of applications**
  - Continuation of the present service
- **Set the groundwork for the migration to a distributed model based on coordination at the European level of National Grid Infrastructures**
  - This is the challenge for the next 2 years, to do this without breaking the 1st goal (continuation of reliable service)

With the constraints:

- 2 years
- Significant y less effort

Centralized vs. distributed

What is our present (EGEE II) model?

- **Grid management**
  - Central coordination for all of the tasks, in many cases localised at CERN (team is called OCC: operations Coordinations Centre)
- **Grid operations and support**
  - In general, problem monitoring (SAM) and reporting done centrally by the COD is not well integrated with the daily operations and monitoring carried out at each site
  - Best effort/informal coordination of operations tool and requirements gathering. Most tools are deployed centrally (main instance run in one region serving the whole infrastructure)
- **Support to VOs, Users, Applications**
  - Central access point for user support, connected to all the ROCs
- **Grid security**
  - Central coordination, with effort from all ROCs and a broad collaboration from sites

Centralized vs. distributed

What is our target model?

- ROCs (NGIs) are responsible for day to day operations, without a central organization overseeing them. Set of operations tools supporting this
- Central body (OCC) responsible for coordination of cross-regional tasks
- Clear interfaces/targets between OCC and ROCs(NGIs), between ROCs(NGIs) and sites
- Sites with well developed fabric tools that monitor local and grid services in a common way and trigger alarms directly, so most of the issues are solved at this level

Target Operations task distribution

- **OCC:**
  - SLA and metric definition
  - Measure reliability and availability by aggregating data produced by the ROCs
  - Application – Resource Provider Coordination (new VOS across regions)
  - Coordination of operational issues that cross regional boundaries
  - Operations tools coordination (gather requirements, avoid duplication)
- **ROCs/NGIs**
  - Oversight and management of grid operations in the region
  - Provision of operations tools
  - SLA monitoring and follow up with sites
  - First line support (helpdesk, roll out of new releases, user support)
- **Sites**
  - Responsible to provide a reliable service, according to the SLA and the level of service required by the supported VOs

How to get there: general

- **Define work plans for all tasks**
- **Questions we need to answer for all the tasks:**
  - status of each task at the end of EGEE III
  - How will it be run? Centrally, distributed to all sites/NGIs/ROCs?
  - How will it be managed? Does it need central control in a central organisation, irrespective of how it will run? Can it be run by a group of peers coordinating without central control? Does it need coordination at all?
  - How will it be funded? Does it need central funding? National funding? VO funding?
  - Will it be stable? Or will it still have a plan for future development? For example one might plan a service that will still need central management at the end of EGEE III but could be completely distributed one year later.
- **When we have answered these questions there will be two benefits:**
  - We will see the constituent parts of SA1 that need the most change during EGEE III. This will let us write more detailed project plans for them.
  - We will have a view of the state of the infrastructure of EGEE III that will transition into an EGI. The EGI-DS desperately needs this input. So do the NGIs as they need to plan their future and will need to have a view on what they are being asked to take

(from John Gordon's proposal for Operations transition meeting next week)

How to get there: OAG

Operations Automation Group (chair: James Casey):

- Improve site reliability by wider deployment of fabric management tools at sites
- Devolve central monitoring systems, where possible, to regional systems
- Create architecture for new shared infrastructure required to support the operational tools
- Measure and improve the availabity and reliability of the operational tools themselves
- Design SLA compliance tools (availability and reliability calculation)
- Collection of usage and accounting information for CPU/Disk/Network
- Provide vizualization of the state of infrastructure for site administration, regional operators and project managers
- Provide reporting tools for the OCC and project management

How to get there: COD

New mandate for the Grid operator on Duty, COD (chair: Helene Cordier), being defined, to:

- Move from central COD model to regional COD model; goal: central COD as we understand it today does not exist at the end of EGEE III; all ROCs being responsible for the COD tasks by the end of the project
- Define what a “thin central COD layer” would do, and if needed introduce it in the new COD model
- Coordination between regional CODs and central thin layer COD
- Share of expertise between regional CODs

How to get there: SLAs

Definition, monitoring and enforcement of Service Level Agreements (chair: John Shade)

- **Measure service level in view of improving it**
  - EC review comment: “The measures of robustness and reliability of the production infrastructure are still very rudimentary.”
- **Formalize the responsibilities of both parties**
  - Avoid misunderstandings
  - Improve relationships between both parties
- **Understand what must be supplied**
- **Understand what is the minimum acceptable**
- **Identify service parameters**
  - Availability, Reliability, Ticket response times

Conclusion

- **Moving to a fully distributed model; we have some experience with this, EGEE operations is partially distributed already**
- **Challenge to do this with less effort and in 2 years; no place for duplication, loose initiatives**
- **Collaboration is essential; we need an agreed vision as input to EGI, and we need to work together towards this vision**
- **Site responsibility for daily operations is the best way of saving effort and simplifying operations at all levels!**
  - We need to provide the tools to facilitate this
  - We need more site involvement in Working Groups
  - Site and ROC/NGI partnership should be reinforced