

WLCG Collaboration Workshop
21 - 25 April 2008, CERN

¿Future Operations?

Jamie.Shiers@cern.ch



Worldwide LHC Computing Grid
Distributed Production Environment for Physics data Processing



Overview

- It was originally foreseen that Les would give a talk on future challenges / opportunities for WLCG
- But he wisely 😊 decided not to be at CERN this week...
- What follows are some **thoughts** - most of which have been aired before in various fora - that will hopefully **stimulate** some discussion...



EGI / WLCG / EGEE III

- We heard yesterday about the goals and timelines of EGI:
 - The European Grid Initiative Design Study - EGI_DS
 - *Dieter Kranzlmüller*
- as well as requirements from WLCG :
 - On the Transition to EGI - Requirements from WLCG and Related Projects
 - *Ian Bird*
- and finally an overview of EGEE III operations
 - Overview of EGEE III operations
 - *Maite Barroso Lopez*

¿What is EGI Operations?

- To answer this question, we need a much better idea of what “the EGI Grid” will be...

Is it:

- ¿ A large-scale, production Grid infrastructure – build on National Grids that interoperate seamlessly at many levels, offering reliable and predictable services to a wide range of applications, ranging from “mission critical” to prototyping and research?
- ¿ A loosely coupled federation of NGIs with little or no cross-grid activity, heterogeneous and sometimes incompatible middleware stacks, no cross-grid accounting, no need for coordinated operations or management
- ¿ A bit of both?
- ↳ **The conclusion of the Rome EGI_DS Workshop was the former...**

EGI Operations Is To Support:

- A large-scale, production Grid infrastructure – build on National Grids that **interoperate** seamlessly at many levels, offering **reliable** and **predictable** services to a wide range of **applications**, ranging from “**mission critical**” to prototyping and **research**
- It is also understood that there are different types of Grid usage – it is **posited** that these can **co-exist** to each other’s **mutual benefit...**

[Legendary] Grid Classification...

- Grid Computing (potentially) offers value to a wide range of applications, broadly classified as follows:
 - **Provisioned**
 - Large scale, long term “Grand Challenge”
 - e.g. LHC (“space microscopes”), space telescopes,
 - **Scheduled**
 - Require large resources for short periods
 - Far too expensive to provision for a single ‘application’
 - ☛ Not (always) time critical – disaster response?
 - **Opportunistic**
 - Which includes the above but also other areas which are less “real time”
- You can find **numerous** examples of “Mission Critical” applications in each of these categories (e.g. EGEE User Forum!)
 - “Mission Critical” as in “Life or Death”

What's Special about "the Grid"

From the viewpoint of a "large consumer":

- ☺ Grids have proven to be an excellent way of federating resources across computer centres of varying sizes into much larger quasi-homogeneous infrastructures.
- ☺ This matches well with the needs of international science, allowing resources at participating institutes to meet the needs of the entire collaboration.
- ☺ This in turn adds value to the individual sites, leading to a positive feedback situation.

And the EGI Added Value?

- In order to be both attractive and maintainable, Grids need to have the following attributes:
 1. **Low cost of entry;**
 2. **Low cost of ownership.**both in terms **operations** as well as **application** and **user** support
- The basic principles of ***reliability and usability*** must be designed in from the start – adding them later is not consistent with the goals of **low cost of ownership.**

How is this achieved?

- 👉 We should not forget one of the key features of the Grid – **resilience** to failure / **scheduled** downtime of **individual** components and / or sites
- This significant advantage can only be realised through a sufficient degree of **interoperability** & **interoperation**
- 😊 But gives individual NGLs much more **freedom** & **flexibility!**

Key Issues

- **Non-disruptive** & **timely** transition from current Operations scenarios to EGI+NGIs
- Ensuring “**value-for-money**”:
 - **Applications Communities;**
 - **NGIs;**
 - **Funding agencies;**

must all be convinced that any money involved is not only **well** but also **optimally** spent!

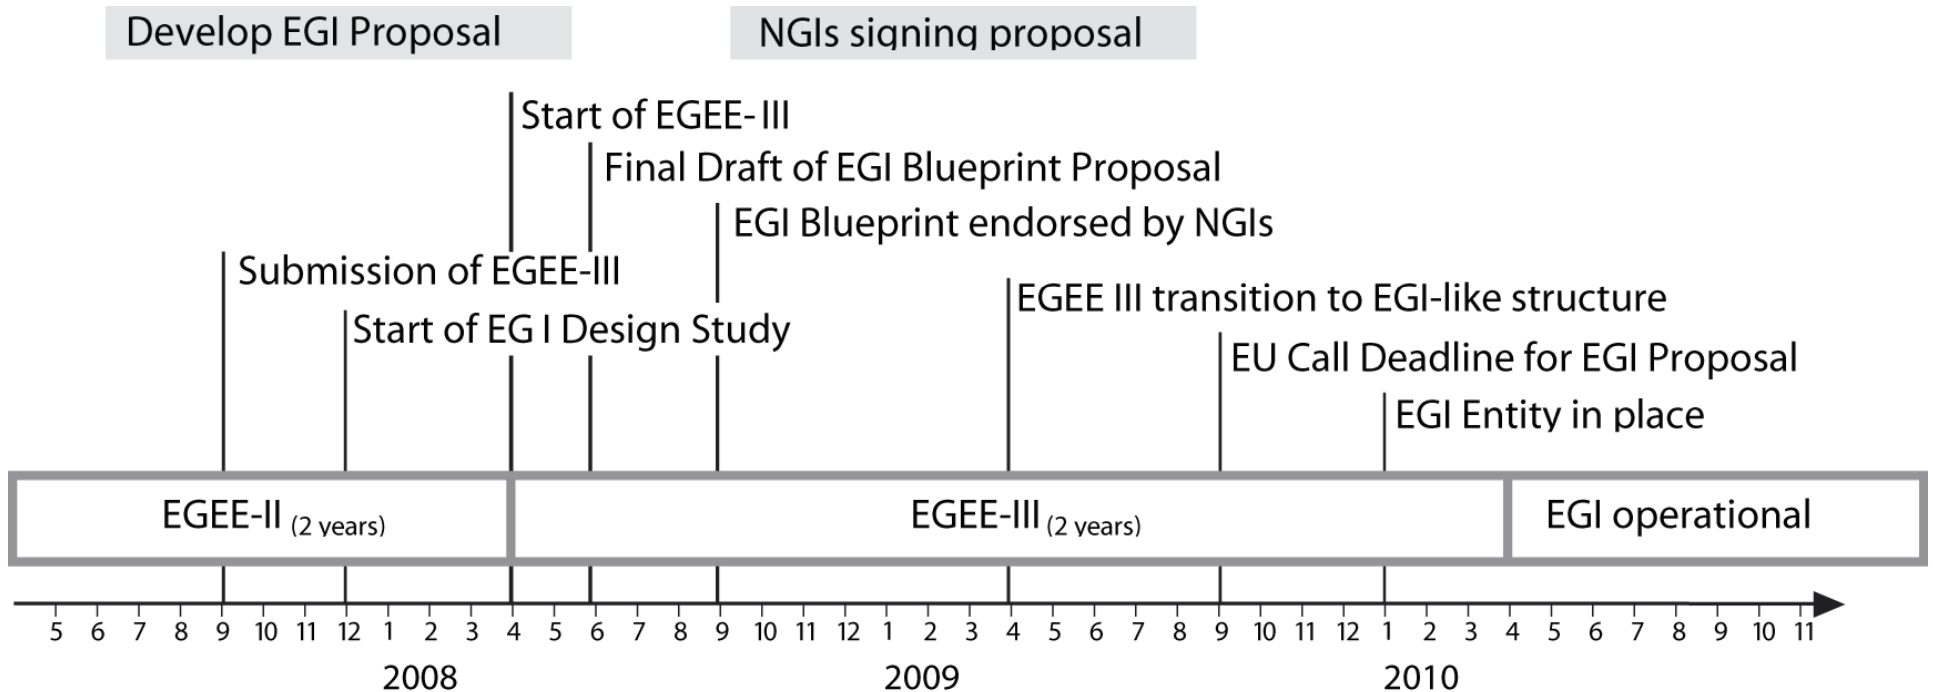
Timeline

- In 2010, the LHC will reach design luminosity
- In 2010, EGEE III will terminate
- **It is inconceivable that we:**
 - a. Don't run the LHC machine
 - b. Run the LHC machine without a computing infrastructure
 - c. Run the computing infrastructure without Grid operations
- This is required for other mission critical applications that are dependant on this infrastructure
- ↳ **The transition to the new scenario must be**
 - a. On time
 - b. Non-disruptive
- This is a fundamental requirement – it is not an issue for discussion

From the DoW...

↳ The establishment of EGI is guided by two basic principles:

1. Build on the experience and successful operation of EGEE and related projects
2. Make EGI operational before EGEE III ends



Key EGI Assumptions

- EGI is **small** – significantly **smaller**, say, than what EGEE (I, II, III) has funded at CERN
- It is **not** located at CERN
- The EGI does **not** run / provide large-scale Grid services / resources
- It is **not** co-located at or near a centre of gravity for any particular application community



The WLCG Experience

- There can be **no** doubt that the close physical proximity / inter-mingling people from the different projects (LCG, EGEE, related, ...) has been **extremely** beneficial during the deployment and hardening phases of WLCG
- This is clearly not **scalable** to large numbers of application communities and may well be in **contradiction** with a “sustainable (long-term?) e-instructure”



The role of EGI

- **Coordination across the NGIs**
 - Operations – overall SLAs, reporting, accounting, reliability, etc.
 - **Cross NGI operations issues should be an agreed process for the NGIs (EGI should broker these processes)**
 - Brokering of resources for applications with the NGIs
 - Operational security coordination – e.g. Incident response
 - Common policy brokering

↳ Support for international VO's (like WLCG) – should they really negotiate with 35 NGIs?

- **Integration/certification/testing of middleware**
 - Whatever this means – many different stacks will be existing
 - Work on “interoperability” is difficult and slow, but running parallel middleware stacks on a site is also very costly

Key Issues

- **Non-disruptive** & **timely** transition from current Operations scenarios to EGI+NGIs
- Ensuring “**value-for-money**”:
 - **Applications Communities;**
 - **NGIs;**
 - **Funding agencies;**must be convinced that any money involved is not only **well** but also **optimally** spent!

EGI Transition Proposal

- The EGEE Grid is currently used for large-scale production by a number of scientific VOs.
- It will be **unacceptable** to them to have a **disruptive** transition to a different operational grid in EGI.
- The two options are:
 1. Define the EGI model very quickly to allow a smooth transition during EGEE III
 2. **Assume that Day One EGI Operations follow the EGEE model and any subsequent change is evolutionary**
- Given the experience in previous Grid projects, it is presumably too late for the first so **we propose a working assumption of the second.**

Adapted from proposal by John Gordon, hence focus on EGEE. Requirement for smooth and timely transition equally valid for other production Grids!

How to achieve this?

- The EGEE Operational model has three levels: EGEE-wide, Regional, and National
 - Don't forget we already have national duties like CA management.
- The migration to EGI will involve a migration of duties down towards NGIs
- The migration from Central to Regional has started in EGEE III
- **Our proposal is that responsibility for the balance between Regional and National be left to the group of NGIs that make up each existing Region.**
- **They have the joint duty to continue the existing EGEE service in their region.**
- They have the freedom to deliver this any way they choose
 - at one extreme they may decide to continue with the existing ROC and organise its funding internally.
 - at the other they may decide to devolve everything to each NGI
 - More likely is some combination of the two, with some migration from the former to the latter over time.
- Leave this to the regions. They can then progress independently as suits regional and national needs and priorities. EGI defines and monitors the operational service definition to ensure a seamless grid for the users.

Adapted from proposal by John Gordon, hence focus on EGEE. Requirement for smooth and timely transition equally valid for other production Grids!

Core Operations Tasks

- Regional Operations coordination;
- Coordination and support for roll out of mw updates;
- Grid security and incident response coordination;
- Interoperations (OSG, EU related projects);
- Weekly operations meetings and operations workshops;
- Support from mw resident service experts;
- Middleware release support;
- VO Membership Service;
- Service Availability Monitoring;
- User support coordination and the global Grid user support (GGUS);
- Certification authority for various VOs;
- Monitoring;
- Pre-production coordination;
- Triage of incoming problems and assignment of tickets to second line support units

Operations Resources

Core Functions	FTEs
Operation of a reliable Grid infrastructure	13
Accounting, reporting & Monitoring	3
VO Management	1
Pre-production coordination	1
Help Desk	3
Network Co-ordination	2

- Resource estimation from draft document for EGI_DS deliverable 5.1
- Needs to be compared with “reality” - i.e. what was on the floor in EGEE II / III
- Transition can be expected to require additional resources!


And the EGI Added Value?

- In order to be both attractive and maintainable, Grids need to have the following attributes:

1. **Low cost of entry;**
2. **Low cost of ownership.**

both in terms operations as well as **application and user support**

- ☹ Currently, adapting an existing application to the Grid environment is a non-trivial exercise that requires an in-depth understanding not only of the Grid computing paradigm but also of the computing model of the application in question.
- One reason for the success of the Application Support team at CERN has been the very close physical proximity of:
“the highest level of middleware expertise”
- 👉 **I would also like to add the enormous enthusiasm and dedication of the people involved!**



Summary

- I do not claim to have resolved *any* issues - or even raised the relevant ones
- However, it is clear that the transition from today's **WORKING** Grid to a *future, multi-disciplinary, low-cost-of-entry/ownership, NGI-based Grid* with a minimal set of functions / responsibilities at the “**EGL level**” will take some time to plan and many years to execute
- **We should be proud of what we have achieved - with a not insignificant amount of effort**
- We should build on this - and evolve it - into a future Grid that provides **more** functionality to **larger** numbers of application communities at **lower** cost
- **And not change for the sake of change...**

CCRC'08