



LCG Deployment – Introduction

Ian Bird

Deployment Sessions
LCG Workshop
23 March 2004



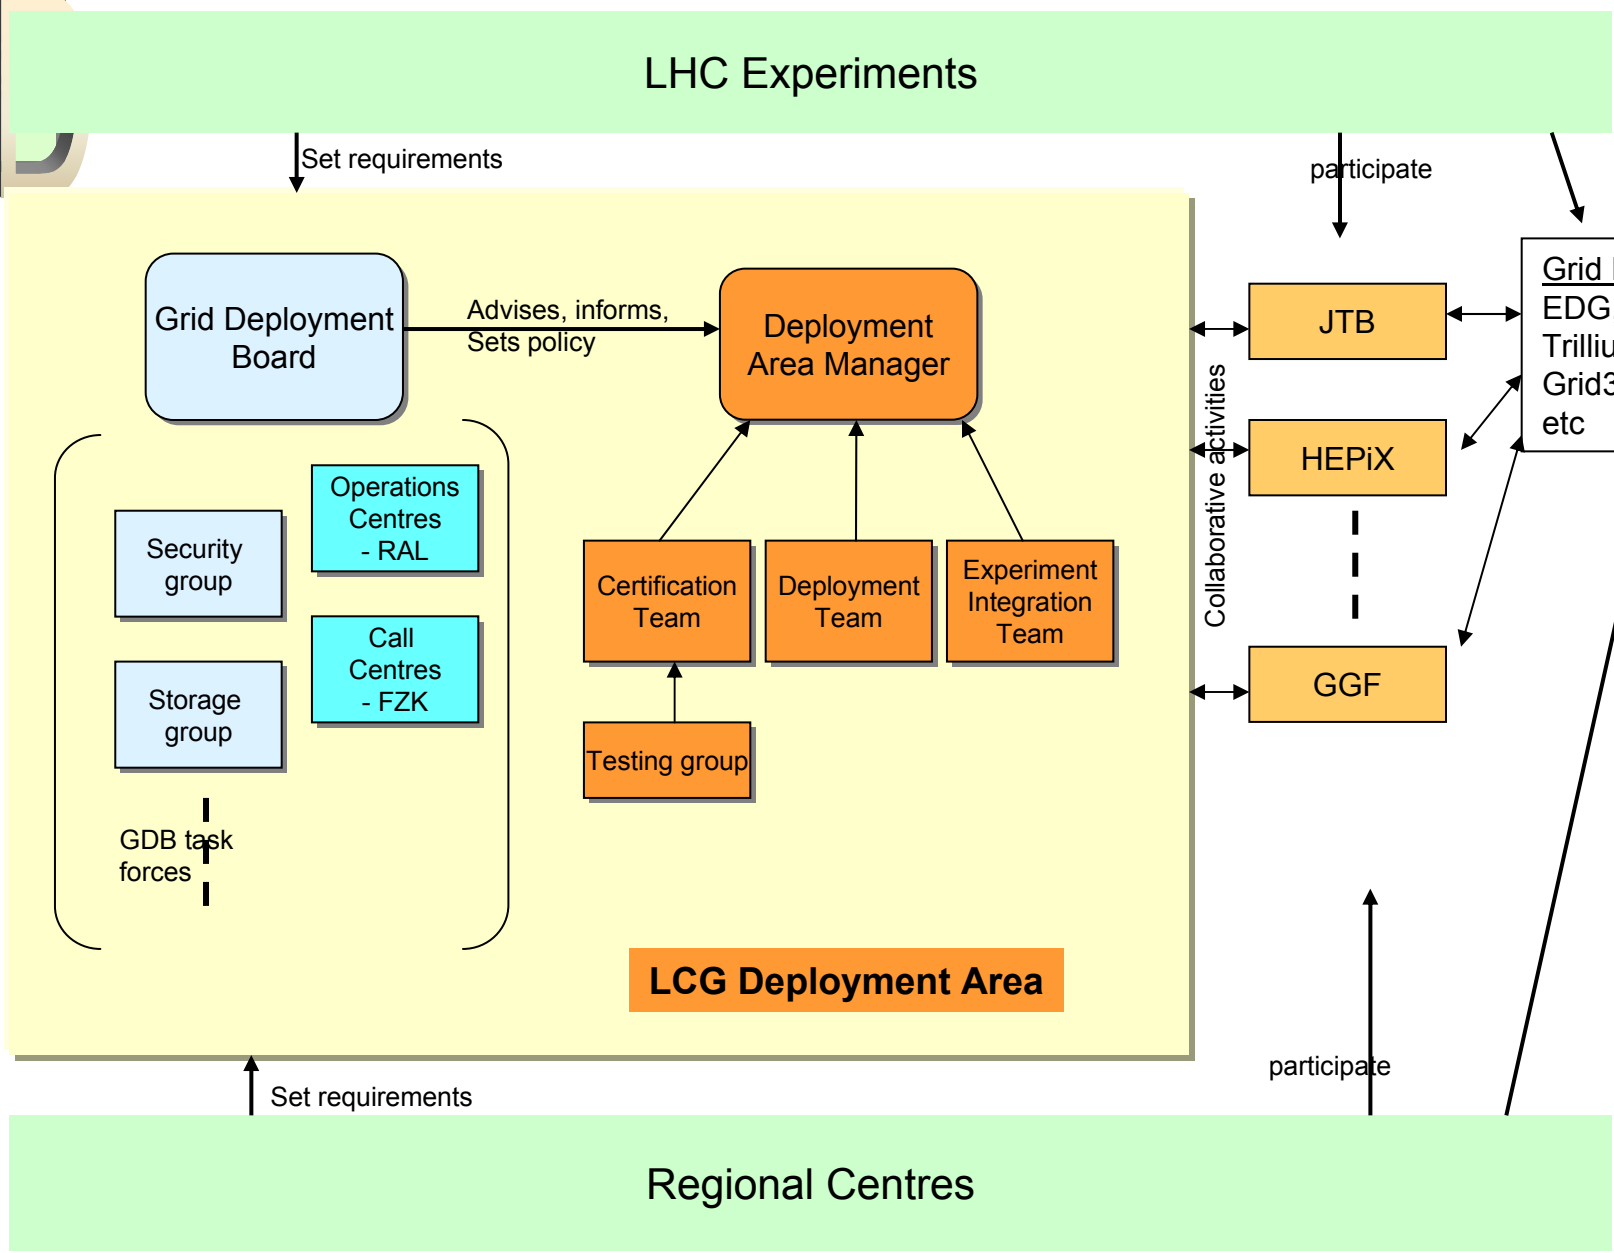
Overview

- Introduction and goals
- LCG Deployment Organisation
- Deployment plan – how we got here
 - WG1 – LCG-0 – LCG-1 – LCG-2
- LCG-2 Core sites – what this means
- Issues
 - Operational effort, installation, packaging, ...




Introduction & Goals

- Timing:
 - At the point of moving from testing to production
 - Data challenges have started;
 - Have a reasonably stable set of middleware
 - Change of focus from integration/certification towards operations
 - Good time to reflect on what needs to be addressed for the future
 - Inconvenient as difficult to organise during DC start-up activities
- Deployment sessions
 - Focus on a few areas that are high priority
 - Bringing sites into LCG-2
 - Storage and data management
 - Information system and monitoring
 - Operations centre
 - Code – release, portability, installation, configuration
- Goals of the workshop for deployment activities:
 - Get more people involved – and gain a common understanding
 - Generate collaborative projects in areas that need effort
 - Bring up issues that need to be discussed more widely



LCG Deployment Organisation and Collaborations



Deployment Activities: Human Resources

<u>Activity</u>	<u>CERN/LCG</u>	<u>External</u>
Integration & Certification	5	External tb sites
Debugging/development/mw support	3	
Testing	3	2 + VDT testers group
Experiment Integration & Support	5	
Deployment & Infrastructure Support	4.5	RC system managers
Security/VO Management	2	Security Task Force
Operations Centres		RAL + GOC Task Force
Grid User Support		FZK + US Task Force
Management	1	
Totals	23.5	

In collaboration

Team of 3 Russians have 1 at CERN at a given time (3 months)

Refer to Security talk

Refer to Operations Centre talk

➤ The GDA team has been very understaffed – only now has this improved with 7 new fellows

➤ There are many opportunities for more collaborative involvement in operational and infrastructure activities



Communication

- Weekly GDA meetings (Monday 14:00, VRVS, phone)
 - Mail-list – project-lcg-gda@cern.ch
 - Open to all – need experiments, regional centres, etc.
 - Technical discussions, understand what priorities are
 - Policy issues referred back to PEB or GDB
 - Experience so far:
 - Experiments join, regional centres don't
 - **NEED** participation of system managers and admins – we need a rounded view of the issues
- Weekly core site phone conference
 - Address specific issues with deployment
- Also at CERN:
 - Weekly DC coordination meetings with each experiment
- GDB meetings monthly
 - Make sure your GDB rep keeps you informed
- Open to ways to improve communication!



Some history

- Reminder – project goals:
 - Prototype computing environment for LHC
 - Focus on building a production-quality service
 - Learn how to maintain and operate a global scale **production** grid
 - Gain experience in close collaboration between regional (resource) centres
 - Understand how to integrate fully with existing computing services
- Learn how to move from test-beds to production services
- Building on the results of earlier research projects;
 - Deploy existing US and EU middleware
- Jan 2003 – GDB agreed to take VDT and EDG components
- March 2003 – LCG-0 – first cut with existing middleware while waiting for EDG 2.0 in Spring
- September 2003 – LCG-1: 3 months late with reduced functionality than expected
- December 2003 – LCG-2: Middleware needed for DC's – but with classic SE
 - Deployed in January; Data challenges started in February;
 - Testing in production!
- March 2004 – Full SRM-based SE



Boundaries of LCG

- Ideally (!):
 - Provide defined grid interfaces to a grid site:
 - Storage, compute clusters, etc.
 - Integration with local systems is site responsibility
 - Middleware layered over existing system installations
- But (real life):
 - Interfaces are not well defined (SRM maybe a first?)
 - Lots of small sites require a packaged solution
 - Including fabric management (disk pool managers, batch systems)
 - That installs magically out of the box
- Strive for the first view, while providing the latter
 - But – “some assembly is required” – it costs effort
- Constraints:
 - Packaging and installation integrated with some middleware
 - Complex dependencies of middleware packages
 - Etc.
- Result: everyone is unhappy




LCG-2 core sites

- Experience with LCG-1 showed:
 - Many sites had configuration problems – cost much effort to resolve – needed more site commitment to operational problems – led to experiment frustration
 - Problem sites caused problems for the whole system
- For LCG-2 GDB agreed:
 - Start with set of core sites, that:
 - Could commit significant compute and storage resources
 - Would guarantee to make the service stable
 - ... in order to get the data challenges started and stabilised
 - Then bring in other sites – one-by-one as each is verified to be correctly configured and stable
 - Leave LCG-1 in place while LCG-2 was stabilised
- Status today:
 - Ready to migrate LCG-1 sites to LCG-2, include new sites that have come since December



LCG-2 Core sites (22 March – 11:00)

- Core sites joined with basic commitments
 - Ramped up since GDB
- At several sites
 - LCG fully integrated with large clusters
 - E.g. at CERN LXBatch is part of LCG-2, available fraction will increase
 - UI is installed as standard user tool in LXplus



Site	LCG-2 CPU	January commitment
CERN	324	400
FZK	144	100
PIC	160	100
FNAL	4	10
CNAF	715	400
Nikhef	250	124
Taipei	98	60
RAL	146	140
Total	1841	1334



Expected Developments in 2004

- General:
 - LCG-2 will be the service run in 2004 – aim to evolve incrementally
 - Goal is to run a **stable** service
- Some functional improvements:
 - Extend access to MSS – tape systems, and managed disk pools
 - Distributed vs replicated replica catalogs – with Oracle back-ends
 - To avoid reliance on single service instances
- Operational improvements:
 - Monitoring systems – move towards proactive problem finding, ability to take sites on/offline; experiment monitoring
 - Continual effort to improve reliability and robustness
 - Develop accounting and reporting
- Address integration issues:
 - With large clusters, with storage systems
 - Ensure that large clusters can be accessed via LCG
 - Issue of integrating with other experiments and apps



Release process

- Monthly coordinated releases
 - Which may or may not be deployed at that moment
 - Cannot stay in the mode where everything is a showstopper – problems must be prioritised
 - Gradually reaching a more stable situation now
- Priorities for future releases – agreed in GDA meetings based on:
 - Experiment experience, problems, and needs
 - Operational experience, stability, robustness, performance



Issues

- Code portability
- Middleware dependencies (OS, packages, 3rd party)
- Build real 24x7 (or ...) operational support
 - GOC at RAL is a start – but – is it enough?
- Operational tools – what real monitors do we need?
 - Move towards proactive behaviour – not just firefighting
- Have to move our focus towards a stable production operation with real support
 - Get away from testbeds
 - Data challenges are running now, but use will grow
- Need to get regional centres more involved in providing and supporting a service