



Grid Deployment Introduction and Overview

Ian Bird
LCG Deployment Area Manager

LHCC Comprehensive Review

22nd November 2004



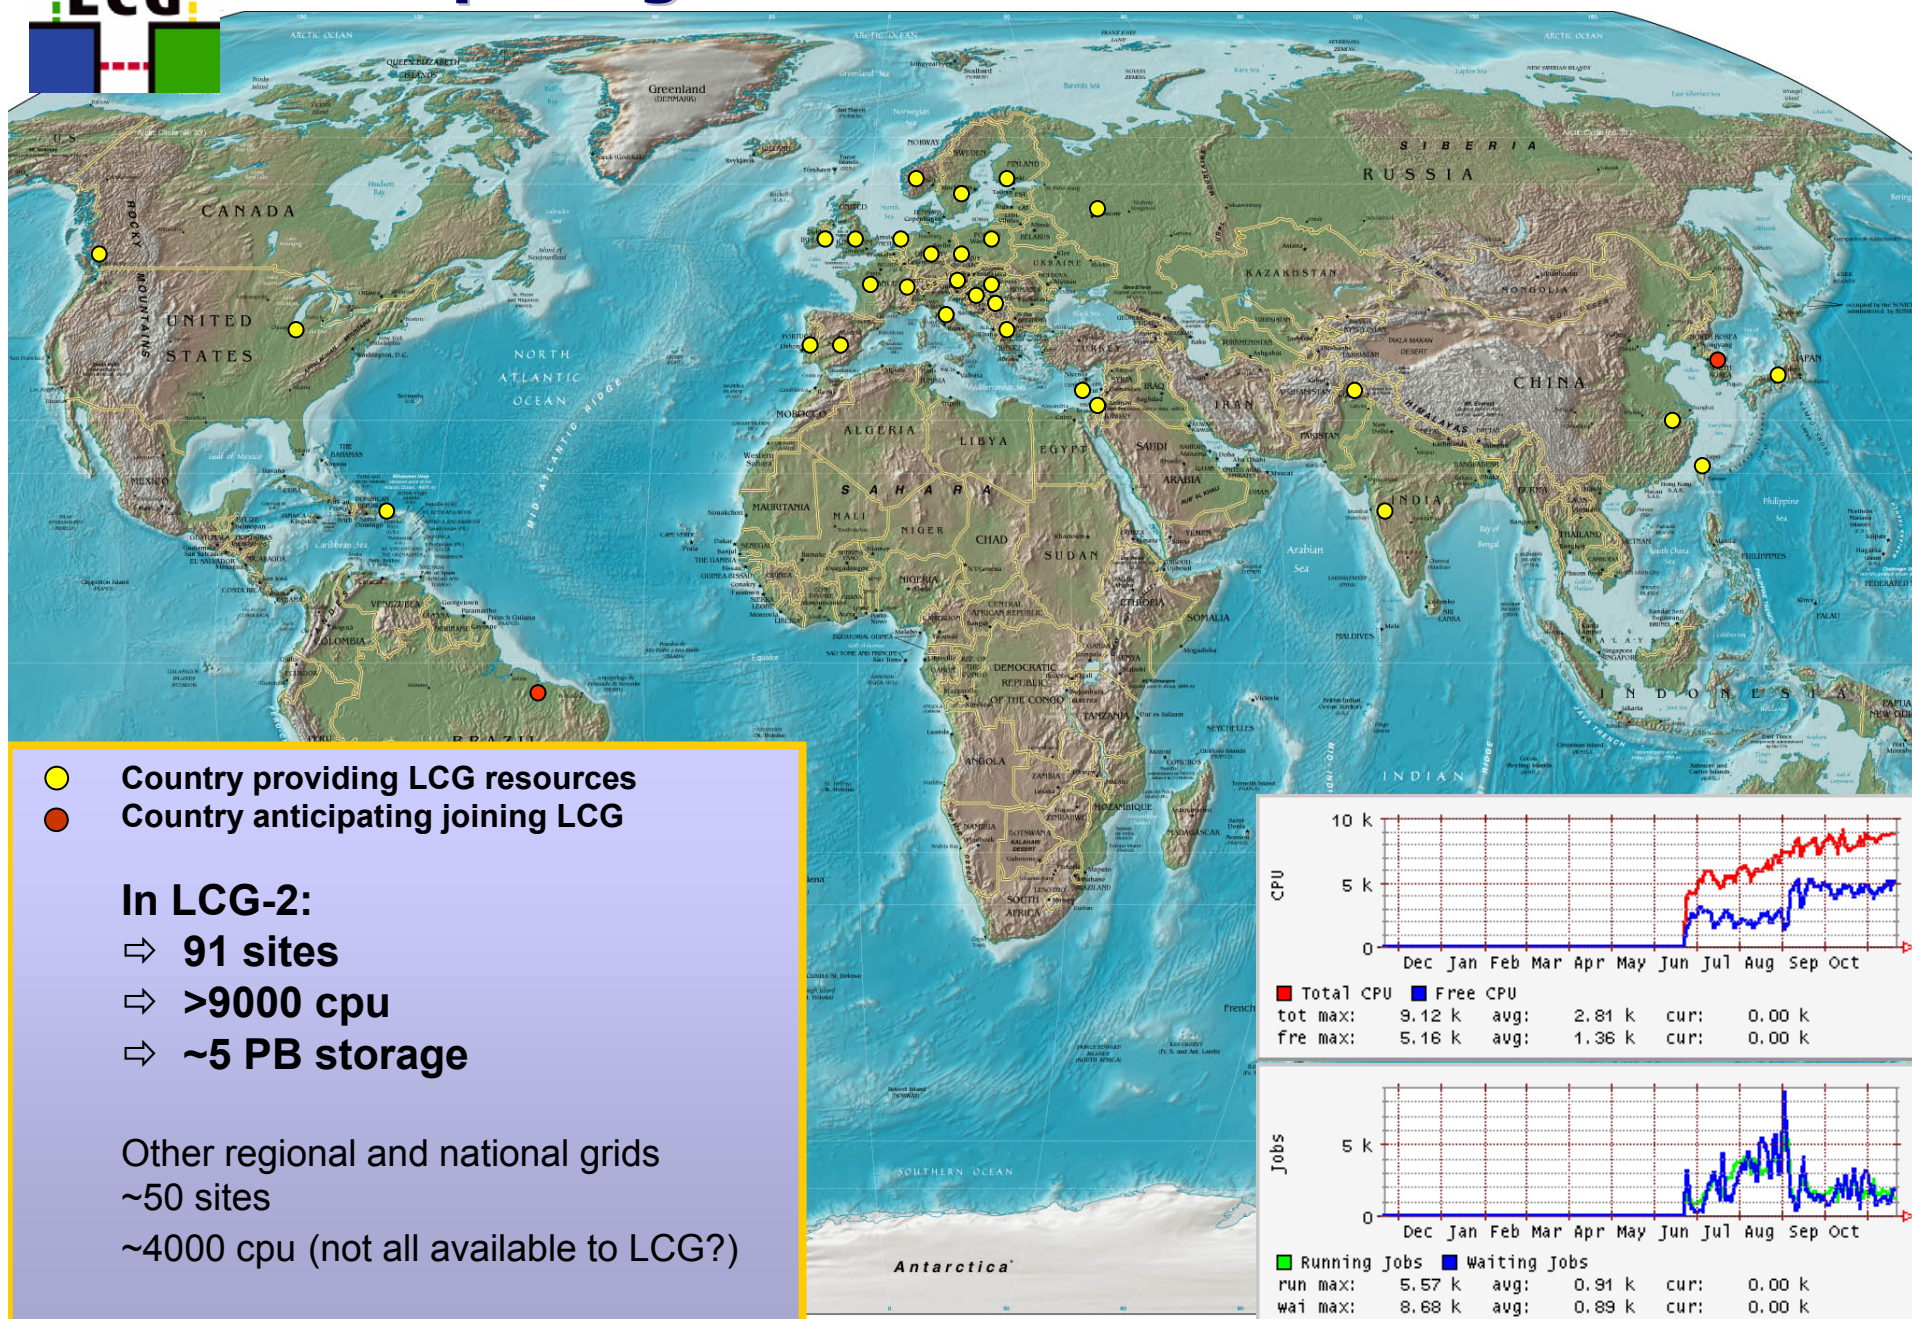


Outline

- General status
- Follow up to 2003 review
- Goals for 2004
- Milestones in 2004
- Activities, Services, migration to EGEE
- Staffing
- Other talks:
 - Operations experience; data challenges and experiment experience; lessons learned and future directions



Computing Resources in LCG – Nov 2004



- Country providing LCG resources
- Country anticipating joining LCG



In LCG-2:

- ⇒ 91 sites
- ⇒ >9000 cpu
- ⇒ ~5 PB storage

Other regional and national grids

~50 sites

~4000 cpu (not all available to LCG?)





Overall status summary

- **Certification, testing and deployment**
 - Has demonstrated how essential this is – key to building a stable service
 - Teams: testing, certification, m/w debugging and fixing, deployment and support
 - Implemented EIS test-bed for experiment verification of releases before production
- **Operations and user support**
 - Was slow in ramping up – lots of issues seen in the data challenges; now becoming a managed process
 - User support was not well defined – relied on mail lists. Now a clear process has been agreed.
- **Security**
 - Active group – joint across LCG/EGEE and OSG



Status – 1

- LCG runs at many sites (>90) now – both large and small
 - Large sites – existing infrastructures – need to add-on grid interfaces, use existing tools, etc.
 - Small sites want a completely packaged, push-button, out-of-the-box installation (including batch system, etc)
 - Satisfying both simultaneously is hard – requires very flexible packaging, installation, and configuration tools and procedures
 - A lot of effort had to be invested in this area
- There are many problems – but in the end we are quite successful
 - Middleware is relatively stable and reliable
 - Many functional issues (e.g. data management) but have progressed to a stable set that is used at many more sites than expected could be supported
 - System is used in production
 - System is reasonably easy to install now – >90 sites
 - Now have a basis on which we can incrementally build essential functionality
- This infrastructure forms the basis of the initial EGEE production service



Status – 2

- Level of complexity anticipated for LHC → $O(100)$ sites
- We already have >90!
- We probably see most of the operations issues already now
- Operations workshop (Nov 2-4) addressed many of these issues
 - And proposed a way forward
- Data challenges over the past 10 months:
 - Probably the first time such a set of large scale grid productions has been done



Data challenges – summary

- Significant efforts invested on all sides – very fruitful collaborations
 - Unfortunately, DCs were first time the LCG-2 system had been used
 - Adaptations were essential – adapting experiment software to middleware and vice-versa – as limitations/capabilities were exposed
 - Many problems were recognised and addressed during the challenges
- Middleware is actually quite stable now
 - But missing functionality, performance
- But – job efficiency is not high – for many reasons
- Started to see some basic underlying issues:
 - Of implementation (lack of error handling, scalability, etc)
 - Of underlying models (workload management, etc.)
 - Perhaps also of fabric services – batch systems ?
- But – single largest issue was lack of stable operations



Follow up from 2003 review

- “Concern that existing m/w is too complex and underdeveloped ... and main risk seems to be lack of product delivery”
 - GDA took over source code for all LCG-2 components other than VDT
 - GDA provides 1st level m/w support team;
 - agreements with original developers and VDT team
 - This proved essential during DC’s to have on-the-spot reactions
- “LHCC considers it very important for the m/w project to ensure tight links ... to the US ... and better collaboration with regional centres”
 - Close relationship between GDA and VDT – agreed use of NSF funding
 - GDB set up group to look at inter-operability with Grid3. Later GDA has activities with Grid3 and is talking to ARC developers



Follow up – 2

- “Concern over resources... regional centres be queried on how ... funds will become available to achieve require capacity”
 - GDB members agreed to provide information 2 quarters in advance. This information was never really provided by all centres.
 - MoU task force has better information on the longer term plan for phase 2.
- “GDB should ensure more detailed technical discussions”
 - GDA weekly coordination meeting.
 - Focus meetings with each experiment weekly during the DC’s.
 - Weekly operations meetings.
- “Installation is too complex”
 - LCG-2 had WN installation independent of any tools.
 - Now all is simplified and scripted



Milestones in 2004

- High-level milestones were:
 - LCG-2 Operational at Core Sites (1 Feb)
 - This was achieved – LCG-2 was installed at the 8 core sites
 - LCG-2 Operational at 30 sites (1 May)
 - By May there were over 30 sites contributing to LCG-2 in May and 60 (with 6000 cpu) by the end of June
 - Now there are over 80 sites involved with ~9000 cpu
 - 50% prototype available (December)
 - The intent was to demonstrate 50% complexity of one of the large experiments
 - In terms of CPU (100K of today's → 25K in 2008)
 - Real complexity is number of sites involved – we are already at the full scale expected (80 vs O(100))
 - However, this is for batch only – analysis has not been addressed
 - Additional complexity comes from multiple grid flavours



Goals and achievements for 2004

- Run and support the data challenges
 - Use experience and feedback to improve the service
- Build experience in operations and user support
 - Learn about running and operating a large scale grid infrastructure
 - Understand issues with integrating grid services into existing production environments
- Set up EGEE operations infrastructure
 - Expand the Grid operations and support
- Middleware:
 - Port to OS other than RH 7.3
 - Simplify installation and configuration
 - Respond to issues raised in the Data Challenges and deployment/operation

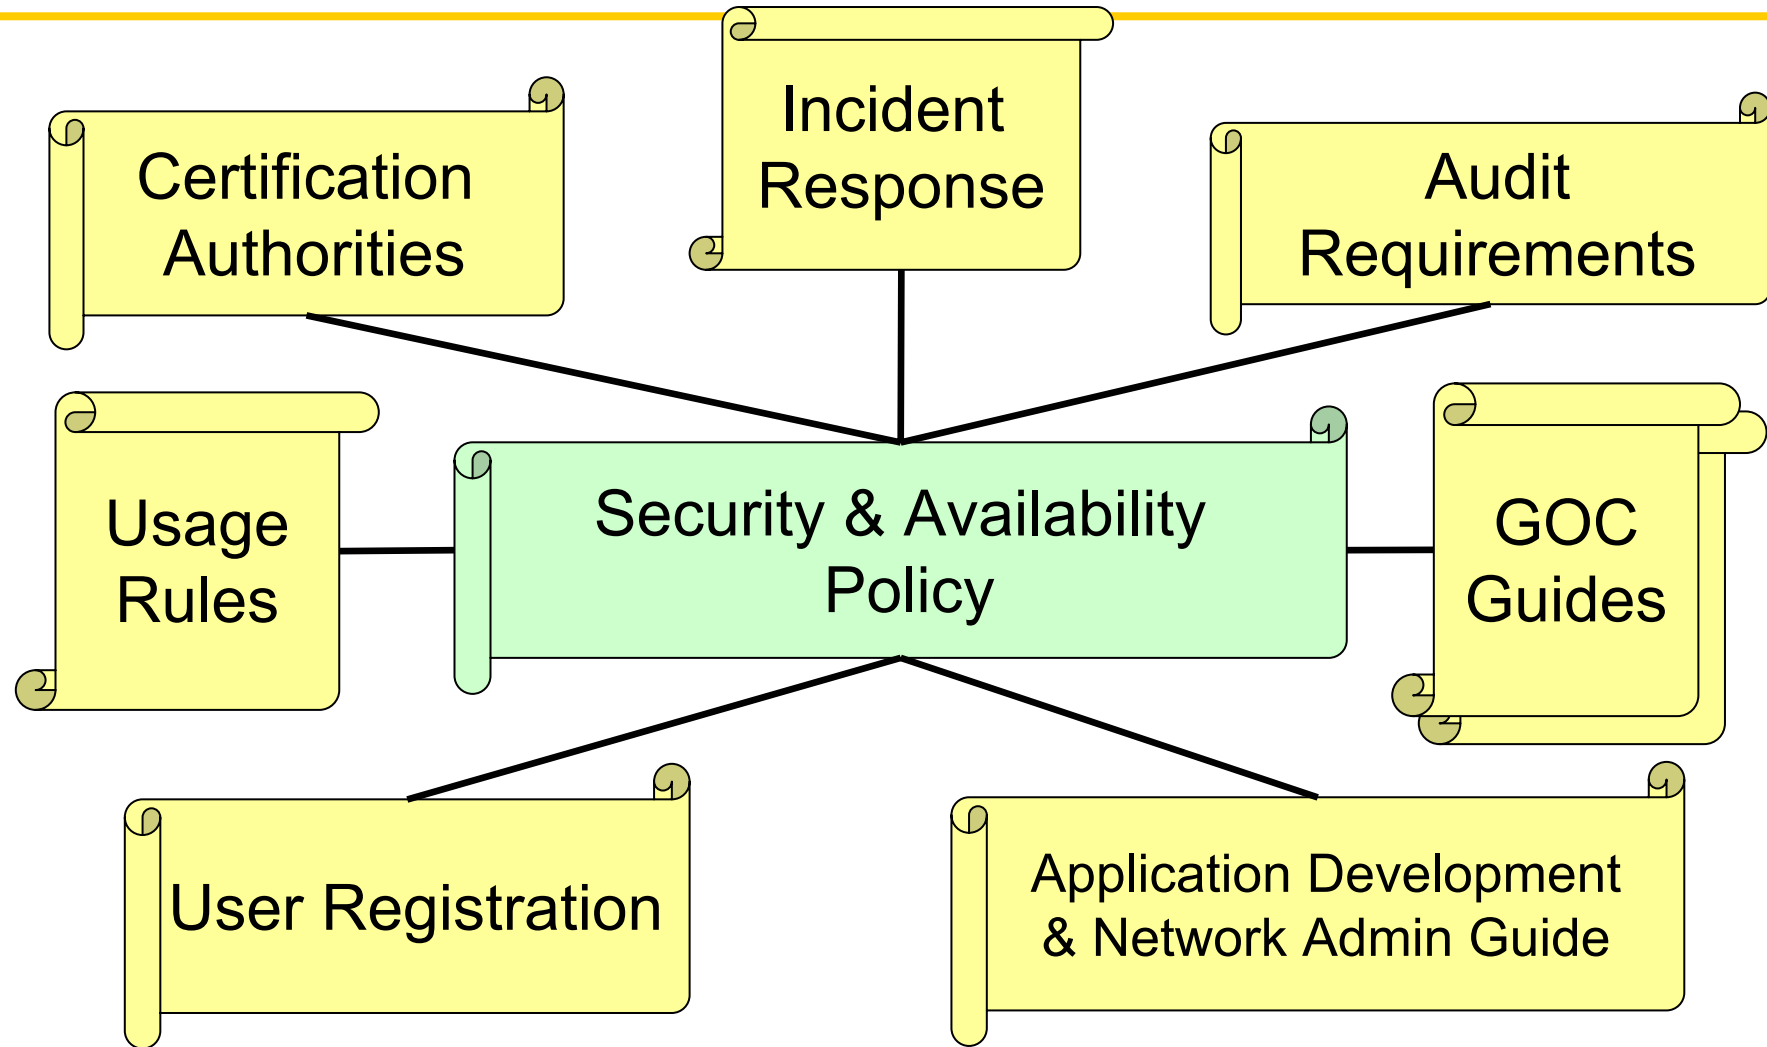


Security achievements


- CA related
 - EUGridPMA charter agreed – now 30 members
 - LCG catch-all established as RA of DOEGrids CA. CNRS as catch-all for EGEE
- Policy group
 - LCG security group → Joint (LCG,EGEE,OSG) policy group
 - Revised VO membership management (user registration), and process using organisational DB (e.g. HRDB at CERN) agreed
 - Revision of acceptable use policy – discussed as a more general AUP for grids
 - First draft of site registration requirements/process
- Operations
 - Operational security coordination effort started
 - Cooperation with OSG on incident handling
 - Planning “service challenges” – walk through of response procedures




Policy – Joint Security Group



<http://cern.ch/proj-lcg-security/documents.html>



LCG → EGEE in Europe



- Operations Management Centre
- Core Infrastructure Centre
- Regional Operations Centre

- **Core Infrastructure Centres (4)**
 - CICs build on the LCG GOC at RAL
 - Also run essential infrastructure services
 - Provide support for other (non-LHC) applications
 - Provide 2nd level support to ROCs
- **Coordination:**
 - At CERN (Operations Management Centre) and CIC for HEP

- **Regional Operations Centres (9)**
 - Act as front-line support for user and operations issues
 - Provide local knowledge and adaptations
- **User Support Centre (GGUS)**
 - In FZK – manage PTS – provide single point of contact (service desk)

- **Taipei provide operations centre, and 2nd instance of GGUS**
 - → start to build round-the-clock coverage
- **Discussions with Grid3/OSG on how to collaborate on ops support**
 - Share coverage?



LCG-2 → EGEE

- **Operations infrastructure**
 - Expand available effort
- **Build pre-production service**
 - In order to have early deployment and experience with new middleware – gLite
 - Allows site managers also to be involved
 - EIS testbed was found as essential – this expands that
 - Certification process expanded to include pre-production feedback and testing
 - Initially PPS will run LCG-2 and deploy gLite components as they arrive
 - Understand compatibility, migration, etc. issues
- **Overhead activities**
 - EU deliverables, processes, etc. take considerable effort



Interoperation

- Reality is that experiments need to use LCG-2/EGEE, Grid3, and NorduGrid
 - Security and policies – has been joint group from the start
 - Information system – try and align Grid3/LCG-2 IS schema
 - Both use GLUE schema – but extensions and interpretations differ
 - Goal of job submission across infrastructures
 - Harder for Nordugrid as their schema is very different
 - Storage interfaces – agreed on SRM as standard – Grid3 and LCG-2 storage elements are compatible and demonstrated
 - Agreement on common accounting interfaces
 - Share many ideas and problems on operations support – see similar issues
 - LCG and OSG operations workshops
 - Canada has successfully built a gateway from Triumpf into Grid Canada and West Grid
 - One way only – jobs from LCG-2 can run there



Interoperation – 2

- In addition other grid activities (SAM-Grid, Babar, ...) need to co-exist at LCG-2 sites
 - Much effort was invested to simplify the WN software and installation (and will become simpler still)
 - Coexistence demonstrated at many sites (Nikhef, FZK, RAL, other UK, Italy Tier 2's)
 - Limitation of compatibility of OS, compilers etc. required by the experiments
- Issues of interoperation are addressed by the GDB
 - High level working group set up last year did not start, but activities are happening at technical level



Staffing

- Changes in 2004:
 - 8 EGEE-funded staff – of which 2.5 FTE are project overhead
 - 2 INFN-funded LCG fellows left (certification team)
 - 1 more departure expected in February

Planned Human Resources required at CERN
Assumes EGEE phase 2 provides staff at current level

	2004	2005	2006	2007	2008
M/W development and support	4	3.2	3	3	3
M/W test, certification and deployment	11.2	11.2	12.2	12	9
Experiment Integration & Support	5	5	5	5	5
Infrastructure coordination and operations	10.6	12.6	12.6	13.6	10