



Operational Issues from Data Challenges

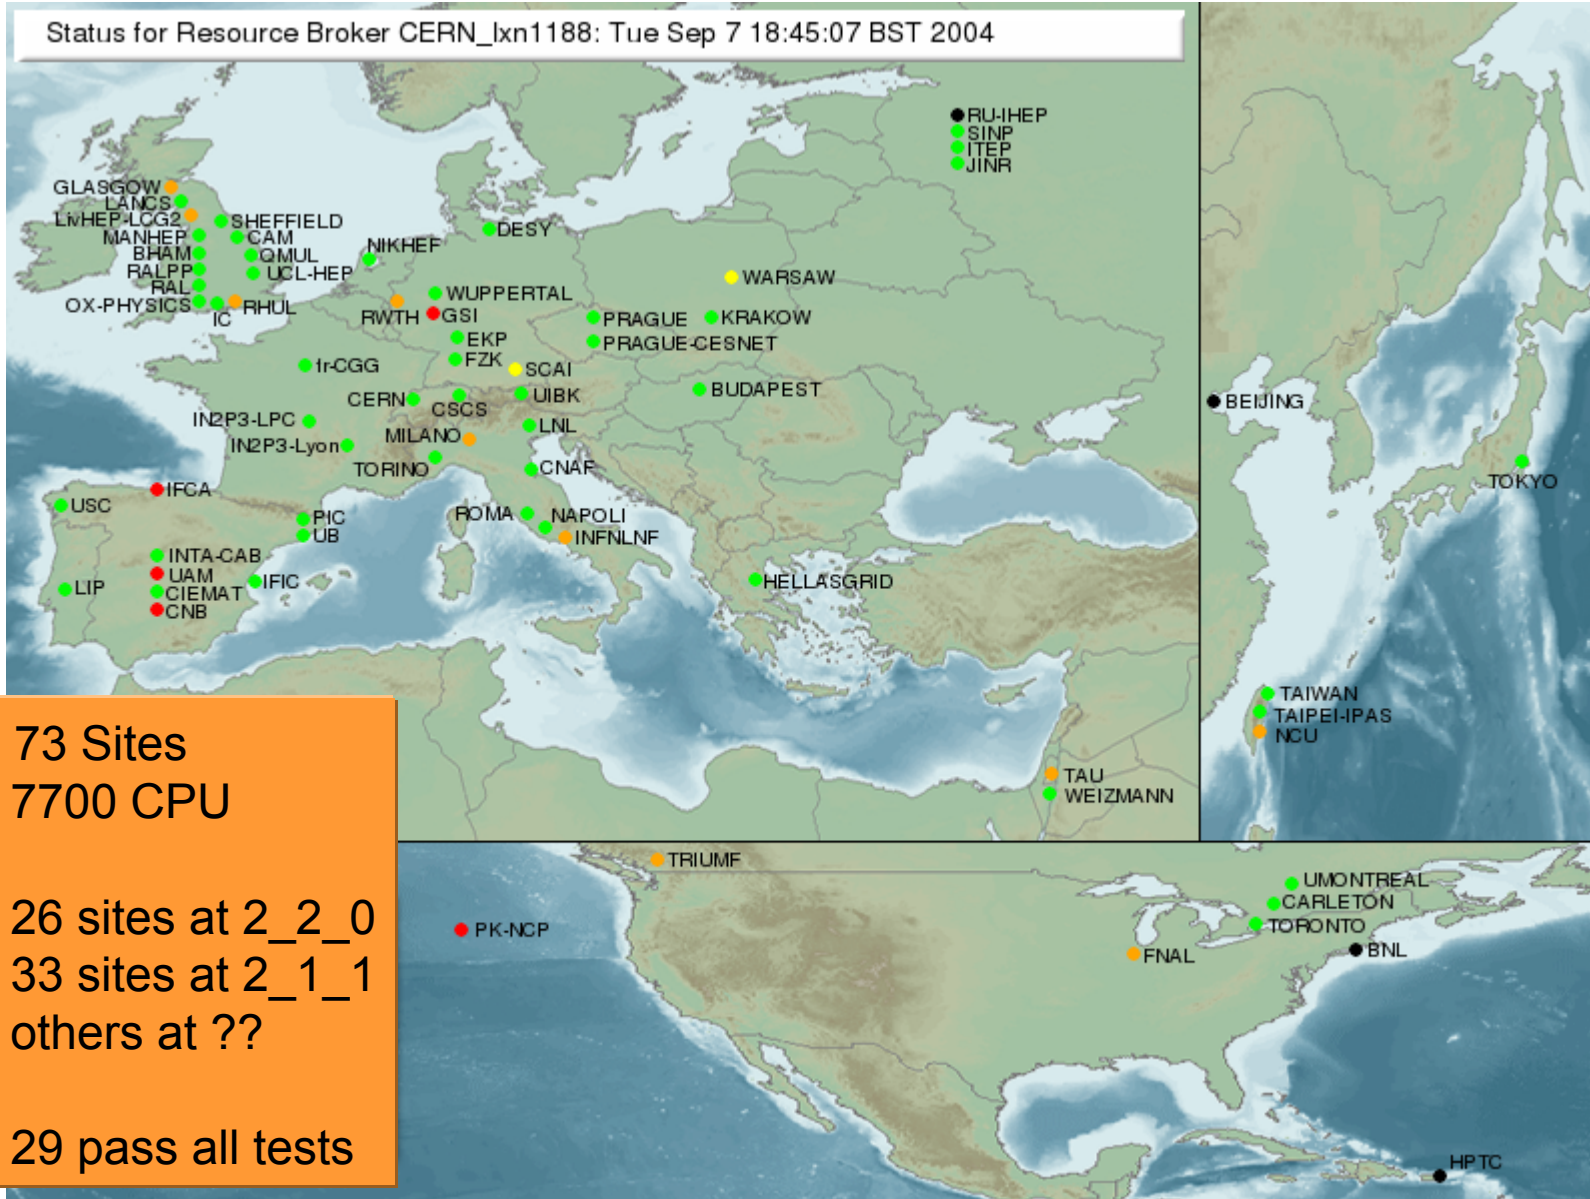
Ian Bird
IT Department, CERN

GDB – CERN
8 September 2004



Current LCG-2 sites: 7/9/04

Status for Resource Broker CERN_lxn1188: Tue Sep 7 18:45:07 BST 2004



- 73 Sites
- 7700 CPU
- 26 sites at 2_2_0
- 33 sites at 2_1_1
- others at ??
- 29 pass all tests



Outstanding middleware issues

➤ See

- <https://edms.cern.ch/file/495809/0.4/Broker-Requirements.pdf>
- GAG summary document (ref?)
- Provides summaries of middleware related issues – much has been discussed previously
- Important: 1st systematic confrontation of required functionalities with capabilities of the existing middleware
 - Some can be patched, worked around, but most has to be direct input as essential requirements to gLite and future developments
 - Some are fundamental problems with underlying models and architectures

➤ Middleware:

- Not perfect but quite stable
- Much has been improved during DC's – a lot of effort still going into improvements and fixes
 - Big hole is missing space management on SE's

➤ Largest problem now is stable operations and providing status information and useful tools to users



Operational issues (selection)

- Slow response from sites
 - Upgrades, response to problems, etc
 - Problems reported daily – some problems last for weeks
- Lack of staff available to fix problems
 - All on vacation, ...
- Misconfigurations (see next slide)
- Lack of configuration management – problems that are fixed reappear
- Lack of fabric management
 - Is it GDA responsibility to provide solutions to these problems? If so, we need more available effort (see slide on workshops etc)
- Lack of understanding (training?)
 - Admins reformat disks of SE ...
- Firewall issues –
 - often no good coordination between grid admins and firewall maintainers
- PBS problems
 - Are we seeing the scaling limits of PBS?
- Forget to read documentation ...



Site (mis) - configurations

- Site mis-configuration was responsible for most of the problems that occurred during the experiments Data Challenges. Here is a non-complete list of problems:
 - – The variable VO <VO> SW DIR points to a non existent area on WNs.
 - – The ESM is not allowed to write in the area dedicated to the software installation
 - – Only one certificate allowed to be mapped to the ESM local account
 - – Wrong information published in the information system (Glue Object Classes not linked)
 - – Queue time limits published in minutes instead of seconds and not normalized
 - – /etc/ld.so.conf not properly configured. Shared libraries not found.
 - – Machines not synchronized in time
 - – Grid-mapfiles not properly built
 - – Pool accounts not created but the rest of the tools configured with pool accounts
 - – Firewall issues
 - – CA files not properly installed
 - – NFS problems for home directories or ESM areas
 - – Services configured to use the wrong BDII
 - – Wrong user profiles
 - – Default user shell environment too big



Addressing operations

- Weekly operations meeting
 - Evolution of GDA meeting
 - Expect (templated) written reports from ROCs, Tier 1's, hopefully also from applications
 - This is EGEE/LCG operations – hope for input from Grid3
- Grid operations ½ day at HEPiX in October
 - Address common issues, experiences
 - Provide some input to:
- Operations and Fabric Workshop
 - CERN 1-3 Nov
 - A goal is to agree an operations model for the next year, and understand what 24x7 operation means in an 8x5 “best-effort” world
 - N.B. EGEE has promised 24x7 operations!
 - Hope to get senior site managers present to agree this model



Operations effort

- The available effort for operations from EGEE is now ramping up:
 - LCG GOC (RAL) → EGEE CICs and ROCs, + Taipei
 - Hierarchical support structure
 - Regional Operations Centres (ROC)
 - One per region (9)
 - Front-line support for deployment, installation, users
 - Core Infrastructure Centres (CIC)
 - Four (+ Russia next year)
 - Evolve from GOC – monitoring, troubleshooting, operational “control”
 - “24x7”
 - Also providing VO-specific and general services
- This is where main focus of effort is going now



Status of scientific linux port

- Worker node port is in next public release
 - Already passed certification
- Full SL3 port will be finished by end September
- Need to be able to support
 - Service nodes on RH73, WN on SLC3 (now)
 - Service nodes on SLC3, WN on RH73 (expect migration need)
 - Addresses security issue, allows farm migrations
- IA64 port has been done by openlab
 - Integrating their work into distribution
- NB: LCFGng is not ported to SLC3
 - Concentrate on manual installation (with automating scripts)
 - Provide quattor components built by CERN and others



Summary

- LCG-2 services have been supporting the data challenges
 - Many middleware problems have been found – many addressed
 - Middleware itself is reasonably stable
- Biggest outstanding issues related to providing and maintaining stable operations
- Has to be addressed in large part by management buy-in to providing sufficient and appropriate effort at each site
- Future middleware has to take this into account:
 - Must be more manageable, trivial to configure, etc
 - Management and monitoring must be built into services from the start