
# Lessons from Data Challenges; Planning for next steps

Ian Bird

GDB - Barcelona
18 May 2004


# Overview

- Situation with LCG-2 now
- Data challenges – what was learned?
- Planning for 2004
  - Service challenges
  - Resources
  - Management (this afternoon)

# LCG-2 today

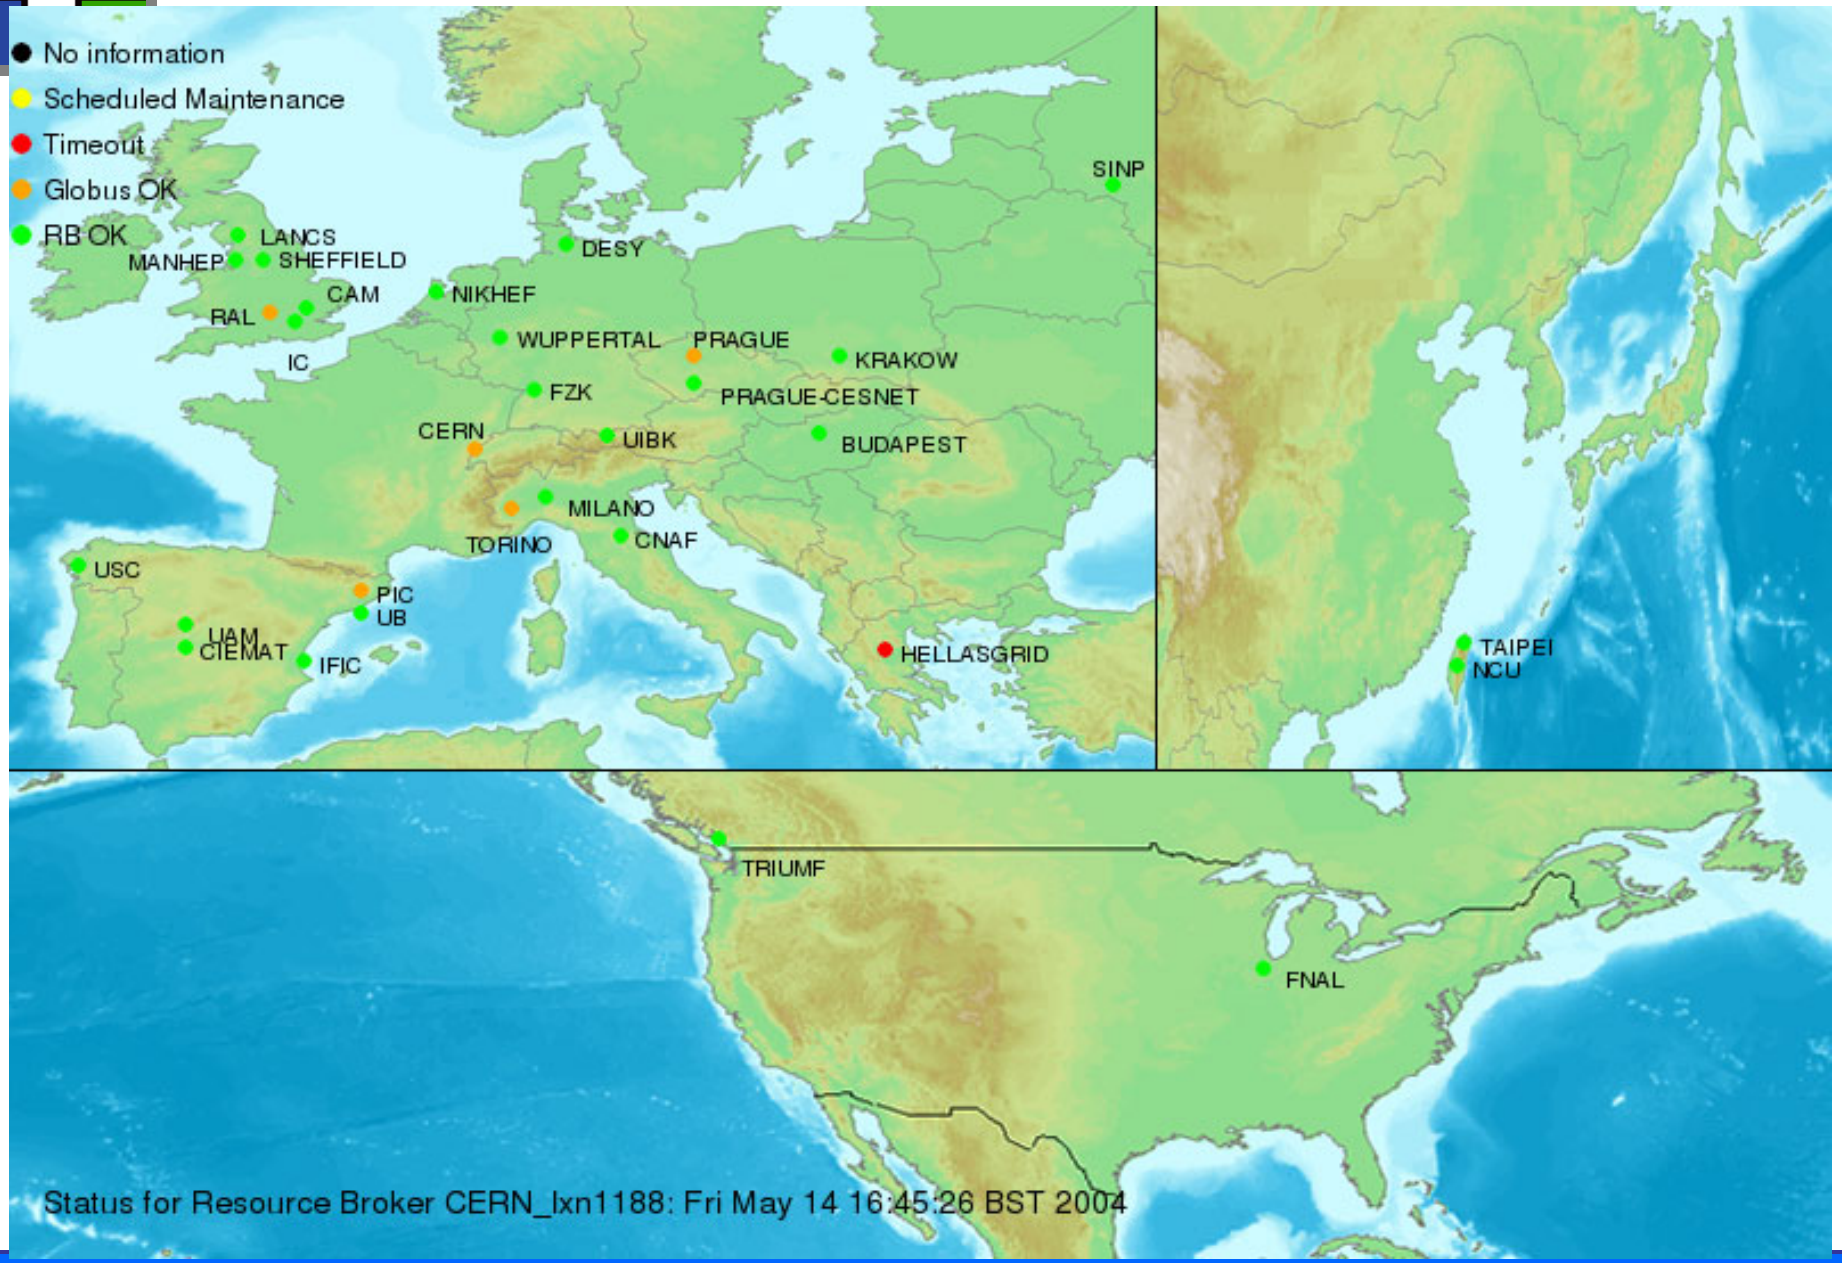
No information
Scheduled Maintenance
Timeout
Globus OK
RB OK
LANCS
MANHEP
SHEFFIELD
CAM
RAL
IC
NIKHEF
DESY
SINP
WUPPERTAL
PRAGUE
KRAKOW
FZK
PRAGUE-CESNET
CERN
UIBK
BUDAPEST
MILANO
TORINO
CNAF
USC
PIC
UB
UAM
CIEMAT
IFIC
HELLASGRID
TAIPEI
NCU
TRIUMF
FNAL
Status for Resource Broker CERN_lxn1188: Fri May 14 16:45:26 BST 2004

# Regional Centres Connected to the LCG Grid

*07-May-04*

| *country* | *centre* |
|---|---|
| Austria | UIBK |
| Canada | TRIUMF, Vancouver |
| | Univ. Montreal |
| | Univ. Alberta |
| Czech Republic | CESNET, Prague |
| | University of Prague |
| France | IN2P3, Lyon** |
| Germany | FZK, Karlsruhe |
| | DESY |
| | University of Aachen |
| | University of Wuppertal |
| Greece | GRNET, Athens |
| Holland | NIKHEF, Amsterdam |
| Hungary | KFKI, Budapest |
| Israel | Tel Aviv University** |
| | Weizmann Institute |
| Italy | CNAF, Bologna |
| | INFN, Torino |
| | INFN, Milano |
| | INFN, Roma |
| | INFN, Legnaro |
| Japan | ICEPP, Tokyo** |
| Poland | Cyfronet, Krakow |
| Portugal | LIP, Lisbon |
| Russia | SINP, Moscow |
| Spain | PIC, Barcelona |
| | IFIC, Valencia |
| | IFCA, Santander |
| | University of Barcelona |
| | Uni. Santiago de Compostela |
| | CIEMAT, Madrid |
| | UAM, Madrid |
| Switzerland | CERN |
| | CSCS, Manno** |
| Taiwan | Academia Sinica, Taipei |
| | NCU, Taipei |
| UK | RAL |
| | Cavendish, Cambridge |
| | Imperial, London |
| | Lancaster University |
| | Manchester University |
| | Sheffield University |
| | QMUL, London |
| USA | FNAL |
| | BNL** |

** not yet in LCG-2

**> 40 sites**
**> 3,100 CPUs**

Core sites:
~1800-2000 CPU

**Centres in process of being connected**

| *country* | *centre* |
|---|---|
| China | IHEP, Beijing |
| India | TIFR, Mumbai |
| Pakistan | NCP, Islamabad |

*Hewlett Packard to provide "Tier 2-like" services for LCG, initially in Puerto Rico*

# Data challenges – so far

- Resources
  - CPU available – Alice could not fully utilise – storage limitations
  - Disk available – mostly very small amounts
    - Need:
      - Plan space vs cpu at a site
      - Ensure that commitments are provided
    - To some extent not requested – delay in dcache SE – asked not to commit all to classic SE's as expected/worried about migration
  - Alice and CMS –
    - Number and size (small) of files:
      - limitations of existing Castor system, also problems in Enstore/dCache
- CPU is mostly in cores sites
  - At the moment (most of) the other sites have relatively few cpu assigned

# Data challenges – 2

- Services:
  - LCG-2 services (RB, BDII, CE, SE etc) have been extremely reliable and stable
    - Even RLS was stable (other issues)
  - BDII has been extremely reliable
    - Provided to experiments – allowed them to define a view of the system
- Software deployment system works
  - Needs some improvement – esp for sites with no shared filesystem
- Information system
  - Schema does not match batch system functionality
  - Information published (job slots, ETT, etc.) does not reflect batch system
  - Solve with CE per VO, need to improve/adapt schema (?)

# RLS Issues

- **RLS performance was much discussed during the data challenge**
- **In general, inserting information was slow**
  - Troublesome for both T0 and replica registrations
  - We more or less survived, higher data rates would be troublesome
  - Several workarounds or solutions produced in-course
    - EDG-based tools, POOL improvements, parallelism
- **Looking up file information by GUID seems sufficiently fast**
- **Queries, including bulk queries by GUID, take a long time**
- **We never really exercised replicated RLS**
- **Several discussions with developers**


# RLS issues

- RLS performance was biggest problem
- Many fixes made during challenge:
  - CLI tools based on C++ API in place on Java tools
  - Added support for non-SE entries
  - Additional tools (register with existing guid)
  - Case sensitivity
  - Performance analysis – usage of metadata queries
  - Lack of bulk operations
  - No support for transactions
- Still unresolved service performance issue (see degradation) –
  - seems to be server related
- No data loss or extended service downtime
- Replication tests with CNAF
  - Not really tested by CMS

# RLS – cont.

- Many of above issues addressed in version currently being tested
- Preparing a note describing proposed improvements for discussion: e.g.
  - Combine RMC and LRC into single db to allow db to optimise and join
  - Resolve issues found in data challenges
  - Model for replicated/distributed catalogs?
- Is the model of metadata appropriate?
  - Experiment vs POOL vs RLS
- With DB group continue to investigate Oracle replication

# What is missing? (functionally)

- A full storage element
  - dCache has had many problems

    - Nice features of SRM gave users a lot of convenience:
      **- auto directory creation;**
    - We were able to continue improving our setup during DC04:

      **- The biggest performance gain was: Michael and his team in DESY developed a new module that reduces the delegated proxy's modulus size in SRM and speeds up the interaction between SRM client and server 3.5 times;**

      **(From CMS FNAL team, based on work done by deployment group)**

  - Packaged - to be deployed
  - Is dCache sufficient/the only solution?
- Demonstrated integration of Tier 1 MSS's

# What is missing? (operationally)

- Assumption of real operational management by GOCs
  - A lot of work on basics has been done – but need problem management
- User call centre
  - Lack of take-up
  - Propose FZK team come to CERN for 2 days to really sort this out
- Accounting:
  - Critical – we have no information about what has been used during the DC's – important for us and for the experiments
- Monitoring:
  - Grid: lack of consistency in what is presented for each site
  - Experiments: we must put R-GMA in place (at least)

# How to proceed?

- Set up service challenges
  - Involves development of services
  - In conjunction with EGEE developers – but not waiting for a new service
- Set up management group within GDA
  - Managers from large centres
  - Provide resources

# Service Challenges

- Purpose
  - Understand what it takes to operate a real grid service – run for days/weeks at a time (outside of experiment Data Challenges)
  - Trigger/encourage the Tier1 planning – move towards real resource planning – based on realistic usage patterns
    - How does a Tier 1 decide what capacity to provide?
    - What planning is needed to achieve that?
    - Where are we in this process?
  - Get the essential grid services ramped up to needed levels – and demonstrate that they work
  - Set out milestones needed to achieve goals during the service challenges
- NB: This is focussed on Tier 0 – Tier 1/large Tier 2
  - Data management, batch production and analysis
- By end 2004 – have in place a robust and reliable data management service and support infrastructure and robust batch job submission


# Service challenges – examples

- Data Management
  - Networking, file transfer, data management
  - Storage management and interoperability
  - Fully functional storage element (SE)
- Continuous job probes
  - Understand limits
- Operations centres
  - Accounting, assume levels of service responsibility, etc
  - Hand-off of responsibility (RAL-Taipei-US/Canada)
- "Security incident"
  - Detection, incident response, dissemination and resolution
- User support
  - Assumption of responsibility, demonstrate staff in place, etc
- VO management
  - Robust and flexible registration, management interfaces, etc
- Etc.

# Data Management – example

- Data management builds on a stack of underlying services:
  - Network
  - Robust file transfer
  - Storage interfaces and functionality
  - Replica location service
  - Data management tools

# Data management – 2

- Network layer:
  - Proposed set of network milestones already in draft
    - Network and fabric groups at CERN – collaborate with (initially) “official” Tier 1’s
    - Dedicated private networks for Tier 0 → Tier 1 “online” raw data transfers

- File transfer service layer:
  - Move a file from A to B, with good perfomance and reliability
  - This service would normally only be visible via the data movement service
    - Only app that can access/schedule/control this network
  - E.g. of this is gridftp, bbftp, etc.
  - Reliability – the service must detect failure, retry, etc.
  - Interfaces to storage systems (SRM)

- The US-CMS/CERN “Edge Computing” project might be and instance of this layer (network + file transfer)

# Data management – 3

- Data movement service layer:
  - Builds on top of file transfer and network layers
  - To provide an absolutely reliable and dependable service with good performance
  - Implement queuing, priorities, etc.
  - Initiates file transfers using file transfer service
  - Acts on application's behalf – a file handed to the service will be guaranteed to arrive
- Replica Location Service:
  - Makes use of data movement
  - Should be distributed:
    - Distributed/replicated databases (Oracle) with export/import to XML/other db's?
    - RLI model?


# Job probes – example

- Continuous flood of jobs
  - Fill all resources
  - Use as probes – test if they can use the resources
    - Data access, cpu, etc
  - Understand limitations, bottlenecks of the system
    - Baseline measurement, find limits, build and improve
- This might be a function of the GOC
  - Overseen by RAL-Taipei-+ collaboration ?
- A challenge might run for a week
  - Outside of experiment data challenges
  - In parallel (or part of) data management or other challenges


# Summary

- Service challenges:
  - Understand what it really takes to operate reliable and performant services
  - Put in place underpinnings of a reliable infrastructure by the end of the year
- Requires:
  - Agreed milestones
  - Commitment of resources and people


# Evolution of GDA Management

Ian Bird

GDB - Barcelona
18 May 2004

# Introduction

- Many different activities must satisfy goals during the next year
  - These necessarily involve many participants – especially Tier1's and large Tier 2's
- Not just "middleware/application" services, but also:
  - Operations centre functionality
  - Accounting
  - User support/call centre functionality
  - VO management services
  - Etc.

- These must be managed within GDA, but with real commitment:
  - To deliver results
  - To provide resources (hardware and people)

# Proposal

- Form a group from among the Tier1 and large Tier 2 centres, to:
  - Write service challenge plan
    - Milestones, functional and performance goals
  - Monitor progress of the plan and associated service challenges
    - Hold post-mortems – summarize problems
    - Set targets and analyse what/why they were not met
  - Provide resources committed to fulfilling the plan
    - Nominate “Data Challenge” leaders/coordinators at each centre
    - Ensure system managers understand priorities

# Makeup of the group & reporting

- Members should be project managers
  - Responsible for and committed to making the LCG service succeed in their centre/region
  - Who have resources to commit!
- The group would be part of the Deployment Area
- Report back to the PEB and GDB as appropriate

- Meet every 2 weeks by phone
  - In person if convenient

- Needs to be in place very quickly