



LCG Service Challenges

Ian Bird

LCG PEB

8 June 2004



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 for phase 2 – 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

- **Priority for now on data management:**
 - Understand bottlenecks and performance limitations in large data transfers between CERN and Tier 1's
 - Understand this now – well in advance of data taking
 - Get a reliable data transfer service in place this year
 - Understand and resolve network, storage/storage interface and transfer problems
 - Test interoperability between grids in use by LHC experiments –
 - Data transfers
 - Data management – eventually including replicated file catalogs
 - Job submissions
- **Set up GDA management group**
 - Scope and schedule of challenges



Service challenges – examples

- Data Management
 - Networking, file transfer, data management
 - Storage management and interoperability
 - Fully functional storage element (SE)
- Continuous job probes
 - Understand limits
- "Security incident"
 - Detection, incident response, dissemination and resolution
- Interoperability
 - Between LCG and Grid3, LCG and NorduGrid



Service milestones

- Also to be managed by same management group:
- IP connectivity
 - Milestones to remove (implementation) need for outbound connection from WN
 - Software installation, data access, write data remotely, publish information/bookkeeping, database access
- Operations centres
 - Accounting, assume levels of service responsibility, etc
 - Hand-off of responsibility (RAL-Taippei-US/Canada)
 - Should be involved in managing the service challenges from an operational viewpoint
- User support
 - Assumption of responsibility, demonstrate staff in place, etc
- VO management
 - Robust and flexible registration, management interfaces, etc



Data Management – data transfer

➤ Goal

- Build up a reliable and performant end-to-end data transfer service:
- **Eventual goal (for example):**
 - Data transfer between CERN and a Tier 1 at 500 MB/s sustained over 2 weeks. This should include resilience against failure, guarantee that a file will be transferred, potentially queuing and prioritisation of requests from different users. This should be achieved by the end of 2004.
 - Secondary milestones: 1-2 days sustained transfers at significant rates to all other Tier 1 (and large Tier 2?) sites



Data transfer service

- To achieve this, underlying milestones would be:
 - High performance network in place
 - Might include specifically routed gridftp traffic onto dedicated/private networks
 - Demonstrate disk-disk sustained transfer rates – simple use based on gridftp
 - Plans to start this work are in hand → Bernd Panzer/ David Foster together with “official” Tier 1 sites
- Higher level services:
 - Demonstrate SRM-SRM copies at sustained transfer rates (show for all Tier 1’s)
 - Have to demonstrate working implementations of SRM at each Tier 1
 - Understand and resolve firewall issues at each (like CERN HTAR)
 - Provide a service that accepts requests, queues them (prioritises), schedules transfer, recovers etc
 - This is a service that has to be written
 - Might be 2 layers – basic reliable service; scheduler
 - This service should guarantee delivery or notification
 - Gradually increase rates, periods of sustained transfers
 - Collections within file catalog



Data transfer – 2

- Once the basic data transfer service is in place
 - Test-bed in which to understand transfer protocols – gridftp etc – and improve/replace as needed without change to applications
 - Can build services on top – while improving the underlying implementation
- Interface LCG data management tools to this service
- Subsidiary work needed:
 - Understand how to do load-balancing gridftp at each site
 - Ensure SRM interface in place and functional
 - Build load generator:
 - Suggested performance implies 250K 2GB files



Replication service

- Work ongoing to propose RLS/RM improvements based on DC experiences
 - Note in draft listing ideas for improvement and simplification
 - How to handle metadata, optimising performance, etc.
 - In line with JRA1 architecture but allows us to prototype some ideas quickly now (and provide better performance)

- Replica Location Service:
 - Adapt RLS to use the underlying data transfer service
 - Proposal from DB group and others (US Atlas, CMS, etc) to understand replication strategy:
 - 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
 - Max jobs/day vs sites/nodes, max I/O etc
 - Use real jobs that exercise all systems, including data management
 - Behaviour with interruptions to services – resilience/recovery
- This might be a function of the GOC
 - Overseen by RAL-Taippei++ collaboration ?
- A challenge might run for a week
 - Outside of experiment data challenges
 - In parallel (or part of) data management or other challenges



Management group – Proposal

- Form a group from among the Tier1 and large Tier 2 centres, to:
 - Write detailed 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
 - Coordinate with experiments and regional centres to schedule the challenges



Makeup of the group & reporting

- Members should be project or computer centre managers
 - Tier 1 or large Tier 2 managers
 - Responsible for and committed to making the LCG service succeed in their centre/region
 - Who control resources and are able to commit them to work on these challenges and milestones
- The group would be part of the Deployment Area
- Report back to the PEB and GDB as appropriate
- Meet weekly or every 2 weeks by phone
 - In person if convenient
- Needs to be in place very quickly and set out milestones and challenges



Other coordination

- Service challenge group
 - People responsible for actually carrying out the challenges
 - Coordinated by Bernd
- Network coordination
 - David Foster – network managers at each Tier 1 site
- With EGEE data management and other groups
- With Grid3 and Nordugrid
 - Must not have parallel developments – must bring these groups together to work and and solve the common problems
 - Data management and replication, file catalogs highest priority
- Can/should these challenges be linked to experiment DC's?



Timescale

- **By mid-June**
 - Have management team agreed
 - **Waiting for responses**
 - Start weekly phone conferences – to flesh out draft of milestones and goals
 - Produce plan for mid-July
- **November/December**
 - Demonstrate sustained reliable data transfer service
 - **Should be usable for LCG-2 and EGEE underlying services**
- **Other milestones and challenges should be scheduled in the plan**