



Grid Deployment Area Meeting

Ian Bird

2 Feb 2004



Agenda

- Introduction
 - Purpose of meeting, ...
- LCG-2 Status
- Issues:
 - Storage Element
 - RLS compatibility (RM, POOL, GFAL)
 - SRM compatibility Castor vs dCache
 - Disk-based SE
 - GridICE (new requirement from CMS)
 - R-GMA as a service (new requirement from CMS)



Introduction and Goals

- Weekly coordination meeting
- Experiment computing and grid coordinators and users
- Open meeting
 - Invite regional centre administrators/ system managers
- Coordination between GDA, experiments, regional centres, etc.
 - Bring all the issues out into the open
- Goal:
 - Improve collaboration, communication between all stakeholders
 - Resolve, address operational issues
 - Understand where effort needs to be focussed
- Agenda:
 - General status – GDA, experiments, sites
 - Address specific issues (e.g. RLS/POOL)



Status of LCG-2



Deployment Plan for LCG-2

- Pick a set of core sites - ~6-8
- Deploy initially to those
 - Avoid configuration and stability issues
 - Sites to commit sufficient support effort and compute resources
- Aim to have 700-800 CPU available in core sites
- Aim at sites essential to experiments
 - Have push from experiments for sites to commit
 - Experiments to request resources be provided through LCG-2
 - Rough correspondence with Tier 1 sites
- Target a rapid deployment at these sites for Alice and CMS data challenges initially
- Not exclude other sites or Atlas, LHCb
 - Slightly longer timescales
- Process and core group ratified by GDB on Jan 13



Core sites and commitments

Site	Immediate	Later
CERN	200	1200
CNAF	200	500
FNAL	10	?
FZK	100	?
Nikhef	124	180
PIC	100	300
RAL	70	250
Taipei	60	?
Russia	30	50
Prague	17	40
Budapest	100	?
Totals	864(+147)	>2600(+>90)

Initial LCG-2 core sites

Other firm commitments

Will bring in the other 20 LCG-1 sites as quickly as possible



Schedule and status

- Dec 20 2003: LCG-2 middleware release certified
 - Done, but SRM interface to replica manager untested
- Jan 5 2004: Finalise deployment preparation
 - Done, installation and deployment instructions, release notes
- Jan 12 2004: Begin deployment to core sites
 - 2 day delay due to rearrangement of Cern computer centre
- Jan 19 2004: Begin ramp up of nodes at core sites
 - Starting
- Alice DC starts 1 Feb – agreed to delay by 1 month
- CMS – starts March



LCG-2 functionality

- General
 - CondorG –
 - new grid manager (critical, now in official VDT)
 - gahp-server (critical, local, with Condor team now)
 - scheduler, memory usage (with Condor team)
 - Globus -
 - RM wouldn't work behind the firewall
 - prevent occasional hangs of CE
 - number of errors in the handling of return status from various functions
 - Note: we refrained from putting all fixes into the current 2.2.x we are running on LCG-2 knowing that they will be included in 2.4.3 we are to test as of next week.
 - RB – new WP1 fixed number of LCG-1 problems (reported by LCG)
 - above this we fixed (with WP1 team) memory leaks in
 - Interlockd
 - network server
 - filelist problem
 - CE – memory leaks
- Installation
 - WN installation independent from LCFGng (or other tools)
 - Still required for service nodes
- Still require outbound IP connectivity from WN's
 - Work to be done to address in Replica Manager
 - Add statement to security policy to recognise the need – but limit it – applications must not rely on this



LCG-2 Storage Element

- Plan:
 - SRM interface to all storage, Disk pool manager software
 - Accessible from RM, GFAL, (gridFTP), ...
- Storage Resource Manager (SRM)
 - Why SRM?
 - Grid interface to storage – location independence
 - Managed storage for disk-only SE's



LCG-2 SE status

- Mass Storage access – to tape
 - SRM interfaces exist for Castor, Enstore/dCache, HPSS
 - Castor – CERN, CNAF, PIC
 - dCache – FNAL, (FZK)
 - Other sites – Lyon (HPSS), RAL
- Disk-based SE's
 - (Packaged version of Castor disk pool manager – no longer available)
 - Packaged version of dCache – being tested
 - Provide both as options for sites that require cache manager
 - Existing MSS/SRM systems need to deploy a GRIS and corresponding info provider
 - Described in installation notes
- GFAL included in LCG-2



SE Status – today

- SE (day 1 – now): available via gridftp
 - Castor (gridftp) at CERN to tape
 - Configured as local SE to all sites
 - Could have local (classic) disk-only at sites
 - No space management – has to be watched and done by hand, no idea of available space
 - Data will be lost when install managed SRM-SE
 - Needs GRIS at CERN (today)
 - Test RM in this configuration – today?
 - Need to verify RB also – today
- Will also be at CNAF and PIC (once this works)

- SE – SRM
- Castor@CERN, CNAF, PIC:
 - Need to test RM against this (!)
- EnStore@FNAL:
 - Need to test RM against this (!)
- Disk-only SRM-SE's:
 - dCache: needs to be tested
 - Have fixes, re-package (RAL), test
 - Castor: not available



RLS and POOL

- Case-sensitivity
 - In LCG-1 (EDG 1.4) RLS attributes were case sensitive
 - In EDG 2.x WP2 changed DB schema:
 - Attribute mapped to a column name
 - SQL does not No longer case sensitive!
 - Only found when using POOL as RLS client (now)
 - Short term solution – make case-preserving, works for Oracle and MySQL
 - Longer term – solutions suggested – e.g. intermediate table, but would require data migration
 - Propose: implement case-preservation now, understand longer term solutions but implementation should be seen in terms of other needed RLS/RM work



RLS Naming – affects all clients

- File naming (srm://, srb://, sfn://, rfio://, etc.)
 - Original paper (2001, CERN, FNAL, LBNL, etc) proposed common syntax
 - LFN:, GUID:, SRM:, etc.
 - GFAL & RM both implement (and require!) this
 - Expects that user provides file name always with a prefix even for a “trivial” local file name (i.e. LFN:)
- Propose POOL implement SURL → TURL (which gets passed to ROOT which does not recognise srm:),
 - Use existing routine from GFAL
- All names stored in RLS would have a prefix
 - Replica Manager and GFAL already consistent



Other issues

- ✓ Storage Element
- ✓ RLS compatibility (RM, POOL, GFAL)
- ✓ SRM compatibility Castor vs dCache
- ✓ Disk-based SE
- GridICE (new requirement from CMS)
- R-GMA as a service (new requirement from CMS)