Status of LCG-2 Deployment

Ian Bird

LHCC Referees Meeting
26 January 2004

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

- Selected with experiments
- Requested sites to guarantee levels of commitment
  - >50 nodes, sufficient support staff
- Other sites in LCG-1 and new sites
  - will migrate to LCG-2 and be included in stable core when they are stable
- If any site causes problems will be removed from core group
- We do not intend to exclude sites, but focus on building up a stable core first

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
- Expect to be ready for:
  - 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 occassional 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 functionality

- Storage Element(s) – based on SRM interfaces
  - Packaged version of Castor disk pool manager – available
  - Packaged version of dCache – needs mods – this week
  - 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
  - Initial SRM-enabled MSS installations
    - CERN, FNAL, PIC, CNAF
    - Other sites addressed one by one
  - GFAL included in LCG-2

Status of LCG-1 sites

- Sites that are in LCG-1 (28)
  - BNL, Budapest, CERN, CNAF (+4 Tier 2), CSCS, FNAL, FZK (+Krakow), Lyon, Moscow, PIC (+6 Tier 2), Prague, RAL (+2 Tier 2), Taipei, Tokyo, Triumf
  - Preparations: IHEP-Beijing, Pakistan
- Xmas productions
  - Productions for CMS and Atlas were run on LCG-1 over Christmas holidays
  - Work by US-Atlas and LCG succeeded in running jobs on LCG from Grid3 sites
    - Basic policy issues to be resolved
  - Some problems with mis-configured sites caused frustration for experiment testers

LCG-1 use over Christmas

- CMS
  - Ran for 9 days on LCG-1 (20-23/12/2003 – 7-12/1/2004)
  - In total 600,000 events were produced (ORCA) using UIs in Padova and Bari
  - Sites used were mainly in Italy and Spain; NIKHEF and FZK excluded because it was not possible to send e-mails from WNs needed to update CMS REFDB.
- EU Atlas
  - Ran on LCG-1 facility 75 jobs of 10 hours each.
  - In total 14,000 events were produced using UI in Milan.
  - Sites used were mainly in Italy: CNAF, Turin and Milan. CERN was also used but had to be excluded because of slow response (jobs in waiting status). CERN RB showed problems before Xmas.
  - Only 5 jobs over 75 failed.
- US ATLAS
  - Sent requests for job execution to LCG-1 from the US Grid3 infrastructure.
  - Details can be found in the report:
    - http://agenda.cern.ch/askArchive.php?base=agenda&categ=a04162&id=a04162s1t71%2Fmoreinfo%2Fchimera-lcg1.pdf
  - After modification to Pegasus/Chimera, 77 events were successfully generated using LCG-1 sites CERN, Turin and Brookhaven with the output data staged at the University of Chicago and registered in the Globus RLS

Plan for 2004

- Services
  - Production service:
    - Plan on running LCG-2 during 2004
    - Put effort into bug fixing, robustifying, etc
    - Development where essential – data management (RLI, RM-proxy, etc), GFAL
  - Development service:
    - Based on EGEE/ARDA prototype – run in parallel with LCG2 at EGEE sites

- Middleware & support, issues
  - LCG will manage source code repository of components layered on top of VDT
  - Support:
    - WP1 – INFN
    - WP2 – Cern
    - Other components – GD
    - VDT/Globus – continue good relationship with VDT
  - Address:
    - Porting to other compilers and architectures
    - Rationalise external and inter-package dependencies
    - Simplify packaging and installation
    - Make use of SPI infrastructure as far as possible

Summary

- Initial LCG-2 deployment to small set of committed core sites
  - Approved by GDB and on schedule
- Commitment by experiments
  - ALICE commitment clear – delayed start by 1 month
  - CMS commitment to be clarified in scope
- SRM SE deployment delayed but will be ready so that Alice can use it to move data, others can start testing now
  - Interaction with replica manager had to be made to work
- Will be ready for Alice and CMS DC's
  - And for Atlas, LHCb
- Take ownership of support process to avoid problems that we had in 2003