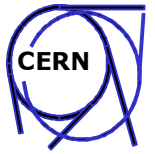

Status of Deployment & tasks

Ian Bird
LCG & IT Division, CERN

GDB – FNAL
9 October 2003

LCG-1 Deployment Status

- Up to date status can be seen here:
  - http://www.grid-support.ac.uk/GOC/Monitoring/Dashboard/dashboard.html
    - Has links to maps with sites that are in operation
    - Links to GridICE based monitoring tool (history of VO's jobs, etc)
      - Using information provided by the information system
    - Tables with deployment status

- Sites that are currently in LCG-1 (here) expect 18-20 by end of 2003
  - PIC-Barcelona (RB)
  - Budapest (RB)
  - CERN (RB)
  - CNAF (RB)
  - FermiLab. (FNAL)
  - FZK
  - Krakow
  - Moscow (RB)
  - RAL (RB)
  - Taipei (RB)
  - Tokyo

Total number of CPUs ~120 WNs

Sites to enter soon

BNL, (Lyon)

Several tier2 centres in Italy, Spain, UK

Sites preparing to join

Pakistan, Sofia, Switzerland

Users (now):

Loose Cannons
Deployment Team
Experiments starting (Alice, ATLAS,..)

Getting the Experiments on

- Experiments are starting to use the service now
  - Agreement between LCG and the experiments
    - System has limitations, testing what is there
    - Focus on:
      - Testing with loads similar to production programs (long jobs, etc)
      - Testing the experiments software on LCG
    - We don't want:
      - Destructive testing to explore the limits of the system with artificial loads
        - » This can be done in scheduled sessions on C&T testbed
  - Adding experiments and sites rapidly in parallel is problematic
    - Getting the experiments on one after the other
    - Limited number of users that we can interact with and keep informed

History

- First set of reasonable middleware on C&T Testbed end of July (PLAN April)
  - limited functionality and stability

middleware was late

- Deployment started to 10 initial sites
  - Focus not on functionality, but establishing procedures
  - Getting sites used to LCFGng
- End of August only 5 sites in
  - Lack of effort of the participating sites
  - Gross underestimation of the effort and dedication needed by the sites
    - Many complaints about complexity
    - Inexperience (and dislike) of install/config Tool
    - Lack of a one stop installation (tar, run a script and go)
    - Instructions with more than 100 words might be too complex/boring to follow
- First certified version LCG1-1_0_0 release September 1st (PLAN in June)
  - Limited functionality, improved reliability
  - Training paid off -> 5 sites upgraded (reinstalled) in 1 day
  - Last after 1 week….
- Security patch LCG1-1_0_1 first not scheduled upgrade took than 24h.
- Sites need between 3 days and several weeks to come online
  - None in not using the LCFGng setup (status Thursday)
- Now: 11 in and several Tier 2 sites waiting, 2 of the original 10 missing

Difficulties

- Sites without LCFGng (even using lite) have severe problems getting it right
  - We can't help too much, dependencies depend on base system installed
  - The configuration is not understood well enough (by them, by us)
  - Need one keystroke "Instant GRID" distribution (hard..)
  - Middleware dependencies too complex
- Debugging a site
  - Can't set the site remotely in a debugging mode
  - The glue status variable covers the LRM's state
  - Jobs keep on coming
  - Discovery of the other site's setup for support is hard
- History of the components, many config files
  - No tool to pack config and send to us
  - Sites fight with FireWalls
- Some sites are in contact with grids for the 1st time
  - There is nothing like "Beginners Guide to Grids"
- LCG is on many sites not a top priority
  - Many sysadmins don't find time to work for several hours in a row
  - Instructions are not followed correctly (short cuts taken)
- Time zones slow things down a bit

Stability-Operation

- Running jobs has now greatly improved
  - "Hello World" jobs are about 95% successful
  - Services crash with much lower rate (some bug fixes already on C&T)
  - Some bugs in LCG1-1_0_x already fixed on C&T
  - Grid services degrade gracefully
- So far the MDS is holding up well
- Focus in this area during the next few months
  - Long running jobs with many jobs
  - Complex jobs (data access, many files,…)
  - Scalability test for the whole system with complex jobs
  - Chaotic (many users, asynchronous access, bursts) usage test
  - Tests of strategies to stabilize the information system under heavy load
    - We have several that we want to try as soon as more Tier2 sites join
  - We need to learn how the systems behave if operated for a long time
    - In the past some services tended to "age" or "pollute" the platforms they ran on
  - We need to learn how to capture the "state" of services to restart them on different nodes
  - Learn how to upgrade systems (RMC, LRC…) without stopping the service
    - You can't drain LCG1 for upgrading

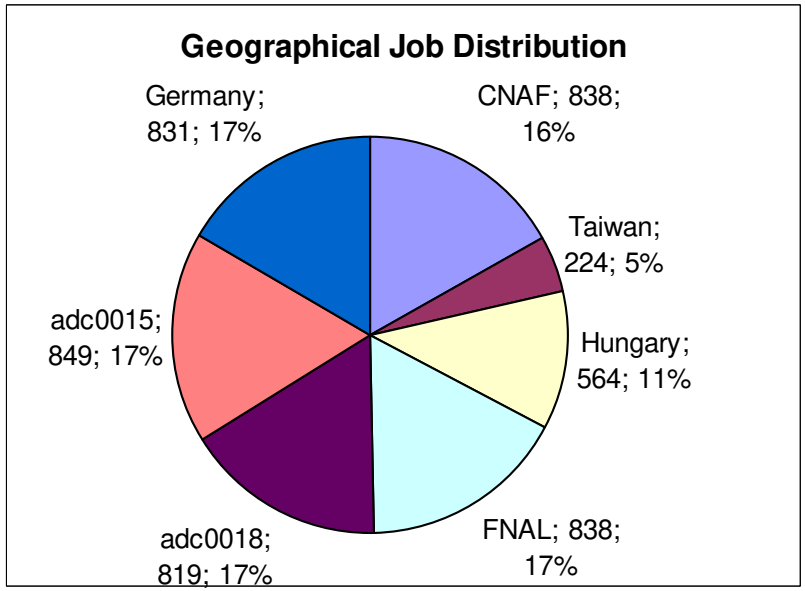

LCG 1.0 Test (19./20. Sept. 2003):

- 5 streams
- 5000 jobs in total
  - Input and OutputSandbox
  - Brokerinfo query
  - 30 sec sleep

Ingo Augustin: Loose Cannon testing

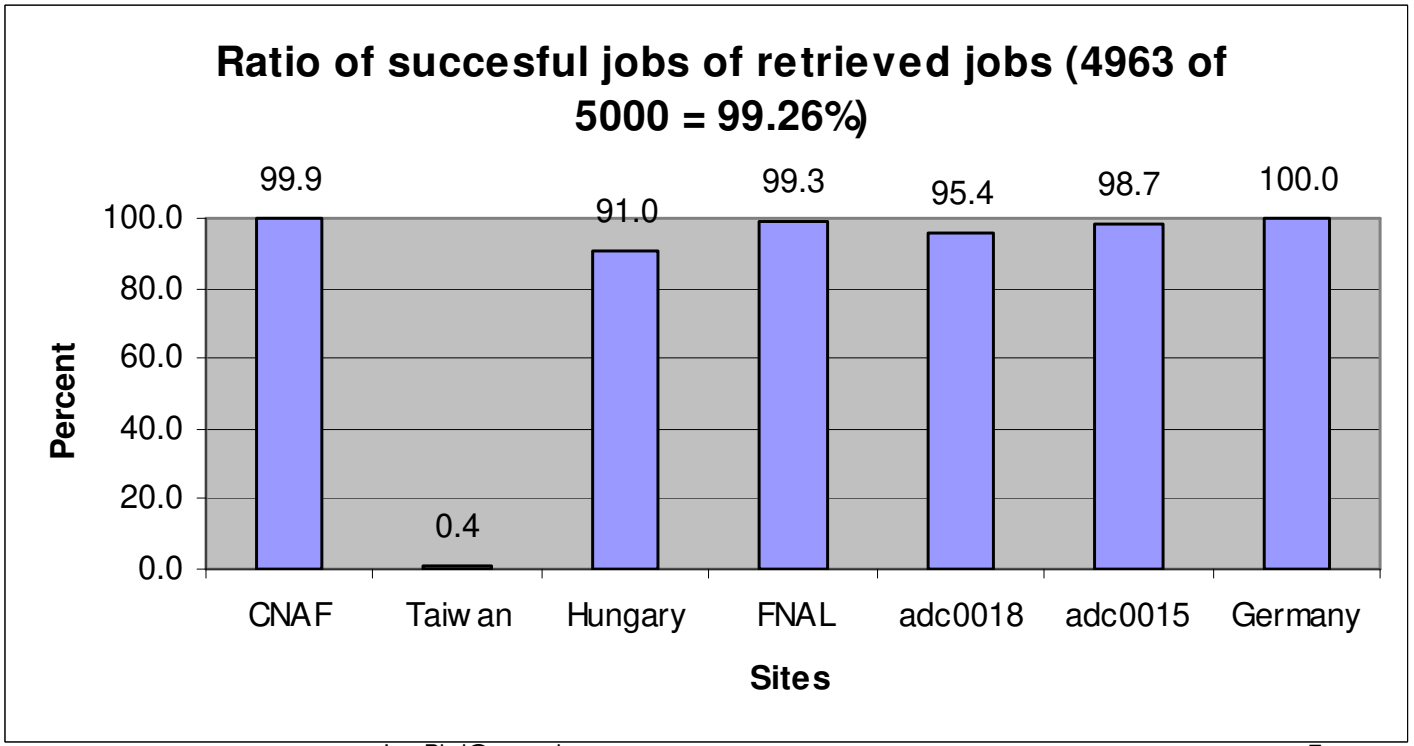

LCG-1 Middleware

- LCG-1 is:
  - VDT (Globus 2.2.4)
  - EDG WP1 (Resource Broker)
  - EDG WP2 (Replica Management tools)
  - GLUE 1.1 (Information schema) + few essential LCG extensions
  - LCG modifications:
    - Job managers to avoid shared filesystem problems
    - MDS – BDII LDAP
    - Globus gatekeeper enhancements ((adding some accounting and auditing features, log rotation, that LCG requires)
    - Many, many bug fixes to EDG and Globus/VDT
- Not yet tested:
  - MDS system under heavy load on the full distributed system
    - Several strategies for improvement if there are problems
  - Scalability, reliability of full LCG-1. This is where we will focus over next few months

Timescale for LCG-2

- Original goal was to have upgraded middleware deployed for large-scale testing by end October
  - This is no longer realistic given delays in LCG-1 availability and deployment
- Functionality integrated in EDG 2.0 and 2.1 is very much reduced from that originally expected
  - That integration is very late – just finished now
- Goal is to have LCG-2 ready to deploy by November 20
  - 1 month to deploy before end of year – needs this time
  - Absolute latest that we can hope to have it ready before 2004
- Time is tight – must be conservative in what we aim for, but nevertheless must have a useful system

Basic upgrades

- Switch to official EDG 2.x distribution
  - Has all LCG found bugs fixed + other improvements
  - Uses gcc 3.2 - *required!*
  - To be verified that this is as good as the LCG-1 distribution
  - This starts next week (when EDG 2.1.x is frozen after Heidelberg)

- Upgrade Globus/VDT to Globus 2.4.x
  - Working on certifying Globus2.4.1 now
  - Want to move to supported Globus (for next year)
  - Want to have a common VDT version with US efforts
  - We will test this upgrade once we have the EDG 2.x tested

Added functionality – 1

- VOMS – role-based authentication/authorization
  - Service is running at Nikhef (for a while) – stable and has been used by US groups
  - LCG has installed service on certification tb
  - 1st step is to use VOMS to generate gridmap files (US usage)
  - 2nd step is test integration with gatekeeper – use LCAS/LCMAPS
    - EDG has integrated this functionality
  - 3rd step is to integrate with file access
    - Initially we will not integrate with storage access – see below

Added functionality – 2

- POOL

- Is currently being installed and tested with LCG-1 on certification test-bed
- Will be repeated against LCG-2
  - Should be simpler as compiler/library issues will be resolved

Added functionality – 3

- Storage access
  - Currently only have simple disk-based SE only
  - Have in hand components of a real solution:
    - GFAL, SRM, Replica Management tools
  - GFAL:
    - Communicates with SRM, Replica Management tools
    - Has been tested stand-alone with Castor and Enstore
    - This week testing in certification tb – then test against Castor, Enstore in LCG-1
  - SRM:
    - Implemented for Castor and Enstore
  - Expect to have integrated SRM/GFAL/RLS providing access to CERN and FNAL
    - Other sites will require SRM implementations against their MSS – or some real effort from WP5 to make a true SRM-compliant SE
  - GFAL will use VOMS to implement file access controls

RLS/RLI

- Issues:
  - Two mutually incompatible versions of RLS exist – one being used in US, EDG version is part of LCG-1; this is a problem for the experiments
  - We have to live in a world of interoperability – between different grid infrastructure
  - POOL currently only works with EDG version
- LCG Plan was:
  - To deploy RLI as soon as possible to remove reliance on single central service
  - Implement a replica management proxy service to avoid need for outgoing connectivity on WNs
- But:
  - EDG RLI is very late – still being tested/integrated now
  - Development effort for proxy not available on required timescale?

RLS/RLI proposal

- Aim to have a common solution as soon as possible
  - But this will not be ready until May/June 2004 – after DCs
- For DCs continue with central solution based on EDG RLS
  - Bulk updates of catalogs from batch production
  - Copies of catalogs possible; but update central master and then push back out
  - Requires pre-job replication is done to ensure data is at site where job runs
  - This model avoids need for proxy on this timescale – all updates can be done from a defined place (e.g. CE) that has connectivity
  - If RLI is demonstrated to work, we may still be able to deploy it on this timescale
  - US Tier 2 sites may be using Globus RLS; POOL will provide tools to allow catalog cross-population (these basically exist already)

RLS common solution

- US will assist in port of POOL to use both RLS implementations
- Globus RLS is being ported to use Oracle
- Roadmap agreed by both EDG and Globus groups:
  - Agree the APIs for RLS and RLI. This discussion should include agreement on the syntax of filenames in the catalog.
  - Implement the Globus RLI in the EDG RLS, make the EDG LRC talk the "Bob" protocol.
  - Implement the client APIs
  - Define and implement the proxy replica manager service
  - Update the POOL and other replica manager clients (e.g. EDG RB)
- Timescale is May 2004; and is joint effort between US and CERN
  - Overseen by LCG and Trillium/Grid3 via JTB

- Confirmed this week with Globus – will commit effort
  - Start next week by re-assessing plan and work

R-GMA?

- It was our intention to deploy R-GMA – at least in parallel with MDS to enable a direct comparison
- Currently R-GMA is now in EDG 2.1 – being tested
- Timescale will not allow testing as a replacement for MDS – but still want to make comparison on the deployed system
- Must inter-operate with US grid infrastructure – for now that means MDS until such time as R-GMA can demonstrate its benefits

Timeline - Preparation

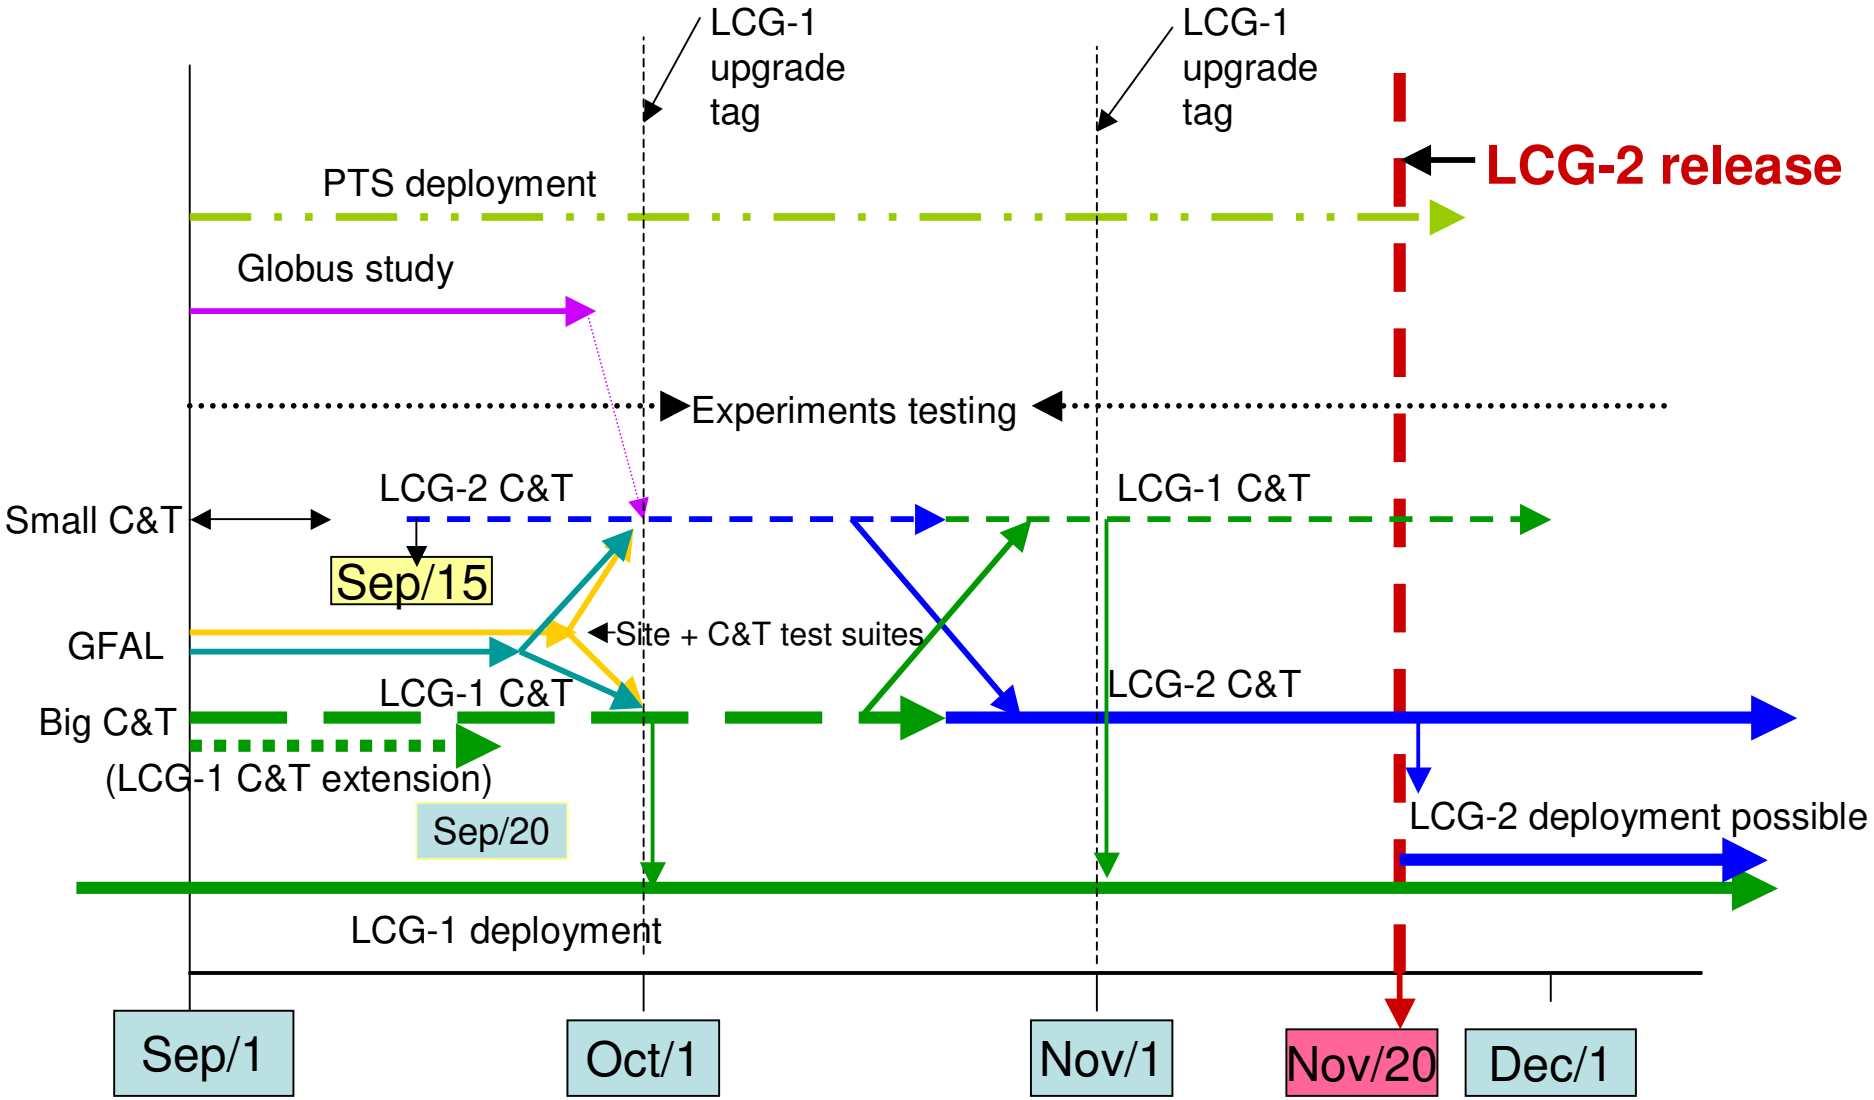

Tasks – need effort now

- Installation issues
  - Produce manual installations – wrappers, alternative tools
  - LCFGng, lite, …
  - Working on "manual" installation procedures
  - Need a collaborative team to improve this process – especially for Tier 2 installation
- Need SRM wg
  - Set LCG requirements on deploying SRM at sites
  - Will need effort at sites with MSS to integrate
- Experiment sw installation
  - Proposal made – based on experiment requirements – need to implement
- IP connectivity – solutions for NAT,
  - General solutions, experiment specific needs
- Local integration issues
  - Batch systems – migrate to LSF, Condor, FBSng, etc.
  - Mass storage access – need solutions – SRM implementations

Tasks 2

- Usage accounting
  - Need a basic system in place
  - RAL have effort – is it sufficient?
- Software maintenance
  - Move away from EDG monolithic, interdependent distribution
  - Separate out WP1, WP2 and other pieces we need – info providers, gatekeeper
  - Set up build system
  - Negotiate maintenance of code (WP1, WP2, gatekeeper) – INFN & CERN
  - LCG will take over maintenance of info providers
  - Need an info-provider librarian
    - Work with VDT, coordinate with monitoring
- VOMS/VOX – workshop in December
  - Need to arrive at agreement on operation next year

Tasks 3

- Inter-operation
  - Initially with US Tier 2 sites – based on Grid3
  - Need small set of goals to achieve basic inter-operation
  - Start discussion tomorrow – need dedicated effort?

- Operational stability issues
  - Long running jobs with many jobs
  - Complex jobs (data access, many files,…)
  - Scalability test for the whole system with complex jobs
  - Chaotic (many users, asynchronous access, bursts) usage test
  - Tests of strategies to stabilize the information system under heavy load
    - We have several that we want to try as soon as more Tier2 sites join
  - We need to learn how the systems behave if operated for a long time
    - In the past some services tended to "age" or "pollute" the platforms they ran on
  - We need to learn how to capture the "state" of services to restart them on different nodes
  - Learn how to upgrade systems (RMC, LRC…) without stopping the service
    - You can't drain LCG1 for upgrading
  - Should this be subject of a working group?

Tasks

- All sites, but especially Tier 1s will need to have effort available to address specific integration issues
  - SRM
  - Publishing accounting info
  - Integrating batch systems
  - Experiment sw installations
  - …

Communication

- Mis-communication, mis-understandings
  - We need now a technical forum for LCG system admins
    - To coordinate integration tasks
  - Hoped that HEPiX would fill this role (but timescale)
  - Perhaps need to now have a workshop for this – with at least the primary sites sending their primary LCG admins
- What is appropriate communication mechanism:
  - To system admins
  - To site/system managers
  - To users
  - How to ensure GDB discussions get transmitted to the sites and the admins
- Regular email newsletter(s)?

Effort

- Many tasks need to be done and coordinated
- Has to be collaborative efforts now
  - Deployment team too small to do it all and support deployments
- Need to now see dedicated effort to join in this collaborative work