

LCG-1 Status & Planning

Ian Bird
LCG & IT Division, CERN

GDB Meeting, CERN
9 September 2003




Overview



- LCG-1 Current Status
 - Loose cannon report
 - Alice usage
- Some Issues – centres (not) participating
- Resource issues
 - Staffing and support
 - Must report in quarterly reports
- Priorities
 - LCG-1 vs Oracle
- RLS issues
- Experiment and general software deployment
- Planning for next deployments
- Experiments use of LCG-1



Middleware Status



LCG-1 is:

- VDT 1.1.8-9 (Globus 2.2.4)
 - Information System (MDS)
- Selected software from EDG 2.0:
 - Workload Management System (RB)
 - EDG Data Management (RLS, LRC, ...)
- GLUE Schema 1.1 + LCG extensions
- LCG local modifications/additions/fixes, such as:
 - Special job managers (LCGLSF, LCGPBS, LCGCONDOR) to solve the problem of sharing home directories
 - Gatekeeper enhancements (adding some accounting and auditing features, log rotation, that LCG requires)
 - Number of MDS fixes (also coming from NorduGrid)
 - Number of misc. Globus fixes, most of them included now in the VDT version LCG is using

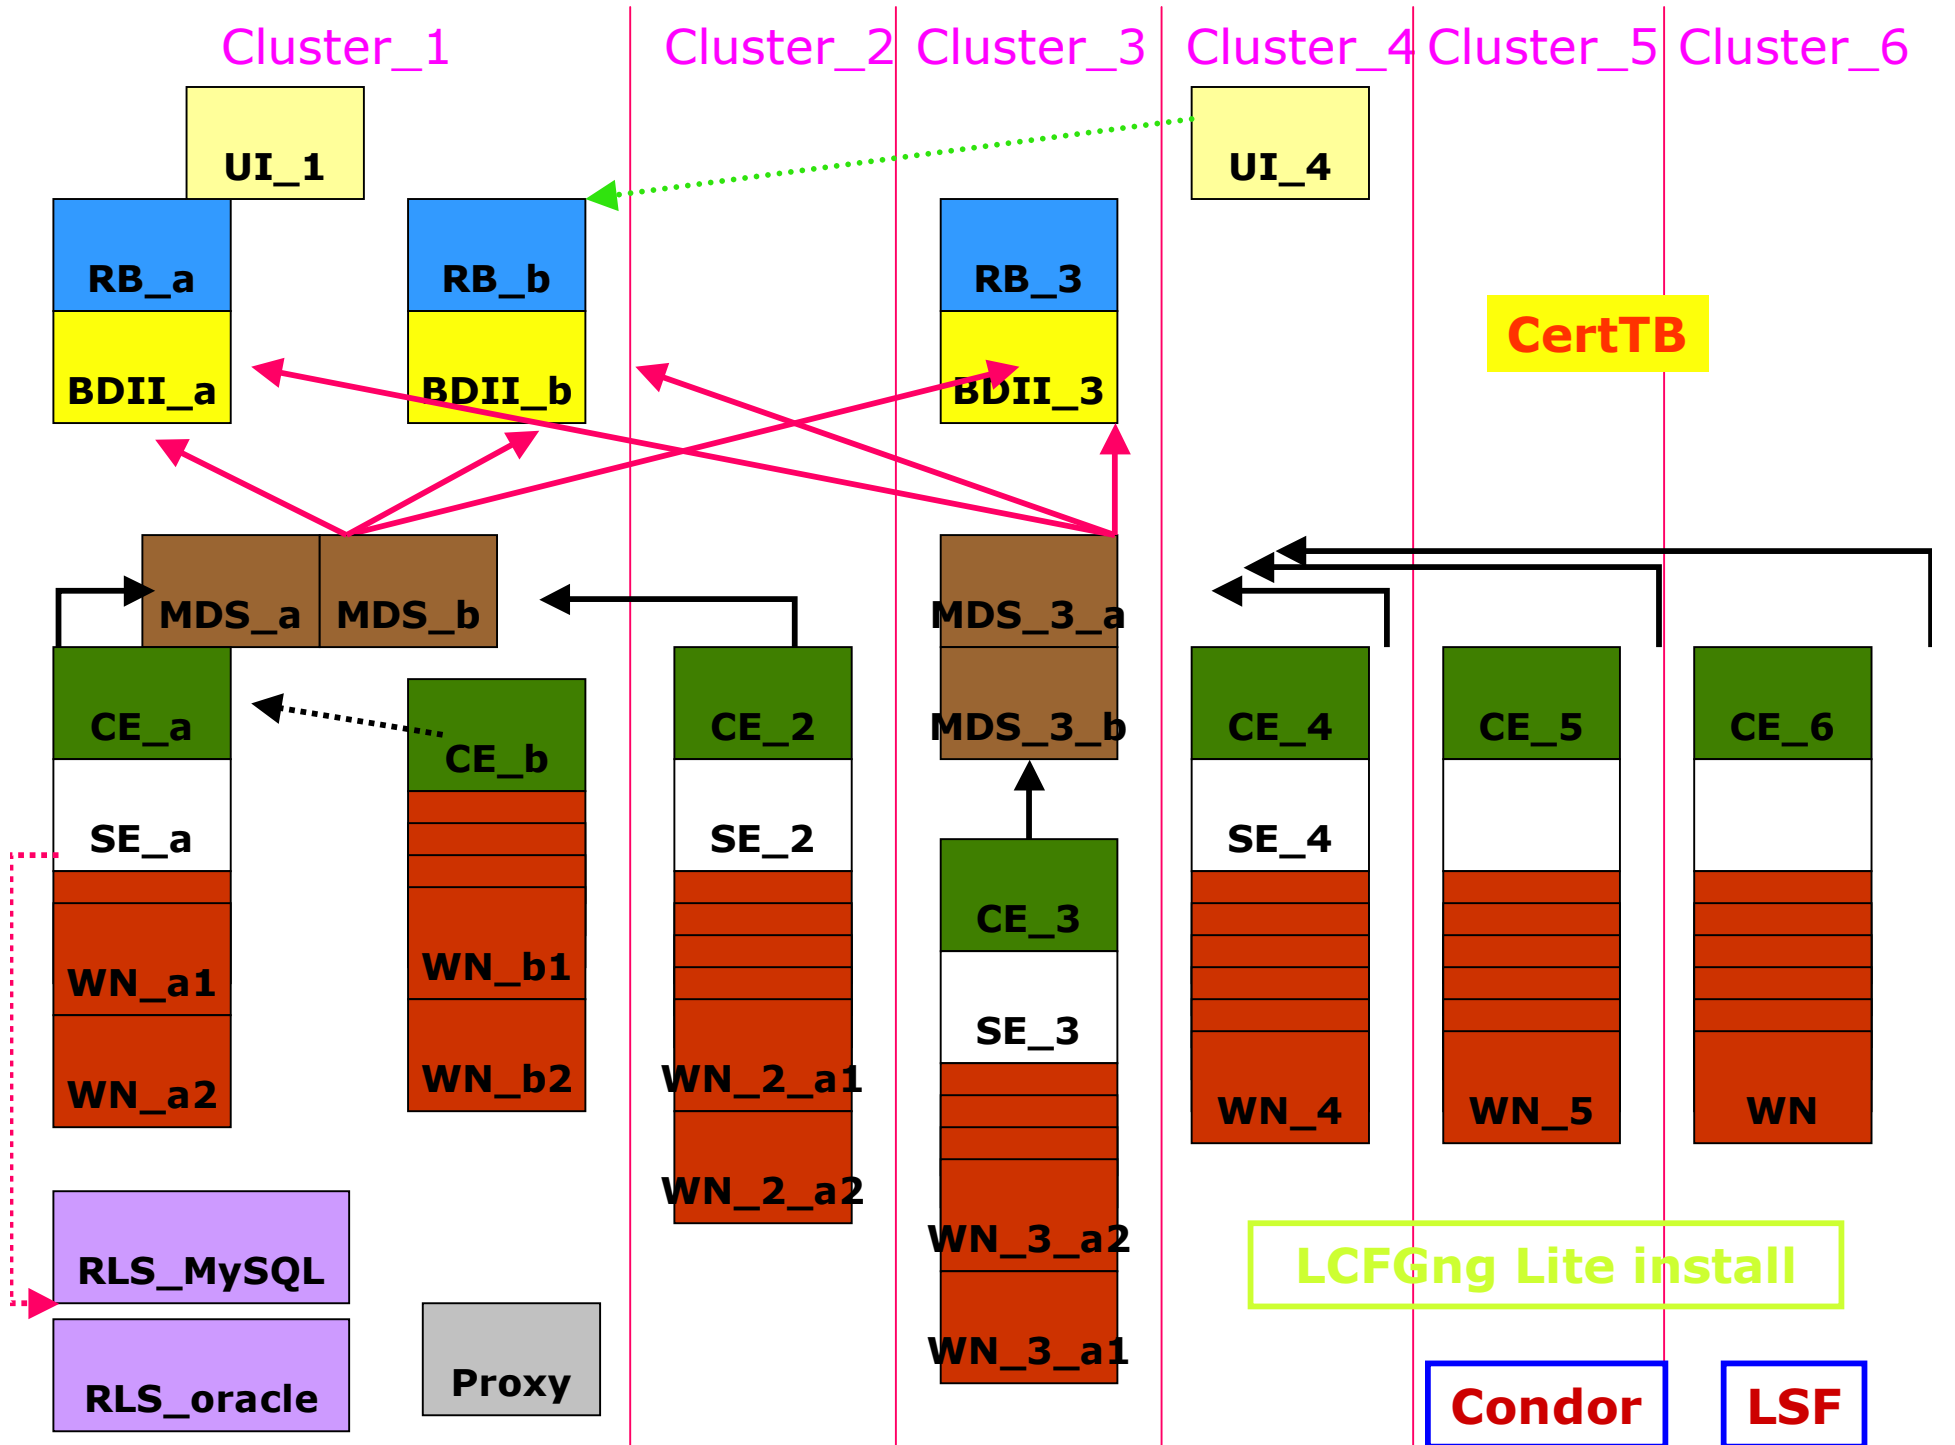


Status - cont



Current release:

- Status report is sent regularly since about 2 months
 - Allowed us to have a clear view of outstanding problems and their status on a day-to-day basis
- No more showstoppers or critical bugs at this point
- LCG will be working on the C&T to improve stability and scalability of the system.
- Scalability will become the next big issue.
- LCG will issue every 3-4 weeks an update to the middleware system based on the work on the C&T testbed
 - Deployment group will decide to deploy or not
- Have a fairly comprehensive suite of tests that has allowed us to achieve a reasonable stability
- EDG Loose cannons have tested system – see report in following slides



Overview on Loose Cannons tests on LCG



WP8 Technical Working Group
Bologna, September 4-5, 2003



Mario.Reale@roma1.infn.it

on behalf of EDG WP8 : I.Augustin, JJ.Blaising, S.Burke, F.Harris, J.Templon, MR

Test code sources

for LC july-aug 03 LCG-CT testing:



- ◆ Full general purpose generic-HEP-application testsuite from Jean-Jacques : (PERL)
 - Gets info on GRID current status (GRID config file) and creates JDL accordingly
 - Submits jobs, monitor status, retrieve output, report results
- ◆ Various LC testing JDLs and scripts (/bin/bash, PERL)
 - general and intensive stress tests
- ◆ Official EDG WP6-SC (edg-site-certification) PERL-oo test suite
 - Basic job submission, output retrieve
 - Basic Data Management tests, Filename registration and file copy
 - MyProxy tests
 - Match-Making specific tests

(http://marianne.in2p3.fr/datagrid/TestPlan/TESTSTATUS/EDG_20_TEST_STATUS.html)

WMS tests and results (1/2)



- ◆ Tested basic job submission and output retrieve "HelloWorld"
[OK]
- ◆ Test submission of 50 jobs with Input and Output sandboxes
with and without retryCount=0
[OK]
- ◆ Test of 250 jobs with long sleep, no resubmission (MyProxy)
[OK]
- ◆ 5 job with 2GB Input&Output sandboxes, no resubmission,
with and without parallel streams
[OK]
- ◆ Match-making tests:
 - Three files on one SE, job matches to associated close CE
[OK]
 - Three file on one SE, add application tag, job matches close CE
[OK]
 - Access a file from a SE from job with file protocol; compute
checksum
[OK]



WMS tests and results (2/2)

- ◆ Submit JJ-HEP generic application on all available GRID CEs, retrieve output, register files on iteam VO LRC via RLS from the running script on the WN, copy them on close SE.
[Resources previously discovered by query to InfoSys]

[OK]

- ◆ Submit a numeric iterative calculation on all available GRID CEs, use BrokerInfo to find CloseSE and mountpoint, copy files there and register into iteam LRC through RLS

[OK]



DMS tests and results (1/3)

- ◆ globus-url-copy : transfer large files to available SEs
- ◆ Upload of 2GB file, register in RLS, delete [OK]
- ◆ Register 1000 file with file names longer than 80 characters. [OK]
- ◆ Accessing files on SE : protocols
 - File [OK] (3 SEs out of 4)
 - gridFTP [OK]
 - RFIO [works only locally – Not connected to CASTOR]
- ◆ 10 jobs producing files on WN, using BrokerInfo to discover close SE, copying files there, registering output files on the LRC [OK]

DMS tests and results (2/3) : tested commands



- ◆ **edg-rm** copyAndRegister [OK]
- ◆ **edg-rm** registerFile [OK]
- ◆ **edg-rm** listGUIDS [OK]
- ◆ **edg-rm** listReplicas [OK]
- ◆ **edg-rm** replicateFile [OK]
- ◆ **edg-rm** getBestFile, listBestFile [OK]
- ◆ **edg-rm** printInfo, getVersion, listBestFile, list [OK]
- ◆ **edg-rmc** mappingsByGuid [OK]
- ◆ **edg-rmc** mappingsByAlias [not OK]
(Unhandled RemoteException : nullNo such operation 'getMappingsByAlias')
- ◆ **edg-lrc** pfnsForGuid [OK]
- ◆ **edg-lrc** guidForPfn [OK]

DMS commands (3/3): tested on UI only



- ◆ edg-rm bulkCopyAndRegister [Not OK]
- ◆ edg-rm getTurl [OK]
- ◆ edg-rm registerFile [OK]
- ◆ edg-rm registerGUID [OK]
- ◆ edg-rm unregisterFile [OK]
- ◆ edg-rm addAlias [OK]
- ◆ edg-rm removeAlias [OK]

-Not tested (not foreseen in LCG-1)

- ◆ getAccessCost

RB Stress Tests by Massive Job Submission



- ◆ RB never crashed
- ◆ ran without problems at a load of 8.0 for several days in a row 20 streams with 100 jobs each (typical error rate ~ 2 % still present)
- ◆ RB stress test in a job storm of 50 streams , 20 jobs each :
 - 50% of the streams ran out of connections between UI and RB. (configuration parameter – but machine constraints)
 - Remaining 50% streams finished normal (2% error rate)
 - Time between job-submit and return of the command (acceptance by the RB) is 3.5 seconds. (independent of number of streams)
- ◆ NOTE:

RB interrogates all suitable CE's : wide area delay-killer (interactive work) ?

Preliminary full simulation and reconstruction tests with ALICE



- ◆ Aliroot 3.09.06 (including HBT correl.) fully reconstructed events
- ◆ CPU-intensive, RAM-demanding (up to 600MB ,160MB average) ,long lasting jobs (average 14 hours)
- ◆ Outcome:
 - > 95 % successful job submission, execution and output retrieval in a lightly loaded GRID environment
 - ~ 95 % success (**first estimate**) in a highly job-populated testbed with concurrent job submission and execution (2 streams of 50 AliRoot jobs and concurrent 5 streams of 200 middle-size jobs)
 - MyProxy renewal successfully exploited [OK]

Further Tests outcome & EDG bugs opening



- ◆ Incorrect publication of **GlueSAStateAvailableSpace**:

now fixed

[OK]

- ◆ **edg-find-resources -doCE -doSE** not working

(opt/edg/lib/perl/ should be included in @INC perl libs)

- ◆ **gang-match** doesn't work (bug 1442) because of SAobject

[Not OK]

- ◆ You can currently block the RB system if you submit the proper (wrong) JDL file using the old non-glue-scheme-compliant classAds

[Not OK]

(i.e.: other.CEId==".." instead of
other.GlueCEUniqueID=="..")

[Not OK]

Currently opened EDG bugs effecting LCG (no show-stoppers)

<http://marianne.in2p3.fr/datagrid/bugzilla/>



- ◆ 1675 : problems deleting a file if the register failed
- ◆ 1676 : the RM doesn't delete automatically if register failed
- ◆ 1681 : bad RM behaviour if the CESEBind is missing
- ◆ 1685 , 1687 : metadata and wildcards are missing from edg-rm
- ◆ 1723 : no rfio TURL
- ◆ 1748 : rfio not properly configured
- ◆ 1756 : you can get multiple replicas on one SE
- ◆ 1762 : gridFTP doesn't lock a file while it's being written
- ◆ 1770 : strange RM error message during stress test
- ◆ 1840 : matchmaking gives bizarre results with two close SEs
- ◆ 1841 : **getBestFile does not take multiple close SEs into account**



Conclusions

- ◆ Some further possible tests can be performed:
 - Further investigation on DMS functionality/commands
 - Further tracking of (non show stoppers) bug fixing as EDG 2.0 evolves into 2.1

- ◆ Impressive **improvement on Stability** w.r.t. old 1.x EDG releases and corresponding testbeds, as stress tests confirmed

- ◆ Room for further improvements, also naturally driven by porting to newer EDG releases

- ◆ As LCG-CT will expand, **Scalability** may become the next relevant issue



Next m/w steps



Next LCG-1 upgrades:

- The same software recompiled with gcc 3.2
- New VDT? Based on Globus 2.4? LCG will work on this issue.
- Add VOMS
- Add RLI? – see discussion of RLS
- R-GMA now seems to be off the table
- Working group to resolve data access issues:
 - Components exist: SRM, GFAL, RLS, gridFTP;
 - need to make coherent based on use-cases and integrate with MSS's

For more features we would like to apply the simple rule:

Add it if and only if

- it is proven to work (by EDG, LCG, together)
- it adds some desirable feature or feature requested by users
- it makes the user's application setup significantly simpler or practical
- it is required by new user applications

Bug fixing will always have the highest priority.

Target release date for 2004 system: November



LCG-1 Deployment



- Deployment started (with pre-release tag) in July, to original 10 Tier 1 sites
 - CERN, BNL, CNAF, FNAL, FZK, Lyon, Moscow, RAL, Taipei, Tokyo
- Installation was finished at:
 - CERN, CNAF, FZK, RAL, Taipei, Tokyo
 - FNAL, BNL, Lyon, Moscow seemed to have problems with resources (people)
- Following request at July GDB
 - Budapest expressed willingness to test deployment and successfully installed
- Followed by Barcelona (PIC)
- Prague also expressed interest but have not yet installed ...

- Official “certified” LCG-1 release (tag LCG-1.0.0) was available on 1 September at 5pm CET
 - Was installed at CERN, Taipei, CNAF, Barcelona, Tokyo 24 hours later(!)




Deployment progress



Site	LCG-0	LCG-1 pre	LCG-1.0.0
CERN	Done	16-7-03	1-9-03
CNAF	Done	25-7-03	2-9-03
Taipei	Done	30-7-03	2-9-03
Tokyo	Done	08-8-03	2-9-03
Barcelona		29-8-03	2-9-03
FZK	Done	23-7-03	Vacation?
RAL	Done	30-7-03	Today?
FNAL	Done	x	8-9-03
BNL	x	x	This week?
Lyon	x	x	??
Moscow	x	28-8-03	3-9-03
Budapest		24-8-03	4-9-03
Prague		-	??




Some comments



- Some sites are very responsive and effective:
 - CNAF, Taipei, Tokyo, Barcelona, Budapest
 - FZK – but only 1 person who is now on vacation
- Some completely unresponsive and react very slowly:
 - Lyon, BNL
- Some clearly have other priorities:
 - RAL, FNAL – admin clearly working on something and very helpful on mail list but system not up
- Most clearly do not have resources to provide a production service (i.e. that works during vacations)
- This could be a problem!
- The upgrade from the LCG-1 pre-release to LCG-1 was very effective – 5 sites within 24 hrs, 2 more within 3 days, but the pre-requisite was the infrastructure being in place
 - Prague is just starting, FZK the 1 admin is on vacation, RAL unknown



LCG Resource Commitments – 1Q04



	<i>CPU (kSI2K)</i>	<i>Disk TB</i>	<i>Support FTE</i>	<i>Tape TB</i>
<i>CERN</i>	700	160	10.0	
<i>Czech Republic</i>	60	5	2.5	
<i>France</i>	420	81		
<i>Germany</i>	207			62
<i>Holland</i>	124			12
<i>Italy</i>			16.0	100
<i>Japan</i>		45	5.0	100
<i>Poland</i>		9	5.0	28
<i>Russia</i>	120	30	10.0	40
<i>Spain</i>	220	30	4.0	120
<i>Sweden</i>	150	30	4.0	100
<i>Switzerland</i>	179	40	2.0	40
<i>UK</i>	26	5	2.0	40
<i>USA</i>	1656	226	17.3	295
<i>USA</i>	801	176	15.5	1741
Total	5600	1169	120.0	4223

How realistic are these numbers – especially FTE?
Double, triple counting?




Major C&T Tasks



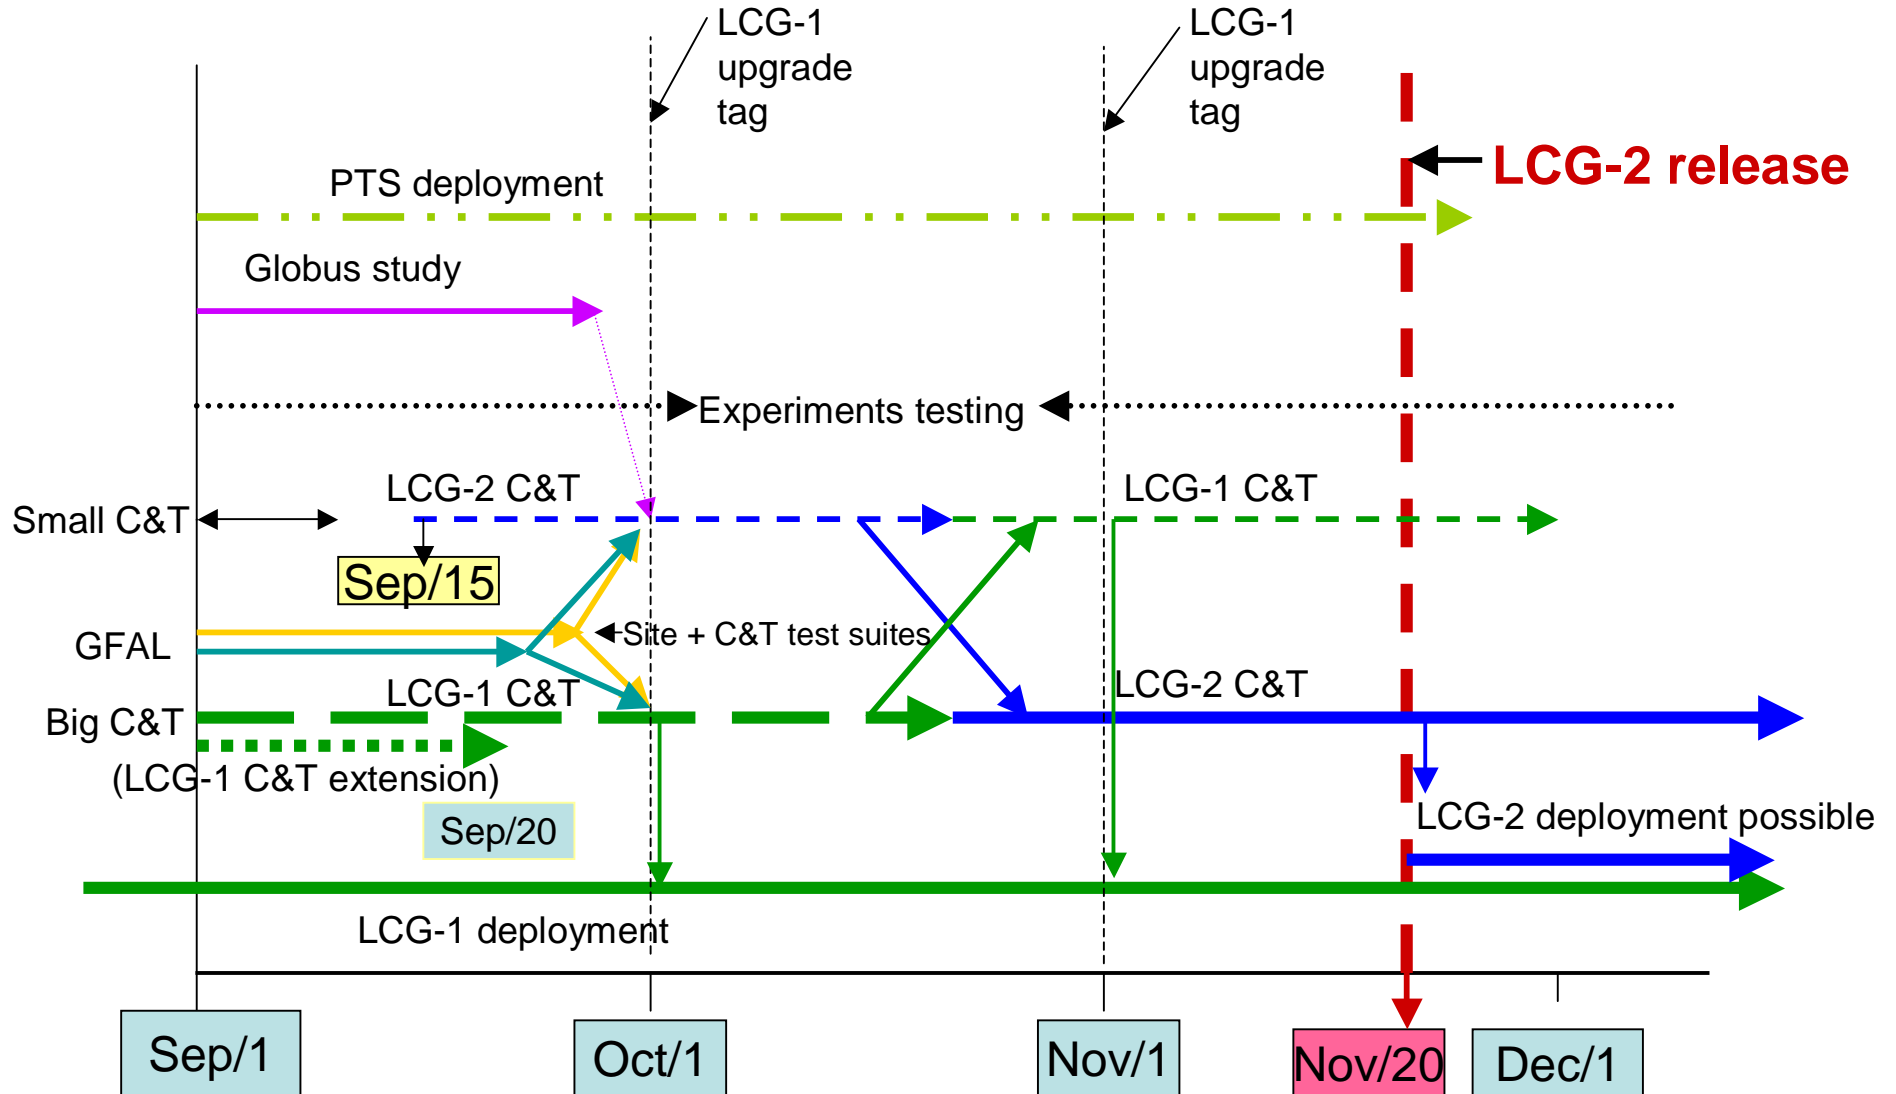
Major tasks until end of the year:

- Support LCG-1 (currently deployed version, gcc2.95)
- Execute test matrix to get certification baseline (ongoing)
- Support experiment/POOL testing on LCG-1
- Add **site verification suite** to the LCG-1 distribution
- Develop and integrate the **C&T verification suite**
- Integrate GFAL
- Evaluate VDT toolkits, find version of Globus to be used by LCG-2
- Install/develop LCG-2 (this is LCG-1 with gcc3.2.2, VOMS, etc)
- Enhance the C&T testbed for outside connections (UofWisc, Budapest, etc...), max 5 sites if we can find volunteers
- Implement PTS usage in all LCG activities

- This level of effort will stretch our available staff – who can assist?



Timeline





Deployment Planning



- After last GDB – sent out summary of requests for information (as well as summary in the minutes)
 - Regional centres ready to join & when? + support contact names
 - Q on Oracle deployment
 - Confirm security incident response procedures
- Got responses from
 - Germany – expected follow up from GSI – not received
 - Spain
 - Czech Rep. – ready in Aug (now working)
 - UK – not conclusive
 - Netherlands – not conclusive
 - Poland
 - !!!!!!**
 - Later approached by Hungary, Pakistan, Spain
- We need now responses – we must now bring in Tier 2 sites, but they must be supported primarily by their Tier 1 (or regional primary site)




Experiments and LCG-1



- Would like to get experiments testing LCG-1 now
 - Discussions with experiments
 - Scheduling – we need to monitor the system as the tests run to understand the problems
 - Migrate CMS LCG-0 to LCG-1
 - Atlas, US_Atlas (want to demonstrate interoperability)
 - ALICE – continue with tests started by Loose Cannons
 - LHCb ?
 - We are scheduling these tests now



Operations tasks



- To make a real production system – now we have a basis in place – many tasks to be done – we need help
- High priority (needed now)
 - Experiment sw distribution mechanism
 - Monitors to watch essential system resources on essential services (/tmp, etc)
 - System cleanup procedures
 - System auditing – must ensure procedures are in place
 - Need basic usage accounting in place
 - Need tool independent WN installation procedure – also for UI
 - Integration with MSS (setting up task force)
 - NB sites with MSS will need to implement interfaces
 - Integration with LXBatch (and others)
 - Standard procedures – we will start but needs a team from sites and GOC
 - for setting Runtime Environments
 - Change procedures
 - Operations
 - Incident handling
- Need people to work on these – many are GOC responsibilities, but where is the staffing?
- LCG deployment team is 2 FTE + 1 security




RLS



- Issue – 2 non-interoperable implementations
- **Proposed Strategy – development**
- Integrate POOL with Globus RLS, so that it is able to communicate with both RLS implementations (but not from the same process).
 - Assuming all the above conditions are met the earliest this work could be completed would be November 2003.
 - This requires close collaboration between the POOL group and someone (a developer?) very familiar with the Globus RLS.
 - Also might require some additions/changes in the RLS API?
- POOL group provide the tools to enable cross population of POOL file catalogs between RLS implementations. These tools are basically available now.
- Globus RLS is ported to Oracle as the database back-end.
- In parallel we work on the interoperability roadmap, with a target date of May 2004 for this to be available:
 - 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)




RLS Proposed Strategy – services



- Now: Run EDG RLS at CERN and at US Tier 1 sites – at least until Globus RLS is running with Oracle.
 - The EDG RLS in this scenario is the LRC only – we will not deploy the RLI.
- By January 2004: LCG service is provided by local EDG LRCs or Globus RLS, and cross-population tools to enable catalog updates. This is the minimal solution for the 2004 Data Challenges.
 - Most batch production work will in any case use bulk updates of the catalogs rather than file-by-file updates from the job.
 - Suggest initially CERN catalog gets all updates and then pushes back out. This implies that pre-job file replication is required so that data is already at site where the job will run.
 - This model removes the requirement for a proxy replica manager since the bulk updates can be run from externally visible nodes.



RLS Outstanding issues:



- Need experiment buy-in
- Can US groups provide help for POOL implementation for Globus RLS?
- Globus RLS (LRC+RLI) must be ported to Oracle
- Need to understand EDG obligations for rest of 2003 (i.e. can they drop EDG RLI?)



Summary of Issues



- Resources – need honest estimates that we can rely on
 - Otherwise planning is irrelevant
- Priorities
 - LCG-1 priority over Oracle/RLS installation tests
- Priorities
 - Need clear statements – administrators have been confused, GDB agreements have not been communicated to them
- Many tasks need to be worked on – need people to work with deployment group – should be under auspices of GOC
 - This is urgent – but not happening?
- Certification – support and essential tasks will stretch team – need help