



# Deployment & Experiment Integration

**Flavia Donno & Markus Schulz**  
LCG

LCG Review  
24 November 2003



# History



- First set of reasonable middleware on C&T Testbed end of July (PLAN April)
  - limited functionality and stability
- Deployment started to 10 initial sites
  - Focus on establishing procedures
  - Training sites (sent personnel to 2 sites)
- End of August only 5 sites in
  - Underestimation of the effort and dedication needed
    - Complexity of the middleware, installation and configuration
    - Lack of experience with install/config tool
- 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 only 24h.
- Sites need between 3 days and several weeks to come online
  - All sites using the fabric management tool

**middleware was late**

**Estimation of duration of the deployment process was correct**



# Release History



Overview: 5 releases up to now

Tag	Date	Comment
lcg1_20030717_1455	17 Jul	Pre-release:CNAF, CERN
LCG1-1_0_0	01 Sep.	EDG pre 2.0, several fixes by LCG
LCG1-1_0_1	19 Sep.	Security Patch (10sites)
LCG1-1_1_0	24 Oct.	Fixes, new WLM (17 sites)
LCG1-1_1_1	05 Nov.	Experiment SW
LCG1-1_1_2	Xx Nov.	Experiment SW (23 sites)



# LCG-1 Deployment Status



- Up to date status can be seen ([here](#)) expect >20 by end of 2003

- PIC-Barcelona** (RB)
  - IFIC Valencia (RB)
  - Ciemat Madrid
  - UAM Madrid
  - USC Santiago de Compostela
  - UB Barcelona
  - IFCA Santander (RB)
- BNL**
- Budapest** (RB)
- CERN** (RB)
- CNAF** (RB)
  - Torino
  - Milano

- FNAL**
- FZK**
  - Krakow
- Moscow** (RB)
- Prague**
- RAL** (RB)
  - Imperial C.
  - Cavendish
- Taipei** (RB)
- Tokyo**

**Total number of CPUs ~150**  
(current focus on # of sites)

## Sites to enter soon

CSCS Switzerland, Lyon, NIKHEF  
Several tier2 centres in Italy

## Sites preparing to join

Pakistan, Sofia,

CPUs added on experiments request

## Users:

EDG: Experiment independent testers  
Experiments: Alice, Atlas, CMS LHCb



# LCG-1 Deployment Status Overview



Site's Status Page	CVS	Site Tag in CVS	LCG tag	Installation Status	Operation	Comment
<a href="#">PIC-BARCELONA (pic.ifae.es)</a>	Yes	LCG1-1_0_1_PIC_20030923_1836	LCG1-1_1_1	Installed	S Running	LCG-ES Coordinator South MDS
<a href="#">IFIC-VALENCIA (ific.uv.es)</a>	Yes	LCG1-1_1_1_IFIC_20031105_1800	LCG1-1_1_1	Installed	S Running	LCG-ES
<a href="#">CIEMAT-MADRID (ciemat.es)</a>	Yes	LCG1-1_1_1_CIEMAT_20031105_1800	LCG1-1_1_1	Installed	S Running	LCG-ES
<a href="#">UAM-MADRID (uam.es)</a>	Yes		LCG1-1_1_1	Installed	S Running	LCG-ES
<a href="#">USC-SANTIAGO-DE-COMPOSTELA (usc.uscgs.es)</a>	Yes	LCG1-1_1_1_USC_20031105	LCG1-1_1_1	Installed	S Running	LCG-ES
<a href="#">UB-BARCELONA (bcm.ub.es)</a>	Yes	LCG1-1_1_1_UB_20031106_0001	LCG1-1_1_1	Installed	S Running	LCG-ES
<a href="#">IFCA-SANTANDER (ifca.unicon.es)</a>	No		LCG1-1_1_1	Installed	S Running	LCG-ES
<a href="#">BNL (bnl.gov)</a>	Yes	LCG1-1_1_1_BNL_20031107_1525	LCG1-1_1_1	Installed	W Running	-
<a href="#">BUDAPEST (kfti.hu)</a>	Yes	LCG1-1_1_1_BUDAPEST_20031110_1535	LCG1-1_1_1	Installed	E Running	
<a href="#">CERN (cern.ch)</a>	Yes	LCG1-1_1_1_CERN_20031107_0857	LCG1-1_1_1	Installed	E Running	East MDS
<a href="#">CSCS-SWITZERLAND (cscs.ch)</a>	initial		LCG1-1_1_0	Installing	E Offline	
<a href="#">CNAF (cnaif.infn.it)</a>	Yes	LCG1-1_1_1_CNAF_20031105_1630	LCG1-1_1_1	Installed	S Running	South MDS
<a href="#">INFN-TORINO-LCG1 (to.infn.it)</a>	No		LCG1-1_1_1	Installed	S Running	S-site
<a href="#">INFN-MILANO-LCG1 (mi.infn.it)</a>	No		LCG1-1_1_1	Installed	S Running	S-site
<a href="#">FNAL (fnal.gov)</a>	Yes	LCG1-1_1_1_FNAL_20031106_1130	LCG1-1_1_1	Installed	W Running	
<a href="#">FZK (fzj.de)</a>	Yes	LCG1-1_1_1_FZK_20031105_1907	LCG1-1_1_1	Installed	E Running	Primary site
<a href="#">KRAKOW (cyf-kr.edu.pl)</a>	Yes	LCG1-1_1_1_CYFRONET_20031118_1436	LCG1-1_1_1	Installed	E Running	Secondary site
<b>LYON</b>	Yes		LCG1-1_1_1	Installed	E Offline	Missing host certificates, web page
<a href="#">MOSCOW (snp.msu.ru)</a>	Yes	LCG1-1_0_0_MOSCOW_20030903_1946	LCG1-1_1_1	Installed	E Running	
<a href="#">PRAGUE (bmm.particle.cz)</a>	Yes	LCG1-1_1_1_PRAGUE_20031107_1317	LCG1-1_1_1	Installed	E Running	
<a href="#">RAL (rdpp.rl.ac.uk)</a>	Yes	LCG1_1_1_1_RAL_20031106_1437	LCG1-1_1_1	Installed	W Running	West MDS P-site
<a href="#">IMPERIAL (hep.ph.ic.ac.uk)</a>	Yes	LCG1_1_1_1_IMPERIAL_20031110_1715	LCG1-1_1_1	Installed	W Running	S-site, dead link
<a href="#">CAVENDISH (hep.phy.cam.ac.uk)</a>	Yes	LCG1-1_1_1_CAVENDISH_20031111_1220	LCG1-1_1_1	Installed	W Running	S-site, wrong format web page
<a href="#">TAIPEI (grid.sinica.edu.tw)</a>	Yes		LCG1-1_1_1	Installed	E Running	East MDS
<a href="#">TOKYO (lscgpn.s.u.tokyo.ac.jp)</a>	Yes	LCG1-1_1_0_TOKYO_20031027_2012	LCG1-1_1_1	Installed	E Running	



# LCG-1 Site Config in CVS

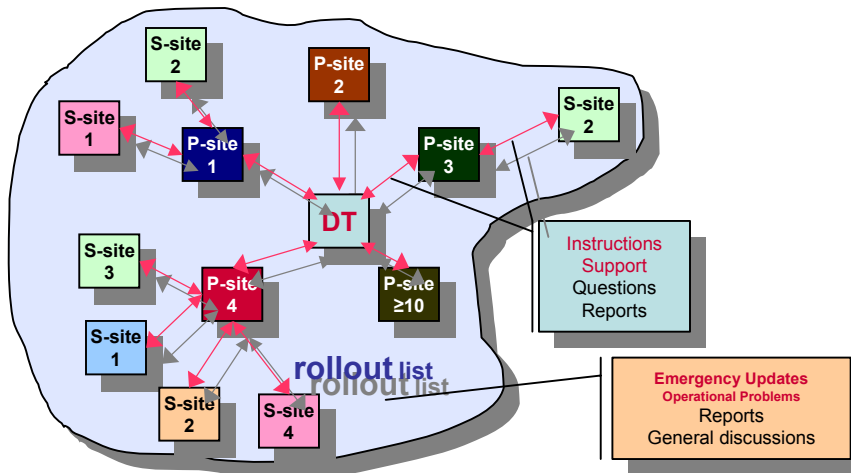


File	Rev.	Age	Author	Last log entry
<a href="#">Attic/</a> <a href="#">[show contents]</a>				
<a href="#">budapestWorkerNode.h</a>	1.2	2 months	gdebrecz	First version using LCG1-1_0_0. Simple LDAP queries, gridftp and simple job su...
<a href="#">cfgdir-cfg.h</a>	1.2	2 months	gdebrecz	First working configuration !
<a href="#">do_mkxprof.sh</a>	1.1	2 months	emanuele	First version of the config files for Budapest site
<a href="#">grid100</a>	1.5	2 hours	gdebrecz	Slight modifications in UI config.
<a href="#">grid109</a>	1.2	2 months	gdebrecz	First version using LCG1-1_0_0. Simple LDAP queries, gridftp and simple job su...
<a href="#">grid110</a>	1.4	2 hours	gdebrecz	Slight modifications in UI config.
<a href="#">grid111</a>	1.2	2 months	gdebrecz	First version using LCG1-1_0_0. Simple LDAP queries, gridftp and simple job su...
<a href="#">grid112</a>	1.2	2 months	gdebrecz	First version using LCG1-1_0_0. Simple LDAP queries, gridftp and simple job su...
<a href="#">grid113</a>	1.2	2 months	gdebrecz	First version using LCG1-1_0_0. Simple LDAP queries, gridftp and simple job su...
<a href="#">grid114</a>	1.2	2 months	gdebrecz	First version using LCG1-1_0_0. Simple LDAP queries, gridftp and simple job su...
<a href="#">grid115</a>	1.2	2 months	gdebrecz	First version using LCG1-1_0_0. Simple LDAP queries, gridftp and simple job su...
<a href="#">grid116</a>	1.2	2 months	gdebrecz	First version using LCG1-1_0_0. Simple LDAP queries, gridftp and simple job su...
<a href="#">grid117</a>	1.1	2 hours	gdebrecz	New working node.
<a href="#">local-cfg.h</a>	1.2	2 months	gdebrecz	First working configuration !
<a href="#">nfsmount-cfg.h</a>	1.3	2 months	gdebrecz	First working configuration !
<a href="#">site-cfg.h</a>	1.5	2 hours	gdebrecz	Slight modifications in UI config.

Show files using tag:

- Keeping track of site configuration
- **Central** CVS repository at CERN
- **ALL** sites, **ALL** configuration, **ALL** versions
- Helps in problem tracking
- First version of config provided by CERN or tier1 center

- Well established procedure (C&T presentation)
  - Software handed to the Deployment Team by C&T
    - Adjustments in the configuration
    - Update of documentation and templates (in CVS)
    - Final installation tests
- How do we deploy?
  - Service Nodes (RB, CE, SE ...)
    - LCFGng (fabric management tool from EDG),
    - We provide for new sites config files based on a questionnaire
  - Worker and User interface nodes (by tool and manual instruction)
- **Communication:** list [LCG-rollout@rl.ac.uk](mailto:LCG-rollout@rl.ac.uk) (~10 mails/day)

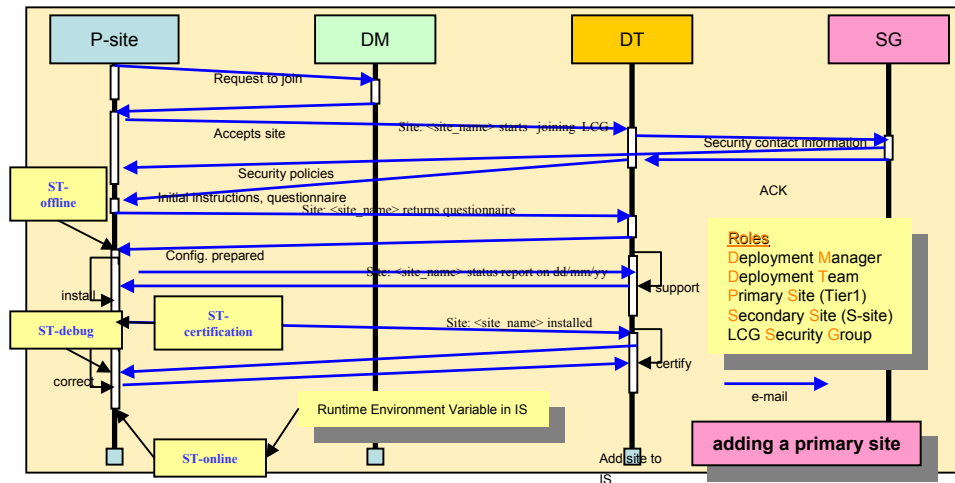
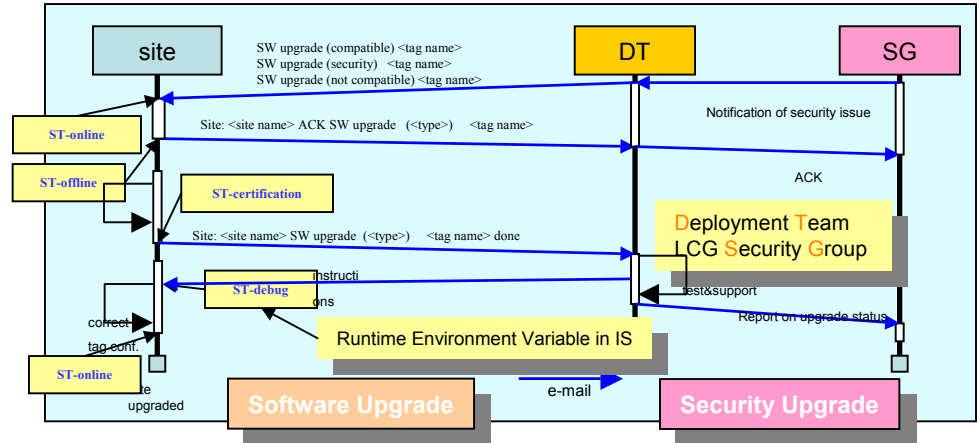


Roles	Description
<b>Deployment Team</b>	Prepares releases, deploys first, supports P-sites and handles escalated problems
<b>Primary site</b>	Experienced site with resources to support some S-sites
<b>Secondary site</b>	Inexperienced site (can have more <u>computing</u> than P-site)

## Procedures defined for:

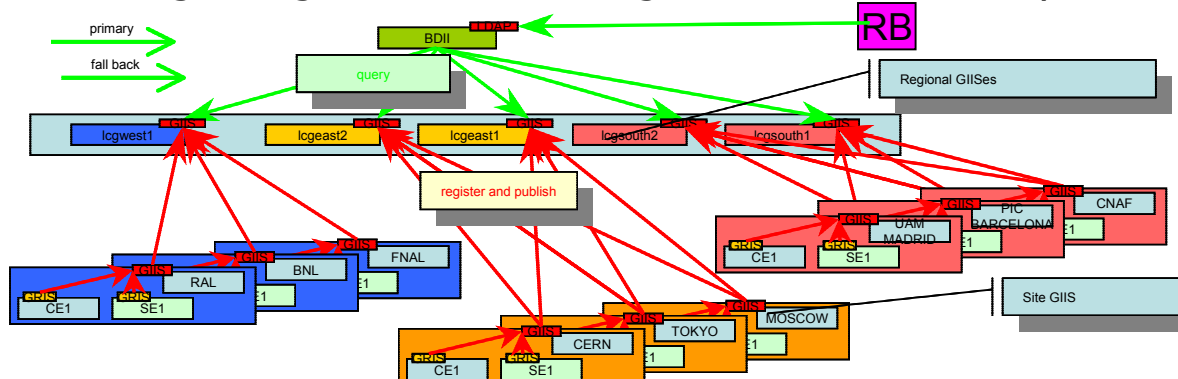
- adding primary/secondary sites
- software upgrades
- security upgrades

- <http://cern.ch/markusw/JoiningLCG.doc>
- <http://cern.ch/markusw/JoiningLCG.pdf>
- <http://cern.ch/markusw/JoinLCG.html>





- The Information System (IS) is the nervous system of LCG
  - Used by almost all services (RB, Replica Manager, RLS, ...) to
    - Discover resources and their properties (*static and dynamic*)
  - Based on Globus MDS
  - Know scalability problems with MDS
    - Number of sites, Amount of data
    - Fatal handling of failures that propagate through the hierarchy
- Ongoing effort to make MDS more robust
  - Modified EDG-BDII replacing top level MDS
    - BDII == database + LDAP server + perl script to query MDS
    - LCG improved version: no stale information, redundant sources
    - Partitioning in regions: less load/region, confinement of problems





# Security (related to deployment)



- LCG Security Group (Dave Kelsey (RAL))
  - Defines policies
- Deployment Team implements policies
  - LCG registration <http://lcg-registrar.cern.ch/>
  - CERN Certification Authority
    - [http://lcg-registrar.cern.ch/pki\\_certificates.html](http://lcg-registrar.cern.ch/pki_certificates.html)
  - Tools for VO management
  - Host VO for LCG1 users and dteam
    - Experiment VOs are run at **NIKHEF**
  - Distribution of security policies to sites
  - Maintains security contacts



# Experiment Integration



## Goal:

Help experiments integrating their production and analysis environment with LCG

- One person assigned to each experiment. But work on global scope.
- Guides and manuals for users and developers
  - **LCG-1 User Guide** (<https://edms.cern.ch/file/412777/1/LCG-1-UserGuide.pdf>)
  - Interface definition for Workload, Data Management and POOL software (<https://edms.cern.ch/file/384019/0.4/WP1-WP2.doc>)
  - The LCG-1 Information System ([https://edms.cern.ch/file/384587/0.2/LCG-1\\_Information\\_System.pdf](https://edms.cern.ch/file/384587/0.2/LCG-1_Information_System.pdf))
  - Experiment Software Installation on LCG-1 (<https://edms.cern.ch/file/412781/1/SoftwareInstallation.pdf>)

Start Here



# Experiment Integration



- Providing assistance/testbed to exercise/integrate new middleware features
- Ongoing activities:
  - ALICE: AliEn tests on LCG-1  
(<https://wwwlistbox.cern.ch/earchive/alice-support-lcg-eis>)
  - Significant effort to create the CMS LCG-0 testbed:
    - **real production done and produced 2 million events**(<http://cmsdoc.cern.ch/cms/LCG/LCG-0/>)
  - ATLAS exercises with software installation via PACMAN in the new proposed Experiment Software Installation LCG Tools; Ongoing testing with Grid3  
(<https://wwwlistbox.cern.ch/earchive/support-eis>)
  - CMS integration with POOL. Exercise with usage of catalogues.  
(<http://server11.infn.it/archive-cms-lcg-edt/>)
  - Distribution, installation and configuration of experiment software
    - Tools provided and under test
    - Some open issues
- Identifying missing functionality
  - Aggregating the experiments requirements

Direct channel for the experiments into the LCG deployment



# Experiment Status on LCG



- Alice, Atlas, CMS, and LHCb on LCG-1
  - Basic functionality tested
- Extensive list of problems reported
  - many configuration related
  - helps defining a better validation procedure for sites
- Preparing for Data Challenges
- Providing first level user support
  - until grid user support at FZK is in full operation



# Problems (deployment)



- Installation too complex
  - Components are too interdependent
  - Manual install procedures only supported for WNs and Uis (now)
  - Better installation guide needed
- Debugging site configurations
  - Discovery of the remote site's setup is hard
  - Changes take a long time
  - Misleading error messages during installation
  - Site certification procedures not adequate
- Some sites are in contact with grids for the 1st time
  - "Beginners Guide to Grids" needed
- Time zones slow down the propagation of changes



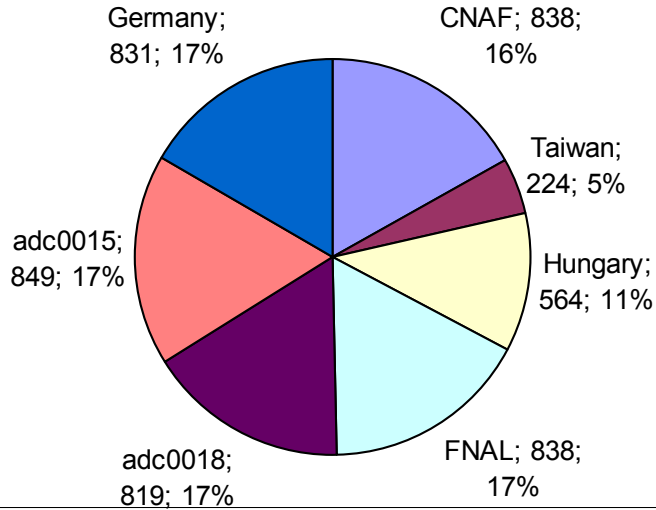
# Stability & Operation



- Stability greatly improved with each release
- Only few MDS related problems (improvements under test)
- Focus now on: Hardening services for production
  - Jobs with realistic workload
  - Chaotic usage test
  - Integration with local production fabric
  - Operate services for extended periods
    - Do they “age” or “pollute” the platforms they are running on?
  - Capture “state” of services to restart them without losing active jobs
  - Learn how to upgrade services (RMC, LRC...) without stopping
    - LCG1 can't be drained for upgrading
  - Integration of new components (GFAL, MSS, managed storage)
  - Work together with EIS on providing user level software installation mechanisms.



### Geographical Job Distribution



# Test



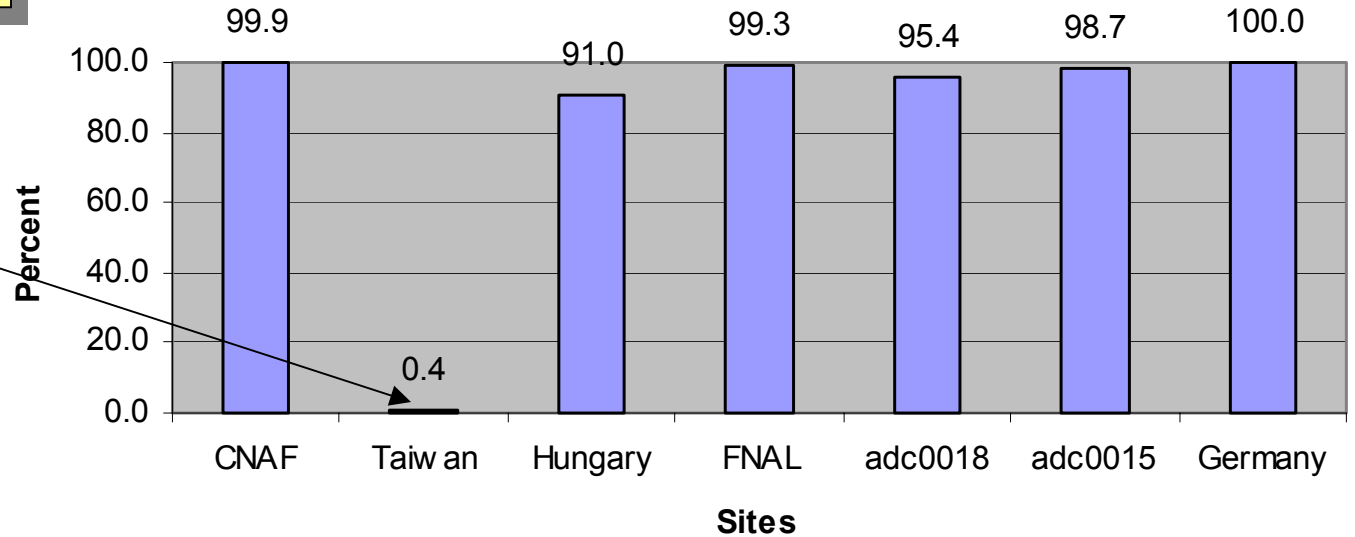
## LCG 1.0 Test (19./20. Sept. 2003):

- 5 streams
- 5000 short jobs

Ingo Augustin  
from EDG WP8

Taipei changed configuration during the test.  
Not a failure of the site

### Ratio of succesful jobs of retrieved jobs (4963 of 5000 = 99.26%)







# Summary



- Middleware was 3 months late
  - Less: functionality, tests, experience with operation
- Number of sites now at scale foreseen (23 sites)
  - Deployment process seems to work
  - Need better site certification
- Experiments are testing the system
  - good end user documentation
  - discover problems (config. errors)
  - SW- distribution process implemented, needs testing/acceptance
- Very little time to turn this into a real production system
  - Critical components are just arriving (SE)
  - Has to be done incrementally on the running service
- Deploying the software at new sites is not always easy
  - various reasons (attitude, complexity, priorities, acceptance of tools)