



22/03/2012

LCG-France Tier-1 & AF

Réunion mensuelle de coordination

Pierre Girard

Pierre.girard@in2p3.fr

dapnia
ceci
saclay

CNRS
CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE



Plan



- Les rats quittent le navire...
- Nouvelles de LCG
 - CR du dernier CP de LCG-France
 - CR du dernier T1SCM
 - CR du dernier GDB
- Disponibilité du site
- Dossiers en cours
- Evénements



Les rats quittent le navire...



■ Mode “passage de main”

- Responsable du projet LCG@CC.IN2P3.FR
 - Quoi
 - Représentation auprès de LCG France
 - Représentation auprès de WLCG
 - Activités WLCG au CCIN2P3
 - Comment transmettre les connaissances
 - Oralement (ok qd qq'un à qui parler)
 - Wiki ? EDMS ?
- Administrateur grille (sysgrid)
 - Discuté avec le chef de l'équipe Infrastructure
- Validation certificats GRID2-FR du CC
 - Fait
 - Nouveaux contacts
 - Jérôme Bernier (Primaire)
 - Benoît Delaunay (Secondaire)

Nouvelles de LCG

CR du dernier CP de LCG-France



■ CP du 27 février 2012

- « *L'activité des sites en 2011 indique que l'infrastructure composée de 9 Tier-2 Tier-3 (en tenant compte de l'AF du CC) autour du Tier-1 répond avec efficacité aux besoins de calcul des expériences.* »
- « *Compte tenu des baisses continues des budgets alloués par les TGIR pour le CCIN2P3 et LCG, et donc des notifications décroissantes aux T2/T3, il est décidé de préparer une **concertation sur l'avenir du calcul en général et celui du LHC en particulier avec tous les partenaires, expériences et directeurs des laboratoires hébergeant des sites LCG-France**. L'objectif est de décider une politique claire de l'IN2P3 avec priorités sur les 3-4 prochaines années au plus tard à la fin de l'été 2012.* »



CR du dernier T1SCM



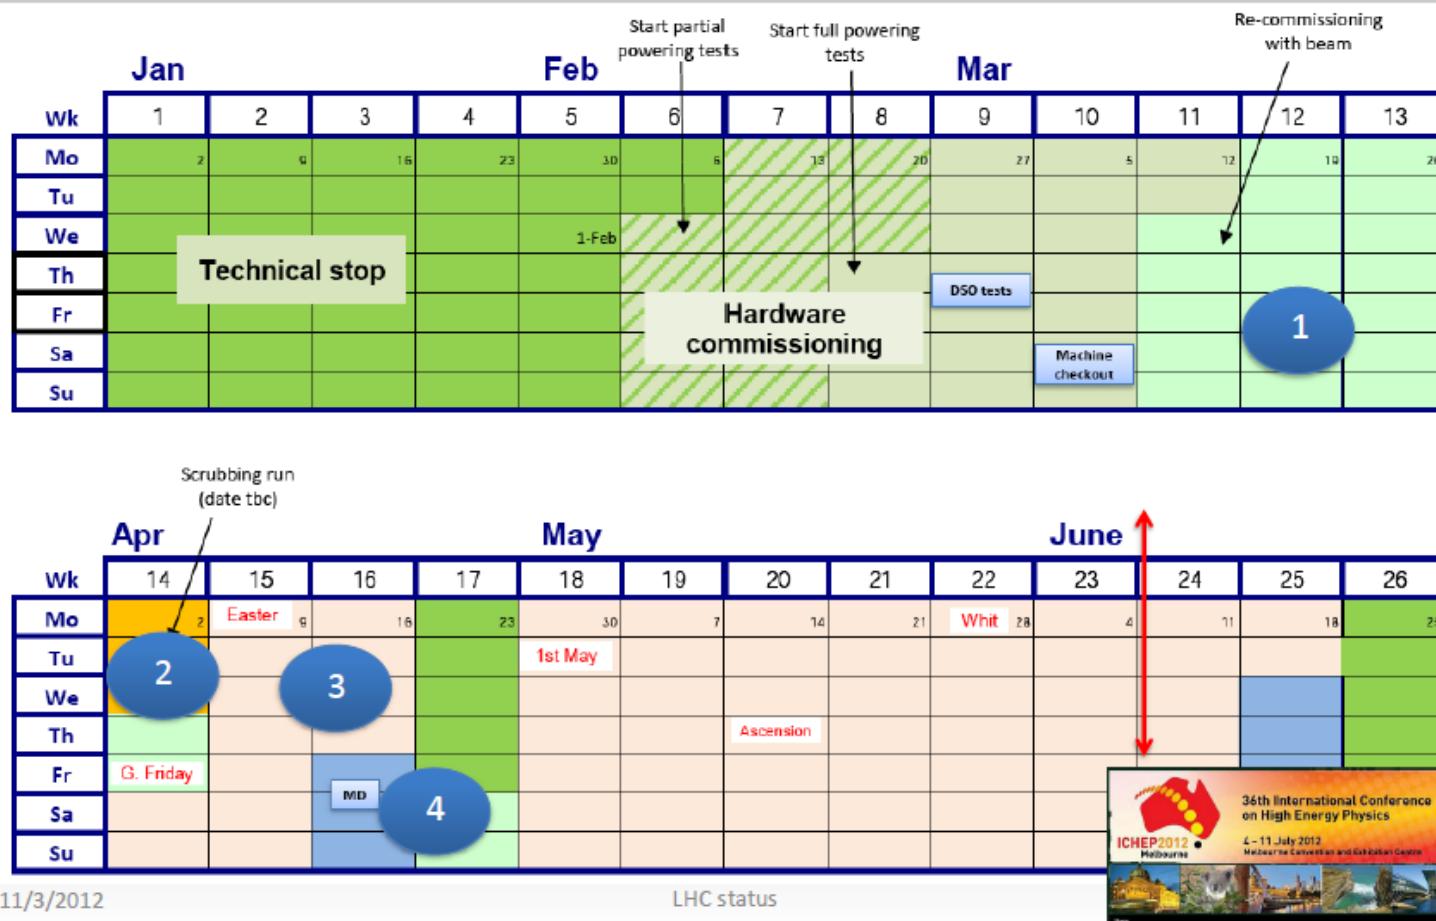
- T1 Service Coordination Meeting du 15 avril 2012
 - <http://indico.cern.ch/conferenceDisplay.py?confId=181610>

LHC Machine Plans for 2012

Mike Lamont (CERN)



2012 LHC schedule Q1/Q2



LHC Machine Plans for 2012

Mike Lamont (CERN)



2012 LHC schedule Q3/Q4

	July						Aug				Sep				
	27	28	29	30	31	32	33	34	35	36	37	38	39		
Mo	2	8		16	↓	21	30	6	10	20	27	1	10	17	24
Tu															
We															
Th											J. Geneva				
Fr															
Sa															
Su															

	Oct				Nov				Dec				
	40	41	42	43	44	45	46	47	48	49	50	51	52
Mo	1	8	15	22	29	5	12	19	26	3	10	17	24
Tu													Xmas
We													
Th													
Fr													
Sa													
Su													

- Technical Stop
- Recommissioning with beam
- Machine development
- Ion run
- Ion setup

Special runs (TOTEM etc.) to be scheduled

5

High beta* runs for ALFA and TOTEM
Van der Meer scans

FTS 2.2.8 deployment at Tier1s

FTS Transfer Dashboard



- Migration vers FTS 2.2.8 poussée par ATLAS
 - Planning établi entre tous les T1s
 - <https://docs.google.com/spreadsheet/ccc?key=0AthhzXLQok7XdFpUeDBfLXE2S1RDZE4zcHp6QWVpUFE#gid=0>
 - CCIN2P3
 - Lundi 26 mars 2012
 - David ?
- Cette version intègre des sondes de monitoring
- Dashboard en préparation
 - Présentation:
 - http://indico.cern.ch/getFile.py/access?contribId=10&resId=0&materialId=s_lides&confId=181610
 - Démo
 - <http://dashb-wlcg-transfers.cern.ch/ui/>



CR du dernier GDB



■ GDB du 21 mars 2012

- <http://indico.cern.ch/conferenceDisplay.py?confId=155066>

Election du nouveau chairman du GDB



- And the winner is ...



**MICHEL
JOUVIN**

CR du dernier GDB

Today

- Vidyo
- CERN Infrastructure
- GE
- Experiment Operations
- Middleware



VLCG
Worldwide LHC Computing Grid



Science & Technology
Facilities Council

GDB

Milestones

July 2012

- 10% resources in CC managed by IaaS system (1000 nodes) including installation, configuration, etc.

January 2013

- 50 % resources managed (over 3 sites)

July 2013

- 95% resources managed

Initial focus on service/server consolidation use cases

Modernisation of the CERN Computing Infrastructure (Ian Bird)



- Agile Infrastructure project has 3 aspects
 - Management tool chain
 - Use of virtualisation
 - Monitoring
 - Possible utilisation de Hadoop pour le stockage central et l'analyse
- Cloud technology
 - Choix d'Openstack
 - Participation du CERN à “Openstack Foundation”
- Cloud investigations
 - Use of commercial clouds: issues of policy, overall feasibility, data management, cloud models, cost
 - Helix Nebula project is one step in this direction
 - ATLAS est un use case
 - Will consider other direct investigations with other providers
 - CERN Openlab

Role of Helix Nebula: The Science Cloud



ATLAS use case

- Simulations (~no input) with stage out to:
 - Traditional grid storage vs
 - Long term cloud storage
- Data processing (== “Tier 1”)
 - This implies large scale data import and export to/from the cloud resource
- Distributed analysis (== “Tier 2”)
 - Data accessed remotely (located at grid sites), or
 - Data located at the cloud resource (or another?)
- Bursting for urgent tasks
 - Centrally managed: urgent processing
 - Regionally managed: urgent local analysis needs
- All experiences immediately transferable to other LHC (& HEP) experiments

Extension of the Tier 0

- In 2006 we first recognised the need for an extension of the Tier 0
- Eventually it was decided that this should not be at CERN
- Several years of investigations and proposals
- Last year – call for proposals
 - Followed by detailed technical follow ups
 - Led to call for tender
- Results adjudicated last week in the Finance Committee (Council)
- Winner is Budapest, Hungary (Wigner Inst.)

Timescales

- Prototyping in 2012 hopefully
- Testing in 2013
- Production in 2014

Experiment Operations



■ ALICE (Latchezar Betev)

- <http://indico.cern.ch/materialDisplay.py?contribId=0&sessionId=5&materialId=slides&confId=155066>
- The LEGO Analysis Trains
 - Polling together many user analysis tasks (wagons) in a single set of Grid jobs (the train)
 - Managed through a web interface by a Physics Working Group conductor (ALICE has 8 PWGs)
- SLC6 notes
 - Many applications will see a large increase in virtual memory usage (but not resident memory)when running on RHEL 6 vs RHEL 5.
- Re-write of proxy-renewal service (IT/ES – M.Litmaath)
 - Etude d'un déploiement sur nos VO boxes (toutes VOs)
 - Contact avec Maarten pour faisablilité

Experiment Operations



■ LHCb (Philippe Charpentier)

- <http://indico.cern.ch/getFile.py/access?contribId=1&sessionId=5&esId=0&materialId=slides&confId=155066>
- **Many non-pledged sites**
- **1.4 PB free on LHCb-Disk space tokens**
 - Parce que des sites ont déployé leurs pledges 2012 en avance
- **Tape recall and cache management:**
 - We need to know the actual size of the cache read buffers
 - Not available with SRM for dCache
- **4 GB virtual memory now OK**
 - Limit per process, not process group (**5GB in this case!**)
 - If possible only kill the offending process, not the job!
- **SL6 being certified, but not foreseen in production during 2012 data taking**
 - Requires MW python bindings for 2.7

Experiment Operations



■ CMS (Stephen Gowdy)

- <http://indico.cern.ch/materialDisplay.py?contribId=2&sessionId=5&materialId=slides&confId=155066>
- Intéressant organigramme des « computing operations » de CMS
- Software deployment
 - 2 solutions for sites: CVMFS or CRON installation
- Multi-Core CMSSW jobs
 - Cf. présentation
 - Clear memory gain (~20%) with multi-core processing
 - Lazy download produces a large startup overhead for multi-core jobs
 - Small overhead if Lazy download turned off
 - Tier-1 sites will not like to move parts of their resources to multi-core usage

Experiment Operations



■ ATLAS (Alessandro Di Girolamo)

- <http://indico.cern.ch/materialDisplay.py?contribId=3&sessionId=5&materialId=slides&confId=155066>
- Memory requirements (for high pileup MC12)
 - “*there will be jobs which need vmem > 3.5GB*”
 - “*we ask sites to have limits on vmem of 4GB. SWAP should be set to 2*RAM size. Physical memory requirement not changed (2 GB).*”
- LFC migration
 - Done for IN2P3-CC
 - Thanks to Emmanouil, David and Osman
- FTS 2.2.8 functionality
 - ~60% of FTS servers already upgraded, other Tier1s have planned/are planning
 - Issue observed on Tier0 FTS Server
 - Gridftp resume transfers in case of particular failures condition
 - Fix already provided in a nonEMI release, installed on CERN pilot and Tier0 FTS server. Will be included in the next EMI rel.
- CVMFS progress (90 / 150 sites)
- Helix nebula test
 - <https://indico.cern.ch/getFile.py/access?resId=0&materialId=slides&contribId=30&sessionId=3&subContId=3&confId=169695>
- Automation of Operations / Evaluation of sites
 - Site usability for analysis activity
 - could be injected in SAM/NAGIOS

■ CREAM

- **BLAH, [v.1.16.5](#) (12 April 2012)**

- Fixes the job registry corruption caused by the purge of the registry under heavy load

■ UI/WN

- **GFAL/lcg_util, [v. 1.12.0](#)** – minor release

- Glue 2.0 support in gLite data clients
 - Return estimatedWaitTime for files in an asynchronous SRM request in Python
 - Bug fixes:
 - **checksum is case sensitive in lcg-get-checksum/lcg-rep**
 - ...
 - lcg-rep segmentation fault with dcache endpoints

■ Extension of security support of gLite 3.2 glite-WN and glite-UI

- To September 30 2012
- Means that sites can keep gLite until EMI-2 !!!

■ UI/WN tarball support

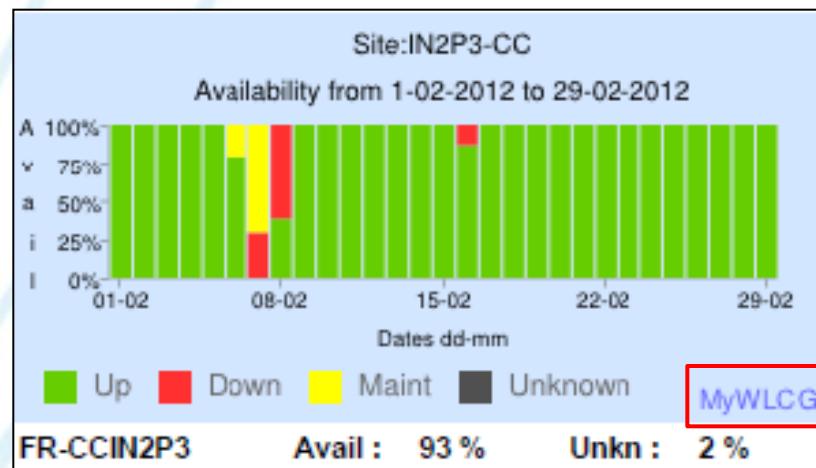
- First testing version available
- Documentation: <https://twiki.cern.ch/twiki/bin/view/EMI/EMIuiwnTar>
- Beta testing

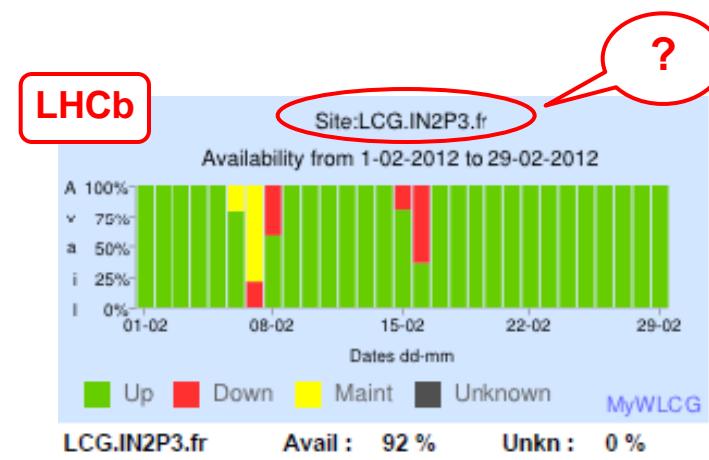
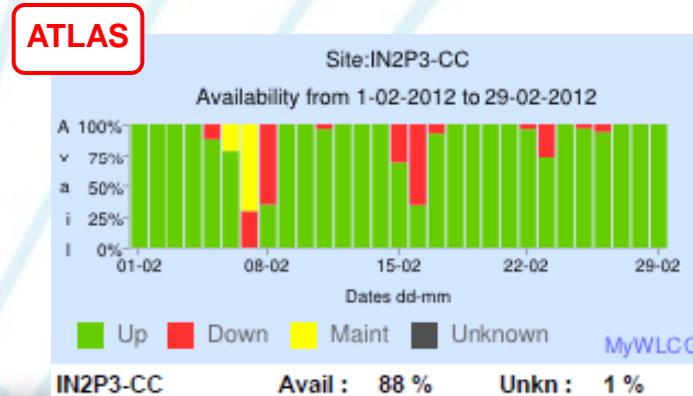
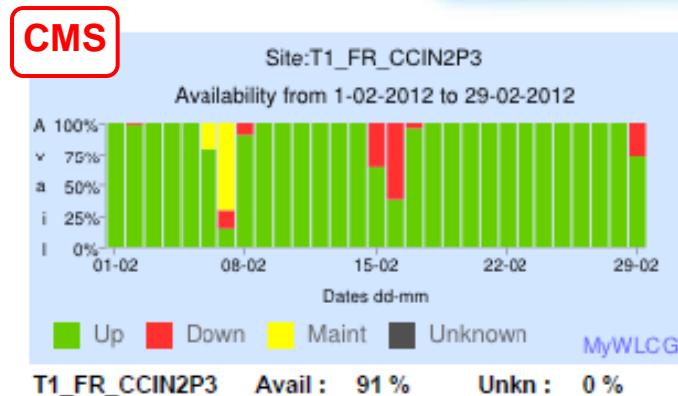
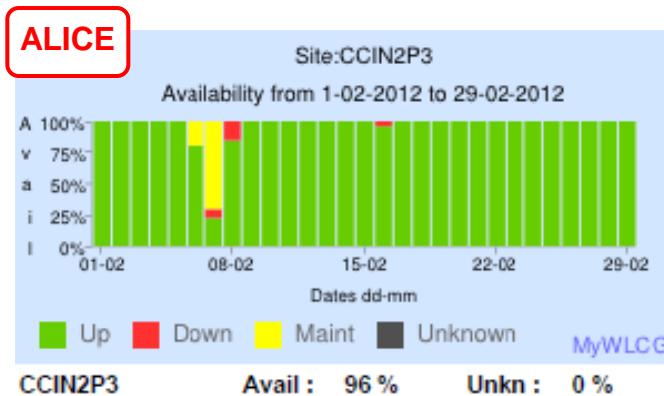
Résultats du site

Résultats Janvier 2012



- MyWLC IN2P3-CC
 - <http://grid-monitoring.cern.ch/mywlcg/sa/?group=1&site=457&graph=1&profile=15>
- LCG Office
 - https://espace.cern.ch/WLCG-document-repository/ReliabilityAvailability/Tier-1/2012/tier1_reliab_0212/Tier1_Summary_Feb2012.pdf





Evénements



- A venir (à mettre à jour)
 - EGI User Forum, Munich, 19-23 March
 - HEPIX Spring Meeting, Prague, 23-27 April 2012
 - EMI AHM, DESY, 8-10 May
 - WLCG Workshop, NYC, 19-20th May 2012
 - CHEP2012, NYC, 21-25th May 2012
 - LCG France, Nantes (SUBATECH), 6-7 June 2012
 - <http://indico.in2p3.fr/conferenceDisplay.py?confId=6455>
 - EGI Technical Forum, Prague, 17-21 September 2012
 - HEPIX Fall Meeting, IHEP, 15-19 October 2012
 - JI IN2P3/IRFU, La Londe les Maures, du 22 au 25 octobre 2012