



13/10/2010

LCG-France Tier-1 & AF

Réunion mensuelle de coordination

Pierre Girard

Pierre.girard@in2p3.fr

dapnia
ceo
saclay

CNRS
CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE



Plan



- Nouvelles de LCG
 - Planning LHC
 - Prochain arrêt du CC
 - CR du derniers GDB
 - Technical Evolution Groups
- Disponibilité du site
- Dossiers en cours
- Evénements



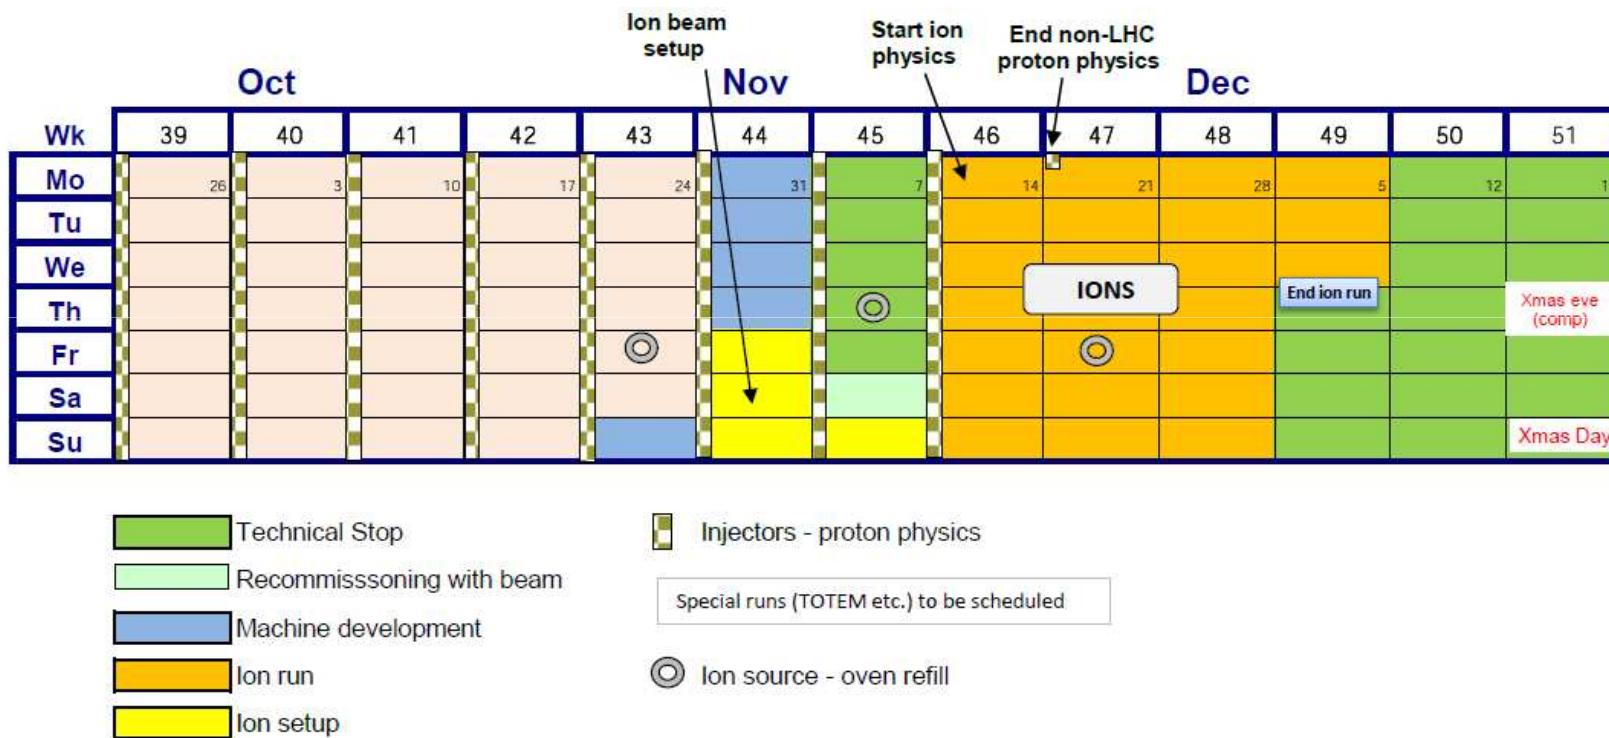
Nouvelles de LCG



Planning du LHC



Source: https://espace.cern.ch/be-dep/BEDepartmentalDocuments/BELHC_Schedule.pdf





Planning du LHC



Source: https://espace.cern.ch/be-dep/BEDepartmentalDocuments/BE/Draft_LHC_Schedule_2012.pdf

ML

14/9/2011

V0.4

2012 LHC Schedule

DRAFT!!!

		Jan	Feb												Mar	Close ring	Re-commissioning with beam
Wk		52	1	2	3	4	5	6	7	8	9	10	11	12			
Mo			26		2	9	16	23	30	6	13	20	27	5			
Tu																	
We																	
Th																	
Fr																	
Sa																	
Su			1														

Technical Stop

Hardware commissioning

Machine checkout

Special runs (TOTEM etc.) to be scheduled

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



Arrêts du CC



- <http://cc.in2p3.fr/Arrets-de-production>

- **6 décembre 2011**
 - 1 jour avant l'arrêt du LHC
 - **7 février 2012**
 - **Update/Upgrade de dCache possible**



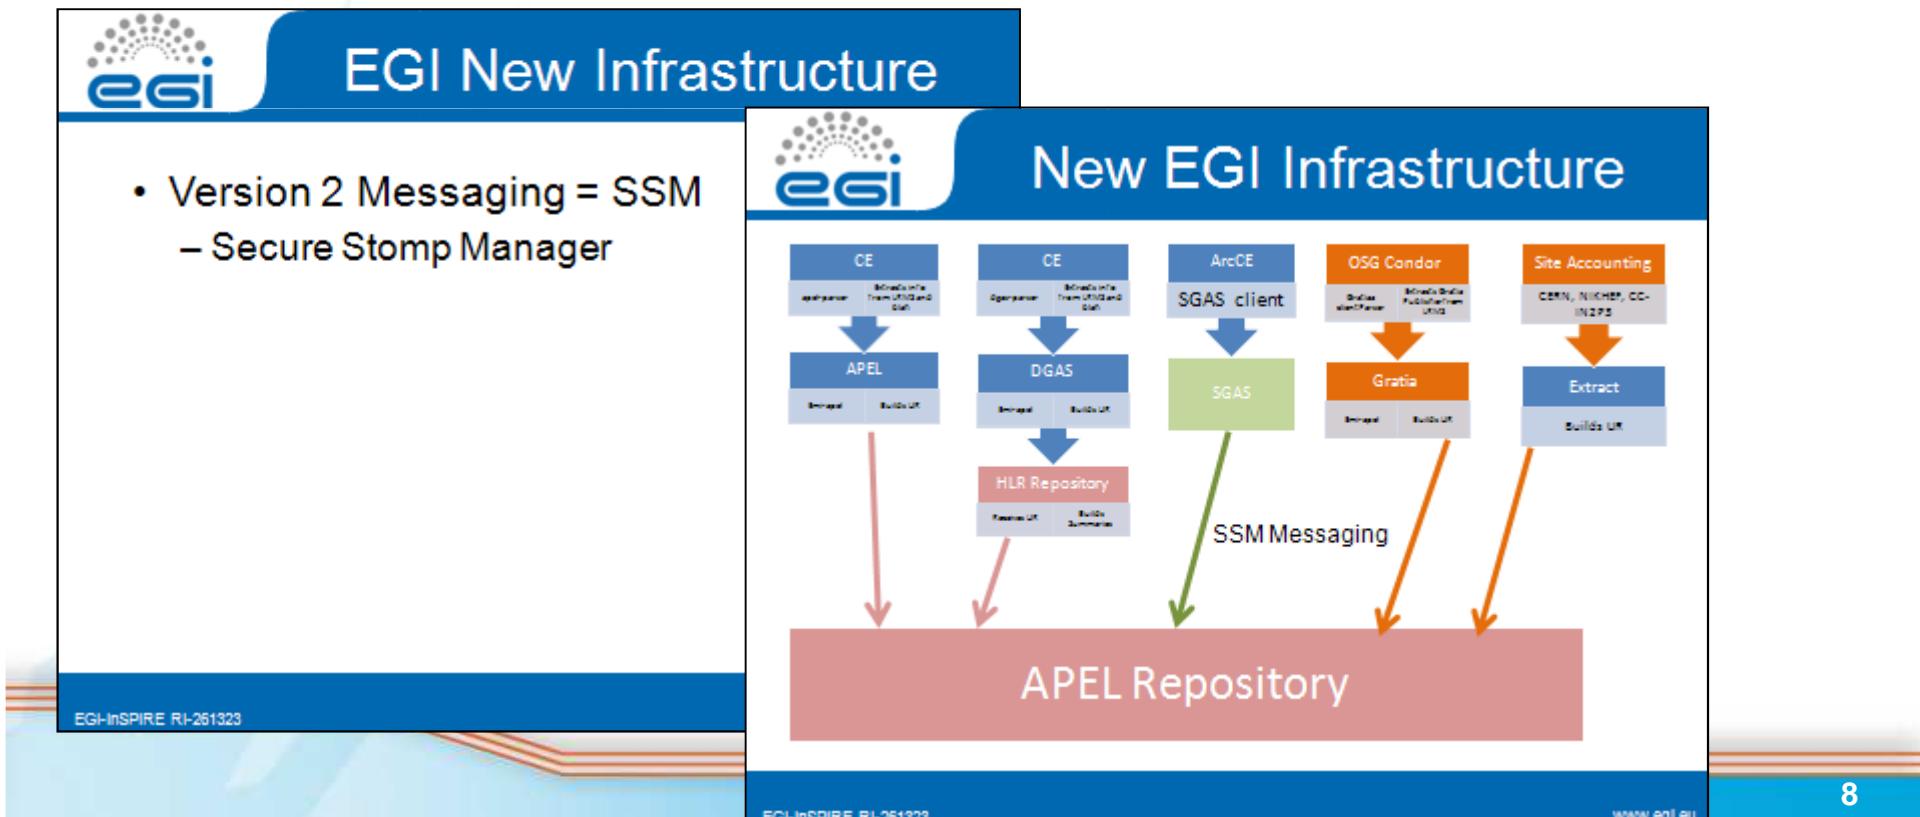
CR du dernier GDB



Accounting



- EGI New Infrastructure
 - [http://indico.cern.ch/materialDisplay.py?sessionId=4&materi alId=1&confId=106650](http://indico.cern.ch/materialDisplay.py?sessionId=4&materialId=1&confId=106650)



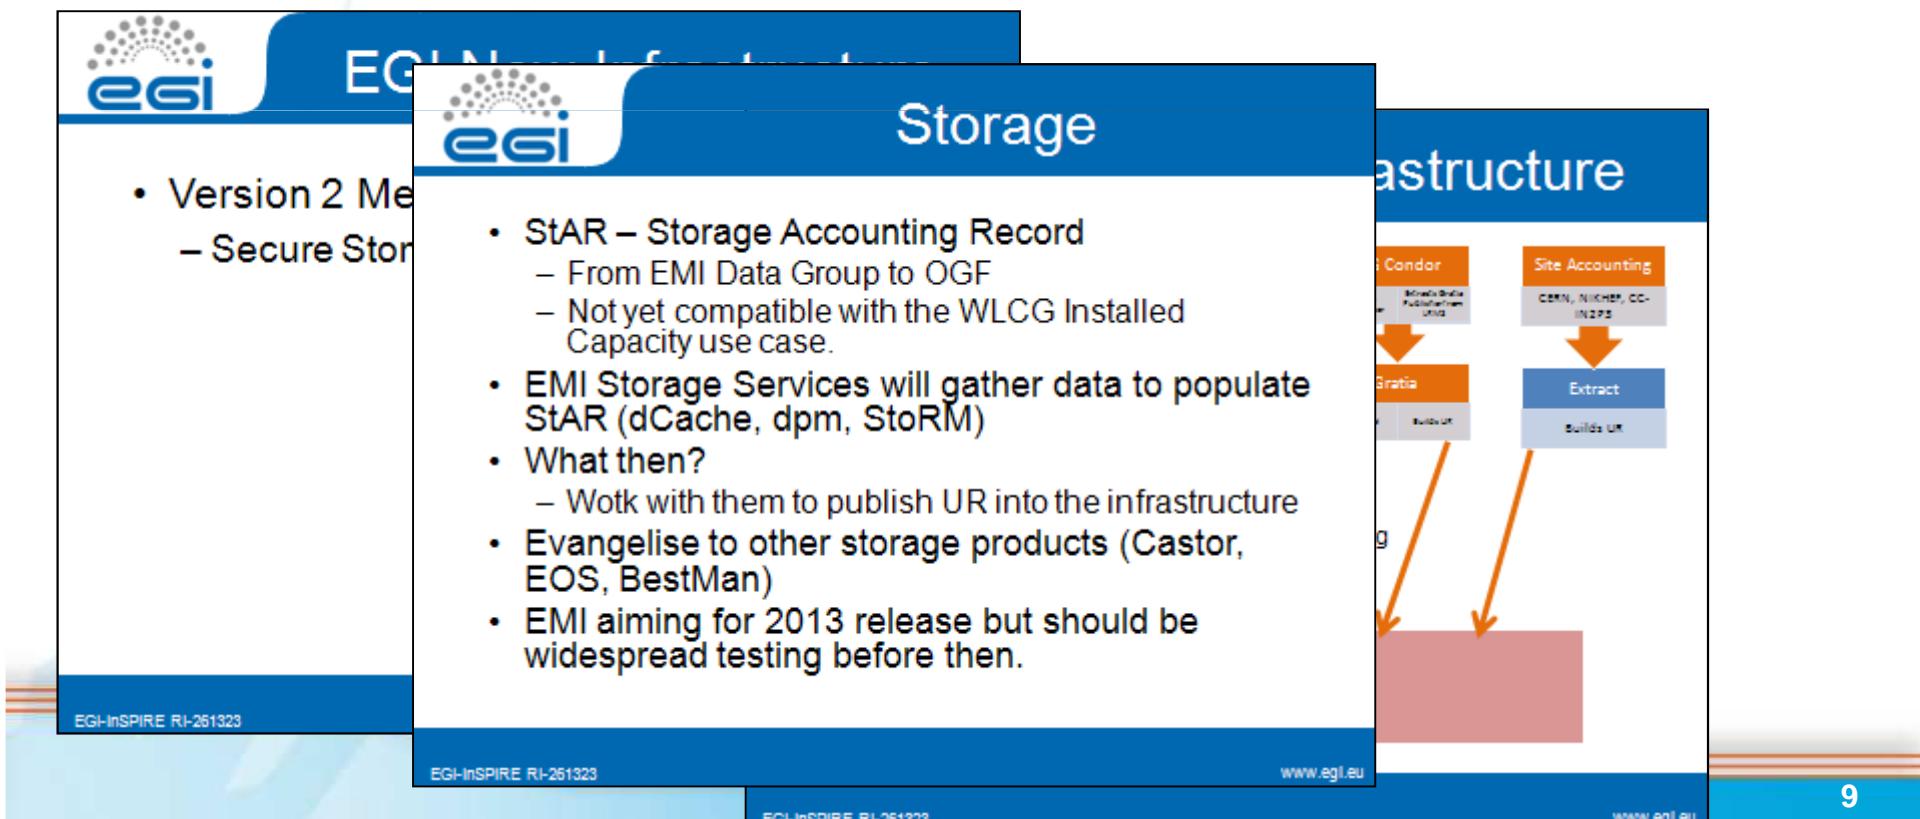


Accounting



■ EGI New Infrastructure

- [http://indico.cern.ch/materialDisplay.py?sessionId=4&materi alId=1&confId=106650](http://indico.cern.ch/materialDisplay.py?sessionId=4&materialId=1&confId=106650)





Middleware/GLite



- Maria Alandes Pradillo (CERN)
 - <http://indico.cern.ch/materialDisplay.py?contribId=4&sessionId=1&materialId=slides&confId=106650>

The slide has a blue header bar with the text 'Grid Technology' on the left, 'End of support for standard updates' in the center, and 'CERN IT Department' on the right.

The main content area contains the following text:

- No more standard updates in gLite 3.1/3.2 since 31.10.2011
 - The following services will only have security updates support until April 2012:
 - gLite 3.2: glite-APEL, glite-ARGUS, glite-BDII_site, glite-BDII_top, glite-CLUSTER, glite-CREAM, glite-GLEXEC_wn, glite-LB, glite-LSF_utils, glite-MPI_utils, glite-SCAS, glite-SE_dcache_info, glite-SE_dcache_nameserver_chimera, glite-SE_dcache_pool, glite-SE_dcache_srm, glite-SGE_utils, glite-TORQUE_client, glite-TORQUE_server, glite-TORQUE_utils, glite-UI, glite-VOBOX, glite-VOMS_mysql, glite-VOMS_oracle, glite-WN.
 - gLite 3.1: glite-CLUSTER, glite-LSF_utils, glite-SGE_utils, glite-TORQUE_utils, glite-WMS and lcg-CE.
 - User support is still provided for all services.

Middleware Update - M. Alandes Pradillo



Middleware/EMI



- Doina Cristina Aiftimiei (INFN)
 - <http://indico.cern.ch/materialDisplay.py?contribId=3&sessionId=1&materialId=slides&confId=106650>

EMI 1 Updates – General Links

- EMI 1 planned products updates
 - http://bit.ly/all_active_tasks
- Near/next Update
 - <http://bit.ly/EMI1UpdatesReady>
- EMI 1 products releases under preparation
 - <http://bit.ly/EMI1UpdatesOpen>

17/11/2011 WLCG 005, 9th November 2011

Next Updates

- Emergency Release:
 - APEL publisher [v. 3.2.8](#)
 - Security Vulnerability – [RT #504](#)
- Major Release:
 - CREAM (S)GE module [y. 1.0.0](#)
 - together with Revision Release of BLAH, v. [1.16.3](#)
 - Both certified
 - Under testing in few volunteer sites (details in next slide)

17/11/2011 WLCG 005, 9th November 2011 4

Status of CREAM SGE Support

- SGE support in CREAM finalised
- Internal certification finished on 07/11
- Currently being validated by a number of WLCG sites to avoid unnecessary delays after official release
 - LIP, Imperial College (confirmed)
 - NGI-Fr and QMUL (available, but TBC)
- As soon as external tests are finished, will be available from the EMI production repository

11/11/11 WLCG 005, 9th November 2011 5

?

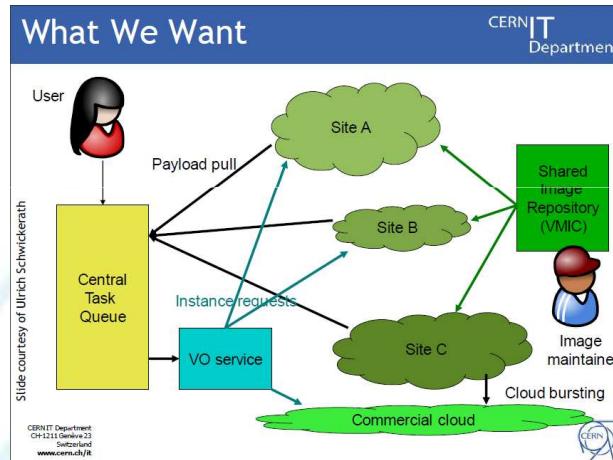


Autres



■ Virtualisation (Tony Cass)

- <http://indico.cern.ch/materialDisplay.py?contribId=2&sessionId=2&materialId=slides&confId=106650>



■ CPU benchmarking (Helge Meinhard)

- <http://indico.cern.ch/getFile.py/access?contribId=1&sessionId=2&resId=2&materialId=slides&confId=106650>

Conclusion:

HS06 no longer very accurate, but probably “good enough”

- *Better carry on with it as defined rather than introducing a new benchmark now or (even worse) re-define existing one*



Technical Evolution Groups



GDB du 09/11/2011



- <http://indico.cern.ch/conferenceDisplay.py?confId=106650>

The screenshot shows a presentation slide with a black header bar containing the word "Today". To the left of the slide is a vertical sidebar featuring the WLCG logo (Worldwide LHC Computing Grid) and the LCG logo (Science & Technology Facilities Council). The main content area lists several topics:

- Technical Evolution Groups
- Accounting
- Middleware
 - Staged Rollout
- HEPiX
 - Benchmarking
 - Virtualisation

In the bottom right corner of the slide, there is a large, stylized "GDB" logo.



WARNING WARNING WARNING...



WLCG
Worldwide LHC Computing Grid



Disk pledges

- Floods in Thailand are seriously affecting disk availability and price
 - Costs already risen by 25-30%
 - http://www.channelregister.co.uk/2011/11/03/disk_drive_prices_may_double/
 - Problem is likely to be long-lasting: 6-9 months?
- Affecting ongoing procurements at CERN and several (all?) Tier 1s, and presumably at Tier 2s
 - Can affect both disk server and CPU server procurements...
- Risk/likelihood:
 - 2012 disk installation may be very(?) late, and potentially reduced wrt pledges
 - Situation still evolving ... but not clear how it will get better very soon
- Experiments need to be prepared for this situation...



LCG CE / CREAM SGE



LCG-CE/CREAM

- Sites no longer required to run an LCG-CE.
- Proposed removing from Availability calculation in October
 - Provided ACE worked OK (it did)
 - But reliable CREAM for SGE not yet rolled out
- MB discussed yesterday and decided to wait longer for the anticipated release of CREAM for SGE.
 - News this afternoon



Science & Technology
Facilities Council





TEG Workload Management



The WM TEG

- Suggested scope (from the TEG WG document v3):
 - Pilots and frameworks
 - CE's – what is needed from them?
 - WMS? Required
 - Clouds
 - Virtualization
 - Security model (MUPJs etc.)
 - Whole node scheduling
 - How to ask for / use GPUs etc.
 - How to make use of parallelism
 - Information services
 - What is still the requirement for non-private IP for the WN's?
- Deliverables:
 - An assessment of the current situation, covering middleware, operation, support
 - A strategy document defining the needs, and a plan to satisfy these needs for the next 2 to 5 years



TEG Workload Management



Organization

- Mailing list, wlcg-teg-workload-mgmt@cern.ch
 - 33 people currently subscribed; representation from experiments, EMI, EGI and about 10 sites; still missing reps from some T1s, though (e.g. IN2P3, NL, RAL)
 - Contact Torre/Davide if interested to join (you can also self-subscribe to the list)
- Indico page,
<https://indico.cern.ch/categoryDisplay.py?categId=3768>
- Wiki page,
<https://twiki.cern.ch/twiki/bin/view/LCG/WorkloadManagementTechnicalEvolution>
- Cross-participation to our meetings from other TEGs is encouraged
 - We will also try to be present at other TEG meetings



The E stands for Evolution

- We should not start by collecting bug reports or feature requests for existing systems
- We should rather step back and look at main concepts behind recent strategy changes/additions on the experimentside?
 - Eg Data placement and Federation
- Which concepts allow to implement Storage Federation?
 - What are the implicit assumptions: eg. namespace congruence, read-only files
 - Upcoming Federation workshop in Lyon will help TEG to collect input
- How and where can new technology come in and which of our “special needs” is preventing this today?
 - Clustered file systems at smaller sites
 - but we need SRM for disk?
 - Cloud storage
 - but we need a global namespace?



The E

- We should not start by changing existing systems
- We should rather step by step introduce strategy changes/additions
 - Eg Data placement and migration
- Which concepts allow us to do this
 - What are the implications?
 - Upcoming Federations
- How and where can new technologies "fill the gaps" in our needs" is preventing to move forward
 - Clustered file system
 - but we need SRM for disk placement
 - Cloud storage
 - but we need a global view

DM Model and Evolution

- A few key questions to guide input collection
 - Is the archive<->disk split an agreed strategy?
 - In other words: can we drop HSM?
 - If yes, several simplifications can be done and benefits can be evaluated. Eg Why SRM for disk-only pools?
 - Do we see "naked" clustered file systems or cloud storage coming in during the reviewed time scale?
 - Can we take our abstract DM model and fill its conceptual boxes with technology options for a give class of sites?
 - Are the experiment data services prepared for this?
 - What role do we see for Federation vs Data Push?
 - Eg Federation is complementing Push at small sites ($T>1$) but Push is still required to maintain control
 - Eg Federation/Push is far better and will be the strategic target



TEG Operations



How we divided the work

- **WG1: Monitoring and Metrics**
 - Editors: Simone Campana and Pepe Flix
 - Contributors: Costin Grigoras , Andrea Sciaba, Alessandro Di Girolamo; (sites) Ian Collier , Xavier Espinal, Vera Hansper , Alexandre Lossent
WG1 Sub-Wiki: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTegOperationsWG1>
- **WG2: Support tools + Underlying Services+ WLCG Operations**
 - Editors: Andrea Sciaba, Maria Dimou (support tools), Lionel Cons (underlying services) and Stefan Roiser (WLCG Operations)
 - Contributors: Simone Campana, Andrea Sciaba, Alessandro Di Girolamo, Joel Closier, Pepe Flix, John Gordon, Alison Packer; (sites) Alexandre Lossent, Xavier Espinal
WG2 Sub-Wiki: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTegOperationsWG2>
- **WG3: Application Software Management**
 - Editor: Stefan Roiser
 - Contributions: Marco Cattaneo, Steve Traylen; (sites) Ian Collier, Alexandre Lossent, Alessandra Forti
- **WG4: Operational Requirements on Middleware**
 - Editors: Maarten Litmaath, Tiziana Ferrari
 - Contributors: Maria Dimou, Laurence Field; (sites) Alexandre Lossent, Jeff Templon, Vera Hansper, Anthony Tiradani, Paolo Veronesi
- **WG5: Middleware Configuration, Deployment and Distribution**
 - Editors: Oliver Keeble, Rob Quick
 - Contributors: Cristina Aifitimi, Simone Campana, ; (sites) Ian Collier, Alexandre Lossent, Pablo Fernandez, Felix Lee
- **Horizontal Representation from OSG (R.Quick), EGI (T. Ferrari), and EMI (C. Aifitimi, L. Field)**





TEG Operations



How we divided the work Where we are headed

- “Provocative Question”: how can we reduce the manpower needed for “operations” to zero?
- Starting point: ask three related questions
 - What works well?
 - What are the three main problems?
 - Where could we save the most manpower?
- These questions were asked to each TEG member for each sub-area.



TEG Operations



How we divided the work with the TEG

Issues (from MB)

- The membership of the group – e.g. missing stakeholders? Site representation is thin, only 5 of 10 listed on the wiki have responded.
- Other issues? It is a huge amount of work, we are thankful for the help of the subgroup editors, we will do our best to deliver within the stated time frame.

Input still welcome, but hurry ... Nov 28th is less than three weeks away!!!!



Participation du CC



- RAPPEL: description des TEGs sur le wiki LCG France
 - <http://lcg.in2p3.fr/wiki/index.php/TEG>
- TEG Storage
 - Yvan Calas
 - Loïc Tortay
 - Fabio Hernandez
- TEG Data Management
 - Jean-Yves ?
 - Qq'un HSM (HPSS) ?
- TEG Workload Management
 - Qq'un du groupe système pour les aspects cloud/whole nodes ?
- TEG Operations
 - Qq'un des opérations du CC ?
- TEG Databases
 - Un DBA du CC ?



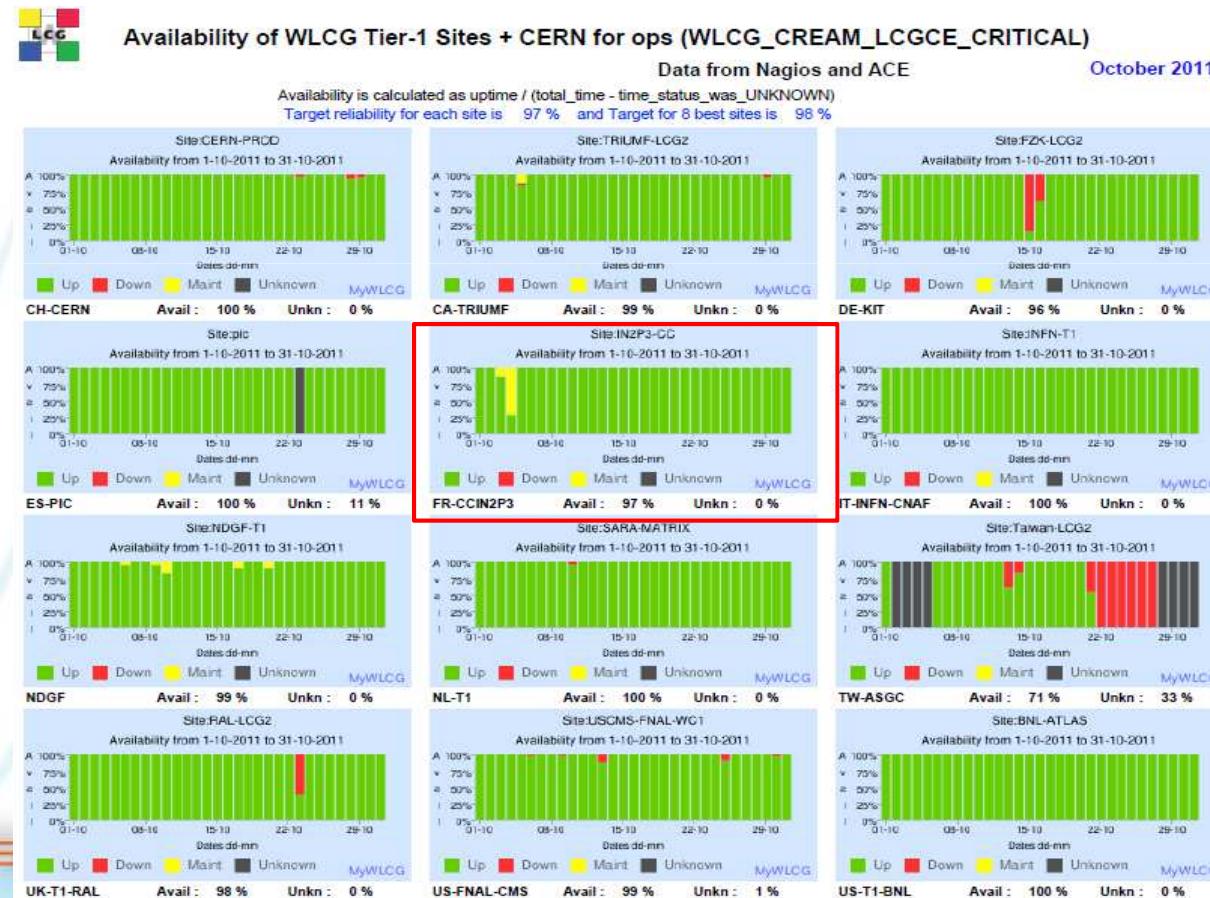
Résultats du site



OPS/Octobre 2011



- http://gvdev.cern.ch/GIRDVIEW/downloads/Reports/201110/wlcg/Tier1_Summary_Oct2011.pdf



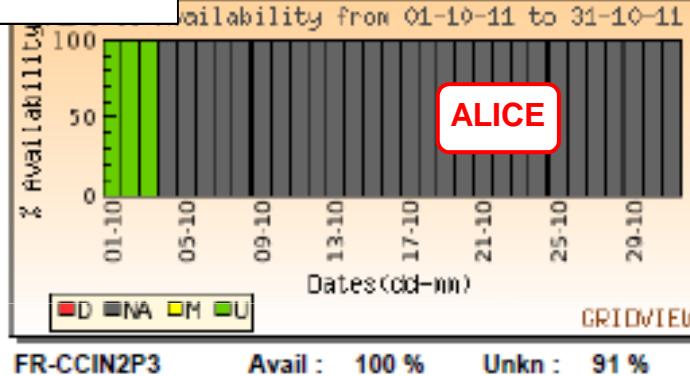


Disponibilités des VOs pour septembre 2011

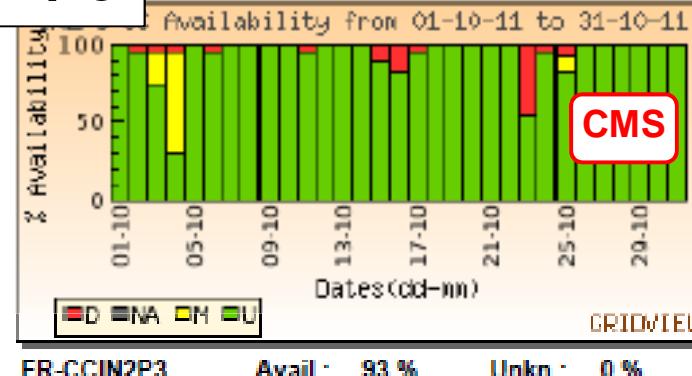


Source: <http://gvdev.cern.ch/GRIDVIEW/downloads/Reports/>

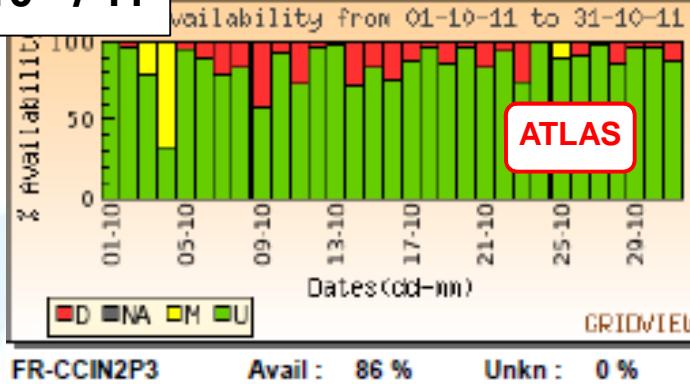
1st / 7



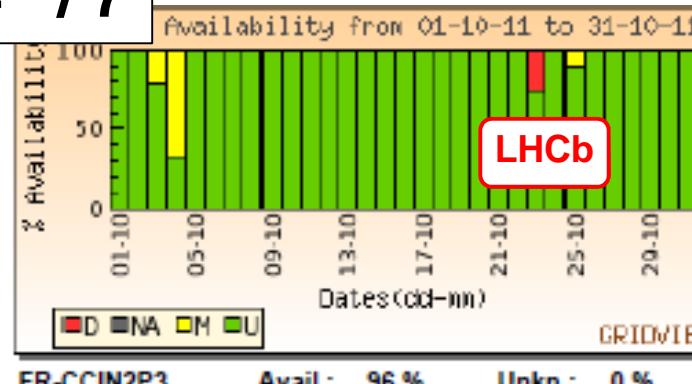
7th / 8



10th / 11



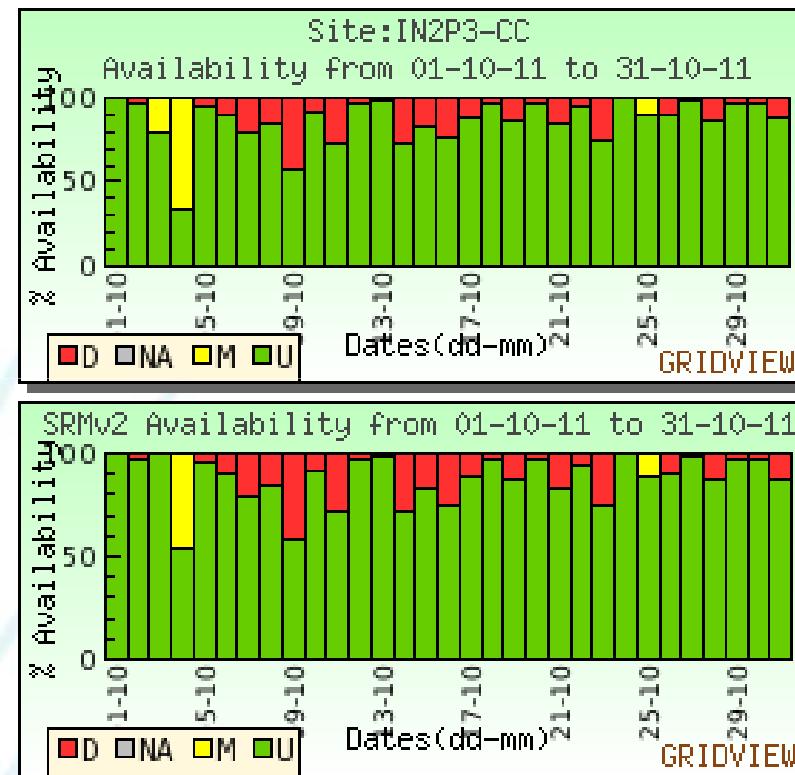
4th / 7



OK!!



Détails de la disponibilité d'Atlas



“Get timeout”
probablement à cause
des transferts lents. Pb
de réseau en cours
d’investigations.



Dossiers en cours



Chantiers en cours



- Arrêt de BQS le 6/12
- CEs
 - Decommissionning des CEs BQS
 - Plus de LCG CEs pour les VOs LHC
 - Plus de CREAM BQS pour le T2
 - **Arrêt du dernier CREAM BQS du T1 le vendredi 25/11/2011**
 - Les pledges CPUs seront assurés sur GE à cette date.
 - 2 CREAM GE par site pour les 4 VOs LHC
 - Récente mise à jour avec le dernier “patch” officialisé
 - Pb du statut des jobs résolu ? (à vérifier)
 - Proposition de donner l'accès au support à l'historique des jobs des CREAM
 - Fonctionnalité du CREAM
 - Christelle a fourni un exemple à Ghita et Emmanouil
 - Feedback ????

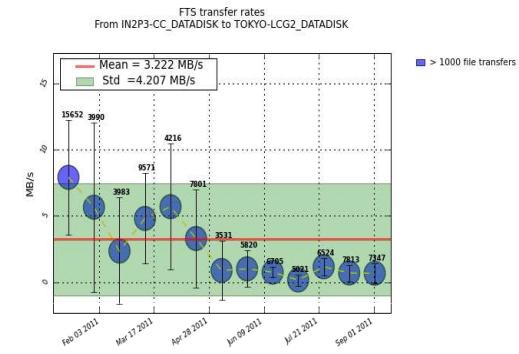


Chantiers en cours



■ Problèmes réseaux

- Avec les sites étrangers qui ne sont pas sur LHCOPN
- Impact sur ATLAS et CMS
 - https://gqus.eu/ws/ticket_info.php?ticket=75983
 - https://gqus.eu/ws/ticket_info.php?ticket=75919
 - https://gqus.eu/ws/ticket_info.php?ticket=75397
 - https://gqus.eu/ws/ticket_info.php?ticket=74268
 - https://gqus.eu/ws/ticket_info.php?ticket=71864
- Procédure WLCG nous désigne comme responsable du suivi de l'incident
 - Telecom suit l'affaire
 - Plusieurs problèmes détectés avec RENATER





Chantiers en cours



- “Whole node” job
 - CMS et maintenant ATLAS
 - Mise en place de workers dédiés en production
 - 5 workers 8 coeurs normaux
 - 5 workers 8 coeurs virtuels
 - 1 CREAM de test
 - cccreamceli01.in2p3.fr
 - A finaliser
 - Wiki à mettre en place ?



Autres sujets



■ Partial Downtime Management

- Ex: downtime sur dCache, impact sur les jobs.
- https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTier1ServiceCoordinationMinutes111103#Partial_Downtime_Management

■ Prochaine T1/AF

- Avenir de la LAF
- Point sur les Pledges 2011



Evénements



■ En cours/A venir

- SC11, Seattle, 14-17 November
- Workshop "Creating Federated Data Stores For the LHC", CCIN2P3, 21-22 novembre
- LCG France, CCIN2P3, 30/11 et 01/12
- EGI Information Services Workshop 1 December
- LHCONE Meeting, Amsterdam, 1-2 December
- ISGC, Taipei, 26 February - 2 March 2012
- EGI User Forum, Munich, 19-23 March
- WLCG Workshop, NYC, 19-20th May 2012
- CHEP2012, NYC, 21-25th May 2012