



LCG-France Tier-1 & AF

Réunion mensuelle de coordination

Fabio Hernandez
fabio@in2p3.fr

Lyon, 25 mars 2010



l r f u
cea
saclay



Calendrier LHC

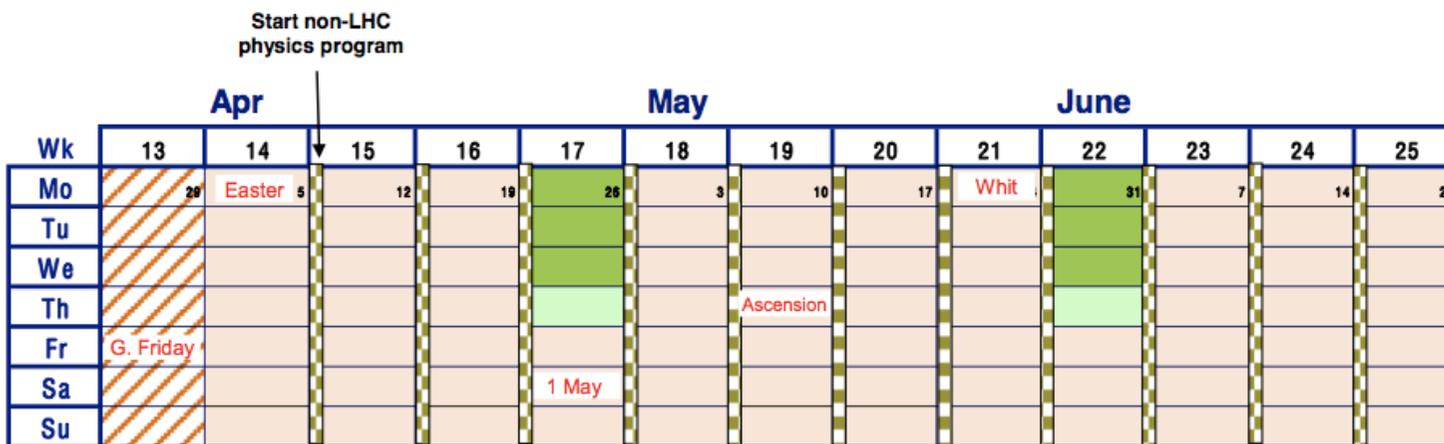
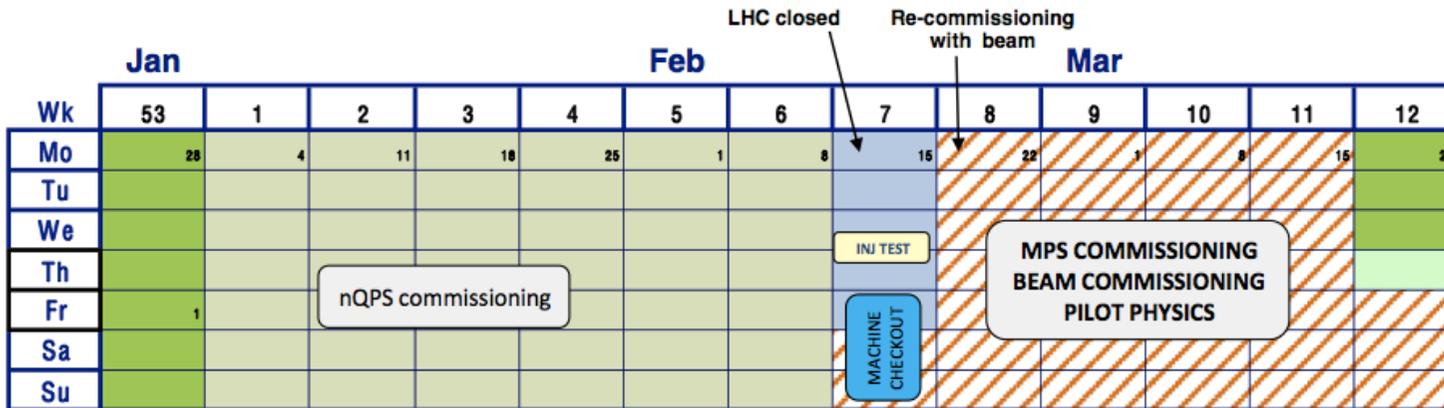
Calendrier LHC: 2010



ML

2010 LHC Schedule

5/2/2010
V1.4



- Technical Stop
- Recommissioning with beam
- SPS et al - physics

Source: http://lhc-commissioning.web.cern.ch/lhc-commissioning/planning/2010/2010-LHC-schedule_v1.4.pdf

Calendrier LHC: 2010 (suite)



	July				Aug				Sep				
Wk	26	27	28	29	30	31	32	33	34	35	36	37	38
Mo	28	5	12	19	26	2	9	16	23	30	6	13	20
Tu													
We													
Th											Jeune G		
Fr													
Sa													
Su													

Ion Beam to SPS

	Oct			Nov				Dec					
Wk	39	40	41	42	43	44	45	46	47	48	49	50	51
Mo	27	4	11	18	25	1	8	15	22	29	6	13	20
Tu													
We													
Th													
Fr													
Sa													
Su													

ION Beam Setup

Start Ion Physics

IONS (approx 4 weeks)

Xmas Day

- Technical Stop
- Recommissioning with beam
- SPS et al - physics

▶ Collisions



- Mardi 30 mars, collisions attendues
 - 2 faisceaux à 3,5 TeV
- Événement presse organisé à cette occasion

▶ Signatures des publications LHC

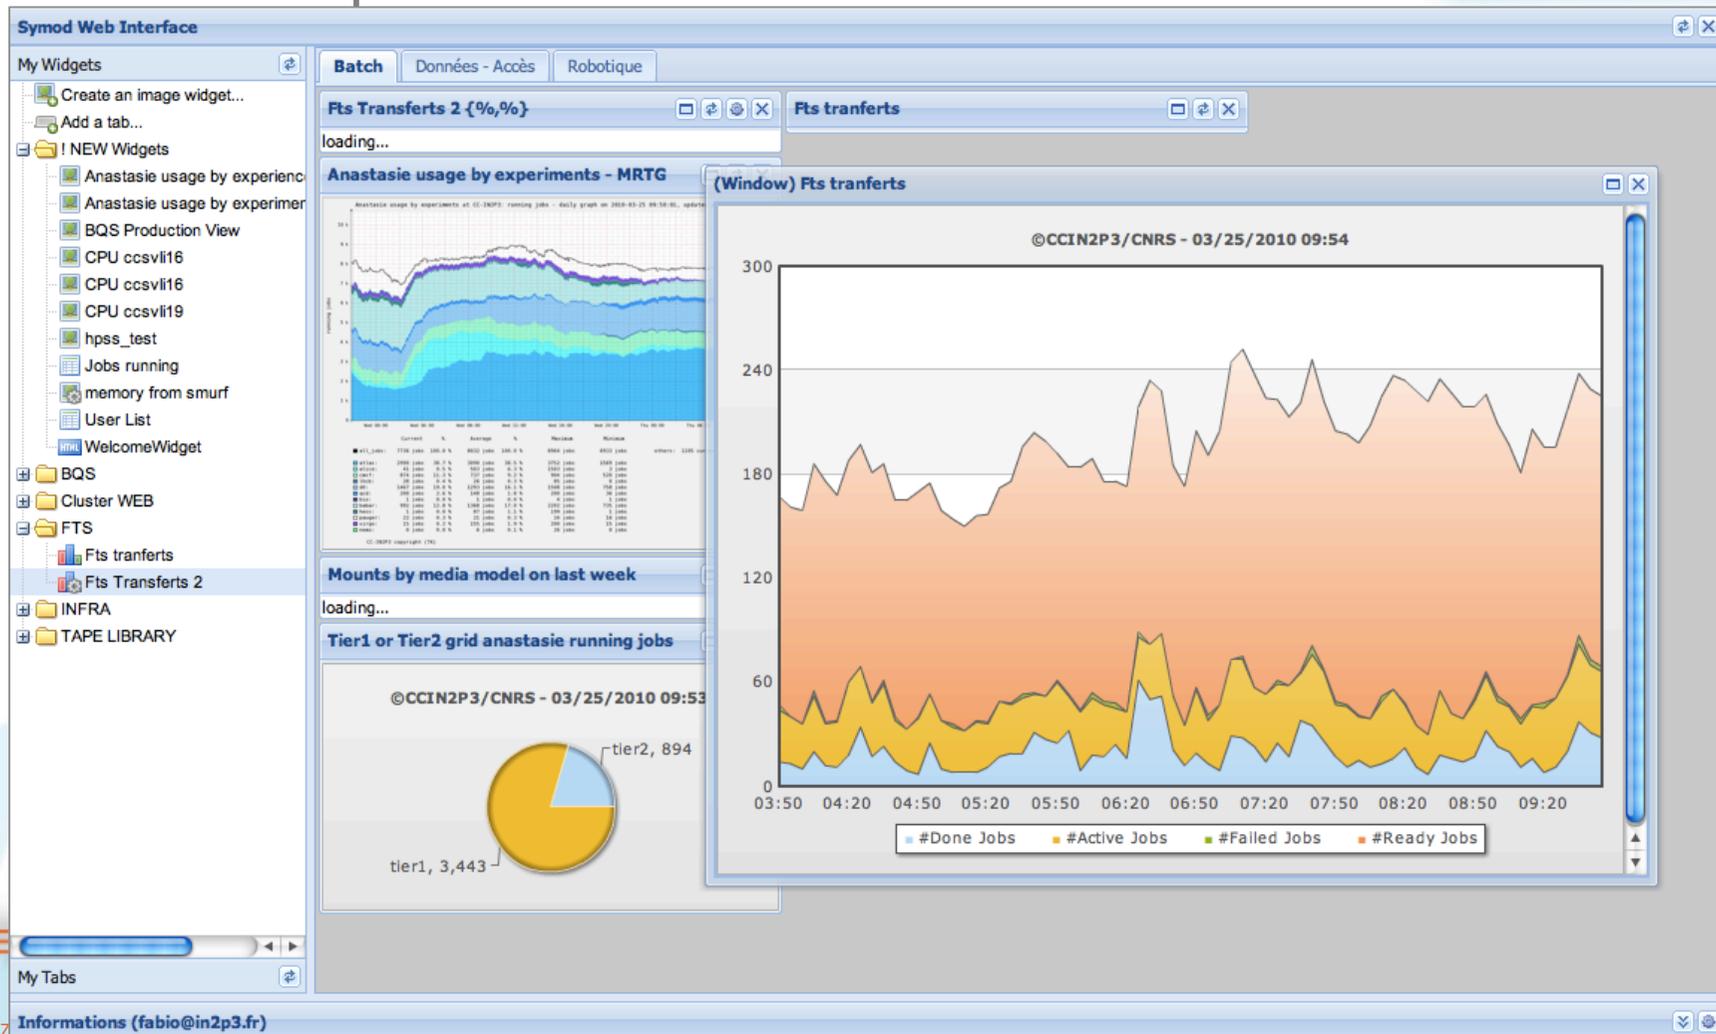


- Depuis le 19/03/2010 le CC-IN2P3 est membre de la collaboration CMS
 - Vote au collaboration board le 19/03/2010 à 94% des voix
- Désormais, CC-IN2P3 formellement membre des 4 expériences LHC
 - Signature des publications avec l'affiliation CC-IN2P3
 - Important pour l'attractivité du site pour les post-docs et chercheurs physiciens



Avancement des chantiers

- Entrée en production le 4 mars



► Système d'information du site



- Collecte par Gstat des informations publiées par le site
 - Modifications faites dans la publication du volet stockage et CPU
 - Des détails encore à finaliser

Name ▲	Status ▼	CPUs		Online Storage Space (GB)		Nearline Storage Space (GB)		Grid Jobs		
		Physical ▼	Logical ▼	TotalSize ▼	UsedSize ▼	TotalSize ▼	UsedSize ▼	Total ▼	Running ▼	Waiting ▼
CGG-LCG2	WARNING	140	140	4,787	3%	0	0%	0	0%	0%
GRIF	CRITICAL	0	0	1,463,260	70%	0	0%	0	0%	0%
IN2P3-CC	N/A	1,143	4,572	3,387,944	62%	20,987,084	39%	6,750	64%	56%
IN2P3-CC-T2	CRITICAL	1,023	4,092	233,472	71%	0	0%	2,024	13%	73%
IN2P3-CPPM	CRITICAL	158	484	51,313	30%	0	0%	629	29%	76%
IN2P3-IPNL	CRITICAL	158	616	57,010	47%	0	0%	752	0%	100%
IN2P3-IRES	CRITICAL	256	1,216	311,013	77%	0	0%	774	365%	99%
IN2P3-LAPP	OK	220	752	182,318	38%	0	0%	820	26%	75%
IN2P3-LPC	CRITICAL	380	1,320	141,641	54%	0	0%	832	20%	66%
IN2P3-LPSC	CRITICAL	156	516	190,064	41%	0	0%	241	6%	87%
IN2P3-SUBATECH	WARNING	115	460	45,569	92%	0	0%	58	0%	100%

Source: Gstat 2.0, 24/03/2010
<http://gstat-prod.cern.ch/gstat/summary/Country/France/>

▶ Interruption des services



- Arrêt programmé le 02/03/2010
 - Changement matériel des serveurs centraux de BQS
 - Intervention sur la base de données HPSS
 - Intervention SPS
 - Reconfiguration ACSLS pour intégration du 4^{ème} silo
 - Détails:
<https://cctools.in2p3.fr/operations/wiki/doku.php?id=arrets:arret02032010>

- Mise à jour de FTS: passage à v2.2.3
 - Réalisée le 18/03
 - Détails: <http://cctools2.in2p3.fr/elog/operations-FTS/56>

- 17/02: Modification de la configuration de la zone disk-only du tier-2 CMS pour introduire des quotas
 - Objectif: éviter que les jobs d'analyse puissent déborder de l'allocation prévue, ce qui impacterait l'activité merge du tier-1 partageant la même zone
 - Space token: CMST2
 - Suite à la réunion CMS du 15/02/2010:
<http://indico.in2p3.fr/conferenceDisplay.py?confId=2805>

► Configuration ferme batch



- Augmentation de la limite mémoire maximum pour la classe G et T
 - Passage à 3GB lors de l'arrêt du 02/03/2010

▶ VO Boxes



- Installation de cclcgatlas03 en SL5
 - Est-elle en production?

▶ ATLAS: FroNTTier/Squid



- Service FroNTTier/Squid pour ATLAS en production
 - Accès optimisé à des données dans Oracle au tier-1, en particulier pour les tier-2s distants
 - Tests initiaux réalisés
 - Tests approfondis pour quantifier l'amélioration vont suivre

▶ ATLAS: dCache



- Renommage des groupes des pools
 - pgroup-atlas-**import-disk**
 - pgroup-atlas-**aod-read**
 - pgroup-atlas-**permanent-disk**
 - pgroup-atlas-**read**
 - pgroup-atlas-**localexport-tape**
 - pgroup-atlas-**xferout**
 - pgroup-atlas-**export-tape**
 - pgroup-atlas-**tapebuffer**
 - pgroup-atlas-**import-tape**

▶ ATLAS: Reprocessing



- Campagne de reprocessing de données

▶ Reprocessing campaign 2/2010

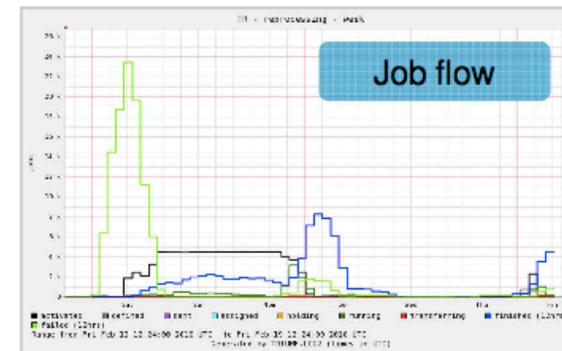


■ 2009 data reprocessing

- 2/2/2010 @ 22:00: Beginning of reprocessing
- 2/2/2010 Change of reprocessing strategy: all jobs sent on the « very-long » queue to avoid memory overflows: (4 GB memory instead of 2.4-2.8 GB for “long”)

■ Direct Consequences

- Very slow job flow
- Cannot provide the expected quota of simultaneous running jobs
- Long lasting reprocessing: more than 6 days (2 in 12/2009)

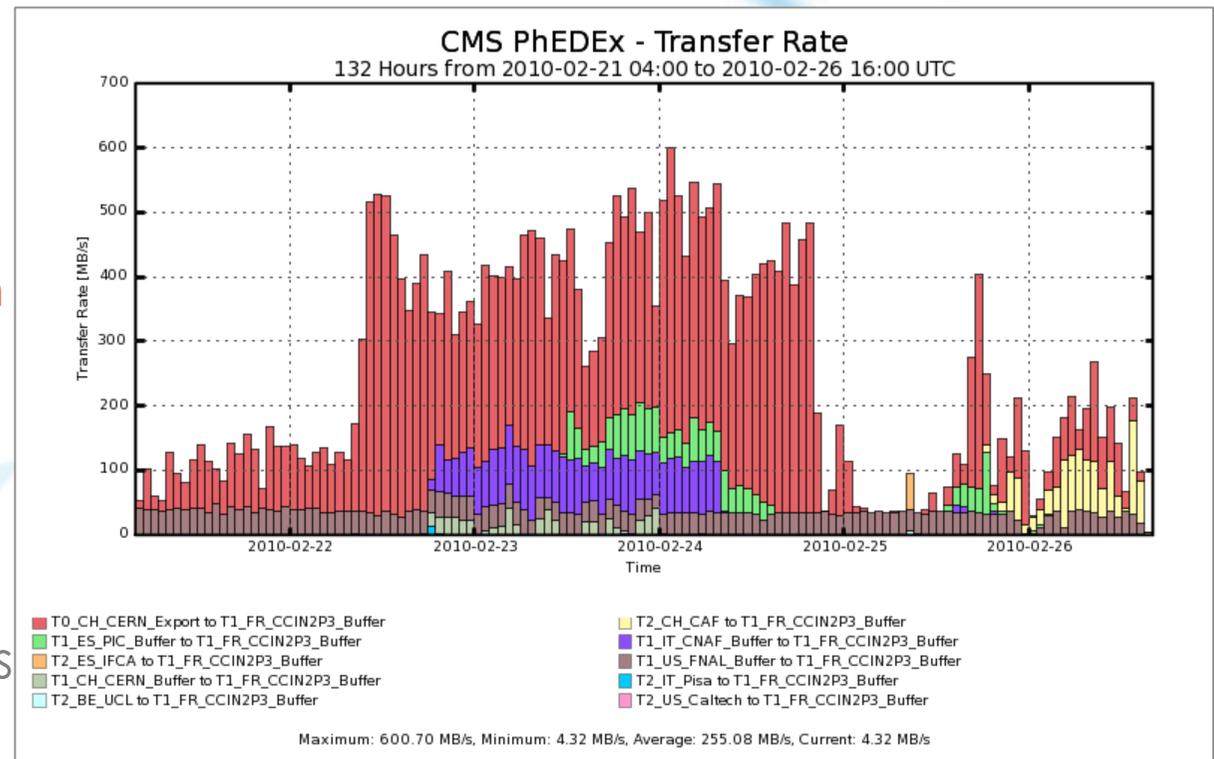


Source: G. Rahal
ATLAS February 2010 reprocessing post-mortem, 18/03/2010
<http://indico.cern.ch/conferenceDisplay.py?confId=87335>

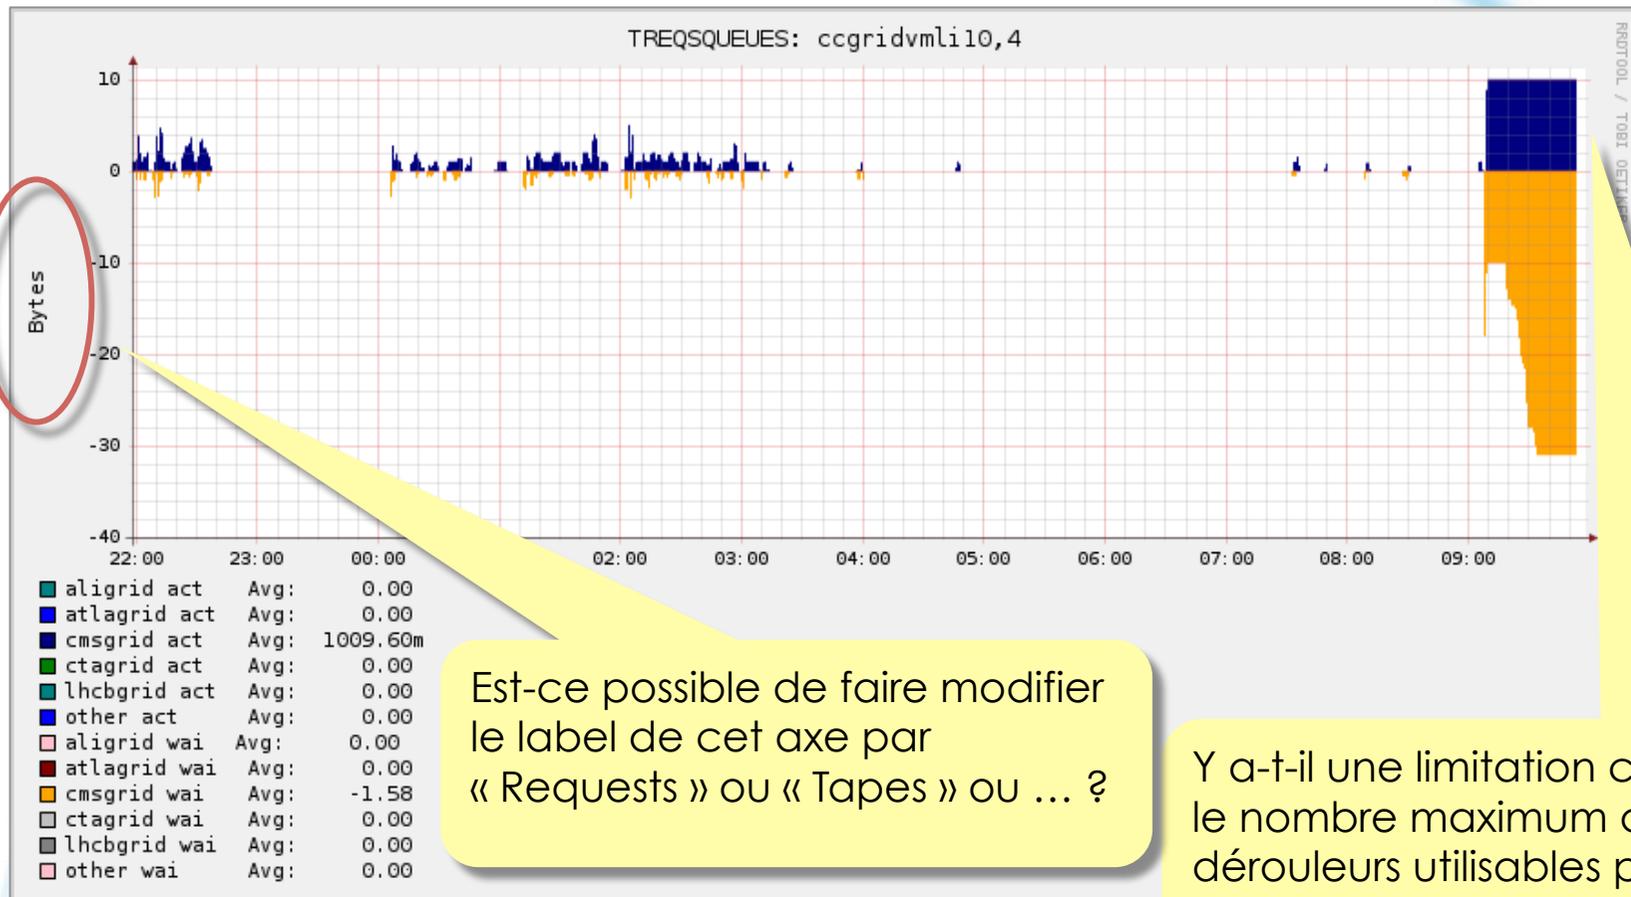
► CMS: Importation de données



- Forte charge en écriture observée sur HPSS avec des données CMS
 - Campagne d'importation de données pour repeupler le site, suite à l'incident de nov. 2009
- Document sur les débits nominaux d'entrée/sortie du MSS par les expériences LHC en cours de préparation
 - Des fluctuations importantes sont inhérentes à l'activité des expériences



TReqS



Est-ce possible de faire modifier le label de cet axe par « Requests » ou « Tapes » ou ... ?

Y a-t-il une limitation concernant le nombre maximum de dérouleurs utilisables pour le staging de fichiers par TReqS?



Incidents

▶ Incidents



- 22/02 (samedi): Problème d'alimentation électrique d'une armoire où sont localisées plusieurs machines de service (site et Top BDII)
 - Détails: <http://cctools2.in2p3.fr/elog/astreinte/566>
- 25/02: Problème d'accès au LFC ATLAS
 - Serveur non configuré pour accepter des proxies émis par lcg-voms.cern.ch
 - Détails: https://gus.fzk.de/ws/ticket_info.php?ticket=55931
- 05/02, 10/02, 15/02 Instabilité SRM: échecs des tests SAM pour des certificats CERN
 - Détails: <http://cctools2.in2p3.fr/elog/operations-dCache/167>

► CMS: tickets



- Les tickets Savannah concernant les sites CMS tier-1 et tier-2 au CC-IN2P3 sont consultables ici:

- Tier-1:
<http://tinyurl.com/ccin2p3-cms-t1>
- Tier-2:
<http://tinyurl.com/ccin2p3-cms-t2>

LCG Savannah

Login Status:
Not Logged In
Login
New User

This Page
Clean Reload
Printer Version
Related Recipes:
Why log in?

Search
in Projects
Search

Hosted Projects
Register New Project
Full List
Contributors Wanted
Statistics

Site Help
User Docs: Cookbook
User Docs: In Depth Guide
Get Support
Contact Us

Links
GNU/Savane

CMS Computing Infrastructure Support - Support: Browse Items

Group: [Main](#) [Docs](#) [Support](#) [Source Code](#) [Tasks](#) [News](#)

(-) Display Criteria

Browse with the **Basic** query form and **Simple** selection. Apply

Category: **Any** Assigned to: **cmscompinfrasup-t1fccin2p3** Open/Closed: **Any** Status: **Any**

Additional constraint **deactivated**:
Any modified since **24** **March** **2010**

Show: **50** items at once ; items with a spam score lower than **5**.

Click a column heading to sort results (up or down), or sort by priority or reset sort. You can also activate multicolumn sort. Currently, results are sorted by Item ID.

Begin Previous Results 158 matching items - Items 1 to 50 Next Results End

Item ID	Summary	Submitted		Assigned To	Submitted By
		On			
#113472	Tape families for Spring10 acquisition era	2010-03-24 12:02		cmscompinfrasup-t1fccin2p3	vdutta
#113373	Custodial LFN's for Summer09 MC requests	2010-03-18 10:55		cmscompinfrasup-t1fccin2p3	vdutta
#113335	Use FTS2.2.3 in PhEDEx T0 FileDownload agent	2010-03-17 11:13		cmscompinfrasup-t1fccin2p3	magini
#113309	Consistency check at T1_FR_CCIN2P3	2010-03-15 17:47		cmscompinfrasup-t1fccin2p3	magini
#113251	Tape families for MinimumBias 2010 data	2010-03-12		cmscompinfrasup-	ceballos



GDB & MB

Jobs pilote multi-utilisateurs



Proposed course of action (1/2)

- Sites should configure glxec with SCAS/Argus/GUMS backend and advertise it with a CE capability:
 - GlueCECapability: glxec
 - Capability was favored over a run-time environment tag
 - No details about the mode of operation or supported VOs
- Experiments should try out their glxec workflows at such sites as they become available
 - First as a background activity, like SAM tests
- When an experiment finds a particular site reliable with glxec, ideally all its MUPJ for that site would start using it there

Maarten Litmaath (CERN)

16

Proposed course of action (2/2)

- Some sites might start with log-only and switch to setuid later
 - CMS and ATLAS representatives have argued that the log-only mode should be disallowed right away
- Some of the experiments would like to require the setuid mode in the medium term
 - Sites that do not support it would no longer receive MUPJ
- Other technologies may affect medium-term strategies
 - When each job is launched in a fresh VM, the risks are much reduced and more setups may be acceptable for MUPJ
- The WLCG policies on MUPJ are expected to become stricter as more experience has been gained by sites and experiments

Maarten Litmaath (CERN)

17

Source: M. Litmaath, GDB, 24/03/2010
<http://indico.cern.ch/conferenceTimeTable.py?confId=84636>

▶ CREAM



Enabling Grids for E-science

Status



- **CREAM CE v. 1.5 now in production**
 - In gLite 3.1 (sl4_i386)
 - In gLite 3.2 (sl5_x86_64)
- **CREAM CE v. 1.6 for gLite 3.2 / sl5_x86_64 certified 2 weeks ago**
 - Certified according to the new certification model
 - While trying it at UCSD last week (in the context of the OSG CREAM evaluation) bug #64516 in trustmanager was found
 - Issue not in CREAM, but affecting (also) CREAM
 - It affects certificates signed by the DoE CA
 - → Needed to rebuild against a fixed version of trustmanager, when ready

CREAM (suite)



Enabling Grids for E-science

CREAM CE 1.6 (cont.ed)

- **Support of SGE in BLAH (as contributed by CESGA/LIP)**
- **Limiter to protect CREAM when the machine is overloaded**
 - New job submissions are blocked when this happens
 - Taking into account load, memory usage, # of file descriptors, etc.
 - Very similar to the limiter used in the WMS
- **Proxy purger**
 - To clean from the delegationDB and from the file system expired delegations
- **Support of file transfers from/to gridftp servers started with user credentials**
 - Asked by Condor
- **Improved CREAM startup**
 - It could take a while if needed to get the status of jobs in a non-terminal status

INFSO-RI-508833

GDB - Amsterdam, March 24, 2010

5



Enabling Grids for E-science

CREAM CE 1.6 (cont.ed)

- **Improved proxy renewals when there are multiple jobs sharing the same delegationid**
 - Typical use case considering at least the submissions from WMS and from Condor
- **Improved performance of some DB queries**
- **Several bug fixes**
 - User 'tomcat' not added anymore to VO groups and glxexec group
 - Glxexec's lcms conf file fixed
 - It could happen that you get mapped to a user different than the one mapped by gridftp
 - GLITE_WMS_RB_BROKERINFO properly set in CREAM jobwrapper
 - ...
- **Configuration changes already communicated to M. Jouvin for Quattor QWG templates**

INFSO-RI-508833

GDB - Amsterdam, March 24, 2010

6

Source: M. Sgaravato, GDB, 24/03/2010
<http://indico.cern.ch/conferenceTimeTable.py?confId=84636>



Chantiers en cours ou à venir

▶ Stockage



- Protection de données des expériences contre les effacement involontaires
- Limiter les possibilités de staging aux seuls utilisateurs autorisés
- Activation de la possibilité de staging pour les storage elements xrootd de ALICE

▶ Interventions Programmées



- Re-configuration robotique et HPSS
 - 13 avril 2010
 - Détails: <https://cctools2.in2p3.fr/operations/wiki/doku.php?id=arrets:arret13042010>



Disponibilité du site

Disponibilité tier-1

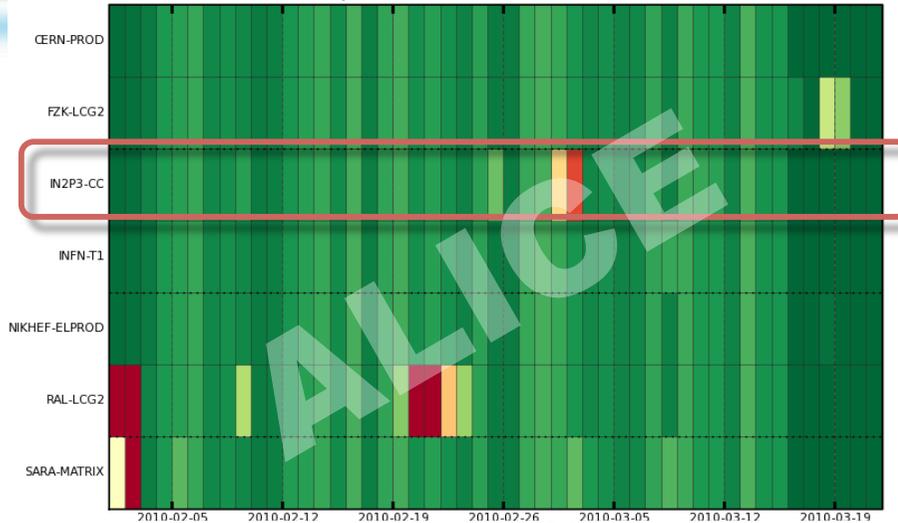
Période: Fév 1 – Mar 22 2010

<http://dashboard.cern.ch/>

Source: LHC experiments dashboard

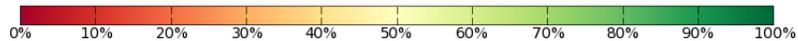
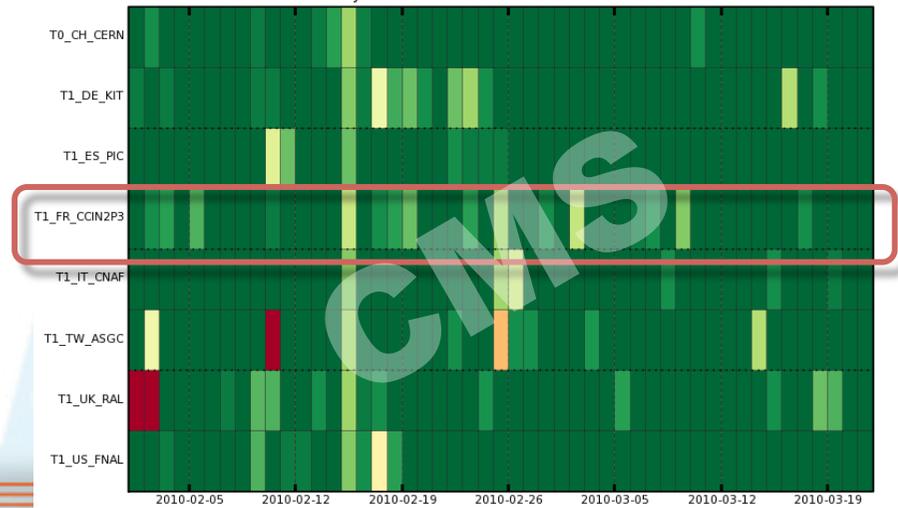
Site Availability using WLCG Availability (FCR critical)

49 Days from Week 05 of 2010 to Week 12 of 2010



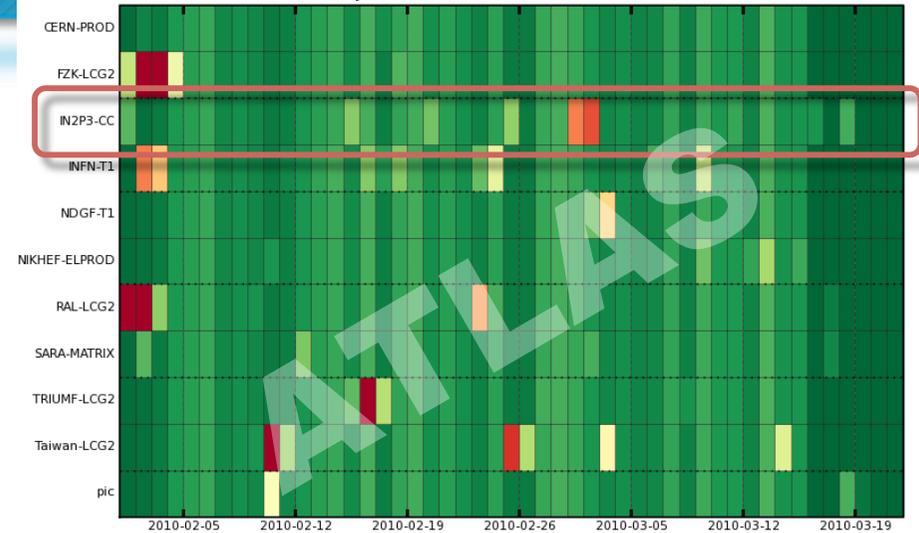
Site Availability

49 Days from Week 05 of 2010 to Week 12 of 2010



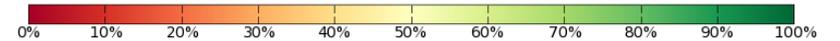
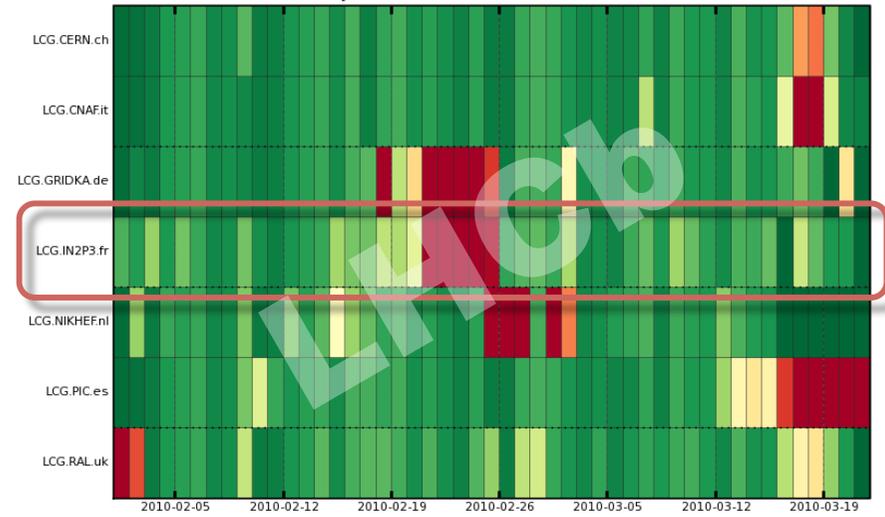
Site Availability using WLCG Availability (FCR critical)

49 Days from Week 05 of 2010 to Week 12 of 2010



Site Availability using LHCb Critical Availability

49 Days from Week 05 of 2010 to Week 12 of 2010



CMS: site readiness status



Tier-1

T1_FR_CCIN2P3																							
Site Readiness Status: W R R R R R R R R R R R R R R W R																							
Daily Metric: O O O O O O E O O O O O O O O O O O E O																							
Maintenance:	Up	Up																					
Job Robot:	100%	100%	100%	100%	100%	100%	100%	n/a*	100%	100%	100%	100%	100%	100%	100%	100%	99%	100%	100%	99%	n/a		
SAM Availability:	100%	100%	100%	96%	92%	100%	76%	100%	100%	100%	100%	100%	100%	100%	92%	100%	100%	100%	100%	84%	100%		
Good T1 links from T0:	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1	1/1		
Good T1 links from T1s:	7/7	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	7/7	6/6	6/6	4/6	6/6	4/6	6/6	6/6		
Good T1 links from T2s:	33/33	32/32	32/32	31/31	31/32	31/32	31/32	32/32	31/32	31/32	31/32	31/32	31/32	32/32	32/32	19/32	32/32	19/31	31/31	32/32	32/32		
Good T1 links to T1s:	6/7	7/7	7/7	7/7	6/7	6/7	6/7	6/7	7/7	7/7	5/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7		
Good T1 links to T2s:	42/47	42/47	41/47	40/47	40/47	39/47	41/47	36/45	44/48	44/47	43/47	43/46	34/47	35/47	42/47	41/47	43/47	42/47	44/47	44/47	44/47		
Active T1 links from T0:	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Active T1 links from/to T1s:	7(d)- 7(u)																						
Active T1 links to T2s:	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45		
	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
	Mar																						

Report made on 2010-03-24 02:30:01 (UTC)

Tier-2

T2_FR_CCIN2P3																							
Site Readiness Status: R R R R R R R R R R R R R R R R																							
Daily Metric: O																							
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	
Job Robot:	100%	100%	99%	100%	99%	100%	100%	n/a*	100%	99%	99%	100%	99%	100%	100%	100%	99%	100%	99%	n/a			
SAM Availability:	100%	100%	100%	100%	88%	100%	84%	100%	100%	100%	100%	100%	100%	100%	100%	92%	100%	100%	100%	92%	100%		
Good T2 links from T1s:	8/8	8/8	7/8	8/8	8/8	8/8	7/8	7/8	7/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8		
Good T2 links to T1s:	8/8	8/8	8/8	7/8	7/8	7/8	7/8	7/8	8/8	8/8	6/8	7/8	7/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8		
Active T2 links from T1s:	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8		
Active T2 links to T1s:	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8		
	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
	Mar																						

Report made on 2010-03-24 02:30:01 (UTC)



Petites Annonces

▶ Événements à venir



- Comité de Pilotage LCG-France
 - Paris, 2 avril 2010
- Hepix Spring 2010
 - Lisbonne, 19-23 avril 2010
 - <http://www.lip.pt/hepixspring2010/index.php>
- Réunion des sites LCG-France
 - Marseille, 24-25 juin 2010
 - <http://indico.in2p3.fr/conferenceDisplay.py?confId=3760>
- WLCG Collaboration Workshop
 - Londres, Juillet 7-9 2010
 - <http://indico.cern.ch/conferenceDisplay.py?confId=82919>
- CHEP 2010
 - Taiwan, 18-22 octobre 2010
 - <http://event.twgrid.org/chep2010/index.html>

▶ Aujourd'hui et à venir



- Aujourd'hui
 - TReqS nouvelle génération
 - Surveillance via NAGIOS
 - CREAM: bilan d'utilisation à 3 mois
- Réunions premier semestre 2010
 - ✓ 14 janvier
 - ✓ 11 février
 - ✓ 25 mars
 - 15 avril
 - 27 mai
 - 17 juin
 - 22 juillet
 - Agendas: <http://indico.in2p3.fr/categoryDisplay.py?categId=102>

▶ Questions/Commentaires

