



15/12/2011

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



- Nouvelles de LCG
 - CR du dernier GDB
 - Technical Evolution Groups
- Disponibilité du site
- Dossiers en cours
- Evénements

Nouvelles de LCG



CR du dernier GDB



■ GDB du 14 décembre 2011

- <http://indico.cern.ch/conferenceDisplay.py?confId=106651>

CR du dernier GDB



Today

- Technical Evolution Groups
- Experiment Operations
- Middleware
- Identity Tracing



VLCG
Worldwide LHC Computing Grid



Science & Technology
Facilities Council

GDB

CR du dernier GDB



January

- Short Meeting
- Technical Evolution Groups?
- OSG
- CREAM SGE
- Site Dashboard?



Science & Technology
Facilities Council

GDB
9

EMI - Beta & Acceptance Testing



- Early availability – once the code is ready:
 - <https://twiki.cern.ch/twiki/bin/view/EMI/EMIBetaAcceptanceTesting>
 - access to interested parties to earlier stages/phases of the software development:
 - *beta phase* - through participation in **EMI Beta Testing** activity, using the testing/ repository
 - *release candidate phase*- through participation in **EMI Acceptance Testing** activity, using the deployment/ repository

EMI INF50-RH-261611

WLCG Technical Evolution Groups

TEG Storage Management

SM/DM-TEG topics

Topics in the **SM-TEG** (*as from the original mandate*):

- ◆ [SM.1] Experiment IO/usage patterns (and so performance requirements for storage)
- ◆ [SM.2] Requirements on future storage systems, and how storage will evolve
- ◆ [SM.3] Separation of archives and disk pools/caches
- ◆ [SM.4] Storage system interfaces to Grid (SRM status?)
- ◆ [SM.5] Filesystems/protocols (standards?)
- ◆ [SM.6] Security/access controls (TBD: to be discussed with the Security TEG)
- ◆ [SM.7] (as site-run services:) storage management interfaces, performance measurements, monitoring, etc

◆ Topics in the **DM-TEG** (*as from the original* mandate*):

- ◆ [DM.1] Review of DM demonstrators from Summer 2010
- ◆ [DM.2] Dataset management and Data placement (policy-based or dynamic)
- ◆ [DM.3] Data federation strategies
- ◆ [DM.4] Transfers and WAN access protocols (HTTP, xrootd, gsiftp)
- ◆ [DM.5] Data transfer management (FTS)
- ◆ [DM.6] Understanding data accessibility and security requirements/needs
- ◆ [DM.7] POOL
- ◆ [DM.8] ROOT, PROOF frameworks
- ◆ [DM.9] Management of namespaces
- ◆ [DM.10] Management of catalogues (LFC, future directions)

Rapprochement entre Storage et Data Management.
Ces 2 TEGs ont beaucoup de sujets communs.

(*) roughly, since DM already presented and re-elaborated the list

TEG Storage Management

Affectation des responsabilités (TEG leader) pour les sujets communs

Reshuffling of {SM,DM}-TEG topics

DM

DM.2

DM topic, mainly

DM.7

CERN-IT / ATLAS discussion

DM.10

Pure DM topic

DM.1

Some demonstrators have high SM implications

DM.3

DM should take the lead though, and SM steps in later. SM encouraged DM to clarify things early here (Lyon helped)

SM.4

Good to overlap

DM.4

DM can take the lead though. SM steps in on implication on the site storage.

DM.5

FTS dev should be contacted and invited to an ad-hoc discussion

SM.6 / DM.6

To be discussed with Security TEG

DM.8

Overlap on storage implications

DM.9

Good to overlap

SM

SM.1

Crucial input from exps. Some overlap with DM.8

SM.2

Pure SM topic

SM.3

If agreed that the exps' data management systems need it, then the topic has to be followed by SM

SM.5

SM topic, mainly

SM.7

Pure SM topic

What's done and what's next

- ◆ digestion and approval of the mandate
- ◆ definition of membership, guarantee a careful balance
- ◆ organizational bits (twiki, Indico, etc), plus collection of personal profiles of the members
- ◆ definition of (brief) set of general questions on SM personal viewpoints (see "list 1")
- ◆ collection and digestion, before the kick-off meeting
- ◆ kick-off meeting
- ◆ definition of boundaries with DM-TEG, definition of work plans with them
- ◆ Lyon workshop on data federation
- ◆ definition of a unique list of joint DM/SM questions to LHC experiments (see "list 2")
- ◆ information gathering phase from **ALICE**, ad-hoc meeting
- ◆ information gathering phase from **CMS**, ad-hoc meeting
- ◆ information gathering phase from **LHCb**, ad-hoc meeting
- ◆ information gathering phase from **ATLAS**, ad-hoc meeting
- ◆ digestion of the experiments' feedback, and synthesis
- ◆ on each sub-topic, the "responsible" TEG elaborates a way to proceed
- ◆ potentially large (to be parallelized) homework for different subsets of members
- ◆ from each subset: feedback collection, digestion, synthesis
- ◆ work on coherence, first steps towards a draft report, including risk management analysis
- ◆ aim to have a statement on "where we are" by/at the Amsterdam TEGs week
- ◆ post-Amsterdam wrap-up: work on "the future" over Feb-Mar
- ◆ ...

We are here





Risks

Threat	Likelihood	Impact	Risk
Misused identities	5	5	25
Compromised identities	5	4	20
Attack propagation between WLCG sites	4	5	20
Exploitation of a serious OS vulnerability	4	4	16
Threats originating from trust services	3	4	12
Negative publicity on a non-event	2	4	8
Insecure configuration leading to undesirable access	3	2	6
Insufficient protection of information leading to sensitive data leakage	3	2	6
Incidents on resources not bound by WLCG policies	1	5	5
Exploitation of a serious VO/middleware software vulnerability	2	2	4
Data removal/corruption/alteration	1	3	3
DoS from an external organisation	1	1	1



Publicity & Traceability

- Publicity
 - Aggravating factor for all previously mentioned risks
 - Media interest needs to be handled with care (and training!)
 - Negative impact is not necessarily linked with severity of incident
 - The accuracy of the information being reported often not relevant
- Traceability
 - Fine-grained traceability is necessary and an essential component of WLCG security.
 - It is only possible to achieve a reasonable level of traceability by implementing a form of Unix identity switching
 - The AAI on the WN subtask will review different technical possibilities and implementations, in terms of security, performance impact, development and maintenance costs.
 - The result of this work will be presented during GDBs.

Authentication & Authorization Infrastructure (AAI) on the WN



Traceability (2)

- User payloads should be separated on the WN
 - Avoid interference → allow for quick identification of involved DN
 - Otherwise need to rely on time-based circumstantial evidence
 - May need to ban multiple users
 - Need to be able to exclude Trojan horses or time bombs
 - Data ownership?
- Possible technologies
 - Glexec → currently needs a proxy
 - Investigate if that could be relaxed for the time being
 - ALICE plugin will check payload signature instead
 - Sudo → how to determine the target account?
 - Virtual machines → when will most sites be ready for that?
 - One account per job slot → only Condor supports it
 - SELinux → not evident



Longer term

- Use of general-purpose proxies on the WN is questionable
 - Cf. AliEn plans
- Relation between payload and user?
 - Payload may have been tampered with
 - Signed payloads would be verifiable
 - Cf. AliEn plans
- Data ownership, restrictions?
 - Cf. AliEn token model

The issue of full delegation vs. accountability/traceability

Why do we need the user proxy for a pilot payload ?

- ▶ Because we don't want to fully trust the VO ?
- ▶ But we already rely on the VO to provide the correct proxy!
- ▶ How much does the presence of a proxy actually prove?
- ▶ At the same time the presence of a proxy is a risk...

We rather need something better.

A proposal: Explicit/Definite Delegation

Signed Grid jobs (in ALICE we are working on that)

This should be...

- ▶ a lot easier
 - ▶ in development / maintenance (one concept from users to sites)
 - ▶ in operation (no callbacks or add-ons / externals)
- ▶ explicit, so there is proof in the end what was submitted
- ▶ without any other permission delegated/passed on
- ▶ gLExec could easily support this (generally green light)

Introduction

- Five areas of work identified
 - WG1: Monitoring
 - WG2: Support tools, Underlying Services and WLCG Operations
 - WG3: Application Software Management
 - WG4: Operational Requirements on Middleware
 - WG5: Middleware Configuration, Deployment and Distribution

TEG Operations and tools



Priorities from the MB

- Discussions at the last MB have asked *"Operations TEG to provide a clear strategy on Availability Reporting which takes account of the needs of all the stakeholders (at least experiments, sites and funding agencies) and reduces the duplication which currently exists"*
- *"The Experiments agreed to continue their contributions in the Operations TEG who will provide an update at the GDB meeting on 14 Dec 2011"*

TEG Operations and tools



Priorities from the MB

- Discussions at the last MB have asked *"Operations TEG to provide a clear strategy on Availability Reporting which takes account of the needs of all the stakeholders (at least experiments, sites and funding agencies) and reduces the duplication which currently exists"*
- *"The Experiments agreed to continue their contributions in the Operations TEG who will provide an update at the GDB meeting on 14 Dec 2011"*

TEG Operations and tools



Sitview - Gridmap Site Status Board

dashb-ssb.cern.ch/dashboard/request.py/siteview?site=T1_FR_CCIN2P3#fullscreen=false

Help Login Site Status for the CMS sites, v0.14.0_rc5

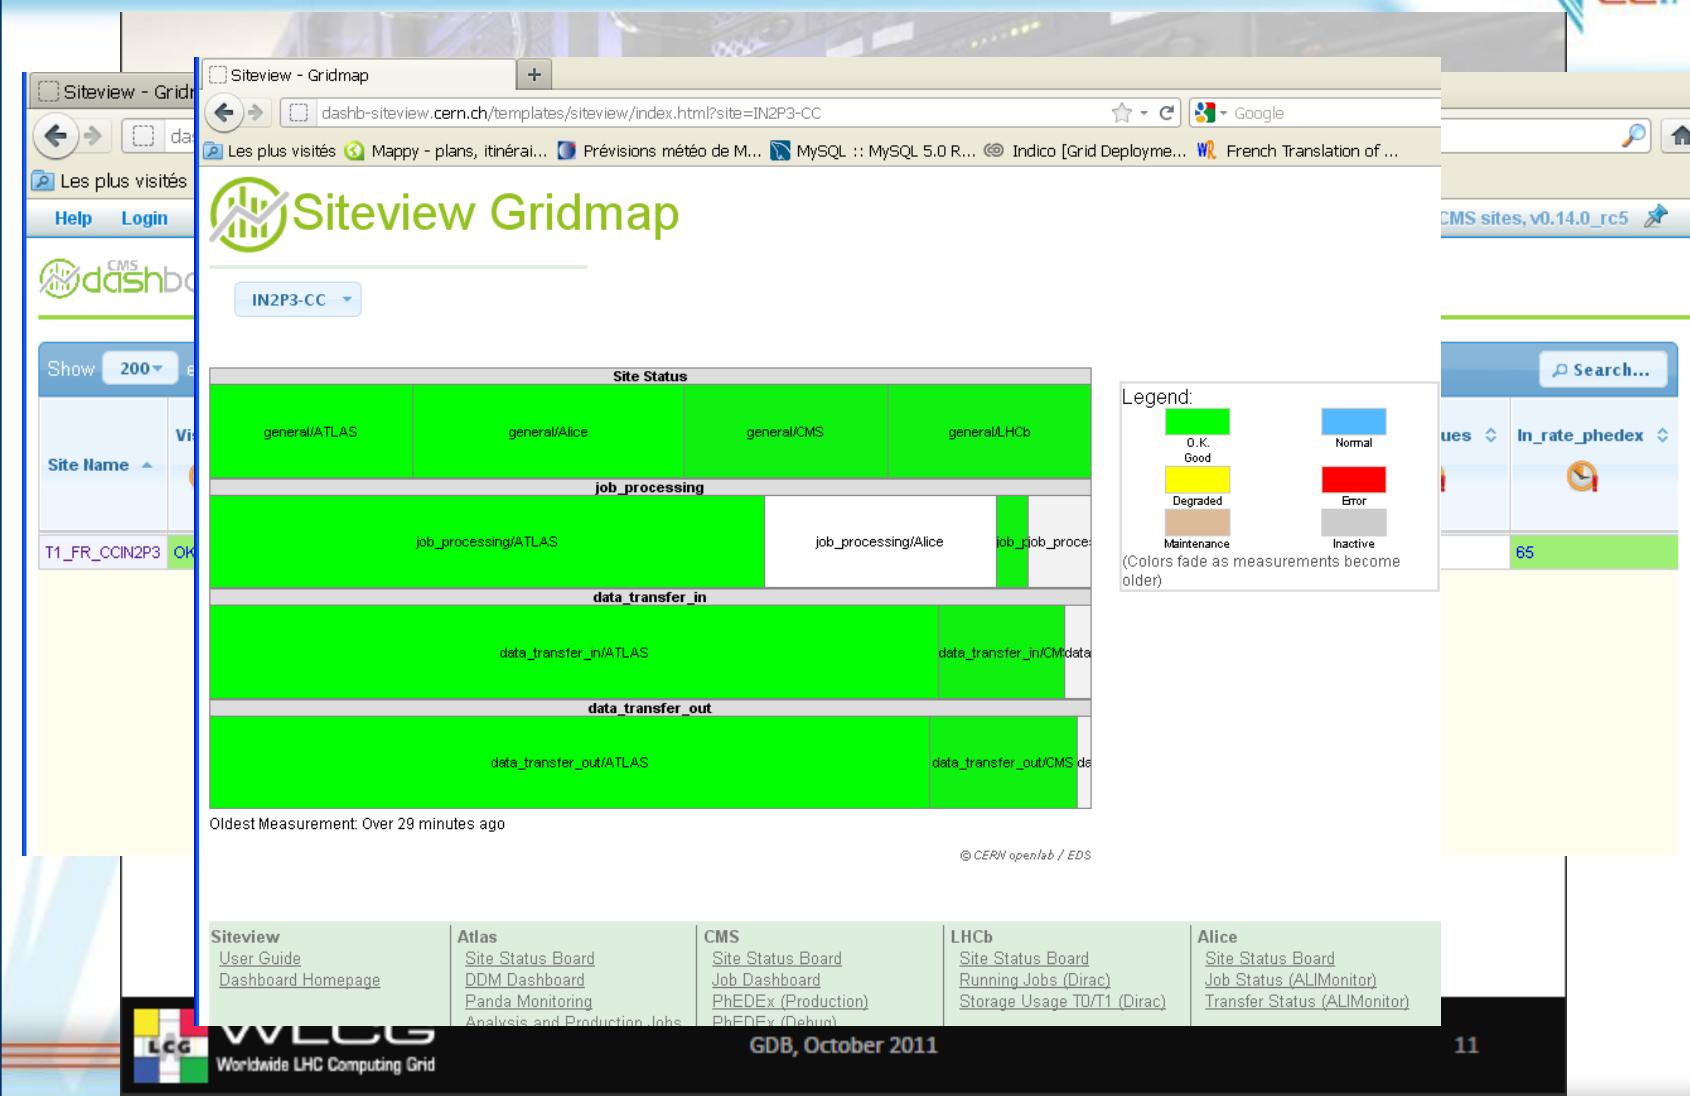
CMS dashboard Index Expanded Table

Site Name	Visible		JobRobot		SAM TESTS		Production		Analysis		Site usage		Commissioned Links (expand this column)	Under investigation	SiteIssues	In_rate_phedex
	CE	SRM	CE	SRM	CE	SRM	Running	Pending	CE	SRM	Running	Pending				
	OK	100%(200)	OK	OK	100%(335)	100%(3)	110	130	3/5 combined	mark	info	65				
T1_FR_CCIN2P3	OK	100%(200)	OK	OK	100%(335)	100%(3)	110	130	3/5 combined	mark	info	65				

ex.html

production but would need to be validated

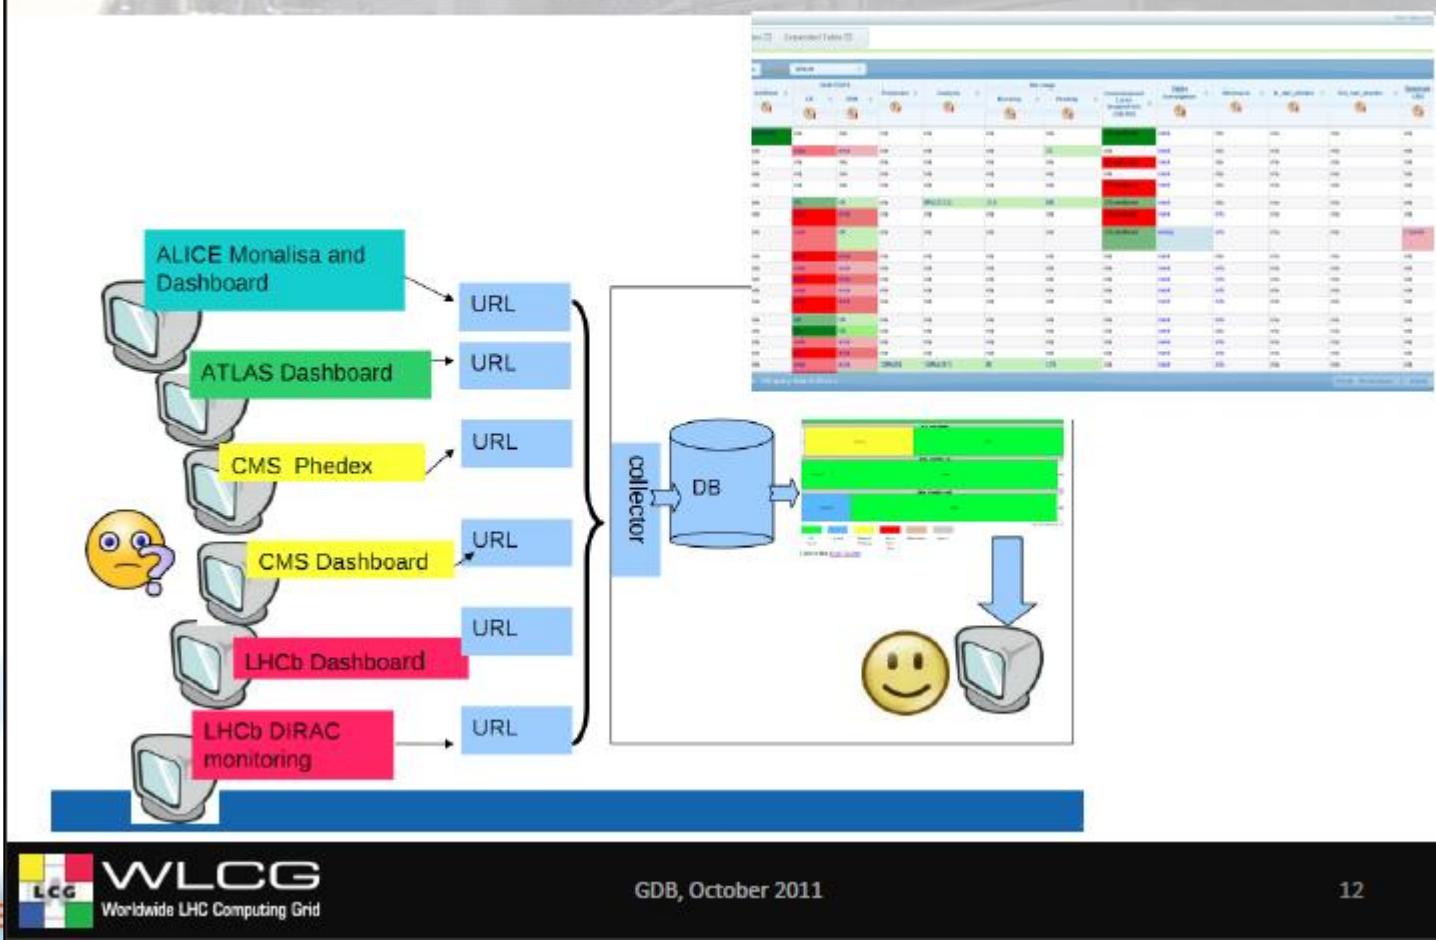
TEG Operations and tools



TEG Operations and tools



SSB and SiteView



Site Monitoring: proposal

- We miss the equivalent of the today's SSB experiment views tailored for sites
- **Proposal** to use the SSB framework to provide this functionality as well
 - Advantages: many metrics already in the SSB for ATLAS and CMS
 - No duplication of effort nor issues of consistency
 - Need to agree on a few common metrics between experiments
 - Relevant from a site perspective
 - Some development needed in SSB to facilitate the visualization
 - Some commitment needed from experiment and sites
 - Experiment (support): define and inject metrics, validation
 - Sites: validation, feedback

Participation du CC



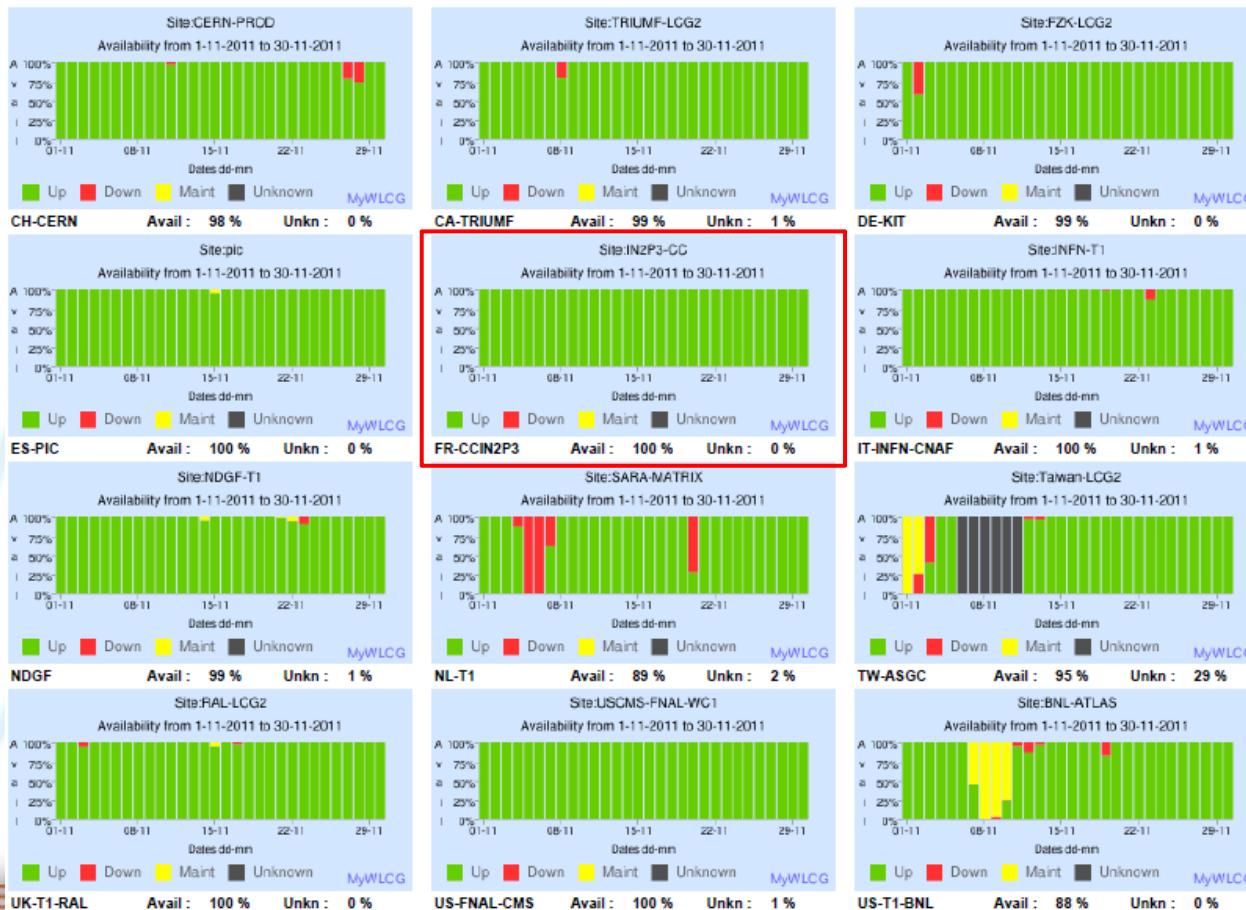
- RAPPEL: description des TEGs sur le wiki LCG France
 - <http://lcg.in2p3.fr/wiki/index.php/TEG>
- TEG Storage
 - Yvan Calas
 - Loïc Tortay
 - Fabio Hernandez
- TEG Data Management
 - Jean-Yves ?
 - Qq'un HSM (HPSS) ?
- TEG Workload Management
 - Pierre Girard (Cloud/Virtualisation)
 - Qq'un des opérations du CC ?
- TEG Databases
 - Osman Aidel

Résultats du site

OPS/Novembre 2011



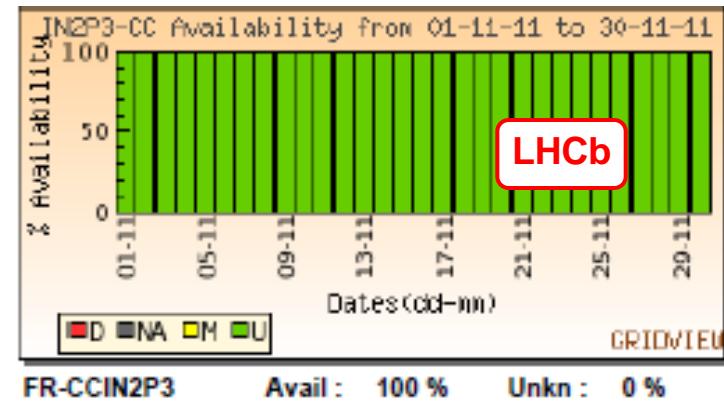
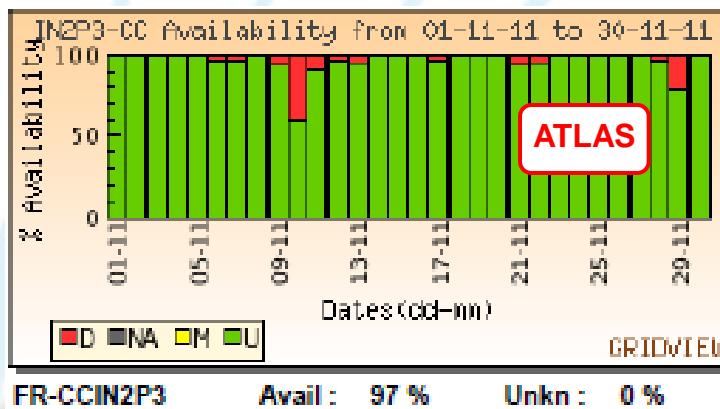
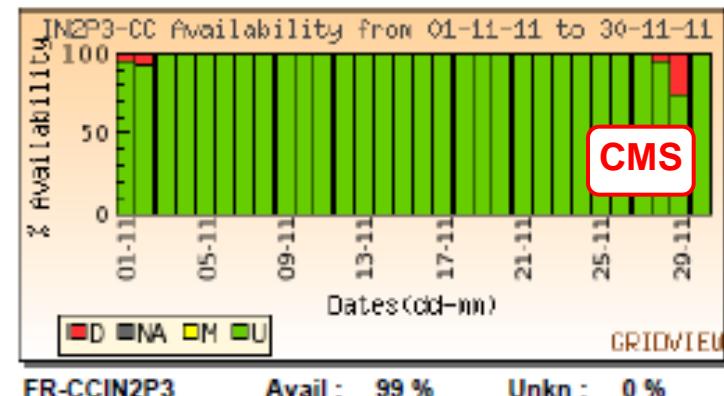
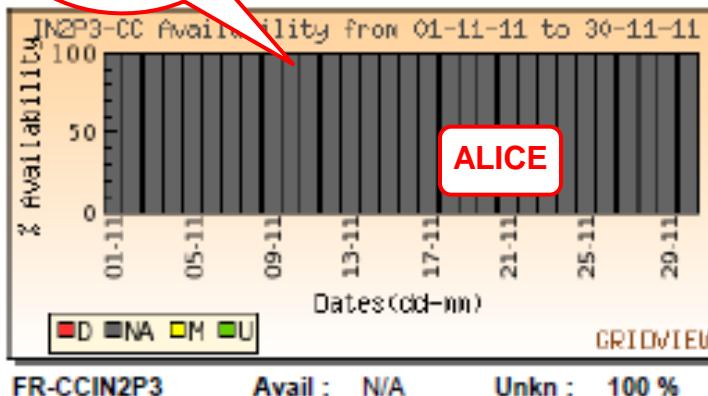
■ http://gvdev.cern.ch/GIRDVIEW/downloads/Reports/201111/wlcg/Tier1_Summary_Nov2011.pdf



Disponibilités des VOs pour novembre 2011

Fin des LCG-
CEs et pas
de SE

Les CREAM ne sont
pas pris en compte



Dossiers en cours

Chantiers en cours



■ CREAM GEs

- Quelques problèmes identifiés
 - Purge des jobs, table MySQL avec 4 000 000 d'entrées, passage des job requirements à GE
- Statut des jobs
 - Tout le support doit y avoir accès
 - Christelle a fait passer la demande des DNs à inclure
 - Interrogation des CEs depuis une UI
 - Glite-ce-job-list <CE>
 - Glite-ce-job-status <job-id>
 - Glite-ce-job-status –all ...

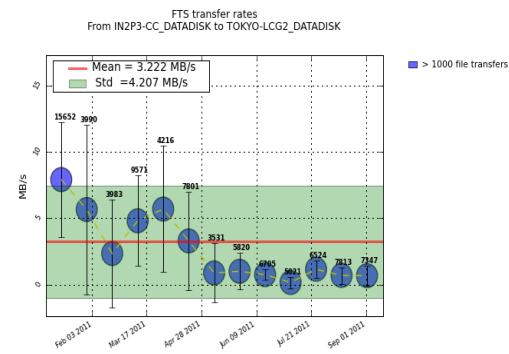
```
girardpi@ccage019(0)# glite-ce-job-status -a -L 0 -s DONE-OK --from '2011-12-15 11:25:00' --endpoint  
cccreamceli10.in2p3.fr  
2011-12-15 11:44:43,169 WARN - No configuration file suitable for loading. Using built-in configuration  
  
***** JobID=[https://cccreamceli10.in2p3.fr:8443/CREAM014725926]  
Status      =[DONE-OK]  
ExitCode    =[0]  
  
***** JobID=[https://ccreamceli10.in2p3.fr:8443/CREAM535953290]  
Status      =[DONE-OK]  
ExitCode    =[0]  
  
***** JobID=[https://ccreamceli10.in2p3.fr:8443/CREAM617789867]  
Status      =[DONE-OK]  
ExitCode    =[0]  
  
***** JobID=[https://ccreamceli10.in2p3.fr:8443/CREAM894582743]  
Status      =[DONE-OK]  
ExitCode    =[0]
```

Chantiers en cours



■ Problèmes réseaux

- Avec les sites étrangers qui ne sont pas sur LHCOPN
- Impact sur ATLAS et CMS
 - https://ggus.eu/ws/ticket_info.php?ticket=75983
 - https://ggus.eu/ws/ticket_info.php?ticket=75919
 - https://ggus.eu/ws/ticket_info.php?ticket=75397
 - https://ggus.eu/ws/ticket_info.php?ticket=74268
 - https://ggus.eu/ws/ticket_info.php?ticket=71864
- EN COURS
 - CCIN2P3/RENATER ont entamé la procédure d'ouverture de ticket à GEANT
 - Travail d'investigation avec Fabio depuis IHEP





Chantiers en cours



- “Whole node” job
 - CMS et 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
 - Suivi par Sébastien (?)

- Aujourd'hui
 - Statut et avenir de la LAF
- Prochaine T1/AF
 - Infrastructure de virtualisation de production
 - Jacques Garnier
 - Point sur les Pledges 2011
 - Promis juré craché...

■ A venir

- LHCOPN/LHCONE, LBL, 30-31 January 2012
- OGF 34, Oxford, 11-14 March 2012
- HEPiX Spring Meeting, Prague, 23-27 April 2012
- 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
- EGI Technical Forum, Prague, September 2012
- HEPiX Fall Meeting, IHEP, October 2012