



LCG-France Tier-1 & AF

Réunion de Coordination

Fabio Hernandez
fabio@in2p3.fr

Lyon, 7 février 2008

dapnia
ceci
saclay

CNRS
CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE



Table des Matières



- Réunions CCRC, GDB et MB de février 2008
- Chantiers en cours
- Disponibilité, fiabilité, efficacité du site
- Thème du jour



GDBs & MBs



- Agendas:
 - GDB: <http://indico.cern.ch/categoryDisplay.py?categId=31181>
 - MB: <http://indico.cern.ch/categoryDisplay.py?categId=666>
 - CCRC: <http://indico.cern.ch/categoryDisplay.py?categId=1613>
- Principaux sujets traités
 - Middleware pour l'autorisation
 - Accounting CPU & stockage
 - Traceability & logging
 - Taille des fichiers sur bande
 - Pilot jobs
 - Middleware natif 64bits
 - Déploiement du middleware

Migration vers SL5



Recommendation

CERN IT
Department

- IT/FIO & IT/GD recommend
 - development of SLC5 services at CERN during 2008,
 - introduction of production batch and interactive services in September, and
 - phasing out SLC4 batch and interactive services during Q1 2009.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Options for SLC @ CERN - 6



Source: Tony Cass, [MB 05/02/2008](#)

Etude des mécanismes d'autorisation



Goal of the Study

- Task by C.Grandi to look into authorization (authZ) in gLite with the goal
 - review authZ from end-to-end
 - to specify design for "authorization service" work item in EGEE-II/-III
 - EGEE-III proposal: authZ service: Nikhef, UvA, CNAF, SWITCH
- Should specify the work in 2008 / early 2009
 - Comment: should be fully deployed within lifetime of EGEE-III
- Deliverable is a proposal with clear recommendations based on input of many people (experiments, SAx, JRA1) to be accepted/rejected by TCG

EGEE-II INFSO-RI-031688

GDB 6.2.2008

GDB 6.2.2008 26

Pierre et Sylvain pourraient être intéressés par cette présentation



Summary

- Documents <https://edms.cern.ch/document/887174/1> :
 1. MJRA1.7 document: authorization in gLite:
 - Description of authZ mechanisms in gLite
 - milestone Jan 31, 2008
 2. Document with recommendations
 - recommendations_authZ.doc
 3. Document with proposal decoupling FQANs - shares
 - recommendation_fqans_shares_v1.0.doc
 4. Forthcoming: Design for new authZ service: in progress
- Feedback needed !
 - Particularly on proposal decoupling FQANs - shares
 - Deadline end of February
- TCG will review situation in March

EGEE-II INFSO-RI-031688

GDB 6.2.2008 26

Source: Christoph Witzig, [GDB 06/02/2008](#)



Job priorities



Tests performed

- Goals
 - Verify that the jobs are submitted to the queues where the FQAN in the proxy is authorized
 - Verify that the number of waiting and running jobs are correctly reported in the VOView information
 - Verify that the WMS takes the information used to calculate the rank from the correct VOView
- Method
 - Submit jobs with different VOMS FQANs and check

Results

- All checks *eventually* succeeded
- It took some time to make things work
 - Some misconfigurations
 - The experience will come useful to prepare the documentation

Source: Andrea Sciaba, [GDB 06/02/2008](#)

WN pour SL 64bits



Status

SL4 64bit WN

- Build done
- On certification testbed since 2 months
 - Seems to work
- Issues
 - One configuration change pending
 - Ld library path
 - **Works only with 64bit user code**



Status

SL4 32/64bit WN

- Build done
 - Packing problems
 - Minor, affects only lcg_util
 - 32 or 64 binary may be selected / Both should work
- In Certification
- Issues
 - Meta packages
 - RPM can't identify the architecture of its dependent packages
 - Solution: YUM Groupinstall
 - In certification
 - Ld_library path YAIM fix in preparation
 - Release to PPS (can't be fast tracked)

18

Source: Markus Schulz, [GDB 06/02/2008](#)

Deployment du middleware



Problems

Propagation of new clients

- Clients need several months to reach all sites
 - Scheduled by sites
 - Examples:
 - Still > 100 sites on non gLite-3_1_x (40%)
 - Released June 2007
 - 52 sites installed Update 07 or better
 - December 6th
 - User can't select on fine grained versions
 - Like: lcg_util-XXX version
 - Common use case
 - No possibility for a pre-view
 - Downgrading is very difficult
 - Slow and careful updates

Source: Markus Schulz, [GDB 06/02/2008](#)



Suggestion

Make new, current and old clients available

- Everywhere
- Instantaneously
- How?
 - Use LHCb's software distribution mechanism
 - Tar balls in shared file system space of ops VO
 - Pushed by SAM test
 - Advertise: actual version of WN libs available
 - Logical selection by:
 - Absolute version
 - Symbolic version (New, Prod, Old)
- WN installed version??
 - If nothing or a matching version is requested in the jdl
 - Site provided version is used



Traceability & Logging Policy



Joint Security Policy Group

Grid Security Traceability and Logging Policy

<https://edms.cern.ch/document/428037>

Quel est l'état
d'implémentation de cette
politique sur notre site?

4 Production and Retention of Logging Data

In order to satisfy the traceability requirements, software deployed in the Grid **MUST** produce sufficient and relevant logging. The software **SHOULD** follow the Grid security guidelines on logging but **MUST** include the ability for a Site to collect logs centrally.

The level of the logging **MUST** be configured to produce the required information which **MUST** be retained for a minimum of 90 days. Grid Security Operations **MAY** define longer periods of retention for specific services and/or operational requirements. The logs **MUST** be put into a safe repository.



CPU accounting



Why sites don't publish UserDN?

- To publish the UserDN, sites have to set the `publishGlobalUserName` option to “yes”.
- Some sites may wish to suppress the DN for reasons of *personal privacy*. Once they have a Policy, WLCG may mandate its sites to publish
- By default, the DN is suppressed from publication.
- Addition of new variable in YAIM (testing in PPS): `APEL_PUBLISH_USER_DN`. If it is set to “yes”, it will enable UserDN encryption. The default is “no”.
- Normally, sites don't change the default values, so if we want to obtain the UserDN information the default value should be set to “yes”.

Sommes-nous prêts
techniquement et
léggalement à envoyer
l'information nominative
d'accounting CPU à
APPEL (UserDN et
FQAN)?

Accounting stockage (encore)



Information system as data source

- Every SE runs a **generic information provider** (GIP) which publishes information according to GLUE schema (v1.2 → v1.3).
- Concept of storage areas (SAs) in GLUE.

```
GlueSiteName GlueSEArchitecture GlueSAStateUsedSpace  
GlueSAStateAvailableSpace GlueSAType GlueSAPath
```

- Harvest these attributes by querying a top level BDII → MySQL DB.
- **Nothing extra has to be installed on the SE.**

Greig Cowan

GDB

Feb 2008

Portal: <http://goc02.grid-support.ac.uk/accountingDisplay/index.php>

Source: Greig Cowan, [GDB 06/02/2008](#)



Web portal

<http://goc02.grid-support.ac.uk/accountingDisplay/index.php>

- Allows site admins, production managers and ROC managers to view the data.
- **Historical plots** of the used and available space.
- **Table** of current used and allocated space numbers.
 - Download latest numbers in CSV format.

Allocated space for VO = VO_GlueSAStateUsedSpace + VO_GlueSAStateAvailSpace

Greig Cowan

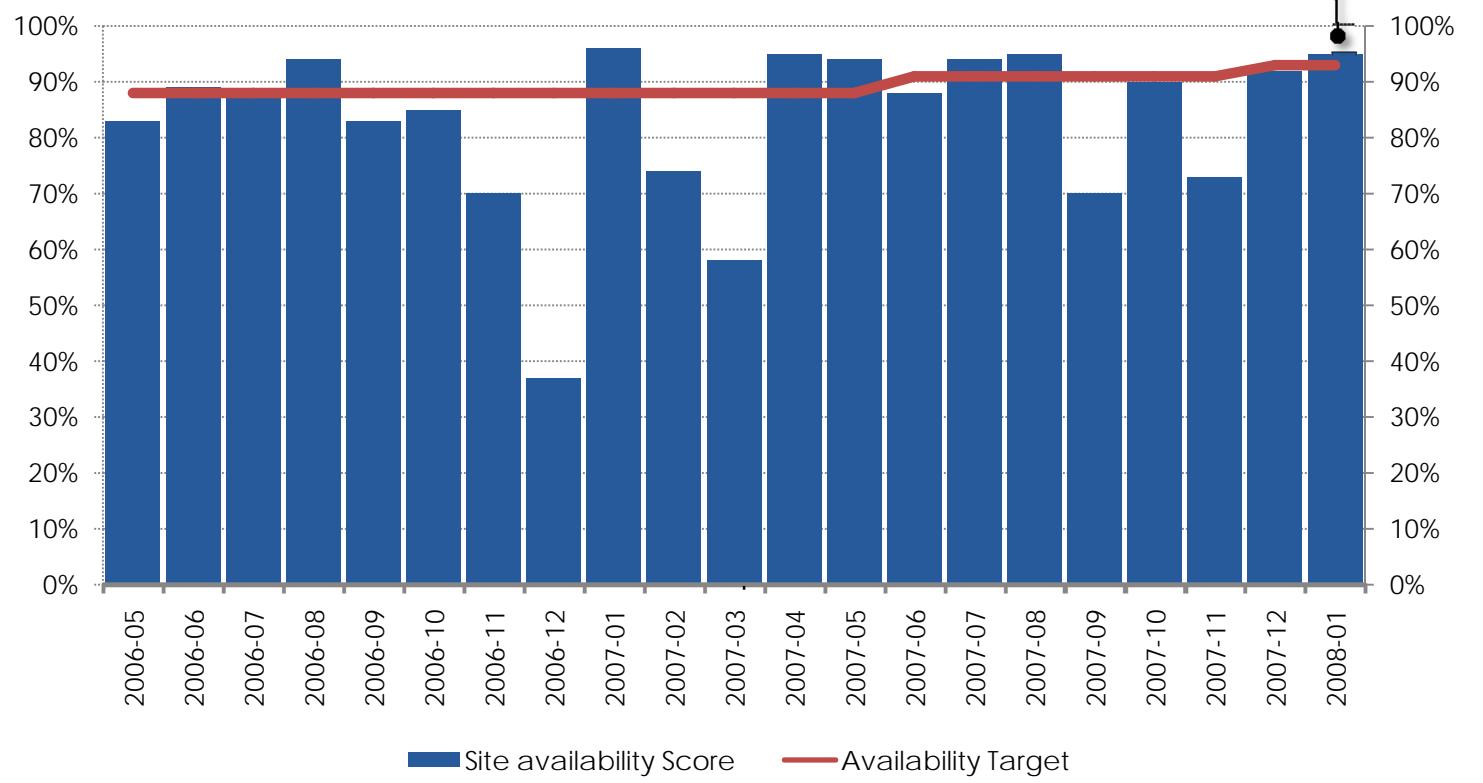
GDB

Feb 2008

Disponibilité du site



CC-IN2P3 Tier-1: monthly availability score
(VO OPS)



En janvier 2008,
l'objectif de
disponibilité
a été à nouveau
atteint



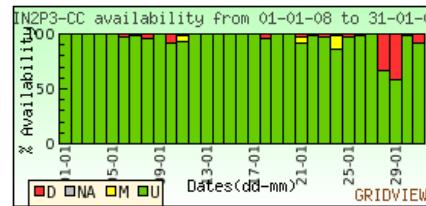
Disponibilité du site (1/3)



- Janvier 2008

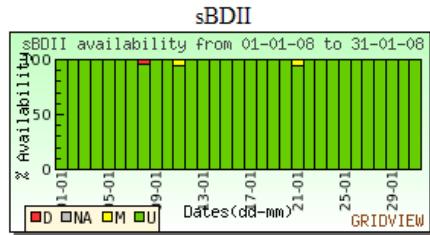
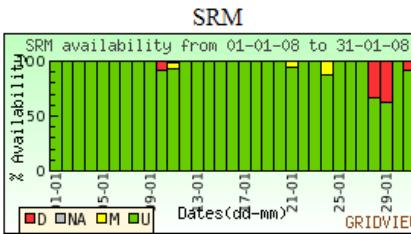
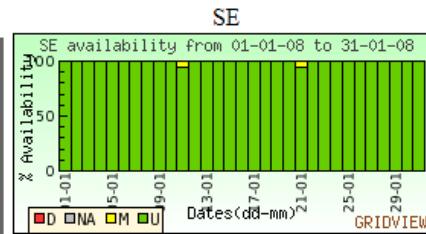
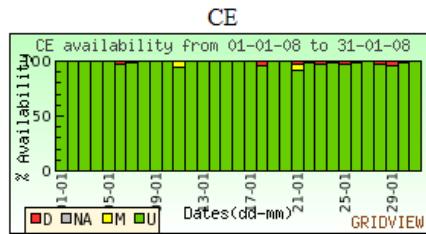
Score: 95%

Overall Service Availability for Site:IN2P3-CC VO:OPS (Daily Report)



Indisponibilité liée
au service SRM

Individual Service Availability for site:IN2P3-CC VO:OPS (Daily Report)



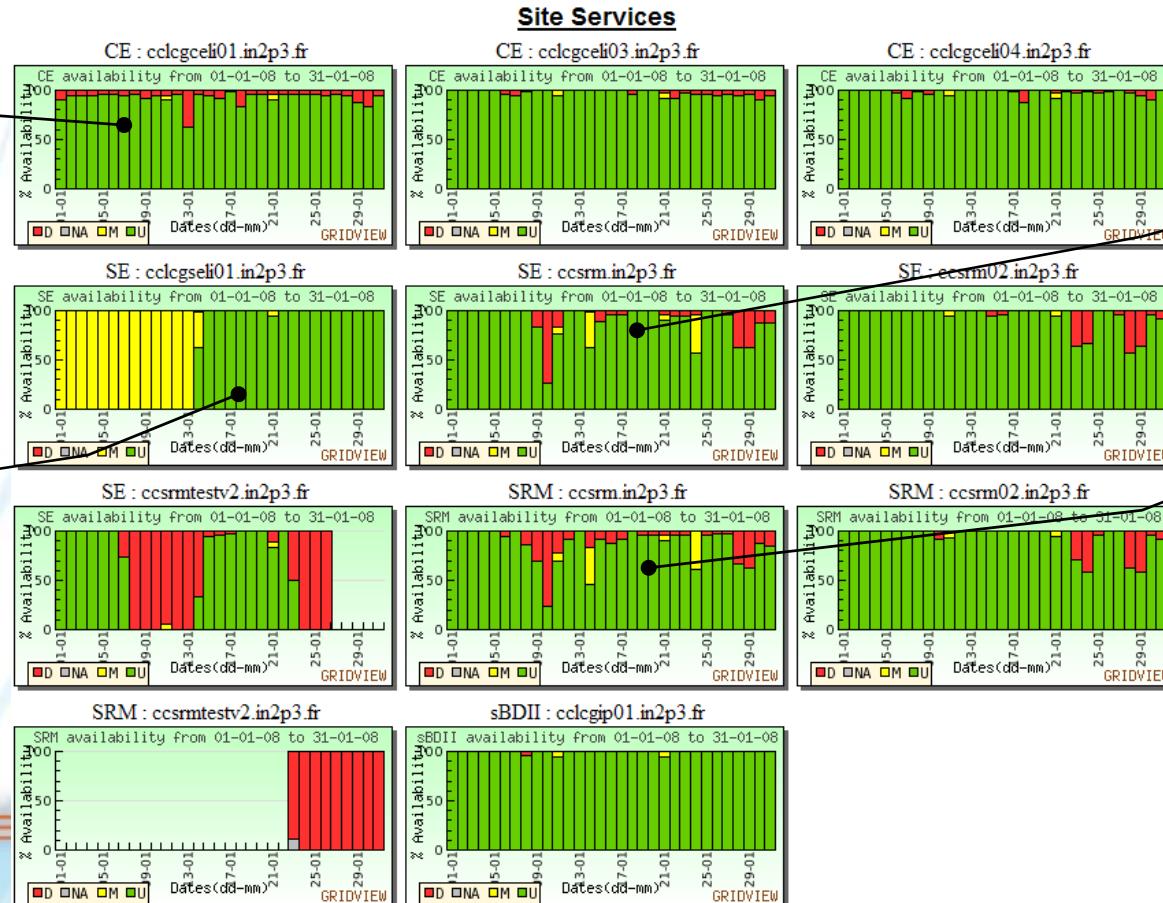


Disponibilité du site (2/3)



- Janvier 2008 (suite)

Service Instance Availability for site:IN2P3-CC VO:OPS (Daily Report)



Tel qu'en décembre 2007, ce CE s'est montré plus indisponible que les autres. Des raisons particulières?

Réapparition du classic SE?

Les raisons de ces indisponibilité sont-elles comprises?

Quelle est la différence des tests pour le service SE et ceux pour le service SRM?

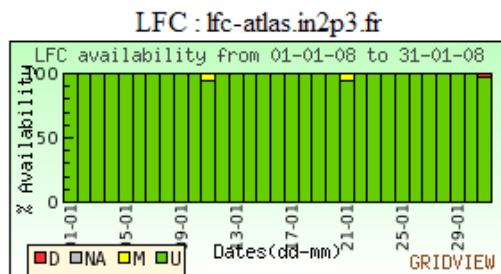
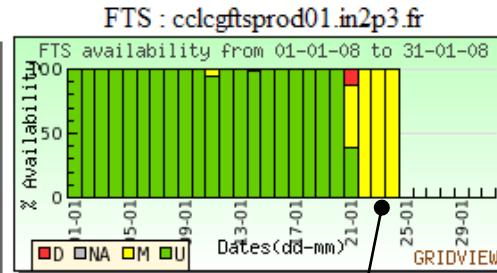
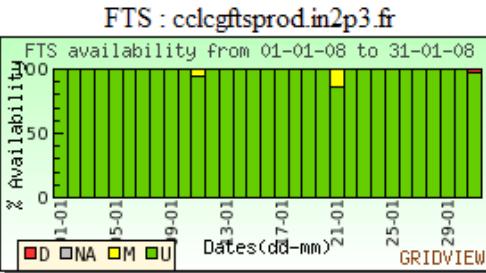
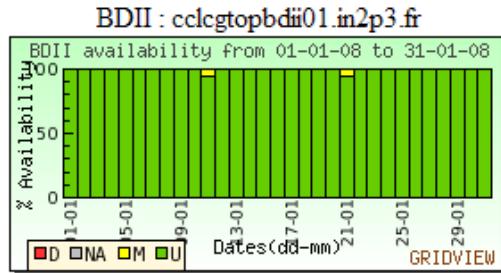
<http://gridview.cern.ch>

Disponibilité du site (3/3)



- Janvier 2008 (suite)

Central Services



Arrêt définitif du service FTS 1.5



VO Boxes: rappel bis



- Travail systématique nécessaire pour identifier
 - Configuration requise de la machine
 - Mémoire, partitionnement du disque
 - Niveau de redondance nécessaire
 - Service système à démarrer
 - gsiSSH, PNFS, ...
 - Services spécifiques VO à démarrer
 - Tâches périodiques
 - Rotation des logs, backups, ...
 - Monitoring
 - Système et services VO
 - Alertes
- Documentation du processus



Revue des chantiers en cours



- Intégration
 - Tests de LCG-CE sous SL4
 - Développement de l'interface BQS pour CREAM-CE
 - Tests d'impact du changement d'utilisateur UNIX (par glexec) pour un job BQS
- Exploitation
 - Consolidation des services grid: LFC, FTS, CE, système d'information
 - Intégration des VO boxes à l'exploitation standard
 - Procédures d'exploitation services grid
 - Plate-forme de monitoring basée sur NAGIOS
 - Séparation des la consommation CPU tier-1 et tier-2 dans la comptabilité interne
 - Rédaction des SLAs pour les VO boxes
 - ✓ Déploiement de LFC-RO pour LHCb
 - ✓ Arrêt de FTS v1.5

Base versions of services and client tools for CCRC'08 (February 2008)

We list here the base versions of the services and client tools that are needed for CCRC'08. Please, make sure your site is up to date.

[gLite 3.1 release notes](#)

[gLite 3.0 release notes](#)

| Service | SubComponent | Base Version | Status |
|---------|----------------|---|---|
| FTS | Server/SL3 | 2.0 PATCH:1589 | Available, inc PATCH:1589 (gLITE 3.0 update 39) |
| dCache | Server | 1.8.0-12 with p1, p2 and p4. | 1.8.0-12-p1,p2 and p4. |
| CASTOR | Core Services | 2.1.4-10 or 2.1.6 | Available |
| CASTOR | Stager | 2.1.4-10 or 2.1.6+GridFTP2 | Available |
| CASTOR | SRM 2.2 Server | ≥ 1.3-9 (1.3-12) | Available |
| DPM | Server | 1.6.7-2 (SL3 or SL4) PATCH:1369 | Released 16th of January ('-1' on SL4) |
| STORM | Server | 1.3.19 | Available the 31st of January 2008 - see release plan |
| LFC | Server | 1.6.8-1 (SL4 32&64 bit) PATCH:1615 , PATCH:1617 | In production 24th Jan |
| CE | Server | LCG-CE 3.0 (SL3) or 3.1(SL4) | In production |
| WN | Service/SL4 | gLITE 3.1 | In production |
| WMS | Server/SL3 | glite-WMS-2.4.6-0 | Released 14th of January |
| LB | Server/SL3 | glite-LB-2.3.4-0 | Released 14th of January |
| RB | Server/SL3 | 3.0.19-1 | In production |
| ARC | Server | 0.6 | In production |

| DM Client Tools | Base Version | Status |
|-----------------|---|--|
| GFAL | 1.10.6-2 (SL4) PATCH:1587 | Released to production Mon 21 Jan. |
| lcg_util | 1.6.5-1 (SL4) PATCH:1587 | Released to production Mon 21 Jan. |
| GFAL | 1.10.7-2 (SL4) PATCH:1641 | 1.10.7-1 ; Production Thursday 24th. Is the '-2' requirement real? |
| lcg_util | 1.6.6-1 (SL4) PATCH:1641 | Production Thursday 24th. |

Source: <https://twiki.cern.ch/twiki/bin/view/LCG/GSSDCCRCBaseVersions>



CCRC: journal

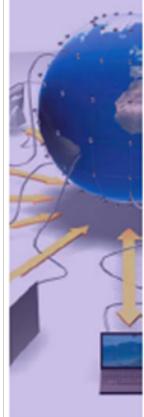


GS

Problems with the twiki

CERN IT
Department

- Hard to generate reports from a twiki
- Statistics extraction is manual
 - Messages/Incidents per day, per site, ...
- Everyone has to poll
 - No feeds
- No categorization
- No threading
- Want it to be write-once, read-many
 - No changing history !



CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

GS

Solution

CERN IT
Department

We believe elog gives us these features

- Let's use CCRC'08 to test it
 - Fallback solution could be a standard blog
- I'd encourage everyone to use it
 - Also secretary of CCRC'08 daily meeting will add items of interest that arise

...Demo...

<https://prod-grid-logger.cern.ch/elog/CCRC'08+Observations/>

RSS feed : <https://prod-grid-logger.cern.ch/elog/CCRC'08+Observations/elog.rdf>

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it





CCRC: journal



| CCRC'08 Logbook, Page 1 of 1 | | | | | | | | | | | | |
|--|--------------------------------|----------------|------------------------|--------------------|--------------|------------------|------------|--|-----------------|-----------------|------|--|
| List Find Select Login Help Full Summary Threaded | | | | | | | | | | | | |
| ID | Date | Author | Email | VO Affected | Severity | Log Stage | Site | MoU Area | GGUS Ticket URL | Original Author | Or D | |
| 19 | Wed, Feb 06, 2008, 17:21 +0100 | Flavia Donno | Flavia.Donno@cern.ch | ATLAS | Top Priority | Site Fixed | US-T1-BNL | Acceptance of data from the Tier-0 centre | | | | |
| 18 | Wed, Feb 06, 2008, 17:15 +0100 | Jan van Eldik | Jan.van.Eldik@cern.ch | CMS | Urgent | Site Fixed | CERN-PROD | Event reco or distribution of data to Tier-1 Centres | | | | |
| 17 | Wed, Feb 06, 2008, 13:06 +0100 | Flavia Donno | Flavia.Donno@cern.ch | CMS | Urgent | Site Acknowledge | CERN-PROD | Event reco or distribution of data to Tier-1 Centres | | | | |
| 16 | Wed, Feb 06, 2008, 12:44 +0100 | Flavia Donno | Flavia.Donno@cern.ch | ATLAS CMS LHCb | Top Priority | Site Acknowledge | FR-CCIN2P3 | Acceptance of data from the Tier-0 centre | | | | |
| 15 | Wed, Feb 06, 2008, 11:47 +0100 | Simone Campana | Simone.Campana@cern.ch | ATLAS | Urgent | New Problem | NDGF | T1 centres outside data taking | | | | |
| 14 | Wed, Feb 06, 2008, 10:55 +0100 | Flavia Donno | Flavia.Donno@cern.ch | ATLAS CMS LHCb | Top Priority | New Problem | FR-CCIN2P3 | Acceptance of data from the | | | | |

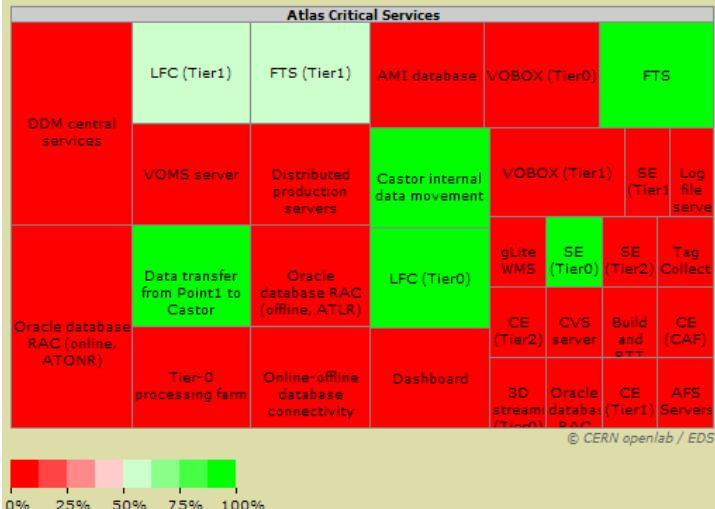
<https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook>

CCRC: visualisation de l'état des services



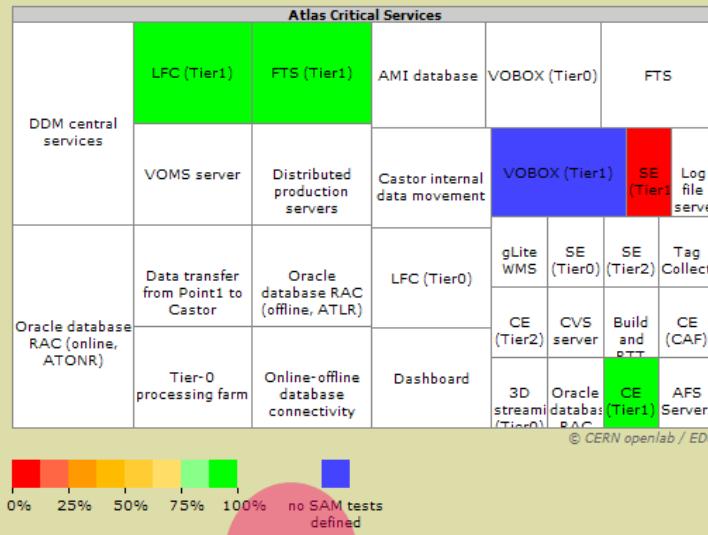
WLCG CCRC'08 Critical Services "GridMap"

Ticklist Status (updated manually)



Alice Atlas CMS LHCb

Test Status (live data)



CERN-PROD TRIUMF-LCG2 IN2P3-CC FZK-LCG2 INFN-T1 SARA-MATRIX NDGF-T1 pic Taiwan-LCG2 RAL-LCG2 BNL-LCG2
USCMS-FNAL-WC1

<http://gridmap.cern.ch/ccrc08/servicemap.html>



Aujourd'hui



- Préparation et avancement du CCRC
- Monitoring basé sur NAGIOS
- Prochaine réunion
 - Jeudi 13 mars 14h salle 202
- Agendas de toutes les réunions
 - <http://indico.in2p3.fr/categoryDisplay.py?categoryId=102>



Questions/Commentaires

