



# LCG-France Tier-1 & AF

*Réunion de Coordination*

**Fabio Hernandez**  
fabio@in2p3.fr

Lyon, 16 octobre 2008

dapnia  
cea  
saclay

**CNRS**  
CENTRE NATIONAL  
DE LA RECHERCHE  
SCIENTIFIQUE

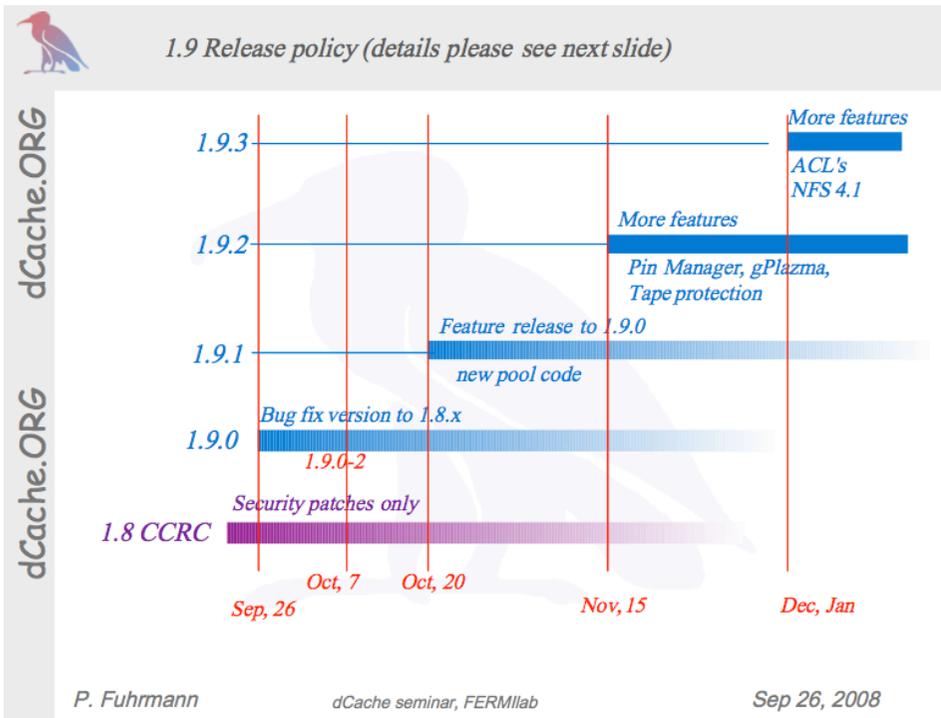
# ▶ Table des Matières



- Réunions MB de septembre et octobre 2008
- Disponibilité, fiabilité, efficacité des sites tier-1 et tier-2
- Chantiers en cours
- Thème du jour

- Agendas:
  - GDB: <http://indico.cern.ch/categoryDisplay.py?categId=31181>
  - MB: <http://indico.cern.ch/categoryDisplay.py?categId=666>
- Principaux sujets traités
  - Mises à jour middleware
  - Fonctionnement des tier-2s
  - Activité récente des expériences

# Stockage



 Release policy

- > The 1.8.x CCRC release series will only get security fixes from now on.
- > 1.9.0 is available by Oct 8.
- > With 1.9.0 we provide detailed 'release notes' in addition to the already known 'change logs'. Please check the dCache.org download area.
- > 1.9.0 is a low risk upgrade from 1.8.x because it's just a bug fix version from 1.8.x
- > 1.9.1 and higher will contain the actual new features. (see following slides)
- > Simplified release naming convention : 1.9.x-y is the y patch to release 1.9.x. (Shorter names)
- > We encourage sites to upgrade to the 1.9.0 at least.

*As soon as we have proven that 1.9.0 is stable and behaves well, support for 1.8 will be terminated.*

P. Fuhrmann dCache seminar, FERMIlab Sep 26, 2008

Source: Flavia Donno, [Pré-GDB, 07/10/2008](#)

# ▶ Mises à jour middleware



## Middleware planning



- VDT1.10
  - Status - The SL4 release of gLite is based on VDT1.6. The SL5 release will be based ultimately on VDT1.10. We could upgrade the SL4 version to 1.10 too. This would require a rebuild and retest of all components.
  - Target - upgrade of gLite3.1/SL4 to VDT1.10. ( **this might be less important since critical improvements are back ported**)
- FTS/SL4
  - Status - A problem has been found in the recent SL4 release (undeployed) so a fixing iteration will be necessary.
  - Integration will set up an SL5 build to get an idea of its potential but its deployment is not currently the plan.
  - Target - full deployment of 2.1 on SL4

## Middleware Planning



- WN/SL5
  - Status - FIO now has a first installation at CERN, which will be tested by the experiments.
  - Target - available on the infrastructure in parallel to SL4
  - We should also continue to pursue the python2.5 and alternative compiler stuff, but this can be added subsequently.
- Multiple parallel versions of middleware available on the WN
  - Status - at the moment it is not easy to install or use multiple parallel versions of the middleware at a site. While the multi middleware versions and multi compiler support are not disruptive, they require some changes on the packaging side and a small adaptation on the user side.
  - Target - it seems advisable to introduce this relatively shortly after the bare bone WN on SL5.

Source: Oliver Keeble, [MB, 07/10/2008](#)

# ► Mises à jour middleware (suite)



## Middleware Planning



### ■ WMS

- Status: Patched WMS ( fixing LHCb's proxy mixup problem) within 1 week
- Target: This patch should be deployed
- ICE to submit to CREAM
  - Not required for certification
  - ICE will be added in a subsequent update ( but better before Feb. 2009)

### ■ Other things we could consider;

- gridftp2 patches
  - These are being back ported to VDT1.6
  - Important or dCache and FTS

### ■ CREAM

- Unlikely to be deployable as an lcg-CE replacement on this timescale, but we can continue with rollout in parallel.

- Here we should be more aggressive:
  - LCG-CE inherently problematic for analysis

## Middleware Planning



### ■ glexec/SCAS

- Status - still nothing delivered to certification.
- Target - enabling of multi-user pilot jobs via glexec. This could conceivably be via another means than SCAS, but this would have to be decided asap.

- News Flash: NIKHEF promised a Patch within one day from now

### ■ Glue2

- Status - Glue2 is awaiting final validation at OGF, expected November.
- Target - we should try to get the new schema deployed to the BDII's so we can iron out initial deployment glitches, leaving us with a working but unpopulated Glue2 infosys in parallel to 1.3. Info providers could subsequently be upgraded gradually, as could clients.

### ■ CE publishing

- Status - A set of changes to rationalise publishing of heterogeneous computing resources is envisaged. A full roadmap will be published by Steve Traylen this week. The first phase will be the deployment of the new tools, enabling simply the current situation. Subsequent phases then take advantage of the new tools.
- Target - the first phase as described above.

Source: Oliver Keeble, [MB, 07/10/2008](#)

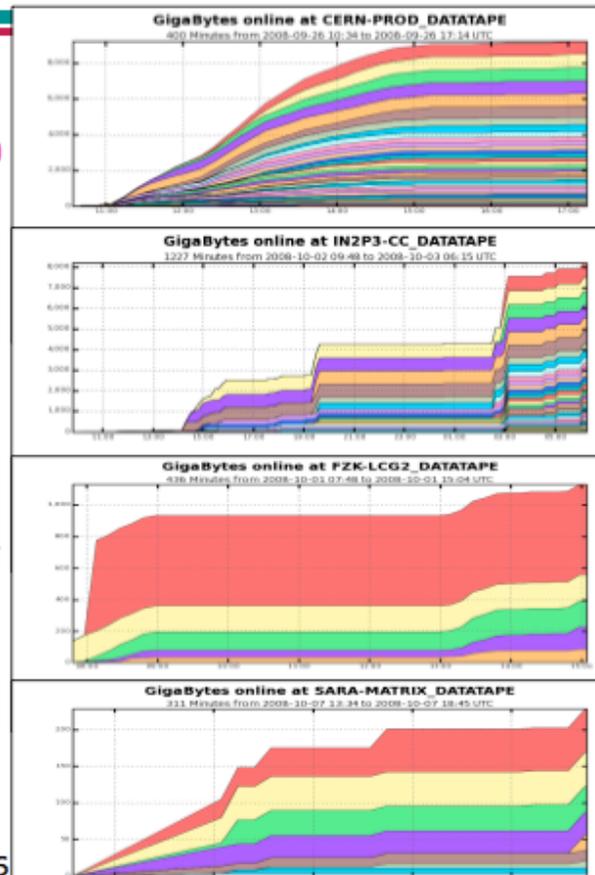
# ▶ Activités des expériences: ATLAS



WLCG MB - 14 October 2008

## Prestaging tests

- We started during the summer prestaging tests at all Tier-1s
  - Recalling whole datasets at a time (up to 10 TB)
- Performance varies a lot as tape back-ends are different at each site
  - After a few tries, most sites are mostly OK
- Outstanding (different) problems at PIC, FZK and SARA
  - There is no point in having 1000s of processing cores if they cannot be fed at the correct rate with data
  - Example:
    - Our reprocessing tasks consume 1.6 MB of raw data every ~7 real seconds
    - One needs a total read rate from tape of 400-500 MB/s to keep 1000 cores busy
      - Including x2 contingency



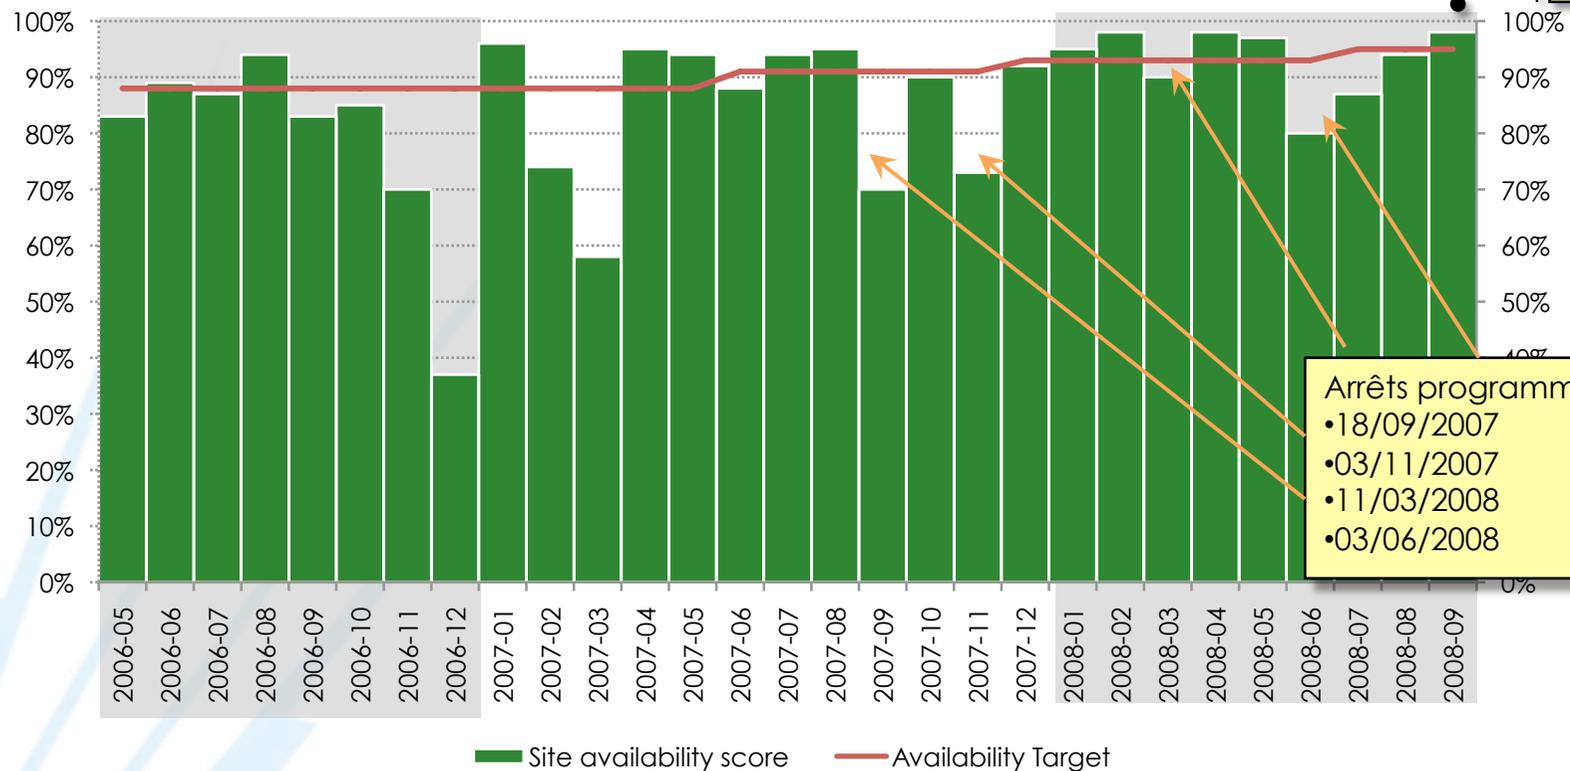
Dario Barberis: ATLAS

Détails des tests et résultats: <https://twiki.cern.ch/twiki/bin/view/Atlas/PreStageTests>

# ▶ Disponibilité du tier-1



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



Disponibilité septembre 98%

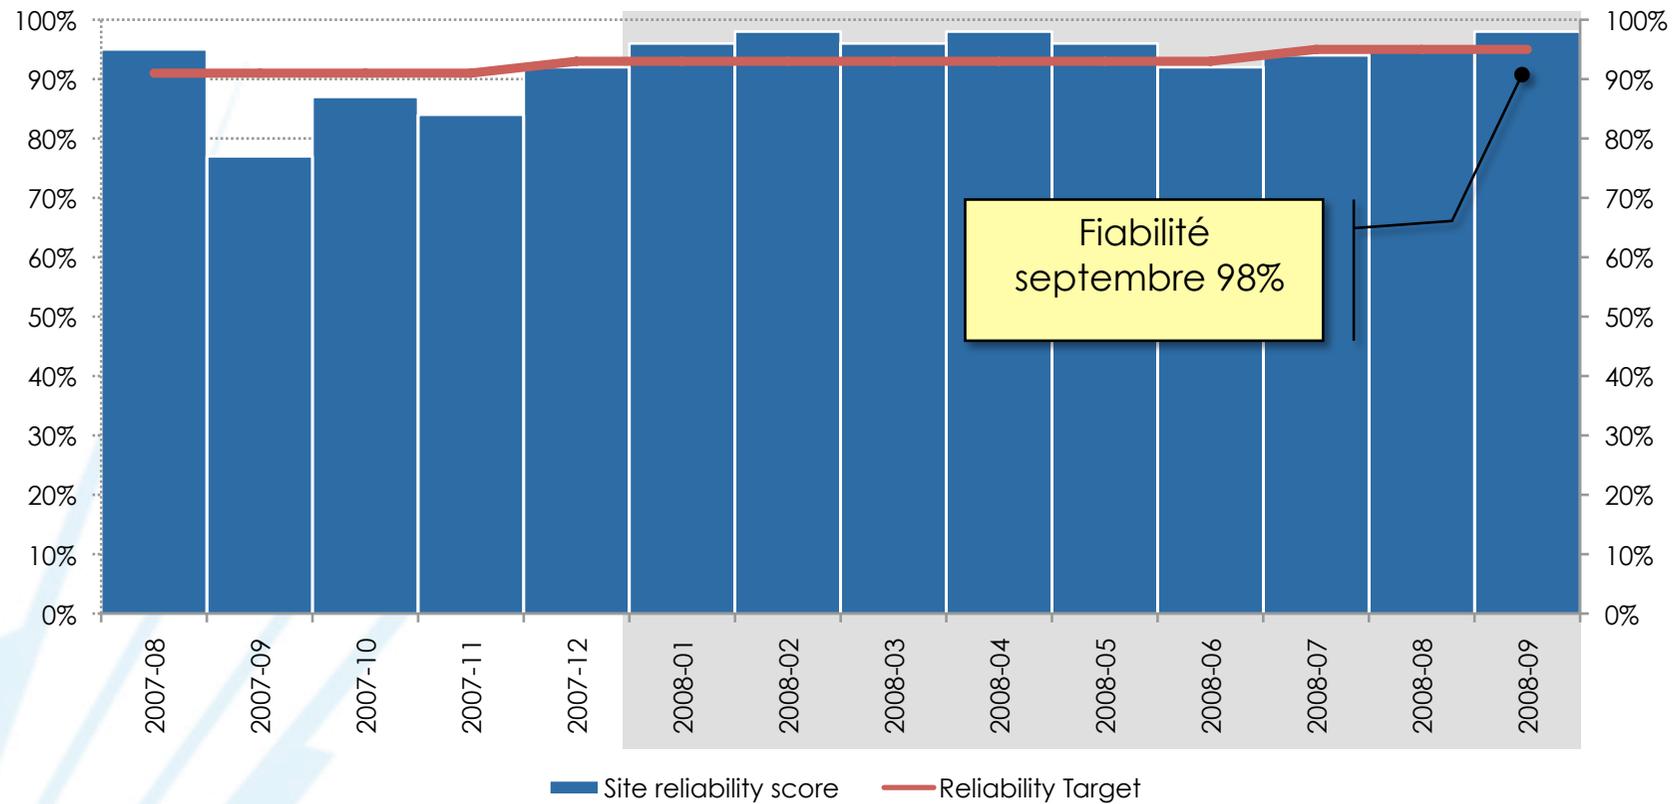
Arrêts programmés:  
 • 18/09/2007  
 • 03/11/2007  
 • 11/03/2008  
 • 03/06/2008

Source: [WLCG T0 & T1 Site Reliability Reports](#)

# ▶ Fiabilité du tier-1



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



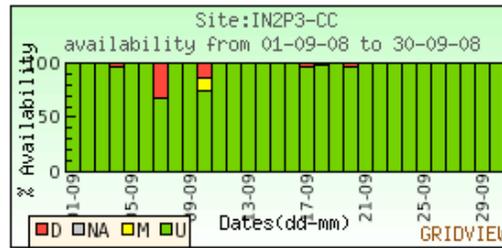
Source: [WLCG T0 & T1 Site Reliability Reports](#)

# ▶ Disponibilité des sites: tier-1 (1/5)



- Septembre 2008

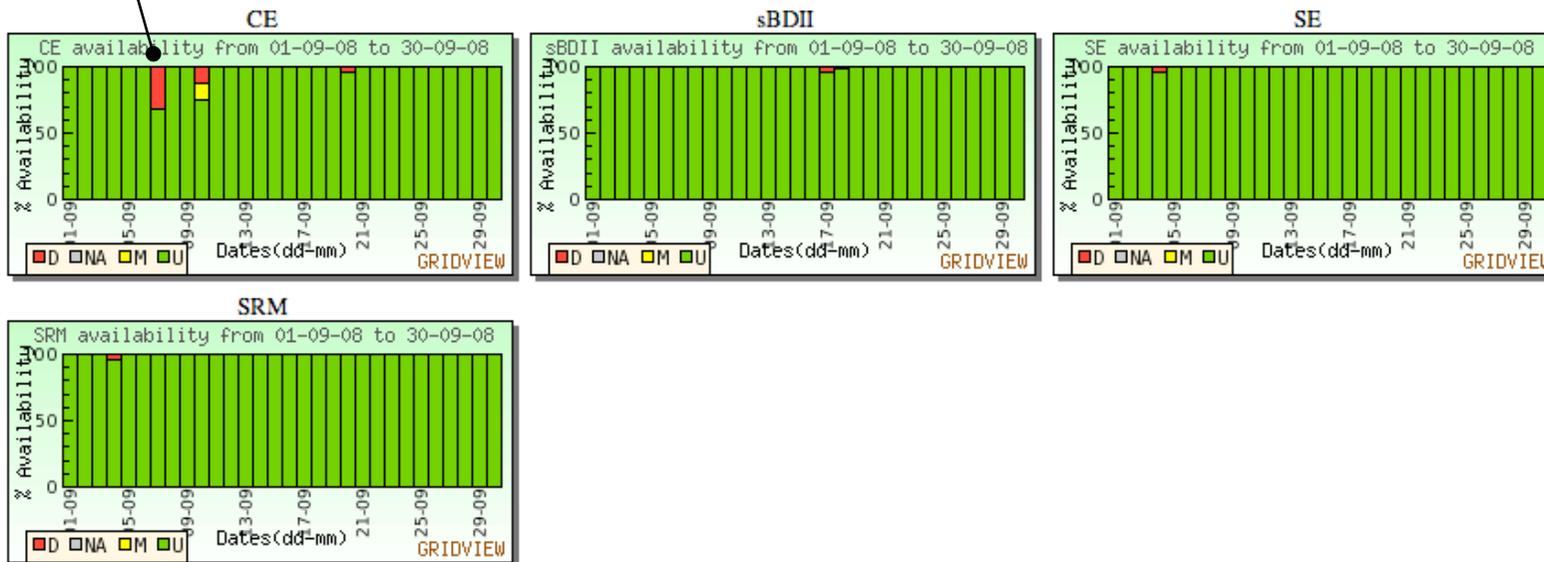
**Overall Service Availability for Site:IN2P3-CC VO:OPS (Daily Report)**



Score: 98%

????

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



Source: <http://gridview.cern.ch>



# Disponibilité des sites: tier-1 (2/5)

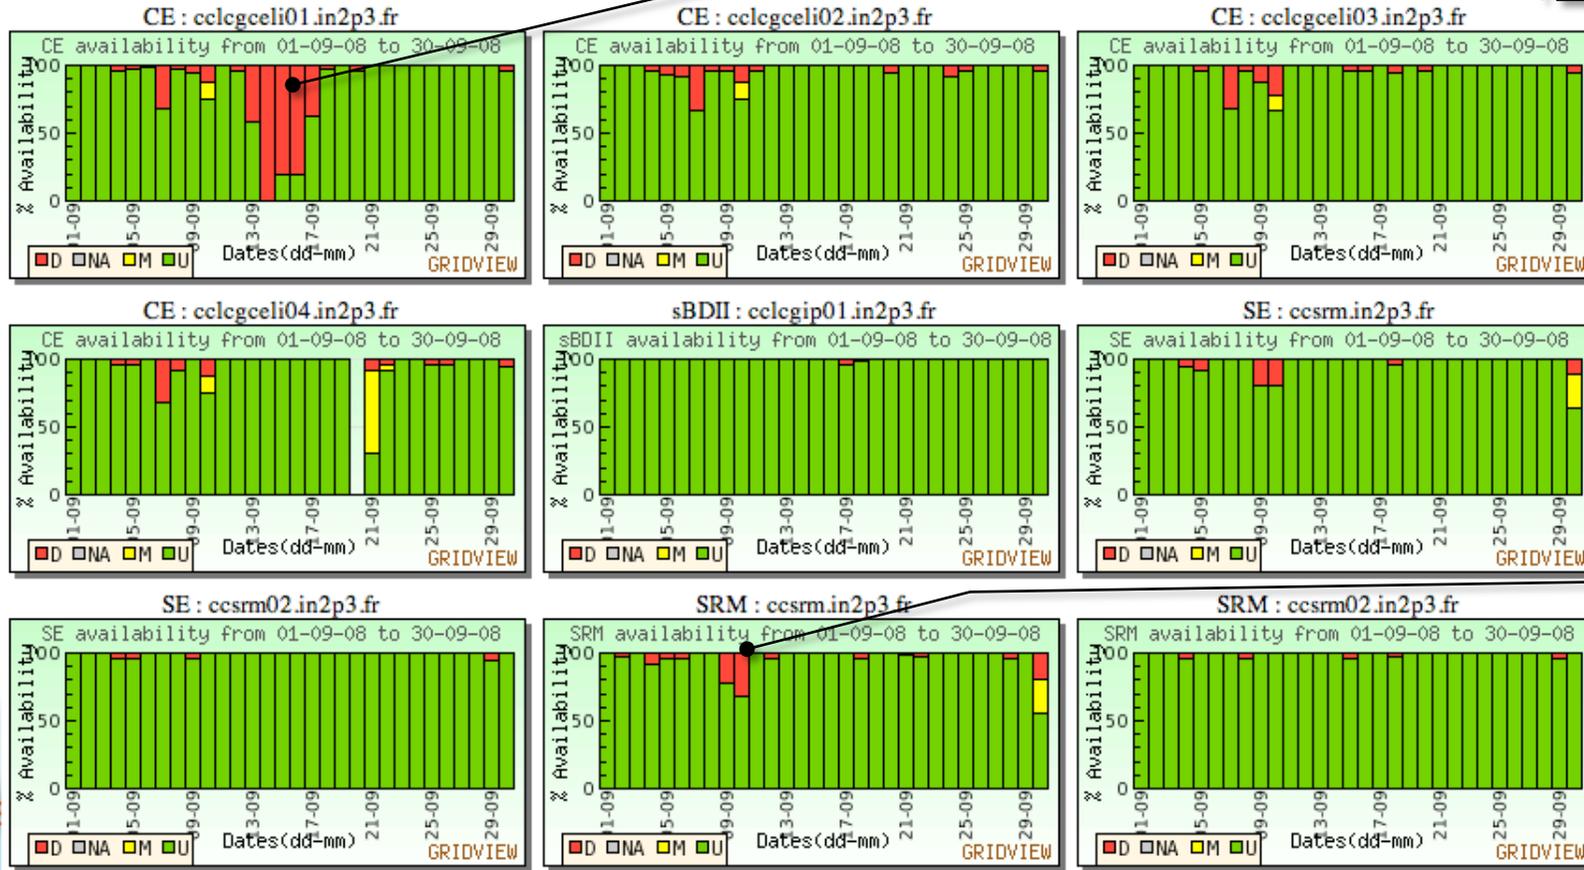


## • Septembre 2008 (suite)

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

Problème matériel?

#### Site Services



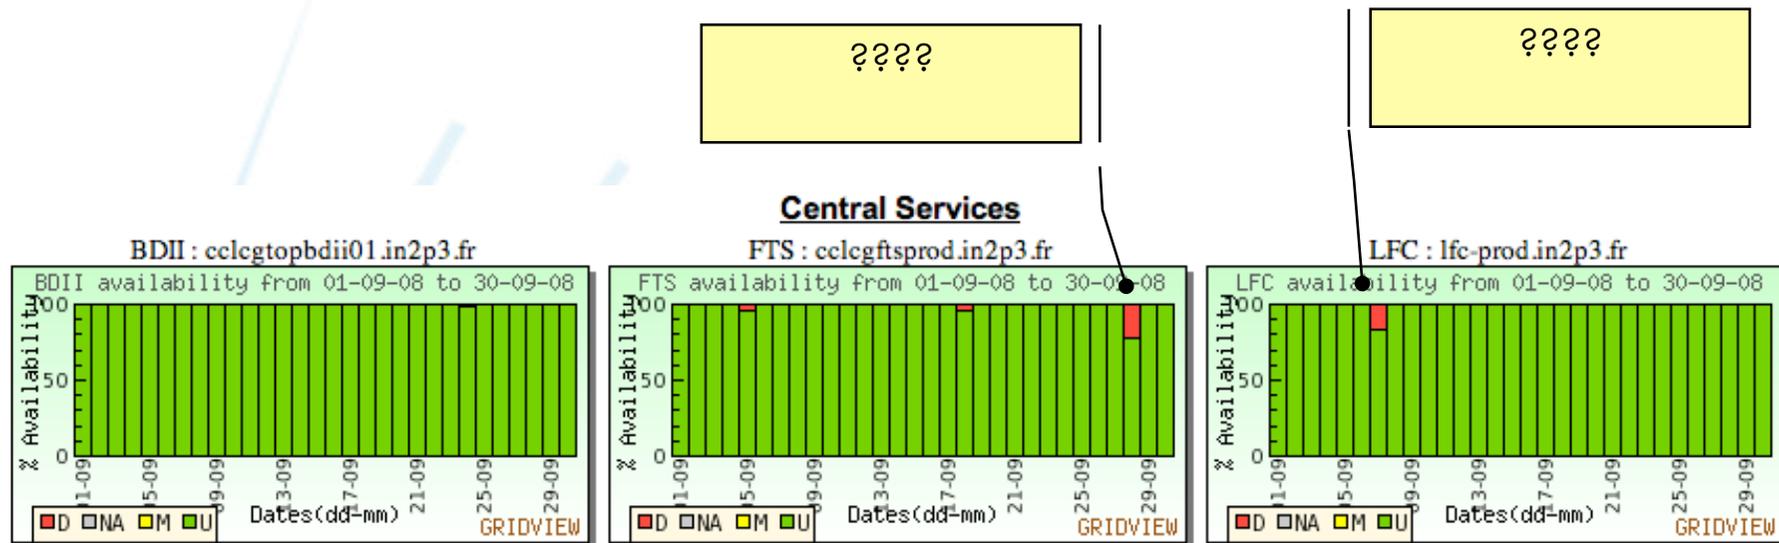
???

Source: <http://gridview.cern.ch>

# ▶ Disponibilité des sites: tier-1 (3/5)



- Septembre 2008 (suite)

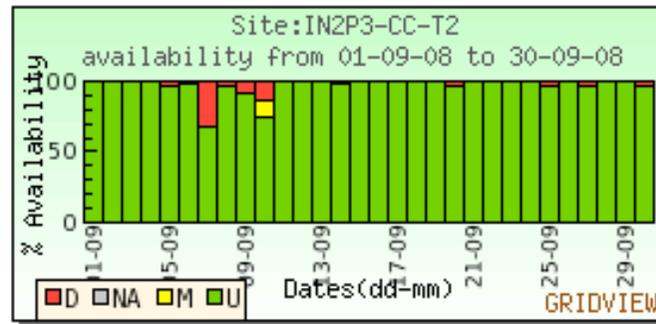


# ▶ Disponibilité des sites: tier-2 (4/5)



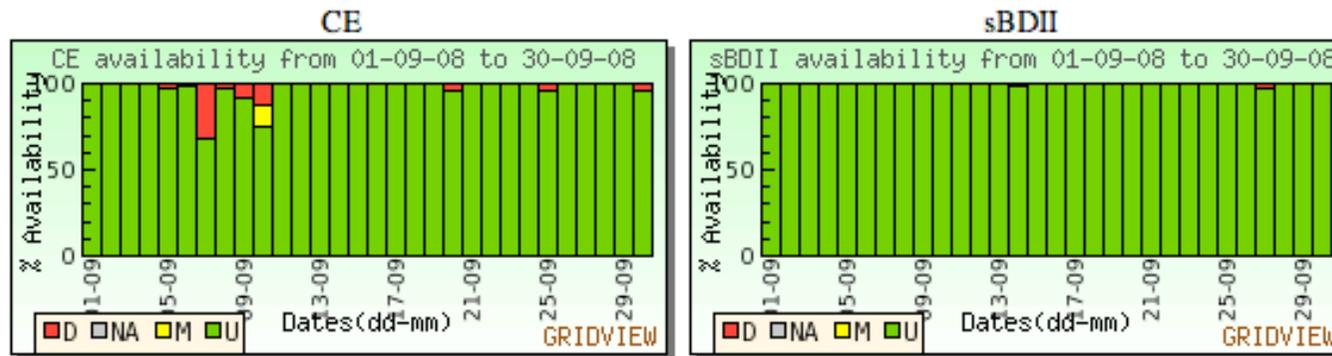
- Septembre 2008

## Overall Service Availability for Site:IN2P3-CC-T2 VO:OPS (Daily Report)



Disponibilité: 94%  
(Fiabilité: 95%)

## Individual Service Availability for site:IN2P3-CC-T2 VO:OPS (Daily Report)



Source: <http://gridview.cern.ch>

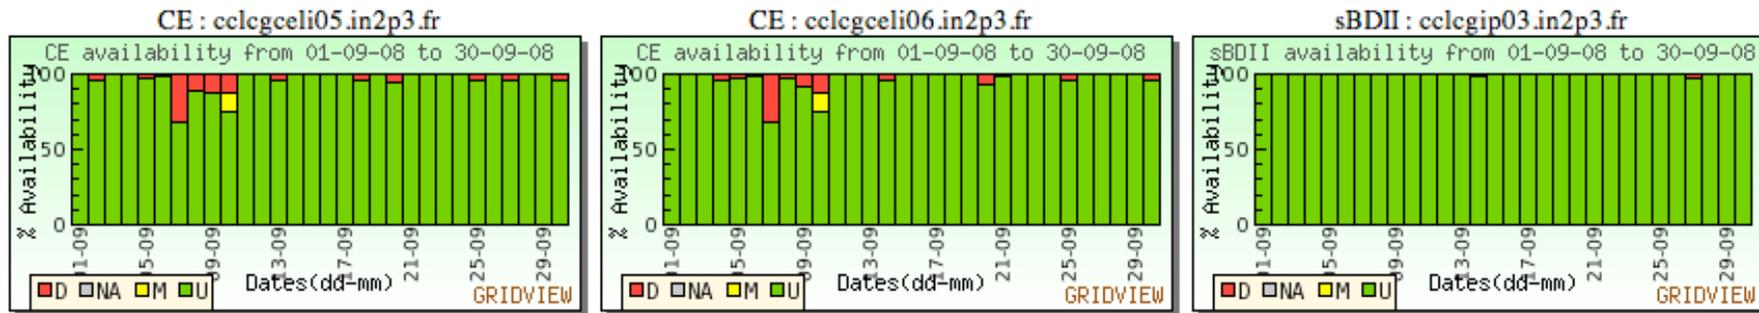
# ▶ Disponibilité des sites: tier-2 (5/5)



- Septembre 2008

## Service Instance Availability for site:IN2P3-CC-T2 VO:OPS (Daily Report)

### Site Services

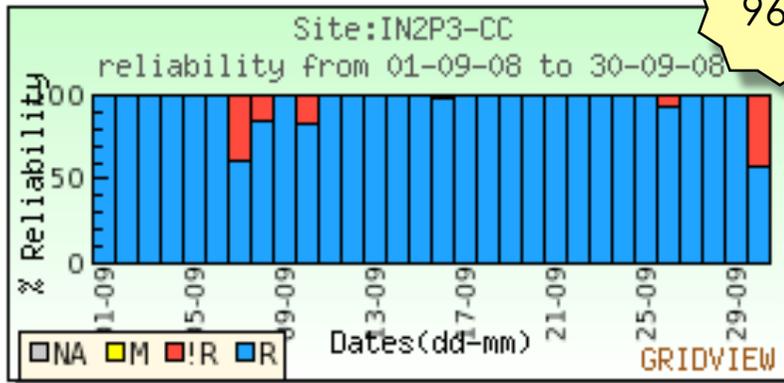


Source: <http://gridview.cern.ch>

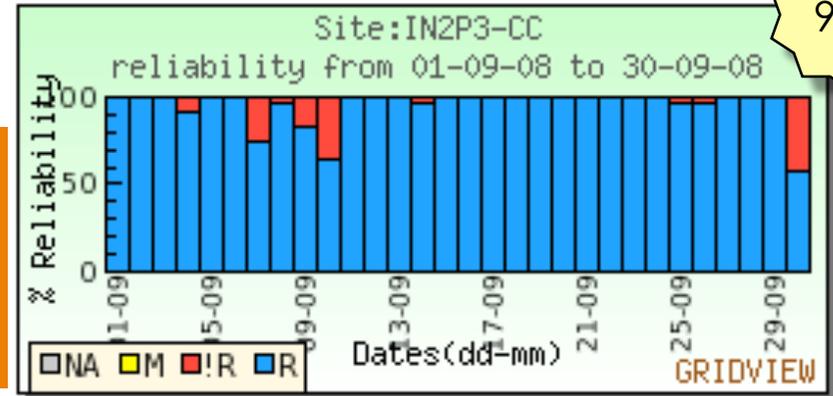
# ▶ Fiabilité vue par les VO



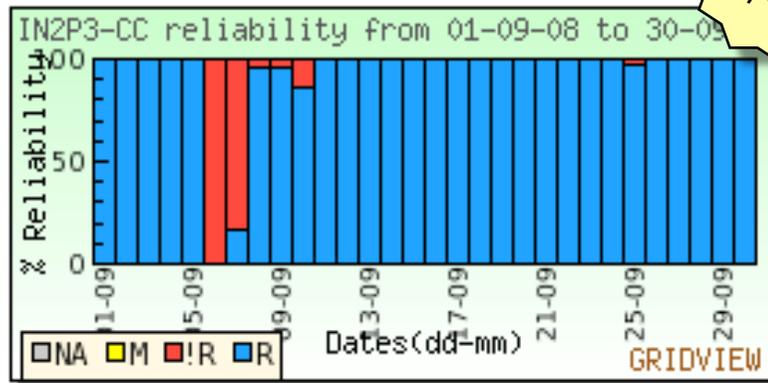
ALICE



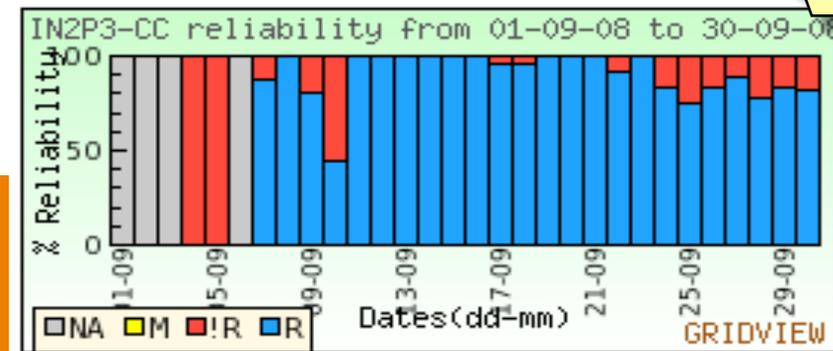
ATLAS



CMS



LHCb



Source: <http://gridview.cern.ch>

# ▶ EGEE France – disponibilité & fiabilité



- Août et Septembre 2008

Region	Avail-ability	Reli-ability
Russia	92 %	93 %
France	90 %	94 %
GermanySwitzerland	88 %	89 %
Italy	87 %	92 %
CentralEurope	86 %	92 %
UKI	86 %	94 %
SouthWesternEurope	85 %	91 %
NorthernEurope	81 %	88 %
SouthEasternEurope	80 %	86 %
CERN	77 %	80 %
AsiaPacific	73 %	78 %

Region	Avail-ability	Reli-ability
France	90 %	92 %
CentralEurope	86 %	89 %
UKI	86 %	92 %
Russia	85 %	85 %
Italy	84 %	89 %
NorthernEurope	84 %	88 %
SouthWesternEurope	80 %	86 %
SouthEasternEurope	77 %	80 %
CERN	76 %	79 %
GermanySwitzerland	75 %	81 %
AsiaPacific	74 %	79 %

Source: <https://edms.cern.ch/document/963325>



# EGEE France – disponibilité & fiabilité (suite)



## Septembre 2008

Region	Site	Avail- ability	Reli- ability	Availability History		
				Jun-08	Jul-08	Aug-08
France ( France )						
	AUVERGRID	99 %	99 %	96 %	98 %	100 %
	CGG-LCG2	86 %	94 %	99 %	99 %	95 %
	GRIF	99 %	100 %	100 %	98 %	99 %
	IBCP-GBIO	72 %	72 %	85 %	56 %	69 %
	IN2P3-CC	98 %	98 %	82 %	87 %	94 %
	IN2P3-CC-T2	97 %	97 %	82 %	87 %	94 %
	IN2P3-CPPM	97 %	97 %	90 %	94 %	99 %
	IN2P3-IPNL	99 %	99 %	99 %	97 %	99 %
	IN2P3-IRES	97 %	97 %	98 %	95 %	99 %
	IN2P3-LAPP	82 %	86 %	99 %	97 %	82 %
	IN2P3-LPC	97 %	98 %	99 %	88 %	99 %
	IN2P3-LPSC	94 %	97 %	98 %	87 %	86 %
	IN2P3-SUBATECH	99 %	99 %	92 %	89 %	96 %
	IPSL-IPGP-LCG2	99 %	99 %	95 %	97 %	98 %
	SN-UCAD	32 %	45 %	N/A	23 %	45 %

Source: <https://edms.cern.ch/document/963325>

# ▶ Revue des chantiers en cours



- Intégration
  - Tests de priorisation des jobs basée sur le rôle VOMS
  - Déploiement généralisé de SL4 64bits avec bibliothèques de compatibilité 32bits sur les workers nodes
  - Développement de l'interface BQS pour CREAM-CE
  - Tests de la nouvelle version de glexec visant un déploiement en mode *logging only*
  - Activation de l'ordonnancement intra-VO disponible dans BQS
  - Amélioration de la fiabilité des services grid: LFC, FTS, CE, SRM, système d'information
  - Déploiement de CRAB (moteur de soumission de jobs d'analyse) pour CMS
  - Consolidation des outils de monitoring de l'échange de données entre dCache et HPSS
- Exploitation
  - Intégration des VO boxes à l'exploitation standard
    - *Service Level Agreement formellement accepté par ALICE*
  - Procédures d'exploitation services grid
  - Plate-forme de monitoring basée sur NAGIOS (et intégration avec SAM)
  - Séparation de la consommation CPU tier-1 et tier-2 dans la comptabilité interne

- WLCG 2009 Data-Taking Readiness Planning Workshop
  - CERN, 13-14 novembre
  - Agenda: <http://indico.cern.ch/conferenceDisplay.py?confId=32660>
- Réunion des sites LCG-France
  - LPSC Grenoble, 27-28 novembre
  - Agenda: <http://indico.in2p3.fr/conferenceDisplay.py?confId=875>

# ▶ Aujourd'hui et à venir



- Aujourd'hui
  - Opérations des services
  - Fiabilité du service LFC
- Sujets prochaines réunions
  - Surveillance et alertes NAGIOS
  - Point sur le niveau de fiabilité de chacun des services (LFC, FTS, dCache, CEs, VOMS, BDII, ...)
- Dates prochaines réunions
  - Novembre 13: collision avec le workshop WLCG
    - *Nouvelle date sera confirmée par e-mail*
  - Décembre 18
  - Agendas: <http://indico.in2p3.fr/categoryDisplay.py?categId=102>



# Questions/Commentaires

