



LCG-France Tier-1

Réunion de Coordination

Fabio Hernandez
fabio@in2p3.fr

Lyon, 3 mai 2007

► Table des Matières




- Disponibilité du site
- GDB de Mai 2007
- Actions en cours
 - Priorités des jobs basée sur rôles/groupes VOMS
 - Accounting stockage
 - Tests SRM v2.2
- SL4
- Milestones LCG
- Thème du jour
- Thème de la prochaine réunion

Disponibilité du site (1/3)




- Avril 2007

Overall Service Availability for site IN2PCC : Daily Report



Individual Service Availability for site IN2PCC : Daily Report



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


Disponibilité du site (2/3)




Service Instance Availability for site IN2PCC : Daily Report

Site Services


CE : cclcgceli01.in2p3.fr




CE : cclcgceli02.in2p3.fr




CE : cclcgceli03.in2p3.fr




SE : cclcgceli01.in2p3.fr




SE : cclcgceli02.in2p3.fr




SE : ccsrm.in2p3.fr




SE : ccsrm02.in2p3.fr




SRM : ccsrm.in2p3.fr



SRM : ccsrm02.in2p3.fr




sBDII : cclcgip01.in2p3.fr




Central Services

BDII : cclcgtopbdii01.in2p3.fr



FTS : cclcgftsprod01.in2p3.fr



► Disponibilité du site (3/3)



- Impact du déploiement de la TopBDII régionale/locale?
 - Premières impressions
- Y a-t-il eu des alertes remontées pour les problèmes du 27 au 30/04?
 - Y compris pour l'astreinte?

- Agenda
 - <http://indico.cern.ch/conferenceDisplay.py?confId=8472>
- Sujets
 - État d'avancement du middleware
 - Rapport du group de travail sur les systèmes de fichiers (Hepix)
 - 5 problèmes principaux soulevés par chaque expérience

WMS status



- 2 weeks ago achieved:
 - 15,000 jobs per day for 7 days - 320 jobs (0.3%) did not run
 - These 320 were later restarted and succeeded - still should fix the underlying issue
 - The restart can be automated
 - All jobs in final states
 - Stress test is not yet complete - do not yet know the limit of scalability (30k jobs per day?)
 - Proxy renewal is working
 - L&B can scale much higher than WMS - not a limiting factor
 - N.B. to achieve the goals, DAG-based collections were removed and a new collection mechanism introduced
- WMS (3.1) now in standard certification process
 - Meanwhile gLite WMS nodes at CERN are updated to this version for production use
 - Next release with this version will allow full deployment and replacement of old LCG-RB

Source: Ian Bird, [GDB 02/05/2007](#)

CE Status



- Current status of gLite CE
 - Close to 100% success of job submission - after resolving a number of timing issues with Condor
 - Submissions of 6000 jobs to a CE (max ~3000 at any time)
 - Do not know scaling of Condor-C instances yet
 - Version with Condor-C instance per VO being set up - not tested yet
 - Several Condor issues were found - not yet clear on a timescale for achieving the criteria
- Fallback proposal
 - Keep the LCG-CE "as-is" - there is no effort to port to SL4 (which implies GT4 and potentially many issues)
 - Deploy either on SL3 nodes (or SL4 with Xen/SL3)
- Should set up a CREAM instance in parallel and subject it to the same testing procedure



Top 5 Experiment Issues


Experiment	ALICE	ATLAS	CMS	LHCb
Issue #1	xrootd-CASTOR2	CASTOR@CERN	CASTOR: functionality & performance	Data Access from T1 MSS
Issue #2	xrootd-DPM	Integration of DDM/FTS/SRM/ GridFTP/LFC etc	SRM I/F with functionality & performance	glexec usage
Issue #3	FTS Service	(Lack of) SRM 2.2	FTS Service	File management
Issue #4	gLite WMS	Data Storage Management Tools	Workload management	Deployment procedure
Issue #5	VOMS	Stability of the Information System	Information system	Information system

Source: Jamie Shiers, [GDB 02/05/2007](#)

Storage Accounting



- Portal
 - <http://goc02.grid-support.ac.uk/storage-accounting/view.php?queryType=storage&ExecutingSite=IN2P3-CC>
- Situation de notre site au 02/05/2007 16h31

Disk Used and Allocated per VO (Units are TB)								
Org Unit	Site	Disk	VO				Total (TB)	Percentage
France			alice	atlas	cms	lhcb		
	IN2P3-CC	Used	0.005	0.000	0.002	0.001	0.008	100 %
		Allocated	1.815	1000.000	1.812	1.811	n/a	n/a
Table Summaries								
		Used (TB)	0.005	0	0.002	0.001	0.008	
		Percentage	62.5%	0%	25%	12.5%	100%	
 CSV dump of this table								

Scientific Linux 4



Enabling Grids for E-science

Status 2nd May 2007

- **Building all components (CE, WMS,)**
 - 99% success – sl4_32
 - 82% success – sl4_64
- **Integration & Certification**
 - UI – Still one packaging bug (gLite-WMS client, #25819)
 - Patch expected today
 - No outstanding runtime problems (PPS might show some)
 - WMS – Testing
 - Still installation problems (responsive JRA1, slow progress)
 - Ready to Test
 - LB, MON (except APEL), gLite-CE, (DM waiting for FTP2 patch)
 - In testing: BDII, VOMS-Oracle
- **PPS**
 - WN – Released 02/04/2007 ----> discovered problems

Note: the experimental gLite-WMS-3.1 is NOT SL4

Source: Markus Schulz, GDB.02/05/2007

► Priorités des jobs



- A finaliser
 - Publication l'information dans le système d'information du site (VOViews)

Tests de SRM v2.2 (1/2)



SRM v2.2 Status

- Good support for srnCopy in push mode (CMS)

Summary of S2 SRM v2.2 cross test - Wednesday 2 May 2007 07:34am CEST

In these tests the srnCopy function is exercised. This function should be implemented by all available Storage System by the end of the 3Q of 2007. dCache is required to implement this function as of now. Therefore, it is OK to have red columns for all SRM endpoints except for dCache. However, it is not OK to have red rows since this means that a file cannot be copied between SRMs with simple get and put operations.

SRM function	CERN C2	DESY dCache	FNAL dCache	CERN DPM	LBNL BeStMan	CNAF StoRM
Copy Tests in PUSH mode						
CopyToCERNCASTOR	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToFNALDCACHE	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToDESYDCACHE	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToCERNDPM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToLBNLDRM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToSTORM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
Copy Tests in PULL mode						
CopyFromCERNCASTOR	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromFNALDCACHE	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromDESYDCACHE	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromCERNDPM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromLBNLDRM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromSTORM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log

CASTOR and DPM
Do not provide
srnCopy for the
moment.
Foreseen for
end of 2007.

dCache@DESY
Authorization/
Connection errors
(dCache developers
will look into the
problem
this coming week)

StoRM does not
provide srnCopy
in PULL mode

Tests de SRM v2.2 (2/2)



- Sites publiant des end-points SRM v2.2

The screenshot shows a web browser window with the address bar displaying <https://twiki.cern.ch/twiki/bin/view/LCG/GSSDendpoints>. The browser's menu bar includes 'Information Technology', 'Rerefence', 'Technology', 'Press', 'Search Engines', 'Photo', 'Voitures', 'Banque', 'To Read', and 'Personnal Interest'. Below the menu bar is a search bar with 'Jump' and 'Search' buttons. The page content shows a breadcrumb trail: 'You are here: TWiki > LCG Web > GSSD > GSSDendpoints'. The title of the page is 'List of available SRM v2.2 endpoints'. Below the title, it states: 'This is the list of available SRM v2.2 full endpoints for the **dteam** VO.' The page is divided into sections for different SRM endpoints: **CASTOR**, **dCache**, and **DPM**. Each section contains a list of endpoints with their respective URLs, publishers, and close-to information.

CASTOR

- **CERN**: <srm://srm-v2.cern.ch:8443/castor/cern.ch/grid/dteam> (CASTOR2 v2.1.1-9); Publisher: CERN_PPS; Close to: lxb2006 (CERN_PPS)
- **CNAF**: <srm://srm-v2.cr.cnaf.infn.it:8443/srm/managerv2?SFN=/castor/cnaf.infn.it/grid/lcg/dteam/test/srm2> (CASTOR2 v2.1.1-9); Publisher: PPS-CNAF; Close to: [cert-ce-01.cnaf.infn.it\(PPS-CNAF\)](#)

dCache

- **FNAL**: <srm://fapl110.fnal.gov:8443/srm/managerv2?SFN=/pnfs/fnal.gov/data/testers/NULL> (dCache v1.8Beta); Publisher: ???; Close to: ???(???)
- **DESY**: <srm://srm-devel.desy.de:8443/srm/managerv2?SFN=/pnfs/desy.de/data/dteam> (dCache v1.8Beta); Publisher: DESY-PPS; Close to: [tb004.desy.de\(DESY-PPS\)](#)
- **BNL**: <srm://dct00.usatlas.bnl.gov:8443/srm/managerv2?SFN=/pnfs/usatlas.bnl.gov/data/dteam> (dCache v1.8Beta); Publisher: ???; Close to: ???(???)
- **FZK**: <srm://pps-srm-fzk.gridka.de:8443/srm/managerv2?SFN=/pnfs/gridka.de/dteam/srmtester> (dCache v1.8Beta) **Part of pre-production**; Publisher: FZK-PPS; Close to: [pps-ce-fzk.gridka.de\(FZK-PPS\)](#)

DPM

- **CERN**: <srm://lxdpm102.cern.ch:8446/dpm/cern.ch/home/dteam> (DPM v1.6.4); Publisher: CERN_PPS; Close to: lxb2006 (CERN_PPS)

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

Milestones LCG (1/2)



WLCG-07-01	Feb 2007	24x7 Support Definition Definition of the levels of support and rules to follow, depending on the issue/alarm	Complete	The 24x7 service for the site is defined and operational since many years, based on a on-call service during out of office hours. The work currently underway is the integration of the core grid services into this framework.
WLCG-07-08	Mar 2007	Accounting Data published in the APEL Repository The site is publishing the accounting data in APEL. Monthly reports extracted from the APEL Repository.	Complete	See comments for milestone IN-07-Q1-09. Some discrepancies between the information sent to the APEL repository and the information published through the EGEE accounting portal have been regularly observed: manual interventions to selectively republish the accounting records is necessary too often. Experts from the site are in regular contact with the GOC accounting portal on this issue. An automatic way to check that the information sent actually matches the information received and stored by the repository is highly desirable.
WLCG-07-09	Mar 2007	3D Oracle Service in Production Oracle Service in production, and certified by the experiment(s)	Complete	The hardware and software infrastructure is in place. Tests performed by the experiment and by the LCG 3D team are satisfactory.

Summary of Progress

General Status	During this quarter we suffered a severe unscheduled power outage. As a result, the site was not completely operational during several days. Some corrective actions were taken for preventing this kind of incidents, including taking a fraction of the compute farm machines out of the UPS circuits to maintain under control the power consumption in those circuits. Several difficulties were also observed for doing the necessary works for upgrading the power infrastructure while maintaining the site in a fully operational state.
Middle ware Services	The dCache/SRM hardware was upgraded and the whole system reconfigured. Additional more powerful machines were setup for supporting the components of the SRM service. Since then, an improved stability of the service has been observed.
CPU Farm	Nothing special to report. Continuous work for improving the support of grid-specific features directly in the site batch system, BQS.

Extrait du
rapport du
CC-IN2P3
pour 2007Q1

Milestones LCG (2/3)



WLCG-07-02	Apr 2007	24x7 Support Tested Support and operation scenarios tested via realistic alarms and situations		
WLCG-07-03	Jun 2007	24x7 Support in Operations The sites provides 24x7 support to as standard operations		
WLCG-07-04	Apr 2007	VOBoxes SLA Defined Sites should propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes		
WLCG-07-05	May 2007	VOBoxes SLA Implemented VOBoxes service implemented at the site, and tested by the Experiment		
WLCG-07-06	Apr 2007	Job Priorities Available at Site Mapping of the Job priorities on the batch software of the site completed and information published		
WLCG-07-07	Jun 2007	Job Priorities of the VOs Implemented at Site Configuration and maintenance of the jobs priorities as defined by the VOs. Job Priorities in use by the VOs.		
WLCG-07-10	May 2007	3D Conditions DB in Production Conditions DB in operations for ATLAS, CMS, and LHCb. Tested by the experiments.		
WLCG-07-11	Depl Date + 30d	SL4 Operational at Site (for WN and UI nodes) This has to happen within 30 days after the release from GD.		It is difficult to plan for a milestone without a clear associated schedule. Our capacity to meet it will depend on when the middleware will be released for SL4: summer period is a difficult one in France for such a major change in the site.
WLCG-07-12	Jun 2007	Site Reliability above 91% Considering each 11 Tier-1 sites		
WLCG-07-13	Jun 2007	Best 8 Sites above 93% Eight sites should reach a reliability above 93%		

- Accent sur l'exploitation des services
- Etat d'avancement passé en revue mensuellement

Milestones LCG (3/3)



WLCG High Level Milestones - 2007														
ID	Date	Milestone	ASGC	CC IN2P3	CERN	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF	BNL	FNAL
WLCG-07-01	Feb 2007	24x7 Support Definition Definition of the levels of support and rules to follow, depending on the issue/alarm												
WLCG-07-02	Apr 2007	24x7 Support Tested Support and operation scenarios tested via realistic alarms and situations												
WLCG-07-03	Jun 2007	24x7 Support in Operations The sites provides 24x7 support to users as standard operations												
WLCG-07-04	Apr 2007	VOBoxes SLA Defined Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes												
WLCG-07-05	May 2007	VOBoxes SLA Implemented VOBoxes service implemented at the site, and tested by the Experiments												
WLCG-07-06	Apr 2007	Job Priorities Available at Site Mapping of the Job priorities on the batch software of the site completed and information published												
WLCG-07-07	Jun 2007	Job Priorities of the VOs Implemented at Site Configuration and maintenance of the jobs priorities as defined by the VOs. Job Priorities in use by the VOs.												
WLCG-07-08	Mar 2007	Accounting Data published in the APEL Repository The site is publishing the accounting data in APEL. Monthly reports extracted from the APEL Repository.												
WLCG-07-09	Mar 2007	3D Oracle Service in Production Oracle Service in production, and certified by the Experiments												n/a
WLCG-07-10	May 2007	3D Conditions DB in Production Conditions DB in operations for ATLAS, CMS, and LHCb. Tested by the Experiments.												
WLCG-07-11	Depl Date + 30d	SL4 Operational at Site (for WN and UI nodes) This has to happen within 30 days after the release from GD.												

Dashboard avec la vue d'ensemble de l'état de tous les tier-1s

- Etat de l'intégration à l'exploitation des services grille (suite)
- Suite des travaux sur les outils de monitoring des services

► Prochaine réunion



- Etat de l'intégration à l'exploitation des services grille (suite)
- Outils de monitoring de services (suite)
- Etat de la consolidation des services grille

► Questions/Commentaires

