



# WLCG Service Report

[Jamie.Shiers@cern.ch](mailto:Jamie.Shiers@cern.ch)

[Maria.Girone@cern.ch](mailto:Maria.Girone@cern.ch)

[Harry.Renshall@cern.ch](mailto:Harry.Renshall@cern.ch)

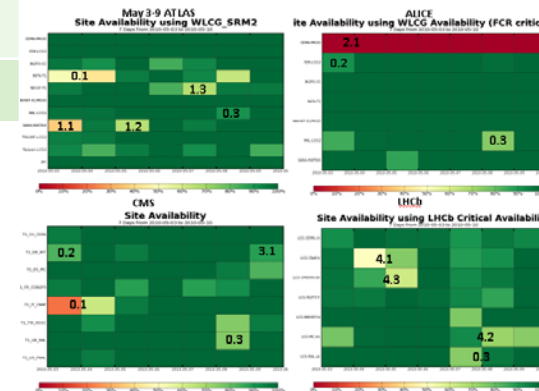
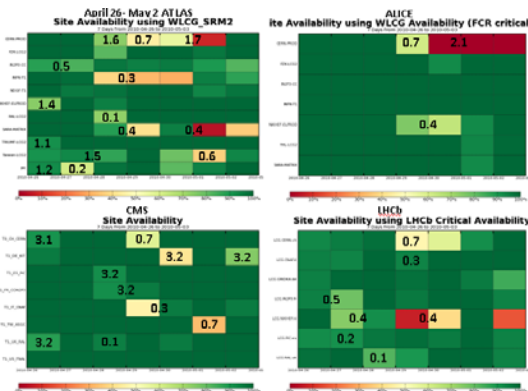
~ ~ ~

WLCG Management Board, 11<sup>th</sup> May 2010

# WLCG Operations Report – Summary

KPI	Status	Comment
GGUS tickets	<b>5 “real” alarms</b> , 1 test alarm; normal # team and user tickets.	<b>Drill-down on real alarms;</b> <b>comment on tests.</b>
Site Usability	<b>Minor issues</b>	<b>Drill-down (hidden)</b>
SIRs & Change assessments	Four incident reports	Drill-down summary.

VO	User	Team	Alarm	Total
ALICE	10	0	0	10
ATLAS	32	200	2	234
CMS	7	4	1	12
LHCb	1	27	4	32
Totals	50	231	7	288



# Alarm tickets

- There were test ALARM tickets with Triumph, the last action from the GGUS Release of April 14<sup>th</sup> to solve the hiccup with a beeper not being activated.
- To address the problem of sudden disappearance of 'Emergency email address' values from GOCDB a meeting was decided by the Tier1 coordination community. Agenda  
<http://indico.cern.ch/conferenceDisplay.py?confId=93347>
- Drills of real ALARMS since last MB follow.

# (Real) Alarm Summary

Ticket	VO	Site	Status	Date	Solved in	Info
57560	CMS	CERN	Verified	22-04	☺.8 hours	Tape migration delays
57771	ATLAS	CERN	Verified	28-04	☺ 3 hours	FTS failing for files GT 4GB
57524	ATLAS	TRIUMF	Verified	21-04		Test alarm
57665	LHCb	KIT	Verified	26-04	☺.3 hours	Space Token Full
57996	LHCb	CNAF	Solved	5-05	☺ 2 hours	LHCb STORM backend down
58063	LHCb	PIC	In Progress	08-05	?	Dirac mysql monitor down. How critical in question.
58065	LHCb	PIC	Solved	08-05	☺ 1 hour	Closed as duplicate of above

# CMS ALARM -> TO - CASTOR

What time	What happened
2010/04/22 11:37	GGUS ALARM ticket opened, automatic email notification to cms-operator-alarm@cern.ch AND automatic assignment to CERN-ROC AND automatic PRMS ticket creation.
2010/04/22 11:45	CERN-ROC re-classifies the PRMS ticket to CASTOR
2010/04/22 12:27	CASTOR team solves the problem. Tape library hardware intervention took more time than planned.

- [https://gus.fzk.de/ws/ticket\\_info.php?ticket=57560](https://gus.fzk.de/ws/ticket_info.php?ticket=57560)

# ATLAS ALARM->TO-CASTOR

What time	What happened
2010/04/28 16:25	GGUS ALARM ticket opened, automatic email notification to atlass-operator-alarm@cern.ch AND automatic assignment to CERN-ROC AND automatic PRMS ticket creation.
2010/04/28 16:43	CERN operator informs piquet
2010/04/28 19:51	CASTOR expert provides work-around. Checksumming turned off for files larger than 4GB.

- [https://gus.fzk.de/ws/ticket\\_info.php?ticket=57771](https://gus.fzk.de/ws/ticket_info.php?ticket=57771)

# LHCB ALARM->FZK-SPACE

What time	What happened
2010/04/26 13:00	GGUS ALARM ticket opened, automatic email notification to de-kit-alarm@scc.kit.edu AND automatic assignment to ROC-DECH
2010/04/26 13:13	Site records in the ticket a misuse of the ALARM concept.
2010/04/26 13:15	Site gives space and closes ticket.

- [https://gus.fzk.de/ws/ticket\\_info.php?ticket=57665](https://gus.fzk.de/ws/ticket_info.php?ticket=57665)

# LHCB ALARM > CNAF-STORM

What time	What happened
2010/05/05 18:06	GGUS ALARM ticket opened, automatic email notification to t1-alarms@cnafe.infn.it
2010/05/05 19:38	Site marks ticket in progress. Working on LHCb STORM backend (gpfs).
2010/05/05 20:15	Site reports gpfs problem fixed and closes ticket.

• [https://gus.fzk.de/ws/ticket\\_info.php?ticket=57996](https://gus.fzk.de/ws/ticket_info.php?ticket=57996)



# LHCB ALARM->PIC-DIRAC

What time	What happened
2010/05/08 03:54	GGUS ALARM ticket opened, automatic email notification to tier1-alarms@pic.es AND automatic assignment to ROC_SW
2010/05/08 04:09	Site acknowledges and starts work on the problem.
2010/05/08 04:24	Site queries criticality of Dirac monitor and puts ticket 'solved' informing will be left till Monday.
2010/05/08 04:27	Re-opened by the submitter – LHCb running blind at PIC. Still 'in progress' on 2010/05/10.

• !

# LHCB ALARM->PIC-DIRAC

What time	What happened
2010/05/08 05:56	GGUS ALARM ticket opened, automatic email notification to tier1-alarms@pic.es AND automatic assignment to ROC_SW
2010/05/08 06:52	Site closes ticket as duplicate of <a href="https://gus.fzk.de/ws/ticket_info.php?ticket=58063">https://gus.fzk.de/ws/ticket_info.php?ticket=58063</a> But marks it 'solved'!?

- [https://gus.fzk.de/ws/ticket\\_info.php?ticket=58065](https://gus.fzk.de/ws/ticket_info.php?ticket=58065)

# SIRs Received

- **CNAF** 28 and 29 April: 9 hours & 12 hours STORM SRM blockage (hardware) followed by separate MCDISK full and STORM bug  
<https://twiki.cern.ch/twiki/pub/LCG/WLCGServiceIncidents/SIR-CNAF--AtlasSRMoutage-April-2010.pdf>
- **IN2P3** 26 Apr: 17.5 hours AFS Distributed File System (AFS) crashed after server overload. Batch also affected.  
<https://twiki.cern.ch/twiki/pub/LCG/WLCGServiceIncidents/SIR-IN2P3-CC-AFSoutage-2010-04-26.pdf>
- **IN2P3** 24 Apr: 17 hours Batch services location service stopped responding to requests blocking most batch system commands  
<https://twiki.cern.ch/twiki/pub/LCG/WLCGServiceIncidents/SIR-IN2P3-CC-BatchOutage-2010-04-24.pdf>
- **IN2P3** 20 Apr: 9 hours & 5 days Grid downtime notifications were impossible after two consecutive incidents  
<https://twiki.cern.ch/twiki/pub/LCG/WLCGServiceIncidents/SIR-IN2P3-CC-OperationsPortal-2010-04-22v2.pdf>

# Other Service Issues (1 of 2)

- FTS bugs escalated by ATLAS on 4 May:
  - Long standing bug of FTS erroneously deleting source files after failed transfers and now many transfers becoming stuck and needing manual clearance was escalated by ATLAS to the daily meeting as seriously disrupting their data processing. A case-sensitivity bug in checksumming was also reported. Follow-up between ES and GT revealed fixes for all three were available and GT group quickly froze a new release and entered it into accelerated certification.
- WMS user limit hit by CMS on 29 April:
  - The CERN CMS WMS hit a hard-coded limit in the WMS Condor back-end of 100 different concurrent users with unfinished jobs, making the back-end continually exit and restart. The limit is removed in a later Condor version and this was attempted on Sunday but gave other problems. Another WMS was already being used with the updated Condor back-end to stress-test the fix for another issue and one CMS WMS was installed with this. This worked well and the remaining three have been re-installed.

# Other Service Issues (2 of 2)

- CNAF ATLAS jobs hitting vmem limit: started with 2GB physical memory then ramped up to 4GB on vmem, combined physical+swap and this is now being reached, they thought in a transient stage during end of job processing. Have temporarily removed limit. It was agreed this needs to be investigated by ATLAS.
- RAL failed ATLAS diskserver: had a faulty disk and card on 28 May followed by long weekend. They were surprised to be blacklisted by ATLAS due to this single diskserver failure. ATLAS agreed this was not usual but there had been a huge number of errors and they had no time to discuss this internally at the weekend.
- NL-T1 dcache access problems: dcache transfers hanging over the weekend of 1/2 May were due to a faulty port on a 4-port switch.

# Summary

- 5 alarms, one being questioned as to criticality.
- Alarm mechanism: test showed a new problem of an emergency email address disappearing from Gocdb. Meeting held covering the full alarm chain.
- 4 SIR received for this period – two more are in the pipeline (CNAF: ATLAS disk full STORM bug and STORM TSM problems).
- Many real problems at many sites needing serious effort to resolve but the effort was there, timely (within limits) and successful.
- Will continue to follow-up on these and other longer term issues at WLCG T1SCM