

Table of Contents

2007.....	1
Week 2007-12-10: Antonio Retico.....	1
Week 2007-11-19: Remi Mollon.....	1
Week 2007-11-05: David Collados.....	2
Week 2007-10-22: Konstantin Skaburskas.....	3
Week 2007-10-15: Antonio Retico.....	3
Week 2007-10-08: Piotr Nyczyk.....	3
Week 2007-10-01: Judit Novak.....	3
Week 2007-09-24: Rolandas Naujikas.....	3
Week 2007-09-17: Remi Mollon.....	4
Week 2007-09-10: Yvan Calas.....	4
Week 2007-09-03: David Collados.....	5
Week 2007-08-27: Konstantin Skaburskas.....	6
Week 2007-08-20: Steve Traylen.....	7
Week 2007-08-13: David Collados.....	7
Week 2007-08-05: Antonio Retico.....	11
Week 2007-07-30: Piotr Nyczyk.....	11
Week 2007-07-23: Yvan Calas.....	11
Week 2007-07-16: Steve Traylen.....	12
Week 2007-07-09: Remi Mollon.....	13
Week 2007-07-02: David Collados (Mon-Tue) & Judit Novak (Wed-Sun).....	13
Week 2007-06-25: Rolandas Naujikas.....	14
Week 2007-06-18: Yvan Calas.....	14
Week 2007-06-11: Diana Bosio.....	15
Week 2007-06-04: Konstantin Skaburskas, Remi Mollon.....	16
Week 2007-05-28: Remi Mollon.....	16
Week 2007-05-21: Piotr Nyczyk.....	17
Week 2007-05-14: Antonio Retico.....	17
Week 2007-05-07: David Collados.....	17
Week 2007-04-23: Yvan Calas.....	18
Week 2007-04-16: Diana Bosio.....	19
Week 2007-04-09: David Collados.....	20
Week 2007-04-02: Steve Traylen, Konstantin Skaburskas.....	21
Week 2007-03-26: Steve Traylen, Konstantin Skaburskas.....	21
Week 2007-03-19: Rolandas Naujikas.....	21
Week 2007-03-12: Maite Barroso.....	21
Week 2007-03-05: Remi Mollon.....	22
Week 2007-02-26: Judit Novak.....	22
Week 2007-02-19: Piotr Nyczyk.....	23
Week 2007-02-12: Antonio Retico.....	24
Week 2007-02-05: David Collados.....	25
Week 2007-01-29: Maria Dimou (Friday: Ian Neilson).....	25
Week 2007-01-22: Konstantin Skaburskas.....	26
Week 2007-01-15: Rolandas Naujikas.....	27
Week 2007-01-08 : Steve Traylen.....	27
2006.....	28
Days 2006-12-31 to 2007-01-07: Robert Harakaly.....	28
Days 2006-12-28 to 2006-12-30: Konstantin Skaburskas.....	30
Days 2006-12-22 to 2006-12-27: Yvan Calas.....	31
Week 2006-12-18: David Collados (Maite as backup).....	32
Week 2006-12-11: Antonio Retico.....	33
Week 2006-12-04: Yvan Calas.....	34
Week 2006-11-27: Maite Barroso.....	35

Table of Contents

2006

Week 2006-11-20: Piotr Nyczyk.....	36
Week 2006-11-13: Maria Dimou.....	36
Week 2006-11-06 : Naujikas Rolandas.....	37
Week 2006-10-30 : Skaburskas Konstantin.....	38
Week 2006-10-23 : Yvan Calas.....	39
Week 2006-10-16 : Antonio Retico.....	41
Week 2006-10-09 : Diana Bosio.....	42
Week 2006-10-02 : Judit Novak.....	43
Week 2006-09-25 : Antonio Retico.....	43
Week 2006-09-18 : Ian Neilson.....	44
Week 2006-09-11 : Maite Barroso.....	45
Week 2006-09-04 : David Collados.....	45

2007

Week 2007-12-10: Antonio Retico

- Friday 14/12
 - ◆ Morning meeting :
 - ◇ Nothing to report
 - ◆ Broadcasts :
 - ◇ CERN SAM 'dteam' submissions will be stopped 17th Dec
- Friday 13/12
 - ◆ Morning meeting :
 - ◇ Nothing to report
 - ◆ Broadcasts :
 - ◇ none
- Wednesday 12/12
 - ◆ Morning meeting :
 - ◇ Remark by Manuel about tickets from ROC not being displayed on the screen
 - ◇ Color codes not used in ROC reports
 - ◆ Broadcasts :
 - ◇ none
- Tuesday 11/12
 - ◆ Morning meeting :
 - ◇ Nothing to report
 - ◆ Broadcasts :
 - ◇ none
- Monday 10/12
 - ◆ Morning meeting :
 - ◇ Nothing to report
 - ◆ Broadcasts :
 - ◇ none

Week 2007-11-19: Remi Mollon

- Friday 23/11
 - ◆ Morning meeting :
 - ◇ ce109: OPERATING_SYSTEM - Rebooted, filesystems checked, back in production
 - ◇ cd112: GRID_BDII_WRONG - Checking
 - ◆ Broadcasts :
 - ◇ *"LFC-writefile test not critical"*
- Thursday 22/11
 - ◆ Morning meeting :
 - ◆ Broadcasts :
 - ◇ *"Problems with CMS LCG RB rb128.cern.ch"*
- Wednesday 21/11
 - ◆ Morning meeting :
 - ◇ lcgui003: NO_CONTACT
 - ◆ Broadcasts :
 - ◇ *"SAM fake alarms"*
 - ◇ *"SAM is back"*
- Tuesday 20/11
 - ◆ Morning meeting :
 - ◇ wms101: HWSCAN_WRONG - Raid rebuilt
 - ◇ rb128: DGLOGD_HIGH

◇ ce107: NO_CONTACT

◆ Broadcasts :

• Monday 19/11

◆ Morning meeting :

◇ wms101: RAID_TW_CTLR

◇ bdii109: NO_CONTACT - EXTFS_WARNING - NMI_RECEIVED

◆ Broadcasts :

◇ *"How to get rid of the voms servers" host certificates"*

Week 2007-11-05: David Collados

• Friday 09/11

◆ Broadcasts:

◇ URGENT: gLite3.0 Update35: fix for known issue in accounting was missing in release notes

Dear Site Administrator,

we realised that a known issue observed in certification dealing with the APE

The issue causes the accounting records not to be published, so it should be

The manual fix is the following:

After reconfiguration on the MON box, the file

/opt/glite/etc/glite-apel-publisher/publisher-config-yaim.xml

will contain a typo

- inspectTable should be

inspectTables.

Please edit the file appropriately.

A fix (to the template file in the glite-apel-publisher rpm) will be made ava

The release notes have been updated accordingly.

We apologise for this incident.

Regards,

The gLite Middleware release team.

• Thursday 08/11

◆ Broadcasts:

◇ Upgrade of CASTORPUBLIC service at CERN from 07:30h UTC, November 14th

Dear all,

Downtime will begin at: 07:30h UTC, November 14th (08:30h Swiss local time).

Downtime will end at: 11:00h UTC, November 14th (12:00h Swiss local time).

Services affected:

CASTORPUBLIC

VOs affected:

OPS

DTeam

GmodJournal < LCG < TWiki

We will upgrade the CASTORPUBLIC service to the latest version of the Castor.
During the intervention no requests to the CASTORPUBLIC service will be handled.

Sorry for the inconveniences,
Grid Manager on Duty

- Tuesday 06/11

- ◆ Broadcasts:

- ◇ Scheduled downtime of SAM server at CERN from 11:30h UTC, November 7th

Dear all,

In order to apply software upgrades, the SAM server node will be rebooted.

Downtime will begin at: 11:30h UTC, November 7th (12:30h Swiss local time).

Downtime will end at: 11:40h UTC, November 7th (12:40h Swiss local time).

Services affected:

- 1) SAM Portal
- 2) SAM Programmatic Interface
- 3) SAM Publishing web service (there might be a few test results that will be affected)

Sorry for the inconveniences.
SAM Team

- Monday 05/11

- ◆ Morning Meeting:

- ◇ wms102: renewd_wrong - Machine set to maintenance

Week 2007-10-22: Konstantin Skaburskas

- Monday 22/10

- ◆ Morning meeting:

- ◇ "Remedy" was down.

- ◆ TODO: keep an eye on possible consequences of the re-boot of the network switch (9.30) see broadcast sent last Friday (CERN-PROD - intervention on a network switch: several grid services affected) - DONE. No problems were noticed.

Week 2007-10-15: Antonio Retico

Week 2007-10-08: Piotr Nyczyk

- Monday 08/10

- ◆ TODO: broadcast: "Upgrade of SAM DB on 2007-10-08" (see request in it-dep-gd-gmod@cernNOSPAMPLEASE.ch mailing list)

Week 2007-10-01: Judit Novak

- Monday 01/10

- ◆ Morning meeting :

Week 2007-09-24: Rolandas Naujikas

- Monday 24/09

- ◆ Morning meeting :

Week 2007-09-17: Remi Mollon

- Friday 21/09
 - ◆ Morning meeting :
 - ◆ Broadcasts :
 - ◇ *"Scheduled intervention on CASTORALICE MSS at CERN on Wed, Sept 26 2007 at 9.00am CEST"*
- Thursday 20/09
 - ◆ Morning meeting :
 - ◆ Broadcasts :
- Wednesday 19/09
 - ◆ Morning meeting :
 - ◆ Broadcasts :
 - ◇ *"Alice WMS rb116.cern.ch removed from production on 25 September 2007"*
 - ◇ *"Atlas WMS rb101.cern.ch and rb110.cern.ch and LB rb201 removed from production on 25 September 2007"*
- Tuesday 18/09
 - ◆ Morning meeting :
 - ◇ lfc102: LFC_DB_ERROR
 - ◆ Broadcasts :
- Monday 17/09
 - ◆ Morning meeting :
 - ◇ ce106: GRID_BDII_WRONG
 - ◇ ce107: GRID_BDII_WRONG, HIGH_LOAD
 - ◇ rb107: SMART_SELFTEST
 - ◆ Broadcasts :
 - ◇ *"Scheduled network intervention at CERN on Thu, September 20th between 12.00am and 1.00pm CEST"*
 - ◇ *"New version on the PRODUCTION CERN AFS UI: 3.1.0-1 (UPDATE 03)"*

Week 2007-09-10: Yvan Calas

- Friday 14/09
 - ◆ 9:00 Daily morning meeting:
 - ◇ grvw001: *var_full*. Service manager contacted.
 - ◆ PRMS tickets: xx
 - ◆ Broadcasts:
 - ◇ Scheduled downtime of LHCb and LFC downstream database on Wednesday 19 September.
 - ◇ Job submission blocked on CMS WMS wms102.cern.ch on Friday 14 September.
- Thursday 13/09
 - ◆ 9:00 Daily morning meeting
 - ◇ wms102: *wm_wrong* still under investigation.
 - ◇ rb122: *wm_wrong*. High job submission rate on this machine this week (in general on CMS nodes).
 - ◇ lb102: *bk_serverd_wrong*. Probably service restarted by the developer.
 - ◇ lxdpm101: *var_full*.
 - ◇ monb002: *high_load*.
 - ◆ Broadcasts: none
- Wednesday 12/09
 - ◆ 9:00 Daily morning meeting

- ◊ wms102: *wm_wrong*. There is a problem with the workload manager service on this node. This is under investigation by the developers.
 - ◊ Many alarms on ce110, ce123.
 - ◆ CERN-PROD tickets:
 - http://egee-docs.web.cern.ch/egee-docs/ROC_CERN/CERN-PROD_tickets/CERN-tickets-20070911.x
 - ◆ 10:00 WLCG-SC meeting: **no meeting this week.**
 - ◆ 15:00 CCSR Meeting: [minutes](#):
 - ◆ Broadcasts:
 - ◊ Scheduled intervention on vomms.cern.ch.
-
- Tuesday 11/09
 - ◆ 9:00 Daily morning meeting:
 - ◊ rb109: *raid_tw* alarm.
 - ◊ fts106: *fta_wrong* alarms. Service manager contacted.
 - ◊ monb002: *spma_error* alarm. To be checked.
 - ◆ Broadcasts : none
 - Monday 10/09
 - ◆ 9:00 Daily morning meeting
 - ◊ rb105 and rb109: *smart_selftest* alarm.
 - ◊ lfc101 and lfc102: *lfc_db_error* alarm. Service manager contacted.
 - ◊ Many alarms on ce101, ce117 and ce119.
 - ◊ fts106: *fta_wrong* alarms. Service manager contacted.
 - ◆ WLCG/EGEE/OSG operations meeting: [see minutes here](#).
 - ◆ Broadcasts: none

Week 2007-09-03: David Collados

- Friday 07/09
 - ◆ Morning Meeting:
 - ◊ lfc103: *smart_selftest* (vendor call)
 - ◊ rb102: *interlogd_wrong*
- Thursday 06/09
 - ◆ Broadcasts:
 - ◊ Scheduled intervention of SAM at CERN from 7h UTC, Friday 6 September

Dear all,

Next Friday September 6th at 7h UTC (9h Geneva time) we will replace the scri

This intervention should be transparent to all users.

Best regards,
GMOD @ CERN

- Wednesday 05/09
 - ◆ Morning Meeting:
 - ◊ lb102: *hwscan_wrong*
 - ◊ lfc101: *lfc_db_error*. Daemon restarted.
 - ◊ rb201: *no_contact*. Machine unreachable. Rebooted.
- Tuesday 04/09
 - ◆ Morning Meeting:
 - ◊ wms101: *raid_tw*
 - ◊ rb109: *smart_failing*
 - ◆ Broadcasts:
 - ◊ Service degradation of CASTOR at CERN from 06:00 UTC, Wednesday 5 Sept

GmodJournal < LCG < TWiki

Dear all,

This is a reminder of the Castor tape libraries intervention scheduled for to

In order to equalise the load distribution on our CASTOR tape libraries, we h

We will do this move on:

Wednesday, 5 September 2007, from 6:00 until 15:00 (UTC)

Wednesday, 5 September 2007, from 8:00 until 17:00 (Geneva local time)

Users needing to read data from the tapes in those robots will experience lon

You can determine whether your data might be affected by the intervention by

```
# nsls -T /castor/cern.ch/user/m/myuserid/mycastorfile
```

```
- 1 1 I05684 62 00002c8a 49858560 100 /castor/cern.ch/user/m/myuserid/mycasto
```

If the 3rd column of the command output shows a tape volume starting with I (

Best regards,

GMOD

- Monday 03/09

- ◆ Morning Meeting:

- ◇ ce109: swap_full - Transient alarm

- ◇ ce114: vm_kill - gatekeeper restart, machine OK (Piquet call)

- ◇ lfc101, lfc102, lfc111: lfc_db_error

- ◇ rb102: interlogd_wrong, rgma_wrong - vendor call

- ◇ rb109: raid_tw - vendor call

- ◇ rb113: wm_wrong

- ◇ voms101: ipmi_power - vendor call

Week 2007-08-27: Konstantin Skaburskas

- Friday 31/08

- ◆ Alarms: wms104 (wm_wrong. According to the procedure : SMS state changed to maintenance (waiting for response...#wms104: OK). Mail sent to WMS.Support@cernNOSPAMPLEASE.ch.).

- ◆ Weekly CERN-PROD RC report

- ◇ Checked. No downtime for the site. Only SAM downtime for 11h starting 23.00 Thursday till Friday morning. Added to report.

- Thursday 30/08

- ◆ nothing to report

- Wednesday 29/08

- ◆ Alarms: wms104 (wmproxy_wrong. State changed to maintenance).

- ◆ Broadcasts:

- ◇ Abbreviations for Site-Names in GridView - CLARIFICATION.

- Tuesday 28/08

- ◆ Alarms: wms102/104 (wmproxy_wrong. SMS state changed to maintenance. Mail sent to WMS.Support@cernNOSPAMPLEASE.ch), rb108.

- ◆ Broadcasts:

- ◇ SAM/FCR server intervention will happen today (reminder).

- ◇ Aug 28 2007 at 16hrs CEST scheduled interruption of the CERN vomrs services (reminder).

- Monday 27/08
 - ◆ Alarms: wms102 (wmpoxy_wrong machine set to maintenance. Service manager contacted wms.support@cernNOSPAMPLEASE.ch), rb102 (interlogd_wrong. For a while.), rb107 (smart_selftest), fts108 (hwscan_wrong), fts114 (smart_selftest).
 - ◆ "LCG Node Status" page was not on-line for some time in the morning - seems like network glitch.
 - ◆ Corrected the link to the "LCG Node Status" page on GMOD description page. It was pointing to the old location.

Week 2007-08-20: Steve Traylen

- Friday 24/08
 - ◆ rb102 errors for seconds, message sent to service manager.
- Thursday 23/08
 - ◆ wms alerts raised due to re installation without putting into maintenance mode.
- Tuesday 21/08
 - ◆ **TODO:** Send a reminder about broadcast sent last Friday (Network and LSF batch interruptions at CERN from 06.00 CEST August 22)
- Monday 20/08

Week 2007-08-13: David Collados

- Friday 17/08
 - ◆ Morning Meeting:
 - ◇ Nothing to report.
 - ◆ Weekly CERN-PROD RC report
 - ◇ Checked. Only 1 downtime of 2h on Tuesday 14th already explained by Jan, so nothing to report from me.
 - ◆ Broadcasts:
 - ◇ Network and LSF batch interruptions at CERN from 06.00 CEST August 22

Dear all,

Between 06.00 and 08.00 CEST (04.00 to 06.00 UTC) on Wednesday 22 August, a t

Connectivity to the LCG Tier 1 sites via the optical private network will be be
castor etc.) will be disrupted.

CERN local and grid LSF batch jobs will be suspended from 05.30 and fresh job

This intervention will be followed by an intervention on the CASTORPUBLIC sta

Sorry for the inconvenience,
Grid Manager on Duty @ CERN

- Thursday 16/08
 - ◆ Morning Meeting:
 - ◇ ce106: high_load
 - ◇ lcgui005: spma_error
- Wednesday 15/08
 - ◆ Morning Meeting:
 - ◇ ce120: high_load
 - ◇ lfc102: lfc_noread
 - ◇ rb102: hwscan_wrong, rgma_wrong
 - ◇ lcgui: spma_error
 - ◆ WLCG SCM Meeting:

◇ Nothing to report.

◆ CCSR Meeting:

◇ There will be a CastorPublic intervention next Wednesday.

◆ Broadcasts:

◇ Scheduled downtime of CASTORLHCB instance at CERN from 07:00h, Thursday 16th August (UTC)

Dear all,

We would like to upgrade the CASTORLHCB instance next Thursday 16th.

Downtime will begin at: 07:00h, 16th August (UTC)

Downtime will end at: 09:00h, 16th August (UTC)

Services affected: srm.cern.ch

VOs affected: LHCB only.

A maintenance release of the Castor software will be installed, fixing several

During the intervention no new Castor requests will be handled. Existing conn

Sorry for the inconvenience,

Grid Manager on Duty @ CERN

◇ Last update on unscheduled downtime of several links between CERN and Tier1 sites from 13:50h (UTC) Monday 13th August

Dear all,

The problem described yesterday on several links in the LHC OPN between CERN

Service degradation began at: 17:00h, Sunday 12th of August (UTC)

Service degradation ended at: 06:30h, Wednesday 15th of August (UTC)

for all Tier1s except ASGC.

For ASGC, service degradation ended around: 10:00h, Wednesday 15th of August

Services affected:

- CERN link to TRIUMF (primary and backup link down)

- CERN link to other Tier-1s affected (primary link down but working on backu

The problem was transparent to the T1s affected except for TRIUMF, as it had

Sorry for the inconvenience,

Grid Manager on Duty @ CERN

• Tuesday 14/08

◆ Morning Meeting:

◇ Nothing to report

◆ Broadcasts:

◇ Scheduled downtime of four CEs at CERN from 16:04h, Monday 13th August (UTC)

Dear all,

There is a scheduled downtime on four CERN LCG CEs.

Downtime began at: 16:04h, Monday 13th of August (UTC)

Downtime will end at: 16:04h, Monday 27th of August (UTC)

Services affected:

cel15.cern.ch,

cel16.cern.ch,

GmodJournal < LCG < TWiki

cell17.cern.ch,
cell18.cern.ch

These services are being drained for moving from SLC3 to SLC4/64 submission h

Please DO NOT USE these services during the downtime period.

Thanks in advance,
Grid Manager on Duty @ CERN

◇ Update on unscheduled downtime of several links between CERN and Tier1 sites from 13:50h (UTC) Monday 13th August

Dear all,

As announced yesterday, there is a service degradation on several links in th

Downtime began at: 17:00h, Sunday 12th of August (UTC)

Downtime should end at: 12:30h, Tuesday 14th of August (UTC)

Services affected:

- CERN link to TRIUMF (primary and backup link down)
- CERN link to Tier-1s affected (primary link down but working on backup link)

The problem is transparent to the T1s affected (except for TRIUMF) as along a they don't exceed an aggregated bandwidth of ~10G, that is the maximum capaci

Alcatel confirmed that the failure is due to a failed amplifier board on the

A new broadcast will be sent when the problem is fixed.

Sorry for the inconvenience,
Grid Manager on Duty @ CERN

• Monday 13/08

◆ Morning Meeting:

- ◇ ce123 - swap_full
- ◇ rb101 - smart_selftest

◆ Broadcasts:

- ◇ Reminder: Short network interruptions at CERN between 06.00-06.30 UTC 14 August

Between 06.00 and 06.30 UTC (08.00 to 08.30 CET) on Tuesday 14 August a CERN router providing the primary OPN network link to ASGC, BNL, TRIUMF, FZK, PIC and RAL will be rebooted for a software upgrade lasting 5-10 minutes. All sites will be switched to their backup paths except RAL which does not have one yet. At the same time a computer centre router linking the central services (DNS, lxplus, lxbatch, nice) will be upgraded with an expected 5 minute downtime. Grid and local batch jobs at CERN will be suspended during these interventions.

The Grid Manager on Duty @ CERN-PROD

◇ Scheduled downtimes of the CERN link to ASGC on Wednesday, 15th August

Dear all,

On Wednesday the 15th of August 2007 there will be two maintenance interventions

1st downtime will begin at 08:00h, Wednesday 15th August (UTC)
1st downtime will end at 09:00h, Wednesday 15th August (UTC)

2nd downtime will begin at 14:00h, Wednesday 15th August (UTC)
2nd downtime will end at 15:00h, Wednesday 15th August (UTC)

Services affected:

The intervention will concern the routing to ASGC. During the interventions there will be a few minutes of downtime for a few minutes. All the other Tier1 sites won't be affected.

Sorry for the inconvenience,
Grid Manager on Duty @ CERN

◇ UNSCHEDULED DOWNTIME of several links between CERN and Tier1 sites from 13:50h (UTC) Monday 13th August

Dear all,

Several links in the LHC Optical Private Network are down at this very moment.

Sorry for the inconvenience,

Grid Manager on Duty @ CERN

◇ Scheduled downtime of CASTORCMS instance at CERN from 07:00h, Wednesday 15th August (UTC)

Dear all,

We would like to upgrade the CASTORCMS instance next Wednesday 15th.

Downtime will begin at: 07:00h, 15th August (UTC)

Downtime will end at: 09:00h, 15th August (UTC)

Services affected: srm.cern.ch

VOs affected: CMS only.

A maintenance release of the Castor software will be installed, fixing several bugs.

During the intervention no new Castor requests will be handled. Existing connections will be handled.

Sorry for the inconvenience,
Grid Manager on Duty @ CERN

◇ Scheduled transparent intervention of the LCGR production database at CERN from 07:00h, Wednesday 15th August (UTC)

Dear WLCG,

Following the successful deployment of the 'CPU July 2007' patch on the integration database,

As proposed on 3rd August IT morning meeting, we plan to apply the patch as a rolling upgrade.

Wednesday August 15th from 07:00 UTC to 09:00 UTC (9am to 11am CEST).

As the upgrade is rolling the service will be continuously available during the intervention.

Following our security and patching policy, the intervention on production is planned to be transparent.

The services running on LCGR are:

lcg_fcr, lcg_fts, lcg_fts_t2, lcg_fts_t2_w, lcg_fts_w, lcg_gridview, lcg_lfc,

For additional information you can contact Miguel.Anjo@cern.ch

Best regards,
Grid Manager on Duty @ CERN

◆ WLCG/EGEE/OSG Operations Meeting

◇ Nothing to report

Week 2007-08-05: Antonio Retico

- Friday 10/08
 - ◆ Morning Meeting
- Thursday 9/08
 - ◆ Morning Meeting
 - ◇ wms_wrong wms104 : remedy ticket
 - ◇ network intervention

On Tuesday 14th of August between 8.00 AM and 8.30AM CET one of the two CERN LHCOPN router will be rebooted for a software upgrade. The router will be down for 5-10 minutes.

Impact:

The primary link of ASGC, BNL, TRIUMF, GRIDKA, RAL, PIC will be down. The traffic for these sites will be re-routed to their backup paths; RAL will be unreachable since it has no backup connectivity.

- ◆ Broadcasts: *Short network interruptions at CERN between 06.00-06.30 UTC 14 August
- Wednesday 8/08
 - ◆ Morning Meeting
 - ◇ NTR
- Tuesday 7/08
 - ◆ Morning Meeting
 - ◇ NTR
- Monday 6/08
 - ◆ Morning Meeting
 - ◇ NTR

Week 2007-07-30: Piotr Nyczyk

gstat stopped monitoring CERN-PROD and some other sites during the week-end results disappeared from SLS. This is now an issue for service availability because Gridview did not seem to suffer. Under investigation with gstat

Week 2007-07-23: Yvan Calas

- Friday 27/06
 - ◆ 9:00 Daily morning meeting: nothing to report.
 - ◆ PRMS tickets: 0
 - ◆ Broadcasts:
 - ◇ LCG RB rb105.cern.ch dedicted to Alice back in production.
- Thursday 26/07
 - ◆ 9:00 Daily morning meeting
 - ◇ rb122: *dglogd_high* because of high job submission rate. Ask CMS to slow down this rate. Problem fixed in the afternoon.
 - ◆ PRMS Tickets: 0
 - ◆ Broadcasts:
 - ◇ 2 new SLC4/64 submission CEs in production (ce113 and ce114).
- Wednesday 25/07
 - ◆ 9:00 Daily morning meeting

- ◇ wm104: *wm_wrong* because of a bug in the middleware.
- ◇ fts110, fts111, fts112 and fts113: *spma_error*. Service manager contacted.
- ◇ ce113 and ce114: *grid_bdii_wrong* error due to a mistake made by the service manager.
- ◇ rb106: *wm_wrong* due to a bug in the proxy renewal service. This problem also affects rb121 (these two machines are Atlas nodes).
- ◇ Atlas database intervention tomorrow in Taiwan.
- ◆ CERN-PROD tickets:
 - http://egee-docs.web.cern.ch/egee-docs/ROC_CERN%5CCERN-PROD_tickets%5CCERN-tickets-20
- ◆ 10:00 WLCG-SC meeting: **no meeting this week.**
- ◆ 15:00 CCSR Meeting: minutes [☞](#):
- ◆ PRMS Tickets: 3
 - ◇ #445707: *wm_wrong* alarm on rb106. Bug in the middleware. Fixed.
 - ◇ #445958: *dglogd_high* alarm on rb122 caused by high jobs submission rate. CMS contacted and the backlog disappeared thereafter.
 - ◇ #445818: rb problem(rb107,rb119,rb122). Due to the high jobs submission rate.
- ◆ Broadcasts:
 - ◇ Castoralice upgrade POSTPONED again.
 - ◇ LCG RB rb105.cern.ch dedicted to Alice removed from production.
- Tuesday 24/07
 - ◆ 9:00 Daily morning meeting:
 - ◇ rb122: *high_load* for few minutes. CMS users are highly using their RBs, nothing to do.
 - ◇ rb113, rb115 and rb120: *ntp_wrong* alarm due to the upgrade of the ntp package (service restarted and node not in maintenance during the upgrade).
 - ◇ ce111: *gridftp_wrong* for less than 15 mn.
 - ◇ ce122: *high_load*, *swap_full* and *no_contact* alarms. Unable to connect. Machine rebooted by operator.
 - ◆ PRMS Tickets: 2
 - ◇ #445561: Problem with the workload manager on wms104. Bug in the middleware. This machine will remain in maintenance for the time being.
 - ◇ #445606 (GGUS ticket #25100): unable to download output sandbox from wms101. Under investigation.
 - ◇ Broken old CERN CA CRL file on AFS UI. Mail sent directly hep-project-grid-cern-testbed-managers and support-eis@cernNOSPAMPLEASE.ch...
 - ◆ Broadcasts :
 - ◇ Castoralice upgrade on on Wednesday, July 25 2007.
 - ◇ Castoralice upgraded POSTPONED to 26 July 2007.
- Monday 23/07
 - ◆ 9:00 Daily morning meeting
 - ◇ lb101: *bkserved_wrong* for few minutes. Service restarted automatically.
 - ◇ Many alarms on some CE nodes.
 - ◆ PRMS Tickets: 1
 - ◇ #445193: Problem with the bkserved service on lb101. I modified the alarm parameters to avoid this problem in the future.
 - ◆ WLCG/EGEE/OSG operations meeting: see minutes here [☞](#).

Week 2007-07-16: SteveTraylen

- Monday 2007-07-16
 - ◆ Morning Meeting

◇ wms101, various alarms, Yvan looks to have intervened.

- Tuesday 2007-07-17
 - ◆ The wms rb111 (Alice), rb125 (CMS) and rb126 (Atlas) will be renamed as wms103, wms104 and wms105, and move to cluster/subcluster grid/gridwms. The middleware installed is the 3.1 one with the latest patches available.". FYI: a broadcast has been sent to CMS by Andrea, no broadcast were needed for Atlas and Alice. I hope that these nodes will be reay tonight, I will let you know the broadcast to send to these 3 VOs. Yvan
 - ◆ It is time to migrate more CEs to SLC4. I will start to drain ce113 and ce114 for this purpose, and schedule a one week down time for these nodes. Ulrich.
- Wednesday 2007-07-18
 - ◆ Just two CERN-PROD tickets, both unexciting.
 - ◆ CSSR Meeting
 - ◇ An Oracle security patch is now being tested that must in time be deployed everywhere. It can be installed as a rolling update. A schedule will be produced next week.

Week 2007-07-09: Remi Mollon

- Friday 13/07
 - ◆ Morning Meeting
- Thursday 12/07
 - ◆ Morning Meeting
 - ◇ rb102: SPMA_ERROR
- Wednesday 11/07
 - ◆ Morning Meeting
 - ◇ ce101: NO_CONTACT
 - ◆ Broadcasts
 - ◇ *CERN PPS AFSUI gLite 3.1/update01*
- Tuesday 10/07
 - ◆ Morning Meeting
 - ◇ rb107: NO_CONTACT
 - ◆ Broadcasts
 - ◇ *Downgrade of prod-bdii at CERN*
- Monday 09/07
 - ◆ Morning Meeting
 - ◇ lfc101: LFC_DB_ERROR
 - ◇ lfc113: LFC_NOREAD
 - ◆ Broadcasts
 - ◇ *Scheduled network intervention at CERN on Wed, July 18th 7.00am UTC*

Week 2007-07-02: David Collados (Mon-Tue) & Judit Novak (Wed-Sun)

- Monday 02/07
 - ◆ 9:00 Daily morning meeting:
 - ◇ lfc113 lfc_noread.
 - ◇ rb122 high load.
 - ◇ rb123: no contact.
 - ◆ WLCG/EGEE/OSG operations meeting

Week 2007-06-25: Rolandas Naujikas

- Friday 29/06
- Thursday 28/06
 - ◆ Broadcast "SAM using GOCDB3" was send.
 - ◆ Broadcast "SAM problem and fix" was send.
- Wednesday 27/06
 - ◆ Planed CASTORALICE upgrade.
- Tuesday 26/06
- Monday 25/06
 - ◆ Remedy is unstable after upgrade - downgrade is in progress.
 - ◆ Broadcast "CASTORALICE upgrade on Wednesday, June 27" was send.

Week 2007-06-18: Yvan Calas

- Friday 22/06
 - ◆ 9:00 Daily morning meeting:
 - ◇ Castor CMS upgrade went fine.
 - ◇ Question related to the SAM failure last Wednesday (tomcat failure on the SAM server, Judit wrote a cron job to restart this service if needed, and the two new SAM servers will be quattor-managed, so alarms will be triggered and tomcat will be restarted automatically by the actuator).
 - ◆ PRMS tickets: 1
 - ◇ #437024: libtar problems.
- Thursday 21/06
 - ◆ 9:00 Daily morning meeting
 - ◇ rb115 back in production after a problem with one of the RAID disks.
 - ◇ rb125: *wmproxy_wrong* alarm. Fixed by service manager.
 - ◇ Castor CMS upgrade started this morning.
 - ◇ FTS intervention of yesterday went fine.
 - ◆ PRMS Tickets: 0
 - ◆ Broadcasts:
 - ◇ SAM was unavailable last night.
- Wednesday 20/06
 - ◆ 9:00 Daily morning meeting
 - ◇ rb125: *wmproxy_wrong* and *interlogd_wrong* alarms. Fixed by service manager.
 - ◇ ce106: *no_contact* and *+gridgris_status_error_* alarms.
 - ◇ rb107 and rb119: *high_load* alarms due to high job submission rate from CMS.
 - ◇ rb115: *raid_tw* alarm. Vendor call.
 - ◇ rb118: *no_contact* alarm, nothing found in the core dump. Machine rebooted by operator, some services did not restart automatically.
 - ◆ CERN-PROD tickets:
 - http://egee-docs.web.cern.ch/egee-docs/ROC_CERN%5CCERN-PROD_tickets%5CCERN-tickets-20
 - ◆ 10:00 WLCG-SC meeting: [minutes](#) ☞.
 - ◆ 15:00 CCSR Meeting: [minutes](#) ☞.
 - ◆ PRMS Tickets: 0
 - ◆ Broadcasts:

- ◇ OPN link between RAL-LCG2 and CERN down (broadcast sent by Derek Ross).
- ◇ CERN FTS Back in Service after Database DeFragmentation + Observed Problems (broadcast sent by Steve).
- ◇ Release of UPDATE 26 to gLite 3.0 (broadcast sent by Nick).

• Tuesday 19/06

- ◆ 9:00 Daily morning meeting:
 - ◇ ce102 and ce109: *gridftp_wrong* alarm. These nodes have been reinstalled from scratch yesterday evening.
 - ◇ lfc101 and lfc102: *lfc_db_error* alarm. Service restarted by service manager.
 - ◇ rb122: *high_load* alarm for few minutes. CMS is using intensively this node.
- ◆ PRMS Tickets: 1
 - ◇ #435913: srmv1 failure at SE dpm0001.m45.ihep.su (see ggus ticket #22538)
- ◆ Broadcasts :
 - ◇ FTS Service Intervention on prod-fts-ws.cern.ch on Jun-20-2007 (broadcast sent by Steve).

• Monday 18/06

- ◆ 9:00 Daily morning meeting
 - ◇ ce102 and ce109: *gridftp_wrong* alarm. To be checked by service managers.
 - ◇ lfc101 and lfc102: *lfc_db_error* and *lfcdaemon_wrong* alarms. To be checked by service managers.
 - ◇ rb102: *interlogd_wrong* alarm. Machine checked and back in production.
 - ◇ rb122: *high_load* alarm. High job submission rate for this machine dedicated to CMS.
 - ◇ LHCb Castor upgrade started this morning.
 - ◇ FTS service intervention started this morning.
- ◆ PRMS Tickets: 0
- ◆ Broadcasts :
 - ◇ Traffic interruption between BNL and CERN on Jun-19-2007.
 - ◇ CASTORCMS upgrade on Jun-21-2007.
 - ◇ Downtime of ce109 and ce112. Reinstallation from scratch of these two nodes.
- ◆ WLCG/EGEE/OSG operations meeting: see minutes here [↗](#).

Week 2007-06-11: Diana Bosio

• Friday - June 15th

- ◆ Morning meeting
 - ◇ rb106 (ticket 184377): *wm_wrong*. According to the procedure machine set in maintenance.
 - ◇ monb001 (several tickets): *high_load*.
- ◆ EGEE BROADCAST:
 - ◇ CERN-PROD: Scheduled downtime for CASTORLHCB - Monday June 18 from 07:00 UTC

• Thursday - June 14th

- ◆ Morning meeting
 - ◇ rb117 (several tickets): *swap_full*, *vm_kill*
 - ◇ monb001 (several tickets): *high_load*. * Wednesday - June 13th
- ◆ Morning meeting
 - ◇ rb102 (several tickets): *interlogd_wrong*, *logs_full*.
- ◆ EGEE BROADCAST:
 - ◇ " REMINDER: unavailability of CASTORATLAS service at CERN - today"

• Tuesday - June 12th

- ◆ Morning meeting
 - ◇ rb127 (ticket 183309): put to maintenance.

- ◊ rb107 (ticket 183499): bkserverd_wrong.
- ◊ rb201 (ticket 183495): xinetd_wrong.
- ◆ EGEE BROADCAST:
 - ◊ "CERN PROD: LFC catalogues will be upgraded today to 1.6.3"
 - ◊ " REMINDER: unavailability of CASTOR2 service at CERN - today"
 - ◊ "CERN-PROD: CASTORATLAS upgrade tomorrow Wednesday June 13th"
 - ◊ "CERN-PROD: Warning CASTOR intervention today June 12th affects SAM tests"
- Monday - June 11th
 - ◆ Morning meeting
 - ◊ rb113 has a no_contact alarm.
 - ◊ rb116 has a raid_tw alarm.
 - ◊ monb001 has a no contact alarm

Week 2007-06-04: Konstantin Skaburskas, Remi Mollon

- Friday - June 8th (Konstantin Skaburskas):
 - ◆ ...
- Thursday - June 7th (Konstantin Skaburskas):
 - ◆ Broadcasts:
 - ◊ "unavailability of CASTOR2 service at CERN - Tuesday 12 June"
 - ◊ "unavailability of ATLAS CASTOR services at CERN - Wednesday 13 June"
- Wednesday - June 6th (Konstantin Skaburskas):
 - ◆ ...
- Tuesday - June 5th (Remi Mollon):
 - ◆ ...
- Monday - June 4th (Konstantin Skaburskas):
 - ◆ Morning Meeting: Remi Mollon attended the morning meeting.
 - ◆ Broadcasts:
 - ◊ "occasional SAM unavailability"

Week 2007-05-28: Remi Mollon

- Friday - June 1st
 - ◆ Broadcasts:
 - ◊ "LCG RB rb107.cern.ch dedicated to CMS back in production" (by Yvan Calas)
 - ◊ "IMPORTANT: problem with SAM tests causing all sites to fail"
 - ◊ "Unscheduled Intervention on ce101.cern.ch"
 - ◆ Morning Meeting:
 - ◊ gridce problem: 8 new boxes added, but still high load problem!
- Thursday - May 31st:
 - ◆ Broadcast: "LCG RB rb107.cern.ch dedicated to CMS down"
- Wednesday - May 30th:
 - ◆ Morning Meeting:
 - ◊ High Load, GridGris_Status_Error and no_contact on several gridce nodes, to be investigated!
 - ◆ Broadcast: "CERN LCG CEs problem due to high activity"
- Tuesday - May 29th :
 - ◆ Morning Meeting :
 - ◊ high load on castradgrid nodes
- Monday - May 28th: Vacation

Week 2007-05-21: Piotr Nyczyk

- Thursday - May 24th:
 - ◆ REMINDER: lcg-voms.cern.ch certificate will be changed today!
- Tuesday - May 22nd
 - ◆ EGEE Broadcast:
 - ◇ lcg-voms.cern.ch certificate change broadcasted (David)
- Monday - May 21st
 - ◆ Morning Meeting:
 - ◇ high load on Atlas VOBOXes
 - ◇ gridbdii cluster - Maarten worked on machines without taking them offline which caused alarms for the OPS

Week 2007-05-14: Antonio Retico

- Friday - May 18th:
 - ◆ Morning Meeting:
- Thursday - May 17th:
 - ◆ VACATION!
- Wednesday - May 16th:
 - ◆ Morning Meeting:
 - ◇ Castor intervention in progress
 - ◇ Gavin suggested to give a summary of FTS, SRM tickets, instead of the full list to cern prod administrators.
 - ◆ EGEE Broadcast:
- Tuesday - May 15th:
 - ◆ Morning Meeting:
 - ◇ Jan Van Eldik complained for srm-v2.cern.ch receiving alarms.
 - ◇ I checked and indeed they wer publishing the node by themselves, so no problems in SAM. As ROC I turned the monitoring off for that service in order for the alarms to stop.
 - ◆ EGEE Broadcast: WARNING: lcg-voms.cern.ch certificate will be changed on May 24th!
 - ◇ I sent a news about this intervention to be inserted also in the service status board
- Monday - May 14th:
 - ◆ Morning Meeting: nothing to report
 - ◆ EGEE Broadcast: FTS test change (dteam -> ops)

FTS test change (dteam -> ops)

Week 2007-05-07: David Collados

- Monday - May 07th:
 - ◆ Morning Meeting:
 - ◇ Cooling incident on May 5th in computer room at 7.00. Re-established at 8.00. Several cpu and disk servers were switched off to avoid damage which meant services were slightly degraded during the morning. This affected different services like srm, Castor, RBs, etc.
 - ◇ gridbdii: raid_tw and smart_self test on bdii111. Vendor call opened.
 - ◇ gridfts: service managers contacted for fts110.
 - ◇ gridlfc: lfc201 to be checked.
 - ◇ gridrb: rb109 and rb115 affected due to cooling problem.
 - ◇ gridce: no_contact on ce114. Being checked.
 - ◇ gridmon: high_load on monb001.

- ◊ Scheduled interventions today:
 - Castor LHCb move to new oracle certified hardware between 09.00 and noon.
 - Decommission of the old HTMLDB/Apex web server - dbnetra04
- ◆ EGEE Broadcast: SAM test "host-cert" becomes critical, and OPS vo will submit FTS tests
- Tuesday - May 08th:
 - ◆ Morning Meeting:
 - ◊ gridbdii: smart_selftest in bdii111
 - ◊ gridfts: high_load on fts110. Service manager contacted and problem fixed.
 - ◊ gridrb: high_load on rb122.
 - ◆ EGEE Broadcast: SAM test "host-cert" not yet critical for SE/SRM services
- Wednesday - May 09th:
 - ◆ Morning Meeting:
 - ◊ gridce: GridGris_StatusError in ce101. Service Manager contacted.
 - ◊ srm: many high loads. Main problem in srm101 affecting other srm client machines.
 - ◊ gridmon: high_load in monb001 due to 'too many tomcat processes'. Service Manager contacted.
 - ◆ EGEE Broadcast: SRM endpoint srm.cern.ch severely degraded
 - ◆ EGEE Broadcast: Service on srm.cern.ch reopened at 09:00h UTC
 - ◆ EGEE Broadcast: gLite WMS dedicated to CMS (rb102.cern.ch) removed from production
- Thursday - May 10th:
 - ◆ Morning Meeting: nothing to report
- Friday - May 11th:
 - ◆ Morning Meeting:
 - ◊ gridce: high load, swap_full, vm_kill alarms on ce107
 - ◊ gridfts: high_load on fts114
 - ◊ gridlfc: lfc_noread on lfc113
 - ◊ gridmon: high_loads on monb991
 - ◆ EGEE Broadcast: CASTORPUBLIC intervention next Wednesday, May 16

Week 2007-04-23: Yvan Calas

- Friday 27/04
 - ◆ 9:00 Daily morning meeting:
 - ◊ A lot of tickets for WM_WRONG alarms on rb107, rb119 and rb122 machines (LCG RBs dedicated to CMS).
 - ◆ PRMS Tickets: 36
 - ◊ 36 GGUS tickets related to the workload manager crash on rb107, rb119 and rb122 machines (LCG RBs dedicated to CMS).
 - ◆ Broadcasts:
 - ◊ CMS LCG RBs rb107, rb119 and rb122 temporary removed from production.
- Thursday 26/04
 - ◆ 9:00 Daily morning meeting * alarm *var_full* on monb001. Fixed by service manager.
 - ◆ PRMS Tickets: 0
 - ◆ Broadcasts: none
- Wednesday 25/04
 - ◆ 9:00 Daily morning meeting
 - ◊ A lot of machines (172 log only tickets) went in NO_CONTACT (for a while) yesterday at 17:03. This is apparently not due to a network failure. LAS team contacted to know if something went wrong on their side.
 - ◊ Maintenance intervention on the primary path of the CERN firewall tomorrow (Thursday the 26th of April 2007 between 8:00AM and 18:00). Should be

transparent.

◊ Upgrade of the Alice database planned tomorrow.

◊ CERN-PROD tickets:

http://egee-docs.web.cern.ch/egee-docs/ROC_CERN\CERN-PROD_tickets\CERN-tickets-2007-04-16

- ◆ 10:00 WLCG-SC meeting: [minutes](#).
- ◆ 15:00 CCSR Meeting: [minutes](#): nothing to report.
- ◆ PRMS Tickets: 0
- ◆ Broadcasts:
 - ◊ Maintenance on the CERN firewall on Thursday 26 April 2007.
 - ◊ Intervention on Castor Alice on 2007-05-02.

- Tuesday 24/04

- ◆ 9:00 Daily morning meeting: nothing to report
- ◆ PRMS Tickets: 1
 - ◊ #421569: Access to the Grid.
- ◆ Broadcasts :
 - ◊ Broadcast sent by Yvan: rb109 back in production after hardware intervention.

- Monday 23/04

- ◆ 9:00 Daily morning meeting
 - ◊ bdii102, bdii102, bdii102, bdii102 and bdii102: several *high_load* alarms for few minutes during the week-end (Saturday 21 April).
 - ◊ monb002: alarm *no_contact* (last Sunday). Machine rebooted by sysadmins.
- ◆ PRMS Tickets: 3
 - ◊ #421198: problem during the middleware upgrade on rb112 (gLite WMS 3.1 for LHCb).
 - ◊ #421201: problem during the middleware upgrade on rb112 (gLite WMS 3.1 for LHCb).
 - ◊ #421316: problem during the middleware upgrade on rb112 (gLite WMS 3.1 for LHCb).
- ◆ Broadcasts :
 - ◊ Broadcast sent by Yvan: SAM history limit (max 14 days) -- cf mail sent by Judit.
- ◆ WLCG/EGEE/OSG operations meeting: see [minutes here](#).

Week 2007-04-16: Diana Bosio

- Friday - April 20th:

- ◆ Morning meeting
 - ◊ rb107 (ticket 174224+174231) has a *no_contact* alarm.
 - ◊ rb109 has a *raid_tw* alarm.
 - ◊ ce102 (ticket 174279)
 - ◊ ce106 (ticket 174456) *high_load*. Was rebooted.
 - ◊ The transparent intervention on the OPN routers went fine.
- ◆ EGEE BROADCAST: CERN-PROD: Unscheduled intervention: rb109 (CMS glite 3.0 WMS) will be put in maintenance this morning.

- Thursday - April 19th:

- ◆ Morning meeting
 - ◊ rb107 (ticket 174224+174231) has a *no_contact* alarm.
 - ◊ Remedy has NICE login problems which are being looked at.
 - ◊ There is an transparent intervention on the OPN routers tomorrow at 6:00 UTC (8:00 CET). A module needs to be changes, the routers will be on reduced power supply for 5 mins.

- Wednesday - April 18th:

- ◆ Morning meeting

- ◊ rb109 (ticket 173894) had a transient raid_tw alarm
- ◊ rb117 (ticket 173917) had a dependency problem, service manager has been contacted.
- ◊ CMS is having a lot of problems with castor. Service is in degraded mode.
- ◊ CMS castor DB will be moved tomorrow April 19.
- ◊ Remedy was down yesterday from 11:55 to 13:47
- ◊ ce101 has been drained.
- ◆ EGEE Broadcasts:
 - ◊ CERN-PROD: tomorrow, April 19th at 07:00h UTC, scheduled transparent intervention of the CERN SAM database
 - ◊ CERN-PROD: tomorrow, April 19th at 07:00h UTC, scheduled transparent intervention of the CMS castor database
- ◆ CCSR meeting:
 - ◊ New ORACLE vulnerability has been released last night. Patch is available and being evaluated. This requires an intervention with a downtime of about 30 mins.
- Tuesday - April 17th:
 - ◆ Morning meeting
 - ◊ rb109 had a transient raid_tw alarm
 - ◊ ATLAS castor stager is down
 - ◆ EGEE Broadcast: CERN-PROD: today, April 17th at 07:00h UTC, scheduled transparent intervention of the CERN vomrs and voms services (voms103)
- Monday - April 16th:
 - ◆ Morning meeting
 - ◊ rb114 had an alarm

Week 2007-04-09: David Collados

- Friday - April 13th:
 - ◆ Morning Meeting:
 - ◊ gridrb: hwscan_wrong on rb117
 - ◆ EGEE Broadcast: CERN-PROD: next Monday, April 16th at 07:00h UTC, scheduled transparent intervention of the CERN vomrs and voms services
- Thursday - April 12th:
 - ◆ Morning Meeting:
 - ◊ gridfts: fts_stuck on fts114. Briefly
- Wednesday - April 11th:
 - ◆ Morning Meeting:
 - ◊ gridce: high_load on ce107
 - ◊ gridce: no_contact on ce101. Disk failure. On vendor call.
 - ◆ WLCG SCM Meeting - Canceled.
 - ◆ CCSR Meeting:
 - ◊ CMS rb sandboxes almost full. Filled by 1 user.
- Tuesday - April 10th:
 - ◆ Morning Meeting:
 - ◊ gridce: high load on nodes (ce113, ce115). Machines rebooted.
 - ◊ gridlfc: no_read on nodes (lfc108, lfc111). LFC Support contacted.
 - ◊ gridrb: rb_wm_wrong on nodes (rb107, rb119). Problems of logging. Ask Rolandas about error.
 - ◊ gridrb: raid_tw on rb115.

Week 2007-04-02: Steve Traylen, Konstantin Skaburskas

- Friday 05/04: Steve Traylen
- Thursday 04/04: Steve Traylen
- Wednesday 03/04: Steve Traylen
- Tuesday 03/04: Konstantin Skaburskas
- Monday 02/04: Konstantin Skaburskas
 - ◆ Numerous CERN Interventions: availability of nodes was checked with Lemon; service managers for BDII, RB/WMS, VOMS, FTS were contacted to check the status of services.
 - ◆ Broadcast -- "completed -- Numerous CERN Interventions - Monday April 2nd"

Week 2007-03-26: Steve Traylen, Konstantin Skaburskas

- Friday 30/03: Konstantin Skaburskas
 - ◆ Preparations to "Numerous CERN Interventions - Monday April 2nd" - message to GD Service Managers (FTS(Paolo), Rolandas(RB/WMS), Maria(VOMS)), talked to Thorsten (CEs) and Veronique (BDIIs)
 - ◆ Broadcast -- "REMINDER -- Numerous CERN Interventions - Monday April 2nd"
- Thursday 29/03: Konstantin Skaburskas
 - ◆ Broadcast -- Emergency circuit maintenance (Start Date / Time: 03/29/07 03:00 GMT -- End Date / Time: 03/29/07 09:00 GMT)
- Wednesday 28/03: Konstantin Skaburskas
 - ◆ Broadcast -- Planned USLHCNET Maintenance on March 29th (I did not notice that Steve already sent it yesterday :))
- Tuesday 27/03: Steve Traylen
- Monday 26/03: Steve Traylen

Week 2007-03-19: Rolandas Naujikas

- Friday 23/03:
- Thursday 22/03:
- Wednesday 21/03:
 - ◆ ce103, lfc107, lfc108, rb107 - alarms
 - ◆ ce112 - GGUS#19892
 - ◆ Broadcast: RBs gdrbxx at CERN will be removed from production on 2007-03-23
- Tuesday 20/03:
- Monday 19/03:

Week 2007-03-12: Maite Barroso

- Friday 16/03:
 - ◆ Broadcast: LHCOPN: Maintenance on Tuesday 20/3/2007, the connections from CERN to RAL and BNL will be interrupted for 5 minutes to allow the replacement of a module in one of the CERN's LHCOPN router
 - ◆ Operators were not able to apply a procedure for an RB alarm because they did not have access to the node. Yvan and Romain are following this.
- Thursday 15/03:
 - ◆ Broadcast: changes to CERN production CEs
 - ◆ Broadcast: [LHCOPN] network intervention today 15/3/2007 at 10:30 CET, to configure the new 5Gbps link to TRIUMF
- Wednesday 14/03:
 - ◆ Broadcast: OPN instabilities T0-T1s are solved

- Tuesday 13/03:
 - ◆ Broadcast about OPN instabilities T0-T1s; mail exchange during the day with Edoardo Martelli to get status updates. The problem seemed to be solved at the end of the day.
- Monday 12/03:
 - ◆ Request to service managers to review threshold for high load alarms
 - ◆ Short unscheduled power cut (4 min) at ~9.00 am in the Meyrin site

Week 2007-03-05: Remi Mollon

- Friday 09/03:
 - ◆ SLS reported that Worldwide LHC Computing Grid was degraded because several Grid Tier-1 sites are down
- Thursday 08/03:
 - ◆ Network intervention is a success
 - ◆ SLS error reported on the morning meeting: Worldwide LHC Computing Grid was not available
 - ◆ Broadcast about 'All T0<->T2 traffic to move to new FTS instance tiertwo-fts-ws.cern.ch'
- Wednesday 07/03:
 - ◆ Broadcast about Scheduled network intervention at CERN. Due to urgent work on the network routers to resolve the problems of network instability currently being experienced, tomorrow morning March 8th between 7.30 and 8.30 the computer centre network will be unavailable for short 5 minute periods. This will affect ***ALL*** grid services hosted at CERN : BDII, CEs/SEs, UIs, FTS, LFC, VOMS/VOMRS (ALICE, ATLAS, CMS, LHCb, DTEAM, OPS, Sixt, Unosat, Geant4), SAM and Gridview, FCR.
- Tuesday 06/03: due to network problems (see Monday) ...
 - ◆ SLS error reported on the morning meeting: Worldwide LHC Computing Grid was not available; quickly fixed.
 - ◆ Broadcast about SAM test failures. The SAM BDII was temporarily suffering from a high load, which resulted in replication test failures (SRM, SE, CE, gCE). The problem is being investigated.
- Monday 05/03:
 - ◆ Tomcat didn't restart correctly on 'gridvoms' boxes, due to CDB misconfiguration which made SPMA uninstall jdk-1.5; quickly fixed.
 - ◆ CERN Computer Centre was down due to network problems from 10.59pm to ~07.50am

Week 2007-02-26: Judit Novak

- Friday 02/03:
 - ◆ Broadcast about a top-level BDII configuration error. The central configuration file that is downloaded by the BDIIs, contained additional sites to the 'Certified', 'Production' ones. This resulted in a PPS LFC URL being published, which could have caused files to be registered in a wrong catalog. The problem appeared since Monday, it was fixed on Friday.
 - ◆ SLS error reported on the morning meeting. The reason was a SAM DB problem which was fixed shortly.
- Thursday 01/03:
 - ◆ Broadcast about CERN SRM config problem, which caused CERN Prod SRMs not to be published for a few hours in the afternoon
- Monday 26/02:
 - ◆ Oracle intervention affecting LFC for LHCb finished successfully. A broadcast was sent about the news.

LHCb Oracle RAC extension/upgrade. Scheduled and agreed with LHCb responsables (who?). (Email by M.Anjo on 2007-02-20). We need more info about time and service impact to write more details here for this

event.

Week 2007-02-19: Piotr Nyczyk

- Friday 23/02:
 - ◆ evaluation of GMOD announcements triggered by interventions at CERN:
 - ◇ In case of intervention on important services it seems to be better to announce it more in advance because of external dependencies (services running in other sites)
 - ◇ The preparation of an intervention should be divided into three phases:
 1. Fix the time slot for intervention
 2. Prepare the list of affected services
 3. Plan the intervention in details
 - ◇ GMOD shouldn't wait until phase 3 is finished but rather should send EGEE broadcast immediately after phase 2.
 - ◇ GMOD should *translate* internal CERN services affected by intervention into the list of externally visible grid services (for example: LCG_SAME service in Oracle should be announced as SAM monitoring service).
 - ◇ To make the announcement clear and understandable, it shouldn't contain any details not relevant for grid users, namely it shouldn't mention internal CERN services - only the grid services affected. Internal CERN services can be mentioned only if they are used directly by some of the grid users (for example experiments database services?)
 - ◇ It would be useful to have a simple template for writing GMOD announcements and separately the mapping (table?) of internal CERN services (Oracle services, nodes, etc) to grid services (with indication of VOs)
 - ◇ The additional time that GMOD announcement is published before the intervention should be estimated in the way that would allow any external service managers to make additional announcements if their grid services are affected by the CERN grid services.
- Thursday 22/02:
 - ◆ LCGR Oracle RAC intervention:
 - ◇ intervention started and finished successfully according to the initial plan at 11:00
 - ◇ all services except VOMS and FTS restarted quickly
 - ◇ VOMS took more time because of multiple Oracle connection strings that have to be updated in many VOMS config files (VOMS should use global tnsnames.ora pointed by TNS_ADMIN variable to avoid such problems)
 - ◇ VOMS and FTS came back to production just before 12:00
 - ◇ EGEE announcement was made at 12:00 that Oracle intervention has ended and all services are fully operational
- Wednesday 21/02:
 - ◆ Morning meeting:
 - ◇ high load on SAM RBs, additional RB as quick solution, Piotr: follow up with Yvan
- Tuesday 20/02: * EGEE Broadcast: * LCG RAC Intervention - broadcast sent after long discussion by mail
- Monday 19/02:
- Notes from the past:
 - ◆ draft plan for a major intervention to LCG RAC to be delivered on Monday. The best info we have so far comes from an e-mail sent to the gmod-list

Week 2007-02-12: Antonio Retico

- Thursday 16/02:
 - ◆ Morning Meeting:
 - ◇ a relevant number of afs alarms started on 15/2 at ~12, still in progress:
 - ◇ castor: alarms reported since ~23 to ~5
 - ◇ ce: (ce104, ~17, OK) (ce101, 23, gris)(ce111, ~8)
 - ◇ rb: (rb115, 15, high load)(rb113, 8, high load)
 - ◆ EGEE Broadcast:

- Thursday 15/02:
 - ◆ Morning Meeting:
 - ◆ EGEE Broadcast:

- Wednesday 14/02:
 - ◆ Morning Meeting:
 - ◇ rb113: high load for a few mins
 - ◇ several alarms for lfc101 during scheduled downtime
 - ◆ CERN GGUS Tickets (Morning Meeting):
 - ◇ 18192: to be closed. Jan complains because opened by a sysadmin during scheduled downtime
 - ◆ LCG SCM Meeting: *more detailed plan to be drafted in advance before scheduling interventions. The plan should include the roll-back time and be followed-up by an appointed intervention coordinator
 - ◇ two interventions on LCG LCF RAC and lhcb LFC RAC to be scheduled for next week (plan to be drafted on monday)
 - ◇ broadcasts for scheduled interventions should be published in the news.
 - ◆ CCSR Meeting:
 - ◇ intervention on SRM affecting staging (FTS) foreseen (roughly) for 20th of March
 - ◇ major intervention (replacement of network switches) which will cause several racks to be put off-line has to be scheduled (possibly before the end of April)

- Tuesday 13/02:
 - ◆ Morning Meeting:
 - ◆ EGEE Broadcast:
 - ◇ 13.55: LCG RAC intervention being rolled-back
 - ◇ 15.00: LCG RAC intervention finished
 - ◇ 19.06: CERN-PROD:intervention on tape library 20/02/2007

- Monday 12/02:
 - ◆ Morning Meeting:
 - ◇ gridbdii: high load on several nodes for a few minutes (bdii101, bdii102, bdii105, bdii107, bdii108, bdii110, bdii111, bdii112)
 - ◇ gridce:
 - ce107 (swap_full, high_load, gridgris_Status_error, unable to connect)
 - c191 (no contact , unable to login, reset dump done and back in service)
 - ◇ gridrb:
 - rb118 (raid_tw),
 - rb109 (no contact, unable to ping, black_Screen, rebooted and then, ok)
 - ◇ lcgrb: high_load for a while on nodes (rb113, rb115)
 - ◆ EGEE Broadcast: Feb 13th 2007 9-12am CET scheduled Oracle interruption on lcg-voms.cern.ch affects CERN vomrs and voms

Week 2007-02-05: David Collados

- Friday 09/02:
 - ◆ Morning Meeting:
 - ◇ gridbdii: high load on several nodes (bdii111, bdii112, bdii101, bdii107)
 - ◇ lcgrb: high load on rb115
 - ◇ Physics services will be without UPS backup next Monday from 8h to 12h (due to the commissioning of new UPS to increase capacity).
- Thursday 08/02:
 - ◆ Morning Meeting:
 - ◇ Some gridbdii with high_load for a while. Nodes: bdii101, 112, 107, 106, 105 and 111.
 - ◇ ce105 was rebooted (had high_load and were unable to connect)
 - ◇ gridlfc. Several machines with many errors. Problem to be checked. Miguel Anjo said was probably related to the DB cluster, that was rebooted for some reason.
 - ◇ lcgrb with high_load due probably to tests submission. Nodes affected: rb113 and rb115.
 - ◇ gridvoms:
 - voms101 time outs. Problem due to a big Tomcat log file with several GBs.
 - voms103 with spma_error. no contact. able to ping but not to login.
 - ◆ EGEE Broadcast: Intervention on tape library 27/09/2006
 - ◆ EGEE Broadcast: LCGR intervention next Tuesday, 13 Feb, from 9:00h to 12:00h CET
 - ◆ EGEE Broadcast: CMS Oracle Rac intervention next Tuesday, 13 Feb, from 9:00h to 12:00h CET
- Wednesday 07/02:
 - ◆ Morning Meeting:
 - ◇ ce105: high load. Node to be rebooted
 - ◇ rb115: high load for a few mins
 - ◇ CERN SRM is down for maintenance until 12h.
 - ◆ CERN GGUS Tickets (Morning Meeting):
 - ◇ 18192: need more info to debug.
 - ◇ 17668: we should close it. Cannot do much about it.
 - ◆ LCG SCM Meeting: Nothing to report.
 - ◆ CCSR Meeting: A new kernel is being deployed in Scientific Linux CERN 4 (SLC 4X) nodes. Sysadmins will have to reboot the machines shortly.
- Tuesday 06/02:
 - ◆ Morning Meeting:
 - ◇ High load on gridbdii101
 - ◇ CERN SRM intervention will be performed tomorrow morning.
- Monday 05/02:
 - ◆ Morning Meeting: Announce today the external firewall maintenance that may lead to unstable service on 06/02 at 6am CET. (Note by Maria Dimou, previous gmod, heard at the morning meeting).
 - ◆ EGEE Broadcast: Maintenance on the CERN external network, Tuesday 6/2/2007 at 6h00
 - ◆ EGEE Broadcast: (Reminder) Scheduled downtime for CERN SRM on 7 February

Week 2007-01-29: Maria Dimou (Friday: Ian Neilson)

- Friday 02/02: Announcement ref: SRM downtime on Wednesday 7/2 should also be repeated at the ops meeting and **reminder announcement next week (? 6/2)**
- Thursday 01/02: It is easy to make the gmod feel useless. Example: Going to the morning meeting and seeing the following row in the operators' log "lfc_db_error. Unable to do the procedure (E486: Pattern not found: lfc101). See Details Diary." I try to find if some LFC service expert should be

prompted to investigate (James, Sophie, Yvan, Gavin ?). I go to LANDB and see "Responsible for the device: ADMINISTRATOR FS IT ". What sh

- Wednesday 31/01:
 - ◆ Sysadmins asked us at the morning meeting **not** to assign tickets to the relevant expert remedy sub-category (!) but to leave them in "General".
 - ◆ Oracle 10.2.0.3 3-monthly patches will be applied in production systems on Feb. 13th. Detailed discussion at the next LCGSCM. **A 3-hour downtime must be announced.**
 - ◆ The LCGSCM Service Reports' twiki page was suggested by Maria to the CCSR minute-taker.
 - ◆ The power-cut of the 24/01 night is due to a EDF failure and a bad function of the Physics' services' UPS which didn't take over. No report by TS dept. is published so far.
 - ◆ A new Physics Services' UPS will be installed on 12/02 during 8-12 hrs. This should be transparent, unless there is a power cut. **Maybe an announcement in due time? Up to the relevant GMod.**
- Monday 29/01:
 - ◆ Problems with [bdiilrb]10x listed in the Logger service were discussed with Yvan and clarified.
 - ◆ SPMA alarms raised during the week-end from the 3 voms servers were due to the removal of the voms-tools and voms-maint rpms, done by Remi and Maria on 25/01, which are now obsolete.
 - ◆ A broadcast to all production sites about Atlas asking Tier IIs to apply an indispensable DPM and LFC patch was published by Maite.

Week 2007-01-22: Konstantin Skaburskas

- Friday 26/01:
 - ◆ Power cut related:
 - ◇ Postmortem meeting on problems at "Daily Morning Meeting".
 - ◇ LFC problem with "pwd of the Oracle account" on lfc105, lfc106 - solved by service administrators.
 - ◇ PXE problem. Yvan submitted a Remedy ticket: CT398523 ENREGISTRE [Problem with PXE boot on some mid-range servers]
 - ◇ DHCP problem. Rolandas sent a message to Frederic.Chapron@cernNOSPAMPLEASE.ch.

Hi,

DHCP configuration used on gdrb(01-11) and probably on another nodes (after IP renumbering) is no dhclient tries to find DHCP server and after timeout, it use previous IP configuration. But ifup network connectivity by pinging to default gateway, but after some timeout it fails and bring into no more way to recover when network connectivity/services come back. The hostname is also issue. is down, many services don't start when hostname hasn't IP address in /etc/hosts. Possible solution:

1. Configure hosts with static IP address and hostname in /etc/hosts.
2. Patch ifup to do not bring down interface because of inaccessible default gateway.
3. Turn on power in defined sequence.

Regards, Rolandas Naujikas

- ◆ ◇ Problem with Lemon on rb10[4-7], rb11[3-5]. Servers were up and running, but Lemon showed that they were down. Message from Yvan:

..., there is a conflict with the following two packages:

- edg-fabricMonitoring-agent (coming from Lemon)
- edg-fabricMonitoring (coming from the gLite middleware for LCG RB nodes only)

The conflict comes from the /etc/rc.d/init.d/edg-fmon-agent file which is a bash script for Lemon. When I installed the machine using quattor for the OS, edg-fabricMonitoring-agent (from lemon) had then I installed the middleware which install its own /etc/rc.d/init.d/edg-fmon-agent, thus deleting

When the machine has been rebooted, the script `/etc/rc.d/init.d/edg-fmon-agent` coming from the mi instead of the one from lemon. The `no_contact` alarm has therefore been triggered...

Conclusion: there is a conflict of name between the two packages coming from Lemon and the middle. The only (quick) way to solve this problem is to force the reinstallation of the lemon package.

Yvan

- ♦ ◇ VOMS. voms10[1-3] - spma_error. Status not known for now (2.05PM). Maria was contacted.
- Thursday 25/01:
 - ♦ Power cut. Most of the grid services were up and OK in the early afternoon. All problems were identified and documented. Submit Remedy tickets to track the solution of the problems.
- Wednesday 24/01:
 - ♦ Nothing to report.
- Tuesday 23/01:
 - ♦ EGEE broadcasts:
 - ◇ Broadcast sent by Konstantin: January 23-26th 2007 voms/vomrs patch installation and tuning on the CERN VOMS servers
- Monday 22/01:
 - ♦ Nothing to report.

Week 2007-01-15: Rolandas Naujikas

- Friday 19/01:
 - ♦ bdii108 - reinstall, lcg-bdii.cern.ch - high load on all others.
 - ♦ ce106 - high load, swap_full, reset.
- Thursday 18/01:
 - ♦ Nothing to report.
- Wednesday 17/01:
 - ♦ A maintenance intervention on the CERN External Network on Thursday - EGEE broadcast send.
 - ♦ Problems with castoratlas at CERN - EGEE broadcast send.
 - ♦ Problems with castoratlas at CERN - second EGEE broadcast send.
 - ♦ From CCSR meeting: AFS clients on Linux have problems - deleted file is visible for others clients.
 - ♦ From CCSR meeting: firmware on WD disks should be upgraded offline - when ?
- Tuesday 16/01:
 - ♦ Nothing to report.
- Monday 15/01:
 - ♦ Nothing to report.

Week 2007-01-08 : Steve Traylen

- GmodWeek20070108

2006

Days 2006-12-31 to 2007-01-07: Robert Harakaly

FTS GMOD journal is available at [FtsMerryChristmas0607](#). Please note the known FTS service issues.

- Sunday 07/01:
 - ◆ Need to restart job controller on rb114.
 - ◆ Nothing to report.
- Saturday 06/01:
 - ◆ *fta_wrong* alarm on fts101, fts102 and fts105. Procedure done by FIO.
 - ◆ bdii107 (FIO: sysadmin called at home). Problem fixed.
 - ◆ ce103 (FIO: sysadmin called at home)
 - ◆ lfc101 (FIO: sysadmin called at home)
 - ◆ Remedy alarms:
 - ◇ 153420LD [↗](#) rb105: No contact, sysadmin called by FIO. Problem fixed.
 - ◇ 153349LD [↗](#) rb101: No contact, sysadmin called by FIO. Problem fixed but bad performances on this machine. Developers contacted.
 - ◇ 153238LD [↗](#) rb117: No contact, sysadmin called by FIO. Problem fixed.
 - ◇ 153384LE [↗](#) gdrb02: No contact, sysadmin called by FIO. Problem fixed.
 - ◇ Remedy ticket CT392757 concerning rb114: problem with the fans (*no_contact* alarm). This machine is still in production. I need to ask to LHCb (Roberto Santinelli) before to remove this machine from production.
 - ◇ *ncm_cdispd_wrong* on *gdrb** continues. German Cancio has been contacted by Yvan and he will take a look today. Problem fixed on 2007-01-08.
 - ◆ GGUS tickets follow up:
 - ◇ GGUS #17045: ticket closed. The machine has been reinstalled before christmas to WMS.
- ◆ NOTES: I tried to submit simple test jobs to rb101, rb105 and rb117. While for rb101 and rb105 jobs were submitted correctly (submitted and I am checking their state) for rb117 I am getting:

```
glite-job-submit --config-vo testWMS.conf hostname.jdl
```

```
Selected Virtual Organisation name (from proxy certificate extension): dteam
Connecting to host rb117.cern.ch, port 7772
Logging to host rb117.cern.ch, port 9002
**** Error: NS_SUBMIT_FAIL ****
"ProxyRenewalException: Error during Proxy Renewal registration." received when submitting a job
```

```
-bash-2.05b$ glite-job-submit --config-vo testWMS.conf hostname.jdl
```

```
Selected Virtual Organisation name (from proxy certificate extension): dteam
Connecting to host rb117.cern.ch, port 7772
Logging to host rb117.cern.ch, port 9002
**** Error: NS_SUBMIT_FAIL ****
"ProxyRenewalException: Error during Proxy Renewal registration." received when submitting a job
```

networkserver_events.log tells:

```
==> /var/log/glite/networkserver_events.log <==
06 Jan, 10:47:52 -I- "CFSI::crFSM": Creating FSM.
06 Jan, 10:47:52 -I- "CFSI::crFSM": Command Received...: GetMultiAttributeList
06 Jan, 10:47:52 -I- "CFSI::crFSM": Client Version...: 3.0.1
06 Jan, 10:47:52 -I- "CFSI::serializeServer": Serializing Server.
```

GmodJournal < LCG < TWiki

```
06 Jan, 10:47:52 -I- "CFSI::setMVA": Setting MultiValueAttributes.
06 Jan, 10:47:53 -I- "CFSI::crFSM": Creating FSM.
06 Jan, 10:47:53 -I- "CFSI::crFSM": Command Received.: JobSubmit
06 Jan, 10:47:53 -I- "CFSI::crFSM": Client Version...: 3.0.1
06 Jan, 10:47:53 -I- "CFSI::serializeServer": Serializing Server.
06 Jan, 10:47:53 -I- "CFSI::crContext": Creating Context for JobSubmit.
06 Jan, 10:47:53 -I- "CFSI::crContext": Setting LB parameters.
06 Jan, 10:47:53 -I- "CFSI::ckJobSize": Checking Job Size.
06 Jan, 10:47:53 -I- "CFSI::ckUserQuota": Evaluating User Quota.
06 Jan, 10:47:53 -I- "CFSI::insertCertSubj": Inserting Certificate Subject
06 Jan, 10:47:53 -I- "CFSI::setJobPaths": Setting Pathnames.
06 Jan, 10:47:53 -I- "CFSI::crReducedPathDirs": Creating reduced path dirs.
06 Jan, 10:47:53 -S- "CFSI::crReducedPathDirs": /var/glite/SandboxDir/cZ already exists.
06 Jan, 10:47:55 -I- "CFSI::evClientCreateDirs": Evaluating Remote Dirs Creation: Success.
06 Jan, 10:47:55 -I- "CFSI::crStagingDirectories": Creating staging dirs - job: https://rb117.cern.ch
06 Jan, 10:47:55 -I- "CFSI::copyFile": Copying file..
06 Jan, 10:47:55 -W- "CFSI::SDCPostControl": Registering job for proxy renewal.
06 Jan, 10:47:55 -F- "CFSI::LogAbortedJobE": Logging Aborted Job. Reason: Error during proxy renewal
06 Jan, 10:47:55 -W- "CFSI::test&Log": L&B call got a transient error. Waiting 60 seconds and trying again.
06 Jan, 10:47:55 -I- "CFSI::test&Log": Try n. 1/3
06 Jan, 10:48:05 -W- "CFSI::test&Log": L&B call got a transient error. Waiting 60 seconds and trying again.
06 Jan, 10:48:05 -I- "CFSI::test&Log": Try n. 2/3
06 Jan, 10:48:15 -W- "CFSI::test&Log": L&B call got a transient error. Waiting 60 seconds and trying again.
06 Jan, 10:48:15 -I- "CFSI::test&Log": Try n. 3/3
06 Jan, 10:48:25 -E- "CFSI::test&Log": L&B call retried 4 times always failed.
06 Jan, 10:48:25 -E- "CFSI::test&Log": Ignoring.
06 Jan, 10:48:25 -W- "CFSI::ckProxyRen": Proxy renewal: 0
```

After several trials I got job submitted also on rb117. Nevertheless in the logs the error:"L&B call got a transient error" appears frequently.

• Friday 05/01:

- ◆ *fta_wrong* alarm on fts101. Procedure done by FIO.
- ◆ Remedy alarms:
 - ◇ 153153SD [↗](#) rb114 FIO: no_contact, unable to do anything. See d.diary, reset dump done and back in service. Update from Yvan: Need to restart manually services *edg-wl-jc*, *edg-wl-lm* and *edg-fmon-server*. Note that this machine has been rebooted (*no_contact* alarm). All the services are back again.
 - ◇ 153138SD [↗](#) lfc112 FIO: lfc_noread. Mail sent to LFC.Support@cernNOSPAMPLEASE.ch.
 - ◇ Alarm *extfswarning* on gdrb06 due to some RAID disk error discovered by Maarten last night. All the services have been stopped (gridftp services excepted) and the machine has been put in maintenance. MySQL database corrupted. Problem fixed.
 - ◇ Alarm *tmp_full* on rb111. There were indeed 3 huge files in directory */tmp*. I deleted the content of these 3 files without removing the files themselves. Fixed.
 - ◇ Alarm *ncm_cdispd_wrong* on all the gdrbxx nodes. I don't know what is the meaning of this alarm. Need some investigations... Update: problem fixed by executing *ccm-fetch* and *ncm-ncd --configure fmonagent* on all the gdrbxx nodes. No, it is not enough 🤔 I sent an email to the experts...
- ◆ GGUS tickets:
 - ◇ 17058 [↗](#) rb103.cern.ch : Logging and Bookkeeping seems to be BLOCKED
 - ◇ Reply from gridview experts, ticket 17052 solved
- ◆ Note from Yvan: rb114 was down this morning during about 1 hour. The sysadmins rebooted this machine, and nothing special has been found in the logs. I restarted *edg-wl-jc*, *edg-wl-lm* and *edg-fmon-server* services manually and all the services are back again.

• Thursday 04/01:

- ◆ *fta_wrong* alarm on fts101, fts102. Procedure done by FIO.
- ◆ GGUS tickets:

- ◇ 17052 [↗](#) SFT_SAME_PUBLISHER_WSDL not reachable (mail sent to service experts)
- ◇ 17048 [↗](#) Two central LFC servers published for dteam VO (Sent request for clarification to submitter) Ticket closed
- ◇ 17045 [↗](#) Resource Broker rb108.cern.ch is not publishing JobStatusRaw data (mail sent to service experts)
 - 19:24 Rolandas: restart of RGMA servicetool. Submitter contacted to verify.
- Wednesday 03/01:
 - ◆ intervention on ce102 done. SAM shows it green
 - ◆ GGUS ticket 17030 [↗](#) concerning possible problem on LFC. (To be fixed by FIO)
 - ◆ remedy ticket 153032SD [↗](#) on rb113. Update made by Yvan: the machine has been rebooted and the services have been restarted automatically. There was only a problem with the Lemon monitoring. I reinstalled the *edg-fabricMonitoring-agent* package because the monitoring was not running anymore. This monitoring is back again and the LAS ticket is closed. Happy New Year to all the readers of this wiki page 😊 Update by Robert: Thank you, Happy new year to you also 😊. Yes RB works well, I submitted test job it behaves normally ...
- Tuesday 02/01:
 - ◆ 22:26 ce102 : No contact ticket: 152994LD (To be fixed by FIO)
 - ◆ Nothing to report
- Monday 01/01:
 - ◆ Nothing to report
- Sunday 31/12:
 - ◆ gdrb06 back in production.
 - ◆ *fta_wrong* alarm on fts102. Procedure done by FIO.
 - ◆ SAM:
 - ◇ 2006/12/31 - 09:25 : Unable to connect the SAM service. Error message: lcg-sam.cern.ch has received an incorrect or unexpected message. Error Code: -12227. Service experts contacted. I investigate the problem in parallel.
 - ◇ 2006/12/31 - 15:25 : Answer from Judit SAM works. Problem was on my side, expired certificate in my browser.

Days 2006-12-28 to 2006-12-30: Konstantin Skaburskas

FTS GMOD journal is available at [FtsMerryChristmas0607](#). Please note the known FTS service issues.

- Saturday 30/12:
 - ◆ *fta_wrong* alarm on fts101, fts102. Procedure done by FIO.
 - ◆ SAM:
 - ◇ 2006/12/30 - 08:59:24 : SAM tests page does not show anything but header! Host (lcg-sam.cern.ch) is reachable on service port 8443. Service experts contacted.
 - ◇ 2006/12/30 - 16:44:43 : Judit Novak reacted on alarm sent to service experts and SAM page is OK now!
- Friday 29/12:
 - ◆ *fta_wrong* alarm on fts102, fts105. Procedure done by FIO.
 - ◆ Lemon:
 - ◇ follow-up of yesterday's suspicious 50% load on grvw002. Maarten's email reached Kumar Viabhav (kvaibhav@barcNOSPAMPLEASE.gov.in): "I saw a strange problem. There was a java process that is showing 50% CPU utilisation in top. But there was no such process shown in ps -ax o/p with that PID. Any way after I killed the process with that PID(although it does not exist in ps -ax). The top now shows

low cpu utilisation. I think now thing will improve."

- ◆ job controller service restarted on rb107 (Yvan).
- Thursday 28/12:
 - ◆ *fta_wrong* alarm on fts105. Procedure done by FIO.
 - ◆ SAM:
 - ◇ ce105.cern.ch - CERN-PROD (CERN) ERROR (28-Dec-2006 11:16:47) Timeout when executing test CE-sft-lcg-rm after 600 second
<https://lcg-sam.cern.ch:8443/sam/sam.py?func=ShowHistory&sensors=CE&vo=ops&nodena>
 - ◆ Lemon:
 - ◇ suspicious one week 50% load on rb114 and grvw002. Checked with Yvan Calas. Yvan checked rb114 - no problems found. Yvan sent an email to Maarten concerning grvw002.

Days 2006-12-22 to 2006-12-27: Yvan Calas

FTS GMOD journal is available at FtsMerryChristmas0607. Please note the known FTS service issues.

- Wednesday 27/12:
 - ◆ *fta_wrong* alarm on fts101.
 - ◆ Restart job controller on rb106 and rb107.
 - ◆ bdii110. back in production after the change of one of its RAID disks.
 - ◆ Maarten solved a problem with the input.fl file on rb101. More details can be found on the support-exp-glite-rb mailing list.
- Tuesday 26/12:
 - ◆ *fta_wrong* alarm on fts101.
 - ◆ LAS Alarm #761103: Unrecoverable error on gdrb06 (problem with RAID disk */dev/hde*). This machine has been put in maintenance. Since it is still possible to submit and retrieve job output successfully, jobs submissions are still allowed. Note also that Maarten modified the configuration of the UI lxplus by removing gdrb06 from the list of LCG RBs dedicated to CMS. No broadcast has been sent to CMS for the time being.
 - ◆ PRMS Tickets: 1
 - ◇ #392085: *uncorrectable_error* alarm on gdrb06.
- Monday 25/12:
 - ◆ Nothing to report.
- Sunday 24/12:
 - ◆ Restart service *lcg-mon-job-status* on rb106.
 - ◆ Intervention on fts103 to fix (again) unstable information provider.
 - ◆ Restart job controller on rb101 because of the high load of this machine (almost 20). This should prevent an alarm to be triggered. Di applied some patches (#940 and #944), causing the load to decrease even if there is still a high system cpu load).
- Saturday 23/12:
 - ◆ LAS Alarm #760385: *smart_selftest* failure on bdii110 (sam-bdii). Machine put in maintenance and vendor call requested by *ebonfill*. Mail sent to the *it-dep-gd-gmod* mailing list. We must check the load of bdii109 which is the second SAM BDII in production in the next few days.
 - ◆ PRMS Tickets: 3
 - ◇ #391901: redirection of the stderr output to */dev/null* in */etc/cron.d/glite-wms-check-daemons.cron cron job* on rb112. I applied the same modification on rb115, rb116, rb117 and rb108.

- ◇ #391541: put the status of this ticket as fixed (root access for Robert on several cluster enabled).
- ◇ #391581: put the status of this ticket as fixed (root access for Robert on several cluster enabled).

Week 2006-12-18: David Collados (Maite as backup)

FTS GMOD journal is available at [FtsMerryChristmas0607](#). Please note the known FTS service issues.

- Friday 22/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ The Lemon problem was fixed yesterday by Dennis Waldron. James will add the new rpm to swrep and configure it in CDB, so we'll start receiving the LFC alarms again by today.
 - ◇ Problems with FTS information provider (`infosys` group problem). This causes the FTS SAM tests to fail intermittently. Service impact is minimal: all experiments cache the endpoint.
 - ◇ Known site support times during Christmas are listed in [FtsMerryChristmas0607SiteStatus](#).
- Thursday 21/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ LFCs still hanging and not raising alarms. James will keep on working on that today to try to solve the problem. Yesterday evening/night we had the 2 machines under `prod-lfc-shared-central` stuck, what among other things, made all the SAM tests fail.
- Wednesday 20/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ `ce106` had high load yesterday
 - ◇ We're still missing one LFC per day due to the problem reported in the last days. Last night it was `lfc106`. James will ask German to create a new lemon metric to be able to detect the problem asap reading the lemon log. Then, he will review the procedures for the sys admins and serv. managers.
 - ◆ CCSR meeting: Nothing to report.
 - ◆ broadcasts: Temporary harmless warning due to the change of host certificate for `voms.cern.ch`
- Tuesday 19/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ Still some LFC problems. Yesterday on `lfc107` and `lfc108`. Lemon alarms didn't raise, so James will look into them again today and will check if the 'sensor time out' error message appeared yesterday on the lemon logs. Only '`lfc_log_db_error`' appeared but probably this was due to the daemons' restart of Sophie.
 - ◇ IBM robots should be operational since yesterday afternoon.
- Monday 18/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ Some LFCs getting stuck due to infinite loop in memory allocation call. Being investigated.
 - ◇ Problem in `lsfnfs02` which makes gCEs not work. Being fixed.
 - ◆ WLCG/EGEE/OSG operations meeting: Nothing to report.

Week 2006-12-11: Antonio Retico

- Friday 15/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ nothing to report.
 - ◆ PRMS Tickets: xx
 - ◆ broadcasts : Tape library intervention 18/12/2006 and 19/12/2006
 - ◆ Other: Two LFC nodes got stuck in an infinite loop. Resolved after restarting the service. Problem will have to be investigated by LFC developers.

- Thursday 14/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ nothing to report.
 - ◆ PRMS Tickets: xx
 - ◆ broadcasts :
 - ◆ Other:

- Wednesday 13/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ nothing to report.
 - ◆ CCSR meeting: nothing to report
 - ◆ PRMS Tickets: xx
 - ◆ broadcasts : Voms/VOMRS not available due to Tomcat problems.
 - ◆ Other:

- Tuesday 12/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ nothing to report.
 - ◆ PRMS Tickets: xx
 - ◆ broadcasts :
 - ◆ Other:

- Monday 11/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ High load for few minutes on rb113 (LCG RB for SAM). This problem will be solved when the second RB dedicated to SAM will be put in production.
 - ◇ *smartd_selftest* alarm on lfc103 and lfc109.
 - ◇ *no_contact* on lxb1540: this machine was stopped due to the vault temperature.
 - ◆ PRMS Tickets: xx
 - ◆ broadcasts :
 - ◇ Firewall maintenance on Tuesday the 12th of December 2006 between 6:00AM and 7:30AM CET.
 - ◇ CASTORATLAS intervention today at 9h45 a.m.
 - ◇ CASTORCMS intervention today at 9h45 a.m.
 - ◇ CERN-PROD: Service interruption due to Castor maintenance at 3.30 PM
 - ◇ CERN-PROD: Possible slowness of tape libraries at 6h.10 PM
 - ◆ WLCG/EGEE/OSG operations meeting:
 - ◆ Other:
 - ◇ Air-conditioning in the vault of the Computer Centre stopped working on Saturday, December 09, around 13:00. The temperature raised very quickly in the critical area of the Vault. In order to protect our critical nodes, about 150 lxb nodes and a few lxbplus nodes located close to the critical zone were stopped in a hurry to limit the emperature increase. Air-conditioning was fixed at 14:45 and the nodes which have been stopped are being re-started.

Week 2006-12-04: Yvan Calas

- Friday 08/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ High load alarm on rb113 (SAM LCG RB). New machine requested (planned on next Monday).
 - ◆ PRMS Tickets: 1
 - ◇ #388024: rb112 under performing. Ticket assigned to Yvan Calas.
 - ◆ broadcasts:
 - ◇ Withdrawal of lcg-registrar.cern.ch planned on 2006-12-11.
 - ◆ Other:

- Thursday 07/12
 - ◆ 9:00 Daily morning meeting: nothing to report.
 - ◆ PRMS Tickets: 4
 - ◇ #387637: R-GMA calls not needed in SAM-client job submission. Ticket assigned to Piotr Nyczyk.
 - ◇ #387644: No SE found for SAM client. Ticket assigned to Piotr Nyczyk.
 - ◇ #387647: PPS job submit (Admin portal). Ticket assigned to Piotr Nyczyk.
 - ◇ #387652: R-GMA calls not needed in SAM-client job submission. Ticket assigned to Piotr Nyczyk.

- Wednesday 06/12
 - ◆ 9:00 Daily morning meeting:
 - ◇ Plan to move VOBOX nodes managed by GD (lxn1179 and lxb2065) to FIO. In the long term, one VOBOX will be assigned to each VO.
 - ◇ High load for few minutes on rb113 (LCG RB for SAM).
 - ◆ 10:00 WLCG-SC meeting: minutes [↗](#)
 - ◇ 6 new gLite CEs will be given to FIO.
 - ◇ 3 new gLite WMS will be given to GD.
 - ◇ FIO will take the responsibility of the 2 VOBOX machines managed by GD (lxn1179 and lxb2065). New machines will replace the former ones.
 - ◇ New VO *EELA* needs to be configured on rb103 and rb108.
 - ◆ 15:00 CCSR Meeting: minutes [↗](#).
 - ◆ PRMS Tickets: 4
 - ◇ #387452: PPS job submission. Reassigned to Antonio Retico.
 - ◇ #387374: SAM sensors. Reassigned to Piotr Nyczyk.
 - ◇ #387343: PPS job submission. Reassigned to Antonio Retico.
 - ◇ #387326: *lcg_mon_job_status* down on rb107. This was due to the reinstallation of rb107 which is now a LCG RB for CMS.
 - ◆ broadcasts:
 - ◇ Downtime of ce103, ce104, ce105, ce106, ce107 on 2006-12-06.

- Tuesday 05/12 (Yvan sick, replaced by Maite)
 - ◆ 9:00 Daily morning meeting:
 - ◆ PRMS Tickets: xx
 - ◆ broadcasts :
 - ◆ Other: none

- Monday 04/12 (Yvan sick, replaced by Maite)
 - ◆ 9:00 Daily morning meeting:
 - ◆ PRMS Tickets: xx
 - ◆ broadcasts :
 - ◆ WLCG/EGEE/OSG operations meeting:

Week 2006-11-27: Maite Barroso

- Monday 27/11
 - ◆ 9:00 Daily morning meeting:
 - ◆ broadcasts : tape library intervention at CERN on Thursday 30/11/200; rb109 shutdown Tuesday 28 November

- Wednesday 29/11
 - ◆ 9:00 Daily morning meeting:
 - ◆ broadcasts: Firewall pre-deployment verification - Thu 30th November; CERN-PROD: FTS upgrade on Thursday 30th
 - ◆ CCSR: several small updates in the firewall during December, should be transparent (clean up)

- Thursday 30/12
 - ◆ 9:00 Daily morning meeting: ce105 is unreachable for the second time this week; it is showing the GRID_GRIS_WRONG alarm. The restart worked the first time, but one day later the problem is there again. Problem escalated to the CE service experts. Steve Traylen is having a look at it. Thorsten will configure automatic restarts in case of GRID_GRIS_WRONG alarms, so none of us will get any tickets.

- Friday 01/12 - What a busy day for the GMOD!!
 - ◆ 9:00 Daily morning meeting: few alarms:
 - ◇ rb101 swap full
 - ◇ ce105 and ce106 unreachable: turned up to be a network problem, a bad network switch that was replaced in the afternoon
 - ◇ lxb7283 (gridrb) high load
 - ◇ The FTS intervention announced for yesterday did not take place because of the CDB problems; it will happen when these problems are solved.
 - ◆ PRMS Tickets: one for rb112, two for VOMS (voms101 and voms103), and one for rb101. All were solved during the day. Thanks to teh service managers for the quick reaction!
 - ◆ broadcasts: * CERN-PROD: ce106 and ce105 out of production * Several fixes applied to the SAM production service at CERN * CERN-PROD: Network intervention on Saturday December 2 and Monday December 4 * CERN-PROD: scheduled power cuts on Monday December 4 * CERN-PROD: 30 min interruption in production CEs * December 4th 2006 @16hrs [☞](#) CET scheduled vomrs upgrade on lcg-voms.cern.ch
 - ◆ Actions:
 - ◇ Friday Dec. 1st CIC Broadcasts (Notes by Maria Dimou on Nov 30th):
 - Please check page <http://cern.ch/it-servicestatus/> [☞](#) for **two** interventions planned for Dec. 2nd and 4th.
 - VOMRS upgrade announcement with the following text and publication parametres:

News on cic.in2p3.fr: YES

Email to: ROC managers, Site managers, VO managers of ALICE, ATLAS, CMS, LHCb, DTEAM, Geant4 and OPS *only!!*, VO users of ALICE, ATLAS, CMS *only!!*.

Cc: vomrs-grid-support@cern.ch, tlevshin@fnal.gov, support@ggus.org

Title: December 4th 2006 @16hrs CET scheduled vomrs upgrade on lcg-voms.cern.ch

Text:

The lcg-voms.cern.ch vomrs service will be unavailable for, we hope, a maximum period of 2 hours on December 4th 2006 @16hrs CET

Reason: Upgrade to vomrs-1.3.0. Contents of this Release in <https://twiki.cern.ch/twiki/bin/>
Users obtaining vomrs-proxies will not be affected by the interruption. Site admins, users, Rep

the web interface or build the gridmap file are kindly asked to be patient.
This applies to VName = ALICE, ATLAS, CMS, LHCB, DTEAM, OPS, Sixt, Unosat, Geant4

Thank you very much for your collaboration

maria dimou / grid deployment

Week 2006-11-20: Piotr Nyczyk

- Wednesday 22/11
 - ◆ Broadcasts: * CERN-PROD: Scheduled service maintenance
- Tuesday 21/11
 - ◆ Broadcasts: * CERN-PROD: SRM failures during the weekend
- Monday 20/11
 - ◆ 9:00 Daily morning meeting: poor FTS performance reported and commented by James/Gavin
 - ◆ PRMS Tickets: 2 tickets related to rbXXX.cern.ch machines at CERN - assigned to Yvan
 - ◆ broadcasts: 1 broadcast about IP renumbering at CERN
 - ◆ WLCG/EGEE/OSG operations meeting: no outstanding issues
- TODO:
 - ◆ Before you go to the CCSR meeting, submit a written report for the minutes.
 - ◆ On Thursday all PRMS tickets that would be assigned to M.Dimou should go to I.Neilson.
 - ◆ On Friday all PRMS tickets that would be assigned to M.Dimou should go to L.Ma.

Week 2006-11-13: Maria Dimou

- Friday 17/11 (by Ian Neilson GMod backup this week)
 - ◆ 9:00 Morning meeting:
 - ◇ No Grid-related items or from the Logger [↗](#)
 - ◆ 14:55 BROADCAST CASTORPUBLIC intervention for up to 30 mins 8:30UTC 20-Nov (Mails were sent but error reported inserting into archive -Error : Insert in Archive failed. But mails have been sent A email has been sent to the CIC Portal Team to inform them)
- Thursday 16/11
 - ◆ 9:00 Morning meeting:
 - ◇ No Grid-related items.
 - ◆ 9:30 Announced on <https://cic.in2p3.fr> [↗](#) the possible instabilities of the CERN network on Fri Nov. 17th 6-7:30am CET due to Firewall tests.
- Wednesday 15/11
 - ◆ 9:00 Morning meeting:
 - ◇ No entries relevant to Grid Services, except possibly an alarm on GRIDFTP_WRONG [↗](#). It was sent to castor.support but is it our business?
 - ◆ 10:00 WLCGSCM
 - ◆ PRMS Remarks: 5 tickets CT38165x, x=3,4,5,,6,7 [↗](#) were assigned to Maria Dimou as GMOD as a result of email from root@rb10y to rb.support@cernNOSPAMPLEASE.ch. The support address never received this mail.
 - ◆ 15:00 CCSR: More information on the UPS test planned for Dec, 6th.
- Tuesday 14/11
 - ◆ 9:00 Morning meeting:
 - ◇ alarm on database inconsistency raised from voms103 (alias lcg-voms.cern.ch) to Service Manager == GMod (M.Dimou). Resulting required actions being discussed with developers.
 - ◇ clusters c2cms, c2alice, c2atlas: many lines in remedy logger but after discussion with CMS VO Admin (A.Sciaba) decided that they don't relate to the Grid Services and can be ignored. FIO and CMS have direct channels to handle such information.

◇ 14:30 Published on <http://cic.in2p3.fr> 10mins VOMRS interruption scheduled for Nov. 15th at 10am UTC due to Oracle backend intervention. VO Admins and users informed by email in addition.

• Monday 13/11

◆ 9:00 Morning meeting:

- ◇ rb104: a few lines on high load. Email to [rb.support@cernNOSPAMPLEASE.ch](mailto:rb.support@cern.ch)
- ◇ fts110+5: several lines on no_contact.
- ◇ ce101,ce107: vm_kill. Email to [ce.support@cernNOSPAMPLEASE.ch](mailto:ce.support@cern.ch)
- ◇ clusters c2cms, c2alice, c2atlas: many lines on sysadmin intervention for raid_tw.
- ◇ Alice equipment in build. 14 failed all week-end. Data transfers failed.

◆ 14:30 Passed info on lcg cvs planned service interruption to O.Keeble and OPS section.

◆ 16:00 Attended Weekly Operations Meeting. Minutes [↗](#)

Week 2006-11-06 : Naujikas Rolandas

• Monday 06/11

◆ 9:00 Morning meeting:

- ◇ small flood in computer center
- ◇ ce102 - problems
- ◇ problems with LFC servers
- ◇ high load on rb110 and lxb7283
- ◇ VOBOX voalice03 and lxb7281 not yet ready for production
- ◇ CMS: T1 -> T2 transfers
- ◇ ALICE: FTS transfers died last night
- ◇ ATLAS: the same problem
- ◇ FTS: known problem with FTS service - patch in preparation

◆ 13:50 send broadcast for CMS people about lxb1930 downtime on 09/11/2006 from 08:00am to noon CET.

◆ 14:20 send broadcast about CERN firewall maintenance on 08/11/2006 from 6:00am to 7:30am CET.

◆ 16:00 WLCG/EGEE/OSG operations meeting

• Tuesday 07/11

◆ 9:00 Morning meeting:

- ◇ rb109 high load
- ◇ lxbatch several nodes have problems
- ◇ lxplus009 unable to login
- ◇ castor file system problems, in testing
- ◇ FTS: no big problems with transfers

◆ 10:03 #379381 reassigned to Yvan Calas

• Wednesday 08/11

◆ 8:50 #379582 reassigned to Grid-WMS-service-expert

◆ 9:00 Morning meeting:

- ◇ fts102
- ◇ rb109
- ◇ FTS: problems with source files access
- ◇ CMS: FTS transfers failed - Oracle intervention could cause this problem
- ◇ GridView problems
- ◇ LSF upgrade today

◆ 10:00 Missed SCM meeting, didn't receive new schedule

◆ 10:40 Closed #379381 because of known problem

◆ 15:00 CCSR meeting (513-R-070)

◇ Remedy could be used for machines reboot

◇ Castor was downgraded for LHCb and ATLAS because of linkage to old libraries

• Thursday 09/11

- ◆ 9:00 Morning meeting:
 - ◇ ce105 - high load
 - ◇ fts103,fts104 - fts_stuck - restart servise
 - ◇ rb104 - high load
 - ◇ FTS: workaround with connection to Oracle, keepalive tuning
- Friday 10/11
 - ◆ 9:00 Morning meeting:
 - ◇ LFC monitoring was in wrong state
 - ◇ rb104 high load
 - ◇ rb108 gone, hardware problem
 - ◇ FTS: some problems
 - ◇ CMS: glite-WMS: 16000 jobs/day, job output retrieval problems
 - ◇ ALICE: FTS problems, in investigation
 - ◇ castor: new version to deploy, EGEE broadcast to send
 - ◆ 9:45 #380436 reassigned to Grid-WMS-service-manager
 - ◆ 13:26 #380562 reassigned to Grid-RB-service-manager

Week 2006-10-30 : Skaburskas Konstantin

- Friday 03/11
 - ◆ 9:00 Daily morning meeting:
 - ◇ gridlfc - lfc112 - alarms (possible problmes with alarm scripts)
 - ◇ gridrb - lxb7283, rb102 - high loads for a while
 - ◇ gridvw - grvw03
 - ◇ problems with putting gdrb in LHC cluster
 - ◇ new LSF RPMs is under testing
 - ◆ broadcasts :
 - ◇ Broadcast sent by Konstantin: announce - tape library intervention 08/11/2006 at CERN (PROBLEM: Error : Insert in Archive failed. But mails have been sent; A email has been sent to the CIC Portal Team to inform them)
 - ◇ Broadcast sent by Konstantin: Announce - tape library intervention today at CERN (PROBLEM: Error : Insert in Archive failed. But mails have been sent; A email has been sent to the CIC Portal Team to inform them)
 - ◇ Broadcast sent by Konstantin: CERN-PROD - Change of GRID queues for DTEAM and OPS users (PROBLEM: Error : Insert in Archive failed. But mails have been sent; A email has been sent to the CIC Portal Team to inform them)
 - ◇ Broadcast sent by Konstantin: [FIREWALL] Maintenance on the CERN firewall, Wed 8th of November 2006 (PROBLEM: Error : Insert in Archive failed. But mails have been sent; A email has been sent to the CIC Portal Team to inform them)
 - ◆ PRMS Tickets:
 - ◇ #378597: resigned to Third Level
 - ◇ #378607: resigned to Third Level
- Thursday 02/11
 - ◆ 9:00 Daily morning meeting: nothing to report.
 - ◆ PRMS Tickets:
 - ◇ #377... (check with Rolandas) : resigned to Third Level
 - ◇ #377... (check with Rolandas) : resigned to Third Levelrpm
- Wednesday 01/11
 - ◆ 9:00 Daily morning meeting: nothing to report.
 - ◆ broadcasts :
 - ◇ TO ANNOUNCE: "announce tape library intervention 08/11/2006"
- Tuesday 31/10
 - ◆ 9:00 Daily morning meeting:
 - ◇ alarms on rb108 (cluster: lcgrb)

- ◊ IPs renumbering in progress
- ◊ GD Data Management team reported: work on workaround to detect hangs of FTS's WS in connection to DB and take appropriate measures.
- ◆ broadcasts :
 - ◊ Broadcast sent by Konstantin: November 1st 2006 @10am UTC scheduled vomrs interruption on lcg-voms.cern.ch
 - ◊ Broadcast sent by Konstantin: GRID batch jobs are back on CERN-PROD RBs (PROBLEM -- Error : Insert in Archive failed. But mails have been sent; A email has been sent to the CIC Portal Team to inform them)
- ◆ PRMS Tickets:
 - ◊ #377848 : "lxb7283 high_load. Load average: 23.57, 22.98, 21.38.". Action --> reassigned to Third Level Grid-WMS-service-expert
- Monday 30/10
 - ◆ 9:00 Daily morning meeting:
 - ◊ New LFC (LHCb) to appear in production this week.
 - ◊ Castor updated.
 - ◊ IPs renumbering on Tuesday morning affects RBs - care must be taken (send reminder broadcast).
 - ◊ BD security patch on Wednesday - FTS and LFC must be tested.
 - ◊ FTS problems:
 - FTS fts103 - tomcat stuck - restarted.
 - CMS - entire FTS service stopped. Saturday - traffic stopped T0->T1 (if channels are not filled enough with transfers dteam takes the channels over).
 - Atlas - there are no transfer failures but transfer rate is very-very low (now).
 - ◆ Problems with Remedy - was not able to access tickets (Cannot open catalog; Message number = 91 :RPC: Procedure unavailable (ARERR 91)).
 - ◆ broadcasts :
 - ◊ Broadcast sent by Konstantin: CERN-PROD - scheduled downtime of RBs (Reminder)
 - ◊ Broadcast sent by Konstantin: [LHCOPN] Maintenance on 31/10/2006: SARA's link reconfiguration. (PROBLEM -- Error : Insert in Archive failed. But mails have been sent; A email has been sent to the CIC Portal Team to inform them)
 - ◊ Broadcast sent by Konstantin: CERN-PROD - scheduled downtime of RBs - TIME CHANGES. (PROBLEM -- Error : Insert in Archive failed. But mails have been sent; A email has been sent to the CIC Portal Team to inform them)
 - ◆ PRMS Tickets: 1 closed (GGUS #14378), 1 reassigned (#)

Week 2006-10-23 : Yvan Calas

- Friday 27/10
 - ◆ 9:00 Daily morning meeting: nothing to report.
 - ◆ PRMS Tickets: 0
 - ◆ broadcasts :
 - ◊ Broadcast sent by Harry: Loss of exp-bdii and atlas-bdii CERN DNS names for 20 minutes from 8.35 UTC today.
 - ◊ Broadcast sent by Yvan: Scheduled service maintenance for the CERN-PROD site due to the IP re-numbering of production nodes (intervention planned on 31/10/2006). The machines involved are:
 - gdrb01 to gdrb11 (LCG RBs).
 - lxb0726 and lxb0726 (gLite UI for Atlas).
 - lxn1179 (VOBOX for Atlas).
 - lxn1180 (SAM server).
 - lxn1181 (SAM server backup).
 - lxn1182 (SAM client).

· lxn1183 (classic SE).

- ◆ Other:
 - ◇ LFC service is now running on new hardware (lfc101 to lfc113).
- Thursday 26/10
 - ◆ 9:00 Daily morning meeting
 - ◇ lxb1134: alarm *lfc_noread*.
 - ◇ bdii101: alarm *high_load*.
 - ◇ ce103: alarm *gridftp_wrong*.
 - ◇ ce104: alarm *gridftp_wrong*.
 - ◇ Merge of the two EGEE.BDII clusters exp-bdii and lcg-bdii into one load-balancing cluster.
 - ◇ Security update for SLC3 and SLC4 (kernel).
 - ◆ PRMS Tickets: 9
 - ◇ #376527: kernel upgrade to do on gdrb10.
 - ◇ #376528: kernel upgrade to do on gdrb02.
 - ◇ #376529: kernel upgrade to do on gdrb03.
 - ◇ #376530: kernel upgrade to do on gdrb11.
 - ◇ #376531: kernel upgrade to do on gdrb06.
 - ◇ #376532: kernel upgrade to do on gdrb07.
 - ◇ #376534: kernel upgrade to do on gdrb01.
 - ◇ #376535: kernel upgrade to do on gdrb09.
 - ◇ #376633: job status problem on rb105.
 - ◆ broadcasts : none.
 - ◆ Other:
 - ◇ The Certificate Authority that has been run in the GD group until now is being replaced by one run by IT/IS Group. The new CA is now operational: <http://cern.ch/ca>.
- Wednesday 25/10
 - ◆ 9:00 Daily morning meeting
 - ◇ lxb1134: problem with alarm *lfc_noread* solved.
 - ◇ Castor upgrade at CERN done this morning.
 - ◇ CERN-PROD tickets: http://egee-docs.web.cern.ch/egee-docs/ROC_CERN\CERN-PROD_tickets\CERN-tickets-2006-10-25
 - ◆ 10:00 WLCG-SC meeting: minutes [☞](#).
 - ◆ 15:00 CCSR Meeting: minutes [☞](#).
 - ◇ Merge tomorrow the two EGEE.BDII clusters exp-bdii and lcg-bdii into one load-balancing cluster.
 - ◆ PRMS Tickets: 1
 - ◇ #376320: Problem with the lcg-mon-job-status service on gdrb02.
 - ◆ broadcasts: none
- Tuesday 24/10
 - ◆ 9:00 Daily morning meeting
 - ◇ lxb1134: alarm *lfc_noread*: a single user caused excessive load on this machine (transient problem).
 - ◇ rb110 is a new gLite WMS for CMS and has been put in production yesterday afternoon.
 - ◇ 15 new LFC machines today, and 15 more tomorrow.
 - ◇ Request to publish the RBs jobs statistics in *gridview* (follow up by James).
 - ◆ PRMS Tickets: 1
 - ◇ #375920: Test of the RB.support@cernNOSPAMPLEASE.ch mailing list on the LCG RBs machines (apt-autoupdate configuration file).

- ◆ broadcasts :
 - ◇ Broadcast sent by Yvan: Castor upgrade at CERN tomorrow wednesday 25 october at 09:30 CEST (apply Oracle security patch, upgrade Castor software and reboot database servers for new kernel).
 - ◇ Broadcast sent by Nick: Update 8 for gLite 3.0 has been released to the EGEE Grid production service (nodes concerned: gLite-FTA-*, gLite-FTS-*, gLite-UI, gLite-VOBOX, gLite-WN). See <http://glite.web.cern.ch/glite/packages/R3.0/updates.asp>.
 - ◇ Broadcast sent by Nick: UTC is now the standard clock for communications on the EGEE grid. For example, 10:00 UTC is 12:00 Swiss local time in summer, 11:00 Swiss local time in winter.
- ◆ Other:
 - ◇ GGUS system down from 14:00 to 16:00 UTC today.
- Monday 23/10
 - ◆ 9:00 Daily morning meeting
 - ◇ bdii103: alarm *grid-bdii-wrong* triggered (EGEE.BDII status wrong). Fixed.
 - ◇ ce106: alarms *no_contact* and *out_of_memory*. Fixed.
 - ◇ lfc0716: alarms *lfc_db_error*, *lfc_noread* and *lfc_nowrite*.
 - ◇ lxb1134: alarm *lfc_noread*.
 - ◇ lxb1540: alarms *lfc_noread* and *lfc_nowrite*.
 - ◇ lxb7283: alarm *spma_error* (problem with *grid-site-shared* and *grid-site-apache* packages). Fixed by updating CDB.
 - ◇ DB problem which affected gridview (some gridview data have been lost). Investigated by Maarten.
 - ◇ New WMS rb110 for CMS. Will be put in production this morning.
 - ◇ FTS intervention this morning: new patches applied (773, 787, 801, 825, 852).
 - ◆ PRMS Tickets: 2
 - ◇ #375343: *tomcat* alarm on voms102. Ticket assigned to Maria. Cron job related to tomcat didn't restart properly. It was however restarted successfully at the next cron run.
 - ◇ #375721: redirect stderr to */dev/null* for cron job */etc/cron.d/glite-wms-check-daemons.cron* on rb110 (new gLite WMS node put in production today).
 - ◆ broadcasts :
 - ◇ Broadcast sent by Judit: The old FCR web interface is not available anymore. The new version can be found at <https://lcg-fcr.cern.ch:8443/fcr/fcr.cgi>. About top-level EGEE.BDII configuration see broadcast message <http://www.listserv.rl.ac.uk/cgi-bin/webadmin?A2=ind0610&L=lcg-rollout&P=26672>
 - ◇ Broadcast sent by Maria: the CMS users must ignore the VOMRS notification to re-sign the Usage Rules (this message was generated from the test suite of the new VOMRS version which is actually in test).
 - ◆ SECURITY INCIDENT: Several site down last week-end due to the Torque/OpenPBS local root privilege escalation vulnerability.
 - ◆ WLCG/EGEE/OSG operations meeting: see minutes [here](#).

Week 2006-10-16 : Antonio Retico

- Friday 20/10
 - ◆ 9:00 Daily morning meeting
 - ◇ Notifications of Network intervention should be sent also to GMOD mailing list. Request sent to Edoardo Martelli
 - ◆ PRMS Tickets: 1

- ◆ broadcasts :
 - ◇ 10.45: CERN-PROD: Upgrade of the FTS server on Moday
 - ◇ 12.15: CERN-PROD: Change to scheduled service maintenance.
- Thursday 19/10
 - ◆ 9:00 Daily morning meeting
 - ◇ *nothing of note*
 - ◆ PRMS Tickets: 0
 - ◆ broadcasts :
- Wednesday 18/10
 - ◆ 9:00 Daily morning meeting
 - ◇ CERN-PROD tickets:
http://egee-docs.web.cern.ch/egee-docs/ROC_CERN\CERN-PROD_tickets\CERN-tickets-20
 - ◆ 10:00 WLCG-SC meeting
 - ◇ LHCb: priority for set-up of ACLs on LFC gows
 - ◆ 15:00 CCSR Meeting
 - ◇ Change in the schedule of IP re-numbering.
 - ◆ PRMS Tickets: 0
 - ◆ broadcasts :
- Tuesday 17/10
 - ◆ **Reminder:** backup system reorganization possibly affecting the database servers (backup of archive logs). Should not affect database client.
 - ◆ 9:00 Daily morning meeting
 - ◇ RB 'rb109.cern.ch' due to be back this afternoon
 - ◆ PRMS Tickets: 0
 - ◆ broadcasts :
 - ◇ 10:20 : CERN-PROD: Scheduled maintenance for several services
- Monday 16/10
 - ◆ 9:00 Daily morning meeting
 - ◇ network link CERN-FNAL had some alarms on 14th Oct
 - ◇ High load on bdii101/2: The new node should be added today
 - ◇ SRM high load on 13th, 14th Oct
 - ◇ rb109 overheating: Atlas moved to use rb103 as a precaution, but they should be less active this week
 - ◆ PRMS Tickets: 0
 - ◆ broadcasts :
 - ◇ 15.00 : BDII improved at CERN
 - ◇ 14.50 : RB 'rb109.cern.ch' temporarily out of production

Week 2006-10-09 : Diana Bosio

- Friday
 - ◆ rb109 has been taken out of production
 - ◆ It seems that "raid_tw" tickets appear as our services are "class D", which means that a PRMS ticket is open for every FIO standard alarm to contact the service manager. Until procedures for the operators are written and all condtions are satisfied to move the servers to "class E" the service manager needs to open an ITCM ticket for the intervention on the server to happen.
 - ◆ discussed with James how to solve the LFC deadlock: 3 tickets untouched for the week. It seems a procedure is needed.
- Thursday

◆ broadcast sent: "FINAL REMINDER October 16th 2006 stoppage of vomms.cern.ch "

• Wednesday

◆ 9:00 Daily morning meeting

◇ ce107 is waiting for firewall configuration

◇ New procedure to configure the firewall: will be done by cluster and not per node.

You just have to decide which cluster your node belongs to

◆ 10:00 meeting

◇ Simone reported few tickets for ATLAS LFC not being acted upon: started investigating

◇ Maybe the lcg-bdii and the exp-bdii will be merged to a bigger load balancing cluster as it appears that FCR is used also on the lcg-bdii.

◇ Maria highlighted that the "raid_tw" tickets are not followed by the sysadmin/operators if one answers them

◆ broadcast sent: "mon box at CERN mon.cern.ch will be upgraded and rebooted"

• Tuesday

◆ 9:00 Daily morning meeting

◇ new bdii nodes are being prepared

◆ raid_tw tickets started to appear

• Monday 9/10

◆ 9:00 Daily morning meeting

◇ a new CE, ce107 will be put in production later this week

◇ bdii is under heavy load

Week 2006-10-02 : Judit Novak

• Monday

◆ restarting lxb1908 UI (request by Antonio)

◆ EGEE broadcast about short downtime of lxb0725, lxb0726 (atlas UIs)

◆ EGEE Broadcast tool bug -- mails were not sent to all recipients CIC Portal team provided a fix in the afternoon

• Tuesday

◆ ssh was upgraded in the SLC4 repository, which affected many sites

(openssh-3.6.1p2-33.30.9 -> openssh-3.6.1p2-33.30.12) As a result, job submission failed on these sites, as locked (pool) accounts could not be used anymore.

• Friday

◆ following the successful VOMS tests, DB Group got green light to perform patch "

5458753:SENTENCES EXECUTED IN WRONG SCHEMA WITH OCI APPLICATION"
on the production DB.

Week 2006-09-25 : Antonio Retico

• Friday 29/09

◆ 9:00 Daily morning meeting

◇ Harry said that the gLite WMS performances decreases badly using a complex JDL. this is going to be studied at a special session held by Markus at the EGEE06 conference.

◆ PRMS Tickets: 1

◆ broadcasts :

- Thursday 28/09
 - ◆ 9:00 Daily morning meeting
 - ◇ *nothing of note*
 - ◆ PRMS Tickets: 1
 - ◆ broadcasts :

- Wednesday 27/09
 - ◆ 9:00 Daily morning meeting
 - ◇ *nothing of note*
 - ◆ 10:00 WLCG-SC meeting
 - ◇ LHCB wants another "super rb" like rb102
 - ◆ 15:00 CCSR Meeting
 - ◇ backup system reorganization possibly affecting the database servers (backup of archive logs). Should not affect database users. To be done on 17/10. I put a reminder in the gmod journal for that week.
 - ◇ Jan Even says: "gLite software is slowing down the deployment of SLC4. We are going to make an official statement of some sort about it"
 - ◇ Participants to the CCSR meeting are now supposed to post their reports in advance (before Wednesday 2:30). Next week they will explain how and where.
 - ◆ PRMS Tickets: 1
 - ◆ broadcasts :
 - ◇ 11:55 - REMINDER October 16th 2006 stoppage of voms.cern.ch (by Maria Dimou)
 - ◇ 18:15 - SRM endpoint "castorgrid.cern.ch" decommissioned on October the 1st (by Jan Van Eldik)

- Tuesday 26/09
 - ◆ 9:00 Daily morning meeting
 - ◇ James: still problems with the FTS web-service hopefully to be fixed within this week - the new node added to share the load was firewalled
 - ◆ PRMS Tickets: 3

- Monday 25/09
 - ◆ 9:00 Daily morning meeting
 - ◇ rb102, rb 103, rb 109 - /tmp full - fixed manually - FIO is waiting for a procedure from GD
 - ◇ Gavin: overload on the FTS web-service this week - it now runs a certified but not released yet version of the info provider
 - ◆ PRMS Tickets: 1

Week 2006-09-18 : Ian Neilson

- Friday meeting
 - ◆ *nothing of note*
 - ◆ broadcasts :
 - ◇ 18:17 - "FTS web-service at CERN-PROD" (By Gavin McCance)

- Thursday meeting
 - ◆ *nothing of note*
 - ◆ WLCG coordination meeting
 - ◇ *nothing of note*
 - ◆ broadcasts - "ce102.cern.ch problems due to high load"

- Wednesday meeting

- ◆ Remedy 367422 - glite-wms-check-daemons. Closed on open savannah ticket 19929 for yaim config to change.
- ◆ CCSR - intermittent (~each 2 weeks) site firewall blocks needing router reboot. Will be announced by MOD.
- ◆ broadcasts - "CERN FTS Web Server intermittent problems"

- Tuesday meeting
 - ◆ *nothing of note*
 - ◆ broadcasts - "Short SFT interruption"

- Monday 18/09 meeting
 - ◆ Rolandas reported on temporary fix for RB /tmp filling up. Scripts are reporting very verbose mode and are hardwired.
 - ◆ James noted scheduled start of CMS activity and need to closely monitor for gLite RB problems.

- ◆ broadcasts - "Tape Lib intervention 27/9", "Short notice on a LSF restart"

- GmodWeek20060918: Ian Neilson

Week 2006-09-11 : Maite Barroso

- 20060911_week_GMOD_report.htm: GMOD Report on the week of 20060911

Week 2006-09-04 : David Collados

- 20060904_week_GMOD_report.txt: GMOD Report on the week of 20060904
-

This topic: LCG > GmodJournal

Topic revision: r363 - 2007-12-14 - AntonioRetico



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback