

Production Services Work Log (2006)

- 2006-12-31: gdrb06 back in production since the alarm disappeared.
- 2006-12-26: Unrecoverable error on gdrb06 (problem with RAID disk */dev/hde*). This machine has been put in maintenance:


```
Dec 26 18:53:09 gdrb06 kernel: hde: drive_cmd: status=0x61 { DriveReady DeviceFau
Dec 26 18:53:09 gdrb06 kernel: hde: drive_cmd: status=0x61 { DriveReady DeviceFau
Dec 26 18:53:09 gdrb06 kernel: hde: drive_cmd: error=0x04 { DriveStatusError }
Dec 26 18:53:09 gdrb06 kernel: hde: drive_cmd: error=0x04 { DriveStatusError }
```
- 2006-12-24: Upgraded rb101 for patches 900 and 944 asked by ATLAS.
- 2006-12-22: Grant access to the MySQL Database on all the new LCG RBs in production (rb104 to rb107, rb113 to rb115).
- 2006-12-22: Reinstallation from scratch of rb104. This LCG RB now supports the following VOs: *unosat, ops, dteam, sixt, gear, geant4*.
- 2006-12-21: VOMS configuration update on all the machines fully managed by Quattor (rbxxx nodes, voatlas01 and volhcb01).
- 2006-12-20: Reinstallation from scratch of rb108. This machine becomes a gLite WMSLB for VO *ops* (and *dteam*) and has therefore been moved to cluster *gridrb*. This machine has been put in production at the end of the afternoon.
- 2006-12-20: Middleware upgrade (update 11 for gLite 3.0) on all the machines in production.
- 2006-12-20: Job submission blocked on rb104 this morning. This node will be reinstalled next Friday.
- 2006-12-19: monb002 and monb003 are using rb110 as gLite WMS for VO OPS.
- 2006-12-19: Problem with lxn1179. This machine is getting out of memory (*kscand* process running for a while). Simone Campana contacted. Problem fixed by stopping DQ2 site services on lxn1179.
- 2006-12-19: Due to crash on rb114 last night, the MySQL database were corrupted (tables *events* and *states* had some errors). This has been fixed by David.
- 2006-12-19: NO_CONTACT exception triggered on rb114 this morning. This machine has been rebooted by the operators, but there is a problem with the *edg-fmon-agent* service. Under investigation.
- 2006-12-19: SMART_SELFTEST exception triggered on bdii110. There is a read failure problem on disk */c0/p1*. Long smart test launched. Problem fixed by the sysadmins. To check the disks, the following commands should be executed (see */etc/smartd.conf* file):

```
smartctl -l selftest --device=3ware,0 /dev/twe0
smartctl -l selftest --device=3ware,1 /dev/twe0
```

- 2006-12-18: VOMS configuration update on gdrb01 to gdrb11, lxn1183, rb105 to rb108, rb113 to rb115, voatlas01, volhcb (see the following wiki page for more details: <https://uimon.cern.ch/twiki/bin/view/LCG/Trash.LCGVOConfigForSites>).
- 2006-12-18: Reconfiguration of VO *compass* on gdrb01 and gdrb03. The host certificate of the VOMS server for this VO can be found in */etc/grid-security/vomsdir/dgrid-voms.fzk.de*. New parameters are:

```
GRIDMAP_AUTH="ldap://lcg-registrar.cern.ch/ou=users,o=registrar,dc=lcg,dc=org ldap://gridldap1.f
VO_COMPASS_SW_DIR=$VO_SW_DIR/compass
VO_COMPASS_DEFAULT_SE=$CLASSIC_HOST
VO_COMPASS_STORAGE_DIR=$CLASSIC_STORAGE_DIR/SE01/compass
VO_COMPASS_QUEUES="compass"
VO_COMPASS_USERS=ldap://gridldap1.fzk.de/ou=compass,dc=gridka,dc=de
VO_COMPASS_VOMS_SERVERS="vomss://voms.fzk.de:8443/voms/compass?/compass/"
VO_COMPASS_VOMSES="compass dgrid-voms.fzk.de 15010 /O=GermanGrid/OU=FZK/CN=host/dgrid-voms.fzk.de
```

- 2006-12-18: Reinstallation from scratch of rb105. This LCG RB is dedicated to VO *Alice* now.
- 2006-12-16: Package *gd-auth* upgraded to version 1.1 on all the nodes in production and not managed by Quattor.
- 2006-12-15: The configuration of the two UIs monb002 and monb003 has been changed:
 - ◆ In file */opt/edg/etc/edg_wl_ui_cmd_var.conf*, line beginning by *LoggingDestination* has been commented out.
 - ◆ In file */opt/edg/etc/\${VO}/edg_wl_ui.conf* (where VO can be *alice*, *atlas*, *cms*, *dteam* and *ops*), lines beginning by *NSAddresses* and *LBAddresses* have been replaced on monb002 and monb003 respectively by:

```
NSAddresses = {"rb113.cern.ch:7772", "rb115.cern.ch:7772"};      (resp. NSAddresses = {"rb115.cern.ch:7772", "rb113.cern.ch:7772"});
LBAddresses = {"rb113.cern.ch:9000"}, {"rb115.cern.ch:9000"};  (resp. LBAddresses = {"rb115.cern.ch:9000"}, {"rb113.cern.ch:9000"});
```

- 2006-12-15: No new job submissions are now allowed on rb105. This machine will be reinstalled from scratch next Wednesday 13 december.
- 2006-12-14: gLite WMS rb116 and rb117 put in production for VO *Alice* and *LHCb* respectively.
- 2006-12-14: rb115 put in production for SAM job submission. monb003 has been reconfigured to submit jobs on rb115.
- 2006-12-13: Add rb115 entry on bdii103 and bdii104 (site-level BDII's).
- 2006-12-13: Installation and configuration of the second LCG RB rb115 for SAM.
- 2006-12-13: Reinstallation from scratch of rb106. This LCG RB is now dedicated to VO *Atlas* mainly.
- 2006-12-13: Middleware upgrade (update 10 for gLite 3.0) on rb112, rb111, rb110, rb109, rb103, rb102 and rb101.
- 2006-12-12: VO *EELA* supported on rb108 and on all the CEs at CERN. Note that a package named *eela-vomscerts* containing the host certificate of the VOMS server for this VO must be installed.
- 2006-12-12: LB server rb201 put in production again after reinstallation and reconfiguration. This machine supports the following VOs: *atlas*, *alice*, *cms*, *lhcb*, *gear*, *geant4*, *unosat*, *sixt*, *dteam* and *ops*.

WorkLog2006 < LCG < TWiki

- 2006-12-11: Update 10 for gLite 3.0 released. Machines involved are:
 - ◆ LCG RBs from 3.0.4-0 to 3.0.5-0: rb104 to rb108, rb113 and rb114.
 - ◆ Classic SE from 3.0.4-0 to 3.0.5-0: voatlas01, volhcb01 and lxn1183.
 - ◆ UI from 3.0.9-0 to 3.0.10-0: lxb0725, lxb0726, lxb1930 and lxb2007.
 - ◆ VOBOX from 3.0.10-0 to 3.0.11-0: lxn1179.
- 2006-12-11: No job submissions are now allowed on rb106. This machine will be reinstalled from scratch next Wednesday 13 december.
- 2006-12-10: *bkservd* process in infinite loop on rb112. Service restarted and problem fixed.
- 2006-12-08: Installation of the *gd_auth* package on lxb2004 and lxb2008.
- 2006-12-07: Installation of the *gd_auth* package on all the *gdrbxx* nodes, lxb0725, lxb0726 and lxb1930.
- 2006-12-07: lxb2008 down this morning. No particular message on the screen. Machine rebooted and back in production.
- 2006-12-07: New alias *sam-bdii* which is now pointing to *bdii109* and *bdii110* (SAM BDII's) in a load-balanced way. *rb113* is now configured to use *sam-bdii*.
- 2006-12-06: Reinstallation from scratch of *rb107*. This machine will be assigned to CMS.
- 2006-12-05: Two BDII's *bdii109* and *bdii110* put in production for SAM (alias with load-balancing: *sam-bdii*).
- 2006-12-04: Jobs submission on *rb107* blocked (see rule *GD_RB_BLOCKED*). This machine will be reinstalled in few days
- 2006-12-02: *rb111* put in production as gLite WMSLB for VO Alice.
- 2006-12-01: *rb112* put in production as gLite WMSLB for VO LHCb.
- 2006-11-29: Package *glite-wms-ism* is upgraded to version 1.5.15-1 to fix the CE disappearing bug on all gLite WMS's.
- 2006-11-29: *rb109* is back in production. According to sysadmins, the air flow is apparently working fine now.
- 2006-11-28: gLite UI and SAM client installed on *monb002*. Need to have the same configuration on *monb003*.
- 2006-11-28: Big Yvan's mistake on *rb106* (machine rebooted by error).
- 2006-11-27: reinstallation of *rb201* (bad partition disk configuration).
- 2006-11-27: Add *rb113* and *rb114* entries on *bdii103* and *bdii104* (site-level BDII's)
- 2006-11-27: *rb113* put in production as a LCG RB for SAME and SFT.
- 2006-11-27: *rb114* put in production as a LCG RB for LHCb.
- 2006-11-27: It seems that there is an airflow problem on *rb109*. This machine will be shutdown

WorkLog2006 < LCG < TWiki

tomorrow Tuesday 28 november and a vendor call will be opened.

- 2006-11-25: High temperature on the RAID controller detected on rb109 (some tickets opened for this case). I asked to the sysadmins team to check it.
- 2006-11-24: Late in the evening, rb108 crashed once again. It was perhaps due to the high temperature on the battery of the RAID controller. Machine rebooted, and services back in production.
- 2006-11-24: Need to restart service *lcg-mon-job-status status* on gdrb02, gdrb09, gdrb10, rb106 and rb108.
- 2006-11-23: Change `GLITE_WMS_QUERY_TIMEOUT` from the default value, 300, to 480 in `/etc/glite/profile.d/glite_setenv.(c)sh` on rb101 (requested by ATLAS).
- 2006-11-23: New machines assigned to GD:
 - ◆ rb201: new Logging and Bookkeeping node dedicated to experiments (*Atlas* and *CMS*).
 - ◆ rb111: gLite WMSLB dedicated to VO *Alice*.
 - ◆ rb112: gLite WMSLB dedicated to VO *LHCb*.
 - ◆ rb113: LCG RB dedicated to VO *ops*.
 - ◆ rb114: LCG RB dedicated to VO *LHCb*.
- 2006-11-22: Add "ExpiryPeriod = 21600;" in the WM section of `glite_wms.conf` and "Dagmanloglevel = 5;" in the JC section of `glite_wms.conf` on rb101 (requested by ATLAS).
- 2006-11-21: Add user *gianelle* for interactive access on rb104.
- 2006-11-21: Package *lcg-fw* upgraded on gdrbxx nodes, on rb104 to rb108, lxb2003 and lxn1183
- 2006-11-21: Package *lcg-fw* installed on lxb0725, lxb0726, lxb1930, lxb2004 and rlx2008 (UIs for experiments).
- 2006-11-20: Need to restart service *lcg-mon-job-status* on rb106, gdrb01 and gdrb06. Fixed.
- 2006-11-15: Middleware upgrade on rb104 to rb108 for lcg-RB 3.0.3-2 to 3.0.4-0.
- 2006-11-15: Middleware upgrade on lxb0725, lxb0726, lxb1930 and lxb2007 for glite-UI 3.0.8-0 to 3.0.9-0.
- 2006-11-15: Middleware upgrade on lxn1179 for glite-UI 3.0.8-0 to 3.0.9-0, and glite-VOBOX 3.0.9-0 to 3.0.10-0.
- 2006-11-15: Middleware upgrade on voatlas01, volhcb01 and lxn1183 for glite-SE_classic 3.0.4-0 to 3.0.5-0.
- 2006-11-12: rb106 down for some unknown reason. The sysadmin made a reset using the Ctrl+E sequence keys. Machine back in service now.
- 2006-11-10: rb108 back in production (memory module exchanged).
- 2006-11-09: NO_CONTACT alarm on rb108. Black screen and unable to reboot it. VO *gear* supported by rb105.
- 2006-11-09: IP re-numbering successfully done on lxb1930 this morning. Back in production.

WorkLog2006 < LCG < TWiki

- 2006-11-09: Installation of a top-level EGEE.BDII lxb2005 used by gdrb02 for SAM. The bdi service has been stopped on gdrb02 because of timeout errors when gdrb02 was loaded.
- 2006-11-06: lxn1182 blocked (*out of memory* message on the screen). Need to reboot this machine. Fixed.
- 2006-11-04: Hard disk *hda* dead on lxb2001.
- 2006-11-04: Problem on lxn1183 due to the IP re-numbering. The former IP address was in file *etc/hosts*. Fixed.
- 2006-11-03: New gLite UI and top-level BDII lxb0728 in production dedicated to SRMv2 tests.
- 2006-11-02: Castor client upgrade (2.1.1-1 to 2.1.1-4) on volhcb01 and lxb2004.
- 2006-10-31: Castor client upgrade (2.1.1-1 to 2.1.1-4) on lxb1930, lxb2003, lxb2007 and lxb2008.
- 2006-10-31: lxb7283 in an infinite loop (HIGH_LOAD alarm triggered) due to a bug in the gLite middleware (processes glite-lb-bkserverd and glite-lb-logd). All the related services have been restarted. Fixed.
- 2006-10-31: Kernel upgrade on gdrbxx nodes, lxn1183, lxn1179, lxb0725 and lxb0726.
- 2006-10-31: Castor client upgrade (2.1.1-1 to 2.1.1-4) on gdrbxx nodes, lxn1183, lxn1179.
- 2006-10-31: IP re-numbering for some machine in production (gdrbxx, lxn11xx and lxb07xx) this morning.

CondorG was unable to restart on gdrbxx because of the wrong IP address for each host in file */etc/hosts*.

- 2006-10-30: rb108 crashed. Need to reboot this machine. Fixed.
- 2006-10-28: Kernel upgrade on lxb2003.
- 2006-10-26: Kernel upgrade on lxb2004.
- 2006-10-26: Cleaning of all the big files on rb102. Fixed by Di.
- 2006-10-25: kernel upgrade on lxn1176 failed. Need to check.
- 2006-10-25: same problem (than for rb106 yesterday evening) on gdrb08. Fixed,
- 2006-10-24: Problem on rb106 due to a single job submission with JDL requirements that cannot be handled by the WM (see bug #20973 [in Savannah](#). It was also the same problem on gdrb08 yesterday). Fixed.
- 2006-10-24: Minor upgrade on the gLite UI (lxb0725, lxb0726, lxb1930) and the VOBOX lxn1179.
- 2006-10-24: Change the value of the variable APTMAILTO found in file */etc/sysconfig/apt-autoupdate* to RB.Support@cernNOSPAMPLEASE.ch on all the LCG RBs.
- 2006-10-24: Minor middleware upgrade on all the LCG RBs (package glite-rgma-api-python 5.0.3-1 to 5.0.4-1).

WorkLog2006 < LCG < TWiki

- 2006-10-23: Sandbox partition on rb102 full. Ask to CMS before to remove the big files.
- 2006-10-23: rb110 is a new gLite WMS for CMS and is now in production.
- 2006-10-23: Production gridfts cluster updated to latest gLite patch (773, 787, 801, 825, 852).
- 2006-10-23: Problem on gdrb08 due to a single job submission with JDL requirements that cannot be handled by the WM (see bug #20973 [in Savannah](#)). Fixed by Maarten.
- 2006-10-21: Deployment of a patch on the nodes in production to fix a vulnerability (Torque/OpenPBS local root privilege escalation vulnerability).
- 2006-10-20: alarm *spma_error* on lxb7283 (problem with *grid-site-shared* and *grid-site-apache* packages). Fixed by updating CDB.
- 2006-10-20: Major security incident: Torque/OpenPBS local root privilege escalation vulnerability. A lot of sites have been switched down during the week-end.
- 2006-10-20: Installation and configuration of a top-level BDII (without FCR) on gdrb02. gdrb02 does not query lcg-bdii anymore.

```
[root@gdrb02 root]# cat /opt/bdii/etc/bdii.conf
.....
BDII_AUTO_MODIFY=no
BDII_UPDATE_LDIF=
.....
[root@gdrb02 root]#
```

- 2006-10-19: CAs upgraded to version 1.10-1 on all nodes in production.
- 2006-10-18: Alias myproxy-fts points now to a new machine prod-px-fts. Former myproxy-fts (lxb0728) has been removed from production.
- 2006-10-18: rb109 back in production (cooling problem solved).
- 2006-10-18: middleware upgrade on rb104 to rb108 (lcg-RB 3.0.3-1 to 3.0.3-2), and on the special UIs used by the experiments (gLite-UI 3.0.6-0 to 3.0.7-0)
- 2006-10-16: Remove archivers lxn1190, lxn1191 and lxn1193 from production (service *tomcat* stopped and cron job */etc/cron.d/check-tomcat* disable).
- 2006-10-16: Set *BOOTPROTO=dhcp* in file */etc/sysconfig/network-scripts/ifcfg-eth{0,1}* on several machines in production in order to prepare the IP renumbering planned from 2006-10-31 to 2006-11-15:
 - ◆ Date: 31/10/2006 from 08:00am to noon
 - ◇ lxb0725 (gliteUI for Atlas).
 - ◇ lxb0726 (glite UI for Atlas).
 - ◇ lxb0728 (myproxy for FTS - this machine will be replaced soon by a new mid-range server).
 - ◇ gdrb01 to gdrb11 (LCG RBs).
 - ◇ lxn1179: VOBOX for Atlas.
 - ◇ lxn1180: SAM server.
 - ◇ lxn1181: SAM server backup.
 - ◇ lxn1182: SAM client.
 - ◇ lxn1183: Classical SE.

WorkLog2006 < LCG < TWiki

- ◆ Date: 08/11/2006 from 08:00am to noon
 - ◇ lxb1930 (gLite UI for CMS).
- ◆ Date: 15/11/2006 from 08:00am to noon
 - ◇ lxb2003 (Classic SE for LHCb).
 - ◇ lxb2004 (UI for LHCb).
 - ◇ lxb2008 (UI for LHCb).
- 2006-10-14: processes related to job controller in infinite loop on rb106. Need to restart service *edg-wl-jc*. Fixed.
- 2006-10-13: change email adress for variable MAILTO in file */etc/cron.d/glite-wms-check-daemons.cron* on rb101 to rb103. This value should be set to *wms.support@cernNOSPAMPLEASE.ch* for this type of nodes. Need to do it on rb109 as well.
- 2006-10-13: problem with one of the RAID disk on rb104 fixed.
- 2006-10-13: Minor upgrade of the middleware on RBs rb104 to rb109 (lcg-RB 3.0.3-0 to 3.0.3-1).
- 2006-10-11: temperature too high on the RAID disks on rb109:

```
[root@rb109 root]# tw_cli info c0
```

Unit	UnitType	Status	%Cmpl	Stripe	Size(GB)	Cache	AVerify	IgnECC
u0	RAID-1	OK	-	-	149.001	OFF	OFF	OFF
u1	RAID-1	OK	-	-	232.82	OFF	OFF	OFF
u2	RAID-1	OK	-	-	232.82	OFF	OFF	OFF
u3	RAID-1	OK	-	-	232.82	OFF	OFF	OFF

Port	Status	Unit	Size	Blocks	Serial
p0	OK	u0	153.38 GB	321672960	VDBE1BTCE521MP
p1	OK	u1	232.88 GB	488397168	VDB41BT4DE3GZC
p2	OK	u1	232.88 GB	488397168	VDK41BT4DWV9DK
p3	OK	u2	232.88 GB	488397168	VDK41BT4DYP4ZK
p4	OK	u2	232.88 GB	488397168	VDK41BT4DX10JK
p5	OK	u3	232.88 GB	488397168	VDK41BT4DYWUK
p6	OK	u3	232.88 GB	488397168	VDK41BT4DXNVRK
p7	OK	u0	232.88 GB	488397168	VDK41BT4DY0BRK

Name	OnlineState	BBUReady	Status	Volt	Temp	Hours	LastCapTest
bbu	On	No	Fault	OK	TooHigh	0	xx-xxx-xxxx

- 2006-10-09: Partition */tmp* full (no more inodes available) on rb101 and rb103. Fixed.
- 2006-10-08: Partition */tmp* full (no more inodes available) on rb103. Fixed.
- 2006-10-07: Partition */tmp* full (no more inodes available) on rb103. Fixed by deleting all the empty files in this directory.
- 2006-10-06: Reconfiguration of the RAID disks on rb101, rb103 and rb109 in order to increase the number of inodes available on the partitions.
- 2006-10-06: Problem with one of the RAID disks on rb104:

```
[root@rb104 root]# smartctl -l selftest --device=3ware,3 /dev/twa0
smartctl version 5.33 [i686-redhat-linux-gnu] Copyright (C) 2002-4 Bruce Allen
Home page is http://smartmontools.sourceforge.net/
```

WorkLog2006 < LCG < TWiki

```
=== START OF READ SMART DATA SECTION ===
```

```
SMART Self-test log structure revision number 1
```

Num	Test_Description	Status	Remaining	LifeTime (hours)	L
# 1	Short offline	Completed: read failure	20%	4496	1
# 2	Short offline	Completed: read failure	40%	4472	1
# 3	Short offline	Completed: read failure	10%	4448	1
# 4	Short offline	Completed: read failure	40%	4424	1
# 5	Short offline	Completed: read failure	30%	4400	1
# 6	Extended offline	Completed: read failure	90%	4375	1
# 7	Short offline	Completed: read failure	10%	4352	1
# 8	Short offline	Completed: read failure	10%	4328	1
# 9	Short offline	Completed: read failure	10%	4304	1
#10	Short offline	Completed: read failure	40%	4280	1
#11	Short offline	Completed: read failure	40%	4256	1
#12	Short offline	Completed: read failure	40%	4232	1
#13	Extended offline	Completed: read failure	90%	4207	1
#14	Short offline	Completed: read failure	30%	4184	1
#15	Short offline	Completed: read failure	10%	4160	1
#16	Short offline	Completed: read failure	40%	4136	1
#17	Short offline	Completed: read failure	30%	4112	1
#18	Short offline	Completed: read failure	40%	4087	1
#19	Short offline	Completed: read failure	40%	4063	1
#20	Extended offline	Completed: read failure	90%	4038	1
#21	Short offline	Completed: read failure	20%	4015	1

```
[root@rb104 root]#
```

A ticket has been opened for this case:

<http://cern.ch/helpdesk/problem/CT371446&email=sysadmin-team@cern.ch>. The disk will be replaced in the next few days.

- 2006-10-05: Pb with the workload manager and the proxy renewal services on gdrb01. Need to restart these services. Fixed.
- 2006-10-05: Castor client upgraded from version 1.7.1.5-1 to version 2.1.1-1 on volhcb01 (specific LHCb configuration, do not need to do the same on voatlas01 according to Simone).
- 2006-10-05: Reconfiguration of the RAID disks on rb102. There was not enough inodes in one of the partitions.
- 2006-10-05: upgrade of the CASTOR packages on classic SE volhcb01.
- 2006-10-03: manually upgrade of the CAs to version 1.9 on some machines in production. It has not been done automatically because of the change of the APT repository.
- 2006-09-27: We have now many gdrbxx machines in degraded mode:
 - ◆ gdrb03 (failure on hdg)
 - ◆ gdrb04 (failure on hde)
 - ◆ gdrb06 (failure on hde)
 - ◆ gdrb08 (failure on hdg)
 - ◆ gdrb10 (failure on hdg).
 - ◆ gdrb05 (failures on hde and hdg) and not used anymore in production as a RB.
- 2006-09-27: Software RAID-I on gdrb06 in degraded mode (problem on hde). The machine has been rebooted successfully.
- 2006-09-27: All the exceptions found on the gdrbxx nodes have been solved. It was a problem with the lemon packages installed. I did the following thing to solve this problem:

```
for x in ` /usr/bin/seq -w 4 11 `; do
```


WorkLog2006 < LCG < TWiki


```
go gdrb$x "REPO=http://swrep.cern.ch/swrep/i386_slc3 ;
wget $REPO/lemon-host-check-1.1.0-7.noarch.rpm $REPO/lemon-sensor-exception-1.
$REPO/lemon-sensor-sure-1.0.1-2.noarch.rpm $REPO/lemon-sensor-fio-1.2-10.noarch.r
rpm -Uvh lemon-sensor-fio-1.2-10.noarch.rpm lemon-sensor-sure-1.0.1-2.noarch.r
lemon-sensor-exception-1.2.1-2.i386.rpm lemon-host-check-1.1.0-7.noarch.rpm ;
ccm-fetch ;
ncm-ncd --co fmonagent ";
done
```

- 2006-09-27: Fixed the problem with the *smartd_wrong* exception on gdrbxx nodes by putting the following content in file */etc/smartd.conf* (this file was empty):

```
/dev/hda -d ata -a -I 194 -I 7 -s (S/../../[^6]/01|L/../../6/00)
/dev/hde -d ata -a -I 194 -I 7 -s (S/../../[^6]/01|L/../../6/00)
/dev/hdg -d ata -a -I 194 -I 7 -s (S/../../[^6]/01|L/../../6/00)
```

- 2006-09-26: Same problem on rb103 than for rb101 this morning. Fixed.
- 2006-09-26: Middleware upgrade on lxn1179 (VOBOX for Atlas). Current version is now 3.0.6.
- 2006-09-26: There was a problem with the gLite Logging and Bookkeeping processes (infinite loop ?) on rb101. It was impossible to stop cleanly these processes, so I killed them by hand using kill -9. I then stopped and restarted all the services and the load on rb101 is now ok. Fixed.
- 2006-09-25: rb108 highly overloaded due to LHCb.
- 2006-09-25: Need to restart service edg-fmon-agent on gdrb06 and gdrb07. There are a lot of exceptions on the gdrbxx machines. Need to be fixed.
- 2006-09-25: Manually added new CA repository and updated CAs on gdrbxx machines. (vvidic)
- 2006-09-25: VO *unosat* configured and supported on gdrb03, gdrb09 and gdrb10.
- 2006-09-22: Network server on rb102 dead. Need to restart it by hand. Fixed.
- 2006-09-13: update of the CAs to version 1.9.
- 2006-09-06: Security updates done on rb104 to rb108, and on voatlas01 and volhcb01.
- 2006-08-31: False NO_CONTACT alarm on all the gdrbxx machines.
- 2006-08-31: Add AFS access to user *fprelz* on rb107.
- 2006-08-30: job controller restarted on rb106 (2 gaph processes in infinite loop). Fixed.
- 2006-08-30: minor upgrade of the classic SE voatlas01 and volhcb01 (current version is now 3.0.3).
- 2006-08-28: daemon *ntpd* stopped on lxb1930. Need to restart it and to set the runlevel information for this service (was in off state). Fixed.
- 2006-08-28: The gLite startup script failed partially during the boot sequence on rb101. Need to stop and restart all services. Fixed.
- 2006-08-27: Service *edg-wl-wm* stopped on gdrb01. Need to restart it. Fixed.
- 2006-08-25: It is now possible to access to gdrb07, gdrb09, gdrb10 and gdrb11 via Kerberos. The

problem was due to a misconfiguration of the Kerberos database.

- 2006-08-24: End of the kernel upgrade for all the machines in production and managed by GD.
- 2006-08-24: Kernel upgraded on gdrb01, gdrb03, gdrb06, gdrb08, gdrb09, gdrb10, lxb0725, lxb0726, lxb1930, lxn1183, rb105, rb106 and volhcb01.
- 2006-08-23: services *edg-wl-wm* and *edg-wl-proxyrenewal* restarted on rb105. Fixed.
- 2006-08-23: service *rfiod* stopped on gdrb01 to gdrb11.
- 2006-08-22: Kernel upgraded on gdrb02, gdrb04, gdrb07, lxb2003, lxb2004, lxb2008, lxb7026, rb107 and volhcb01.
- 2006-08-21: Migration of all files found on the raid disk sdb to sda (except directory */var/edgwl/SandboxDir/* which is still on sdb) on rb108. This machine also supports VO LHCb now. We would like to compare the performance of rb107 and rb108 because we suspect that the raid disk could be a bottleneck (raid misconfiguration).
- 2006-08-21: Process *edg-wl-wm* (worload manager) not running on gdrb01. Need to restart it manually. Fixed.
- 2006-08-20: Modify manually configuration file */etc/cron.daily/slocate.cron* and */etc/updatedb.conf* used by *slocate* and *updatedb* in order to avoid the indexation of files contained in *rb-state* directory. Need to check if these configuration files are replaced during the update of the package *slocate*.
- 2006-08-17: False NO_CONTACT alarm on all the gdrbxx machines.
- 2006-08-16: False NO_CONTACT alarm on all the gdrbxx machines.
- 2006-08-16: Need to restart some services (*edg-wl-wm* and *edg-wl-proxyrenewal*) on rb105. Fixed.
- 2006-08-16: Need to restart *edg-fmon-agent* on gdrb08. The lemon monitoring was off on this machine during 6 days. Fixed.
- 2006-08-15: Machines *myproxy-fts*, gdrb11, rb108 and rb104 rebooted with the new kernel.
- 2006-08-10: Beginning of the kernel upgrade (kernel *2.4.21-47.EL.cernsmp*) on all the machines in production. Still need to reboot the machines. Ask to EIS when it is possible to do it.
- 2006-08-09: jobs monitoring tool deployed on all the RBs machines (gdrb01 to gdrb11, and rb104 to rb108) in order to monitor the number of running and idle jobs in the condor queue. Results available at the following link (only from inside the Cern): <http://lxb1524.cern.ch/plots.html> .
- 2006-08-09: file systems from the disk servers were not remounted on lxn1183 after the last reboot. Fixed by Maarten.
- 2006-08-08: CAs upgraded to version 1.8-1 on all machines in production.
- 2006-08-08: Remove package *shell-compat* on rb104 to rb108 because there is a conflict between packages *shell-compat* and *cern-compat-locallinks*). Fixed.
- 2006-08-01: Some modifications made by Ulrich on the globus-gridftp startup script on volhcb01 (see mail sent by Ulrich on 2006-08-01):

WorkLog2006 < LCG < TWiki

- ◆ changed startup level from 55 to 99.
- ◆ IP tables setup added in the start up script to make sure that the number of requests accepted is limited.

- 2006-07-31: Mirror broken on volhcb01. Fixed.

- 2006-07-28: major power cut at CERN.

- 2006-07-28: volhcb01 down. Need to reboot it. Fixed.

- 2006-07-27: CERN-CIC site definitely closed. Resources have migrated to CERN-PROD.

- 2006-07-27: WM was dead on gdrb01, input queue 4MB. WM restarted. (vvidic)

- 2006-07-27: New BDIIs in production (bdii105 and bdii106) used by experiments (Freedom of Choice for Resources -FCR- running on them). The aliases used for these machines are: exp-bdii, atlas-bdii and prod-bdii-exp.

- 2006-07-26: Need to reboot lxb2007. Fixed.

- 2006-07-23: Some host certificates expired on several RBs (gdrb07 to gdrb11). These certificates have been replaced and we experienced some problems with the LM and the JC services due to CondorG. Fixed.

- 2006-07-24: Wiki page GmodRoleDescription created with the definition and duties fo the GMOD.

- 2006-07-20: rb102 fully managed by GD now.

- 2006-07-19: lxn1183 back in production this morning.

- 2006-07-18: lxb1133 (alias lfc-lhcb) put in maintenance. Need to check the connexions between this machine and the database.

- 2006-07-18: Problem with a fuse in the CC. lxn1183 is down due to this problem.

- 2006-07-18: VO *unosat* configured on lxn1183.

- 2006-07-12: Need to reboot lxn1179 (VOBox for Atlas). Fixed.

- 2006-07-11: Security incident. Need to remove a ssh public key on all machines in production.

- 2006-07-11: Problem found with the log monitor service on gdrb03. This service opened a lot of CondorG files and exceeded the number of file descriptors dedicated to it (ie. 1024). Other RBs checked since this problem could occur to them as well. Fixed.

- 2006-07-11: Unable to connect to gdrb11. Need to reboot this machine. Fixed

- 2006-07-10: End of the CERN-CIC site (this site will be removed from GOC DB).

- 2006-07-10: lxn1183 (Classic SE used by Atlas, Unosat and Geant4) moved from CERN-CIC to CERN-PROD site.

- 2006-07-10: lxn1179 (VOBox for Atlas) moved from CERN-CIC to CERN-PROD site.

- 2006-07-10: lxn1184 (CE) and lxb2001 (Monbox) removed from CERN-CIC site. Services stopped on lxn1184.
- 2006-07-10: Need to reboot lxn1179 (VOBox for Atlas). *Out of memory* error message on the screen. Fixed.
- 2006-07-04: *http* and *https* servers configured and running on voatlas01.
- 2006-07-04: New classic SE voatlas01 (alias: atlas-logs) in production. This SE will be used by Atlas to store their log files.
- 2006-06-29: *http* and *https* servers configured and running on volhcb01.
- 2006-06-28: Lemon update on gdrb01 to gdrb11. Remove gdrb05 from the GD-RB subcluster (see <http://lemonweb.cern.ch/lemon-status/info.php?cluster=lxnoq%25GD-RB>).
- 2006-06-28: New classic SE volhcb01 (alias: lhcb-logs) in production. This SE will be used by LHCb to store their log files.
- 2006-06-28: Since this morning the sustained load on the three new WMS (rb101 to rb103) is around 20. Need to stop (some processes have to be killed by hand with -9, especially *glite-lb-bkserv* and *glite-lb-logd*) and restart all the services on these nodes. One bug found (see). Fixed.
- 2006-06-27: UI lxplus configured to support the new WMS nodes rb101 to rb103.
- 2006-06-27: rb101 to rb103 (gLite WMS) are now in production. The VOs supported are:
 - ◆ rb101 (alias wms-atlas): *dteam*, *ops* and *atlas*.
 - ◆ rb102 (alias rb-cms): *dteam*, *ops* and *cms*.
 - ◆ rb103 (alias rb-alice, rb-lhcb, rb-shared): *dteam*, *ops*, *gear*, *unosat*, *sixt*, *na48* and *geant4*.
- 2006-06-26: Problem with the configuration of the CEs ce101 and ce102. Special accounts *xxxprd* were not included in the */etc/security/limits.conf* file, generating problem on the RBs which were unable to determine the exit status of the user's job (see for example ggus ticket #9743).
- 2006-06-23: Inconsistency between two rpms on rb104 to rb108:
 - ◆ Package *edg-fabricMonitoring-agent-2.12.1-1* coming from Quattor/Lemon.
 - ◆ Package *edg-fabricMonitoring-2.5.4-4* (coming from the middleware, more precisely metapackage lcg-RB) which provides the client for Lemon and Gridice.

We fixed the situation by **1)** Installing package *edg-fabricMonitoring-2.5.4-4* via *apt-get*; **2)** Removing package *edg-fabricMonitoring-agent-2.12.1-1*; **3)** Reinstalling package *edg-fabricMonitoring-agent-2.12.1-1*.

In this way, we now avoid errors when we are doing an *apt-get upgrade* or an *apt-get dist-upgrade*. See *EdgFabricMonitoringConflictWithRBs* for more details. Fixed.

- 2006-06-23: Upgrade of *lcg-vomscerts-4.2.0-1* not done automatically on rb104 to rb108 (this problem has been discovered thanks to ggus ticket #9476). The auto-update of the middleware packages has not been done before because the package *apt-autoupdate* has not been installed by default by quattor. Fixed.
- 2006-06-22: file server lxfsrk524 in bad shape this morning (services *nfs* and *portmap* down). I restarted these services. Need to reboot lxb2003, lxb2004, lxb2008 and lxn1183 because of the problem with lxfsrk524. Fixed.

- 2006-06-22: Due to the power cut at Cern last night, the apt-autoupdate of the CAs has failed on all the machines in production managed by GD (host grid-deployment.web.cern.ch was unreachable). I updated these packages manually this morning. Fixed.
- 2006-06-22: major power cut last night in the CC. All the machine in production down and restarted this morning. Some services restarted manually on the RBs (rb105, rb106, rb107). The other RBs (gdrbxx) have not been touched by this power cut.
- 2006-06-19: 100 new pool account created on rb104 to rb107 for VO *alice*, *atlas*, *cms* and *lhcb* respectively. Pool accounts *aliceprd*, *atlasprd*, *cmsprd* and *lhcbprd* also created on the same machines.
- 2006-06-19: new rpm *CERN-CC-tmpwatch-1.3-2* installed on rb104 to rb108. The previous version of this rpm caused some problems on these RBs last week (see 2006-06-16).
- 2006-06-17: lxb2003 frozen. It seems to be a problem with AFS (kernel panic on screen). Machine rebooted only on 2006-06-19 due to the week-end. Fixed now. Note that a new machine has been requested to FIO two weeks ago.
- 2006-06-16: Problem detected on all the RBs (rb104 to rb108) due to the bug in the CERN version of the "tmpwatch" system rpm (see 2006-06-13). A backlog of some 8000 logging events had built up (in */var/tmp* on rb107) and was only getting processed very slowly, because there was competition from a continuous stream of new job submissions. Fixed by Maarten.
- 2006-06-14: VO *ops* configured on rb104 to rb108. Account *opssgm* has been also created.
- 2006-06-13: Problem on rb104 to rb108. Symbolic link */var/tmp* (which points to */rb-state/var/tmp*) disappeared and has been replaced by regular file */var/tmp* (due to cron job */etc/cron.hourly/tmpwatch.sh*). Fixed (patch proposed by Maarten).
- 2006-06-12: FTS service stopped for DB and Castor upgrade. Resumed OK. (mccance)
- 2006-06-06: FTA_WRONG_SURE alarm installed on gridfts cluster (mccance)
- 2006-06-06: lxn1183, lxb2005, lxb2007 and lxn1194 upgraded to gLite 3.0.0.
- 2006-06-01: Swap full on gdrb03. Fixed.
- 2006-06-01: lxb2009 (i.e. mon.cern.ch) removed from production. Alias mon.cern.ch will point to monb001 (FIO managed official monbox).
- 2006-06-01: lxb0725 and lxb0726 upgraded to gLite 3.0.0.
- 2006-05-31: Same persistent problem with lxb2003 due to a mis-configuration of the disk server lxfsrk524. This problem seems to be fixed now.
- 2006-05-30: RBs rb104 to rb108 put in production. VOs supported are (VO *ops* not supported yet):
 - ◆ rb104 (alias rb-*alice*): *dteam* and *alice*.
 - ◆ rb105 (alias rb-*atlas*): *dteam* and *atlas*.
 - ◆ rb106 (alias rb-*cms*): *dteam* and *cms*.
 - ◆ rb107 (alias rb-*lhcb*): *dteam* and *lhcb*.
 - ◆ rb108 (alias multi-vo-rb): *dteam*, *gear*, *unosat*, *sixt*, *na48* and *geant4*.
- 2006-05-30: new *myproxy* server named prod-px with alias myproxy. This node is now managed by FIO. Former myproxy server lxn1192 will be retired.

WorkLog2006 < LCG < TWiki

- 2006-05-29: Problem with lxb2003 again (out of memory messages due to gridftp connexions). Need to reboot it. Fixed.
- 2006-05-26: Channel cleanup: Unused Tier-0 to Tier-1 / Tier-1 to Tier-0 / wildcard channels removed from gridfts (mccance)
- 2006-05-26: Added `ops` VO to gridfts (mccance)
- 2006-05-24: Created gLite environment on fts10[6-8] (mccance)
- 2006-05-24: Problem with the rpm database on lxb2003. Need to regenerate this db. Fixed.
- 2006-05-23: problem during the job submission on gdrb09. Should be fixed now.
- 2006-05-23: Partition `/var` full on lxb2003. Fixed.
- 2006-05-23: Need to start service `edg-wl-proxyrenewal` on gdrb03. This service was stopped. Fixed.
- 2006-05-23: Scheduled intervention on FTS to change channel definitions to use GOCDB site names. (mccance)
- 2006-05-22: Pb on gdrb03. Unable to submit jobs. Fixed.
- 2006-05-20: Same problem on lxfsrk524. Fixed.
- 2006-05-19: Unable to write (for normal users) on the disk server lxfsrk524 mounted from lxn1183. Need to reboot the disk server. Fixed.
- 2006-05-19: DTEAM background transfers switched over to FTS validation cluster on fts00[1-5] for Oracle 10gR2 validation (mccance)
- 2006-05-18: FTS history cleanup DBMS job installed on lcg_fts_prod account (mccance)
- 2006-05-17: FTS GridView data collection trigger installed on lcg_fts_prod account (mccance)
- 2006-05-17: All FTS channels switched Active again (mccance)
- 2006-05-16: All FTS channels switched Inactive to avoid draining queue (Castor still recovering from power failure) (mccance)
- 2006-05-16: update for edg-mkgridmap (2.6.1) on all the RBs (gdrb01 to gdrb11). This update fixes a problem which occurs when a VOMS or LDAP server is unavailable at the time the grid-mapfile is created (every 6 hours).
- 2006-05-16: All LHCB jobs are failing registration to LFC since yesterday CERN power problem. Fixed ?
- 2006-05-16: Unable to ping gdrb10 this morning. Bad reboot of the machine after EXTFSWARNING alarm. Fixed
- 2006-05-16: power supply problem on lxn1181 (spare proxy). ITCM ticket generated.
- 2006-05-16: major power cut at Cern. Need to check all the machines in production, especially the RBs.

- 2006-05-16: FTS now publishing load-balanced alias in EGEE.BDII as prod-fts-ws.cern.ch (mccance)
- 2006-05-16: Bad CAs on the machines in production (due to apt-autoupdate). Need to re-install them by hand to version 1.2-1.
- 2006-05-16: Pb with raid disk /dev/hde on gdrb10 this morning. Machine rebooted and no more problem detected. MFixed.
- 2006-05-15: FTS servers on production fts101, fts102, fts103, fts104, fts105 upgraded to 3.0 release. Pilot test cluster (fts001 - fts006) upgraded (mccance)
- 2006-05-15: FTS servers (web-service nodes only) on production fts103 and fts104 upgraded to 3.0 release (mccance)
- 2006-05-09: Partition dedicated to Atlas full on lxn1183 (size of partition: 1.8TB).
- 2006-05-09: Problem with the two CE ce101 and ce102 due to a configuration error. Fixed this morning.
- 2006-05-06: Problems on ce102 (VM_KILL, GRID_GRIS_WRONG, NO_CONTACT, MIRROR_BROKEN). Will be fixed by FIO.
- 2006-05-02: AFS error on gdrb10. Need to reboot this machine. Fixed.
- 2006-04-24: bdii103 and bdii104: Final configuration of LFC - now using GRIS on LFC nodes to publish information (jamesc).
- 2006-04-24: bdii103 and bdii104: entries removed for RLS since the nodes have been taken out of production by IT-PSS (jamesc).
- 2006-04-24: /var filled up on lxb2003. Fixed and run rotate script for /var/log/messages (vvidic)
- 2006-04-24: upgraded lcg-CA and lcg-yaim rpms on lxb2003 (vvidic)
- 2006-04-20: centralized firewall configuration installed on myproxy, myproxy-fts, lxn1181 (spare proxy), lxn1178, lxb2009, lxn1190, lxn1191 and lxn1193.
- 2006-04-17: gdrb08 blocked. Need to reboot it. Fixed.
- 2006-04-14: ce102. blocked. Need to reboot it and to restart globus MDS. Fixed.
- 2006-04-12: srm.cern.ch published by prod-bdii (i.e. bdii103 and bdii104). For this, file *cern-cic-static.sh* updated.
- 2006-04-05: Updated lcgdm-mkgridmap.conf on all LFC nodes to be the same as the one generated by yaim. *James*
- 2006-03-31: VO *ops* configured on all the RBs (gdrb01 to gdrb11).
- 2006-03-30: lfc005 updated via Quattor to upgrade 4 RPMS: LFC-server-oracle, LFC-client, LFC-interfaces and lcg-dm-common to version 1.5.5-2.
- 2006-03-30: VO *ops* configured on gdrb02, gdrb11, lxn1183 and lxn1184. Need to configure the other RBs.

- 2006-03-30: kernel upgraded in lfc001 to version: 2.4.21-40.EL.cernsmp
- 2006-03-30: kernel upgrade on all production nodes.
- 2006-03-30: lfc001 updated via Quattor to upgrade 4 RPMS: LFC-server-oracle, LFC-client, LFC-interfaces and lcg-dm-common to version 1.5.5-2.
- 2006-03-30: Move all services registered from CERN-SC to CERN-PROD (i.e. LFC nodes, myproxy-fts, castorgridsc and prod-fts-ws).
- 2006-03-29: kernel upgrade on all RBs (gdrb01 to gdrb11).
- 2006-03-29: update of the site BDII's bdii103 and bdii104 to add the `ops` VO to `lfc-shared.cern.ch` (updated by James).
- 2006-03-29: update of the RBs gdrb01 to gdrb11 to patch #701 [↗](#).
- 2006-03-27: alarm LFC_DB_ERROR triggered on lfc009.
- 2006-03-27: High load (> 64 🤖) on lxb2003 due to a lot of connections via gridftp. Need to reboot it. Fixed.
- 2006-03-23: CROND_WRONG alarm triggered on gdrb04. Need to kill some processes hanging related to cron daemon `edg-mkgridmap`. Fixed.
- 2006-03-23: centralized firewall configuration installed on lxb2003.
- 2006-03-23: VO `ops` configured on lfc-shared, lfc-dteam-test and lfc001 (file `/opt/lcg/etc/lcgdm-mapfile` modified).
- 2006-03-22: after the reboot of lfc-shared (due to kernel upgrade), configuration for VO `unosat` disappeared. Files `edg-mkgridmap.conf` and `lcgdm-mkgridmap` updated by James and Maarten. Fixed.
- 2006-03-22: service `pbs_mon` stopped on lxb2003 and lxn1183 (not needed for this type of nodes).
- 2006-03-22: centralized firewall configuration installed on lxn1184 and lxn1183.
- 2006-03-22: kernel upgrade done on all the LFC nodes (lfc001 to lfc011).
- 2006-03-22: all the LFC nodes (lfc001 to lfc011) have been upgraded to LFC 1.5.4.
- 2006-03-21: service `mysql` stopped on lxn1184 (not needed for this type of node).
- 2006-03-20: Current version of LFC installed on the LFC nodes (latest version is LFC 1.5.4):
 - ◆ lfc001: LFC 1.4.1.
 - ◆ lfc002 (lfc-atlas-test): LFC 1.5.4.
 - ◆ lfc003 (lfc-cms-test): LFC 1.5.4.
 - ◆ lfc004 (lfc-atlas): LFC 1.4.1.
 - ◆ lfc005 (lfc-dteam-test): LFC 1.4.5.
 - ◆ lfc006 (lfc-shared or lfc-dteam): LFC 1.5.4.
 - ◆ lfc007 (lfc-alice): LFC 1.4.1.
 - ◆ lfc008 (lfc-atlas): LFC 1.4.1.
 - ◆ lfc009 (lfc-cms): LFC 1.4.1.
 - ◆ lfc010 (lfc-lhcb): LFC 1.4.1.

◆ lfc011 (lfc-lhcb-ro): LFC 1.4.1.

- 2006-03-17: Kernel needs to be upgraded on all the LFC nodes (lfc001 to lfc011). Planed for next week.
- 2006-03-17: Upgrade of LFC on lfc-shared to version 1.5.4.
- 2006-03-16: Misconfiguration of the SE name on lxn1184. File *site-info.def* modified and yaim reexecuted. Fixed.
- 2006-03-15: Need to restart maui service on lxn1184 (Job submission via SFT failed for this reason). Fixed.
- 2006-03-15: Start to deploy the new centralized firewall configuration on the production nodes (gdrbxx, lxn1194, lxb2005 and lxb2007).
- 2006-03-15: VO *gear* configured and now supported on gdrb01, gdrb03 and lxn1183.
- 2006-03-14: Update of the *edg-mkgridmap.conf* on all nodes in the LFC cluster (lfc001 to lfc011) to use VOMS (they were just using LDAP and were missing new people).
- 2006-03-13: New alias lhcbui created which points to lxb2004 (UI for LHCB experiments).
- 2006-03-13: VO *gear* configured and now supported on lfc-shared.
- 2006-03-11: lxb2008 installed as a new UI for LHCB experiments.
- 2006-03-11: Problem with AFS on lxb2004 (machine blocked, kernel panic on the screen). Need to reboot it. Fixed.
- 2006-03-10: Upgrade of LFC on lfc-atlas-test and lfc-cms-test to version 1.5.4.
- 2006-03-09: Installation of the grid log retention tool (<https://twiki.cern.ch/twiki/bin/view/LCG/LogRetention>) on gdrb01 to gdrb10 (except gdrb05).
- 2006-03-08: Installation of new package *lcg-mon-job-status-2.0.7-1_sl3.noarch.rpm* on all the RBs in production (Patch #690).
- 2006-03-08: Bug found in proxy configuration. This affects myproxy and myproxy-fts (cf. file */etc/init.d/myproxy-generate-config.pl*). Fixed by Maarten.
- 2006-03-07: Installation of the grid log retention tool (<https://twiki.cern.ch/twiki/bin/view/LCG/LogRetention>) on gdrb11.
- 2006-03-07: installation of LCG 2.7.0 finished for all the production nodes.
- 2006-03-06: VO Atlas supported on lxn1183 (lxf5592 mounted on this machine).
- 2006-03-06: Problem with the PBS server on lxn1184. Need to shutdown the gatekeeper, mds and bdi. These services have been restarted successfully. Fixed.
- 2006-03-06: LCG 2.7.0 installed on myproxy-fts and lxb2037. UI for Atlas lxb0725 and lxb0726 upgraded too.

WorkLog2006 < LCG < TWiki

- 2006-03-03: Almost all the machines in production have been upgraded to LCG 2.7.0. Only some UIs for experiments (lxb0725, lxb0726, lxb1930, lxb2004 and lxb2037) and myproxy-fts have still LCG2.6.0 installed.
- 2006-02-28: gdrb04, gdrb06, gdrb07 and gdrb08 upgraded to LCG 2.7.0.
- 2006-02-28: The following nodes have been switched off and will be reinstalled from scratch (done the 2006-03-08):
 - ◆ lxn1177 (Prod RB)
 - ◆ lxn1186 (Prod RB)
 - ◆ lxn1188 (Test zone RB)
 - ◆ lxn1185 (CMS RB)
 - ◆ lxb2008 (EGEE.BDII for LHCB)
 - ◆ lxn1187 (EGEE.BDII for CMS)
 - ◆ lxn1189 (Test zone EGEE.BDII)
- 2006-02-24: Need to restart two daemons on gdrb08. Fixed.
- 2006-02-24: Raid disk full on gdrb08. Fixed.
- 2006-02-24: gdrb01, gdrb02, gdrb03, gdrb09 and gdrb10 upgraded to LCG 2.7.0.
- 2006-02-23: Beginning of the migration to LCG 2.7.0 on all the machines in production (gdrb11 updated).
- 2006-02-22: monb001 (New monbox managed by FIO) tested. and is ok. Go in production.
- 2006-02-20: VO Compass supported by gdrb01.
- 2006-02-13: VO Compass supported by gdrb03.
- 2006-02-09: LCG 2.7.0 installed on bdii103 and bdii104 (alias prod-bdii).
- 2006-02-09: LCG 2.7.0 installed on bdii101 and bdii102 (alias lcg-bdii).
- 2006-02-08: monb001 installed and configured as a new monbox (managed by FIO) with LCG 2.7.0.
- 2006-02-06: raid disk full on gdrb01 and gdrb03. Fixed.
- 2006-02-03: Problem with afs authentication on lxb1930. Need to restart ntpd service. Fixed.
- 2006-02-02: gdrb04 is now using lcg-bdii as a BDII (instead of lxb2008).
- 2006-02-02: Connection/routing problems with some sites in Taiwan and IN2P3. Fixed.
- 2006-02-02: Alias lcg-bdii points now to bdii101 and bdii102.
- 2006-01-27: Bug (known) in the myproxy-server daemon fixed on myproxy (the myproxy-server daemon was in a deadlock this morning).
- 2006-01-26: RAID disk in degraded mode on gdrb03 (hdg dead ?).
- 2006-01-26: Restart nfs service on lxb2003 (SE for LHCB).

WorkLog2006 < LCG < TWiki

- 2006-01-25: Restart nfs service on lxb2004 (UI for LHCB).
- 2006-01-25: Power supply on lxn1181 (myproxy.cern.ch) failure. Swith this service to node lxn1192.
- 2006-01-25: Power cut in the CERN CC last night.
- 2006-01-25: Need to restart services */edg-wl-lm*, *edg-wl-jc* and *edg-fmon-server* on gdrb03.
- 2006-01-24: lxn1192 (archiver) removed from production. Reinstalled from scratch and put it as free.
- 2006-01-22: New CE available for CERN-PROD: ce101. All the RBs will point to it when LCG 2.7.0 will be installed.
- 2006-01-20: /var partition almost full on gdrb04 due to a very huge file (/var/wtmp.1). Fixed by bzip2'ing it.
- 2006-01-19: Hard disk changed on lxn1194. Reinstalled from scratch and put it as free.
- 2006-01-19: Memory changed on gdrb10. This machine goes back in production.
- 2006-01-13: Problem with the interlogd daemon on gdrb04. Fixed by David.
- 2006-01-13: There was some trouble with the aliases lcg-bdii and prod-bdii but it has been fixed this morning. To sum up:
 - ◆ lcg-bdii points to: bdii001 and bdii002.
 - ◆ prod-bdii points to: bdii103 and bdii104.
 - ◆ Alias site-bdii should point to prod-bdii (not yet).
- 2006-01-12: bdii103 and bdii104 are now prod-bdii machines (it corresponds to site-bdii).
- 2006-01-10: lxn1178 and lxn1192 blocked. Need to reboot these two machines. Fixed.
- 2006-01-06: Problem with the raid disk on gdrb03. Machine rebooted and raid disks checked. Fixed.
- 2006-01-06: Serious problem with hda on lxn1194 (site-bdii.cern.ch). Machine removed from production and ITCM ticket generated.
- 2006-01-06: gdrb10 freezes at random time. Machine removed from production and ITCM ticket generated.
- 2006-01-06: Problem with the raid disk on gdrb03. Need to reboot the machine. No more error detected.
- 2006-01-05: Emergency power cut of the computing center affecting all services. Services back at the end of the afternoon.
- 2006-01-03: IO errors with the raid disk on gdrb06. Services stopped. Machine rebooted in the evening, and no more error detected.
- 2006-01-03: R-GMA developers have now root access to all the production RBs (gdrbxx).

This topic: LCG > WorkLog2006

Topic revision: r4 - 2008-08-05 - SteveTraylen



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)