

Table of Contents

Week of 110829.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	5
Thursday.....	8
Friday.....	10

Week of 110829

Daily WLCG Operations Call details

To join the call, at 15.00 CE(S)T Monday to Friday inclusive (in CERN 513 R-068) do one of the following:

1. Dial +41227676000 (Main) and enter access code 0119168, or
2. To have the system call you, [click here](#)
3. The scod rota for the next few weeks is at [ScodRota](#)

WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments

VO Summaries of Site Usability				SIRs, Open Issues & Broadcasts			Change assessments
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	WLCG Service Open Issues	Broadcast archive	CASTOR Change Assessments

General Information

General Information				GGUS Information	LHC Machine Information
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	WLCG Blogs	GgusInformation	Sharepoint site - Cooldown Status - News

Monday:

Attendance: local(Maarten, Ignazio, Mike, Hong, Oli, Dirk, Maria, Fernando, Ulrich, Massimo, Edoardo, Alessandro, Simone, MariaD);remote(Michael, Joel, Rolf, Gonzalo, Ulf, Jhen-Wei, Catalin, Ron, Lorenzo, Dimitri, Rob).

Experiments round table:

- ATLAS reports -
- Central Services
 - ◆ LFC db migration on Monday - CERN-PROD/analysis Panda queues were turned off on Sunday. Data transfers to/from CERN were stopped Monday morning.
 - ◆ Castoratlas db patch and defragment
 - ◆ rolling intervention on ATONR db
- T1 sites
 - ◆ IN2P3-CC finished the last part of reprocessing jobs on Sat. indicating the end of the phase-I data reprocessing. Next phase will be launched around 8 Sept.
 - ◆ INFN-T1 SRM problem Saturday morning. Issue resolved after alarm ticket.
 - ◆ hurricane in US - BNL announced emergency downtime Sat. evening. ADC activities in US clouds were turned off and users were notified. Sunday evening TW reported that the 10 Gb

link between Chicago-Amsterdam went down affecting data transfers to TW.

- ◆ TW castor upgrade on Monday and Tuesday.
- ◆ BNL (Michael) : Lab reopened at 8am for personnel. Restoring the services. Should take 3-4 hours. Hope to be fully functional by 6pm CERN time. An announcement will be sent out.

- T2 sites: ntr.

- CMS reports -

- CERN / central services
 - ◆ Friday saw degradations of CASTORCMS default service class (GGUS:73848) due to a single user doing something funny in her job (opening all files before doing some processing). This then blocked all slots in Castor and the system degraded. We killed the jobs of the user in the end, after Massimo tried all tricks and opened up as many limits as possible.
 - ◆ Sunday: noticed transfer problems from CERN with source errors "SOURCE error during TRANSFER_PREPARATION phase: [INTERNAL_ERROR] Invalid SRM version [] for endpoint []", GGUS:73870. Could be related to FTS. Experts contacted.
 - ◆ This week:
 - ◇ Tuesday, Aug. 30., 9 AM CERN time, migrate myproxy.cern.ch to new service, transparent for users (users registered with the old service will stay registered with the new service due to sync)
 - ◇ Planned: Wednesday, Aug. 31., complete SL5 migration of CERN T2 FTS, 8 AM to 10 AM UTC, transfer requests during this time will be queued, will be confirmed in today's WLCG call
 - ◇ Planned: Wednesday, Aug. 31., morning, move backing NFS store of batch system to better hardware, requires stop of batch system, running jobs keep on running, requested feedback from CMS: OK, go ahead

- T1 sites:
 - ◆ next week
 - ◇ Monday/Tuesday, Aug. 29./30., ASGC downtime to upgrade Castor
 - ◇ Wednesday, Aug. 31., PIC at risk due to network operations on firewall testing a patch that solves some problems with multicast traffic

- T2 sites:
 - ◆ NTR

- ALICE reports -

- T0 site
 - ◆ Nothing to report

- T1 sites
 - ◆ IN2P3: Since Saturday there are no jobs running at the site. Several problems:
 - ◇ Clients outside CERN were not able to contact the main seeder for certain file.
 - ◇ The AliEn installation was done in the wrong location
 - ◇ After fixing those things, there are still no jobs running. Under investigation

- T2 sites
 - ◆ Usual operations

- LHCb reports -

- Experiment activities:
 - ◆ Processing and stripping is finished. Nothing to report.

- New GGUS (or RT) tickets:
 - ◆ T0: 0
 - ◆ T1: 1
 - ◆ T2: 0
- Issues at the sites and services
 - ◆ T0
 - ◆ T1

Sites / Services round table:

- BNL nta (see ATLAS report)
- NDGF: ntr
- IN2P3: ntr
- PIC: ntr
- ASGC: Problem with the 10Gbit link with Chicago. working on it.
- FNAL:ntr
- NL-T1: ntr
- KIT: ntr
- CNAF: ntr
- OSG: ntr
- Grid Services: Interventions on Tuesday and Wednesday (see above the CMS report). The ATLAS LFC intervention went smoothly. Prod LFC fully in production on 3 virtual machines - can add more if needed. Consolidation: removed some unused services.
- Network: two links to Chicago down. No time estimation for recover yet. Firewall maintenance tomorrow morning from 7am to 8am. Could possibly cause short glitches.
- Storage services: CASTOR ATLAS db defragmentation + patch to 2.1.11-2. Done. Tomorrow CASTOR LHCb and CASTOR ALICE on Thursday.

AOB: (MariaDZ) Concerning the 2.5 hrs delay to broadcast SARA's unscheduled downtime last Thursday, the CIC portal developers answered: "After investigation we have identified some problems in the insertion of data into the DB. Since the Oracle problem in July there is some slowness and it affects the global system. With the help of the Oracle expert we have added different indexes and the DB has been tuned to be more efficient and, hopefully, avoid such problems."

Tuesday:

Attendance: local(Fernando, Jamie, Daniele, Maarten, Ignacio, Mike, Steve, Gavin, Alessandro, Jacek, Uli);remote(Ulf, Michael, Gonzalo, Catalin, Xavier, Rolf, Paco, Shu-Ting, Lorenzo, Rob).

Experiments round table:

- ATLAS reports -
- Tier0/Central Services
 - ◆ myproxy.cern.ch server migration
 - ◆ Network intervention on backbone router and firewall
 - ◆ Rolling database intervention on ATLR, ADCR and ATLDSC
- T1 sites
 - ◆ BNL-OSG2 - came back yesterday ~19:00 CERN time from the emergency downtime
 - ◇ dCache upgrade foreseen for tomorrow is postponed one week
 - ◆ TAIWAN-LCG2 - 10Gbps link between Chicago-Amsterdam was restored this morning and is fine now
- T2 sites

Tuesday:

◆ ntr.

- CMS reports -

(Daniele Bonacorsi CRC for this week)

- LHC / CMS detector
 - ◆ Machine is in "Technical STOP", no beam operation and no data taking
- CERN / central services
 - ◆ **myproxy not working from outside CERN**
 - ◇ After the end of the scheduled intervention today on myproxy.cern.ch, CMS test users at CERN reports it's ok, but CMS test users outside CERN report it does not work (e.g. "myproxy-init -d -n Unable to connect to 128.142.202.247:7512 Unable to connect to myproxy.cern.ch Connection refused"). Firewall issues again? The impact is potentially high (on all CMS analysis users not at CERN), and first users are indeed complaining -> **ALARM (GGUS:73926)**. Ulrich commented promptly, work in progress. CRC updated it at ~12.30: also at CERN it fails, for both users ("myproxy-info -d -s myproxy.cern.ch" -> "Error reading: failed to read token. Broken pipe") and CMS services (VomsProxyRenew in PhEDEx -> "ERROR from myproxy-server (myproxy.cern.ch)"). Gavin asked DNS. Daniele/CRC provided the PhEDEx case one. [Screw-up. Issues that affects CMS users not yet understood.]
 - ◆ bad transfer quality from T1_CH_CERN to almost all T2 sites (GGUS:73870)
 - ◇ getting bad again since this night at 2am UTC. The error changed to "AGENT error during TRANSFER_SERVICE phase: [INTERNAL_ERROR] no transfer found for the given ID. Details: error creating file for memmap /var/tmp/glite-url-copy-edguser/CERN-STAR__2011-08-30-0555_rD3pIy.mem: No such file or directory". May need to closely check the state of all channels on the CERN FTS server serving CERN-T2 traffic?
 - ◆ PhEDEx graph server triggered an alarm
 - ◇ Aug-29, ~4pm: "cmsweb_phedex_graphs_is_not_responding" alarm detected on vocms132 by Computer.Operations, notified to the cms-service-webtools list, checks done: server processes running, service listening to its port, IT restarted the service anyway, good because in the log there was a time-out at 29/Aug/2011:15:46:20, coming back OK only after being restarted at 29/Aug/2011:16:11:06 by the IT operator. Cause not clear. There was the CMSR DB intervention yesterday, maybe something occurred related to that, even though it was meant to be rolling transparent upgrade? DB operators actually warned that session disconnections were likely - maybe the PhEDEx graph server did not handle those gracefully. Keeping an eye if the same alarms occur again.
 - ◇ CASTORCMS_DEFAULT: SLS low availability -> down to 0 for ~6 hrs [Ignacio - correlated with load. No available slots.]
- T1 sites:
 - ◆ ASGC: scheduled downtime for Castor upgrade [downtime closed, performed in 1 1/2 days - well done!]
- T2 sites:
 - ◆ Scheduled downtimes
 - ◇ T2_UA_KIPT: site maintenance (Aug-19 12:00am - Aug-31 06:00pm) GOCDB
 - ◆ Open issues
 - ◇ T2_TW_Taiwan also in downtime of course (see ASGC T1)
 - ◇ T2_US_Purdue: 1. nSRMget errors (no ticket), 2. 100% JobRobot jobs fail (Savannah:123202)
 - ◇ T2_BE_UCL: 1. nCEmc error (no ticket); 2. nSRMput errors (no ticket); 3. 62% JobRobot jobs fail (Savannah:123188)

Tuesday:

- ◇ T2_ES_CIEMAT: 60% JobRobot jobs fail
- ◇ T2_TR_METU: transfer to PIC are failing (Savannah:123204)
- ◇ T2_US_Nebraska: transfer from PIC are failing (Savannah:123205)

- Scheduled interventions for tomorrow:
 - ◆ Wednesday, Aug-31 09:00-11:00 GVA time - Batch service NFS store replacement
 - ◇ Question: should we switch T0 off? [A: jobs which have been submitted and dispatched should be ok. Pending jobs will not be dispatched - do not switch off]
 - ◆ Wednesday, Aug-31 10:00-12:00 GVA time - Upgrade of CERN T2 FTS service
 - ◇ Question: anything was done to prepare for this that could explain GGUS:73870? [Steve - problem on SL4 T2 service, same problem that happened 3 months ago due to changing of user ids at CERN. Intervention should make this go away. Will update ticket.]
-

- ALICE reports [↗](#) -
 - T0 site
 - ◆ GGUS: 73928. Sites had problems contacting myproxy server after the update. There was a firewall issue solved by now
 - T1 sites
 - ◆ FZK: 32 TB of data not available due to an error in the filesystem
 - ◆ IN2P3: Still problems at the site. Under investigation
 - T2 sites
 - ◆ Usual operations
-

- LHCb reports - Processing and stripping is finished. Nothing to report. Castor intervention finished at 11:15, mail to contacts plus info in status board, all went all.

Sites / Services round table:

- NDGF - ntr
- BNL - intervention to move from pnfs to Chimera postponed, will send out a new announcement and will likely start next Wednesday
- PIC - ntr
- FNAL - ntr
- KIT - nta
- IN2P3 - ntr
- NL-T1 - ntr
- ASGC - Upgraded CASTOR to 2.1.11; removed LSF scheduling.
- CNAF - ntr
- OSG - ntr

AOB:

- CERN DB: we are in progress with patching production DBs, some yesterday, some today and tomorrow finished. Small incident - 3rd node of ADCR was rebooted by accident. Not available for about 10' but services relocated to surviving nodes.

Wednesday

Attendance: local(Ignazio, Luca, Gavin, Ulrich, Steve, Fernando, Daniele, Alessandro, Maria, Mike, Maarten, MariaD); remote(Michael, Rolf, Onno, Catalin,Joel, Jhen-Wei, Rob, Gareth, Gonzalo).

Experiments round table:

- ATLAS reports -
- T0/Central Services
 - ◆ CERN: "Invalid SRM version [] for endpoint []" error (GGUS 73918) has been affecting T2-CERN transfers, but disappeared this morning - presumably fixed with the CERN FTS T2 OS upgrade.

Steve: Fixed by the upgrade.

- T1 sites
 - ◆ NTR
 - T2 sites
 - ◆ NTR
-

- CMS reports -
- LHC / CMS detector
 - ◆ Machine is in "Technical Stop", no beam operation and no data taking
- CERN / central services
 - ◆ update on "myproxy not working from outside CERN" (GGUS:73926): now fixed (a painful incident, but thanks for the excellent support). Phedex will be moved back to myproxy.
 - ◆ bad transfer quality from T1_CH_CERN to almost all T2 sites (GGUS:73870)
 - ◇ looks better, but waiting also for the scheduled intervention on FTS to finish before re-checking and hopefully close this as well
 - ◆ Minor observations
 - ◇ CASTORCMS_CMSCAF: % free space, 2 drops in SLS, cannot correlate with anything, probably glitches, only FYI, no ticket
- T1 sites:
 - ◆ ASGC: a ticket as a reminder to check migration speed (slower than expected) after the successful Castor upgrade (Savannah:123229).

Jhen-Wei: working on the issue.

- T2 sites
 - ◆ NTR
-

- ALICE reports -
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ FZK: problem reported yesterday is partly solved. The vendor has given a solution that has not been implemented yet, data is not accessible until the export-import of the data to the new file system has been done
 - ◇ IN2P3: Still problems at the site. Under investigation. Looks like the explanation could be the absence of the user proxy on the WN (i.e. not forwarded by CREAM to the batch system); there was such a bug in older versions of CREAM
 - Note added after the meeting: the job proxy looks OK on the WN
 - ◆ T2 sites
 - ◇ Usual operations
-

- LHCb reports -

Experiment activities:

- Processing and stripping is finished. Nothing to report.

New GGUS (or RT) tickets:

- T0: 0
- T1: 3
- T2: 0

Issues at the sites and services

- T1 : getting the storage dumps from your dCache system for the LHCb files
 - ◆ PIC : (GGUS:73957)
 - ◆ GRIDKA : (GGUS: 73958)
 - ◆ IN2P3 : (GGUS: 73959)
 - Work ongoing for CVMFS at IN2P3. An update will be given tomorrow at the daily.
-

Sites / Services round table:

- IN2P3: we were obliged to kill a lot of LHCb jobs; stopped 100 WNs yesterday and they ran LHCb jobs which had not terminated this morning. Need machines for maintenance and switching to new batch system so had to kill these jobs. Sorry for this. Joel - could we be warned a few days before so that we can avoid this type of issue? Rolf - not that easy always to plan - depends a lot on availability of people: will try. NTA
- BNL: we talked about this yesterday; planned migration from pnfs to Chimera. Yesterday said would be postponed by one week but in view of ATLAS' 2nd phase of reprocessing will postpone this to next technical stop in November.
- NL-T1: ntr
- FNAL: ntr
- RAL: ntr
- ASGC: ntr
- PIC: ntr
- CNAF: ntr
- KIT: ntr
- NDGF: ntr
- OSG: ntr
- CERN DB: have a problem with ATLAS streams replication to tier1s. Intervention today on downstream capture DB for ATLAS / LHCb. LHCb went ok but ATLAS not. A gap has been created after system has been restarted after patch and trying to fill this to gap. Further updates will follow.

Replication to ATLAS T1s stopped.

- CERN Grid: batch intervention this morning went fine - down about 30'; myproxy issue: real bug in server introduced a couple of versions ago. CMS were first to run into it. Very lucky that main developer had time to fix immediately and deliver a patch. Avoided having to rollback.
- CERN storage: CASTOR ALICE Intervention tomorrow; network people need to replace some switches in CC which affect 9 D/S for ATLAS and 11 for CMS. Tomorrow morning or Friday? Short disconnect. ATLAS: mainly T0 ATLAS, CMS mainly CMS CAF USER.

AOB:

Thursday

Attendance: local(Jamie, Joel, Uli, Daniele, Jamie, Maria, Fernando, Manuel, Steve, Luca, Andrea, Mike, Lola, MariaDZ, Maarten);remote(Michael, Gonzalo, Kyle, Ulf, Felix, Jhen-Wei, Paco, Giovanni, John, Pavel, Catalin, Rolf).

Experiments round table:

- ATLAS reports -
- T0/Central Services
 - ◆ Status of yesterday's problem with ATLDSC intervention [Luca- problem started ~12:30 and replication down until 16:30. Some missing transactions in queue captured by downstream capture. Some messages missing - manually DBAs developed procedure to apply missing transactions. Applied first to BNL and then all others as ok. Root cause unknown]
 - ◆ Schedule of network interruption affecting T0ATLAS servers planned for today
- T1 sites
 - ◆ NTR
- T2 sites
 - ◆ NTR

- CMS reports -
- LHC / CMS detector
 - ◆ Machine is in "technical stop", no beam operation and no data taking
- CERN / central services
 - ◆ update on "myproxy not working from outside CERN" (GGUS:73926): yesterday it was supposedly fixed and being kept under observation, but:
 - ◇ At 3am today, PhEDEx operators experience again problems with proxy renewals on the cmsweb backends. Yesterday after the WLCG Ops call PhEDEx moved back to use myproxy.cern.ch. It seems something changed overnight? Steve confirmed no intervention/incidents have happened since then. Now, credentials manually re-uploaded to the backend machines as of 10am GVA time today, but tests seem to indicate the issue persist. We had better let Maarten/Steve/Gavin browse and parse the logs (see GGUS for details). Comment: we *need* external debugging: the client logging is too poor for us to do it ourselves.
 - ◆ update on "bad transfer quality from T1_CH_CERN to almost all T2 sites" (GGUS:73870)
 - ◇ I confirm that transfer quality is indeed much more green now, only some CERN-T2 links have issues from time to time, which could just be normal life. GGUS closed and verified.
 - ◆ yesterday afternoon we had **no transfers from P5 to the Castor t0streamer pool**. The rfc processes seemed blocked. Opened a TEAM GGUS ticket to Castor then raised to ALARM (GGUS:73965). Promptly commented and understood (thanks). The rfcs were killed and the

new copy processes worked afterwards. Ticket closed and verified.

◊ also, rfcop size-0 file issues (being discussed on the Storage Manager Ops HN, is being discussed by Stephen Gowdy with CERN-IT since the issue is indeed getting worse than it used to be). No updates today, maybe in next days?

◆ Minor observations

◊ some CASTORCMS SLS observations, but busy and could not check further yet

• T1 sites:

◆ ASGC: check migration speed after the successful Castor upgrade (Savannah:123229). Felix: some problematic tapes, stuck in busy state not properly mounted, labelled those read-only and added new tapes to the pool, ticket closed, investigations on problematic ones continue though

• T2 sites:

◆ ticketing as usual (mainly SAM/JR failures)

• ALICE reports [↗](#) -

• T0 site

◆ Nothing to report

• T1 sites

◆ IN2P3: Still problems at the site. Network problems have been discarded. Under investigation

• T2 sites

◆ Usual operations

• LHCb reports -

◆ T1 :

◊ IN2P3 : no answer to GGUS ticket from yesterday.. Any reason ? (GGUS:73959)

◊ No more space on shared area to install s/w - can we have an increase? (Will open ticket).

Sites / Services round table:

• BNL - ntr

• PIC - ntr

• NDGF - we have a disk problem on one pool in Sweden - acting strangely and will be emptied by tomorrow. Probably ATLAS data - will know then if we have lost data or not. Will know more tomorrow.

• ASGC - ntr

• NL-T1 - ntr

• CNAF - ntr

• RAL - ntr

• FNAL - ntr

• IN2P3 - don't have access to ticket for moment; CERNVMFS - we are about to install this for WNs with 24 cores - process is ongoing; abit complicated as we have to resize several disk areas and so have to take out WNs and reconfigure and bring back in etc. To avoid multiple actions like this we do this when transferring WNs from BQS to GridEngine so those actions are related. About 20% of GridEngine cluster is already under CERNVMFS and hope to have 50% by Monday.

• OSG - ntr

• CERN DB - in addition apply for conditions of ATLAS at TRIUMF has stopped. Contacted dba - probably restart will be needed.

• CERN myproxy - investigation continuing. Hope new service is ok 😊 Have updated ticket. Daniele - 2nd patch? What should that fix? Maarten - not convinced by commit log that this would solve

problem seen. Should apply patch. Steve - 17:00 today

AOB:

Friday

Attendance: local(Daniele, Jamie, Fernando, Eddie, Dawid, Ivan, Claudio, Steve, Maarten);remote(Elizabeth, Roger, Michael, Gonzalo, Onno, Jhen-Wei, Catalin, Rolf, Tiju, Andreas, Joel, Rob, Giovanni).

Experiments round table:

- ATLAS reports -
 - T0/Central Services
 - ◆ Automatic Site Cleaning did not restart after the reboot of the machine few days ago. Noticed yesterday and service running again.
 - T1 sites
 - ◆ NTR
 - T2 sites
 - ◆ NTR
-

- CMS reports -

LHC / CMS detector

- Machine is in "technical stop", no beam operation and no data taking
 - CERN / central services
 - ◆ update on "myproxy not working from outside CERN" (GGUS:73926): no more errors since the PhEDEx ProxyRenewal cron-ed at 8am today. CMS expert and WLCG experts sitting together to boost a solution, or agree on a plan for the weekend.
 - T1 sites:
 - ◆ ASGC: SAM SRM test failures ("failed to contact SRM"), Felix reports they found SRM was suffering SOAP errors but only from time to time, debugged and traced to some of the disk-servers CRLs that expired since last intervention, they should be fixed by now, indeed CMS shifters confirm it's back to green (Savannah:123263 -> CLOSED).
 - ◆ CNAF: PhEDEx agents down for a quick intervention, done now (Savannah:123269, bridged to GGUS:74015 -> CLOSED)
 - T2 sites:
 - ◆ T2_CH_CSCS: Phedex Agent FileDownload is off -> fixed
 - ◆ T2_CN_Beijing: SAM TESTS CE error for >12 hrs -> fixed
 - ◆ T2_AT_Vienna: bdII error for a>18 hours (Savannah:123242, open)
-

- ALICE reports [↗](#) - NTR - issue at Lyon still being debugged, main alien developer expects to be given access to WN to see what's going wrong there, no further news since this morning.
-

- LHCb reports -

Experiment activities:

Friday

- Processing and stripping is finished. Nothing to report.

New GGUS (or RT) tickets:

- T0: 0
- T1: 1
- T2: 0

Issues at the sites and services

- T1 :
 - ◆ IN2P3 : AFS space full (GGUS:73986)
-

Sites / Services round table:

- NDGF - disk problem discussed yesterday: one disk has hard errors leading to some failure in stored meta data. No data should be lost and should be available R/O. Scheduled downtime on Monday for dCache updates.
- BNL - ntr
- PIC - yesterday evening had an incident that affected some ATLAS production jobs. Misconfig in DNS alias of dcap doors - human error. About 1/4 of production jobs affected but fixed this morning
- NL-T1 - ntr
- ASGC - ntr
- FNAL - ntr
- RAL - ntr
- KIT - ntr
- CNAF - router had some connectivity problems, affected traffic with NDGF and ASGC and non-LHC experiments. After some ACL change. Now stable and in contact iwht CISCO tech support. AT RISK from 17-18 UTC on Sep Thursday next week
- IN2P3 - we will reduce capacity by 1/3 during w/e to prepare for installation of CERNVMFS and for some recabling on Monday
- OSG - ntr; Monday is a holiday so we will be holiday procedures this coming Monday

- CERN Myproxy: CMS were using a reg exp during my proxy init step and parentheses were not interpreted correctly due to bug. Changed reg exp in client code and all seems ok. WIll follow up and get bug fixed. Only CERN machines were affected; remote PhEDExes were fine. Conclusions in ticket. This area of code was known to have changed; old version also did something wrong for which a workaround was also needed. Propose to close SNOW ticket which will close GGUS (alarm).

AOB:

-- JamieShiers - 18-Jul-2011

This topic: LCG > WLCGDailyMeetingsWeek110829

Topic revision: r15 - 2011-09-02 - MaartenLitmaath



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