

Table of Contents

LHCb Production Operations Procedures.....	1
Setting Production environment.....	1
CERN - lxplus.....	1
non CERN machines.....	1
Building,Deploying & Installing (and Core Software).....	1
Building the DIRAC binary distributions.....	1
Installing DIRAC on lxplus.....	2
Installing on non CERN machines.....	2
"user guide" on how to take advantage of the CMT setup of.....	2
Deployment of LHCb Software on the Grid.....	2
Workload management.....	2
Primary job states in DIRAC.....	2
Job Management Operations.....	3
Running Jobs Locally With DIRAC.....	3
Production management.....	3
Launching productions from the production request page.....	3
How to derive a production.....	3
How to rerun a production job starting from a log SE link.....	3
Dealing with Production IDs, Production Job IDs and WMS Job IDs.....	4
Checking the Status of Files in the Production Database.....	4
How to Validate a new production.....	4
How to deal with Stalled jobs.....	4
Submitting test jobs.....	4
Production workflow.....	4
User workflow.....	4
Procedures in case of site failures.....	4
Production API Notes.....	4
Getting production to 100%.....	4
Closing a production.....	5
Pilots monitor.....	5
Creating "mini productions".....	5
Dealing with gaudi applications commands.....	6
Data Management.....	7
How to fix screwed up replication transformations.....	7
Transfer PIT - CASTOR.....	7
Job Data Access Issues.....	7
Staging request blocked.....	7
Changing the Data manager.....	7
File recovery.....	8
How to recover replicas that are lost even if SRM reports they are existing.....	8
Changing the Default Protocols List for a given Site (Tier-1).....	8
Banning a SE if a Site is in downtime or full.....	8
Banning, unbanning, and re-directing the.....	9
Checking if a file is cached.....	9
Generating a POOL XML slice for some LFNs.....	9
Determining Versions of LCG Utils, GFAL and the Default SRM2 Protocols List in DIRAC.....	9
Checking the throughput from the pit to Castor (during data taking).....	9
DIRAC Configuration Service.....	9
Adding new users.....	9
Restarting the configuration service.....	9
Getting information from BDII.....	9
DIRAC (and not) Services monitoring.....	10
Getting the List of Ports For DIRAC Central Services (and how to ping them).....	10
SLS sensors.....	10

Table of Contents

LHCb Production Operations Procedures

Overview.....	10
CERN Core Services monitoring.....	10
Sites Management.....	10
Banning and allowing sites.....	10
Sites rank: Unspecified Grid Resources Error.....	10
Site Availability Monitoring (SAM) tests.....	10
Sites Troubleshooting.....	10
Investigating failed jobs.....	10
SQLite hint.....	10
Site Problems Follow up.....	11
Getting in touch with sites: Tickets and Mails.....	11
GGUS.....	11
Site Mail contacts.....	12
Daily Shifter Checklist.....	12
End of production Checklist.....	12
Miscellaneous.....	12
Feature Requests and Bug Reports.....	12
DIRAC3 service problems on 4 October 2008.....	12
Mailing lists.....	13
Questions and comments to the experts.....	13
Bookkeeping System.....	13
Set the data quality flag.....	13
Web portal.....	14
Web portal stuck: how to restart it.....	14
Documents.....	15

LHCb Production Operations Procedures

This is the LHCb Production Operations Procedures page which contains procedures for Grid Experts and Grid Shifters to follow. More information about Production Operations and some useful links can be found at the Production Operations page. The draft template for defining production procedures is available here in the following formats: Template (.doc), Template (.rtf), Template (.txt). Each procedure should be posted with a link to the documentation and the desired audience e.g. Grid Experts, Grid Shifters or both.

An attempt to categorize common procedures has been made below, feel free to update this list as appropriate.

Setting Production environment

CERN - lxplus

After you login on lxplus, you have to run the two following commands to get the production environment

- SetupProject LHCbDirac
- lhcb-proxy-init -g lhcb_prod

non CERN machines

to install outside CERN, you should download and install LHCbDirac (for example v5r3)

- setenv MYSITEROOT /my/location/to/install/dirac
- setenv CMTCONFIG <LCG_TAG>
- cd \$MYSITEROOT
- wget -p http://cern.ch/lhcbproject/dist/install_project.py
- python install_project.py -p LHCbDirac -v v5r3 -b
- source \$MYSITEROOT/LbLogin.csh or . \$MYSITEROOT/LbLogin.sh
- SetupProject LHCbDirac

To use LHCbDirac on a machine without lcg stuff, you need to add in the \$DIRACROOT/etc/dirac.cfg the following lines

```
Resources
{
  FileCatalogs
  {
    LcgFileCatalogCombined
    {
      Status = InActive
    }
  }
}
```

to deactivate the LFC checking.

Building, Deploying & Installing (and Core Software)

Building the DIRAC binary distributions

The instructions for how to build the binaries for DIRAC are here. (Grid Experts)

Installing DIRAC on lxplus

- If you are developing tools for DIRAC, it probably makes sense for you to install your own version of LHCbDirac.

Installing on non CERN machines

For installing LHCbDirac on your local machine, you should download and install LHCbDirac, specifying the version <version>

- `setenv MYSITEROOT /my/location/to/install/dirac`
- `setenv CMTCONFIG <LCG_TAG>`
- `cd $MYSITEROOT`
- `wget -p http://cern.ch/lhcbproject/dist/install_project.py`
- `python install_project.py -p LHCbDirac -v <version>`

Then in your login script you should include:

- `source $MYSITEROOT/LbLogin.csh` **or** `. $MYSITEROOT/LbLogin.sh`

and to set up the LHCbDirac environment (beware if you use ganga this is not needed as done internally by ganga)

- `SetupProject LHCbDirac`

If you see an error message like "Warning : Cannot add voms attribute /lhcb/Role=user to proxy Accessing data in the grid storage from the user interface will not be possible. The grid jobs will not be affected." then try doing `chmod 644 $DIRAC_VOMSES/lhcb-voms.cern.ch`. You will also need to set `$X509_CERT_DIR` and `X509_VOMS_DIR`. Refer to `lxplus` for default settings, or take a look at the Dirac tool `dirac-admin-get-CAs` available in Diracs later than `v4r19`. However you do it, if you make a local copy of these two directories, you will need to keep that copy up-to-date. -- WillReece - 2009-10-06

"user guide" on how to take advantage of the CMT setup of

Summary of commands to be used for taking advantage of the LHCbDirac installation using CMT.

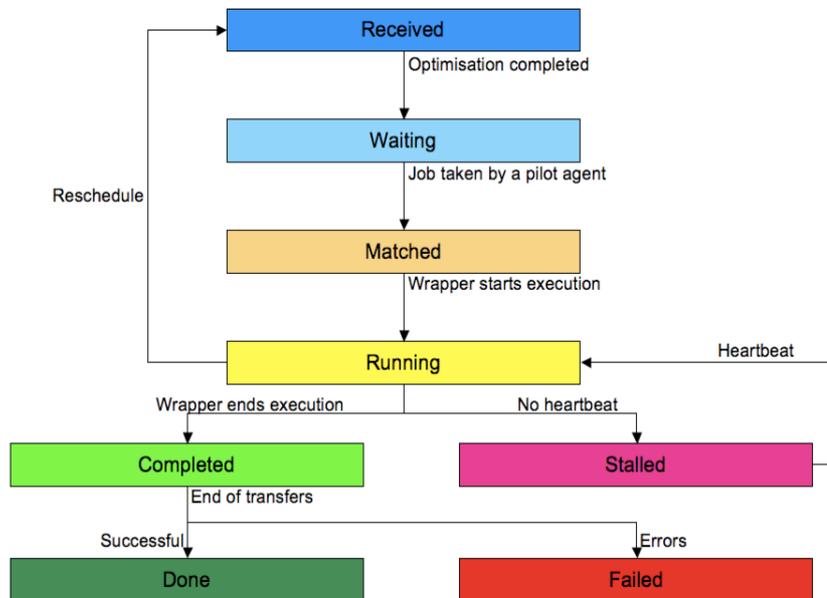
Deployment of LHCb Software on the Grid

- Installing LHCb software on the Grid.
- Deploying LHCb Software on the Grid (Grid Experts)

Workload management

Primary job states in DIRAC

ProductionProcedures < LHCb < TWiki



Job Management Operations

The procedures for investigating production jobs (suitable for Grid Shifters and Grid Experts) is coming.

The procedure for search string in output for selected jobs (suitable for Grid Shifters and Grid Experts).

A low level investigation on LSF to check why LHCb jobs do not start at CERN.

Running Jobs Locally With DIRAC

There are three submission modes associated with DIRAC: default WMS submission, local execution of a workflow and finally execution of a workflow in the full agent machinery. This procedure explains the steps for running jobs locally with DIRAC.

Production management

Launching productions from the production request page

An example of how to launch a simulation production is available. Note that the prerequisite for launching productions is having the `lhcb_prmgr` role.

How to derive a production

Typically, when a new application version is released and should substitute the current one, it's very useful to derive a production

How to rerun a production job starting from a log SE link

Frequently individual jobs can fail with an error that should be investigated by applications experts. The following guide on how to rerun a job can be circulated in case of questions by the applications expert.

Dealing with Production IDs, Production Job IDs and WMS Job IDs

This simple guide shows how to obtain production IDs from WMS job IDs and vice versa (Grid Shifters and Grid Experts)

Checking the Status of Files in the Production Database

A brief guide describing how to check the status of given productions in the Production DB is available (Grid Shifters and Grid Experts)

How to Validate a new production

Each time a new production is created, the Grid Shifter should check that a few jobs are successful prior to submitting the entire batch.

How to deal with Stalled jobs

Jobs moving into the Stalled state is OK, as they can recover again and complete successfully. However, if the job has been stalled for many days, action should be taken.

Submitting test jobs

Test jobs should be submitted to help debug problems. This can be done through DIRAC to test the full chain of Grid submission or could be running an LHCb application directly on the site WN (if you have relevant access permissions) if you know that the problem is confined to the site.

Production workflow

- Application Log File Analysis
- Production Job Finalization

User workflow

Section to describe general policies for user jobs in DIRAC.

- User Output Files (under construction)
- User Job Finalization (under construction pending module)

Procedures in case of site failures

- SiteFailureProcedures

Production API Notes

The following link gives an introduction and examples using the Production API. (Grid Shifters and Grid Experts)

Getting production to 100%

Many times, it seems that is very easy for a production to reach 95%, but what is difficult is to reach 100%. A list of cases can be found in this link. (Mostly for Grid Experts and Production Manager, but Grid shifters can still grasp useful information)

Closing a production

It is very cumbersome to keep in the production system old production, may be still active generating unduly load on various component of the Production System like for example the BookkeepingWatchAgent that will also loop on these not longer useful productions stretching the time to create tasks for effectively active productions. At this link a procedure to pick up and close not longer alive production is provided.

Pilots monitor

If jobs are not being submitted for a long time, you can check first of all if pilots are submitted, and then if they are actually matched. First, you can look in the portal in the "Pilot monitor" page, to see if there pilots running or submitted. Then, with the command

```
dirac-admin-get-job-pilots jobID
```

you check if pilots are submitted, for you job queue. This will print the logs for the pilots in the queue. If you don't see a line with

```
'Status': 'Submitted'
```

then it might be that there is a problem.

Also, through the pilot monitor page you can see the pilot output for the "done" pilots, that can contain useful information of why the pilots might not be matched.

Creating "mini productions"

So called "mini productions" are sometimes necessary to process a small set of files with a new production (i.e. improved application)

- Create a new production, if necessary modify steps and create a new request from those steps,
- Launch the production as usual (I propose to set the run range to a set which is not used anywhere else - see later - e.g. the run/s which are concerned by this prod
- After the production has been launched two steps need to be done quickly !!!
 - ◆ Stop the production in the Dirac production monitor
 - ◆ Delete the BkQuery and Params for this production in the Database (ProductionDB). The BkQueryID can be found in the "Additional Params" of the production on the Dirac Production Monitor page
 - ◆ Delete BkQuery

```
mysql> select * from BkQueries where BkQueryID = 7590 ;
```

BkQueryID	SimulationConditions	DataTakingConditions	ProcessingPass	FileType	E
7590	All	Beam3500GeV-VeloClosed-MagUp	Real Data	RAW	9

```
1 row in set (0.00 sec)
```

```
mysql> delete from BkQueries where BkQueryID = 7590 ;
```

```
Query OK, 1 row affected (0.02 sec)
```

```
mysql> select * from BkQueries where BkQueryID = 7590 ;
```

```
Empty set (0.00 sec)
```

- ◆ Delete Additional Parameter from Production

```
mysql> select * from AdditionalParameters where TransformationID = 16309 and ParameterName = "BkQ
```

```
+-----+-----+-----+-----+
| TransformationID | ParameterName | ParameterValue | ParameterType |
+-----+-----+-----+-----+
|           16309 | BkQueryID    |       7590     | StringType    |
+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

```
mysql> delete from AdditionalParameters where TransformationID = 16309 and ParameterName = "BkQue
Query OK, 1 row affected (0.01 sec)
```

```
mysql> select * from AdditionalParameters where TransformationID = 16309 and ParameterName = "BkQ
Empty set (0.00 sec)
```

- Add the needed files to the production via the script below, provide production ID, run number and list of files (Certificate role needs to be lhcb_prod)

```
transID = 16310
run = 102140
listOfFiles = ['/lhcb/LHCb/Collision11/SDST/00012938/0000/00012938_00006376_1.sdst']

from DIRAC.Core.Base.Script import parseCommandLine
parseCommandLine()
from DIRAC.Core.Utilities.List import sortList
from DIRAC.Core.DISET.RPCClient import RPCClient

tsClient = RPCClient( 'Transformation/TransformationManager' )

res = tsClient.addFilesToTransformation( transID, sortList( listOfFiles ) )
if not res['OK']:
    raise Exception, res['Message']
res = tsClient.addTransformationRunFiles( transID, run, sortList( listOfFiles ) )
if not res['OK']:
    raise Exception, res['Message']
```

Dealing with gaudi applications commands

There are a set of options with which the gaudi applications are run. Some of the flags can be set in the CS. All these option can be made dependent from the setup. The example that follows are for the LHCb_Production setup.

For what regards install_project.py:

```
Operations->LHCb-Production->GaudiExecution->installProjectOptions
```

can be used for setting the flags of install project, when really installing the project. I remind you that this action is triggered only if the project is not (yet) installed. If such option is not set, the default is to run install_project.py with "-b" flag.

Instead, the option

```
Operations->LHCb-Production->GaudiExecution->checkProjectOptions
```

can be used for setting possible flags for checking if the project is already installed. This is done running with the default flags "-b --check". In case you want to override such behavior, by setting this option in the CS, do not forget to always add at least "--check".

The option

```
Operations->LHCb-Production->GaudiExecution->removalProjectOptions
```

is instead for removing application. It's by default '-r'.

It is also possible to modify the `install_project` location (the script is downloaded from the web server), setting:

```
Operations->LHCb-Production->GaudiExecution->install_project_location
```

which, by default, points at <http://lhcbproject.web.cern.ch/lhcbproject/dist/>

Data Management

How to fix screwed up replication transformations

Using the `dirac-transformation-debug`, instructions here

Transfer PIT - CASTOR

The Data transfer between the PIT and CASTOR for the RAW is handle on the machine `lbidirac.cern.ch` by the user `lhcbprod`. The `dirac` installation is done under `/sw/dirac/data-taking`. The transfer itself is managed by the Agent `/sw/dirac/data-taking/startup/DataManagement_transferAgent`. This python process should run **MaxProcess** processes and each process can start a new process for each transfer (**MaxProcess** can be found in `/sw/dirac/data-taking/etc/DataManagement_TransferAgent.cfg`). If you don't see too many processes, you can look at the log `/sw/dirac/data-taking/DataManagement_TransferAgent/log/current`. A typical behaviour can be seen here.

You can also look at this web page [to spot a potentiel problem](#) if you see that the rate decrease. In principle in normal condition of data taking period, it means that one or several processes are stuck. you can find them with **strace -f -pid _PID_**. As soon as you find it you can kill it **kill -9 _PID_**. If it has no effect, you can stop the agent in a proper way **touch**

/sw/dirac/data-taking/control/DataManagement/TransferAgent/stop_agent. If it does not produce any effect, you can finally try **runsvctrl t /sw/dirac/data-taking/startup/DataManagement_TransferAgent**. As last resort, you will have to kill it by hand **kill -9 _PID_**

You can apply the recipe for the `RemovalAgent`.

Job Data Access Issues

The following document is meant to give Grid shifters a few hints on things they can check when a job has problem accessing data on a site. (Grid Shifters and Grid Experts). Check list to debug `dcache` [and](#) `CASTOR` [issues](#) are available also.

Staging request blocked

If there are some `STAGEIN` request blocked you can follow the recipe (<http://lblogbook.cern.ch/Operations/4647>) to recover the situation

Changing the Data manager

It happens more frequently than one expects the need of swapping the identity of the LHCb Data Manager. In this procedure the steps to accomplish smoothly this operation are described.

File recovery

- Recovery Of Files Lost By SE (Grid Experts)

How to recover replicas that are lost even if SRM reports they are existing

This can happen. The file is physically lost but SRM (`lcg-ls`) reports the file is there, see this [GGUS](#). This replica is totally lost from tape and disk:

```
> lcg-ls -l srm://gridka-dCache.fzk.de/pnfs/gridka.de/lhcb/data/2011/RAW/FULL/LHCb/COLLISION11/9
-rw-r--r--  1      2      2 3145768992          NEARLINE /pnfs/gridka.de/lhcb/data/2011/RAW/FUL
    * Checksum: 51e2fc3d (adler32)
    * Space tokens: 39930230
```

You have then to remove the lost replicas and then copy them over again from other another site:

```
$ dirac-dms-remove-lfn-replica /lhcb/data/2011/RAW/FULL/LHCb/COLLISION11/98298/098298_0000000077.
$ dirac-dms-replicate-lfn /lhcb/data/2011/RAW/FULL/LHCb/COLLISION11/98298/098298_0000000077.raw
```

If there is only one replica and the corresponding file at the site has been lost completely, then you need to use `dirac-dms-remove-files` to remove the entry in the replica catalogue. You need to double check this is really the case, as this command will remove all replicas of the given file!

Changing the Default Protocols List for a given Site (Tier-1)

The order of the list of protocols supplied to SRM can be changed e.g. testing root at NIKHEF means root is prepended to the list. This guide explains how to change the protocols list for a given site. This operation is restricted to those with the `diracAdmin` role. (Grid Experts)

Banning a SE if a Site is in downtime or full

You (as Grid Expert or Data Manager) can use the command `dirac-admin-ban-se / dirac-admin-allow-se` to disable or enable a SE in case of problem or downtime affecting a SE.

Example to ban all the SEs at RAL in writing.

```
dirac-admin-ban-se -c LCG.RAL.uk
```

Example to ban one SE at RAL

```
dirac-admin-ban-se RAL-DST
```

Example to ban one SE in writing at CNAF

```
dirac-admin-ban-se -w CNAF-USER
```

Also keep in mind that NIKHEF and SARA have different SE. LCG.NIKHEF.nl SE are: NIKHEF-RAW, NIKHEF-RDST

While the others are in fact based at SARA: NIKHEF-DST NIKHEF-FAILOVER NIKHEF-USER NIKHEF_M-DST NIKHEF_MC-DST NIKHEF_MC_M-DST

Banning, unbanning, and re-directing the

The actual implementation is using the Configuration Service to store the connection strings of the oracle DB at the the site. There are plans for switching to another technology, but for the moment this is how it is implemented.

In the CS, in section Resources/CondDB there is a subsection for each of the T1 sites, with connection strings and status. If the status is "Active", such database is used. If "InActive" is used instead (or anything but "Active"), a second connection string between the available ones is used, chosen with a random shuffle. Anyway, sometimes it's better to specify a fixed redirection.

A "trick" is used instead for the specific redirection: just save the original, real connection in a section (call it for example "LCG.CNAF.ir.REAL") and set the status as "InActive". Then, if you, for example, wants to redirect to CERN, just copy the CERN section and call it LCG.CNAF.it.

Remember that SARA and NIKHEF share the same DB.

Checking if a file is cached

Simple four steps to know how to check if a file is cached or not (Grid Shifters and Grid Experts)

Generating a POOL XML slice for some LFNs

The following is a simple description of the replacement for genCatalog in DIRAC3. This uses the standard LHCb input data resolution policy to obtain access URLs.

Determining Versions of LCG Utils, GFAL and the Default SRM2 Protocols List in DIRAC

There is a simple non-intrusive script to obtain the external package versions in DIRAC (Grid Shifters and Grid Experts)

Checking the throughput from the pit to Castor (during data taking)

The link band-with is 10Gbit. Expected rate (beginning of 2012) is about 280 MB/s, some more details here.

DIRAC Configuration Service

Adding new users

- How to add users to the DIRAC configuration service (Grid Experts)

Restarting the configuration service

- The DIRAC configuration service sometime goes down and has to be restarted. (Grid Experts)

Getting information from BDII

- Getting information from BDII(Grid Experts)

DIRAC (and not) Services monitoring

Getting the List of Ports For DIRAC Central Services (and how to ping them)

The ports for DIRAC central services for a given setup can easily be checked (useful for hardware requests). The procedure for port discovery as well as how to ping a service is available (Grid Shifters and Grid Experts).

SLS sensors

Overview

CERN Core Services monitoring

Sites Management

Banning and allowing sites

- When to ban a site.

Sites rank: Unspecified Grid Resources Error....

- In this Rank procedure some guidelines to debug understand why a site is not running payload.

Site Availability Monitoring (SAM) tests

SAM tests are used by LHCb to check that basic functionality of each Grid site is working. They run a few times each day at each site. The SAM framework has been updated in DIRAC3 and now runs as a specialized workflow with tailored modules for each functional test. Running the SAM tests is restricted to Grid Experts having the legadmin VOMS role.

- Description of the SAM framework.
- How to submit a SAM test (Grid Experts).
- What to do when a site fails a SAM test (Grid Shifters/Experts).
- LHCb SAM Test documentation.

Sites Troubleshooting

Investigating failed jobs

If many jobs start to fail at a site they should be immediately investigated.

SQLite hint

The following SQLite is meant to provide a template description to feed a GGUS requests of investigation about one of the most recurrent problems.

Site Problems Follow up

Grid jobs (=pilot jobs in DIRAC terminology) are often failing. In this link an encyclopedic summary of all known issues concerning (mainly) grid jobs issues is reported

Getting in touch with sites: Tickets and Mails

How to deal with GGUS follows

GGUS

In this practice-guide we aim to provide few clear rules to the operators/GEOCs/experts to submit GGUS tickets and a quick introduction to the ticketing system. In Grid a *problem is not a problem if a GGUS has not been open*. With that clear in mind we want to present and analyze the best way to submit a GGUS ticket to a site (for a site specific problem). In the early days there was the mail as unique tool to contact sites. It was quick and also efficient. GGUS ticketing system came in the game bringing much more functionality but also a less fast way to get in touch with sites. While indeed it was a good tool to track problems and also to accumulate know-how about problems, the path of the ticket was not always straight to the experts that had to fix the problem on the remote sites. We noticed (years of experience) that a ticket GGUS + direct mail to support mailing list at the site was the most efficient and quick way to get in touch with the site. The responsiveness increases if a LHCb local T1 contact person was also put in the loop. (please note that contact mailing address for T1s are available at #Mailing_lists)

The recent releases of GGUS introduced new features that matched both the need for a quick contact with sites in a old mail fashion and the robustness typical of a ticketing system. More in details together with the usual USER ticket GGUS is offering the possibility to submit both *TEAM* tickets and *ALARM* tickets.

- A **TEAM** ticket is a special ticket that matters production activities and is targetted to problems at sites that have to be followed by a crew as a whole rather than a single person (ex. production team). Any problem at every each site that everyone in the operations team is potentially demanded to follow and intervene, must be spawned via a **TEAM** ticket. A **TEAM** ticket either can allow for a direct routing of the problem to the site (in which case the submitter must put from a drop-down menu the GOC name of the site affected) or can go through the usual path TPM/ROC/SITE with the unavoidable lost of time. **A TEAM ticket is not a top priority ticket**. Submitter has the possibility to select the severity from the web form. The only difference is that everybody in the same **TEAM** could modify and interact with the ticket that is owned by the **TEAM** and not the user. GGUS knows about the member of the **TEAM** via VOMS. All people entitled to dress the Role=production (now) or Role=team (coming soon) are part of the **TEAM** and recognize to act on the ticket.
- An **ALARM** ticket is another special ticket that is meant really to generate **ALARMS** on the interested sites. The implementation of the **ALARM** at site level is different and different the support each site decided to put in place. Mails to special site mailing list (alarming mailing list) internally on the site, might in turn trigger procedures to react quickly to a problem even **outside working hour**. SMS, phone calls, operators, control rooms, remedy tickets...Everything is left behind the scenes. T1's, as per MoU, are demanded to react in less than 30 minutes, 24X7 to **ALARM** tickets. What matters here is: VOs are guaranteed to have a answer in at least 30 minutes **but the solution is not necessarily guaranteed in such short time!** Only a very well restricted number of people inside the VO (read: alarmers) are entitled to submit **ALARM** tickets. This limitation is clearly a need to avoid that non-experts had the possibility to wake up someone else for a fake problems. Soon GGUS will retrieve authorized alarmers from VOMS (Role=alarm). A list of authorized **ALARMERS** for LHCb is today available here.

Would you please mind that - despite those new tools are extremely useful and important **we warmly recommend to not abusing about them**. The net effect indeed is a lost of credibility that would relax the **ALARM** threshold. I propose below some suggestions about typical problems and action to be taken.

1. If the problem is a show stopper the shifter has to call the GEOC. The Experts has then to investigate whether the problem is really a show stopper and in case submit the ALARM. A show stopper here is mainly a problem that prevents to continue with the activity on the site. In the GGUS portal for ALARM ticket, there is available a list of identified MoU activities that may give origin to an alarm. It's worth to remind however that at CERN a *show stopper only matters data*. When submitting please put in cc also lhcb-grid@cernNOSPAMPLEASE.ch mailing list and open an entry in the e-logbook [↗](#)
2. If the problem affects severely one of the services at T1's and compromises one of the activities on the site, a TEAM ticket with "Top Priority" or "Very Urgent" is recommended. We leave up to the GEOC to decide but also an entry in the e-logbook must be filed.
3. If the problem interests the production activity at other sites the GEOC or the shifter must open a TEAM ticket with a severity that ranges from "Less Urgent" to "Top Priority" depending on how the problem impacts the (T2) site. If just few jobs have problems and the rest is running happily (let's say less than 10%) it may be just a (few) WN problem (Less Urgent). If the site is acting as a black hole and compromise also the activities somewhere else by attracting and failing jobs that otherwise may reach other sites, the site must be banned and the TEAM ticket deserves a Top Priority level.
4. Normal users can also get in touch with sites via Standard ticket. Severity is again matter of personal feelings. We discourage however to always think "my problem is always more important than any other else". In WLCG soon 5K users will start doing their activities.

Ticket escalation:

Since Jan 21st GGUS allows for an escalation of ticket slowly taken by the support units or unresponsive sites. Please read this escalation procedure document [↗](#).

Site Mail contacts

WLCG "conventional" site-support mailing list are available here.

Daily Shifter Checklist

- Each day the shifter should routinely check the following items to ensure the smooth running of distributed computing for LHCb.

End of production Checklist

- At the end of a production, the shifter should follow this procedure.

Miscellaneous

Feature Requests and Bug Reports

- Submitting Feature Requests and Bug Reports

DIRAC3 service problems on 4 October 2008

The problem and its resolution

Mailing lists

- **lhcb-dirac** : all people involved in development of DIRAC
- **lhcb-gridshifters** : all Computing shifters
- **lhcb-ganga** : LHCb ganga developers
- **lhcb-lcgcontact** : LHCb people interacting with LCG
- **lhcb-datamanagement** : LHCb support for Datamanagement
- **lhcb-bookkeeping** : LHCb support for the bookkeeping
- **lhcb-sam** : LHCb support for SAM test suite
- **lhcb-production-manager** : LHCb people in charge of the management of the production
- **lhcb-gridresources** : LHCb people in charge of the LHCb GRID resources
- **lhcb-cern-contact** : LHCb contact at CERN
- **lhcb-cnaf-contact** : LHCb contact at CNAF
- **lhcb-gridka-contact** : LHCb contact at GRIDKA
- **lhcb-in2p3-contact** : LHCb contact at IN2P3
- **lhcb-nikhef-contact** : LHCb contact at NIKHEF
- **lhcb-pic-contact** : LHCb contact at PIC
- **lhcb-ral-contact** : LHCb contact at RAL
- **lhcb-wms** : LHCb WorkloadManagement System for Operations
- **lhcb-grid-alarms** : lhcb-wms + lhcb-gridteam
- **lhcb-grid** : lhcb-grid-alarms + lhcb-gridshifters
- **lhcb-paste** : lhcb-grid + lhcb-ganga
- **lhcb-gridteam** : lhcb-lcgcontact + lhcb-datamanagement + lhcb-bookkeeping + lhcb-sam + lhcb-production-managers + lhcb-gridresources + lhcb-cern-contact + lhcb-cnaf-contact + lhcb-gridka-contact + lhcb-in2p3-contact + lhcb-nikhef-contact + lhcb-pic-contact + lhcb-ral-contact
- **lhcb-production** : lhcb-grid + lhcb-ppg
- **lhcb-distributed-analysis** : lhcb-grid + all members of LHCb interested

Questions and comments to the experts

- Questions from shifters to the experts
- Comments from shifters to the experts

Bookkeeping System

Set the data quality flag

- You know the data quality flag you can use: `dirac-bookkeeping-setdataquality-run` or `dirac-bookkeeping-setdataquality-files`. The commands without input parameter shows the available quality flags and how to use this command.

```
(DIRAC3-user) zmathe@pclhcb43 /scratch/zmathe/dirac[406]>dirac-bookkeeping-setdataquality-run
Available data quality flags:
UNCHECKED
OK
BAD
MAYBE
Usage: dirac-bookkeeping-setdataquality-run.py <RunNumber> <DataQualityFlag>
```

- The following commands you can use to set the data quality flag:

dirac-bookkeeping-setdataquality-run

The input parameters is the run number and the data quality flag, if you want to know the data quality flag,

you have to use this command without input parameters.
for example:

```
(DIRAC3-user) zmathe@pclhcb43 /scratch/zmathe/dirac[406]>dirac-bookkeeping-setdataquality-run
Available data quality flags:
UNCHECKED
OK
BAD
MAYBE
Usage: dirac-bookkeeping-setdataquality-run.py <RunNumber> <DataQualityFlag>
```

The data quality flag is case sensitive. Set data quality a given run:

```
(DIRAC3-user) zmathe@pclhcb43 /scratch/zmathe/dirac[408]>dirac-bookkeeping-setdataquality-run 207
Quality flag has been updated!
(DIRAC3-user) zmathe@pclhcb43 /scratch/zmathe/dirac[409]>
```

dirac-bookkeeping-setdataquality-files

The input is a logical file name or a file. This file contains a list of lfns. Set the quality flag one file:

```
(DIRAC3-user) zmathe@pclhcb43 /scratch/zmathe/dirac[413]> dirac-bookkeeping-setdataquality-files
['/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/44026/044026_0000000002.raw']
Quality flag updated!
(DIRAC3-user) zmathe@pclhcb43 /scratch/zmathe/dirac[414]>
```

Set the quality flag a list of file:

```
(DIRAC3-user) zmathe@pclhcb43 /scratch/zmathe/dirac[416]> dirac-bookkeeping-setdataquality-files
['/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/44026/044026_0000000002.raw', '/lhcb/data/2009/RAW/EXPRES
Quality flag updated!
(DIRAC3-user) zmathe@pclhcb43 /scratch/zmathe/dirac[417]>
```

The lfns.txt contains the following:

```
/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/44026/044026_0000000002.raw
/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/43998/043998_0000000002.raw
/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/43995/043995_0000000002.raw
/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/43994/043994_0000000002.raw
/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/43993/043993_0000000002.raw
/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/43992/043992_0000000002.raw
/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/43989/043989_0000000002.raw
/lhcb/data/2009/RAW/EXPRESS/FEST/FEST/43987/043987_0000000002.raw
lfns.txt (END)
```

Web portal

Web portal stuck: how to restart it

First try to restart the Paster with runit:

```
runsvctrl t runit/Web/Paster
```

if this is not enough, then it will be necessary to get all Paster processes ('ps faux | grep -i web_paster') and do a 'kill -9' of them.

Documents

- Procedure Template (.doc)
 - Procedure Template (.rtf)
 - ProcedureTemplate (.txt)
 - Recovery Of Files Lost By SE
 - LHCb SAM Test documentation
 - Deploying LHCb Software on the Grid
 - Submitting Feature Requests and Bug Reports
-

This topic: LHCb > ProductionProcedures

Topic revision: r90 - 2017-08-01 - FedericoStagni



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
Ideas, requests, problems regarding TWiki? Send feedback