Production Shift Notes

June 2, 2009
DRAFT 1.1

LHCb Production Shift Notes

Sturt Paterson[1], Paul Szczypka[2]

[1] CERN
[2] EPFL

Abstract

Notes on operational procedures taken during the week beginning 18^{th} August 2008.

Contents

1 Introduction

This document describes the frequently-used tools and procedures available to Grid Shifters when managing production activities. It is expected that the current Grid Shifters should update this document by incorporating the stable procedures available on the LHCb Production TWiki pages [1] when appropriate. All files and document templates are available via CVS.

Section 3 gives some brief information about the various Grid sites and their backend storage systems. Section 4 details the jobs types a Grid Shifter is expected to encounter and provides some debugging methods. The methods available to manage and monitor productions are described in section 5. Section 6 describes the main features of the Production Monitor webpage.

A chronological guide through a production shift, from beginning to end, is presented in section 7. Section 8 outines the situations for which the submission of an ELOG is appropriate. Finally, section 9 details the well-established procedures for Grid Shifters.

A number of quick-reference sections are also available. Appendices A and B list the available DIRAC 3 scripts and commonly-used acronyms respectively.

2 Updating this Guide

To update this guide please use CVS. Set the following environment variables:

```
export CVSROOT=:ext:YourCERNUserName@isscvs.cern.ch:/local/reps/dirac
```

```
export CVS_RSH=ssh
```

2.1 Checkout

Then checkout the guide using the command:

```
cvs co Documentation
```

2.2 Updating and Editing

Ensure you update whatever file you propose to edit before you actually edit it:

```
cvs update <aFile>
```

You may now perform your edit.

2.2.1 Adding New Files and Directories

To add new files or directories use:

```
cvs add <myNewFile>
```

for each new directory. You may then add all the files in the directory using:

```
cvs add *
```

once you're in the new directory.

2.3 Commiting your Changes

Make the guide to ensure that everything works:

```
./make.sh
```

Assuming everything is ok, then clean up the current directory:

```
rm notes.{a,b,g,i,l,o,p,to}*
```

and perform one last update to check no one has also edited the file you are editing.

```
cvs update <aFile>
```

Finally perform a commit:

```
cvs ci
```

.

3 Grid Sites

Jobs submitted to the Grid will be scheduled to run at one of a number of Grid sites. The exact site at which a job is executed depends on the job requirements and the current status of all relevant grid sites. Grid sites are grouped into two tiers, Tier-1 and Tier-2. Cern is an exception, because it is also responsibe for processing and archiving the RAW experimental data it is also referred to as a Tier-0 site.

3.1 Tier-1 Sites

Tier-1 sites are used for Analysis, Monte Carlo production, file transfer and file storage in the LHCb Computing Model [2].

1. `LCG.CERN.ch`
2. `LCG.CNAF.it`
3. `LCG.IN2P3.fr`
4. `LCG.NIKHEF.nl`
5. `LCG.PIC.es`
6. `LCG.RAL.uk`
7. `LCG.GRIDKA.de`

3.2 Tier-2 Sites

There are numerous Tier-2 sites with sites being added frequently. As such, it is of little worth presenting a list of all the current Tier-2 sites in this document. Tier-2 sites are used for MC production in the LHCb Computing Model [2].

Backend Storage	Tier-1 Site
Castor	`CERN, CNAF, RAL`
dCache	`IN2P3, NIKHEF, GridKa, PIC`

Table 1: Backend storage technology employed at the Tier-1 sites.

3.3 Backend Storage Systems

Two backend storage technologies are employed at the Tier-1 sites, Castor and dCache. The Tier-1 sites which utilise each technology choice are summarised in table 1.

4 Jobs

The number of jobs created for a productions varies depending on the exact requirements of the production. Grid Shifters are generally not required to create jobs for a production.

4.1 Job ID's

A particular job is tagged with the following information:

- Production Identifier (ProdID), e.g. 00001234 - the 1234^{th} production.
- Job Idetifier (JobID), e.g. 9876 - the 9876^{th} job in the DIRAC system.
- JobName, e.g. 00001234_00000019 - the 19^{th} job in production 00001234.

4.2 Job Status

The job status of a successful job proceeds in the following order:

1. Received,
2. Checking,
3. Staging,
4. Waiting,
5. Matched,
6. Running,
7. Completed,
8. Done.

Jobs which return no heartbeat have a status of "Stalled" and jobs where any workflow modules return an error status are classed as "Failed".

The basic flowchart describing the evolution of a job's status can be found in figure 1. Jobs are only "Grid-active" once they have reached the "Matched" phase.

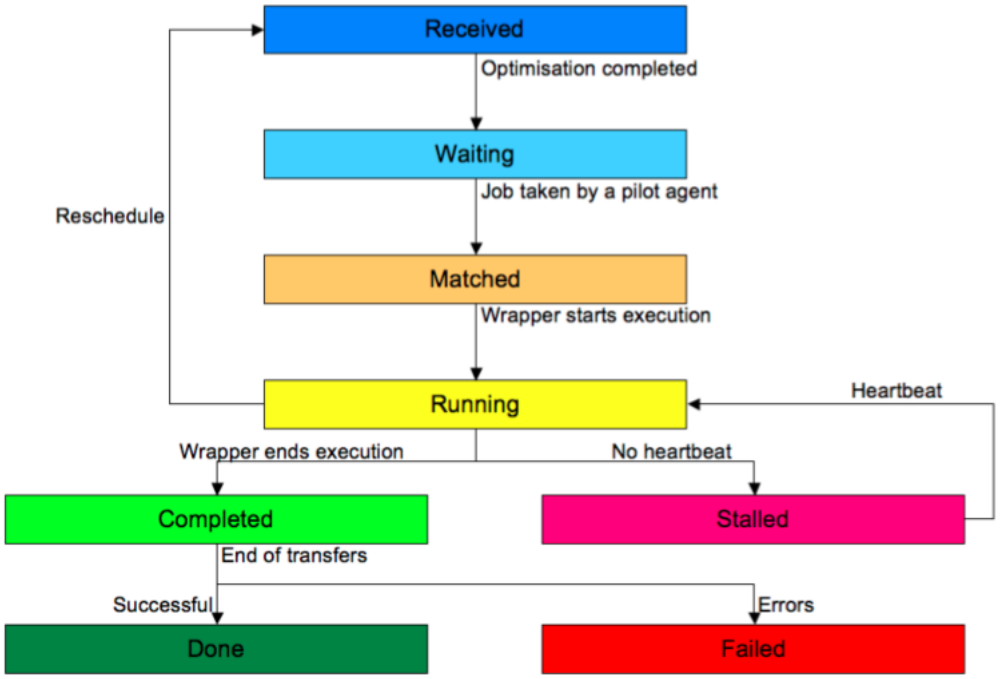

Figure 1: Job status flowchart. Note that the "Checking" and "Staging" status are omitted.

4.3 Job Output

The standard output and standard error of a job can be accessed through the API, the CLI and the webpage via a global job "peek".

4.3.1 Job Output Via the CLI

The `std.out` and `std.err` for a given job can be retrieved using the CLI command:

```
dirac-wms-job-get-output <JobID> |[<JobID>]
```

This creates a directory containing the `std.out` and `std.err` for each JobID entered. Standard tools can then be used to search the output for specific strings, e.g. "FATAL".

To simply view the last few lines of a job's `std.out` ("peek") use:

```
dirac-wms-job-peek <JobID> |[<JobID>]
```

4.3.2 Job Output Via the Production Monitoring Webpage

There are two methods to view the output of a job via the Job Monitoring Webpage [3]. The first returns the last 20 lines of the `std.out` and the second allows the Grid Shifter to view all the output files.

To "peek" the `std.out` of a job:

- Navigate to the Job Monitoring Webpage.

Figure 2: Peek the `std.out` of a job via the Job Monitoring Webpage.

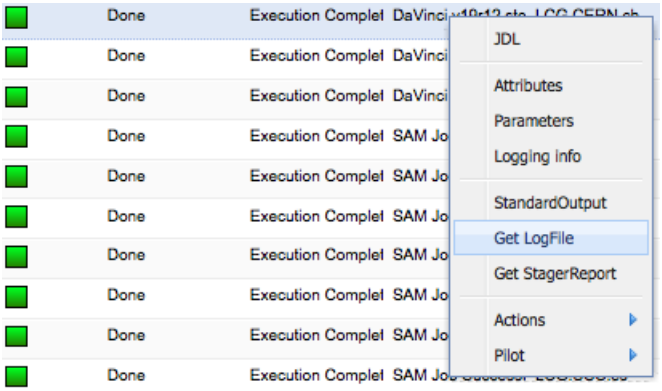

Figure 3: View all the output files of a job via the Job Monitoring Webpage.

- Select the relevant filters from the left panel.
- Click on a job.
- Select "StandardOutput" (Fig. 2).

Similarly, to view all output files for a job:

- Navigate to the Job Monitoring Webpage.
- Select the relevant filters from the left panel.
- Click on a job.
- Select "Get Logfile" (Fig. 3).

This method can be particularly quick if the Grid Shifter only wants to check the output of a selection of jobs.

4.4 Job Pilot Output

The output of the Job Pilot can also be retrieved via the API, the CLI or the Webpage.

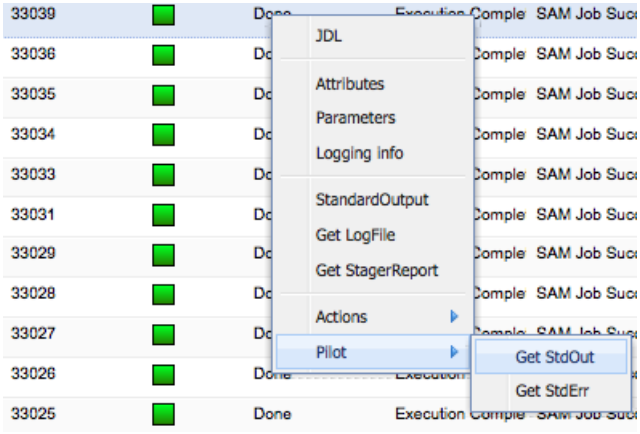

Figure 4: View the pilot output of a job via the Job Monitoring Webpage.

4.4.1 Job Pilot Output via the CLI

To obtain the Job Pilot output using the CLI, use:

```
dirac-admin-get-pilot-output <Grid pilot reference> [<Grid pilot reference>]
```

This creates a directory for each JobID containing the Job Pilot output.

4.4.2 Job Pilot Output via the Job Monitoring Webpage

Viewing the `std.out` and `std.err` of a Job Pilot via the Job Monitoring Webpage is achieved by:

- Navigate to the Job Monitoring Webpage.
- Select the relevant filters from the left panel.
- Click on a job.
- Select "Pilot" then "Get StdOut" or "Get StdErr" (Fig. 4).

4.5 Operations on Jobs

The full list of scripts which can be used to perform operations on a job is given in (App. A.19). The name of each script should be a clear indication of it's purpose. Running a script without arguments will print basic usage notes.

5 Productions

As a Grid Shifter you will be required to monitor the official LHCb productions. Each production is assigned a unique Production ID (ProdID). These consist of Monte Carlo

(MC) generation, data stripping and CCRC productions. Production creation will generally be performed by the Production Operations Manager and is not a duty of the Grid Shifter.

The current list of all active productions can be obtained with the command:

```
dirac-production-list-active
```

The command also gives the current submission status of the active productions.

5.1 Starting a Production

The submission of a production can be started once it has been formulated and all the required jobs created. Grid Shifters should ensure they have the permission of the Production Operations Manager (or equivalent) before starting a production (Sec. 5.1.1). Production jobs can be submitted manually (Sec. 5.1.2) or automatically (Sec. 5.1.3).

The state of a production can also be set using:

```
dirac-production-change-status <Command> <Production ID> |<Production ID>
```

where the available commands are:

```
'start', 'stop', 'manual', 'automatic'
```

5.1.1 Starting and Stopping a Production

The commands:

```
dirac-production-start <Production ID> |<Production ID>
```

and

```
dirac-production-stop <Production ID> |<Production ID>
```

are used to start and stop a production. Grid Shifters may stop a current production if a significant number of jobs are failing.

5.1.2 Manual Submission

A production is set to manual submission by default. To reset the submission status of one or more productions, use the command:

```
dirac-production-set-manual <Production ID> |<Production ID>
```

A small number of test jobs should be manually submitted for each new production. In the case of stripping or CCRC productions, a small number of test jobs should be sent to all the Tier 1 sites and closely monitored.

To manually submit jobs to a selected site, use the following command:

```
dirac-production-site-submit <ProdID> <Num Jobs> <Site>
```

Note that the full site name string must be entered, e.g. to submit a job to CERN you must type:

```
dirac-production-site-submit <ProdID> 1 LCG.CERN.ch
```

Any observed problems or job failures should be investigated and an ELOG entry submitted. Assuming there are no problems in all the test jobs, the production may be set to automatic submission.

5.1.3 Automatic Submission

When started, a production set to automatic submission will submit all jobs in the production in quick succession.

A production can be set to automatic submission once you are satisfied that there are no specific problems with the production jobs. To set a production to automatic submission use:

```
dirac-production-set-automatic <Production ID> |<Production ID>
```

5.2 Monitoring a Production

Jobs in each production should be periodically checked for failed jobs (Sec. 5.2.1) and to ensure that jobs are progressing (Sec. 5.2.2).

When monitoring a production, a Grid Shifter should be aware of a number of issues which can cause jobs to fail:

- Staging.
- Stalled Jobs.
- Segmentation faults.
- DB access.
- Software problems.
- Data access.
- Shared area access.
- Site downtime.
- Problematic files.
- Excessive runtime.

5.2.1 Failed Jobs

A Grid Shifter should monitor a production for failed jobs and jobs which are not progressing. Due to the various configurations of all the sites it is occasionally not possible for an email to be sent to the `lhcb-datacrash` [4] mailing list for each failed job. It is therefore not enough to simply rely on the number of `lhcb-datacrash` emails to indicate if there are any problems with a production. In addition to any `lhcb-datacrash` notifications, the Grid Shifter should also check the number of failed jobs in a production via the CLI or the Production Monitoring Webpage [5].

Using the CLI, the command:

```
dirac-production-progress [<Production ID>]
```

entered without any arguments will return a breakdown of the jobs of all current productions. Entering one or more ProdIDs returns only the breakdown of those productions.

A more detailed breakdown is provided by:

```
dirac-production-job-summary <Production ID> [<DIRAC Status>]
```

which also includes the minor status of each job category and provides an example JobID for each category. The example JobIDs can then be used to investigate the failures further.

5.2.2 Non-Progressing Jobs

In addition to failed jobs, jobs which do not progress should also be monitored. Particular attention should be paid to jobs in the states "Waiting" and "Staging". Problematic jobs at this stage are easily overlooked since the associated problems are not easily identifiable.

5.2.3 Non-Starting Jobs

Jobs arriving at a site but then failing to start have multiple causes. One of the most common reasons is that a site is due to enter sceduled downtime and are no longer submitting jobs to the batch queues. Jobs will stay at the site in a "Waiting" state and state that there are no CE's available. Multiple jobs in this state should be reported.

5.3 Ending a Production

Ending a completed production is handeled by the Productions Operations Manager (or equivalent). No action is required on the part of the Grid Shifter.

5.4 Operations on Productions

All CLI scripts which can be used to manage productions are listed in (App. A.16). Running a script without arguments will return basic usage notes. In some cases further help is available by running a script with the option `--help`.

6 Web Production Monitor

Production monitoring via the web is possible through the Production Monitoring Webpage [5]. A valid grid certificate loaded into your browser is required to use the webpage.

6.1 Features

The Production Monitoring Webpage has the following features:

- Site Downtime Calendar [6].
- Various Plot Views [7]
- Production Monitor [5]
- Job Monitor [3]

6.1.1 Site Downtime Calendar

The calendar [6] displays all the sites with scheduled and unscheduled downtime. Calendar entries are automatically parsed through the site downtime RSS feed and added to the calendar.

Occasionally the feed isn't parsed correctly and Grid Shifters should double-check that the banned and allowed sites are correct. Useful scripts for this are:

```
dirac-admin-get-banned-sites
```

and

```
dirac-admin-get-site-mask
```

6.2 Plots

The Production Monitoring Webpage has the capacity to produce various plots. Many of which are extremely useful to monitor the performace of the production system.

Links to useful plots can be found on the DIRAC System Monitoring Pages [8]. These plots should be monitored at three times daily.

6.3 Buglist and Feature Request

The procedure to submit a bug report or a feature request is outlined in section 9.

7 Shifts

Grid Shifters are required to monitor all the current LHCb productions and must have a valid Grid Certificate and be a member of the LHCb VO.

7.1 Before a Shift Period

The new shifter should:

- Ensure their Grid certificate is valid for all expected duties (Sec. 7.1.2).
- Create accounts on all relevant web-resources (Sec. 7.1.2).
- Subscribe to the relevant mailing lists (Sec. 7.1.3).

7.1.1 Grid Certificates

A Grid certificate is mandatory for Grid Shifters. If you don't have a certificate you should register for one through CERN LCG [9] and apply to join the LHCb Virtual Organisation (VO).

To access the production monitoring webpages [5] you will also need to load your certificate into your browser. Detailed instructions on how to do this can be found on the CERN LCG pages [10].

7.1.2 Web Resorces

Primary web-based resources for DIRAC 3 production shifts:

- ELOG [11]
- Production Monitoring [5]
- Production Twiki [1]

7.1.3 Mailing Lists

The new Grid Shifter should subscribe to the following mailing lists:

- `lhcb-datacrash` [12, 4, 13].
- `lhcb-dirac-developers` [13].
- `lhcb-dirac` [13].
- `lhcb-production` [13].

Note that both the `lhcb-datacrash` and `lhcb-production` mailing lists recieve a substantial amount of mail daily. It's suggested that suitable message filters and folders are created in your mail client of choice.

7.1.4 Productions Operations Key

The new shifter should obtain the Production Operations key (TCE 5) from the LHCb secretariat or the previous Grid Shifter.

7.2 During a Shift

During a shift Grid Shifters are expected to monitor all current productions and be aware of the current status of the Tier 1 sites. A knowledge of the purpose of each production is also useful and aids in determining the probable cause of any failed jobs.

7.2.1 Daily Actions

Grid Shifters are expected to carry out the following daily actions for sites used in the current productions:

- Trigger submission of pending productions.
- Monitor active productions.
- Check transfer status.
- Verify that the staging at each site is functional.
- Check that there is a minimum of one successful (and complete) job.
- Confirm that data access is working at least intermittently.
- Report problems to the operations team.
- View the plots in the DIRAC System Monitoring Webpage [14]
- Monitor the current performance of the lhcb-production instance 7.2.2.
- Submit a summary of the job status at all the grid sites to the ELOG 8.

7.2.2 Performance Monitoring

Grid Shifters should view the plots accessible via the DIRACSystemMonitoring page [8] at least three times a day and investigate any unusal features present.

7.2.3 Production Operations Meeting

A Production Operations Meeting takes place at the end of the morning shift and allows the morning Grid Shifter to highlight any recent or outstanding issues. Both the morning and afternoon Grid Shifter should attend. The morning Grid Shifter should give a report summerising the morning activities.

The Grid Shifter's report should contain:

- Current production progress, jobs submitted, wating etc.
- Status of all Tier 1 sites.
- Recently observed failures, paying particular attention to previously-unknown problems.

7.3 Ending a Shift

At the end of each shift, morning Grid Shifters should:

- Pass on the key (TCE 5) for the Production Operations room to the next Grid Shifter.
- Prepare a list of outstanding issues to be handed over to the next Grid Shifter and discussed in the Production Operations meeting.
- Submit an ELOG report summarising the shift and any ongoing or unresolved issues.

Similarly, evening Grid Shifters should:

- Place the key (TCE 5) to the Productions Operations room in the secretariat key box.
- Submit an ELOG report summarising the shift and any ongoing or unresolved issues.

7.4 End of Shift Period

At the end of a shift period the Grid Shifter may wish to unsubscribe from the various mailing lists (Sec. 7.4.1) in addition to returning the Production Operations room key, TCE 5 (Sec. 7.4.2).

7.4.1 Mailing Lists

Unsubscribe from the following mailing lists:

- `lhcb-datacrash` [12, 4, 15].
- `lhcb-dirac-developers` [15].
- `lhcb-dirac` [15].
- `lhcb-production` [13].

7.4.2 Miscellaneous

Return the key for the Production Operations Room (TCE 5) to the secretariat or the next Grid Shifter.

7.5 Weekly Report

A weekly report should be prepared by the Grid Shifter at the end of each week. The report should contain information on all the processed production and user jobs, the respective failure rates and some basic analysis of the results.

The report should be compiled on the last day of the shift and contain information about the previous seven full days of operation, i.e. it should not include information from the day the report is compiled.

The weekly reports are to be uploaded to the Weekly Reports Page [16] on the LHCb Computing tWiki. Grid Shifters should use the template [17] provided when compiling a report.

7.5.1 Base Plots

The following plots should always be included in the report:

- Total number of Jobs by Final Major Status
- Daily number of Jobs by Final Major Status
- Done—Completed Jobs by User Group
- Done—Completed Production Jobs by JobType
- Failed Jobs by User Group
- Failed Production Jobs by Minor Status
- Failed User Jobs by Minor Status
- Done—Completed Production Jobs by Site
- Done—Completed User Jobs by Site

From these plots the Grid Shifter should then create a number of further plots to analyse the causes and execution locations of failed jobs.

7.5.2 Specific Plots

On analysis of the failed jobs, the Grid Shifter should produce plots of the breakdown by site of all failed jobs with the three or four main job "MinorStatus" results.

- Failed Production Jobs by FinalMinorStatus
  - Failed Production Jobs (FinalMinorStatus 1) by Site
  - Failed Production Jobs (FinalMinorStatus 2) by Site
  - Failed Production Jobs (FinalMinorStatus 3) by Site
  - ...
- Failed User Jobs by FinalMinorStatus
  - Failed User Jobs (FinalMinorStatus 1) by Site
  - Failed User Jobs (FinalMinorStatus 2) by Site
  - Failed User Jobs (FinalMinorStatus 3) by Site
  - ...

7.5.3 Machine Monitoring Plots

Monitoring of the LHCb VO boxes is vital to maintaining the efficient running of all Grid operations. Particular attention should be paid to the used and free space on the various disks, network and CPU usage.

Reports on the state of the following boxes should be constructed:

- `vobox01`
- `vobox06`
- `vobox09`
- `vobox10`

For each machine, save and then upload the plots for:

- CPU utilization
- Network utilization
- Partition Used
- Swap Used

Note: Mac users may find that the suggested name when saving the plots does not follow the format "∗.gif.png" and they should take care to either rename the saved files or edit that week's report page accordingly.

7.5.4 Analysis and Summary

A summary of each group of plots should be written to aid the next Grid Shifter's appraisal of the current situation and to enable the Grid Expert on duty to investigate problems further.

8 ELOG

All Grid Shifter actions of note should be recorded in the ELOG [11]. This had the benefits of allowing new Grid Shifters to familarise themselves with recent problems with current productions. ELOG entries should contain as much relevant information as possible.

8.1 Typical ELOG Format

Each ELOG entry which reports a new problem should include as much relevant information as possible. This allows the production operations team to quickly determine the problem and apply a solution.

8.1.1 ELOG Entry for a New Problem

A typical ELOG entry for a new problem contains:

- The relevant ProdID or ProdIDs.
- An example JobID.
- A copy of the relevant error message and output.
- The number of affected jobs.
- The Grid sites affected.
- The time of the first and last occurance of the problem.

8.1.2 Subsequent ELOG Entries

Once a problem has been logged it is useful to report the continuing status of the affected productions at the end of each shift.

If a Grid Shifter is unsure whether a problem has been previously logged then they should submit a fresh ELOG following the format outlined in (Sec. 8.1.1) .

8.2 When to Submit an ELOG

Submit an elog in the following situations:

- Jobs finalise with exceptions.

8.2.1 Exceptions

Jobs which finalise with an exception should be noted in the ELOG. The ELOG entry should contain:

- The production ID.
- An example job ID.
- A copy of the relevant error messages.
- The number of jobs in the production which have the same status.

8.2.2 Crashed Application

Submit an example error log for the crashed application.

8.2.3 Datacrash Emails

The Grid Shifter should filter the datacrash emails and determine if the crash reported is actually due to one of the applications. If so, then the Grid Shifter should submit an ELOG describing the problem and including an example error message. The Grid Shifter should ensure the "Applications" radio button is selected when submitting the ELOG report since this means that the relevant experts will be alerted to the problem.

8.2.4 ELOG Problems

If ELOG is down, send a notification email to `lhcb-production@cern.ch`.

9 Procedures

If a problem is discovered it is very important to escalate it to the operations team. Assessing the scale of the problem is very important and Grid Shifters should attempt to answer the questions in section 9.1.1 as soon as possible.

9.1 On the Discovery of a Problem

Once a problem has been discovered it is important to assess the severity of the problem. Section 9.1.1 provides a checklist which the Grid Shifter should go through after discovering a problem. Additionally, there are a number of Grid-specific issues to consider (Sec. 9.1.2) .

9.1.1 Standard Checklist

On the discovery of a new problem, attempt to provide answers to the following questions as quickly as possible:

- How many jobs does the problem affect?
- Are the central DIRAC services running normally?
- Are all jobs affected?
- When did the problem start?
- When did the last successful job run in similar conditions?
- Is it a DIRAC problem?
  - Can extra redundancy be introduced to the system?
  - Is there enough information available to determine the error?

Figure 5: Browse current support issues.

9.1.2 Grid-Specific Issues

- Was there an announcement of downtime for the site?
- Is the problem specific to a single site?
  - Are all the CE's at the site affected?
- Is the problem systamatic across sites with different backend storage technologies? (Sec. 3.3)
- Is the problem specific to an SE?
  - Are there any stalled jobs at the site clustered in time?
  - Are other jobs successfully reading data from the SE?

9.2 Feature Requests

Before submitting a feature request, the user should:

- Identify conditions under which the feature is to be used.
- Record all relevant information.
- Identify a use-case for the new feature.

Once the user has prepared all the relevant information, they should:

- Login at the DIRAC Savannah Webpage.
- Browse the existing support issues and ensure the feature request has not been previously submitted (Fig. 5).

Assuming the feature request has not been previously submitted, the user should then:

- Navigate to the "Support" tab at the top of the page (Fig. 6) and click on "submit".
- Ensure that the submission webform contains all relevant information (Fig. 7).
- Set the severity option to "wish".
- Set the privacy option to "private".
- Submit the feature request.

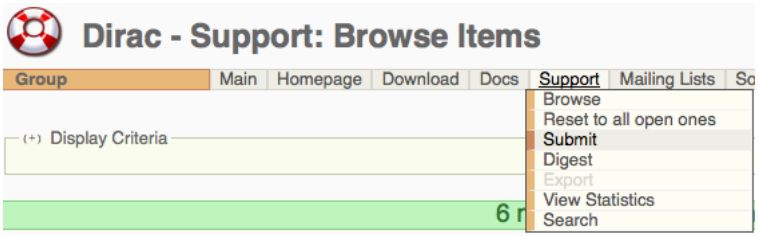

Figure 6: Savannah support submit.

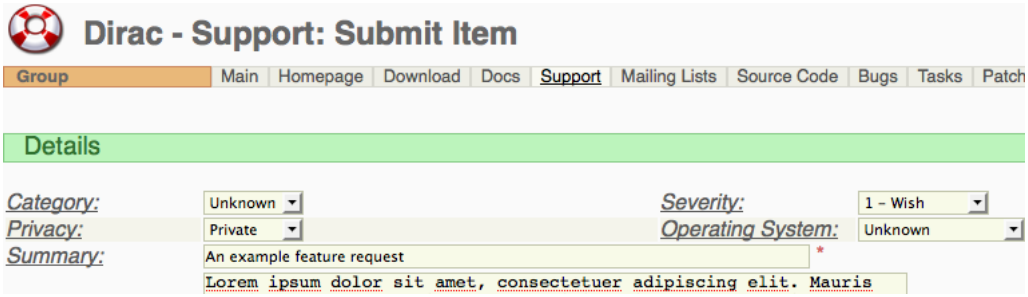

Figure 7: Savannah support submit feature request.

9.3 Bug Reporting

Before submitting a bug report, the user should:

- Identify conditions under which the bug occurs.
- Record all relevant information.
- Try to ensure that the bug is reproducible.

Once the user is convinced that the behaviour they are experiencing is a bug, they should then prepare to submit a bug report. Users should:

- Login at the DIRAC Savannah Webpage.
- Browse the existing bugs and ensure the user's bug has not been previously submitted (Fig. 8).

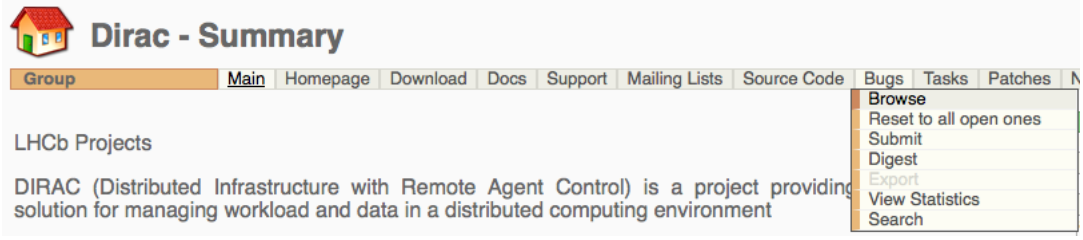

Figure 8: Browse current bugs.

Assuming the bug is new, the procedure to submit a bug report is as follows:

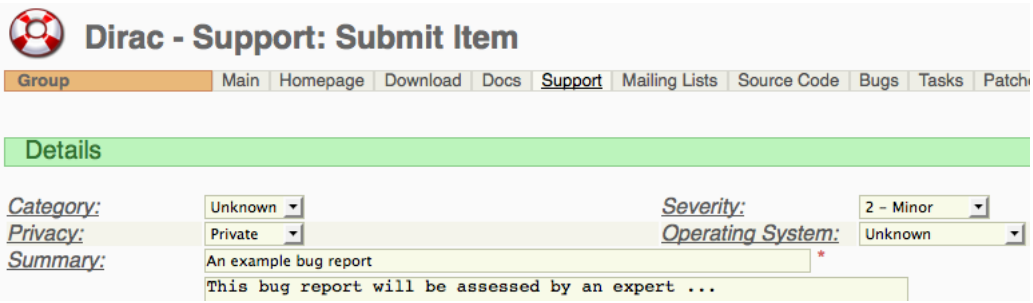

Figure 9: Example bug report.

- Navigate to the "Support" tab at the top of the page (Fig. 6) and click on "submit".
- Ensure that the submission webform contains all relevant information (Fig. 9).
- Set the appropriate severity of the problem.
- Write a short and clear summary.
- Set the privacy option to "private".
- Submit the bug report.

9.4 Software Unavailability

Symptom: Jobs fail to find at least one software package.

Software installation occurs during Service Availability Monitoring (SAM) tests. Sites which fail to find software packages should have failed at least part of their most recent SAM test.

Grid Shifter actions:

- Submit an ELOG report listing the affected productions and sites.
- Ban the relevant sites until they pass their SAM tests.

A DIRAC 3 Scripts

A.1 DIRAC Admin Scripts

- dirac-admin-accounting-cli
- dirac-admin-add-user
- dirac-admin-allow-site
- dirac-admin-ban-site
- dirac-admin-delete-user
- dirac-admin-get-banned-sites
- dirac-admin-get-job-pilot-output
- dirac-admin-get-job-pilots
- dirac-admin-get-pilot-output
- dirac-admin-get-proxy
- dirac-admin-get-site-mask
- dirac-admin-list-hosts
- dirac-admin-list-users
- dirac-admin-modify-user
- dirac-admin-pilot-summary
- dirac-admin-reset-job
- dirac-admin-service-ports
- dirac-admin-site-info
- dirac-admin-sync-users-from-file
- dirac-admin-upload-proxy
- dirac-admin-users-with-proxy

A.2 DIRAC Bookkeeping Scripts

- dirac-bookkeeping-eventMgt
- dirac-bookkeeping-eventtype-mgt
- dirac-bookkeeping-ls
- dirac-bookkeeping-production-jobs
- dirac-bookkeeping-production-informations

A.3 DIRAC Clean

- dirac-clean

A.4 DIRAC Configuration

- dirac-configuration-cli

A.5 DIRAC Distribution

- dirac-distribution

A.6 DIRAC DMS

- dirac-dms-add-file
- dirac-dms-get-file
- dirac-dms-lfn-accessURL
- dirac-dms-lfn-logging-info
- dirac-dms-lfn-metadata
- dirac-dms-lfn-replicas
- dirac-dms-pfn-metadata
- dirac-dms-pfn-accessURL
- dirac-dms-remove-pfn
- dirac-dms-remove-lfn
- dirac-dms-replicate-lfn

A.7 DIRAC Embedded

- dirac-embedded-external

A.8 DIRAC External

- dirac-external

A.9 DIRAC Fix

- dirac-fix-ld-library-path

A.10 DIRAC Framework

- dirac-framework-ping-service

A.11 DIRAC Functions

- dirac-functions.sh

A.12 DIRAC Group

- dirac-group-init

A.13 DIRAC Jobexec

- dirac-jobexec

A.14 DIRAC LHCb

- dirac-lhcb-job-replica
- dirac-lhcb-manage-software
- dirac-lhcb-production-job-check
- dirac-lhcb-sam-submit-all
- dirac-lhcb-sam-submit-ce

A.15 DIRAC Myproxy

- dirac-myproxy-upload

A.16 DIRAC Production

- dirac-production-application-summary
- dirac-production-change-status
- dirac-production-job-summary

- dirac-production-list-active
- dirac-production-list-all
- dirac-production-list-id
- dirac-production-logging-info
- dirac-production-mcextend
- dirac-production-manager-cli
- dirac-production-progress
- dirac-production-set-automatic
- dirac-production-set-manual
- dirac-production-site-summary
- dirac-production-start
- dirac-production-stop
- dirac-production-submit
- dirac-production-summary

A.17 DIRAC Proxy

- dirac-proxy-info
- dirac-proxy-init
- dirac-proxy-upload

A.18 DIRAC Update

- dirac-update

A.19 DIRAC WMS

- dirac-wms-job-delete
- dirac-wms-job-get-output
- dirac-wms-job-get-input
- dirac-wms-job-kill
- dirac-wms-job-logging-info

- `dirac-wms-job-parameters`
- `dirac-wms-job-peek`
- `dirac-wms-job-status`
- `dirac-wms-job-submit`
- `dirac-wms-job-reschedule`

B Common Acronyms

Glossary

ACL Access Control Lists, 26
API Application Programming Interface, 26
ARC Advance Resource Connector, 26
ARDA A Realisation of Distributed Analysis, 26

BDII Berkeley Database Information Index, 26
BOSS Batch Object Submission System, 26

CA Certification Authority, 26
CAF CDF Central Analysis Farm, 26
CCRC Common Computing Readiness Challenge, 6
CDF Collider Detector at Fermilab, 26
CE Computing Element, 26
CERN Organisation Européenne pour la Recherche Nucléaire: Switzerland/France, 2
CNAF Centro Nazionale per la Ricerca e Svilupponelle Tecnologie Informatiche e Telematiche: Italy, 2
ConDB Conditions Database, 26
CPU Central Processing Unit, 26
CRL Certificate Revocation List, 26
CS Configuration Service, 26

DAG Directed Acyclic Graph, 26
DC04 Data Challenge 2004, 26
DC06 Data Challenge 2006, 26
DCAP Data Link Switching Client Access Protocol, 26
DIAL Distributed Interactive Analysis of Large datasets, 26
DIRAC Distributed Infrastructure with Remote Agent Control, 26
DISET DIRAC Secure Transport, 26
DLI Data Location Interface, 26
DLLs Dynamically Linked Libraries, 26
DN Distinguished Name, 26
DNS Domain Name System, 26
DRS Data Replication Service, 26
DST Data Summary Tape, 26

ECAL Electromagnetic CALorimeter, 26
EGA Enterprise Grid Alliance, 26
EGEE Enabling Grids for E-sciencE, 26
ELOG Electronic Log, 1
ETC Event Tag Collection, 26

FIFO First In First Out, 26
FTS File Transfer Service, 26

Ganga Gaudi / Athena and Grid Alliance, 26
GASS Global Access to Secondary Storage, 26
GFAL Grid File Access Library, 26
GGF Global Grid Forum, 26
GIIS Grid Index Information Service, 26
GLUE Grid Laboratory Uniform Environment, 26
GRAM Grid Resource Allocation Manager, 26
GridFTP Grid File Transfer Protocol, 26
GridKa Grid Computing Centre Karlsruhe, 2
GriPhyN Grid Physics Network, 26
GRIS Grid Resource Information Server, 26
GSI Grid Security Infrastructure, 26
GT Globus Toolkit, 26
GUI Graphical User Interface, 26
GUID Globally Unique IDentifier, 26

HCAL Hadron CALorimeter, 26
HEP High Energy Physics, 26
HLT High Level Trigger, 26
HTML Hyper-Text Markup Language, 26
HTTP Hyper-Text Transfer Protocol, 26

I/O Input/Output, 26
IN2P3 Institut National de Physique Nucleaire et de Physique des Particules: France, 2
iVDGL International Virtual Data Grid Laboratory, 26

JDL Job Description Language, 26
JobDB Job Database, 26
JobID Job Identifier, 3

L0 Level 0, 26
LAN Local Area Network, 26
LCG LHC Computing Grid, 26

LCG IS LCG Information System, 26
LCG UI LCG User Interface, 26
LCG WMS LCG Workload Management System, 26
LDAP Lightweight Directory Access Protocol, 26
LFC LCG File Catalogue, 26
LFN Logical File Name, 26
LHC Large Hadron Collider, 26
LHCb Large Hadron Collider beauty, 26
LSF Load Share Facility, 26

MC Monte Carlo, 6
MDS Monitoring and Discovery Service, 26
MSS Mass Storage System, 26

NIKHEF National Institute for Subatomic Physics: Netherlands, 2

OGSA Open Grid Services Architecture, 26
OGSI Open Grid Services Infrastructure, 26
OSG Open Science Grid, 26

P2P Peer-to-peer, 26
Panda Production ANd Distributed Analysis, 26
PC Personal Computer, 26
PDC1 Physics Data Challenge, 26
PFN Physical File Name, 26
PIC Port d'Informació Científica: Spain, 2
PKI Public Key Infrastructure, 26
POOL Pool Of persistent Objects for LHC, 26
POSIX Portable Operating System Interface, 26
PPDG Particle Physics Data Grid, 26
ProdID Production Identifier, 3, 6
PS Preshower Detector, 26

R-GMA Relational Grid Monitoring Architecture, 26
RAL Rutherford-Appleton Laboratory: UK, 2
RB Resource Broker, 26
rDST reduced Data Summary Tape, 26
RFIO Remote File Input/Output, 26
RICH Ring Imaging CHerenkov, 26
RM Replica Manager, 26
RPC Remote Procedure Call, 26
RTTC Real Time Trigger Challenge, 26

SAM Service Availability Monitoring, 20
SE Storage Element, 18, 26
SOA Service Oriented Architecture, 26
SOAP Simple Object Access Protocol, 26
SPD Scintillator Pad Detector, 26
SRM Storage Resource Manager, 26
SSL Secure Socket Layer, 26
SURL Storage URL, 26

TCP/IP Transmission Control Protocol / Internet Protocol, 26
TDS Transient Detector Store, 26
TES Transient Event Store, 26
THS Transient Histogram Store, 26
TT Trigger Tracker, 26
TURL Transport URL, 26

URL Uniform Resource Locator, 26

VDT Virtual Data Toolkit, 26
VELO VErtex LOcator, 26
VO Virtual Organisation, 11, 26
VOMS Virtual Organisation Membership Service, 26

WAN Wide Area Network, 26
WMS Workload Management System, 26
WN Worker Node, 26
WSDL Web Services Description Language, 26
WSRF Web Services Resource Framework, 26
WWW World Wide Web, 26

XML eXtensible Markup Language, 26
XML-RPC XML Remote Procedure Call, 26

References

[1] LHCb Collaboration, LHCb Production TWiki, Website, https://twiki.cern.ch/twiki/bin/view/LHCb/LHCbComputing.

[2] N. Brook, CERN Report No. LHCb-2004-119. CERN-LHCb-2004-119, 2004 (unpublished).

[3] LHCb, DIRAC Job Monitoring Webpage, Website, https://lhcbweb.pic.es/DIRAC/jobs/JobMonitor/display.

[4] Various, LHCb Datacrash Mailing List Archive, Mailing List, https://mmm.cern.ch/public/archive-list/l/lhcb-datacrash/.

[5] LHCb, DIRAC Production Monitoring, Website, https://lhcbweb.pic.es/DIRAC/jobs/ProductionMonitor/display.

[6] DIRAC, DIRAC: calendar, Website, https://lhcbweb.pic.es/DIRAC/web/External/display?sit

[7] LHCb, DIRAC Plot Views, Website, https://lhcbweb.pic.es/DIRAC/systems/monitoring/plot

[8] CERN, DIRAC System Monitoring, Website, https://twiki.cern.ch/twiki/bin/view/LHCb/DIRACSystemMonitoring .

[9] CERN, LCG Users Registration, Website, http://lcg.web.cern.ch/lcg/users/registration/registration.html .

[10] CERN, Loading certificates into a browser, Website, http://lcg.web.cern.ch/lcg/users/registration/load-cert.html .

[11] LHCb Collaboration, ELOG operations, Website, http://lblogbook.cern.ch/Operations/.

[12] Various, LHCb Datacrash Simba Page, Website, https://websvc03.cern.ch/listboxservices/simba2/listeditor.aspx?list=lhcb-datacrash.

[13] LHCb , Simba - Listbox Services: Subscribe, Website, https://websvc03.cern.ch/listboxservices/simba2/free/subscription.aspx/.

[14] LHCb Collaboration, LHCb DIRAC System Monitoring Plots, Website, https://twiki.cern.ch/twiki/bin/view/LHCb/DIRACSystemMonitoring.

[15] LHCb, Simba unsubscribe, Website, https://websvc03.cern.ch/listboxservices/simba2/free/uns

[16] LHCb Collaboration, LHCb Production TWiki, Website, https://twiki.cern.ch/twiki/bin/view/LHCb/DIRACWeeklyReport.

[17] LHCb Collaboration, LHCb Production TWiki Weekly Report Template, Website, https://twiki.cern.ch/twiki/bin/view/LHCb/DIRACWeeklyReportTemplate.