

# Software Quality, Testing & Validation and CI at LHCb



Chris Burr  
27th October 2020



- Testing and validation means many different things
  - Crashes
  - Performance
  - Obviously doing the wrong thing ⚠
  - Subtly doing the wrong thing! 💀
- Different use cases => multiple CI systems in use
  - Standard GitLab CI
  - “Nightly” tests in Jenkins (currently being migrated to a custom alternative)
  - Customised GitLab CI





General CI



- Most LHCb code is kept in a CERN-hosted instance of GitLab
  - Large number of “standard” shared GitLab CI runners available
  - A few specialised GitLab CI runners (privileged or additional CPU+RAM)
  - Tend to avoid dedicated GitLab CI runners - poor resource utilisation
- GitLab CI is used by almost every repository in LHCb
  - Check for code formatting and copyright notices
  - Unit tests: All core software packages
  - Some larger scale testing: LHCbDIRAC and the new GPU HLT application (Allen)
  - Some analyses are also using CI (from unit tests to “run the entire analysis on every commit”)
    - Potentially interesting for analysis preservation





# CI for Physics Applications



## Nightly builds

- Run nightly builds and tests for all Physics applications
  - Large matrix of projects, platforms and branches
  - Thousands of jobs are ran each day

lhcb-head - build: 2750 (2020-10-26)

head of everything against Gaudi/master and LCG\_97a

available on: **cvmfs**

Compare with previous build

Compare with other slot

Rebuild

Deploy on CVMFS

Browse files

		x86_64-centos7-gcc9-opt  Completed at 06:56:03 <div>Platform diff</div>		x86_64-centos7-gcc9-dbg  Completed at 11:56:23 <div>Platform diff</div>		x86_64-centos7-clang8-opt  Completed at 04:04:35 <div>Platform diff</div>		x86_64-centos7-clang8-dbg  Completed at 08:20:56 <div>Platform diff</div>		x86_64-centos7-gcc9-do0  Completed at 09:00:14 <div>Platform diff</div>		x86_64+avx2+fma-centos7-gcc9-opt  Completed at 09:47:28 <div>Platform diff</div>		skylake_avx512+vecwid256-centos7-gcc9-opt  Completed at 08:06:56 <div>Platform diff</div>		x86_64-centos7-gcc9+py3-opt  Completed at 04:09:22 <div>Platform diff</div>		x86_64-centos7-gcc9+py3-dbg  Completed at 07:33:15 <div>Platform diff</div>	
Project 	Version	build		build		build		build		build		build		build		build		build	
PARAM	None	build		build		build		build		build		build		build		build		build	
DBASE	None	build		build		build		build		build		build		build		build		build	
LCG		97a																	
Gaudi	master	build (4)	tests	build (4)	tests	build	tests	build	tests	build (4)	tests	build (4)	tests (1)	build (4)	tests	build (4)	tests	build (4)	tests
Online	HEAD	build	tests	build	tests	build	tests	build	tests (1)	build	tests	build	tests (1)	build	tests (3)	build	tests (6)	build	tests (6)
Detector	v0-patches	build (1)	tests (0)	build (1)	tests (0)	build (1)	tests (0)	build (1)	tests (0)	build (1)	tests (0)	build (1)	tests (0)	build (1)	tests (0)	build (1)	tests (0)	build (1)	tests (0)
LHCb	HEAD	build	tests	build	tests	build	tests	build	tests (1)	build	tests	build	tests	build	tests	build	tests	build	tests
Lbcom	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests
Boole <span>PR</span>	HEAD	build	tests (10)	build	tests (10)	build	tests (10)	build	tests (10)	build	tests (10)	build	tests (5)	build	tests (5)	build	tests (10)	build	tests (10)
Rec	HEAD	build	tests	build	tests	build	tests	build	tests (1)	build	tests	build	tests	build	tests	build	tests	build	tests
Brunel <span>PR</span>	HEAD	build	tests (8)	build	tests (8)	build	tests (8)	build	tests (8)	build	tests (8)	build	tests (8)	build	tests (8)	build	tests (8)	build	tests (8)
Phys	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests
Allen	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests
Moore <span>PR</span>	HEAD	build	tests (14)	build	tests (14)	build	tests (14)	build	tests (14)	build	running	build	tests (14)	build	tests (14)	build	tests (14)	build	tests (14)
Analysis	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests
DaVinci	HEAD	build	tests	build	tests	build	tests	build	tests (2)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests	build	tests
Panoramix	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests (1)	build	tests (1)
Bender	HEAD	build	tests (71)	build	tests (72)	build	tests (71)	build	running	build	tests (72)	build	tests (71)	build	tests (71)	build	tests (89)	build	tests (89)
MooreOnline	HEAD	build (4)	tests (1)	build (4)	tests (1)	build (8)	tests (1)	build (8)	tests (1)	build (4)	tests (1)	build (4)	tests (1)	build (4)	tests (1)	build (5)	tests (1)	build (5)	tests (1)
Panoptes	HEAD	build (5)	tests	build (5)	tests	build (5)	tests	build (5)	tests	build (5)	tests	build (5)	tests	build (5)	tests	build (2)	tests	build (2)	tests
Alignment	HEAD	build	tests (1)	build	tests (4)	build (4)	tests (1)	build (4)	tests (4)	build	tests (4)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (4)
Geant4	HEAD	build	tests (0)	build	tests (0)	build (8)	tests (0)	build (8)	tests (0)	build	tests (0)	build	tests (0)	build	tests (0)	build	tests (0)	build	tests (0)
Gauss	HEAD	build (1)	tests (66)	build (4)	tests (65)	build (6)		build (6)		build (47)	tests (54)	build (48)	tests (20)	build (125)	tests (21)	build (15)		build (15)	
Urania	HEAD	build	tests	build (1)	tests	build	tests	build (2)	tests	build (1)	tests	build	tests	build	tests	build	tests	build (1)	tests
Castelao	HEAD	build	tests	build	tests	build (37)	tests	build (37)	tests	build	tests	build	tests	build	tests	build (10)	tests	build (10)	tests
Noether	HEAD	build (2)	tests (5)	build (2)	tests (5)	build (4)	tests (5)	build (4)	tests (5)	build (2)	tests (5)	build (2)	tests (5)	build (2)	tests (5)	build (2)	tests (5)	build (2)	tests (5)
Kepler	HEAD	build	tests (2)	build	tests (2)	build (62)	tests (2)	build (62)	tests (2)	build	tests (2)	build	tests (2)	build	tests (2)	build (1)	tests (2)	build (1)	tests (2)
AlignmentOnline	HEAD	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)
Lovell	HEAD	build (6)		build (6)		build (6)		build (6)		build (6)		build (6)		build (6)		build (6)		build (6)	
Orwell	HEAD	build (15)		build (15)		build (15)		build (15)		build (15)		build (15)		build (15)		build (15)		build (15)	
LHCbIntegrationTests	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests (2)	build	tests (2)
MooreAnalysis	HEAD	build	tests	build	tests	build	tests	build	tests (5)	build	tests (3)	build	tests	build	tests	build	tests (7)	build	tests (7)

## List of slots

lhcb-2016-patches / 1447	lhcb-2017-patches / 1065	lhcb-2018-patches / 924
lhcb-coverity / 530	lhcb-dd4hep / 671	lhcb-digi14-patches / 692
lhcb-g4-dev / 629	lhcb-gaudi-head / 2764	lhcb-gauss-conf / 4
lhcb-gauss-dev / 2555	lhcb-gauss-fast / 113	lhcb-gauss-gen-dev / 539
lhcb-gauss-gen2-dev / 111	lhcb-gauss-lamarr / 168	lhcb-gaussino / 826
lhcb-head / 2750	lhcb-head-2 / 48	lhcb-hlt2011-patches / 622
lhcb-hlt2012-patches / 626	lhcb-hlt2016-patches / 444	lhcb-lcg-dev3 / 1431
lhcb-lcg-dev4 / 1445	lhcb-master / 1235	lhcb-master-mr / 1477
lhcb-master-ref / 972	lhcb-new-cmake / 207	lhcb-prerelease / 2572
lhcb-reco14-patches / 1089	lhcb-reco15-patches / 573	lhcb-run2-gaudi-head / 12
lhcb-run2-patches / 849	lhcb-run2-patches-dev4 / 492	lhcb-run2-prerelease / 64
lhcb-run3-cleanup / 51	lhcb-sanitizers / 705	lhcb-sim09 / 1424
lhcb-sim09-cmake / 1153	lhcb-sim09-upgrade / 1354	lhcb-sim10 / 584
lhcb-stripping21-firstpass-patches / 1029	lhcb-stripping21-patches / 1029	lhcb-stripping24-patches / 895
lhcb-stripping24r2-28r2-patches / 193		



- Builds can be triggered from merge request comments
- Many options for changing the build (platforms, other MRs, ...)
- Bot automatically comments with the test results

Robert Andrew Currie 🌱 @rcurrie · 4 days ago Developer

/ci-test --merge


1

▼ Collapse replies

Lhcb Software @lhcbsoft · 4 days ago Developer  
Started [reference](#) and [integration test](#) builds. Once done, check the [comparison](#) of build and test results.

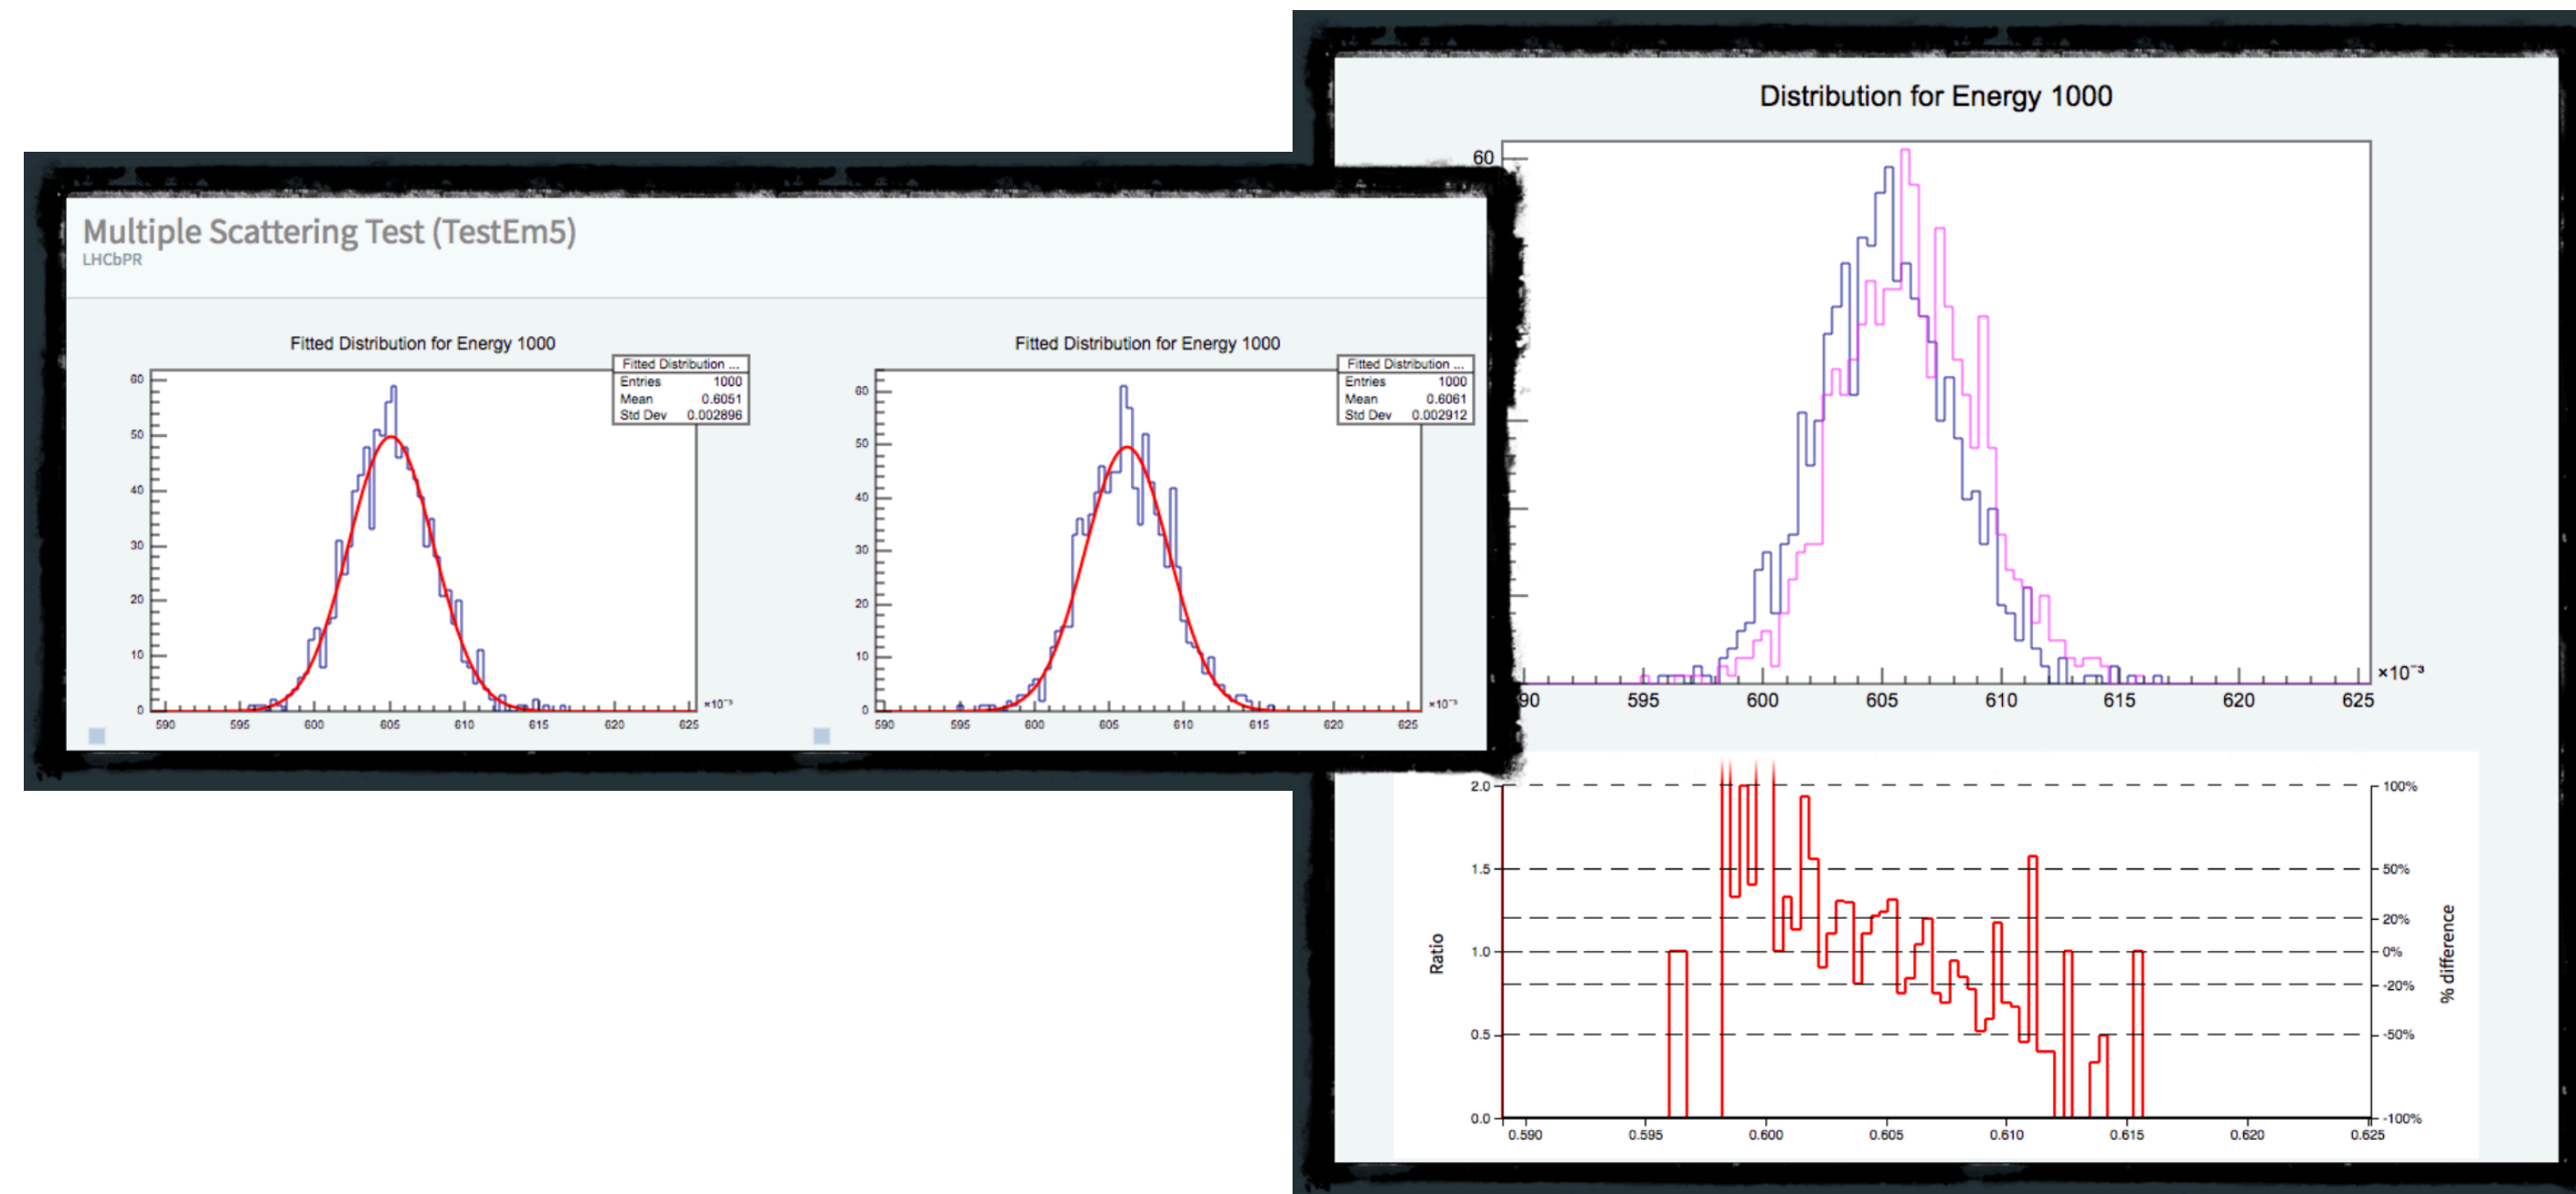
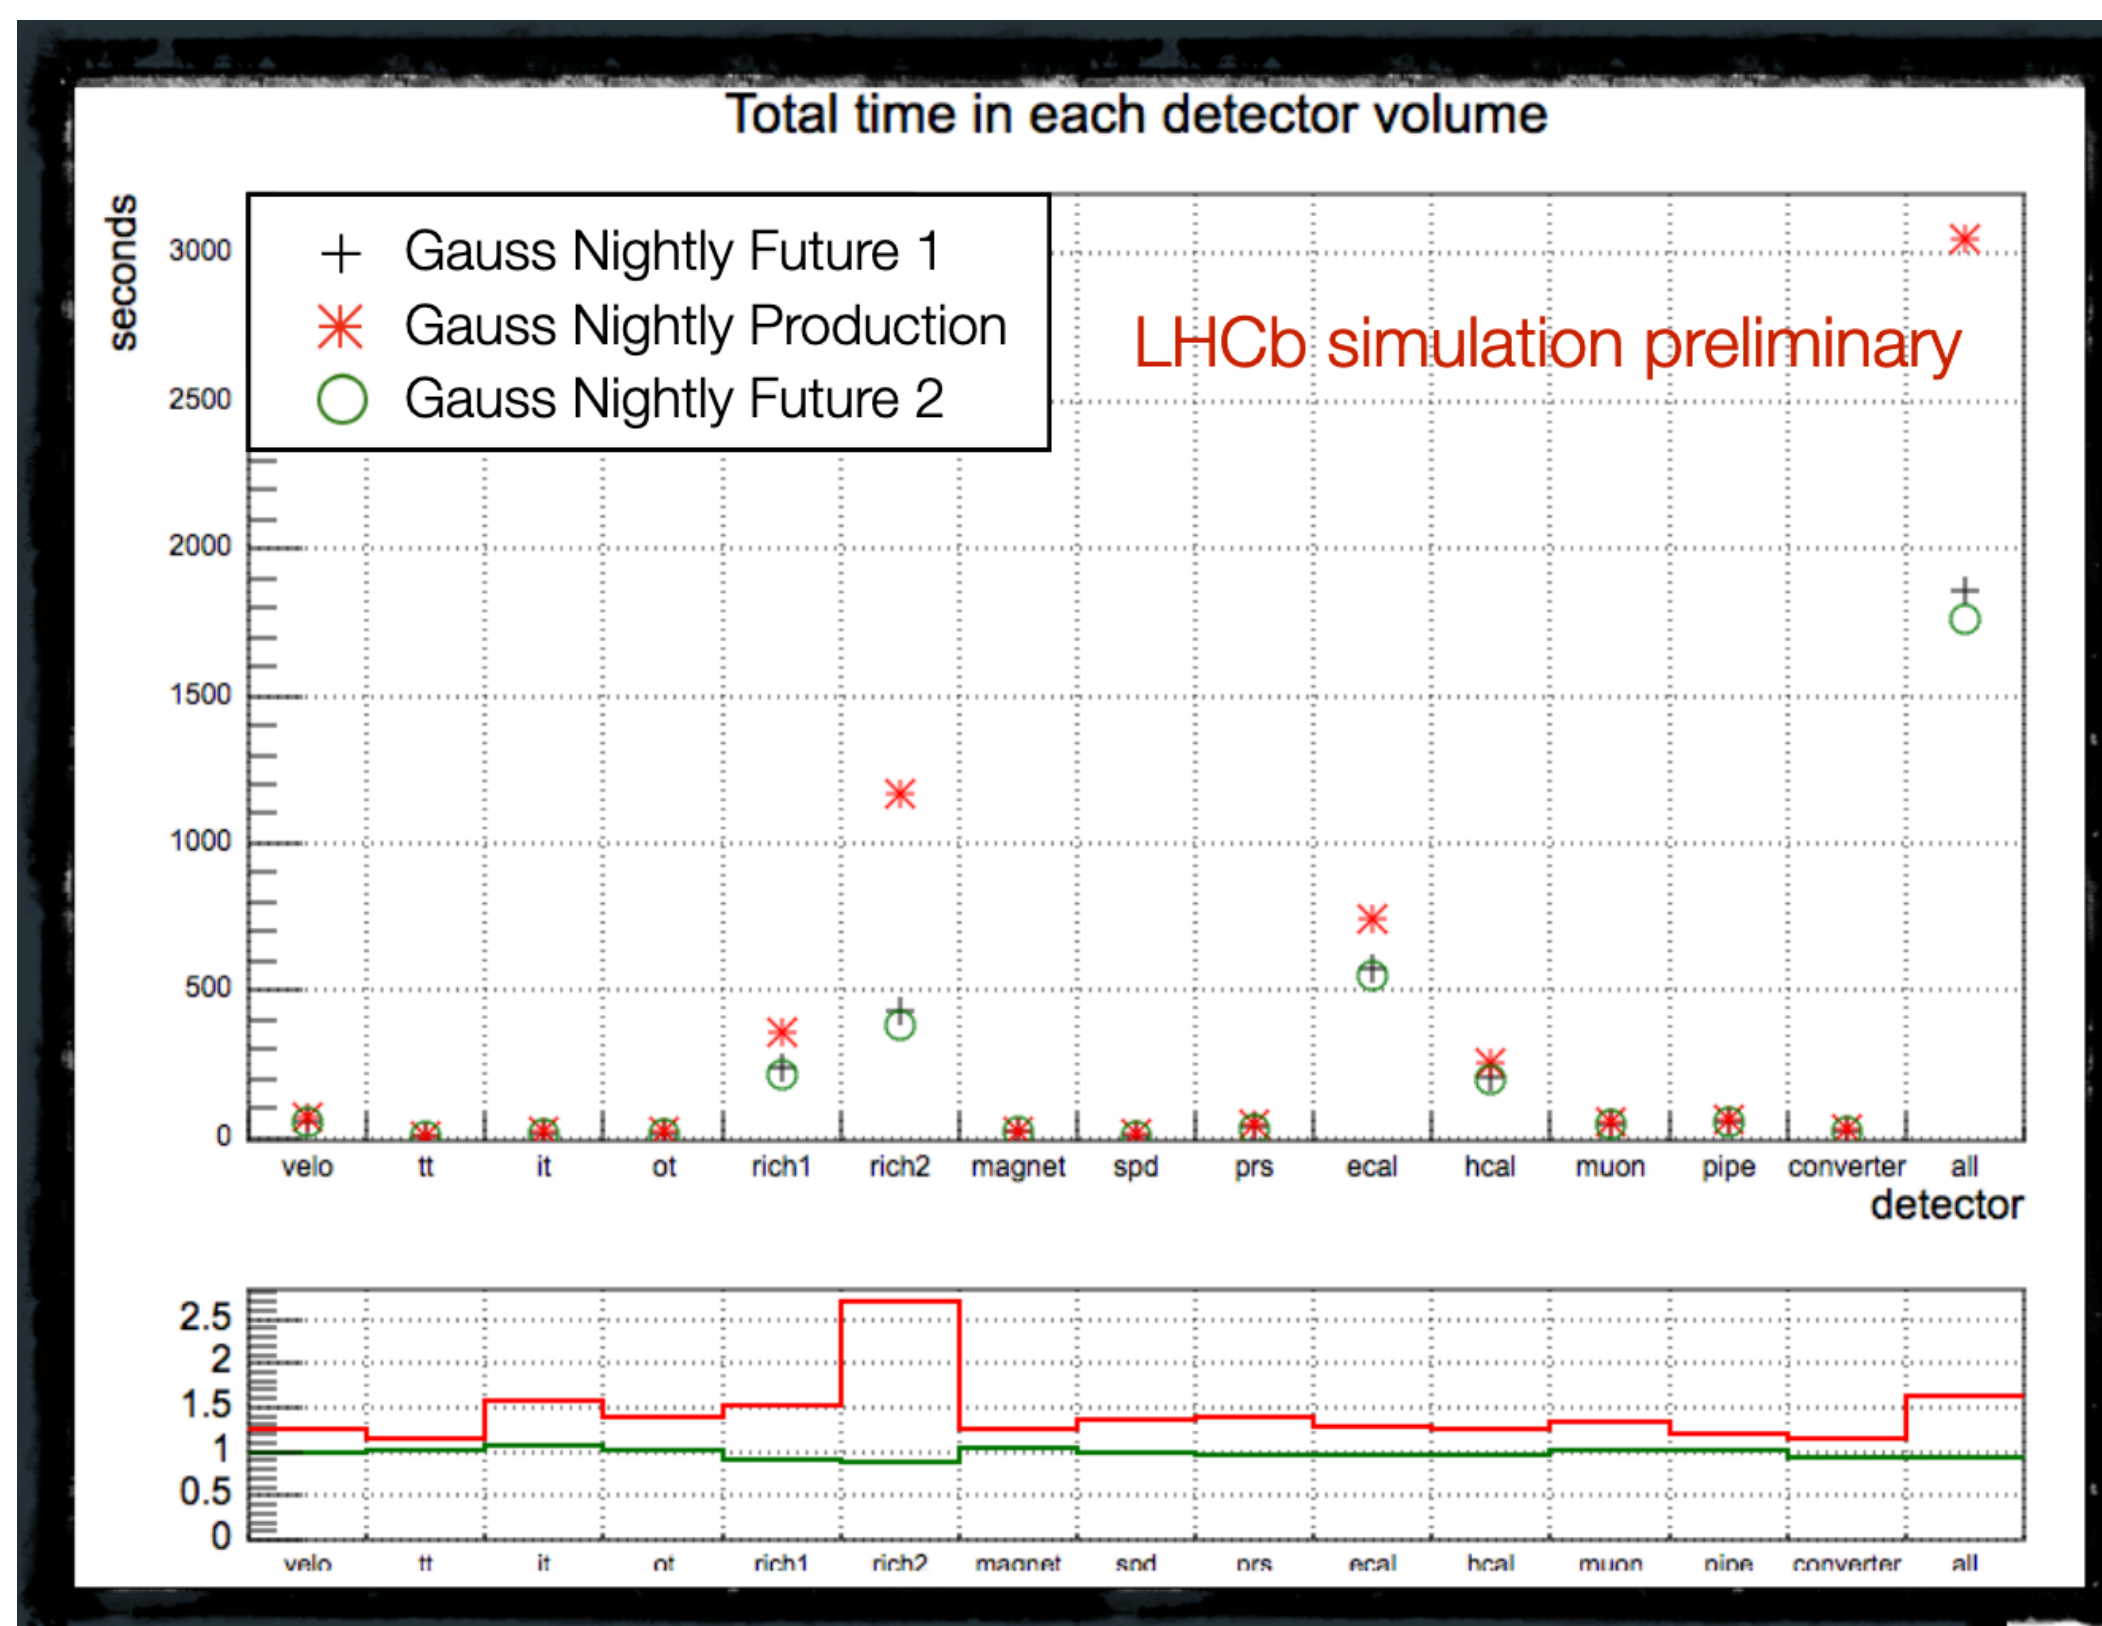
Lhcb Software @lhcbsoft · 4 days ago Developer  
Throughput Test [Moore\\_hlt1\\_pp\\_default](#): 29591.7 Events/s -- change of 0.08% vs. [reference](#)

Lhcb Software @lhcbsoft · 3 days ago Developer  
Throughput Test [Moore\\_hlt2\\_reco\\_baseline](#): 135.7 Events/s -- change of -0.38% vs. [reference](#)





- Larger tests can also be ran on demand
  - Bridge the gap between nightly tests ( $\mathcal{O}(5)$  events) and full validation campaigns ( $\mathcal{O}(1e6)$  events)

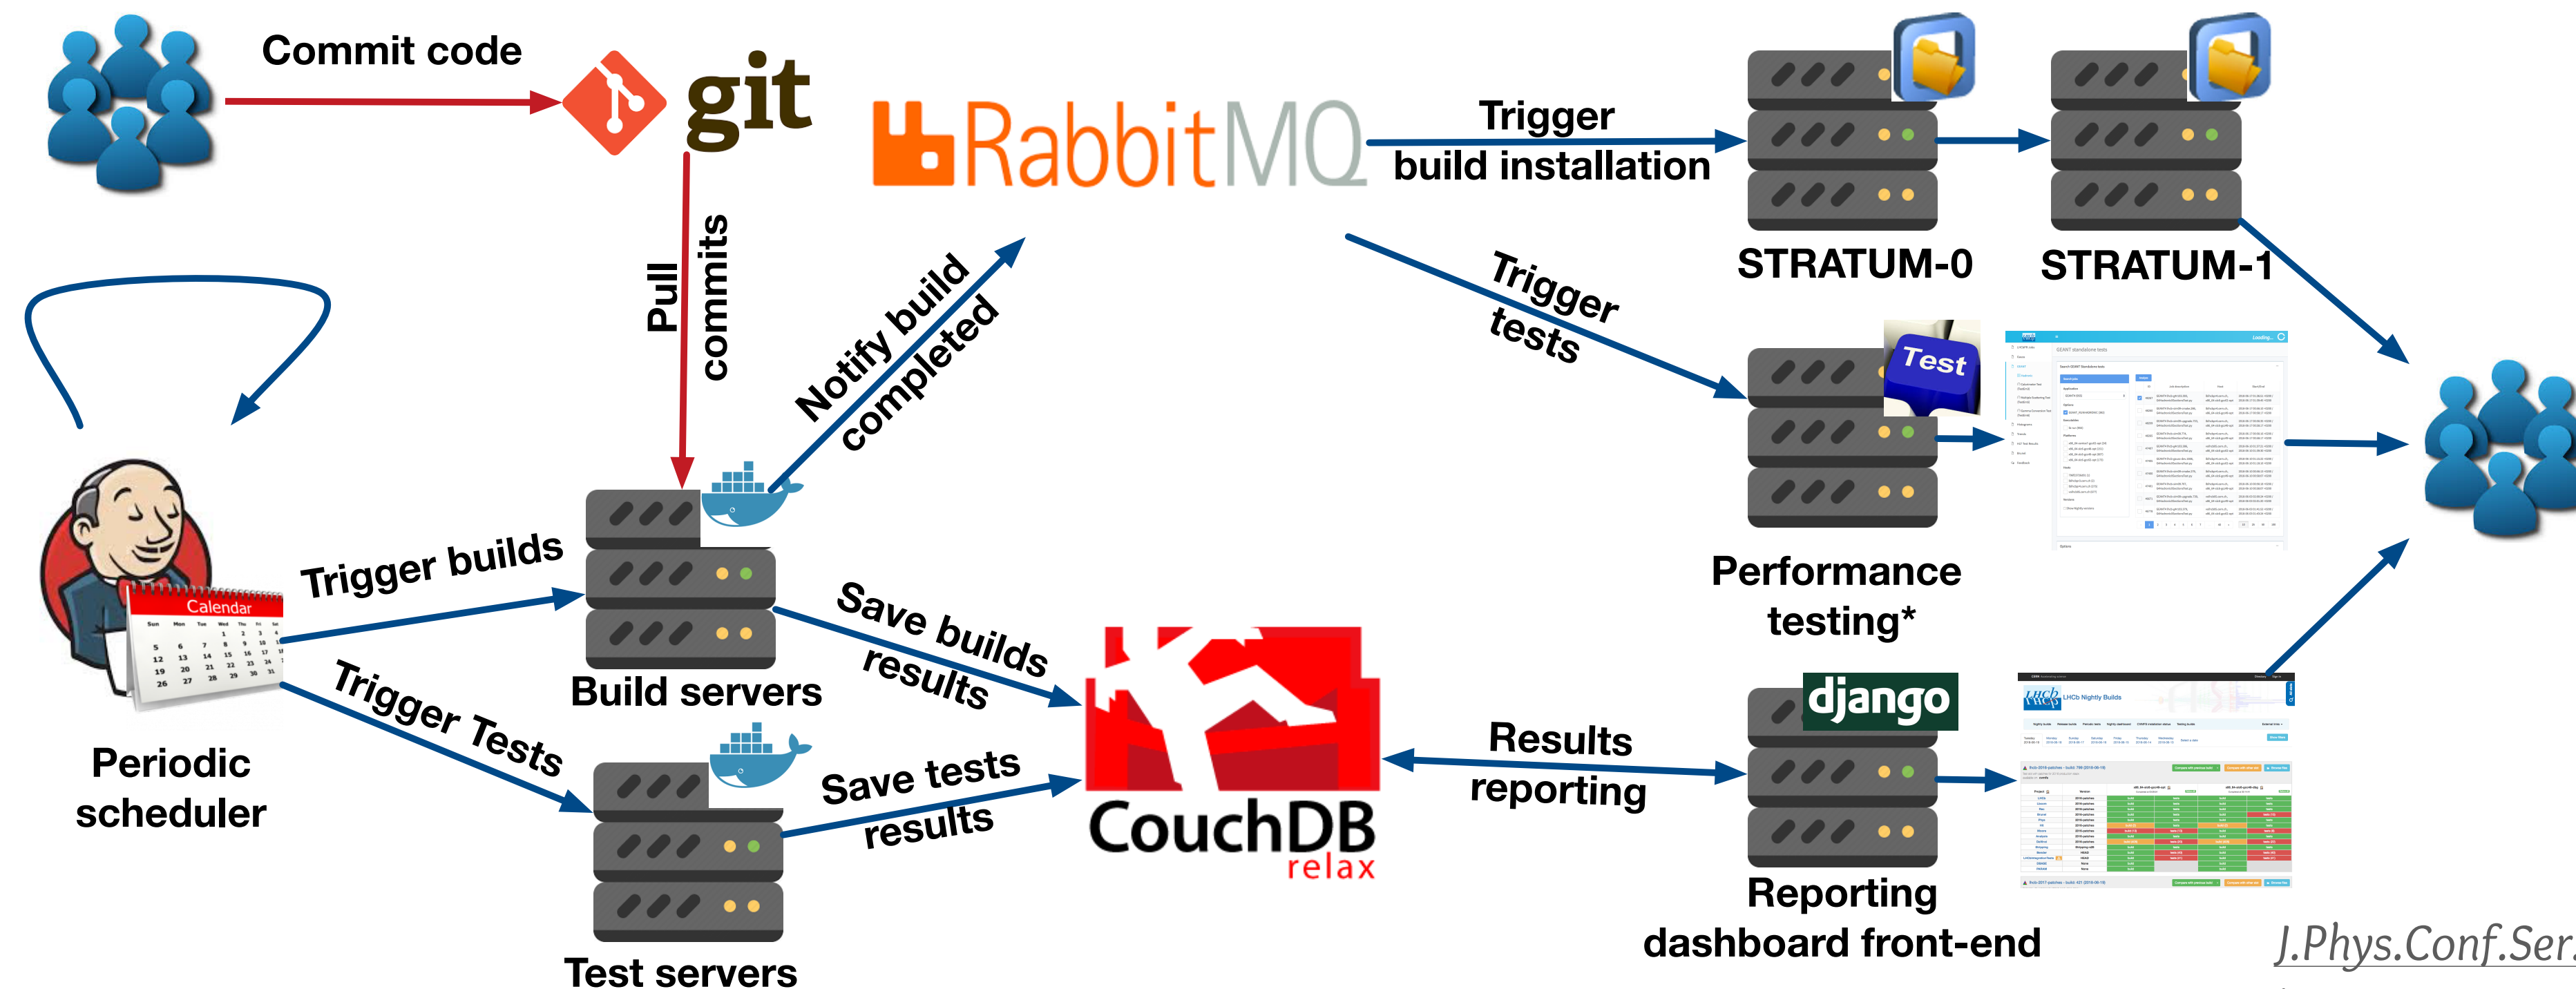


EPJ Web of Conferences 214, 02043 (2019)

EPJ Web of Conferences 214, 05014 (2019)



- Currently using Jenkins to manage the build machines
  - Manages over 700 CPU cores
  - Also manages release builds
- Investigating the use of a new custom solution based on Luigi + Celery



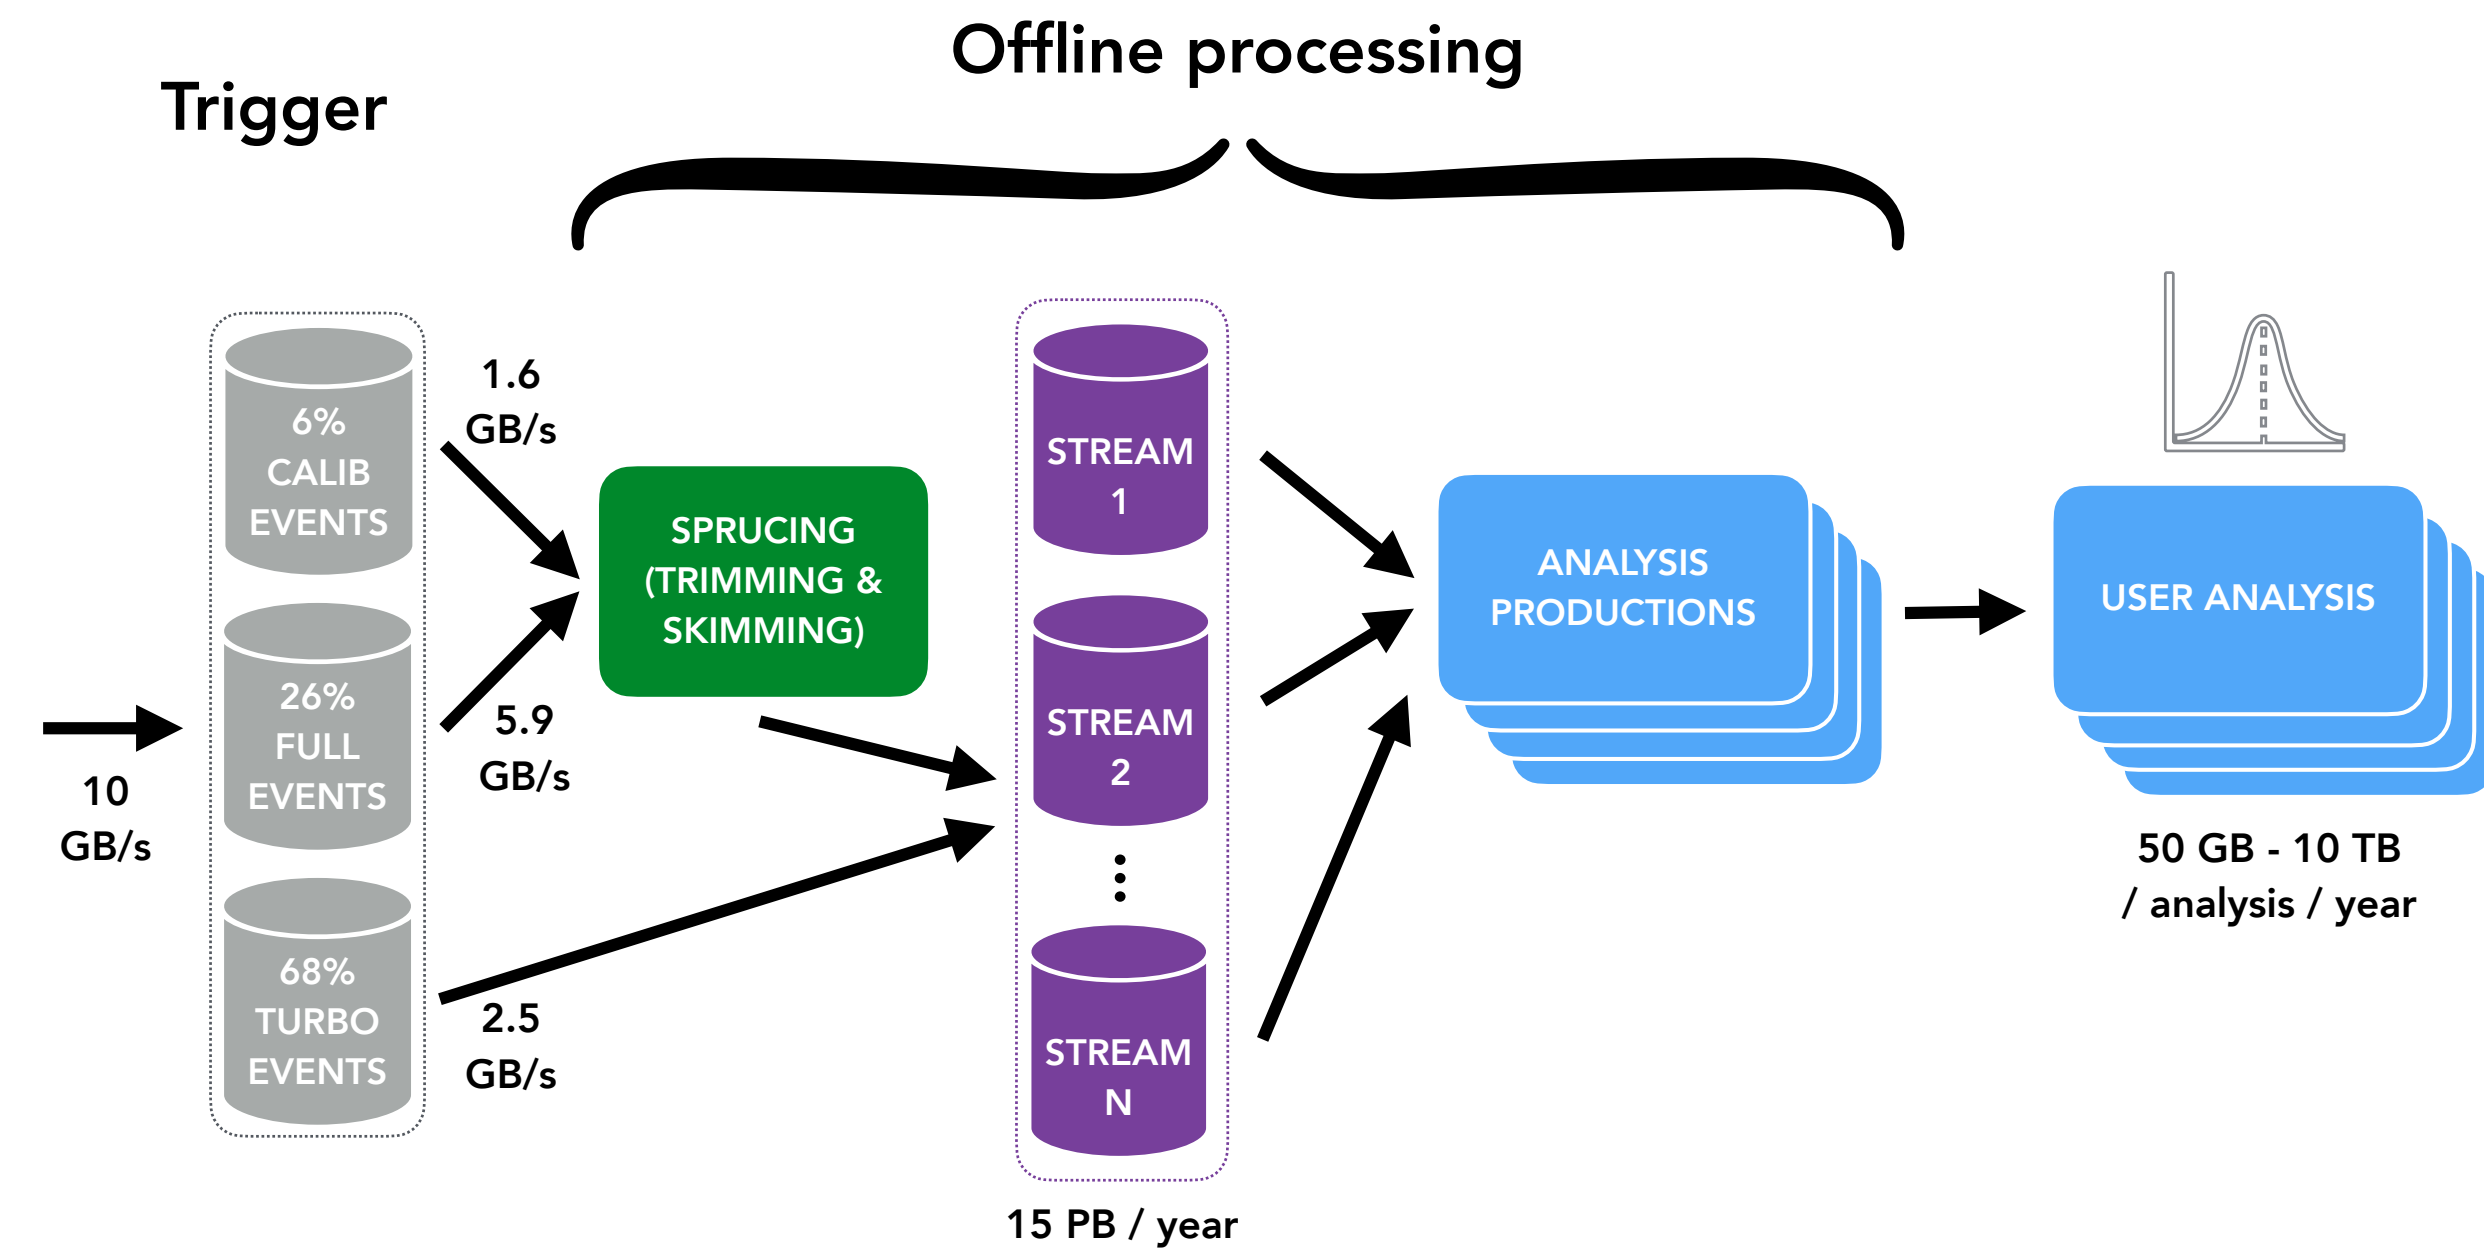
*J.Phys.Conf.Ser. 513 (2014) 052007*  
*EPJ Web of Conferences 214, 05001 (2019)*



The background is a complex, abstract pattern. It features a dense network of thin, dark green lines that swirl and intersect across a bright orange-yellow field. Interspersed among these lines are various blue and purple shapes, including small dots, larger blotches, and some star-like or cross-like symbols. The overall effect is one of organic, almost cellular or molecular complexity. A semi-transparent white rectangular box is centered horizontally, containing the text 'Custom GitLab runners' in a dark purple, serif font.

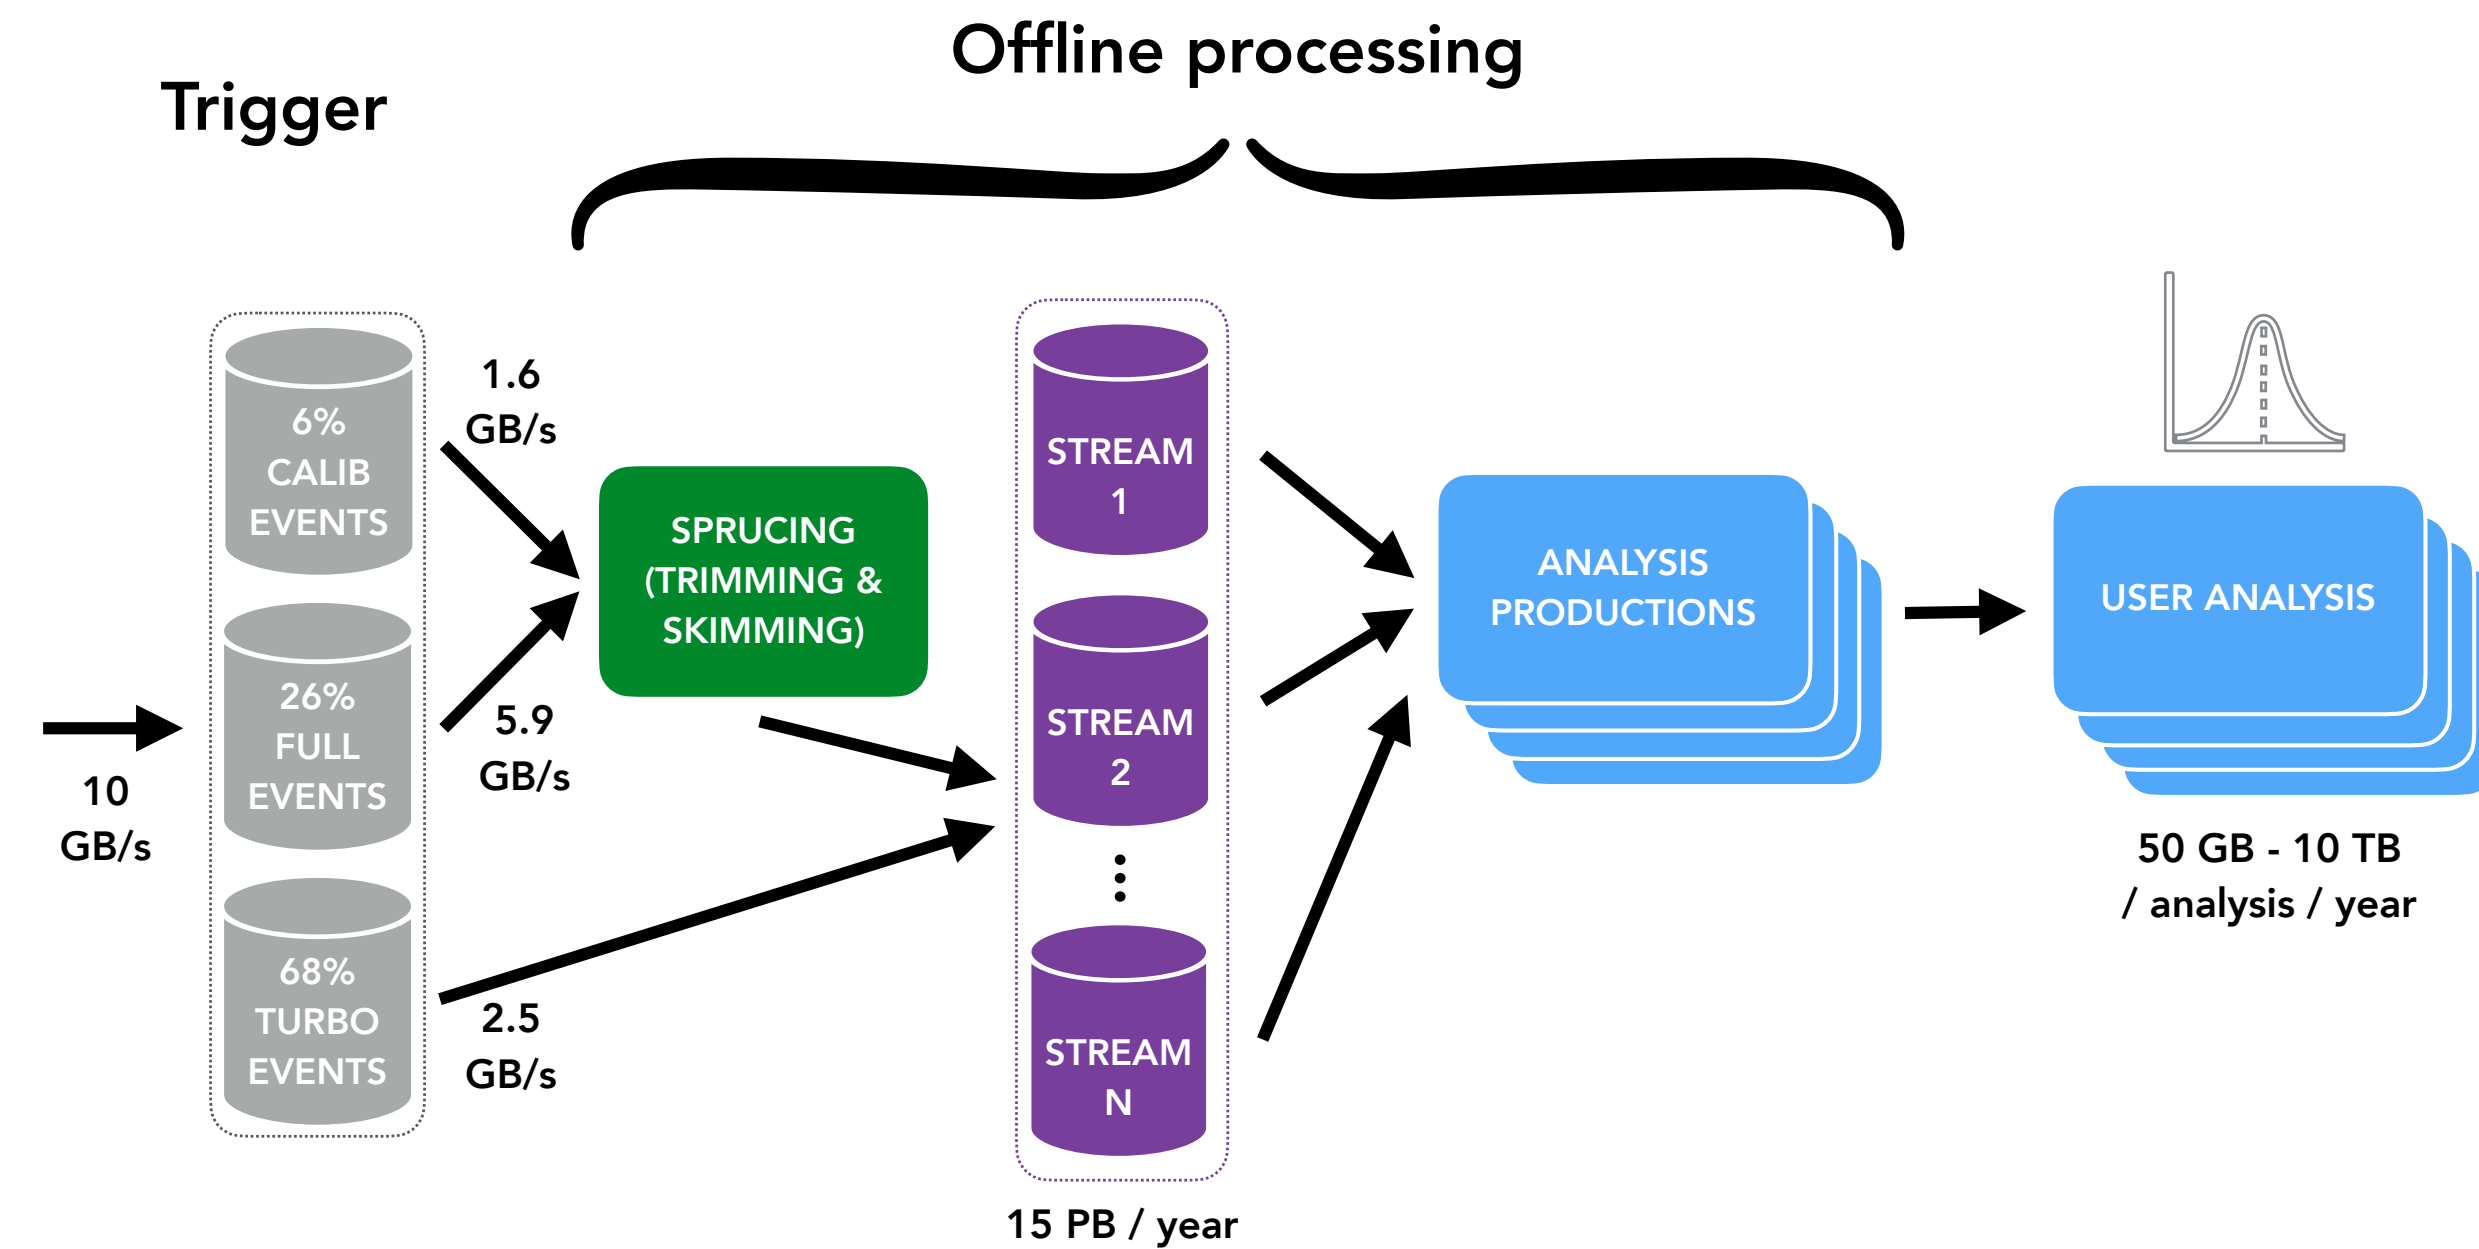
# Custom GitLab runners





- During LHC Run 1, most users filter datasets with “user jobs” on the grid
  - Often 10,000+ jobs required - extremely time consuming to manage
- In Run 3 (pictured above), analysts will need to filter 15PB per data taking year
  - Not feasible to continue with the current method
  - Instead make use of the DIRAC transformation system to declaratively submit processing campaigns





- Already used for all MC generation and centralised processing campaigns
  - Very powerful, but also potentially hazardous if broken transformations are submitted
- Solution: New transformation type with automatic testing
  - “Analysis Productions”



## ► Develop options files locally

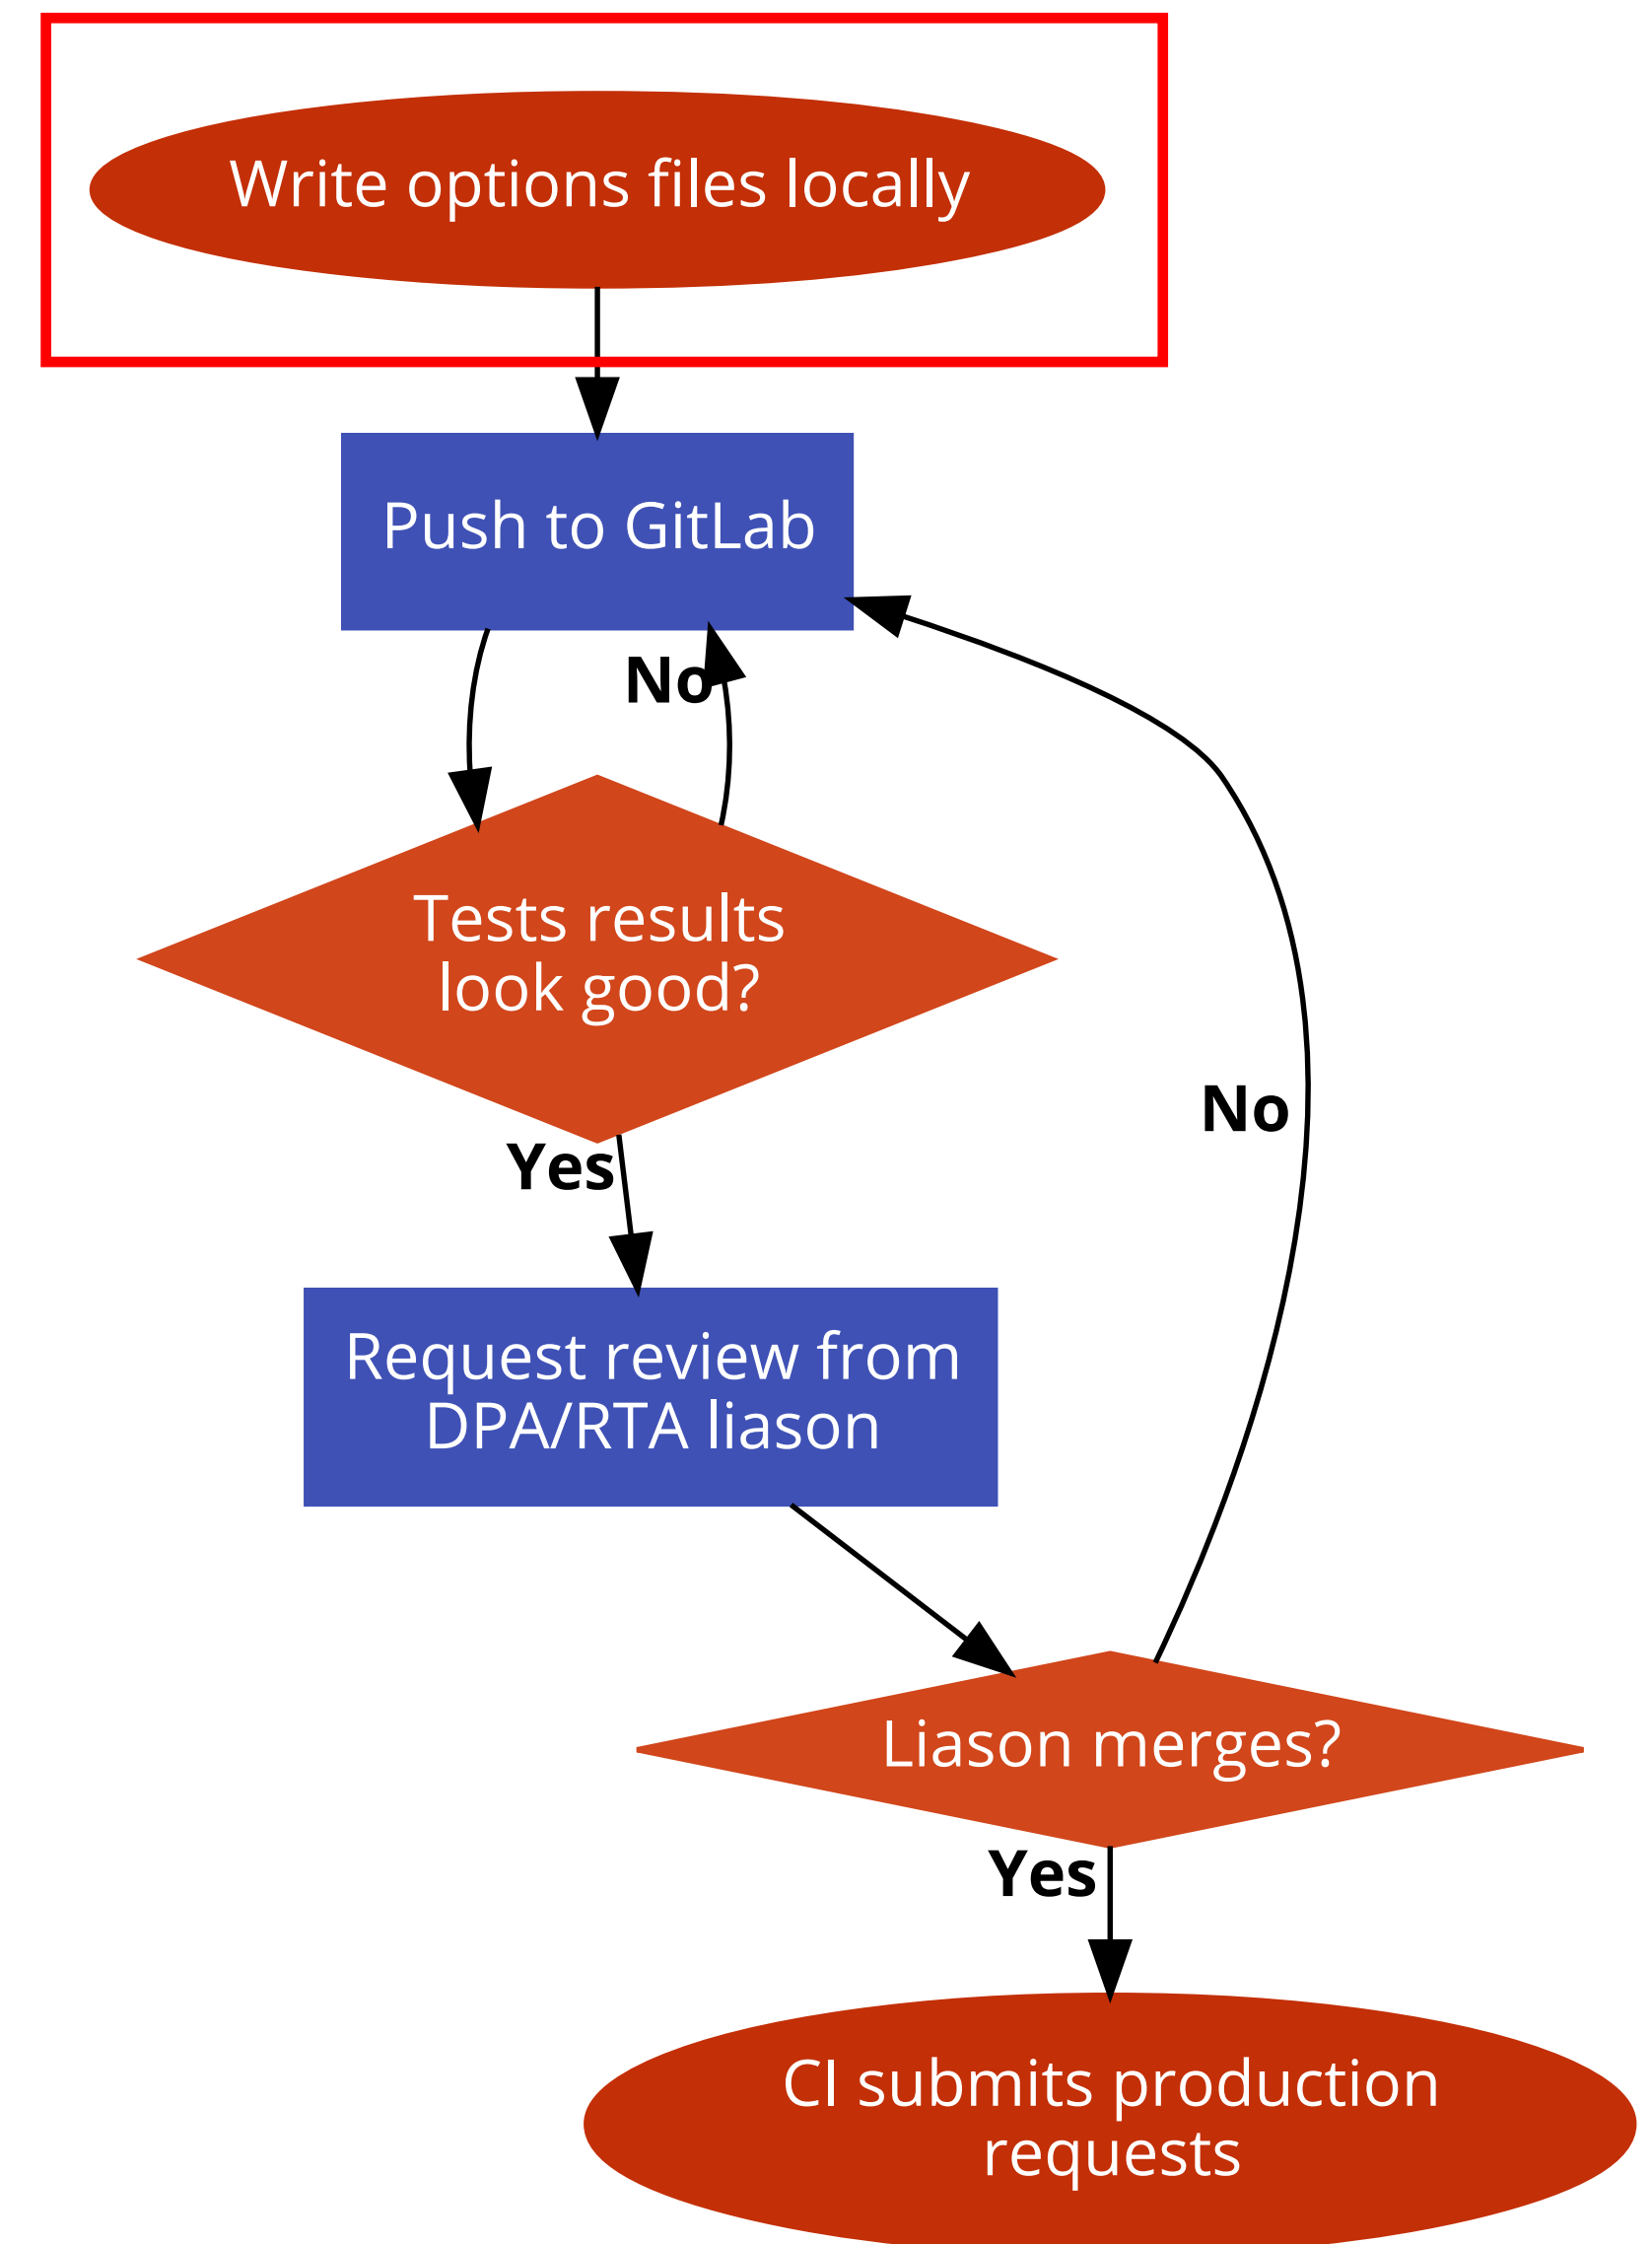
```
cburr@arch-desktop ~/D/D/AnalysisProductions (master) [2]> lb-ap test D02HH_Starterkit 2016_MagDown_PromptMC_D02KK
Validating environment
Running tests in /home/cburr/Development/DPA/AnalysisProductions/local-tests/D02HH_Starterkit-2020-09-29-23-22-03
Found existing dynamic dir pointing to /home/cburr/Development/DPA/AnalysisProductions/local-tests/BsToJpsiPhi-2020-09-19-12-26-09/dynamic, un
Setting CMAKE_PREFIX_PATH to /home/cburr/Development/DPA/AnalysisProductions/local-tests/D02HH_Starterkit-2020-09-29-23-22-03
Successfully logged in to https://lhcb-analysis-productions.web.cern.ch/
Generating pool XML catalog
Starting lb-run with: lb-run --unset=LD_LIBRARY_PATH --unset=PYTHONPATH --unset=XrdSecPROTOCOL --siteroot=/cvmfs/lhcb.cern.ch/lib/ --allow-con
dirun.py -T '$ANALYSIS_PRODUCTIONS_DYNAMIC/D02HH_Starterkit/2016_MagDown_PromptMC_D02KK_autoconf.py' '$ANALYSIS_PRODUCTIONS_BASE/D02HH_Starter
WARNING:lb-run:Decided best platform to use is x86_64-centos7-gcc9-opt
WARNING:lb-run:Decided best container to use is singularity
# setting LC_ALL to "C"
# Restarting with LD_PRELOAD='libtcmalloc.so'
# --> Including file '/home/cburr/Development/DPA/AnalysisProductions/local-tests/D02HH_Starterkit-2020-09-29-23-22-03/dynamic/D02HH_Starterki
# <-- End of file '/home/cburr/Development/DPA/AnalysisProductions/local-tests/D02HH_Starterkit-2020-09-29-23-22-03/dynamic/D02HH_Starterkit/2
# --> Including file '/home/cburr/Development/DPA/AnalysisProductions/D02HH_Starterkit/ntuple_options.py'

2020-09-29 21:23:19 UTC TimingAuditor.T... INFO * L0DUPromRaw | 1.111 | 0.698 | 0.234 | 2
2020-09-29 21:23:19 UTC TimingAuditor.T... INFO * Hlt1DecReportsDecoder | 0.476 | 11.963 | 0.027 | 749
2020-09-29 21:23:19 UTC TimingAuditor.T... INFO * Hlt2DecReportsDecoder | 6.190 | 6.354 | 0.185 | 367
2020-09-29 21:23:19 UTC TimingAuditor.T... INFO -----
2020-09-29 21:23:19 UTC NTupleSvc INFO NTuples saved successfully
2020-09-29 21:23:19 UTC ApplicationMgr INFO Application Manager Finalized successfully
2020-09-29 21:23:19 UTC ApplicationMgr INFO Application Manager Terminated successfully

Summary of log messages:
Found 2 ERROR messages
* 1 instances of "SubstitutePIDTool:: NO substitutions pi- : (Xc -> (D0 -> ^{(K-) K+) Meson) StatusCode=FAILURE"
* 1 instances of "SubstitutePIDTool:: NO substitutions pi+ : (Xc -> (D~0 -> ^{(K+) K-) Meson) StatusCode=FAILURE"
Found 1 WARNING messages
and 1 others (1 unique), pass "--suppress=0" to show all messages

General explanations
* Line: 351
Histograms are not being saved as no filename has been specified for storing them. This
message is harmless and normally ignored.

Success! Output can be found in /home/cburr/Development/DPA/AnalysisProductions/local-tests/D02HH_Starterkit-2020-09-29-23-22-03/output
```

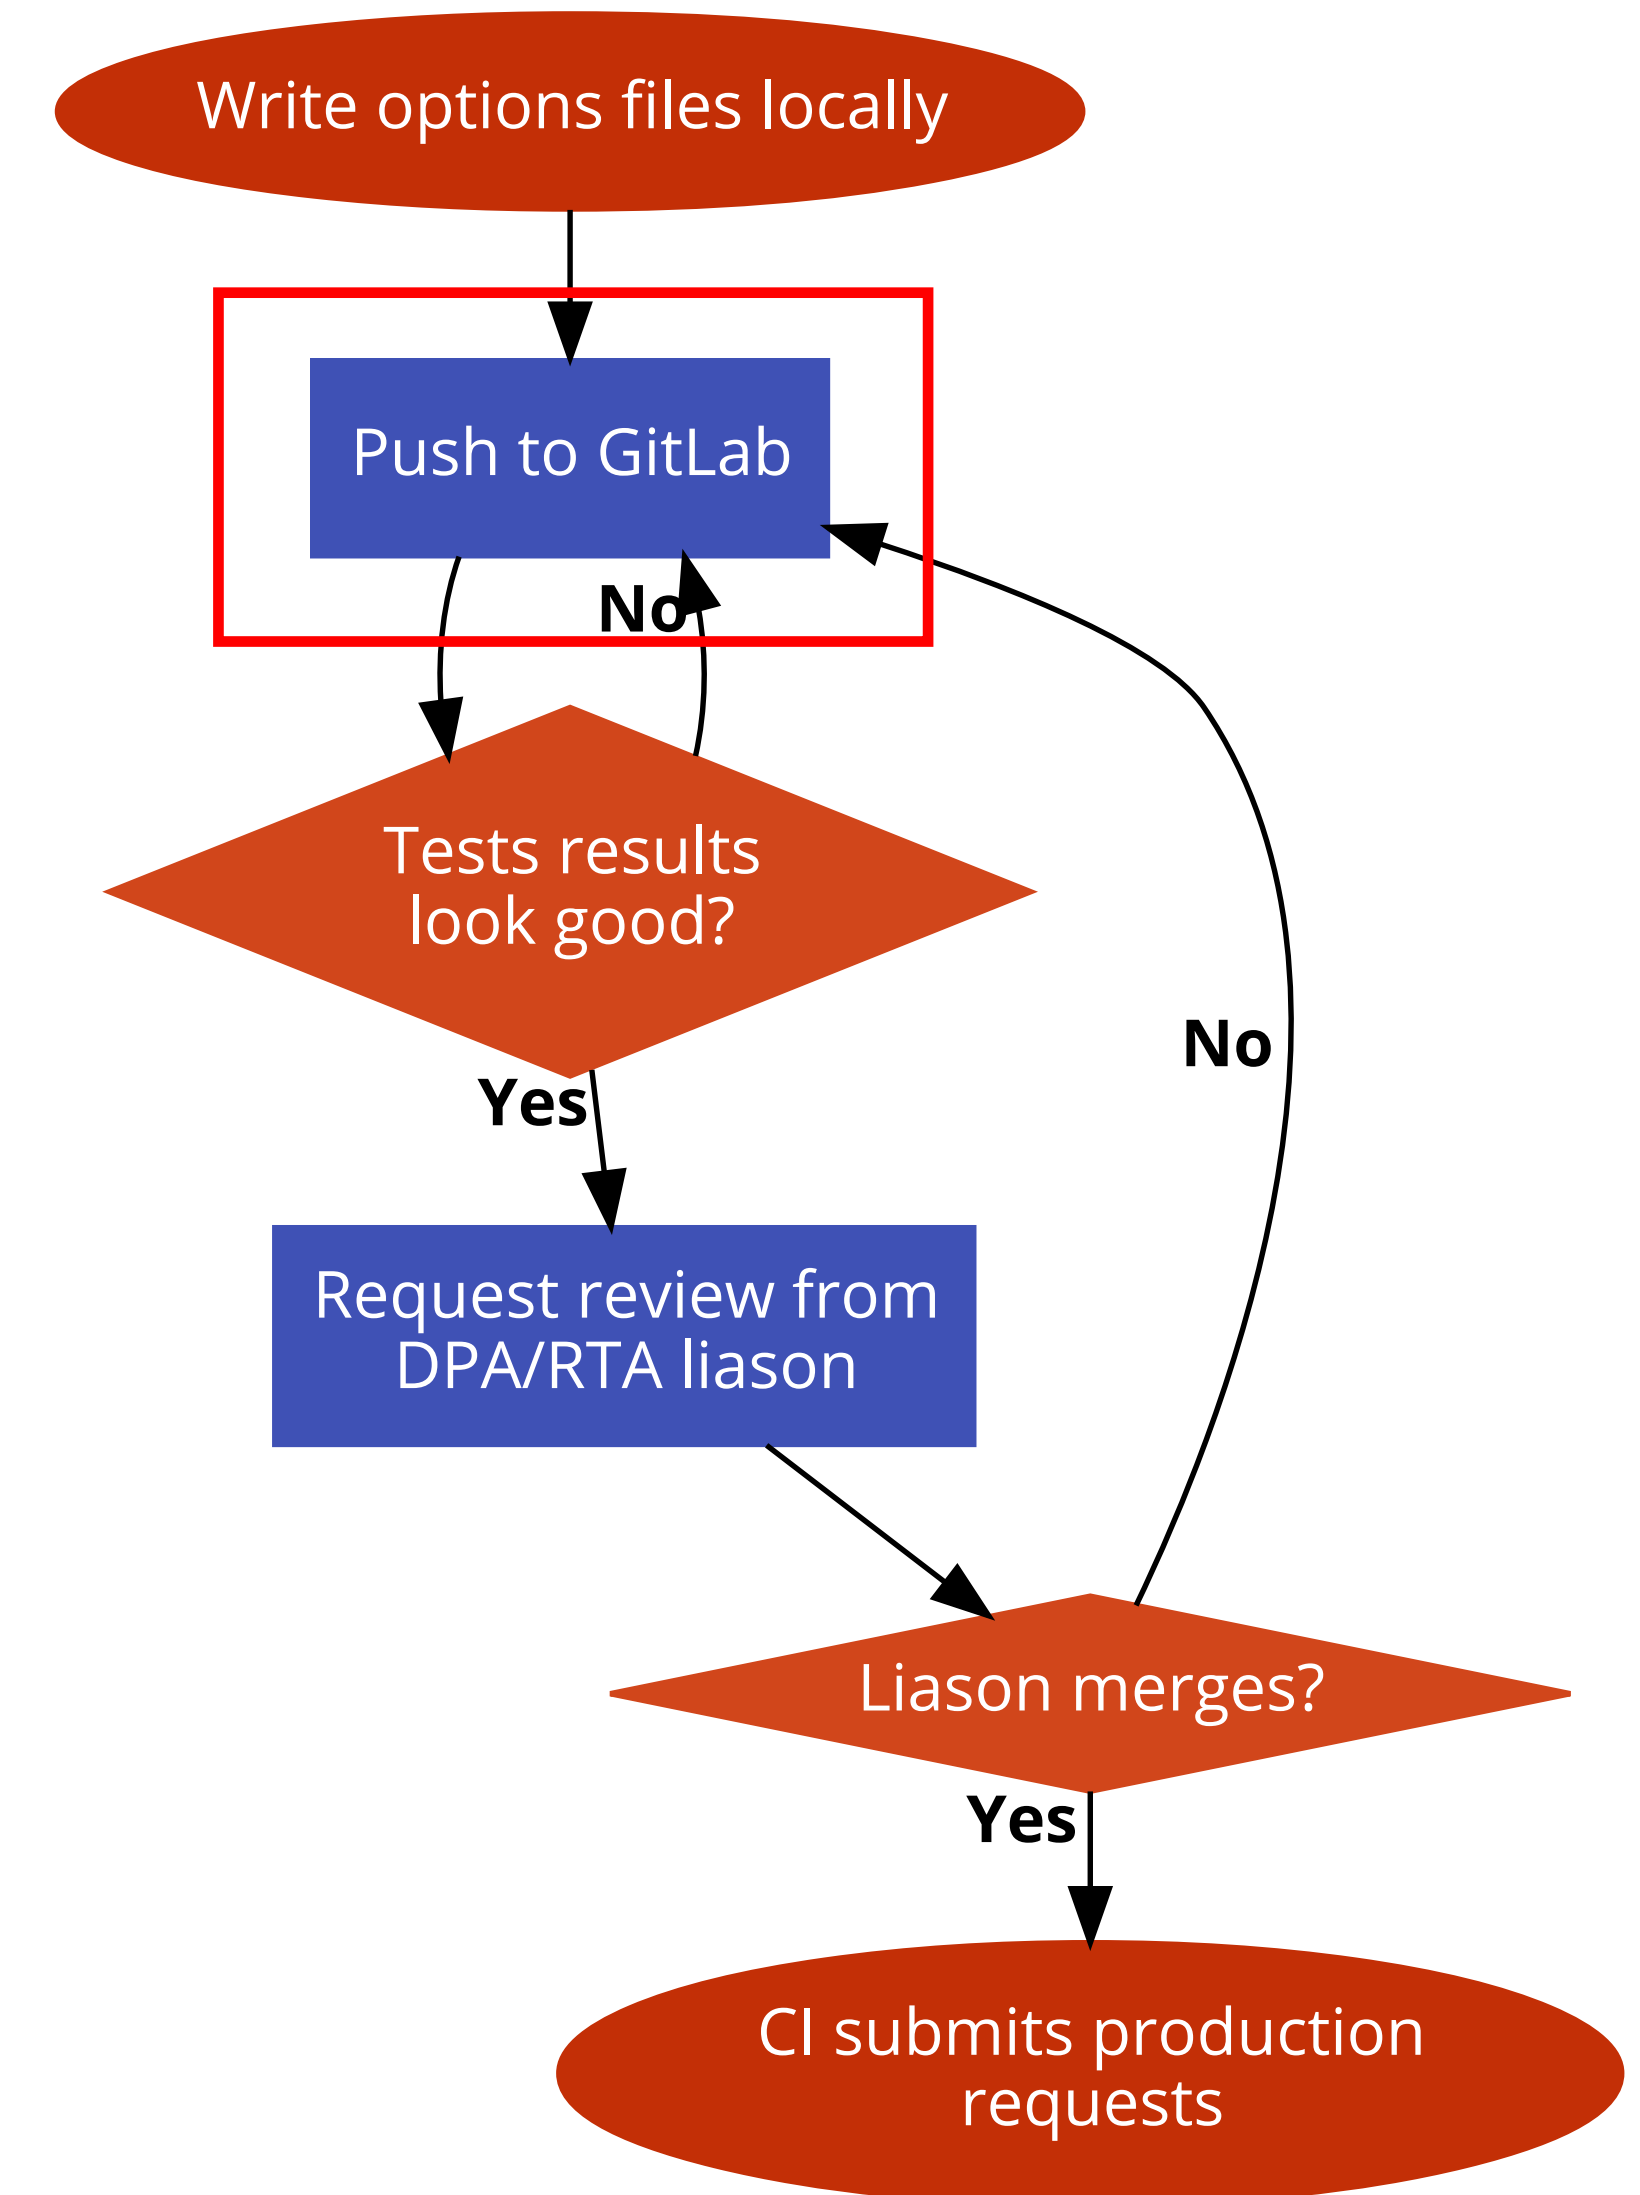


For details see: <https://gitlab.cern.ch/lhcb-dpa/analysis-productions/lbaplocal#lbaplocal>



## ➤ Push to GitLab

The screenshot displays the GitLab web interface for the 'AnalysisProductions' project. The left sidebar shows the 'CI / CD' section with 'Jobs' selected. The main content area shows a pipeline run for 'test' with a status of 'passed'. The job log on the left lists 24 steps, including creating a pipeline, generating configuration options, and submitting jobs. The right sidebar shows the job details for 'test', including its duration (89 minutes 58 seconds), timeout (1d), and runner (#5453). A 'Retry' button is visible next to the job name.





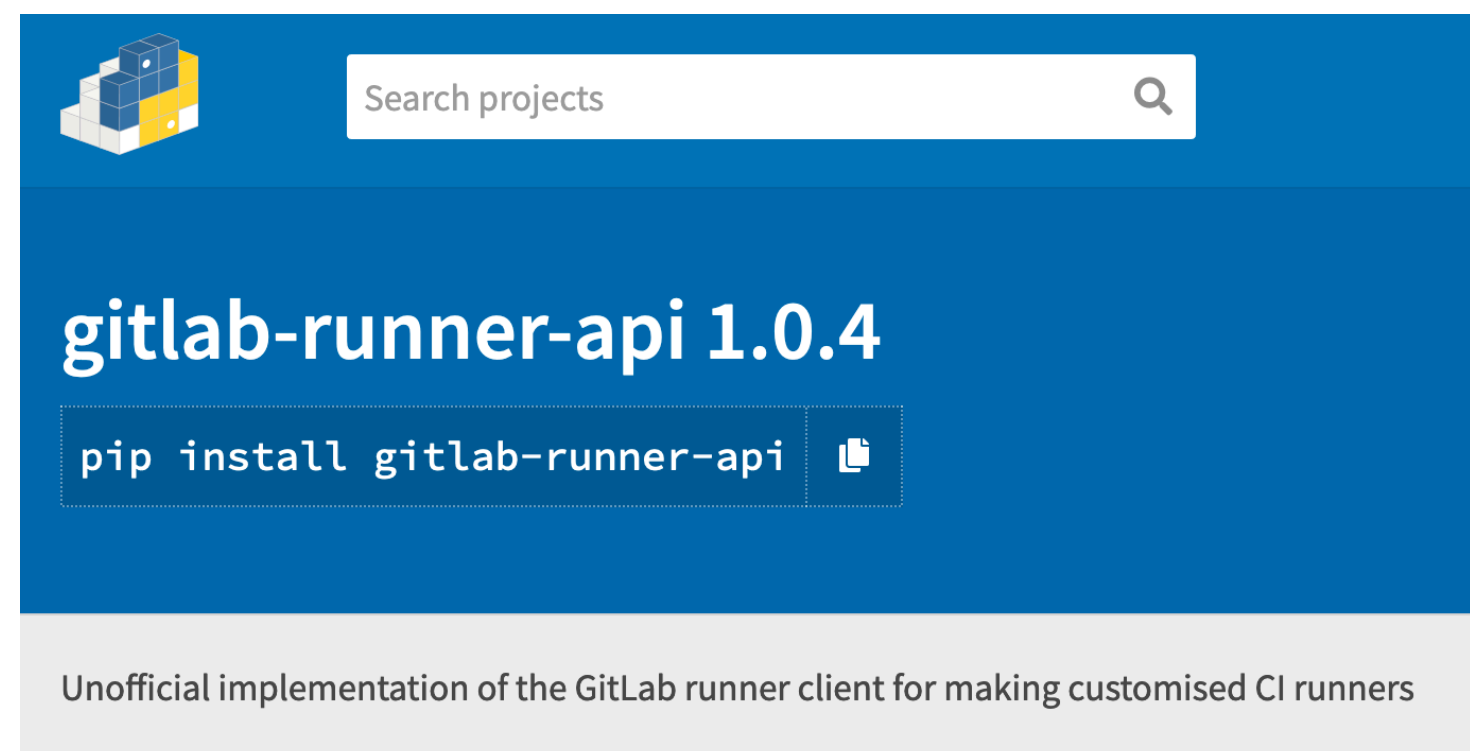


## Why not use a conventional runner?

- Batches of productions vary in scale (a lot!)
  - Sometimes: 2 productions that each take a few minutes
  - Sometimes: 100+ productions each taking multiple hours
  - Common to have many analysts working at the same time
- Very inefficient to have hundreds of cores assigned to CI runners
  - These would be unused for 99% of the time
- Also limited in the human-resources that can be used for development
  - No time to develop complex systems for displaying CI logs and status



- GitLab uses a simple REST interface to communicate with runners
  - Pull-based so no need for inbound connectivity on runners
- Reverse engineered this API
  - Implemented in a simple Python package - available from PyPI and conda-forge
- No support for executing jobs directly
  - Query job metadata and reporting results back to GitLab



<https://pypi.org/project/gitlab-runner-api/>

```
from gitlab_runner_api import Runner
runner = Runner.load("my-runner-data.json")
runner.check_auth()
if job := runner.request_job():
    print("Received a new job, starting executor")
    my_job_executor(job)
else:
    print("No jobs are currently available")
```

*Requesting a job*

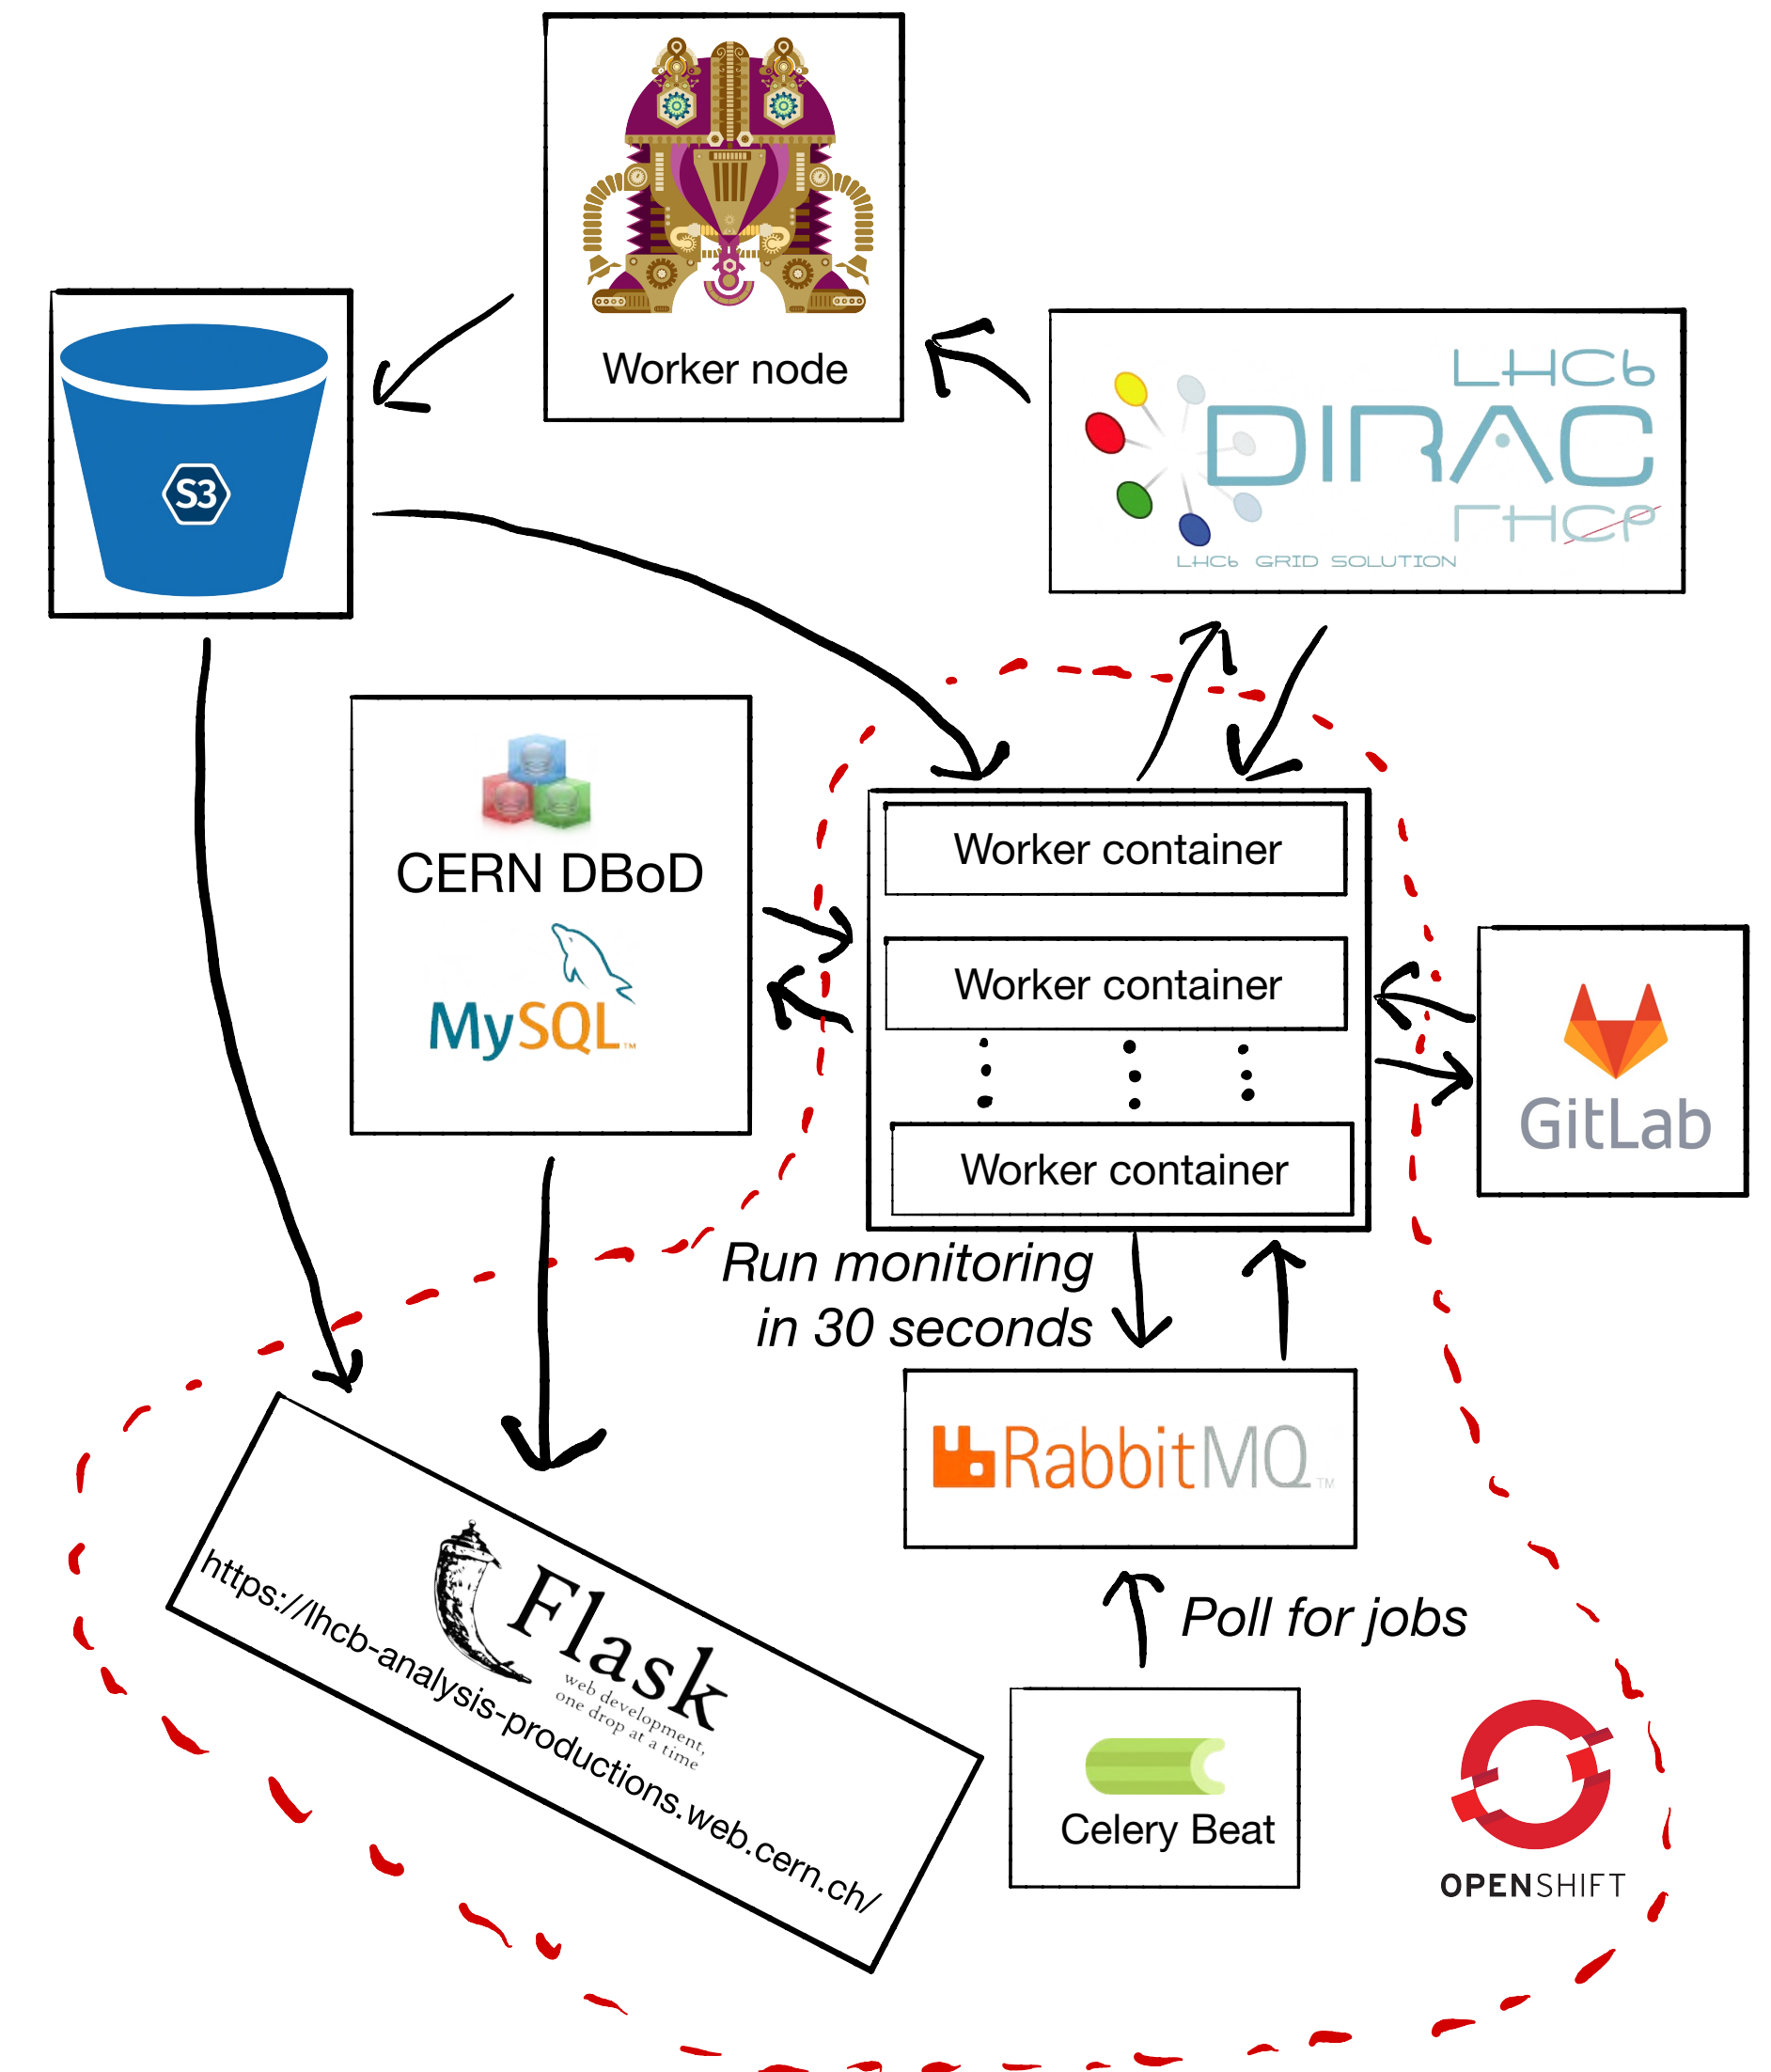
```
from gitlab_runner_api import failure_reasons

job.log += f"Starting job with id {job.id} for branch {job.ref}\n"
do_clone_and_checkout(job.repo_url, job.commit_sha)
success = run_tests(job)
if success:
    job.log += "All tests ran successfully\n"
    job.set_success()
else:
    # ANSI formatting codes can be used to enhance the CI logs
    job.log += "\u001b[31mJob failed!!!\u001b[0m\n"
    job.set_failed(failure_reasons.ScriptFailure())
```

*Executing a job with logs/status sent to GitLab*



- For “Analysis Productions” use a custom runner to submit tests to the grid
  - Effectively unlimited parallelisation available
- Credentials for data access don't need to be stored as a secret
  - GitLab provides trusted identity so Jobs can be executed as the user who pushed
- Use Celery and RabbitMQ for managing long-running jobs
- A summary is sent to the GitLab CI log
  - Simple flask front-end for exposing detailed results





# Reporting results to analysts

GitLab

Projects ▾

Groups ▾

More ▾

+

Search or jump to...

Q

8

3

72

?

A

AnalysisProductions

Project overview

Repository

Issues0

Merge Requests0

CI / CD

Pipelines

**Jobs**

Schedules

Operations

Packages & Registries

Analytics

Members

Settings

LHCb Data Packages > AnalysisProductions > Jobs > #10108707

passed

Job #10108707 triggered 6 days ago by Chris Burr

1 Running with gitlab\_runner\_api 0.1.dev42+g457dbd5

2 INFO:Creating new pipeline for ID 1958460

3 ALWAYS:Results will be available at https://lhcb-analysis-productions.web.cern.ch/1958460/

4 INFO:Creating production D02HH\_Starterkit

5 INFO:Generating configuration options for D02HH\_Starterkit 2016\_MagDown\_PromptMC\_D02KK (1.1/2)

6 INFO:Generating configuration options for D02HH\_Starterkit 2016\_MagUp\_PromptMC\_D02KK (2.1/2)

7 ALWAYS:Submitting jobs for D02HH\_Starterkit

8 INFO:Submitting test for D02HH\_Starterkit 2016\_MagUp\_PromptMC\_D02KK (2/2)

9 INFO:Submitting test for D02HH\_Starterkit 2016\_MagDown\_PromptMC\_D02KK (1/2)

10 INFO:Submitted DIRAC job for D02HH\_Starterkit 2016\_MagUp\_PromptMC\_D02KK with ID 405432666

11 INFO:Submitted DIRAC job for D02HH\_Starterkit 2016\_MagDown\_PromptMC\_D02KK with ID 405432667

12 ALWAYS:2 jobs still running at 2020-09-23T12:45:07.464027

13 ALWAYS:2 jobs still running at 2020-09-23T12:45:55.646284

14 ALWAYS:2 jobs still running at 2020-09-23T12:46:36.042122

15 ALWAYS:2 jobs still running at 2020-09-23T12:47:16.577207

16 ALWAYS:2 jobs still running at 2020-09-23T12:48:01.112186

17 ALWAYS:2 jobs still running at 2020-09-23T12:48:40.594081

18 ALWAYS:2 jobs still running at 2020-09-23T12:49:23.145484

19 ALWAYS:2 jobs still running at 2020-09-23T12:50:09.824533

20 ALWAYS:2 jobs still running at 2020-09-23T12:50:54.779824

21 ALWAYS:2 jobs still running at 2020-09-23T12:51:52.586056

22 ALWAYS:2 jobs still running at 2020-09-23T12:52:33.424767

23 ALWAYS:2 jobs still running at 2020-09-23T12:53:15.346984

24 ALWAYS:2 jobs still running at 2020-09-23T12:54:02.530262

test

Retry

Duration: 89 minutes 58 seconds

Timeout: 1d (from project)

Runner: (#5453)

Tags: lbap\_runner

Commit 53210696 in !16

Minimise repetition

Pipeline #1958460 for djwhite/starterkit-practice

test

→ test

## Analysis production for pipeline #1958460


### D02HH\_Starterkit

Commit: 53210696  
Message: Minimise repetition  
Colour code: Running Success Failed Unknown

Raw YAML Show

Rendered YAML Show

Click on a job for more information



Job Name	Test job statistics			Log messages			Estimated total	
	Events	Runtime	Size	Warn	Error	Fatal	Size	Kept
2016_MagDown_PromptMC_D02KK	2548	0:00:34	0.2 MB	1	2	0	0.1 GB	True
2016_MagUp_PromptMC_D02KK	2916	0:00:34	0.2 MB	1	2	0	0.1 GB	True

For support see the [DPA WP2 Analysis Productions](#) channel on mattermost.

LbAnalysisProductions 0.0.2.dev2+g2875ced · LbAPCommon 0.0.4



Analysis production job for [D02HH\\_Starterkit](#)

2016\_MagDown\_PromptMC\_D02KK

Status	Commit	Requested	Processed	Runtime	Kept
Success	<a href="#">53210696</a>	-1 events	2548 events	0:00:34	True

Input

Name	Size	Total
/lhcb/MC/2016/ALLSTREAMS.DST/00070793/0000/00070793_00000206_7.AllStreams.dst	544.8 MB	0.4 TB

Output

Name	Size	Total (estimated)
<a href="#">00012345_00006789_1.d02kk.root</a>	0.2 MB	0.1 GB

Browse output file [Show](#)

Reproduce on Ixplus:

```
1 | lb-ap reproduce 1958460 "D02HH_Starterkit" "2016_MagDown_PromptMC_D02KK" 0
```

Explain job log [Show](#)

Job log [Show](#)

DIRAC log [Show](#)

If the logs show "Error 404 while fetching file: NOT FOUND" it is sign the test job has not started ;  
For support see the [DPA WP2 Analysis Productions](#) channel on mattermost.  
LbAnalysisProductions 0.0.2.dev2+g2875ced · LbAPCommon 0.0.4

Browse output file [Hide](#)

Read a ROOT file

JSROOT version 5.8.1 7/05/2020

[open all](#) | [close all](#) | [clear](#)

00012345\_00006789\_1.d02kk.root

TupleDstToD0pi\_D0ToKK

DecayTree;1

Dstar\_ENDVERTEX\_X

Dstar\_ENDVERTEX\_Y

Dstar\_ENDVERTEX\_Z

Dstar\_ENDVERTEX\_XERR

Dstar\_ENDVERTEX\_YERR

Dstar\_ENDVERTEX\_ZERR

Dstar\_ENDVERTEX\_CHI2

Dstar\_ENDVERTEX\_NDOF

Dstar\_ENDVERTEX\_COV\_

Dstar\_OWNPV\_X

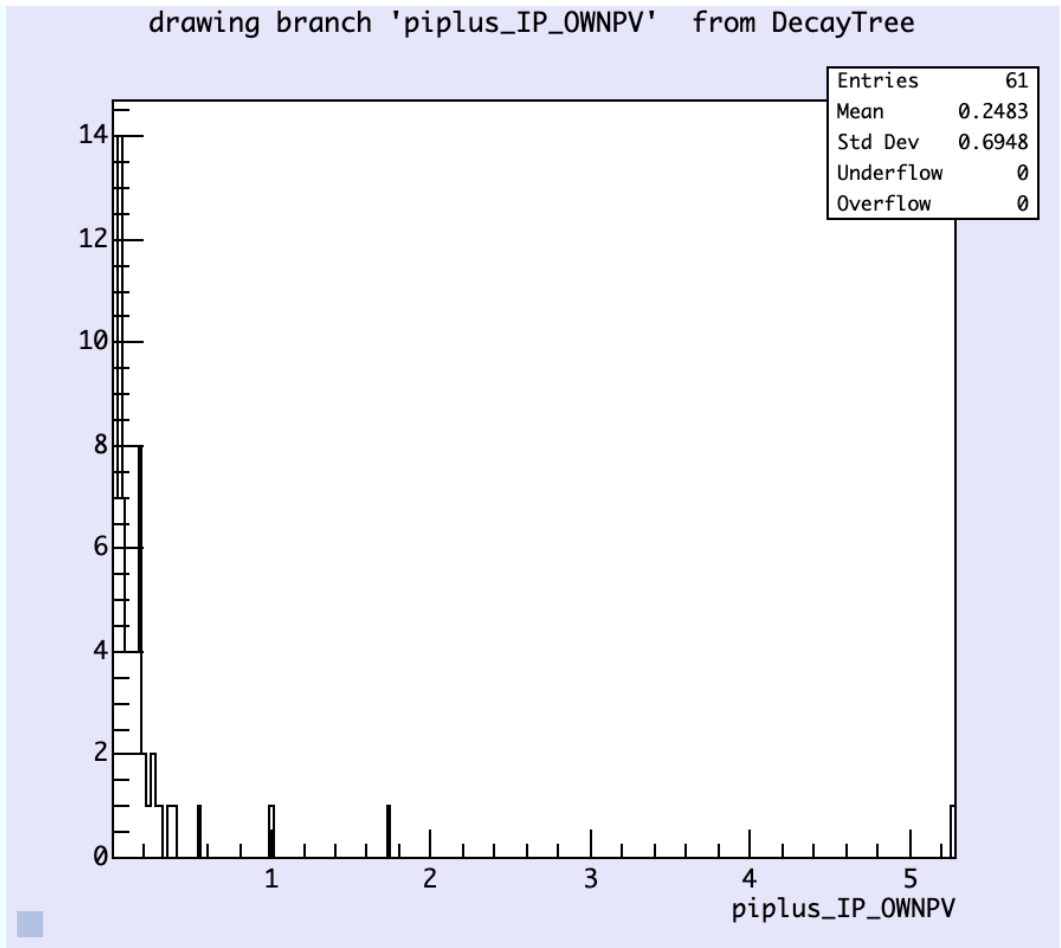
Dstar\_OWNPV\_Y

Dstar\_OWNPV\_Z

Dstar\_OWNPV\_XERR

Dstar\_OWNPV\_YERR

Dstar\_OWNPV\_ZERR



Analysis production for pipeline [#1958460](#)


D02HH\_Starterkit

Commit: [53210696](#)  
Message: Minimise repetition  
Colour code: [Running](#) [Success](#) [Failed](#) [Unknown](#)

Raw YAML [Show](#)

Rendered YAML [Show](#)

Click on a job for more information



Job Name	Test job statistics			Log messages			Estimated total	
	Events	Runtime	Size	Warn	Error	Fatal	Size	Kept
2016_MagDown_PromptMC_D02KK	2548	0:00:34	0.2 MB	1	2	0	0.1 GB	True
2016_MagUp_PromptMC_D02KK	2916	0:00:34	0.2 MB	1	2	0	0.1 GB	True

For support see the [DPA WP2 Analysis Productions](#) channel on mattermost.

LbAnalysisProductions 0.0.2.dev2+g2875ced · LbAPCommon 0.0.4

Explain job log

Fatal errors:

Suggestions:

Other explanations:

- Line 356: Histograms are not being saved as no filename has been specified for storing them. This message is harmless and normally ignored.

Job log [Hide](#)

```
1 | WARNING:lb-run:Decided best platform to use is x86_64-centos7-gcc48
2 | WARNING:lb-run:Decided best container to use is None
3 | # setting LC_ALL to "C"
4 | # Restarting with LD_PRELOAD='libtcmalloc.so'
5 | # --> Including file '/scratch/dir_37441/LOBLDmSB8exnXNrZMnAHaOWqZnKwcmABFKDm3xgaDmABFKDmRw0yXn/DIRAC_t7dkzUpilot/405432667/AnalysisProductions/dynamic/D02HH_Starterkit/2016_MagDown_PromptMC_D02KK.root'
6 | # <-- End of file '/scratch/dir_37441/LOBLDmSB8exnXNrZMnAHaOWqZnKwcmABFKDm3xgaDmABFKDmRw0yXn/DIRAC_t7dkzUpilot/405432667/AnalysisProductions/dynamic/D02HH_Starterkit/2016_MagDown_PromptMC_D02KK.root'
7 | # --> Including file '/scratch/dir_37441/LOBLDmSB8exnXNrZMnAHaOWqZnKwcmABFKDm3xgaDmABFKDmRw0yXn/DIRAC_t7dkzUpilot/405432667/AnalysisProductions/D02HH_Starterkit/ntuple_options.py'
8 | # <-- End of file '/scratch/dir_37441/LOBLDmSB8exnXNrZMnAHaOWqZnKwcmABFKDm3xgaDmABFKDmRw0yXn/DIRAC_t7dkzUpilot/405432667/AnalysisProductions/D02HH_Starterkit/ntuple_options.py'
9 | # --> Including file '/scratch/dir_37441/LOBLDmSB8exnXNrZMnAHaOWqZnKwcmABFKDm3xgaDmABFKDmRw0yXn/DIRAC_t7dkzUpilot/405432667/AnalysisProductions/dynamic/use-jsroot-compression.py'
10 | # <-- End of file '/scratch/dir_37441/LOBLDmSB8exnXNrZMnAHaOWqZnKwcmABFKDm3xgaDmABFKDmRw0yXn/DIRAC_t7dkzUpilot/405432667/AnalysisProductions/dynamic/use-jsroot-compression.py'
11 | # --> Including file '/scratch/dir_37441/LOBLDmSB8exnXNrZMnAHaOWqZnKwcmABFKDm3xgaDmABFKDmRw0yXn/DIRAC_t7dkzUpilot/405432667/output/Local_q6gJU7_JobDir/prodConf_DaVinci_00012345_00006789_1.py'
12 | # <-- End of file '/scratch/dir_37441/LOBLDmSB8exnXNrZMnAHaOWqZnKwcmABFKDm3xgaDmABFKDmRw0yXn/DIRAC_t7dkzUpilot/405432667/output/Local_q6gJU7_JobDir/prodConf_DaVinci_00012345_00006789_1.py'
13 | # applying configuration of ProdConf
14 | ProdConf version unknown
15 | # /***** User ProdConf/ProdConf *****/
16 | # I-NOFEvents = -1 (default: -1)
17 | # I-XMLSummaryFile = 'summaryDaVinci_00012345_00006789_1.xml'
18 | # I- (default: 'summary%(Application)s_%(OutputFilePrefix)s.xml')
19 | # I-HistogramFile = ''
20 | # I-DDDBTag = ''
21 | # I-CondDBTag = ''
22 | # I-ProcessingPass = ''
23 | # I-FirstEventNumber = 0
24 | # I-OptionFormat = 'WGProd' (default: '')
25 | # I-AppVersion = 'v45r5' (default: '')
26 | # I-Application = 'DaVinci' (default: '')
27 | # I-TCK = ''
```



- CI is useful but nothing beats debugging locally when things fail
- Analysis Productions includes a command for interactively reproducing tests

Browse output file [Show](#)

Reproduce on Ixplus:

```
1 | lb-ap reproduce 1958460 "D02HH_Starterkit" "2016_MagDown_PromptMC_D02KK" 0
```

Explain job log [Show](#)

- Starting to use pre-commit\* to run formatting and basic (i.e. quick) tests before committing to git



\* <https://pre-commit.com/>



- Software from large HEP experiments is complex
- Continuously investing in testing and validation is very valuable
- CI best practices are valuable for all software *including data analysis*
- Investigating if CI can prove valuable for later stages of analysis
  - Selections, Efficiencies, Fitting, Systematics, ...
- Making custom GitLab CI runners for a community is relatively easy
  - Useful when you have specialised requirements
  - Check out [gitlab-runner-api](#) if this is interesting to you!





Questions?