
LHCb operations



Federico Stagni, CERN

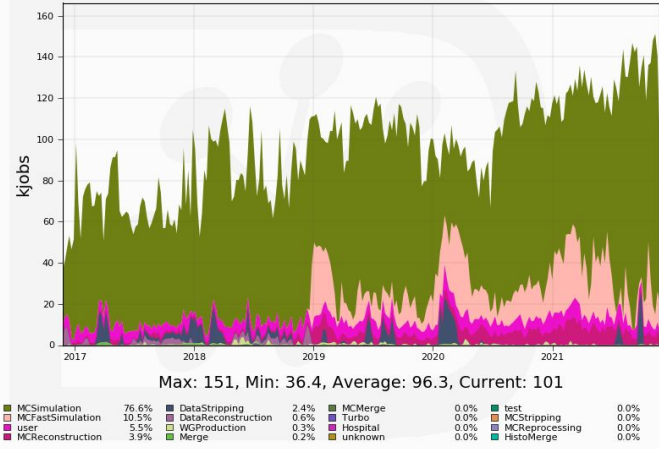
For the LHCb distributed computing team

[pre-GDB on operations effort](#), Thursday 24 Feb 2022




Operations (historic)

of running jobs, last 5 years
261 Weeks from Week 47 of 2016 to Week 47 of 2021

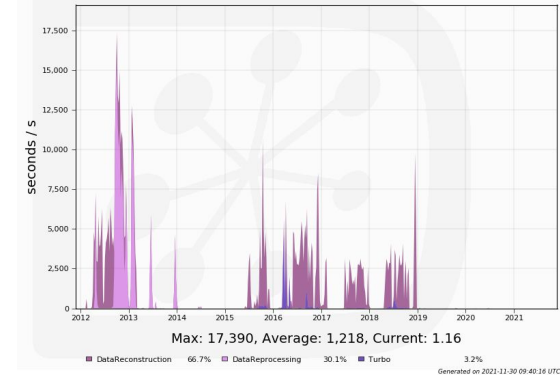


CPU time, by JobType, last 10 years
522 Weeks from Week 47 of 2011 to Week 47 of 2021



Computing work dominated by MC production (95%), simulating about 150 million events per day

CPU time, by JobType, last 10 years
522 Weeks from Week 47 of 2011 to Week 47 of 2021



Distributed computing usage and evolution “almost independent” from LHC(b) schedule. Nevertheless, within operations and developments Run3 is a milestone.

Computing model and Computing resources

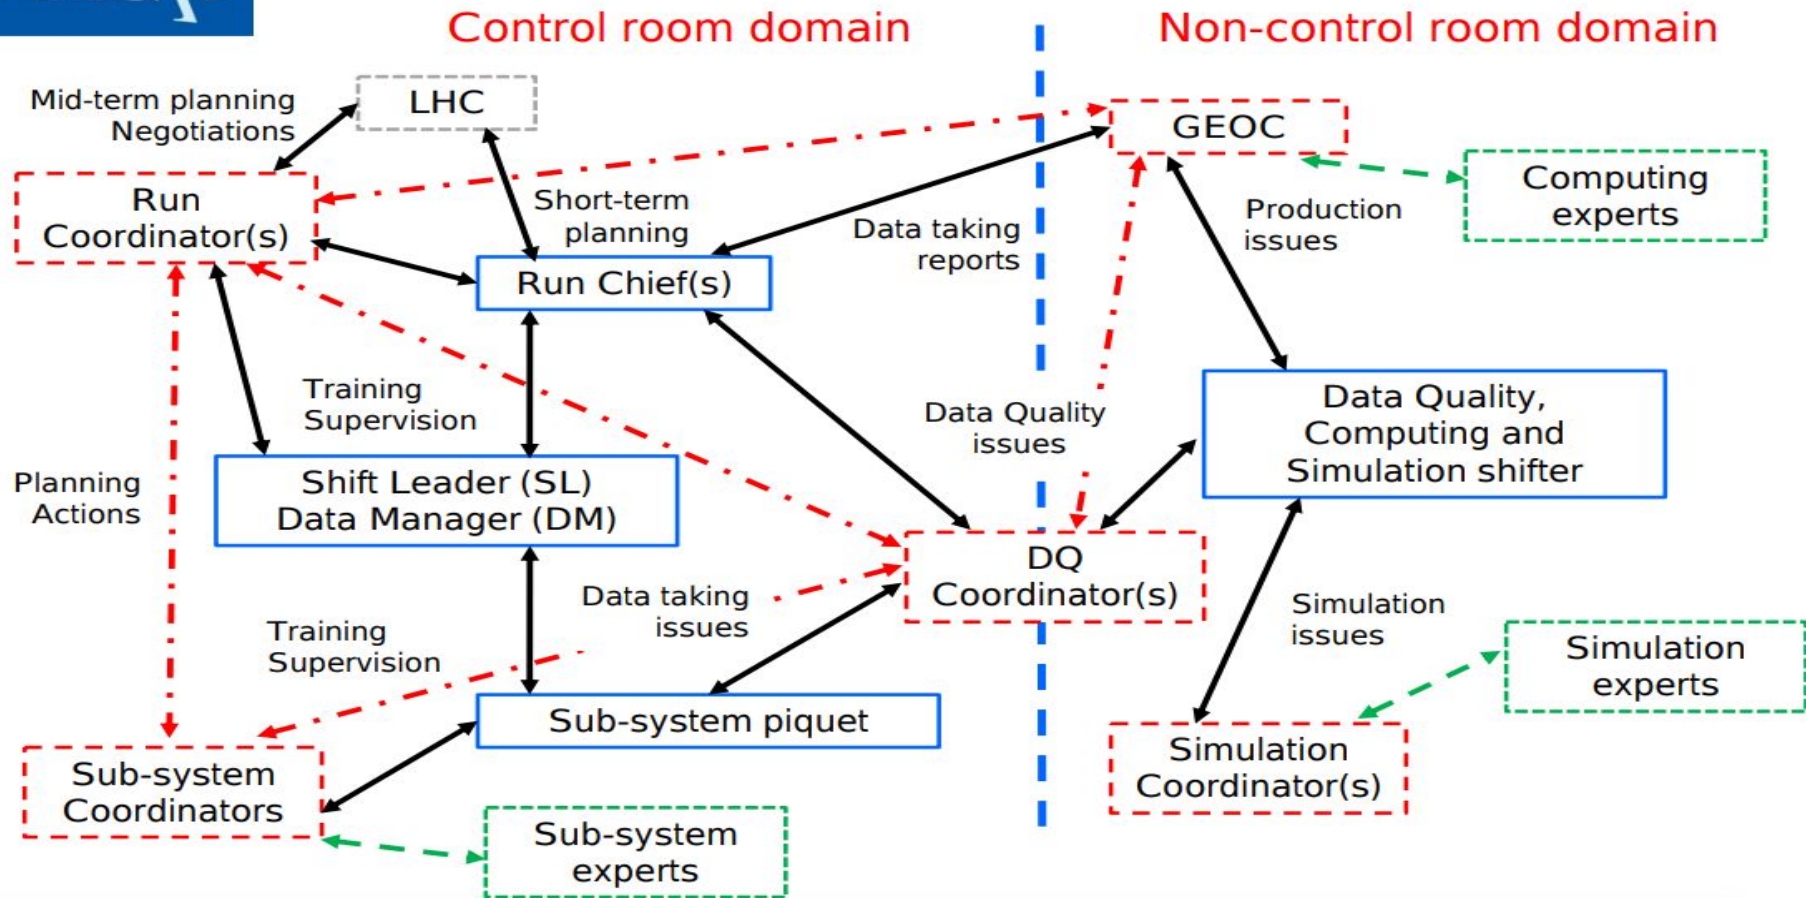


Run3 won't call for a revolution to the computing model: the goal is to stay flexible, and run “where possible”

- MC keeps being the workflow that we can more easily run
- It's becoming anyway increasingly difficult to significantly expand the computing resources available through the “standard” WLCG resources
 - ARC and HTCondor Computing Elements
- “opportunistic resources” is now a catch-all term for whatever is not pledged to WLCG
 - “Grid” sites but mostly off-grid clusters and HPC. And of course the LHCb HLT farm.
- expansion can't happen (much) further without the use of non x86_64 architectures
 - (LHCb)DIRAC can be installed on aarch64 and ppc64le platforms
 - but LHCb production workflows must be validated
- DIRAC-based central monitoring is a must



Lines of communication



Distributed computing ops team

- Grid Expert (changing every quarter)
 - only 4 or 5 experts around
- Weekly shifter: DQCS
 - Will have also Data Quality responsibilities when data taking starts
 - Could be “anyone”
- Production manager(s)
 - currently 3
- Data Manager (1)
- Sites manager(1)
- Tier2D coordinator
- T1s contacts (at least, a few)
- CERN IT liaison
- DIRAC experts [developments and dev-ops]
 - (doing the roles above too)

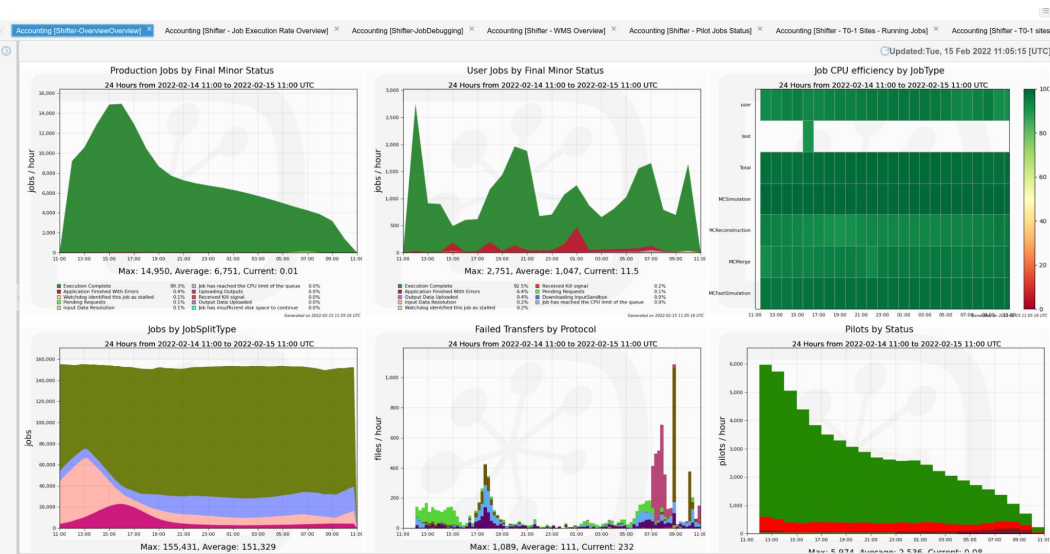
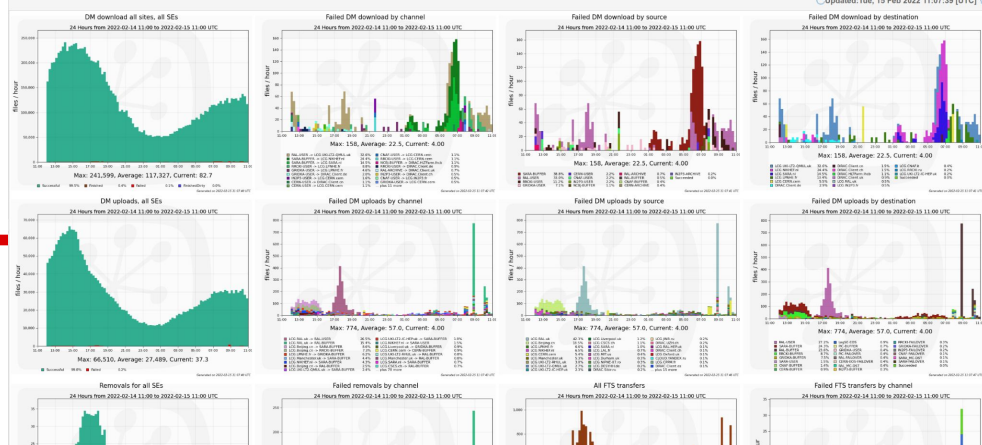
a grand total of
about 3.5 FTE!

Ops team work

- Ops meetings 3 times per week (~30 mins)
 - run by the DQCS, quasi-full team in attendance
- “Usual” channels:
 - mattermost, emails, logbook
 - GGUS
- The current operations scheme is rather solid, no need for big changes
 - usual cry for more people!

Dashboards (for shifters)

- DIRAC-based dashboards used in production (what the shifters look at)
- Backends: MySQL, ElasticSearch/OpenSearch



Pros:

- “everything in one”
- internally maintained

Cons:

- static-ish
- not fully open to non-LHCb members
- internally maintained

More Dashboards (not for shifters)



DIRAC and non-DIRAC ones, like:

- [FTS3 monitor](#)
- [ETF/CheckMK](#)
- MONIT [for DIRAC hosts](#)
- (MONIT) [DIRAC services responsiveness](#)

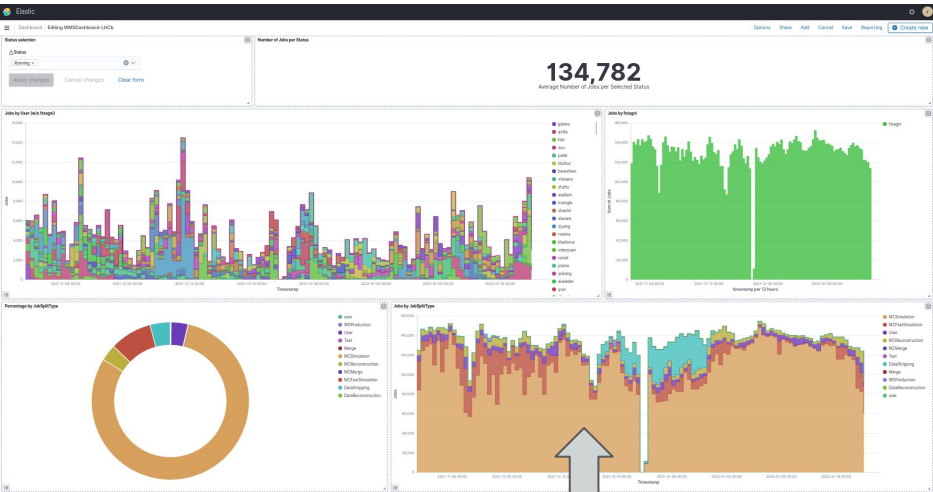
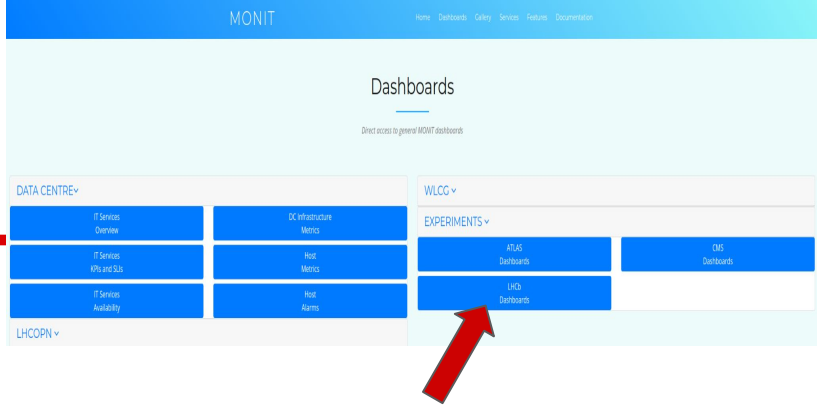
Monitoring developments



Moving towards “standards”:

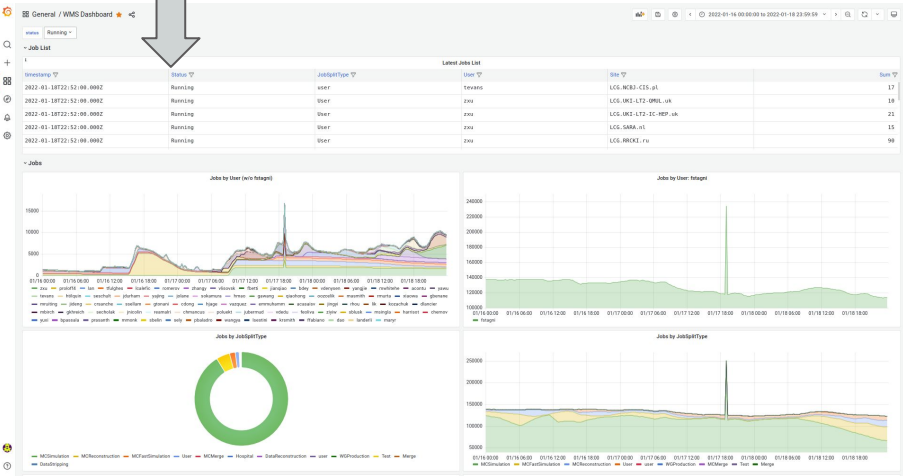
- MySQL → ES for real time monitoring
 - and use “modern” visualization tools
 - historic accounting will stay in MySQL
- Aim to have ES for all real time monitoring of LHCb distributed computing
 - trying to keep it “standard” (maybe influxDB a day)
- Looking at operational intelligence developments

ES Dashboards for LHCb



WMS dashboard in Kibana

WMS dashboard in MONIT-Grafana



Joint Optimization



- Where you see a potential for saving effort or optimization
- Possible areas where you think the effort can be shared across experiments
- Try to *standardize* monitoring activity across experiments
 - MONIT serves as common infrastructure
 - Documentation for grafana is not always top-notch
 - Implement standardized “ready-to-use” monitoring dashboards to be available for use for all experiments/users?
 - at least, they would serve as examples
 - Job Monitoring from a ES datasource is already being used by LHCb, ATLAS and CMS. Any chance of trying to provide a standardized interface across experiments?

Q?