

LHCb
ГЛСР



10th DIRAC Users' Workshop
June 19th 2024

Vladimir Romanovskiy, Federico Stagni

List of news

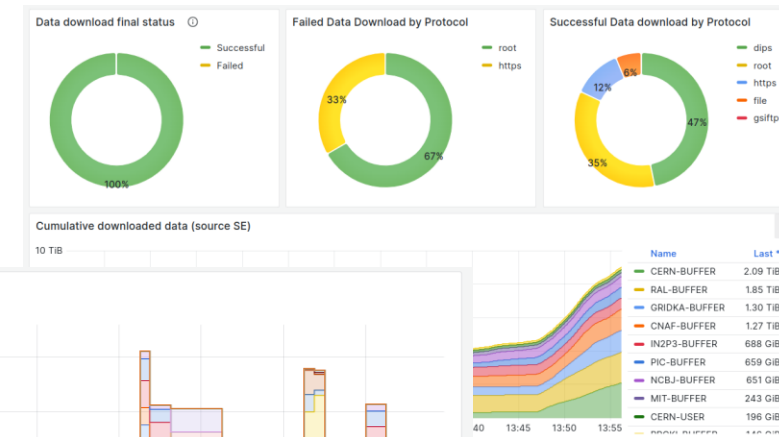
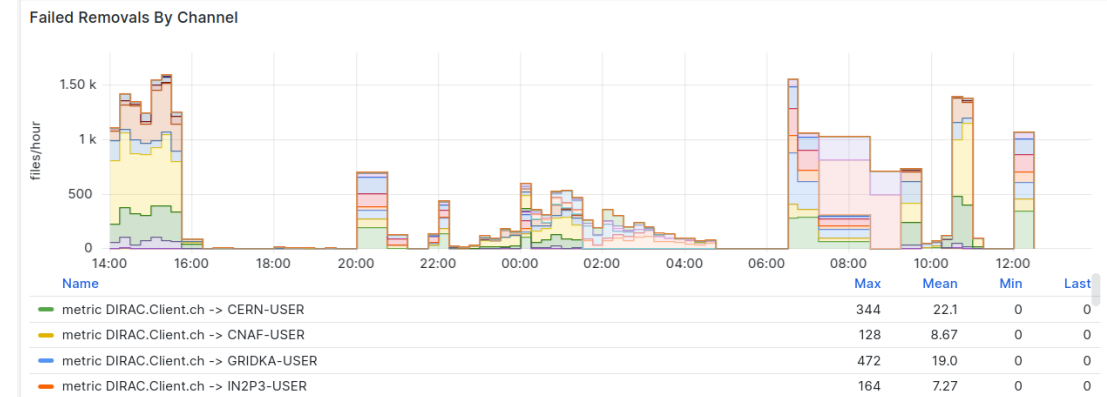



- From EL7 to EL9
 - Servers:
 - Productions: Almost finished
 - Certification: **Done**
 - VOMS inside IAM: **Done**
 - Containers: Alma9 as base for SingularityCE: **Done**
 - OpenDistro to OpenSearch: almost **Done**
- WLCG Grid: ARM Farm at Glasgow Tier2
 - Tested by LHCb, but we are not ready to use it at scale
 - DIRAC OK, software not yet fully OK
 - ATLAS would be OK to have ARM resources pledged to them (starting from 2025)

Monitoring



- Distributed and core software activities monitoring fully ported to [grafana](#) (CERN MONIT infrastructure)
 - This includes shifters' monitoring (computing part)



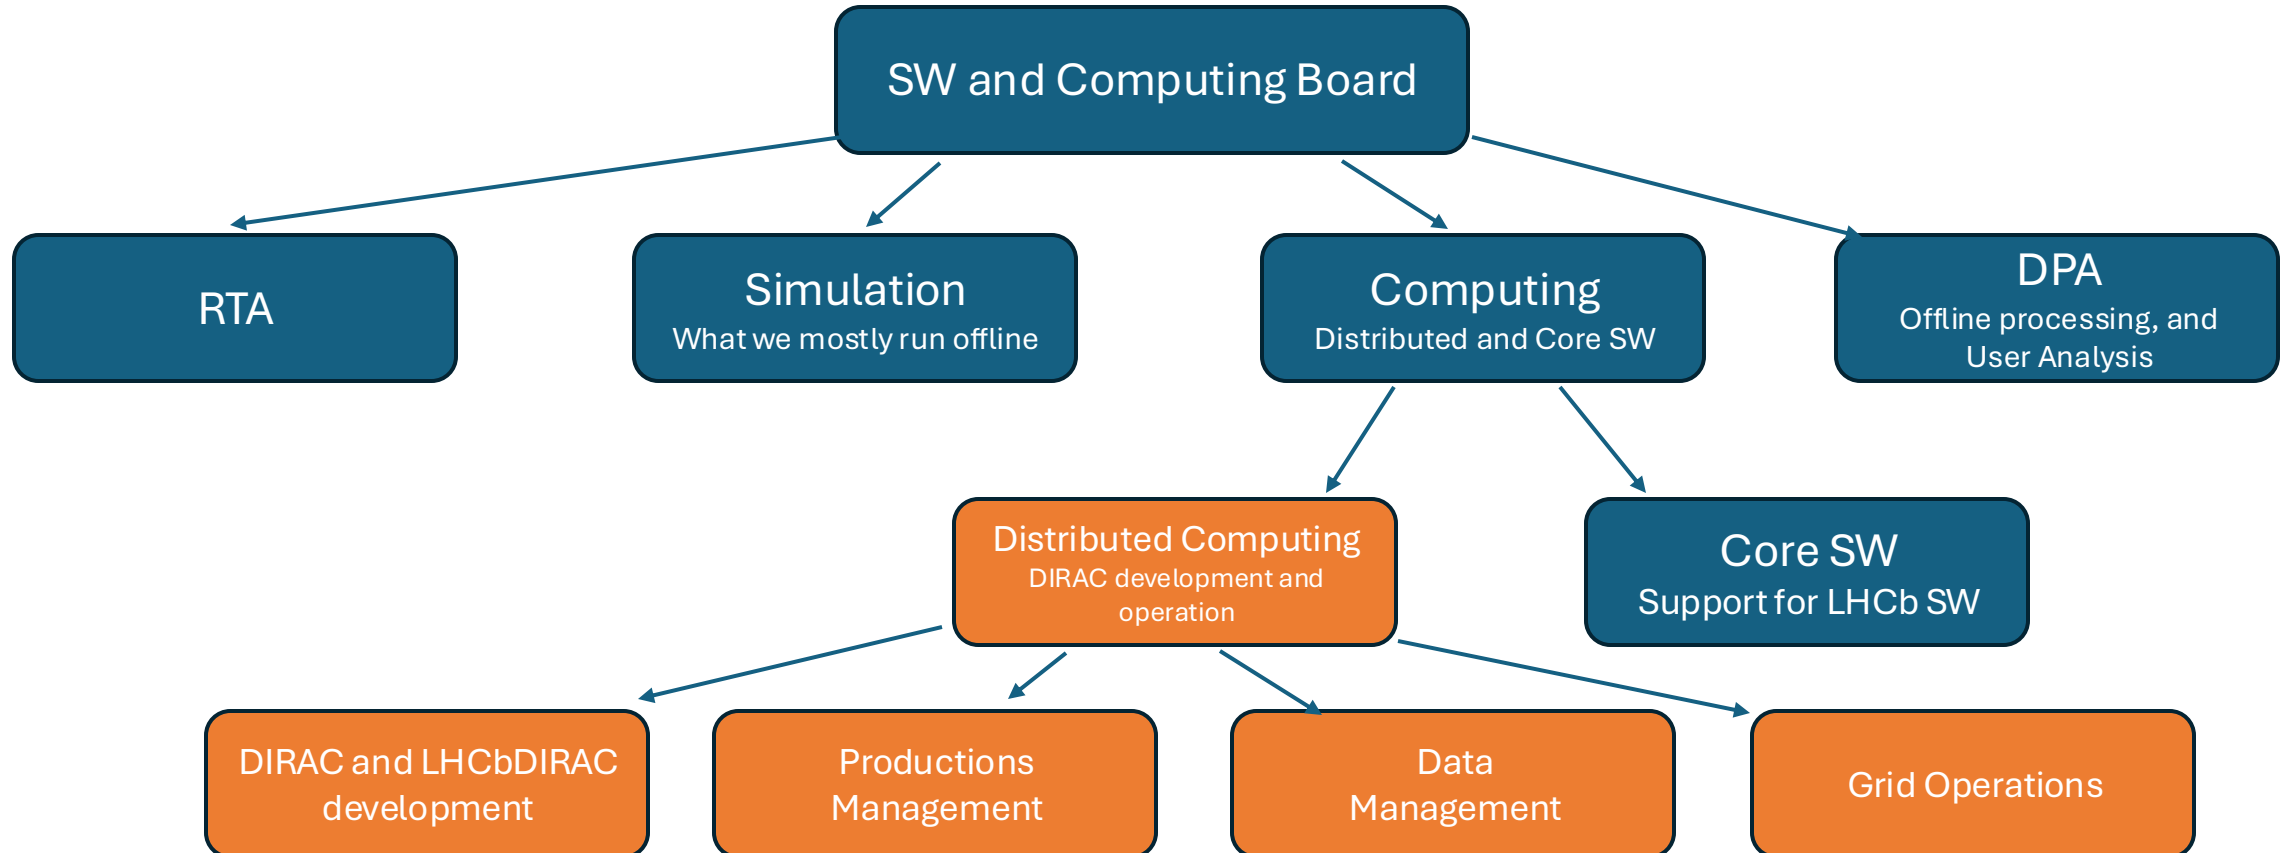
LHCb Operations

Because Federico believes it might be useful for the organization of other VOs

Because it *might* be that some development choices of DIRAC fit better our way we do operations

We are **not** saying you should do like LHCb!

High-level organization view



Roles, and people

/1




- Data Manager(s): Christophe++
- Production Managers:
 - Real Data productions: Raja
 - MonteCarlo productions: Vladimir
 - Analysis productions: Chris B

NB: all these are stable roles – they need experience – no shifters!

Roles, and people

/2



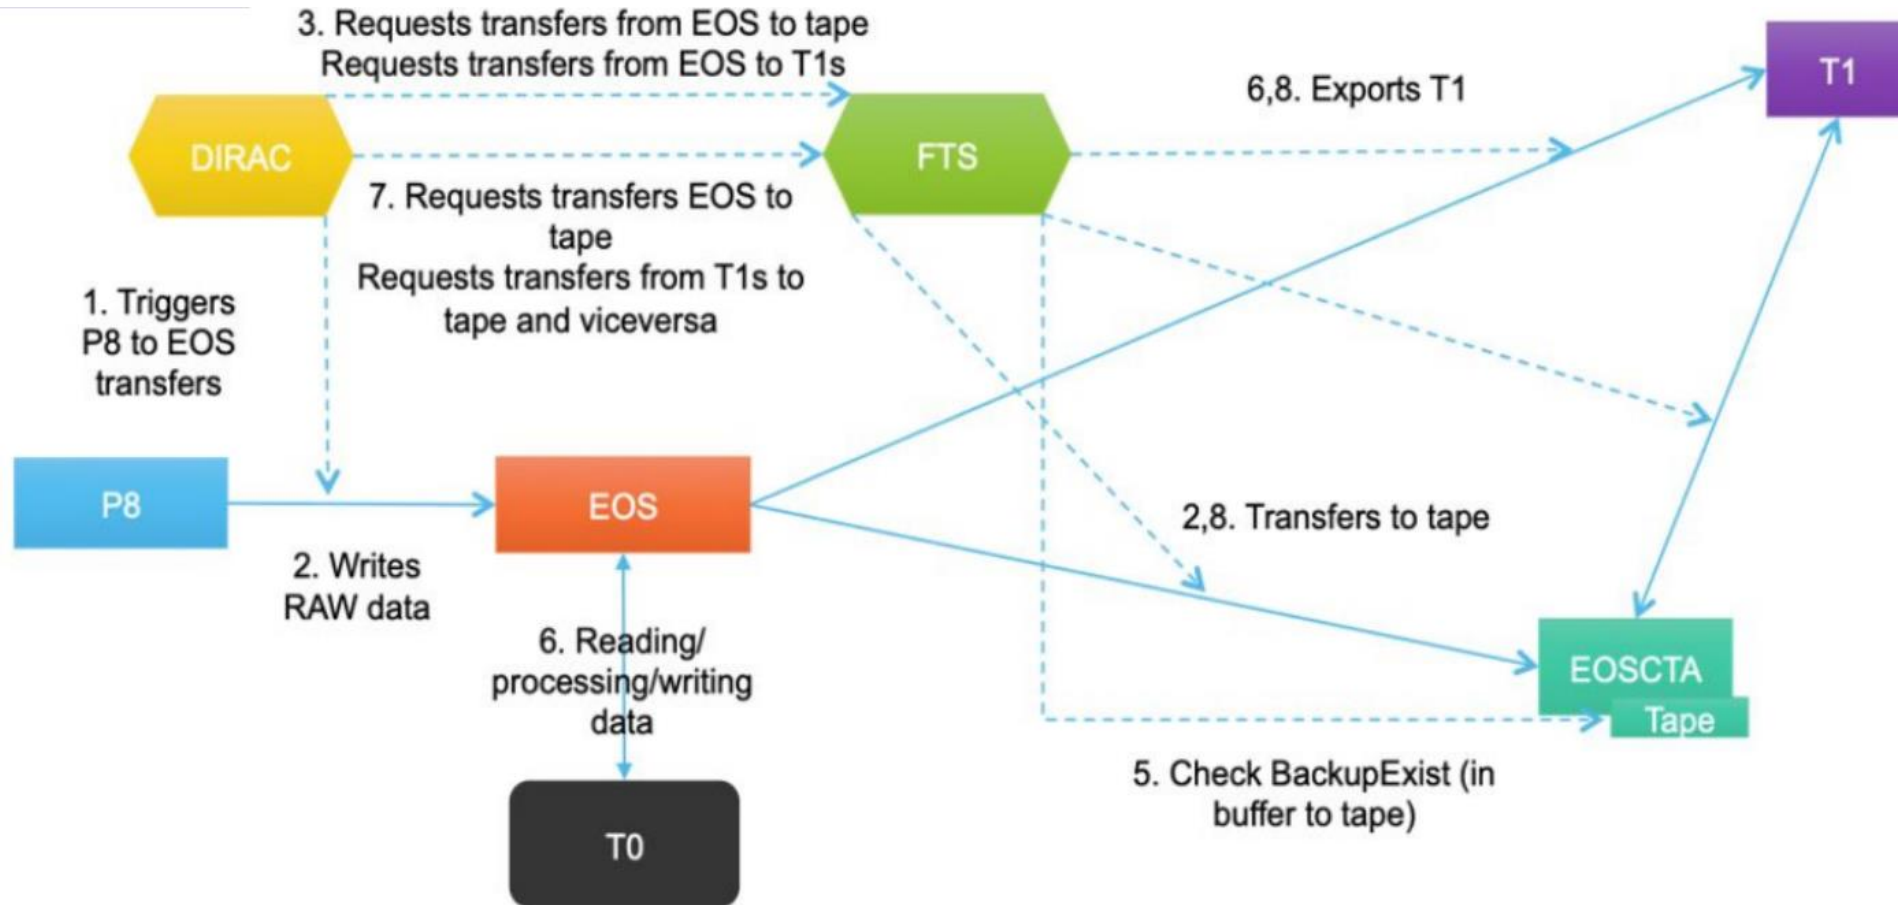
- Grid operations:
 - Sites managers: Vladimir, Henryk
 - Not changing over time
 - Grid Experts: Concezio, Federico, Mark, Henryk
 - Each covers the role for 3 months per year
 - Shifters: large-ish group of shifters
 - Non-experts, including students
 - Covering the role for 1 week
 - Basically: they look at plots and they report possible issues. Plus they do some minor administrative tasks
 - We try to give them precise instructions
 - and TBH, we do not expect too much

Data Management bits



- We have storage elements almost only at Tier0 and Tier1
 - Tier0 is CERN, and there are 8 Tier1s. Also exposing Tape SEs.
 - We have a few selected Tier2D (Tier2 with disk).
- Users can not move datasets around.
- Older MC samples are removed from disk (stay on tape). Users (or working groups) who would like to run on these old datasets have to ask data managers to first stage them.
 - We do run some very simple data popularity analysis
- For real data productions, we stage data ahead of the processing, and remove them from disk ("buffer") as soon as it is processed.

Pit export



User analysis



- We encourage users to run analysis through the Analysis Productions System.
 - Which is a client of the Transformation System
- Most of the user jobs are submitted through Ganga.
- We occasionally put users who "misbehave" in "low priority" group.

Summary



- Moving from EL7 to EL9: almost done
- ARM processors: we tested them
- Monitoring is consolidating on CERN/MONIT infrastructure
 - Meaning grafana
- The transformation system is not directly exposed to the users
 - No productions, no dataset manipulation