LHCb and DIRAC

Concezio Bozzi, INFN Ferrara
For the LHCb distributed computing team

2nd Virtual DIRAC workshop, May 10th 2022

THE INTERWARE

What do you use DIRAC for, and which DIRAC functionalities you don't use, and why?

LHCb uses DIRAC for all its distributed computing activities, including:

- WMS (pilots, jobs)
  - + productions management
- DMS (replica catalog, storage management)
  - + dataset management
- Accounting, Monitoring

As of today, we are running v7.3 in production

- Fully migrated to python 3

We don’t use DIRAC for:

- metadata and provenance catalog (LHCb bookkeeping)
- Production System (we have an LHCb version)
- Interacting with Clouds, because as right now we have no clouds!
  - we have used *vcycle* in the past, we’ll use DIRAC for the purpose in the future

Do you have a DIRAC extension? Why?

- LHCbDIRAC (DIRAC extension)
- LHCbWebDIRAC (WebAppDIRAC extension)

The main reasons behind these extensions are:

- ➢ The LHCbDIRAC Bookkeeping (a metadata and provenance catalog (backend: Oracle)
- ➢ The LHCbDIRAC Production system
  - including some LHCb-specific TransformationPlugins

- LHCbPilot (Pilot3 extension)

  for getting LHCbDIRAC from LHCb CVMFS, and adding LHCb-specific tags

Do you think some of the extensions could become part of the vanilla projects?

- Probably not, as they are all LHCb specific developments.

- LHCb developers are the main DIRAC developers and maintainers, and try as much as possible to develop in vanilla DIRAC.

What is your biggest frustration with DIRAC?

There are no “biggest frustrations”, some “annoyances” that are being tackled e.g.

- “Ageing” Web application:
  - Monitoring and accounting plots are slow
  - Loading selector can also be slow
  - Selection conditions are sometimes cumbersome to apply
    - “NOT” feature no longer available…
- Better pilot monitoring
  - Some functionalities are there but they do not work
  - e.g. pilots outputs and errors
- “Out-of-the-box” monitoring of services and agents

You can magically add one feature to DIRAC, what is it?

- Pilot logs for every computing resource
  - that don’t disappear in 24 hours
- Auto-magical configuration of opportunistic resources (e.g. HPCs…)

Any notable operations incident in the last year?

none

To support your "Grid", do you have to use other systems than DIRAC?

WLCG environment:

- VOMS → tokens (IaM)
- FTS3
- BDII
- GOCDB
- GGUS

We also provide info to/from:

- MONIT (Cern monitoring infrastructure) ← some recent progress
- ETF (successor of “SAM”, for site testing)
- CRIC (Grid InfoSys -- was ATLAS)
  - the “new BDII”, but not really

but we don’t fully rely on them

How would you rate the communication?

Excellent :-)

In the last two years, what has been the DIRAC usage in terms of jobs ran, CPU (or wall time) used, and data transfers?

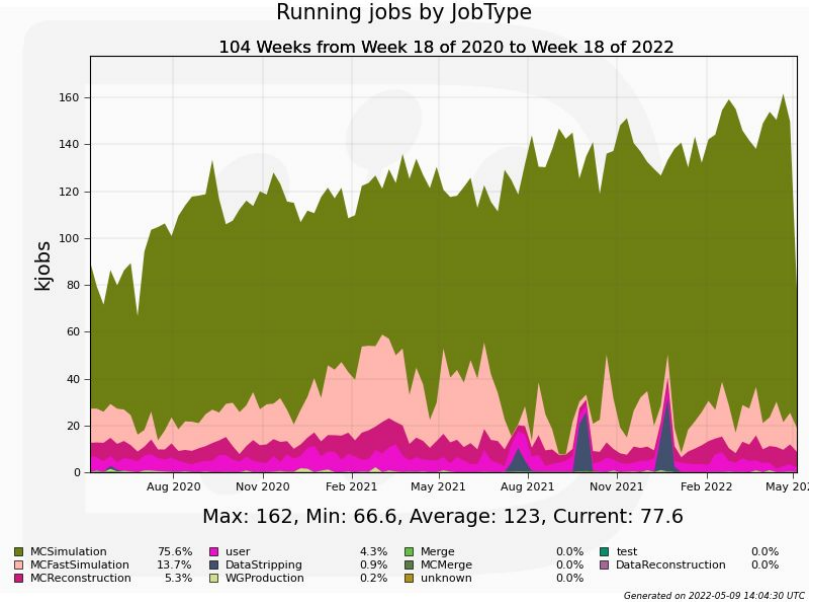

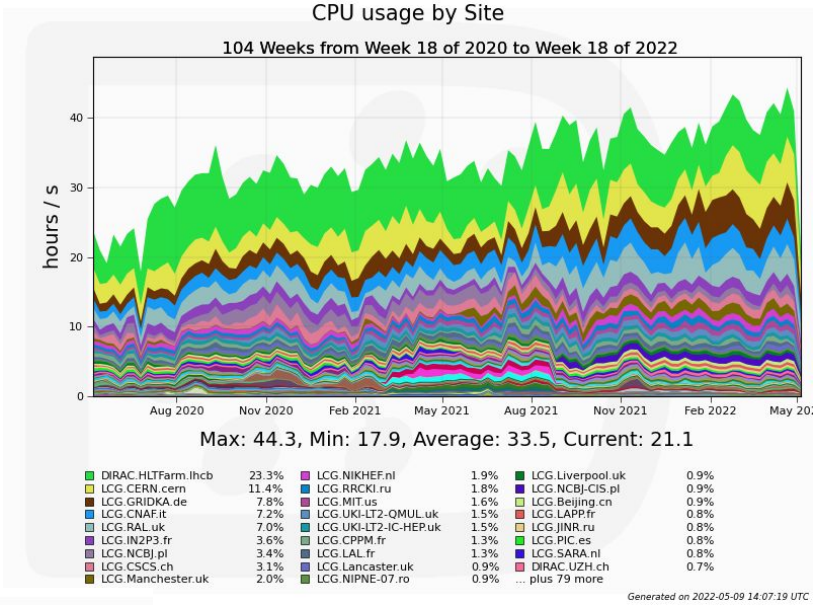

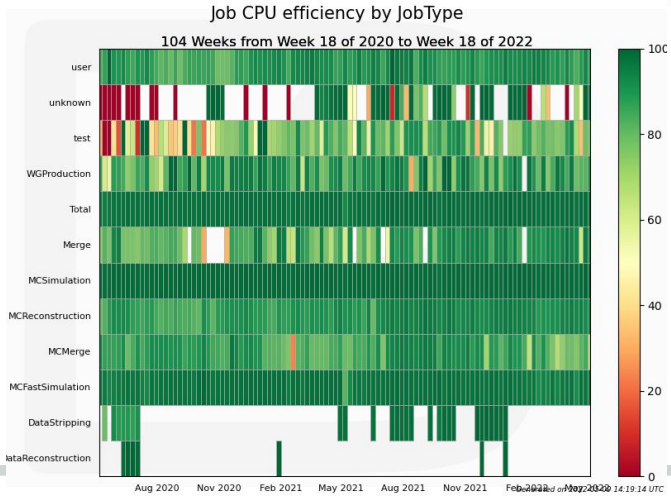

In the last two years, what has been the DIRAC usage in terms of jobs ran, CPU (or wall time) used, and data transfers?

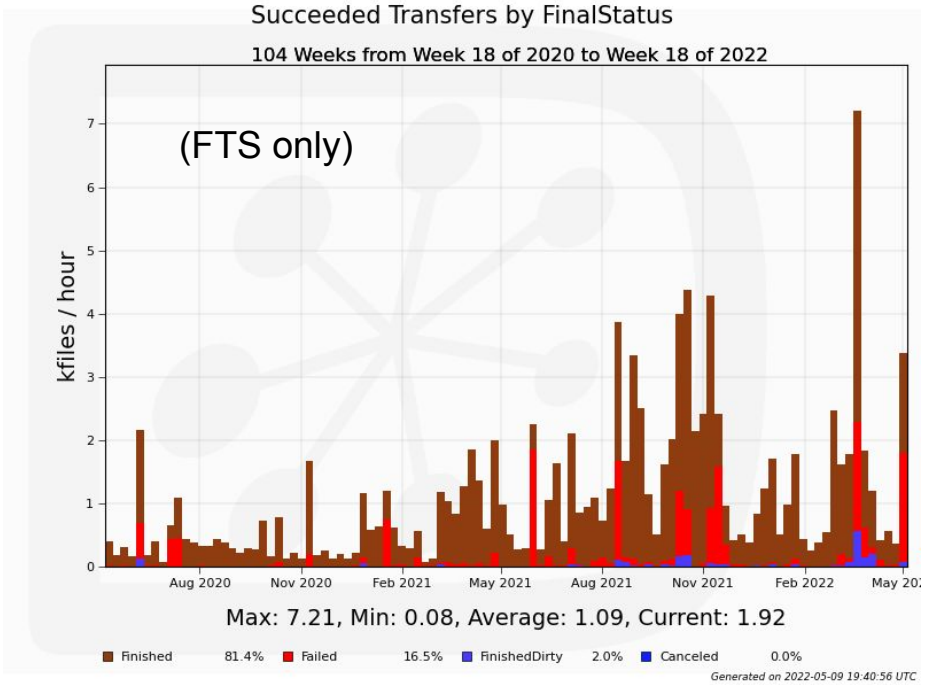

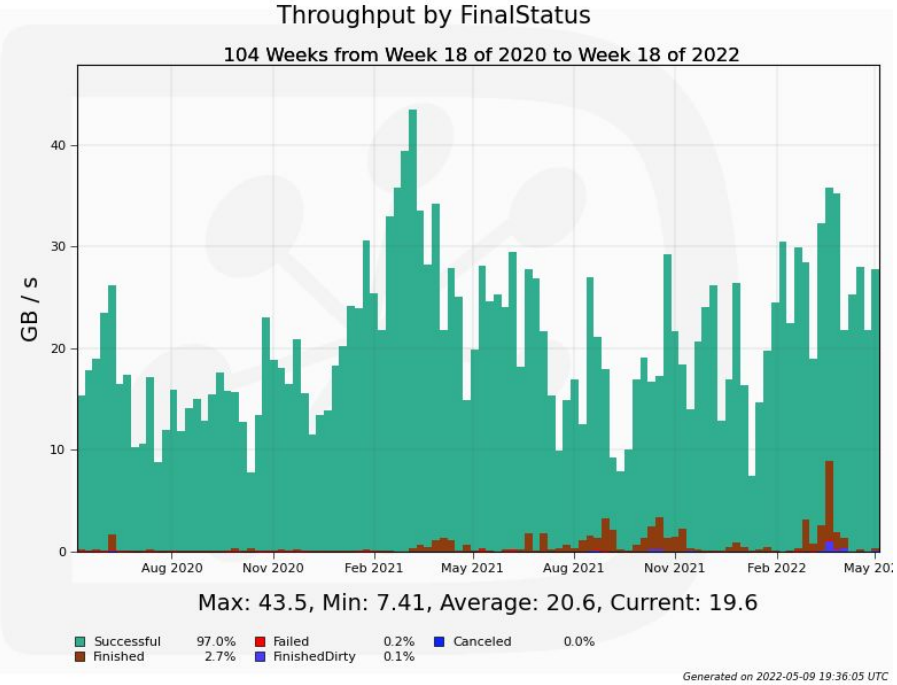

Resources

- Computing:
  - LCG resources: ARC and HTCondor CEs
  - opportunistic: clusters with DIRAC SSH CE, HLT farm, HPCs
  - Inner CEs:
    - We use PoolCE for all multiprocessor WNs
      - but we run almost exclusively single processor jobs
    - We tried to set SingularityCE
      - only activated on a few selected CEs
      - monitoring issues
- Storage: using CTA at CERN and Antares at RAL (replacement of Castor)
- Transfers: using https for TPC for all disk storage (except RAL)
  - Commissioning data transfers at Run3 rates

Data challenges

CERN disk → T1 disk → T1 tape

Site	expected Speed (GB/s)	Target achieved
CERN	11	Yes
CNAF	1.72	~Yes
GRIDKA	2.23	Yes
IN2P3	1.25	Yes
NCBJ	1.32	No
PIC	0.2	Yes
RAL	2.96	~No
RRCKI	0.25	Yes
SARA	1.07	Yes

1)Test 2) ~~Fail~~ Learn
3)Change 4)Go to 1

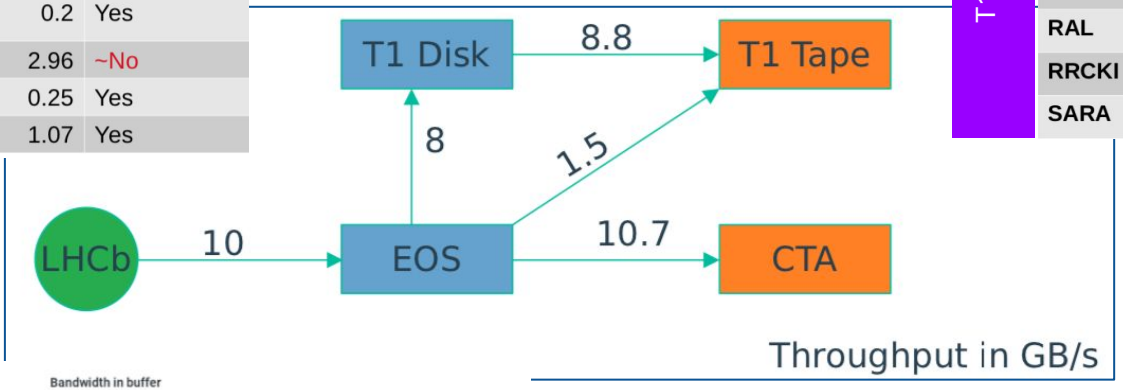

T1 tape → T1disk

Site	Expected Speed (GB/s)	Success
CNAF	1.35	Yes
GRIDKA	1.36	Yes
IN2P3	0.98	Yes
NCBJ	0.91	No
PIC	0.17	Yes
RAL	1.93	No
RRCKI	0.21	Yes
SARA	0.74	Yes

LHCb → CERN disk → CERN tape

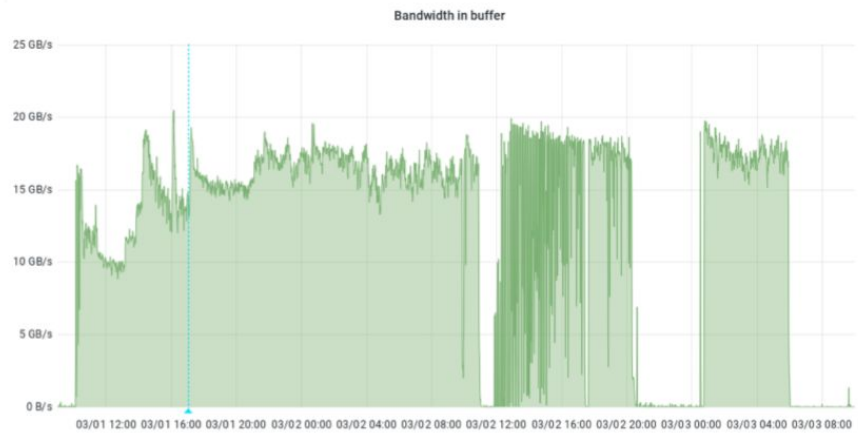

We are on good tracks...

... but not there yet

2022 should be okay though

Recent HPC developments: MareNostrum@BSC

- Some implementations mentioned by Alexandre yesterday were needed to use HPC centers
  - No external network connectivity
  - No CVMFS
  - Many-core nodes
- MareNostrum at Barcelona SC Center as test case
- CVMFS shrink-wrap
- PushJobAgent
- PoolCE

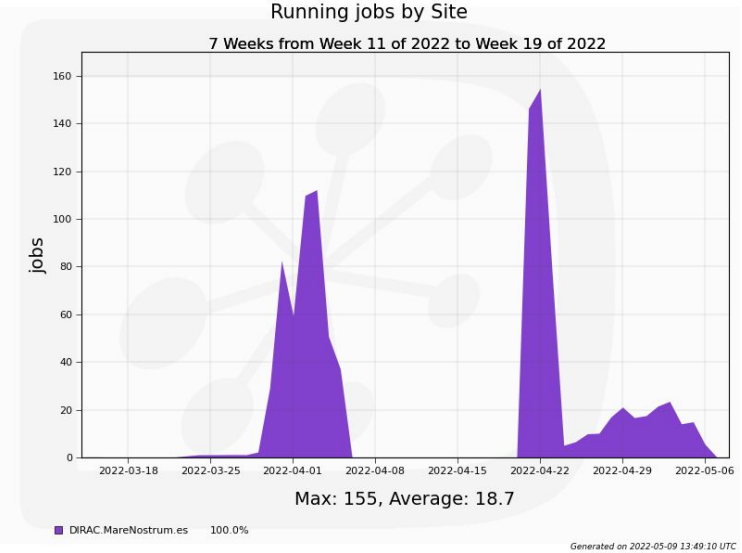

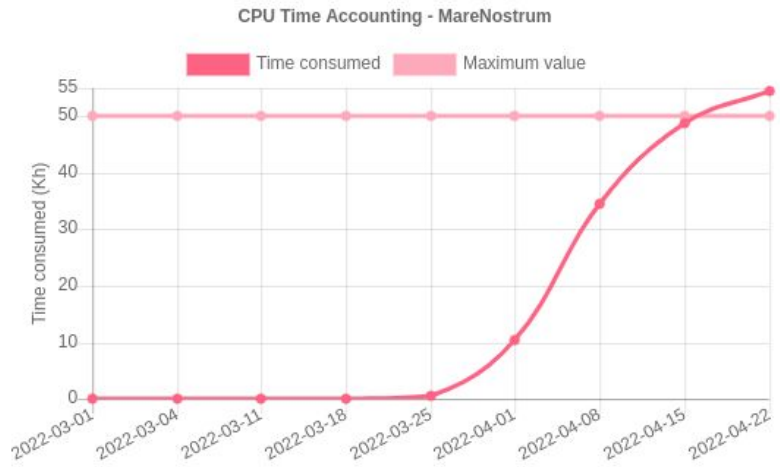

Grafana WMS dashboard

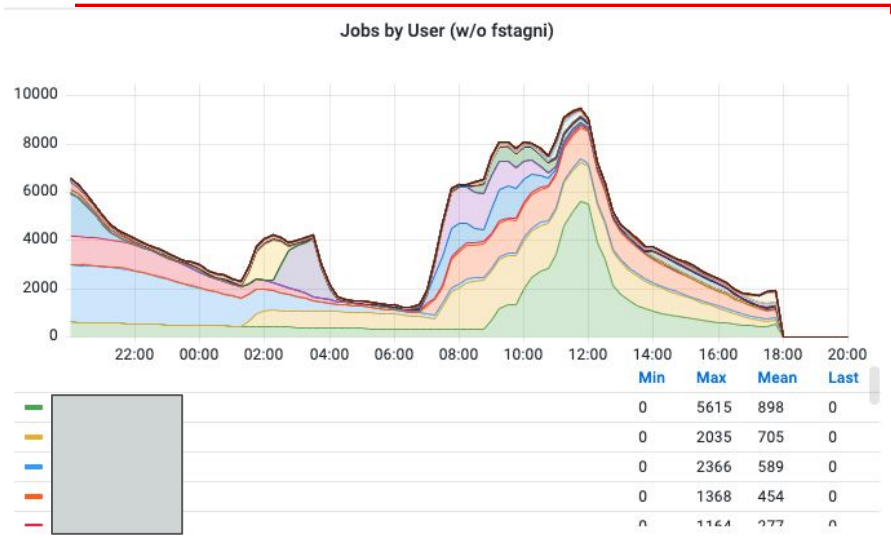

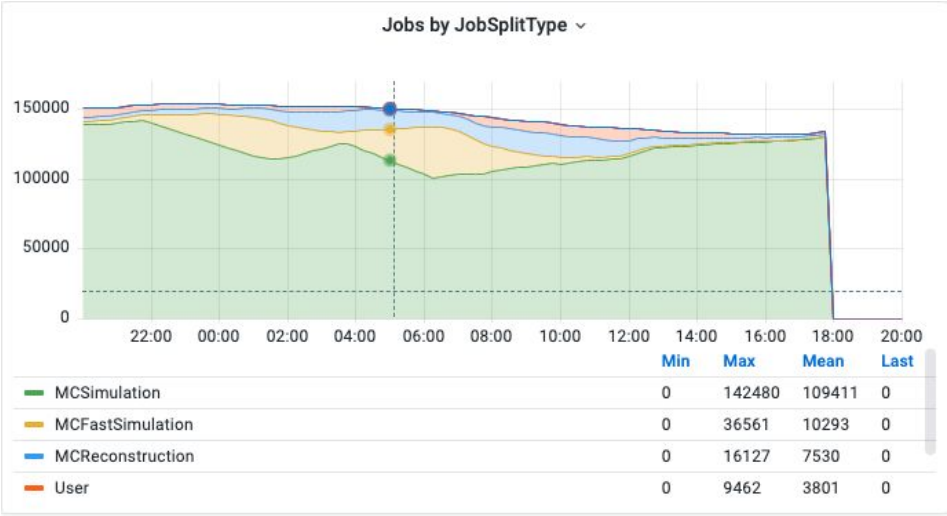

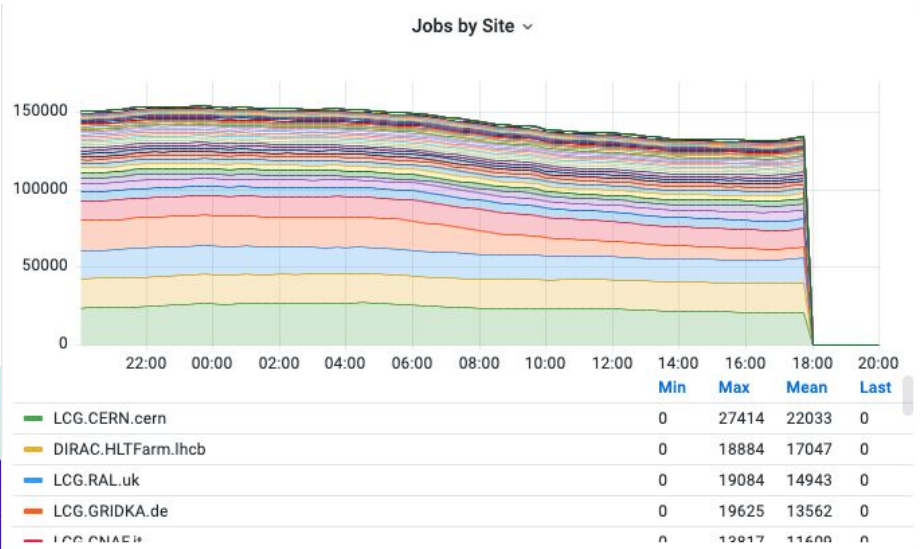

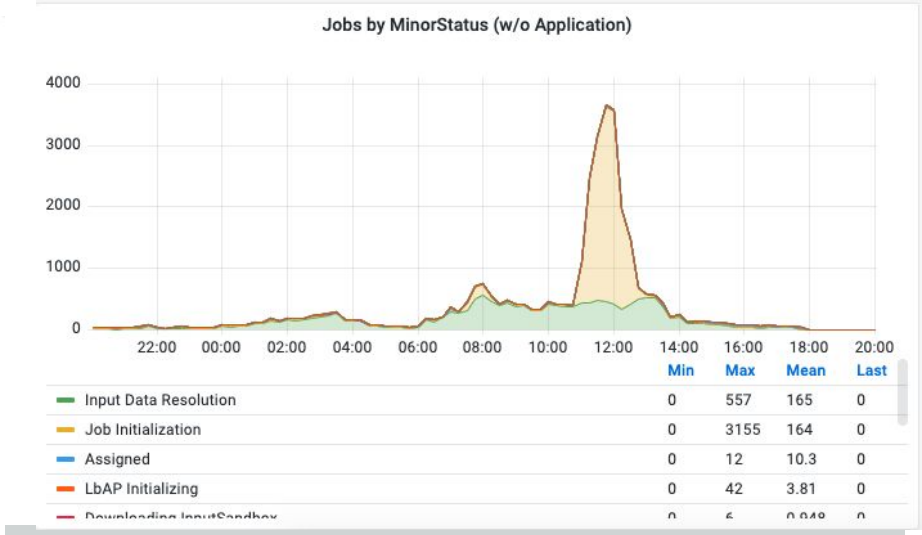

Grafana WMS failure dashboard

Grafana DIRAC services dashboard

LHCb-DIRAC Services Status (History)

Probe Time (UTC)	Host	Services	Probe Status	SvcRespTime
2022-05-09 17:18:14	lbvobox304.cern.ch	BookkeepingManager	OK	1
2022-05-09 17:33:36	lbvobox304.cern.ch	BookkeepingManager	OK	2
2022-05-09 17:48:24	lbvobox304.cern.ch	BookkeepingManager	OK	2
2022-05-09 18:03:44	lbvobox304.cern.ch	BookkeepingManager	OK	2
2022-05-09 18:18:33	lbvobox304.cern.ch	BookkeepingManager	OK	1
2022-05-09 18:34:02	lbvobox304.cern.ch	BookkeepingManager	OK	1
2022-05-09 18:49:05	lbvobox304.cern.ch	BookkeepingManager	OK	1
2022-05-09 19:03:53	lbvobox304.cern.ch	BookkeepingManager	OK	1
2022-05-09 19:19:04	lbvobox304.cern.ch	BookkeepingManager	OK	1
2022-05-09 19:33:18	lbvobox304.cern.ch	BookkeepingManager	OK	2

1 2 3 4 5 6 7 8 9

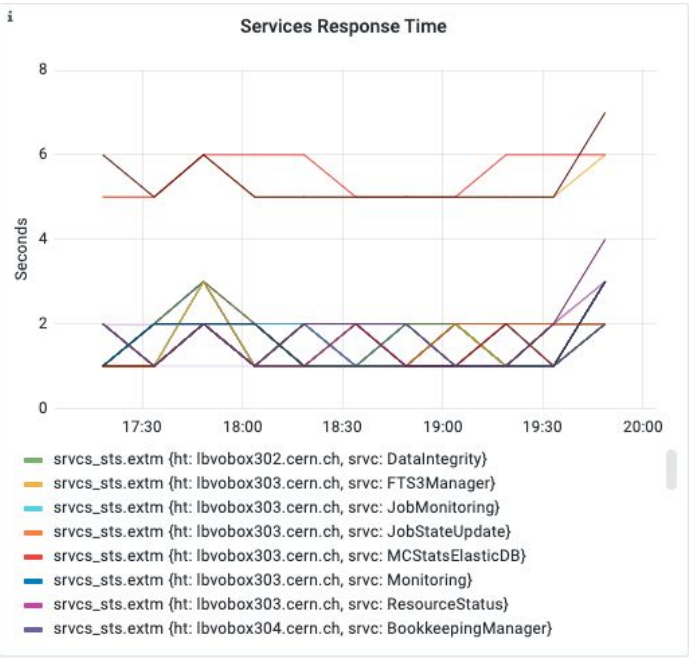

?