ATLAS BigPanDA monitoring
BelleII DDM

S. Padolski (NPPS)

BigPanda Monitoring

Introduction

Reference talks:

- Alexei (https://indico.bnl.gov/event/6290/)
- Tadashi (https://indico.bnl.gov/event/6333/)

Instances: Atlas (https://bigpanda.cern.ch), Compass, EC2 (LSST, LQCD)
http://pandamonitor.org

Current content

- A window into the PanDA system
- > 100 different views
- From production dashboards to logs
- Covers scope in range 1...10^{11} events

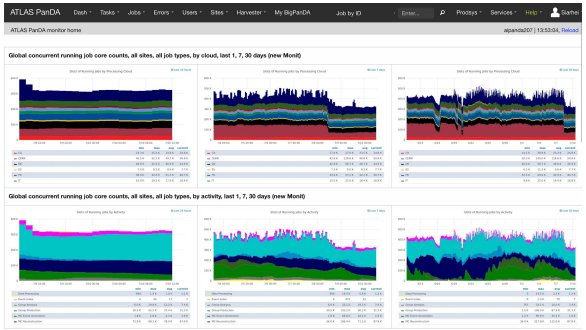

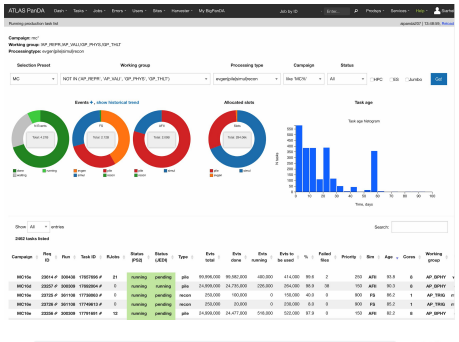


```
2019-07-10 01:23:07|143203|pUtil.py      | current dir:
/home/tmp/atlas_jBm5m8QO/Panda_Pilot_143203_1562721787
2019-07-10 01:23:07|143203|pilot.py       | Pilot options:.....................................................
2019-07-10 01:23:07|143203|pilot.py       | appdir: None
2019-07-10 01:23:07|143203|pilot.py       | debugLevel: 0
2019-07-10 01:23:07|143203|pilot.py       | jobrec: False
2019-07-10 01:23:07|143203|pilot.py       | jobRequestFlag: True
2019-07-10 01:23:07|143203|pilot.py       | jobSchedulerId: harvester-CERN_central_A
2019-07-10 01:23:07|143203|pilot.py       | maxjobrec: 20
2019-07-10 01:23:07|143203|pilot.py       | maxNumberOfRecoveryAttempts: 15
2019-07-10 01:23:07|143203|pilot.py       | pilotId: https://aipanda184.cern.ch/condor_logs_2/19-07-
09_22/grid.2148787.16.out
2019-07-10 01:23:07|143203|pilot.py       | pshttpurl: https://pandaserver.cern.ch
2019-07-10 01:23:07|143203|pilot.py       | psport: 25443
2019-07-10 01:23:07|143203|pilot.py       | queuename: BNL_PROD_UCORE
2019-07-10 01:23:07|143203|pilot.py       | rmwkdir: None
2019-07-10 01:23:07|143203|pilot.py       | sitename: BNL_PROD_UCORE
2019-07-10 01:23:07|143203|pilot.py       | stageinretry: 2
2019-07-10 01:23:07|143203|pilot.py       | stageoutretry: 2
2019-07-10 01:23:07|143203|pilot.py       | uflag: managed
2019-07-10 01:23:07|143203|pilot.py       | workdir: /tmp/atlas_jBm5m8QO/Panda_Pilot_143203_1562721787
2019-07-10 01:23:07|143203|pilot.py       | logFileDir:
2019-07-10 01:23:07|143203|pilot.py       | .....................................................................
2019-07-10 01:23:08|143203|pUtil.py      | getSiteInformation: got experiment=ATLAS
2019-07-10 01:23:08|143203|ATLASExperim| Executing command: .
/cvmfs/atlas.cern.ch/repo/ew/arc/client/latest/slc6/x86_64/setup.sh;arcproxy -i vomsACvalidityLeft
2019-07-10 01:23:08|143203|ATLASExperim| ec=0 output=279137
2019-07-10 01:23:08|143203|ATLASExperim| Voms proxy verified (279137s)
2019-07-10 01:23:08|143203|ATLASExperim| Voms proxy verified using arcproxy
2019-07-10 01:23:08|143203|pilot.py       | Collecting WN info from: /tmp/atlas_jBm5m8QO
2019-07-10 01:23:08|143203|pilot.py       | Got max memory limit: 64000 MB (from queuedata)
2019-07-10 01:23:08|143203|pilot.py       | Executing command: ulimit -a
2019-07-10 01:23:08|143203|pilot.py       |
core file size          (blocks, -c) 4194303
data seg size           (kbytes, -d) unlimited
scheduling priority             (-e) 0
file size               (blocks, -f) unlimited
pending signals                 (-i) 514006
```

Current usage

17000 json requests a day

- 6.5 (+3%) user queries a day
- From 1 to 626 pages a day per user
- 1110 monthly active users
- 342 daily users

Is a primary tool ATLAS wide for shifters, experts and ADC in general

Architecture

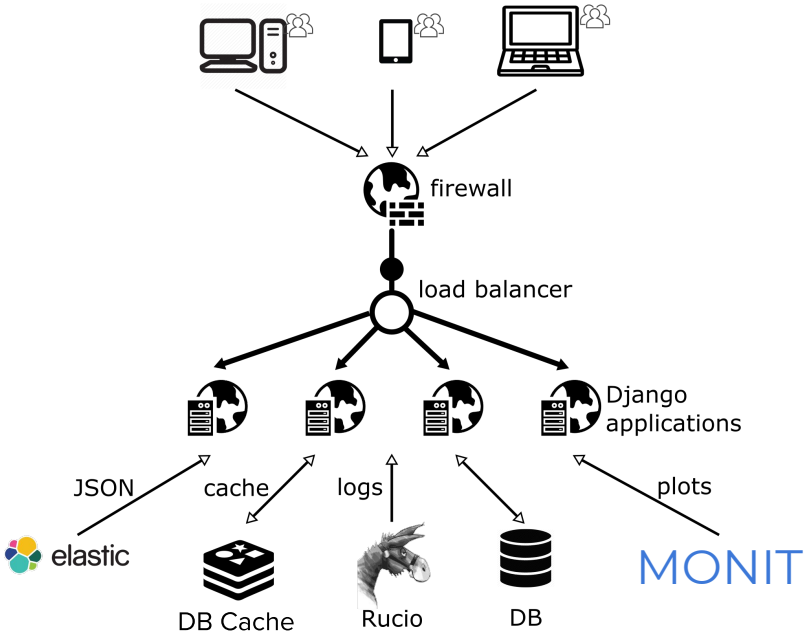

Data-flow diagram of the BigPanDA monitor

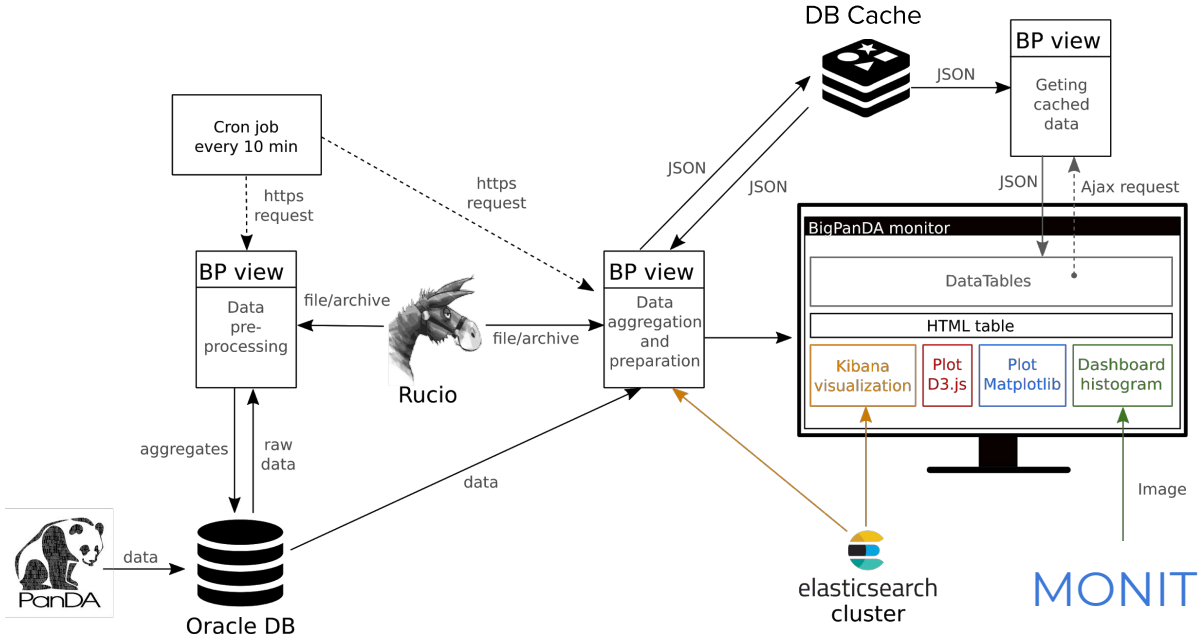

Self monitoring

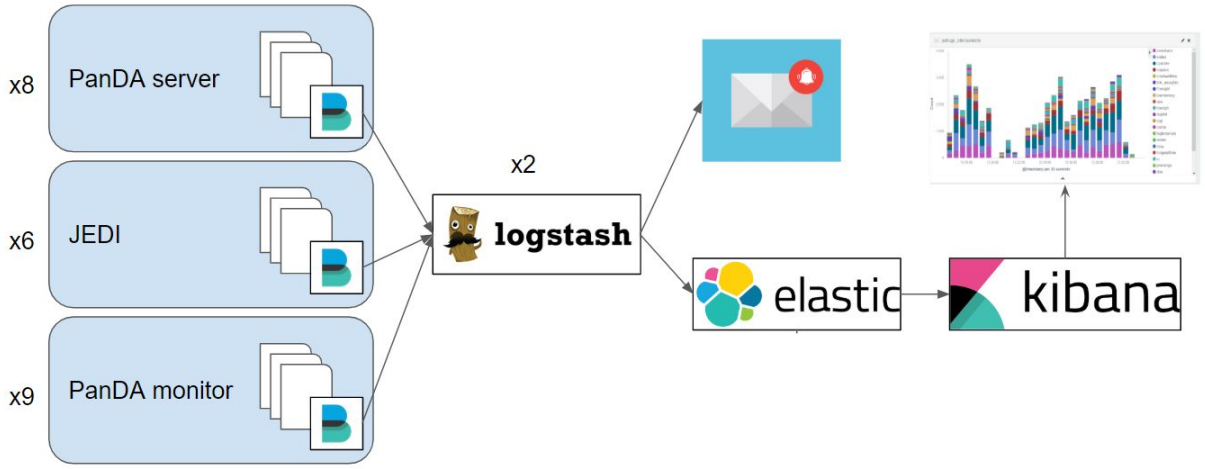

Operation Intelligence

- A lot of routine work is required for exposing and digging into problems
- We develop a mechanization such of work
  - BigPanDA attempts to provide information in the user friendly, quickly accessible way
  - OI attempts to process this information as a user would do

Operation Intelligence

- A practical case is to provide reasons, at particular level, why a number or jobs failed in a task
- A prototype:
  - Builds online failure model for a task (less than a second to train)
  - Use job definition as set features
  - Extracts from the model most important factors (and its combinations) which leads to failing jobs at particular conditions
  - Provide these factors as an outcome
- A computation engine separated from monitoring itself
- Now we work on wrapping this prototype into MVP

Forecasting

- GRID computing is a big queue
  - When is my order to get served?
  - When I do receive needed result?
  - How much can I order in principle to get result in reasonable time?
- Campaign is the ProdSys2 object which unites large computing activity (10^{10} events) and involves different parties (physicists, managers, shifters)
- A “Hot” model was developed
- Currently it is getting to production

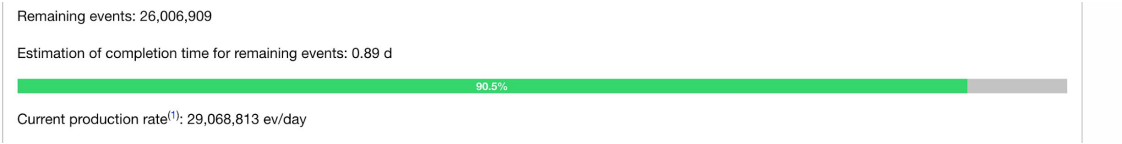

Usability R&D

Aims:

- Raise UP level of BigPanDA monitoring usability
- Build the whole system at one presentation technology stack
- Make development easy, from well defined bricks
- Technology assessment is finishing (Tatiana is leading this effort)

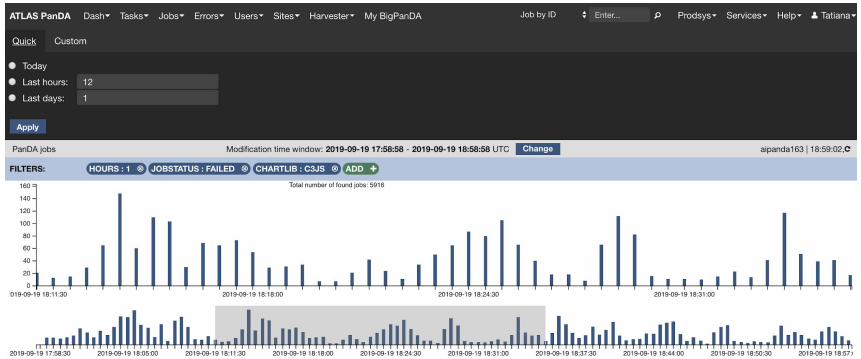

BelleII DDM

Introduction

Reference talks: Paul (https://indico.bnl.gov/event/6124/)

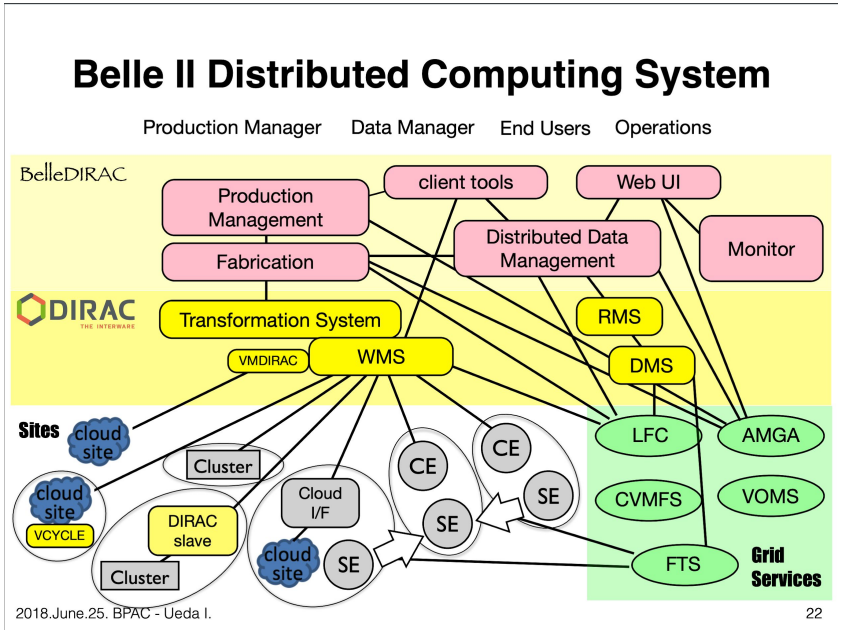

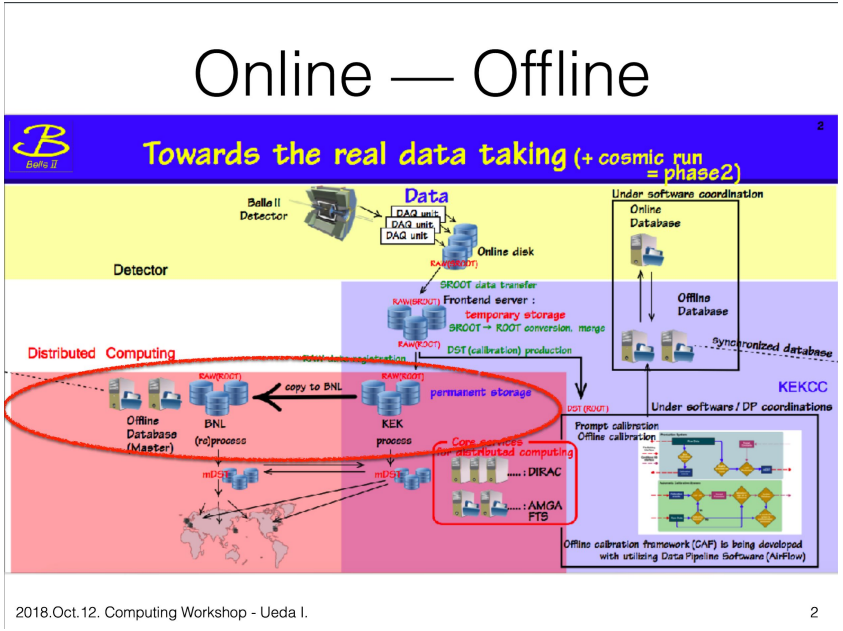

Pictures taken from Ueda I talks on B2 meetings

Distributed computing

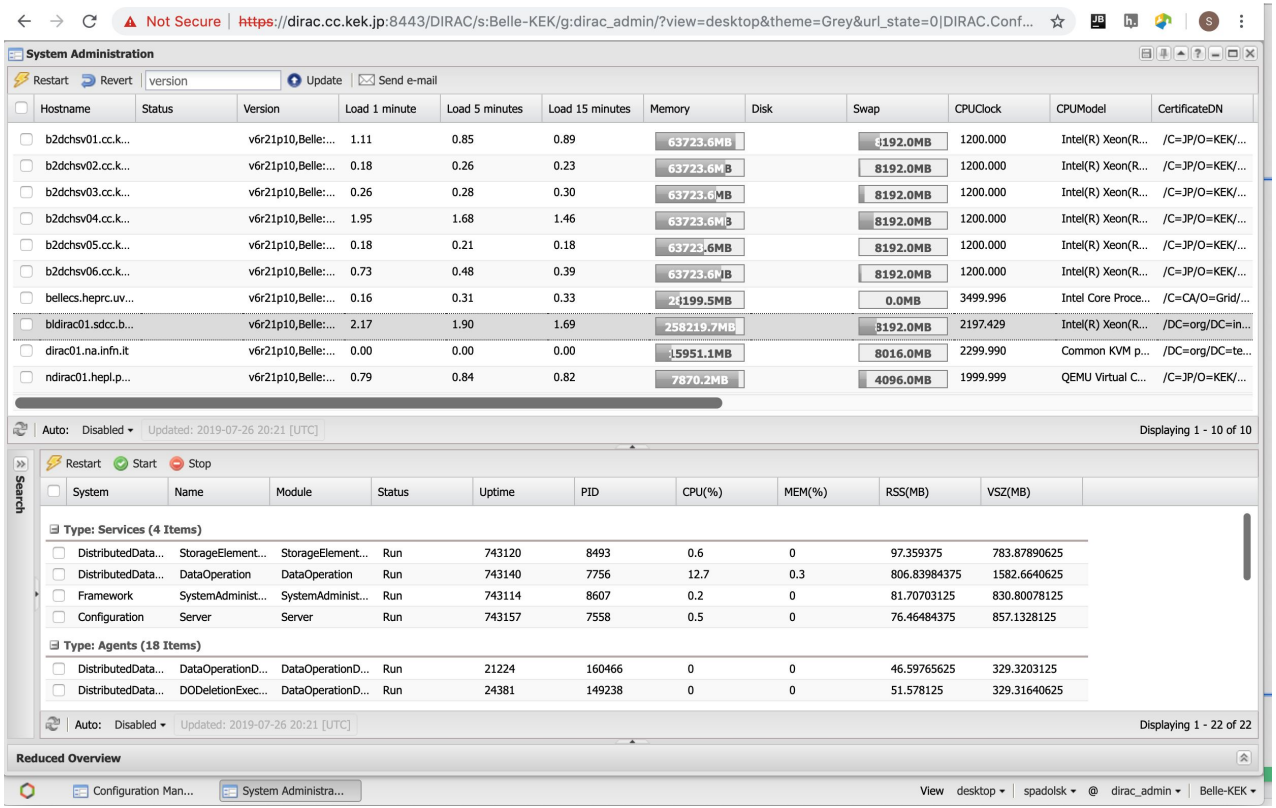

Distributed Data Management

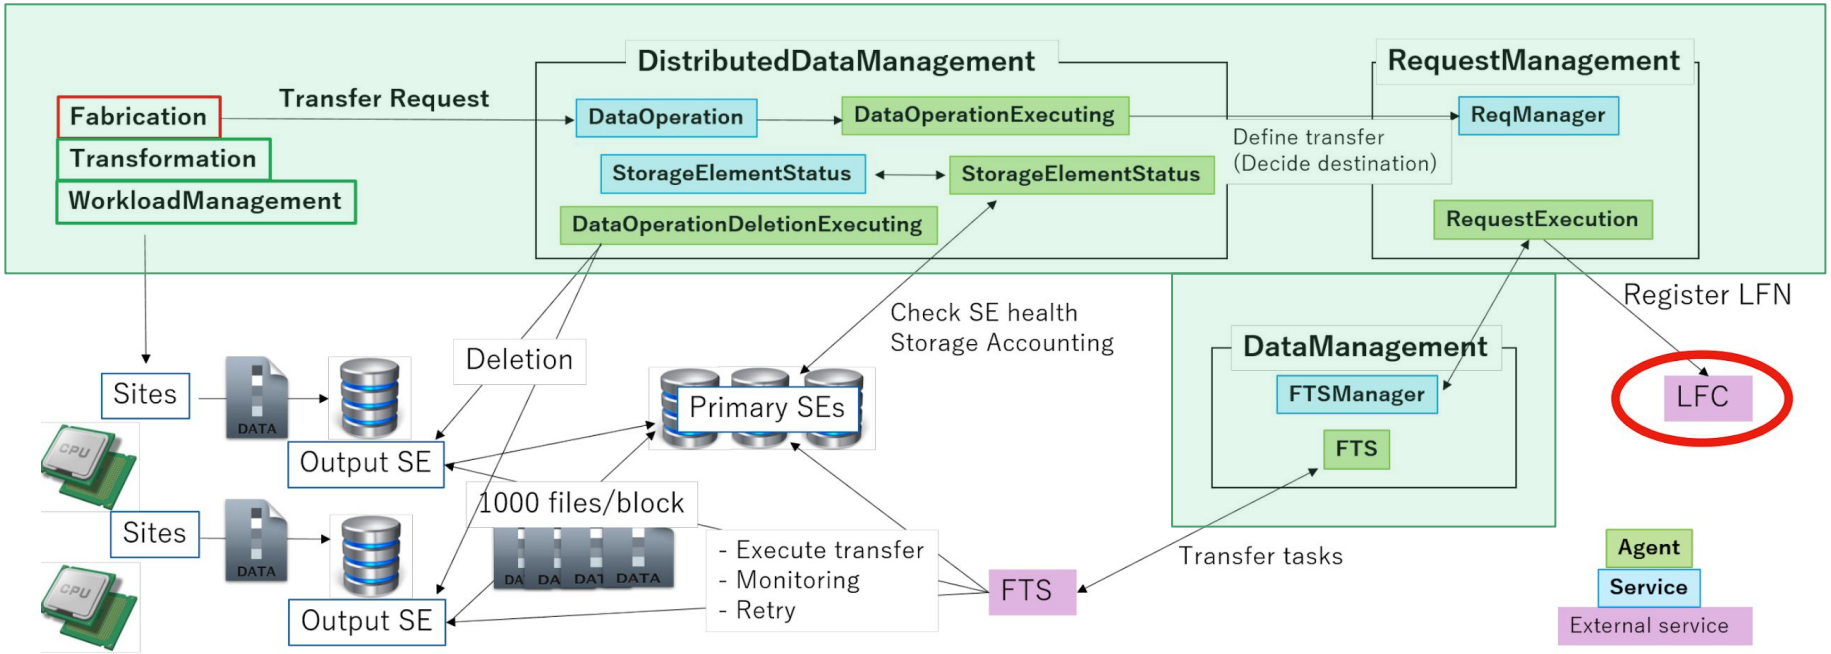

Initial picture taken from Paul introductory talk

Distributed Data Management Scale

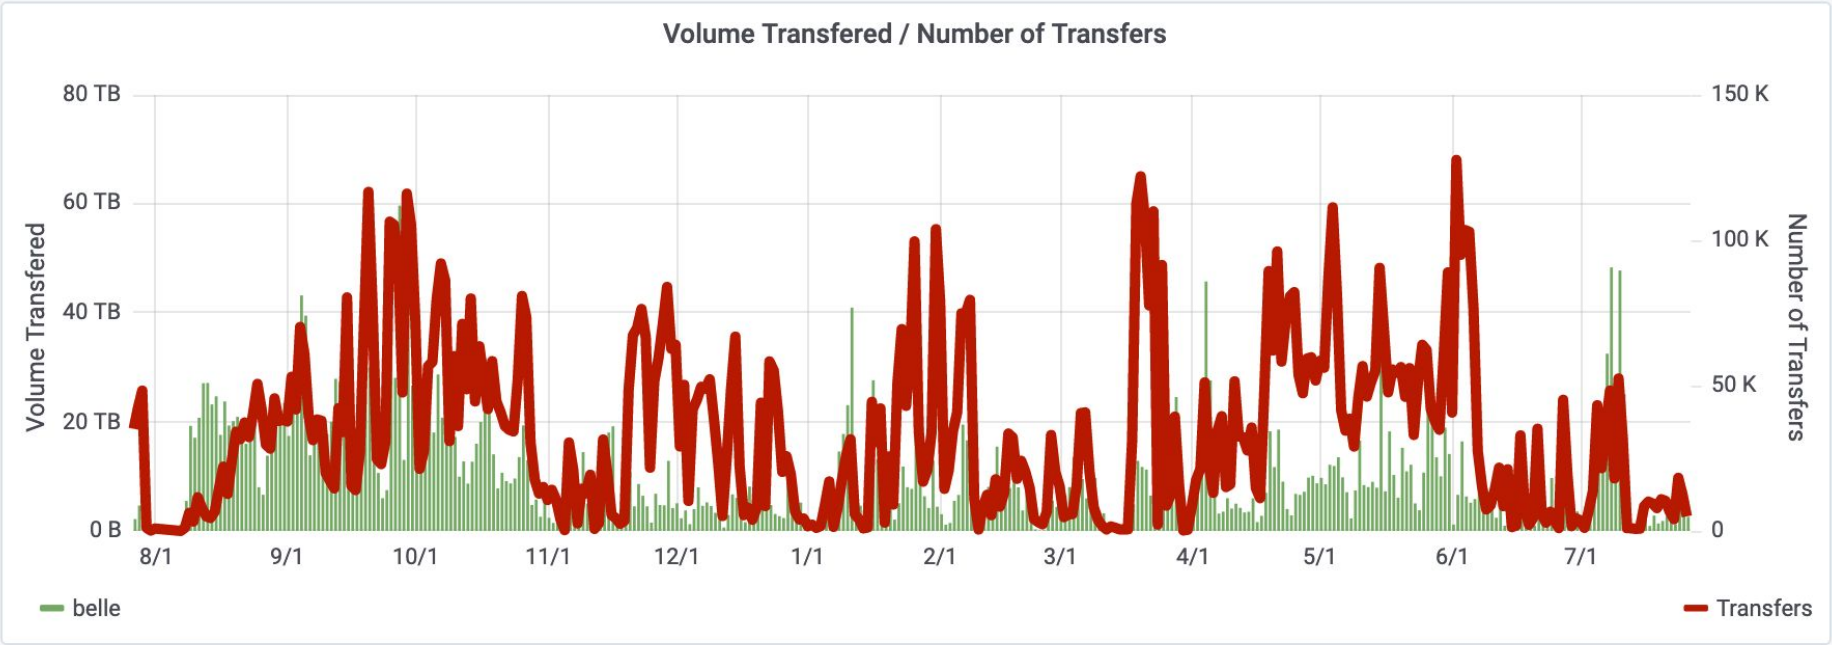

Daily transfers

BNL time Developments

- Data mover scripts (PNNL->BNL)
- Functional tests subsystem
- Tider integration with DIRAC Resource system
- Advanced transfers load balancer. Developed to prevent possible data transfers stucks. Implements accurate submissions to the FTS served respecting:
  - SE current performance
  - Links quality
  - Recent experience of completed operations
  - Current queue state
  - Operations importance
- Deletion at scales (bulk operations, parallel processing)
- SQL queries analysis and optimization
- Simple DDM monitoring

\+ Operational support

Operational experience: 2 DDM (and whole DIRAC) down due to server disk spaces exhausting and one due to RMS system failure, one failed deployment. Successful raw files distribution during data taking