Healthy GPU use together

Roel Aaij, Emily Kooistra

Inventory

- 4x V100
- 20x L40S

Goal

- High utilisation of the GPUs
- High job throughput
- More science per user
- More science across the board

GPUs are popular!

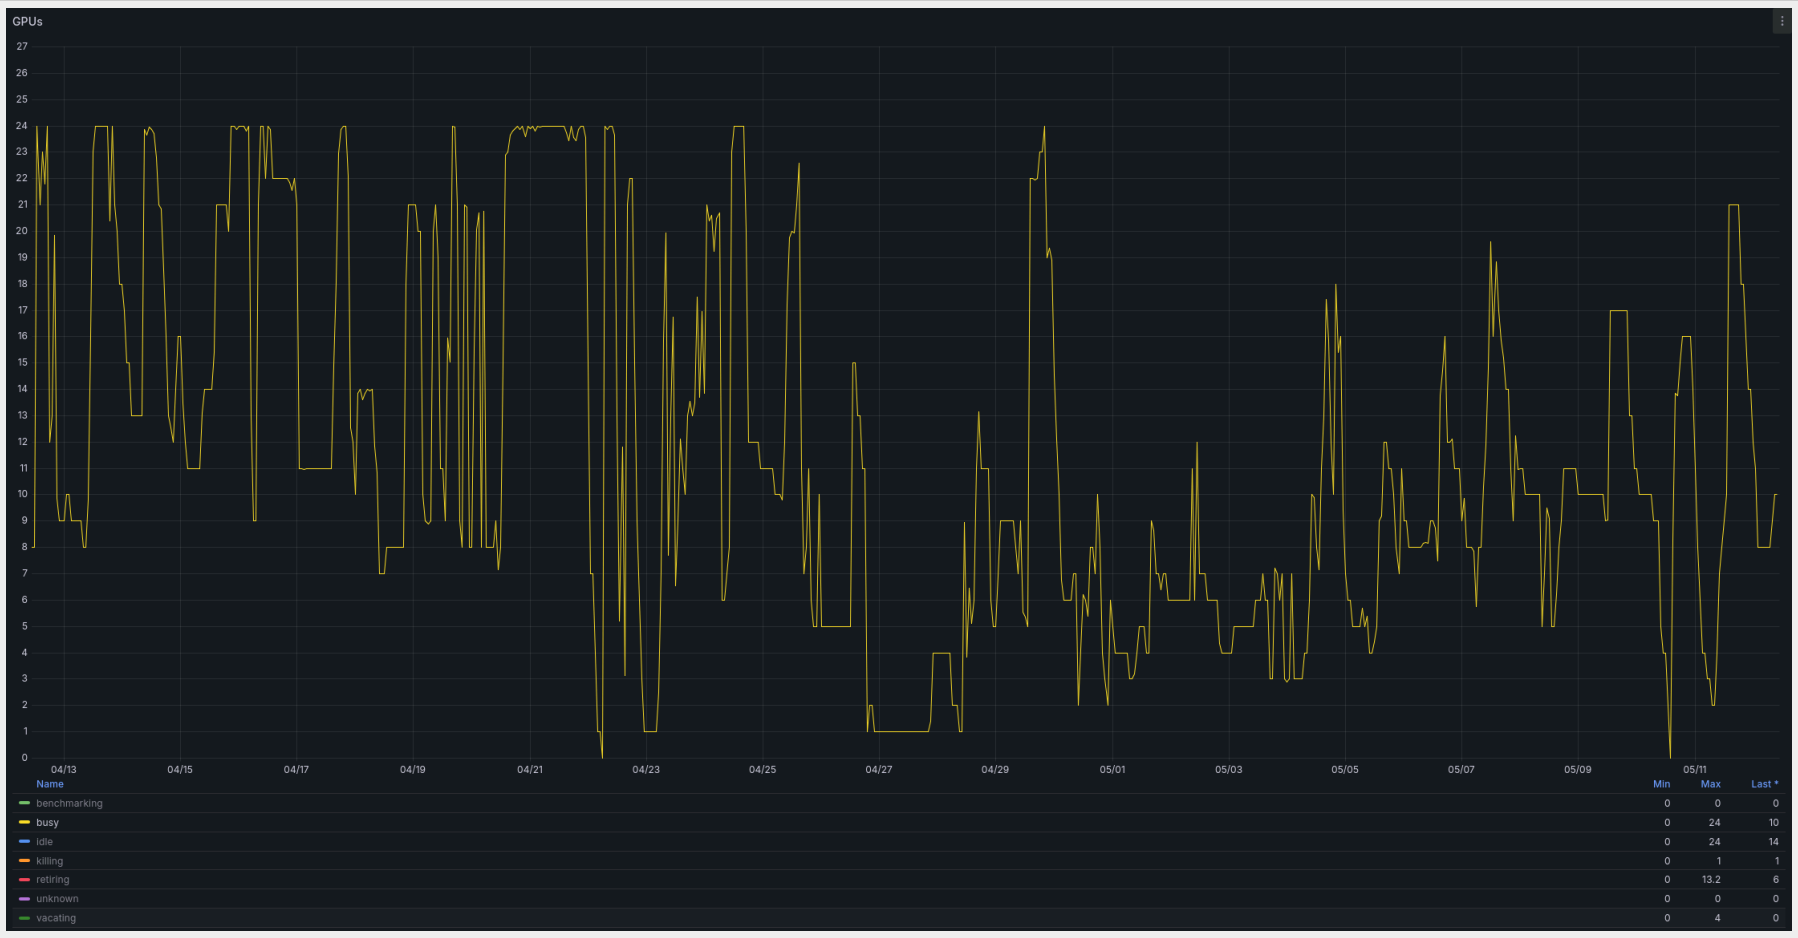

Recurring issues

- Low GPU compute utilisation
- Low GPU memory footprint
- Suboptimal datamovement

Compute utilisation

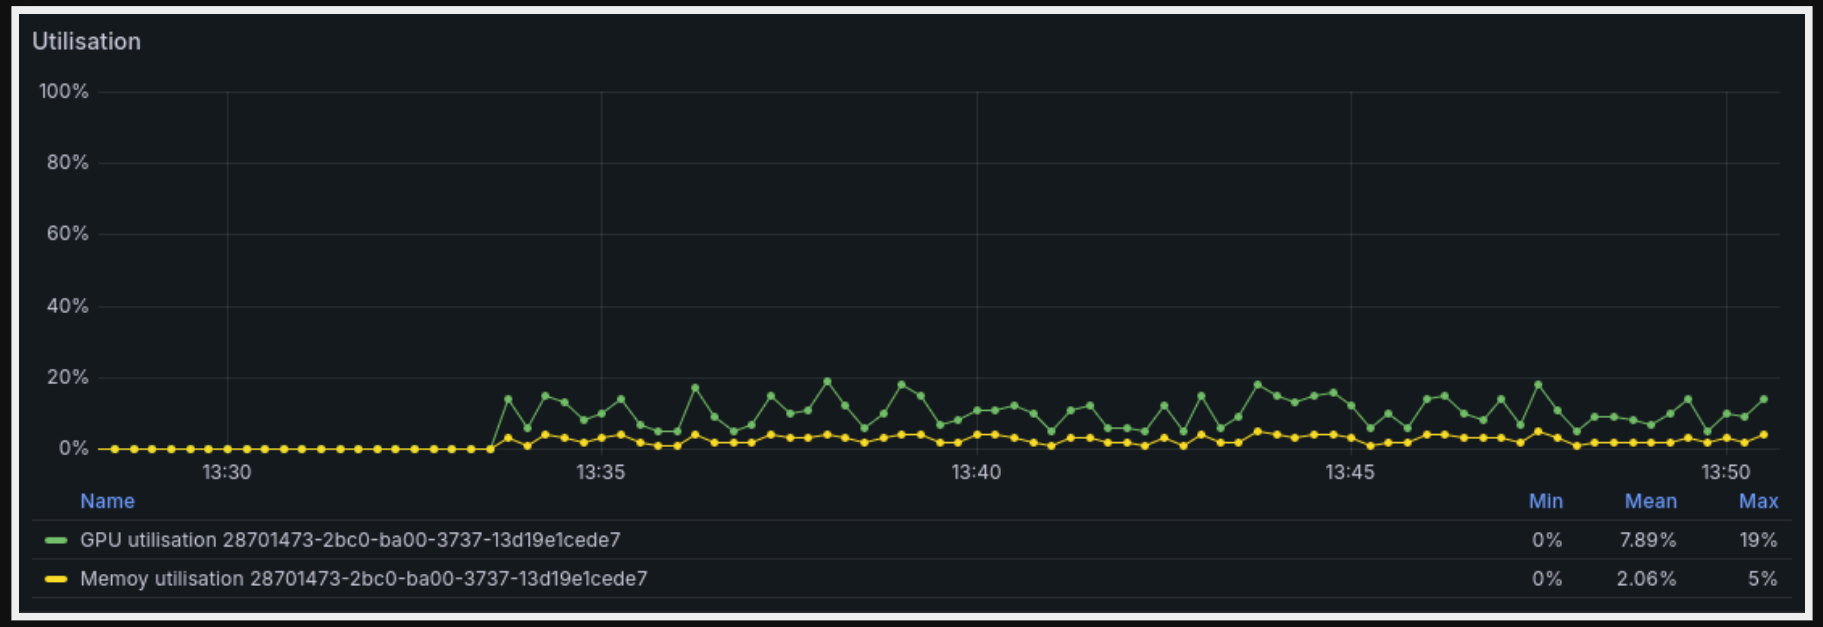

Compute utilisation, what to do?

- Preprocess data in separate jobs
- Use a DAG if this is needed for every submission
- Run multiple GPU workloads in a single job
- Have a look at optimisation tips for e.g. PyTorch, JAX or Tensorflow

Memory utilisation

Memory utilisation, what to do?

- *Low* memory utilisation is not really a problem in itself
- Allows multiple workloads per GPU per job, which gives higher throughput
- High memory usage can lead to lower compute utilisation
- Optimise memory allocations in e.g. PyTorch, JAX or Tensorflow

Data Movement

- We see a lot of load on /data from GPU jobs
- Unnecessary load affects everybody
- dcache is underutilised
- Files bigger then 1GB -> dcache
- Avoid small files: bundle output as archives
- Use $TMPDIR whenever possible
- Avoid writing logs (directly) to storage

IO /data

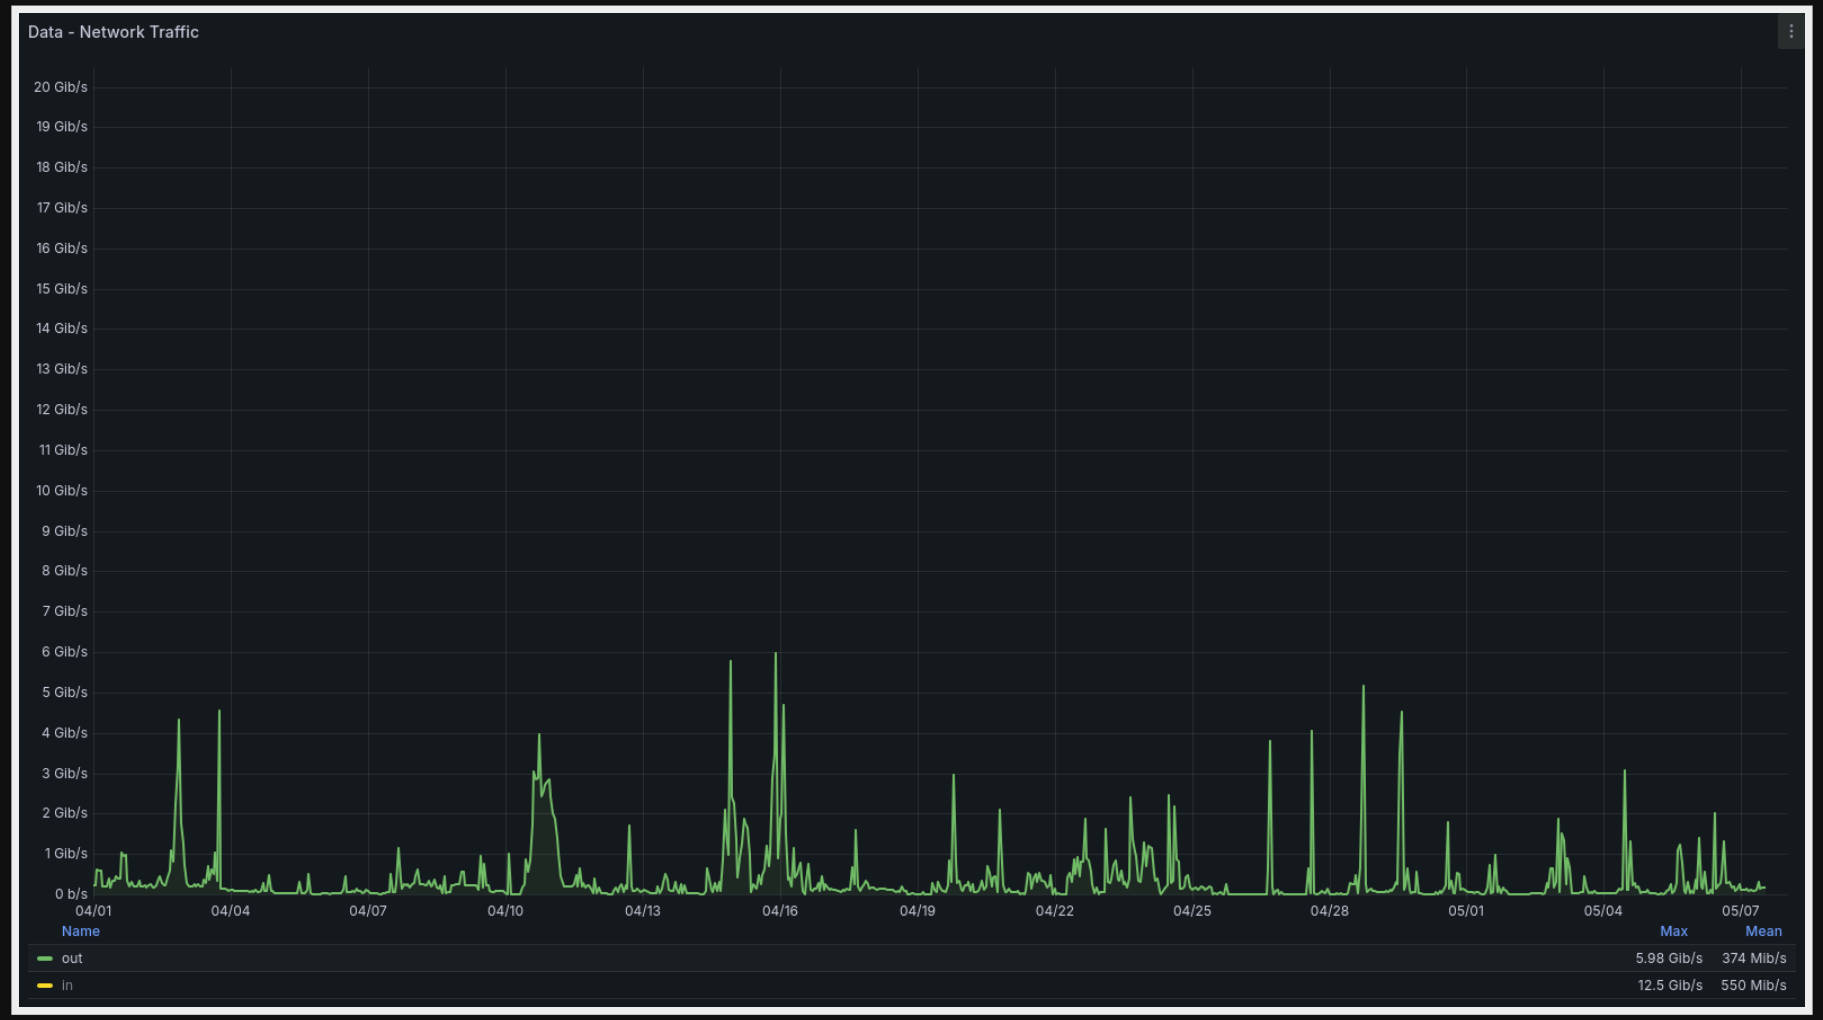
Data - Network Traffic
20 Gib/s
19 Gib/s
18 Gib/s
17 Gib/s
16 Gib/s
15 Gib/s
14 Gib/s
13 Gib/s
12 Gib/s
11 Gib/s
10 Gib/s
9 Gib/s
8 Gib/s
7 Gib/s
6 Gib/s
5 Gib/s
4 Gib/s
3 Gib/s
2 Gib/s
1 Gib/s
0 b/s
04/01
04/04
04/07
04/10
04/13
04/16
04/19
04/22
04/25
04/28
05/01
05/04
05/07
Name
Max
Mean
out
5.98 Gib/s
374 Mib/s
in
12.5 Gib/s
550 Mib/s

IO /dCache

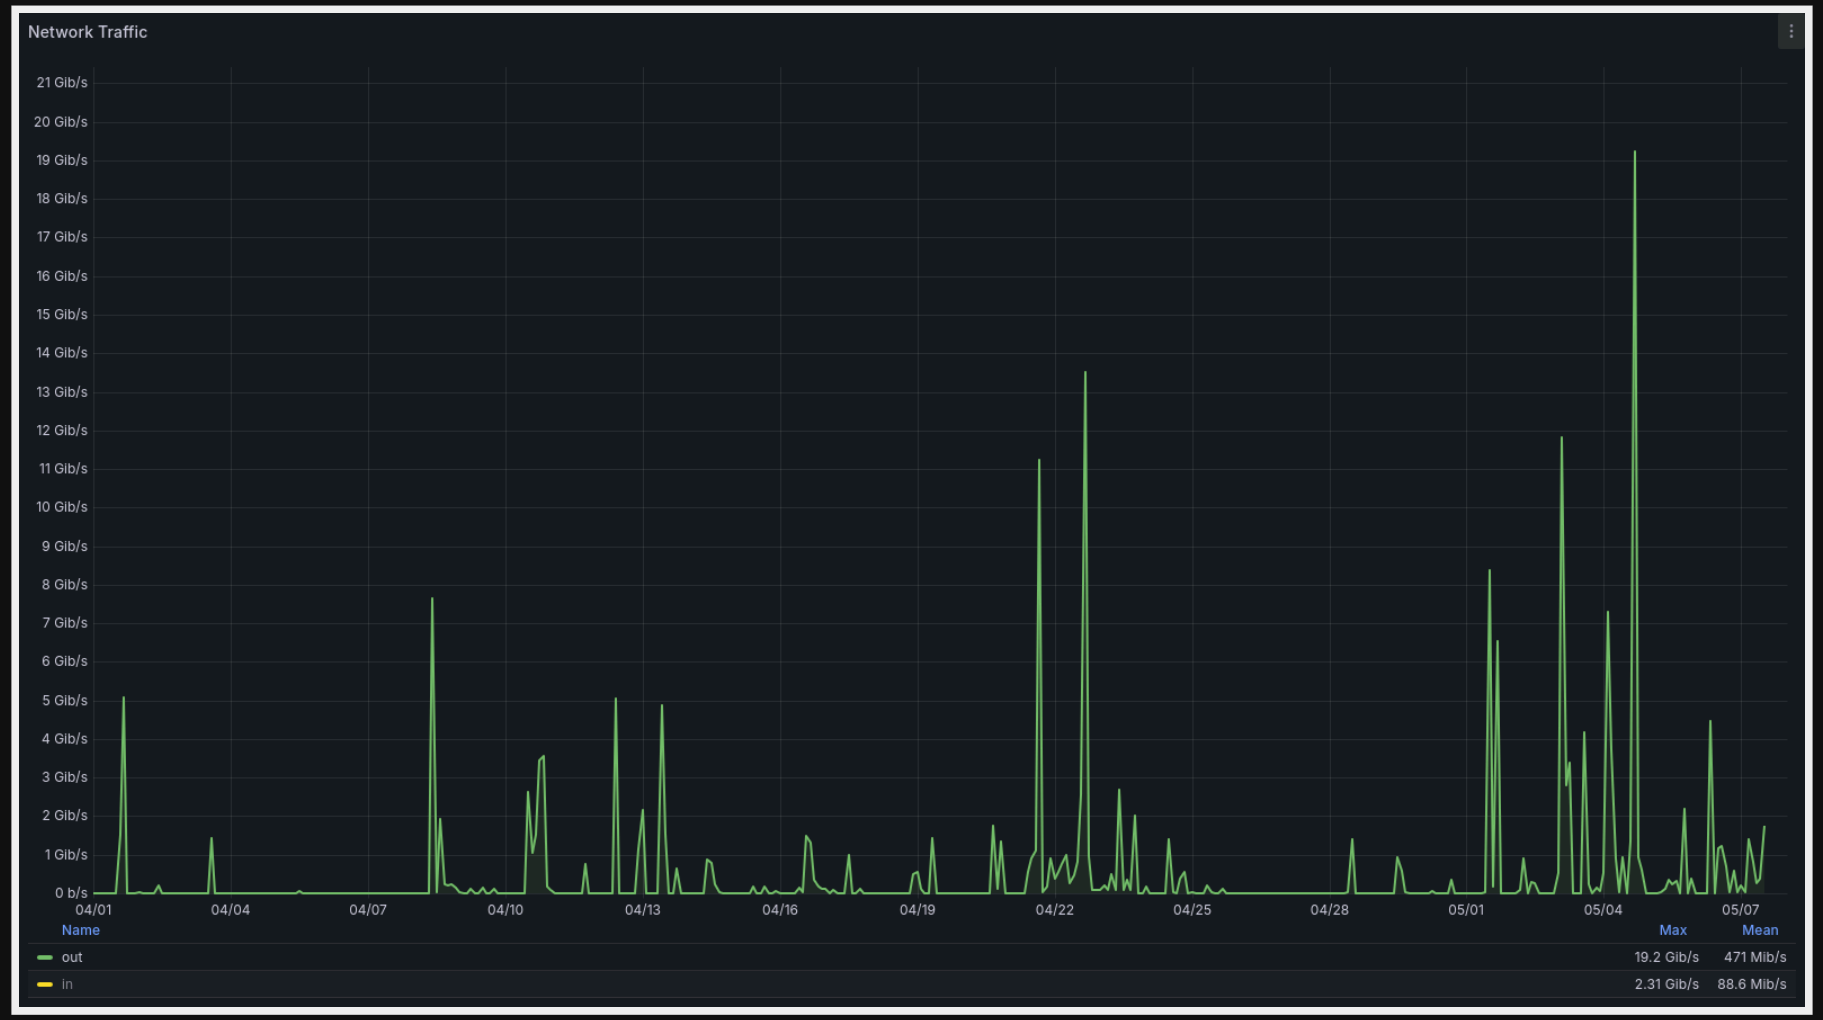
Network Traffic
21 Gib/s
20 Gib/s
19 Gib/s
18 Gib/s
17 Gib/s
16 Gib/s
15 Gib/s
14 Gib/s
13 Gib/s
12 Gib/s
11 Gib/s
10 Gib/s
9 Gib/s
8 Gib/s
7 Gib/s
6 Gib/s
5 Gib/s
4 Gib/s
3 Gib/s
2 Gib/s
1 Gib/s
0 b/s
04/01
04/04
04/07
04/10
04/13
04/16
04/19
04/22
04/25
04/28
05/01
05/04
05/07
Name
Max
Mean
out
19.2 Gib/s
471 Mib/s
in
2.31 Gib/s
88.6 Mib/s

Storage usage

- Huge increase of data volume on /data
- On average 70% of all data is never used again after 7 days
- Many groups are even in the 80-90% range
- Help your group members: clean up or archive

Talk to us

- Workflows with data bigger then 1TB?
- GPU workflows
- GPU optimisation
- Data processing pipelines
- Accessing data from CERN/other experiments
- https://go.nikhef.nl/stbc-slot-usage