

WHO ATE MY IOPS? IDENTIFYING HEAVY HITTERS USING JOBSTATS

James Beal, Principal Systems Administrator

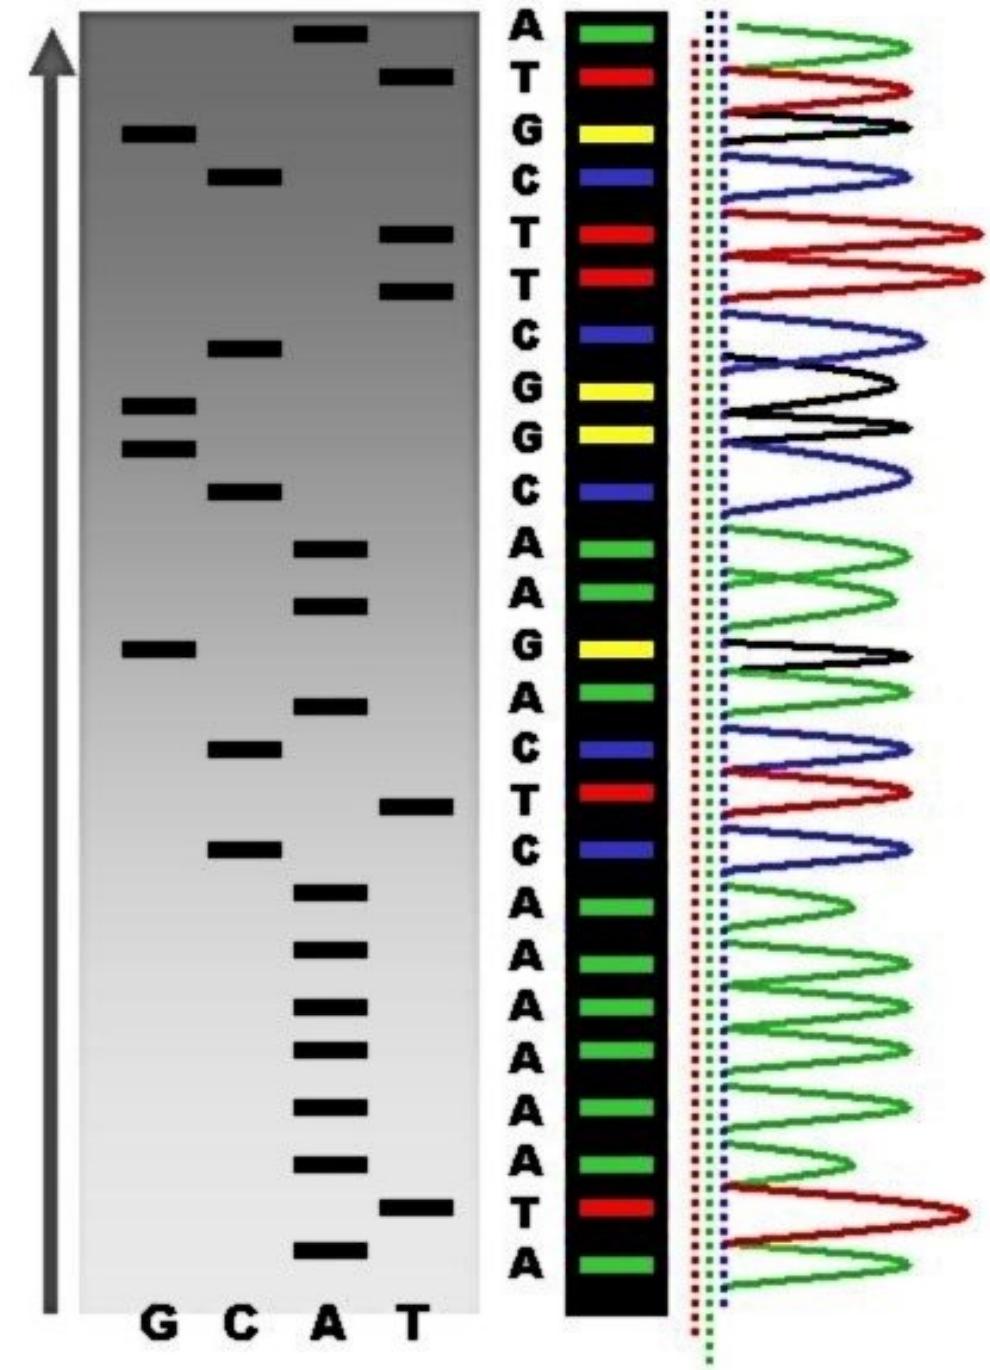
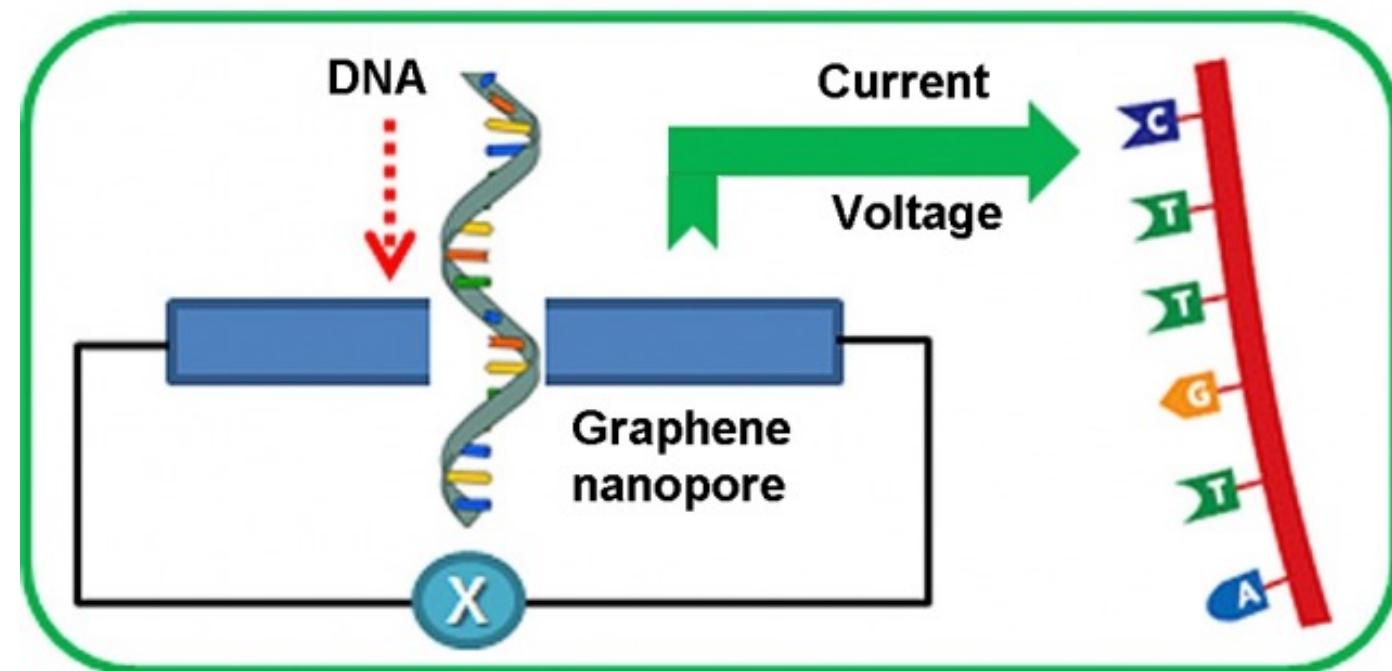
Dave Holland, Principal Systems Administrator

V09 2022/11/30

INTRODUCTION



DNA SEQUENCING



THE PAST

Capillary sequencer (very slow)

- ***2 Mega (2×10^6) bases a day.***
- Read length of about 1000 base pairs
- A single human genome is about 6.4 gigabases



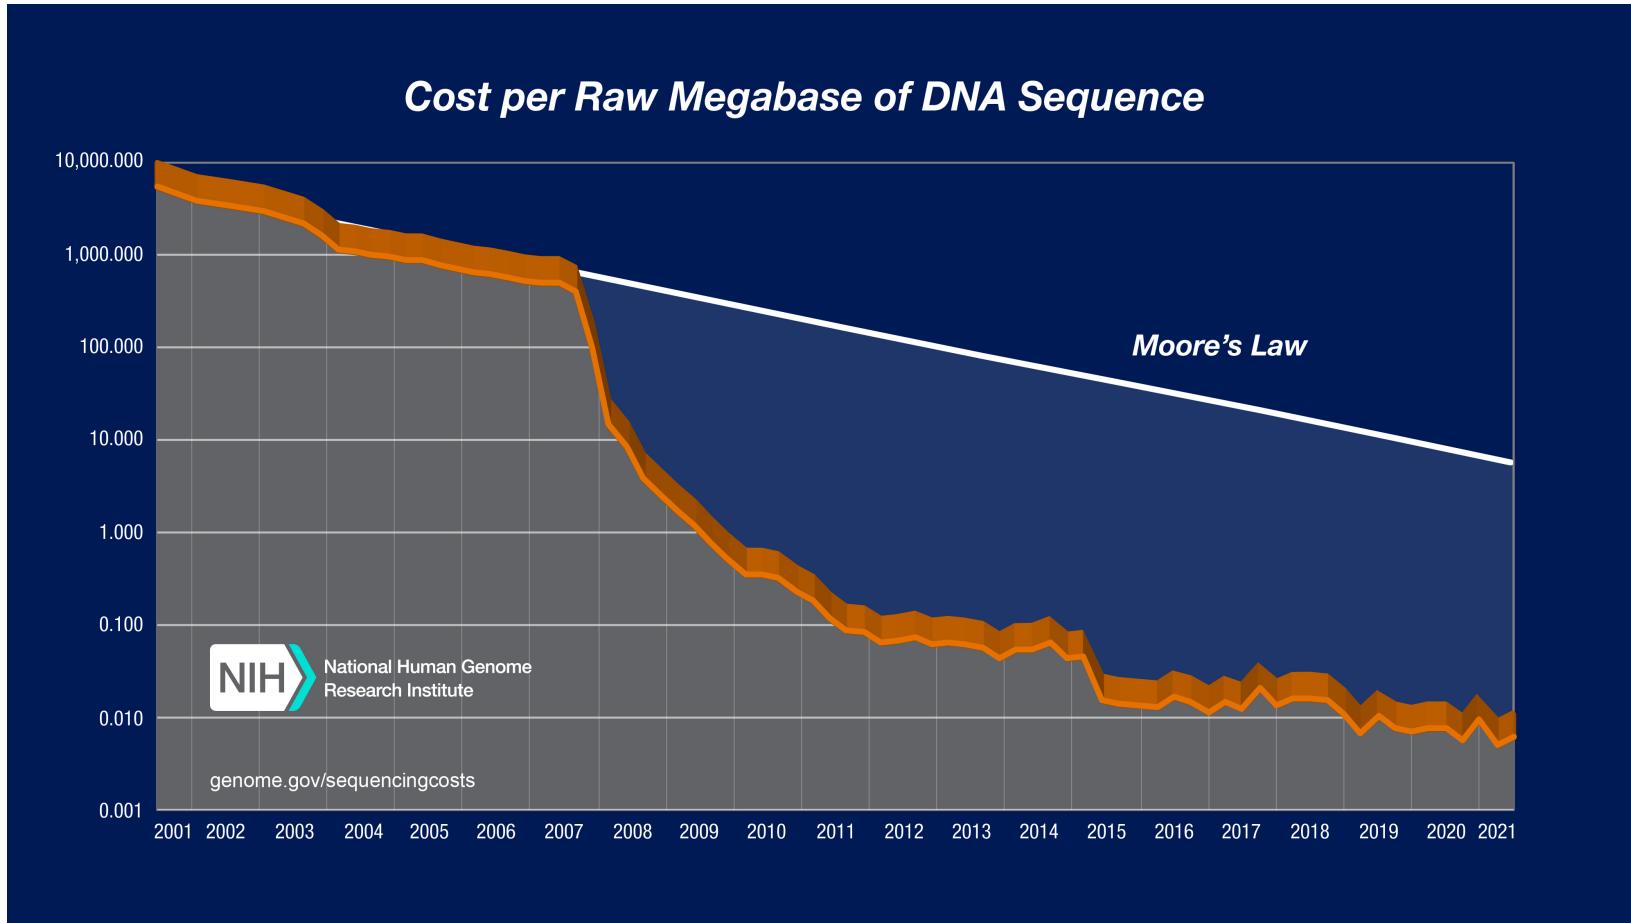
THE PRESENT

SBS (sequencing by synthesis, Novaseq X).

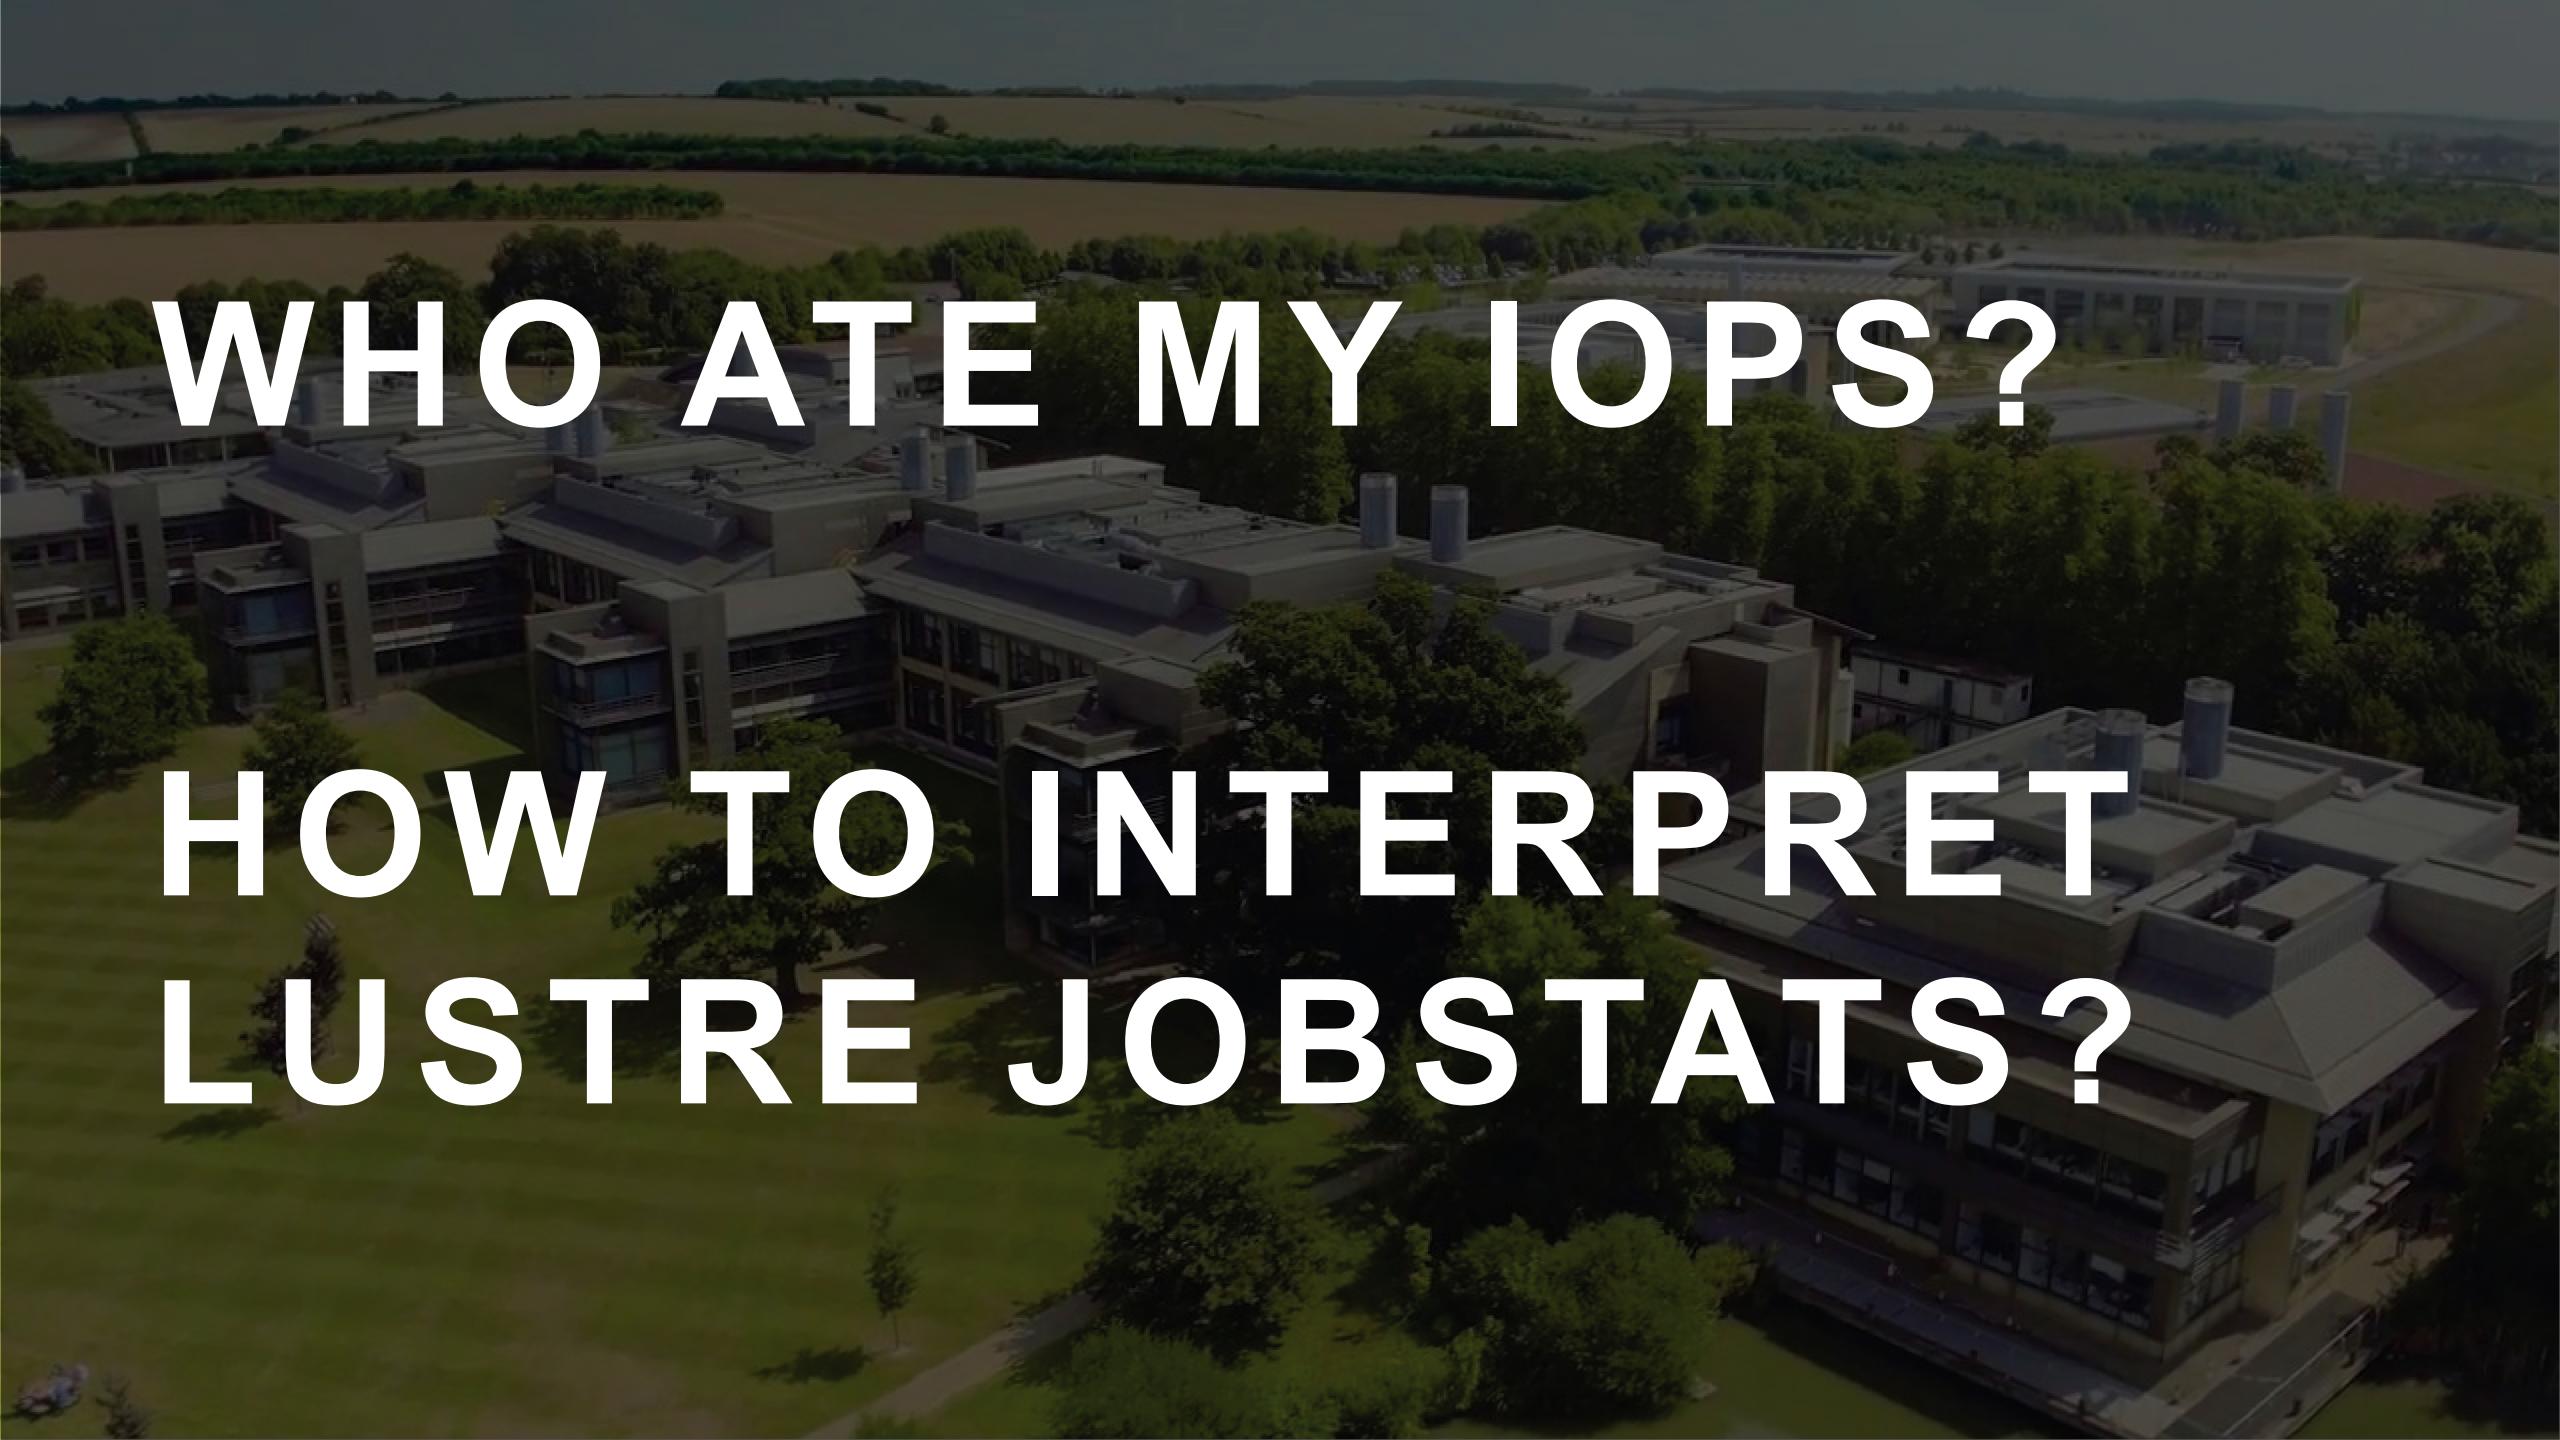
- ***8 Terabases (8×10^{12}) a day per sequencer.***
- 2.5 bytes a base or 20TB per sequencer per day working space.
- 150 base pair read length
- We have 20 of the previous generation currently in production.

THE FUTURE

Cost of sequencing is falling rapidly, HUGE impact on compute and storage!



<https://www.genome.gov/about-genomics/fact-sheets/DNA-Sequencing-Costs-Data>

The background image shows an aerial view of a large, modern building complex, likely a supercomputer facility, situated in a rural area with green fields and trees. The buildings are multi-story structures with various roof types and some solar panels visible. The overall scene is bright and sunny.

WHO ATE MY IOPS?

HOW TO INTERPRET LUSTRE JOBSTATS?

REQUIREMENTS

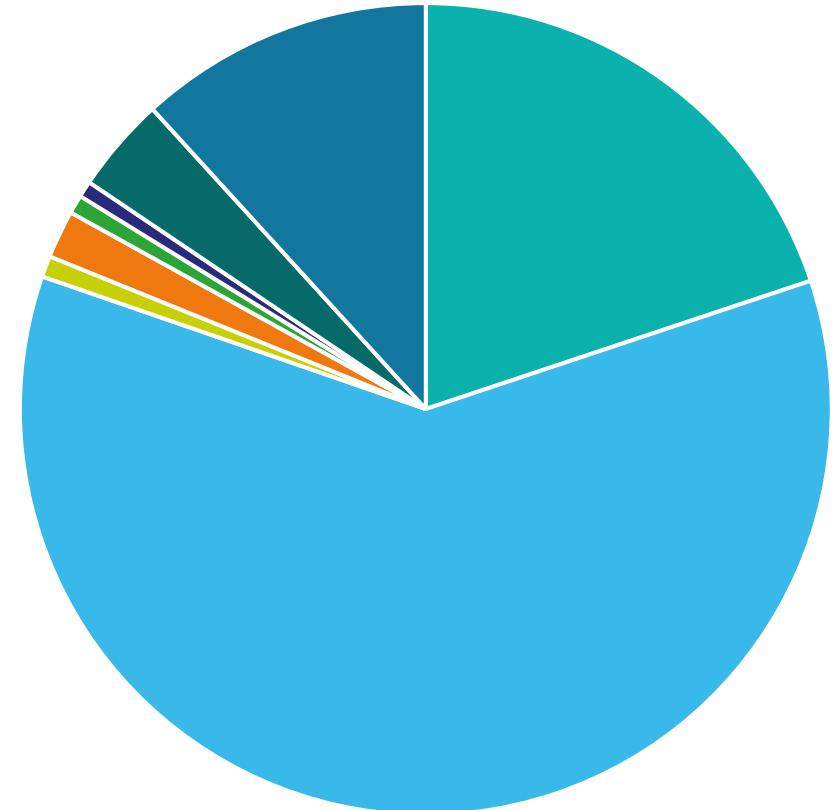
- User name
- Group name
- LSF project
- Compute cluster
- A pony



COMPUTE CLUSTERS

- A number of very similar configuration LSF clusters of different sizes. Separate clusters are used to ringfence resources.
- All* lustre filesystems are mounted everywhere.
- Lustre filesystems are not generally dedicated to a single group.

- Cancer and somatic mutation
- Main compute cluster
- General purpose and training
- Impute service
- Pathogens
- Sequencing research
- Sequencing Production
- Tree of Life



```
jb23@seqfarm3-head1:~$ df -H /lustre/*
Filesystem      Size  Used Avail Use% Mounted on
172.17.220.18@tcp:172.17.220.19@tcp:/lus17  3.1P  2.2P  886T  71% /lustre/scratch117
172.17.220.28@tcp:172.17.220.29@tcp:/lus18  3.1P  1.8P  1.3P  59% /lustre/scratch118
172.28.64.102@tcp:172.28.64.103@tcp:/lus19  2.5P  2.2P  254T  90% /lustre/scratch119
10.177.252.5@tcp:10.177.252.4@tcp:/lus20   3.1P  1.2P  1.9P  38% /lustre/scratch120
10.160.32.5@tcp:10.160.32.4@tcp:/lus23   5.3P  2.8P  2.5P  53% /lustre/scratch123
10.160.36.4@tcp:10.160.36.5@tcp:/lus24   5.3P  1.8P  3.5P  34% /lustre/scratch124
10.160.40.4@tcp:10.160.40.5@tcp:/lus25   5.3P  229T  5.1P  5% /lustre/scratch125
10.160.42.4@tcp:10.160.42.5@tcp:/lus26   5.3P  327T  5.0P  7% /lustre/scratch126
```

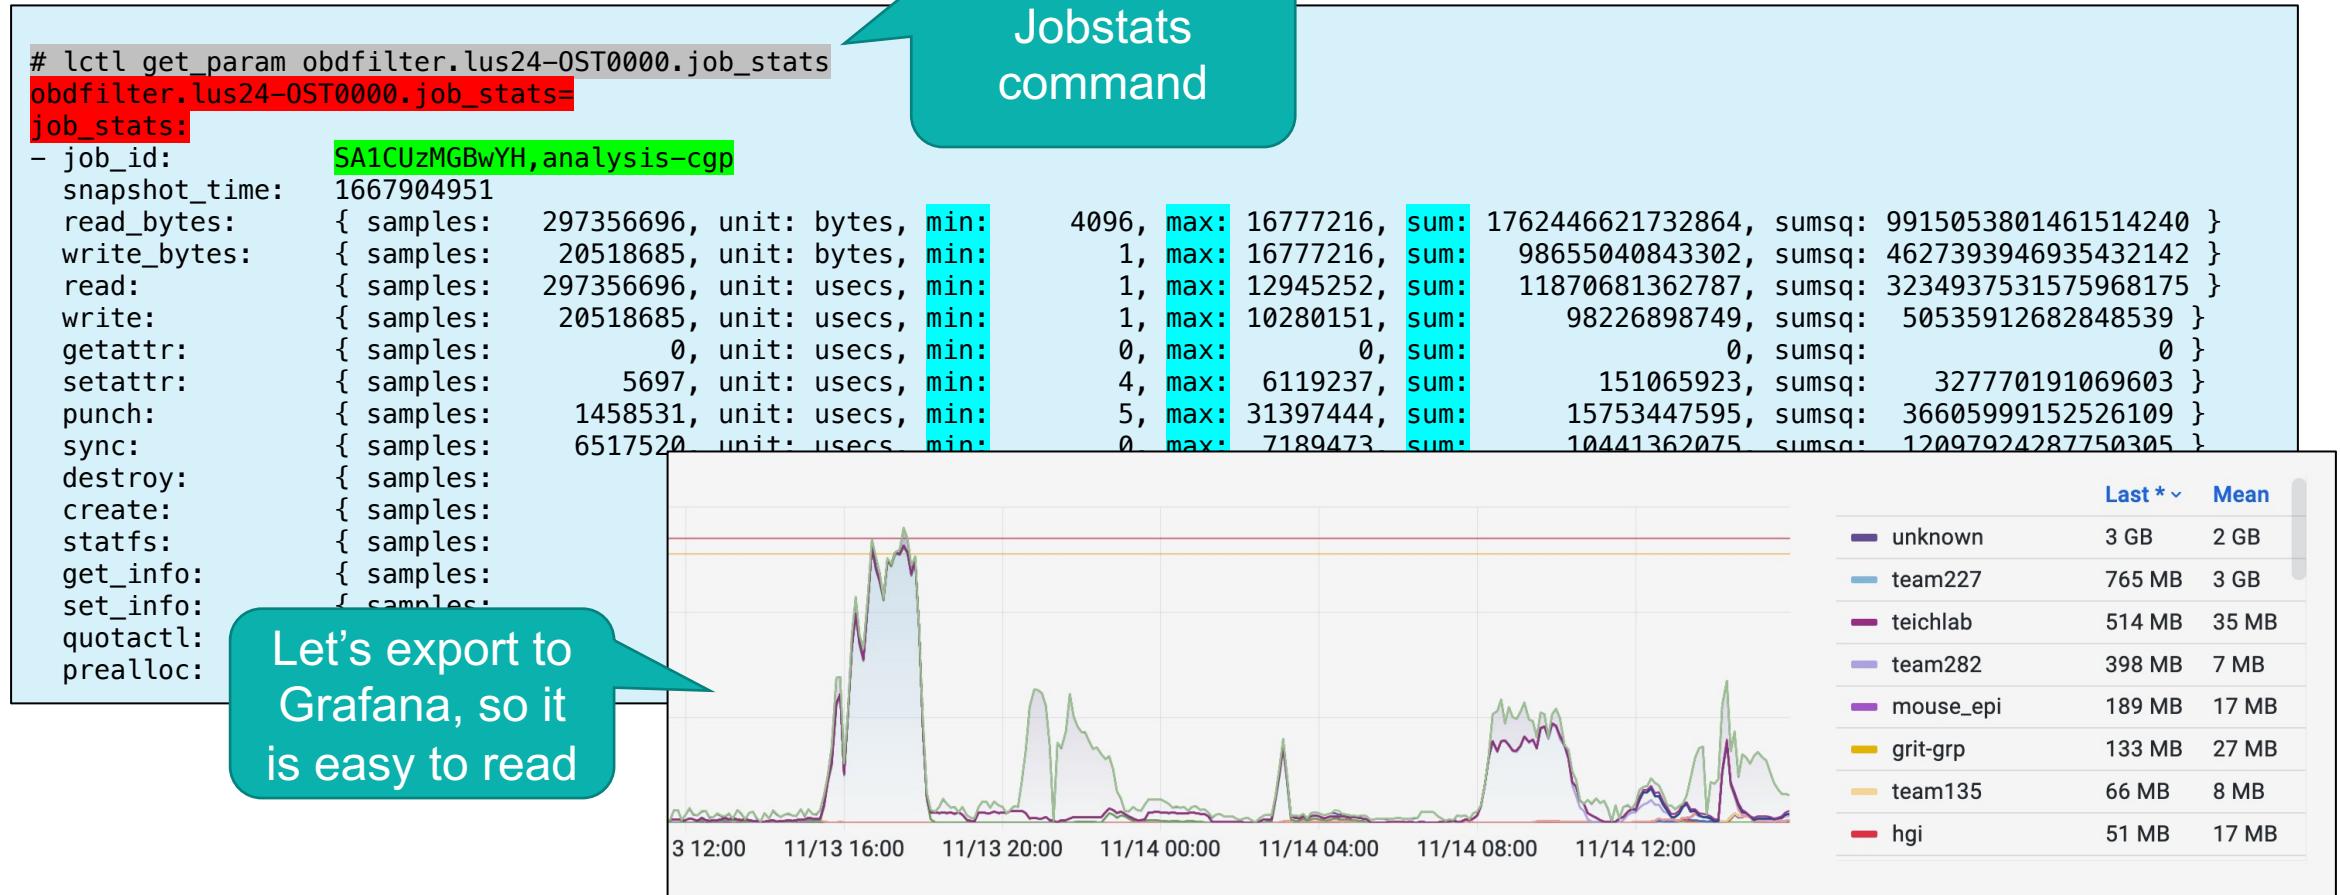
JOB STATS

- https://doc.lustre.org/lustre_manual.xhtml#jobstats
- A method of tagging I/O with additional metadata.
- Introduced in Lustre 2.3 (in 2011)
- A 32 character tag.
- The default tag is procname_uid
 - executable_name.uid
- A filesystem or a compute node can be configured to read a environment variable as the tag.



RAW JOBSTATS OUTPUT

All the information we need is in “jobstats” but is difficult to interpret



EXAMPLE 1





filesystem

lus23

type

Enter variable value

**Select
filesystem**byfarm
bygroup
bylsf_project
byusername**Lustre "spikes" are due to a Lustre server bug and should be disregarded.**internal.sanger.ac.uk/d/000000203/lustre-stats?orgId=1

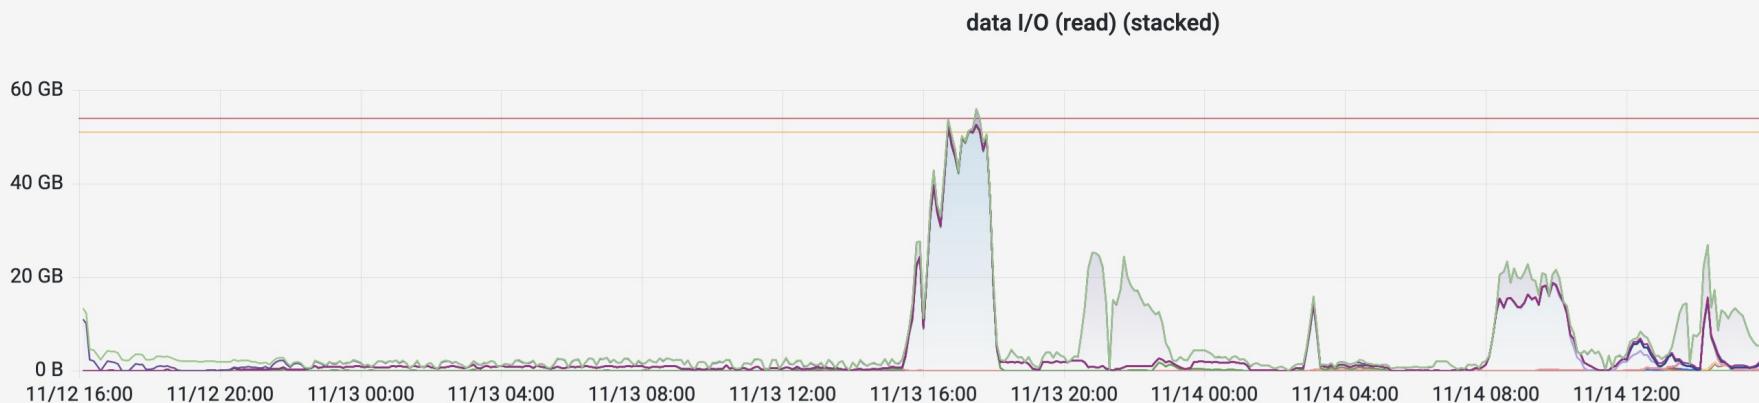
These spikes are temporary. The yellow line indicates 85%. Around these levels, the filesystem will feel slow to interactive users.



▼ Read/Write



Bytes per second



▼ Stats

Bytes per second





filesystem

lus23

type

Enter variable value

byfarm

bygroup

bylsf_project

byusername

Select how to aggregate data

NB occasional u

Compare with overall I/O showr

The red line indicates 90% of te

lustre server bug and should be disregarded.

▼ Read/Write

data I/O (read) (stacked)



Last * Mean

data I/O (write) (stacked)



Last * Mean

▼ Stats



filesystem

lus23

type

Enter variable value

byfarm

bygroup

bylsf_project

byusername

NB occasional u

Compare with overall I/O showr

The red line indicates 90% of te

ata "spikes" are due to a Lustre server bug

internal.sanger.ac.uk/d/000000203/lustre-stats?orgId=1

Observable
window

ty. The yellow line indicates 85%. Around these levels, the filesystem will feel slow.

▼ Read/Write

data I/O (read) (stacked)



Last * Mean

data I/O (write) (stacked)



Last * Mean

▼ Stats



filesystem

lus23

type

Enter variable value

byfarm

bygroup

bylsf_project

byusername

NB occasional u

Compare with overall I/O showr

The red line indicates 90% of te

ata "spikes" are due to a Lustre server bug and should be disregarded.

internal.sanger.ac.uk/d/000000203/lustre-stats?orgId=1

ty. The yellow line indicates 85%. Around these levels, the filesystem will feel slow to interactive users.

▼ Read/Write



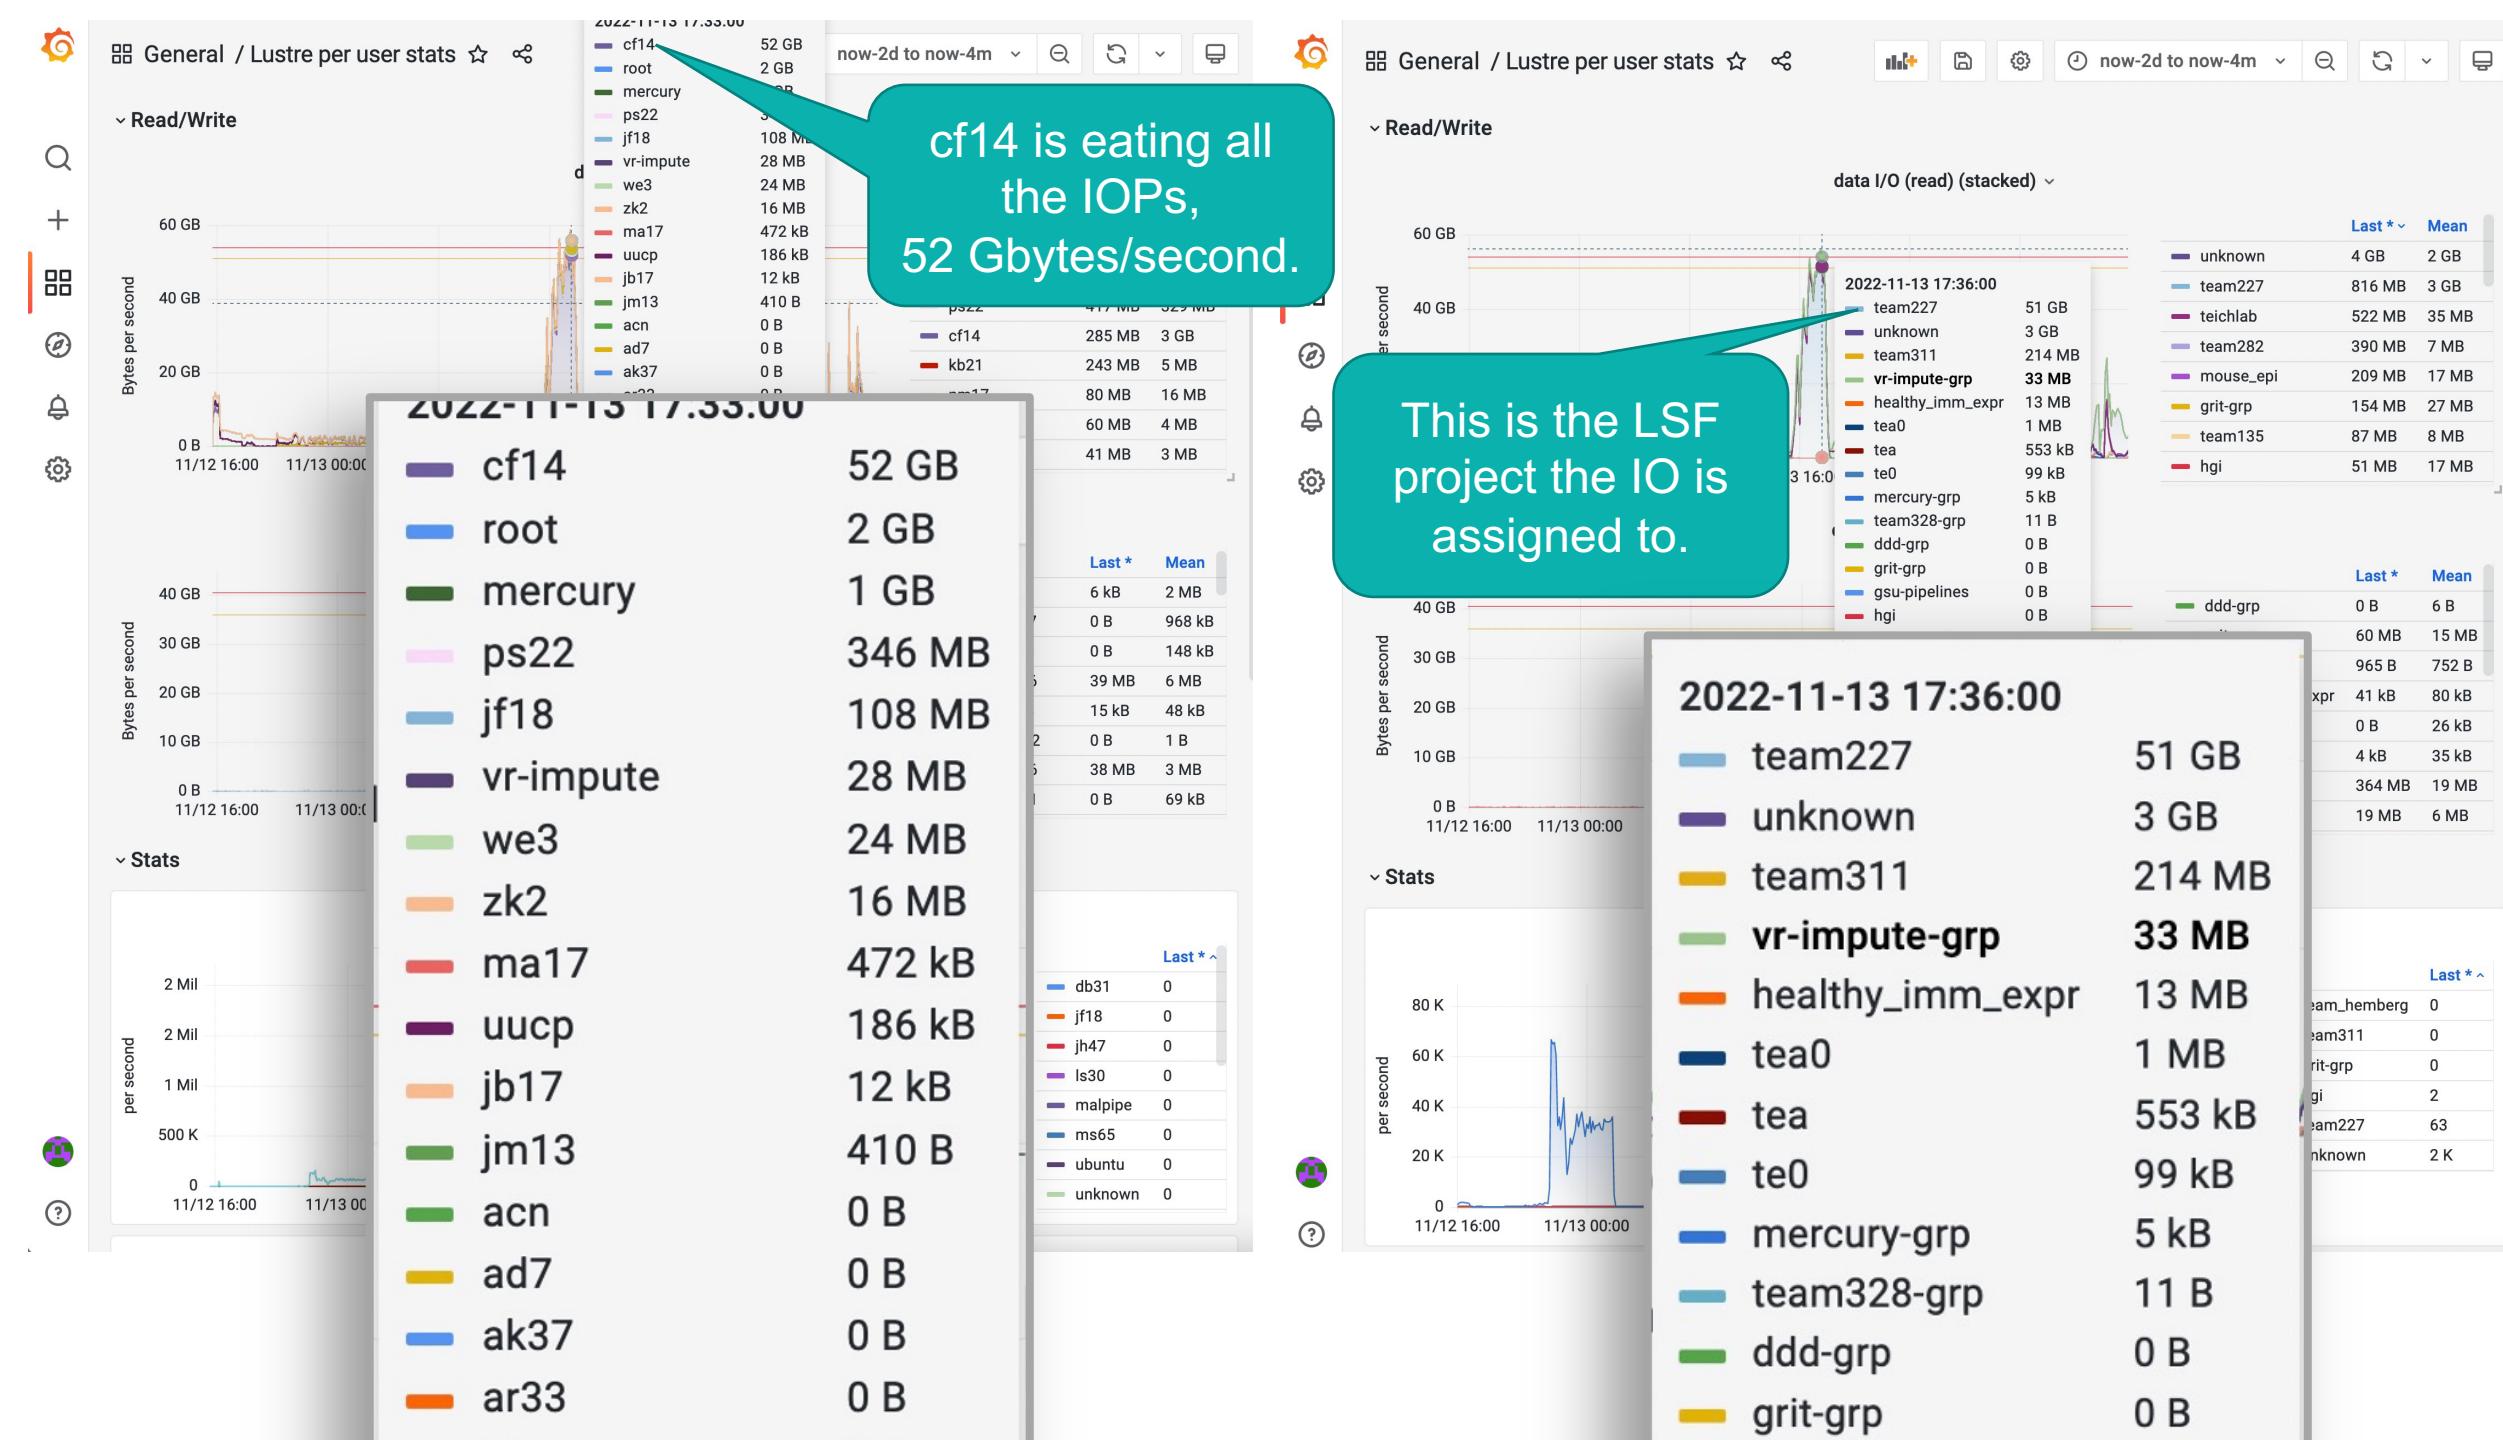
Bytes per second



▼ Stats

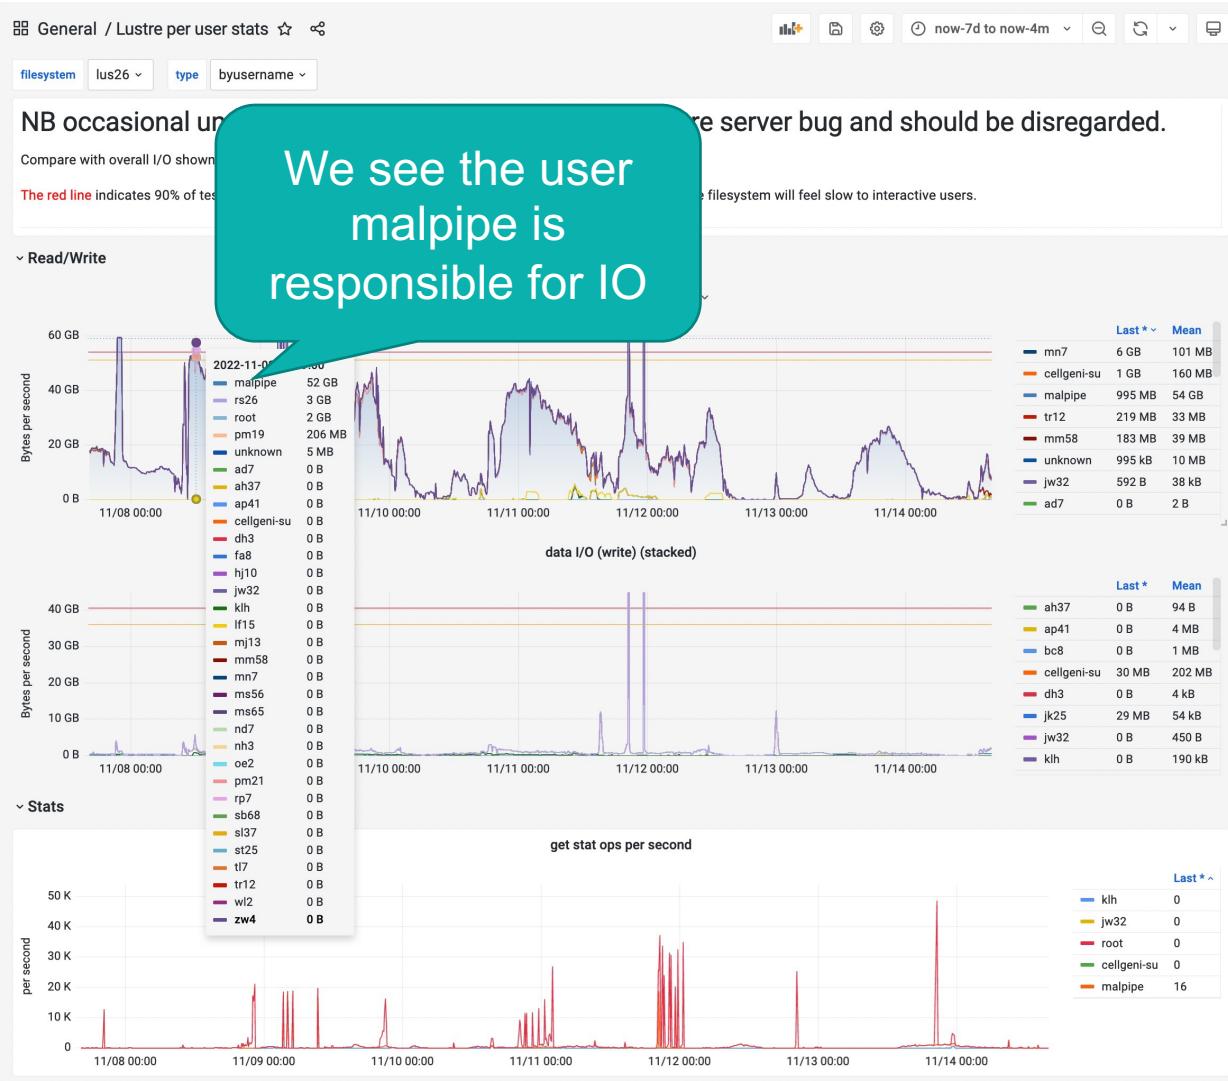
Bytes per second



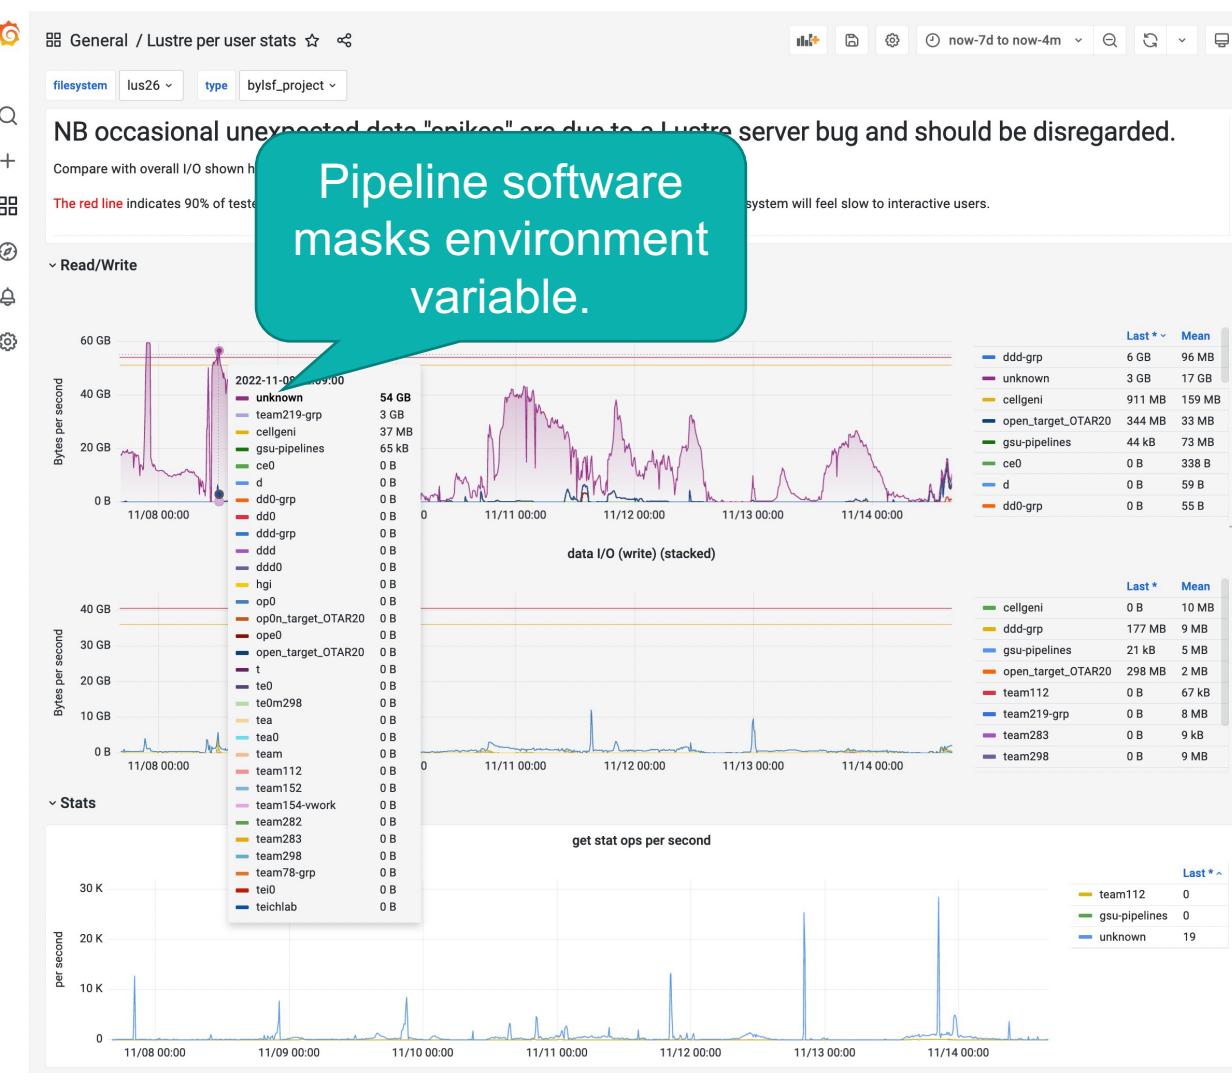


EXAMPLE 2

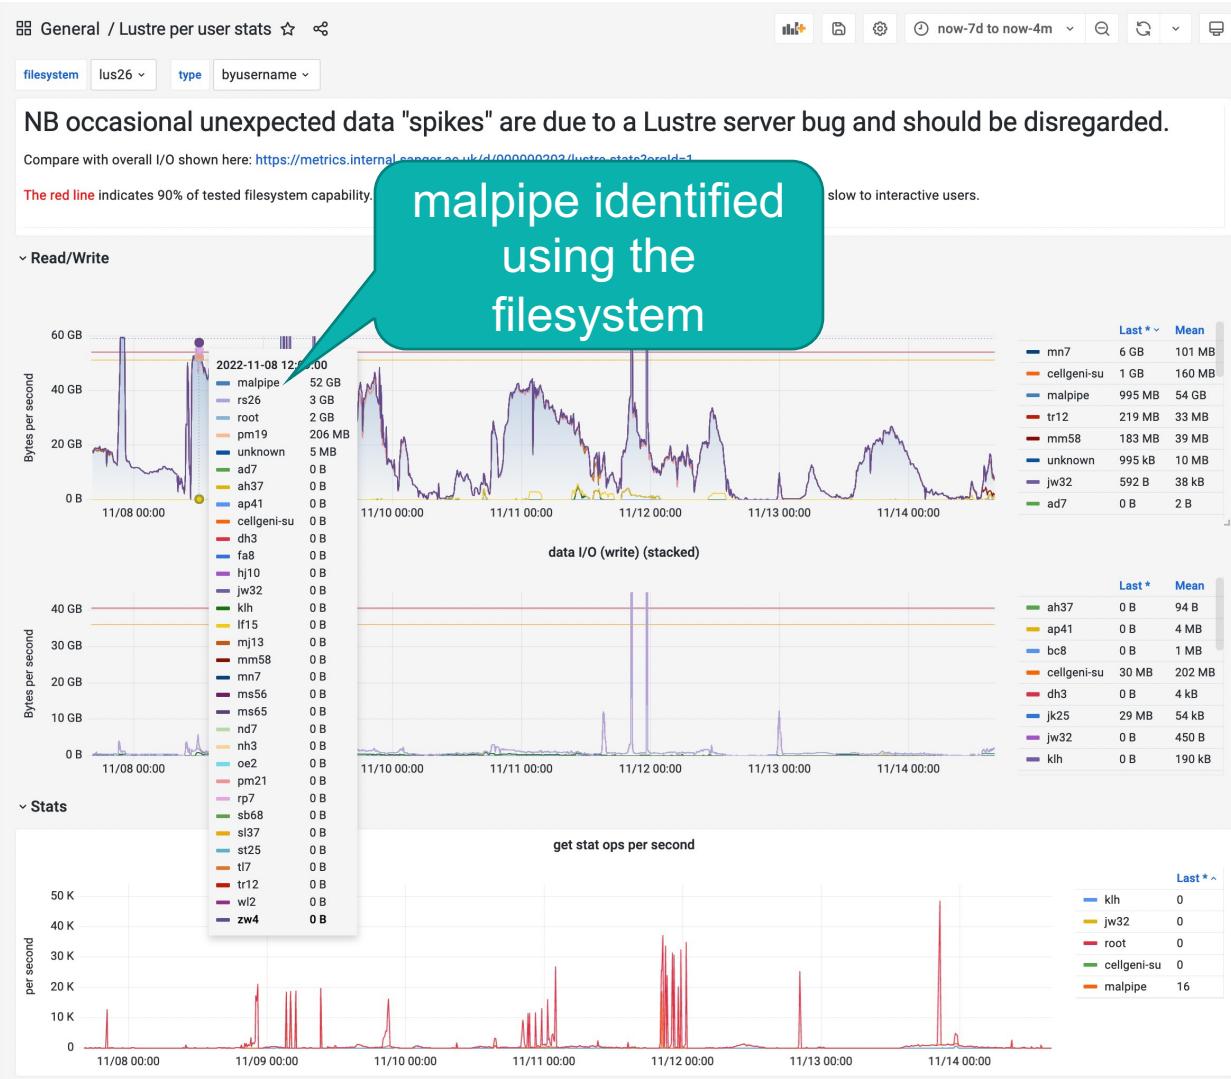




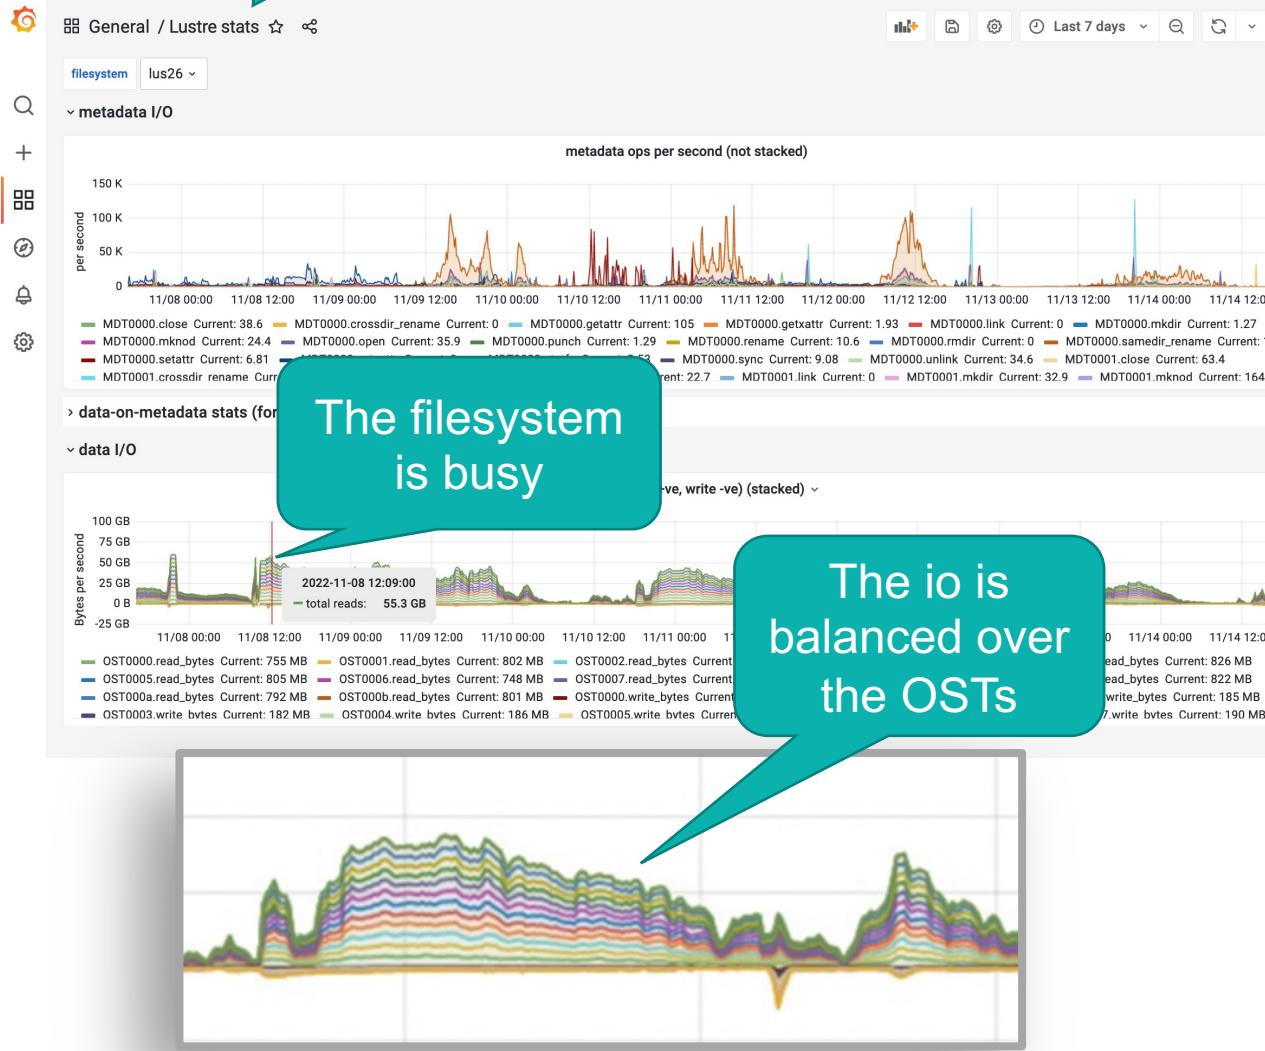
We see the user malpipe is responsible for IO



Pipeline software masks environment variable.



Switching to
filesystem stats



HOW IT WORKS

An aerial photograph of a large, modern industrial or research complex. The buildings are multi-story with grey roofs and light-colored walls, some featuring glass facades. The facility is nestled among lush green trees and is situated in a rural area with rolling hills and agricultural fields in the background under a clear sky.

ARCHITECTURE

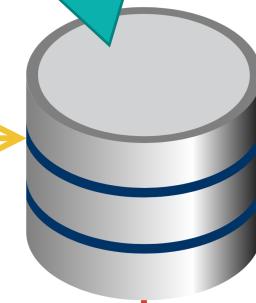
Exported and graphed



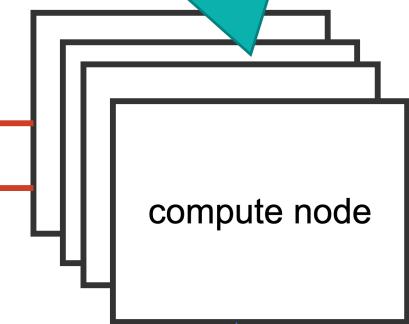
Polls each Lustre server



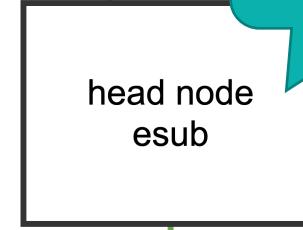
Server receives
RPC with I/O
operation



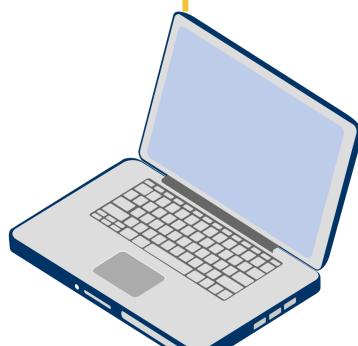
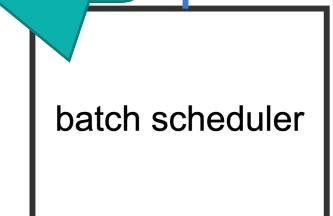
Enriched info
exported to
\$LSF_JOBID



User job
submission



Job scheduled



users/operator terminal

https

ssh

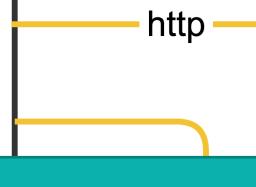
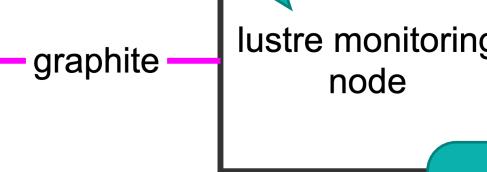
lustre

lustre filesystem

LSF

lustre filesystem

LSF



EXPORT JOB STATS TO GRAPHITE

Parse jobstat output, including enriched information on user, group and more

1. Simple web app for Lustre servers – Return JSON structure with jobstat info

```
def do_GET(self):
    data={}
    self._set_headers()
    osts=map( lambda x: x.replace("/proc/fs/lustre/obdfilter/", "").glob.glob("/proc/fs/lustre/obdfilter/lus*"))
    mdts=map( lambda x: x.replace("/proc/fs/lustre/mdt/", "").glob.glob("/proc/fs/lustre/mdt/lus*"))
    for ost in osts:
        data[ost]=run_lctl('lctl get_param obdfilter.{0}.job_stats'.format(ost).split())
    for mdt in mdts:
        data[mdt]=run_lctl('lctl get_param mdt.{0}.job_stats'.format(mdt).split())
    self.wfile.write(json.dumps(data))
```

2. Parse the jobstat output to create our data structure

```
fix_jobname = re.compile(r'^- job_id:(\s*)(\S(?:.*\S)?)(\s*)$')

def process_line(line):
    # Quote the job_id
    line = fix_jobname.sub(r'- job_id:\1"\2"',line)
    # Ensure valid yaml with a space
    return line.replace("max:", "max:").replace("min:", "min:").replace("sum:", "sum:").replace("samples:", "samples:").replace("sumsq:", "sumsq: ")

def run_lctl(cmd):
    p = Popen(cmd, stdout=PIPE)
    # skip the first two headers lines
    # https://stackoverflow.com/questions/15571137/yaml-scanner-scannererror-while-scanning-a-directive
    data_stream = "\n".join(map(process_line,p.communicate()[0].split("\n")[2:]))
    return yaml.load(data_stream)
```

3. Encode & Decode the user, group, and other extra info into the Job ID

```
if var_has_contents("LSB_SUB MODIFY_ENVFILE"):
    with open(os.environ["LSB_SUB MODIFY_ENVFILE"], "a") as modfile:
        header = "SA1"
# if
```

4. Decode the Job ID to recover user, group, and other extra info

```
def decode_jobname(name,store):
    store['project']="unknown"
    store['uid']="unknown"
    store['pgid']="unknown"
    store[store]
```

5. Example JSON output, including enriched data – import directly into Grafana!

```
curl -s http://lus23-oss1-mgmt.internal.sanger.ac.uk:8080 | jq .
```

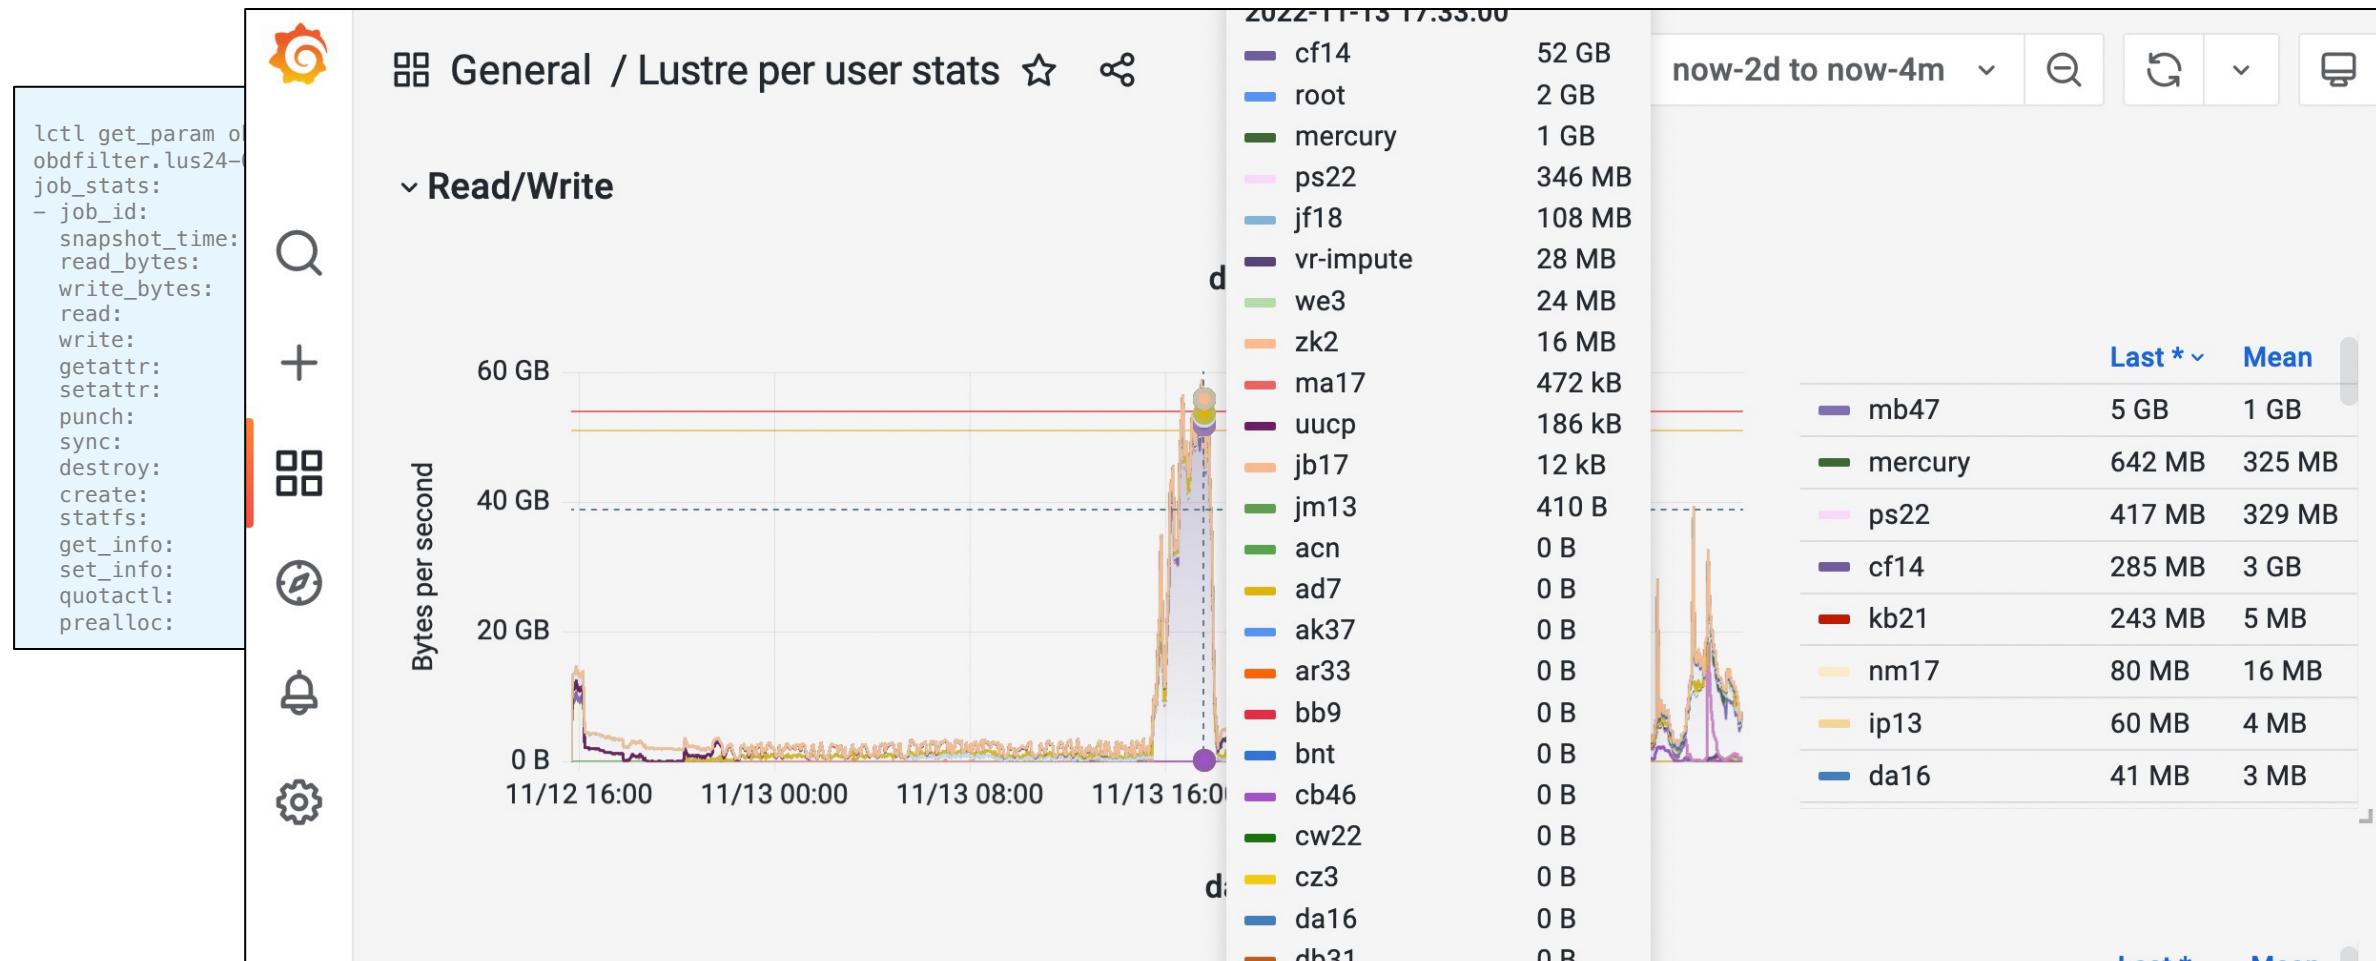
```
{
    "lus23-OST001": [
        {
            "write_bytes": {
                "max": 16777216,
                "sum": 1277816778962,
                "samples": 413738,
                "unit": "bytes",
                "min": 4
            },
            "pgid": "15152",
            "job_id": "SA1TwDkw0zA7,team135",
            "uid": "14784",
            ...
            "project": "team135",
            "gid": "15152",
            "farm": "tol",
            "read_bytes": {
                "max": 16777216,
                "sum": 51752874618880,
                "samples": 22381014,
                "unit": "bytes",
                "min": 4096
            }
        }
    ]
}
```

Enriched info exported to Graphite

| Series | Last | Mean |
|-----------|--------|-------|
| unknown | 0 B | 2 GB |
| team227 | 765 MB | 3 GB |
| teichlab | 514 MB | 35 MB |
| team282 | 398 MB | 7 MB |
| mouse.api | 189 MB | 17 MB |
| grit.grp | 133 MB | 27 MB |
| team135 | 66 MB | 8 MB |
| hg1 | 51 MB | 17 MB |

EXPORT JOBSTATS TO GRAPHITE

Parse jobstat output, including enriched information on user, group and more



JOB STATS: CARE NEEDED

- Raw data is fetched in grey with lctl get_param
- **LU-16110, LU-13857** – jobs_stats not valid yaml
- **LU-15870** – job_stats is sometimes corrupt
- **LU-15826** – job_id not quoted

Skip me

The command we run

This space is important

```
lctl get_param obdfilter.lus24-OST0000.job_stats
obdfilter.lus24-OST0000.job_stats=
job_stats:
- job_id: SA1CUzMGBwYH,analysis-cgp
  snapshot_time: 1667904951
  read_bytes: { samples: 297356696, unit: bytes, min: 4096, max: 16777216, sum: 1762446621732864, sumsq: 9915053801461514240 }
  write_bytes: { samples: 20518685, unit: bytes, min: 1, max: 16777216, sum: 98655040843302, sumsq: 4627393946935432142 }
  read: { samples: 297356696, unit: usecs, min: 1, max: 12945252, sum: 11870681362787, sumsq: 3234937531575968175 }
  write: { samples: 20518685, unit: usecs, min: 1, max: 10280151, sum: 98226898749, sumsq: 50535912682848539 }
  getattr: { samples: 0, unit: usecs, min: 0, max: 0, sum: 0, sumsq: 0 }
  setattr: { samples: 5697, unit: usecs, min: 4, max: 6119237, sum: 151065923, sumsq: 327770191069603 }
  punch: { samples: 1458531, unit: usecs, min: 5, max: 31397444, sum: 15753447595, sumsq: 36605999152526109 }
  sync: { samples: 6517520, unit: usecs, min: 0, max: 7189473, sum: 10441362075, sumsq: 12097924287750305 }
  destroy: { samples: 0, unit: usecs, min: 0, max: 0, sum: 0, sumsq: 0 }
  create: { samples: 0, unit: usecs, min: 0, max: 0, sum: 0, sumsq: 0 }
  statfs: { samples: 0, unit: usecs, min: 0, max: 0, sum: 0, sumsq: 0 }
  get_info: { samples: 0, unit: usecs, min: 0, max: 0, sum: 0, sumsq: 0 }
  set_info: { samples: 0, unit: usecs, min: 0, max: 0, sum: 0, sumsq: 0 }
  quotactl: { samples: 0, unit: usecs, min: 0, max: 0, sum: 0, sumsq: 0 }
  prealloc: { samples: 0, unit: reqs }
```

JOBSTATS: CONFUSION

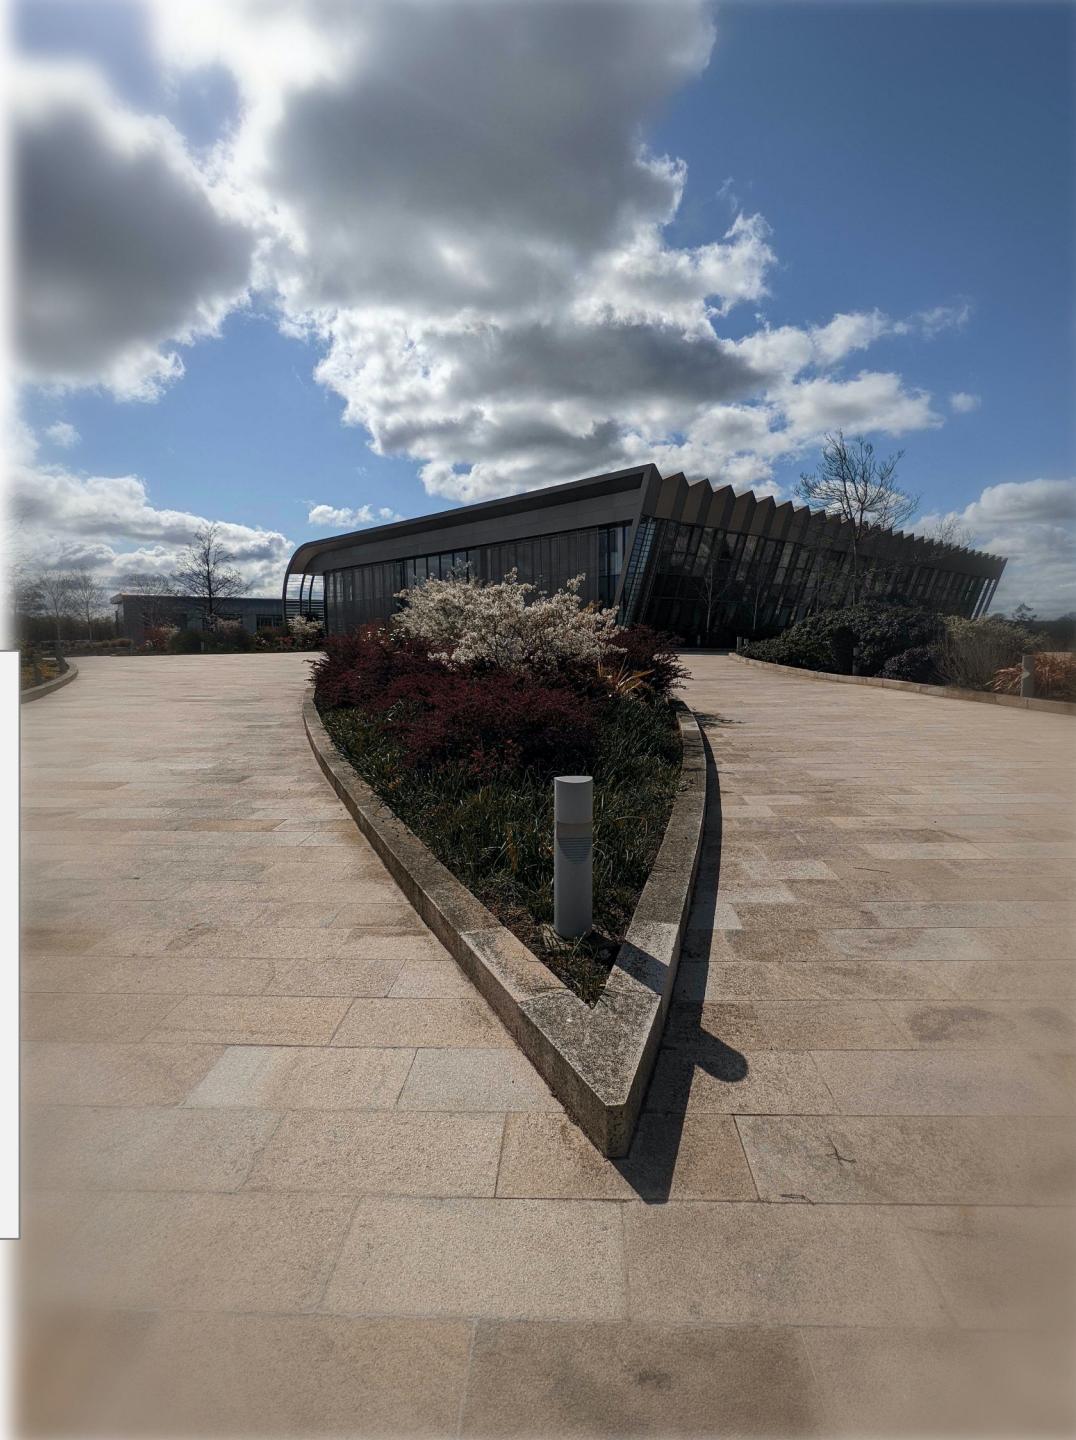
- We see corruption in the job id, LU-15870, LU-16251
- We see kworker recorded as the executable with root.
- We see thread-pool with no uid.

```
- job_id: kworker/u518:6eamtrynka
snapshot_time: 1668009798
read_bytes: { samples: 1, unit: bytes, min: 663552, max: 663552, sum: 663552 }
write_bytes: { samples: 0, unit: bytes, min: 0, max: 0, sum: 0 }

- job_id: kworker/u130:0
snapshot_time: 1668009796
read_bytes: { samples: 21459, unit: bytes, min: 4096, max: 2297856, sum: 160104448 }
write_bytes: { samples: 0, unit: bytes, min: 0, max: 0, sum: 0 }

- job_id: thread-pool-14113912
snapshot_time: 1668009775
read_bytes: { samples: 2, unit: bytes, min: 50804, max: 983040, sum: 1019904 }
write_bytes: { samples: 0, unit: bytes, min: 0, max: 0, sum: 0 }

[root@lus23-oss1 ~]# lctl get_param obdfilter.lus23-OST0000.job_stats | grep job_id | grep thread-pool | wc -l
126
[root@lus23-oss1 ~]# lctl get_param obdfilter.lus23-OST0000.job_stats | grep job_id | grep kworker | wc -l
155
[root@lus23-oss1 ~]# lctl get_param obdfilter.lus23-OST0000.job_stats | grep job_id | wc -l
330
```



SECURE LUSTRE

- A method of delivering lustre filesystems to OpenStack projects using
 - Sub directory mounts
 - And user id mappings
- <https://hpc-news.sanger.ac.uk/2021/05/20/update-on-secure-lustre-lustre-users-group-may-2020/>
- Jobid is inserted on the client, the uid inserted may and often does not exist on the outside
- The nodemap the request is received on is unavailable in the jobstats data
- Our current aggregation program appears to drop this traffic.



CONCLUSIONS

- Small change on Lustre server
- Not perfect
- Still useful
- Released under Apache 2 license
- <https://github.com/wtsi-ssg/lustre-jobstats>

FOR THE FUTURE

- We need users to have a “fair share” for I/O
- Is NRS the answer to that question?





ACKNOWLEDGMENTS

- ISG (Photo is missing Anna, Bruno, Fabio and Sai) (And includes Brett who is no longer with us).

THANK YOU



@wellcomegenomecampus



@wellcomegenome



Wellcome Genome Campus



Wellcome Genome Campus