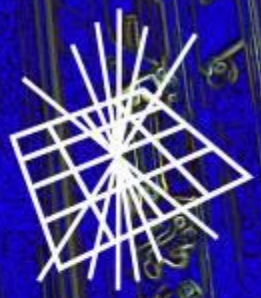

GridPP

UK Computing for Particle Physics

Ceph at RAL

an update

James Adams, RAL Tier 1

Ian Collier

GridPP 35, Liverpool, 11/09/2015

Background

- Ceph is an awesome object store
- Interest all over, Tier 1, SCD, STFC and beyond
  - ISIS using CephFS in production
  - SAGE project with CCFE Culham
  - JASMIN interested in high-performance Ceph
  - BNL running pre-production Ceph instance
- Tier1 Aiming to replace CASTOR disk with Ceph
  - CASTOR tape works well and will stay
    - Looks like it may become a layer on top of Ceph anyway
  - Trying to provide thinnest possible interface layer
    - GridFTP & Xrootd servers that talk RADOS

Background

- Tier 1 running three clusters
  - Dev
    - Development environment for gateways (see later slides)
    - Currently being used to test upgrade to latest Ceph release
  - Cloud
    - No major changes
    - Running smoothly
    - Starting to support additional larger RBDs for VMs (eg. CVMFS data disks)
  - Grid
    - Rebuilt on all new hardware
    - Lots of development and testing work
    - Focus of rest of talk

Grid Cluster Timeline

- 2015
  - Q1 : Hardware delivered
  - Q2 : Cluster deployed, developing plugins for GridFTP & Xrootd
  - Q3 : Internal stress testing, VO functional testing
  - Q4 : Launch as pre-production service with a large VO
- 2016
  - Q1 : Launch as production service, start migrating VOs from CASTOR

Grid Cluster Hardware

- Now have “real” hardware
  - Three physical monitors
    - Dell R420 with fast disks
  - Three gateway machines
    - Dell R430 with quad 10Gbps Ethernet
    - Problems deploying these… support issues with vendor…
  - 48 storage nodes
    - 100TB-120TB with dual 10Gbps Ethernet, 64GB RAM, etc…

Deployment

- All deployment using Quattor (aquilon)
- Ceph component developed with University of Ghent
  - Manages configuration files
  - Drives ceph-deploy
- Can now deploy RADOSGW with Quattor
- Still no solution for secure distribution of keys
- Adding new nodes from bare metal is now only a few steps
  - Add host to aquilon
  - Set to install
  - Reboot machine

  …wait…
  - Set personality to Ceph storage node

Erasure Coding

- Aiming to hit a similar storage overhead to CASTOR
  - ~83% of raw storage usable
- Erasure coding decisions
  - ***k*** data chunks + ***m*** parity chunks
  - We want ***m*** >= 3
  - We want ***k*** as large as possible without making performance stinky
  - Aspiring for 14+3 or 16+3
    - For comparison… Yahoo using 8+3, Facebook using 10+4
  - Currently testing lots of combinations
    - Watch this space (presentation at HEPiX)
- Possible risk with patents and jerasure library

Pools & Placement Groups

- One or two pools per WLCG VO
- Still thinking about approach for “small” Vos
  - One big pool or individual smaller pools?
  - ACLs are one of the issues here
  - Would really rather they used S3 (or Swift)
- Placement groups are “per pool”
- Need many placement groups
  - Overhead on cluster a factor
  - Too few and the cluster won’t balance properly
  - Too many and the cluster will never recover
  - Thousands per pool as a baseline
  - Testing this now

Gateways

- RADOS gateway
  - Mostly interested in S3 support
    - Sub-URLs rather than subdomains (we don't control our DNS)
  - IPV6 works well (yay!)
- GridFTP gateway
  - GSI proxies working
  - Working on VOMS proxies now
  - Being packaged for release soon™
- Xrootd gateway
  - Part of Xrootd 4.2 release

Monitoring

- Dashboards
  - Build by Nuffield placement student Ignacy Debicki
  - InfluxDB + Grafana
  - Metric collection scripts on GitHub
    - https://github.com/stfc/ceph-InfluxDB-metricsCollector
  - Can now watch cluster doing its thing

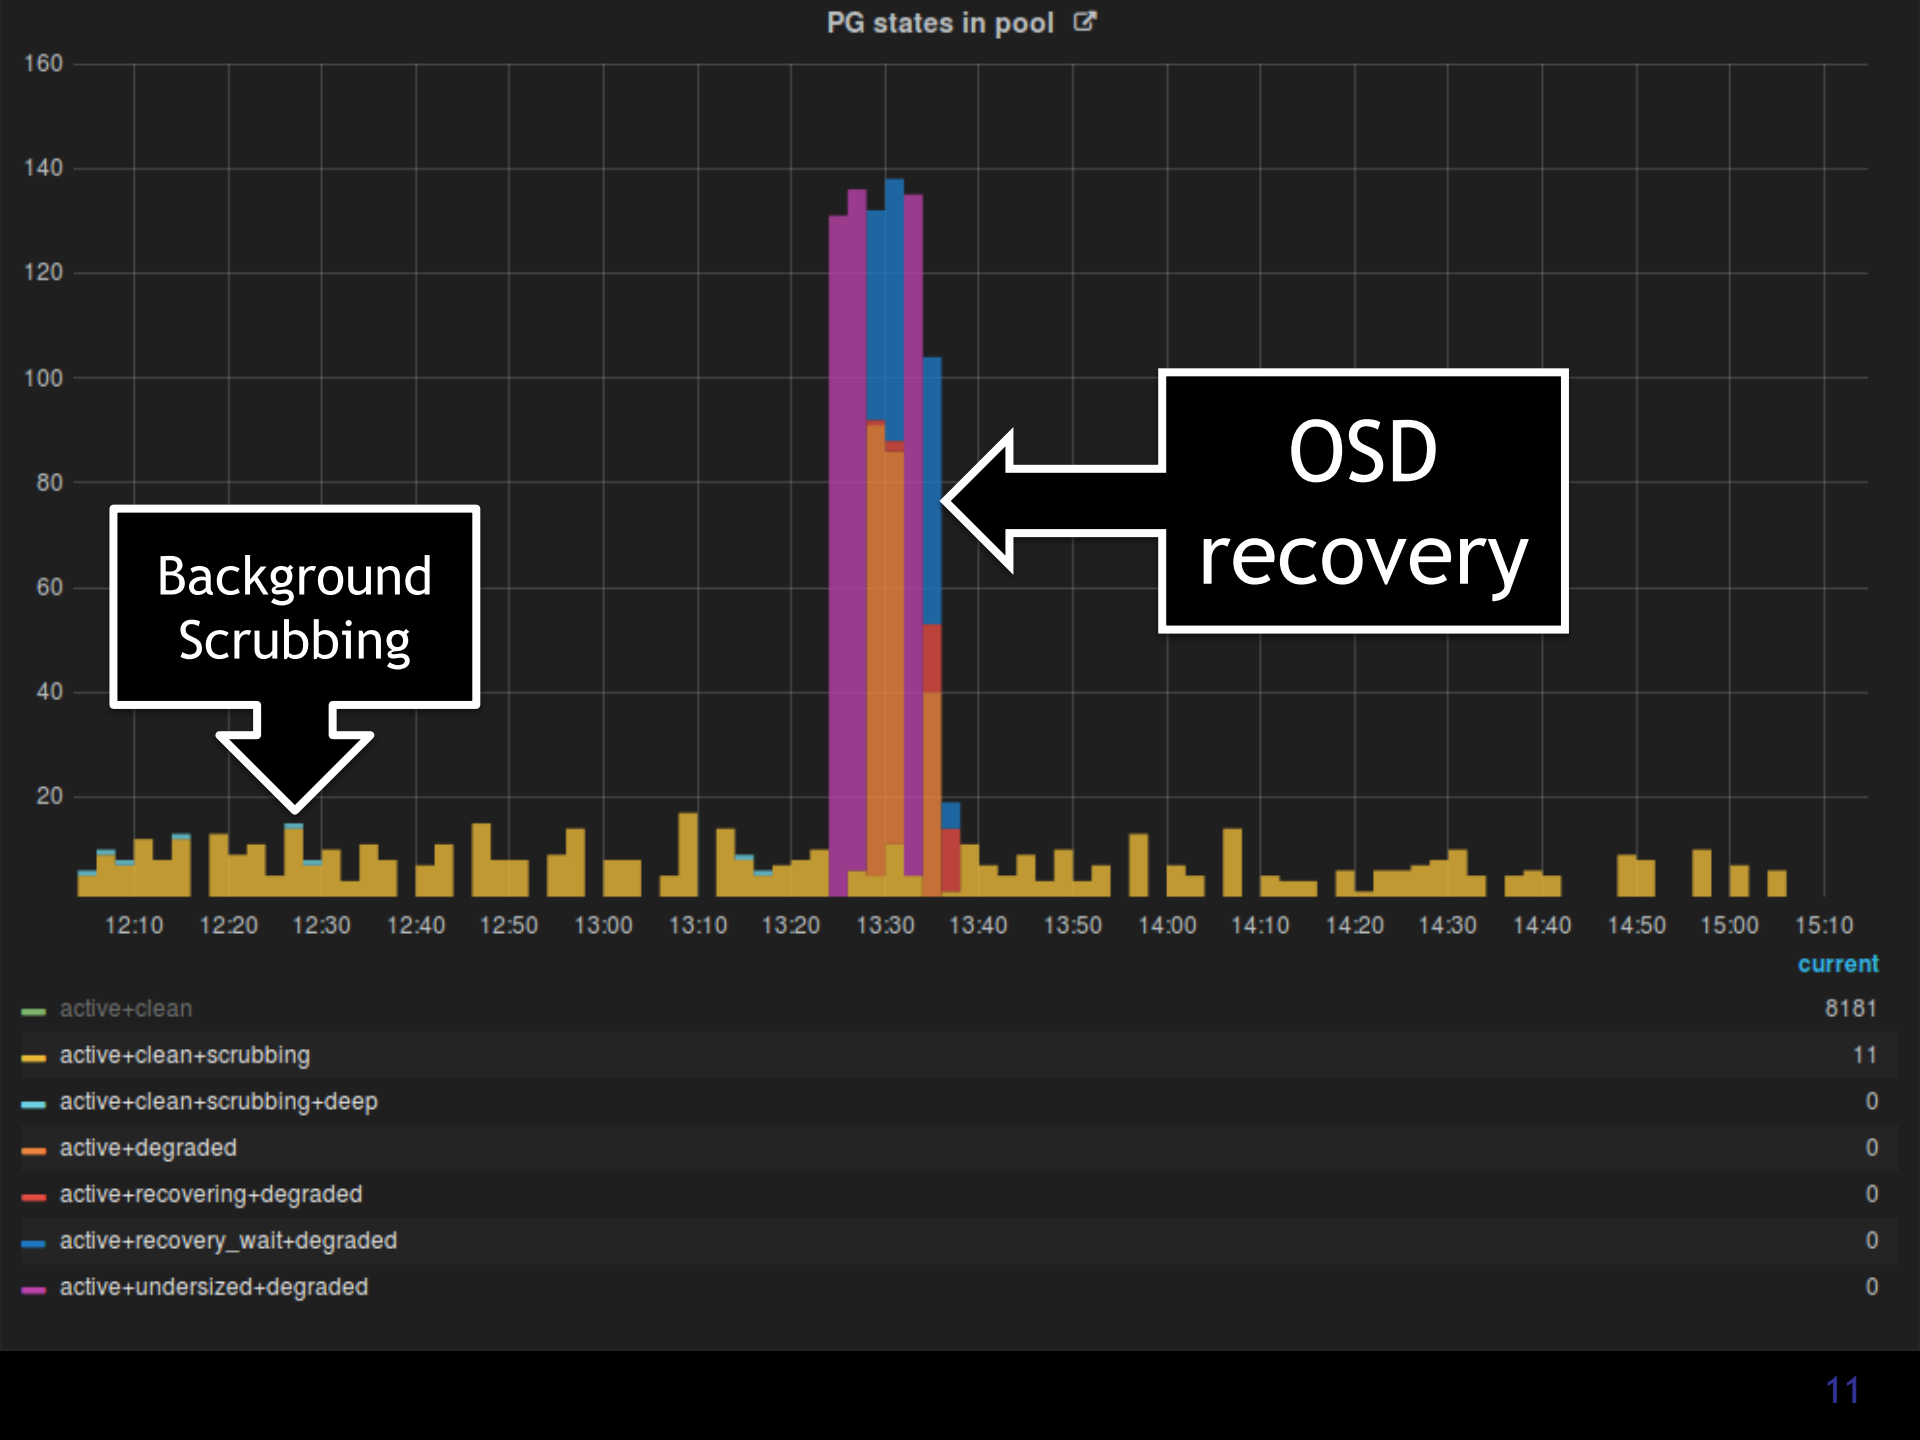
PG states in pool
160
140
120
100
80
60
40
20
12:10
12:20
12:30
12:40
12:50
13:00
13:10
13:20
13:30
13:40
13:50
14:00
14:10
14:20
14:30
14:40
14:50
15:00
15:10
Background Scrubbing
OSD recovery
current
active+clean 8181
active+clean+scrubbing 11
active+clean+scrubbing+deep 0
active+degraded 0
active+recovering+degraded 0
active+recovery_wait+degraded 0
active+undersized+degraded 0

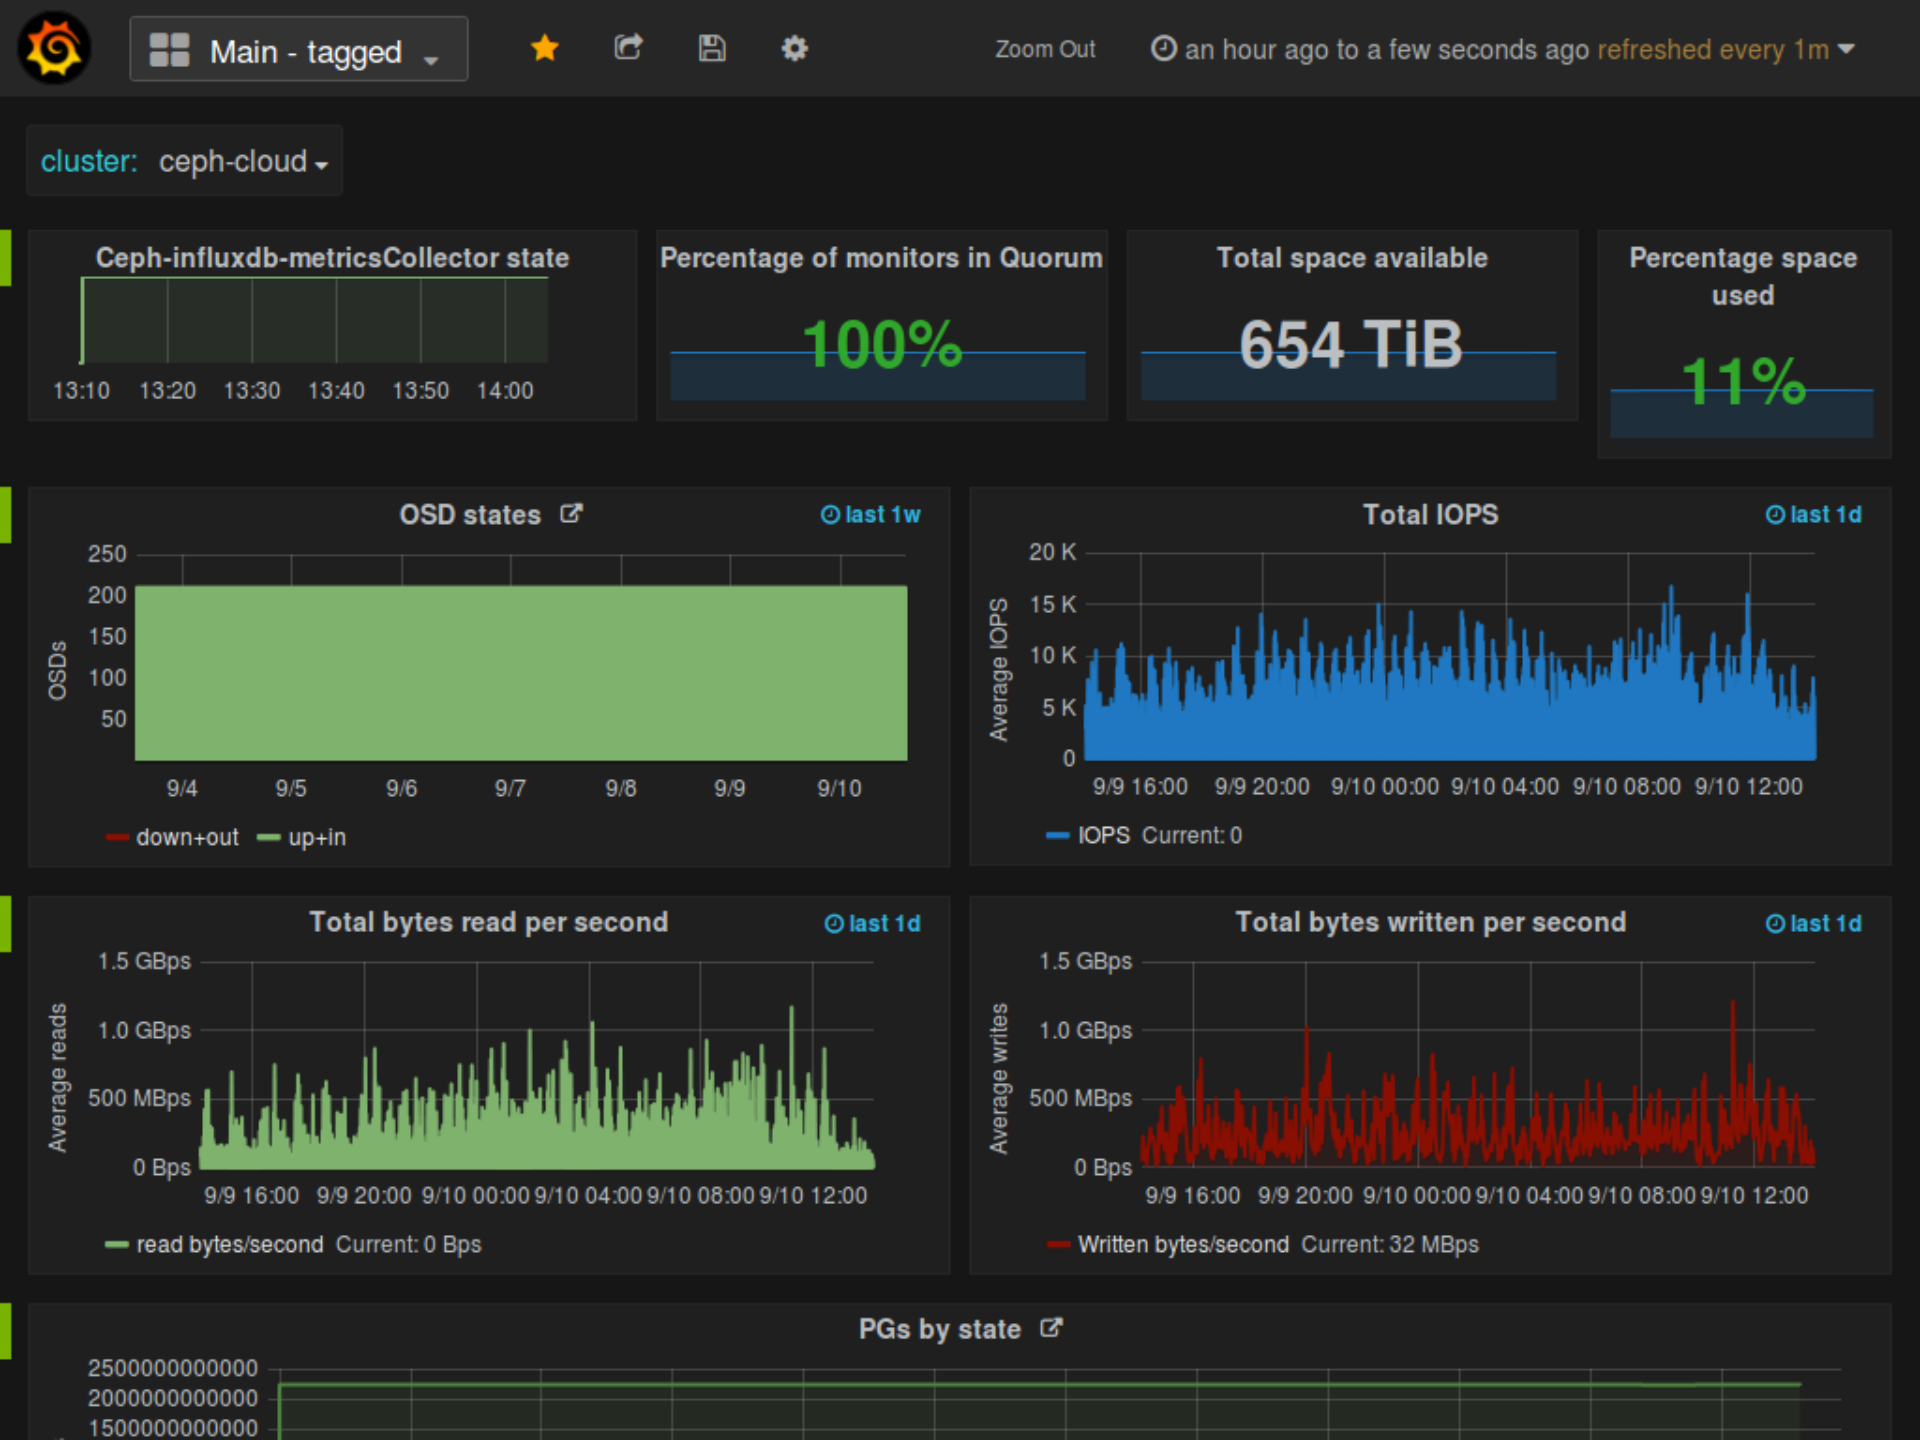
Main - tagged
Zoom Out
an hour ago to a few seconds ago refreshed every 1m
cluster: ceph-cloud
Ceph-influxdb-metricsCollector state
13:10 13:20 13:30 13:40 13:50 14:00
Percentage of monitors in Quorum
100%
Total space available
654 TiB
Percentage space used
11%
OSD states
last 1w
OSDs
250 200 150 100 50
9/4 9/5 9/6 9/7 9/8 9/9 9/10
down+out up+in
Total IOPS
last 1d
Average IOPS
20 K 15 K 10 K 5 K 0
9/9 16:00 9/9 20:00 9/10 00:00 9/10 04:00 9/10 08:00 9/10 12:00
IOPS Current: 0
Total bytes read per second
last 1d
Average reads
1.5 GBps 1.0 GBps 500 MBps 0 Bps
9/9 16:00 9/9 20:00 9/10 00:00 9/10 04:00 9/10 08:00 9/10 12:00
read bytes/second Current: 0 Bps
Total bytes written per second
last 1d
Average writes
1.5 GBps 1.0 GBps 500 MBps 0 Bps
9/9 16:00 9/9 20:00 9/10 00:00 9/10 04:00 9/10 08:00 9/10 12:00
Written bytes/second Current: 32 MBps
PGs by state
2500000000000
2000000000000
1500000000000

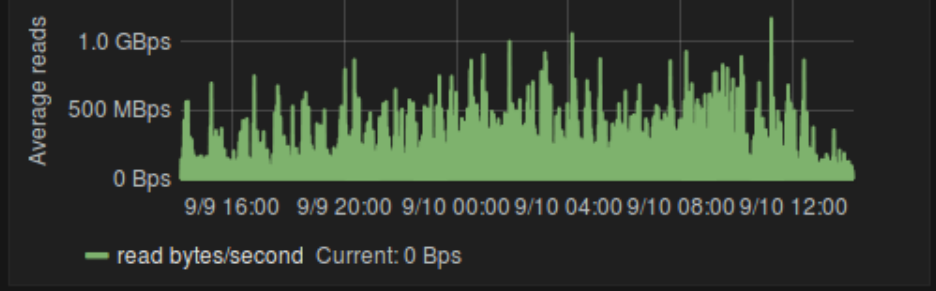
Average reads
1.0 GBps
500 MBps
0 Bps
9/9 16:00
9/9 20:00
9/10 00:00
9/10 04:00
9/10 08:00
9/10 12:00
read bytes/second Current: 0 Bps

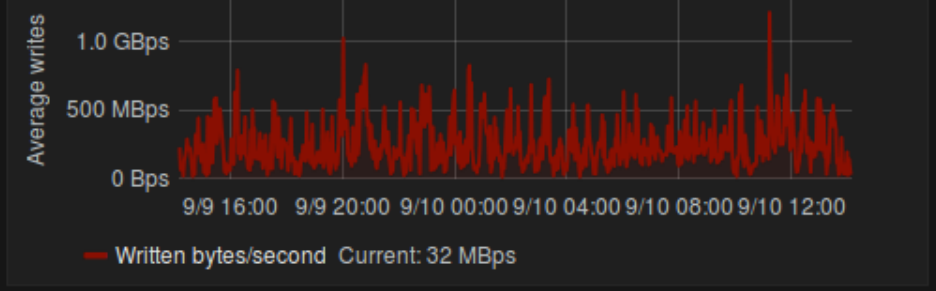
Average writes
1.0 GBps
500 MBps
0 Bps
9/9 16:00
9/9 20:00
9/10 00:00
9/10 04:00
9/10 08:00
9/10 12:00
Written bytes/second Current: 32 MBps

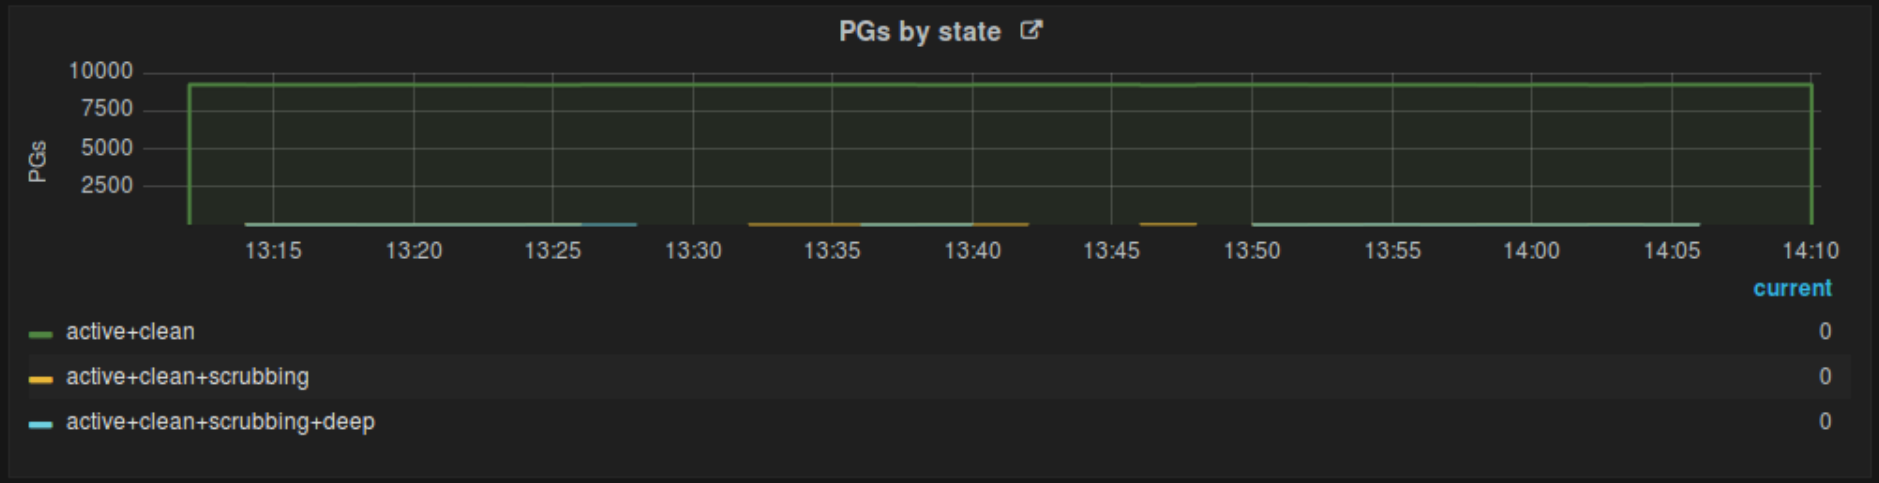
PGs by state
PGs
10000
7500
5000
2500
13:15
13:20
13:25
13:30
13:35
13:40
13:45
13:50
13:55
14:00
14:05
14:10
current
active+clean 0
active+clean+scrubbing 0
active+clean+scrubbing+deep 0

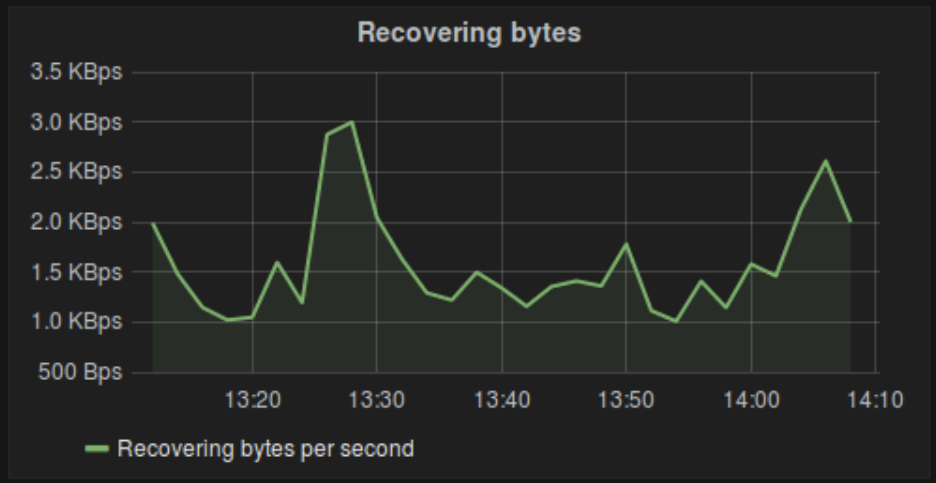
Recovering bytes
3.5 KBps
3.0 KBps
2.5 KBps
2.0 KBps
1.5 KBps
1.0 KBps
500 Bps
13:20
13:30
13:40
13:50
14:00
14:10
Recovering bytes per second


Recovering objects
3.5 K
3.0 K
2.5 K
2.0 K
1.5 K
1.0 K
500
0
13:20
13:30
13:40
13:50
14:00
14:10
Recovering objects per second

HELP

+ADD ROW

Summary

- Grid cluster is main focus of work now
  - Lots of testing to determine optimal configuration
- Start testing with VOs in October
  - Functional testing
- Production service early next year
  - Migration VO by VO
- Cloud storage cluster is operational now