

# Operations Working Group Summary

## 2017-12-11

Present: Andy, Grant, Jon, Tom, Matt  
Apologies: Emilie

### Disk space on katla

The main disk array on katla, the API database replica hosted by Bytemark in York, is filling up. It is currently configured for higher performance at the expense of disk space, but it was decided that, given we now have another, faster API database replica hosted by UCL in Slough and donated by ARC, that a reconfiguration to make more space available would be appropriate.

Actions: Move to RAID6. Drop the current array, re-image, and re-sync the database (Grant & Tom -- [DONE](#)).

After, consider if it's worth moving eddie (replica currently at Slough) to York, or turning Slough into a full site -- currently it's missing web & services machines.

### Disk space for the rendering database on scorch.

The rendering database on scorch, the tile rendering server donated by OVH, is a single 800GB disk, which is reaching capacity. We can use the disks in a way which frees up more space, at the cost of redundancy. However, given we have multiple rendering servers, the possibility of one failure might be acceptable.

Option 1: Ask OVH if we can have more disks. Timeout 1 week. (Tom?)

Option 2: Put all disks into a RAID0 array

### Whether to continue providing sarel for gosmore.

Sarel, a gosmore routing service hosted by UCL, has extremely old data on it and is not currently maintained. It might be a useful service, but not with such old data, and not if there's no one willing to maintain it.

Actions: Andy to follow up with current maintainer. If there's no movement there, follow up on the operations tracker.

## Provo tile caches (nadder-01/02) which have been down for several months.

The tile caches donated by Bluehost in Provo, Utah, USA have been down for several months. The person who originally donated them has been uncontactable.

### Actions:

1. Ping Bluehost support to see if we can get a response. (Grant?)
2. Else, shut down nadders - they're dead and our contact is gone.
3. Put out a call for NA tile caches.
4. If all else fails, move taginfo somewhere else and use stormfly-01 (bandwidth limit permitting).

## 2018 Budget.

Need to decide what to do, not actually work through the details, so things like what big purchases to we see, what do we want to get contractors in to do, etc

1. Set aside money for hosting: £12k.
2. ½ year of IC: £5k.
3. 3x FE, 3x BE + storage for Slough: £15k.
4. Contracting: XML replication: £10k.
5. 36x new 2TB disks for Ceph: £5k.
6. New render machine: £3k (360g8 128GB + 8xSSD)
7. Contingency and consumables costs.

Actions: Andy to write up and circulate for agreement.

## Imperial College data centre shut down

The Imperial College data centre will be shut down soon, and they have politely reminded us that we need to figure out a new home for those machines - whether that's a new DC, more space in an existing DC, etc...

### Actions:

1. Add 3x FE, 3x BE + storage (smaug?) to Slough to make it a full site (Matt, Grant)
2. 3rd site:
  - a. Put request for proposals up on the operations website (Matt)
  - b. Make a blog post about it
  - c. Reach out to existing partners.

## We're stuck on PG 9.5

Until we can move away from `xid`-based replication. There has been some work on an “osm-logical” plugin, but it is far from being a complete, working system.

See [work done by Jon](#).

### Actions:

1. Write a spec sheet for what we want the result to be. (Jon?)
  - a. Performance
  - b. How it's deployed. Danger to internals of postgres?
  - c.
2. Contracting?

## Ceph cluster.

We've been wanting to get one running for a while. There are some public cookbooks, but these aren't simple to integrate due to the structure of our internal Chef repositories. Andy and Grant both already run Ceph clusters, so it should be possible to get this working. We want to try and get others involved, perhaps via a hack weekend.

### Recommendations:

- Ceph setup
  - 3 monitors
  - 2 gateways
  - As many OSDs as we can get
  - Erasure coding
  - Need version 12.2.2 (Luminous)
- Object storage only (i.e: no CephFS)
- First test: user profile images

### Actions:

- Schedule a hack weekend to get something working (Grant?)