

EOS at Edinburgh



Intro

Unfortunately we've had a busy month in Edinburgh so this talk is a lot less about what we've done/tried to do and a lot more about what I would like to do over the next few months.

Please correct me, point out where I'm wrong, criticise me and point out a better alternative I'm approaching this from a pragmatic perspective I'm not trying to be a salesman.

Why?

In the UK Birmingham have made the move. This was largely due to supporting ALICE.

DPM is shifting towards DPM+DOME which makes DPM look more and more like an xrootd plugin.

After SRM/gridftp become deprecated is there a strong requirement for DPM at storage sites?

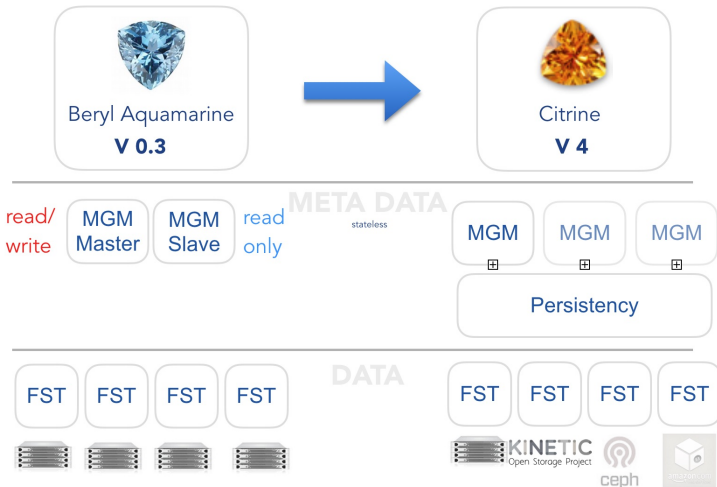
In the simplest case if a site can just support `root://` does it need anything other than a clever xrootd server?

What about this new thing called datalakes?

EOS vs DPM

EOS	DPM
Pros:	
Supports S3	Better Documentation
Deployable via Docker	Supports SRM
Supports CEPH	Proven
Can be distributed	Community experience
Has industry partnership	
Cons:	
Non-industry standard	Only supports HTTPS
Documentation is lacking	1 head node per site

What



So what do I hope to do?

My simple plan is to set up the most minimal EOS installation I can.

1 Head node (MGM) + 1 storage node (FST).

The same as our DPM+DOME test node:
`srm-test.gridpp.ecdf.ed.ac.uk`

This could be useful as a test endpoint and I'd like to make comparisons between the 2.

What I want to test

Questions I want to try and answer:

- How difficult is it to setup by comparison to a simple DPM?
- Can it be used as a drop-in replacement for DPM or does it require significant reworking of site storage?
- Does it support `https/root/s3` transfers and authentication out of the box?
- Is there anything new offered by EOS that can be called a killer feature?
- Can it talk to xroot as hosted by an existing DPM head node as an FST?

So far what have I done?

- Deployed an EOS instance via Docker.
This required updating docker on CentOS and pulling some rather large images.
- Played with the EOS setup.
Have verified that the docker instance does indeed setup multiple FTS to 1 MGM.
- Worked out how to add disks or a disk node?
Not yet, reading up on this atm.

EOS vs CEPH

I've decided to try and play with Docker rather than CEPH due to there being experience available through other people already with CEPH from other projects.

Looks like a smaller investment possibly to setup a CEPH test instance, but this has been done many times from a quick google.

I would be interested in can EOS be deployed with multiple types of FTS and without directly managing the disks.

If this is possible then can EOS manage a pool of disks + CEPH at one site as a local disk pool?

If this all goes wrong I may just play with/evaluate CEPH.

Potential Future Work

EOS has been shown to support distributed storage.

This works for large Tier1 sites but does it work for Tier2 resources and is it reliable?

If I can expand the simple case of 1 storage node + 1 master node to include multiples of each is this a simple mesh storage network?

From this could we build an EOS Scotgrid Datalake?

Summary

Apologies there's not a lot to actually show on this.

I've a plan for how to go about simple testing but getting even that far will likely need input from mailing lists/twikis.

Hopefully setting up a simple test instance shouldn't be too difficult.