



Science and
Technology
Facilities Council

RAL Storage for non WLCG users

What we Run

- Our usual audience is WLCG people...
 - So we spend a lot of time talking about Echo
 - (See previous talks)
- Our team has other disk storage services

Echo S3

- Echo provides an S3 endpoint as well as XrootD.
 - This is the bit we spend less time talking about.
- This provides S3 to users that request it, and is also exposed via Swift to the users on the openstack cloud.
- Use cases include easily web accessible data, and bulk storage with a non-grid access protocol.
- Big users are Diamond and RFI.
 - RFI are looking at more interactive use of the S3 buckets that involves watching for data being placed into it and starting a workflow when this happens.

Sirius

- Pure NVMe cluster built to provide underlying VM storage for the SCD Cloud
 - Block storage for Openstack
 - Originally built from HDDs
 - ...but those ran out of IOPS very very badly
 - Cloud service requires maximum performance, and in principle, not too much storage
 - Triple replication, so ~700 TiB raw, 230 TiB usable
- The SCD Cloud is a STFC-wide resource used by many STFC projects
- •Diamond, ISIS, Euclid...
- Intended use case was for the system disks of VMs – users were not expected to keep data there.

Aried

- Sirius had a problem
 - Users weren't supposed to use our very very expensive NVMe storage for their data...
 - ...but they did anyway ☹
- Tried telling the users off, didn't work...
- So, a new cluster
 - 1.2PiB raw SATA SSD, 400 TiB usable
- Aim to lure users away from storing data locally on VMs with new features
 - The ability to provision filesystems via a Web UI
 - Shareability
 - Not having to make a new VM every time they need more storage
 - Not being told off!
- Same users as Sirius
- Fully deployed last year on the Ceph side, currently in user testing by Cloud
- Exposed via Openstack Manila

Deneb

- Original idea:
 - Q: We have a load of HDD servers which we had to remove from Sirius, what should we do with them?
 - A: Make a CephFS cluster and see if anyone will find it useful
 - Turns out they did
- Now, 7.9PiB raw, 5.7PiB usable (8+3 EC)
 - Uses ~8TB HDDs with SSDs for the OSD caches (metadata ops)
 - Sharable bulk filesystem space for local facilities
 - RFI, ISIS, CLF (exposed as NFS/Samba)
- Future investigation – can we use the same hardware as Echo for the service?
 - Performance might not be good enough, but big £saving if it is...



Science and
Technology
Facilities Council

Thank you

Facebook: Science and
Technology Facilities Council

Twitter: @STFC_matters

YouTube: Science and
Technology Facilities Council