

Backup and DR with Large Datasets

SWITCH

Fergus Kerins

fergus.kerins@switch.ch

Copenhagen 27/01/2020

Our customers



Higher education

- Cantonal universities
- ETH domain with research institutions
- Universities of applied sciences
- Universities of teacher education

University-related organizations

- Spin-Offs
- Research institutions
- eLearning Center
- University hospitals

Our offer

- SWITCHengines
- SWITCHdrive



SWITCHdrive



Owncloud 10.3.2

Infrastructure

- CEPH
- Openstack
- 2 Datacenters / 2 Clusters / 2 Regions
- RBD volumes
- 6 NFS Servers
- 85 x 5TB disks

Storage

- *280 TB of Data*
- 237 Million files
- 20 K unique Users a month

Our Challenge



- We would like to provide offsite Backup and DR for our Customers
- Our Dataset is quite large and not really suited to the classic backup, of a full backup weekly and a daily Incremental

Options



CEPH

- rbd-mirror

3rd Party

- “scripts” rclone
- backy2 / benji

Plan



- We have done some basic testing and seen some positive results with using rclone to backup to s3 objectstore
- Due to our Data (relatively static)

Call to the Community

- Any Experiences or Comments would be very welcome

