

Ceph and XRootD Deployment and Monitoring at Lancaster

GridPP48, 31 August 2022

Gerard Hand, Steven Simpson, Matt Doidge

The Story So Far

Lancaster decided to move to a CEPHFS + XROOTD for well documented reasons (see our talk at GridPP47)...

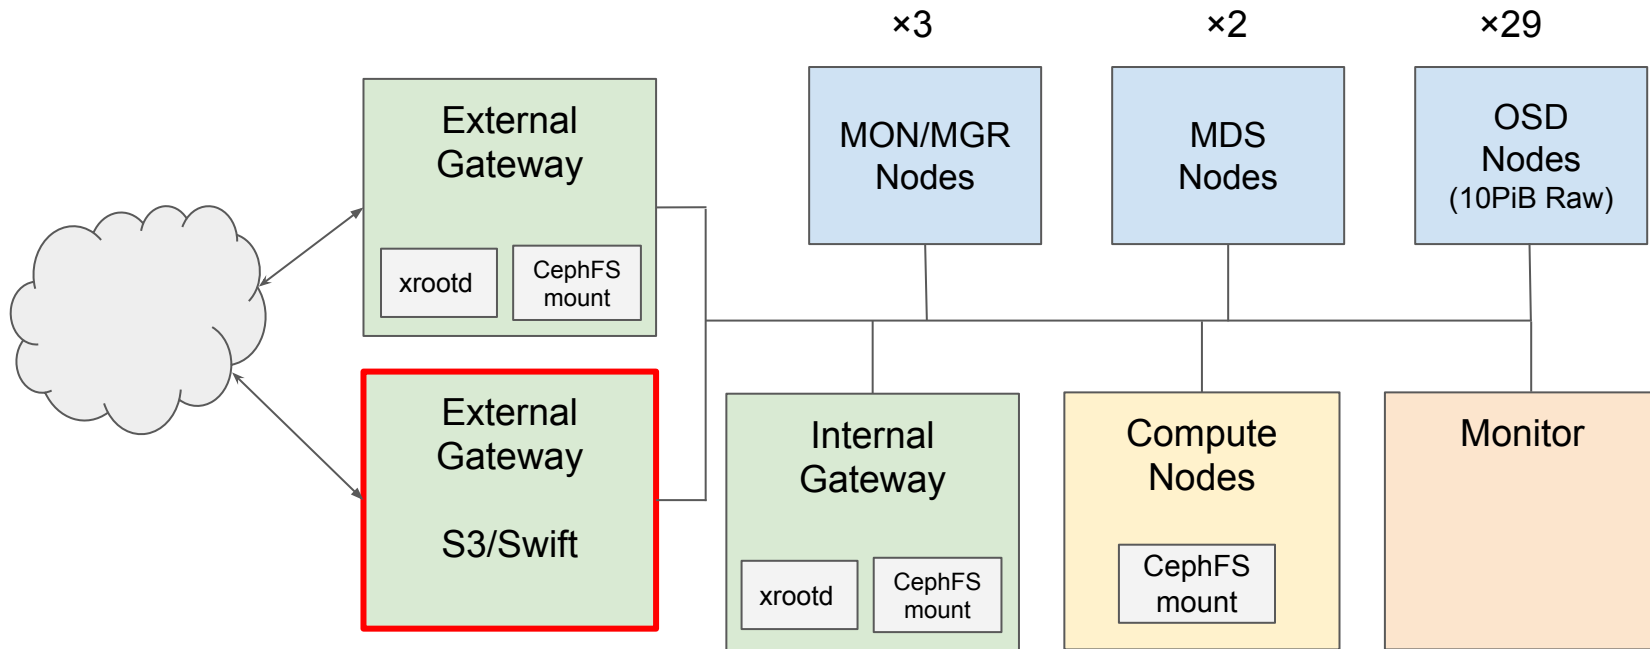
...and it's gone quite well.

- No disasters (touch wood).
- Moved into production without a lot of people noticing.
- The placeholder setup (of a single standalone xroot server fronting the whole thing) held up well.
 - And we've snuck a redirector in again without anyone really noticing.

But it's not been a journey not without a few gotchas and lessons learnt.

Ceph architecture

- An S3/Swift gateway has been added.




Lessons learned/Gotcha's Experienced

- Using Pacific instead of Octopus has proved a good choice as updates/patches are not getting back ported to Octopus.
- Creating an S3/Swift gateway was straightforward.
- Cephadm + Ceph Orchestrator have worked well. We have rebooted all the nodes in the cluster after system updates with no downtime.
- Recent problems with scrubbing not being performed in time required changes to default configuration settings.
- Running a Pacific cluster and using Octopus to mount CephFS on the client machines has worked without problems (No support for Pacific+ on Centos 7). There are currently 985 CephFS clients connected to the cluster.
 - All the odd jobs NFS used to do, like shared directories, CEPHFS now does.
- We have noticed that intermittently the Active MGR drops out of the cluster when running the Ceph Dashboard.


XRootD Xperiences

- Moved to a two-node redirector setup (see dodgy pic right)
 - Would rather have a 3 node setup (dedicated redirector + 2 dedicated servers), but this was easier to slot into place.
 - Discovery: Redirection very sensitive to IPv6 problems.
 - No problems noticed (so far).
 - A standalone server actually did quite well, the reasonably specced 25Gb-connected box coped with about 60% of our load.
- It was a race between rolling out the redirector or having the rucio “SymLink” plugin in place.
 - This won, as the rucio symlink hit a problem when it was attempted to be rolled out and had to be rolled back.
- Found that xrootd logging levels seem to be almost binary: “Too Much” or “Nothing Useful”.
 - Almost ran out of disk space when we left the logging on “all -debug”.
- Had some problems as we have overlooked testing xrootd TPC.
 - We thought no one really used this, we were wrong.
 - Turned out you need the “-f”
- We were hit again by the rhel8 “assertion error” problem during our testing.
 - This should be fixed in the upcoming xroot 5.5.0
 - But Steven provided fixed code several minor versions ago.



Redirector Balance

- I/O
 - Some oscillation over short periods (30min)
 - Flattens out over-all
- Requests
 - Even, even at small granularity
- Load
 - Based on getrusage
 - Even, even at small granularity
 - Slightly less load on shared host
 - Would have expected the other way around.
- This is all with the default settings.

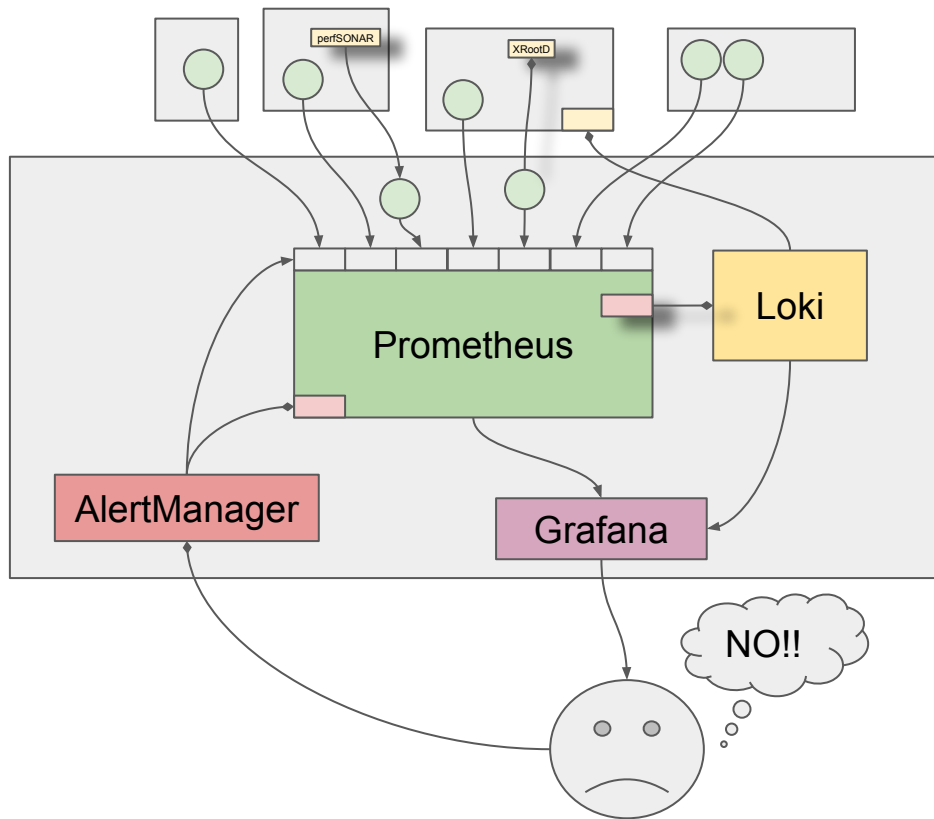


XRootD NeXt Steps

- How many 25Gb-connected xrootd boxes does it take to fill our 40Gb bandwidth?
<insert lightbulb joke that doesn't get us sent to HR>
 - I think the answer will be higher than 40/25
 - Whatever the number is, this is the minimum number of redirectors we will need.
- Tokens!
 - Scitokens have been enabled since day 1, but we're not in the testbed (yet).
 - Enabling ESCAPE tokens for (e.g.) SKA looks like it should be "simple".
 - The main blocker with rolling out tokens is testing - we're well versed in grabbing a voms proxy and firing off `gfal-*/xrd*` commands. Not so with `oidc-*/eldritch` curl invocations.
- Using some old DPM nodes for an xrootd testbed - useful for trying out Rocky8/xroot 5.5.X/anything weird.
- Ultimate aim is an HA redirector in front of (enough + 1) xroot servers, with none pulling double duty as a redirector and data server.
- Our configs can be seen at: <https://github.com/mdoidge/lancsxroot>



Monitoring updates

- Use Loki in anger
 - Easy to knock over
 - Too much logging from XRootD during debug
 - Slow queries in Grafana; many fail
 - Mitigated by recordings pushed into Prometheus
- Overhaul of metric labelling
 - Distinguish between host and interface
 - Eases correlation of metrics from diverse sources
- Custom exporters
 - XRootD
 - xrd.report push XML over UDP
 - pulled as time-stamped OpenMetrics
 - perfSONAR
 - periodic pull JSON from esmond
 - pulled as time-stamped OpenMetrics



Balancing the metric sources

An ongoing task is figuring out the cause for differences between the xroot and system metrics, such as in the data in/out plots below*. The differences seem too large to just accounted for by CEPHFS traffic - unless there's unexpected activity.



*external/internal rates are stacked.

FIN!