



Contribution ID: 72

Type: **Presentation**

Dismantling the Monolith

Tuesday, 28 January 2020 17:00 (20 minutes)

Every year at CS3 we all come together to talk about the things we've built and how they've grown - more users, more files, more shares, more storage used than in past years, more features we've added. Last year, we introduced a particularly interesting feature to the AARNet CloudStor ecosystem: S3 gateways as a means of convenient, high-speed data transfer directly to our backend storage. This year, we're going to talk about how that effort led us to a turning point in CloudStor's history, and where we're going from here.

Our earlier (per 2019) deployments of the S3 gateways revealed a critical issue; resource contention between different access pathways to the backend storage resulted in multiple outages across the entire CloudStor ecosystem. This experience highlighted a different, perhaps more pressing problem - that we'd inadvertently built a monolith, and that one component of the system could take out everything else.

Through the next few months, we addressed this issue by splitting worked to split out the S3 backend storage from the CloudStor Prime backend storage, and going one step further, we decided to shard the new environment we were building. In the new model, institutions or groups of institutions are allocated to a separate storage shard, greatly reducing the blast radius of an outage - even if one shard is experiencing issues, customers on the other shards remain unaffected. Additionally, leveraging both Kubernetes as well as the new QuarkDB namespace for EOS, we've managed to cut outage/upgrade downtime from approximately an hour down to a matter of seconds.

This new model has worked so well that we're looking to apply it to the rest of CloudStor, which will be a significant amount of work, but worth the effort. It's been a great run, but perhaps it's time to dismantle the monolith that CloudStor has become, and transform it into something more robust, scaleable, and modular.

Primary author: Ms CHUA, Crystal (AARNet Pty Ltd)

Co-authors: Mr JERICHO, David (AARNet Pty Ltd); Mr USHER, Michael (AARNet Pty Ltd)

Presenter: Ms CHUA, Crystal (AARNet Pty Ltd)

Session Classification: Scalable Storage Backends for Cloud, HPC and Global Science

Track Classification: Scalable Storage Backends for Cloud, HPC and Global Science