

Fast for the Curious: Speeding up CloudStor's research storage performance

Gavin Kennedy¹, Chris Myers¹

¹Gavin Kennedy, ¹Chris Myers, ¹David Jericho, ¹Crystal Chua, ¹Michael Usher, ¹Michael D'Silva

¹AARNet, Australia,

gavin.kennedy@aarnet.edu.au, chris.myers@aarnet.edu.au, david.jericho@aarnet.edu.au,
Crystal.Chua@aarnet.edu.au, michael.usher@aarnet.edu.au, michael.dsilva@aarnet.edu.au

Summary

In this presentation, Gavin Kennedy and Chris Myers from AARNet's Cloud Services team will give a joint presentation on the research and development effort that has gone into lifting CloudStor(ref1) from a ubiquitous file sync and share tool, to a high speed data storage and collaboration platform, and our roadmap for lifting the performance take advantage of AARNet's dual Gigabit per second data superhighway and to match the demands of the current and future data generators and their collaborators who need to access, move, share and consume high volume research data.

Abstract

Marking its 10 year anniversary this year, we will present on how AARNet's CloudStor has evolved from an initial file sharing service (the popular CloudStor FileSender(ref2)), to a ubiquitous file sync and share tool(ref3), and then on to become a high speed data storage and collaboration platform trusted by Australian universities and research and cultural institutions. In this time we have had over 70,000 registered users across over 400 domains collectively storing, sharing and moving hundreds of petabytes of data. In addition AARNet have implemented valuable services on top of CloudStor to support research collaborations, including collaboration tools and analysis tools. We have also worked to ensure research data is accessible to compute services by providing gateways and APIs, as well as putting in interconnects between CloudStor and HPC infrastructure including NCI, as well as interconnect with commercial clouds. All of this effort and investment has been driven by AARNet's vision of a globally networked data-sharing ecosystem, in which CloudStor is a consistent and reliable connector.

But much of our storage infrastructure has operated at modest speeds, often below 1Gbps for users accessing CloudStor services across 10Gb AARNet links. As universities start to adopt AARNet's 100Gbps service (ref4), the disparity between what our network can deliver and what our storage can achieve is even greater. The reasons for this are many and lie in a mix of organic growth of the service and evolution of the applications that CloudStor is built upon. But a confluence of advancements in the applications and operational software, combined with improved hardware, has given us the opportunity to deliver a service that pushes the boundaries of speed.

In this talk we will go through the improvements that we have made to CloudStor, including:

- Re-architecting (sharding) CloudStor's storage to create purpose built, independent CloudStor storage instances;
- Adopting Kubernetes to orchestrate our docker containers to deliver a faster and more innovative microservice orientated infrastructure;
- Create a faster network pathway to Cloudstor nodes right down to the container with host based routing with improved security for container with service based overlay networking and multi-path BGP;
- Implementation of an FST service per disk architecture to increase write concurrency and decrease file locking;
- Working with CERN to upgrade our EOS file storage system to utilize their QuarkDB fast storage cluster;
- Working with ownCloud on their Nexus project to implement the CERN developed REVA storage layer between ownCloud and EOS.
- Rolling out NVME storage devices as we upgrade our storage platforms.

A lot of this sounds like devops buzzword bingo, but the end result should be considered by all users of research data infrastructure. Our goal is to deliver an 80Gbps ingest rate to our on-net customers, and to deliver an internal transfer rate of between 100Gbps and 200Gbps, so that services like SWAN Jupyter Notebooks will interoperate efficiently with user research data. This performance also becomes critical for the sharing of data with international collaborators, using the cloud services mesh approach being developed by an EU consortium partnering with AARNet.

Finally we will present our Cloud Services roadmap and identify the key new services and service enhancements that we are planning, as well as the benefits and opportunities they deliver.

REFERENCES

1. <https://cloudstor.aarnet.edu.au/>
2. AARNet takes large files to the cloud, PC World Australia, 19 November, 2009.
https://www.pcworld.idg.com.au/article/326880/aarnet_takes_large_files_cloud/
3. Reid, Alex. AARNet's cloudstor+ cloud storage initiative: Cloud-based value-added services [online]. *Journal of Telecommunications and the Digital Economy*, Vol. 2, No. 2, Jun 2014: [35.1]-[35.8]. Availability:
<https://search.informit.com.au/documentSummary;dn=422791934454423;res=IELBUS> ISSN: 2203-1693. [cited 20 Jun 19].
4. <https://news.aarnet.edu.au/university-of-queenslands-new-100-gigabit-per-second-aarnet-data-superhighway-goes-live/>