



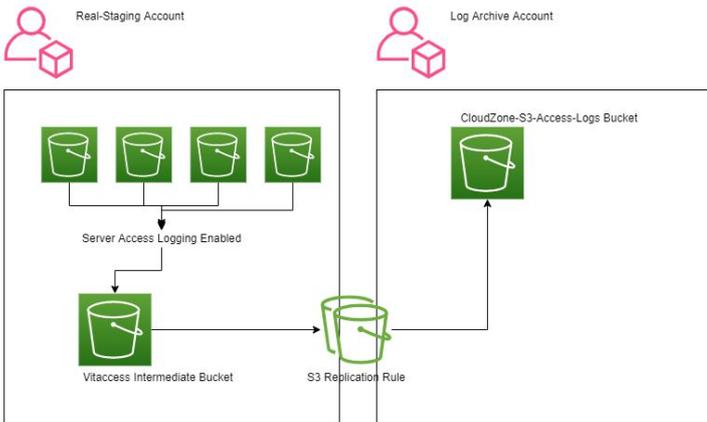
Vitaccess is a multi-award-winning digital healthcare research scale-up and strategic consultancy. They offer a world-class real-world evidence (RWE) data portfolio and data science research via MyRealWorld®. Vitaccess are a digital business born in the cloud, leveraging cloud technology using AWS native services.

Cloud Challenge

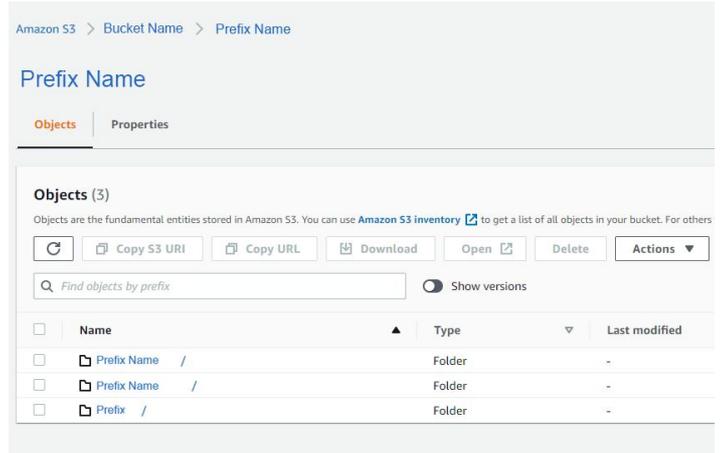
Vitaccess recently made use of the **CloudZone** Log Archive centralised Amazon S3 Access Logs Bucket. This bucket, until now, has been used to log the server access logs from the main **CloudZone** Log Archive centralised log bucket. But **Vitaccess** needed, in addition, to organise & centralise their access logs from multiple Amazon S3 bucket locations from potentially multiple different regions and AWS accounts under their Organisation. Finally, they needed to maintain the solution cost efficient ensuring they avoid duplicating data.

CloudZone Engagement

Through use of the **CloudZone** Amazon S3 Replication feature and setting up an intermediate bucket in other accounts, **Vitaccess** were able to use the Log Archive Amazon S3 Access Logs Bucket as their main centralised location for access logs for multiple Amazon S3 Buckets throughout their organisation.



Vitaccess created an access logs bucket within one of their AWS Accounts, which will act as an Intermediate bucket. The purpose of this bucket was to use it as a local location for Amazon S3 server access logs from other buckets in this account. By local location, we mean other Amazon S3 Buckets only within this AWS account. We can simply point all of their server access logging location to the same intermediate bucket.

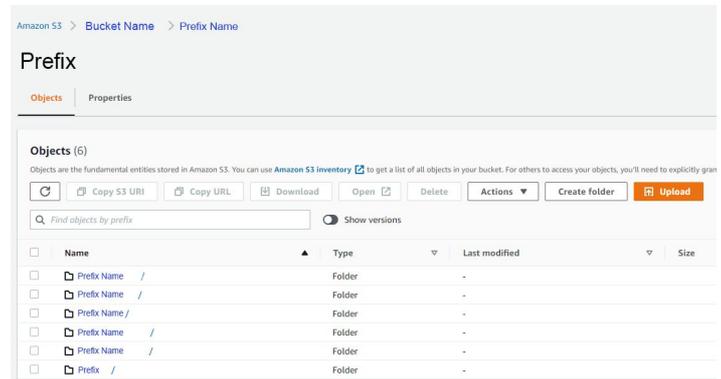


This would work nicely with the prefix added from each server access logging setup applied from **Vitaccess**, so that when you list the contents of the intermediate bucket, you would first see the prefix, this is to help identify where each log originated from. From an auditability perspective, you can easily locate server access logs by prefix this way.

This solved the first part of what **Vitaccess** needed, but they also had to centralise their process. Through the power of **CloudZone**'s Amazon S3 Replication feature, we are now able to replicate the contents of this logging bucket into the Log Archive Account's Central Bucket.

Tag Key/Value:
S3Replication = Target-Bucket-Name-Here-To-Replicate-Into

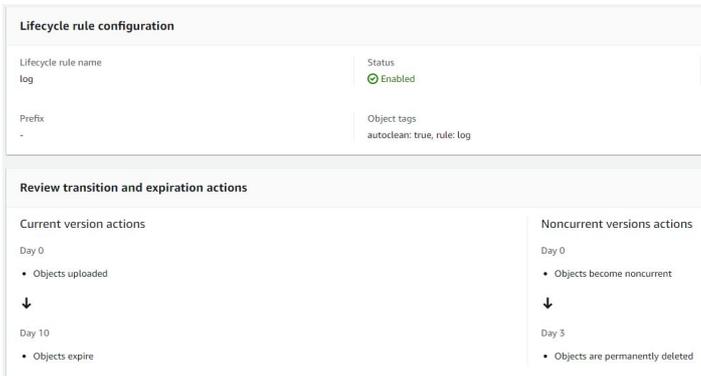
With this in place after the **CloudZone** feature was enabled, we send the contents of this bucket to the destination listed as the Tag Value.





The last part of the challenge was to ensure the solution was cost efficient by avoiding duplicate data. This may happen due to the growing amount of data in multiple buckets. If this solution was replicated across many AWS accounts within the organization, we could have a lot of access log data stored in both the intermediate buckets, and replicated to the central S3 Access Log Bucket.

To combat this, a lifecycle rule was placed onto the intermediate bucket, which is set to expire the current version of any objects across the entirety of the intermediate bucket after they reach 10 days old.



Objects deleted this way are not removed from the central log access bucket post-replication. As the current AWS Documentation on this particular combination of action confirms:

“However, if Amazon S3 deletes an object due to a lifecycle action, the delete marker is not replicated to the destination buckets”

[Source Documentation Article](#)

Tools Used

- Amazon S3
- Amazon S3 Replication feature
- S3 lifecycle rule
- Landing zone Log Archive Account: Log Archive Account’s Central Bucket



CloudZone has provided us with a simple and secure way of replicating our application logs into a single secure account without having to worry about constructing the replication policy or having to worry about the AWS IAM policies.

Stephen Rayner, Application Development Lead



Realised Benefits

The final result provided **Vitaccess** with the capability of organised, centralised access logs from multiple Amazon S3 bucket locations from potentially multiple different regions and AWS accounts under their **CloudZone** Organisation, by just configuring an intermediate bucket, a tag added to it, and a lifecycle rule as an optional configuration piece which will ensure storage costs are kept to a minimum.

