On premises file and block data migration to AWS storage services

CHALLENGE

BILL needed to migrate hundreds of millions of files and directories from on premises to the cloud.

This massive dataset needed to appear to an Amazon ECS container as a hierarchical filesystem.

We could not re-tool the application to utilize Amazon S3 or alternate storage types and required a filesystem storage service.

Most of the data could be treated as read-only. However, portions of the data set needed to be accessible for read and write operations.

SOLUTION

Read-only data: AWS DataSync was used to copy hundreds of millions of files and directories from on premises to Amazon Elastic File System (EFS).

Read/write data: Block-based SnapMirror replication was used between an on-premises NetAPP ONTAP storage array and Amazon FSx for NetApp ONTAP.

File access: ECS containers mounted EFS and FSx for ONTAP volumes.

A validation tool was developed to ensure the read only source was not modified during the migration.

OUTCOME

- ✓ **Block-based replication is advised when SnapMirror is available.**
- ✓ **DataSync scaled out to parallelize the file migration with multiple tasks.**
- ✓ **Validation tool verified that the source file system remained read-only, pre and post migration.**

KEY SERVICE(S): AWS DataSync, Amazon EFS, Amazon FSx for NetApp ONTAP, Amazon ECS

© 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.