



Transforming financial security: Powering disaster recovery solutions for a top 3 US bank.

About the top 3 US bank

The top 3 US bank is an American multinational investment bank and financial services holding company headquartered in North Carolina. The organization is a global leader in wealth management, corporate and investment banking, and trading across a broad range of asset classes, serving corporations, governments, institutions, and individuals around the world. The company serves clients through operations across the United States, in U.S. territories, and in approximately 35 countries.

The organization is focused on



Disaster recovery

Objective and challenges

The organization uses an on-premises Cloudera CDP cluster for one of their analytics platforms and a secondary CDP environment for Disaster Recovery (DR). The company has established a 15-minute service level agreement (SLA) for the recovery time objective (RTO) and recovery point objective (RPO) of the DR environment. Their current implementation is not able to meet the SLA requirements.

To achieve the SLAs, they need a solution that can replicate and keep synchronized more than one million transactions per minute. This high volume of transactions means that manual or batch-based tools would not be a viable option, and this has opened the opportunity for IBM Big Replicate, which utilizes

Cirata LiveData Migrator to support real-time data replication. The initial process is replicating data between two on-premises data centers and there is future opportunities to also replicate the data to a public cloud environment.

Solution

Cirata LiveData Migrator (LDM) supports complete and continuous replication of data sets at any scale. With zero disruption or impact to the existing production system, LDM migrates the initial data sets with a single pass through the source storage, eliminating the overhead of repeated scans while also supporting continuous replication of any ongoing changes as they occur.

The organization reviewed multiple approaches before selecting IBM Big Replicate / LDM. The other approaches that were evaluated included DistCp, Replication Manager (which leverages DistCp), and other open-source tools such as Apache NiFi. None of the alternatives were able to meet their SLA requirements because the alternatives operate in a batch/scheduled fashion and do not replicate ongoing changes in real time. This means the alternatives cannot guarantee that changes made within the last 15 minutes have been successfully replicated to the DR environment.

Other critical factors in selecting IBM Big Replicate / LDM were as follows:

- Ease of use of the solution (simple deployment, configuration, and management).
- Ability to replicate Hadoop data as well as Hive metadata.
- Ability to set up and begin replicating data very quickly.