

Large-scale **DATA MIGRATION WITHOUT DOWNTIME**

Migrated large amounts of data from legacy storage hardware to the new storage without disrupting business operations or risking data loss.

— The Client

The client is one of the world's largest financial institutions, serving individuals, small- and middle-market businesses and large corporations with a full range of banking, investing, asset management and other financial and risk management products and services.

— **Executive Summary**

We migrated large amounts of data from old storage hardware to the new storage without disrupting business operations or risking data loss using data management solutions.

We also ensured that:

1. The hardware is maintained and upgraded for optimum performance and cost-effectiveness
2. There are no outages or productivity impacts associated with aging hardware

— Challenge

Even if there is a great system in place that works fast and has low latency, moving petabytes of data just takes time, and involves a great deal of risk of losing data.

The larger the scale of the data migration, the more opportunity for issues along the way. We had to find a way to do the large-scale data migration without any downtime, outside business hours.

— Approach

We used storage-native, third party tools to migrate the large amounts of data. Our approach also allowed applications to continue modifying a source system's data without causing issues between the source and the target.

By engaging with Wavelabs, the client was assured of a seamless hardware refresh process. We built a specialized team for full migration lifecycles around block (SAN) and file (NAS) storage systems, including developers and database experts to implement and support a custom data warehouse.

Project management

to organize and co-ordinate data migration processes

Change control

to manage documentation and approvals

Technical engineers

for development, DBA, analytics, and reporting

Storage engineers

for data replication, and configuration of new systems

— Solution

Storage systems have an average useful lifecycle of four years. Beyond that, the chances of hardware failures increase, as do the maintenance costs. It is advisable to keep hardware updated rather than risk downtime and data loss due to outdated hardware. Each time the system is upgraded, the data must be migrated to the new system.

We designed the data migration process in a way that was frictionless, enabled business continuity and ensured no risk of data loss. We executed this in 4 parts:



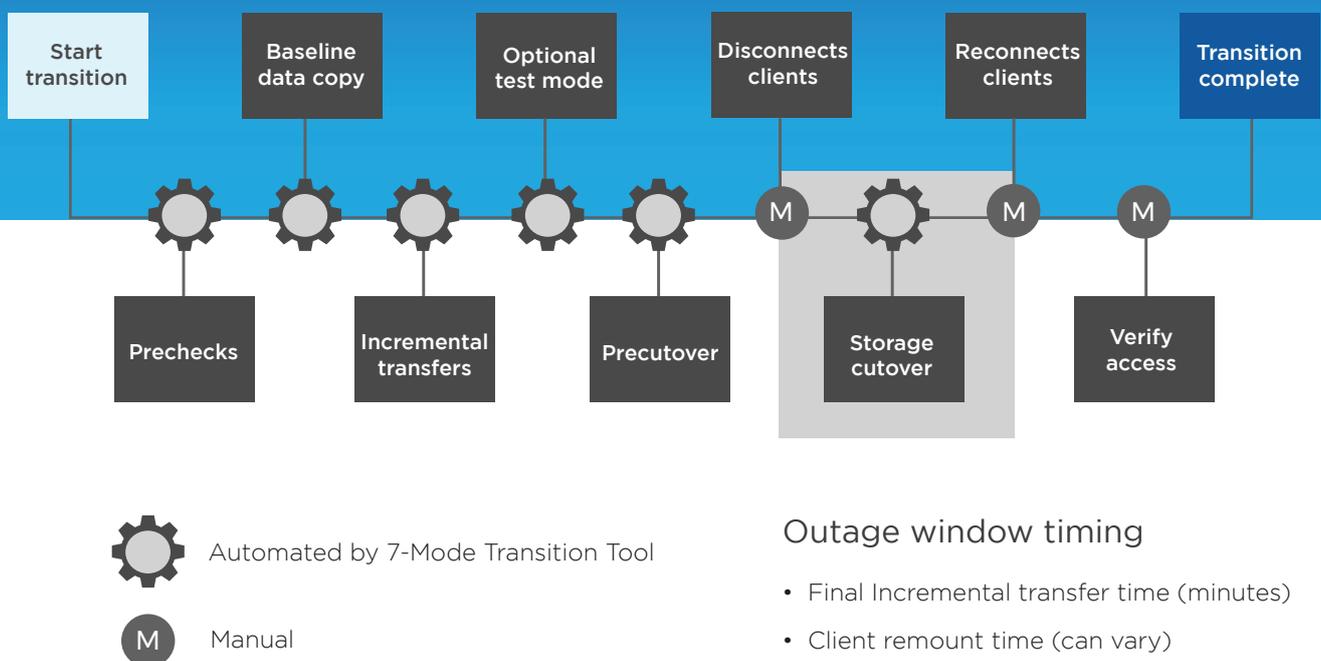
Tech Stacks

1. Snapmirror Technology (License from NetApp)

```
cluster1 :: snapmirror>

about          break          config-replication
create        delete        initialize
initialize-is-set list-destinations modify
policy        promote       protect
quiesce       release       restore
resume        resync        set-options
show          show-history  snapshot-owner
update        update-is-set
```

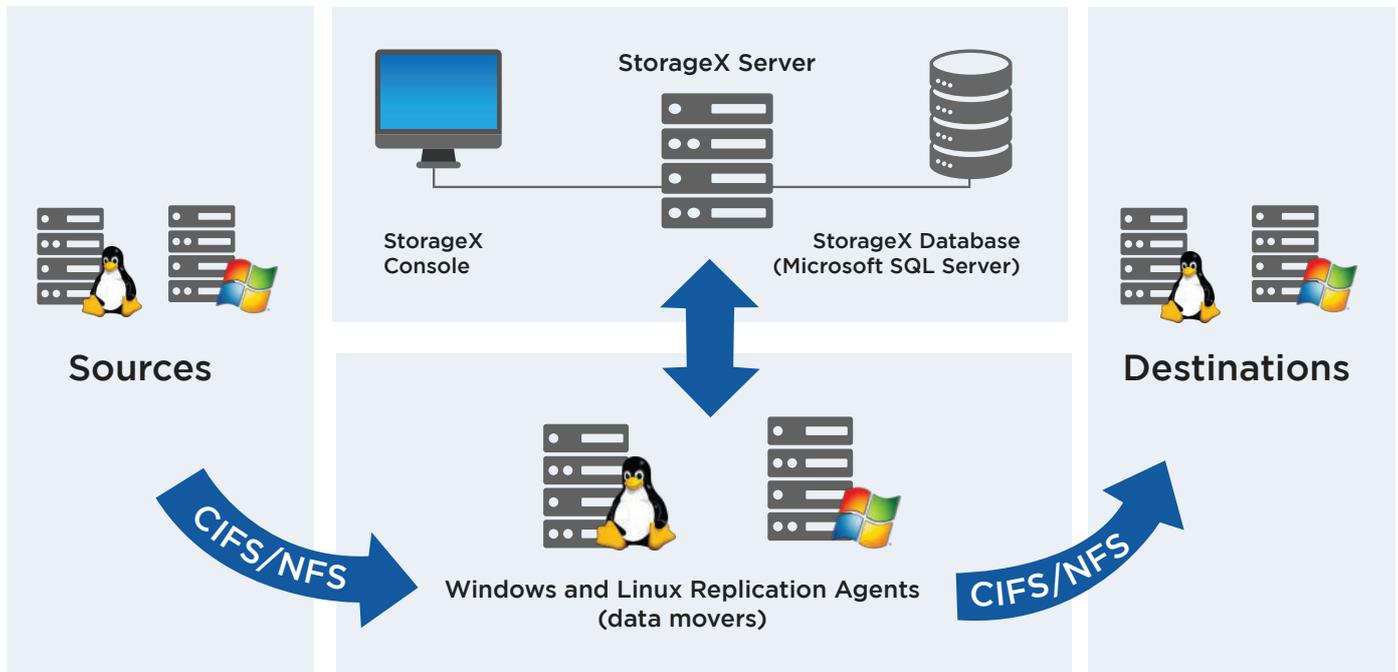
1.1 Screenshot of commands used for Snapmirror CLI Migration



1.2 Snapmirror migration flow

- StorageX tool is used to analyze, move, manage and modernize the infrastructure to drive cost reduction, risk mitigation and policy automation.

The StorageX architecture consists of the StorageX server, StorageX Console, the Storage database, which is a Microsoft SQL Server database, and StorageX replication agents, or data movers.



2. Storage X Tool Architecture

— How did partnering with Wavelabs help?

Having a dedicated team with specialized developers engage with the client ensured that the data migration was executed on time, and without any impact on productivity.

- We ensured proper hardware lifecycle management, avoiding the risks and costs of aging tech.
- We ensured the migration was carried out outside of business hours, eliminating system downtime and conflicts with business priorities, and ensuring that the data was accessible when needed.