

WHITEPAPER

Building a Migration Factory for Google Cloud Platform to Leverage Hybrid and Multi-Cloud Environments



Business Challenge

Enterprises are increasing their use of cloud storage to keep up with the exponential growth of unstructured data, with many companies using multiple cloud providers. While using multiple cloud providers is reasonable for a number of different objectives (risk mitigation, technical requirements for applications, sovereignty, price, etc.), it also adds a significant amount of complexity, often creating a silo effect that limits data movement and diminishing its value to the company. A recent survey¹ of over 200 enterprise IT managers summed up the situation:

“Most companies have a cloud management vision but they do not have a cloud management strategy because they do not have the right tools to deliver their vision”

This is where a cloud migration factory comes in. The efficient movement of unstructured data on an enterprise scale requires a special combination of subject matter expertise, smart technology that leverages automation, analytics to drive better outcomes faster, and a proven architecture that handles the scale and performance requirements necessary to support today’s hybrid cloud environments.

This whitepaper is a high level overview of how to build a cloud migration factory for the Google Cloud Platform to enable the more efficient movement to deliver the right data, to the right location at the right time for hybrid cloud environments.

Cloud Migration Factory Overview

The Process

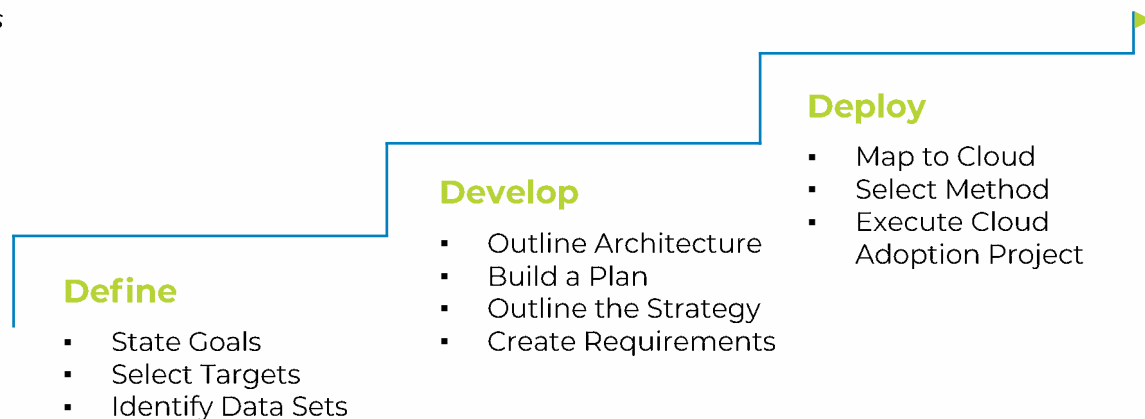


Figure 1– Migration Factory Process

- Define the goals of this cloud migration project and then leverage them to develop the optimal strategy to deliver those goals
- After you understand what the migration strategy, you can start to develop the requirements necessary to support that strategy
- Working with the various stakeholders, identify target data scan and analyze, develop key metrics queries to apply to the target metadata for analysis, and create policies based on classification and tags

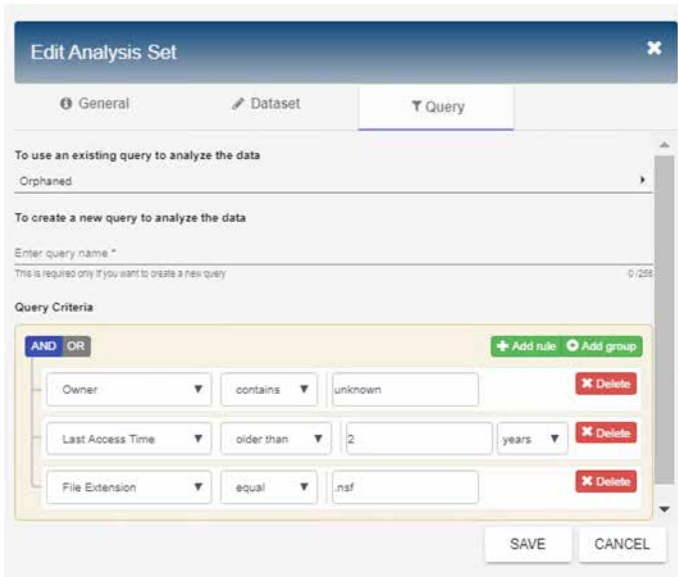


Image 1 – Selection of Source Data

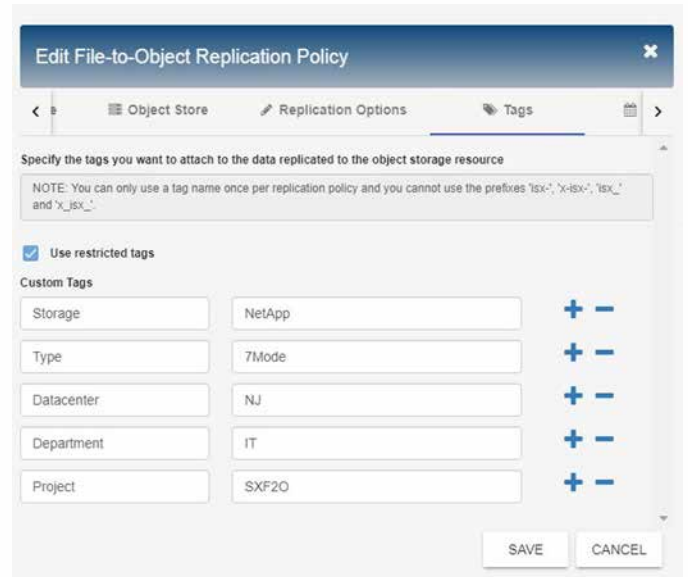


Image 2 – Custom Tagging

- Once you have these, you create a query against the metadata to classify data and tag it for migration.

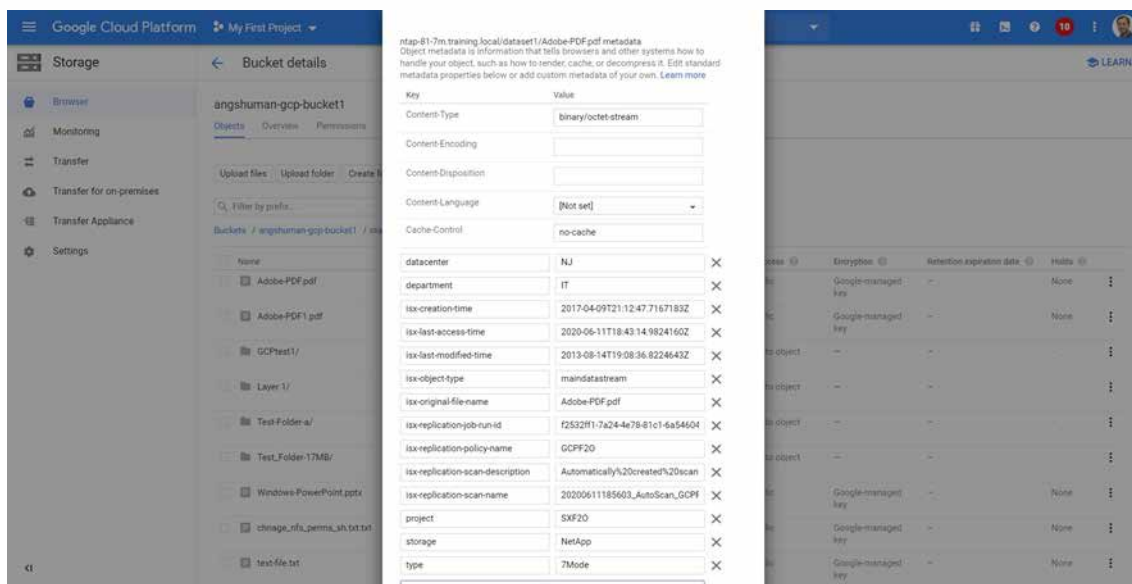


Image 3 – Classification Tagging within GCP

- This classification will be mapped to a migration policy that automates the migration process.
- Within StorageX, there are 3 different methods for moving it to the cloud with policy based movement
 - Archive – move the primary copy of unstructured data from a NAS (NSF, CIFS, SMB) to an object store (S3, Azure Blob).
 - Migration – File to File (F2F) migration from one NAS to another
 - Transform & Sync – File to Object (F2O) migration from NAS to an object store (S3, Azure Blob)

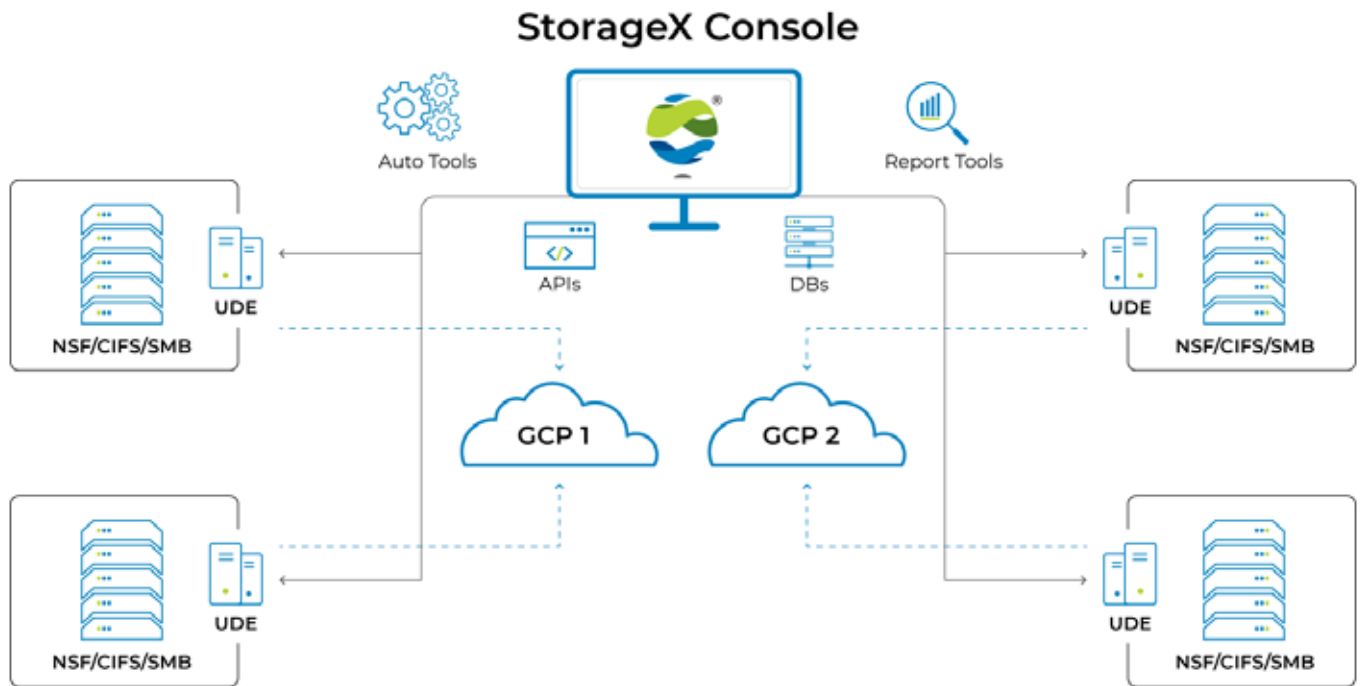








Figure 2 – Example of Migration Factory Architecture

Based upon the information gathered throughout the process, you can assemble the components to deliver the proper migration services.

The StorageX architecture is built using a software defined solution that easily adapts to each project's requirements and allows you to easily scale and move the UDE's (Universal Data Engines) to where they can process the information most efficiently.

The flexibility of the architecture allows you scale up or scale down performance as needed. The UDEs can be QOS based on business requirements to ensure optimal use of network bandwidth. This in conjunction with its multi-threaded architecture, hundreds of policies can be run simultaneously to meet your cloud adoption timelines.

- Define the internal management and reporting processes to be used
- Coordinate scheduling with various stakeholders to minimize conflicts and develop calendar
- Define migration schedules
- Final plan review and approval

Category	Component Symbol	Description
Management		StorageX Console – Single pane of glass management hub that manages all StorageX
		Automation Tool – easy API integration with all major automation tools such as Ansible, Chef, Puppet
		Reporting Tool – easy API integration with all major reporting tools such as Service Now and VRealize
Infrastructure		Database – Analytics leverages an Elastic Database to support metadata
		API – Full set of RESTful APIs provides easy integration with other management tools for customization and delivery of higher SLAs
		UDE – Universal Data Engines provide the migration infrastructure supporting the analytics scan processing and all data movement functions: archive, migrate and transform.



Migration Factory Use Cases

On Premise to public cloud

Data Movement Method	Use Case Description
Migration – F2F	<p>Standard migration of unstructured data files from on premise NAS (NFS/CIFS/SMB) to NAS Storage tier within public cloud.</p> <p>This is used when the goal is to data from on prem to the cloud, eliminating on-prem access.</p> <p>Files are migrated using 3 phase approach:</p> <ul style="list-style-type: none">▪ Initial Replication Phase (typically 85% of data is replicated during this phase)▪ Incremental Replication Phase (additional 10% of data is replicated including incremental changes to existing replicated files)▪ Final Cutover Phase (access to NAS is suspended, incremental replication captures remaining 5% of data, and all access controls are mapped to the new location)
Transform & Sync – F2O	<p>Replication and synching of data from an on premise NAS (NFS/CIFS/SMB) to an S3 Object Storage tier within public cloud.</p> <p>This is used when the goal is to provide a second copy of the on premise file within a public cloud Object Storage tier.</p> <p>Transform and Sync maintains a “File Path is Key” structure within the public cloud Object Storage target so that data appears to maintain original folder structure for the easy location of data within the public cloud bucket.</p> <p>Additionally, classification tagging can be applied to add key search and data management functionality to the public cloud Object Storage target.</p>
Archive – F2O	<p>This method provides a movement of the primary copy of the data from the on premise NAS (NFS/CFS/SMB) to a public cloud Object Storage target for long term retention and resiliency.</p> <p>This is used when the goal is to move inactive and other data off of expensive, on premise Tier 1 storage and cull backup data sets to active data that really does require backup protection for versioning and archive data to a public cloud Object Storage Tier that does not.</p>

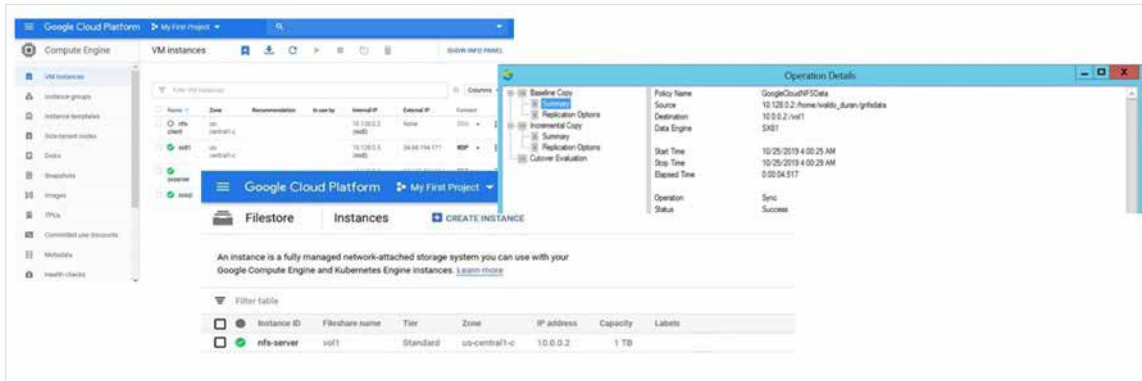


Image 4 – F2F Migration to public cloud

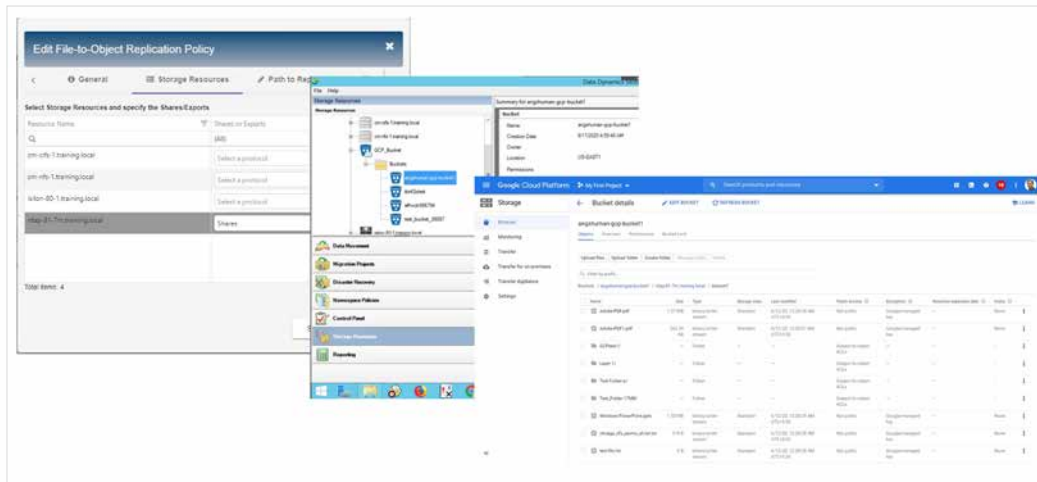


Image 5 – F2O Transform & Sync to Google Cloud Platform

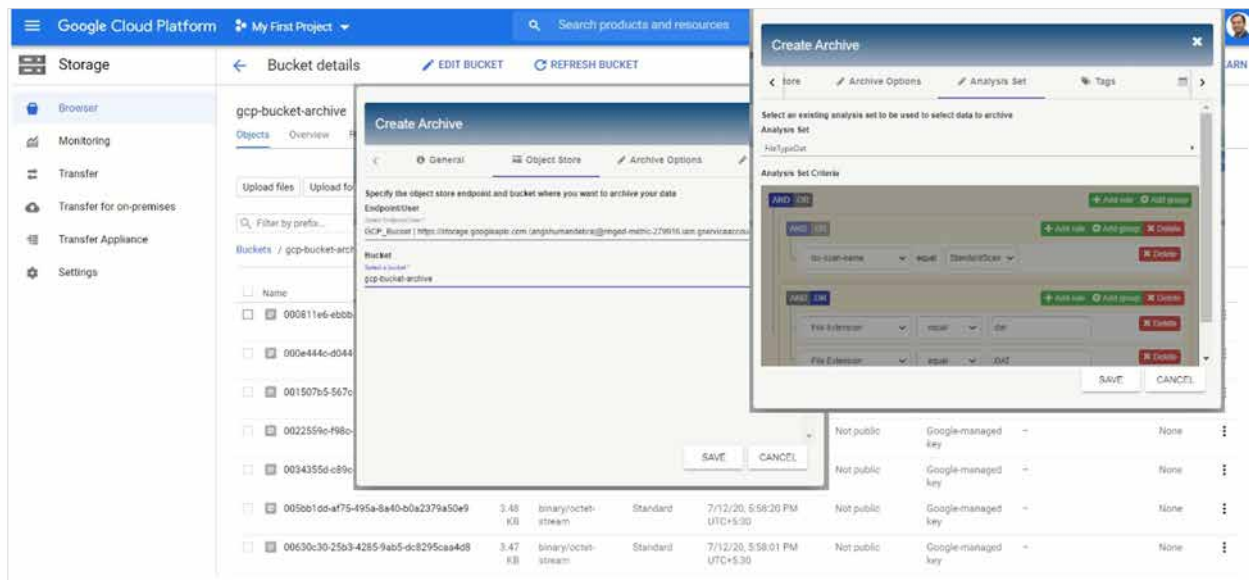


Image 6 – Archive to GCP's Cloud Storage

Benefits of a StorageX Migration Factory for Hybrid Cloud Environments

Benefit	Description
Analytical Driven Process	Many of the issues related to hybrid cloud environments concern moving the wrong data to the wrong location. Analytics takes the guesswork out and utilizes proven, efficient methods to identify the right data and move it to the right location within the hybrid cloud environment
Scalability	A software solution that utilizes a scale out architecture that allows you to spin up or down depending on the current migration requirements of the factory. Our UDE design allows the factory to easily move data movement and analytics services to where they are needed, when they
Performance	The solution also provides the ability to scale up depending on what performance is required to
Use Case Coverage	Broad use case coverage makes the StorageX migration factory more valuable as an integral part of data management for all unstructured file data within the hybrid cloud environment.
Integration	StorageX provides a robust set of APIs to drive integration with key orchestration and reporting tools that are a must for enterprise deployments. These integrations deliver a true factory functionality by automating many of the tasks involved to drive efficiency and reduce human
Data Mobility	For hybrid clouds to be efficient the management of data requires enterprise performance when moving massive amounts of data without disrupting day to day operations. A StorageX migration factory provides the right infrastructure to move that data intelligently.
Licensing	StorageX subscription based capacity license is designed to support a hybrid cloud environment, providing a true cost basis that does not limit your ability to scale up or down

Summary

The word 'cloud' has become ubiquitous in the digital world we live in but very few enterprises are leveraging it in a manner that drives the most value out of their data. To be more effective and take advantage of the services that a hyper scaler such as public cloud can offer, there needs to be an understanding of the existing data, both its metadata and content.

By empowering enterprises with this information, the value from the use of the cloud is garnered day one! This in conjunction with a need to meet the scale required for an infrastructure transformation to the cloud is extremely vital.

Enterprises today sit with tens of petabytes of data, and the ROI from a transformation must be realized for senior executives to continue to support the initiative. Showcasing cost savings whilst leveraging a better understanding and usage of the data using public cloud's native analytics capabilities will create a snowball effect on the speed with which the transformation transpires.

The StorageX migration factory enables, empowers and automates the execution of this journey and is critical to the success of a digital transformation to the cloud.

Evaluator Group Survey, "Trends in Multi-Cloud Data Management", November 2019

Data Dynamics, a global leader in enterprise data management, stands at the forefront of the industry-wide shift towards Digital Trust & Data Democracy. Trusted by 300+ organizations, including 25% of the Fortune 20, the company is recognized for its commitment to creating a transparent, unified, and empowered data ecosystem. Whether addressing data risk, privacy, sovereignty, optimization, sustainability, or facilitating seamless, policy-driven data migration across hybrid and multi-cloud environments, the company is ushering in a new era where data ownership, control, & actionability reside with the data owners.



Contact Sales

Book a Demo