

Best Practices for Orchestrating Hybrid- and Multi-Cloud Architectures

Note: This document was reviewed by Gartner

Content:

Market & Trends of the Cloud

Cloud Data Ecosystem: Data Management in Multi- and Hybrid-Cloud Architectures

Role of Data Integration (Data Virtualization) in Cloud Architectures

Trends Are Leading to More Use Case Centric Approaches

Use Case: Clouds for Financial Services Conclusion

Cloud platforms have become immensely popular in recent years and are now an integral part of the IT strategy for most enterprises. Some enterprises believe in a single cloud and fully rely on solely one cloud platform infrastructure and build their architecture around it. But the dependence on a single cloud platform comes with several risks that will be discussed later. The majority of companies are rather moving to a hybrid- or multi-cloud architecture approach in order to benefit from flexibility, security, and potential cost optimization that these modern architectures offer.

The cloud data ecosystem¹ is an essential part of all cloud strategies and here are three reasons why it actually needs more attention:

- To build the architecture that fits your specific needs, not the one that the cloud platforms of your choice impose on you.
- To fill the gaps that cannot be filled solely by one single cloud provider such as data integration, metadata management, data governance, and data quality.
- To avoid another silo built in the infrastructure of the cloud platform.

Following, we will elaborate further the importance of a healthy cloud data ecosystem and how you can benefit from it.

¹ The cloud data ecosystem is an interconnected infrastructure that combines numerous data sources and adds value through the interaction of the processed data. It enables different cloud architectures and attributes more importance to the cloud.

● Market & Trends of the Cloud

The cloud market is dominated by the major providers Amazon, Google, and Microsoft with the respective Infrastructure-as-a-Service (IaaS) platforms Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform. A recent Gartner study showed that 81% of enterprises stated that they plan with more than one cloud provider in their future roadmap. While 19% reported to use only one public cloud, 40% stated 2 public clouds. 41% even said three or more different clouds. The adoption for a multi- or hybrid-cloud architecture already emerges.

When deploying advanced business intelligence, analytics, and data science platforms, public cloud architectures are preferred (27%). However, there is still a long way to go before such modern platforms will have completely moved to the cloud. A staggering 34% of companies state that they want to keep their data warehouses and/or data lakes "on premise" for the time being (in the medium/long term).

● Cloud Data Ecosystem: Data Management in Multi- and Hybrid-Cloud Architectures

Coming back to the question single cloud vs multi-cloud, why do multi-cloud architectures have decisive advantages over single clouds:

- Mainly to avoid dependencies on individual providers.
- To address security aspects or regulatory reasons can make the use of hybrid- or multi-cloud indispensable.
- To utilize a "best-of-breed" approach and thus use the best available solution for each use case making your company more flexible.

The use of multi- and hybrid cloud architectures certainly comes with various challenges. And a critical success factor is a well-thought-through strategy for the entire cloud data ecosystem as:


- Most cloud platforms miss certain important data management capabilities in their overall infrastructure e.g. data virtualization, data fabric, data governance, data quality, and metadata management.
- Most cloud platforms like Google Cloud Platform don't have a neutral data virtualization layer, yet.

Independent Software Vendors (ISV) specialized in the different areas of data management such as data integration, data quality, data governance, master data management (MDM), and metadata management help to plug in the gaps of the cloud platforms and ensure independence from the specific architecture choice driven by the cloud platform.

Multi-Cloud Adoption

Respondents Using Public Cloud

Number of Public Cloud
Platforms Used



Customer example:

- One of the largest property development and investment companies in the UK built a multi-cloud architecture with AWS and Azure to gain more flexibility and agility.
- The 5th largest pension fund world wide benefitting from a modern architecture in a hybrid environment with Snowflake.
- One of the biggest reinsurers worldwide that built an agile hybrid-cloud architecture with Snowflake to meet today's and tomorrow's increasing regulatory requirements and digital needs.

● Role of Data Integration (Data Virtualization) in Cloud Architectures

Data integration is the process of bringing together data from multiple disparate sources and plays a crucial role for both architectures, multi- as well as hybrid-cloud. Depending on the underlying use case, different types of data integration are the better fit.

ETL (extract, transform, load) and ELT (extract, load, transform) i.e. extracting, cleansing, and transforming the data, enable data movement and ingestion in the cloud infrastructure and thereby provide the following benefits:

- Cost considerations/incentives
- Make use of the various capabilities and services that the different cloud platforms provide
- Bypass restrictions or reservations posed on specific cloud use (e.g. company policy)

Data virtualization (DV) plays a crucial role in the implementation of any cloud architecture. Neither the path to the cloud nor the subsequent operation in the cloud is conceivable without data virtualization as it brings:

- Flexibility
- Scalability
- Cost optimization
- Change management
- Semantic layer delivery

By using data virtualization, data integration is extremely simplified. Data does not have to leave the storage location and companies can retain agility and flexibility of their „data ecosystem“: the data is thus more secure and it is easier to change providers. Furthermore, companies can also work with sensitive data in a GDPR compliant manner. New data sources can be connected easily and new requirements and ideas can be tested quickly.

● Trends Are Leading to More Use Case Centric Approaches

As use cases often evolve over time, you may not want to restrict yourself with one data integration approach while you conceptualize/build your cloud data ecosystem. The Logical Data Warehouse (LDW) is one concept that was primarily driven by the fact that individual data integration technologies such as data replication (ETL/ELT) and data virtualization alone could not serve all use cases. The LDW's hybrid approach allows, on the one hand, to move data where necessary and, on the other hand, to combine data without physically moving it. Thus, the LDW not only enables a fully integrated data landscape, but also strengthens the transparency, governance and security aspects for a compliant system. This approach enables companies to implement flexible and sustainable data strategies for increasingly complex environments.



The latest trends such as Data Fabric, Data Mesh, Unified Data and Analytics Platform, etc. also follow the fundamental idea of the Logical Data Warehouse: the understanding that the business use cases cannot be served with a single monolithic approach such as a data warehouse or data lake, but that you have to combine several technologies and distributed architectures, also across clouds to achieve good results.

They all address the following key market trends:

- The data landscape is becoming increasingly complex
- Agility is becoming increasingly important
- Use cases are becoming more diverse.

The integrative data landscape of these concepts embrace hybrid- and/or multi-cloud architectures to improve data visibility and insights.

Data Virtuality already addresses many aspects of a data fabric. By combining data virtualization and data replication (ETL/ELT) with various data management features such as metadata management, AI-based optimizers, etc., it conceptually covers 60 to 70% of the data fabric.

Also the requirements of the underlying data platform of a data mesh can be fulfilled with Data Virtuality. This platform enables a fast development of data products, allowing the data to gain commercial value.

The benefits are profound:

- Just one solution (instead of several) which leads to
 - Less development and maintenance work
 - High performance also for large data sets
- Specialized staff isn't needed as everything is SQL based
 - ⇒ Customers report 5-times faster time-to-market and total cost savings of up to 80%

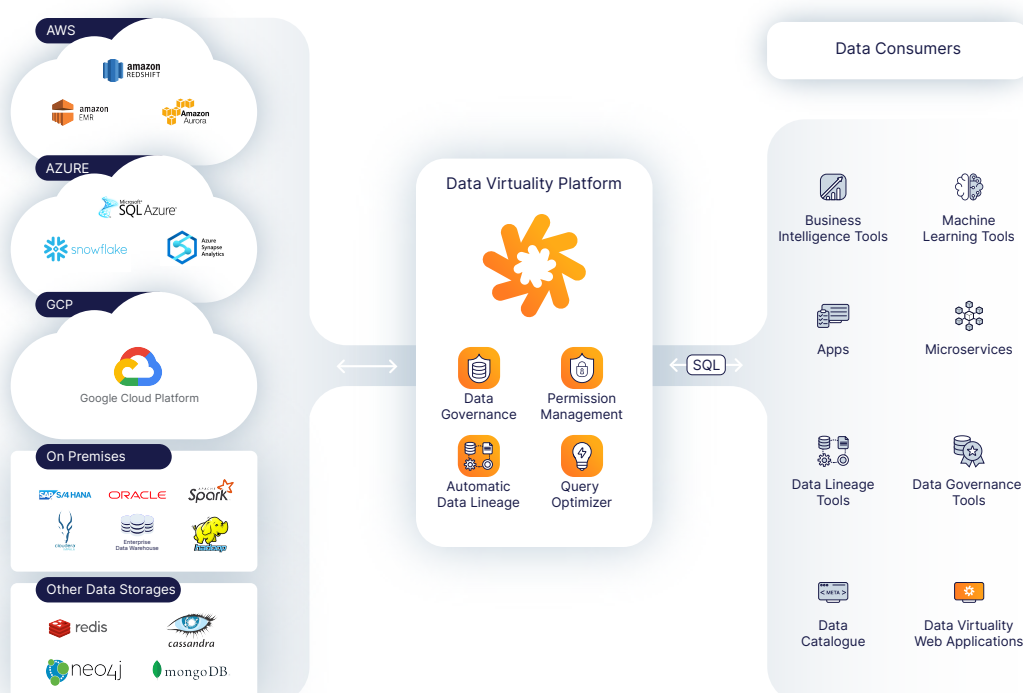
● Use Case: Clouds for Financial Institutions

Financial institutions mostly rely on hybrid cloud architectures as they have many legacy systems in place. A growing list of global financial institutions, including Bank of America, BNP Paribas, Santander and Royal Bank of Canada have adopted hybrid- or multi-cloud models. Hybrid clouds allow institutions to scale up and down as needed, providing both immediate and long-term flexibility that enables institutions to adapt to ever-changing business conditions.

A sophisticated hybrid-cloud strategy offers a number of benefits for financial institutions:

- **Reduce costs**
Hybrid-cloud helps banks scale their data needs in real-time and avoid the expensive shelfware/unused digital capacity.
- **Increased performance**
In a hybrid-cloud architecture, data doesn't have to unnecessarily move to the cloud. It can be accessed where it resides and yet be integrated with the workloads that are already in the cloud. The results are better performance of workloads and jobs, better cost parameters, and superior time to integrated data delivery.
- **Make banking more efficient**
Hybrid cloud enables banks to quickly move digital resources to where they are needed, responding to changing customer demands.
- **Drive innovation**
Hybrid-cloud accelerates innovation by not being limited to a specific geographic location or even a specific organization.
- **Ensure security**
Security threats are constantly changing, and the hybrid-cloud gives banks access to AI-driven tools that will be critical to identifying and combating cyber threats.

How Data Virtuality Supports the Cloud Data Ecosystem



Financial institutions have always been early adopters of data management technologies. Cloud options are no different. Legal, regulatory, and compliance requirements in the financial sector are basically forcing these companies to adopt multi- or hybrid-cloud architectures. Moreover, many banks have recognized and are leveraging the potential of the cloud.

● Conclusion

Clouds offer companies new ways to manage and work with their data. Old concepts of data management such as data lake or data warehouse often cannot reflect the new realities. Therefore, innovative approaches to enterprise-wide data management are needed, considering the opportunities and risks of cloud deployment. Once you have decided to move to the cloud, the strategy for the cloud data ecosystem is crucial for your success.

Several aspects to consider for your cloud journey:

1. Is a single cloud or a multi-/hybrid-cloud architecture better for your business? For sustainable, agile, and flexible data management, we believe the answer is clear.
2. For optimal use of hybrid-/multi-cloud architectures, specialized tools from ISVs play a crucial role. The data integration/virtualization, metadata, data quality, and data governance capabilities that are offered by cloud providers are mainly focused on the data that is already in their cloud data stores. ISVs liberate companies from having to move ALL their data to a single cloud.
3. The importance for the organization to be independent of individual providers and to have flexibility through a "best- of-breed" approach.

● About Data Virtuality

- Data Virtuality:
Solution provider for data integration and management technologies.

- Smart Data Virtualization Solution:
Data Virtuality Platform - Smart Data Virtualization for Flexible Data Architectures By uniquely combining data virtualization and data replication, Data Virtuality Platform gives data teams an increased flexibility in this fast-paced world of data. It is an enabler for Data Mesh and Data Fabric by providing the self-service capabilities and data governance features that are indispensable for these frameworks. Enterprises around the world, such as BSH, PGGM, PartnerRe, Crédit Agricole, and Vontobel use the Data Virtuality Platform to build modern data architectures that meet today's and tomorrow's business requirements.

- Data Replication Solutions:
Data Virtuality Pipes Professional
Data Virtuality Pipes

- Founded:
2012 by Nick Golovin (PhD) in Leipzig, Germany after 8 years of research

- Offices:
Leipzig, Munich, San Francisco

- Awards:



Message: info@datavirtuality.com
Visit: datavirtuality.com
Request Demo: demo@datavirtuality.com

