



Secure Multi-Cloud Data Orchestration Frameworks for Digital Banking Ecosystems

Anumandla Mukesh

Independent Researcher, India

ABSTRACT: Secure Multi-Cloud Data Orchestration Frameworks for Digital Banking Ecosystems reflect the increasing demand for data services and complex workloads in the digital banking ecosystem. Service and data integration across clouds is a prominent requirement but sharing and cross-cloud geo-replication of sensitive data introduce security and privacy risks that control dispositions across cloud service providers do not adequately address. In this work, Security Data Orchestration in Multi-Cloud Banking Ecosystem proposes an architecture and a Data Orchestration in Multi-Cloud Banking Ecosystem Methodology for achieving a cloud-agnostic Control Plane. The comprehensive control set satisfies independent compliance frameworks and enables secure access to sensitive data while satisfying data policies, residency requirements, and regulatory considerations.

Given the interest and demand for Banking-as-a-Service from within the finance industry, a clear implementation roadmap and criteria enable banking institutions to plan, design, and secure a Service “Orchestration” and/or “Data Orchestration” capability. Security Data Orchestration in Multi-Cloud Banking Ecosystem present the Trade-Offs for Centralized, Federated, and Hybrid Architectures of Data Orchestration in Multi-Cloud Banking Ecosystem and Information and Data Aspects of Control Plane.

KEYWORDS: Clarify security, orchestration, IAM, data protection, and governance in support of secure Banks-orchestration customers-Banks-multi-cloud-data-orchestration-ecosystems.

I. INTRODUCTION

Digital banking and multi-cloud orchestration highlight the migration of company-sensitive data to multi-cloud service providers. Although industry players have increased funding in these two key strategic areas, several gaps remain. First, the sensitivity of financial data necessitates a careful assessment of the security, privacy, and compliance issues involved in exploiting multi-cloud offerings. A failure to consider distributed data management and orchestration for analytics, machine learning, and business-intelligence workloads exposes institutions to significant risk. Second, regulatory bodies such as the Basel Committee on Banking Supervision have begun to issue specific guidance on the management of technology risks, particularly regarding the use of public-cloud services. Although distributed orchestration is a key aspect of addressing these risks, no security design principles or building blocks linking these two requirements in a cloud-agnostic manner have been proposed. This study aims to fill these gaps, thereby helping financial institutions strike an optimal balance between value creation, pragmatic deployment, resilient operations, and regulatory compliance as they adopt multi-cloud strategies.

A secure multiparty control-and-data-sharing framework that enables data orchestration, reduces risk exposure, enforces security and audit controls, supports data cataloguing, and provides interoperable visibility across different data-sharing schemes in a multi-cloud environment can provide a useful foundation. The proposed framework incorporates security strategies aligned with supervisory authority guidelines and industry best practices. By enabling and securing cross-cloud data-sharing transactions, moving sensitive data onto data services rather than the applications themselves, and maintaining cross-cloud visibility across multi-cloud deployments, such a framework can act as a reference architecture for the risk-averse adoption of data-sharing services across multiple service providers in a cloud-distrust context.

1.1. Background and Significance

Seamless integration of local, private, exploratory, and multi-cloud resources into a single coherent framework, scalable data processing, multi-cloud operational resilience, and multifaceted analytics result in the Swiss army knife effect. Digital banking is at the forefront of using cloud technologies. By, cloud solutions are expected to finally become the building blocks of financial services, with IT infrastructure in Europe being mostly still on-premise,



moving day-by-day to the cloud as cloud services become cloud-native and compliance, handling and data segregation issues become addressed.

Seamless orchestration across clouds is crucial for enabling the sharing of the cloud ecosystem between organizations—in particular, the sharing of the ecosystem data without duplicating it. Cybersecurity resilience demands access from anywhere. However, by distributing the architecture into different organizational infrastructures, security, governance, and integrity issues become very important. Security must also apply to cross-cloud, cross-organization access control. Indeed, sharing the multi-cloud resources requires a strong identification and authorization process, and an access-control model that grants access only to the data required to perform the operation is crucial; hence the least-privilege model must be supported.

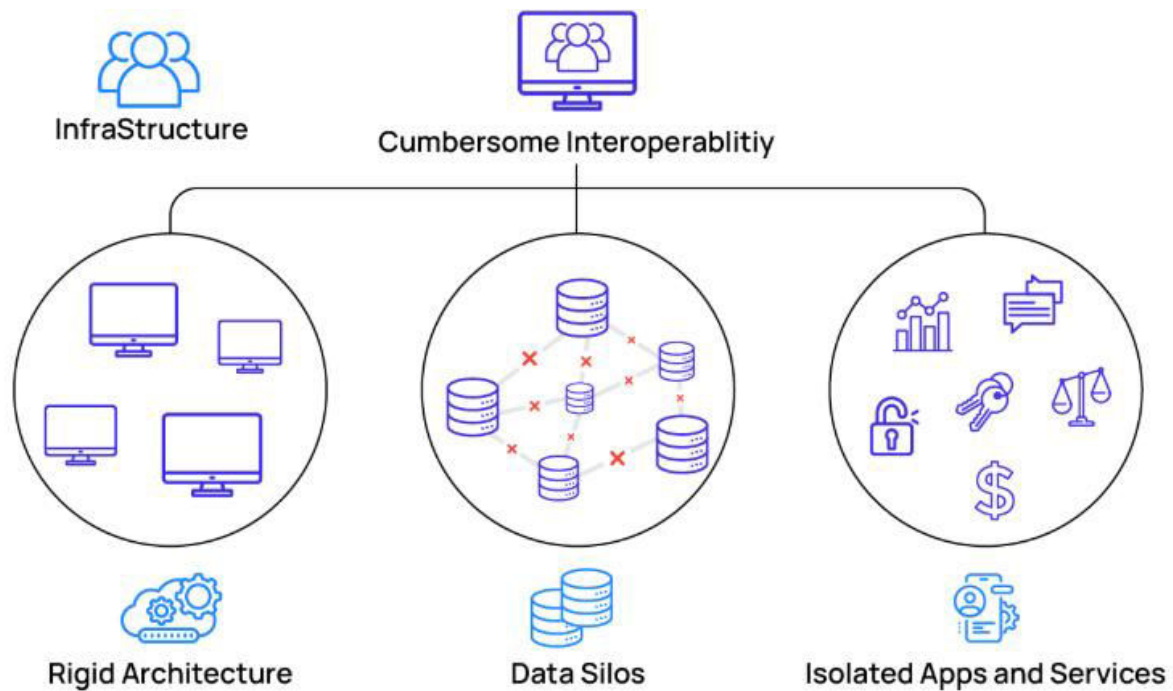


Fig 1: Multi-cloud Data Management for Financial Services

1.2. Research design

The study begins by addressing the extant research gaps pertaining to multi-cloud data orchestration and extending the analysis to a digital banking ecosystem. Subsequently, the conclusions provide an actionable roadmap for multi-cloud data orchestration in digital banking, together with the critical design criteria required to implement secure solutions.

Currently in a transitory phase, digital banking is evolving within an internal ecosystem of banking service providers and an external ecosystem of third-party service providers and consumers. The continual emergence and rapid evolution of new platforms and functions have created considerable friction among regulatory authorities, external data service providers, and the banks themselves. Multi-cloud orchestration supports continual and rapid development and deployment by orchestrating components across cloud providers. The alignment with digital banking is therefore indirect; the alignment with regulatory design criteria is more important, as aligning closely with risk, oversight, and internal control arouses less internal resistance.

Equation 1: Availability equation for centralized vs federated/hybrid

Let each independent cloud region/provider have availability:

$$A_i = P(\text{site } i \text{ is up})$$

So site failure probability is:

$$P(\text{site } i \text{ down}) = 1 - A_i$$

2.2 “At least one is up” availability (active-active / failover)

Assume independence (common in first-cut models):

- Probability all *N* sites are down:



$$P(\text{all down}) = \prod_{i=1}^N (1 - A_i)$$

- Therefore probability the system is **available** (at least one up):

$$A_{\text{total}} = 1 - P(\text{all down}) = 1 - \prod_{i=1}^N (1 - A_i)$$

II. FOUNDATIONS OF MULTI-CLOUD DATA ORCHESTRATION

Digital Banking Ecosystem. Multi-Origin and Multi-Location Representation offer more advantages for Data Orchestration than services provisioning in a single Cloud Provider. But, it brings new challenges e.g., Services offering and location Management, Data Distribution and Sharing, Operation Transparency and Visibility, Providers Operation Interoperation. The concept of Data Orchestration is new and unique concept that addresses these challenges. It provides data aware decision, Data routing, Marked Data Storage, Operation Timing.

Digital banks must be ready to respond promptly to changes affecting the financial markets, the needs of customers and other market participants. To this end, Digital Banks are implementing Multi-Cloud strategy: In order to be agile and competitive it is advisable to implement a Multi-Cloud strategy by taking advantage of the characteristics of the various Cloud Providers. The Use of Multi-Cloud Services of various Cloud Service Providers facilitates these key points: transparency and visibility of operations; efficient operation of the Bank; timely communication of important market information; introduction of new products and creating new opportunities for customers; strategic use of Data and Algorithms; continuous Offering of Data Services to Customer; Mitigation of Counter-party Risk by diversification. The Digital Banking Ecosystem often does not have control and visibility of the Banks Operations, nor the important information that should communicated on-time.

The Multi-Cloud Concept is not restricted to the Operating Model of Data Services or Data Products, it includes Operating Model of Events and Alerts or Messages. Data Orchestration in Digital Banking Ecosystem can be understood as the visual, logical and optional representation of the services/responses/resources offered by the various Cloud Providers, on a central data service vertically/horizontally through a single Cloud Provider, on a combination of Cloud Providers, using combination of business rules, service routing control.

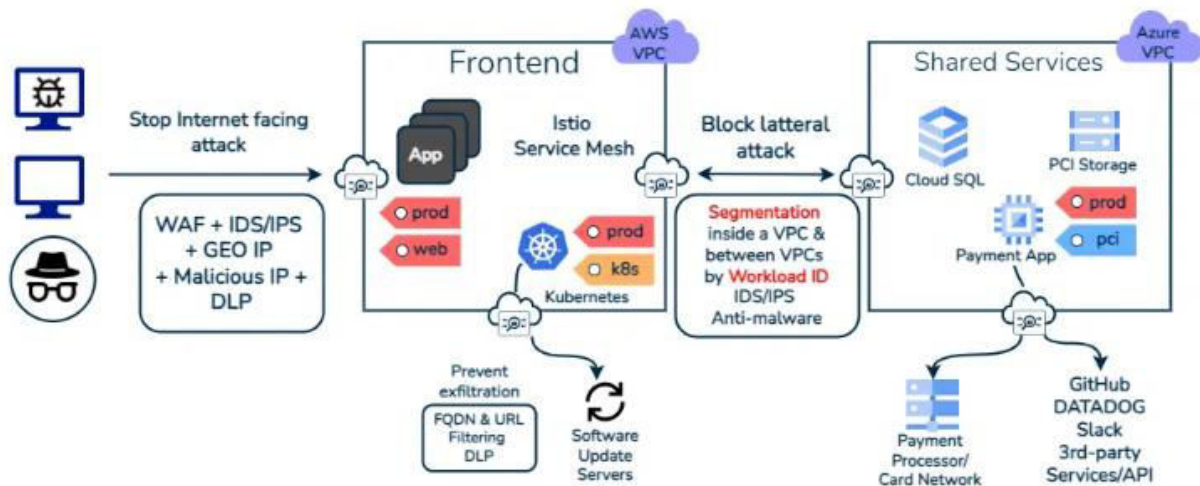


Fig 2: Multi-Cloud Data Orchestration

2.1. Definitions and Scope

Demand for agile, tailored, digital-first experiences in the banking domain is rising. One-fifth of global consumers, IIoT initiatives, pandemic-induced changes, and cost-cutting aspirations bolster churn and NPLs. Digital-first transformations are vital to mitigate these factors. Bank refusal to abandon legacy systems risks losing market share to better-positioned disruptors. Multi-cloud environments sidestep vendor lock-ins, combine capabilities of multiple



providers, and assure business continuity—yet the core idea is difficult to realize in practice. Such environments are not fully exploited for business, security, or governance reasons. This study seeks to define a secure multi-cloud data orchestration framework for digital banking processes supporting diverse patterns, fostered through a central control plane.

Multi-cloud arrangements must fulfil an ever-growing list of requirements addressed through a move from the digital presence to the digital ecosystem concept. Domain-driven design, microservices architecture, and new patterns of access control such as zero trust and concepts around the self-sovereign identity principle play a vital role. Research gaps exist even in terms, definitions, and distinctions at the high-boilerplate level; multi-cloud, data orchestration, governance, security, and interoperability must first be defined. Centralized, federated, and hybrid models must also be compared to refine and propose an architecture that answers the original challenge in the specific case of the digital banking ecosystem.

Equation 2: Identical sites (simplifies plotting)

If $A_i = A_{\text{site}}$ for all i :

$$A_{\text{total}} = 1 - (1 - A_{\text{site}})^N$$

That's exactly what I graphed in the line plot: "Availability gain from multi-region/provider redundancy: $A_{\text{total}} = 1 - (1 - A_{\text{site}})^N$ ".

Interpretation mapped to the paper

- **Centralized** $\approx N = 1$ (single deployment \rightarrow more "single point of failure" exposure).

Secure Multi-Cloud Data Orchest...

- **Federated/Hybrid** increase N , raising A_{total} (but governance complexity rises, which the paper discusses via control plane / hooks).

2.2. Architectural Models

Available architectural models for organizations implementing multi-cloud setups include centralized, federated, and hybrid designs. Each option has specific trade-offs that must be weighed against the organization's unique needs. The underlying requirements of digital banking ecosystems, for example, can benefit from particular properties of federated architectures.

In a centralized model, all multi-cloud-aware services and administrative functions are consolidated into a single cloud deployment. This setup offers simplicity, consolidation of management and governance operations, and maximum homogeneity in services. Its centralized nature, however, can negatively impact the availability and performance of the services and features provided to other clouds, especially those that are far removed from the central deployment. Additionally, it may present a single point of failure and become a target for malicious actors seeking to disrupt operations or obtain sensitive data. Hence, a centralized model is only suitable for multi-cloud environments where the risks associated with such disadvantages are acceptable.

III. SECURITY AND COMPLIANCE IN MULTI-CLOUD ENVIRONMENTS

Access control, access management principles, mechanism, and mechanisms implemented in the multi-cloud data flow play a vital role in the security of any multi-cloud environment. For example, the banker needs access to invoices that are being moved from a narrow cloud in the European Union and shared with a cloud hosted in Australia for tax purposes and should therefore have no role in assessing risks associated with a new range of credit cards that a cloud in Asia-Pacific would like to launch.

Protecting Data in Transit because data is shared among different multi-cloud service providers it can be intercepted during its transit such as the customer credit information. Hence it is necessary to take precautions before it travels such as data masking and tokenization. The data can be encrypted while remaining in use so that it is in its readable state only for the desired user reducing risk exposure while remaining usable. Data protection while at rest primarily deals with encryption of important data. Encryption keys should be kept in a different cloud provider. Furthermore digital banking ecosystems being responsible for data residency, data sovereignty, data origination, cross-border transfers and compliance check must ensure whether the shared country specific data should comply with the regulatory requirements of the hosting cloud region.

Implementation of control plane in multi-cloud data orchestration involves the four steps of orchestration and policy enforcement, visibility across multiple cloud providers, and cross-cloud compatibility making cloud propagation possible when needed. Cross-cloud data sharing enables sharing of data stored in any cloud provider with another cloud provider and removes the dependency on data replication. Full compatibility of security controls used in the different clouds can be kept in the governance hooks where the security controls of the specific cloud resides.

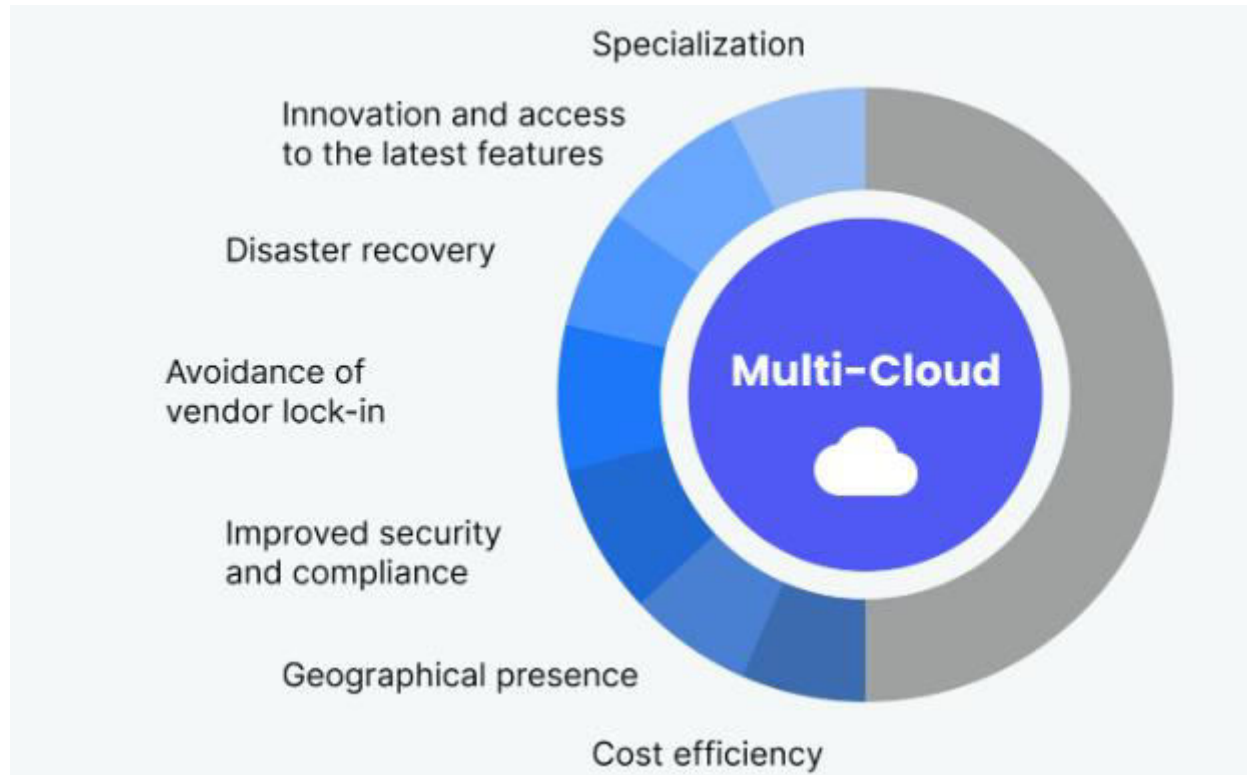


Fig 3: Multi-Cloud Security Strategies

3.1. Identity and Access Management

Identity and Access Management (IAM) is paramount in multi-cloud environments, safeguarding data flow among clouds and with users. IAM principles direct access policy updates in response to cloud provider identity, region, or data classification changes. The framework assigns a data manager role to create data and grant read access to a target cloud platform, with operational configuration control consigned to IAM-aware data management services. Access policies employ the principle of least privilege, and IAM role-based access control enforces job-function-related controls. The architecture applies the Zero Trust philosophy, verifying all users and components regardless of locality. An identity repository supports data governance via IAM provider APIs or connectors to external repositories. Data protection is integral to IAM policy definition and runtime workflows, mitigating security breaches through comprehensive log records. These structure cloud provider identity information, action details, and access control responses, enabling audit trail verification, compliance assurance, and access policy improvement.

3.2. Data Protection and Encryption

Data protection in multi-cloud environments requires defining policies specifying how the data is protected according to its classification, ensuring that sensitive data is encrypted before being sent to the cloud. Policies must also specify the cryptographic key on which the encryptions should be performed and whether additional forms of data protection, such as tokenization or masking, should be used. Encryption in transit is mandatory for data travelling across security domains. Bou-Rezeq et al. suggest using an encryption-as-a-service (EaaS) model to centralize the implementation of encryption services. Support for EaaS should be integrated into the data orchestration process, allowing clouds that do not offer in-transit data encryption to make use of an EaaS service during the data transfer.

Encryption at rest is also advisable, since it can mitigate the risks against unauthorized access to the data. However, the fact that the encryption keys are also stored in the provider's infrastructure diminishes its security aims. As a countermeasure, key management implementation could involve not creating the key to encrypt sensitive data or creating, encrypting, and sharing these keys by means of the company's employees and processes. However, this approach add extra costs and complexity to the solution. Another alternative involves employing technology that allows sensitive data to be divided in portions; some of these portions are encrypted/sent on their original format while others are non-sensitive data or tokenized data that can be shared without any protection.



Equation 3: Break latency into components

Let:

- *RTT*= round-trip/network handshake latency (seconds)
- *S*= payload size (bytes)
- *B*= effective throughput (bytes/sec)

Baseline transfer time:

$$T_{net} = RTT + \frac{S}{B}$$

IV. DATA FLOWS AND ORCHESTRATION PATTERNS

Concepts, Patterns and Orchestration of the Data Flow in the Multi-Cloud Environment in the Banking Context

The master data flow process covering data ingestion from data sources, transformation processes, data flow routing in upstream and downstream data services and data lineage is identified. Components of the flow are mapped in accordance with the following concepts that enable interoperability across the different cloud providers: standards for serialization formats and data exchange schemas; metadata formats; and techniques that verify the conformance of data exchanged between clouds. Supplementary requirements are addressed: Data Residency, Data Sovereignty, Cross Border Data Flow and Data Residency.

A well-defined, orchestrated Data Flow process is imperative in data-driven banking due to external data integration, data engineering, analytics-driven advance banking services and customer demand for Data Sharing Services. Without explicit orchestration and monitoring, banks run the risk of exposing sensitive data in Data Sharing Services and exposing their reputation through unmonitored Data Sharing Services. Data-driven banking implies a well-defined master data flow process covering data ingestion from Data Sources, Data Transformation Processes, routing of the Data Flow in Upstream and Downstream Data Services and Data Flow Lineage. The absence of a well-orchestrated and monitored Data Flow process increases operational risk, potential data snooping and reputational risk from inappropriate cross-cloud Data Sharing Services.

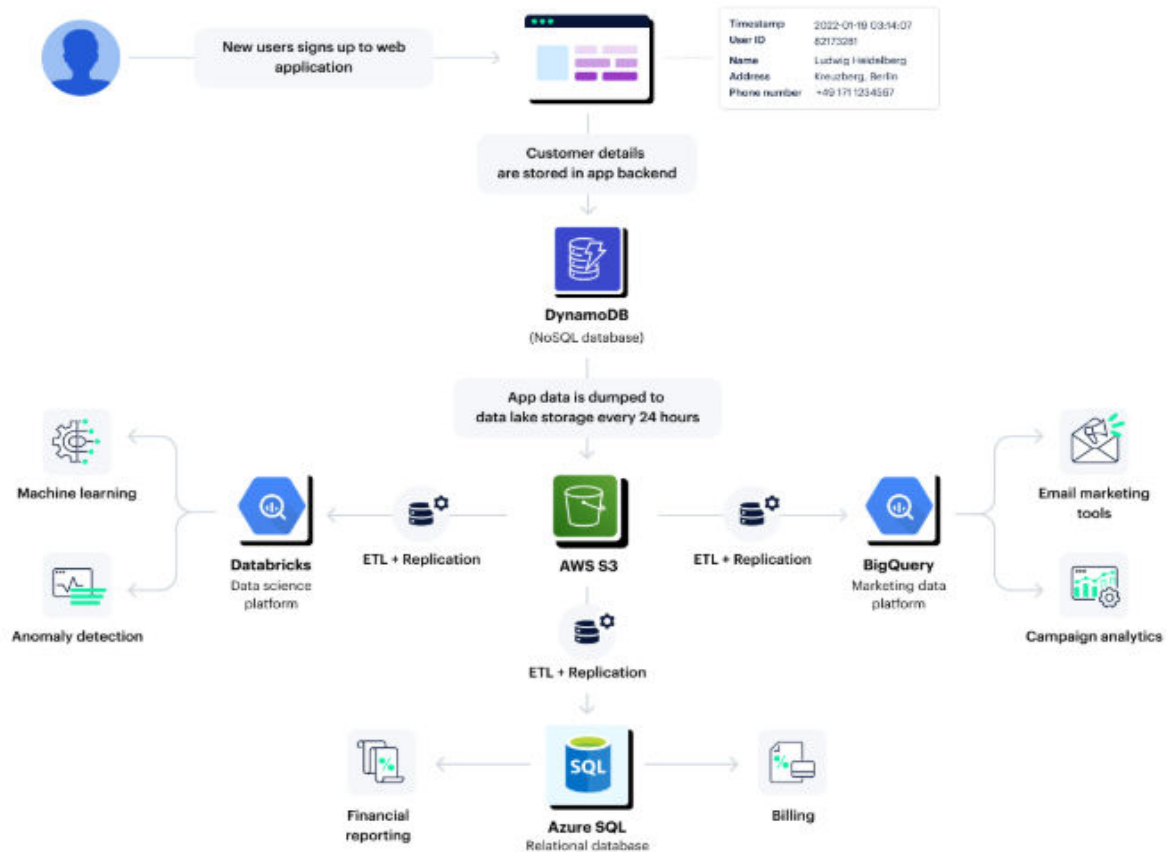


Fig 4: Data Flow Security



4.1. Data Ingestion and Transformation

A data orchestration framework spans all orchestration and governance activities, including data processing, routing, transformation, and lineage across clouds. Map the logical flow of data as it enters the ecosystem, describing how incoming data lines are detected and logged, specifying any processes that monitor for schema evolution, and mapping how the incoming data are routed through any data pipelines that will conduct transformation or enrichment before use. Data sources can be heterogeneous, with accompanying variations in data formats, metadata, timing, and other considerations. Specify how any discrepancies will be resolved. Define any default transformations to filter or fill data gaps and detect data types. Describe the mechanisms employed to record data lineage, standardize data formats, and maintain up-to-date metadata on the data in the system. The result should provide a complete mapping of data ingestion and transformation flows, including details about critical data sources and governance considerations.

For banking ecosystems, centralization risks crippling business continuity; losses incurred due to outages in a central service cannot be mitigated with load balancing and fail-over measures. Regulatory requirements often restrict cross-border data flows, and sovereign clouds may lack a local presence for key providers. Banking regulators increasingly promote risk-sharing principles for banking supervision; cooperative arrangements delegating authority from one agency to another can enable focus on bank conduct and prudential oversight of non-banks. A federated model enables greater data locality and offers the possibility of regulatory reciprocity, although mission-specific federations may require the strengthening of governance arrangements controlling day-to-day operations. Hybrid architectures further improve data locality, as storage and processing capabilities can be controlled by the data provider.

Equation 4: Add encryption overhead

If encryption costs e seconds per byte (CPU or service overhead), then:

$$T_{\text{enc}} = S \cdot e$$

Total secure transfer latency:

$$L = RTT + \frac{S}{B} + S \cdot e$$

This is what I plotted in:

- “Cross-cloud transfer latency vs payload size (illustrative model)” and tabulated in:
- “Latency model table (illustrative)”

4.2. Data Residency and Sovereignty

Banking regulations often specify that institutions must store data in their home jurisdiction, and governments may restrict the transfer of personal data outside national borders. In federated multi-cloud orchestration, compliance with residency and sovereignty requirements can be particularly challenging because a single cloud provider will not own all the regions required by the client. Distribution of external elements among public clouds creates privacy concerns. Finally, computational operations (for instance AI processing) must not be executed across data borders.

From the sovereignty point of view, operating data remaining in internal data centers reduces exposure related to national legislation (as the banking system remains under the country's government protection). Within edges, AI operations can be realized without route to central data centers; partitions own sharp latency decrease and higher parallelization feasibility.

V. ARCHITECTURAL COMPONENTS OF A SECURE FRAMEWORK

A multi-cloud digital banking ecosystem entails a multitude of components necessary for sustaining a growing data interaction and orchestration architecture. These components can be grouped into the following architectural planes: (i) multi-cloud control plane, consisting of the orchestration, policy enforcement, visibility, and interoperability components; and (ii) data planes, made up of the actual data ingestion, transformation, routing, cross-cloud sharing, and lineage aspects, as well as the sharable data services, data catalogs, governance hooks, and security controls for data that needs to be shared, served, or stored outside its primary cloud.

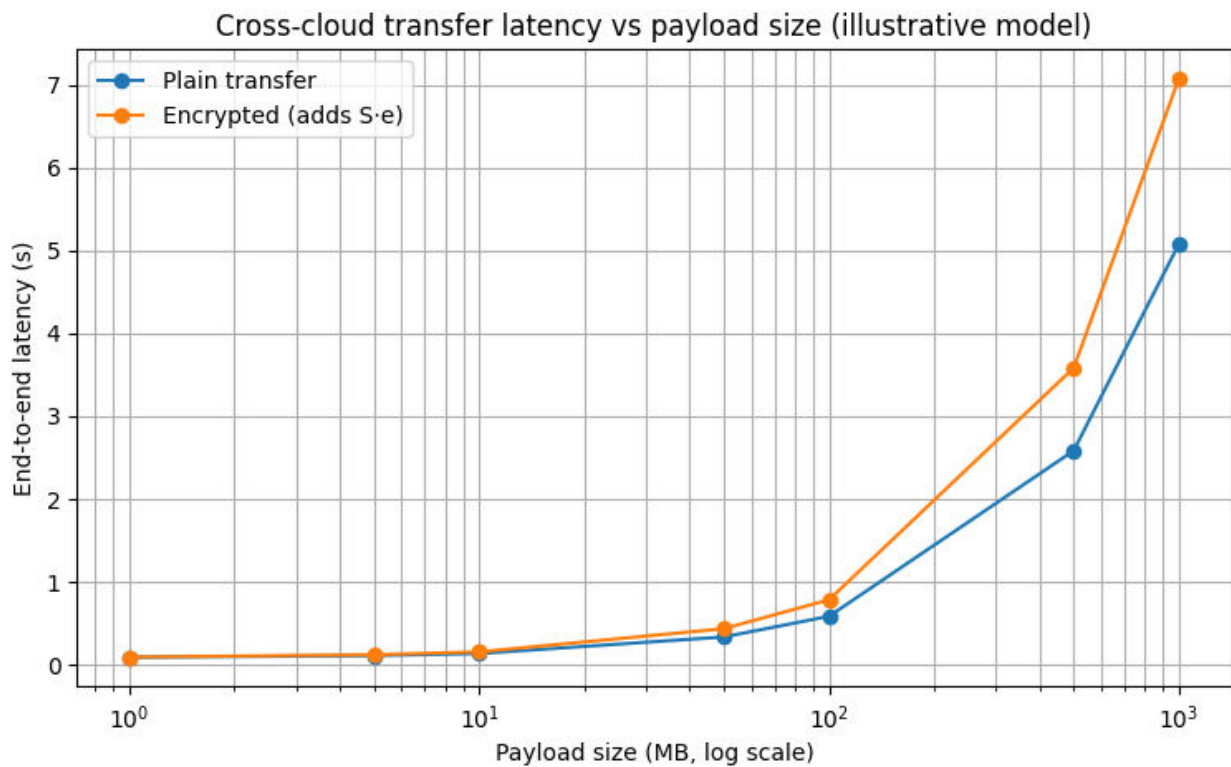
The multi-cloud control plane is paramount to the functioning of the data planes. Unlike traditional cloud frameworks, the control plane needs to cater to all the orchestrated providers, not just a single one, and ensure a seamless integration of the various security and governance requirements and constraints. On the data side, the focus is primarily on cross-cloud interactions, although certain controls may extend to data regions that are confined to a single provider. All data movement within such a bank remains governed by the standard controls of the respective providers, which is a highly satisfactory state of affairs.



5.1. Multi-Cloud Control Plane

A secure multi-cloud data orchestration framework for digital banking ecosystems is embraced through a multi-cloud control plane comprising orchestration, policy enforcement, visibility, and interoperability across cloud providers. Orchestration entails coordinating multiple cloud providers to ensure intended infrastructure set-up and maintenance percolate the deployment native environment with virtual, co-located resources within a given provider or region. Delivering numerous decentralized services without endorsing single providers is essential, thereby ensuring multi-cloud resiliency. Each service across multi-clouds is monitored, security polices laid down by the authority are enforced, and incidents are communicated to the governance layer. Key management, encryption, and masking services can be segregated across clouds; hence these common services are maintained in a sharable horizontal plane. Data can easily flow across clouds and the orchestration services can deploy rules to monitor the transfer and enforce security policies laid down by data owners. Data catalogs can embed additional business and regulatory metadata as required for each location. Data governance hooks determine approval requests, while permissions for the underlying data movement are granted by the cross-cloud governance plane. The control plane is equipped to pull audit information from other providers to provide visibility and ensure compliance.

The control plane deploys auto-scaling policies using threshold values laid down by authority. A data service plane is provisioned with services which can effectively share data across clouds with ease. The purpose of defining dedicated planes is to maintain resilience across clouds. If any plane is in a state of failure, only that service deployment will be affected; services hosted with other providers will be functioning well. The implementation of services with appropriate granularity assists in providing scalability and resilience for the end-user services, thereby provisioning a high level of service. The control plane provides necessary visibility and monitors the operations across the provider with basic governance aspects. The series of sharable data services forms the data service plane. These services are primarily used by sharable dashboards and report engines which need to access multiple clouds. Data can be shared across clouds seamlessly with these services taking care of the security requirements on data sharing. Data owners can catalog the data sets which are available for sharing with other clouds. The catalogs also capture, wherever applicable, the internal cross-cloud data transfer policies laid down by the data governance authority.



5.2. Data Planes and Sharable Data Services

Five deployment planes constitute a Secure Multi-cloud Orchestration Framework. The control plane orchestrates the cloud ecosystems across different service providers, while the deployment-specific planes provide services, data, or management capabilities associated with that plane. The control plane ensures compliance by constructing control paths in the plane that correspond to controls mapped from the Risk Assessment Framework, Threat Modelling Framework,



and Design and Implementation Consideration. Services in the infrastructure management plane include resource management, performance and capacity management, and IAM. The two deployment-specific planes in a banking context are the data plane and the application plane. The application plane offers Application-as-a-Service (AaaS) capabilities, while the data plane provides sharable data services such as data storage, data catalog, Governance Hooks, and security controls for sharable data.

Data residency and security concerns necessitate sharing sensitive data between service providers while deploying banking applications and services across them, thereby resulting in data plan. Sensitive data such as customer details, transaction details, and transaction history maintained by respective data controllers during business operations must be retained within the same jurisdiction (data-residency requirement). Remaining data may cross jurisdictions along with user consent (data-sovereignty requirement), while data without user identification may cross boundaries without limitation (data-aggregation requirement). The data plane provides securely sharable data services and ensures compliance with these regulations.

Equation 5: Policy enforcement as constraints (data residency/sovereignty) (step-by-step)

- Let D be datasets.
- Let R be regions/cloud locations.
- Let $Allowed(d) \subseteq R$ be regions allowed by policy for dataset d .
- Decision variable:

$$x_{d,r} = \begin{cases} 1 & \text{if dataset } d \text{ is stored/processed in region } r \\ 0 & \text{otherwise} \end{cases}$$

$$\sum_{r \in R} x_{d,r} = 1 \forall d \in D$$

$$x_{d,r} = 0 \forall d, \forall r \notin Allowed(d)$$

Equivalently:

$$\sum_{r \in Allowed(d)} x_{d,r} = 1 \forall d$$

If $C_{d,r}$ is cost (or expected latency), choose:

$$\min \sum_{d \in D} \sum_{r \in R} C_{d,r} x_{d,r}$$

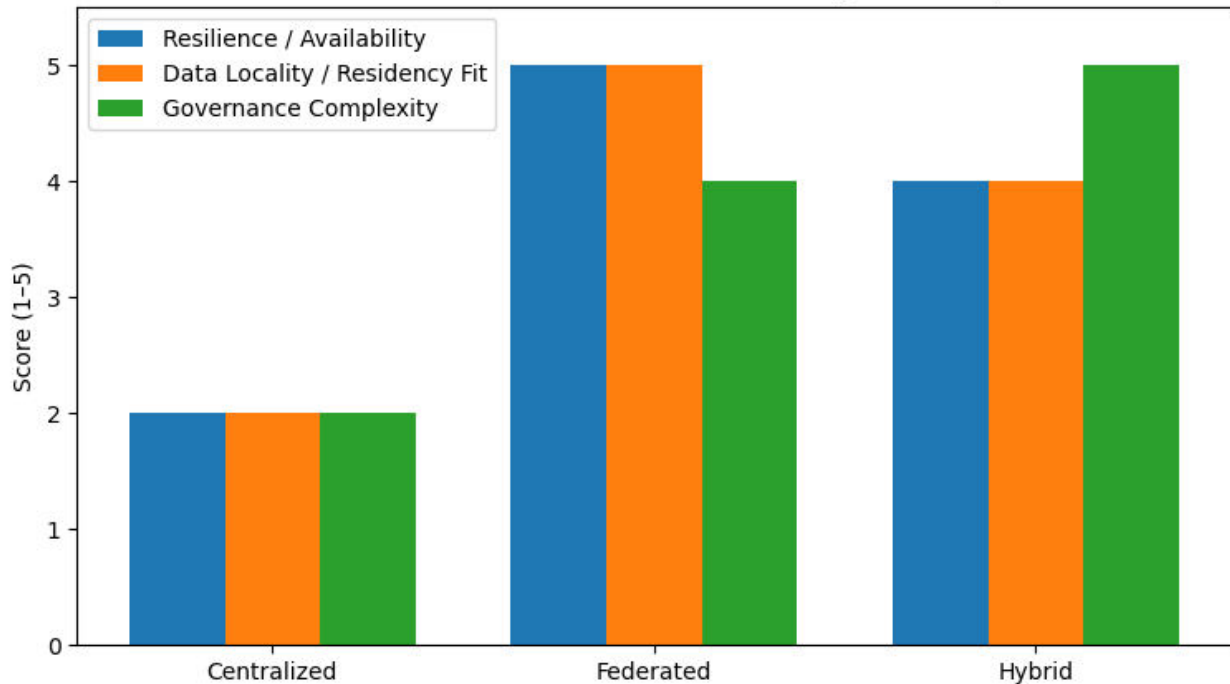
VI. GOVERNANCE, RISK, AND COMPLIANCE (GRC) IN DIGITAL BANKING

Banking governance, risk, and compliance (GRC) faces numerous complex challenges, especially in terms of technology adoption. Digital transformation, net-zero transition, and the application of emerging technologies, together with the advent of new business models, remain key priorities. Nevertheless, associated risks also appear greater than ever, particularly with regard to operational risk and cyber-risk. Many institutions realize they are failing to embrace GRC in a cohesive fashion. Their GRC functions are often poorly aligned, leading to inconsistent, duplicative, or conflicting activities across the organization.

The public cloud is a significant part of many banks' infrastructure strategies. Nevertheless, few banks have fully integrated it into their GRC frameworks and risk assessment processes, an oversight that may enable serious errors in execution and create risks undermining the benefits. Insufficient attention to managing banks' GRC arrangements, considered just a box-ticking exercise, let many institutions down during the COVID-19 crisis; banks that embedded GRC considerations into their decisions generally fared better. A loss event linked to GRC vigilance can be several times the level of the initial expected loss. The level of responsibility, relationships, and governance structure should be clear at the outset, especially as questions related to data sharing with state authorities and third parties re-emerge.



Trade-offs across architectural models (illustrative)



6.1. Risk Assessment Methodologies

Understanding risks and threats is a prerequisite for implementing effective security controls. Banking environments are characterized by high residual risks, and consequently, substantial investment is often required in control mechanisms. Control design and investment must therefore be aligned with the bank’s risk appetite, in keeping with the principles of Risk Based Information Security. A formal risk assessment methodology enables risk analysis by allowing threat modeling, prioritization of threats and enumeration of the assets the banks holds. The risk control catalog is then consolidated with control mappings based on widely accepted frameworks, ensuring cloud-agnostic applicability.

The scope of the risk assessment focuses on security controls. Control mapping is intended to either report that a significant control is implemented in some or just a few cloud services, or to highlight the need to implement the specific control, providing a residual risk analysis when necessary. Kentico, a Content Management System (CMS) for Websites is used as a testbed for cloud-agnostic control mapping, supporting the risk assessment analysis. branca, Counter Argument—a model that addresses cloud service residual risks—serve as inspiration for performing the analysis and response with an increasing adoption of cloud Computing services, especially IaaS.

6.2. Regulatory Mapping for Banking Services

The Regulatory Mapping component details the mapping of Regulatory Control Specifications to Cloud Agnostic Control Implementations to record Regulatory adherence, Compliance evidence with Residual Risk Profiling and Audit Traceability Record of Cloud Agnostic Control Specifications. Every Cloud Environment used for Digital Banking services implements all Regulatory controls expectations as specified by Basel III/IV, FFIEC, GDPR and other maintainance Regulatory authorities. Whenever hints for Clouds or Cloud Providers are used, Regulatory Control Mapping serves as evidence of proper cloud planning and implementing a set of services/functionalities which take care of all the mentioned aspects. The approach is based on Principle-Based Regulation and not rule-based regulation, hence every principle is taken into account into consideration for building a particular Cloud Environment IAM Model. Supporting Banks in understanding Requirements must implement a Risk Assessment Methodology, undertake Threat Modelling for Cloud Services, and identify Controls to address Bank, Service Provider & third-party Residual Risk across Banking Services using Cloud. Residual risk is then linked back to Risk Assessment Results and Auditable Traceability. All this process thus serves as an endorsement for selected Cloud Services, and is mapped to Basel III & IV, GDPR, FFIEC for Evidence for Audit. Using a multi-cloud management strategy, Regulatory Mapping for each Cloud service Provider is done and the attention goes towards how these cloud services are consuming & growing for Banks.



VII. CONCLUSION

The implementation of secure multi-cloud data orchestration frameworks enables the data of a digital banking ecosystem to be ingested and managed across multiple cloud providers while governances and security capabilities are enforced at all times and from a single vantage point. The proposed design and validation is the first of its kind and covers the design of control planes that provide visibility and support the orchestration of policies across clouds, and the design of the data planes that provide the data services needed to support the ecosystem.

To facilitate future highly available, reliable, and resilient deployments; to prove that the design is cloud service and technology agnostic; and finally, to lay the foundation for the risk- and compliance-oriented design of control planes that monitor the governances and security capabilities of the data being shared, a roadmap is proposed. The roadmap indicates that the decision of the banks' geographic presence on the cloud services is critical for a secure implementation of the design, and a foundation for cloud residency, sovereignty, cross-border data flow, and risk- and compliance-oriented control planes has been laid.

Payload (MB)	Latency plain (s)	Latency with encryption overhead (s)	Extra encryption time (ms)
1	0.085	0.087	2.0
5	0.105	0.115	10.0
10	0.13	0.15	20.0
50	0.33	0.43	100.0
100	0.58	0.78	200.0
500	2.58	3.58	1000.0

Table : Latency model table (illustrative)

7.1. Future Trends

The recent pandemic has accelerated a shift towards digital banking, with banks reporting increases in transaction volumes without similar increases in headcounts that would typically be expected. Banks and financial service providers have been increasingly relying on several cloud providers and, in many instances, large multi-cloud setups across their centralized data lakes, more so than ever, owing to various reasons. However, orchestration of controlled, secure, and risk-compliant cross-cloud data flow across service and data-redacted perspective across these multi-cloud setups has yet to be addressed, given the distributed nature of data across regions.

Some of the orchestration challenges faced are complex control-plane terminating flows for seamless visibility and remote control access verification across major cloud platforms, separation of data redaction and user-location-based data flow control from the multiple cloud locations to the user that address data-residency requirements, risk-compliant cross-cloud data-sharing capability, orchestration and policy-control options to share as well as consume common services, policy-based orchestration for data-sharing across the clouds that forces cloud data platforms to meet the data-sharing requirements of user-cloud vendors, and ability to map user transactions across the multiple-cloud locations. Addressing these challenges while meeting the external regulations governing cloud banks remains paramount.

REFERENCES

- [1] Adegbite, M. A. Data privacy and data security challenges in digital finance. *Journal of Digital Security and Forensics*, 2(1), 6–19.
- [2] Goutham Kumar Sheelam, "Semiconductor Innovation for Edge AI: Enabling Ultra-Low Latency in Next-Gen Wireless Networks," *International Journal of Advanced Research in Computer and Communication Engineering (IJARCCE)*, DOI: 10.17148/IJARCCE.2022.111258
- [3] Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 734–749.
- [4] Davuluri, P. N. (2020). Improving Data Quality and Lineage in Regulated Financial Data Platforms. *Finance and Economics*, 1(1), 1-14.
- [5] Arasu, A., & Kaushik, R. (2014). Data cleansing: A context dependent approach. *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, 135–146.
- [6] Rongali, S. K. (2020). Predictive Modeling and Machine Learning Frameworks for Early Disease Detection in Healthcare Data Systems. *Current Research in Public Health*, 1(1), 1-15.



- [7] Armbrust, M., Das, T., Davidson, A., Ghodsi, A., Or, A., Rosen, J., Stoica, I., Wendell, P., Xin, R., & Zaharia, M. (2021). Delta Lake: High-performance ACID table storage over cloud object stores. *Proceedings of the VLDB Endowment*, 13(12), 3411–3424.
- [8] Inala, R. Advancing Group Insurance Solutions Through Ai-Enhanced Technology Architectures And Big Data Insights.
- [9] Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50–58.
- [10] Chava, K., Chakilam, C., & Recharla, M. (2021). Machine Learning Models for Early Disease Detection: A Big Data Approach to Personalized Healthcare. *International Journal of Engineering and Computer Science*, 10(12), 25709–25730. <https://doi.org/10.18535/ijecs.v10i12.4678>
- [11] Babcock, J., Chaudhuri, S., & Das, G. (2004). Dynamic sample selection for approximate query processing. *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, 539–550.
- [12] Sriram, H. K. (2022). Advancements in Credit Score Analytics using Deep Learning and Predictive Modeling Techniques. Available at SSRN 5255128.
- [13] Bifet, A., & Gavaldà, R. (2007). Learning from time-changing data with adaptive windowing. *Proceedings of the 2007 SIAM International Conference on Data Mining*, 443–448.
- [14] Muthusamy, S., Kannan, S., Lee, M., Sanjairaj, V., Lu, W. F., Fuh, J. Y., ... & Cao, T. (2021). Cover Image, Volume 118, Number 8, August 2021. *Biotechnology and Bioengineering*, 118(8), i-i.
- [15] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- [16] Vadisetty, R., Polamarasetti, A., Guntupalli, R., Raghunath, V., Jyothi, V. K., & Kudithipudi, K. (2021). Privacy-Preserving Gen AI in Multi-Tenant Cloud Environments. Sateesh kumar and Raghunath, Vedaprada and Jyothi, Vinaya Kumar and Kudithipudi, Karthik, Privacy-Preserving Gen AI in Multi-Tenant Cloud Environments (January 20, 2021).
- [17] Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications*, 19(2), 171–209.
- [18] Dwaraka Nath Kummari. (2022). Fiscal Policy Simulation Using AI And Big Data: Improving Government Financial Planning. *Kurdish Studies*, 10(2), 934–945. <https://doi.org/10.53555/ks.v10i2.3855>
- [19] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- [20] Gadi, A. L. The Role of Digital Twins in Automotive R&D for Rapid Prototyping and System Integration.
- [21] Das, T., Zhu, A., Li, S., Narayanamurthy, S., & Bhat, P. (2013). Distributed and fault-tolerant streaming computation in Spark. *Proceedings of the ACM Symposium on Cloud Computing*, 1–12.
- [22] Siva Hemanth Kolla. (2022). Knowledge Retrieval Systems for Enterprise Service Environments. *International Journal of Intelligent Systems and Applications in Engineering*, 10(3s), 495–506. Retrieved from <https://ijisae.org/index.php/IJISAE/article/view/8037>
- [23] Dean, J., & Ghemawat, S. (2008). MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1), 107–113.
- [24] Paleti, S. (2022). Financial Innovation through AI and Data Engineering: Rethinking Risk and Compliance in the Banking Industry. Available at SSRN 5250726.
- [25] DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., Sivasubramanian, S., Vosshall, P., & Vogels, W. (2007). Dynamo: Amazon’s highly available key-value store. *Proceedings of the 21st ACM Symposium on Operating Systems Principles*, 205–220.
- [26] Sriram, H. K., ADUSUPALLI, B., & Malempati, M. (2021). Revolutionizing Risk Assessment and Financial Ecosystems with Smart Automation, Secure Digital Solutions, and Advanced Analytical Frameworks.
- [27] Dwork, C. (2008). Differential privacy: A survey of results. *Proceedings of the 5th International Conference on Theory and Applications of Models of Computation*, 1–19.
- [28] Varri, D. B. S. (2021). Cloud-Native Security Architecture for Hybrid Healthcare Infrastructure. Available at SSRN 5785982.
- [29] Elmagarmid, A. K., Ipeirotis, P. G., & Verykios, V. S. (2007). Duplicate record detection: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 19(1), 1–16.
- [30] Dwaraka Nath Kummari. (2022). Machine Learning Approaches to Real-Time Quality Control in Automotive Assembly Lines. *Mathematical Statistician and Engineering Applications*, 71(4), 16801–16820. Retrieved from <https://philstat.org/index.php/MSEA/article/view/2972>
- [31] Fader, P. S., Hardie, B. G. S., & Lee, K. L. (2005). “Counting your customers” the easy way: An alternative to the Pareto/NBD model. *Marketing Science*, 24(2), 275–284.
- [32] Inala, R. (2022). Engineering Data Products for Investment Analytics: The Role of Product Master Data and Scalable Big Data Solutions. *International Journal of Scientific Research and Modern Technology*, 155-171.
- [33] Davuluri, P. N. (2020). Improving Data Quality and Lineage in Regulated Financial Data Platforms. *Finance and Economics*, 1(1), 1-14.



- [34] Meda, R. Enabling Sustainable Manufacturing Through AI-Optimized Supply Chains.
- [35] Ghemawat, S., Gobiuff, H., & Leung, S. T. (2003). The Google file system. Proceedings of the 19th ACM Symposium on Operating Systems Principles, 29–43.
- [36] Varri, D. B. S. (2022). A Framework for Cloud-Integrated Database Hardening in Hybrid AWS-Azure Environments: Security Posture Automation Through Wiz-Driven Insights. *International Journal of Scientific Research and Modern Technology*, 1(12), 216-226.
- [37] Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. *Universal Journal of Business and Management*, 1(1), 1–13. Retrieved from <https://www.scipublications.com/journal/index.php/ujbm/article/view/1357>
- [38] Gottimukkala, V. R. R. (2022). Licensing Innovation in the Financial Messaging Ecosystem: Business Models and Global Compliance Impact. *International Journal of Scientific Research and Modern Technology*, 1(12), 177-186.
- [39] Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.
- [40] Vadisetty, R., Polamarasetti, A., Guntupalli, R., Raghunath, V., Jyothi, V. K., & Kudithipudi, K. (2022). AI-Driven Cybersecurity: Enhancing Cloud Security with Machine Learning and AI Agents. Sateesh kumar and Raghunath, Vedapada and Jyothi, Vinaya Kumar and Kudithipudi, Karthik, *AI-Driven Cybersecurity: Enhancing Cloud Security with Machine Learning and AI Agents* (February 07, 2022).
- [41] Hellerstein, J. M., Haas, P. J., & Wang, H. J. (1997). Online aggregation. Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data, 171–182.
- [42] Garapati, R. S. (2022). Web-Centric Cloud Framework for Real-Time Monitoring and Risk Prediction in Clinical Trials Using Machine Learning. *Current Research in Public Health*, 2, 1346.
- [43] Hu, Y., Koren, Y., & Volinsky, C. (2008). Collaborative filtering for implicit feedback datasets. Proceedings of the 2008 IEEE International Conference on Data Mining, 263–272.
- [44] Amistapuram, K. (2022). Fraud Detection and Risk Modeling in Insurance: Early Adoption of Machine Learning in Claims Processing. Available at SSRN 5741982.
- [45] Davuluri, P. S. L. N. (2021). Event-Driven Compliance Systems: Modernizing Financial Crime Detection Without Machine Intelligence. *Journal of International Crisis and Risk Communication Research*, 339–354. <https://doi.org/10.63278/jicrcr.vi.3636>
- [46] Meda, R. (2022). Integrating Edge AI in Smart Factories: A Case Study from the Paint Manufacturing Industry. *International Journal of Science and Research (IJSR)*, 1473-1489.
- [47] Jagadish, H. V., Gehrke, J., Labrinidis, A., Papakonstantinou, Y., Patel, J. M., Ramakrishnan, R., & Shahabi, C. (2014). Big data and its technical challenges. *Communications of the ACM*, 57(7), 86–94.
- [48] Segireddy, A. R. (2020). Cloud Migration Strategies for High-Volume Financial Messaging Systems.
- [49] Khatri, V., & Brown, C. V. (2010). Designing data governance. *Communications of the ACM*, 53(1), 148–152.
- [50] Amistapuram, K. (2021). Digital Transformation in Insurance: Migrating Enterprise Policy Systems to .NET Core. *Universal Journal of Computer Sciences and Communications*, 1(1), 1–17.
- [51] Kleppmann, M. (2017). *Designing data-intensive applications*. O'Reilly Media.
- [52] Nagabhyru, K. C. (2022). Bridging Traditional ETL Pipelines with AI Enhanced Data Workflows: Foundations of Intelligent Automation in Data Engineering. Available at SSRN 5505199.
- [53] Lahiri, M., & Venkatasubramanian, S. (2013). Robust record linkage. Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, 101–112.
- [54] Rongali, S. K. (2021). Cloud-Native API-Led Integration Using MuleSoft and .NET for Scalable Healthcare Interoperability. Available at SSRN 5814563.
- [55] Leskovec, J., Rajaraman, A., & Ullman, J. D. (2014). *Mining of massive datasets* (2nd ed.). Cambridge University Press.
- [56] Rongali, S. K. (2022). AI-Driven Automation in Healthcare Claims and EHR Processing Using MuleSoft and Machine Learning Pipelines. Available at SSRN 5763022.
- [57] Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1), 76–80.
- [58] Meda, R. (2021). Digital Infrastructure for Predictive Inventory Management in Retail Using Machine Learning. *International Journal of Advanced Research in Computer and Communication Engineering (IJARCCE)*, DOI, 10.
- [59] Lin, J., Kolcz, A., & Szymanski, B. K. (2012). Large-scale machine learning at Twitter. Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, 793–804.
- [60] Sheelam, G. K. Power-Efficient Semiconductors for AI at the Edge: Enabling Scalable Intelligence in Wireless Systems. *International Journal of Innovative Research in Electrical, Elec-tronics, Instrumentation and Control Engineering (IJIREEICE)*, DOI, 10.



- [61] Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C., & Byers, A. H. (2011). Big data: The next frontier for innovation, competition, and productivity. McKinsey Global Institute.
- [62] Vadisetty, R., Polamarasetti, A., Guntupalli, R., Rongali, S. K., Raghunath, V., Jyothi, V. K., & Kudithipudi, K. (2021). Legal and Ethical Considerations for Hosting GenAI on the Cloud. *International Journal of AI, BigData, Computational and Management Studies*, 2(2), 28-34.
- [63] Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *Proceedings of the International Conference on Learning Representations*, 1–12.
- [64] Ramesh Inala. (2022). Cross-Domain MDM Integration Using AI-Driven Data Governance: A Case Study In Financial Technology Architecture. *Migration Letters*, 19(2), 280–304. Retrieved from <https://migrationletters.com/index.php/ml/article/view/11982>
- [65] Montoya, D. Y., Neto, A. M., & da Silva, A. S. (2016). A survey of entity resolution in big data. *Journal of Big Data*, 3(1), 1–22.
- [66] Aitha, A. R. (2021). Optimizing Data Warehousing for Large Scale Policy Management Using Advanced ETL Frameworks.
- [67] Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S., & Stoica, I. (2010). Spark: Cluster computing with working sets. *Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing*, 1–7.
- [68] Varri, D. B. S. (2022). AI-Driven Risk Assessment and Compliance Automation in Multi-Cloud Environments. Available at SSRN 5774924.
- [69] Zaharia, M., Das, T., Li, H., Shenker, S., & Stoica, I. (2012). Discretized streams: Fault-tolerant streaming computation at scale. *Proceedings of the 24th ACM Symposium on Operating Systems Principles*, 423–438.
- [70] Segireddy, A. R. (2021). Containerization and Microservices in Payment Systems: A Study of Kubernetes and Docker in Financial Applications. *Universal Journal of Business and Management*, 1(1), 1–17.
- [71] Zhai, C., & Massung, S. (2016). Text data management and analysis: A practical introduction to information retrieval and text mining. ACM & Morgan Claypool.
- [72] Davuluri, P. N. (2020). Event-Driven Architectures for Real-Time Regulatory Monitoring in Global Banking.
- [73] Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5, 135–146.
- [74] Keerthi Amistapuram, "Energy-Efficient System Design for High-Volume Insurance Applications in Cloud-Native Environments," *International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJREEICE)*, DOI 10.17148/IJREEICE.2020.81209
- [75] Goutham Kumar Sheelam. (2022). Reconfigurable Semiconductor Architectures For AI-Enhanced Wireless Communication Networks. *Kurdish Studies*, 10(2), 1027–1040. <https://doi.org/10.53555/ks.v10i2.3867>
- [76] Batarseh, F. A., & Yang, R. (2019). *Federal data science: Transforming government and society*. Academic Press.
- [77] Gottimukkala, V. R. R. (2021). *Digital Signal Processing Challenges in Financial Messaging Systems: Case Studies in High-Volume SWIFT Flows*.
- [78] Bhasin, H., & Bhatia, P. (2020). Clickstream data mining for web analytics and customer behavior modeling: A review. *ACM Computing Surveys*, 53(6), 1–34.
- [79] Kolla, S. H. (2021). Rule-Based Automation for IT Service Management Workflows. *Online Journal of Engineering Sciences*, 1(1), 1–14. Retrieved from <https://www.scipublications.com/journal/index.php/ojes/article/view/1360>
- [80] Uday Surendra Yandamuri. (2022). Cloud-Based Data Integration Architectures for Scalable Enterprise Analytics. *International Journal of Intelligent Systems and Applications in Engineering*, 10(3s), 472–483. Retrieved from <https://ijisae.org/index.php/IJISAE/article/view/8005>
- [81] Abedjan, Z., Golab, L., & Naumann, F. (2016). Profiling relational data: A survey. *The VLDB Journal*, 24(4), 557–581.
- [82] Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. *International Journal of Scientific Research and Modern Technology*, 1(12), 227–237. <https://doi.org/10.38124/ijsrmt.v1i12.1111>
- [83] Dwaraka Nath Kummari. (2022). AI-Driven Audit Frameworks For Enhancing Compliance In Modern Manufacturing Systems. *Migration Letters*, 19(S8), 2150–2177. Retrieved from <https://migrationletters.com/index.php/ml/article/view/11912>
- [84] Davuluri, P. N. Event-Driven Compliance Systems: Modernizing Financial Crime Detection Without Machine Intelligence.
- [85] Baesens, B., Van Vlasselaer, V., & Verbeke, W. (2021). *Fraud analytics using descriptive, predictive, and social network techniques: A guide to data science for fraud detection (2nd ed.)*. Wiley.



- [86] Avinash Reddy Aitha. (2022). Deep Neural Networks for Property Risk Prediction Leveraging Aerial and Satellite Imaging. *International Journal of Communication Networks and Information Security (IJCNIS)*, 14(3), 1308–1318. Retrieved from <https://www.ijcnis.org/index.php/ijcnis/article/view/8609>
- [87] Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning*. fairmlbook.org (Book manuscript).
- [89] Garapati, R. S. (2022). AI-Augmented Virtual Health Assistant: A Web-Based Solution for Personalized Medication Management and Patient Engagement. Available at SSRN 5639650.
- [90] Gottimukkala, V. R. R. (2020). Energy-Efficient Design Patterns for Large-Scale Banking Applications Deployed on AWS Cloud. *power*, 9(12).
- [91] Ahmad, M. A., Eckert, C., & Teredesai, A. (2018). Interpretable machine learning in healthcare. *Proceedings of the ACM Conference on Health, Informatics, and Data Science*, 1–10.
- [92] Aitha, A. R. (2022). Cloud Native ETL Pipelines for Real Time Claims Processing in Large Scale Insurers. Available at SSRN 5532601.
- [93] Aljabre, A. (2019). Cloud computing security in healthcare. *Journal of King Saud University – Computer and Information Sciences*, 31(1), 10–18.
- [94] Kolla, S. K. (2021). Architectural Frameworks for Large-Scale Electronic Health Record Data Platforms. *Current Research in Public Health*, 1(1), 1–19. Retrieved from <https://www.scipublications.com/journal/index.php/crph/article/view/1372>
- [94] Akanfe, O. A. (2022). *Advancing digital financial inclusion: Data privacy, regulatory compliance, and cross-country cultural values in digital payment systems use* (Doctoral dissertation, The University of Texas at San Antonio).
- [95] Avinash Reddy Segireddy. (2022). Terraform and Ansible in Building Resilient Cloud-Native Payment Architectures. *International Journal of Intelligent Systems and Applications in Engineering*, 10(3s), 444–455. Retrieved from <https://www.ijisae.org/index.php/IJISAE/article/view/7905>
- [96] Kothapalli Sondinti, L. R., & Syed, S. (2022). The Impact of Instant Credit Card Issuance and Personalized Financial Solutions on Enhancing Customer Experience in the Digital Banking Era. *Universal Journal of Finance and Economics*, 1(1), 1223. Retrieved from <https://www.scipublications.com/journal/index.php/ujfe/article/view/1223>
- [97] Crisanto, J. C., Leuterio, C. B., Prenio, J., & Yong, J. Regulating AI in the financial sector: Recent developments and main challenges. *FSI Insights on Policy Implementation*, (63).