Kurdish Studies
Dec 2022
Volume: 10, No: 2, pp 1273-1283
ISSN: 2051-4883 (Print) | ISSN 2051-4891 (Online)
www.KurdishStudies.net

DOI: 10.53555/ks.v10i2.4069

# Cloud-Native Data Platform Modernization for Regulatory Compliance in Global Banking

**P S L Narasimharao Davuluri**[1*]

[1*]Assosciate Principal Data Engineering pslnarasimharao.davuluri@ieee.org, ORCID ID: 0009-0009-0820-8184

**Abstract**
A global banking organization migrated data platforms from on-premises architectures to public cloud services enabling scale, elasticity, and cost efficiency. Data residency and support for anti-money laundering monitoring were identified as key regulatory requirements impacting data residency strategy. Therefore, data platform services must store copies of specific data sets in defined jurisdictions. Requirements and implementation considerations were defined for the modernization of the shared data lake and data warehouse platform services. Enabling exposure required significant improvement across multiple capabilities.

Capability deficiencies were assessed across three distinct dimensions: governance-supporting processes and controls; data lifecycle management processes required to generate and maintain exposure evidence; and observability and metadata management requirements, including integration with existing data-sharing platforms and compliance risk assessment tooling. A compliance-by-design governance framework was created to support future security, privacy, and performance isolation requirements. Security policies enforcing data residency at rest were enhanced for data transmission and for supporting encrypted queries across data-sharing services.

**Keywords :** Regulatory compliance; cross-border data flow; financial crime compliance; data governance; metadata management; legal obligations; data lifecycle; risk assessment; operating model.

## 1. Introduction

Digital transformation in the banking industry is primarily driven by new technologies and customer needs, yet regulatory requirements also play a crucial role. Cloud computing enables banks to modernize their data platforms as part of business transformation, yet many cloud-native designs do not meet the requirements of data regulatory and governing bodies. This is especially true for banks with operations in multiple countries and regions, which must meet multiple regulatory requirements, such as those related to data residency, financial crime, and anti-money laundering. Data platforms that deviate from regulatory requirements increase the cost of data management, as additional processes and controls are needed to ensure compliance.

Regulatory compliance is often not considered in cloud-native platform designs. Business and regulatory requirements must be integrated into a holistic governance framework centred on regulatory compliance provisioned as part of the data platform, leading to a cloud-native platform design where compliance is considered in the workflow and business logic. A case study in global banking is used to illustrate the need for regulatory compliance in platform design. A compliance-by-design framework enables integration of the processes and controls needed for constant, global compliance into data pipelines and workflows. The result is a data platform architecture that meets regulatory requirements by design, with business value and regulatory compliance mutually reinforcing rather than conflicting objectives.

### 1.1. Background and Significance

Data governance and compliance enforcement have emerged as key requirements for global data platform facilities supporting analytics projects. A recent project for a large, global bank led to the design of privacy, performance, reliability, risk management, and cross-border compliance mechanisms for a cloud-native data platform facility. Regulatory considerations cognizant of the environment's wide reach include data sovereignty and residence, financial crime processes, reliability and privacy standards, and quality of service measures. The bank's data governance and compliance requirements informed a design approach, including a compliance-by-design governance method supported by risk assessments and controls associated with on-boarded data assets. Regulatory and legal stakeholders are required in up-front design nominations, and process metadata capture enables automated risk assessment and control mapping. Control execution and ownership remain with project teams, who leverage the compliance tools as compliance-enforcing data lifecycle capabilities.

The requirements of the global regulatory landscape for banking highlighted the need for enablement therefore within a cloud-native data platform facility. Regulatory considerations cognizant of the environment's wide reach include data sovereignty and residence, financial crime processes, reliability and privacy standards, and quality of service measures. The bank's data governance and compliance requirements informed a design approach that included a compliance-by-design governance method supported by risk assessments and controls associated with on-boarded data assets. Regulatory and legal stakeholders are required in up-front design nominations, and process metadata capture enables automated risk assessment and control mapping. Control execution and ownership remain with project teams, who leverage the compliance tools as compliance-enforcing data lifecycle capabilities.

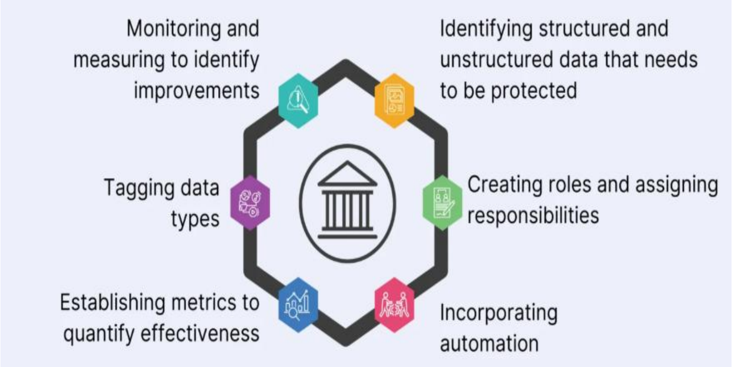
Monitoring and measuring to identify improvements
Identifying structured and unstructured data that needs to be protected
Tagging data types
Creating roles and assigning responsibilities
Establishing metrics to quantify effectiveness
Incorporating automation

**Fig 1: Cloud-Native Data Platform Modernization**

### 1.2. Research design

Case study research and design science were used to develop a cloud-native data platform architecture that aligns with data compliance requirements for cross-border data residency and AML financial regulations in a large, international bank. Compliance, legal, and IT experts operated as a collaborative design team, supported by data architects and engineers with cloud platform implementation experience. The outcome comprises an architectural blueprint, a compliance-by-design governance framework that integrates business and technology controls, a metadata management plan for regulatory observability, and a prioritized compliance risk assessment that identifies responsibilities for implementing technology controls. The contribution serves as a practical governance guide for data platform operations and enables compliance with additional regulatory requirements or controls for other lines of business.

Rapid digital technology adoption enables new business models and services for banking, capital markets, and insurance. Key trends include data, analytics, cloud, and regulation. Customers reside in many countries, with operations regulated by domestic and cross-border legislation. Security, fraud, crime, and consumer protection remain highly regulated for a stable financial system. In regulated markets, data compliance requirements on cross-border data residency and anti-money laundering financial crime laws are especially challenging. Data must be stored, processed, transferred, and disclosed in accordance with data regulations undertaken by local entities and by global organizations using cloud services. Data platforms that act as technology intermediaries have specific challenges in meeting regulatory requirements.

## 2. The Regulatory Landscape in Global Banking

Regulatory constraints are a key consideration for organizations in response to the governance and compliance risks of using public cloud services to deliver a cloud-native data platform. These risks can be grouped into two broad categories: (1) compliance with the residency and sovereignty of cross-border data transfer, storage, and processing in cloud-based storage services provided by non-resident cloud service providers and (2) conformity with banking and financial crime regulations, which mandate the set of technologies for anti-money laundering, anti-terrorism financing, fraud detection, and similar capabilities needed for defending against financial crime. Although no single organization can satisfy the whole set of regulatory requirements of different jurisdictions, a global bank needs to comply with the regulations of the jurisdictions in which it operates. Compliance can be achieved by selecting the appropriate data that must adhere to specific regulatory requirements and hosting this data on the cloud platform in a manner that facilitates monitoring through controls and processes.

Regulatory compliance can be accomplished administratively by implementing approval processes for data transfer, storage, and transfer technology selection. However, the increasingly rapid pace of business change imposed by customers and the growing volume of data necessitate the automation of these processes by embedding regulatory controls into data platform services, thus establishing a compliance-by-design approach. The absence of an automated compliance framework may lead to the unintentional breach of privacy or regulatory controls, resulting in significant fines for the organization, substantial reputational damage, and a loss of customer trust. Consequently, the demand for cloud-native data platforms is outpacing the available supply, driven by the volume and velocity of data. Automation of compliance for the high-velocity data produced by customer interactions may not be achievable unless the detection capability becomes self-sufficient and the risk is borne by the data owner.

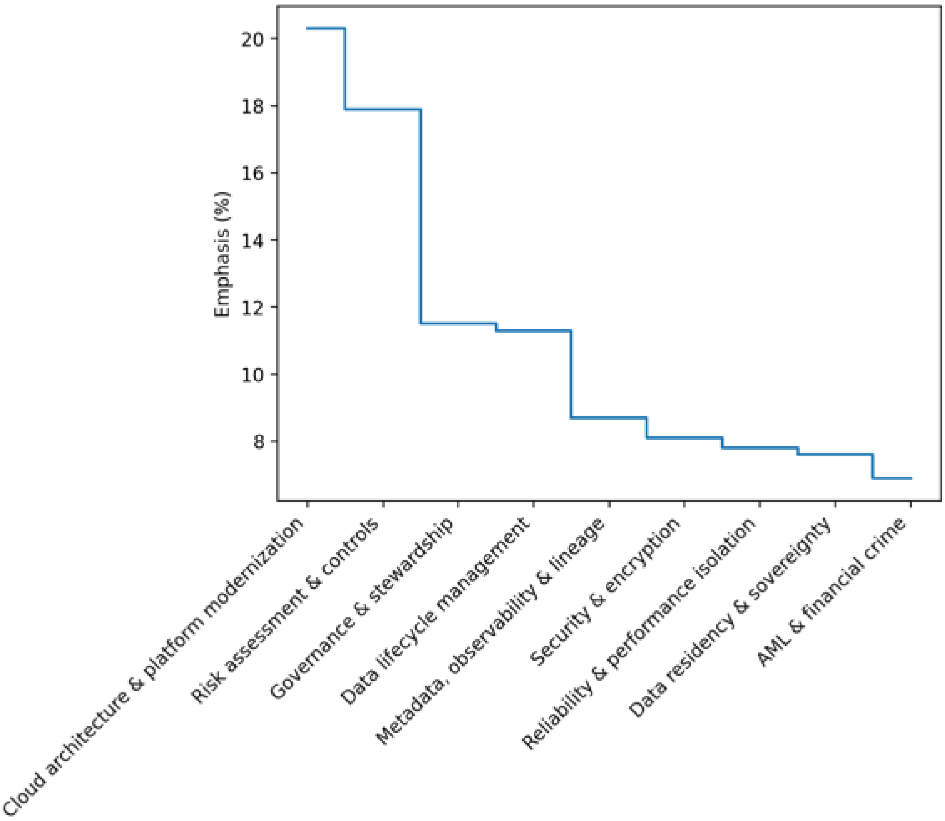

**Fig 2: Discrete Variation of Thematic Emphasis Across Modernization Domains**

### 2.1. Cross-Border Data Residency and Sovereignty

In North America and Europe, cross-border data residency requirements demand that certain types of sensitive customer data—including personally identifiable information and information identifying the customers' ultimate beneficial owners or controllers—be stored in dedicated on-premises infrastructure located within the jurisdiction. In Asia/Pacific and Latin America, legislation in various countries imposes similar but distinct data residency legislation, requiring that at least specific types of sensitive customer data remain in on-premises infrastructure located within the jurisdiction. The observability features of modern cloud-native data-management platforms provide the metadata associated with data-assets storage for all forms of risk, including cross-border data-residency, that are derived from geographic location. These features also make it possible to define and enforce separate-forensic environments providing testing and troubleshooting capabilities without any ability to extract data from the environment, thus enabling companies in heavily regulated sectors to operate in the cloud. Across global financial institutions, the recent growth of clouds in China and Russia has raised questions concerning both data-residency and data-sovereignty risk, highlighting the consideration that must be applied to public-cloud and SaaS service choices in these regions. In regions where cross-border data-residency and data-sovereignty risk must be considered, secure connectivity to on-premises data is the first line of defence. This is complemented by the availability of an adequate set of dedicated regional cloud services to meet data needs, thereby enabling the cloud-region selection to be governed by factors such as price and performance. Regions with a limited service footprint often end up being secondary markets for global finance, offering limited product choice at higher prices. For products that either do not reside in those regions or have not been integrated with local data residencies, private-cloud or on-premises offerings are required. For data and analytics workloads not constituting cross-border data-residency risk, Cloud and SaaS services in these jurisdictions can and should be deployed for everything except trading-related production.

### 2.2. Financial Crime and Anti-Money Laundering Compliance

Implementing effective systems and controls to detect and prevent financial crime across multiple markets and jurisdictions is one of the largest challenges that the global banking sector faces in its efforts to meet diverse and national laws and regulatory obligations. Banks and other financial institutions are required to implement not only "know your customer" (KYC) standards, but also "know your customer's customer" (KYCC) and "know your customer's customer's customer" standards, which, especially in Asia, may extend across the full network of subsidiary and branch relationships. The shadow economies of certain jurisdictions generate large volumes of illicit proceeds, which create an enormous testing and monitoring burden for banks as they seek to satisfy the requisite levels of transaction surveillance. These controls must operate effectively not only for transacting customers, but also for counterparties, other banks, remittances and payments from high-risk jurisdictions, and all aggregates routed through gaming houses and casinos.

Detecting and preventing money laundering are not only operational challenges; they are seemingly insatiable regulatory demands. Criminal networks employ complex multi-jurisdictional channel chains for value transfer and real economy facilitation, using a variety of techniques, including moving value through the blockchain, and dark web user liquidity zoekt, which deter detection and reduce regulatory compliance effectiveness. In addition, the complex architecture of multi-entity/branch financial institutions poses challenges in maintaining transaction monitoring rules that are comprehensive and consistent across all jurisdictions.

**Equation 1: Provenance integrity with cryptographic hashing**

Let $x$be the raw extracted data (bytes). A hash is:

$$h = H(x)$$

To validate integrity later, recompute:

$$h' = H(x')$$

Integrity holds if:

$$h' = h$$

**Step-by-step**

1. At ingestion, capture raw data $x$(paper: "raw form, without transformation"
Cloud-Native Data Platform Mode…
).
2. Compute a cryptographic digest $h$.
3. Store $h$in metadata / catalog alongside provenance fields (source, timestamp).
4. On audit, hash the retrieved data $x' \rightarrow h'$.
5. If $h' = h$, the dataset has not changed since ingestion.

**3. Cloud-Native Data Platform Architecture for Compliance**

Public clouds have rapidly become critical infrastructure for global enterprises, for whom regulatory compliance is paramount. A cloud-native data platform accelerates innovation and cost-efficiency, yet regulatory demands require an organisational architecture that systematically embeds controls into the design and processing lifecycle.

Compliance is inherently a governance process. Regulatory needs emerge at different lifecycle stages and thus require design-and-build, run, and decommission controls. A centralised compliance-by-design framework follows a process-based analysis of the data governance model, regulatory risk assessment, and control mapping. Controls then audit actual data and platform usage. Such a model ensures the observability of every data asset, providing a complete register of who did what, where and when, with capability—and evidence—to convince a regulator of the ability to control risk and comply with policy on demand. Regulators are increasingly interested in understanding the installed base, data ownership, and usage terms before approving a cloud provider. Data-stewardship roles must therefore be established and empowered at a working level and then escalated to decision-making and adjudication bodies for the use of sensitive assets. Data origin, in-market obligation, and ownership are key. For example, in compliance, production data must remain in market, whilst non-production workloads can move to a lower-cost region. Key resources and data-asset importance must also be clear to the reliability and risk teams.

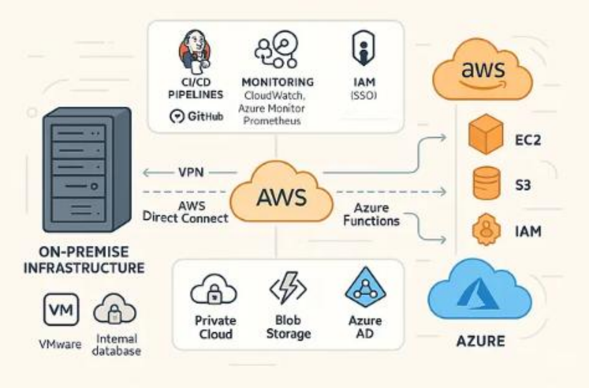

**Fig 3: Hybrid Cloud Architecture Explained**

**3.1. Data Governance and Stewardship in the Cloud**

The accelerating adoption of cloud-based services is shifting the production and operational footprints of businesses. A bank's public cloud platform may process data from its Retail business around the globe, contributing expertise without serving customers in the cloud region. Migration of dedicated business functions and online banking may also replicate such compliance footprints. As businesses and technology operations move rapidly to the cloud and data processing cross border boundaries, ensuring business-sensitive cloud-based regulatory compliance is increasingly challenging.

A global bank's enterprise data platform ensures cloud-based compliance for business functions. Compliance requirements for financial crime-related data are often clearer, as are data governance requirements for cross-border data and data associated with retail customer activities. Modern Data Platforms offer new capabilities to ingest, structure, share, and disseminate data at scale and to monitor what happens between input and output. Regulatory compliance by design can identify compliance requirements and derive processes and controls mapped to the relevant cloud-based data lifecycle stages.

### 3.2. Metadata Management and Observability

Achieving regulatory compliance with data assets in a cloud data platform requires sufficient metadata to satisfy the needs of regulators as well as internal stakeholders such as risk, legal, compliance, fraud, and operations teams. A cloud-native approach to metadata discovery and management reduces the effort for all data platform users and operators and helps the organisation comply with laws and regulations such as the General Data Protection Regulation (GDPR), Health Insurance Portability and Accountability Act (HIPAA), and cross-border data transfer restrictions. From a compliance-by-design perspective, compliance-enabling metadata should therefore be captured and managed as a first-class entity by the data platform, with the use of a data catalog for discovery, sharing, and governance.

Compliance-related metadata can be generated automatically. Integrations with software development kits (SDKs) and application programming interfaces (APIs) for dataingestion flows can capture metadata on data provenance, through the use of cryptographic hashing. Streaming services can automatically capture metadata on patterns for fraud detection and on suspect transaction monitoring. Data engineering processes can track transformation logic and sensitive data masking to cater for data privacy regulations. Access management systems can manage consent for sensitive data sharing, enforce security zones for different data classifications, track user access requests, approvals, and notifications, and provide transparency for access requests made by other parties such as law enforcement agencies. A data cataloging capability can support discovery and sharing of sensitive-data metadata among data platform users.

## 4. Regulatory Compliance by Design: Processes and Controls

Evidence of regulatory compliance is a prerequisite for operating a global banking business. Designing regulatory compliance into a cloud-native data platform requires a comprehensive framework of processes and controls. A compliance-by-design governance framework identifies the execution-related roles and responsibilities, the risks to be managed by functional controls, and the controls that require evidence of execution for business operations. The framework informs a regulatory-compliance risk assessment that maps the business process control statements to platform data assets and platform processes.

Compliance-by-design governance framework. Regulatory compliance cannot be an afterthought. Evidence of compliance is usually a prerequisite for operating a global banking business. Still, many organisations struggle to demonstrate their ability to fulfil their regulatory obligations when requested. The planned cloud-native data platform provides a governmental data-centric view of the organisation and serves as a basis for continuous compliance "by design" through a comprehensive range of processes and controls that meet the data and analytic needs of global operations. A compliance-by-design governance framework identifies the execution-related roles for these processes and the risks that need management by functional controls. The framework also identifies the controls that require evidence of execution for business operations. These processes and controls ensure that evidence of compliance is readily available.

### 4.1. Compliance-by-Design Governance Framework

Compliance-by-design within a cloud-native data platform entails a compensation-by-design governance framework with responsibilities, processes, and technologies for implementing compliance controls in regulatory-sensitive areas such as data protection, financial crime, and anti-money laundering. Governance lenses raise compliance-related considerations during all phases of the data lifecycle and across all types of data support by the data management platform. In line with the bank's risk management framework, data protection impact assessment (DPIA) for a new data product involves not only privacy and technology compliance teams but also risk managers, data governance councils, legal advisors, and wider stakeholders potentially affected by the data product (e.g., fraud and AML teams). The DPIA exercise defines potential risks and mitigation measures that regular observability can track.

Compliance-by-design governance lens enable regulatory-sensitive data to be ingested following country-specific regulations, with systems supporting ingestion fully aware of data provenance. On data use, data recipe and data product creators are responsible for interpreting regulatory requirements related to privacy or financial crime. Data products containing risk-rated attributes have minimum standards for provisioning and are subject to business validation and acceptance procedures that include risk managers, data protection officers, and financial crime compliance specialists.

### 4.2. Risk Assessment and Control Mapping

A risk assessment and control-mapping exercise, leveraged by the platform-level data controls, provides a foundation for generating risk assessments and control mappings for regulated workloads—including controls such as data encryption, logging, and monitoring, as well as controls supporting data residency and sovereignty needs. With a logical separation between the platform and workload-level needs, technologies catering to demand isolation enable data privacy and compliance without impacting the performance of independent workloads or their business services.

The controls identified through this mapping underpin the Compliance-by-Design governance framework, effectively establishing a set of predefined controls for workloads on the shared regulatory fabric. In addition, the mapping of risks to controls is an important enabler of regulatory audit readiness efforts. By automatically maintaining an audit evidence repository for data stored outside the Data Store Technology Domain, the combination of compliance-by-design processes with continual data discovery automates collection for a key set of audit requirements, breaking down the burdens placed on security and compliance functions.

**Equation 2: Risk assessment & control mapping (quantitative form)**

Let:

- Likelihood $L$(e.g., 1–5)

- Impact $I$(e.g., 1–5)

$$R = L \times I$$

**Residual risk after controls**

Let control effectiveness $E \in [0,1]$(0 = useless, 1 = perfect):

$$R_{\text{res}} = R \times (1 - E)$$

**Step-by-step**

1.Start with inherent risk $R$as a product of likelihood and impact.
2.Controls reduce risk by a fraction $E$.
3.Remaining fraction is $1 - E$.
4.Multiply to get residual risk.

## 5. Data Lifecycle Management for Regulatory Needs

Strong data lifecycle management is essential across several regulated industries. Consequently, the data platform architecture must include processes, controls, and tooling to support specific data lifecycle aspects for regulated workloads.

Data Ingestion and Provenance

Regulated workloads often require evidence of where data came from and proof that it is unaltered. Consequently, a regulated workload platform ensures data ingestion tools automatically check for and capture all product data, as well as key metadata, as they enter the cloud-native data platform. For these workloads, data extraction tools are configured to grab data in its raw form, without transformation or aggregation. Provenance information, such as the source and time of extraction, is automatically captured, recorded, and retained with the data.

Data Processing and Transformation with Lineage

Regulated workloads often have rigorous processing requirements that include construction of new or aggregate datasets from incoming product data feeds. Such transformations need to be implemented and executed in a manner that guarantees data quality and integrity over the entire transformation process. A regulated workload platform provides reliable execution environments, whether batch or streaming, for the implementation of such transformations. Reliability is enabled both through platform reliability engineering and through cloud-native data platform logging and observability.

Improved Data Quality
Better Data Security
Reduced Storage Costs
Better Regulatory Compliance
Increased Operational Efficiency
Simplified Data Management
Improved Decision Making

**Fig 4: Data Lifecycle Management**

### 5.1. Data Ingestion and Provenance

A cloud-native data platform in a regulated environment must formally define how data enters the environment—from data sources such as databases, APIs, or third-party data provider services—along with the configuration that enables this process. The solution must ensure that the data, once ingested, retains enough lineage information to enable authorized users to trust its origin and source.

A common solution to data ingestion is an accelerator that drives the creation of ingestion definitions with appropriate metadata. The engine allows a business user to define a connection to a source database via an API form while ensuring the parameters are filled appropriately and not missing.

The accelerator starts the data movement process in a manner that respects the organization's rules regarding the locality of the data at rest. For example, if the user-defined connection points to a source database in a county where data may not leave its borders, the definition is forwarded to the ingestion team responsible for the ingestion of this data in full compliance with the laws and regulations governing the location.

### 5.2. Data Processing and Transformation with Lineage

In support of regulatory requirements, cloud-native data platforms must manage data throughout the end-to-end data lifecycle. This includes implementing production-grade processing routines that are resilient and performant for the high throughput required in the capital and transaction systems. Unit-tested processing must be orchestrated with telemetry monitoring in place for alerts and corrective actions. Data storage must provide isolation across sensitive workloads such as private client banking. Lineage tracking and observability must support antifinancial crime, internal investigation, discovery, and associated operations.

Compliance with the data elimination duties in connection with legal and regulatory requests must be built into the platform processes at ingestion, processing, and storage. Data that is subject to data residency or sovereignty risks must be rejected at ingestion time or de-identified before departure from the control boundary. Data mandated for deletion within predefined timeframes must be purged by the processing or storage jobs responsible for their generation. Lineage integration must extend to these elimination routines, with access requested through conventional legal or regulatory channels.

**Table 1: Regulatory Drivers Impacting the Cloud-Native Data Platform**

| Regulatory Domain | Key Requirements | Platform Impact | Why It Matters |
|---|---|---|---|
| **Cross-Border Data Residency** | Sensitive data must remain within specific jurisdictions | Regional storage policies, geo-fencing, residency enforcement controls | Avoids violations of national data sovereignty laws |
| **Data Sovereignty** | Control over where data is processed and who can access it | Region-aware processing, encryption, access zoning | Prevents unlawful foreign access |
| **Financial Crime (AML, KYC, KYCC)** | Monitoring of transactions, relationships, and beneficial ownership | High-throughput analytics, lineage tracking, fraud detection pipelines | Mandatory for preventing money laundering and terrorist financing |
| **Privacy Regulations** | Protection of PII and sensitive customer information | Data masking, consent tracking, access approval workflows | Prevents privacy breaches and regulatory penalties |
| **Operational Resilience** | Continuous availability and risk monitoring | SLO-driven reliability engineering, active/passive architectures | Ensures regulatory acceptance of cloud use |

## 6. Platform Reliability, Privacy, and Performance

The cloud-native banking data platform must provide high availability and security within data regulation and privacy requirements. Achieving Service Level Objectives (SLOs) for reliability, data protection, and performance in cloud-native architectures requires collaboration in platform definition, workload implementation, and validation for Quality of Service (QoS). Close collaboration with reliability engineering and application operation teams leads to early detection and mitigation of potential availability issues.

The service delivery model must isolate reliability and privacy for regulated workloads, ensuring that critical processes attain adequate SLOs. Opting for dedicated resources rather than shareable pools is recommended to meet the needs of regulated workloads. This option not only facilitates the identification and resolution of potential performance bottlenecks based on the real workload footprint, but also enables the application of additional zoning and protection measures (network, storage, compute) to ensure regulated workloads adhere to defined QoS levels without placing the confidential or regulated nature of the data at risk and impact users accessing the information.

**Equation 3: Reliability: SLO availability equation**

Let total period time $T$and downtime $D$:

$$A = 1 - \frac{D}{T}$$

**Step-by-step derivation**

1. Uptime $U = T - D$
2. Availability is uptime fraction: $A = \frac{U}{T}$
3. Substitute: $A = \frac{T-D}{T} = 1 - \frac{D}{T}$

### 6.1. Reliability Engineering for Regulated Environments

High reliability engineering provides a structured approach to building reliable systems through failure models, fault trees, and two-dimensional reliability assessment matrices. High availability, which minimizes service downtime to maximize reliability, is critical for operational performance and revenue protection. For business-critical applications, organizations may adopt active–active architectures; besed for banking applications, active–passive architectures are common.

Despite not being mission-critical for bank customers, the customer experience is visible to users and so must be managed. Governing bodies, including the bank's board and senior executive committee, approve user experience levels given the expected level of spend. The bank must remove barriers to positive user experience, add hooks to detect issues early, and ensure integration with social media for rapid decomposing of issues. For platforms serving multiple lines of business, integration testing and production readiness clearance become especially important. Beyond sign-off by each line of business, full integration testing is essential for any production change.

### 6.2. Performance Isolation and Quality of Service

To ensure compliance with data privacy regulations and support the growth of financial services based on cloud-native architectures in a strict cross-border regulatory environment, consistent Quality of Service (QoS) must be maintained during ingestion and preparation of transactional data. Such workloads will run on multi-tenant infrastructures, where privacy-sensitive data from more than one country and jurisdiction is physically co-located. This situation exists because certain data-driven services require minimal latency to be useful and must be placed as close as possible to the user.

Multi-tenancy of these infrastructures introduces risks to the stability and reliability of workloads that process privacy-sensitive data. The ingestion, processing, and storing of data used for analytics and thus possibly exposed to data privacy regulations cannot be reliably predicted; even a single overload incident in one component could impact other components on the same infrastructure. The solution to avoid this risk relies on traffic control: an elastic, scalable proxy controller with QoS support ensures the preservation of the latency-bound SLA for all services involved in privacy-sensitive data transactions. The requirement is threefold: the injected delays must be controllable by the SLA management without modifying the business components.

## 7. Stakeholder Roles, Collaboration, and Change Management

Proper implementation of a compliance-by-design framework requires active collaboration among all stakeholders, enabling the identification and mitigation of compliance risks in a timely, structured manner. Such collaboration must continually foster an understanding of compliance requirements across the organization and create well-defined execution roles. Governance bodies must consider compliance in their change-management decisions and ensure that all changes—from a simple new data-source introduction to a global cloud service migration—are examined for compliance effect. Regulatory stakeholders must, at least periodically, meet with the data-usage community to foster discussion and confirm that any perceived compliance risk is properly understood by all parties. A key enabler is IT's identification of the compliance risks all regulatory stakeholders face and its promotion of tooling and processes that would make it easier to meet the demands.

The need for an independent control layer across data products often results in changes being backlogged, adding further regulatory burden. Agile methodologies and a compliance-by-design culture within product-development teams can assist, but if product-delivery teams do not share in regulatory-expectation ownership, it becomes difficult to drive change. A more effective long-term solution is risk assessment and control mapping for ingestion, transformation, algorithmic management, feedback incorporation, and data-timeliness processes within the data-lifecycle-management framework. With the respective risks and control owners identified, product-delivery teams can harmoniously address compliance-by-design governance.

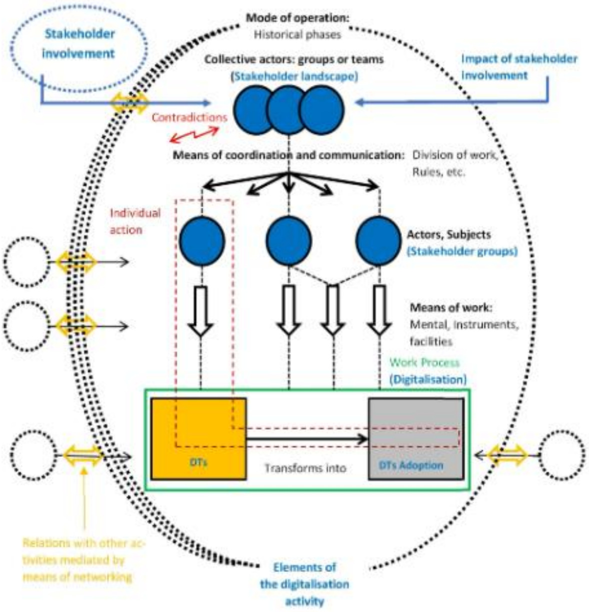
Stakeholder involvement
Mode of operation:
Historical phases
Collective actors: groups or teams
(Stakeholder landscape)
Impact of stakeholder involvement
Contradictions
Means of coordination and communication: Division of work, Rules, etc.
Individual action
Actors, Subjects
(Stakeholder groups)
Means of work:
Mental, instruments, facilities
Work Process
(Digitalisation)
DTs
Transforms into
DTs Adoption
Relations with other activities mediated by means of networking
Elements of the digitalisation activity

**Fig 5: The Influence of Stakeholder Involvement in the Adoption of Digital Technologies**

### 7.1. Governance Bodies and Decision Rights

The combinatorial nature of regulatory and compliance requirements necessitates a comprehensive cloud governance framework that incorporates the perspectives of all involved parties. Such decisions can only be made by a relevant decision-making body — for example, a runtime compliance board mandated with the responsibility to enable sound decision-making at runtime. Any action required by the compliance-by-design governance board must be made at a senior level, which is the only level at which decisions can be effectively managed and mitigated.

This approach differs from pre-cloud governance models, where decisions are made sequentially. In a fully cloud-native environment, the decision process is multiplicative and therefore has a very wide upper surface. Although care is taken in decisions regarding the implementation of steps online and in support of compliance-by-design platforms, because these decisions can be quickly evaluated, the emphasis is on an efficient and effective management process rather than formal decision-making bodies. The compliance-by-design governance board supports all cloud areas across a bank — that is, the board does not make the decision but simply informs the decision process and the action to be taken. The board provides a much-needed area of support and oversight, without slowing down the speed of innovation.

### 7.2. Legal, Compliance, and IT Collaboration

A dedicated team is required to manage the dynamic regulatory landscape. Legal and compliance experts help to interpret incoming regulations and assess their relevance and impact, while the IT organization is responsible for executing the associated changes to the cloud-native data platform. The governing bodies outlined earlier in this document are required to enable cloud capabilities in support of compliance by design.

Formalized interaction points across the domains are essential to minimize risk exposure and the response time for urgent requests. The need for regulatory support capabilities may originate from cloud governance bodies, regulators, business controls, compliance with partners or external customers, data stewardship boards, or even from within the legal and compliance teams themselves. An example of a common governance input is System and Organization Controls (SOC) reporting from third-party audits of cloud service providers. Representation and involvement of key geographic stakeholders in these activities helps facilitate efficient collaboration.

## 8. Conclusion

Emerging technologies such as artificial intelligence (AI), blockchain, and 5G are fit-for-purpose in regulatory compliance, enabling financial institutions to apply the concept of security-by-design in support of compliance-by-design initiatives. AI offers regulatory authorities enriched capabilities for monitoring accounts and activities for anti-money laundering (AML) and anti-terrorist financing. Anti-money laundering-regulatory technology (AML-RegTech) harnesses AI for these solutions. AI presents financial institutions AML-RegTech solutions that reduce the burden of their obligations through advanced detection of dubious transactions and strengthens the detection and prevention of regulatory breaches. By leveraging RegTech solutions, institutions extend their reporting universe.

Cloud, AI, and 5G convergence presents an opportunity, coupled with the broader shift to the digital economy, for regulatory authorities to rewire the way compliance requirements are met. The emergence of a shared compliance infrastructure spreads across sectors and across borders. In the future, banks execute centralized Know Your Customer (KYC) via a multi-party Cloud-and-5G ecosystem founded on fully secure(transaction privacy) and fully trusted (identity authenticity) distributive-ledger-technology and strong AI capabilities. Banks implement anti-money laundering with minimal, or even zero, step-in costs, outsourcing the creation of the system to an AML-RegTech player. Banks that are late-cycle adopters will continue to rely on traditional AML processes, burdened by higher step-in costs.

### 8.1. Emerging Technologies and RegTech Trends

Emerging technologies such as artificial intelligence (AI) and machine learning are already in use to enhance the speed and consistency of compliance processes. Providers of regulatory technology (RegTech) products and services are embracing these technologies to help regulated organisations improve their efficiency, execute compliance activities at lower cost, and reduce the opportunity for human error in areas such as financial crime detection. These techniques are also being used to automate the testing of transactions for compliance with various regulatory requirements. Machine learning can train algorithms to identify expected patterns, flagging those that deviated from the norm and prompting further investigation.

Conversely, the honeypot approach pursued by many of the large technology companies, whereby data are freely collected from consumers and business partners to drive improvements in services and monetisation opportunities, is coming under increasing regulatory scrutiny. As commercial providers of cloud technology seek to expand their global services, they are investing heavily in regions with the largest data-protection concerns. These investments highlight that the compliance challenges associated with operating in such geographies, and the corresponding natural barriers to entry, have historically prevented the creation of Uber-like models for cloud technology. As a result, the luxury of being able to process data in any location is certainly not its intended, natural use case.

## 9. References

[1] Nagabhyru, K. C. (2022). Bridging Traditional ETL Pipelines with AI Enhanced Data Workflows: Foundations of Intelligent Automation in Data Engineering. Available at SSRN 5505199.

[2] Basel Committee on Banking Supervision. (2021). *Principles for operational resilience* (BCBS 516). Bank for International Settlements.

[3] 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.

[4] Basel Committee on Banking Supervision.*Principles for the effective management and supervision of climate-related financial risks*. Bank for International Settlements.

[5] 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

[6] European Parliament & Council of the European Union. (2022). *Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA)*. Official Journal of the European Union.

[7] Amistapuram, K. (2022). Fraud Detection and Risk Modeling in Insurance: Early Adoption of Machine Learning in Claims Processing. Available at SSRN 5741982.

[8] European Banking Authority. (2021). *EBA report on outsourcing to cloud service providers*. European Banking Authority.

[9] European Central Bank. (2020). *Guide on outsourcing cloud services to cloud service providers*. European Central Bank.

[10] 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.

[11Prudential Regulation Authority. (2021). *Outsourcing and third party risk management* (Supervisory Statement SS2/21). Bank of England.

[12] Monetary Authority of Singapore. (2021). *Outsourcing guidelines*. Monetary Authority of Singapore.

[13] Varri, D. B. S. (2022). AI-Driven Risk Assessment And Compliance Automation In Multi-Cloud Environments. Available at SSRN 5774924.

[14] Federal Reserve. (2021). *SR 21-16: Third-party relationships: Interagency guidance on third-party relationships: Risk management*. Board of Governors of the Federal Reserve System.

[15] 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.

[16] Financial Stability Board. (2021). *Effective practices for cyber incident response and recovery*. Financial Stability Board.

[17] 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

[18] Committee on Payments and Market Infrastructures & International Organization of Securities Commissions. (2022). *Operational resilience of critical financial market infrastructures*. Bank for International Settlements.

[19] Gottimukkala, V. R. R. (2021). Digital Signal Processing Challenges in Financial Messaging Systems: Case Studies in High-Volume SWIFT Flows.

[20] National Institute of Standards and Technology. (2020). *Security and privacy controls for information systems and organizations* (NIST Special Publication 800-53, Rev. 5). U.S. Department of Commerce.

[21] 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. Retrieved from https://www.scipublications.com/journal/index.php/ujbm/article/view/1352

[22] National Institute of Standards and Technology. (2022). *Supply chain risk management practices for federal information systems and organizations* (NIST Special Publication 800-161, Rev. 1). U.S. Department of Commerce.

[23] 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.

[24] Cloud Security Alliance. (2022). *Security guidance for critical areas of focus in cloud computing (v4.0)*. Cloud Security Alliance.

[25] Aitha, A. R. (2022). Cloud Native ETL Pipelines for Real Time Claims Processing in Large Scale Insurers. Available at SSRN 5532601.

[26] International Organization for Standardization. (2022). *ISO/IEC 27002:2022 Information security, cybersecurity and privacy protection—Information security controls*. ISO.

[27] Garapati, R. S. (2022). AI-Augmented Virtual Health Assistant: A Web-Based Solution for Personalized Medication Management and Patient Engagement. Available at SSRN 5639650.

[28] European Union Agency for Cybersecurity. (2022). *Threat landscape for supply chain attacks*. ENISA.

[29] 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. Retrieved from https://www.scipublications.com/journal/index.php/ujcsc/article/view/1348

[30] Nambiar, R., Poess, M., Masland, A., Taheri, A., & Bhandarkar, M. (2021). Benchmarking modern cloud data warehouses. *Proceedings of the VLDB Endowment, 14*(12), 3165–3177.

[31] 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.

[32] Zhamanov, A., Sakhiyeva, Z., Suliyev, R., & Yessenova, Z. (2021). Review of data governance: Data quality, lineage, and metadata management. *Procedia Computer Science, 193*, 169–176.

[33]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

[34] Al-Ruithe, M., Benkhelifa, E., & Hameed, K. (2019). Data governance taxonomy: Cloud versus non-cloud. *Sustainability, 11*(1), 95.

[35] 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.

[36] Bonifati, A., Chiticariu, L., Cudré-Mauroux, P., Dong, X. L., Lin, I., Mazumder, P., & Stoyanovich, J. (2021). Data management for AI: Challenges and opportunities. *Proceedings of the VLDB Endowment, 14*(12), 3218–3228.

[37] Inala, R. Advancing Group Insurance Solutions Through Ai-Enhanced Technology Architectures And Big Data Insights.

[38Giebler, C., & Haux, R. (2022). Data quality management and governance in regulated environments: A systematic review. *International Journal of Information Management, 63*, 102438.

[39] 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.

[40] Fernández, A., Mondéjar, R., & Celdrán, A. H. (2021). Privacy-preserving data publishing in the era of big data: A survey. *Information Fusion, 70*, 1–12.

[41] Meda, R. Enabling Sustainable Manufacturing Through AI-Optimized Supply Chains.

[42] Alasmary, H., Alhaidari, F., & Mitra, K. (2022). Secure data processing in public clouds using trusted execution environments: A survey. *ACM Computing Surveys, 55*(6), 1–36.

[43 De, S., Jones, R., & Kolla, H. (2022). Uncertainty Propagation in Dynamical Systems via Stochastic Collocation on Model Dynamics (No. SAND2022-10601C). Sandia National Lab.(SNL-CA), Livermore, CA (United States).

[44] Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology, 10*(2), 1–19.

[45] 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, Vedaprada and Jyothi, Vinaya Kumar and Kudithipudi, Karthik, AI-Driven Cybersecurity: Enhancing Cloud Security with Machine Learning and AI Agents (February 07, 2022).

[46] Mühle, A., Grüner, A., Gayvoronskaya, T., & Meinel, C. (2018). A survey on essential components of a self-sovereign identity. *Computer Science Review, 30*, 80–86.

[47] National Cyber Security Centre. (2021). *Cloud security principles*. NCSC (UK).

[48] Rongali, S. K. (2022). AI-Driven Automation in Healthcare Claims and EHR Processing Using MuleSoft and Machine Learning Pipelines. Available at SSRN 5763022.

[49] Reserve Bank of India. (2022). *Cyber security framework in banks: Updates and supervisory expectations*. RBI.

[50] Financial Action Task Force. (2022). *Targeted financial sanctions related to terrorism and terrorist financing: Best practices paper*. FATF.

[51] 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.

[52] Arner, D. W., Barberis, J., & Buckley, R. P. (2017). FinTech, RegTech, and the reconceptualization of financial regulation. *Northwestern Journal of International Law & Business, 37*(3), 371–413.

[53]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.

[54] Gozman, D., Liebenau, J., & Mangan, M. (2018). The innovation mechanisms of FinTech start-ups: Insights from SWIFT's Innotribe competition. *Journal of Management Information Systems, 35*(1), 145–179.

[55] Jonnalagadda, A., Kulkarni, S., Rodhiya, A., Kolla, H., & Aditya, K. (2022). A study of the fourth order joint statistical moment for dimensionality reduction of combustion datasets. Bulletin of the American Physical Society, 67.

[56] Wolff, J. (2021). Model risk management in the era of machine learning: Governance, controls, and compliance. *Journal of Risk Management in Financial Institutions, 14*(3), 220–235.

[57]Banna, H., Hassan, M. K., & Rashid, M. (2021). Digital transformation and banking risk: Evidence from emerging markets. *International Review of Financial Analysis, 77*, 101838.