Sudhir Vishnubhatla, 2018, 6:1
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

# From Risk Principles to Runtime Defenses: Security and Governance Frameworks for Big Data in Finance

**Sudhir Vishnubhatla**

**Abstract- Financial institutions are simultaneously among the most data-intensive and the most heavily regulated industries. The rise of big data platforms, distributed file systems, event-driven ingestion backbones, and advanced analytics engines has created extraordinary opportunities for fraud detection, customer insight, and regulatory reporting. Yet the same platforms magnify risks around privacy, security, and governance. This article reviews the evolution of security and governance frameworks for big data in regulated finance from 2000 through early 2018. Drawing on international regulations (Basel III, PCI DSS v3.2, GDPR, FFIEC, FCA FG16/5), industry best practices (NIST SP 800-53, NIST Big Data Interoperability Framework), and emerging open-source governance tools (Apache Ranger, Apache Atlas), we propose a layered control architecture. Three illustrative figures—the NIST Big Data Reference Architecture, its security/privacy overlay, and the container security lifecycle—demonstrate how regulated financial institutions can align technical implementations with supervisory mandates.**

**Keywords: Big data governance, financial regulation, BCBS 239, PCI DSS v3.2, GDPR compliance, NIST reference architecture, container security.**

## I. INTRODUCTION

Financial institutions today operate under a dual imperative: the relentless drive to maximize the strategic value of data while simultaneously complying with some of the strictest supervisory frameworks in the world. Unlike other sectors where data innovation can proceed in relatively permissive environments, banks, insurers, and securities firms must reconcile their appetite for advanced analytics with regulatory regimes that demand transparency, integrity, and accountability at every step.

Since the early 2000s, regulators have recognized that the financial system's systemic risk is inseparable from its information flows. This recognition produced a steady progression of mandates designed to enhance data governance. For instance, the Basel Committee on Banking Supervision's BCBS 239 principles on risk data aggregation (2013) codified requirements for banks to demonstrate accuracy, completeness, timeliness, and adaptability in their reporting processes. At the same time, sector-specific standards such as the Payment Card Industry Data Security Standard (PCI DSS v3.2), which took effect on February 1, 2018, imposed granular obligations for cardholder data protection, including encryption, key management, network segmentation, and continuous auditing. These mandates reflected a broader regulatory trend: moving from general exhortations about "sound risk management" to explicit, testable controls embedded in IT infrastructures.

In parallel, governments have legislated privacy and data protection frameworks with sweeping implications for financial services. The European Union's General Data Protection Regulation (GDPR, 2016) stands out as the most consequential, introducing principles of privacy by design, rights of access and erasure, and strict breach notification requirements. Similar guidance emerged worldwide, from the Federal Financial Institutions Examination Council (FFIEC) in the United States, which emphasized third-party risk in cloud arrangements, to the UK Financial Conduct Authority's FG16/5 guidance (2016) and the Monetary Authority of Singapore's Technology Risk Management Guidelines (2013). Collectively, these frameworks ensured that the adoption of emerging big-data technologies could not occur in a regulatory vacuum but had to be coupled with governance models enforceable by both internal audit and external supervisors.

At the same time, the technical substrate of financial IT underwent a profound transformation. For decades, relational databases and highly controlled

data warehouses served as the backbone for batch-driven reporting and business intelligence. These systems, while reliable, were inherently rigid and poorly suited for modern requirements such as real-time fraud detection, intraday liquidity monitoring, or millisecond-level customer personalization. By the mid-2010s, the convergence of open-source ecosystems and cloud innovation introduced distributed, cloud-native platforms capable of handling unprecedented scale and heterogeneity. Technologies such as Hadoop for distributed storage and batch processing, Apache Kafka for event streaming, Apache Spark for real-time analytics, and hyperscaler object stores like Amazon S3 or Google Cloud Storage for durable archival rapidly became common backbones of financial IT pipelines.

This shift fundamentally altered the risk and compliance landscape. Where legacy systems allowed for centralized governance through tightly coupled architectures, distributed platforms fragmented data flows across clusters, clouds, and microservices. Security controls could no longer be enforced only at the perimeter; they had to permeate every layer of ingestion, transformation, storage, and consumption.

The result is an operational environment where the "three Vs" of big data volume, velocity, and variety directly intersect with non-negotiable governance mandates. Each terabyte ingested, each millisecond latency reduction, and each new data source connected must be evaluated not only for its business utility but also for its compliance posture and resilience under regulatory scrutiny.

Thus, by early 2018, the financial industry stood at a crossroads. On one hand, distributed, cloud-native ecosystems promised unprecedented agility and analytical power. On the other, regulators demanded a rigor of governance that these ecosystems were not originally designed to deliver. Bridging this divide became the central challenge for architects of security and governance frameworks for big data in regulated financial industries.

## II. REGULATORY DRIVERS OF BIG DATA GOVERNANCE

The foundation of governance in big-data platforms within financial services is not voluntary best practice but rather the outcome of a dense and evolving web of regulatory mandates. These mandates arose in response to systemic crises, technological shifts, and growing public concerns about privacy, and they have progressively hardened expectations for how financial institutions manage their information lifecycles.

One of the most influential frameworks is the Basel Committee on Banking Supervision's BCBS 239 principles (2013). Issued in the wake of the global financial crisis, BCBS 239 sought to address the fragmentation and unreliability of risk data aggregation across large banks. It demanded that institutions demonstrate accuracy (data must be reconciled and free from material errors), completeness (all relevant risk exposures must be captured), timeliness (reports must be produced quickly enough to support decision-making during stress events), and adaptability (systems must generate new types of reports in response to supervisory requests). For big-data pipelines, these requirements translate into strong lineage, metadata management, and reconciliation processes—ensuring that distributed event streams and analytics outputs can withstand regulatory scrutiny in both normal and stressed conditions.

The Payment Card Industry Data Security Standard (PCI DSS v3.2, 2016) reinforced a different dimension of governance: the protection of sensitive financial data. By mandating end-to-end encryption of cardholder data, strict key management practices, robust network segmentation to isolate sensitive environments, and continuous audit logging, PCI DSS defined a baseline of technical safeguards that every financial institution handling card payments had to implement. With its February 1, 2018 enforcement deadline for new requirements, PCI DSS v3.2 served as a forcing function for banks to extend governance principles beyond policy and into the fine-grained configuration of their big-data infrastructures ensuring that Hadoop clusters, Kafka topics, and object storage buckets handling card

data were subject to the same rigor as traditional cardholder environments.

In parallel, privacy regulation gained unprecedented strength. The European Union's General Data Protection Regulation (GDPR, 2016) articulated sweeping new obligations: privacy by design required that data protection be engineered into systems from the outset; data subject rights (such as access, rectification, and erasure) mandated operational workflows to honor individual requests; and auditability required that organizations prove compliance through logs, records, and documented controls. For big-data environments where information is replicated, transformed, and persisted across multiple services the GDPR underscored the necessity of governance layers capable of not only controlling access but also providing transparency and traceability to regulators and customers alike.

Finally, supervisory guidance on outsourcing and third-party risk emphasized that governance does not end at the organizational boundary. The FFIEC's 2012 Cloud Computing Statement in the United States cautioned banks to retain oversight of vendors, asserting responsibilities for audit rights, continuity planning, and contractual clarity. The UK FCA's FG16/5 guidance (2016) made explicit that outsourcing to cloud providers did not absolve firms of accountability for data security and governance, requiring due diligence, risk assessments, and documented exit strategies.

Similarly, the Monetary Authority of Singapore's Technology Risk Management Guidelines (2013) reinforced expectations for board-level oversight, access control, and incident response. Collectively, these supervisory statements embedded governance into the very fabric of cloud adoption, reminding financial institutions that regulatory compliance is a shared responsibility that extends into vendor relationships.

Taken together, these frameworks converged on a clear message: security and governance are not discretionary enhancements but core design imperatives for big-data adoption in financial services. Each regulation, whether focused on risk reporting, payment security, data protection, or outsourcing oversight, contributes a layer of expectations that shape how distributed data platforms must be architected. The result is a governance landscape where compliance is not a separate activity but an operational principle embedded into every ingestion pipeline, storage cluster, and analytics engine.

## III. REFERENCE ARCHITECTURE FOR BIG DATA SECURITY

To translate complex regulatory mandates into engineering practice, financial institutions rely on standardized reference models that articulate the major actors, flows, and control points in a big-data ecosystem. Among these, the NIST Big Data Reference Architecture (NBDRA) has emerged as a widely recognized framework.

Figure 1 illustrates the NBDRA, showing the canonical roles of Data Providers, Data Consumers, Big Data Application Providers, Big Data Framework Providers, and the System Orchestrator. Data providers supply raw information transaction logs, market feeds, or customer interactions—that moves through stages of collection, curation, analytics, visualization, and access before reaching data consumers such as compliance officers, business analysts, or regulatory bodies.

Beneath this application layer sits the framework provider domain, which includes the essential components for processing and managing massive datasets:

- Processing engines (batch, interactive, and streaming) that enable fraud detection models, real-time liquidity monitoring, or retrospective risk analyses.
- Data organization and distribution platforms, including indexed storage and distributed file systems (e.g., HDFS, S3), that manage both archival and low-latency access needs.
- Infrastructure resources, spanning virtualized computing, networking, and physical hardware, which underpin the performance and resilience of the entire stack.

At the top, the System Orchestrator provides coordination, ensuring that all actors adhere to

agreed policies, interfaces, and workflows. Crucially, surrounding all components is a Security & Privacy Management layer, symbolizing that governance is not localized to a single function but is pervasive across the architecture. This layer addresses identity management, encryption, access control, logging, audit, and compliance enforcement controls that regulators expect to be demonstrated comprehensively, not piecemeal.

For financial institutions, the power of the NBDRA lies in its ability to map regulatory obligations onto technical responsibilities. For example, BCBS 239's requirements for accuracy and completeness can be tied to lineage and reconciliation services in the data organization layer; PCI DSS encryption mandates align with infrastructure and storage protections; GDPR's auditability and privacy-by-design provisions map onto the orchestration and access services. By embedding these obligations into a reference architecture, banks can create pipelines that are not only technically efficient but also regulatorily defensible.

In essence, Figure 1 is more than a conceptual diagram: it is a blueprint for operationalizing compliance within distributed data ecosystems, ensuring that every interaction whether between providers and consumers, or between storage and analytics engines occurs within a fabric of governance and security.

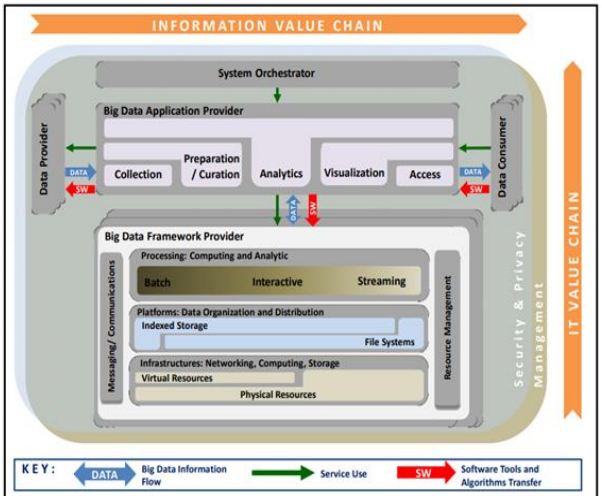


Figure 1: NIST Big Data Reference Architecture

## IV. EMBEDDING SECURITY AND PRIVACY BY DESIGN

If Figure 1 defines the structural actors of a big-data ecosystem, Figure 2—the NIST Security and Privacy Fabric Overlay— illustrates the control mesh that must enwrap them. The overlay translates high-level governance mandates into a layered set of technical and procedural safeguards, ensuring that security and privacy are not adjunct features but continuous properties of the system's design.

At its foundation, the fabric incorporates identity and access management (IAM), where policies define not just who can access data but under what contextual conditions. This extends beyond simple authentication to include role-based, attribute-based, and even policy-based encryption mechanisms, ensuring that sensitive financial datasets can only be decrypted or queried by authorized actors.

Above this lies the key management and encryption layer, addressing requirements from PCI DSS and GDPR for data confidentiality. Whether through symmetric encryption of transaction logs, asymmetric controls for cross-institutional data exchange, or advanced approaches such as fully homomorphic encryption, the overlay demonstrates how financial institutions can secure data without compromising analytical utility.

The next dimension is policy enforcement and monitoring. Here, controls automate governance principles: for instance, enforcing GDPR's "right to be forgotten" by ensuring that data deletion cascades across distributed storage layers, or guaranteeing that HIPAA-protected medical-financial data is never routed to unauthorized consumers. Policies become executable artifacts rather than static documents, monitored in real time to detect violations.

The overlay also emphasizes logging, auditing, and provenance tracking. For financial firms subject to BCBS 239's accuracy and adaptability requirements, this ensures that every data transformation is

traceable, every access logged, and every lineage verifiable. Such capabilities create an audit trail that regulators can examine and organizations can use to demonstrate accountability.

Importantly, the privacy-by-design ethos of this fabric builds directly upon Ann Cavoukian's pioneering framework (2009/2010), which argued that privacy must be proactive, embedded, and end-to-end. By operationalizing these principles in the context of distributed, high-volume, regulated financial systems, the overlay allows institutions to reconcile the competing imperatives of agility, analytics, and regulatory compliance.

In short, while Figure 1 provided a map of actors and flows, Figure 2 defines the protective sheath that ensures those flows occur within the bounds of law, ethics, and institutional trust. It makes explicit that in financial big-data systems, governance is not merely about controlling infrastructure but about ensuring that trust, compliance, and resilience are woven into the very fabric of information exchange.

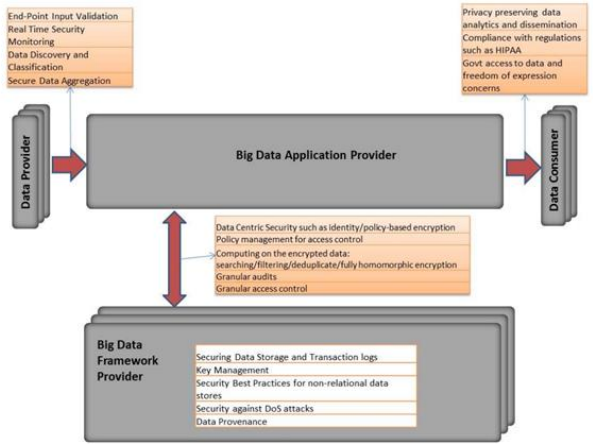


Figure 2: NIST Security and Privacy Fabric Overlay

## V. SECURING CONTAINERIZED BIG DATA SERVICES

By 2017, many financial institutions had begun migrating portions of their big-data platforms into containerized environments, often orchestrated by Kubernetes, Mesos, or Docker Swarm. This transition promised elastic scaling, workload portability, and faster development lifecycles, but it also introduced an entirely new attack surface. In a sector where regulatory compliance is non-negotiable, the security of containerized deployments became an urgent governance priority.

Figure 3, adapted from the NIST SP 800-190 Application Container Security Guide, illustrates the tiered architecture and lifecycle phases of containerized environments. It maps the flow of containerized workloads from initial development and image creation, through testing and accreditation, to storage and retrieval in registries, and finally into deployment and orchestration across clustered hosts.

The first phase, image creation and testing, involves developers packaging application code, dependencies, and configurations into container images. In financial services, this phase must be subject to rigorous code review, vulnerability scanning, and accreditation processes, ensuring that sensitive workloads such as transaction monitoring or anti-money laundering engines are built on trusted components.

The second phase, image storage and retrieval, introduces both internal registries (controlled by the institution) and external registries (third-party or vendor-provided). At this stage, key management, signing of images, and policy enforcement are essential. For example, a registry could enforce that only signed, verified images are retrievable, preventing the accidental deployment of malicious or untested builds.

The third phase, orchestration and runtime management, is handled by platforms like Kubernetes, which dynamically allocate containers to hosts, scale them under load, and ensure availability. While orchestration simplifies operations, it also creates a control plane that is highly attractive to attackers. Misconfigured orchestrators, weak admin credentials, or insufficient network segmentation can lead to lateral movement across container clusters, jeopardizing sensitive datasets.

Finally, the host layer—physical or virtual machines running the containers—represents the execution surface where runtime monitoring, intrusion detection, and granular audit logging are indispensable. Compliance frameworks such as PCI DSS v3.2 and GDPR demand real-time visibility into what processes are running, how data is accessed, and whether unauthorized exfiltration attempts are occurring.

By framing these lifecycle stages in a structured way, **Figure 3 makes clear that container security is not a single control but a chain of interdependent safeguards spanning development, registries, orchestration, and runtime. For financial institutions, this means that governance frameworks must integrate with DevOps pipelines ("shift-left" security), enforce strict registry hygiene, monitor orchestration platforms continuously, and validate runtime compliance with regulatory policies.

In practice, the adoption of this lifecycle model enables institutions to balance the speed and flexibility of containerized big-data platforms with the robust security posture required by regulators and auditors. It operationalizes the principle that agility and compliance can coexist—if the full container lifecycle is treated as a governance domain in its own right.

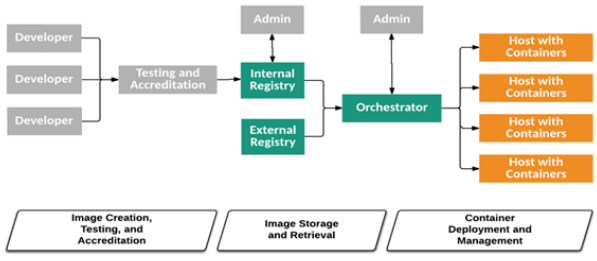


Figure 3: NIST SP 800-190 Container Architecture Tiers and Lifecycle

## VI. GOVERNANCE PRACTICES IN PRACTICE

While technical architectures and security overlays provide the scaffolding for compliant big-data systems, true governance in regulated financial industries cannot be reduced to technology alone. It must be embedded into organizational structures, decision-making processes, and day-to-day practices. Without institutional accountability, even the most sophisticated technical safeguards risk becoming "check-the-box" exercises rather than living systems of control.

Data Stewardship is the cornerstone of this organizational dimension. In a financial institution, data cannot be treated as an amorphous corporate asset; it must have clearly designated owners who are accountable for its quality, security, and lifecycle. Assigning named stewards for each dataset ensures traceability of responsibility, reduces disputes between business units, and provides regulators with a concrete governance trail. For example, the transactional dataset supporting anti-money laundering monitoring may be formally stewarded by the compliance division, while customer interaction data may be owned by the digital banking unit. These named stewards act as the bridge between technical custodians and regulatory expectations, embedding governance into organizational roles rather than leaving it solely to IT teams.

Policy Harmonization addresses the challenge of aligning external regulatory mandates with internal bank controls. Global frameworks such as NIST 800-53 control families (covering access control, incident response, system integrity, and more) provide granular security baselines. However, banks often maintain internal policies developed over decades. Harmonization involves systematically mapping NIST controls to existing bank policies and standards, identifying overlaps, gaps, and redundancies. The outcome is a unified control framework that avoids duplication, satisfies regulators, and gives internal teams a coherent set of requirements to implement. This harmonization is particularly vital for institutions operating in multiple jurisdictions, where reconciling BCBS 239, PCI DSS, GDPR, and local supervisory requirements into a single internal playbook is essential for consistency and defensibility.

Cross-Cloud Oversight has emerged as a pressing organizational necessity with the rise of hybrid and multi-cloud strategies. As banks increasingly distribute workloads across AWS, GCP, and private cloud environments, traditional governance models rooted in single-platform control are no longer sufficient. Instead, institutions must design control plane abstractions—centralized oversight layers that enforce policies, identity management, and monitoring uniformly across heterogeneous infrastructures. This oversight ensures that encryption keys, access roles, and logging standards do not fragment along provider lines, preventing the very silos that BCBS 239 sought to eliminate. From an organizational standpoint, cross-cloud oversight requires not only technology (e.g., unified IAM systems, observability platforms) but also governance councils that span IT, compliance, and business functions to set uniform cross-provider standards.

Incident Response Integration ties governance directly to operational resilience and regulatory accountability. Financial regulators worldwide now require not only rapid containment of incidents but also timely and transparent reporting of breaches, outages, and data integrity failures. Building playbooks that align with these obligations means that incident response cannot remain a purely technical domain, it must be institutionalized as a multi-stakeholder process. For instance, a playbook may stipulate that within 72 hours of detecting a GDPR-relevant breach, IT security triggers both technical containment procedures and compliance notification workflows to supervisory authorities. Similarly, PCI DSS requires that evidence of incident handling be retained for audit, while the FFIEC stresses communication with banking regulators during service disruptions. Embedding these requirements into structured playbooks ensures that response is both operationally effective and regulatorily defensible.

Taken together, these organizational practices demonstrate that governance is a socio-technical construct. Technology provides the control surface, but accountability, harmonization, oversight, and coordinated response transform those controls into a living governance system. In regulated financial industries, where the stakes of failure are measured not only in financial losses but also in reputational damage and supervisory sanctions, this holistic view of governance is indispensable.

## VII. CHALLENGES AND OPEN QUESTIONS

Despite considerable progress in developing security and governance frameworks, financial institutions continue to grapple with challenges that are structural, regulatory, and operational in nature. These challenges highlight the friction between the promise of cloud-native big-data platforms and the non-negotiable constraints of regulated industries.
Vendor Lock-In vs. Portability:

The rapid growth of hyperscale platforms such as AWS, GCP, and Azure has provided financial firms with access to highly scalable services ranging from streaming ingestion to serverless analytics. Yet, this abundance of proprietary tools creates the risk of vendor lock-in, where critical workloads become tied to one provider's ecosystem. Proprietary APIs, unique data storage formats, and differentiated identity management systems can make migration or multi-cloud adoption costly and operationally complex. For governance, this lock-in translates into a potential loss of bargaining power with vendors and reduced resilience if one provider suffers a disruption or regulatory sanction. Balancing feature-rich managed services with standards-based portability (e.g., containerization, open-source frameworks, and cross-cloud orchestration) remains an unresolved tension.

**Data Residency and Sovereignty Constraints:**
With the enforcement of GDPR in 2018 and the strengthening of national data protection laws across jurisdictions, financial firms must contend with strict data residency rules. These laws often require that customer data remain within specific geographic boundaries or be subject to local supervisory oversight. For global banks that operate across dozens of markets, this introduces significant architectural complexity: data lakes may need to be fragmented by jurisdiction, encryption keys

segregated by region, and cross-border data flows tightly monitored. Even when providers offer "regionalized" services, regulators increasingly scrutinize metadata handling, control-plane functions, and access by foreign entities, making compliance more than a matter of physical storage location.

**Operational Maturity Gaps:**
Even with advanced orchestration tools like Kubernetes and governance overlays like Apache Ranger or Atlas, many financial institutions face maturity gaps in implementation. Orchestration platforms may be deployed without robust role-based access control (RBAC), leaving administrative consoles vulnerable. Key management services may exist on paper but lack the automation, rotation policies, or integration needed for enterprise-scale resilience. Similarly, monitoring systems may provide logs but not the real-time analytics necessary for proactive threat detection. These gaps underscore that governance is not simply about adopting the right frameworks but about cultivating the operational discipline to implement, monitor, and evolve them effectively.

**Evolving Threat Landscape:**
The attack surface of big-data systems is continuously expanding. Real-time pipelines, prized for their agility, can also be exploited for rapid data exfiltration or manipulation before anomalies are detected. Meanwhile, machine learning models, increasingly embedded in fraud detection and credit scoring, present new governance challenges: adversaries can probe models to infer sensitive data (model inversion attacks) or manipulate training datasets to introduce bias (data poisoning). These threats expose the limitations of traditional perimeter-based defenses and demand governance frameworks that extend into the domains of algorithmic integrity, adversarial testing, and continuous validation of data inputs.

Together, these challenges demonstrate that while frameworks such as NIST's Security Fabric Overlay or SP 800-190 provide blueprints for compliance, financial institutions must still adapt them to a reality characterized by geopolitical constraints, organizational inertia, and adversaries who evolve as quickly as the technologies themselves. Governance, therefore, is not a destination but an ongoing practice of adaptation, negotiation, and resilience-building.

## VIII. CONCLUSION

By February 2018, security and governance frameworks for big data in financial services had evolved from scattered best practices into a layered, interdisciplinary discipline. No longer could governance be seen as a peripheral compliance activity bolted onto technical systems; it had become an architectural principle.

This maturation was anchored in regulatory mandates such as BCBS 239, PCI DSS v3.2, and GDPR, which forced financial institutions to confront issues of accuracy, auditability, confidentiality, and accountability. These mandates provided the legal and supervisory scaffolding that ensured innovation in data systems would not come at the expense of resilience or consumer protection.

At the same time, frameworks like the NIST Big Data Reference Architecture offered conceptual clarity, mapping the diverse actors and flows within distributed ecosystems. Figure 1 underscored that governance cannot be abstract: every actor, from data provider to consumer, must be integrated into a coherent system of responsibility.

The NIST Security and Privacy Fabric Overlay (Figure 2) operationalized this principle by showing that governance requires a mesh of controls identity management, encryption, key handling, auditing, and policy enforcement that cut across all stages of the data lifecycle. This embedding of "privacy by design" into architecture was a decisive step in reconciling distributed big-data systems with regulatory imperatives.

Finally, the NIST SP 800-190 Container Security Lifecycle (Figure 3) highlighted the next frontier: securing containerized deployments where big-data services increasingly resided. Containers promised agility and portability, but they demanded

governance structures that extended into DevOps pipelines, registry management, orchestration, and runtime monitoring.

Taken together, these figures and frameworks narrate the progression from conceptual architectures to operational defenses, demonstrating how financial institutions were beginning to reconcile rapid innovation with non-negotiable compliance. By embedding governance at every layer from regulation to architecture, from orchestration to runtime banks could pursue data-driven transformation while maintaining the trust of regulators, customers, and markets.

The trajectory remains unfinished. Emerging domains such as federated analytics, confidential computing, and AI governance will demand new frameworks. Yet, the foundational lesson by early 2018 was unmistakable: security and governance are not barriers to big-data innovation in finance they are its enabling conditions.

## REFERENCES

1. Basel Committee on Banking Supervision, Principles for effective risk data aggregation and risk reporting (BCBS 239), Bank for International Settlements, Jan. 2013.
2. PCI Security Standards Council, Payment Card Industry Data Security Standard v3.2, Apr. 2016.
3. European Union, General Data Protection Regulation (GDPR), Official Journal of the European Union, Apr. 2016.
4. Federal Financial Institutions Examination Council (FFIEC), Outsourced Cloud Computing Guidance, Jul. 2012.
5. UK Financial Conduct Authority (FCA), FG16/5: Guidance for firms outsourcing to the "cloud" and other third-party IT services, Jul. 2016.
6. Monetary Authority of Singapore (MAS), Technology Risk Management Guidelines, Jun. 2013.
7. C. Lynch, J. Chang, et al., NIST Big Data Interoperability Framework, Volume 6: Reference Architecture, NIST Special Publication 1500-6, 2015.
8. F. Bohn, K. Bennett, et al., NIST Big Data Interoperability Framework, Volume 4: Security and Privacy, NIST Special Publication 1500-4, 2015.
9. M. Souppaya, J. Morello, and K. Scarfone, NIST SP 800-190: Application Container Security Guide, National Institute of Standards and Technology, Sep. 2017.
10. A. Cavoukian, Privacy by Design: The 7 Foundational Principles, Information & Privacy Commissioner of Ontario, 2009.
11. ENISA, Privacy by Design in Big Data: An Overview of Privacy Enhancing Technologies in the Era of Big Data Analytics, European Union Agency for Network and Information Security, Dec. 2015.