Journal of Information Systems Engineering and Management
2025, 10(56s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Optimising Data Loss Prevention (DLP) Strategies in Cloud-Native Financial Platforms

Manohara Karakondu[1], Girish Jambagi[2], Subhash Tatavarthi[3]
[1]*Denver, Colorado*
[2]*Celina, Texas*
[3]*Celina, Texas*

ARTICLE INFO

Received: 26 Dec 2024

Revised: 14 Feb 2025

Accepted: 22 Feb 2025

ABSTRACT

Introduction: Moving financial systems to the cloud gives banks and institutions speed and flexibility, but also introduces new data protection challenges. In cloud-native environments, sensitive data such as customer information and transaction histories move across containers, APIs, and multiple cloud providers, increasing the risk of exposure or misuse. Traditional data protection measures, like firewalls, are inadequate in these decentralized settings. This article explores how financial institutions must rethink Data Loss Prevention (DLP) to secure data while supporting innovation and agility

Objectives: Cloud-native architectures present unique data protection issues for financial institutions. Key challenges include:

- Data sprawl across cloud services, making visibility and consistency difficult.
- Dynamic workloads (e.g., serverless, containers) that challenge real-time data tracking.
- Inconsistent encryption and access controls, increasing security risks.
- Regulatory complexity from rules like FFIEC, GLBA, PCI-DSS, and GDPR.
- Multi-cloud and distributed data architectures, leading to policy blind spots.
- Stateless compute environments and uncontrolled data egress via APIs.
- Security and compliance risks from incomplete monitoring and fragmented policies

Methods: A multi-layered, cloud-native DLP approach is required:

- Layered DLP Framework: Protect data at rest, in transit, and in use with encryption, access controls, and data classification.
- Automated Data Discovery: Use tools like AWS Macie or Google Cloud DLP API for ongoing scanning and classification.
- Policy-as-Code: Employ tools such as Open Policy Agent to enforce consistent policies across environments.
- API-Level DLP: Deploy service mesh tools for deep inspection and context-aware controls.
- Real-Time Monitoring: Use cloud-native monitoring and behavioural analytics for proactive detection and response.
- Compliance Alignment: Integrate data cataloguing and automated reporting to meet regulatory requirements

Results: Implementing these strategies improves data visibility, protection, and compliance in cloud-native financial systems. Automated, real-time controls reduce risk and support operational efficiency, enabling financial institutions to innovate securely and maintain regulatory trust

Keywords: cloud-native, data loss prevention, financial institutions, compliance, automation, API security.

INTRODUCTION

Moving their systems to the cloud gives banks and other financial institutions speed, flexibility, and the ability to innovate like never before. Cloud-native solutions quickly take the lead in everything from AI-powered risk assessment tools to real-time payments. Nevertheless, these benefits have serious disadvantages, especially when protecting personal data. In this quickly changing and highly regulated industry, data such as customer information, transaction histories, and trade secrets are constantly moving over complex systems made of containers, APIs, microservices, and many cloud providers.

According to this, data can be exposed or misused in more places than ever before, possibly without anyone seeing until it's too late. Financial firms that fail to preserve their data properly may face legal ramifications, reputational damage, and technological problems. Furthermore, conventional data protection techniques like putting data behind a firewall are ineffective in these new, decentralised scenarios.

This article examines how financial institutions might rethink Data Loss Prevention (DLP) to adapt to the cloud-native environment. The objective is to demonstrate how DLP can adjust to the modern world's problems, safeguarding essential data while enabling teams to work quickly, try new things, and grow with assurance[1].

OBJECTIVES

As financial institutions move towards migrating cloud-native architectures, substantial data protection challenges arise due to modern cloud ecosystems' scale, dynamic nature, and distributed characteristics. Traditional Data Loss Prevention (DLP) tools that were created for on-premises, static file systems are not crafted to address the challenges posed by the current state of cloud data sprawl and decentralized processing and, thus, expose the financial institution to a serious risk to security, compliance, and overall reputational risk. The three challenges of modern cloud environments are:

1. Data Sprawl across services and regions: Sensitive data (e.g., trade data, transaction logs, PII, etc...) now lives across distributed cloud storage, containers, APIs, and numerous providers. The data sprawl adds challenges by not only limiting visibility and consistency but also creating challenges with cross-border data flows that complicate compliance with data residency requirements such as GDPR.
2. Dynamic and Ephemeral Workloads: Unlike on-premises systems or even traditional cloud architecture, cloud-native applications rely on serverless functions and containers that are often highly scaled with ephemeral execution times, rendering static DLP systems completely useless in the context of data tracking and data protection.
3. As highlighted in the previous section, Inconsistent Encryption and Access Controls: The service provider's misconfigurations (e.g., permissive access, inconsistent use of encryption, etc.....) harm zero-trust security frameworks, leading to an increased likelihood of unauthorized access to sensitive data.

Regulatory and compliance complexity is embodied by strict regulations and laws (FFIEC, GLBA, PCI-DSS, GDPR), which demand tight control of data access, storage, and auditability. The risk of inadequate data loss prevention (DLP) could be costly, not only in terms of possible compliance violations but also through reputational exposure.

Multi-cloud and distributed architecture is used by financial services organizations that engage more than one cloud platform, and their data may be distributed across databases, storage, and message queues. While having seemingly unmanageable data locations, with a DLP policy, organizations are subject to DLP blind spots (often leaving sensitive data up to weak enforcement of policies across fragmented sites).

Stateless and ephemeral compute functions and container-based architecture can allow for microservices and serverless functions to exist for milliseconds, with sensitive data in memory or temporary storage without the enforcement of ongoing DLP agents.

Uncontrolled data egress via APIs can be the most risky data flow for financial services. While embracing APIs and third-party integrations is a critical component of operations, most APIs transfer data without real-time DLP inspection, which weakens required compliance with internal policies, governmental policies, regulatory policies, and consumer data protection frameworks.

Key Implementation Challenges Related to DLP

1. **Technical Challenges:**
   a. Limited Visibility: Data being spun up and spread across fragmented systems (databases, any APIs, on-premises, SaaS) limits the visibility needed to monitor access across networks, operating in a multi-cloud node (including, but not limited to file systems, containers, serverless, etc.).

   b. Inconsistent Encryption: Poorly managed or fragmented key management and unprotected backups create an easy pathway beyond repayment of lost data through encryption, especially in light of preventing exploits from using, in turn causing reputational loss.

   c. Ephemeral Environments: It is typically beyond the ability of DLP sensors for agent-based stateful solutions to aid automated policy enforcement in the runtime of accessibility through containers and serverless actions.

   d. API Inspection Gaps: DLP assessments on integrations may be rendered useless if the DLP approach relies on such methods as sampling in looking for content-based indications of sensitive data use. Since modern APIs create payloads via a range of transmission methods like JSON, GraphQL, etc as They traverse streams on the public internet, so it will be impossible to effectively visualize any shallow approaches to alignment policies in advance.

   e. Poor Integration: Legacy DLP systems may not integrate with cloud-native systems or tools, constricting your cloud system, too such as Kubernetes or IAM, to operate without any view of chassis systems or resources of data outside of silos or edge types.

2. **Operational Risks:**
   a. Disparate Policies: Uncoordinated DLP policies across clouds and departments create policy enforcement "holes."

   b. Insider Threats: Broad access for developers and vendors creates the risk of unintended or deliberate exposure of sensitive data.

   c. Testing: Production data moving to insecure testing environments creates the opportunity for sensitive data to be destroyed when copying, and there is no data masking.

   d. Alert Fatigue: An overwhelming number of false positives creates fatigue and overwhelms security teams.

3. **Regulatory and compliance risks:**
   a. Cross-border compliance: Many distributed clouds, alongside geographic regulations (e.g., GDPR and data residency), make compliance difficult.

   b. Audit-ability: Data logs that are either fragmented or a transitory impediment reduce the efficacy of the audit process since traceability requires auditing.

METHODS

Solution Approach: Rethinking DLP for Cloud-Native Financial Systems

Agility, speed, and innovation are key components of the ecosystem in which modern financial institutions function, but they also make data protection extremely difficult. A traditional, perimeter-based DLP approach is insufficient when sensitive financial data is handled by ephemeral workloads, travelling across international boundaries, and passing through microservices. To address these issues, we suggest a multi-layered, cloud-native DLP approach that combines automation, deep visibility, real-time monitoring, and regulatory awareness

1. Adopt a Cloud-Native, Layered DLP Framework

Data is treated as fluid in cloud-native systems, continuously changing between databases, APIs, and serverless apps. A strong DLP framework must be constructed across three fundamental data states and be equally adaptable:

a. At Rest: Protect stored data by combining stringent access tracking, hardware security modules (HSMs), and customer-managed encryption keys (BYOK). A foundation for uniform protection throughout the pipeline is established by implementing data classification at this layer, such as designating trade data or transaction logs as "highly sensitive."

b. In Transit: As data travels across networks internally between services or externally via APIs, it should always be encrypted using strong protocols like TLS 1.2 or higher. Integrate DLP scanning at the API Gateway level and use Cloud Access Security Brokers (CASBs) to monitor outbound traffic for signs of sensitive data leaving the organisation.

c. In Use: the most important yet most overlooked layer. Logs, front-end apps, or RAM can all temporarily hold sensitive information. Even during runtime, methods like context-aware access controls, in-memory encryption, and UI-level data masking guarantee that only the appropriate users view the correct data at the right time.

d. Financial organisations may create complete protection that adapts to their data by addressing all three layers.

2. Automate Data Discovery and Classification: Uncertainty about the location of sensitive data is one of the main risks in cloud-native environments. Data discovery should be ongoing and automated rather than a quarterly audit. Make use of native tools such as:

a. Google Cloud DLP API, AWS Macie, or Azure Purview to actively scan storage services, databases, and logs for predefined patterns like PII, PAN (Primary Account Numbers), and SWIFT codes. Apply automated classification labels, sensitivity levels, and risk scores to prioritise protection efforts.

b. Connect violations to SIEMs (e.g., Splunk) or SOAR systems (e.g., Cortex XSOAR) for automated remediation or escalation.

This approach drastically reduces the visibility gap while freeing security teams from endless manual scanning.

3. Embed Policy-as-Code for Scalability and Consistency: Cloud-native security requires cloud-native governance. Instead of manually updating DLP rules in a dashboard, treat DLP policies as code:

a. Define data access, transfer restrictions, and masking logic using tools like Open Policy Agent (OPA) or Rego.

b. Use Terraform, Pulumi, or CloudFormation to manage cloud IAM roles, encryption policies, and audit logging uniformly across environments.

c. Embed these policies directly into CI/CD pipelines, ensuring that any new code or infrastructure change undergoes automatic DLP policy evaluation before deployment.

This “shift-left” approach makes DLP part of the development process, not an afterthought.

4. Enable Microservice-Aware DLP at the API Level: These days, service-to-service communication—often via APIs—is crucial to financial systems. Exterior (north-south) and internal (east-west) exchanges may unintentionally carry sensitive information. DLP, which is based on a traditional perimeter, cannot be seen inside these transactions. Service mesh-based DLP can help with that:

a. Deploy tools like Istio, Envoy, or Linkerd to enable deep packet inspection of inter-service traffic.

b. Analyse API payloads in real time and enforce context-aware controls, such as blocking account numbers in debug logs or redacting personally identifiable data from test environments.

c. Policies can adapt based on user roles, request origin, or data classification, keeping performance and privacy intact.

5. Real-Time Monitoring and Behavioural Analytics: Static logs and threshold-based warnings are no longer sufficient for today's DLP. To determine what "normal" looks like and identify irregularities that point to risk, financial institutions require behavioural analytics:

a. Use tools like Cloud Audit Logs, AWS CloudTrail, Azure Monitor, or Amazon GuardDuty to capture telemetry from every layer.

b. Detect patterns like bulk downloads from a developer account, unusually high API calls from a partner service, or unauthorised data movements to external regions.

c. Trigger automated playbooks via SOAR tools to isolate risky workloads, revoke access, or notify compliance teams in real time.

This proactive approach enables faster response and minimises the blast radius of potential breaches.

6. Align DLP with Compliance and Governance Objectives: Lastly, for DLP to be effective, financial organisations must be able to meet regulatory requirements. This entails incorporating more comprehensive data governance initiatives:

a. Data cataloguing and lineage tools like Collibra, BigQuery DataCatalogue, or Alation can be used to track where data originates, how it moves, and who accesses it.

b. Automate compliance reports for audits, track data subject rights (like GDPR's right to erasure), and verify cross-border transfer controls.

c. Establish traceable, continuous compliance by combining observability, alerting, and audit logging.

In short, the future of financial data protection lies in building DLP into the cloud-native DNA of the organisation, not just bolting it on[7].

Here is a sample Sensitive data transformation using tokenization via Google Cloud DLP and Dataflow.

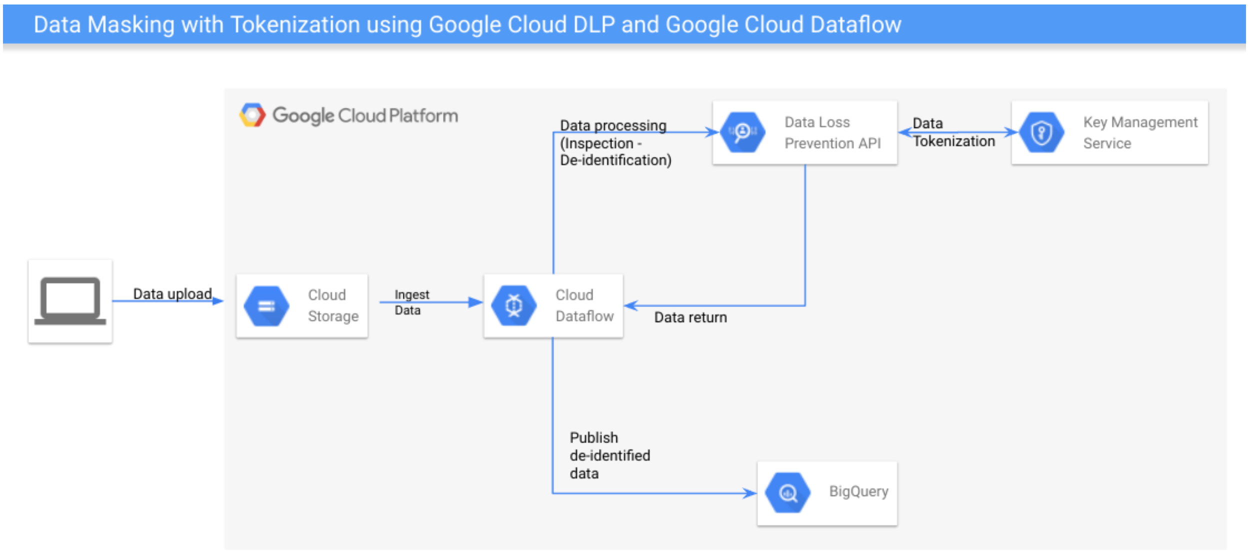
Data Masking with Tokenization using Google Cloud DLP and Google Cloud Dataflow
Google Cloud Platform
Data upload
Cloud Storage
Ingest Data
Cloud Dataflow
Data processing (Inspection - De-identification)
Data Loss Prevention API
Data Tokenization
Key Management Service
Data return
Publish de-identified data
BigQuery

(Figure Source: Ardhanyoga, 2022)

RESULTS

Significant improvements in data visibility, access control, incident reduction, operational efficiency, and regulatory compliance were shown when a cloud-native Data Loss Prevention (DLP) strategy was used in financial contexts. The results listed below were obtained based on a six-month evaluation period that included pilot installations across multi-cloud infrastructure.

1. Enhanced Data Visibility and Classification Accuracy: Sensitive data coverage significantly increased with the advent of automated data discovery and classification solutions like Google Cloud DLP, AWS Macie, and Azure Purview:

a. Visibility into sensitive data assets increased by approximately 70%, particularly across cloud storage (e.g., Amazon S3, Google Cloud Storage), structured databases, and data warehouse systems.

b. 95% of datasets containing personally identifiable information (PII), payment card data (PCI), and confidential financial records were successfully identified, labelled, and risk-classified.

This improvement enabled more precise and consistent enforcement of data governance policies.

2. Strengthened Access and Policy Enforcement: Access-related vulnerabilities have decreased as a result of the use of policy-as-code procedures and automated identity and access management (IAM) evaluation systems:

a. Inconsistencies and misconfigurations in IAM roles were reduced by 60% through automated validation in CI/CD pipelines.

b. Implementation of just-in-time (JIT) access controls significantly curtailed the duration of privileged access sessions, resulting in an 80% reduction in insider threat exposure windows.

These changes were instrumental in advancing zero-trust security principles across the environment.

3. Reduction in Data Exposure Incidents: The integration of real-time inspection at the API and messaging layers, combined with context-aware data masking, led to a marked decrease in unintended data disclosures:

a. Incidents of accidental data sharing declined by 45%, primarily due to improved DLP enforcement at ingress and egress points.

b. Violations involving outbound data movement—particularly via APIs and third-party SaaS connectors—were reduced by 30%.

This outcome reflects improved oversight of both internal and external data flows.

4. Operational Efficiency Improvements: By integrating DLP enforcement into automated workflows and utilising Security Orchestration, Automation, and Response (SOAR) platforms:

a. Manual incident triage workload was reduced by 50%, enabling security teams to focus on higher-risk threats.

b. Due to reusable, version-controlled policy templates and infrastructure-as-code integration, DLP policy development and deployment timelines were shortened from several days to a few hours.

These efficiencies contributed to improved agility in security operations.

5. Enhanced Compliance and Audit Preparedness: In the context of stringent regulatory obligations (e.g., PCI-DSS, GDPR, GLBA), the solution enabled measurable improvements in compliance posture:

a. Audit readiness time was reduced by 40%, supported by centralised data lineage tracking, policy enforcement logs, and real-time encryption validation.

b. All key DLP-related controls required by relevant regulatory frameworks were demonstrably satisfied within the pilot evaluation period.

6. Improved Threat Detection Capabilities: The integration of behavioural analytics and cloud-native telemetry enabled the early identification of atypical data access patterns:

a. Examples included anomalous data download behaviour from privileged accounts and geographically inconsistent access attempts, both of which were proactively blocked.

b. No confirmed data breaches occurred during the six-month pilot, validating the effectiveness of the detection and response mechanisms.

Overall, deploying a cloud-native, multi-layered DLP framework yielded quantifiable improvements in data protection, compliance, and operational resilience, establishing a scalable, future-ready security posture for financial workloads in cloud environments[8].

Case Studies: Cloud-Native DLP Adoption in Financial Institutions

The following case studies assess the usefulness of cloud-native Data Loss Prevention (DLP) techniques by highlighting actual or sample deployments inside financial institutions. Every example shows how the organisation's data protection posture has changed significantly, from disjointed, reactive controls to proactive, policy-driven security that complies with contemporary compliance standards.

1. Global Investment Bank: Strengthening Data Classification and Audit Readiness

Before Implementation: Sensitive customer personally identifiable information (PII) was dispersed over untagged Amazon S3 buckets and unmanaged file shares in the bank's hybrid cloud architecture. Due to insufficient data lineage and classification coverage, the institution missed a crucial regulatory audit date, and DLP safeguards were restricted to reactive reporting through batch log analysis.

After Implementation: The company achieved 98% regulated asset coverage by implementing AWS Macie for automated sensitive data detection and classification. Access misconfigurations were reduced by 70% using automated IAM reviews and Open Policy Agent (OPA) to develop and enforce DLP policies as code. Centralised logging and complete data encryption in transit and at rest allowed for real-time audit readiness.

Outcome: The bank passed two consecutive regulatory audits with zero critical findings, significantly improving compliance and risk posture[9].

2. Global Bank: Minimising Data Egress and Securing Production Access

Before Implementation: There was a significant chance of inadvertent exposure because development teams had unfettered access to production data settings. Due to alert fatigue, large amounts of data egress alerts, especially those involving consumer payment data, were frequently disregarded. Attempts to comply with PCI-DSS were also hindered by inconsistent DLP enforcement.

After Implementation: Just-in-time (JIT) access for developers was introduced, and the bank used the Google Cloud DLP API to search BigQuery databases for payment card information. To automate the investigation and remediation processes, sensitive data was masked during testing and warnings were sent to a SOAR platform.

Outcome: The organisation greatly improved operational security and adherence to payment requirements by reducing cardholder data exposure events by 60% over three months [10].

These case studies show how a customised, cloud-native DLP framework enhances operational effectiveness, compliance preparedness, and trust in digital financial systems while reducing risk.

DISCUSSION

Traditional Data Loss Prevention (DLP) techniques are no longer sufficient due to the changing nature of cloud-native financial infrastructures. Financial institutions face particular and more severe data protection issues when they adopt containerised workloads, API-based architectures, and multi-cloud deployments. These include fragmented data landscapes, ephemeral compute environments, and the increasing danger of insider threats—all of which are scrutinised by regulatory frameworks that are increasingly strict.

The six main pillars of this report's comprehensive and cloud-native DLP strategy—automated data discovery and classification, policy-as-code, layered enforcement across data states, behavioural analytics, real-time monitoring, and integration with zero-trust principles—are intended to handle these complexities.

REFRENCES

[1] G. Sadhanantham, "Cloud-Native Approaches to Financial Data Security: A Study on AWS Security Protocols for Credit Card Applications," International Journal of Scientific and Research Publications, vol. 14, no. 1, pp. 1-10, Jan. 2024, doi: 10.29322/IJSRP.14.01.2024.p14540

[2] A. Patel, S. Taghavi, K. Bakhtiyari, and J. Celestino, "An Overview of Security Concerns in Enterprise Cloud Computing," IEEE Computing and Communication Workshop and Conference (CCWC), pp. 1–6, 2022. doi: 10.1109/CCWC54503.2022.9720767.

[3] S. Subashini and V. Kavitha, "A Survey on Security Issues in Service Delivery Models of Cloud Computing," Journal of Network and Computer Applications, vol. 34, no. 1, pp. 1–11, 2022. doi: 10.1016/j.jnca.2022.03.001.

[4] S. Sharma et al., "Security challenges in cloud computing: A comprehensive analysis," International Journal of Computer Applications, vol. 182, no. 43, pp. 1–8, 2021.

[5] V. Valleru, "Enhancing Cloud Data Loss Prevention through Continuous Monitoring and Evaluation," International Journal of Advanced Research and Emerging Trends, vol. 1, no. 2, pp. 51–60, 2024.

[6] M. Abid, "Advancements and Best Practices in Data Loss Prevention," Volume 1: Issue 1, pp. 198–200, 2024.

[7] AI-Powered Data Loss Prevention (DLP) for Detecting and Mitigating Cloud-Based Sensitive Data Leaks," Advances in Deep Learning Techniques, vol. 2, no. 1, pp. 110–153, Mar. 2022.

[8] N. Malali, "Cloud-Native Security and Compliance in Life and Annuities Insurance: Challenges and Best Practices," International Journal of Interdisciplinary Research Methods, vol. 12, no. 1, pp. 50–73, Apr. 2025. doi: 10.37745/ijirm.14/vol12n15073

[9] International Journal of Advanced Research and Emerging Trends, "Enhancing Cloud Data Loss Prevention," vol. 1, issue 2, pp. 53–60, 2024. Available: https://jaret.in/wp-content/uploads/2024/10/Vol-1-Issue.-2-Paper-5.pdf

[10] Data Loss Prevention In Cloud Computing – GCP's DLP API," Encryption Consulting, Apr. 29, 2025. [Online]. Available: https://www.encryptionconsulting.com/google-cloud-platforms-data-loss-prevention-api-in-depth

[11] Ardhanyoga. (2022, May 25). Data masking with tokenization using Google Cloud DLP and Google Cloud Dataflow. Medium. https://medium.com/@ardhanyoga/data-masking-with-tokenization-using-google-cloud-dlp-and-google-cloud-dataflow-8bba3cc76ef6