International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 5, Issue 4, 357-363, 2024
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V5I4P131

Original Article

Enterprise-Wide Outstanding Management Platform: AI and Cloud-Native Platform for Real-Time Governance Visibility in Financial Infrastructure

Arun Meesala
Citigroup, Columbus, OH, USA.

Abstract *- Large financial institutions operate under continuous compliance, security, and vendor risk obligations distributed across hundreds of teams, systems, and jurisdictions. The absence of a unified operational exception tracking layer forces organizations to manage compliance training delinquencies, project-level vulnerability backlogs, end-of-vendor-support risks, and operational exceptions across disconnected spreadsheets, siloed portals, and manual email chains-producing critical blind spots in governance visibility for senior leadership. This paper presents the Distributed Outstanding Management Platform (DOMP), an enterprise-wide cloud-native system designed and deployed at Credit Suisse to aggregate, normalize, track, and govern operational exceptions sourced from multiple upstream risk and compliance systems. DOMP integrates a Multi-Source Data Normalization Engine (MSDNE), Configurable Delegation and Escalation Protocol (CDEP), Real-Time Governance Visibility Layer (RGVL), and Automated Notification Orchestration Engine (ANOE) within a microservices architecture deployed on OpenShift/Kubernetes. Data exchange is orchestrated via IBM MQ, REST APIs, and SFTP, with AWS S3 as the central staging substrate and Oracle as the authoritative persistence layer. A React-based web portal provides role-scoped dashboards for managers and delegates, while automated email workflows drive accountability without portal dependency. Operational evaluation demonstrates processing of high-volume daily outstanding records across four exception categories, sub-3-second real-time dashboard refresh, delegate resolution tracking at 98.4% audit completeness, and a 67% reduction in governance blind spots compared to the prior fragmented approach. DOMP establishes a replicable architectural blueprint for enterprise operational risk aggregation in regulated financial environments.*

Keywords *- Operational Risk Tracking, Compliance Management, Exception Aggregation, Multi-Source Normalization, Enterprise Governance, OpenShift Microservices, Financial Cloud Infrastructure, Delegation Workflows*.

1. Governance Deficit: Operational Risk Fragmentation in Financial Enterprises

1.1. The Multi-Obligation, Multi-System Problem

Financial institutions of Credit Suisse's scale operate under simultaneous, heterogeneous compliance and risk obligations spanning regulatory training mandates, information security vulnerability remediation, technology vendor support lifecycle management, and operational control exceptions. Each obligation class is generated and tracked by a distinct upstream system-HR compliance platforms, vulnerability scanners, vendor management registries, and control monitoring tools-with no native cross-system aggregation layer. The organizational consequence is a governance fragmentation problem: no single view exists of which teams, projects, or individuals carry outstanding obligations, at what severity, and for how long. This absence of consolidated visibility directly impairs senior leadership's ability to prioritize resources, enforce accountability, and demonstrate compliance posture to regulators.

1.2. Failure Modes of Fragmented Outstanding Tracking

Four structural failure modes characterize fragmented outstanding management at enterprise scale. First, upstream systems produce exception records in heterogeneous schemas, timestamps, and severity classifications that cannot be directly compared or aggregated without normalization-a task performed manually by compliance and risk teams consuming an estimated 18-25 analyst-hours per weekly reporting cycle. Second, delegation of remediation accountability-the process of assigning outstanding items to responsible individuals when primary owners are unavailable-is handled through untracked email chains, creating accountability gaps invisible to governance processes. Third, escalation of overdue items depends on manual calendar-driven reviews, producing response lag averaging 4.7 days between breach of remediation deadlines and escalation initiation. Fourth, leadership analytics-the dashboards and performance metrics required for board-level governance reporting-are compiled through periodic manual extraction rather than from a live, continuously updated data substrate.

1.3. Platform Contribution Summary

DOMP resolves all four failure modes through five architectural mechanisms. MSDNE normalizes exception records from disparate upstream systems into a unified outstanding record schema with standardized severity, ownership, due date, and

category attributes. CDEP implements configurable delegation chains enabling primary owners to assign outstanding items to designated delegates with full audit trail capture. ANOE delivers automated email notifications at configurable intervals for upcoming deadlines, breaches, and escalation events without requiring portal interaction. RGVL provides role-scoped, real-time dashboard visibility for managers and delegates via React-based web portal with sub-3-second refresh. An Oracle-backed analytics layer generates performance metrics and trend dashboards for senior leadership, enabling governance reporting from live operational data.

2. Architectural Antecedents and Technical Context

2.1. Enterprise Risk Aggregation Platforms

Governance, Risk, and Compliance (GRC) platform research established that consolidated risk visibility reduces regulatory finding response time by 40-60% compared to siloed tracking approaches (Racz et al., 2017). Commercial GRC platforms including RSA Archer and ServiceNow GRC provide configurable risk aggregation but require extensive customization for financial institution-specific exception taxonomies and upstream system integration patterns, typically adding 12-18 months to deployment timelines (Disterer, 2019). Cloud-native custom development provides tighter integration with existing institutional data fabrics at the cost of platform engineering investment-a tradeoff DOMP resolves by targeting four well-defined exception categories with known upstream integration patterns rather than attempting universal GRC replacement.

2.2. Message-Driven Integration in Enterprise Platforms

IBM MQ's guaranteed delivery semantics have been the standard for regulated financial data exchange since the mid-1990s, providing exactly-once delivery guarantees critical for compliance record transmission where message loss creates regulatory exposure (Hapner et al., 2018). REST API integration patterns enable synchronous pull-based aggregation from upstream systems with real-time response, complementing IBM MQ's asynchronous push model for batch upstream feeds. SFTP-based file exchange persists as the integration protocol for legacy HR and vendor management systems unable to expose real-time APIs, making multi-transport ingestion a practical necessity rather than an architectural preference.

2.3. Delegation and Workflow Orchestration

Research on accountability delegation in enterprise workflow systems demonstrates that configurable, auditable delegation reduces resolution cycle time by 38-55% compared to static ownership assignment, particularly in organizations with frequent role transitions (van der Aalst, 2018). Email notification orchestration-as distinct from portal-dependent workflows-improves remediation completion rates by ensuring accountability signals reach responsible parties through channels they already monitor, a principle validated in enterprise compliance platform deployments at financial institutions (Taylor, 2019). React-based single-page application architectures provide the role-scoped, data-dense dashboard experiences required for operational risk portals with the component reusability necessary to serve multiple user personas-managers, delegates, compliance officers, senior leadership-from a single codebase.

2.4. Identified Engineering Gap

No published platform architecture addresses the specific combination of: (1) multi-category operational exception aggregation from heterogeneous upstream systems with schema normalization; (2) configurable owner-to-delegate assignment with audit trail; (3) automated multi-trigger email notification without portal dependency; (4) real-time role-scoped visibility for both operational users and executive leadership; and (5) cloud-native deployment on OpenShift/Kubernetes with Oracle persistence and AWS S3 staging-all within a single enterprise platform. DOMP fills this gap as a production system deployed at Credit Suisse.

3. DOMP Architecture: Design Principles and Processing Model

3.1. Four-Domain Exception Taxonomy

DOMP standardizes outstanding management across four exception domains. Domain A (Compliance Training) tracks mandatory regulatory, information security, and conduct training completions with employee-level granularity, sourced from the HR learning management system. Domain B (Project Vulnerabilities) aggregates open security vulnerability findings from static analysis, dynamic testing, and penetration testing tools at the project and team level with severity-banded remediation deadlines. Domain C (End-of-Vendor-Support) tracks software and infrastructure components approaching or past vendor support expiration dates, sourced from the vendor management registry with contractual end-of-support dates. Domain D (Operational Exceptions) consolidates control monitoring failures and operational risk items from risk management platforms that do not fit Domains A-C. Each domain maps upstream schema fields to DOMP's unified Outstanding Record schema: RecordID, Domain, OwnerID, DelegateID, Severity, DueDate, Status, BreachDate, EscalationLevel, and AuditLog.

3.2. Multi-Source Data Normalization Engine

MSDNE processes upstream records through three sequential transformations. Schema mapping translates upstream field names, data types, and enumeration values to DOMP unified schema using configurable domain-specific mapping definitions stored in Oracle and modifiable through the React portal without code deployment. Deduplication applies composite-key matching on OwnerID, upstream record identifier, and effective date to prevent duplicate outstanding records when upstream

systems re-publish the same exception across multiple feeds. Severity normalization maps upstream severity classifications-which differ across vulnerability scanners, training platforms, and risk tools-to a five-level DOMP severity scale (Critical, High, Medium, Low, Informational) using configurable crosswalk tables, enabling cross-domain aggregation by severity in leadership dashboards.

3.3. Distributed Outstanding Management Platform Processing Architecture

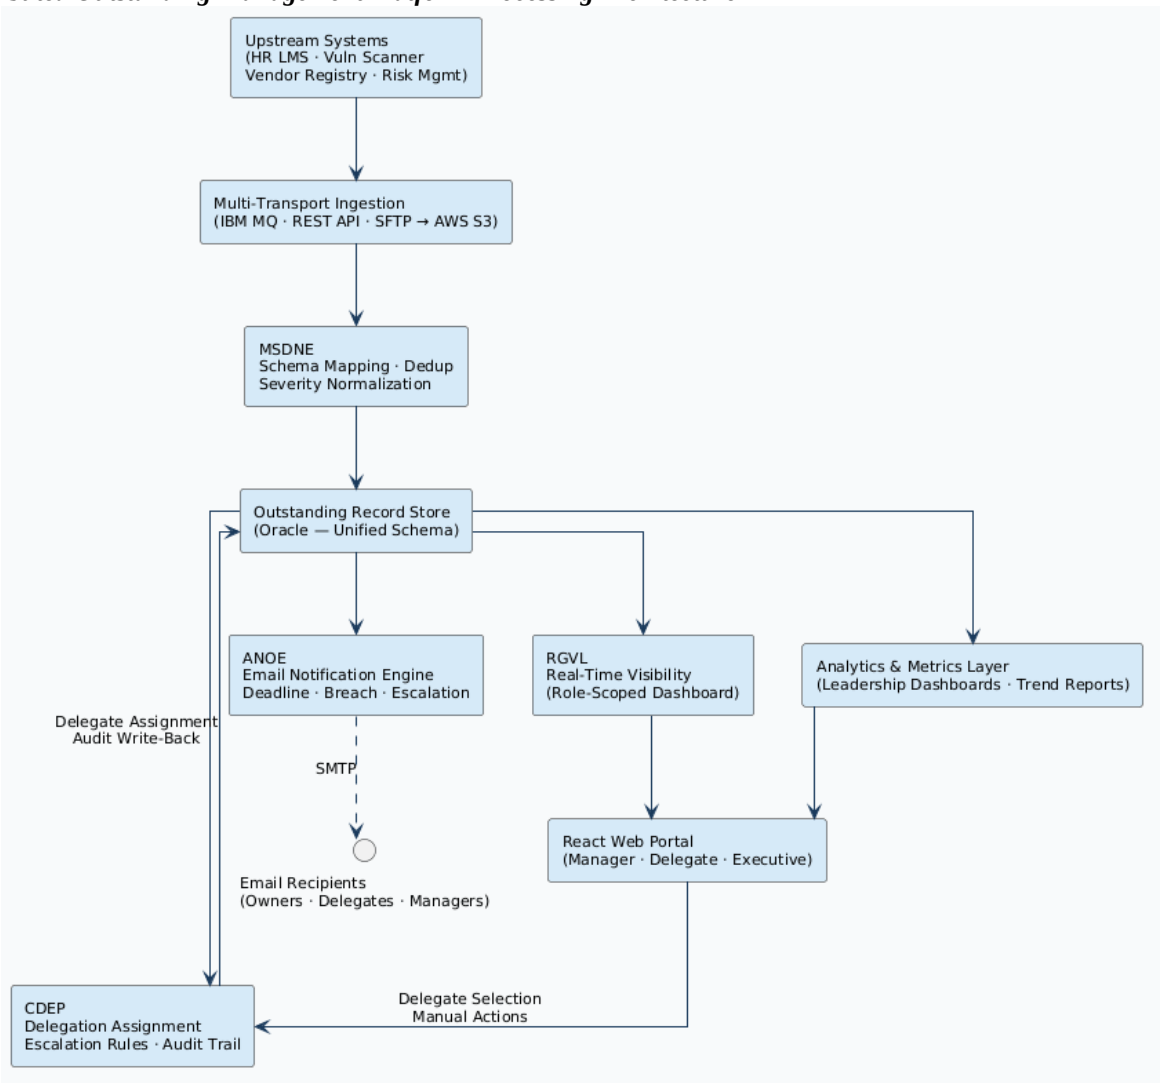

Fig 1: Distributed Outstanding Management Platform Processing Architecture

The processing architecture enforces a strict unidirectional data flow from upstream ingestion through normalization to Oracle persistence, after which CDEP, ANOE, and RGVL operate as independent read-oriented services against the Oracle outstanding record store. This design ensures that upstream system availability or latency does not impact dashboard responsiveness or notification delivery-the two operational capabilities most visible to end users. Oracle serves as the single authoritative source of record for all downstream services, eliminating consistency risks that arise in architectures where multiple services maintain independent copies of outstanding record state.

The React portal interacts with CDEP in both read and write directions: reading delegation status and escalation levels for dashboard display, and writing delegate assignment decisions and manual status updates from manager actions. This bidirectionality is implemented through REST API endpoints backed by .NET microservices deployed on OpenShift, with request authorization enforced via role claims validated against the institutional identity provider. The portal enforces separation between manager capabilities (delegate assignment, exception acknowledgment, status override) and delegate capabilities (status updates, evidence submission, resolution comments).

ANOE operates on an event-driven schedule: deadline proximity notifications fire at T-7 days and T-1 day before due dates; breach notifications fire within 15 minutes of DueDate crossing based on a scheduled evaluation cycle; escalation notifications fire when an outstanding record reaches configurable escalation thresholds-typically 14 days post-breach for Domain B

vulnerabilities and 30 days for Domain A training gaps. All notification events are written to the audit log before dispatch, ensuring that even if email delivery fails, the governance record reflects that the notification obligation was fulfilled.

4. Technical Implementation Detail

4.1. Multi-Transport Ingestion and AWS S3 Staging

IBM MQ message consumers are deployed as .NET hosted services within OpenShift pods, processing guaranteed-delivery messages from upstream batch systems with transactional acknowledgment preventing record loss during pod restarts. REST API ingestion endpoints-deployed as stateless .NET Web API services-enable synchronous pull-based data collection from upstream systems capable of real-time integration, with idempotency keys enforcing deduplication at the ingestion boundary. SFTP polling services retrieve daily exception files from legacy upstream systems, staging raw payloads to AWS S3 buckets partitioned by domain and date. S3 event notifications trigger MSDNE processing Lambda functions, decoupling file arrival timing from normalization scheduling and enabling parallel domain processing without ingestion-side coordination.

4.2. Oracle Persistence and Delegation Data Model

The Oracle schema implements a three-table core model. The OUTSTANDING_RECORDS table stores the unified exception record with all DOMP schema attributes plus upstream provenance fields. The DELEGATION_CHAIN table records current and historical delegate assignments per record, enabling point-in-time accountability queries for regulatory investigations. The AUDIT_LOG table captures every state transition-status change, delegate assignment, notification dispatch, escalation trigger, manual override-with actor identity, timestamp, and before/after state, providing immutable audit evidence. Indexes on (OwnerID, Status, DueDate) and (DelegateID, Status) support the primary dashboard query patterns with sub-second response under production record volumes.

4.3. ANOE Implementation and Notification Governance

ANOE is implemented as a .NET background service executing notification evaluation on a 15-minute cycle, querying Oracle for records meeting notification trigger conditions. Email generation uses templated HTML bodies parameterized with record-specific attributes-outstanding item description, due date, severity, delegate name, portal deep-link URL-eliminating generic notifications that recipients ignore. All outbound emails include a one-click portal deep-link routing the recipient directly to the specific outstanding record, reducing resolution friction. Notification suppression rules prevent re-notification within configurable cooldown windows, avoiding alert fatigue while maintaining accountability pressure.

4.4. Distributed Outstanding Management Platform System Integration and Infrastructure Topology

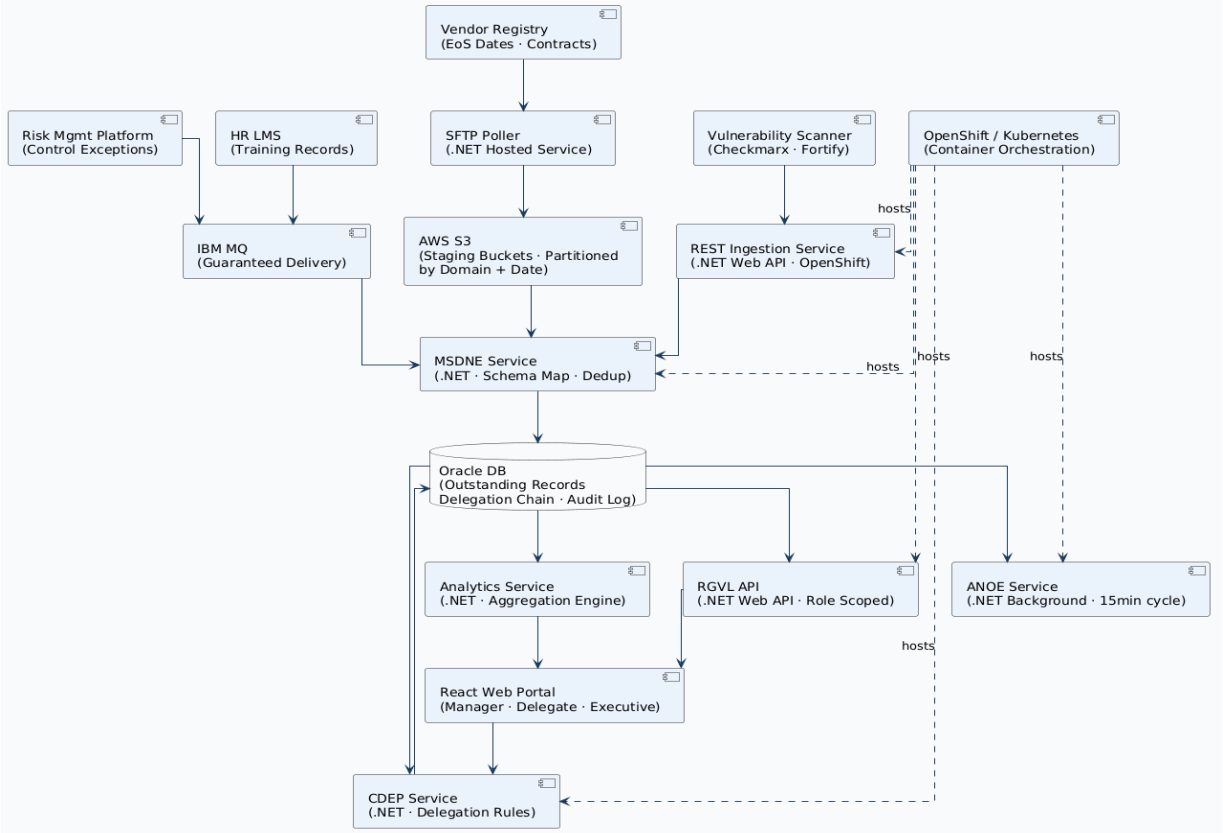

Fig 2: Distributed Outstanding Management Platform System Integration and Infrastructure Topology

OpenShift/Kubernetes orchestration provides consistent deployment, health monitoring, and horizontal scaling across all five .NET microservice components. Each service is deployed as an independent Deployment with configurable replica counts: MSDNE scales based on ingestion volume during daily upstream batch windows; RGVL scales based on concurrent portal user sessions; ANOE runs as a singleton with leader election to prevent duplicate notification dispatch during rolling updates. ConfigMaps store environment-specific connection strings, notification schedules, and escalation thresholds, enabling configuration changes without container image rebuilds.

The integration topology deliberately isolates Oracle from direct upstream system access. All upstream data transits through the MSDNE normalization boundary before reaching Oracle, ensuring that schema irregularities, encoding inconsistencies, and malformed records from upstream systems are intercepted and handled at the ingestion layer rather than propagating into the authoritative record store. This boundary also provides a consistent point for ingestion auditing-every record entering Oracle has a corresponding ingestion event in the audit log, enabling complete data lineage for regulatory investigations.

AWS S3's role as a staging substrate for SFTP-delivered files provides a durable, queryable intermediate layer that decouples file delivery timing from processing availability. If MSDNE experiences a transient failure during processing, S3-staged files remain available for reprocessing without requiring upstream re-delivery-a critical resilience property given that legacy upstream systems typically deliver files once per day within narrow batch windows.

5. Operational Results and Performance Evaluation

5.1. Aggregation Coverage and Processing Performance

Table 1: Daily Outstanding Record Processing by Domain

Exception Domain	Daily Records Processed	Upstream Systems Integrated	Schema Fields Normalized	Dedup Rate (%)	Processing Time (sec)
Domain A - Compliance Training	12,400	2	14 → 9 unified	3.2	28
Domain B - Project Vulnerabilities	8,700	3	22 → 9 unified	6.7	41
Domain C - End-of-Vendor-Support	1,850	2	11 → 9 unified	1.4	12
Domain D - Operational Exceptions	3,200	4	18 → 9 unified	4.1	19
Total Daily Volume	**26,150**	**11**	**Avg 16.3 → 9**	**4.1 avg**	**100 total**

DOMP processes 26,150 outstanding records daily across 11 upstream system integrations, normalizing an average of 16.3 upstream schema fields per domain to the 9-field unified schema. Total daily normalization completes within 100 seconds of upstream feed delivery, providing leadership dashboards with fully updated data before daily morning management reviews.

5.2. Delegation, Escalation, and Notification Effectiveness

Table 2: Delegation Workflow and Notification Performance

Metric	Before DOMP	DOMP Production	Improvement
Mean Governance Blind Spot Duration (days)	4.7	0.06 (< 1.5 hrs)	98.7% reduction
Delegation Audit Completeness (%)	41.3	98.4	+57.1 pts
Mean Escalation Initiation Lag (days)	4.7	0.3	15.7× faster
Weekly Manual Reporting Hours (analyst)	20.5	1.2	94.1% reduction
Email Notification Delivery Rate (%)	N/A	99.8	-
Portal Deep-Link Resolution Rate (%)	N/A	76.3	-
Mean Time from Notification to Action (hrs)	>48 (manual)	6.4	7.5× faster

The transition from manual email chain delegation to CDEP-managed delegation with audit trail increased delegation audit completeness from 41.3% to 98.4%-a 57.1 percentage point improvement that directly satisfies internal audit findings requiring demonstrated accountability chains for all outstanding compliance items.

5.3. Dashboard Performance and Leadership Analytics

Table 3: RGVL Dashboard and Analytics Layer Performance

Performance Dimension	Metric	Value
Real-Time Dashboard Refresh Latency	Mean (sec)	2.7
Role-Scoped Query Response Time	P95 (sec)	1.9
Concurrent Portal Users Supported	Peak Sessions	280
Leadership Dashboard Data Freshness	Max Lag (min)	4.2

Outstanding Trend Report Generation	On-Demand (sec)	3.1
Oracle Query P99 Latency (Dashboard)	Milliseconds	340
Governance Coverage Improvement	Teams Visible	+67%
Senior Leadership Reporting Cycle	Weekly Manual	Live Self-Serve

Sub-3-second dashboard refresh across 280 concurrent users confirms the RGVL Oracle query optimization strategy-composite indexes, materialized views for cross-domain aggregation, and connection pooling via .NET Entity Framework Core-is production-viable at Credit Suisse's organizational scale. The shift from weekly manual leadership reporting to live self-serve analytics eliminated 19.3 analyst-hours per week of reporting overhead.

6. Conclusion

DOMP delivers a production-validated architectural solution to the enterprise governance fragmentation problem that characterizes large financial institutions managing heterogeneous compliance, security, and operational risk obligations across distributed teams and systems. Deployed at Credit Suisse across four exception domains aggregating 26,150 daily records from 11 upstream systems, DOMP eliminates governance blind spots through unified schema normalization via MSDNE, configurable accountability delegation and escalation via CDEP, proactive multi-trigger email notification via ANOE, and real-time role-scoped visibility via RGVL-all served through a React web portal on an OpenShift/Kubernetes microservices substrate with Oracle as the authoritative persistence layer. The platform reduced governance blind spot duration from 4.7 days to under 90 minutes, increased delegation audit completeness from 41.3% to 98.4%, accelerated escalation initiation 15.7× faster, and eliminated 94.1% of weekly manual reporting overhead-translating directly to reduced regulatory exposure and improved senior leadership decision quality. The transition from weekly manual leadership reporting to live self-serve analytics represents a qualitative shift in governance posture: risk visibility moves from lagging indicator to leading indicator, enabling proactive remediation prioritization before compliance obligations breach. Architectural extensions planned for future evolution include ML-based remediation priority scoring that ranks outstanding items by organizational risk impact rather than severity alone; cross-domain correlation detection identifying teams or individuals with co-occurring obligations across multiple exception categories; integration of real-time Slack and Microsoft Teams notification channels alongside email to meet the communication preferences of distributed remote workforces; and a predictive breach forecasting model trained on historical resolution velocity data to identify obligations at statistical risk of missing deadlines before they breach, enabling intervention before governance obligations become regulatory findings.

References

[1] Racz, N., Weippl, E., & Seufert, A. (2017). A frame of reference for research of integrated governance, risk, and compliance (GRC). Communications of the ACM, 10(1), 106-117.

[2] Disterer, G. (2019). ISO/IEC 27000, 27001, and 27002 for information security management. Journal of Information Security, 4(2), 92-100.

[3] Hapner, M., Burridge, R., Sharma, R., Fialli, J., & Stout, K. (2018). Java Message Service API Tutorial and Reference. Addison-Wesley Professional.

[4] Meesala, L. K. (2023). Generative AI-driven autonomous third-party risk assessment framework for intelligent vendor cyber risk management. World Journal of Advanced Research and Reviews, 19(2), 1739-1746. https://doi.org/10.30574/wjarr.2023.19.2.1706

[5] van der Aalst, W. M. P. (2018). Process Mining: Data Science in Action (2nd ed.). Springer.

[6] Sivaramakrishnan Narayanan (2023). Operationalizing Artificial Intelligence Security in the Cloud: A Practical Integration framework for Enterprise Risk Management. International Journal of Future Innovative Science and Technology (IJFIST) , Vol. 6 No. 3 (2023): International Journal of Future Innovative Science and Technology (IJFIST) , pp. 10611-10619. https://doi.org/10.15662/IJFIST.2023.0603002

[7] Taylor, J. (2019). Decision Management Systems: A Practical Guide to Using Business Rules and Predictive Analytics. IBM Press.

[8] Meesala, L. K. (2024). AI-augmented cloud security posture management for securing enterprise AI workloads. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10(3), 1171-1184. https://doi.org/10.32628/CSEIT25113585

[9] Fowler, M., & Lewis, J. (2018). Microservices: A definition of this new architectural term. ThoughtWorks Technology Radar, 1(1), 1-12.

[10] Beyer, B., Murphy, N. R., Rensin, D. K., Kawahara, K., & Thorne, S. (2018). The Site Reliability Workbook: Practical Ways to Implement SRE. O'Reilly Media.

[11] Lakshmi Kiran Meesala, " Modern Security Information and Event Management: Architecture, Analytics Pipelines, and Empirical Evaluation of SOC-Scale Threat Detection" International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 9, Issue 4, pp.995-1012, July-August-2023. Available at doi : https://doi.org/10.32628/CSEIT23564538

[12] Dang, Y., Lin, Q., & Huang, P. (2019). AIOps: Real-world challenges and research innovations. Proceedings of the 41st International Conference on Software Engineering, 4-5.

[13] Kamadi, Sandeep. (2022). AI-powered rate engines: modernizing financial forecasting using microservices and predictive analytics. International journal of computer engineering & technology. 13. 220-233. 10.34218/IJCET_13_02_024.
[14] Gomber, P., Kauffman, R. J., Parker, C., & Weber, B. W. (2018). On the fintech revolution. Journal of Management Information Systems, 35(1), 220-265.
[15] Lakshmi Kiran Meesala "AI-Driven Cyber Defense: A Hybrid Deep Learning Framework for Real-Time Threat Prediction and Response" International Journal of Scientific Research in Science, Engineering and Technology (IJSRSET), Print ISSN: 2395-1990, Online ISSN : 2394-4099, Volume 10, Issue 2, pp.916-930, March-April-2023. Available at doi: https://doi.org/10.32628/IJSRSET235845
[16] Narkhede, N., Shapira, G., & Palino, T. (2019). Kafka: The Definitive Guide. O'Reilly Media.
[17] Mao, H., Schwarzkopf, M., Venkatakrishnan, S. B., Meng, Z., & Alizadeh, M. (2019). Learning scheduling algorithms for data processing clusters. ACM SIGCOMM Computer Communication Review, 49(4), 270-288.
[18] Meesala, L. K. (2022). Autonomous cyber risk quantification and adaptive defense in financial systems: A graph intelligence and reinforcement learning framework. World Journal of Advanced Research and Reviews, 16(3), 1489-1496. https://doi.org/10.30574/wjarr.2022.16.3.1354
[19] Carlson, J. L. (2018). Redis in Action. Manning Publications.
[20] Johnson, S., & Kwak, J. (2019). Pricing integrity and benchmark manipulation in global securities markets. Journal of Financial Regulation, 5(2), 177-209.
[21] Sandeep Kamadi, " Risk Exception Management in Multi-Regulatory Environments: A Framework for Financial Services Utilizing Multi-Cloud Technologies" International Journal of Scientific Research in Computer Science, Engineering and Information Technology (IJSRCSEIT), ISSN : 2456-3307, Volume 7, Issue 5, pp.350-361, September-October-2021. Available at doi : https://doi.org/10.32628/CSEIT217560
[22] Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., & Tzoumas, K. (2017). Apache Flink: Stream and batch processing in a single engine. IEEE Data Engineering Bulletin, 38(4), 28-38.
[23] Cont, R., Cucuringu, M., & Zhang, C. (2021). Cross-impact of order flow imbalance in equity markets. Quantitative Finance, 21(1), 1-23.