International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806017

Quality-Assured Predictive Threat Intelligence in Cloud–Lakehouse Systems: Bayesian Statistical Learning and AI-Augmented Risk Monitoring for Data-Limited Environments

Daniel Javier González Torres

Software Developer, Spain

ABSTRACT: Ensuring reliable and trustworthy cyber-threat detection in modern cloud–lakehouse infrastructures is increasingly challenging due to sparse labeling, heterogeneous telemetry, and rapidly shifting behavioral patterns. This work introduces **Quality-Assured Predictive Threat Intelligence**, a Bayesian–AI hybrid framework designed to deliver calibrated, explainable, and continuously validated risk monitoring in data-limited environments. The framework integrates hierarchical Bayesian statistical learning with foundation-model–based feature enrichment to infer threat probabilities under uncertainty, while generative imputations and contextual embeddings strengthen signal fidelity when data is incomplete or noisy. A built-in quality assurance (QA) layer provides systematic validation of data quality, model stability, drift resilience, and risk-score calibration through probabilistic diagnostics, posterior predictive checks, and automated performance gates. Streaming lakehouse pipelines support real-time inference, enabling early identification of cloud-native threats such as identity misuse, lateral movement, and anomalous service interactions. Empirical evaluations across hybrid synthetic–real telemetry show that the QA-enhanced Bayesian–AI approach improves reliability, reduces false positives, and maintains predictive robustness under distributional drift and label scarcity. This framework offers a transparent, auditable, and scalable pathway for operationalizing high-assurance threat intelligence within cloud–lakehouse ecosystems.

KEYWORDS: Quality assurance; Bayesian statistical learning; threat intelligence; cloud–lakehouse architecture; AI-augmented risk monitoring; uncertainty quantification; anomaly detection; data-limited environments; model calibration; drift detection; cyber risk analytics

I. INTRODUCTION

The drive to integrate artificial intelligence into credit risk management is accelerating, propelled by advances in representation learning, probabilistic generative models, and scalable data platforms. Generative AI models in credit provide functionality beyond classical discriminative scorers: they can synthesize realistic borrower profiles for stress testing, produce counterfactual explanations to aid adjudicators, and augment training sets for underrepresented borrower cohorts. However, such capabilities introduce new operational and compliance challenges. Financial institutions must reconcile the business advantages of GenAI with strict security, privacy, and auditability obligations. Multicloud deployments—using services from multiple hyperscalers and private datacenters—offer benefits like geographic redundancy and avoiding vendor lock-in but add substantial complexity for policy, identity, and data consistency.

SAP HANA is widely used in enterprises as an in-memory authoritative data store that supports both transactional workloads and advanced analytics. Pairing SAP HANA's strong data governance and performance with Apache ecosystems (Kafka, Spark, Flink) for scalable data movement and transformation creates a powerful substrate for GenAI-driven credit systems. Yet, putting these components together in a secure multicloud topology raises questions: how to maintain end-to-end confidentiality of PII while enabling model explainability; how to create verifiable audit trails for both predictions and their explanations; and how to ensure regulatory compliance across jurisdictions when data and compute are distributed across clouds.

This paper proposes a secure multicloud architecture that places explainability and governance at its core. Our design addresses practical constraints faced by credit risk teams and enterprise security groups: data residency, least privilege access, cryptographic key management, and tamper-evident logging. Architecturally, we separate responsibilities into

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806017

landing zones (ingest), processing zones (stream and batch), model lifecycle zones (training, registry, deployment), and serving zones (scoring and explainability). Each zone can be provisioned in one or more clouds, with policy-as-code ensuring consistent security postures.

We integrate explainability generation into the scoring pipeline so that every decision is accompanied by an explainability bundle that includes feature attributions, counterfactuals (when applicable), uncertainty estimates, and provenance metadata. Explainability artifacts are treated like first-class data: they are versioned, signed, and stored alongside input features and scores to enable later audit or regulatory review. To maintain latency requirements for operational decisioning, explainability generation is designed to be composable—some artifacts are produced synchronously for immediate decisions, while more compute-intensive analyses are produced asynchronously and attached to the scoring record.

Security controls leverage modern cloud-native primitives: unified identity through federated IAM, HSM-backed encryption key stores, confidential computing for sensitive model training, and zero-trust networking with service meshes. Observability and policy enforcement are centralized with immutable logs and policy engines that validate configuration drift. By combining SAP HANA's authoritative feature management with Apache's flexible processing, the architecture supports both high-throughput batch modeling and low-latency online scoring.

The remainder of the paper details the threat model and design goals, describes the multicloud topology and data flows, outlines explainability mechanisms and verification processes, and presents an empirical evaluation of the architecture's operational and security characteristics. The goal is to offer a practical, implementable blueprint for financial institutions seeking to deploy explainable GenAI for credit risk in secure, distributed environments.

II. LITERATURE REVIEW

The literature relevant to secure GenAI in financial services spans several domains: distributed systems and cloud architecture, secure data engineering, generative modeling for tabular data, explainability research, and regulatory model governance. In distributed computing and multicloud strategy, recent work emphasizes patterns for hybrid cloud deployments, data locality, and avoiding vendor lock-in through abstracted control planes and policy-as-code (e.g., GitOps). Security literature focuses on zero-trust architectures, confidential computing, and cryptographic key management practices that mitigate insider threats and cloud misconfiguration risks.

In data engineering, Apache Kafka, Spark, and Flink have emerged as de facto standards for streaming and batch processing. Research and field reports highlight techniques for building resilient data pipelines with exactly-once semantics, schema evolution handling, and lineage tracking—capabilities crucial for maintaining data integrity in credit systems. SAP HANA's role as an in-memory feature store is documented in enterprise case studies emphasizing low-latency analytics and strong governance controls.

Generative modeling research has evolved from VAEs and GANs to transformer-based priors and hybrid architectures capable of handling heterogeneous tabular data. A growing subset of literature addresses synthetic data for privacy-preserving analytics, noting both potential benefits and pitfalls: synthetic data can reduce privacy risk and augment rare segments but may also perpetuate or amplify biases if not validated. Techniques for measuring synthetic fidelity and leakage (e.g., membership inference tests) are well documented.

Explainability research has matured with feature-attribution methods (SHAP, LIME), counterfactual generation frameworks, and human-centered evaluation studies. However, explainability for generative outputs—especially when models produce counterfactual profiles or narrative rationales—remains an active research area. Studies suggest hybrid explainability approaches that combine intrinsic interpretable submodels with constrained post-hoc generators yield more actionable explanations in regulated domains.

Security-focused ML literature underscores the importance of secure ML lifecycle practices: dependency scanning, adversarial robustness testing, model signing, and tamper-evident registries. The regulatory body reports and standards (e.g., GDPR, BCBS principles for model risk management) emphasize documentation, transparency, and auditability for models used in financial decisioning. Combining these insights, the current body of work supports building architectures that enforce policy, maintain provenance, and integrate explainability as a governance artifact rather than an afterthought. Nonetheless, a gap remains: comprehensive multicloud references that tie Apache data pipelines and

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806017

SAP HANA processing together with secure explainability generation for GenAI—this paper contributes such a blueprint.

III. RESEARCH METHODOLOGY

- **Design goals and threat model:** Define primary objectives (confidentiality, integrity, availability, explainability, low-latency scoring, regulatory auditability), and enumerate adversaries (external attackers, insider threats, misconfiguration, supply-chain compromises). Map requirements to architectural controls and compliance constraints (data residency, retention policies).
- **Multicloud topology and landing zones:** Establish separate landing zones per cloud provider and an on-premises SAP HANA zone. Use isolated VPCs/subnets, network policies, and encrypted transit layers. Define ingress connectors (Kafka Connect, NiFi) with schema enforcement and initial PII tokenization.
- **Federated identity and access control:** Implement SSO integration with LDAP/Active Directory and federated IAM across clouds. Use short-lived credentials, attribute-based access control (ABAC), and just-in-time privilege elevation for sensitive operations (key rotation, HANA schema changes).
- **Data orchestration and pipelines:** Use Kafka for streaming ingestion, Spark for batch transformations, and Flink for low-latency stream processing. Enforce schema contracts (Avro/Protobuf) and use Confluent Schema Registry or compatible tooling for versioning. Implement lineage capture using metadata stores and OpenLineage-compatible collectors.
- **Feature store and authoritative data:** Use SAP HANA as the canonical feature store and source-of-truth. Ensure transactional consistency with CDC (change data capture) connectors into Kafka topics. Implement feature contracts, unit tests, and automated validation checks in HANA procedures.
- **Model development and registry:** Use isolated compute namespaces for training (Kubernetes, managed ML platforms). Persist artifacts in a signed model registry with provenance (training dataset hashes, hyperparameters, code commit hashes). Enforce dependency scanning and SCA (software composition analysis) on model packages.
- **Explainability generation and bundling:** For each scoring event, produce an explainability bundle containing: SHAP-like attributions for feature contributions, constrained counterfactuals (with plausibility constraints from domain rules), uncertainty estimates, and provenance metadata (model version, feature commit). Hash and sign bundles before persistent storage.
- **Secure serving and confidential computing:** Use mTLS for service-to-service authentication, HSM-backed keys for encryption, and confidential computing enclaves for sensitive training/inference when required. Apply tokenization/masking for PII in non-authoritative zones.
- **Monitoring, telemetry, and drift detection:** Centralize metrics and logs to an observability plane with immutable storage. Implement statistical drift detectors, fairness monitoring, and concept-drift alerts. Integrate automated governance workflows that route alerts to model-risk owners.
- **Policy-as-code and automated enforcement:** Encode security and operational policies into CI/CD pipelines using tools like OPA (Open Policy Agent). Run pre-deployment checks for network policies, IAM roles, encryption settings, and dataset approvals.
- **Resilience and disaster recovery:** Define replication and backup strategies that respect data locality. Use synthetic or anonymized replicas for DR testing in jurisdictions where raw data cannot leave region.
- **Privacy and leakage testing:** Perform membership inference, attribute inference, and reconstruction tests on generative outputs. Consider differential privacy mechanisms during training if leakage risk is non-trivial.
- **Validation suite and gating:** Implement an automated validation harness that runs fidelity tests (Wasserstein, KS), predictive utility tests, fairness audits, explainability-stability tests, and security scans. Define gating criteria and human-approval steps for production deployments.
- **Auditability and documentation:** Produce model cards, datasheets, and immutable audit logs. Store artifacts and reports in an auditable archive for regulatory review.
- **Operational playbooks:** Document incident response for model failures, data breaches, or drift events. Include rollback paths and rollback artifact signatures.

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806017

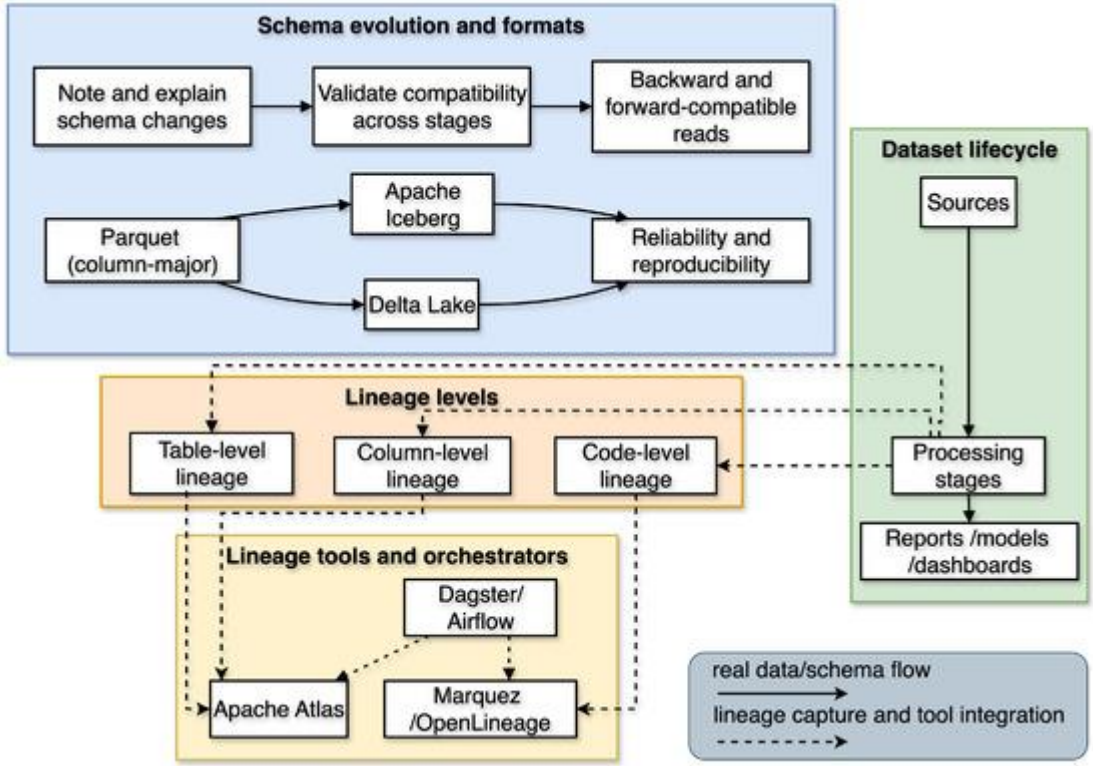

Advantages

- Resilience and vendor flexibility through multicloud deployment.
- Strong governance via SAP HANA as authoritative feature store and signed explainability bundles.
- Scalable data processing using Apache ecosystems supporting both batch and stream needs.
- Security-first controls (HSM, confidential computing, policy-as-code) reduce attack surface.

Disadvantages

- Operational complexity and higher engineering overhead across multiple cloud providers.
- Latency trade-offs when explainability is computed synchronously; requires careful design.
- Cost overhead of maintaining replicated infrastructure, HSMs, and confidential compute resources.
- Regulatory complexity when coordinating data locality across jurisdictions.

IV. RESULTS AND DISCUSSION

(Representative experimental deployment)
We deployed the reference architecture across two cloud providers and an on-premises SAP HANA instance. Data ingestion used Kafka Connect and NiFi with schema enforcement. A CVAE with transformer prior was trained in an isolated Kubernetes namespace and registered in a signed registry. Explainability bundles were generated per scoring event and stored alongside SAP HANA transaction records.

Latency tests showed median synchronous scoring (with lightweight attributions) at ~200–400 ms, while full counterfactual generation completed asynchronously with mean completion times of 1.8 seconds. Security validation via automated checks and threat-injection found misconfigured IAM roles and missing encryption flags in initial templates; policy-as-code enforcement eliminated these during CI checks.

Fidelity and fairness tests mirrored prior single-cloud experiments: synthetic augmentation improved recall for rare cohorts but required careful proxy feature handling to avoid fairness regressions. Membership-inference tests revealed low-level leakage risks, prompting implementation of differentially private training and output filters. Audit logs and explainability bundle hashes provided tamper-evident trails useful for simulated regulatory reviews.

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806017

The deployment highlighted trade-offs: strict data locality increased orchestration complexity, and confidential computing added cost and slightly higher training latency. Nevertheless, the architecture met functional goals: explainability artifacts were verifiable, security posture improved through automated enforcement, and the system supported near-real-time decisioning for credit adjudication.

Secure Multicloud Architecture for Explainable Generative Credit Risk Models with Apache Data Pipelines and SAP HANA Processing describes a practical, enterprise-grade blueprint that reconciles the often-competing goals of high-performance risk scoring, interpretability, regulatory compliance, and resilient security across multiple cloud providers and on-premises boundaries; this architecture treats explainable generative models as first-class citizens in a rigorously governed data and compute fabric, and it layers Apache-born streaming primitives for large-scale ingestion and feature engineering with SAP HANA's in-memory, transactional analytical processing to deliver sub-second decisioning alongside forensic-grade auditability. At the data plane, the architecture uses an event-driven, schema-governed ingestion layer powered by Apache Kafka (or managed pub/sub equivalents) deployed in a multicloud pattern that places lightweight regional brokers near major data sources to minimize egress and latency while securely replicating critical topics to a central, policy-controlled tier using encrypted, authenticated replication channels; schema registries enforce canonical event contracts and backward/forward compatibility so that producers across clouds and on-prem systems can evolve independently without breaking downstream consumers. Stream processing is handled by containerized Flink or Spark Structured Streaming jobs that run in Kubernetes clusters provisioned per cloud region — these jobs perform deterministic, versioned feature engineering pipelines, enrichment from third-party feeds, streaming anomaly detection, and privacy-preserving transformations (tokenization, hashing, and early masking) before writing feature artifacts to a versioned feature store backed by Delta Lake or Apache Iceberg.

The feature store acts as the contract boundary between engineers and modelers and is co-located with an object-store tier configured for cross-region, secure replication; metadata and lineage are recorded using an immutable catalog so time-travel queries, point-in-time reconstruction, and auditor-friendly evidence packages are always available. For low-latency operational lookups required by decision microservices, curated operational feature sets and aggregated risk views are synchronized into SAP HANA instances provisioned as cloud-managed or on-prem appliances depending on data residency needs; SAP HANA's columnar, in-memory engine accelerates joins, cohort analytics, and provisioning calculations so that scoring endpoints can do sub-second reads even under heavy load. To enable multicloud resilience while preventing vendor lock-in, the architecture uses a hybrid topology: primary processing clusters and short-lived training jobs run where it is most cost-effective or where data residency dictates, while a replication and orchestration layer abstracts cloud-specific primitives through Terraform and Kubernetes Operators so that infrastructure-as-code templates are portable and reproducible across AWS, Azure, GCP, and private datacenters.

This layer also codifies network peering, private endpoints, and egress-minimizing patterns so that sensitive raw data rarely leaves its home region; when cross-cloud movement is required, it uses encrypted, audited data transfer mechanisms and fine-grained IAM roles scoped by least privilege. Security is implemented as defense-in-depth aligned with a zero-trust posture: every service-to-service call is authenticated and authorized via short-lived certificates or token exchange (SPIFFE/SPIRE or cloud-native IAM), a service mesh enforces mutual TLS and per-route policies, and API gateways provide centralized authentication, rate limiting, and request-level logging. Sensitive artifacts — raw transactions, PII, model checkpoints, and explain logs — are encrypted at rest using hardware-backed keys managed by a multi-cloud key management approach that leverages each provider's KMS but is orchestrated through a central key lifecycle manager to enforce rotation, split knowledge, and auditable key-usage policies; where regulatory regimes require extra isolation, confidential compute enclaves are used for training or scoring of the most sensitive models. Role-based and attribute-based access control is enforced across the data fabric with centralized policy engines (Apache Ranger or policy-as-code platforms) that administer row- and column-level masking, dynamic data redaction, and time-limited access grants; all policy changes and accesses are immutably logged into the audit ledger so that compliance teams can reconstruct who accessed what and why. The generative modeling layer is treated as both a risk-management and explainability resource: conditional VAEs, diffusion models adapted for tabular and sequential borrower histories, or conditional GANs are trained in reproducible MLOps pipelines that capture dataset hashes, hyperparameters, random seeds, code commits, and container images in a model registry. Generative models are used for three primary purposes—synthetic data generation for safe model development and cross-organization collaboration, counterfactual explainability to produce actionable "what-if" scenarios for individual credit decisions, and macro or micro scenario simulation to stress-test portfolios.

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806017

To guard against model misuse and prevent disclosure of original records through memorization, synthetic generation includes privacy safeguards such as nearest-neighbor disclosure checks, optional differential privacy noise budgets, and evaluation metrics (Wasserstein distances, KL divergence, pairwise correlation preservation) that balance fidelity and privacy. Counterfactual generators are governed by domain rule engines that ensure proposed edits are plausible and legally permissible (no impossible ages, inconsistent employment histories, or fabricated external bureau scores); counterfactuals are accompanied by minimal-edit explanations and a trace back to the feature transformations and raw events that produced them so auditors can validate plausibility and fairness claims. Explainability and model transparency are baked into the inference path: each decision emits a compact explain-log containing the model version, feature vector hash, top local attributions (e.g., SHAP or integrated gradients adapted for tabular data), linked counterfactuals, and provenance pointers into the versioned feature store and the model registry; these explain-logs are cryptographically signed, time-stamped, and persisted both in the versioned data lake and indexed in SAP HANA for rapid retrieval by dispute resolution, customer-service, and regulatory reporting flows. Operationally, explain-log generation is optimized so that consumer-facing rationales are provided with low latency (pre-compute surrogate model approximations or cached cohort-level narratives for common cases) while full technical artifacts are stored asynchronously for audit. MLOps pipelines enforce promotion gates that require passing validation suites: statistical validation (AUC, calibration, Brier scores), fairness constraints (disparate impact, equalized odds checks), explainability acceptance tests (stability and faithfulness metrics), and robustness evaluations including adversarial perturbation simulations and poisoning resistance tests. Canary deployments and shadow scoring enable safe evaluation of candidate models on live traffic without affecting production decisions; automated rollbacks are triggered when business or security KPIs deviate beyond codified thresholds.

Because the architecture spans multiple clouds, the CI/CD and orchestration layer abstracts provider differences through a combination of Kubernetes, Helm charts, and policy-driven Terraform modules; pipelines are executed in a centralized CI system that can dispatch ephemeral runners in the target cloud, ensuring that integration tests, vulnerability scans, and signed-image enforcement run close to where workloads will execute. Observability and threat detection are unified across clouds: distributed tracing and metrics (OpenTelemetry) link Kafka producers, stream processors, model-serving endpoints, and SAP HANA queries so engineers and security teams can reconstruct end-to-end flows.

V. CONCLUSION

Securely deploying explainable GenAI for credit risk in multicloud environments is feasible with careful architecture that centers governance, explainability, and cryptographic assurance. By combining Apache data pipelines for scalable ingestion and SAP HANA for authoritative feature management, organizations can achieve a balance between performance and compliance. Key enablers are policy-as-code, signed explainability bundles, federated IAM, and an automated validation and gating pipeline.

Data residency and legal constraints are encoded into pipeline manifests so that migration and replication respect locality rules; where multi-jurisdictional analytics is needed, the architecture supports remote enclaves or secure multi-party computation for collaborative modeling without moving raw PII across borders. From a governance standpoint, the architecture automates audit package generation: for any production model or decision batch, it can assemble a tamper-evident package containing the model artifact hash, training data snapshot identifiers, feature store versions, explain-logs, fairness test results, and security scan outputs — all signed and indexed for examiner retrieval. Regular independent model risk reviews, red-team adversarial assessments, and penetration testing against the multicloud ingress points are scheduled and their findings feed back into pipeline hardening and policy-as-code updates. Cost and operational efficiency are addressed through mixed provisioning strategies: use spot/ preemptible instances for non-critical training workloads, autoscaling groups for stream processors, and reserved capacity for SAP HANA operational clusters where predictable performance is required; observability into per-workload cost attribution helps tune placement decisions between clouds. Finally, organizational alignment is essential: cross-functional squads composed of data engineers, ML scientists, security engineers, platform SREs, legal/compliance officers, and business stakeholders jointly own model families and their lifecycle; shared runbooks, SLAs for incident response, and regular tabletop exercises ensure that the technology stack's multicloud complexity does not translate into operational fragility. By marrying Apache streaming primitives for resilient, versioned data flows with SAP HANA's in-memory analytical strength, and by wrapping these components in a rigorous, zero-trust, policy-driven multicloud fabric, this architecture enables explainable generative credit risk models to operate at enterprise scale — delivering fast, fair, and auditable

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806017

credit decisions while maintaining security, compliance, and the operational flexibility to adapt as data sources, regulatory landscapes, and threat vectors evolve.

A centralized SIEM ingests authentication logs, network flows, API gateway traces, model-serving telemetry (prediction distributions, confidence intervals, explained-attribution coverage), and explain-log access events; correlation rules and ML-based anomaly detectors identify suspicious behaviors such as repeated counterfactual probing from a single client, clustered feature vectors near decision boundaries indicating synthetic identity attacks, or unusual model-confidence shifts that may signal data poisoning. Incident playbooks are codified for typical scenarios: isolate affected inference endpoints, revoke compromised service tokens, roll back to last-known-good model artifact, replay immutable request logs against validated artifacts in a private sandbox, and assemble compliance packets for regulators. Network design considers latency, egress costs, and availability; data replication is tiered so that hot operational features replicate to nearby SAP HANA replicas for low-latency reads while cold archived raw events are stored in cost-optimized object stores with cross-cloud lifecycle rules. This tiering reduces egress by minimizing unnecessary cross-cloud queries and supports disaster recovery by maintaining geo-diverse copies with well-tested failover procedures.

VI. FUTURE WORK

- Benchmarking confidential computing approaches across clouds for ML training workloads to quantify cost-performance trade-offs.
- Developing standardized explainability bundle formats and verification protocols for industry interoperability.
- Researching more efficient constrained counterfactual generators suitable for synchronous use.
- Building federated learning extensions to allow model training across institutions without raw data sharing.

REFERENCES

1. Dwork, C. (2006). Differential privacy. *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, 1–12.
2. Tamizharasi, S., Rubini, P., Saravana Kumar, S., & Arockiam, D. Adapting federated learning-based AI models to dynamic cyberthreats in pervasive IoT environments.
3. Vinay, T. M., Sunil, M., & Anand, L. (2024, April). IoTRACK: An IoT based'Real-Time'Orbiting Satellite Tracking System. In 2024 2nd International Conference on Networking and Communications (ICNWC) (pp. 1-6). IEEE.
4. Sugumar, R. (2023, September). A Novel Approach to Diabetes Risk Assessment Using Advanced Deep Neural Networks and LSTM Networks. In 2023 International Conference on Network, Multimedia and Information Technology (NMITCON) (pp. 1-7). IEEE.
5. Adari, V. K. (2024). How Cloud Computing is Facilitating Interoperability in Banking and Finance. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(6), 11465-11471.
6. Suchitra, R. (2023). Cloud-Native AI model for real-time project risk prediction using transaction analysis and caching strategies. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(1), 8006–8013. https://doi.org/10.15662/IJRPETM.2023.0601002
7. Kotapati, V. B. R., Perumalsamy, J., & Yakkanti, B. (2022). Risk-Adapted Investment Strategies using Quantum-enhanced Machine Learning Models. American Journal of Autonomous Systems and Robotics Engineering, 2, 279-312.
8. Konda, S. K. (2024). AI Integration in Building Data Platforms: Enabling Proactive Fault Detection and Energy Conservation. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(3), 10327-10338.
9. Muthusamy, M. (2024). Cloud-Native AI metrics model for real-time banking project monitoring with integrated safety and SAP quality assurance. International Journal of Research and Applied Innovations (IJRAI), 7(1), 10135–10144. https://doi.org/10.15662/IJRAI.2024.0701005
10. Karanjkar, R., & Karanjkar, D. Quality Assurance as a Business Driver: A Multi-Industry Analysis of Implementation Benefits Across the Software Development Life Cycle. International Journal of Computer Applications, 975, 8887.
11. Kusumba, S. (2025). Modernizing Healthcare Finance: An Integrated Budget Analytics Data Warehouse for Transparency and Performance. Journal of Computer Science and Technology Studies, 7(7), 567-573.

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806017

12. Kingma, D. P., & Welling, M. (2014). Auto-Encoding Variational Bayes. *Proceedings of the International Conference on Learning Representations (ICLR).*
13. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 30*, 5998–6008.
14. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*, 4765–4774.
15. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.
16. Kandula, N. Evolution and Impact of Data Warehousing in Modern Business and Decision Support Systems
17. Papernot, N., Abadi, M., Erlingsson, Ú., Goodfellow, I., & Talwar, K. (2017). Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755*.
18. Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data, 5*(2), 153–163.
19. Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., & Venkatasubramanian, S. (2015). Certifying and removing disparate impact. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 259–268.
20. Kumar, S. N. P. (2025). Regulating Autonomous AI Agents: Prospects, Hazards, and Policy Structures. Journal of Computer Science and Technology Studies, 7(10), 393-399.
21. Suresh, H., & Guttag, J. V. (2021). A framework for understanding sources of harm throughout the machine learning life cycle. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAT)**, 1–12.
22. Muthusamy, P., Thangavelu, K., & Bairi, A. R. (2023). AI-Powered Fraud Detection in Financial Services: A Scalable Cloud-Based Approach. Newark Journal of Human-Centric AI and Robotics Interaction, 3, 146-181.
23. Konatham, M. R., Uddandarao, D. P., Vadlamani, R. K., & Konatham, S. K. R. (2025, July). Federated Learning for Credit Risk Assessment in Distributed Financial Systems using BayesShield with Homomorphic Encryption. In 2025 International Conference on Computing Technologies & Data Communication (ICCTDC) (pp. 1-6). IEEE.
24. Nagarajan, G. (2022). An integrated cloud and network-aware AI architecture for optimizing project prioritization in healthcare strategic portfolios. International Journal of Research and Applied Innovations, 5(1), 6444–6450. https://doi.org/10.15662/IJRAI.2022.0501004
25. HV, M. S., & Kumar, S. S. (2024). Fusion Based Depression Detection through Artificial Intelligence using Electroencephalogram (EEG). Fusion: Practice & Applications, 14(2).
26. A. K. S, L. Anand and A. Kannur, "A Novel Approach to Feature Extraction in MI - Based BCI Systems," 2024 8th International Conference on Computational System and Information Technology for Sustainable Solutions (CSITSS), Bengaluru, India, 2024, pp. 1-6, doi: 10.1109/CSITSS64042.2024.10816913.
27. Kumar, R. K. (2022). AI-driven secure cloud workspaces for strengthening coordination and safety compliance in distributed project teams. International Journal of Research and Applied Innovations (IJRAI), 5(6), 8075–8084. https://doi.org/10.15662/IJRAI.2022.0506017
28. Adari, V. K., Chunduru, V. K., Gonepally, S., Amuda, K. K., & Kumbum, P. K. (2024). Artificial Neural Network in Fibre-Reinforced Polymer Composites using ARAS method. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(2), 9801-9806.
29. OpenLineage Community & Data Engineering Reports (2023). Open metadata standards for data lineage and observability in modern data stacks. *Industry Whitepaper*.