



From Trade to Compliance a Cloud-Native AI Architecture for Real-Time Derivatives Pricing

Dhanaraj Sathiri

Independent Researcher, Malaysia

dhanrajsathiri@gmail.com

Natalie Cooper

Graduate Research Assistant, Lakeshore Institute of Technology, United States

nataliecoopeerr@outlook.com

Publication History: Received: 30.06.2026; Revised: 28.07.2026; Accepted: 03.08.2026; Published: 14.08.2026.

ABSTRACT: The growing demand for low-latency computing in derivatives markets challenges the traditional approach of monolithic applications implemented in dedicated data center clusters. Cloud-native architectures offer a promising solution due to the flexibility, scalability, and resilience of the cloud. Nevertheless, for latency-sensitive workloads, latency is not only a requirement; cloud-native design principles can lead to improvements in the computing model itself. These include breaking down monolithic applications into microservices deployed in a service mesh, taking advantage of on-demand and pay-per-use resources to optimize costs for off-peak traffic, and improving computing efficiency through loading-balancing across cloud availability zones.

Furthermore, the demonstrably increasing risk of operational resilience in the cloud at a suitably architected cloud level requires organisations and business partners to reconsider their outsourcing risk appetite. Even where a loss occurs, insurance can cover events, as long as the incident doesn't lead to the organisation ceasing to trade for an extended period. For many services, intelligent brokerage can ensure continued service. Examined in more depth, these principles may also strengthen the fundamental efficiency and performance of the models beyond a simple scaling gain.

KEYWORDS: Low Latency Computing, Derivatives Market Infrastructure, Cloud Native Architecture, Microservices Decomposition, Service Mesh Design, Latency Sensitive Workloads, Real Time Trading Systems, Elastic Compute Scaling, Pay Per Use Resource Optimization, Load Balancing Across Availability Zones, Cloud Resilience Engineering, Operational Resilience Risk, Outsourcing Risk Appetite, Intelligent Service Brokerage, High Availability Architectures, Cloud Based Market Systems, Performance Optimized Computing Models, Financial Market Infrastructure Modernization, Scalable Trading Platforms, Resilient Cloud Ecosystems.

I. INTRODUCTION

Market risk, valuations, capital reserves, and liquidity all are governed by financial institutions addressing the derivatives markets, which also remain fundamental to managing risk and enabling speculation. Market participants rely on a spectrum of vendors to price, hedge, and process transactions involving the underlying. Regulatory requirements for transaction submission and trade-related data — driven by the 2008 financial crisis and the collapse of major financial institutions — must be met for all parties and in many cases are accompanied by real-time monetization. A cloud-native approach to product, price, risk, and regulatory governance management is requisite because of the high stakes involved. But a market- and data-driven cloud-native architecture is essential. Products need to be priced accurately, reliably, and with the minimal latency possible. From a computational point of view, the processes can be deep and prone to errors, and there can be a high volume of instant price and risk request traffic.

A cloud-native architecture can deliver the requisite fault tolerance on the pricing side by enabling the underlying price models to be computed easy-to-access cacheable elements in the future via a data lakehouse architecture. The resulting product-price-risk plus transaction-processing storm needs to be processed directly via cloud-native principles in a cloud-native way. And on-the-fly processing to serve the real-time monetization market is required. Such an approach enables a market vendor to underpin real-time price, risk, and transaction processing, offer pricing models for the pricing tree



during that market, and allow clients simple latency querying for instantaneous transaction marking, risk updating, and executions during their functional and financial storm period.

1.1. Overview of the Study

Advances in cloud-native technology enable new paradigms for building and operating scalable, resilient systems to improve performance through wider distribution. Market-makers, enables, and liquidity-providing clients are looking for fast execution, which requires low-latency-local solutions that push risk controls as close to the pricing engines as possible. Deploying parts of a pricing or risk computation engine in the cloud, closer to the demands on the asset, may be the only way to fulfil customer demand—and compliance requirements, business continuity, and operational risk—at an acceptable cost.

The market for regulations is large, and the response is slow, creating compliance risk. Trading and security-regulatory compliance are separate functions that rarely share any common data. The approval of new regulations follows public scrutiny with lengthy feedback loops and often contradictory outcomes among multiple regulators aiming to reach the same goal. Detecting security and trading infraction transactions often comes too late, as multiple regulators may investigate parallel without real-time intelligence. The traditional waterfall process is too slow and costly for regulatory compliance and real-time detection. Cycles of continuous and discrete audit and control checks are necessary to govern new transactions. An architecture that builds compliance features into the trading system itself, rather than leaving compliance to a post-trade activity, lowers compliance creation costs and increases speed to market for compliance with new rules. By treating compliance as true revenue rather than a support function, the costs can be capitalized and subsequently recovered.

This work outlines a modular cloud-native architect as a foundation for an end-to-end AI-powered solution that rapidly addresses compliance regulations in the derivatives space. The building blocks for a real-time derivatives-processing pipeline are described, and the pipeline is augmented with a trading compliance and security system that meets future-market risk and regulatory demands, with the rapid implementation of new or amended regulatory compliance demands that frequently arise in both areas.




Fig 1: Cloud-Native Modular Architectures for Real-Time Derivatives Compliance: Shifting from Post-Trade Auditing to Embedded Regulatory Intelligence


II. BACKGROUND AND MOTIVATION

Many financial institutions have expressed interest in cloud-native architecture for real-time derivatives pricing. In the pricing domain, stringent latency requirements, given the size of the trade, have often meant on-premise solutions to minimize latency, despite cloud-native architecture being easier to build, scale and operationalize. This also reduces



vendor concentration risk for resilience as cloud-native microservice-based architectures are managed by different vendors and can be loaded on a per-region basis. Reducing vendor concentration risk for pricing is not only business but also country risk management. This paper models derivatives processing in cloud-native architecture, handling distribution of both pricing and risk computation, with the aim of scaling up compute and also reducing region-based vendor concentration risk on pricing with lower latency for large trades.

However, many compliance and regulatory governance requirements require audit trails, data lineage and security setting that are inherently large overheads. For example, market and data compliance regulations (such as MI-FID II, CCAR, FINRA, DGS, FRTB, BCBS 239, GETS, CGETS) demand a complete data and model lineage for years. These regulatory requirements have to be built into the architecture and processes so that the complete audit trace is captured with minimum or no human intervention. The complete data model lineage of both market data and integrity/completeness of the market data is often as large, if not more than, the distribution price computation. Audit trail, security, and access of the business are also as large in overhead as product pricing on a derivatives processing system. Thus, Chuck et al. (Marvel) proposed a compliance rule-based Governance (G) architecture integrated into AI to remove bias and control the AI model audit requirements with minimal overhead in model development and parameter tuning.



2.1. Significance of Cloud-Native Approaches in Derivatives

A cloud-native architecture supports high-speed transaction processing and ultra-low-latency risk management; it addresses other key needs apparent from growing demand for operational resilience and regulatory control.

Rapid growth in the global derivatives market has heightened the importance of low-latency solutions for pricing, risk computation, and operational resilience: Cloud,—and especially cloud-native—architectures are being proposed to respond to these challenges. But derivatives markets have special demands that can make the application of cloud technology in this area tricky. For example, cloud-native technologies are often applied to the construction of operationally scalable infrastructures that facilitate the building and running of many concurrent systems across the financial services industry within a defined Region. However, the need for very low-latency transaction processing, ultra-low-latency risk computation, and operational resilience—where even single-digit sub-millisecond latencies in processing are vital—mean that for such systems cloud infrastructure must be provisioned regionally, within a single data centre, to be effective. Even then, these localised environments may not always have the 3*2 specifications required for optimal performance. Rather, such environments must be shortened further, to become a 1*2*3 zone with extremely low-latency routing.

Operational resilience is now required not only for exchange Centre but throughout the ecosystem. The need for these services to be resilient from cloud failure is also important, especially as end-user final-line solution-run teams migrate operations to cloud. Anti-fragility is thus becoming a driver for these solutions. At the same time, demand for robust data governance and management control is gaining new momentum due to the operational demands of evolving Data Protection Acts and Data Protection Regulation, GDPR, MSSP, MiFid II, and many other related regulations that now govern financial services...exponential growth is expected across the industry over the coming years, but also in relation to signalling and alerting.



Metric	Value
Products N	345,000
Time budget T	10 s
Required throughput $R = N/T$	34,500 products/s

Equation 1: Pricing throughput equation

Step-by-step derivation

Let:

- $N = 345,000$ products
- $T = 10$ seconds (deadline)

Required throughput

$$R = \frac{N}{T}$$

Substitute:

$$R = \frac{345,000}{10} = 34,500 \text{ products/s}$$

Concurrency (workers/cores) needed

If average compute time per product is t seconds/product (single-threaded worker), then one worker's service rate is $\approx 1/t$. To sustain throughput R , the required concurrency W is:

$$W = R \cdot t$$

Example with $t = 5 \text{ ms} = 0.005 \text{ s}$:

$$W = 34,500 \cdot 0.005 = 172.5$$

III. ARCHITECTURAL PRINCIPLES FOR CLOUD-NATIVE AI IN DERIVATIVES

Successfully capturing the advantages of cloud-native computing and deploying a highly automated AI-driven architecture for real-time derivatives transactions should pragmatically follow a set of guiding principles. The intention is to establish a neurocloud-based shared and modular business ecosystem to maximize the flexibility and scalability of pricing, model management, risk computation, distribution, and compliance. Latency-sensitive pricing forms key prerequisites for practical adoption, creating a need for distributed on-demand pricing models with idempotency guarantees.




Fig 2: NeuroCloud: A Modular, Event-Driven Architecture for Automated AI-Enabled Derivatives Liquidity and Compliance

The ensuing architecture must combine microservices to enforce modular boundaries for testing, supports open and maintainable tooling based on stateless services, adheres to the event-driven processing paradigm, provides precise observability, and aligns with all relevant market compliance requirements.



Core building blocks comprise serving microservices and their monitoring, management, governance, and quantification plugins deployed within a service mesh. A hub-and-spoke approach implements a data lakehouse with a feature store to capture and expose query features across the entire business. Data provisioning, governance, semantics, and quality encompass a complete and semantically consistent view of all properties, including the management of data model life cycles to support valid and compliant real-time analytics.

3.1. Core Design Principles for Cloud-Native AI in Derivatives


Core design principles for cloud-native AI in derivatives are articulated, encompassing microservices boundaries, stateless processing, event-driven workflows, idempotency, observability, and governance controls. Modularity allows individual service components to be specified, optimized, maintained, and deployed independently. Scalability requires elastic resource provisioning in multiple cloud regions. Fault tolerance accommodates failures in any component, detecting errors and making appropriate corrections. Security embeds auditing, traceability, and regulatory governance across both architecture and processes.

Architectural development is guided by established cloud-native principles, with six design considerations extending beyond general cloud-native qualities to address AI-specific needs: modularity, scalability, fault tolerance, security, observability, and idempotency. Modularity is a central tenet of microservices, enabling specific services to be optimized and developed independently. Latency-sensitive applications benefit from provisioning in multiple cloud regions, allowing automatic scaling that responds to demand elasticity. AI models for pricing, risk, and other derivatives functions often require intensive compute resources, necessitating careful management to minimize delays. However, resilience to runtime failure is even more critical. Derivatives are valuable assets typically held for only a few days: corresponding trades last a few seconds or minutes, while processing times should aim for milliseconds. Such short-lived, high-volume traffic patterns create the potential for custom-built hardware but do not justify the cost of such an investment; operational services need only fail gracefully.

IV. REAL-TIME DERIVATIVES PROCESSING PIPELINE

The derivatives processing pipeline implements the end-to-end specification. Data ingestion from trading, reconciliation, and market feed systems supports ingestion, validation, and normalization, with effectivity and lineage tagging, feature resolution, and schema-lifecycle controls. Ingestion-enriched events drive event-driven architecture: option-pricing and risk-computation services model compute as features, supporting real-time data science. Cache coherence meets sub-millisecond latency targets, enabling low-latency function publication. Model-serving and feature-store caches portrait feature replays.

The violation of capital and risk limits by dealers causes a volatility surge in derivatives prices while demand rises from clients. These events lead to risk latency, limits violation, and put dealers in arbitrage risk. The background to the architecture allows introducing a microservice development approach using Azure AI ML to provide a real-time option-pricing pipeline serving a low-latency production environment. Integration is sought across multiple cloud-natives and the services computing nodes while providing production-monthly-use change. Microservice principles supporting production-quality components enhance configurability and scalability. Network-load-management patterns and the Kubernetes environment finally provide a multi-environment load balancing and scaling solution.





Equation 2: End-to-end latency equation (sub-millisecond target + cache)

Step-by-step derivation: latency decomposition

Let end-to-end latency be the sum of stage latencies:

$$L_{e2e} = L_{ingest} + L_{normalize} + L_{feature} + L_{route} + L_{compute} + L_{audit}$$

If the system uses caching, average compute latency depends on cache hit ratio h :

- Cache-hit path latency: L_{hit}
- Cache-miss path latency: L_{miss}

Average:

$$E[L] = h \cdot L_{hit} + (1 - h) \cdot L_{miss}$$

And you still need the compliance/audit emission overhead:

$$L_{e2e} = E[L] + L_{audit}$$

4.1. Data Ingestion and Normalization

The real-time derivatives processing pipeline operates on market, reference data, and event feeds that drive all aspects of trading operations from pricing and risk management to detection of market anomalies. All feeds are ingested into a dedicated staging area in Cloud Storage, which permits automation and simple parallelism. Data is then consolidated, normalized, and made available to downstream applications in a structured format. An event-driven architecture detects changes and updates interested consumers asynchronously.

Incoming events are validated against quality specifications—such as acceptable ranges for pricing curves or surface normals—and assigned a freshness score based on market close timestamps, ensuring stale data does not propagate. Operations teams receive alerts for flagged deviations, and data is rejected if failing thresholds for critical data sets. Ingested data includes market state and trading positions, ultimately normalized to an open, high, low, and close schema. Limit orders—either volume-weighted intakes for primary, visible orders or implicit orders for non-explicit trading—are built for pricing and risk computation. The IPC schema facilitates association with side information and loading into a temporal data lake for future analysis.

Auditability and lineage are guaranteed by explicit naming patterns in Cloud Storage, which reflect original schema, producer, data classification, source, and freshness. Structural and semantic models enforce quality expectations as part of a broader regulatory compliance initiative, deterring price and market manipulation through the introduction of explicitly orchestrated disorder.

4.2. Event-Driven Computation and pricing

Event-driven computations respond to normalized reference data ingested from disparate sources and reservoirs (Figure 2). Internal pricing services retrieve normalized data for staging feature pipelines—data cleansing, emerging market directions, and product-specific model variables. Prices are computed with separated models for swaps, swaptions, and cross-currency basis. Inputs converge into newly introduced concept services—option tans for Bermuda-style transactions, lumpsum amounts for notify/confirm segments—and enforce organization-specific features. Market needs dictate high-quality pricing latency, set for 10 seconds with a product count of 345,000. An internal query service retrieves valid market points, identifying dangerously outdated segments. Feature engineering and prediction for the AI-based Risk charge follow the same approach; supporting concepts appear in the third section.

Stability (X) is ensured with pre-computed hot caches for main market segments. Hot caches are useful with stable market texture but may vaporize during a volatility spike or a historic pickup in new issued volumes. Rarely, with longer timeframes, predictions for products with 5-year outlier time variability may be refreshed more slowly. When market conditions call for predictions, a dedicated service invokes fresh computations, and a dedicated replacement process engages when both stability checks fail. Severe market uncertainty heightens a bank's risk of excessive consumerism; during such periods, a bank rushes to plead with the market-makers to conceive a moral treatise ruling out risky positions through modeled consideration.



Avg compute time per product (ms)	Required throughput (products/s)	Approx concurrent workers needed
1	34,500	34.5
2	34,500	69.0
5	34,500	172.5
10	34,500	345.0
20	34,500	690.0
50	34,500	1,725.0

V. COMPLIANCE AND REGULATORY GOVERNANCE FRAMEWORK

A compliance and regulatory governance framework for data and analytics is proposed, aligned with market regulations, built-in governance, and observability. It provides a comprehensive mapping of regulatory requirements to organizational and architectural controls, ensuring that adequate processes, tools, and technologies creating data and analytics follow compliance regulations, providing audit trails and traceability. Key aspects include a security by design plan combined with an operational model to ensure separation of duties and role segregation for compliance workload processing. The end-to-end data and analytics flow provides data lineage for market data and pricing resulting from the derivative transaction and compliance process.

Market data created to provide near real-time updates to positions created by the transaction processing system at the exchange or brokerage level and triggering pricing models in stations close to market data sources is monitored through a combination of rules over the pipeline data processed and alerting based on pipeline observability, ensuring the needed trust in this computation. Data observability integrated into the data-flow notifies possible data quality issues. Components processing tagged data for compliance analytics are monitored through these same concepts assuring their availability and reliability to follow the demand of the workload.




Fig 3: Integrated Governance by Design: A Multi-Layered Observability and Compliance Framework for High-Frequency Market Data Ecosystems



5.1. Regulatory Requirements Mapping

Regulatory governance extends beyond compliance mandates, necessitating formalized processes that respond to internal and external parties, encompassing risk appetite, monitoring, and follow-up. Data and operational pipelines are also subject to legal scrutiny, driven by GDPR, user agreements, and financial regulations. These further conditions determine the lifecycle of data sources and features, with controls for data quality and access.

These requirements represent specific compliance controls within the broader governance framework. Each API needs associated audit trails and traceability. Compliance mandates derive from sensitivity analyses on derivatives pricing. Requirements mapping ensures incorporation into architecture, services, and flows, confirming the deployment of related governance controls.

5.2. Monitoring, Reporting, and Notification


While risk governance is a priority, compliance with other regulations is also critical. A detailed mapping of the relevant regulations ensures that controls are applied to all relevant processes. Because these controls impact nearly all parts of the architecture and organization, they are not separately enumerated here, presenting a risk of underestimating their scope. Nevertheless, compliance with these requirements is paramount.

Reporting and notification requirements arise from many regulations, including anti-money-laundering policies and the Securities and Futures Ordinance. Audit trails are part of many market regulations, such as MIFID-II, SRB, and the Hong Kong Code on Unit Trusts and Mutual Funds. Lineage information, including the schematic details of data sources, is necessary for other regulations, such as BCBS 239 and DFSA Rulebook. Thus separation of duty needs to be implemented, either within the architecture or within specific services, particularly for functions related to compliance and reporting.

VI. CLOUD-NATIVE ARCHITECTURE AND INFRASTRUCTURE

An end-to-end cloud-native architecture with features supporting microservices, security, reliability, and observability patterns completes the discussion. Such a cloud-native architecture facilitates the deployment of an event-based, modular application. A robust service mesh enables secure service discovery, observability, and cross-cutting functionality. Data is managed within a designated cloud account, while optimized infrastructure patterns reduce overhead. The previously described processing, pipeline, and governance capabilities are realized using an event-driven microservices architecture running on a managed Kubernetes service. Microservices boundaries defined by functional domains, application microservices grouped using logical sensible domains. A dedicated service mesh provides underlying support for service discovery, communication, and security. Observability requirements are fulfilled using open-source projects. Critical auditing, incident-response and enabling the observability gap with trifecta of metrics logging and tracing.

Processing uses a lakehouse approach that unifies a data warehouse with a data lake, having the combined capabilities of the two in a single store. Supporting the accessibility for analytics without being duplicated. The governance of the lake is aided by a data governance tool that maintains a catalogue of metadata and offers enhancements like data-lineage tracing impact analysis classification metadata-based discovery. The governance tool is further supported by a semantic model that adds a business glossary. The feature stores supports the consumption of big-data in machine-learning workload at a faster pace by caching the intermediate features to minimize the processing time. Policy-based management further maintains the sustainable use of the environment.





Equation 3: “Freshness score” equation (data quality gate)

Step-by-step derivation: freshness from time-delta

Let:

- t_{close} = market close timestamp for the relevant feed/instrument
- t_{now} = current processing timestamp
- $\Delta t = t_{now} - t_{close}$ (age)

You need a function mapping age $\rightarrow [0,1]$.

A common, stable choice is **exponential decay** with half-life H :

$$F(\Delta t) = \exp\left(-\ln(2) \cdot \frac{\Delta t}{H}\right)$$

Properties (step-by-step):

- If $\Delta t = 0$: $F(0) = e^0 = 1$ (perfectly fresh)
- If $\Delta t = H$: $F(H) = e^{-\ln 2} = 1/2$ (half as fresh)
- As $\Delta t \rightarrow \infty$: $F \rightarrow 0$

Freshness gate (reject stale data)

If the system rejects data below a threshold τ :

$$Accept \Leftrightarrow F(\Delta t) \geq \tau$$

Equivalently solve for max age allowed:

$$F(\Delta t) \geq \tau \Rightarrow e^{-\ln(2)\Delta t/H} \geq \tau \Rightarrow -\ln(2)\Delta t/H \geq \ln(\tau) \Rightarrow \Delta t \leq -\frac{H}{\ln 2} \ln(\tau)$$

6.1. Microservices and Service Mesh

The proposed architecture adheres to a microservices style, with inferencing and auxiliary functions encapsulated as independent, self-sufficient services that share no state or data. Service contracts and APIs use open standards (gRPC, REST, JSON) to support cross-platform service composition. Normalization and feature-pipeline services are implemented as stateless components that do not introduce state or side effects to the overall processing workflow, aside from the creation of specialized event records. Latency-sensitive computation and serving workloads are typically run on a dedicated cloud region with the lowest round-trip latency to trading venues; pipelines for order- and credit-line monitoring and reporting follow suit. Processing is triggered by multiple events through a cloud-native workflow system, with independent, distributed workloads for pricing and computation.

A service mesh abstracts much of the operational complexity of microservices deployments, directing network traffic through a dedicated data plane of sidecar proxies that automatically inject security policies (e.g., authentication, authorization, encryption) and visibility features (traffic monitoring, logging, tracing, alerting) within the data plane. Requests and responses are also automatically routed and authenticated via an API gateway that exposes a uniform API surface for external clients.

6.2. Data Lakehouse and Feature Stores

A lakehouse encapsulates the high-level properties required to govern a compliant data asset for multiple service consumers, and thus, is common across services. The data asset is organized as a semantically ground-truthed, versioned, and governed pool of both raw and processed data, residing in a cloud-native data storage layer that is optimized for both the speed of hot analytic queries (Parquet) and the cost of cold workloads (ORC, iceberg tables). Apache Hudi governs the lifecycle of raw data stored in the data lake such that hot data is kept available for cross-service exploration and avoided during cold-query planning. Support for external tools (PowerBI, Tableau) is enabled by semantic models for approved sets of assets. Transformed data is versioned into a structured source-of-truth model for all potential consumers and enabled with machine-readable quality and sensitivity policies along specific data axes.

Feature stores offer a real-time, privacy-compliant, and operationally resilient solution for powering ML-based functions along event-driven workflows for analytics-augmented transaction processing. The feature pipeline orchestrates the extraction of the features from the lakehouse, fuses them with real-time data in accordance with parameterizable slicing profiles, and serves pre-computed feature sets under hierarchical partitioning rules that automate data-flowing across the zones of a cloud-native data lake. It acts both as a source of truth for the features applied in model inference and model retraining, and serves features across both online and batch ML workloads for complex derivatives. Sensitivity analysis is integrated to guide governance policies and information controls on features used in production models.



VII. CONCLUSION

The architecture presented enables the processing and storage of derivatives transactions and events in a cloud-native environment, moving beyond transaction security checking to support additional data-driven services and regulatory compliance. The next steps include scaling to production volumes, supporting the storage and retrieval of derivative features, implementing tools for visualizing the results of machine-learning capabilities, and integrating regulatory governance support. Furthermore, the scope of the architecture is being broadened to investigate additional cloud-native AI services for risk control and compliance checking.

The transition from transaction to compliance requires an integrated view, encompassing both the ongoing governance of transactions and the additional and historical data underpinning compliance checks. An additional challenge is that support services must be able to process transaction volumes of tens and hundreds of millions of transactions each day while also providing low-latency pricing for derivative instruments. By moving toward these aims, a robust prototype is being developed to guide a validated production deployment.




Fig 4: Functional Scope: Transaction to Compliance

7.1. Final Thoughts and Future Directions

Interest in cloud-native architectures for latency-sensitive operational processes—such as pricing, risk computation, and trade emission into transactions or clearing formats—has grown. The COVID-19 pandemic significantly increased demand for cloud-native solutions, revealing the considerable operational risks of highly concentrated, home-hosting operations for swap dealers. Such concentration exposes institutions to multiple correlated risks: the risk of losing the assets themselves, of being unable to meet their own operating and transactional requirements, and of being unable to support customers and counterparties. Solutions involving infrastructure and services hosted by third parties acting within a shared ethical framework that ensures compliance with market and data regulations can therefore facilitate compliance, scalability, and operational resilience in such latency-sensitive markets by distributing operational resources more equitably while also providing a more sustainable use of global resources.

The final thoughts and future directions consider a representative cloud-native AI architecture designed for a regulatory framework-compatible end-to-end processing pipeline for real-time derivatives transactions. Such a pipeline is required to enforce compliance with audit and traceability controls throughout the processing. Unplanned latency in the generation of features and labels, required for the successful serving of predictive machine-learning models, should also be kept within a feasible-boundary. The resulting architecture thus enables not only the end-to-end processing of transactions in real-time, but also the compliance assessment of the expected services within a well-understood latency tolerance.

REFERENCES

1. Polineni, T. N. S., Yasmeen, Z., Ganti, V. K. A. T., Nirvin, P., & Vignesh, P. (2025, February). Climate Change Perception and Adaptation Strategies for Crop Growth with CNN and BiGRU Models. In 2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS) (pp. 1-6). IEEE.
2. Maguluri, K. K., Chandrasekar, M., Gorshetwar, S. D., Prasanna, S., Anusuya, M., & Vallarasu, K. (2025, February). Applications and advancements of nanoparticle-based drug delivery in lung cancer treatment. In 2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS) (pp. 1-6). IEEE.



3. Challa, K., Singireddy, J., Pamisetty, A., Garapati, R. S., Kannan, S., & Sriram, H. K. (2025, December). Harnessing Agentic AI and IT Infrastructure in Banking to Drive Consumer Insights, Operational Excellence, and Intelligent Financial Innovation. In 2025 3rd International Conference on IoT, Communication and Automation Technology (ICICAT) (pp. 1-7). IEEE.
4. Burugulla, J. K. R., Kannan, S., Malempati, M., Challa, H. A. H. S. R., Pandugula, C., & Goma, T. (2025, April). Federated Learning Based Cloud Computing Solutions with Privacy Preservation for Intelligent Utilities. In 2025 4th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
5. Chava, K., Challa, S. R., Sriram, H. K., Chakilam, C., Kannan, S., & Annapareddy, V. N. (2025, July). Utilizing Query Performance in NoSQL Databases for Applications Based on Big Data. In 2025 2nd International Conference on New Frontiers in Communication, Automation, Management and Security (ICCAMS) (pp. 1-6). IEEE.
6. Sirisha, N., Kumar, D. S., Ilanchezian, P., Komaragiri, V. B., Sri, G. K., & Thamizhkani, B. (2025, May). Energy Aware Artificial Intelligence Models for Scalable Digital Interactions Efficient Resource Management and Carbon Neutral Sustainable Technology Innovations. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 576-586). Atlantis Press.
7. Ande, R., Pandugula, C., Mehta, D., Vankayalapati, R., Birbal, P., Verma, S., ... & Nanavati, N. (2025). Understanding climate change impacts on streamflow by using machine learning: case study of Godavari Basin. *Water*, 17(8), 1171.
8. Rani, P. S., Amistapuram, K., Pamisetty, V., Singireddy, S., Kummari, D. N., & Sheelam, G. K. (2025, November). Hybrid Knowledge Graph-Deep Learning Framework for Automated Exception Handling and Investigation in Complex Insurance Claims. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
9. Vadisetty, R., Nuka, S. T., Kalisetty, S., Pandugula, C., Burugulla, J. K. R., & Annapareddy, V. N. (2025, February). Generative AI for Advanced Recycling Processes in Polyethylene and Polypropylene Manufacturing. In International Ethical Hacking Conference (pp. 269-284). Singapore: Springer Nature Singapore.
10. Amistapuram, K., Pandiri, L., Raju, V. R., Paleti, S., Singireddy, S., & Sheelam, G. K. (2025, December). AI-Based Cloud Infrastructure and MLOps Frameworks for Scalable Data Engineering Across Banking and Insurance. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 186-192). IEEE.
11. Sheelam, G. K., & Komaragiri, V. B. (2025). Self-adaptive wireless communication: Leveraging ML and agentic AI in smart telecommunication networks. *Metallurgical and Materials Engineering*, 1381-1401.
12. Suura, S. R., Chava, K., Chakilam, C., Malempati, M., & Nuka, S. T. (2025, April). Revolutionizing Retail Supply Chains with AI and Cloud Computing: A Big Data Approach to Demand Forecasting and Logistics Optimization. In International Conference on Smart Computing and Informatics (pp. 271-283). Cham: Springer Nature Switzerland.
13. Chabokpour, J., Kalisetty, S., Malempati, M., Challa, K., Mandala, V., Kumar, B., & Azamathulla, H. M. (2025). Integrating Hydrological and Hydraulic Approaches for Adaptive Environmental Flow Management: A Multi-Method Approach for Adaptive River Management in Semi-Arid Regions. *Water*, 17(7), 926.
14. Vadisetty, R., Kannan, S., Vankayalapati, R. K., Sriram, H. K., Chakilam, C., & Malempati, M. (2025, February). AI-Driven Analysis of Urban Heat Islands for Sustainable City Planning. In International Ethical Hacking Conference (pp. 131-147). Singapore: Springer Nature Singapore.
15. Kumar, B. H., Nuka, S. T., Malempati, M., Sriram, H. K., Mashetty, S., & Kannan, S. (2025). Big data in cybersecurity: Enhancing threat detection with ai and ml. *Metallurgical and Materials Engineering*, 31(3), 12-20.
16. Nagamani, A., Nuka, S. T., Iyyanar, P., Suzana, A. A., Nayak, M., & Kumar, S. S. (2025, May). Artificial Intelligence optimized citizen engagement platforms for resilient urban futures inclusive policy making and smart community driven governance. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 221-231). Atlantis Press.
17. Pamisetty, A. (2025). The Psychology of Addiction: Understanding Behavioral Patterns. *SCIENTIFIC CULTURE*, 11(1), 1-15.
18. Ceasar, M. J., Karthik, N. L., Paleti, S., & Devaru, S. D. B. Entrepreneurial risk-taking and cultural values: A global behavioral perspective.
19. Annapareddy, V. N. (2025). The Intersection of Big Data, Cybersecurity, and ERP Systems: A Deep Learning Perspective. *Journal of Artificial Intelligence and Big Data Disciplines*, 2(1), 45-53.
20. Adusupalli, B., Malempati, M., Paleti, S., Mashetty, S., & Singireddy, J. (2025). Integrated financial ecosystems: AI-driven innovations in taxation, insurance, mortgage analytics, and community investment through cloud, big data, and advanced data engineering. *Journal of Information Systems Engineering and Management*, 10, 1103-1117.
21. Sheelam, G. K., Meda, R., Pamisetty, A., Nuka, S. T., & Sriram, H. K. (2025). Semantic Negotiation Among Autonomous AI Agents: Enabling Real-Time Decision Markets for Big Data-Driven Financial Ecosystems. *Metallurgical and Materials Engineering*, 31(4), 587-598.



22. Komaragiri, V. B. (2025). Agentic ai for autonomous network orchestration: A new frontier in telecommunications. *eksplorium-buletin pusat teknologi bahan galian nuklir*, 46(1), 221-241.
23. Endla, P., Kannan, S., Adhiyaman, M., Ramulu, K. P., Kumar, B. B. C., & Hemamalini, M. (2025, May). Sustainable Artificial Intelligence Orchestrated Urban Planning Models with Digital Twin Technology for Smart Cities and Green Development. In *International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024)* (pp. 1044-1058). Atlantis Press.
24. Sadhukhan, S. K., Jaishankar, R., Adusupalli, B., Revathi, V. M., Kaliappan, S., & Ponnusamy, M. (2025, October). Healthcare Insurance Fraud Detection Using a Blockchain–Deep Neural Network–AdaBoost Hybrid Model. In *2025 2nd Asian Conference on Intelligent Technologies (ACOIT)* (pp. 1-6). IEEE.
25. Manikandan, K., Pamisetty, V., Challa, S. R., Komaragiri, V. B., Challa, K., & Chava, K. (2025). Scalability and efficiency in distributed big data architectures: a comparative study. *Metallurgical and Materials Engineering*, 31(3), 40-49.
26. Koppolu, H. K. R., Gadi, A. L., Motamary, S., Dodda, A., & Suura, S. R. (2025). Dynamic Orchestration of Data Pipelines via Agentic AI: Adaptive Resource Allocation and Workflow Optimization in Cloud-Native Analytics Platforms. *Metallurgical and Materials Engineering*, 31(4), 625-637.
27. Nandal, N., Tyagi, K., Kumar, L., & Gadi, A. L. (2025, November). AI-Driven CV Analysis and Job Matching in HRM Using a Hybrid Fuzzy Logic-Deep Neural Network Approach for Higher Education. In *2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN)* (pp. 1-6). IEEE.
28. Kumar, S. S., Garapati, R. S., Segireddy, A. R., Kalisetty, S., Inala, R., & Nagabhyru, K. C. (2026, March). Hybrid Deep Neural Network–DevOps Pipeline Optimization for Risk Prediction in Cloud-Native Workers' Compensation Platforms. In *2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON)* (pp. 1-6). IEEE.
29. Mohan, A. A., Maguluri, K. K., Yasmeen, Z., & Pandugula, C. (2025). Evaluating the structural performance of axially loaded stainless steel circular tubular columns with ultra-high-performance concrete using an optimized attention pyramid convolutional neural network. *Structural Concrete*.
30. Challa, K., Challa, S. R., Pamisetty, A., Kaulwar, P. K., & Koppolu, H. K. R. (2025, December). Transforming Payments: The Role of AI and Big Data in Fraud Alerts, Credit Monitoring, and Secure Transactions. In *2025 IEEE International Conference on Communication Networks and Computing (CNC)* (pp. 1406-1414). IEEE.
31. Syed, S., Nampalli, R. C. R., Nikam, M., Krishnan, T., & Perada, A. (2025, February). IoT-Driven Environmental Pollution Monitoring with a Deep Attentional Hybrid Transformer Model. In *2025 International Conference on Emerging Systems and Intelligent Computing (ESIC)* (pp. 356-361). IEEE.
32. Komaragiri, V. B., Keerthana, G., Fathima, S. K., Vats, S., Priya, R. V., & Sivagami, S. (2025, May). Quantum Computing Integrated Digital Interaction Platforms for Secure Public Engagement Decentralized Data Exchange and Policy Optimization. In *International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024)* (pp. 1989-2001). Atlantis Press.
33. Recharla, M., & Nuka, S. T. (2025). Translational Approaches To Commercializing Neurodegenerative Therapies: Bridging Laboratory Research With Clinical Practice. *South Eastern European Journal of Public Health*, 121-144.
34. Priyanka, R. P., Annapareddy, V. N., Yenugu, B. C., Nagabhyru, K. C., & Kapila, D. (2025, September). Optimization of Battery Management Systems Using Machine Learning. In *2025 International Conference on Computing and Communications (COMPUTINGCON)* (pp. 1-6). IEEE.
35. Recharla, M., Chakilam, C., Kannan, S., Nuka, S. T., & Suura, S. R. (2025). Revolutionizing healthcare with generative AI: enhancing patient care, disease research, and early intervention strategies. *American Journal of Psychiatric Rehabilitation*, 28(1), 98-111.
36. KOTHAPALLI, S. L., SEENU, A., DILEEP, V., & YASMEEN, Z. (2025). THE FUTURE OF CUSTOMER ENGAGEMENT IN RETAIL BANKING: EXPLORING THE POTENTIAL OF AUGMENTED REALITY AND IMMERSIVE TECHNOLOGIES. *INTERNATIONAL JOURNAL*, 73(1), 72-79.
37. Paleti, S., Kummari, D. N., Garapati, R. S., Sheelam, G. K., Adusupalli, B., & Singireddy, J. (2025, December). Building a Cyber-Resilient Payment Infrastructure: Transforming Payment Security with Zero Trust Architecture. In *2025 3rd International Conference on IoT, Communication and Automation Technology (ICICAT)* (pp. 1-7). IEEE.
38. Krishnan, M., Aitha, A. R., Amistapuram, K., Nandan, B. P., Kaulwar, P. K., & Singireddy, J. (2025, November). Human-in-the-Loop Hybrid Neuro-Symbolic AI Model for Reliable Data Engineering in High-Stakes Industrial Systems. In *2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN)* (pp. 1-7). IEEE.
39. Challa, S. R., Kaulwar, P. K., Koppolu, H. K. R., Adusupalli, B., & Suura, S. R. (2025, April). Big Data and AI in Wealth Management: Leveraging Cloud Computing for Secure and Scalable Financial Infrastructure. In *International Conference on Smart Computing and Informatics* (pp. 307-318). Cham: Springer Nature Switzerland.
40. Vankayalapati, R. K. (2025). Zero-trust security models for cloud data analytics: Enhancing privacy in distributed systems. Available at SSRN 5121185.



41. Consumer, R. E. P. (2025, May). Artificial Intelligence Infused Real Time Interaction Systems for Next Generation Retail Experiences Personalized Consumer Insights and Augmented Shopping Journeys. In Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24) (p. 208). Springer Nature.
42. Singireddy, S., Pandiri, L., Motamary, S., Kummari, D. N., Somu, B., & Lakkarasu, P. (2025, August). Blockchain-Powered Claims Validation for Enhancing Trust in Health and Auto Insurance Ecosystems. In International Conference on Artificial Intelligence: Theory and Applications (pp. 26-38). Cham: Springer Nature Switzerland.
43. Sriram, H. K., Gadi, A. L., Challa, K., & Singireddy, S. (2025). Leveraging AI, ML, and Gen AI in Automotive and Financial Services: Data-Driven Approaches to Insurance, Payments, Identity Protection, and Sustainable Innovation. Anil Lokesh and Challa, Kishore and singireddy, Sneha, Leveraging AI, ML, and Gen AI in Automotive and Financial Services: Data-Driven Approaches to Insurance, Payments, Identity Protection, and Sustainable Innovation (March 25, 2025).
44. Koti¹, V. M., S. A. J., Vinoth, M., & Bordoloi, D. (2025, May). AI-Driven Gamification Models for Long-Term Educational Engagement. In Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24) (p. 73). Springer Nature.
45. Kannan, S., Raghavendra, C., Kumar, J. S., Narayanasamy, S., Balaram, A., & Aravindh, S. (2025, May). Edge Computing and Artificial Intelligence Powered Dynamic Environmental Conservation Strategies for Sustainable Ecosystem Management and Predictive Climate Analysis. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 477-490). Atlantis Press.
46. Inala, R. (2025). A Unified Framework for Agentic AI and Data Products: Enhancing Cloud, Big Data, and Machine Learning in Supply Chain, Insurance, Retail, and Manufacturing. EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR, 46(1), 1614-1628.
47. Ganti, S., Vankayalapati, R. K., Krishnamoorthy, P., Thakare, P. S., Nayak, U. A., & Vignesh, P. (2025, February). Enhancing IoT-Driven Smart Home Security and Automation with a GCN Model. In 2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS) (pp. 1-6). IEEE.
48. Kummari, D. N., Challa, S. R., Pamisetty, V., Motamary, S., & Meda, R. (2025). Unifying temporal reasoning and agentic machine learning: A framework for proactive fault detection in dynamic, data-intensive environments. Metallurgical and Materials Engineering, 31(4), 552-568.
49. Chava, K., Suura, S. R., Burugulla, J. K. R., Malempati, M., Nuka, S. T., & Komaragiri, V. B. (2025, July). Machine Learning Enabled Financial Statements in Assessing a Business's Performance. In 2025 2nd International Conference on New Frontiers in Communication, Automation, Management and Security (ICCAMS) (pp. 1-6). IEEE.
50. Mashetty, S., & Valiki, S. (2025). AI-Powered Data Engineering for Intelligent Retail Stock Management. International Journal of Research and Applied Innovations, 8(6), 13096-13109.
51. Suura, S. R., Chava, K., Chakilam, C., Nuka, S. T., Maguluri, K. K., & Goma, T. (2025, April). Blockchain-Based Secure and Scalable Models for Healthcare Network Traffic Monitoring and Optimization. In 2025 4th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
52. Dodda, A., Polineni, T. N. S., Yasmeen, Z., Vankayalapati, R. K., & Ganti, V. K. A. T. (2025, January). Inclusive and Transparent Loan Prediction: A Cost-Sensitive Stacking Model for Financial Analytics. In 2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI) (pp. 749-754). IEEE.
53. Kalisetty, S., Lakkarasu, P., Singireddy, S., Burugulla, J. K. R., Challa, K., & Gadi, A. L. (2025, August). Next-Gen Payment Gateways: Leveraging Federated Learning for Fraud Detection in Cross-Border Transactions. In International Conference on Artificial Intelligence: Theory and Applications (pp. 200-211). Cham: Springer Nature Switzerland.
54. Pote¹, X. R., Pamisetty, A., Karthikeyan, G., & Gupta¹, D. (2025, May). Artificial Intelligence Enabled Smart Energy Conservation Systems for Intelligent Resource Management and Sustainable Future Power Grids. In Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24) (p. 196). Springer Nature.
55. Chakilam, C., Kannan, S., Recharla, M., Suura, S. R., & Nuka, S. T. (2025). The impact of big data and cloud computing on genetic testing and reproductive health management. American Journal of Psychiatric Rehabilitation, 28(1), 62-72.
56. Vellela, S. S., Purimetla, N. R., Rao, P. V., Daniel, V. A. A., Koppolu, H. K. R., & Janani, B. (2025). AI-Enabled Wearable Hemodynamic Monitoring System for Early Identification of Thrombotic Events. Vascular and Endovascular Review, 8(16s), 321-336.
57. Paleti, S., Kummari, D. N., Garapati, R. S., Sheelam, G. K., Adusupalli, B., & Singireddy, J. (2025, December). Building a Cyber-Resilient Payment Infrastructure: Transforming Payment Security with Zero Trust Architecture. In 2025 3rd International Conference on IoT, Communication and Automation Technology (ICICAT) (pp. 1-7). IEEE.



58. Mahadevan, S. (2024). Clouding the Future: Innovating Towards Net-Zero Emissions. *International Journal of Computing and Engineering*, 6(2), 17-23.
59. Garapati, R. S., Adusupalli, B., Kaulwar, P. K., Gadi, A. L., Annapareddy, V. N., & Challa, K. (2025, December). The Evolution of Digital Payments: A Study on AI-Powered Transaction Monitoring Systems. In 2025 3rd International Conference on IoT, Communication and Automation Technology (ICICAT) (pp. 1-8). IEEE.
60. Mashetty, S. (2025). LEVERAGING DEEP LEARNING, NEURAL NETWORKS, AND DATA ENGINEERING FOR INTELLIGENT MORTGAGE LOAN VALIDATION. *INTERNATIONAL JOURNAL OF SOCIAL SCIENCE & INTERDISCIPLINARY RESEARCH* ISSN: 2277-3630 Impact factor: 8.036, 14(04), 51-65.
61. Mahadevan, S. (2024). Empowering Manufacturing: Generative AI Revolutionizes ERP Application. *Int. J. Innov. Sci. Res*, 9, 593-595.
62. Kumar, M. V. K., Kannan, S., Annapareddy, V. N., Adusupalli, B., Paleti, S., & Challa, S. R. (2025, July). Transforming Underground Electric Cable Management with AI in Smart Cities. In 2025 2nd International Conference on Computing and Data Science (ICCDs) (pp. 1-6). IEEE.
63. Pandugula, C., Ganti, V. K. A. T., Kalisetty, S., Ande, R., & Mehta, D. (2025). Modeling the Som Kamla Amba Sub-Watershed Using Artificial Neural Networks (ANN) and the SWAT Tool. Available at SSRN 5341614.
64. Annapareddy, V. N., Kollam, M., Challa, K., Vankayalapati, R. K., Malempati, M., & Yasmeen, Z. (2025). Building resilience in water resource management using artificial intelligence and interdisciplinary strategies for sustainable water security and future governance. *Discover Water*, 6(1), 7.
65. Annapareddy, V. N., Challa, K., Komaragiri, V. B., Sriram, H. K., Kalisetty, S., & Goma, T. (2025, April). Explaining AI Techniques such as SHAP LIME and RISE in Limited Sample Size Neuroimaging Studies. In 2025 4th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
66. Ganti, V. K. A. T. (2025). Federated edge computing for privacy-preserving analytics in healthcare and IoT systems. *International Journal of Computer Trends and Technology*.
67. Narasareddy Annapareddy, V., Pamisetty, A., Malempati, M., Kaulwar, P. K., & Bhardwaj Komaragiri, V. (2025, April). Enhancing Solar Power System Efficiency Through AI-Driven Predictive Maintenance and Cloud-Based Infrastructure Stability Solutions. In International Conference on Smart Computing and Informatics (pp. 328-337). Cham: Springer Nature Switzerland.
68. Challa, S. R., Burugulla, J. K. R., Pamisetty, A., Challa, K., & Paleti, S. (2025, April). AI and ML-Powered Cybersecurity Strategies for Cloud Computing: Ensuring Infrastructure Stability in Financial and Retail Sectors. In International Conference on Smart Computing and Informatics (pp. 315-327). Cham: Springer Nature Switzerland.
69. Maguluri, K. K., Ganti, V. K. A. T., Yasmeen, Z., & Pandugula, C. (2025, January). Progressive GAN framework for realistic chest X-ray synthesis and data augmentation. In 2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI) (pp. 755-760). IEEE.
70. Chakraborty, S., Pamisetty, A., Chandana, N., & CS, B. (2025, October). Depth-Wise Temporal Convolutional Networks with Layer Normalization for Waste Food Prediction. In 2025 2nd International Conference on Software, Systems and Information Technology (SSITCON) (pp. 1-6). IEEE.
71. Koppolu, H. K. R., Munnangi, A. S. M., Reddy, M. B., & Ravulapalli, L. T. (2025). Artificial Intelligence-Based Evaluation of ApoA-I's Anti-Inflammatory and Anti-Fibrotic Roles in Experimental Pulmonary Injury. *Vascular and Endovascular Review*, 8(16s), 299-313.
72. Kannan, S. (2025). Transforming community engagement with generative ai: harnessing machine learning and neural networks for hunger alleviation and global food security. *Cuestiones de Fisioterapia*, 54(4), 953-963.
73. Meda, R., Motamary, S., Dodda, A., Kummari, D. N., Kaulwar, P. K., & Gadi, A. L. (2025, August). Edge-AI Integration for Accelerated Decision-Making in Industrial Robotics and Manufacturing Plants. In International Conference on Artificial Intelligence: Theory and Applications (pp. 332-344). Cham: Springer Nature Switzerland.
74. Nandan, B. P., Kumar, M. V. K., Garapati, R. S., Bandi, V. D. V. K., Davuluri, P. N., & Mangalampalli, B. M. (2026, March). AI-Enhanced Semiconductor Yield Optimization Using Hybrid Deep Learning and Edge Data Analytics. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-6). IEEE.
75. Kumar, B. H., Nuka, S. T., Recharla, M., Chakilam, C., Suura, S. R., & Pandugula, C. (2025, July). Addressing Ethical Challenges in AI-Driven Health Predictions. In 2025 2nd International Conference on Computing and Data Science (ICCDs) (pp. 1-6). IEEE.