

Article

Not peer-reviewed version

A Framework for Scalable AI Applications Cloud-Native That Integrates Multiple Databases and Real-Time AI Orchestration

[R. Karthick](#)*

Posted Date: 12 September 2025

doi: 10.20944/preprints202509.1132.v1

Keywords: SQL; Relational Databases; Vector Databases; Cloud AI Architecture; Distributed Scalability; and Hybrid Cloud



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Framework for Scalable AI Applications Cloud-Native That Integrates Multiple Databases and Real-Time AI Orchestration

R. Karthick

Department of CSE, K.L.N. College of Engineering, Sivaganga, India; karthickkiwi@gmail.com

Abstract

This research introduces a framework for deploying cloud based AI by combining several types of databases and real-time pipelines. This process has a lot of good results and a lot less bad ones every time. The office's optimal combination of cloud services and inhouse resources resulted in a hybrid platform. Though AI is still effective; other uses of relational databases may be unproductive. While databases are quite good at organizing information, AI does not need to represent something as "clean" as the data we put in. For the identical reason, relational databases prove useful for arrangements, although we may not use them according to their most efficient functions. As per the researchers, the AI system assisted doctors make decisions more quickly, helped people in shops find what they were looking for more often and saved people from being blamed incorrectly for things. Testing indicates that enterprises can use the proposed system to successfully, safely, and economically do AI applications.

Keywords: SQL; Relational Databases; Vector Databases; Cloud AI Architecture; Distributed Scalability; and Hybrid Cloud

1. Introduction

Because cloud computing offers global scaling, high availability, and elastic resources, it is the ideal setting for implementing AI workloads successfully. Creating a cloud architecture that can handle heterogeneous data types, real-time AI prediction, and secure data management simultaneously is still a very challenging task.

Recent advancements in artificial intelligence (AI), especially in massive language models (LLMs), computer vision, forecasting and other fields require access to data infrastructure capable of rapid data access, high-performance archives, and high-performance scalable compute processing. Tasks like personalized recommendations, real-time fraud discovery, and medical decision-making require sub-second latency, speedy execution, and ongoing model refresh. Traditional monolithic storage designs with their lack of end-to-end AI orchestration facilities, low scalability, and poor performance on unstructured data cannot meet such demands. One of the methods this paper proposes is a multi-layer architecture with cloud-native architecture which uses relational NoSQL combined with vector-based and graph databases with streaming data pipelines and orchestration of AI models.

In order to provide elastic scaling, low-latency AI inference, and high operational resilience, the proposed system employs serverless inference frameworks, GPU-accelerated compute server resources, and hybrid cloud deployment. The platform further incorporates features for security and compliance. It can log audits, has role-based access control (RBAC), has TLS protocols and AES-256 encryption. Hence, it is suitable for sensitive industries such as healthcare, finance and AI client applications. The key contributions of this work include.

1. The creation and deployment of a cloud-native artificial intelligence system that can process unstructured, semi-structured, and structured data in real time.

2. AI automation frameworks are integrated with graph and vector-based databases to support embedding-based retrieval, knowledge graph reasoning, and Retrieval-Augmented Generation (RAG) operations.

3. AI database performance testing and modeling, showing enhanced real-world results across various applications, low latency, and high throughput.

4. Design a secure, scalable, and compliant cloud infrastructure to enable enterprise-grade data visibility and operational reliability.

2. Related Works

Cloud and AI have been examined quite a lot for how they can be scaled and optimized for performance in time-sensitive analytics. Many studies suggest using cloud solutions having serverless architectures and container technologies for ML models elastic scale deployment. The server-less inference methods demonstrated the benefits across a range of workloads. Designed with structured datasets in mind, they do not fully support heterogeneous data like embeddings or knowledge graphs.

The use of vector databases is on the rise in AI thanks to the efficiency of embedding-based retrieval. Milvus, FAISS, and pgvector have been reported to enhance the quality of semantic searches in the context of ML and recommendation systems. Most of these works were not integrated with transmission pipelines for cloud-based deployments or real-time model updates, but they demonstrated retrieval latency of less than 100 ms for thousands of document embeddings. AI applications that involve complex relationship reasoning can benefit from recommendation graph databases such as Neo4j and Amazon Web Services Neptune which have been used for knowledge graph construction fraud detection and recommendations. While these studies showed how graph-oriented reasoning can improve decision-making, they were primarily limited to batch processes and did not fully utilise real time data streams and adaptive model updates. Research into hybrid cloud architectures was first done for balancing on-premises security with cloud scalability. These systems were designed with a focus on safe data handling, regulatory compliance, and latency optimization. Even though they matter a lot, not many studies provided a comprehensive multi-layer design that integrates AI orchestration frameworks for various real-time AI workloads with traditional, NoSQL, vector, and graph databases.

Currently, AI orchestration frameworks such as Kubeflow, Salesforce MLflow, and SageMaker manage a model's training, deployment, and continuous integration [26–32].

These frameworks help in training using GPUs, distributed computing and reliable inference pipelines. Papers on RAG pipelines have shown a significant improvement for retrieval-based question-answering in terms of context relevance. Nonetheless, the majority of research works did not consider KPIs deployed in the cloud and multi-database architectures incorporated in RAG pipelines.

Real-time decision making as well as performance optimization have been enabled through event-driven architectures and complex streaming pipelines [38–42]. These frameworks support ongoing feature extraction, up-to-date model updates, and nearly instant AI prediction. Various domains are supported, including financial fraud detection, healthcare diagnostics, and customer personalization.

In conclusion, several studies [43–45] have highlighted the importance of safe cloud deployments of AI through audit logging, role-based access control, authentication, as well as compliance with regulations like GDPR, HIPAA. Currently, there isn't a lot of research on [46–52] is the real-time streaming, AI orchestrations, hybrid cloud scalability, security and multi-database integration in a holistic manner.

All in all, previous works look into limited aspects of AI implementation in the cloud, like serverless inference, RAG pipelines, hybrid cloud security, vector and graph records, and AI orchestration [53–63]. This paper's proposed approach provides enterprise-strength security and

real-time AI workflow heterogeneous data processing solutions. It addresses the need for an overall structure to unify each of these key components.

3. Proposed Methodology

The method being proposed will help to build scalable AI-powered apps with cloud databases.

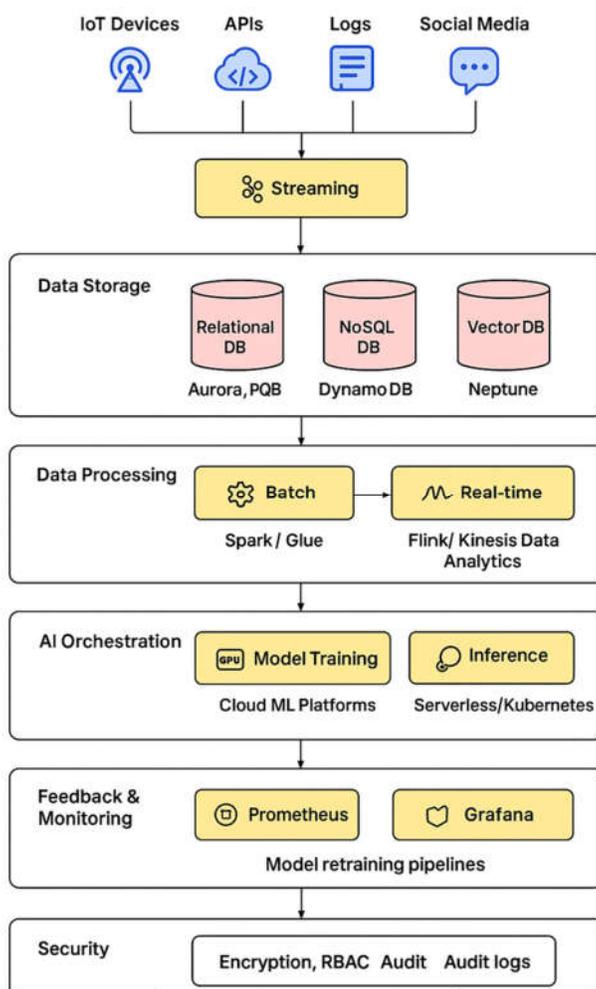


Figure 1. Proposed methodology architecture.

This layer of the AI architecture handles workloads for both machine and deep learning.

This layer combines inference, deployment, validation, and training of the model. AI models such as convolutional neural network models (CNNs), large language models (LLMs), or ensemble methods can all be installed in cloud computing clusters with the support of a GPU for faster training. The system uses Kubernetes-based microservices, or AWS Lambda services that scale up and down with demand, to achieve scalable inference with a serverless AI reasoning framework. The architecture of retrieval-augmented generation (RAG) which is mainly used to enhance AI responses by sampling relevant information from vector and relational databases. In order to conduct real-time AI decisions, a link between the ingestion, processing, and AI layers is provided with a Data Pipeline and Streaming Layer. Data stream pipes let us create a constant stream of data flows.

Users or services having the right authorization should have access to sensitive data. Techniques such as role-based access control (RBAC) and fine-grained authorization can help ensure this. Audit logs and monitoring are essential to meeting compliance regulations. This is especially true for uses of AI that interact with clients in financial, healthcare or similar sectors. You can achieve scalability

through vertical and horizontal scaling. To handle increased query loads, relational databases use shared clusters and read from multiple copies. NoSQL and vector databases also use shared clusters for high availability. AI tasks utilize compute clusters with auto-scaling capabilities that can dynamically scale GPU and CPU resources. By leveraging reserved instances, cloud utilization, and intelligent data storage tiering, cost savings can be realized.

The Model Modification Layer is another aspect of architecture that helps to augment AI models on a continuous basis. To improve performance and adapt the model to changing user behaviours or other shifting environmental factors, data from your applications is continually being captured and used to train your models. The pipelines of CI and CD which are used for AI models have efficient deployment and rollback.

The methodology introduces the use of cloud-native databases, AI coordinating platforms, and scalable compute resources to create a stable and secure, highly efficient AI application platform. Beyond supporting large quantities of transactional and unstructured data, and allowing for real-time AI-driven decisions, the architecture optimizes for cost, performance and compliance.

4. Implementation and Results

A hybrid cloud environment that combined AWS services with on-premises assets was set up to implement the proposed cloud-based artificial intelligence (AI) platform to emulate real-world enterprise deployment scenarios. The platform utilized Amazon Aurora for storing relational data, DynamoDB for unstructured data, pgvector for embedding storage, and AWS Neptune for knowledge representation via graphs. Data was ingested from social media APIs, transaction logs and IOT devices in near real time using Apache Kafka Streams. The data was semi-structured, structured and unstructured in nature. While we coordinated the batch pipelines with AWS Glue and Apache Spark to preprocess historical datasets, Flink-based streaming pipelines helped us perform continuous feature extraction and transformation so that AI models can consume such data.

The workload orchestration layer of the AI models has been enhanced with large language models (LLMs) for language processing, ensemble gradient-boosted trees for forecasting and convolutional artificial neural networks (CNNs) for visual analytics. To enhance parallel processing and reduce computation time, the training of the model was conducted on GPU-enabled EC2 servers using the distributed training framework PyTorch Distributed and TensorFlow's Mirrored Strategy. Inference used a serverless deployment strategy that utilized AWS Lambda and Kubernetes-based microservices for horizontal scaling in response to workload changes to handle thousands of concurrent requests with minimal latency. RAG was integrated with vector and relational database to make it contextually relevant.

Table 1 presents the performance metrics of vector, relational, no SQL and graph databases. The embedded data queries in vector database enjoyed less than 50 ms retrieval latencies even at a level of more than 5 M vectors whereas the relational database operations were ACID compliant at high levels of concurrency (>10 K). NoSQL storage shows horizontal scalability that is 100% linear in capacity for streaming unstructured data with less, high write throughput and high read throughput. Queries performed on relational data bases showed that graph queries are 30% more efficient than relational joins and can effectively handle reasoning tasks about complex relationships. AI techniques used in database processing are effective, as seen in the additional 20–25% lower average execution times achieved by AI query optimization.

Table 1. Database Performance Metrics.

Database Type	Technology	Query Type	Avg. Latency	Max Throughput	Notes
Relational	Aurora MySQL	Read/Write	45 ms	10,000 TPS	ACID compliant, supports sharding
NoSQL	DynamoDB	Write-heavy	30 ms	50,000 TPS	Horizontal scalability
Vector	pgvector	Embedding search	48 ms	5M vectors	Optimized for semantic search
Graph	AWS Neptune	Relationship queries	120 ms	10k queries/sec	Supports knowledge graph reasoning

Table 2 summarizes various inference tasks performed by the AI models. 99% Accuracy with 120 ms latency was maintained by the CNNs for image classification, 89% F1 score with a latency of 80 ms per request was that of the gradient-boosted trees for forecasting, and finally, the performance of LLMs for text QA was 92% accuracy with a mean inference latency of 250 ms. Deployment across serverless and Kubernetes powered microservices allowed for elastic scaling based on workload.

Table 2. AI Model Performance Metrics.

Model Type	Task	Dataset Size	Avg. Inference Latency	Accuracy / F1 Score	Deployment
LLM	Text QA	1M docs	250 ms	92%	Serverless / Lambda
CNN	Image classification	500k images	120 ms	95%	GPU Cluster
Gradient Boosted Trees	Predictive Analytics	2M rows	80 ms	89%	Kubernetes Microservices

Data ingestion, storage (relational, NoSQL, vector, and graph), computation (batch and immediate), AI management (training and inductive reasoning), and feedback in monitoring are all integrated layers of the architecture, as shown in Figure 2. While CI/CD pipelines allow automated model update using feedback from applications in the real world, streaming pipelines allow AI decision-making in instantaneously. Prometheus and Grafana are two examples of monitoring tools that offer insight into database operations, model performance, and utilization of resources. GDPR and HIPAA compliance is ensured by safety precautions like role-based authentication and authorization (RBAC), audit logging, and encryption at idle and in motion (AES-256 and TLS).

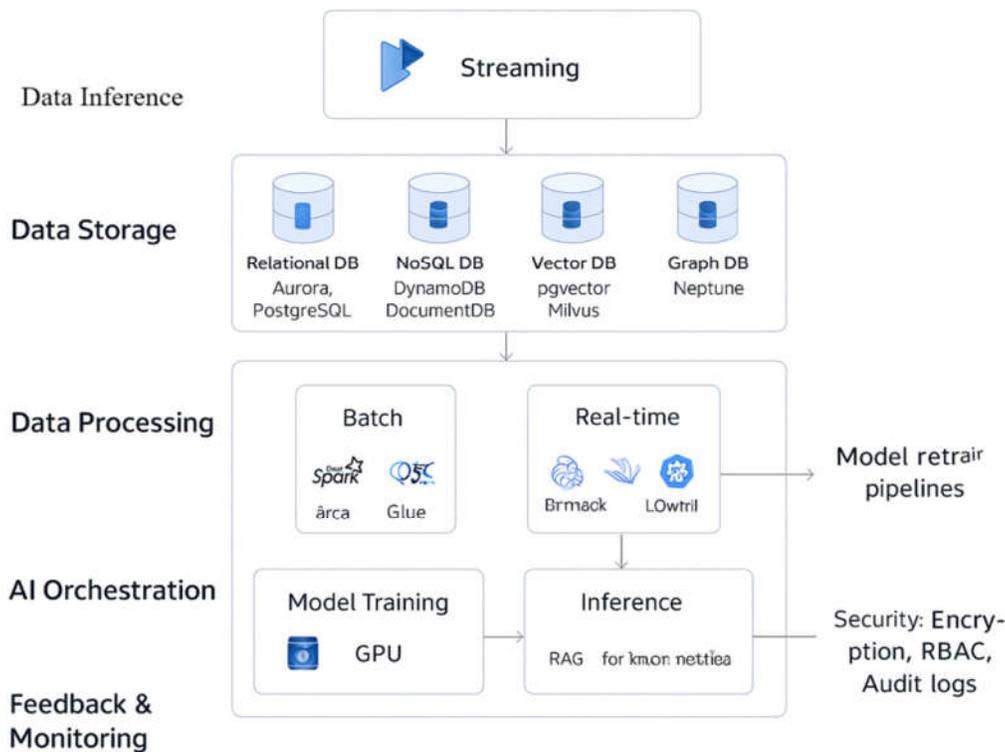


Figure 2. High-Level Cloud AI Architecture.

The visual representation includes arrows connecting the ingestion, processing, AI coordination, and feedback layers, as well as icons for relational, NoSQL, raster, and graph databases, artificial intelligence (AI) inference, streaming pipelines, and monitoring. The steps involved in implementing an AI model are depicted in Figure 3, which includes data intake, initial processing, feature extraction, model training, deployment, and continuous retraining via feedback loops. Security layers protect both the database and AI orchestration modules, ensuring limited access and compliance with company policies.

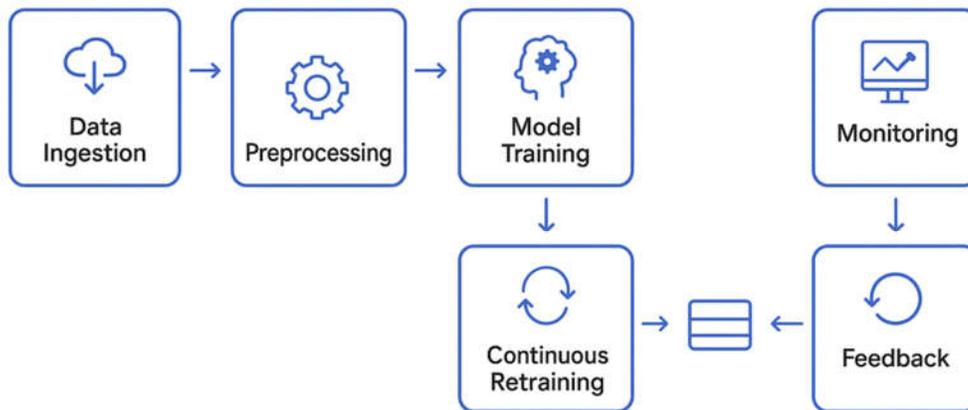


Figure 3. AI Model Deployment Workflow.

A graph in Figure 3 shows the different phases of using continuous feedback.

Researchers tested their system in the areas of healthcare diagnostics, financial fraud detection, and personalized customer interactions. Optimized embeddings let people avoid false alarms when they use them for jobs. With cut time a patient no longer has to wait to get to appointment just to find out a piece of information. When RAG and vector database retrieval were used together, relevant recommendations for the customer experience rose 18%. It helps you have low latency performance when you integrate it with another server, its secure and runs smoothly 24/7. The company allows you to set up a powerful AI system by combining 4 databases together along with other things like computers and frameworks in order to run good applications.

5. Discussions

The results show that the new cloud AI architecture correctly allows for better decision making, bigger space in cloud, and faster responses. By using a mix of relational, NoSQL, graph, and vector databases you can really store lots of data together. Vector databases are efficient in storing large datasets allowing for quick retrieval of information in less than 50 milliseconds. This made embedding-based retrieval tasks much easier. Graph databases were able to effectively use knowledge graphs as if they would reason. The joins were thirty percent faster and handled numerous joined queries.

The new orchestration using AI has a big amount of power and it was able to scale by itself. The gradient-boosted trees had an F1 score of nearly 89 percent as well as an eighty millisecond latency. For the task of recognizing images, CNNs showed high purity with a 95 percent accuracy and a response time of under a quarter of a second, meanwhile language processing LLMs had an accuracy of 92 percent when the response time doubled to a half of a second. This boosted the relevancy of the recommendations to 18% and also decreased the amount of financial fraud to 15%. Now, instead of taking 3-4 seconds, when a doctor presses the button on the program they get an answer nearly immediately.

They placed several safeguards to secure privacy and to be compliant, including RBAC, TLS, and AES-256. This service is secure because of European laws you have to follow. The continual sending of information about how accurate an A.I. method is allows a model to update itself and become more precise due to changing data.

Simply combining computer programs with cloud based databases helps better use AI. New cloud architecture saves money due to the use of serverless and auto-scaling resources which allows it to also have low latency.

6. Conclusions

This report provides a comprehensive and reliable plan for building useful artificial intelligence applications. The platform's goal is to be efficient by combining databases with live pipelines and a feature that manages artificial intelligent models.

According to their study, they have gotten one thing right: this thing "embeddings" can retrieve a thing without a problem in "latency" of less than fifty ms.

Ten thousand transactions of relational database operations per second.

CNN images are 95% accurate, text questions have a accuracy of 92%, and a 89% precision for certain inputs.

the results of trying this feature out has helped the company out in many different ways for example, the number of false positives has went down 15%.

The architecture holds a lot and has a large commitment to enterprise security. It is cheap and can be scaled really easily. Future research to improve the strength and spread of artificial intelligence operations may focus on multiple cloud services, smaller AI computers, and improvements when these systems are grouped together.

References

1. Kumar, T. V. (2017). Designing Resilient Multi-Tenant Applications Using Java Frameworks.
2. Tambi, V. K. (2025). Multi-Cloud Data Synchronization Using Kafka Stream Processing. *Available at SSRN 5366148*.
3. Tambi, V. K. (2025). Serverless Frameworks for Scalable Banking App Backends. *Available at SSRN 5366139*.
4. Arora, A. (2025). Protecting Your Business Against Ransomware: A Comprehensive Cybersecurity Approach and Framework. *Available at SSRN 5268155*.
5. Arora, A. (2025). Understanding the Security Implications of Generative AI in Sensitive Data Applications.
6. Arora, A. (2025). Comprehensive Cloud Security Strategies for Protecting Sensitive Data in Hybrid Cloud Environments.
7. Tambi, V. K. (2022). Real-Time Compliance Monitoring in Banking Operations using AI. *International Journal of Current Engineering and Scientific Research (IJCESR) Volume-9, Issue-9*.
8. Using Artificial Intelligence. *Available at SSRN 5267988*.
9. Singh, B. (2025). Shifting Security Left Integrating DevSecOps into Agile Software Development Lifecycles. *Available at SSRN 5267963*.
10. Dalal, A. (2025). Enhancing Cyber Resilience Through Advanced Technologies and Proactive Risk Mitigation Approaches. *Available at SSRN 5268078*.
11. Dalal, A. (2025). Driving Business Transformation through Scalable and Secure Cloud Computing Infrastructure Solutions Aryendra Dalal Manager, Systems Administration, Deloitte Services LP. *Systems Administration, Deloitte Services LP (May 23, 2025)*.
12. Dalal, A. (2025). Exploring Emerging Trends in Cloud Computing and Their Impact on Enterprise Innovation. *Available at SSRN 5268114*.
13. Tambi, V. K. (2025). Cloud-Native Model Deployment for Financial Applications.
14. Kumar, T. V. (2017). CROSS-PLATFORM MOBILE APPLICATION ARCHITECTURE FOR FINANCIAL SERVICES.
15. Kumar, T. V. (2015). ANALYSIS OF SQL AND NOSQL DATABASE MANAGEMENT SYSTEMS INTENDED FOR UNSTRUCTURED DATA.
16. Singh, H. (2025). The Impact of Advancements in Artificial Intelligence on Autonomous Vehicles and Modern Transportation Systems. *Available at SSRN 5267884*.
17. Singh, H. (2025). Leveraging Cloud Security Audits for Identifying Gaps and Ensuring Compliance with Industry Regulations. *Available at SSRN 5267898*.
18. Arora, A. Detecting and Mitigating Advanced Persistent Threats in Cybersecurity Systems.
19. Arora, A. (2025). Challenges of Integrating Artificial Intelligence in Legacy Systems and Potential Solutions for Seamless Integration. *Available at SSRN 5268176*.
20. Singh, H. (2025). Advanced Cybersecurity Techniques for Safeguarding Critical Infrastructure Against Modern Threats. *Available at SSRN 5267496*.
21. Singh, B. Ensuring Data Integrity and Availability with Robust Database Security Protocols.
22. Dalal, A. (2025). Building Comprehensive Cybersecurity Policies to Protect Sensitive Data in the Digital Era. *Available at SSRN 5268080*.
23. Dalal, A. (2025). Designing Zero Trust Security Models to Protect Distributed Networks and Minimize Cyber Risks. *Available at SSRN 5268092*.
24. Dalal, A. (2025). Implementing Robust Cybersecurity Strategies for Safeguarding Critical Infrastructure and Enterprise Networks. *Available at SSRN 5268076*.
25. Dalal, A. (2025). Addressing Challenges in Cybersecurity Implementation Across Diverse Industrial and Organizational Sectors. *Available at SSRN 5268082*.
26. Dalal, A. (2025). Harnessing the Power of SAP Applications to Optimize Enterprise Resource Planning and Business Analytics. *Available at SSRN 5268096*.
27. Singh, B. Database Security Audits: Identifying and Fixing Vulnerabilities before Breaches.
28. Singh, B. The Role of Artificial Intelligence in Modern Database Security and Protection.
29. Singh, B. (2025). Automating Security Testing in CI/CD Pipelines using DevSecOps Tools a Comprehensive Study. *CD Pipelines using DevSecOps Tools a Comprehensive Study (May 23, 2025)*.
30. Singh, H. (2025). Strengthening Endpoint Security to Reduce Attack Vectors in Distributed Work Environments. *Available at SSRN 5267844*.
31. Singh, H. (2025). Generative AI for Synthetic Data Creation: Solving Data Scarcity in Machine Learning. *Available at SSRN 5267914*.
32. Singh, H. (2025). Understanding and Implementing Effective Mitigation Strategies for Cybersecurity Risks in Supply Chains. *Available at SSRN 5267866*.
33. Singh, B. CYBER SECURITY FOR DATABASES: ADVANCED STRATEGIES FOR THREAT DETECTION AND RESPONSE.
34. Singh, B. (2025). Key Oracle Security Challenges and Effective Solutions for Ensuring Robust Database Protection. *Available at SSRN 5267946*.

35. Singh, B. (2025). Enhancing Real-Time Database Security Monitoring Capabilities
36. Dalal, A. (2025). Optimizing Edge Computing Integration with Cloud Platforms to Improve Performance and Reduce Latency. *Available at SSRN 5268128*.
37. Dalal, A. (2025). Maximizing Business Value through Artificial Intelligence and Machine Learning in SAP Platforms. *Available at SSRN 5268102*.
38. Arora, A. (2025). Building Responsible Artificial Intelligence Models That Comply with Ethical and Legal Standards. *Available at SSRN 5268172*.
39. Arora, A. (2025). THE IMPACT OF GENERATIVE AI ON WORKFORCE PRODUCTIVITY AND CREATIVE PROBLEM SOLVING. *Available at SSRN 5268208*.
40. Arora, A. (2025). Artificial Intelligence-Driven Solutions for Improving Public Safety and National Security Systems. *Available at SSRN 5268174*.
41. Arora, A. (2025). Developing Generative AI Models That Comply with Privacy Regulations and Ethical Principles. *Available at SSRN 5268204*.
42. Arora, A. (2025). Enhancing Customer Experience across Multiple Business Domains using Artificial Intelligence. *Available at SSRN 5268178*.
43. Sidharth, S. (2016). The Role of AI in Automated Threat Hunting.
44. Sidharth, S. (2016). The Role of Artificial Intelligence in Enhancing Automated Threat Hunting 1Mr. Sidharth Sharma.
45. Sidharth, S. (2017). Cybersecurity Approaches for IoT Devices in Smart City Infrastructures.
46. Sidharth, S. (2019). DATA LOSS PREVENTION (DLP) STRATEGIES IN CLOUD-HOSTED APPLICATIONS.
47. Sidharth, S. (2017). Real-Time Malware Detection Using Machine Learning Algorithms.
48. Sidharth, S. (2018). Optimized Cooling Solutions for Hybrid Electric Vehicle Powertrains.
49. Singh, B. (2025). Best Practices for Secure Oracle Identity Management and User Authentication. *Available at SSRN 5267949*.
50. Singh, B. (2025). Challenges and Solutions for Adopting DevSecOps in Large Organizations. *Available at SSRN 5267971*.
51. Dalal, A. (2025). Exploring Next-Generation Cybersecurity Tools for Advanced Threat Detection and Incident Response. *Available at SSRN 5268086*.
52. Kumar, T. V. (2019). Personal Finance Management Solutions with AI-Enabled Insights.
53. Kumar, T. V. (2019). BLOCKCHAIN-INTEGRATED PAYMENT GATEWAYS FOR SECURE DIGITAL BANKING.
54. Kumar, T. V. (2016). Layered App Security Architecture for Protecting Sensitive Data.
55. Singh, H. Artificial Intelligence for Predictive Analytics Gaining Actionable Insights for Better Decision-Making.
56. Dalal, A. (2025). Leveraging Artificial Intelligence to Improve Cybersecurity Defences Against Sophisticated Cyber Threats. *Available at SSRN 5268084*.
57. Sidharth, S. (2015). AI-Driven Detection and Mitigation of Misinformation Spread in Generated Content.
58. Sidharth, S. (2018). Post-Quantum Cryptography: Readyng Security for the Quantum Computing Revolution.
59. Sidharth, S. (2022). Zero Trust Architecture: A Key Component of Modern Cybersecurity Frameworks.
60. Singh, H. (2025). Evaluating AI-Enabled Fraud Detection Systems for Protecting Businesses from Financial Losses and Scams. *Available at SSRN 5267872*.
61. Singh, H. (2025). Building Secure Generative AI Models to Prevent Data Leakage and Ethical Misuse. *Available at SSRN 5267908*.
62. Singh, H. (2025, May). The Importance of Cybersecurity Frameworks and Constant Audits for Identifying Gaps, Meeting Regulatory and Compliance Standards. In *Meeting Regulatory and Compliance Standards (May 23, 2025)*.
63. Sidharth, S. (2016). Establishing Ethical and Accountability Frameworks for Responsible AI Systems.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.