International Journal of Technology, Management and Humanities Volume 12, Issue 1, 2026

Machine Learning for Early Detection of Operational Risks in Complex Organizations

Ajay Chakravarty

CCSIT, Teerthanker Mahaveer University, Moradabad, Uttar Pradesh, India

Abstract

Operational incidents in complex organisations, such as process failures, service outages and control breakdowns, are rarely instantaneous; they are usually preceded by a build-up of abnormal readings across several monitored metrics. This paper evaluates machine learning for detecting that build-up early enough to intervene. A simulation harness generates daily readings of five process metrics for 60 operational processes over 200 days, with incidents preceded by a gradual drift over three to nine days. Three detection approaches are compared on 20 held-out processes: unsupervised anomaly scoring with an isolation forest, and supervised logistic regression and gradient boosting trained to predict an incident within the next five days from current readings, their day-to-day changes and five-day rolling means. Supervised models detect 82% of incidents within ten days of onset with mean lead times of 4.3 and 5.4 days, respectively, at a false-alert rate of about 8.5% of process-days; the isolation forest detects 62% with a longer mean lead but near-random discrimination at the window level. The paper distinguishes detection rate, lead time and false-alert burden as the three quantities that jointly determine operational value, shows that they trade off differently across methods, and proposes an alerting architecture that combines unsupervised screening with supervised confirmation. Guidance is offered on labelling incidents, choosing the prediction window and governing alert fatigue.

Keywords: operational risk; early warning; anomaly detection; isolation forest; gradient boosting; alert management

International Journal of Technology, Management and Humanities (2026) DOI: 10.21590/ijtmh.12.1.20

Introduction

Complex organisations, including banks, utilities, logistics networks and hospitals, run hundreds of interdependent processes whose failure can be costly, disruptive or dangerous. Operational-risk management frameworks require such organisations to identify, monitor and mitigate the risk of loss from inadequate or failed processes, people and systems [1]. Systems-safety theory adds that accidents in complex systems typically emerge from the interaction of many small deviations rather than from a single failure [2], which means that early warning requires monitoring many metrics jointly rather than any one in isolation.

Traditional monitoring relies on static thresholds set per metric. Thresholds are easy to understand but detect problems only when a metric has already left its normal range, and they generate large numbers of alerts when several metrics drift simultaneously. Machine learning offers two alternatives. Unsupervised anomaly detection flags readings that are unusual relative to the historical distribution without requiring labelled incidents [3], [4]. Supervised learning, where incident labels exist, trains a classifier to predict an incident within a forward window from current and recent readings.

This paper compares these approaches on the quantities that determine operational value: how many incidents are detected, how early, and at what cost in false alerts. Because incident data are scarce and sensitive, a simulation harness with an explicit incident-precursor mechanism is used, which allows lead times to be measured exactly. The results motivate an alerting architecture and practical guidance for organisations building early-warning capability.

Corresponding Author: Ajay Chakravarty, CCSIT, Teerthanker Mahaveer University, Moradabad, Uttar Pradesh, India, e-mail: ajay.chakravarty1@gmail.com

How to cite this article: Chakravarty, A. (2026). Machine Learning for Early Detection of Operational Risks in Complex Organizations. *International Journal of Technology, Management and Humanities,* 12(1), 181-186.

Source of support: None

Conflict of interest: None

BACKGROUND

Operational risk and precursors

Operational risk frameworks [1] distinguish key risk indicators, which are leading measures, from loss events, which are lagging. The premise of early detection is that incidents have precursors visible in the indicators. Studies of major accidents consistently find such precursors in hindsight [2]; the difficulty is detecting them in foresight, when they are mixed with normal variation.

Anomaly detection

Anomaly detection assigns each observation a score reflecting its deviation from normal behaviour [3]. The isolation forest [4] does so by measuring how few random partitions are needed to isolate a point; anomalies are isolated quickly. It requires no labels and scales to many metrics, which makes it attractive for operational monitoring. Its weakness is that it scores point-wise unusualness, not incident likelihood: a gradual drift that remains within historical ranges may never look anomalous.

Supervised early warning

When incident logs exist, early warning can be framed as classification: predict whether an incident occurs within the next h days from features computed at the current day. Gradient boosting [5] and logistic regression are standard choices. Feature engineering matters: changes and rolling statistics capture drift that raw levels do not. Evaluation uses discrimination metrics [6], but operational assessment also requires lead time and false-alert rates, which this paper reports explicitly.

METHODOLOGY

Incident-precursor process

The harness simulates $P = 60$ processes over $T = 200$ days, each with five metrics. Metric j of process p has baseline noise scale $b_{p,j} \sim U(0.5, 1.5)$. On any day after day 20, an incident begins with probability 0.015 (with a 15-day refractory period after each incident). An incident at day e is preceded by a precursor ramp of length $d \sim U\{3, \ldots, 9\}$ days during which the metrics drift upward:

$$x_{p,j,t} = b_{p,j}\,\eta_{p,j,t} + b_{p,j}\,(k/d)\,w_{p,j} \quad \text{for } t = e - d + k,\ k = 0, \ldots, d - 1, \tag{1}$$

with η standard normal and $w_{p,j} \sim U(0.5, 1.2)$ a metric-specific drift weight. On the incident day and the day after, the drift reaches 1.5 $b_{p,j}$. The precursor is therefore gradual and spread across metrics, and its early days are indistinguishable from noise on any single metric. The label for supervised learning is whether an incident begins within the next five days; 5.2% of process-days carry the label. The held-out set comprises 20 processes with 39 incidents.

Detection methods

Fifteen features are used by the supervised models: the five raw metrics, their day-to-day differences and their five-day rolling means. The isolation forest uses the five raw metrics only, as is typical for unsupervised deployment. Models are trained on 40 processes and evaluated on the remaining 20 using scikit-learn [7]. For every method, an alert is raised on a process-day when its score exceeds the 90th percentile of scores on the held-out set, so that all methods operate at the same alert volume.

Operational metrics

Three operational quantities are computed in addition to AUC and average precision at the window level. Detection rate is the share of incidents for which at least one alert is raised in the ten days before the incident. Lead time is the number of days between the first such alert and the incident, averaged over detected incidents. False-alert rate is the share of process-days on which an alert is raised but no incident begins within five days.

RESULTS

Window-level discrimination and operational metrics

Table 1 reports the results. At the window level, logistic regression (AUC 0.710) and gradient boosting (0.694) discriminate moderately, while the isolation forest (0.529) is close to random: the gradual precursor rarely produces readings that are unusual in isolation. The operational picture is more nuanced. Both supervised methods detect 82.1% of incidents within ten days of onset, with mean lead times of 4.3 days (logistic regression) and 5.4 days (gradient boosting). The isolation forest detects 61.5% of incidents but with the longest mean lead time, 6.7 days, because its alerts are triggered by random excursions that happen to fall in the pre-incident window rather than by the precursor itself. False-alert rates are similar across methods, between 8.5% and 9.5% of process-days, as expected under the common alert-volume constraint.



Table 1: Detection performance on 20 held-out processes (39 incidents; alert threshold at the 90th percentile of scores)

Method	*AUC (window)*	*Average precision*	*Detection rate*	*Mean lead time (days)*	*False-alert rate*
Isolation forest	0.529	0.053	61.5%	6.7	9.5%
Logistic regression	0.710	0.109	82.1%	4.3	8.5%
Gradient boosting	0.694	0.097	82.1%	5.4	8.7%

Lead-time distributions

Figure 1 shows the distribution of lead times. The supervised methods concentrate alerts in the two-to-five-day range before an incident, which is where the precursor ramp becomes detectable above noise, with a secondary cluster near the start of the ten-day window reflecting earlier, weaker alerts. The isolation forest's lead times are spread almost uniformly, consistent with the interpretation that its detections are largely coincidental. Gradient boosting's longer mean lead relative to logistic regression arises from its use of the rolling-mean features, which accumulate the drift.

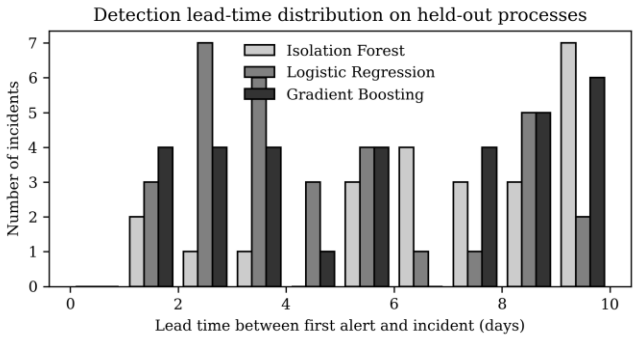


Figure 1: Distribution of lead time between first alert and incident for the three methods on held-out processes.

Trade-offs

The results expose a three-way trade-off. Detection rate and lead time respond to different aspects of the model: detection rate depends on whether the model recognises the precursor at all, lead time on how early it recognises it. False-alert rate is set by the threshold and is the quantity that determines alert fatigue. A method with high detection but short lead may be adequate when interventions are fast, whereas one with longer lead but lower detection may be preferred when interventions take days to mobilise. The choice is therefore operational rather than statistical.

An Alerting Architecture for Early Detection

The findings support a two-stage architecture. The first stage is an unsupervised screen, such as an isolation forest, run across all processes and metrics. It requires no labels, can be deployed immediately and catches gross anomalies that the supervised models may never have seen. Its output should be treated as a candidate list rather than an alert.

The second stage is a supervised confirmation model, trained on the organisation's incident log with drift-sensitive features, that scores the candidates and any process whose recent readings resemble known precursors. Alerts are raised from this stage, with a threshold set on the basis of the acceptable false-alert burden. Because supervised models can only learn incident types present in the log, the first stage provides coverage of novel failure modes while the second provides precision on known ones.

Both stages should feed an alert-management layer that tracks, for every alert, whether it was acted on and whether an incident followed. These records are the training data for the next version of the supervised model and the evidence base for adjusting the threshold.

Practical Guidance

Labelling incidents

Supervised early warning is only as good as the incident log. Organisations should record incident onset dates, not discovery dates, since the gap between the two determines the usable lead time; they should include near-misses, which share precursors with incidents and expand the training set; and they should standardise incident definitions across processes so that models trained centrally transfer.

Choosing the prediction window

The five-day window used here matches the typical precursor length. A window shorter than the precursor

wastes lead time; a window much longer than the precursor dilutes the label with days on which nothing is detectable and lowers precision. The window should be set from the empirical distribution of precursor lengths, which can be estimated by examining metric histories before logged incidents.

Managing alert fatigue

At a false-alert rate near 9% of process-days, an organisation with 500 processes would receive about 45 alerts per day, most of them false. Three mitigations are available: raising the threshold at the cost of detection; escalating only alerts that persist for two consecutive days, which filters random excursions; and routing alerts by severity score so that the highest-scoring cases receive human attention first. The simulation suggests that persistence filtering is particularly damaging to the isolation forest, whose detections are transient, and particularly benign for supervised models, whose scores rise steadily along the precursor ramp.

LIMITATIONS

The precursor mechanism is stylised: all incidents have the same qualitative signature, a monotone ramp across all metrics, which favours supervised models that learn that signature. Real precursors are heterogeneous, and some incidents have none. The incident rate, precursor length and noise scales are design parameters, and the held-out set contains only 39 incidents, so the operational metrics carry substantial sampling variability. All figures are simulation outputs and are not estimates for any organisation.

CONCLUSION

Supervised early-warning models detected 82% of simulated operational incidents with four to five days of lead at a false-alert rate below 10% of process-days, substantially outperforming unsupervised anomaly scoring on detection but not on lead time among detected cases. Because detection rate, lead time and false-alert burden trade off differently across methods, the choice of method is an operational decision that depends on intervention speed and alert capacity. A two-stage architecture that uses unsupervised screening for coverage and supervised confirmation for precision, with disciplined incident labelling and alert management, offers a practical route to early detection in complex organisations.

REFERENCES

[1] Ahuja, R. (2024). Deployment of content security policies in front-end frameworks for blocking cross-site scripting and clickjacking through declarative access restrictions. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 7(11), 17234–17242. https://doi.org/10.15680/IJMRSET.2024.0711132

[2] Chatrath, A. (2018). The impact of containerization on DevSecOps: Security considerations for Docker, Kubernetes, and microservices architectures in continuous deployment pipelines. International Journal of Advanced Research in Education and Technology, 5(6). https://doi.org/10.15680/IJARETY.2018.0506019

[3] Agrawal, N. K., Kumar, S., Bisht, G. S., Srivastav, A., Bansla, V., & Jain, A. (2024, May). Development of a novel quality of service aware video streaming on demand system for the Internet of Thing enabled MANET. In 2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE) (pp. 80–85). IEEE.

[4] Aggarwal, A. (2018). Implementation of cloud security automation in continuous integration pipelines for accelerating DevSecOps adoption and minimizing manual misconfigurations through security-as-code. International Journal of Innovative Research in Science, Engineering and Technology, 7(11). https://doi.org/10.15680/IJIRSET.2018.0711083

[5] Bhaskara Raju Rallabandi. (2024). DIGITAL TWIN–DRIVEN 6G CAMPUS NETWORKS: INTELLIGENT FAULT RECOVERY AND ENERGY OPTIMIZATION. American Journal of AI Cyber Computing Management, 4(3), 30–34. https://doi.org/10.64751/ajaccm.2024.v4.n3.pp30-34

[6] Anchala, S. (2019). Database security challenges in cross-border data flows: Analysing compliance, privacy, and risk in international data transfers. International Journal of Advanced Research in Education and Technology, 6(3), 831–838. https://doi.org/10.15680/IJARETY.2019.0603016

[7] Mahajan, R., Kanabar, K., Dey, R., Jadhav, P. V., Khan, M., & Karimullah, S. (2026). Blockchain complementary model: A new way of devouring smart data by driverless vehicles for next-generation communication. In S. K. Singh, R. B. Jadeja, A. Khanna, P. K. Dutta, & B. Virdee (Eds.), Security paradigms in 6G smart cities and IoT ecosystems: Navigating the future (pp. 115–126). CRC Press. https://doi.org/10.1201/9781003564645-8

[8] Aggarwal, A. (2018). Application of federated learning in medical IoT devices for preserving patient privacy and preventing data breaches through decentralized cyber defense models. International Journal of Innovative Research in Computer and Communication Engineering, 6(3), 2654–2660. https://doi.org/10.15680/IJIRCCE.2018.0603167

[9] Chatrath, A. (2019). A multi-layered DevSecOps framework: Integrating secure coding, automated testing, vulnerability scanning, and incident response into end-to-end software development lifecycles. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 8(3). https://doi.org/10.15662/IJAREEIE.2019.0803102

[10] Rangappa, A. S. (2026). The architecture of agentic AI systems: A technical study on planning, memory, goal-oriented reasoning, and environment interaction for building robust



autonomous agents. Journal of Integrative Science and Societal Impact, 2(1), 39–46. https://doi.org/10.29121/9shxnt36

[11] Mahadasu, S. S. (2021). Telemetry-assisted mobility management for 5G standalone networks with distributed edge user-plane functions. International Journal of Engineering Research and Science & Technology, 17(1), 141–145. https://doi.org/10.62643/ijerst.2019.v17.n1.4406

[12] Devulapalli, N. (2024). Agentic AI in finance and economics: The role of autonomous agents in fraud detection, algorithmic trading, credit risk assessment, and ethical governance of automated financial systems. International Journal of Multidisciplinary and Scientific Emerging Research, 12(4), 1958–1964. https://doi.org/10.15662/IJMSERH.2024.1204057

[13] Bhanushali, B. (2018). Risk-based approaches in anti-money laundering for balancing regulatory compliance costs and maximizing detection of high-risk financial activities. International Journal of Innovative Research in Computer and Communication Engineering, 6(11), 9006–9014. https://doi.org/10.15680/IJIRCCE.2018.0611081

[14] Rangappa, A. S. (2024). Application of AI-augmented log analysis in Dev Sec Ops monitoring for accelerating incident response and root cause analysis through pattern recognition and auto-correlation. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 7(10), 16135–16142. https://doi.org/10.15680/IJMRSET.2024.0710081

[15] Anchala, S. (2018). Adoption of infrastructure as code security validation in cloud-native environments for reducing configuration drifts and compliance violations through automated template scanning. International Journal of Innovative Research in Science, Engineering and Technology, 7(10), 8305–8313. https://doi.org/10.15680/IJIRSET.2018.0710087

[16] Khan, S. (2023). AI-driven speech and voice analytics in CRM platforms: A study on intelligent call center automation, emotion recognition, and service personalization. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 6(2), 510–517. https://doi.org/10.15680/IJMRSET.2023.0602035

[17] Talasila, D. (2024). Implementation of secure DevOps pipelines in cloud-native applications for ensuring code integrity and reducing vulnerability exposure through automated security integration. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 13(11), 2535–2543. https://doi.org/10.15662/IJAREEIE.2024.1311012

[18] Khan, S. (2023). AI-powered innovations in cross-border merchant payment systems: Exploring the role of artificial intelligence in enhancing speed, transparency, and trust in international financial transactions. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 12(5), 1217–1224. https://doi.org/10.15662/IJAREEIE.2023.1205026

[19] Dwivedi, D. (2024). Agentic AI and global governance: Addressing the challenges of regulating autonomous artificial agents across borders in a digitally interconnected world. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 13(11), 2527–2534. https://doi.org/10.15662/IJAREEIE.2024.1311011

[20] Mahajan, R. A., Dey, R., Khan, M., Kanabar, K., & Gupta, S. K. (2025). Real-time vehicular-to-vehicular (V2V) communication and applications for pedestrian safety in AI-optimized VANETs for autonomous vehicles. In H. Maryam, M. M. Malik, I. U. Khan, & S. K. Gupta (Eds.), AI-driven transportation systems: Real-time applications and related technologies (pp. 151–167). Springer. https://doi.org/10.1007/978-3-031-98349-8_8

[21] Kumar, S., Ghai, D., Jain, A., Tripathi, S. L., & Rani, S. (Eds.). (2023). Multimodal biometric and machine learning technologies: Applications for computer vision. John Wiley & Sons.

[22] Gupta, A. (2025). Blockchain-enhanced cloud security: Exploring decentralized trust models, data integrity verification, and secure cloud auditing mechanisms. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 14(1). https://doi.org/10.15662/IJAREEIE.2025.1401019

[23] Jain, A. (2024, December). Energy-efficient signal processing architectures for IoT networks design. In 2024 International Conference on Emerging Technologies and Innovation for Sustainability (EmergIN) (pp. 362–367). IEEE.

[24] Gupta, A. (2026). Bias, fairness, and inclusivity in generative AI systems: A critical examination of algorithmic bias, representation gaps, and the challenges of ensuring equity in AI-generated outputs. Journal of Integrative Science and Societal Impact, 2(1), 23–30. https://doi.org/10.29121/JISSI.v2.i1.2026.38

[25] Bhaskara Raju Rallabandi. (2022). SUB-MICROSECOND 5G SYNC FOR RELIABLE FACTORY ROBOT CONTROL. American Journal of Management and IOT Medical Computing, 1(1), 34–39. https://doi.org/10.64751/ajmimc.2022.v1.n1.pp34-39

[26] Khan, S. (2025). The role of artificial intelligence in cyber security: Leveraging machine learning and predictive analytics for intrusion detection, threat intelligence, and autonomous defense mechanisms. The Research Journal, 11(3), 1–7.

[27] Bhaskara Raju Rallabandi. (2023). NATIVE OBSERVABILITY FOR PRIVATE 5G: A REFERENCE ARCHITECTURE ON COMBINATION OF NWDAF, SMO, AND EDGE ANALYTICS. American Journal of AI Cyber Computing Management, 3(1), 24–28. https://doi.org/10.64751/ajaccm.2023.v3.n1.pp24-28

[28] Talasila, D. (2025). Adoption of dynamic application security testing (DAST) in DevOps lifecycles for identifying runtime vulnerabilities and business logic flaws through continuous fuzzing and penetration testing. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 8(2), 1192–1200. https://doi.org/10.15680/IJMRSET.2025.0802092

[29] Ahuja, R. (2025). The risks and opportunities of generative AI in cyber security: Investigating the dual-use nature of generative models for phishing, malware development, and defence through synthetic data and automated threat detection. International Journal of Innovative Research in Computer and Communication Engineering, 13(2), 1758–1767. https://doi.org/10.15680/IJIRCCE.2025.1302059

[30] Anchala, S. (2018). Deployment of container image scanning in DevSecOps pipelines for mitigating open-source vulnerabilities and ensuring runtime safety through CVE-based detection mechanisms. International Journal of Advanced Research in Electrical, Electronics and

Instrumentation Engineering, 7(3), 1546–1554. https://doi.org/10.15662/IJAREEIE.2018.0703058

[31] Ahuja, R. (2026). Adoption of secure session management in multi-device applications for mitigating session hijacking and replay attacks through token rotation and expiry controls. Journal of Integrative Science and Societal Impact, 2(1), 101–108. https://doi.org/10.29121/JISSI.v2.i1.2026.49

[32] Mahadasu, S. S. (2020). Cloud-native life cycle management of virtualized RAN functions using Kubernetes operators and policy-driven automation. International Journal of Engineering Science and Advanced Technology, 20(4), 21–26. https://doi.org/10.64771/ijesat.2020.v20.i04.pp3846

[33] Dwivedi, D. (2026). Agentic AI and the future of artificial general intelligence: A study on how goal-oriented autonomous agents contribute to the development of more human-like intelligence, reasoning, and decision-making abilities. Journal of Integrative Science and Societal Impact, 2(1), 72–79. https://doi.org/10.29121/JISSI.v2.i1.2026.44

[34] Bhanushali, B. (2019). Graph-based network analysis for identifying complex money laundering typologies, detecting suspicious transaction chains, and tracing illicit financial flows. International Journal of Multidisciplinary and Scientific Emerging Research, 7(4), 2007–2014. https://doi.org/10.15662/IJMSERH.2019.0704047

[35] Devulapalli, N. (2026). A comparative study of predictive AI and agentic AI: Examining key differences in autonomy, goal orientation, human interaction, and system accountability. Journal of Integrative Science and Societal Impact, 2(1), 80–87. https://doi.org/10.29121/JISSI.v2.i1.2026.45

[36] Jain, J., Jain, A., Srivastava, S. K., Verma, C., Raboaca, M. S., & Illés, Z. (2022). Improved security of e-healthcare images using hybridized robust zero-watermarking and hyper-chaotic system along with RSA. Mathematics, 10(7), 1071.

[37] Kanabar, K., Dey, R., Mahajan, R., Jadhav, P. V., Khan, M., & Ghantasala, G. S. P. (2026). The benefits of web and blockchain in decision-making for Industry 4.0. In S. K. Singh, R. B. Jadeja, A. Khanna, P. K. Dutta, & B. Virdee (Eds.), Security paradigms in 6G smart cities and IoT ecosystems: Navigating the future (pp. 289–304). CRC Press. https://doi.org/10.1201/9781003564645-18

[38] Talasila, D. (2026). Agentic AI in multi-agent systems: Exploring the coordination, negotiation, and cooperation of autonomous artificial agents in competitive and collaborative digital ecosystems. Journal of Integrative Science and Societal Impact, 2(1), 64–71. https://doi.org/10.29121/JISSI.v2.i1.2026.43

[39] Dwivedi, D. (2025). A comparative study of public, private, and hybrid cloud security architectures: Evaluating trade-offs in scalability, privacy, cost, and risk management. International Journal of Multidisciplinary Research in Science, Engineering and Technology, 8(1), 485–493. https://doi.org/10.15680/IJMRSET.2025.0801069

[40] Chatrath, A. (2019). DevSecOps in Agile and Lean development: Examining the alignment of security, speed, and scalability in rapid iteration and continuous delivery environments. International Journal of Multidisciplinary and Scientific Emerging Research, 7(3), 1506–1513. https://doi.org/10.15662/IJMSERH.2019.0703046

[41] Rangappa, A. S. (2025). Agentic AI and cyber security: Autonomous threat hunting, intrusion detection, and adaptive defense mechanisms in a world of increasingly sophisticated cyber attacks. Journal of Digital Security and Forensics, 2(1), 128–138. https://doi.org/10.29121/digisecforensics.v2.i1.2025.86

[42] Mahadasu, S. S. (2019). Performance-aware VNF placement for distributed 5G core networks over heterogeneous edge and cloud infrastructure. International Journal of Engineering Science and Advanced Technology, 19(10), 5–11. https://doi.org/10.64771/ijesat.2019.v19.i10.pp3849

[43] Jaiswal, I. A., & Khan, S. (2025). Leveraging cloud-based projects (AWS) for microservices architecture. Universal Research Reports, 12(1), 195–202.

[44] Bhanushali, B. (2019). Integration of biometric authentication with AML systems for strengthening know-your-customer (KYC) processes and reducing identity fraud in financial institutions. International Journal of Advanced Research in Education and Technology, 6(5), 853–860. https://doi.org/10.15680/IJARETY.2019.0605019

[45] Gupta, A. (2025). Agentic AI for national security and defense: An analytical study on autonomous agents in surveillance, strategic decision-making, and military operations with a focus on ethical and legal boundaries. Journal of Digital Security and Forensics, 2(1), 163–175. https://doi.org/10.29121/digisecforensics.v2.i1.2025.84

[46] Burnout in Global Healthcare Systems: Psychosocial Determinants, Regional Realities, and Behavioral Optimization Pathways. [Author/source details missing in supplied list]

[47] Devulapalli, N. (2025). Machine learning in database security: Leveraging anomaly detection, intrusion detection, and predictive analytics for proactive threat prevention. International Journal of Innovative Research in Science, Engineering and Technology, 14(2), 1685–1692. https://doi.org/10.15680/IJIRSET.2025.1402105

[48] Mahadasu, S. S. (2022). GitOps-based continuous delivery of cloud-native 5G network functions with automated configuration validation. International Journal of Data Science and IoT Management System, 1(1), 50–56. https://doi.org/10.64751/ijdim.2022.v1.n1.1286

[49] Bhaskara Raju Rallabandi. (2022). DETERMINISTIC STABILITY ANALYSIS FOR PRODUCTION-LINE CONTROL OVER PRIVATE 5G. American Journal of AI Cyber Computing Management, 2(4), 33–37. https://doi.org/10.64751/ajaccm.2022.v2.n4.pp33-37

[50] Aggarwal, A. (2018). Implementation of quantum cryptography in financial transactions for enhancing data confidentiality and resisting future quantum attacks through secure key distribution mechanisms. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 7(7), 3196–3204. https://doi.org/10.15662/IJAREEIE.2018.0707016

[51] Kumar, S., Sharma, K., Kumar, P. A., Jain, A., Bhagat, S. K., & Singh, P. (2024, September). An improved particle swarm approach for energy-aware location-aided routing in mobile ad-hoc network. In 2024 7th International Conference on Contemporary Computing and Informatics (IC3I) (Vol. 7, pp. 1119–1124). IEEE.

