International Journal of Mathematical Theory and Computer Science Issues
https://academicajournal.com/IJMTCSI
Volume: 3 | Number: 1 (2026) January, 21-44

Artificial Intelligence for Secure Cloud-Native Critical Infrastructure: Improving National Cyber Resilience, Operational Reliability, and Economic Security

Jawad Yaqoob Mir[1]*, Fawad Mir[2]
[1,2]Master of Science in Information Technology, Washington University of Science and Technology (WUST), Alexandria, Virginia, USA
E-mail: jawadmir488@gmail.com

*Correspondence : Jawad Yaqoob Mir
jawadmir488@gmail.com

Received: 05-06-2026
Accepted: 15-07-2026
Published: 10-08-2026



Open Access

Copyright: © 2024 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0/) .

Abstract : Critical infrastructure increasingly depends on cloud-native technologies, including containers, microservices, Kubernetes, application programming interfaces, distributed cloud platforms, and automated software-delivery pipelines. Although these technologies improve scalability and service availability, their dynamic architecture, machine identities, third-party components, and complex dependencies create cybersecurity and operational risks that can cascade across essential national services. Existing approaches often address cloud-native security, AI-based threat detection, operational reliability, business continuity, economic-risk assessment, and governance separately. This fragmentation leaves a gap for a unified, service-centered framework that connects technical security events with mission-critical dependencies, operational consequences, human decision authority, and wider economic impacts. This study proposes the Artificial Intelligence-Enabled Cloud-Native Critical Infrastructure Resilience Framework (AI-CNCI-RF) as a conceptual design-science artifact. The framework integrates cybersecurity governance, critical-service and dependency intelligence, secure cloud-native engineering, zero-trust access control, trusted telemetry, AI-assisted detection and prediction, controlled incident response, trusted recovery, and economic-risk assessment. Its novelty lies in combining technical cybersecurity evidence with service criticality, cascading dependencies, data quality, operational reliability, economic consequences, and risk-controlled automation. AI is positioned as an explainable decision-support capability, while high-impact actions remain subject to policy constraints, safety gates, and human authorization. Theoretically, AI-CNCI-RF unifies previously fragmented domains of cloud-native security and resilience. Practically, it provides a phased implementation and governance roadmap. Operationally, it supports risk prioritization, earlier detection, reversible response, service continuity, and verified recovery. Because the framework is conceptual and not yet empirically validated, future work should test its variables, thresholds, indicators, and effectiveness through prototypes, simulations, cyber exercises, controlled experiments, and sector-specific deployments.

Keywords: artificial intelligence, cloud-native security, critical infrastructure, cyber resilience, operational reliability

Introduction

Critical infrastructure are the systems, services, networks and assets which enable basic functioning of a country. Typically, these include electricity, telecommunications, transportation, banking, healthcare, water supply, food distribution, emergency and government platforms. Such services are intertwined. When a telecom failure happens, they

may be unable to use banking, emergency communication, transport and healthcare. Delayed power outages could impact nearly all aspects of a nation's existence. Resilience of critical infrastructure should, therefore, not be taken as a simple technical matter but rather as a matter of national concern. Digital transformation of critical infrastructure has led to growth of cloud-native computing. Cloud-native systems typically break large applications into smaller services which can be developed, deployed and scaled independently. Containers bundle up applications and their associated software, and orchestration platforms like Kubernetes automate application deployments, communication, scaling, and recovery. Distributed cloud platforms enable organizations to operate services across multiple regions or data centers, while automated software pipelines enable fast distribution of updates. The level of challenge escalates when cloud-native applications enable critical National functions. The breach of security at a regular site could cause the inconvenience of customers. The same error in a health care system, national payment system, energy-management service, or emergency communication network can be life-threatening, cause a disruption in the economy and lead to loss of public trust. Cloud-native critical infrastructure consequently needs to be designed to not only protect against attacks but also to keep the system running when faced with an attack, and then to safely recover from the attack.

AI can be a potential solution to controlling the magnitude and pace of this space. AI systems can interpret identity, application logs, network flows, container activity, system calls, software modification, performance metrics, threat intelligence. They are capable of identifying relationships that might be hard for a human analyst to spot in a timely method. AI can also assist operational teams to anticipate equipment failure and application failure before they can happen. It can detect unexpected service communications, unusual resource usage, latency, erratic workloads or unexpected user activity. A legitimate activity might land in the block list of an AI model or a new or unseen cyber attack might not be classified as such. Also it is susceptible to having limited or low quality or biased or manipulated training data. Furthermore, complex AI models can generate decisions without offering any explanation, leading to a lack of understanding or verification by security teams. In certain cases, these mistakes can truly be dire, potentially interrupting suitable communications or even causing essential systems to automatically conclude. To ensure the safe and reliable operation of AI, it's crucial to incorporate it within a well-governed cybersecurity framework, complete with clearly defined boundaries, human oversight, persistent surveillance and monitoring, and recovery protocols. This paper attempts to present something of a blueprint for mitigating both dangers and opportunities of cloud security with the help of AI. This proposed AI-CNCI-RF framework encompasses technical cyber security, operational continuity, national resilience, and economic security. It does not presume to supplant human professionals with AI. Rather, it specifies AI as a systematic decisions-as-a-service functionality working with particular guidelines, validated information, secured cloud-based infrastructure and human oversight.

A cloud-native application is one that leverages cloud distributed computing resources efficiently. Can be composed of many microservices or a single program. Every service is responsible for a smaller service and communicates with other services using an

API. These services are placed within containers with consistent environments and managed through workload placement, scaling, health checking, networking and replacement via Kubernetes or some other orchestrator. Kubernetes is distinctively significant because it regulates a considerable lot of workloads and correspondence connections. Depending on what permissions the attacker was granted on Kubernetes, the attacker could deploy malicious workloads, acquire sensitive data, modify settings, or hop around services. Protection thus relies on vulnerability scanning, least privilege access, separation of networks, robust authentication, policy enforcement and full audit logging. Cybersecurity is sometimes defined as the act of preventing unauthorized access, attack, or damage to systems. Cyber resilience is broader. It is the capacity to plan for disruption, hold off visitors, keep key functions going, get back secure, and enhance following an attack. Part of being a cyber-resilient organization is recognizing that you will not stop every attack. Thus, it acquires a skill to minimize the effect of successful attacks.

Resilience for critical infrastructure is more than simply the return to a computer system. The essential services range of the organization should be conserved at an acceptable level. A hospital is allowed to offer emergency services even if some digital services are not available. During an electricity operator's investigation of suspicious activity, the electricity operator can lower electricity usage but continue it. Government services might refer citizens to an alternate platform. Resilience, therefore, is not only the technical recovery and the continuity of the operations, but also emergency communication, trusted decision making and coordination with dependent organizations. The continual ability of a system to perform its intended function is operational reliability. Cybersecurity and reliability are quite closely tied since both attack or failure can impact availability or the integrity of the data. While this could be caused by a malicious actor, an incorrect update or certificate expiry, overloaded systems or configuration errors could also cause the same effect.

AI can be used to help with reliability as it can detect signs of failure before actual failure occurs. It could be sensing containers are becoming slow to serve applications, odd memory usage, frequent container restart, faulty network or odd changes in demand [1]. The alerts give engineers time step up. However, over delegating to AI can negatively affect reliability. An automated security application might quarantine a legitimate service, prevent an authorized operator from accessing a service, or prevent critical communications. The framework should therefore assess the security upside as well as the cost of services of each automated response. The uninterrupted function of critical infrastructure is critical for economic security. Electricity, communication, transport, finance and digital government are needed by businesses. The repercussions of such unavailability of one critical service can be cascading throughout supply chains and regions. The economic impact of a cyber incident can include loss of production, disruption of transactions, disruption in transportation, recovery costs, legal costs, loss of customer confidence and harm to suppliers.

Identity is one of the major security perimeters in cloud. Identities are used by human administrators, ordinary users, containers, services, actors of automated deployment systems and software agents. Permissions are given to each identity. An attacker that is able

to steal a privileged account or token can masquerade as an authorized user. Images can be reused to become containers. The issue could impact on all workloads derived from the image, if it is going to impact on individual workload, if the image is containing the malicious code or a vulnerable software package. Containers can also be granted more access than required, which can lead to them accessing other workloads, secrets or host systems. Being cloud-native, they rely on numerous external packages, libraries, tools and images. A software piece can go from a source code repository, build platform, testing service, registry and deployment pipeline to the production environment. In any phase, compromising can result in malicious code being introduced.

Application interfaces enable microservices to communicate with each other. Processing a very large number of internal/external requests is possible on a large cloud-native platform. Attackers can hope for weak authentication, insecure input handling, too much data exposure, lack of rate limits or lack of protection for administrative functions [2]. The operation of AI-driven cybersecurity relies on information gathered across the user, the device, app, cloud platform, and operational devices. This data can include extremely confidential information regarding infrastructure construction, business operations, employees, customers and security controls. Inadequate data security can raise issues of privacy, commerce and/or national security.

Even the AI model used could be targeted. The attackers might try to contaminate training data, modify feedback, replace model files, adjust thresholds for model decision, steal sensitive data of the model, or induce behavior to evade attacks. The language-model components might be subjected to prompt-injection or they can provide an unsupported explanation. It is therefore essential that AI components are protected using access control, monitoring, integrity checking, secure development and controlled change management. There are interdependencies between critical services. Telecommunications need electricity, financial services need telecommunications, healthcare and transport systems depend on a number of external services, and energy is also dependent on telecommunications. Various identity platforms might serve the needs of both cloud providers and organizations. One incident can thus lead to cascading failure. Crippling one shared platform can not only disrupt any number of operators, but technical incidents can lead to operational issues, public concern and economic losses.

Methods

This study is orientated towards conceptual design-science. Design science is appropriate as the purpose of the study is to develop a "structure" to solve a "problem". Therefore, the major research output is not necessarily the training of one particular AI model and the experimental testing of the same model, but the framework AI-CNCI-RF. The process to develop the design started with identifying the three objectives of the system to be designed: national cyber resilience, operational reliability and economic security. In the second step, cloud-native assets and threat conditions that impact such outcomes have been identified. The third step identified the governance, security, data, AI, response and recovery capabilities needed to mitigate risk. The fourth step refined the formulas in such a way that organisations can quantify risk, access level of confidence,

anomalies, availability, resilience and impact on the economy. The last stage developed a roadmap for implementation for phase-in.

The Artificial Intelligence-Enabled Cloud-Native Critical Infrastructure Resilience Framework ties together organisational governance and technical protection, intelligent monitoring, controlled response, trusted recovery and national coordination. It's based on the premise that effective cyber resilience is reliant on human, policy, technology, data, operational planning, and economic decision-making.

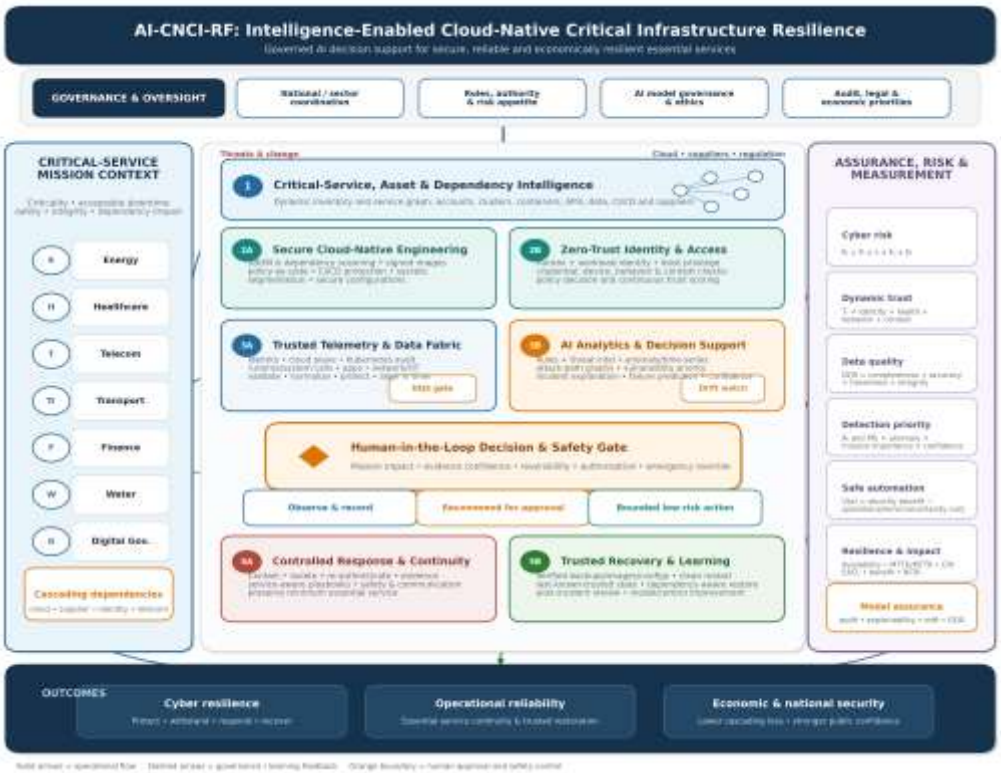


Figure 1. Proposed AI-CNCI-RF Architecture

Governance sets who has responsibility, who has authority to take decisions, and the risk parameters. Cybersecurity governance should be linked to organizational goals, legal liability, the threat from the supply chain and the oversight of top management. Assign responsibility to executive leaders, cybersecurity managers, cloud engineers, critical-service owners, and data specialists and operational personnel, and to business-continuity teams. These groups should be involved together in common governance process, which ensures that tech security decision reflects with operational needs [3]. Use of AI models should be agreed by an AI and cybersecurity governance board, which specifies the objective and allowable use of every model. The board's guidelines should specify what kind of information can be used, how the performance will be assessed, what activities can be performed automatically, and what activities must be manually approved [4]. Additionally, it should consider big changes in models, incidents, failures and ethics issues.

Table 1. Core Components of the Proposed AI-CNCI-RF

Framework component	Primary purpose	Key mechanisms and controls	AI contribution and human control
National and organizational governance	Establish accountability, authority, risk appetite, and acceptable AI use.	Governance board; legal and ethical requirements; model approval; audit; automation limits; emergency override.	AI supports risk analysis and prioritization. Human authorities approve use cases, thresholds, model changes, and high-impact actions.
Critical-service, asset, and dependency intelligence	Identify essential services and map the technologies and external entities supporting them.	Service inventory; dynamic asset discovery; dependency graph; supplier mapping; acceptable downtime and integrity targets.	AI identifies hidden relationships and cascading paths. Service owners validate criticality, dependencies, and mission consequences.
Secure cloud-native engineering	Reduce vulnerabilities throughout the software and infrastructure lifecycle.	Secure coding; SBOM; dependency and image scanning; image signing; secrets management; policy-as-code; CI/CD protection.	AI prioritizes vulnerabilities and flags abnormal modifications. All AI-generated remediation must undergo normal review and testing.
Zero-trust identity and access control	Continuously verify human and machine access requests.	Strong authentication; workload identity; least privilege; device and workload health; context-aware policies; dynamic trust scoring.	AI supplies behavioral and contextual evidence. Policy rules and authorized personnel retain control of privileged or exceptional access.
Trusted telemetry and data management	Provide accurate, timely, protected, and usable data for security decisions.	Multi-source collection; time alignment; validation; normalization; integrity protection; privacy; retention and access controls.	AI processes identity, cloud, Kubernetes, runtime, application, network, and operational data. Human governance determines acceptable data use.
AI detection,	Detect attacks and failures,	Rules; threat intelligence; anomaly	AI detects, predicts, correlates, and

Framework component	Primary purpose	Key mechanisms and controls	AI contribution and human control
prediction, and decision support	prioritize events, and explain recommended actions.	detection; graph analytics; time-series forecasting; alert correlation; explainability; drift monitoring.	recommends. Analysts investigate uncertain outputs and approve recommendations with operational impact.
Controlled response and service continuity	Limit damage without unnecessarily interrupting essential services.	Service-aware playbooks; reversible actions; evidence preservation; re-authentication; enhanced monitoring; bounded isolation.	AI may observe, recommend, or execute approved low-risk actions. Humans authorize high-impact actions and confirm safety and mission effects.
Trusted recovery and resilience measurement	Restore verified services and evaluate security, reliability, and economic outcomes.	Secure backups; validated images, configurations, identities, and data; dependency-aware restoration; post-incident review.	AI assists trusted-state identification and continuing-compromise detection. Humans validate restoration and approve reconnection.

Organizations that don't know their systems can't be protected. A continuous knowledge of critical service, enablers, cloud account, containers, clusters, interfaces, datastores, identities and software pipelines, suppliers and operational dependency is required in the framework. The starting point for the organization should be the service to be provided, and not the technical asset. It should be used to identify the public or business service that is essential, the permissible downtime, the desired level of data integrity and the systems providing support to it. This information relates technical cyber security and operational impact. The dynamic dependency map should demonstrate how a larger application is dependent on each workload and how each application is dependent on an essential service. It should also depict external entities like cloud services, software vendors, Telecoms and identity providers.

$$R_i = P_i \times I_i \times E_i \times D_i$$

R_i here stands for the overall Cyber risk score. P_i is the probability of compromise that we presume could happen, I_i is the impact that the expected operational activity would have, E_i is the level of exposure, and D_i is the importance of dependencies and potential cascading impact. The formula makes sure that a technical weakness is measured not just by its normal severity. A moderate weakness may need to be more

clearly addressed if related to a highly connected national service versus a severe weakness in an isolated test environment. Starting with security during design and development phase, rather than after deployment. That means it calls for secure techniques on source code, software components, container images, infrastructure localization, software build systems, software registries, and deployment pipelines [5]. For source code, scrutiny and testing of vulnerability should be performed. Do not store sensitive credentials in code. To identify software dependencies, use software bill of material and perform software vulnerabilities check. All container images should be derived from approved sources, scanned prior to deployment, signed by authorised systems and verified by the orchestration platform.

AI can help in secure engineering by prioritising vulnerabilities and alerting for any atypical functional modification on software [6]. It can put in a technical seriousness and information about exposure, activity in the workload and importance of services. It can provide suggestions on secure coding or secure testing, but AI-generated modifications need to undergo the same review and testing process as human-generated modifications. Zero trust implies that all accesses have to be vetted or checked before being improperly approved. The decision should take into account: identity, device condition, workload condition, behaviour, location, requested resource, operation context. Hence the framework is based on zero-trust for all human and machine identities. An administrator might need to use robust authentication, a trusted computer system and sometimes elevated privileges. A workload can need verified app Identity, a signed container image, and can only talk to the services it is allowed access to [7]. By identifying behavior changes, AI can help facilitate ongoing trust risks assessments. The dynamic trust score for a request r is represented as follows:

$$T_r = w^I I_r + w^H H_r + w^B B_r + w^c C_r + w^s S_r$$

The T_r in this formula is the combined trust score. I_r is the confidence in identity, H_r is confidence in health of a device/workload, B_r is confidence that behaviour appears normal, C_r is confidence that the context appears normal, and S_r is confidence that the request is appropriate based on the sensitivity of the resource. The factors are represented by the terms w's which define their importance. The final decision should be a policy decision. AI providing behaviour evidence, but rules to stop beyond approved access.

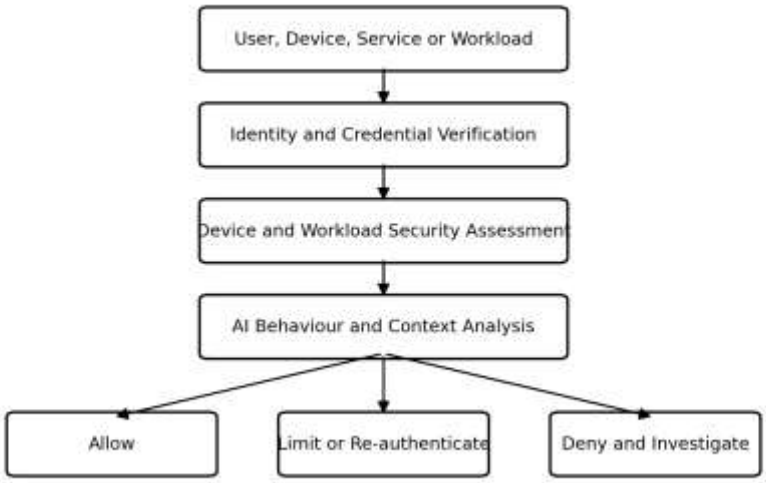


Figure 2. AI-Assisted Zero-Trust Access Process

For cyber security to be effective with AI, it's crucial that visibility is reliable. Whenever possible within its framework, it gathers information from identity systems, cloud control planes, kubernetes audit logs, container runtime tools, backup platforms, vulnerability scanners, software pipelines, service meshes, network devices and operational systems etc.

$$DQS = w^{c}C + w^{A}A + w^{T}T + w^{I}I$$

In this case, DQS is the data-quality score. C is the completeness, A is the accuracy, T is the timeliness and I is the integrity. The weights are assigned based on how important the different quality dimensions are, for the purpose of the intended model. Wherever sources of important information are incomplete or unreliable, it is prudent not to use a model for important decisions.

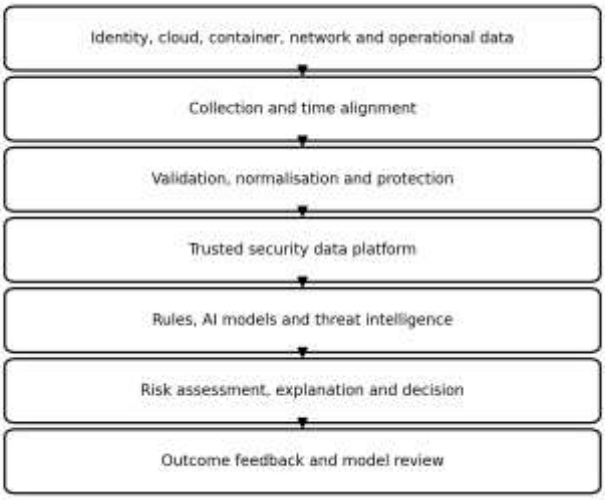


Figure 3. Security Data-to-Decision Flow

The framework constitutes a synthesis of various analyses. Known events like an unsigned production image, access by a blocked identity, and access from a privileged container that was disabled, are all important events that need to be detected using rule-based detection. Rules are simple but they don't necessarily catch all new threats. Relationships can be analysed between multiple identities, workloads, services, data and vulnerabilities with graph based models. It is possible that they can discover an unexpected attack path, or an abnormal relationship between services. Models for time series can forecast workload demand changes, latency, storage utilization, and/or rate of errors [8]. A language model can be used to summarise incidents, to organise the evidence, explain recommended actions in simple vocabulary.

$$A_t = \alpha E_t + \beta F_t + \gamma G_t + \delta Q_t$$

A_t in this formula is the Overall Anomaly Score. E_t for Behaviour-model error, F_t for Deviation from expected performance, G_t for Graph-based abnormality and Q_t for Evidence from security rules or threat intelligence. The relative importance of each source is indicated by the Greek terms.

$$PR_t = A_t \times M_t \times C_t$$

PR_t = event-priority score, A_t = anomaly score, M_t = importance of the service or mission affected, C_t = confidence in the evidence available. This decreases the risk of analysts spending a significant amount of time on low impact and/or unusual events. It focuses on events that are suspicious, and important to essential operation(s). Provide understandable reasons with high-impact AI recommendations. The system needs to provide information on which signals were used for the alert, which services might be impacted, as well as confidence the model felt in the alert and also what situation the model could not address. It should also explain what may happen if the recommended action is or isn't taken [9]. Model changes should be considered wherever there are any changes to the model. Model drift must be monitored, as the application behaviour is not constant and the techniques change over time. If the data quality of a model is not reliable; if there is an increase in the error rate of a model; if significant changes have occurred in the operating environment of a model; or if the breadth and scope of model-based decisions is difficult to explain for the level of risk [10].

Rapid response will minimize damage, but a wrong response will disrupt critical services. As a result, theframework allows several degrees of automation. The lowest level is automated observation and alerting of events by AI. At the next level it suggests actions for man to approve. Bounded Automation enables it to make pre-programmed, reversible, and low-risk moves. Allows for greater speed of automation but only when tried and trusted situations not to delay would cause serious harm.

$$U(a) = B^s(a) - C^O(a) - C^f(a) - C^U(a)$$

U(a) is the anticipated value of the usefulness of the action. $B^s(a)$ denotes its expected security benefit, $C^O(a)$ denotes operational cost, $C^f(a)$ denotes the cost of taking an incorrect action and $C^U(a)$ is uncertainty.

$U(a) > theta$

The organization's lowest level of approved automatic action is gamma. Any action must be automated only if it is deemed to be beneficial above the acceptable safety limit and if it is reversible, tested and backed up by humans.

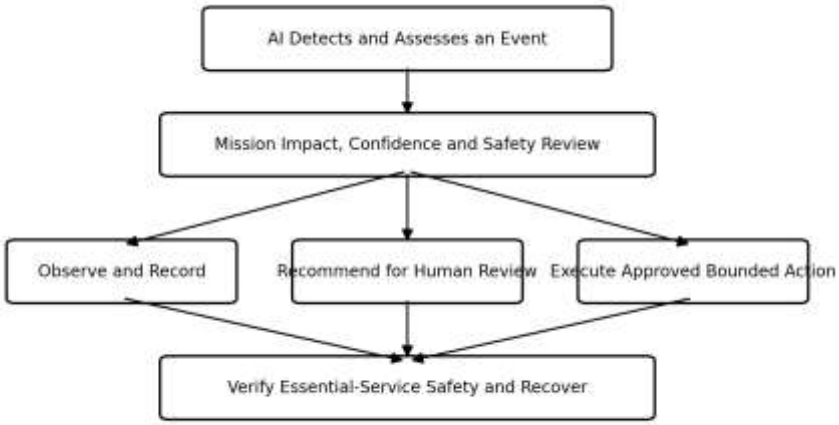


Figure 4. Risk-Controlled Response Process

Recovery should not just get the system up and running fast, but should rebuild an accepted service. Check backups, software images, configurations, identities, and data prior to reconnection. The organization should have secure backups, operating environments available for tests, software objects verified and restart procedures established for each service, taking into account its dependency on other services. AI can help with last-known-trusted state identification, comparing restored configurations with approved configurations, and identifying potential indicators for ongoing compromise [11].

Results

Availability is "the time percentage for during which a service is capable of delivering the required function. Can be represented in terms of mean time between failure and mean time to restore:

$Availability = MTBF / (MTBF+MTTR)$

In this formula MTBF (meaning time between failure) stands for the average time

between failures and MTTR (meaning time to repair/restore service) stands for the on average time needed to repair or restore the service. Increased availability can come from reduced failure rates or from faster restoration times or both. AI could enhance availability by being able to detect warning signs other sooner and assist engineers to pinpoint the cause of the issue. However, if automated actions are performed incorrectly, this can reduce availability, and should be considered in conjunction with the accuracy of the availability estimation and the safety of the response [12].

The speed of detection is improved as shown below:

DI is defined as $DI = (MTTD^B - MTTD^A) / MTTD^B$

DI is an acronym for detection improvement, $MTTD^B$ is the baseline mean time to detection, and $MTTD^A$ is the mean time to detect (MTTD) with the use of AI.

The amount of service losses and the rate at which they are recovered is essential to cyber resilience. Define a function S(t) of the performance of essential services at time t. The cyber-resilience index is calculated as:

$CRI = (1 / T) \int_0^T S(t)\, dt$

CRI = resilience over the period of assessment, S(t) = service performance and T = period of observation. A resilient system will preserve a higher level of service when a disruption occurs and will more rapidly resume to normal operation.

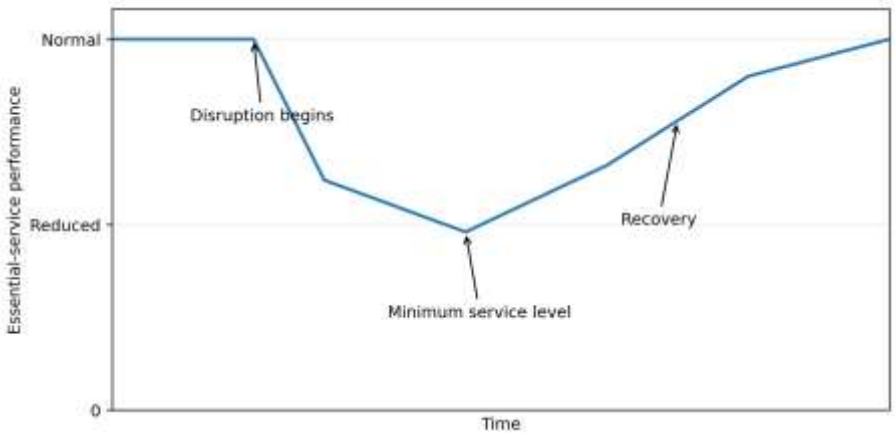


Figure 6. Essential Service Resilience Curve

Precision, recall, and the F1 score should be used to examine the performance of AI detection, not just overall accuracy:

Precision = TP / (TP + FP)

Recall = TP / (TP + FN)
F1 = 2 × [(Precision × Recall) / (Precision + Recall)]

TP: correctly detected threats, FP: normal events classified as threats, FN: missed threats in these formulas. The necessary amount of balance depends on the use of the AI output. If the model is only used by an analyst, then it can be comfortable with more false alarms, but if it is a model that automates production services, then evidence and accuracy must be more stringent.

The annual cyber damage can be predicted by:

$$EACL = \Sigma_{j=1}^{m} p_j L_j$$

EACL is the expected annual loss due to cyber events and p_j is the probability of cyber event j and L_j denotes the total loss from that event j.

$$L_j = L^D + L^O + L^s + L^R + L^P$$

L^D is the direct cost of response and recovery, L^O is the cost of operational interruption, L^s is the cost of the lost value of the supply chain, L^R is the cost of the reputational damage, and L^P is the cost of wider impact to the public or economy.

$$EACL^R = \Sigma_{j=1}^{m} p_j L_j (1 - r_j)$$

$EACL^R$ represents the reduction in expected loss after the framework is implemented and r_j represents the expected reduction in probability or impact.

$$B = EACL^B - EACL^R$$

The variables B, $EACL^B$, and $EACL^R$ stand for expected annual benefit, baseline expected annual loss and, residual expected annual loss, respectively. They bring together security investment and portfolio assessment in decreasing operational and economic risk.

A national resilience indicator should integrate the combined concepts of governance, protection, detection, response, recovery, reliability and economic-risk reduction:

$$NCRI = w^G G + w^P P + w^D D + w^R R + w^O O + w^E E$$

The National Cyber-Resilience Index (NCRI) is. The letters stand for governance maturity (G), protection capability (P), detection performance (D), response/recovery capability (R), operational reliability (O) and economic-loss reduction (E). The weights

are based on the national priorities. Once more the total score should mask a serious weakness – decision makers need to look at each component.

Table 2. AI-CNCI-RF Evaluation Metrics and Decision Indicators

Metric	Purpose in the framework
Cyber-risk score	Prioritizes risks using technical likelihood, service impact, exposure, and interdependence.
Dynamic trust score	Supports continuous zero-trust access decisions while leaving final enforcement to policy.
Data-quality score	Determines whether evidence is reliable enough for important AI-supported decisions.
Overall anomaly score	Combines heterogeneous security and reliability signals into one anomaly indicator.
Event-priority score	Directs analyst attention toward suspicious events that matter to essential operations.
Action-utility score	Determines whether a response is sufficiently safe and beneficial for bounded automation.
Service availability	Measures the proportion of time an essential service can perform its required function.
Detection improvement	Quantifies the reduction in mean time to detect produced by AI assistance.
Cyber-resilience index	Measures service preservation and recovery performance across a disruption.
Detection performance	Evaluates AI detection quality without relying only on overall accuracy.
Expected annual cyber loss	Estimates the annual economic damage associated with cyber incidents.
Expected annual benefit	Connects framework investment with expected economic-risk reduction.
National cyber-resilience index	Provides an integrated national indicator while preserving visibility of component weaknesses.

The proposed framework makes three main contributions. First, it populates a relationship between technical security events and crucially important services these events can impact. Many security systems generate a huge number of alerts, but don't indicate which ones are critical to the national or operational mission [13]. With resource and dependency intelligence AI-CNCI-RF puts each and every one of the events into context. Human oversight is still required. While AI can handle vast amounts of information faster, it doesn't necessarily comprehend emergency scenarios, physical safety needs, public services and political implications. The model might not contain important context that could be in the possession of a hospital operator, an electricity engineer or a government service owner [14]. Therefore, it is not recommended that

safety-critical services be given unrestricted autonomous control in the framework.

Governance should be first with implementation starting. Defined responsibilities for senior leaders, cybersecurity, cloud engineers, operational managers, data professionals, and critical-service owners. The organization should identify the key services, level of acceptable disruption, and should create an AI and cybersecurity governance board. A governance board should establish policies governing model-risk, data demands, automation boundaries and the human review process, and emergency override. It needs to identify critical use cases and determine what is successful in a measurable way when it comes to detection, response, service availability, recovery, safety, and economic-risk reduction. The organization should establish an up-to-date list of accounts, clusters, containers, applications, interfaces, identities, data stores, pipelines, suppliers, and recovery systems in the cloud. Every asset should be a part of and linked to the critical service it serves [15]. A dependency analysis should be carried out and shared systems identified that have the capability to impact multiple services in the event of their failure. Dependency map should be automatically updated on change in cloud-native environment based on discover and communication analysis. Do not progress to high impact AI automation when assets are not fully visible, because if AI can't see something that's important or dependent, it may suggest taking unsafe action.

Table 3. Phased Implementation Roadmap for AI-CNCI-RF

Implementation phase	Main activities	Expected outputs	Illustrative success indicators
Phase 1: Governance and mission definition	Form the AI and cybersecurity governance board; assign responsibilities; identify essential services; define acceptable disruption, safety needs, and automation boundaries.	Governance charter; responsibility matrix; critical-service register; AI-use policy; escalation and emergency-override procedures.	Critical services with assigned owners; approved AI policies; documented authority and escalation routes.
Phase 2: Asset and dependency discovery	Inventory cloud accounts, clusters, containers, APIs, identities, data stores, pipelines, suppliers, and recovery assets; map their service dependencies.	Dynamic asset inventory; dependency graph; supplier-risk register; service criticality classifications.	Asset coverage; assets linked to essential services; unidentified or unverified dependencies.
Phase 3: Security baseline	Improve authentication, least privilege, workload identity, image signing,	Hardened cloud-native environment;	Privileged-access reduction; signing and

Implementation phase	Main activities	Expected outputs	Illustrative success indicators
strengthening	segmentation, secrets management, CI/CD security, backups, and tested response processes.	protected software pipeline; verified backup design; tested incident and recovery playbooks.	scanning coverage; remediation rate; backup and playbook test success.
Phase 4: Trusted data foundation	Collect, protect, time-align, validate, and normalize identity, Kubernetes, runtime, application, network, pipeline, vulnerability, backup, and operational data.	Trusted security-data platform; data-quality controls; privacy, retention, and access policies.	Data-quality score; telemetry coverage; ingestion delay; integrity and protection coverage.
Phase 5: AI-assisted monitoring	Deploy AI for alert correlation, anomaly detection, vulnerability prioritization, configuration drift, incident summarization, attack-path analysis, and failure prediction.	AI-assisted operations dashboard; prioritized alerts; explainable recommendations and confidence estimates.	Precision, recall, F1; false-positive rate; MTTD; analyst effort and decision-time reduction.
Phase 6: Validation and controlled testing	Test models through historical events, simulations, cyber exercises, production trials, normal operational change, and adversarial scenarios.	Validation report; approved thresholds; documented failure cases; tested response and rollback procedures.	Explanation quality; model drift; unsafe recommendations detected; response accuracy and operational impact.
Phase 7: Bounded automation	Automate only predefined, reversible, authorized, and low-risk actions such as evidence collection, stronger authentication, temporary token withdrawal, and enhanced monitoring.	Approved automation catalogue; auditable action logs; rollback procedures; accessible human override.	Successful automated actions; rollback success; unauthorized actions; service disruption caused by automation.

Implementation phase	Main activities	Expected outputs	Illustrative success indicators
Phase 8: Recovery, learning, and national coordination	Verify restored systems; perform post-incident review; improve models and controls; share trusted indicators; conduct cross-sector and national exercises.	Trusted recovery process; lessons-learned repository; updated models and playbooks; coordination and information-sharing procedures.	MTTR; availability; CRI; expected-loss reduction; exercise performance; recurrence of known failure modes.

Before introducing AI advanced security control should be enhanced. The organization should be enhancing authentication, least privilege access, workload identity, container security, network segmentation, secret management, software signing, protection of pipelines, and backup protection [15]. It should also have incident-response and recovery processes which have been tested and included in the security baseline. AI doesn't make up for lax access control, unsecured backups or sloppy software-development processes. The organization should be able to retrieve appropriate data from identity platforms, cloud systems, Kubernetes, applications, networks, software pipelines, backup functionalities, and operational environments. Information needs to be quality controlled, protected, time aligned and normalized. Data governance should be established to determine who should have access to the information, how its tenure is managed, and how its privacy and/or sensitive information is managed. No use of a model must be used in an important decision when the data sources are missing, are delayed or are not reliable.

Initially, AI assistance should be used to assist and not to take command of critical systems. Applications are such as alert correlation, vulnerability prioritization, identity anomaly detection, container-behavior monitoring, incident summarization, configuration-drift detection, and failure prediction [16]. In this phase, the organization ought to validate the AI suggestions with specialist choices. Documented errors should be explained and incorporated on improving the model. The following should be emphasized as much as technical accuracy: Its practical usefulness. Controls production trials, cyber exercises, and simulated attacks, historical events, and normal operation changes should all be used to test AI models. Testing should include accuracy, false alerts and false negative statistics, explanation quality, data drift issues, response time, and operational issues. Response playbooks need to be exercised as well [17]. Each action should be verified to be safe, authorized, reversible and appropriate for the impacted service. It should not be assumed that a model that is capable of functioning well in one environment can be considered to function similarly in another.

Once validated, predefined low risk activities are automatically performed by the organization. These actions can range from collecting evidence to more robust

authentication, temporary token removal, and/or more intensive monitoring or isolation of a non-critical workload [18]. High impact actions should only be under the control of humans unless carefully tested, approved Emergency Playbook. All automated activities should be recorded, auditable and reversible. System should be easily by-passable in case of human personnel. Linkage between organizational monitoring and sectoral and national coordination are expected to take place at the final implementation stage [19]. Operators shall employ common terms for incidents and provide appropriate threat indicators, dependency warnings and lessons learned to trusted authorities. Tests should consider the simultaneous disruption of a number of originations at sector and national levels. The exercises should discuss communication, decision making, continuity of operations, supplier coordination, public information, and recovery priorities [20]. The roadmap should be an on-going project in the form of a cycle. Technology and threats continuously evolve, and governance, asset awareness, data quality, model performance, security controls, recovery plans, and national vulnerabilities should be consistently assessed.

Discussion

A critical limitation of traditional cybersecurity methods is a lack of security controls that are integrated with business continuity planning, economic-risk assessment, and national resilience planning in a unified, cloud-based framework, known as the Cloud-Native Critical Infrastructure Resilience Framework (CNCI-RF), and the proposed Artificial Intelligence-Enabled Cloud-Native Critical Infrastructure Resilience Framework (AI-CNCI-RF) aims to alleviate that concern. It is not enough to assess a cyber-security incident by technical severity, however, in a critical infrastructure. Its significance also relies on the crucial service that it impacts, the quantity of dependent systems, the duration the disruption takes place for, its safety implications, and the potential financial loss. The framework thus links technical security events to the mission-critical services and their operational requirements. A major contribution of AI-CNCI-RF is being service-centered. Traditional security products typically rate vulnerabilities based on technical scores and security monitoring platforms pump out alerts that don't necessarily provide insight into operational impact. The proposed framework enhances this process with connections to assets, identities, containers, APIs, software pipelines, cloud resources, suppliers and external services dedicated to the critical function they serve. For this reason, a seemingly modest vulnerability may be given a higher priority when it is a part of a national payments system, emergency communications system, hospital services, energy-management system or other highly interdependent services. This way will see a more realistic perspective on the cyber risk and potential ripple effect.

A key observation from the framework is that AI is not necessarily a standalone cybersecurity solution. The advantages of AI for analysis include accelerating the correlation of security events, spotting anomaly detection alerts, discovering hidden attack paths, predicting possible system failures, prioritizing vulnerabilities, and summarizing evidence from incidents. These capabilities are extremely beneficial in cloud-native environments where workloads, identities, services, and network relationships constantly

change. Due to the quantity and velocity of telemetry data flowing from Kubernetes clusters, containers, service meshes, APIs, software pipelines, identity platforms and operational systems, it can be challenging for human analysts to process manually. But the quality and reliability of the data sources play a crucial role in the effectiveness of these AI tools. This combination of these factors (incomplete logs, late telemetry, log falsification, incorrect timestamp, disjointed process) can result in incorrect conclusions. This is why it is important that AI-CNCI-RF includes a data-quality score to ensure that models don't get taken for granted in every instance. When the completeness, accuracy, timeliness, or integrity of the AI recommendations is below the approved risk parameters, that information should have less autonomy in the decision-making process. This forms a proportional relationship between Data Governance and the acceptable level of AI-supported automation.

The other key phenomenon of the framework is the "zero-trust" security collaborative integration with AI. Cloud-native systems are made up of human and machine identities such as programming agents, containers, services, deployment pipelines, and users. Having one of the valid credentials does not imply that an access request is secure. Dynamic trust model consists of identity confidence, device/workload health, behavior normalization, operational context and sensitivity of resources. AI can facilitate the assessment of behavioral and contextual evidence, yet gain access is still managed by robust security policies. This ensures that an AI model does not give access beyond the approved access boundaries of the organization and aligns with the philosophy of continuous verification in zero trust architecture. AI-CNCI-RF also understands that cybersecurity is intertwined with operational reliability. An incident or event caused by a malicious attack, software failure, incorrect configuration, overloaded service, expired certificate or defective software update can have similar impact including service unavailability or an inability to protect data integrity. That's why cyber-threat monitoring goes hand in hand with failure prediction and operational monitoring within the framework. Time-series analysis can be used to alert you to escalating workload demands, repeated container restarts, latency abnormalities, running out of resources, and escalating error rates. These alerts can be a way for operators to take steps to prevent a technical problem from turning into a big interruption in service.

Still, swift detection doesn't necessarily warrant speedy autonomous response. A wrong answer could cause greater harm than the original security incident. An automatic operation that is performed on a legitimate hospital application, that targets the electricity system's operator or that disables an emergency communication service could result in significant consequences to the operation or safety of the application. The proposed action-utility function aims to solve the above problem by juxtaposing the expected security gain of taking an action against the operational cost, the error cost, and the uncertainty of the action. This enables the risk-controlled response model where AI can only do bounded and reversible actions, watch, suggest or execute. This is perhaps most critical aspect of the framework, the man in the middle is called "human in the loop" safety gate. Low risk activities like collecting evidence, bumping up monitoring or asking for a token increase might be automated after testing and approval. Actions which may disrupt critical services,

on the other hand, should generally require approval from competent human personnel. This architecture ensures that organizations can reap the advantages of machines' velocity while maintaining control of safety-critical systems at the hands of people.

Trusted recovery is also regarded as more than like-for-like restoration of technical availability. A compromised workload, a compromised configuration, a compromised identity, a compromised backup and/or a container image can restore the attacker with the service. AI-CNCI-RF thus needs verification of software images, configurations, identities, data, backups, and software dependencies before systems are allowed to go back to operation. AI can help by providing the “last-known-trusted” state, comparing restored environments to approved baselines, and identifying indications of persistent compromise. The final decision to reconnect should still be subject to the authorized service owner and authorized recovery personnel. The measure model of the framework gives it greater practical significance. Availability, MTTD, MTTR, precision, recall, F1 score, cyber-resilience index, expected annual cyber loss, and national cyber-resilience indicators offer a number of ways organizations can assess performance. This is crucial since a highly accurate detection rate doesn't necessarily indicate that an AI-powered security system is good. A single model can also find more anomalies but also introduce more nuisance alarms and/or create unsafe automation. The metrics when combined into one helps to drive organizations to consider effectiveness, reliability, recovery time, security continuity and economic impact in combination.

An important element of this is the economic-security aspect, which is directly relevant to critical infrastructure. When it comes to supply chains, business operations, public confidence, and national productivity, any disruptions to electricity, telecommunications, financial, transportation, healthcare, water, and government platforms can have an impact. The cyber loss expected and the cyber loss after the framework will provide a good understanding for decision makers to compare the reduction of cyber risk in relation to the investment for cybersecurity. While technical severity scores are included, these estimates also have their own limits of uncertainty, but give a better idea on which ones to prioritize the security resources. AI-CNCI-RF indicates on the national level that there should be coordination between the infrastructure operators, cloud providers, software suppliers, regulators, emergency authorities and government institutions. Failure in one of these interdependent critical sectors identities, cloud technologies, telecom service, or software can impact multiple organizations at once. By using common incident terminology, trusted information sharing, dependency warnings and cross-sector exercises, and coordinated priorities for recovery, the chances of a local technical incident turning into a national crisis can be decreased.

AI-CNCI-RF is best viewed as a decision support and resilience regime governed as opposed to fully automated security regime. It is primarily known for offering cloud-native security, zero-trust access, trusted telemetry, AI-assisted analysis, human-controlled response, and verified recovery and measuring economic risk. The approach does not imply a definitive prevention of attacks, nor a definitive rightness in the results of the AI. Its approach is to facilitate the recognition of major threats at an earlier stage in the process, to lead organizations to informed decision-making, to facilitate the continuity of minimum

essential functions during interruptive situations and to restore the organizational functions through trusted, accountable processes. The framework is a design-science contribution in the form of a conceptual model, and its later implementation can be tailored and conducted in the particular sectors of energy, healthcare, telecommunications, finance, transportation, water, and digital government or others.

Availability of data is the first big challenging problem. Data in critical-infrastructure incidents are confidential and may be hard to share. Older systems can also result in incomplete history information [21]. Models which have been trained with public or laboratory datasets may not perform exactly the same in real production environment. Synthetic data and cyber exercises can assist, but can't necessarily replicate all conditions of operation. False positive and false negative are problems that have not gone away. Analysts can get swamped or have to redirect their efforts to deal with false positives. In addition, false negatives don't detect that attackers are operating on the system [22]. AI should thus enter into the picture as one element of a defense in-depth strategy, among secure architecture, access control, monitoring rules, threat intelligence, human analysis and recovery capability. This study is conceptual and does not give the results from a complete deployment. The key variable, weights and thresholds in the formulas should be validated. Future research should validate the framework for the specific sectors and give the comparison between their baseline and their performance with the help of AI.

Conclusion

AI can assist organizations in handling these intricate surroundings, by processing vast quantities of security information and operational data. It's able to spot unusual activity, prioritize critical alerts, map out potential attack vectors, anticipate technical issues, aid in incident investigations, and suggest appropriate actions. However, it's important to not look at AI as a stand-alone security solution. There can be additional risks due to incorrect predictions, false alerts, poor data quality, manipulation of models and unsafe automated actions. Therefore, human supervision, transparent management, safe options for data control, model verification, and properly managed automation are required. In this paper an approach has been proposed, called Artificial Intelligence-Enabled Cloud-Native Critical Infrastructure Resilience Framework (AI-CNCI-RF). The framework incorporates governance, mapping of critical services, visibility of assets and dependence, secure cloud-native engineering, Zero-trust Access Control, Trusted Telemetry, AI-driven threat detection, controlled response, consistent recovery, and economic-risk measurement. It's primarily designed to make sure the AI is aiding cybersecurity decision-making without taking away the need for human involvement and without causing disruption in the process. In summary, when used responsibly, AI can help to strengthen cloud-native critical infrastructure, making it more secure, reliable, and resilient. The most powerful advantage will be the facilitation of human expertise to comprehend threats quicker, and make better decisions, protect critical services and recover more effectively. Using secure technology, accountable governance, operational planning and national coordination can help create a framework that can help make a nation more resilient and economically secure in the cyber world.

References

[1] M. Abd Elaziz, L. Abualigah, and I. Attiya, "Advanced optimization technique for scheduling IoT tasks in cloud-fog computing environments," Future Generation Computer Systems, vol. 124, pp. 142–154, 2021, doi: 10.1016/j.future.2021.05.026.
[2] P. Alvaro, K. Andrus, C. Sanden, C. Rosenthal, A. Basiri, and L. Hochstein, "Automating failure testing research at Internet scale," in Proceedings of the Seventh ACM Symposium on Cloud Computing (SoCC '16), 2016, doi: 10.1145/2987550.2987555.
[3] W. Chipidza, "The effect of toxicity on Covid-19 news network formation in political subcommunities on Reddit: An affiliation network approach," International Journal of Information Management, vol. 61, Art. no. 102397, 2021, doi: 10.1016/j.ijinfomgt.2021.102397.
[4] A. Basiri, N. Behnam, R. de Rooij, L. Hochstein, L. Kosewski, J. Reynolds, and C. Rosenthal, "Chaos engineering," IEEE Software, vol. 33, no. 3, pp. 35–41, 2016, doi: 10.1109/MS.2016.60.
[5] M. I. A. Chy and R. A. Ridu, "A next-generation zero-trust security framework for cloud-native microservices powered by AI," Advances in Intelligent Systems Research, pp. 919–931, 2026, doi: 10.2991/978-94-6239-664-7_63.
[6] S. R. Davies, R. Macfarlane, and W. J. Buchanan, "Differential area analysis for ransomware attack detection within mixed file datasets," Computers & Security, Art. no. 102377, 2021, doi: 10.1016/j.cose.2021.102377.
[7] D. Xu, I. Gondal, X. Yi, T. Susnjak, P. Watters, and T. R. McIntosh, "The erosion of cybersecurity zero-trust principles through generative AI: A survey on the challenges and future directions," Journal of Cybersecurity and Privacy, vol. 5, no. 4, Art. no. 87, 2025, doi: 10.3390/jcp5040087.
[8] A. K. Dharmalingam and T. R. Kalai Lakshmi, "Evaluating exposure-driven security governance for proactive risk mitigation of AI-powered cyber attacks in digital infrastructures," Discover Computing, vol. 29, no. 1, 2026, doi: 10.1007/s10791-026-10184-z.
[9] A. M. A. Hamdi, F. K. Hussain, and O. K. Hussain, "Task offloading in vehicular fog computing: State-of-the-art and open issues," Future Generation Computer Systems, vol. 133, pp. 201–212, 2022, doi: 10.1016/j.future.2022.03.019.
[10] C. Wu, E. Haihong, and M. Song, "An automatic artificial intelligence training platform based on Kubernetes," in Proceedings of the 2020 2nd International Conference on Big Data Engineering and Technology, pp. 58–62, 2020, doi: 10.1145/3378904.3378921.
[11] S. Y. Kim, H. Park, H. Kim, J. Kim, and K. Seo, "Technostress causes cognitive overload in high-stress people: Eye tracking analysis in a virtual kiosk test," Information Processing & Management, vol. 59, no. 6, Art. no. 103093, 2022, doi: 10.1016/j.ipm.2022.103093.
[12] T. Theodoropoulos, L. Rosa, C. Benzaid, P. Gray, E. Marin, A. Makris, L. Cordeiro, F. Diego, P. Sorokin, M. Di Girolamo, P. Barone, T. Taleb, and K. Tserpes, "Security in cloud-native services: A survey," Journal of Cybersecurity and Privacy, vol. 3, no. 4, pp. 758–793, 2023, doi: 10.3390/jcp3040034.
[13] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," NIST Special Publication 800-207, 2020, doi: 10.6028/nist.sp.800-207.

[14] S. Yellu, "Resilient cloud security architectures: Combining adaptive anomaly recognition with AI-driven Kubernetes credential governance," SSRN Electronic Journal, 2026, doi: 10.2139/ssrn.6096586.
[15] S. Patel, "Cloud-native AI security architecture for the U.S. electric grid," International Journal of Global Innovations and Solutions (IJGIS), vol. 3, no. 3, 2026, doi: 10.63412/26p8n109.
[16] I. Moya, M. Chica, J. L. Sáez-Lozano, and Ó. Cordón, "Simulating the influence of terror management strategies on the voter ideological distance using agent-based modeling," Telematics and Informatics, vol. 63, Art. no. 101656, 2021, doi: 10.1016/j.tele.2021.101656.
[17] P. R. Nangi, C. K. R. N. Obannagari, and S. Settipi, "A multi-layered zero-trust security framework for cloud-native and distributed enterprise systems using AI-driven identity and access intelligence," International Journal of Emerging Trends in Computer Science and Information Technology, vol. 4, pp. 144–153, 2023, doi: 10.63282/3050-9246.ijetcsit-v4i3p115.
[18] M. Potluri, "AI-driven zero trust security for Kubernetes and multi-cloud deployments," World Journal of Advanced Engineering Technology and Sciences, vol. 15, no. 2, pp. 394–404, 2025, doi: 10.30574/wjaets.2025.15.2.0559.
[19] Z. Liu, Y. Ma, H. Zheng, D. Liu, and J. Liu, "Human resource recommendation algorithm based on improved frequent itemset mining," Future Generation Computer Systems, vol. 126, pp. 284–288, 2022, doi: 10.1016/j.future.2021.08.017.
[20] NIST, "AI risk management framework," Artificial Intelligence Risk Management Framework (AI RMF 1.0), vol. 1, no. 1, 2023, doi: 10.6028/nist.ai.100-1.
[21] H. Liang, "Book review," Technological Forecasting and Social Change, vol. 182, Art. no. 121809, 2022, doi: 10.1016/j.techfore.2022.121809.
[22] Z. Lin, Y. Wu, and X. Xing, "DirtyCred," in Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, 2022, doi: 10.1145/3548606.3560585.
[23] E. O. Afolabi and G. Nwabueze, "Security challenges in AI and cloud infrastructure, including software supply chains and model-deployment pipelines," American Journal of Innovation in Science and Engineering, vol. 5, no. 2, pp. 58–69, 2026, doi: 10.54536/ajise.v5i2.7220.
[24] O. P. Ezeugwa, A. Bello, and T. A. Raheem, "Automated cyber threat detection and incident response using AI-enabled infrastructure analytics," American Journal of Innovation in Science and Engineering, vol. 5, no. 2, pp. 70–82, 2026, doi: 10.54536/ajise.v5i2.7219.
[25] F. Darbandeh, M. Rizwan Asghar, and L. Chen, "Cyber resilience in IoT networks revisited: Frameworks, mechanisms, and future directions," ACM Computing Surveys, vol. 58, no. 13, pp. 1–45, 2026, doi: 10.1145/3818682.
[26] H. M. Khan and S. M. H. Zaidi, "Detecting security system misconfiguration threats in cloud computing environments using AI," American Journal of Innovation in Science and Engineering, vol. 3, no. 3, pp. 31–40, 2024, doi: 10.54536/ajise.v3i3.3272.
[27] Y. Majeed and A. Majeed, "Audit to assurance: Evaluating the impact of regulatory cybersecurity audits on organisational cyber resilience and strategic decision-making,"

American Journal of Innovation in Science and Engineering, vol. 5, no. 1, pp. 86–101, 2026, doi: 10.54536/ajise.v5i1.7223.
[28] A. Montieri, G. Bovenzi, G. Aceto, D. Ciuonzo, V. Persico, and A. Pescapè, "Packet-level prediction of mobile-app traffic using multitask deep learning," Computer Networks, vol. 200, Art. no. 108529, 2021, doi: 10.1016/j.comnet.2021.108529.
[29] J. Y. Mir and F. Mir, "Secure cloud-native platforms for critical service continuity: An AI-driven framework for national cyber and economic resilience," in press.
[30] J. Y. Mir and F. Mir, "Strengthening U.S. digital infrastructure through AI-powered cloud-native software resilience and secure delivery systems," Journal of Computer Science and Technology Studies, in press.
[31] K. R. Muppa, "AI-driven continuous governance for machine identities in cloud-native zero trust environments," 2026, doi: 10.64971/j.cph.eijtem.v13.i2.11.2026.
[32] P. R. Nangi, C. K. R. N. Obannagari, and S. Settipi, "A multi-layered zero-trust security framework for cloud-native and distributed enterprise systems using AI-driven identity and access intelligence," International Journal of Emerging Trends in Computer Science and Information Technology, vol. 4, pp. 144–153, 2023, doi: 10.63282/3050-9246.ijetcsit-v4i3p115.
[33] S. Singh, C. S. Kubam, N. N. Modh, K. M. Zubair, and Ashanand, "A bio-inspired deep learning intrusion detection framework for cyber-physical system security," in 2025 International Conference on Computational Intelligence, Security, and Artificial Intelligence (IntelliSecAI), Al-Khobar, Saudi Arabia, 2025, pp. 1–6, doi: 10.1109/IntelliSecAI66368.2025.11473168.
[34] K. M. Zubair, T. R. Akash, and S. A. Chowdhury, "Autonomous threat intelligence aggregation and decision infrastructure for national cyber defense," Frontiers in Computer Science and Artificial Intelligence, vol. 2, no. 2, pp. 26–51, 2023.