The Eastasouth Journal of Information System and Computer Science
Vol. 4, No. 02, December 2026, pp. 1-18
ISSN: 3025-566X, DOI: 10.58812/esiscs.v4i02

Risk-Aware Federated Cyber Defense for Cloud–Edge Systems Through Privacy, Network Resilience, and Accountable Response

Celeste Yvonne Ramirez
BBW University of Applied Sciences, Berlin, Germany

Article Info

Article history:

Received Oct, 2026
Revised Oct, 2026
Accepted Oct, 2026

Keywords:

Cloud Security;
Cybersecurity Governance;
Edge Computing;
Federated Learning;
Intrusion Detection;
Network Resilience;
Privacy;
Risk-Aware Response

ABSTRACT

Cloud services, edge devices, and organizational information systems expose security teams to a shared problem: a locally accurate detector can still produce a poor operational decision when its recommendations ignore service dependencies, tenant boundaries, privacy constraints, or the consequences of interruption. This paper develops a design-science framework for coordinating these concerns within a single decision process. A purposive corpus of 55 supplied publications informs the requirements; it is treated as a bounded design resource rather than a systematic or exhaustive evidence base. The proposed architecture connects tenant-local telemetry, context-sensitive threat estimation, privacy-preserving collaborative learning, dependency-aware response selection, and verifiable governance records. Its central contribution is a constrained decision formulation that minimizes expected attack loss and response-induced disruption while enforcing explicit limits on privacy expenditure, latency, and service availability. Analytical examination establishes three conditional properties: independent event thresholds can violate shared service constraints, selective abstention can reduce decision loss when review is sufficiently informative, and federated participation does not by itself establish privacy. A transparent worked example demonstrates how a capacity constraint changes which alerts receive an automated response; the example is illustrative and does not represent measured performance. A prospective evaluation protocol specifies chronological and tenant-separated testing, rare-attack assessment, poisoning experiments, privacy accounting, recovery analysis, and paired comparisons against operational baselines. No experiments or production deployments are reported. The framework provides an implementable research specification and testable propositions for examining whether coordinated cyber defense improves operational outcomes beyond predictive accuracy alone.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Name: Celeste Yvonne Ramirez
Institution: BBW University of Applied Sciences, Berlin, Germany
Email: celesteramirez.analytics@gmail.com

1. INTRODUCTION

An intrusion detector does not protect an organization merely by assigning the correct label to a network flow. Protection depends on what happens next: whether a credential is revoked, whether a container is isolated, whether a critical service remains available, and whether another tenant is exposed while the incident is investigated.

These decisions are increasingly distributed across cloud services, edge devices, identity platforms, and management information systems. A response that appears sensible within one component can interrupt a dependency elsewhere. A collaborative model can improve coverage while also creating new channels for inference or poisoning. Security therefore requires a decision process that joins prediction to operational consequences.

Research on cybersecurity frameworks for cloud-integrated management information systems brings detection, data governance, federated learning, and organizational behavior into the same architectural discussion [1].

Work on intelligent cloud data protection provides a related architectural starting point for considering how protective controls can be coordinated across computing environments [2].

The connection between cloud infrastructure, organizational information, and defensive reasoning is also central to cognitive cyber-defense research [3].

These contributions motivate a concrete question: how should a security system choose an action when the strongest local response is incompatible with the availability, privacy, or review capacity of the wider organization? Answering that question requires more than another classification model. It requires an explicit account of the protected service, the decision maker's available information, the action's expected benefit, and the conditions under which automation must stop. It also requires an evaluation that measures harm caused by the defense itself.

The supplied literature includes broader discussion of resilience in the cyber-economic setting, which gives a reason to examine security as an organizational decision problem rather than solely a technical classification task [4].

The present paper develops a design-science artifact: a risk-aware federated cyber-defense architecture and its associated decision rule, assurance conditions, and prospective evaluation protocol. Its contribution is the formal coupling of tenant-local threat estimates with shared operational constraints. The paper does not claim that its component technologies are new. Instead, it specifies how they can interact without treating prediction quality, privacy, resilience, and accountability as interchangeable achievements.

Three research questions guide the design. RQ1 asks how tenant-local and shared threat information can support coordinated decisions without assuming that raw telemetry can be pooled. RQ2 asks how service dependencies, response capacity, and uncertainty should constrain automation. RQ3 asks what evidence would be necessary to demonstrate a practical advantage over centralized classification, ordinary federated detection, or unconstrained response policies. The intended setting is a multi-tenant cloud–edge environment with heterogeneous workloads and a limited incident-review team. Claims about universal deployment suitability are outside the scope.

2. RESEARCH DESIGN AND EVIDENCE BOUNDARIES

2.1 Design-science orientation

The research procedure consists of six activities: defining the operational problem, organizing a bounded literature corpus, deriving requirements, specifying the architecture, examining conditional properties, and preparing a falsifiable evaluation protocol. The unit of analysis is an alert-to-response decision within a tenant and service context. A second unit is a coordinated response batch whose actions may jointly affect a shared dependency. These units differ from a network record used for classification: a single incident may create many records, while one response may affect many services.

The corpus contains the 55 references supplied for this manuscript. Its inclusion rule is user selection, not database retrieval. No search-completeness claim, quality-weighted meta-analysis, pooled effect estimate, or

systematic-review designation is made. Sources are used to identify relevant research problems and design alternatives. Numerical results are not inferred from titles, conference names, or publication venues. Several records lack complete authors or pagination; those limitations remain visible in the reference list.

Each source is assigned a primary role: architectural context, privacy mechanism, detection approach, network or service requirement, governance requirement, or adjacent application. The assignment is interpretive rather than a claim about the source's exclusive contribution. Titles and available publication descriptions support topical placement; they do not establish that a mechanism performs reliably in the proposed setting. Full-text appraisal and bibliographic correction remain necessary before journal submission.

2.2 *Deriving requirements without overstating evidence*

Work connecting AI, big data, and information systems to supply-chain integrity and infrastructure protection motivates a requirement to retain business-service context when assessing threats [5].

The economic treatment of supply-chain vulnerabilities motivates distinguishing a detection error from the loss that the error can cause in an interconnected service environment [6].

Research on AI-assisted protection of critical infrastructure explicitly connects technical detection and response with institutional policy considerations [7].

These themes lead to five requirements. First, every operational prediction must be associated with its tenant, event time, model version, and available evidence. Second, collaborative learning must define its privacy mechanism and adversary instead of relying on data locality alone. Third, response decisions must account for affected services and shared dependencies. Fourth, uncertain or unusually consequential actions must have a review path. Fifth, the system must preserve enough decision evidence to investigate both successful containment and harmful intervention.

Work on advanced cyber threats and cybersecurity innovation provides an additional context for retaining both emerging attack scenarios and organizational responses within the requirements [8].

These are design requirements, not findings about all 55 studies. Their validity will depend on the intended deployment and stakeholder review. For example, a safety-sensitive service may require a stricter interruption limit than a development workload. The framework allows those limits to be specified explicitly rather than embedding them in an unexplained model threshold.

3. RELATED RESEARCH AND THE INTEGRATION PROBLEM

3.1 *Governance and organizational resilience*

Cybersecurity risk management embedded within information-systems governance provides a direct context for linking technical controls to organizational resilience and infrastructure protection [9].

Treating cyber threat intelligence as a management information system makes the route from security evidence to accountable organizational action an explicit subject of study [10].

The application of AI-driven cybersecurity to IT project management introduces another relevant setting in which detection and risk mitigation must operate within delivery and resource constraints [11].

Employee behavior is also a distinct concern in cybersecurity research, as reflected in work examining security training in business environments [12].

The implication for this design is that a response policy cannot be defined only by a model developer. Service owners must identify unacceptable interruptions; analysts must define review capacity; privacy officers must set permitted data uses; and infrastructure operators must specify rollback procedures. A model can estimate threat probability, but it cannot independently determine these organizational commitments. The architecture therefore treats policy inputs as versioned artifacts with named owners.

3.2 Detection, time, and rare events

Big-data analytics for threat detection in management information systems represents one strand of the supplied security literature [13].

AI-enhanced threat detection in management information systems offers another directly relevant application area [14].

Real-time security analytics introduces the question of whether evidence can be processed quickly enough to support an intervention while it remains useful [15].

Temporal cyber-threat prediction is relevant to the separate challenge of distinguishing information available before an incident from information created by the incident itself [16].

The influence of AI on data-system security provides broader context for examining how predictive tools affect protective decisions [17].

Comparisons of transformers and LSTM architectures address model-family choices, although a model comparison alone does not settle deployment requirements [18].

Rare-attack detection in IIoT networks is directly relevant because a model can perform well on common classes while missing infrequent events with substantial consequences [19].

The proposed framework separates three times: event occurrence, feature availability, and decision execution. This distinction prevents features from being assigned to an earlier decision merely because they were later attached to the same incident. It also prevents evaluation from treating many near-duplicate flows as independent successful detections. Detection quality will be assessed at both record and incident levels, with attention to class prevalence, calibration, and alert burden.

3.3 Adaptation and defensive failure

Continual learning for insider-threat detection raises the issue of adaptation when behavior changes during an organization's normal operation [20].

Automated vulnerability assessment introduces asset exposure as a separate information source that can complement observed traffic [21].

Malware-image classification using binary representations offers a modality-specific detection approach that could provide endpoint evidence within a wider incident record [22].

Adversarial machine learning for cyber defense makes deliberate manipulation of the detector part of the design problem [23].

Self-healing security systems using reinforcement-learning agents provide a relevant alternative for selecting or adapting defensive actions [24].

Threat prediction and autonomous response for containerized microservices place these concerns within hybrid-cloud service operation [25].

Adaptation is therefore treated as a controlled change to a deployed decision system. A drift alarm does not directly authorize retraining and rollout. It triggers investigation, candidate development, and a comparison against the current version. Reinforcement learning is an optional policy component, not a prerequisite: an explicitly constrained optimizer provides an interpretable baseline

against which an adaptive policy must justify its additional complexity.

3.4 Privacy and collaborative learning

Federated cyber threat intelligence in cloud-integrated information systems motivates learning across organizations without assuming centralized access to their raw records [26].

Homomorphic federated models introduce encrypted computation as a possible privacy-preserving mechanism [27].

Secure multi-party computation across collaborative cloud tenants provides another candidate mechanism for joint intrusion detection [28].

TrustCloud-FL brings self-healing behavior and privacy-preserving analytics into the multi-tenant hybrid-cloud setting [29].

AEGIS-FL adds auditability to federated threat detection across cloud–edge systems [30].

Efficient homomorphic encryption for public-cloud big-data analytics highlights the practical importance of computational overhead and performance trade-offs [31].

Federated healthcare analytics offers an adjacent application in which data sensitivity and institutional separation are central design constraints [32].

Privacy-aware data protection models for cloud environments provide a further basis for defining tenant-local data controls before selecting a collaborative mechanism [33].

These mechanisms solve different problems. Encrypted transport protects a communication channel. Secure aggregation can restrict visibility into individual updates. Differential privacy, if used with a defined adjacency relation and accountant, can bound specified disclosure risks. Robust aggregation targets malicious updates. None is a substitute for all the others. The design records which protection applies to which message and which adversary it addresses.

3.5 Provenance, blockchain, and data governance

Privacy-preserving information systems combining blockchain, federated learning, and zero-trust models motivate attention to how identity and access decisions interact with collaboration [34].

Cryptographic techniques for resilient information-systems design provide a related foundation for integrity and authenticity controls [35].

Blockchain-based access control and compliance auditing for federated cloud providers directly concerns the evidence available to assess cross-provider decisions [36].

Blockchain integration for cybersecurity and data integrity offers another provenance-oriented application in management information systems [37].

Privacy-aware big-data governance using blockchain motivates separating an audit record's integrity from the sensitivity of the information it contains [38].

A trustworthy data lakehouse combining federated learning and blockchain supplies an adjacent data-platform setting for this separation [39].

The application of blockchain alongside AI and analytics to sustainable business practices broadens the organizational context, although it does not establish security effectiveness for this architecture [40].

The proposed system does not assume that every audit trail needs a blockchain. A signed append-only log is the default baseline. A distributed ledger becomes a candidate only when multiple parties require shared verification without accepting a single record custodian. Raw telemetry, personal information, and secret keys remain outside the ledger. An immutable record of an incorrect decision remains incorrect; provenance supports

investigation, not automatic trust in the underlying judgment.

3.6 Networks, edge services, and operational constraints

Network survivability and traffic-flow importance provide a basis for representing the consequences of losing a shared component [41].

Secure edge intelligence for digital-twin communication systems introduces a setting in which models and communication requirements intersect [42].

Secure zero-touch 6G frameworks extend this context to automated network operation and edge-AI applications [43].

A secure semantic communication architecture for edge–IoT–6G systems provides a further reason to specify which information must cross the network and when [44].

Trust and bio-inspired optimization in digital-twin frameworks introduce optimization and trust as related architectural themes, although the supplied bibliographic record requires title completion [45].

Cybersecure smart-building sensors provide an application in which network attacks can affect services associated with physical environments [46].

Resilient cybersecurity architectures for large distributed systems motivate testing the defense under component failure and communication disruption [47].

These topics support a dependency-aware design. The relevant graph is not simply a diagram of network links. It must distinguish identity dependencies, data dependencies, service failover, and physical-process constraints. A graph assembled from stale configuration data can make the optimizer confidently wrong. Dependency freshness is therefore an assurance condition rather than an incidental engineering detail.

3.7 Infrastructure, sustainability, and adjacent decisions

National cyber-threat intelligence research provides context for connecting local defenses to broader infrastructure protection objectives [48].

AI-enhanced threat detection and response for critical infrastructure supplies another motivation for examining operational consequences beyond classifier output [49].

Real-time cyberattack analytics and proactive mitigation connect threat estimation to the timing of the defensive action [50].

Green and secure data-center research makes resource consumption relevant to the choice of protective mechanisms [51].

Cloud-based information systems for project governance highlight the coordination and ownership issues involved in operating shared infrastructure [52].

Big data and cloud computing in project management similarly provide an organizational setting for measuring performance and decision consequences [53].

Secure financial analytics offers an adjacent setting where a protective decision can carry a direct service or transaction cost [54].

Adversarial robustness for mobile biometric verification supplies a further application in which a security intervention must be evaluated alongside its effect on legitimate users [55].

Across these themes, the integration problem is consistent: protective mechanisms compete for time, compute, communication, and organizational attention. The framework does not claim that one architecture dominates every alternative. It makes the competing objectives explicit so that comparative testing can reveal when coordination is worthwhile and when its overhead is unjustified.

4. PROPOSED ARCHITECTURE

4.1 System model and trust boundaries

Consider K tenants operating cloud services and edge devices. Tenant k maintains a local observation stream and a set of protected assets. A tenant can participate in collaborative learning without granting the coordinator unrestricted access to its records. Local administrators control operational policy within their own boundaries. A shared coordinator distributes approved model versions, receives permitted updates, and proposes cross-tenant learning changes. It does not acquire authority to execute arbitrary tenant actions.

The threat model includes credential theft, compromised endpoints, malicious network traffic, unauthorized lateral movement, malicious model updates, replayed telemetry, and inference attempts against shared learning artifacts. The coordinator may be curious or become unavailable. Some participating clients may be malicious. A compromised signing key or a malicious authorized tenant administrator falls outside any guarantee that depends on that identity remaining trustworthy; such events must still be represented in incident handling. The design assumes authenticated channels, operational key management, and a locally enforceable policy engine. It does not assume honest telemetry merely because a device is enrolled.

The architecture has five layers: local evidence collection, contextual threat estimation, protected collaboration, constrained action selection, and governance with recovery. Interfaces between these layers carry explicit identifiers and timestamps. A prediction record includes the tenant, asset, evidence cutoff, feature-schema version, model version, calibration version, estimated threat probability, uncertainty indicator, and applicable policy. An action record adds the selected intervention, expected affected services, authorization basis, expiration time, and rollback procedure.

4.2 Local evidence and contextual estimation

Collectors transform packet, identity, endpoint, configuration, and workload events into a common local schema. Sensitive identifiers remain tenant-local unless a permitted transformation exists. The system preserves raw evidence according to the tenant's retention policy and records the transformation from evidence to features. Late-arriving events may update an investigation, but they cannot retroactively improve the prediction credited to an earlier decision.

A local detector estimates p_i, the probability that event or incident i represents an actionable threat within a defined response horizon. A separate calibration stage maps model scores to probabilities using validation data that precede the test period. Uncertainty u_i captures a specified reason for caution, such as ensemble disagreement or distance from the training distribution. Neither measure is interpreted as a universal guarantee of correctness. For missing or unreliable evidence, the system can abstain rather than silently insert a reassuring default.

Contextual features include asset criticality, credential privilege, current exposure, service dependencies, maintenance windows, recent response actions, and the availability of failover. These features enter only when they are available at the decision time. Operational costs are not learned exclusively from attack labels; service owners supply initial estimates and plausible ranges. This separation allows the same detector to support different policies without pretending that a universal threat threshold is appropriate for every tenant.

4.3 Protected collaboration

Tenant k trains or updates a local model using its permitted records. An outbound update is clipped, transformed, and protected according to

an approved collaboration configuration. The configuration states whether secure aggregation, differential privacy, encrypted computation, or an approved combination is used. The coordinator rejects updates with invalid signatures, incompatible feature versions, or expired authorization. A valid signature confirms origin and integrity of the message, not the truth of its contents.

The collaboration design includes a privacy–robustness compatibility check. If a server sees only a sum of updates, a robust rule requiring each individual update cannot simply be added as though the updates were visible. Candidate implementations must either use an aggregate-compatible defense, establish a separate protected validation route, or revise the trust model. The evaluation must measure the resulting protection and overhead rather than listing mechanisms whose assumptions conflict.

Participation is conditional on both security and privacy constraints. A client that has exhausted its approved privacy budget stops contributing under that configuration. A disconnected tenant continues with the last approved local model and policy for a bounded interval. Model release requires a signed manifest, compatibility checks, validation evidence, and rollback instructions. Participation rates and exclusion decisions are monitored because repeatedly excluding small or unusual tenants can reduce the relevance of the shared model to those tenants.

4.4 Dependency-aware response

The response engine operates on a service dependency graph G. Nodes represent assets or services; edges describe a specified dependency type. Each proposed action is mapped to its affected nodes and its expected interruption footprint. The engine can select monitoring, rate limiting, credential restriction, session termination, workload isolation, or analyst review. Destructive or difficult-to-reverse actions require an explicit authorization class.

Before execution, a local enforcement point checks the action against the current policy and the latest observed service state. This rechecks matters because conditions may change between optimization and execution. If the state no longer satisfies the action's preconditions, execution pauses or falls back to a permitted lower-impact action. A ledger entry alone does not authorize execution, and a shared recommendation cannot bypass a tenant's enforcement point.

4.5 Accountable execution and recovery

Every executed action has a time limit, an owner, a reason, and an observable completion state. Recovery is treated as a planned step: reconnecting a workload, restoring a credential, removing a rate limit, or escalating a failed rollback. The audit record stores hashes or references to protected evidence, policy and model versions, the estimated decision costs, the selected action, and subsequent outcomes. Access to the record is itself governed.

The resulting accountability is bounded. The record can establish which approved component produced a decision and which policy allowed it. It cannot establish that an unobserved threat was absent or that an explanation fully describes a complex model's internal reasoning. Investigations therefore combine the recorded decision context with independent evidence and operator testimony where appropriate.

5. FORMAL DECISION MODEL

5.1 Event-level expected loss

Let a belong to a finite set A_i of permitted actions for incident i. Let p_i denote its calibrated threat probability. Let $c_i(a)$ be the expected disruption caused by action regardless of whether the event is malicious. Let $d_i(a)$ be the additional loss of applying that action to a benign event. Let $h_i(a)$ be residual

attack harm when the event is malicious. The expected local loss is defined as:

$$L_i(a) = c_i(a) + (1 - p_i)d_i(a) + p_i h_i(a).$$

The terms have a common unit, such as an agreed operational-loss scale or monetary equivalent. Their construction must avoid counting the same service interruption twice. Where reliable conversion is impossible, the framework retains a vector of objectives and reports a Pareto comparison rather than inventing a precise exchange rate between safety, privacy, and money.

For a binary choice between observation and containment, suppose observation has residual harm H_i and no immediate disruption. Suppose containment costs c_i and reduces residual harm to h_i. If benign-specific damage is zero, containment is locally preferred when $p_i(H_i - h_i) > c_i$. This expression explains why a probability threshold depends on consequence and action effectiveness. It does not establish that those quantities are known precisely. Sensitivity analysis is required wherever estimates materially affect the chosen action.

5.2 Joint constraints

Let z_{ia} equal one when action a is selected for incident i and zero otherwise. Each incident receives exactly one action. Let r_{ia} be its resource use, v_{ias} its interruption contribution to service s, and g_{ia} its required review time. A coordinated batch solves:

Minimize $\Sigma_i \Sigma_a z_{ia} L_i(a) + C_{shared}(z, G)$

subject to $\Sigma_a z_{ia} = 1$ for every i;

$\Sigma_i \Sigma_a z_{ia} r_{ia} \le R$;

$\Sigma_i \Sigma_a z_{ia} v_{ias} \le V_s$ for every protected service s;

$\Sigma_i \Sigma_a z_{ia} g_{ia} \le B$;

$z_{ia} = 0$ whenever the action violates local authorization, privacy, or latency policy.

C_{shared} represents dependency effects that cannot be captured by independent event costs. R is the available response-resource capacity, V_s is a service-specific interruption limit, and B is review capacity over the decision interval. A graph-dependent nonlinear effect must be represented directly or bounded conservatively; it cannot be assumed additive merely for convenience. A practical implementation can use mixed-integer optimization for bounded batches, with a documented timeout and feasible fallback.

5.3 Uncertainty and abstention

Uncertainty can affect both the allowed action set and the cost calculation. An action is not automatically permitted because p_i is high when the evidence is outside the detector's validated operating range. A review option has its own cost, delay, and expected information benefit. Review is selected only when it is feasible and preferable to the remaining actions. If review capacity is exhausted, the system follows a preapproved contingency policy; it does not treat waiting as costless.

For uncertain costs, let Θ be a documented set of plausible consequence parameters. A robust version minimizes the largest modeled loss over Θ. A less conservative implementation can use scenario-weighted loss when scenario probabilities are justified. Both alternatives must report how uncertainty changes the action. Robust optimization cannot rescue a consequence model that omits a major harm or uses an implausible uncertainty set.

5.4 Privacy accounting

When differential privacy is used, adjacency must identify whether protection applies to a record, device, person, or tenant. Record-level protection does not automatically become tenant-level protection. A release accountant tracks the complete sequence of approved mechanisms and its stated composition guarantee. Secure aggregation is recorded separately

because concealing individual updates from the coordinator is not equivalent to bounding information in the released model.

An illustrative outbound mechanism is an update norm bound followed by calibrated noise. This expression is a design placeholder until the sampling procedure, clipping unit, sensitivity analysis, noise distribution, and accountant are specified. No numerical privacy guarantee is claimed in this paper. The prospective study must also assess practical inference attacks under a clearly stated attacker view, while keeping formal and empirical privacy claims distinct.

6. ANALYTICAL RESULTS AND WORKED DECISION EXAMPLE

6.1 Proposition 1: Local optimality does not imply joint feasibility

Assume two incidents have individually beneficial containment actions. Suppose each containment interrupts a service dependency by one unit, while the shared service permits at most one unit of simultaneous interruption. Independent thresholds select both actions because each reduces local expected loss. Their combined interruption is two units, which violates the shared constraint. A coordinated feasible policy selects at most one containment. Thus, independently optimal decisions need not form a feasible joint policy.

This is a constructive feasibility result, not evidence that the proposed optimizer improves measured attack detection. Its value is to identify a failure that a detector-only evaluation cannot reveal. The argument applies whenever responses compete for a shared limit, including failover capacity, analyst attention, or a common dependency. If no shared constraints or interactions exist, the joint problem can reduce to independent decisions and coordination may provide no benefit.

6.2 Proposition 2: Review has conditional decision value

Let L_auto be the smallest expected loss among immediate automated actions. Let c_review include the review cost and delay-induced harm. Let L_after be the expected loss of the action selected after receiving the review information. Review is preferable exactly when c_review + E[L_after] < L_auto, provided review is operationally feasible. This inequality does not assume perfect analysts. It requires the information gained to be valuable enough to offset its cost and delay.

The implication is that abstention is a decision option with measurable consequences, not a generic indicator of responsible automation. Poorly targeted review can increase total loss by consuming scarce capacity or delaying containment. The study must therefore compare selective review with no review and with a simple uncertainty-based allocation rule.

6.3 Proposition 3: Data locality is not a privacy guarantee

Consider a federated client whose update encodes a sensitive local value in a parameter visible to the coordinator. The raw record never leaves the client, but the parameter reveals the value. Therefore, retaining raw data locally does not by itself prevent disclosure through updates. A privacy claim requires additional assumptions and a mechanism that addresses the specified attacker view.

This counterexample is deliberately simple. It establishes a logical limitation of locality, not the leakage rate of a particular federated model. Practical experiments must distinguish update exposure, aggregate exposure, model-query access, and collusion. Reporting that a system is federated is insufficient to identify which of these exposures it controls.

6.4 Illustrative capacity-constrained example

Consider three hypothetical alerts sharing a response capacity of one containment action per interval. For transparency, benign-specific damage is set to zero, and costs are expressed in illustrative loss units. Observation has loss pH. Containment has loss c + ph. All values below are chosen examples, not estimates from a dataset or values reported by a cited study.

Table 1. Illustrative Loss Comparison Under a Capacity-Constrained Containment Policy

Alert	Threat probability p	Uncontained harm H	Residual harm h	Containment cost c	Observation loss	Containment loss	Avoided loss
A	0.90	20	5	8	18.0	12.5	5.5
B	0.60	100	10	15	60.0	21.0	39.0
C	0.30	40	10	12	12.0	15.0	−3.0

A probability-ranked policy contains A because it has the largest probability. Its total loss is 12.5 + 60 + 12 = 84.5 units. The constrained expected-loss policy contains B and observes A and C. Its total is 18 + 21 + 12 = 51 units. Observing all three would total 90 units. Independently containing both A and B is locally attractive but infeasible under the one-action capacity limit.

The example shows a mechanism, not an empirical performance gain: a lower-probability alert can deserve priority when its avoidable harm is greater. The conclusion depends on the assumed costs and containment effectiveness. For B, the advantage over observing B remains positive while its containment cost is below 54 units. Rankings can change when those estimates change. A useful deployment must expose this sensitivity rather than presenting 51 units as a precise prediction of actual loss.

6.5 What the analysis does and does not establish

The propositions establish conditional relationships among decisions, shared constraints, and privacy claims. They do not demonstrate real-world superiority, statistical significance, reduced incident frequency, or improved recovery time. The worked example cannot validate the detector, the graph, the privacy mechanism, or the operational cost estimates. Its role is to make the design's logic inspectable and to identify variables that a later experiment must measure.

7. PROSPECTIVE EMPIRICAL EVALUATION

7.1 Data and experimental units

A future empirical study should combine a documented public traffic benchmark with a controlled multi-tenant cloud–edge testbed. A benchmark can assess record-level discrimination, while a testbed can generate response, service, and recovery observations that a static dataset usually lacks. Exact dataset versions, licensing conditions, schema transformations, attack-generation procedures, and tenant construction rules must be recorded before analysis. Synthetic tenant partitions must be labeled as such and cannot support a claim of demonstrated cross-organization generalization.

The testbed should include distinct identity boundaries, service dependencies, background workloads, and an authorized attack environment. The attack schedule should include common and rare incidents, benign workload changes, compromised-client behavior, and communication outages. Attacks must be confined to the controlled environment. Each incident receives an identifier so that all related records can be grouped during splitting and uncertainty estimation.

7.2 *Chronological and tenant-separated validation*

Training, calibration, policy tuning, and final testing use disjoint periods in that order. A gap between periods is chosen according to the maximum feature lookback and incident duration, not an arbitrary convenient interval. Incident groups, duplicated records, and linked sessions remain within a single partition. Feature selection, normalization, resampling, and hyperparameter search use training data only. Calibration and action thresholds are fixed before the final test begins.

Two complementary tests are required. A temporal test uses future incidents from known tenants. A tenant-transfer test withholds entire tenants from shared training and evaluates the predeclared adaptation procedure separately. The target domain must not contribute labels through early stopping or hyperparameter selection. Reported performance must distinguish these settings rather than combining them into one average.

7.3 *Baselines and equal-information comparisons*

The baseline set should include a rule-based detector, a calibrated tabular classifier, a sequence model where the data support temporal modeling, local-only learning, ordinary federated averaging, and the proposed protected collaboration. All receive the same permissible pre-decision information where a comparison is intended to isolate the model or collaboration mechanism. A centralized model can serve as a technical upper-reference setting only when lawful pooling is assumed explicitly; it is not automatically a deployable baseline.

Response baselines include fixed probability thresholds, probability ranking under the same capacity limit, local expected-loss decisions, and dependency-aware constrained decisions. A factorial comparison separates detection changes from response changes. For example, compare two action policies on identical predictions before combining a new detector with a new optimizer. Otherwise, improved outcomes cannot be attributed to the component claimed to produce them.

7.4 *Outcomes and statistical analysis*

The primary operational outcome is incident-level modeled loss under a preregistered consequence schedule, reported alongside its separate components. This schedule must be independently reviewed and accompanied by sensitivity ranges. Supporting outcomes include harmful-response rate, undetected high-impact incidents, time to detection, time to containment, time to recovery, service availability, analyst workload, and SLA violations. Model metrics include precision–recall area, macro-F1, per-class recall, false-positive rate, Brier score, and calibration error. Class prevalence and alert volume accompany every aggregate score.

Confidence intervals use resampling at the tenant or incident level appropriate to the experimental dependence structure. When tenants and incidents have crossed dependence, the analysis must address both or justify a conservative alternative. Repeated random seeds describe algorithmic variability; they do not create independent production incidents. Paired comparisons use the same incident schedules and workloads. Multiple comparisons are limited or corrected, with effect sizes and uncertainty reported alongside any significance test.

7.5 *Stress tests, ablations, and failure reporting*

Stress tests vary malicious-client prevalence, participation rate, packet loss, feature missingness, class imbalance, attack novelty, graph staleness, and calibration drift. Privacy experiments state the attacker's knowledge, access, auxiliary data, and success metric. Poisoning experiments distinguish label manipulation, update manipulation, and compromised enrollment. Defenses are

evaluated against adaptive attackers who know the relevant protection where feasible.

Ablations remove calibration, the dependency graph, uncertainty-based review, robust update handling, and audit verification one at a time. Privacy mechanisms are compared at matched protection assumptions, with compute and communication overhead measured. Ledger-backed provenance is compared with signed append-only storage under the same audit tasks. Energy use is measured per completed workload or protected service interval; a faster experiment does not automatically imply a more efficient deployment.

Failures are part of the results. The study reports optimizer timeouts, infeasible action sets, unsuccessful rollback, incorrect dependency information, privacy-budget exhaustion, and incidents delayed in the review queue. A proposed framework that succeeds only after discarding these cases has not demonstrated operational reliability.

7.6 Falsifiable hypotheses

H1 predicts that dependency-aware action selection reduces modeled operational loss relative to probability ranking under matched capacity and availability constraints. H2 predicts that protected collaboration improves rare-attack detection on held-out tenants relative to local-only learning without exceeding a declared privacy configuration. H3 predicts that selective review reduces harmful automated responses after accounting for review delay and workload. H4 predicts that provenance verification improves reconstruction of authorized actions under tampered-log tests compared with an otherwise identical unverified log.

Each hypothesis can fail. Coordination may add latency without changing actions. A shared model may disadvantage atypical tenants. Review may be too slow. Provenance may improve traceability without improving attack containment. These possibilities are not exceptions to be hidden; they determine the framework's useful operating range.

8. DISCUSSION

8.1 Contribution beyond model accuracy

The central contribution is an explicit connection between a calibrated estimate and a feasible organizational action. Classification metrics remain useful, but they do not reveal whether a defense interrupts a critical dependency or exhausts the review queue. The worked example explains why action ranking can differ from probability ranking even when the probabilities are correct. The shared-constraint proposition explains why independently beneficial responses can be jointly unacceptable.

This changes the interpretation of a model improvement. An increase in rare-attack recall is valuable when it leads to a timely action with acceptable consequences. If the additional alerts cannot be reviewed or contained, their operational value is uncertain. Conversely, a detector with slightly lower aggregate discrimination may support better decisions if its probabilities are better calibrated in consequential cases. The empirical study must measure these possibilities rather than assuming a fixed relationship between accuracy and protection.

8.2 Privacy, trust, and accountability

Privacy protection should be described as a set of bounded claims. Who sees individual updates? Who sees aggregates? What is released to model users? Which identities can collude? What does a compromised client learn? Answers to these questions determine whether a proposed protection applies. A collection of privacy-related technologies is not a substitute for a coherent threat model.

Accountability is similarly specific. A signed decision record can

support attribution and replay of the policy logic, while access controls protect the evidence it references. It does not prove that a detector was accurate or that a policy was ethically appropriate. Organizations must retain responsibility for setting intervention limits and correcting harmful outcomes. A system that records decisions carefully but cannot reverse avoidable disruption remains operationally weak.

8.3 Adoption and complexity

Implementation should proceed in stages. A first stage can attach calibrated probabilities, evidence cutoffs, and model versions to existing alerts. A second can add service ownership and consequence ranges. A third can evaluate constrained recommendations in shadow mode while operators retain execution authority. Protected cross-tenant learning should be introduced only after its privacy, robustness, and operational assumptions are tested independently. This sequence allows the organization to identify whether the added architecture solves a measured problem.

The framework is not intended to justify maximal technical complexity. An accurate local model with a transparent rule may be sufficient for a small environment with independent services. A signed log may be sufficient when one trusted custodian exists. Homomorphic computation may be inappropriate where its latency cannot satisfy the response horizon. The design's value lies in making these choices testable, including the decision to omit a component.

8.4 Research implications

Future work should examine the quality of operational cost estimates, because the optimizer can be only as useful as the decision model supplied to it. Tenant-specific estimates may be sparse, politically contested, or unstable during an incident. Reporting decisions across plausible cost ranges can be more informative than optimizing a single apparently exact number. Methods for eliciting and updating these ranges deserve their own evaluation.

Another research direction is joint assurance of privacy and malicious-client resistance. Secure aggregation can obscure the evidence needed for client-level filtering. Strong filtering can exclude legitimate minority behavior. A useful comparison must report protection, inclusion, and utility together. The right balance may differ across organizations, which argues for explicit operating envelopes rather than a universal configuration.

9. LIMITATIONS

This paper presents an architecture, conditional analytical results, and a prospective evaluation protocol. It does not report a trained model, benchmark experiment, deployed prototype, human-subject study, or measured operational improvement. Its claims must be read at that level. The numerical example is deliberately small and uses selected illustrative costs; it cannot support general performance claims.

The literature corpus is purposive and contains related author groups, mixed publication types, and incomplete bibliographic records. Its composition may narrow the design perspective. The paper therefore does not claim that a documented research gap is absent from the wider literature. A submission-ready version requires full-text appraisal, verification of all reference metadata, and an updated comparison against independently selected relevant work. Seven remaining provisional author records must be completed; one title also appears truncated.

The architecture depends on usable service dependencies, reliable time alignment, enforceable policies, and meaningful consequence estimates. Stale graphs, compromised collectors, and incorrect costs can undermine its decisions. Its optimization model may also become expensive for large batches with complex dependencies. A feasible fallback and honest timeout

accounting are necessary. Finally, cross-tenant collaboration can create participation and representation problems that are not resolved by average performance metrics or cryptographic protection alone.

10. CONCLUSION

Cyber defense in cloud–edge environments is a coordinated decision problem. Threat estimates must be connected to tenant authority, service dependencies, privacy limits, response capacity, and recovery. The proposed framework formalizes those connections through a layered architecture and a constrained expected-loss objective. Its analytical contribution is to show why local thresholds can violate shared limits, when review can be useful, and why federated data locality does not itself guarantee privacy.

The worked example makes the decision mechanism transparent without presenting selected numbers as experimental evidence. The prospective protocol defines the empirical work required to assess the framework, including chronological and tenant-separated testing, matched-information baselines, rare-attack evaluation, privacy and poisoning tests, service-level outcomes, and failure reporting. Whether the framework improves protection remains an open empirical question. Its immediate contribution is a reproducible specification of what to build, what assumptions to examine, and what evidence would justify deployment.

DECLARATIONS

Data availability: No empirical dataset was generated or analyzed. The worked example is fully specified in Section 6.4.

Code availability: No implemented system or experimental code is claimed in this manuscript.

Funding, competing interests, author contributions, and ethics: These statements must be completed by the actual authors before submission.

REFERENCES

[1] K. S. Mudaliyar and S. S. Kumar, "Intelligent cybersecurity frameworks for data protection in cloud-integrated management information systems," *The Eastasouth Journal of Information System and Computer Science, vol. 2*, no. 2, pp. 305–323, 2024. https://doi.org/10.58812/esiscs.v2i02.1165

[2] P. Chakraborty, H. U. Himel, M. Bashir, M. S. Sikder, H. Kaur, M. T. Bin Ansar, J. Kaur, and M. K. Tuhin, "An intelligent cybersecurity framework for data protection in cloud computing environments," *Spectrum of Engineering Sciences, vol. 2*, no. 3, pp. 657–677, 2024. https://doi.org/10.5281/zenodo.21538159

[3] M. D. Hossain, M. S. Sikder, M. S. Uddin, R. M. Ahsan, B. Uddin, and T. Hossen, "Cognitive cyber defense: AI–MIS integration through big data and cloud frameworks for next-generation digital resilience," *The Eastasouth Journal of Information System and Computer Science, vol. 1*, no. 2, pp. 140–152, 2023. https://doi.org/10.58812/esiscs.v1i02.764

[4] R. M. Ahsan, B. Uddin, T. Hossen, and S. Das, "Resilient intelligence: AI and MIS in the cyber-economic era," *The Eastasouth Journal of Information System and Computer Science, vol. 3*, no. 2, pp. 151–163, 2025. https://doi.org/10.58812/esiscs.v3i02.758

[5] M. S. Uddin, M. S. Sikder, M. M. Anwar, and F. Hossain, "AI-driven cybersecurity and big data-enabled MIS frameworks: Strengthening supply chain integrity, energy resilience, and critical infrastructure protection," *Journal of Computer Science and Technology Studies, vol. 7*, no. 9, pp. 223–232, 2025. https://doi.org/10.32996/jcsts.2025.7.9.26

[6] M. A. Goffer, S. N. Hasan, N. Das, J. Kaur, J. Hassan, C. R. Barikdar, and S. Das, "Cybersecurity and supply chain integrity: Evaluating the economic consequences of vulnerabilities in U.S," infrastructure. *Journal of Management World, vol. 2025*, no. 2, pp. 233–243, 2025. https://doi.org/10.53935/jomw.v2024i4.907

[7] N. Das, H. Kaur, K. B. Siddiqa, S. N. Hasan, P. Chakraborty, J. Kaur, H. Rahman, A. Shan-A-Alahi, and R. Hasan, "AI-driven threat detection and response framework for protecting U.S," critical infrastructure from cyberattacks. *International Cybersecurity Law Review, vol. 7*, no. 2, pp. 147–163, 2026. https://doi.org/10.1365/s43439-026-00169-5

[8] J. Kaur, S. N. Hasan, S. M. Orthi, M. A. Miah, M. A. Goffer, C. R. Barikdar, and J. Hassan, "Advanced cyber threats and cybersecurity innovation—Strategic approaches and emerging solutions," *Journal of Computer Science and Technology Studies, vol. 5*, no. 3, pp. 112–121, 2023. https://doi.org/10.32996/jcsts.2023.5.3.9

[9] S. N. Hasan, P. Chakraborty, M. T. Bin Ansar, M. K. Tuhin, M. A. Siam, J. Kaur, J. Hassan, and C. R. Barikdar, "Embedding cybersecurity risk management into information systems governance: A national-scale framework for

organizational resilience, economic stability, and critical infrastructure protection," *International Journal of Applied Mathematics, vol. 39*, no. 1s, 2026. https://doi.org/10.12732/ijam.v39i1s.1623

[10] S. M. Orthi, P. Chakraborty, M. A. Siam, A. Shan-A-Alahi, A. Al Zaiem, S. N. Hasan, J. Kaur, F. Mahmud, and M. A. Goffer, "AI-driven cyber threat intelligence as a management information system: Integrating cybersecurity governance and IT project management for organizational resilience," *The Eastasouth Journal of Information System and Computer Science, vol. 1*, no. 2, pp. 194–213, 2023. https://doi.org/10.58812/esiscs.v1i02.873

[11] F. Mahmud, C. R. Barikdar, J. Hassan, M. A. Goffer, N. Das, S. M. Orthi, J. Kaur, S. N. Hasan, and R. Hasan, "AI-driven cybersecurity in IT project management: Enhancing threat detection and risk mitigation," *Journal of Posthumanism, vol. 5*, no. 4, pp. 23–44, 2025. https://doi.org/10.63332/joph.v5i4.974

[12] A. Shan-A-Alahi, M. Mustafizur, K. M. R. Hossan, A. Al Zaiem, and M. M. Rahman, "Cybersecurity training and its influence on employee behavior in business environments," *Computer Fraud and Security, vol. 2024*, no. 12, 2024. https://doi.org/10.52710/cfs.689

[13] S. N. Hasan, J. Hassan, C. R. Barikdar, P. Chakraborty, U. Haldar, M. A. R. Chy, E. Rozario, N. Das, and J. Kaur, "Enhancing cybersecurity threat detection and response through big data analytics in management information systems," *Fuel Cells Bulletin, vol. 2023*, no. 12, 2023. https://doi.org/10.52710/fcb.137

[14] N. Das et al., "AI-enhanced cyber threat detection: Transforming security frameworks in management information systems," in *2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET)*, pp. 1–6, IEEE, 2025. https://doi.org/10.1109/ICECET63943.2025.11472511

[15] M. K. Tuhin, H. Kaur, H. Imam, A. Hossain, A. Shan-A-Alahi, and P. Chakraborty, "AI-augmented security analytics for real-time threat identification," in *2026 6th International Conference on Intelligent Technologies (CONIT)*, pp. 1–9, IEEE, 2026. https://doi.org/10.1109/CONIT69683.2026.11621644

[16] P. Chakraborty, "Cyber threat prediction using temporal data analysis model," in *2026 6th International Conference on Intelligent Technologies (CONIT)*, pp. 1–10, IEEE, 2026. https://doi.org/10.1109/CONIT69683.2026.11620680

[17] S. N. Hasan, H. Kaur, S. C. Mohonta, K. B. Siddiqa, J. Kaur, U. Haldar, M. A. R. Chy, and M. M. T. G. Manik, "The influence of artificial intelligence on data system security," *International Journal of Computational and Experimental Science and Engineering, vol. 11*, no. 3, pp. 5589–5598, 2025. https://doi.org/10.22399/ijcesen.3476

[18] J. Kaur, M. Prabha, M. Samiun, S. N. Hasan, R. Hasan, H. Esa, M. F. H. Bhuiyan, M. A. Rob, and D. Shahi, "Comparative analysis of transformer and LSTM architectures for cybersecurity threat detection using machine learning," *EAI Endorsed Transactions on AI and Robotics, vol. 4*, 2025. https://doi.org/10.4108/airo.9759

[19] M. Shahiduzzaman, M. S. A. Sourav, M. Kabtia, R. S. Shammah, M. A. Islam, P. Chakraborty, M. M. Rahman, M. Islam, R. B. A. Tonoy, and M. Adnan, "Attention-guided quantum-inspired extreme learning machine with grey wolf optimization for rare-attack detection in IIoT networks," *Discover Artificial Intelligence, vol. 6*, Article 1258, 2026. https://doi.org/10.1007/s44163-026-01704-3

[20] N. Nabi, M. K. Tuhin, M. Bashir, K. Y. Lucky, M. Raihan, and H. Imam, "Continual learning pipelines for detecting insider threats across corporate cybersecurity infrastructures," in *2025 International Conference on Electrical Engineering and Informatics (ICEEI)*, pp. 1–8, IEEE, 2026. https://doi.org/10.1109/ICEEI68459.2025.11330487

[21] H. Kaur, H. U. Himel, M. H. Rijvi, M. S. Sikder, M. Y. Biplob, and H. Imam, "Automated vulnerability assessment using machine intelligence," in *2026 6th International Conference on Intelligent Technologies (CONIT)*, pp. 1–7, IEEE, 2026. https://doi.org/10.1109/CONIT69683.2026.11620731

[22] U. S. Saha, H. Imam, M. Mustafizur, A. Roy, A. A. Zaiem, and M. M. T. G. Manik, "Deep learning based malware image classification using binary image representation for secure management information systems," in *2026 Technology and Innovation Conference (TIC)*, pp. 1–6, IEEE, 2026. https://doi.org/10.1109/TIC68483.2026.11702958

[23] M. Nazmussakib, M. S. Sikder, M. A. K. P. Hemal, A. Al Zaiem, H. U. Himel, and N. Das, "Adversarial machine learning for robust cyber defense systems," in *2026 International Conference on Computing Theory and Wireless Communications (ICCTWC)*, pp. 1–9, IEEE, 2026. https://doi.org/10.1109/ICCTWC68241.2026.11583629

[24] S. N. Hasan, A. Shan-A-Alahi, A. Al Zaiem, H. Kaur, M. T. Bin Ansar, and J. Kaur, "Self-healing cybersecurity systems using RL agents," in *2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT)*, pp. 1–8, IEEE, 2025. https://doi.org/10.1109/ICAFT66710.2025.11452866

[25] A. Shan-A-Alahi et al., "Deep learning-based threat prediction and autonomous response mechanisms for containerized microservices in hybrid cloud deployments," in G.-N. Nguyen, A. Swaroop, & P. Shukla (Eds.), *Proceedings of Fifth International Conference on Computing and Communication Networks: ICCCN 2025* (Lecture Notes in Networks and Systems, Vol. 1859). Springer, 2026. https://doi.org/10.1007/978-3-032-21499-7_42

[26] M. K. Tuhin, M. T. Bin Ansar, M. S. Sikder, H. U. Himel, H. Kaur, H. Imam, and A. Shan-A-Alahi, "Federated machine learning for privacy-preserving cyber threat intelligence in global cloud-integrated MIS platforms," *The Eastasouth Journal of Information System and Computer Science, vol. 2*, no. 2, pp. 284–293, 2024. https://doi.org/10.58812/esiscs.v2i02.1123

[27] N. Das, S. Sultana, M. S. Sikder, H. U. Himel, U. S. Saha, and R. A. M. Rashed, "AI-enhanced privacy preservation using homomorphic federated models," in *2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT)*, pp. 1–8, IEEE, 2025. https://doi.org/10.1109/ICAFT66710.2025.11453096

[28] N. Das et al., "A privacy-preserving federated intrusion detection system leveraging secure multi-party computation across collaborative cloud tenants," in G.-N. Nguyen, A. Swaroop, & P. Shukla (Eds.), *Proceedings of Fifth International Conference on Computing and Communication Networks: ICCCN 2025* (Lecture Notes in Networks and Systems, Vol. 1773). Springer, 2026. https://doi.org/10.1007/978-3-032-14197-2_42

[29] K. S. Mudaliyar and S. Satish Kumar, "TrustCloud-FL: A self-healing and privacy-preserving architecture for secure analytics across multi-tenant hybrid clouds," *European Journal of Applied Science, Engineering and Technology, vol. 4*, no. 5, pp. 151–171, 2026. https://doi.org/10.59324/ejaset.2026.4(5).10

[30] K. S. Mudaliyar and S. S. Kumar, "AEGIS-FL: Auditable federated threat detection for multi-tenant cloud-edge systems," *The Eastasouth Journal of Information System and Computer Science, vol. 4*, no. 1, pp. 101–121, 2026. https://doi.org/10.58812/esiscs.v4i01.1217

[31] S. C. Mohonta et al., "Efficient homomorphic encryption techniques for confidential big-data analytics over public cloud infrastructures: Algorithms, performance, and trade-offs," in G.-N. Nguyen, A. Swaroop, & P. Shukla (Eds.), *Proceedings of Fifth International Conference on Computing and Communication Networks: ICCCN 2025* (Lecture Notes in Networks and Systems, Vol. 1859). Springer, 2026. https://doi.org/10.1007/978-3-032-21499-7_41

[32] S. M. Orthi, M. H. Rahman, K. B. Siddiqa, M. Uddin, S. Hossain, A. Al Mamun, and M. N. Khan, "Federated learning with privacy-preserving big data analytics for distributed healthcare systems," *Journal of Computer Science and Technology Studies, vol. 7*, no. 8, pp. 269–281, 2025. https://doi.org/10.32996/jcsts.2025.7.8.31

[33] H. U. Himel, M. Mustafizur, T. Akter, R. Anzum, Z. A. Tanni, and A. Hossain, "Privacy-aware data protection models for cloud environments," in *2026 6th International Conference on Intelligent Technologies (CONIT)*, pp. 1–10, IEEE, 2026. https://doi.org/10.1109/CONIT69683.2026.11621371

[34] A. Roy, E. Hossain, H. U. Himel, H. Imam, R. A. Md Rashed, and H. Kaur, "Privacy-preserving MIS using blockchain, federated learning and zero-trust models," in A. E. Hassanien, S. Anand, A. Jaiswal, & P. Kumar (Eds.), *Innovative computing and communications: ICICC 2026* (Lecture Notes in Networks and Systems, Vol. 2037). Springer, 2026. https://doi.org/10.1007/978-3-032-30008-9_16

[35] M. T. B. Ansar, M. S. Sikder, M. K. Tuhin, N. Das, P. Chakraborty, and H. Rahman, "Resilient MIS design using cryptographic and cyber-security techniques," in A. E. Hassanien, S. Anand, A. Jaiswal, & P. Kumar (Eds.), *Innovative computing and communications: ICICC 2026* (Lecture Notes in Networks and Systems, Vol. 2037). Springer, 2026. https://doi.org/10.1007/978-3-032-30008-9_24

[36] U. Haldar, S. Sultana, K. B. Siddiqa, E. Rozario, M. A. Miah, H. Rahman, and M. A. R. Chy, "Blockchain-driven access control and compliance auditing framework for federated cloud service providers: Architecture, prototype and evaluation," in G.-N. Nguyen, A. Swaroop, and P. Shukla (Eds.), *Proceedings of Fifth International Conference on Computing and Communication Networks: ICCCN 2025* (Lecture Notes in Networks and Systems, Vol. 1773, pp. 464–487). Springer, 2026. https://doi.org/10.1007/978-3-032-14197-2_41

[37] J. Hassan et al., "Blockchain integration in management information systems a decentralized approach to strengthening cybersecurity and data integrity," in *2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET)*, pp. 1–7, IEEE, 2025. https://doi.org/10.1109/ICECET63943.2025.11472020

[38] S. Bauskar, R. K. Sahoo, S. S. Boda, H. Singhai, M. M. Bakhsh, and M. Adnan, "Privacy-aware big data governance framework using blockchain," in *2025 IEEE International Conference on Emerging Trends in Computing and Communication (ETCOM)*, pp. 1–9, IEEE, 2025. https://doi.org/10.1109/ETCOM66606.2025.11436976

[39] P. Chakraborty, M. A. Miah, M. A. Siam, H. Imam, K. B. Siddiqa, and H. Rahman, "Trustworthy data lakehouse design using federated learning and blockchain," in *2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT)*, pp. 1–8, IEEE, 2025. https://doi.org/10.1109/ICAFT66710.2025.11453041

[40] M. M. Rahaman, M. R. Islam, M. M. R. Bhuiyan, M. M. Aziz, M. M. T. G. Manik, and I. R. Noman, "Empowering sustainable business practices through AI, data analytics and blockchain: A multi-industry perspectives," *European Journal of Science, Innovation and Technology, vol. 4*, no. 2, pp. 440–451, 2024. https://www.ejsit-journal.com/index.php/ejsit/article/view/550

[41] K. B. Siddiqa, H. Rahman, H. Imam, M. T. Bin Ansar, A. Al Zaiem, A. Shan-A-Alahi, and M. M. T. G. Manik, "Assessment of survivability and importance analysis for networks managing intricate traffic flows," *IEEE Communications Standards Magazine*, Advance online publication, 2025. https://doi.org/10.1109/MCOMSTD.2025.3638981

[42] S. R. Varanasi, S. S. S. Valiveti, M. Adnan, M. I. Faruk, M. J. Hossain, and M. M. T. G. Manik, "Cross-domain standardization and secure edge intelligence for real-time digital twin deployments in next-generation communication systems," *IEEE Communications Standards Magazine*, Advance online publication, 2026. https://doi.org/10.1109/MCOMSTD.2026.3662187

[43] M. R. H. Mahin, P. Chakraborty, N. Das, H. Kaur, H. U. Himel, J. Kaur, and A. G. Mohapatra, "Secured and standardized intelligent zero-touch 6G framework for Edge-AI applications," *IEEE Communications Standards Magazine*, Advance online publication, 2026. https://doi.org/10.1109/MCOMSTD.2026.3660159

[44] U. K. R. Gangula, M. A. Miah, K. Mula, M. Dhakan, Q. T. Sadat, and S. Nayak, "A lightweight and secure semantic communication architecture for edge–IoT–6G systems," *IEEE Communications Standards Magazine*, 2026. https://doi.org/10.1109/MCOMSTD.2026.3677038

[45] E. Hossain et al., "AI-driven trust and bio-inspired optimization integration in standardized digital twin frameworks for next-generation," *IEEE Communications Standards Magazine*, 2026. https://doi.org/10.1109/MCOMSTD.2026.3688125

[46] M. A. Siam, K. Y. Lucky, S. N. Hasan, J. Kaur, H. Kaur, M. S. Uddin, and M. M. T. G. Manik, "Cybersecure intelligent sensor framework for smart buildings: AI-based intrusion detection and resilience against IoT attacks," *Sensors, vol. 25*, no. 24, Art. 7680, 2025. https://doi.org/10.3390/s25247680

[47] A. Shan-A-Alahi, M. S. Sikder, H. U. Himel, M. T. Bin Ansar, M. K. Tuhin, and H. Kaur, "Resilient cybersecurity architectures for large-scale distributed systems," in *2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON)*, pp. 1–8, IEEE, 2026. https://doi.org/10.1109/I3CTCON68242.2026.11507768

[48] M. A. Siam, A. Shan-A-Alahi, M. K. Tuhin, E. Hossain, M. Bashir, K. Y. Lucky, S. M. M. Uddin, and A. Al Zaiem, "AI-driven cyber threat intelligence systems: A national framework for proactive defense against evolving digital warfare," *International Journal of Computational and Experimental Science and Engineering, vol. 11*, no. 3, pp. 6126–6140, 2025. https://doi.org/10.22399/ijcesen.3793

[49] M. A. Goffer, M. S. Uddin, J. Kaur, S. N. Hasan, C. R. Barikdar, J. Hassan, N. Das, P. Chakraborty, and R. Hasan, "AI-enhanced cyber threat detection and response advancing national security in critical infrastructure," *Journal of Posthumanism, vol. 5*, no. 3, pp. 1667–1689, 2025. https://doi.org/10.63332/joph.v5i3.965

[50] S. Sultana, M. Uddin, M. A. R. Chy, S. N. Hasan, E. Hossain, H. Kaur, M. N. Khan, and J. Kaur, "AI-augmented big data analytics for real-time cyber attack detection and proactive threat mitigation," *International Journal of Computational and Experimental Science and Engineering, vol. 11*, no. 3, pp. 5639–5647, 2025. https://doi.org/10.22399/ijcesen.3564

[51] M. D. Hossain, M. S. Uddin, M. S. Sikder, T. Hossen, B. Uddin, and R. M. Ahsan, "Green and secure data centers: Balancing energy efficiency with advanced cybersecurity measures," *Journal of Computer Science and Technology Studies, vol. 6*, no. 5, pp. 300–315, 2024. https://doi.org/10.32996/jcsts.2024.6.5.24

[52] F. Mahmud et al., "The role of cloud-based management information systems in enhancing IT project governance and stakeholder collaboration," in J. C. Bansal, P. Jamwal, & S. Hussain (Eds.), *Proceedings of the International Conference on AI and Robotics: AIR 2025* (Lecture Notes in Networks and Systems, Vol. 1629). Springer, 2025. https://doi.org/10.1007/978-3-032-05548-4_1

[53] F. Mahmud, S. M. Orthi, A. S. M. Saimon, M. Moniruzzaman, M. A. Miah, M. K. Ahmed, F. B. Khair, M. S. Islam, and M. M. T. G. Manik, "Big data and cloud computing in IT project management: A framework for enhancing performance and decision-making," *Fuel Cells Bulletin, vol. 2023*, no. 9, 2023. https://doi.org/10.52710/fcb.166

[54] H. Imam, M. Mustafizur, T. Roy, A. Lalafaryan, A. Mkrtchyan, and V. Roy, "Intelligent and secure financial analytics platforms using advanced machine learning," in *2026 IEEE 18th International Conference on Computational Intelligence and Communication Networks (CICN)*, pp. 769–775, IEEE, 2026. https://doi.org/10.1109/CICN70047.2026.11594305

[55] M. Raihan, M. Adnan, M. J. Hossain, K. B. Siddiqa, F. Karim, and S. C. Mohonta, "Adversarial robustness mechanism for safeguarding biometric verification across mobile financial applications," in *2025 International Conference on Electrical Engineering and Informatics (ICEEI)*, pp. 1–7, IEEE, 2026. https://doi.org/10.1109/ICEEI68459.2025.11330502