
Volume 1 Issue 1

Article Number: JBSMR20260005

**A Lightweight Cybersecurity Risk-Scoring Framework for Industrial
Internet of Things Networks**

Abdul Rahim Bin Ridzuan*¹, Thrushna Matharasi², and Nithesh Naik³

¹Institute of Big Data Analytics and Artificial Intelligence, Universiti Teknologi MARA (UiTM),
40450 Shah Alam, Selangor, Malaysia

²University of Bridgeport, Bridgeport, Connecticut 06604, United States of America

³Research and Innovation, European City University, Paris, France

Abstract

Industrial Internet of Things (IIoT) networks join constrained field devices, legacy control protocols, edge gateways, and enterprise services while cyber incidents may produce physical and operational consequences. Security teams therefore need risk prioritization that is timely and intelligible, yet many assessment approaches require extensive modeling, complete evidence, or computationally demanding analytics. This conceptual paper proposes a lightweight, explainable framework that calculates bounded asset and network risk scores using simple arithmetic at an edge gateway or monitoring server. The framework combines asset criticality, vulnerability and configuration weakness, exposure, threat or anomaly evidence, control effectiveness, local propagation potential, and evidence confidence. It explicitly separates measured risk from confidence, applies conservative fallbacks when evidence is missing, and maps 0–100 scores to reviewable response priorities. Original equations define base, mitigation-adjusted, topology-aware, confidence-adjusted, aggregated, and temporally smoothed risk while preserving monotonicity and bounds. A gateway-centered architecture, operational workflow, linear-time scoring algorithm, and symbolic industrial scenario show how the design can support triage without imposing continuous analysis on constrained nodes. The contribution is a transparent synthesis intended for organizational calibration, not a validated standard or a substitute for safety, business-impact, or complete industrial risk assessment. No experiment, simulation, dataset, prototype, or field deployment is reported; expert elicitation and empirical validation remain future work.

Keywords: Industrial Internet of Things; Cybersecurity Risk Assessment; Industrial Control Systems; Edge Security; Vulnerability Assessment; Explainable Security; Risk Scoring

1. Introduction

IIoT connects sensing, actuation, supervision, and analytics across operational technology (OT) and information technology (IT). Its security problem is not merely enterprise confidentiality: loss of integrity or availability can disturb a physical process. Heterogeneous devices, remote locations, bandwidth and energy constraints, and legacy systems complicate protection [1]; industrial communications also impose availability and timing requirements that conventional controls may not preserve [2]. Consequently, an enterprise score dominated by software severity may understate a vulnerable controller's process consequence, while a highly connected but noncritical host may be over-prioritized.

***Corresponding Author:** Abdul Rahim Bin Ridzuan (rahim670@uitm.edu.my)

Received: 25 June 2026; **Revised:** 30 July 2026; **Accepted:** 21 August 2026; **Published:** 30 September 2026

© 2026 The Author(s).

This is an open access article licensed under a Creative Commons Attribution-NonCommercial 4.0 International Licence (CC BY-NC 4.0).

Existing risk frameworks provide useful governance concepts, but IoT complexity and heterogeneity do not fit them without contextual adaptation [3]. Recent surveys describe fragmented IoT risk-management processes [4] and continuing gaps in industrial and consequence-oriented analysis [5]. At the opposite extreme, probabilistic, attack-graph, or machine-learning models may require data, expertise, or compute unavailable for frequent plant-level triage. The problem addressed here is therefore: *how can heterogeneous IIoT evidence be converted into a bounded, explainable priority score without computation on constrained devices or claims of quantitative precision unsupported by evidence?*

The objective is to develop a conceptual, edge-assisted scoring framework. Three research questions guide it: **RQ1**, which IIoT-specific factors are necessary for an actionable asset score; **RQ2**, how can mitigation, topology, and uncertainty be represented with bounded simple arithmetic; and **RQ3**, how can asset scores support network prioritization while remaining auditable? The scope is periodic or event-driven triage inside an identified industrial zone; formal safety assessment, monetary loss estimation, penetration testing, and autonomous remediation are outside scope.

This paper contributes: (1) a modular architecture separating device telemetry from gateway computation; (2) an original monotone model that reports risk and evidence confidence together; (3) a one-hop topology approximation and linear-time aggregation suitable for edge use; and (4) an explainable mapping from factors to response priority. These are design proposals, not empirically validated findings.

2. Background and Related Work

IoT-specific risk factors reported in prior work include organizational behavior, dependency, attack surface, susceptibility, severity, and uncertainty [6]. A vulnerability-oriented process further links weaknesses, vulnerabilities, threat events, and attack scenarios [7]. Network structure is commonly represented more richly: Bayesian models can be derived from attack paths [8]; a review of IoT attack graphs highlights their value and scalability challenges [9]; node-edge scoring correlates vulnerabilities along possible routes [10]; and automated graph construction can combine topology with CAPEC, CWE, and CVE knowledge [11]. Dependency-aware CPS assessment shows why relationships among cyber and physical components matter [12], whereas knowledge-based simulation models propagate cause and effect across complex systems [13].

Time and operational cost create another tradeoff. Dynamic assessment literature emphasizes rapidly changing evidence [14] and the relationship between risk and infrastructure resilience [15]. Edge integration can place heavier monitoring away from field nodes [16]; even basic safeguards must respect constrained-device resources [17]. Intrusion-detection taxonomies show that alerts depend on architecture, layer, and attack coverage [18], while IIoT situation assessment using AHP and XGBoost illustrates the additional data and model path chosen by learning-based work [19]. Threat likelihood is distinct from technical severity, as the data-driven EPSS formulation makes explicit [20]; nevertheless, cyber-risk data are often scarce or uncertain [21].

Dynamic-risk reviews find no single method suitable for every context [22]. Bayesian C2P models [23], ATT&CK-based CPS analysis [24], and OT situational-awareness research [25] provide richer causal reasoning, while dynamic probabilistic assessment may employ Monte Carlo simulation [26]. Such methods are valuable for deep assessment, but their requirements motivate a complementary first-pass score. Table 1 compares representative approaches only on features explicitly documented in those papers.

The precise gap is an edge-calculable IIoT score that jointly includes consequence, weakness, exposure, live evidence, controls, local propagation, and data quality without requiring attack-path enumeration, model training, or complete evidence. The proposed framework fills that narrow triage gap; it does not supersede the richer methods above.

3. Design Requirements and Threat Model

The intended environment is a segmented industrial cell containing sensors and actuators, programmable logic controllers (PLCs), human-machine interfaces (HMIs), engineering workstations, a historian, and an edge gateway. Protected assets include process availability and integrity, safety-relevant control logic, configurations, credentials, telemetry, and the monitoring plane. The gateway collects passive summaries and runs the score; constrained nodes need only expose already available health/configuration state or be observed passively.

The gateway, signed policy, authenticated inventory, and authorized analysts are trusted for calculation, although each evidence item carries a quality value. Device self-reports, unauthenticated discovery, external feeds, and alerts are untrusted until corroborated. Threat actors include remote attackers, malicious or careless insiders, compromised maintenance endpoints, and supply-chain adversaries. Attack surfaces include remote access, edge APIs, wireless links, insecure credentials, firmware, exposed services, removable media, and legacy protocols lacking modern authentication. Assumptions are that a minimally useful inventory and directed communication graph can be maintained, time is

approximately synchronized at the gateway, and asset owners can classify process consequence. Exclusions include unknown physical dependencies not represented in the inventory, quantitative safety integrity, adversary strategy optimization, and guaranteed detection of zero-days.

Table 1: Comparison of representative open-access research.

Study	Environment	Main risk factors	Scoring method	Topology	Resource awareness	Explainability	Limitation relevant here
Kandasamy et al. [3]	IoT/IoMT	Risk vectors, impact, likelihood	Formula and risk rank	No	No	Factor-level	Domain examples and calibrated inputs are needed.
Andrade et al. [6]	General IoT	Organization, behavior, dependency, surface, susceptibility, severity, uncertainty	Factor-based methodology	Attack graphs discussed	Capacity constraint noted	Construct-level	Framework derivation includes expert judgment.
Flores et al. [8]	Smart home IoT	Attack paths and impact	Bayesian network	Yes	No	Causal graph	Smart-home scope; parameters use simulated attacks.
Shin et al. [10]	Enterprise network	CVSS, vulnerability correlation, node/edge risk	Attack-graph scoring	Yes	No	Path and node scores	Requires graph construction and correlation analysis.
Akbarzadeh and Katsikas [12]	CPS/ICS	Dependencies, attack paths, critical components	Dependency-based ranking	Yes	No	Explicit paths	Requires a detailed CPS dependency model.
Phillips et al. [13]	Complex CPS	Threat causes, effects, interdependencies, controls	Knowledge-based simulation	Yes	No	Causal trace	More modeling and simulation than lightweight triage.
Yi and Guo [19]	IIoT network	Security situation hierarchy and attack data	AHP plus XGBoost	Not explicit	Data balancing considered, not node cost	AHP hierarchy	Training data and ML pipeline are required.

4. Proposed Lightweight Risk-Scoring Framework

Fig. 1 presents the original architecture. Device-side activity is limited to normal telemetry and optional status export. The industrial network and edge gateway provide observation points; inventory, vulnerability intelligence, configuration records, topology, and alert evidence feed a scoring engine. A dashboard exposes factor contributions and confidence before an analyst or approved playbook chooses a response.

Table 2 converts these conditions into design requirements. Protocol independence means scoring normalized evidence, not parsing every protocol identically. “Low overhead” refers to the proposed arithmetic, not to the upstream cost of vulnerability scanning or anomaly detection.




Figure 1: Gateway-centered IIoT risk-scoring architecture. Raw or summarized evidence is processed away from constrained nodes; security decisions remain governed and reviewable.

Table 2: Design requirements and mechanisms.

Requirement	Proposed mechanism
Low compute/memory	Additions, multiplications, maxima, and one pass over assets/links at the gateway.
Explainability	Store each factor, weight, control credit, confidence, and band trigger.
Modularity/protocol independence	Evidence adapters emit normalized factor records.
Incomplete evidence	Confidence is reported; missing values blend with conservative policy priors.
Periodic/event updates	Recompute changed nodes and smooth reported scores.
High-consequence priority	Criticality enters base risk and aggregation; low confidence can force review.
Operational compatibility	Passive collection preferred; no continuous scoring on constrained nodes.

The modules are: (1) asset inventory and classification; (2) criticality assessment over safety, production, quality, environmental, and recovery consequences; (3) vulnerability and insecure-configuration assessment; (4) exposure assessment for reachable services, trust zones, and remote paths; (5) threat/anomaly evidence from alerts or observed deviations; (6) verified control effectiveness; (7) a directed one-hop propagation factor; (8) evidence confidence; (9) node score calculation; (10) network aggregation; (11) risk-band interpretation; and (12) response-priority mapping. Table 3 defines the inputs. Values are ordinal-to-continuous policy encodings on $[0, 1]$, not physical probabilities unless their source is explicitly probabilistic.

4.1. Factor Construction and Explanations

An implementation should define a versioned rubric for every factor rather than normalize unlike raw fields together. Criticality should be assigned per operating mode because a maintenance HMI and a production PLC may exchange priority when the process state changes. Vulnerability assessment may combine the most consequential applicable weakness with a capped contribution from additional independent weaknesses; summing every CVE would reward duplicated records and device age. Exposure should distinguish an externally reachable service, a cross-zone management path, and an observed local flow, while respecting direction and firewall policy. Threat evidence should decay with age and increase only when corroboration, specificity, and asset identity are adequate. These transformations are organizational design choices and must be displayed with their source records.

Control effectiveness requires evidence of operation, not a control catalogue checkbox. Preventive controls (hardening and authentication), detective controls (monitoring), containment (segmentation), and recovery (tested restoration) may be combined through a documented rubric, but overlapping controls should not receive repeated credit. Expired tests reduce the corresponding confidence before they reduce the control value; this distinction shows whether the concern is an ineffective control or simply stale assurance. Likewise, a vulnerability feed can be authoritative yet have poor coverage of an unsupported legacy controller.

The explanation object for asset i contains the raw observations, normalized values, per-source confidence, priors used for missing evidence, weight contributions $w_x x_i$, control reduction ρM_i , reachable critical neighbors contributing to P_i , and the rule producing the response category. Two concise counterfactuals are useful: the smallest verified control improvement that would cross the next lower band, and the missing evidence item whose confirmation would most change confidence. These are arithmetic what-if statements, not causal claims. An analyst can therefore challenge a factor or policy assumption without treating the final score as ground truth.

Table 3: Risk-factor definitions, ranges, and interpretations.

Factor and range	Interpretation of larger values	Candidate evidence
C_i Asset criticality [0, 1]	Greater safety, process, environmental, quality, or recovery consequence	Asset owner classification, process/safety records
V_i Vulnerability/configuration [0, 1]	More severe/exploitable weaknesses or insecure settings	Vulnerability database, firmware mapping, configuration audit
E_i Exposure [0, 1]	More reachable services, remote paths, or weak zone boundaries	Passive flows, firewall rules, remote-access and service inventory
T_i Threat/anomaly evidence [0, 1]	Stronger, fresher, corroborated hostile or abnormal activity	IDS alerts, authentication events, anomaly summaries, threat intelligence
P_i Propagation potential [0, 1]	Direct reachability to more critical adjacent assets over usable links	Directed topology and permitted communication graph
M_i Control effectiveness [0, 1]	Stronger verified prevention, detection, containment, and recovery controls	Segmentation tests, hardening records, monitoring and backup evidence
Q_i Evidence confidence [0, 1]	More complete, fresh, trustworthy, and corroborated evidence	Source provenance, age, coverage, consistency checks

The reporting bands in Table 4 are proposed initial policy categories, not validated thresholds. A low-confidence, high-criticality asset is escalated for evidence collection even if its numerical band is lower.

The update sequence and the one-hop topology interpretation are visualized in Figs. 2 and 3, respectively, and detailed in Sections 6 and 7.

Table 4: Proposed score bands and response priorities.

Score	Category	Recommended priority mapping
0–24	Low	Routine monitoring; verify on scheduled review.
25–49	Moderate	Plan remediation; confirm weak evidence.
50–74	High	Expedited analyst review and compensating controls.
75–100	Critical	Immediate governed response; assess safe containment.

5. Mathematical Risk Model

For asset i , let $\hat{X}_i$ be the observed value of $X \in \{V, E, T\}$, q_{Xi} its source confidence, and π_X a conservative organizational prior. Missing or weak evidence is imputed transparently:

$$X_i^* = q_{Xi}\hat{X}_i + (1 - q_{Xi})\pi_X. \quad (1)$$

Every term lies in $[0, 1]$, hence so does X_i^* . A missing observation sets $q_{Xi} = 0$ and yields π_X , rather than zero. Priors are policy assumptions recorded with the score; they are not claimed frequencies.

The base asset risk is

$$B_i = w_C C_i + w_V V_i^* + w_E E_i^* + w_T T_i^*, \quad \sum w_x = 1, \quad w_x \geq 0. \quad (2)$$

Thus $B_i \in [0, 1]$ and increases monotonically in each contributing factor. Criticality is included additively so a low threat signal cannot erase high consequence. The factors are operationally separated: V describes intrinsic weakness, E reachable opportunity, and T current hostile evidence.

Verified controls reduce, but need not eliminate, the base risk:

$$A_i = B_i (1 - \rho M_i), \quad 0 \leq \rho < 1. \quad (3)$$

Here M_i is composite control effectiveness and ρ is the maximum allowed control credit. Consequently $0 \leq A_i \leq B_i \leq 1$; even $M_i = 1$ leaves residual risk when $\rho < 1$. Only implemented and tested controls receive credit.

Let $\mathcal{N}_i$ contain assets directly reachable from i along an allowed or observed directed link, and let $a_{ij} \in [0, 1]$ denote link usability for propagation. A lightweight topology factor and topology-aware risk are

$$\begin{aligned} P_i &= \begin{cases} \frac{\sum_{j \in \mathcal{N}_i} a_{ij} C_j}{\sum_{j \in \mathcal{N}_i} a_{ij}}, & \sum a_{ij} > 0, \\ 0, & \text{otherwise,} \end{cases} \\ G_i &= A_i + (1 - A_i) \beta P_i, \quad 0 \leq \beta \leq 1. \end{aligned} \quad (4)$$

P_i is the weighted mean criticality of directly reachable neighbors, not an attack probability. Equation (4) is bounded because it fills only the remaining distance to one; it never lowers risk. The one-hop approximation avoids path enumeration but cannot represent all multi-stage attacks.

Evidence confidence is aggregated as $Q_i = \sum_X \nu_X q_{Xi}$, where $\nu_X \geq 0$ and $\sum_X \nu_X = 1$. With a conservative fallback $\pi_i^R \in [0, 1]$, the reported pair is

$$\begin{aligned} U_i &= \max\{G_i, \pi_i^R\}, \\ R_i &= 100 [Q_i G_i + (1 - Q_i) U_i], \quad K_i = 100 Q_i. \end{aligned} \quad (5)$$

$R_i \in [0, 100]$ is risk and $K_i \in [0, 100]$ is confidence. Lower confidence cannot reduce R_i because $U_i \geq G_i$. A policy may make π_i^R monotone in criticality, but its calibration remains explicit.

For n assets, network risk combines systemic concentration and weighted average exposure:

$$R_{net} = \lambda \max_i R_i + (1 - \lambda) \frac{\sum_{i=1}^n s_i R_i}{\sum_{i=1}^n s_i}, \quad 0 \leq \lambda \leq 1, \quad s_i > 0. \quad (6)$$

The parameter λ controls sensitivity to the worst asset; s_i are governance weights, commonly linked to criticality. The convex combination stays in $[0, 100]$. It is an index for prioritization, not expected loss.

To damp measurement noise while retaining event responsiveness,

$$\tilde{R}_i(t) = \eta_i(t) R_i(t) + [1 - \eta_i(t)] \tilde{R}_i(t-1), \quad 0 < \eta_i(t) \leq 1. \quad (7)$$

The bounded convex update remains in $[0, 100]$; a governed event rule can raise η_i for a confirmed exploit or topology change. Smoothing must not delay safety-relevant escalation, for which the unsmoothed score is also displayed.

To reduce double counting, each evidence field has one primary factor, P_i contains only downstream criticality and link usability (not local exposure), and control evidence is excluded from V_i . Organizations should inspect factor correlations and merge or down-weight overlapping inputs. Weights w_x , ν_X , ρ , β , λ , priors, and thresholds should be selected through structured expert elicitation involving OT, safety, operations, and security staff, constrained as above, then tested by sensitivity analysis and future cases. No weights are asserted to be optimal.

6. Framework Workflow and Implementation Considerations

Fig. 2 shows the operational cycle: discovery and evidence collection precede normalization; factor and confidence records feed node scoring; aggregation and priority mapping support review; accepted changes update the inventory and policy.




Figure 2: Risk-scoring workflow and governed feedback cycle.

Table 5: Lightweight IIoT scoring.

Step	Lightweight scoring operation
1	Load policy, inventory A , and directed links L .
2	For each asset $i \in A$, collect evidence and provenance.
3	Normalize C_i , $\hat{V}_i$, $\hat{E}_i$, $\hat{T}_i$, M_i ; assign q_{Xi} .
4	Compute X_i^* , B_i , and A_i using (1)–(3).
5	Scan outgoing links once; accumulate $\sum a_{ij}$ and $\sum a_{ij}C_j$.
6	Compute P_i , G_i , R_i , K_i ; retain factor contributions and flags.
7	Compute R_{net} , map bands, sort assets by governed priority.
8	Smooth display score; present raw score, confidence, and rationale for review.

With adjacency lists, steps over assets cost $O(n)$ and the topology scan costs $O(m)$ for m directed links; sorting is $O(n \log n)$, or $O(n)$ when only four band queues are required. Scoring memory is $O(n + m)$ at the gateway. Incremental operation can recompute a changed asset and its in-neighbors, but the stated bound assumes a full refresh. Devices execute no risk model.

Implementation can combine a passive network sensor, vulnerability-to-firmware mappings, configuration-management records, an asset inventory, topology from observed flows and engineering documentation, and existing alert/anomaly summaries. Active probing should be separately safety-approved. Adapters should timestamp and sign records where feasible, preserve raw identifiers for audit, and avoid turning absence of traffic into evidence of safety. The dashboard should show factor values, priors, control credit, topology neighbors, score deltas, and the rule that set the response band.

Table 5 illustrates how provenance governs interpretation. A record is not discarded merely because confidence is low; it is blended through (1), flagged, and routed for collection. Contradictory sources should remain visible rather than being silently averaged. Update triggers include a new vulnerability mapping, configuration drift, changed remote-access rule, topology edge, high-confidence alert, control test, process-mode transition, or analyst correction. A signed audit entry should retain the old and new factor vectors, policy version, actor, and reason. Rollback restores the policy and inventory state, not an attacker-controlled device state.

Table 6: Evidence sources, quality concerns, and handling.

Evidence source	Factors	Quality-aware handling
Asset and configuration records	C, V, M	Check owner, scope, version, and age; flag unsupported firmware mappings.
Passive flow and rules	E, P	Preserve direction and observation window; reconcile documentation conflicts.
Alerts/threat feeds	T	Require asset identity, freshness, specificity, and corroboration.
Control tests	M, Q	Credit only verified scope; expire confidence when assurance becomes stale.

7. Conceptual Application Scenario

Consider the architecture in Fig. 3: an industrial sensor reports to a PLC; an HMI supervises the PLC; an engineering workstation can modify logic; a historian receives process data; and an edge gateway bridges monitoring services. This is a symbolic scenario, not a deployment, incident, dataset, simulation, or measured network.

Qualitatively, a PLC with high process criticality, a known severe weakness, weak segmentation, and reachable engineering access receives high C , V , E , and P contributions. The engineering workstation may rank similarly despite lower direct process criticality because it can propagate to the PLC. Direct external exposure raises the gateway’s E ; strong authentication, segmentation, and verified monitoring increase M and reduce its adjusted score. A sensor with limited reach and effective compensating controls may rank lower, but stale firmware inventory lowers Q and invokes the conservative fallback. The historian may have high data impact yet lower propagation if flows are constrained one way. Each ordering follows visible factor differences; no numerical value or ranking is presented as observed or validated.

The scenario also illustrates event-driven reasoning. A newly mapped workstation vulnerability changes its V record, while verified blocking of the engineering-to-PLC path lowers P without pretending that the vulnerability disappeared. A widened gateway remote-access rule changes E before an alert occurs. An uncorroborated anomaly can raise T while lowering identity confidence; policy determines whether the result triggers collection or response.

Response priority must remain process-aware. For the PLC, “immediate” need not mean disconnecting a running controller; it may mean restricting an upstream path, increasing observation, and scheduling a safe change window with OT and safety owners. The same historian category might permit isolation, while a low-confidence sensor may first require inventory confirmation. A playbook still checks process state, fail-safe behavior, recovery readiness, and authorization.

8. Discussion

The framework’s expected advantages are traceability, computational simplicity, and deployability alongside an IIoT edge gateway. Every score can be decomposed into policy weights and evidence. The model distinguishes weakness, reachable opportunity, active threat, control credit, and propagation instead of collapsing them into an opaque output. It can therefore complement governance frameworks, vulnerability-management processes, and deeper attack-path or safety analyses by deciding where scarce analyst attention should begin.

8.1. Relationship to Existing Methods and Organizational Use

The framework occupies a deliberately limited layer between an asset register and a full risk assessment. Compared with a flat severity list, it adds local consequence, exposure, verified controls, and directed adjacency. Compared with attack graphs, it gives up multi-stage path completeness in exchange for a single edge scan. Compared with Bayesian or dynamic probabilistic models, it does not require conditional-probability tables, distribution fitting, or repeated simulation. Compared with machine-learning situation assessment, it needs no training set and exposes the sign and size of every contribution. These are tradeoffs, not evidence of superiority: a detailed graph or probabilistic study should replace the lightweight approximation when the decision justifies its cost.




Figure 3: Topology-aware conceptual example. A reachable engineering workstation can raise priority because its directed neighborhood includes a high-criticality PLC.

Operationally, asset owners can use factor sheets during change review; vulnerability teams can sort remediation by process context; and defenders can examine unexpected high- P bridges or low-confidence areas. The network score in (6) can summarize a zone, but cross-plant comparison requires aligned rubrics, weights, coverage, and operating modes. A decreasing index is not proof of reduced incident probability, especially if inventory coverage has fallen.

8.2. Sensitivity, Manipulation, and Human Review

Sensitivity is greatest when a large weight multiplies an uncertain factor or an asset lies near a band threshold. Review should vary policy parameters, examine rank reversals, and use joint scenarios for correlated inputs. If small defensible changes produce different actions, the output requires judgment, not another decimal place.

The evidence pipeline is also an attack surface. A compromised device could suppress self-reported weaknesses; a misconfigured collector could omit flows; an asset owner might overstate controls to lower a queue position; and an alert flood could inflate threat evidence. Mitigations include source authentication, separation of duties for policy and asset changes, passive corroboration, append-only audit records, rate limits, and alerts on abrupt factor or confidence changes. These mechanisms protect scoring integrity but do not make evidence trustworthy by default. Analysts must be able to override a priority with a recorded reason, and subsequent validation should study both the override and the original score rather than treating the override as error.

These advantages are conditional. Weight selection can reverse close rankings; incomplete inventories hide both assets and links; uncertain alerts can inflate or suppress T ; and correlated inputs can create false precision. Attackers or asset owners might manipulate self-reported status, so provenance and independent passive evidence are important. Vulnerability and threat evidence changes rapidly, whereas asset criticality may change during operating modes. A one-hop P misses long paths, shared services, and conditional reachability. Controls may be credited on paper but fail operationally. Thus raw evidence, confidence, and score changes require human review, particularly before containment that could disrupt a physical process. The score supports prioritization; it neither guarantees security nor establishes acceptable residual risk.

9. Limitations and Future Validation

This paper proposes a conceptual framework only. It provides no empirical validation, experiment, field trial, simulation, dataset, prototype, benchmark, accuracy result, overhead measurement, or evidence that the proposed bands or weights predict incidents. The normalized factors are organizational judgments unless grounded in a validated source, and the arithmetic does not transform ordinal inputs into true probabilities.

Future work should use structured expert elicitation to define factors and priors; industrial case studies to test usability

and failure modes; sensitivity and correlation analysis; comparison with established qualitative, attack-graph, Bayesian, and dynamic methods; and calibration of thresholds against reviewed decisions rather than convenience. A prototype should measure gateway CPU, memory, update latency, and evidence-collection cost. Longitudinal studies should test ranking stability, alert responsiveness, and analyst overrides across manufacturing, energy, water, transportation, and process industries. Sector-specific validation is necessary before operational reliance.

Validation should first measure disagreement among independent OT, safety, and security experts without treating consensus as ground truth. Historical, de-identified cases could compare priority and explanations while preventing outcome leakage. A prototype should measure scaling and test missing, stale, contradictory, and manipulated evidence; prospective observation should record overrides and adverse operational effects. Thresholds and weights require versioned governance, and no stage establishes universal validity.

10. Conclusion

IIoT risk prioritization must account for cyber-physical consequence, heterogeneous weakness and exposure, live evidence, controls, topology, and uncertainty without burdening constrained devices. This paper developed an edge-assisted framework that produces bounded asset and network scores with simple, explainable arithmetic and a separate confidence indicator. Its main contribution is a transparent synthesis linking factor definitions, conservative missing-evidence handling, local propagation, aggregation, smoothing, and response mapping. The framework may help organize analyst attention, but it is not a standard, complete industrial risk assessment, or validated predictor. Expert calibration, prototype measurement, sectoral case studies, and longitudinal comparison are required before practical adoption.

Author Contributions

Abdul Rahim Bin Ridzuan (A.R.B.R.): Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Supervision; Thrushna Matharasi (T.M.): Conceptualization, Writing – review & editing; Nithesh Naik (N.N.): Writing – review & editing, Validation.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data Availability

No datasets, experiments, or empirical materials were generated or analysed in this conceptual study.

Declaration of Generative AI and AI-Assisted Technologies in the Writing Process

AI-assisted tools were used selectively to accelerate early-stage drafting and improve the efficiency of manuscript preparation. All AI-assisted content was subsequently reviewed, edited, and verified by the authors, who remain fully responsible for the accuracy, originality, interpretation, and final content of the manuscript.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- [1] L. L. Dhirani, E. Armstrong, and T. Newe, “Industrial IoT, cyber threats, and standards landscape: Evaluation and roadmap,” *Sensors*, vol. 21, no. 11, p. 3901, 2021.

- [2] T. Gebremichael, L. P. I. Ledwaba, M. H. Eldefrawy, G. P. Hancke, N. Pereira, M. Gidlund, and J. Åkerberg, "Security and privacy in the industrial internet of things: Current standards and future challenges," *IEEE Access*, vol. 8, pp. 152 351–152 366, 2020.
- [3] K. Kandasamy, S. Srinivas, K. Achuthan, and V. P. Rangan, "IoT cyber risk: A holistic analysis of cyber risk assessment frameworks, risk vectors, and risk ranking process," *EURASIP Journal on Information Security*, vol. 2020, no. 1, p. 8, 2020.
- [4] E. K. Parsons, E. Panaousis, G. Loukas, and G. Sakellari, "A survey on cyber risk management for the internet of things," *Applied Sciences*, vol. 13, no. 15, p. 9032, 2023.
- [5] T. S. AlSalem, M. A. Almaiah, and A. Lutfi, "Cybersecurity risk analysis in the IoT: A systematic review," *Electronics*, vol. 12, no. 18, p. 3958, 2023.
- [6] R. O. Andrade, S. G. Yoo, I. Ortiz-Garcés, and J. J. Barriga, "Security risk analysis in IoT systems through factor identification over IoT devices," *Applied Sciences*, vol. 12, no. 6, p. 2976, 2022.
- [7] M. Beyrouiti, A. Lounis, B. Lussier, A. Bouabdallah, and A. E. Samhat, "Vulnerability-oriented risk identification framework for IoT risk assessment," *Internet of Things*, vol. 27, p. 101333, 2024.
- [8] M. Flores, D. Heredia, R. Andrade, and M. Ibrahim, "Smart home IoT network risk assessment using bayesian networks," *Entropy*, vol. 24, no. 5, p. 668, 2022.
- [9] O. S. M. B. H. Almazrouei, P. Magalingam, M. K. Hasan, and M. Shanmugam, "A review on attack graph analysis for IoT vulnerability assessment: Challenges, open issues, and future directions," *IEEE Access*, vol. 11, pp. 44 350–44 376, 2023.
- [10] G.-Y. Shin, S.-S. Hong, J.-S. Lee, I.-S. Han, H.-K. Kim, and H.-R. Oh, "Network security node-edge scoring system using attack graph based on vulnerability correlation," *Applied Sciences*, vol. 12, no. 14, p. 6852, 2022.
- [11] F. Ö. Sönmez, C. Hankin, and P. Malacaria, "Attack dynamics: An automatic attack graph generation framework based on system topology, CAPEC, CWE, and CVE databases," *Computers & Security*, vol. 123, p. 102938, 2022.
- [12] A. Akbarzadeh and S. K. Katsikas, "Dependency-based security risk assessment for cyber-physical systems," *International Journal of Information Security*, vol. 22, no. 3, pp. 563–578, 2023.
- [13] S. C. Phillips, S. Taylor, M. Boniface, S. Modafferi, and M. Surridge, "Automated knowledge-based cybersecurity risk assessment of cyber-physical systems," *IEEE Access*, vol. 12, pp. 82 482–82 505, 2024.
- [14] R. M. Czekster, P. Grace, C. Marcon, F. Hessel, and S. C. Cazella, "Challenges and opportunities for conducting dynamic risk assessments in medical IoT," *Applied Sciences*, vol. 13, no. 13, p. 7406, 2023.
- [15] A. A. Ardebili, M. Lezzi, and M. Pourmadadkar, "Risk assessment for cyber resilience of critical infrastructures: Methods, governance, and standards," *Applied Sciences*, vol. 14, no. 24, p. 11807, 2024.
- [16] T. Zhukabayeva, L. Zholshiyeva, N. Karabayev, S. Khan, and N. Alnazzawi, "Cybersecurity solutions for industrial internet of things–edge computing integration: Challenges, threats, and future directions," *Sensors*, vol. 25, no. 1, p. 213, 2025.
- [17] V. A. Thakor, M. A. Razzaque, and M. R. A. Khandaker, "Lightweight cryptography algorithms for resource-constrained IoT devices: A review, comparison and research opportunities," *IEEE Access*, vol. 9, pp. 28 177–28 193, 2021.
- [18] K. Albulayhi, A. A. Smadi, F. T. Sheldon, and R. K. Abercrombie, "IoT intrusion detection taxonomy, reference architecture, and analyses," *Sensors*, vol. 21, no. 19, p. 6432, 2021.
- [19] J. Yi and L. Guo, "AHP-based network security situation assessment for industrial internet of things," *Electronics*, vol. 12, no. 16, p. 3458, 2023.
- [20] J. Jacobs, S. Romanosky, B. Edwards, I. Adjerid, and M. Roytman, "Exploit prediction scoring system (EPSS)," *Digital Threats: Research and Practice*, vol. 2, no. 3, pp. 1–17, 2021.
- [21] F. Cremer, B. Sheehan, M. Fortmann, A. N. Kia, M. Mullins, F. Murphy, and S. Materne, "Cyber risk and cybersecurity: A systematic review of data availability," *The Geneva Papers on Risk and Insurance–Issues and Practice*, vol. 47, no. 3, pp. 698–736, 2022.
- [22] P. Cheimonidis and K. Rantos, "Dynamic risk assessment in cybersecurity: A systematic literature review," *Future Internet*, vol. 15, no. 10, p. 324, 2023.

- [23] X. Lyu, Y. Ding, and S.-H. Yang, “Bayesian network based C2P risk assessment for cyber-physical systems,” *IEEE Access*, vol. 8, pp. 88 506–88 517, 2020.
- [24] A. Amro, V. Gkioulos, and S. K. Katsikas, “Assessing cyber risk in cyber-physical systems using the ATT&CK framework,” *ACM Transactions on Privacy and Security*, vol. 26, no. 2, pp. 1–33, 2023.
- [25] M. N. Nafees, N. Saxena, Á. A. Cárdenas, S. Grijalva, and P. Burnap, “Smart grid cyber-physical situational awareness of complex operational technology attacks: A review,” *ACM Computing Surveys*, vol. 55, no. 10, pp. 1–36, 2023.
- [26] X. Diao, Y. Zhao, C. Smidts, P. K. Vaddi, R. Li, H. Lei, Y. Chakhchoukh, B. K. Johnson, and K. L. Blanc, “Dynamic probabilistic risk assessment for electric grid cybersecurity,” *Reliability Engineering & System Safety*, vol. 241, p. 109699, 2024.