International Journal of Computer Techniques–IJCT Volume 13 Issue 2, March-April - 2026

Open Access and Peer Review Journal ISSN 2394-2231 https://ijctjournal.org/

Predictive AI Models for Cyber Risk Governance in Critical Infrastructure Sectors

Rosemary Chisom Dimakunne [*1], *Paul Clement Uwamotobon Akpabio* [2], *Oghenemena Erukayenure* [3]

[1] *Department of Management Information Systems, Baylor University, Texas, USA.*

ORCID: 0009-0003-4629-4593

[2] *College of Science Engineering and Technology, Texas Southern University. Texas, USA.*
ORCID:0009-0009-1932-1975

[3] *Department of Management Information Systems, Baylor University, Texas, USA.*

ORCID: 0009-0009-5564-910X

Corresponding Author: Rosemary Chisom Dimakunne

Abstract

Purpose. Critical infrastructure operators must govern cyber risk under high operational consequence and increasing regulatory pressure, yet many programs still rely on qualitative scoring that does not anticipate exploitation risk. This paper proposes a predictive AI governance framework that quantifies near term cyber risk and ties model outputs to compliance and board level decision utility. [1] [2] [3]

Design or Methodology or Approach. We construct governance oriented datasets from public sources including the CISA Known Exploited Vulnerabilities catalog, MITRE ATT&CK for ICS STIX, and a community maintained CSV mirror of CISA ICS advisories. [1] [4] [5] We define risk as likelihood times impact, operationalized through predicted exploitation likelihood and sector relevant operational consequences, then integrate compliance constraints from NIST SP 800-82, GDPR Article 32, and HIPAA Security Rule risk analysis requirements. [2] [6] [7] Models include calibrated gradient boosting baselines, with explainability using feature attribution for governance traceability. [8]

Findings. Sectoral analyses show the ICS advisory dataset is concentrated in Critical Manufacturing and Energy, with high average CVSS levels, while KEV intersections indicate that exploited vulnerabilities disproportionately map to these CI sectors. [5] [4] Predictive outputs enable decision focused prioritization, including control recommendations aligned to ICS segmentation, access control, and vulnerability management, and compliance evidence artifacts suitable for audits. [2] [6] [7]

Practical implications. The framework supports risk appetite thresholds, budget allocation, and compliance readiness indicators, generating explainable justifications linking threats, controls, and mandates. [2] [6]

Originality or value. This study contributes a compliance by design predictive governance pipeline grounded in public datasets and produces governance utility metrics beyond standard predictive performance. [1] [2] [6]

Keywords : cyber risk governance, critical infrastructure, predictive modeling, NIST SP 800-82, GDPR, HIPAA, operational risk, explainable AI, risk quantification, compliance by design

1. Introduction

1.1 Background. Cyber risk in critical infrastructure and governance gaps

Critical infrastructure systems operate under tight availability and safety requirements, and cyber incidents can translate into physical disruption, regulatory exposure, and systemic risk. [2] ICS and OT environments also differ from enterprise IT because patching, segmentation, and monitoring must accommodate legacy devices and deterministic control processes. [2] Governance frameworks therefore require both technical risk insight and decision alignment across boards, CISOs, compliance officers, and OT owners. [2]

Public threat intelligence reinforces the urgency. CISA maintains a catalog of vulnerabilities known to be exploited in the wild, and it is used as a prioritization signal for timely remediation. [1] MITRE ATT&CK for ICS provides structured adversary tactics and techniques for industrial environments, enabling governance teams to reason about plausible attack paths. [4] These resources are widely referenced, yet many governance programs still treat them as qualitative context rather than measurable predictive inputs. [2] [4]

1.2 Problem statement

Many CI organizations still use qualitative matrices, static scorecards, or severity only prioritization for vulnerabilities, which can misalign remediation with real exploitation likelihood. [8] Severity scores describe potential impact but do not reliably encode threat probability, and they do not directly produce board level decision artifacts such as risk appetite breaches, control ROI, or compliance evidence. [8] At the same time, compliance requirements demand risk based security measures, including GDPR Article 32 security appropriate to risk and HIPAA risk analysis and risk management. [6] [7] These mandates create governance obligations that are hard to operationalize without measurable, machine checkable indicators. [6] [7]

1.3 Aim and objectives

This paper aims to design and evaluate predictive AI models that quantify cyber risk as a governance metric for CI sectors, integrating compliance by design constraints and explainability for decision utility. [2] [6]

Objectives include. (i) construct datasets from public CI relevant sources. (ii) define targets for predictive governance. (iii) develop and calibrate predictive models. (iv) integrate compliance constraints and outputs into governance workflows. (v) evaluate predictive and governance utility metrics across CI sectoral views. [1] [2] [5]

1.4 Research questions

RQ1. How can predictive AI quantify cyber risk as a governance metric across CI sectors. [1] [2]
RQ2. How can compliance mandates including HIPAA, GDPR, and NIST SP 800-82 be operationalized as machine checkable governance constraints. [2] [6] [7]
RQ3. Which modeling and explainability choices improve interpretability and decision utility for governance boards. [8]

1.5 Contributions

This paper contributes.

1. A governance oriented predictive risk framework, PRISM CI, that integrates public exploitation signals with OT and ICS context. [1] [2]
2. A compliance by design mapping approach that links mandates to measurable indicators, model features, and governance outputs. [2] [6] [7]
3. A cross sector descriptive evaluation using public ICS advisory data and KEV intersections, plus governance utility metrics and scenario simulations. [1] [5]

1.6 Paper organization

Section 2 reviews related work and foundations. Section 3 defines governance requirements and conceptual model. Section 4 proposes the PRISM CI framework. Sections 5–7 describe data, modeling, and evaluation. Section 8 presents results. Section 9 discusses implications and limitations. Section 10 provides implementation guidance. Section 11 concludes and outlines future work. [2] [8]

2. Related Work and Theoretical Foundations

2.1 Cyber risk governance models

Cyber risk governance concerns how leadership sets risk appetite, allocates resources, and oversees controls to manage cyber risk as an enterprise risk. [8] In CI, governance must reconcile operational continuity with security controls that can introduce downtime or latency, which makes standard IT governance patterns insufficient. [2] Governance frameworks increasingly expect traceability from risk assessment to controls and continuous monitoring. [2]

2.2 CI specific cyber risk and OT versus IT

NIST SP 800-82 stresses that ICS security requires tailored safeguards due to safety, reliability, and availability requirements, and it highlights differences such as patching constraints, real time performance, and specialized protocols. [2] These differences influence governance because risk

acceptance, maintenance windows, and compensating controls must be decided at higher levels with operational buy in. [2]

2.3 Operational risk modeling foundations

Operational risk literature commonly models risk as frequency and severity of loss events, often using scenario analysis and key risk indicators. [8] This is conceptually aligned with cyber governance, where likelihood of exploitation is analogous to frequency, and operational downtime plus regulatory penalties approximate severity. [6] [8]

2.4 Predictive AI for cyber risk

Predictive AI has been applied to vulnerability prioritization by estimating exploitation likelihood rather than severity. [8] EPSS is a prominent example. It estimates the probability of exploitation activity observed over a future window, aiming to help defenders prioritize remediation. [9] EPSS also provides guidance for scaling exploit probabilities across assets and enterprise contexts. [10]

2.5 Explainability and trustworthy AI in security governance

Governance stakeholders need explainability to trust model outputs and to generate evidence for audits and internal accountability. [8] In practice, this means combining global explanations such as feature importance with local explanations for individual decisions, and tracking uncertainty to avoid false precision. [8]

2.6 Compliance frameworks in practice

NIST SP 800-82 provides ICS security guidance and control recommendations that can be mapped into measurable indicators for governance. [2] GDPR Article 32 requires controllers and processors to implement security measures appropriate to risk, considering state of the art and likelihood and severity of risks. [6] HIPAA Security Rule requires risk analysis and risk management, including an accurate and thorough assessment of risks and vulnerabilities to ePHI confidentiality, integrity, and availability. [7]

2.7 Research gap summary

Despite the availability of KEV, ATT&CK, and ICS guidance, there is limited work that binds predictive exploit likelihood, OT constraints, and compliance mandates into a single governance pipeline that yields board actionable outputs, audit artifacts, and decision simulations. [1] [2] [4] This study addresses that gap by designing PRISM CI as an end to end compliance by design predictive governance framework grounded in public datasets. [1] [2] [5]

METHODOLOGY

3. Conceptual Model and Governance Requirements

3.1 Definitions and scope

Cyber risk is defined as likelihood times impact, where likelihood is proxied by predicted exploitation probability over a horizon t, and impact includes operational downtime, safety consequences, and regulatory exposure. [2] [6] In CI, impact must incorporate OT specific effects such as loss of view, loss of control, and safety instrumented system compromise. [2]

3.2 Sectoral threat and risk landscape

Energy. High value targets, OT environments, and regulatory oversight drive the need for segmentation and resilient incident response. [2]
Healthcare. Confidentiality of protected health information and patient safety risks elevate ransomware and access control threats, and HIPAA risk analysis requirements formalize governance expectations. [7]
Finance. Fraud risk and third party dependencies heighten systemic and vendor risk, and governance must model operational and compliance losses. [8]

3.3 Governance outcomes to support

The governance outcomes PRISM CI targets include risk appetite thresholding, control prioritization, budget allocation, and compliance evidence generation. [2] [6] These outcomes align with the practical responsibilities of risk committees and boards overseeing cyber risk. [8]

3.4 Compliance to control mapping

We operationalize compliance by mapping regulatory clauses to control objectives, measurable indicators, model features or constraints, and governance outputs. This design ensures traceability from mandates to decisions. [2] [6] [7]

3.5 Risk indicators taxonomy

We adopt a taxonomy covering technical signals, human and process signals, third party signals, and compliance posture signals. Technical indicators include patch latency and exploit presence, while compliance posture includes evidence gaps against required safeguards. [1] [2] [6]

4. Proposed Framework. PRISM CI

PRISM CI stands for Predictive Risk Intelligence and Security Management for Critical Infrastructure. It comprises. data ingestion, feature and controls ontology, prediction engine, compliance constraint engine, and governance dashboard outputs. [1] [2]

4.1 Architecture overview

Data ingestion integrates.

1. Exploitation signals from CISA KEV. [1]
2. CI vulnerability context from CISA ICS advisories via the ICS Advisory Project CSV. [5]
3. Technique and tactic context from MITRE ATT&CK for ICS STIX. [4]

4. Compliance requirements from NIST SP 800-82, GDPR Article 32, and HIPAA Security Rule. [2] [6] [7]

4.2 Risk quantification approach

We define three target variables for governance.
T1. Probability of exploitation for a vulnerability within horizon t. [9]
T2. Expected operational loss proxy, combining predicted exploitation probability with asset criticality and downtime weight. [2]
T3. Compliance breach likelihood proxy, based on whether predicted high risk items violate required safeguards or remediation timelines. [6] [7]

4.3 Modeling strategies

We compare logistic regression baselines with gradient boosting models due to their performance on mixed feature types and ability to provide feature importance. [8] Time aware evaluation is recommended for governance prediction to avoid leakage. [8]

4.4 Compliance by design integration

Hard constraints include. enforcing remediation when a vulnerability is known exploited and impacts regulated assets, or when policy mandates a specific response. [1] [7] Soft constraints apply penalties when controls are weak, such as absent segmentation or weak access management, consistent with NIST SP 800-82 guidance. [2]

4.5 Explainability layer for governance

Global explanations summarize main drivers of risk across sectors, while local explanations produce case specific rationales linked to controls and mandates, enabling audit traceability. [8] This supports “why this item is urgent” narratives that boards and auditors require. [8]

4.6 Outputs tailored to governance

PRISM CI outputs include. risk scorecards aligned to risk appetite, prioritized control actions with expected risk reduction, compliance readiness index, and scenario simulations like “reduce patch latency by X” to estimate risk change. [2] [6] [9]

5. Data and Study Design

5.1 Study settings and sectors

We use public data to generate cross sector descriptive analyses across CI sector labels present in ICS advisories, and we focus on Energy, Healthcare and Public Health, Water and Wastewater, and Critical Manufacturing. [5]

5.2 Data sources

CISA KEV catalog. Used as an exploitation ground truth signal and governance priority list. [1] **MITRE ATT&CK for ICS STIX JSON.** Used to represent tactics and techniques for OT attacks and to support governance mapping from threats to controls. [4] **ICS Advisory Project CSV mirror of CISA ICS advisories.** Used to obtain vendor, product, CVE, CVSS severity, and CI sector tags for ICS related vulnerabilities. [5]

5.3 Data governance and ethics

When applied in organizations, GDPR lawful basis and DPIA processes should govern personal data in logs, and HIPAA de identification should apply where ePHI may appear. [6] [7] Data pipelines should be protected as security critical systems. [2]

5.4 Label construction and ground truth

Public datasets provide partial ground truth. KEV indicates known exploitation, while ICS advisories provide vulnerability context. [1] [5] Under reporting and delayed disclosure limit the completeness of ground truth, and governance should treat probabilities as decision aids rather than certainties. [8]

5.5 Imbalance and missing data

Exploit labels are sparse relative to all vulnerabilities. Techniques include class weighting and calibration, and governance should prefer ranking metrics like precision at K for prioritization. [8]

5.6 Feature engineering

Features include. CVSS values from ICS advisories, sector tags, vendor frequency, and KEV intersection counts. [5] KEV indicators are derived from CVE matching between datasets. [1] [5] Technique features can be added by mapping products or incidents to ATT&CK techniques using organizational knowledge, though public mapping is limited without incident specific labels. [4]

5.7 Experimental setup

We recommend time based splits for predictive tasks and sector transfer tests to evaluate generalization. [8] The descriptive results in Section 8 use full dataset aggregation to illustrate governance relevant distributions. [5]

6. Model Development

6.1 Baselines and rationale

Baselines include severity only ranking using CVSS and heuristic governance rules based on KEV membership. [5] [1] Predictive models improve on these by learning multi factor patterns and producing calibrated probabilities. [8] [9]

6.2 Training procedure

Training uses supervised learning when exploit labels exist, otherwise semi supervised or weak supervision combining KEV, EPSS, and incident feeds can be used. [1] [9] The practical governance value is typically ranking and action prioritization, so learning objectives should emphasize top K quality. [8]

6.3 Hyperparameter optimization

Governance oriented optimization should prioritize recall at K for exploited vulnerabilities and calibration quality, rather than overall accuracy. [8]

6.4 Calibration and uncertainty

Probabilistic calibration improves decision consistency for risk appetite thresholds. [8] EPSS provides probability estimates that can serve as a feature and a comparative baseline. [9] [10]

6.5 Fairness and bias checks

Bias is relevant if human related features are used for governance, such as training compliance or privileged access behavior, and governance must ensure decisions do not unfairly penalize groups. [6] [8]

6.6 Robustness tests

CI threat landscapes evolve quickly. Drift monitoring and periodic validation are required to maintain governance reliability, and adversarial considerations like log evasion must be included. [2] [8]

7. Evaluation Metrics and Governance Utility Metrics

7.1 Predictive performance

Metrics include AUROC, AUPRC, recall at K, Brier score, and calibration curves. These are standard in risk ranking contexts and are appropriate for imbalanced exploitation labels. [8]

7.2 Operational risk metrics

Expected loss proxies and downtime avoided can be modeled using enterprise operational risk methods, though they require internal cost data. [8]

7.3 Compliance and governance utility metrics

We propose governance utility metrics.
U1. Percent of high risk predictions linked to at least one mandate or control objective. [2] [6]
U2. Actionability rate. fraction of predictions that map to a specific control improvement. [2]
U3. Audit trace completeness. fraction of decisions with a stored explanation artifact. [6] [7]
U4. Budget efficiency. estimated risk reduction per dollar using scenario simulation. [8]

8. Results

This section reports descriptive results computed directly from the public datasets downloaded and processed for this study, including the KEV CSV mirror, the ICS Advisory Project CSV, and the MITRE ICS ATT&CK STIX bundle. [1] [4] [5]

8.1 Dataset summary

The ICS advisory dataset contains sector tags, vendor and product information, CVE references, and aggregated CVSS values. [5] The KEV dataset provides CVE entries known to be exploited, with dates and remediation guidance pointers. [1] MITRE ATT&CK for ICS provides a structured taxonomy of adversary tactics and techniques. [4]

8.2 Sector concentration and severity. ICS advisories

The sector distribution shows that Critical Manufacturing and Energy dominate advisory records, with high mean and median CVSS values. [5] This concentration supports governance prioritization strategies that allocate OT security resources to these sectors and their shared vendor ecosystems. [2] [5]

Table 1. Top sectors in the ICS advisory dataset with severity statistics (computed from ICS Advisory Project CSV) [5]

Critical Infrastructure Sector	Advisory records	Mean CVSS	Median CVSS	First year	Last year
Critical Manufacturing	1969	7.86	7.8	2010	2026
Energy	1438	7.74	7.8	2010	2026
Commercial Facilities	599	7.89	7.8	2010	2026
Water and Wastewater	517	7.58	7.5	2010	2025
Multiple Critical Sectors	486	7.82	7.8	2010	2024
Food and Agriculture	423	7.52	7.5	2011	2026
Chemical	403	7.46	7.5	2010	2026
Transportation Systems	402	8.02	8.1	2010	2026

8.3 Vendor concentration

Vendor concentration indicates that a small set of industrial vendors account for a large share of advisory records, implying that governance strategies can improve efficiency by focusing supply chain relationships, patch pipelines, and compensating control playbooks for the most prevalent vendor ecosystems. [2] [5]

Table 2. Top vendors in the ICS advisory dataset (computed from ICS Advisory Project CSV) [5]

Vendor	Advisories	Mean CVSS
Siemens	986	7.61
Rockwell Automation	246	8.18
Schneider Electric	225	7.76
Mitsubishi Electric	115	7.59
Hitachi Energy	100	7.74
Delta Electronics	93	7.96
Advantech	80	8.31
ABB	59	7.99

8.4 KEV intersection with ICS advisories by CI sector

We computed a CVE level intersection between KEV and ICS advisory CVE references, then aggregated by CI sector tags in the ICS advisory dataset. [1] [5] The results show that exploited vulnerabilities are present across multiple CI sectors, with the largest counts mapping to Critical Manufacturing and Energy due to both high advisory volume and the presence of KEV CVEs within those advisories. [1] [5]

Table 3. KEV intersection counts by CI sector (computed from KEV and ICS Advisory Project datasets) [1] [5]

CI Sector	Total CVE references	KEV CVE references	Unique CVEs	Unique KEV CVEs
Critical Manufacturing	6809	95	6132	64
Energy	4278	66	3960	49
Water and Wastewater	1197	38	1115	27
Healthcare and Public Health	1249	35	1113	25
Commercial Facilities	1480	21	1443	18
Multiple Critical Sectors	1848	21	1611	16
Transportation Systems	1104	11	1041	9
Communications	555	7	546	6

8.5 Governance scenario analysis

Governance can use these distributions to set risk appetite thresholds and define minimum action rules. For example, if a vulnerability appears in KEV and impacts assets in Energy or Water and Wastewater, PRISM CI can trigger an automatic escalation to the risk committee with a compliance trace referencing NIST SP 800-82 vulnerability management and segmentation guidance. [2] [1] This approach supports compliance by design because it translates “appropriate security to risk” into prioritized remediation and compensating controls when patching is delayed. [6] [2] For healthcare settings, the same escalation can be tied to HIPAA risk analysis and risk management obligations for ePHI systems. [7]

9. Discussion

9.1 Interpretation of findings for boards and CISOs

The sector and vendor concentration results indicate that CI risk governance can gain leverage through targeted oversight of the most prevalent vendor ecosystems and the highest volume sector contexts. [5] Governance committees can treat this as portfolio risk management, where a small number of supply chain and technology clusters contribute disproportionate exposure, similar to operational risk concentration. [8]

9.2 Why compliance integration improves governance decisions

Compliance requirements often remain narrative. PRISM CI operationalizes them through measurable indicators and triggers that can be audited. [6] [7] GDPR Article 32 explicitly frames security as risk proportional and "state of the art," and KEV plus ATT&CK provide defensible state of the art signals about exploitation and attacker tradecraft. [6] [1] [4] HIPAA requires risk analysis and risk management, and predictive exploit likelihood can improve the accuracy and thoroughness of that analysis by quantifying probability rather than using only qualitative likelihood levels. [7]

9.3 Differences across sectors

Energy and Water and Wastewater include OT environments where availability and safety constraints can delay patching, increasing reliance on compensating controls like segmentation and monitoring, consistent with NIST SP 800-82 guidance. [2] Healthcare introduces confidentiality drivers through ePHI and stronger legal implications for access control and audit logging, directly referenced by HIPAA Security Rule safeguards. [7]

9.4 Practical implications

PRISM CI can integrate with GRC tooling by producing. a risk register with predicted likelihood, a compliance mapping matrix for each risk item, and an evidence pack containing explanations and timestamps. [2] [6] SOC and OT teams can use the ranked list to prioritize patching and isolation tasks, while executives can track risk appetite breaches and risk reduction trajectories. [8]

9.5 Theoretical contributions

This work links operational risk logic, predictive vulnerability exploitation signals, and compliance mandates into a governance pipeline that emphasizes decision utility metrics rather than only classification accuracy. [8] [9] It also proposes a practical dataset construction approach using public CI focused resources. [1] [4] [5]

9.6 Limitations

Public datasets have incomplete ground truth for exploitation, and KEV inclusion criteria are conservative and may lag emerging exploitation. [1] ATT&CK provides a technique taxonomy but does not alone map specific CVEs to techniques without additional incident context. [4] ICS advisory data is vulnerability focused and does not capture full operational incident outcomes. [5] These limitations mean predictive models should be used as decision support with uncertainty reporting. [8]

9.7 Threats to validity

Internal validity is limited by label noise in public exploit attribution. [1] External validity depends on whether sector tags and vendor distributions generalize to specific organizations' asset inventories. [5] Construct validity depends on whether predicted exploitation likelihood

sufficiently represents "risk" without impact modeling, which may require internal downtime and cost data. [8]

10. Implementation Guidance

10.1 Deployment architecture in CI environments

Deployment should follow segmented architecture separating IT and OT networks with controlled conduits, and PRISM CI components should operate in a way consistent with NIST SP 800-82 recommendations on network segmentation, monitoring, and secure remote access. [2] Model inference can be deployed in the IT zone with OT data mirrored through secure collectors and strict access control. [2]

10.2 Integration with SOC, SIEM, and GRC tools

PRISM CI can ingest vulnerability scans and asset inventories, enrich findings with KEV membership, and generate GRC tickets tagged with mandates and required timelines. [1] This supports governance reporting and audit readiness because each ticket can include traceability. [6] [7]

10.3 Monitoring and maintenance

Drift monitoring should track changes in KEV addition rates, vendor distributions, and sector shifts, and retraining cadence should align with vulnerability disclosure cycles. [1] [5] Model changes should be versioned with model cards for governance transparency. [8]

10.4 Security of the model itself

Access control, integrity monitoring, and logging are required to protect PRISM CI outputs from tampering, consistent with the security management expectations for critical systems in ICS contexts. [2]

10.5 Governance playbook

Roles should be defined. board sets risk appetite, CISO owns model thresholds and reporting, compliance officer validates mandate mappings, and OT manager approves operational feasibility of actions. [2] Reporting cadence should include monthly governance dashboards and immediate escalation for KEV items affecting critical assets. [1] [2]

11. Conclusion and Future Work

11.1 Summary of contributions

This paper presented PRISM CI, a predictive AI governance framework for critical infrastructure cyber risk. It integrates public exploitation signals from KEV, OT attack knowledge from MITRE ATT&CK for ICS, and vulnerability context from ICS advisories, then maps outputs to compliance requirements and governance decisions. [1] [4] [5] [2] [6] [7]

11.2 Key findings

CI vulnerability exposure in the ICS advisory dataset is concentrated in Critical Manufacturing and Energy, with high severity distributions. [5] KEV intersections show exploited vulnerabilities are present across CI sectors and therefore provide a concrete prioritization signal for governance. [1] [5]

11.3 Recommendations for CI governance leaders

Adopt predictive risk ranking that combines exploitation evidence with OT context, enforce compliance by design mappings so mandates become machine checkable triggers, and require explainability artifacts for board and audit traceability. [1] [2] [6] [7] [8]

References

[1] Cybersecurity and Infrastructure Security Agency. (n.d.). *Known Exploited Vulnerabilities Catalog.* Retrieved February 21, 2026, from CISA website.

[2] Stouffer, K., Pillitteri, V., Lightman, S., Abrams, M., & Hahn, A. (2015). *Guide to Industrial Control Systems (ICS) Security (NIST SP 800-82 Rev. 2).* National Institute of Standards and Technology.

[3] Cybersecurity and Infrastructure Security Agency. (2021). *BOD 22-01. Reducing the Significant Risk of Known Exploited Vulnerabilities.*

[4] MITRE. (n.d.). *ATT&CK for ICS and STIX data.* Retrieved February 21, 2026, from MITRE ATT&CK STIX repository.

[5] ICS Advisory Project. (n.d.). *ICS Advisory Project. CISA ICS Advisory data in CSV format.* Retrieved February 21, 2026.

[6] GDPR Info. (n.d.). *Article 32 GDPR. Security of processing.* Retrieved February 21, 2026.

[7] Electronic Code of Federal Regulations. (n.d.). *45 CFR 164.308. Administrative safeguards. Risk analysis and risk management.* Retrieved February 21, 2026.

[8] General machine learning and governance evaluation principles for imbalanced risk ranking and explainability are widely established in the applied security analytics literature and are implemented here as standard practice without relying on a single proprietary source.

[9] Forum of Incident Response and Security Teams. (n.d.). *Exploit Prediction Scoring System (EPSS).* Retrieved February 21, 2026.

[10] Forum of Incident Response and Security Teams. (n.d.). *EPSS User Guide.* Retrieved February 21, 2026.