Ethical, Social, Regulatory, And Sustainability Dimensions of Zero Trust Security in Cloud Computing: A Critical Review

Nwokpuru Samuel Abafu[1], Chinonso Job[2], Onwe, Festus Chijioke[3]
[1]*Department of Computing, Software Engineering, University of Greater Manchester*
[2]*University of greater Manchester, United Kingdom*
[3]*Information Technology Department, University of Port Harcourt, Rivers State, Nigeria*

***Abstract*—Zero Trust Security (ZTS) has emerged as the dominant architectural response to the collapse of the network perimeter in cloud-native, remote-first computing environments. By replacing implicit, location-based trust with continuous, risk aware verification of every user, device, and transaction, ZTS addresses well-documented technical weaknesses of perimeterbased defence. However, the same mechanisms that make ZTS technically effective—persistent behavioural monitoring, automated risk scoring, and dense telemetry collection—introduce a distinct set of ethical, social, regulatory, and environmental tensions that are not yet systematically synthesised in the literature. This paper presents a structured narrative review of these four dimensions. We show that continuous verification, while reducing the attack surface, raises documented concerns about workplace surveillance, employee autonomy, and trust fatigue; that risk- and trust-scoring algorithms inherit the same bias and opacity failure modes long documented in algorithmic risk-assessment literature; that compliance with the General Data Protection Regulation, the Health Insurance Portability and Accountability Act, and ISO/IEC 27001 is necessary but insufficient to resolve these tensions; and that the computational intensity of continuous, AI-assisted verification creates a measurable sustainability cost that the cybersecurity literature has only recently begun to quantify. We synthesise findings from peer-reviewed sources spanning 2010–2026 and conclude with a structured research agenda directed at closing the gap between Zero Trust's technical maturity and its socio-ethical governance maturity.**

***Index Terms*—Zero Trust Security, cloud computing, algorithmic bias, workplace surveillance, GDPR, HIPAA, sustainable computing, cybersecurity ethics.**

I. INTRODUCTION

The dissolution of the traditional network perimeter is among the most consequential shifts in enterprise cybersecurity of the last decade. Cloud adoption, distributed workforces, and the proliferation of personally owned and managed devices have rendered the assumption underlying classical perimeter defence—that everything inside the network boundary can be implicitly trusted—untenable [1]. Zero Trust Security (ZTS) reframes this problem: trust is never granted implicitly by network location or asset ownership, but is instead continuously evaluated from identity, device posture, behavioural context, and assessed risk before every access decision [1], [2].

The architecture was first articulated by Forrester Research and subsequently formalised by the U.S. National Institute N. S. Abafu is with the Department of Computing, Software Engineering. Manuscript prepared for submission; corresponding author email to be inserted by the author prior to submission of Standards and Technology in Special Publication 800-207, which remains the most widely cited normative reference for Zero Trust Architecture (ZTA) design [1]. Industrial deployments at scale—Google's BeyondCorp [3], Microsoft's Conditional Access [4], and government-level adoption guided by the UK National Cyber Security Centre's Zero Trust design principles [5]—have demonstrated that the architecture is technically mature and operationally deployable across organisations of substantially different scale and risk profile.

What is comparatively underexamined is the second-order consequence of how ZTS achieves its security

guarantees. Continuous verification is, definitionally, continuous monitoring: of login behaviour, device telemetry, geolocation, and increasingly, machine-learned behavioural baselines [2], [6]. This creates a distinctive tension. The same instrumentation that allows a Zero Trust system to detect a compromised credential in real time also constitutes one of the most comprehensive employee-monitoring infrastructures an organisation can deploy. Recent systematic and narrative reviews of ZTA have catalogued its technical components, deployment models, and domain-specific applications in considerable depth [7]– [10], but the ethical, social, regulatory, and environmental dimensions of this trade-off remain comparatively fragmented across disparate literatures—human-computer interaction, labour studies, algorithmic fairness, data protection law, and green computing—that rarely cite one another.

This paper addresses that fragmentation directly. We synthesise findings across these four dimensions specifically as they apply to ZTS in cloud computing contexts, and we conclude with a structured agenda for closing the gap between Zero Trust's technical maturity and its governance maturity.

A. Contributions

This review makes three contributions. First, it consolidates ethical, social, regulatory, and sustainability concerns associated with ZTS into a single analytical frame, rather than treating them as isolated footnotes to technical deployment papers. Second, it identifies specific points of tension between ZTS mechanisms (continuous behavioural monitoring, automated trust/risk scoring, and computationally intensive Realtime verification) and established bodies of law and ethics (data protection regulation, algorithmic fairness research, and workplace privacy scholarship). Third, it proposes a structured agenda of open research questions intended to guide future empirical work bridging Zero Trust engineering and Zero Trust governance.

B. Review Methodology

This is a structured narrative review rather than a formal systematic review with PRISMA-style screening counts. Sources were identified through targeted search of IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, and Google Scholar using combinations of the terms "Zero Trust," "cloud computing," "algorithmic bias," "workplace surveillance," "GDPR," "HIPAA," and "sustainable computing," restricted to publications between 2010 and 2026, with priority given to peer-reviewed journal articles, NIST and ISO standards documents, and systematic reviews already published on adjacent sub-topics (e.g. [7] for ZTA implementation challenges broadly, and [11] for SME cybersecurity adoption specifically). Where a systematic review already existed for a sub-topic, this paper draws on its synthesised findings rather than re-screening primary studies, and cites the systematic review directly. Vendor documentation and case material (e.g., Google, Microsoft) are used only to ground real world examples, not as a substitute for peer-reviewed evidence on ethical, social, or regulatory claims.

The remainder of the paper proceeds as follows. Section II positions Zero Trust against the limitations of perimeterbased defence. Section III examines surveillance, privacy, and algorithmic bias concerns. Section IV examines social and organisational impacts, including the disproportionate burden on small and medium enterprises (SMEs). Section V examines regulatory and compliance obligations. Section VI examines the environmental cost of continuous verification. Section VII synthesises these findings into a structured research agenda, and Section VIII concludes.

II. BACKGROUND: FROM PERIMETER DEFENCE TO CONTINUOUS VERIFICATION

Traditional network security operationalises trust geographically: an actor inside the firewall is trusted by default, and security effort is concentrated at the boundary. This model degrades under three converging pressures that the literature treats as largely settled: pervasive cloud adoption, in which data and workloads are routinely distributed across third-party infrastructure outside any single organisational perimeter; bring-your-own-device and remote work norms, which make the boundary itself difficult to define; and the rise of API-driven, microservice-based application architectures, against which static firewall rules provide limited protection [1], [8].

ZTA responds to these pressures with five interlocking mechanisms documented consistently across the literature: granular, least-privilege access policies evaluated per-request rather than per-session; micro-

segmentation, which partitions the network into small enforcement zones to limit lateral movement; continuous identity verification, rather than onetime authentication at login; User and Entity Behaviour Analytics (UEBA), which establishes behavioural baselines and flags deviations; and dynamic, risk-adjusted authorisation, which can step up authentication requirements or revoke access mid-session based on assessed risk [2], [6], [7].

Table I Structural Comparison of Perimeter-Based and Zero Trust Security Models in Cloud Environments

Dimension	Perimeter-Based Model	Zero Trust Model
Network exposure	Trust assumed once inside the boundary	Every request verified regardless of origin
Perimeter	Static, geography-defined	Dissolved; identity is the new boundary
Modern apps	Firewalls poorly suited to APIs/microservices	Micro-segmentation limits lateral movement
Access control	Static roles, vulnerable to credential theft	Continuous, context- and risk-adjusted
Authentication	Session-based, infrequent re-checks	Continuous re-evaluation per request
Threat response	Reactive, slower integration with modern tooling	Real-time monitoring and automated response

Table I summarises the structural contrast between perimeter-based and Zero Trust models as documented across the surveyed literature.

Industrial-scale validation of this model is well documented. Google's BeyondCorp eliminated internal VPN dependency across its global workforce by substituting device- and usertrust evaluation for network-location trust [3]. Microsoft's Conditional Access evaluates sign-in risk using device health, geolocation, and behavioural signal at the scale of billions of authentication events daily [4]. At a policy level, the UK government's adoption of NCSC Zero Trust design principles for public-sector cloud migration [5], and the U.S. Cybersecurity and Infrastructure Security Agency's Zero Trust Maturity Model for federal agencies, demonstrate that ZTA has moved from vendor proof-of-concept to mandated national infrastructure policy. This maturity is precisely what makes the under-examined governance dimension consequential: ZTS is no longer experimental, and its ethical, social, regulatory, and environmental externalities are already being incurred at scale.

III. ETHICAL CHALLENGES

A. Surveillance and Privacy

The defining mechanism of ZTS—continuous behavioural verification—is mechanically indistinguishable from comprehensive workplace surveillance. This is not an incidental side effect; UEBA components are explicitly designed to build behavioural baselines and flag deviation [2]. The broader workplace-surveillance literature, developed largely independently of the Zero Trust security literature, has documented the consequences of this kind of monitoring in considerable depth. Ebert et al. frame algorithmic workplace monitoring as a human-rights-adjacent due-diligence problem, arguing that the shift from human to algorithmic employment decision-making raises legal and ethical risks that organisations frequently underestimate [12]. Van Zoonen et al. show empirically that algorithmic surveillance shapes worker compliance through perceived fairness and trust, rather than through the deterrent effect of monitoring alone, implying that poorly explained or opaque monitoring can be counterproductive even on the organisation's own terms [13].

Applied to ZTS specifically, this literature implies that continuous verification mechanisms which are not clearly communicated to employees—what is collected, for what purpose, and under what retention policy—risk being perceived as covert or punitive rather than protective, with consequences for morale, trust in the employer, and willingness to selfreport security incidents [14]. Gandy's earlier analysis of automated decision systems is directly relevant here: systems whose decision logic and evaluative

benchmarks are opaque to the people they evaluate generate perceptions of unfairness independent of whether the system is, in a narrow technical sense, accurate [15].

B. Algorithmic Bias in Trust and Risk Scoring

ZTS risk engines assign continuously updated trust or risk scores to users and devices, which directly determine access outcomes [2], [6]. This places ZTS risk scoring in the same technical family as the algorithmic risk-assessment systems that have been most thoroughly critiqued in the fairness literature: pretrial recidivism risk scores, credit scoring, and fraud-detection systems. Angwin et al.'s investigative analysis of recidivism risk scoring demonstrated that a system can be approximately accurate in aggregate while being systematically mis calibrated across demographic subgroups, producing disparate false-positive rates [16]. Kleinberg et al. subsequently formalised why this occurs: when base rates of the outcome being predicted differ across groups, no risk score can simultaneously satisfy calibration and equal false positive-rate fairness criteria except in degenerate cases—the trade-off is mathematically inherent, not merely a data-quality defect [17]. Barocas and Selbst's foundational legal-technical analysis extends this point: disparate impact in automated systems typically arises from proxy variables and historical training data, not from explicit discriminatory intent, which makes it both harder to detect and harder to litigate [18].

These results transfer directly to ZTS trust scoring. A risk engine trained on historical access logs will reproduce whatever access patterns, working hours, or device profiles were historically associated with "legitimate" use, with the structural risk that employees whose work patterns deviate from the historical baseline—part-time staff, employees on non-standard schedules, or users of older or shared devices for reasons of economic constraint—are systematically flagged as higher risk independent of actual threat. Wang et al.'s reinforcement-learning-based dynamic access control model illustrates the broader pattern: the model is explicitly optimised for detection accuracy and response latency, with fairness or disparate-impact auditing not specified as a design objective [2]. Mitigating this requires the same toolkit developed in the algorithmic fairness literature—subgroup-disaggregated accuracy auditing, fairness-constrained training, and human review of high-impact automated decisions—explicitly incorporated into ZTS risk-engine validation, which the current ZTS deployment literature does not yet treat as a standard requirement.

C. Professional and Ethical Frameworks

The ACM Code of Ethics and Professional Conduct obligates computing professionals to avoid harm, respect privacy, and ensure that automated decisions affecting people are fair and contestable [19]. Applied to ZTS deployment, this implies concrete organisational obligations: transparent disclosure of what behavioural data is collected and why; a meaningful optout or contestation pathway where feasible; and active auditing of automated access decisions for disparate impact, rather than relying on aggregate accuracy metrics alone [14]. The gap between this normative expectation and current deployment practice is, at present, large: none of the major vendor case studies reviewed in Section II publish subgroup-disaggregated fairness audits of their risk-scoring engines as standard practice.

IV. SOCIAL IMPACTS

A. Employee Autonomy and Organisational Culture

Zero Trust enforcement mechanisms—blocking access from unrecognised IP addresses, disabling sessions on noncompliant devices, requiring step-up authentication mid-task— are, by design, more intrusive on routine work than perimeterbased controls that, once passed, impose no further friction. The literature on algorithmic management of work indicates this friction has a measurable cost to perceived autonomy and morale, even when employees understand and accept the underlying security rationale [12], [13]. This cost is rarely quantified in ZTS deployment case studies, which tend to report security and latency metrics but not user-experience or attrition-risk metrics.

B. Disproportionate Burden on Small and Medium Enterprises

A substantial and growing body of cybersecurity-adoption literature documents that SMEs face structurally different barriers to security adoption than large enterprises: constrained budgets, limited in-house security expertise, and lower perceived threat salience [11]. These constraints apply with particular

force to ZTS, whose components—identity governance platforms, UEBA tooling, dedicated security operations capacity—carry both licensing and skilled-labour costs that scale poorly for smaller organisations [20]. The consequence, documented across the SME cybersecurity literature more broadly [11], is a widening security-capability gap between large organisations that can fully operationalise continuous verification and SMEs that adopt partial or vendor-managed approximations of it, with corresponding differences in residual risk.

C. Trust Fatigue

Repeated authentication challenges and step-up verification prompts impose a cumulative cognitive burden on users—a phenomenon documented in usable-security research as alert or authentication fatigue, in which users habituate to security prompts and either disengage from them or actively seek workarounds. In ZTS deployments with aggressively tuned risk thresholds, this manifests operationally as elevated helpdesk load from denied-access tickets and false-positive step-up challenges, a pattern several of the case studies discussed in Section II report informally without publishing the underlying ticket-volume data. This represents a clear opportunity for future quantitative work, discussed further in Section VII.

V. REGULATORY AND COMPLIANCE ISSUES

ZTS deployments process exactly the categories of data subject to the strictest contemporary data-protection regimes: continuous behavioural telemetry, device identifiers, and geolocation. Compliance is therefore not an optional governance layer but a structural requirement.

The General Data Protection Regulation restricts the legal bases on which personal data, including behavioural and device telemetry, may be processed, and requires that data subjects be informed of automated decision-making that significantly affects them [21]. The Health Insurance Portability and Accountability Act imposes parallel, sector-specific obligations on the handling of protected health information, directly relevant to ZTS deployments in healthcare cloud environments [22]. NIST SP 800-207 provides the technical reference architecture for Zero Trust but is explicitly non-prescriptive on data-protection law, leaving the compliance burden to be discharged separately [1]. ISO/IEC 27001 defines the information-security-management-system requirements, including access-control governance, against which many organisations' ZTS programmes are externally audited [23].

A further complication, documented in cross-border cloud governance literature, is that cloud services routinely span multiple legal jurisdictions, such that the physical location of a data centre, the location of the data subject, and the location of the controlling organisation may each fall under different and potentially conflicting data-protection regimes. This jurisdictional layering means that a ZTS deployment can be simultaneously compliant with one applicable regime and exposed under another, a risk that large cloud providers typically manage through region-specific policy enforcement and data-residency controls rather than through a single global configuration.

VI. SUSTAINABILITY AND COMPUTATIONAL DEMANDS

Continuous verification is computationally expensive by design: every access request triggers identity verification, device posture checks, and, in AI-augmented deployments, inference against behavioural models. This computational profile places ZTS squarely within the broader environmental-cost-of-AI security literature, which has expanded rapidly in the last two years. Deep-learning-based intrusion detection systems have been shown to consume substantial per-inference energy at scale, creating a direct tension between detection accuracy and energy efficiency [24]; this same tension applies to the behavioural-anomaly models embedded in ZTS risk engines. Industry analyses of AI-driven cybersecurity's environmental footprint similarly document that the energy intensity of largescale, always-on security inference is large enough to motivate dedicated mitigation strategies, including carbon-aware workload scheduling and lightweight model architectures [25].

The cybersecurity research community's response to this tension has thus far concentrated on three mitigation strategies, each with a direct analogue in ZTS deployment: event triggered rather than continuous-poll verification, which limits inference

calls to moments of genuine state change; lightweight or distilled behavioural models, which trade a measured amount of detection sensitivity for substantially reduced per-inference energy cost; and serverless or carbon aware scheduling, which defers non-time-critical verification workloads (e.g. periodic re-baselining of behavioural models) to periods of higher renewable-energy availability [1], [26]. None of these mitigations is yet standard practice in ZTS reference architectures, including NIST SP 800-207 itself, which does not address energy or carbon considerations [1]. This represents a clear gap between the security-engineering and sustainable-computing literatures that future ZTS design should close.

VII. DISCUSSION: TOWARD A CLOSED GOVERNANCE LOOP

The four dimensions reviewed above are not independent. A ZTS deployment that addresses surveillance ethics through transparent disclosure (Section III) but does not audit its risk scoring engine for disparate impact has not actually resolved the underlying fairness concern; one that achieves regulatory compliance (Section V) through extensive logging has, by the same act, increased its energy footprint (Section VI) and its surveillance intensity (Section III). Treating these dimensions as a checklist to be satisfied independently therefore risks solving each one in a way that worsens another.

We propose the following structured research agenda, organised around the gaps identified in each section:

1) Subgroup-disaggregated fairness auditing as a standard ZTS validation step. Current ZTS risk-engine evaluation, as exemplified by [2], optimises for aggregate detection accuracy and latency. Future work should report subgroup-disaggregated false-positive and access denial rates across employment categories (e.g., parttime, contractor, non-standard schedule) as a standard validation metric, analogous to the practice now expected in algorithmic-fairness-audited credit and recidivism scoring [16], [17].
2) Quantified trust-fatigue and help-desk-cost metrics. The case studies reviewed in Section II report security outcomes but rarely publish step-up-authentication denial rates or help-desk ticket volume attributable to ZTS friction. Standardised reporting of these operational costs would allow organisations to tune risk thresholds against a genuine cost-benefit frontier rather than security metrics alone.
3) SME-scaled reference architectures. Given the documented adoption gap [11], [20], there is a clear need for a NIST- or ISO-endorsed lightweight ZTS reference profile explicitly scoped to resource-constrained organisations, analogous to existing lightweight-cryptography standardisation efforts.
4) Energy- and carbon-aware extensions to NIST SP 800-207. The reference architecture's silence on computational sustainability [1] should be addressed through a formal extension or companion guidance document, incorporating the event-triggered and carbona ware scheduling strategies already validated in adjacent AI-security sustainability research [24], [26].
5) Jurisdiction-aware compliance tooling. Given the cross-border data flows inherent to cloud-hosted ZTS, future work should examine automated policy enforcement mechanisms that can reconcile simultaneously applicable, potentially conflicting data-protection regimes at the point of access decision, rather than at deployment-time configuration alone.

VIII. CONCLUSION

Zero Trust Security represents a technically mature and increasingly mandated response to the structural security weaknesses of perimeter-based defence in cloud computing environments. This review has shown, however, that the continuous verification mechanisms responsible for ZTS's security effectiveness simultaneously generate four under-synthesised governance challenges: surveillance and privacy tensions arising directly from continuous behavioural monitoring; algorithmic bias risks in trust and risk scoring that mirror well documented failure modes in other algorithmic risk-assessment domains; regulatory compliance obligations under GDPR, HIPAA, and ISO/IEC 27001 that are necessary but, on their own, insufficient to resolve the underlying ethical tensions; and a measurable, currently unaddressed sustainability cost arising from the computational intensity of continuous, AI assisted verification. Closing the gap between Zero Trust's technical and governance maturity will require treating these four

dimensions as jointly optimised design constraints, not independently satisfied compliance checkboxes, and we have proposed five concrete research directions toward that end.

IX. LIMITATIONS

As a structured narrative rather than a fully systematic review, this paper does not report PRISMA-style screening statistics, and the search strategy described in Section II prioritised depth on four specific dimensions over exhaustive domain coverage. Several claims in Sections IV and VI— particularly regarding trust fatigue and help-desk cost—are presented as plausible, literature-consistent inferences rather than as directly measured findings, because the primary studies that would measure them directly do not yet exist; this gap is itself identified as a research priority in Section VII.

REFERENCES

[1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," *National Institute of Standards and Technology (NIST) Special Publication 800-207*, 2020.

[2] R. Wang, C. Li, K. Zhang et al., "Zero-Trust Based Dynamic Access Control for Cloud Computing," Cybersecurity, vol. 8, no. 1, p. 12, 2025.

[3] Google Cloud, "BeyondCorp Enterprise: A New Approach to Zero Trust," 2021. [Online]. Available: https://cloud.google.com/beyondcorp

[4] P. Borra, "Securing Cloud Infrastructure: An In-Depth Analysis of Microsoft Azure Security," International Journal of Advanced Research in Science, Communication and Technology, vol. 4, 2024.

[5] National Cyber Security Centre, "Zero Trust Architecture Design Principles," 2023. [Online]. Available: https://www.ncsc.gov.uk/collection/zero-trust/architecture-design-principles

[6] K. Jena, "Zero-Trust Security Models Overview," International Journal of Scientific Research in Computer Science, Engineering and Information Technology, vol. 9, no. 6, pp. 70–76, 2023.

[7] T. Sasi, A. H. Lashkari, R. Lu, P. Xiong, and S. Iqbal, "A Systematic Literature Review on the Implementation and Challenges of Zero Trust Architecture Across Domains," Sensors, vol. 25, no. 19, p. 6118, 2025.

[8] E. B. Fernandez and A. Brazhuk, "A Critical Analysis of Zero Trust Architecture (ZTA)," Computer Standards & Interfaces, vol. 89, p. 103832, 2024.

[9] M. A. Azad, S. Abdullah, J. Arshad, H. Lallie, and Y. H. Ahmed, "Verify and Trust: A Multidimensional Survey of Zero-Trust Security in the Age of IoT," Internet of Things, vol. 27, p. 101227, 2024.

[10] P. Phiayura and S. Teerakanok, "A Comprehensive Framework for Migrating to Zero Trust Architecture," IEEE Access, vol. 11, pp. 19487–19511, 2023.

[11] A. Alahmari and R. Duncan, "Cybersecurity Risk Management in Small and Medium-Sized Enterprises: A Systematic Review," arXiv preprint arXiv:2309.17186, 2023.

[12] I. Ebert, I. Wildhaber, and J. Adams-Prassl, "Big Data in the Workplace: Privacy Due Diligence as a Human Rights-Based Approach to Employee Privacy Protection," Big Data & Society, vol. 8, no. 1, 2021.

[13] W. van Zoonen, M. E. von Bonsdorff, and B. I. J. M. van der Heijden, "Algorithmic Surveillance and Workers' Compliance: The Role of Trust, Privacy Concerns, and Fairness in Online Crowdwork," New Technology, Work and Employment, 2025.

[14] A. Pigola and F. d. S. Meirelles, "Zero Trust in Cybersecurity: Managing Critical Challenges for Effective Implementation," Journal of Systems and Information Technology, vol. 27, no. 4, pp. 517–564, 2025.

[15] O. H. Gandy, "Engaging Rational Discrimination: Exploring Reasons for Placing Regulatory Constraints on Decision Support Systems," Ethics and Information Technology, vol. 12, no. 1, pp. 29–42, 2010.

[16] J. Angwin, J. Larson, S. Mattu, and L. Kirchner, "Machine Bias: There's Software Used Across the Country to Predict Future Criminals, and It's Biased Against Blacks," ProPublica, 2016.

[17] J. Kleinberg, S. Mullainathan, and M. Raghavan, "Inherent Trade-Offs in the Fair Determination of Risk Scores," in Proc. Innovations in Theoretical

Computer Science (ITCS), 2017, arXiv:1609.05807.
[18] S. Barocas and A. D. Selbst, “Big Data’s Disparate Impact,” California Law Review, vol. 104, pp. 671–732, 2016.
[19] Association for Computing Machinery, “ACM Code of Ethics and Professional Conduct,” 2018. [Online]. Available: https://www.acm.org/code-of-ethics
[20] D. Dang and T. Vartiainen, “Exploring Socio-Technical Gaps in the Cybersecurity of Energy Informatics for Sustainability,” in Adoption of Emerging Information and Communication Technology for Sustainability. Boca Raton, FL, USA: CRC Press, 2024, pp. 288–304.
[21] European Parliament and Council of the European Union, “Regulation (EU) 2016/679 (General Data Protection Regulation),” Official Journal of the European Union, L 119, 2016.
[22] U.S. Congress, “Health Insurance Portability and Accountability Act of 1996,” Public Law 104-191, 1996.
[23] International Organization for Standardization, “ISO/IEC 27001:2022 – Information Security, Cybersecurity and Privacy Protection – Information Security Management Systems,” 2022.
[24] “Integrating Sustainability into Cybersecurity: Insights from Machine Learning Based Topic Modeling,” Discover Sustainability, vol. 6, 2025.
[25] “AI-Driven Cybersecurity Environmental Impact,” Finovista Industry Analysis, 2026. [Online]. Available: https://finovista.com/environmental-impact-of-ai-driven-cybersecurity/
[26] “Green Security: A Systematic Review on the Link Between AI-Driven Cybersecurity and Sustainability,” Authorea Preprints, 2026.