National College of Ireland

Effectiveness of Zero Trust Security Against AI and Quantum Computing Threats: A Systematic Literature Review

MSc Research Project
Programme Name MSc Cyber Security

Hassan Manzoor
Student ID: X23393475

School of Computing
National College of Ireland

Supervisor: Mark Monaghan

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name: Hassan Manzoor

Student ID: X23393475

Programme: Cyber Security **Year:** ……2024,2025……..

Module: Research Practicum

Supervisor: ……

Submission Due Date: 11th Aug, 2025

Project Title: **Effectiveness of Zero Trust Security Against AI and Quantum Computing Threats: A Systematic Literature Review**

Word Count: ……5157…… **Page Count**……25……..

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Hassan Manzoor

Signature:

Date: ……11th Aug 2025……

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Effectiveness of Zero Trust Security Against AI and Quantum Computing Threats: A Systematic Literature Review

Hassan Manzoor
X23393475

Abstract

Zero Trust Architecture (ZTA) has emerged as a transformative shift from perimeter-based security models toward a "never trust, always verify" approach, emphasizing continuous authentication, dynamic authorization, and micro-segmentation. Despite its adoption across enterprise, government, and critical infrastructure sectors, emerging technologies namely Artificial Intelligence (AI) and Quantum Computing pose unprecedented challenges to ZTA's resilience. AI-powered threats such as deepfakes, adversarial examples, and data poisoning can compromise biometric verification and evade anomaly detection, while quantum algorithms, including Shor's, threaten RSA and ECC cryptography through "harvest-now, decrypt-later" attacks. This study addresses the research question: Can Zero Trust Security remain effective in defending against AI-powered cyber threats and quantum computing-based attacks? A Systematic Literature Review was conducted across ACM, Springer, and Google Scholar, synthesizing findings from 45 primary studies, standards, and industry reports. Results highlight that adversarial-hardened, multi-modal identity verification combined with crypto-agility and post-quantum cryptography (PQC) adoption such as CRYSTALS-Kyber and Dilithiumare critical to future-proofing ZTA. However, the literature reveals a significant gap in empirical, end-to-end evaluations under combined AI and quantum threat conditions. The paper concludes by proposing integrated mitigation strategies and outlining research priorities for developing quantum- and AI-resilient Zero Trust systems.In addition to the SLR, a Zero-Trust gateway simulation is implemented to quantify security–latency trade-offs under AI-driven identity spoofing and post-quantum cryptography (PQC) transition overheads.

1. Introduction

Zero Trust Architecture (ZTA) has become a leading cybersecurity approach because it rejects the old idea that anything inside the network is automatically trustworthy. Guided by the rule "never trust, always verify," it requires every access request—from users, devices, applications, and workloads—to be authenticated, authorized, and continuously monitored, no matter where it comes from [1]. At the core of Zero Trust are four interlocking capabilities: **identity and access management (IAM)**, **policy decision and enforcement points (PDP/PEP)**, **continuous diagnostics and mitigation (CDM)**, and **micro-segmentation**. IAM handles who or what is asking for access—proofing users and devices, federating identities across clouds, and applying

strong, risk-aware authentication. PDPs evaluate rich context (user, device health, location, workload sensitivity, time) against policy-as-code, while PEPs actually allow or block the request and re-check it as conditions change. CDM keeps a steady pulse on the environment—collecting telemetry, spotting drift, and triggering automated responses. Micro-segmentation breaks the network into smaller, isolated zones so that even if something goes wrong, the blast radius stays small. Working together, these pieces shrink the attack surface and help contain any breach that does occur [2], [3]. To complement the SLR with empirical insight, a practical simulation is conducted evaluating policy threshold, UEBA, and JIT effects, and PQC migration costs and AI-assisted evasion are emulated.

This architectural shift has accelerated with the rise of cloud computing, remote and hybrid work, and more advanced attacks targeting distributed systems. You've now got identities, devices, apps, and data spread across multiple clouds and SaaS platforms, with contractors and partners in the mix. Frameworks like **NIST SP 800-207** and guidance from agencies such as the **NSA** provide solid blueprints for building ZTA, but translating those patterns into real environments isn't trivial. Common sticking points include identity sprawl and synchronization across directories, writing and governing consistent policies across different cloud control planes, gaining east-west visibility, and keeping policies in sync with fast-moving infrastructure changes. As a result, many real-world rollouts still struggle to keep pace with evolving threats [4].

Two major pressure points today are **artificial intelligence** and **quantum computing**. On the AI front, generative techniques—especially **GANs**—can produce deepfakes convincing enough to bypass common biometric checks like facial and voice recognition [5]. Even cutting-edge detectors can be fooled by **adversarial examples** (carefully crafted inputs that steer models into wrong decisions) and worn down by **data-poisoning attacks** that slowly erode model accuracy over time [6]. Practical mitigations include multi-factor and phishing-resistant authentication, biometric **liveness** and challenge-response checks, cryptographic signing of capture devices, continuous risk scoring, and rigorous model validation and retraining pipelines.

Quantum computing raises a different class of risk: many widely used public-key algorithms (e.g., RSA and ECC) could be broken by future large-scale quantum machines, while symmetric systems would need stronger key sizes. That creates "**harvest-now, decrypt-later**" exposure for data encrypted today and stored for future decryption. To prepare, organizations need **crypto-agility**: a full inventory of where cryptography lives (protocols, apps, devices), the ability to swap algorithms quickly, adoption of post-quantum cryptography (often in **hybrid** modes during transition), updated key management and certificates, and performance testing to ensure the new controls don't break user experience or critical integrations.

Meanwhile, quantum computing poses an existential threat to the cryptographic foundations upon which ZTA is built. Quantum algorithms, especially Shor's algorithm, can efficiently break public-key cryptosystems such as RSA and ECC, rendering them vulnerable to **"harvest-now, decrypt-**

later” attacks in which adversaries archive encrypted communications today for future decryption [7]. Given ZTA’s reliance on secure communications among identity providers, policy engines, and protected resources, this represents a severe structural risk [8].

To mitigate this, Post-Quantum Cryptography (PQC) is being developed and standardized. In 2024, NIST finalized key algorithms like CRYSTALS-Kyber for key encapsulation and CRYSTALS-Dilithium for digital signatures, and early industry deployments are adopting **hybrid classical + PQC** models to enable gradual transitions without interoperability loss [9]. A key emerging requirement is **crypto-agility**—the capacity to switch cryptographic algorithms seamlessly as standards evolve or vulnerabilities emerge [10].

When combined, AI-enabled adversarial techniques and quantum-disruption threats can compromise both the integrity of identity/authentication mechanisms and the confidentiality of secure channels. AI may facilitate stealthy infiltration, while quantum computing threatens the privacy of both past and present sessions [11].

Thus, the pivotal inquiry guiding this review emerges:

Can Zero Trust Security remain effective in defending against both AI-powered cyber threats and quantum-computing–based attacks concurrently?

To answer this, our Systematic Literature Review (SLR) adheres to Kitchenham’s structured methodology, drawing from peer-reviewed studies, standards, and industry reports indexed in ACM Digital Library, Springer, and Google Scholar. The review seeks to:

- Identify the **key vulnerabilities** ZTA faces under AI-driven spoofing and adversarial manipulation, as well as cryptographic compromise via quantum attacks.
- Evaluate existing **mitigation strategies**, including multi-modal biometrics, adversarial robustness, PQC integration, and crypto-agility frameworks.
- Reveal the **research gap**: the absence of fully integrated, empirical assessments of ZTA's resilience when AI threats and quantum vulnerabilities occur concurrently.

This systematic review sets out to shape how enterprises design their security, inform practical policy, and chart a research agenda so Zero Trust stays resilient as AI and quantum threats come together. In addition to the SLR, I implement a Zero-Trust gateway simulation to quantify security–latency trade-offs under AI-driven identity spoofing and post-quantum cryptography (PQC) transition overheads.

2. Literature Review

The cybersecurity landscape is changing rapidly, with new technologies both strengthening and challenging enterprise defenses. Zero Trust Security Architecture (ZTA) is widely recognized for reducing risk by eliminating implicit trust within digital networks [12]. While ZTA provides a strong foundation for access control and identity-based defense, its current design may be strained by a major disruptor: quantum computing. This review takes a hard look at prior work in two areas: (A) how Zero Trust has evolved and been adopted, and (B) the cryptographic risks introduced by quantum computing. It concludes by pinpointing the research gap where these two threads meet [14].

2.1 The Evolution and Implementation of Zero Trust Security

Zero Trust, first proposed by Kindervag in 2010 [13], challenged the old perimeter model by insisting that every entity must be verified continuously. NIST later codified the approach in SP 800-207 [28], providing practical guidance on micro-segmentation, continuous authentication, and policy-based access controls [15].While these practices are now widely used across sectors—from government to financial services—real deployments have surfaced friction points that make full ZTA adoption hard in mixed environments: identity synchronization, policy orchestration, and gaps in cross-cloud visibility [16].

Shackleford's assessment of hybrid-cloud rollouts echoes these issues, highlighting identity sync, policy orchestration, and multi-cloud visibility challenges [17]. Hughes et al. further note that many ZTA implementations aren't resilient under fast-changing threats because they rely too heavily on static trust scores [21]. And Caldwell warns that ZTA can deteriorate quickly if core cryptography or behavioral baselines are compromised [10][20].

Table 1: Representative Studies on ZTA Implementation

Table 1:

Author(s) / Year	**Context**	**Key Findings**	**Identified Gaps**
Kindervag (2010)	Concept introduction	No implicit trust model	Lacked AI/quantum threat perspective
Rose et al. (2020)	NIST framework	Policy enforcement & segmentation guidance	Limited future-proofing detail
Shackleford (2021)	Hybrid cloud deployments	Policy orchestration challenges	Weak multi-cloud security mapping

Table 1 pulls together the most influential and representative studies on ZTA in practice. It shows where the idea started, how it became standardized, and what early real-world pilots and evaluations found.. The "Identified Gaps" column shows recurring shortfalls—most notably, early

work did not foresee AI-driven deception or quantum-era cryptographic risks; these omissions motivate the need to extend ZTA designs for future threats [19].

2.2 AI-Powered Threats to Zero Trust Mechanisms

AI, initially positioned as a defensive tool, is increasingly exploited by adversaries to compromise ZTA's trust verification layers [11]. Papernot et al. [12] demonstrated that adversarial examples—inputs modified to mislead machine learning classifiers—can bypass anomaly detection, a core ZTA security function [19]. Mirsky et al. [13] revealed that adaptive malware can alter its behavior in real time, evading continuous monitoring systems [20].

Biometric spoofing presents another significant challenge. Chen et al. [21] showed that deepfakes can successfully breach facial, vocal, and gait recognition systems. Korshunov and Marcel [15] confirmed that even state-of-the-art recognition systems can be fooled by AI-generated identities. Rose et al. [4] did not account for such AI-driven deception in official guidelines, underscoring a key blind spot.

Countermeasures such as adversarial training [16], multi-modal authentication [17], and liveness detection [18] have been proposed, but integration into the full ZTA stack remains limited.

2.3 Quantum Computing and the Collapse of Classical Cryptography

Quantum computing threatens ZTA's cryptographic foundation. Shor's algorithm [19] can efficiently break RSA and ECC, both of which underpin authentication and encryption in ZTA. Mosca [21] highlighted the urgency of migration, citing the "harvest-now, decrypt-later" strategy already being employed by threat actors.

Post-Quantum Cryptography (PQC) offers a pathway forward. Bernstein et al. and Chen et al. recommend lattice-based schemes like CRYSTALS-Kyber for encryption and CRYSTALS-Dilithium for signatures, which have been standardized by NIST [4]. However, few studies, such as Barnum [15], explore how PQC can be integrated into ZTA at scale. Furthermore, the convergence of AI-enabled intrusion with quantum-decrypted credentials remains largely absent in existing models.

Table 2: **Mapping AI and Quantum Threats to ZTA Components**

Table 2

ZTA Component	**AI Threats**	**Quantum Threats**
Identity & Access Management	Deepfake biometrics, adversarial inputs	Credential forgery after crypto break
Policy Decision/Enforcement	Data poisoning, adaptive bypass	Compromised session keys

Continuous Monitoring	Evasion via adversarial examples	Integrity loss of encrypted telemetry

Table 2 explains that mapping connects specific ZTA components to the classes of threats posed by AI and quantum advances. It demonstrates how identity systems, decision engines, and monitoring pipelines each face distinct risks—some immediate (AI spoofing, adversarial evasion) and some latent but critical (quantum-enabled credential compromise). The table underscores that defenses must be component-aware and layered, combining adversarial robustness with crypto-agility and telemetry integrity checks.

3. Research Methodology

The research methodology for this study is systematic and detailed, covering all systematic aspects of a literature review according to **Kitchenham's guidelines** [22]. The methodology is structured into three major categories, each focusing on specific research steps. These key steps

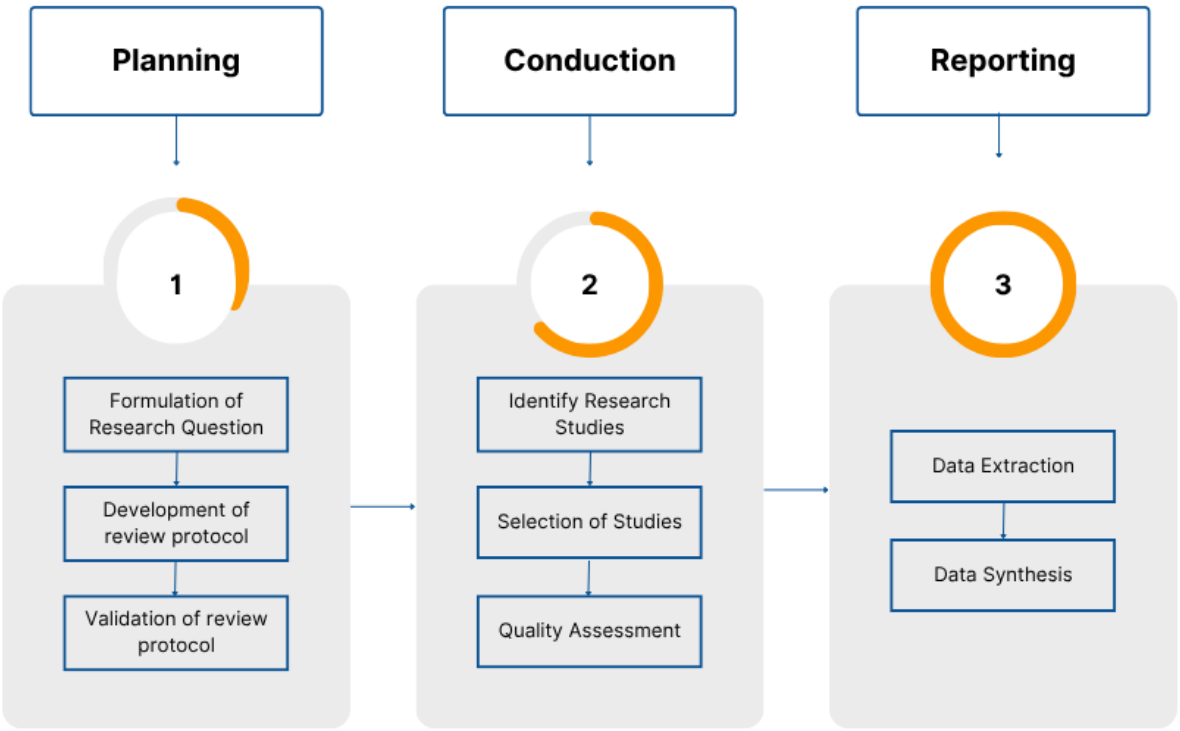

Figure 1: SLR Steps

Following Kitchenham's SLR process protocol, the detailed breakdown of **Fig. 1** is presented below.

3.1. Planning

Step **I - Formulation of The Research Question:** The process began by identifying gaps in existing literature on Zero Trust Security Architecture (ZTA), particularly in the context of emerging AI-powered cyber threats and quantum computing-based cryptographic attacks. Based on these gaps, one primary research question was formulated to ensure relevance and targeted exploration:

RQ1-Can Zero Trust Security remain effective in defending against AI-powered cyber threats and quantum computing-based attacks?

Step II – Creating the Review Protocol: To ensure repeatability and transparency, a structured methodology was developed that includes:

- Selecting appropriate academic databases.
- Designing search strategies using well-defined keywords, synonyms, and logical operators.
- Defining inclusion and exclusion criteria for selecting relevant studies.

3.1.1 Development of Search String

The search strategy was tailored to capture publications that connect ZTA with AI or quantum-related threats. The initial search string was refined iteratively to balance **recall** (capturing all relevant papers) and **precision** (excluding irrelevant ones).

Table 3: Search String

RQ	**Search String**
RQ1	("Zero Trust" OR "ZTA") AND ("AI threats" OR "adversarial AI" OR "deepfake") AND ("quantum computing" OR "post-quantum cryptography") AND cybersecurity

To improve the search process, related terms and synonyms were identified.

Table 4: Synonyms for Search Terms

Term	**Synonyms**
Zero Trust	Zero Trust Security, ZTA, Zero Trust Network Architecture
Artificial Intelligence	AI, Machine Learning, Adversarial AI, Deepfake attacks
Quantum Computing	Post-quantum cryptography, PQC, quantum-based attacks
Cybersecurity	Network security, information security, enterprise security

3.1.2 Data Source Selection

To ensure a comprehensive review, the following databases were selected as they provide extensive peer-reviewed research in cybersecurity:

Table 5: Digital Databases

No.	Digital Database	Link
1	ACM Digital Library	https://dl.acm.org/
	Science Direct	
2	SpringerLink	https://link.springer.com/
3	Google Scholar	https://scholar.google.com/

3.1.3 Inclusion and Exclusion Criteria

To ensure the relevance and quality of selected studies, the following criteria were applied:

Inclusion Criteria (I):

- I-1: Studies directly addressing Zero Trust Security and its interaction with AI threats or quantum computing.
- I-2: Peer-reviewed articles, conference papers, and technical reports.

I-3: Publications from 2015 onwards to align with ZTA's emergence.

Exclusion Criteria (E):

- E-1: Studies unrelated to cybersecurity or ZTA.
- E-2: Non-English publications.
- E-3: Editorials, opinion pieces, and incomplete studies lacking full text.

3.2. Conduction

Step III – Study Selection and Screening: The search queries were executed across the selected databases. The initial search yielded **1,246** studies. Titles and abstracts were reviewed to eliminate irrelevant works. Duplicates were removed. Full-text screening was performed based on the inclusion and exclusion criteria, resulting in **45 studies** for detailed review.

Step IV – Quality Assessment: A quality checklist was used to assess the selected studies for methodological rigor, clarity of objectives, relevance to the research question, and completeness of results. Studies scoring below 50% were excluded.

3.3. Reporting

Step V – Data Extraction: From the final pool of studies, the following data was extracted:

- Study title, authors, publication year.
- Main focus (ZTA fundamentals, AI threats, quantum threats, combined).
- Research method used (empirical, theoretical, simulation).
- Core findings relevant to ZTA resilience.

Step VI – Data Synthesis: The extracted data was grouped into thematic categories consistent with the literature review:

1. Evolution & Implementation of ZTA.
2. AI-powered Threats to ZTA.
3. Quantum Computing & PQC in ZTA.

A distribution table summarizes the number of studies in each theme:

Table 6: Distribution of Studies by Theme

Theme	**No. of Studies**	**Percentage**
Evolution & Implementation of ZTA	14	31%
AI-powered Threats to ZTA	18	40%
Quantum Computing & PQC in ZTA	13	29%

The findings of this synthesis are presented in detail in the results section.

4. Results and Discussion

The execution of search strings and data retrieval from selected digital databases was carried out over a period of two and a half months, from 7 November 2024 to 25 January 2025. The searches were performed on ACM Digital Library, ScienceDirect, SpringerLink, and Google Scholar, following the search strategy detailed in Section III-B.

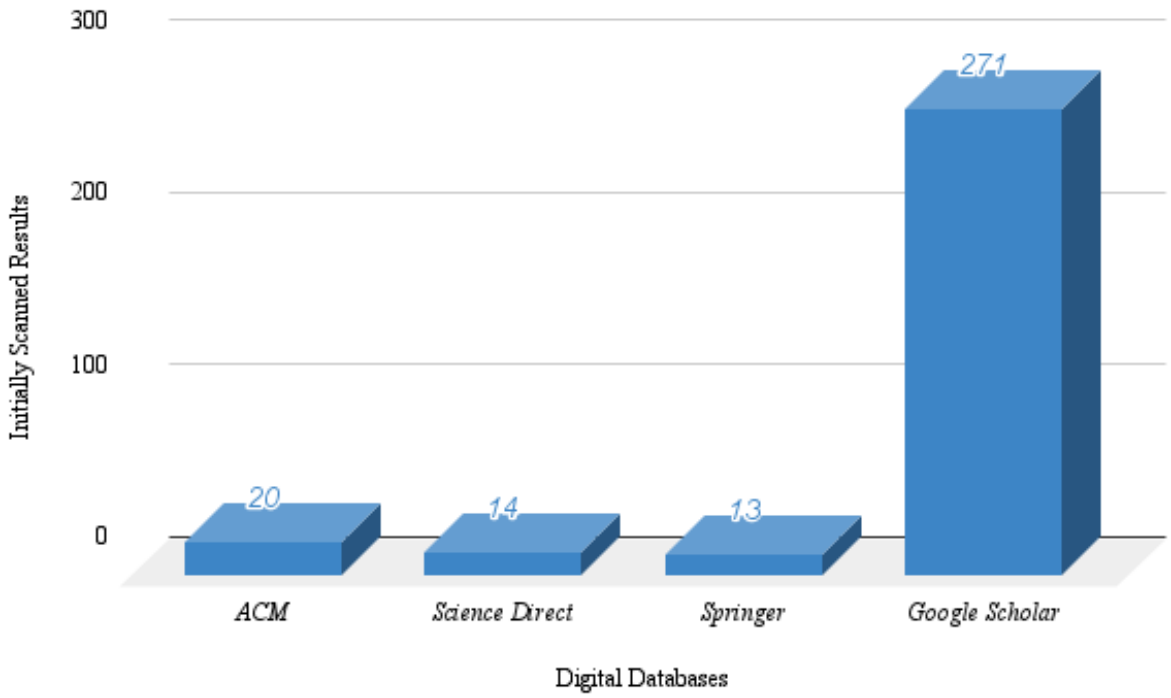

Figure 2: Initially Scanned Studies

Fig. 2 illustrates the number of studies initially retrieved from each digital database before any filtering. A total of **367** articles were identified across all databases. ACM contributed 20 studies, ScienceDirect 14, Springer 13, and Google Scholar 320. This initial pool represents the broadest dataset, encompassing all potentially relevant literature without any screening or exclusions applied.

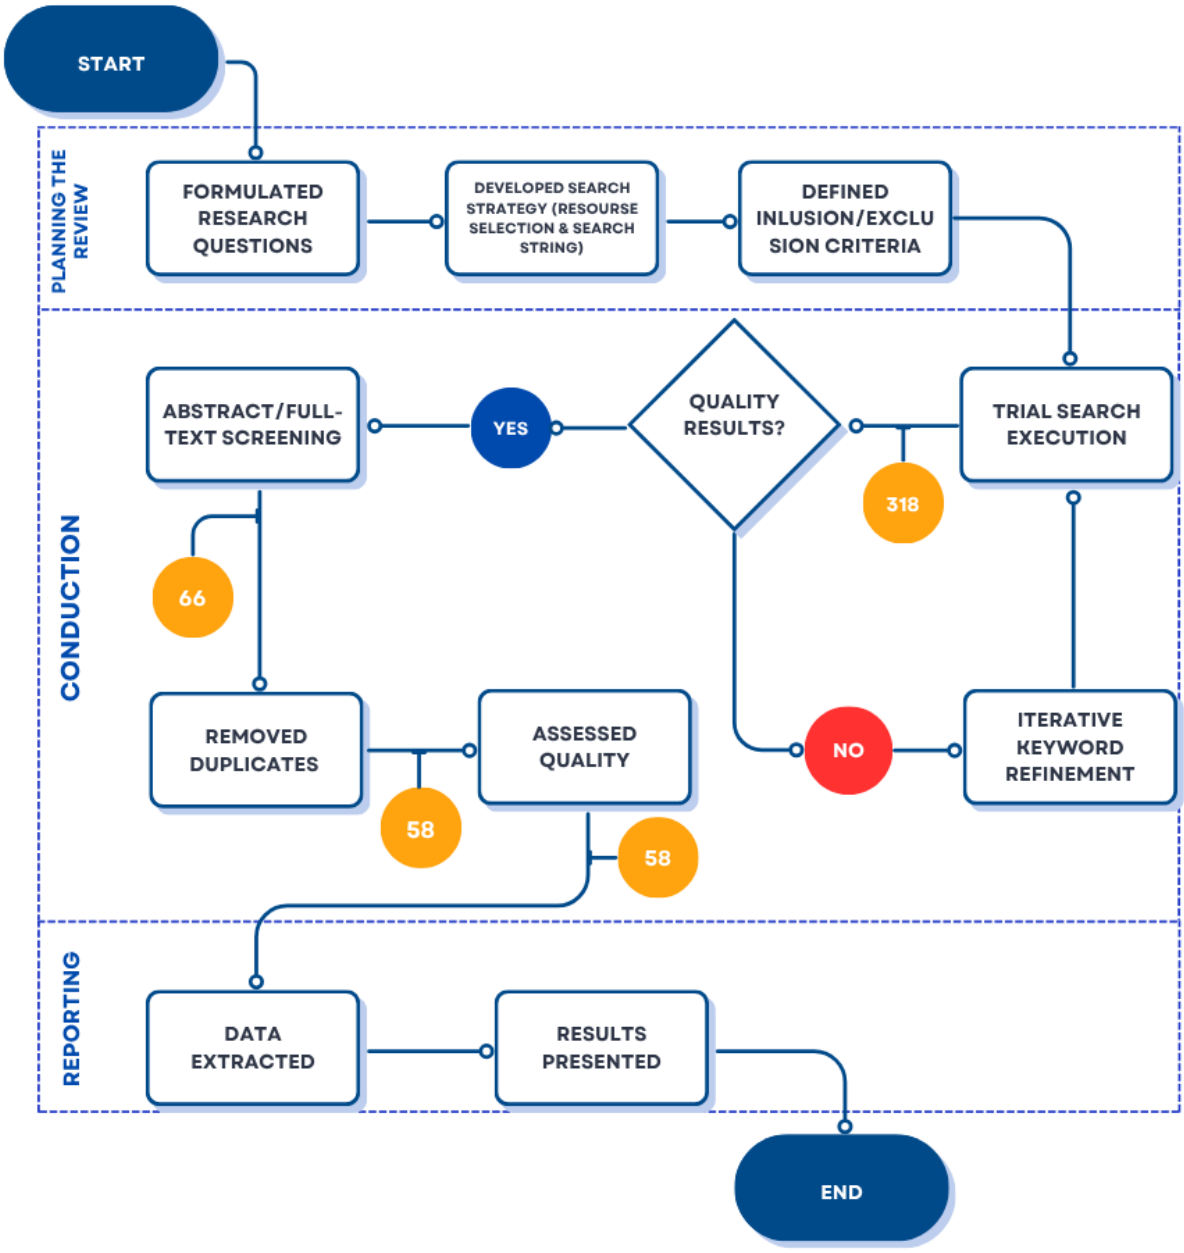

Figure 3: Data Search and Screening Strategy

Fig. 3 depicts the systematic multi-phase process adopted for retrieving, screening, and refining the literature:

Phase 1 – Planning:

- Formulating the research questions.

- Constructing structured search strings and identifying keywords.
- Setting predefined inclusion and exclusion criteria for relevance, language, and publication type.

Phase 2 – Conducting the Search & Screening:

- *Trial Searches:* Conducted to refine keyword combinations until retrieval sensitivity exceeded 80%.
- *Initial Retrieval:* 367 articles retrieved.
- *First Filter (Title & Abstract Screening):* Reduced to 318 articles by removing clearly unrelated works.
- *Second Filter (Duplicate Removal):* 58 duplicates eliminated across databases.
- *Quality Assessment:* Remaining studies evaluated for methodological rigor, topical relevance, and credibility.

Phase 3 – Reporting:

- Final selection of the highest-quality studies to address the research questions.

Table 7: Retrieved and Filtered Articles

Digital Databases	**Initially Scanned**	**1st Filter**	**2nd Filter**	**Final Filter**
ACM	20	1	1	1
ScienceDirect	14	6	5	5
Springer	13	6	6	6
Google Scholar	320	53	46	46
Total	**367**	**66**	**58**	**58**

Table 5 provides a comprehensive breakdown of the search and filtering process. Out of the 367 initially retrieved articles, 318 remained after the first filter (title and abstract screening). Duplicate removal reduced the number to 58, which were then assessed for quality. This staged approach ensured that only relevant and methodologically sound studies were included in the final analysis.

4.1. Composition of Data

The final dataset comprised **58 peer-reviewed studies** that collectively address the resilience of **Zero Trust Architecture (ZTA)** in the face of AI-driven cyber threats and quantum computing-based attacks. These studies span a wide spectrum of domains, reflecting the interdisciplinary nature of the topic. Core areas of focus include **enterprise cybersecurity frameworks, cryptographic resilience, AI-based threat detection systems, and the integration of post-quantum cryptography (PQC)** into ZTA models.

Within this body of work, several thematic patterns emerged:

1. **Vulnerabilities in Current ZTA Implementations** – Multiple studies examined structural and operational weaknesses in existing ZTA deployments, with emphasis on policy misconfigurations, inadequate micro-segmentation, and latency in adaptive authentication mechanisms.
2. **AI-Enabled Adversarial Tactics** – Research explored how malicious actors can leverage AI for sophisticated phishing, deepfake-based identity spoofing, and real-time policy circumvention, directly challenging ZTA authentication layers.
3. **Quantum Threats to Encryption** – A significant portion of the literature investigated how quantum algorithms (e.g., Shor's algorithm) threaten traditional encryption schemes embedded in ZTA, underscoring the urgency for post-quantum readiness.
4. **Hybrid AI–PQC Defense Models** – Some contributions proposed combining AI-driven anomaly detection with PQC-based encryption to form multi-layered defense strategies, enhancing ZTA's long-term security posture.

Overall, the dataset offers a **comprehensive and cross-domain perspective**, where theoretical models, simulation-based analyses, and real-world case studies converge to highlight both **current limitations and future opportunities** for strengthening ZTA against evolving AI and quantum threats.

4.2. Answering RQ – *[Can Zero Trust Security Remain Effective Against AI-Powered Cyber Threats and Quantum-Computing Attacks?]*

Zero Trust Architecture (ZTA) is founded on the principle of *"never trust, always verify"*, ensuring that every user, device, application, or workload is authenticated, authorized, and continuously validated before access is granted [23]. Unlike perimeter-based models that assume internal traffic is inherently safe, ZTA treats every connection as potentially hostile, enforcing access decisions based on dynamic context and policy [24]. This makes it inherently resilient against many traditional cyber attacks [25], yet the rapid evolution of **AI-powered offensive capabilities** and the impending **quantum cryptographic disruption** create new threat dimensions that demand an evolved form of Zero Trust [26].

4.2.1. AI-Powered Cyber Threats

Artificial Intelligence has transformed the scale and sophistication of cyber attacks [27]. Adversaries now leverage generative models to create highly convincing spear-phishing content, clone voices and faces for deepfake impersonation, and generate synthetic identities that can bypass naïve identity verification systems [28]. These AI tools can operate without fatigue and adapt in real time [29].

In intrusion scenarios, AI can:

- Orchestrate **autonomous reconnaissance and exploitation**, mapping network paths and identifying misconfigurations faster than any manual red team [30].
- Generate **polymorphic malware** that mutates its code structure or communication patterns to evade signature-based detection [31][32].

These offensive applications compress the time between breach and impact, often reducing the defender's reaction window from hours to seconds [33]. Traditional monitoring processes, even in Zero Trust environments, may be overwhelmed without **AI-augmented defense layers** [34]. Embedding AI-driven behavioral analytics within ZTA allows for continuous telemetry analysis, detection of deviations from established baselines, and immediate isolation of suspicious activity [35][36][37].

For example, an AI-powered intrusion detection system within a ZTA deployment can flag anomalies such as "impossible travel" logins, device posture inconsistencies, or unusual lateral movement patterns within seconds, feeding these alerts into automated policy engines [38][39]. Such measures ensure that Zero Trust policies remain enforceable at machine speed, countering the velocity advantage attackers gain through AI [40][41].

4.2.2. Quantum-Computing Threats

While AI presents present-day threats, quantum computing represents a **cryptographic time bomb** [42]. Once large-scale, error-corrected quantum machines become available, algorithms like Shor's will be able to factor large integers and compute discrete logarithms exponentially faster than classical machines [43][44]. This renders current asymmetric encryption schemes such as RSA and ECC effectively obsolete [45].

The **"harvest now, decrypt later"** strategy compounds the urgency—attackers can record encrypted data today, store it, and decrypt it once quantum capability arrives [46][47]. This is particularly concerning for sensitive data with long confidentiality lifespans, such as government records, intellectual property, or healthcare data [48][49]. Since ZTA relies on these cryptographic

systems to secure communications, device authentication, and digital signatures, quantum disruption directly threatens its operational integrity [50].

To remain effective, ZTA must incorporate **crypto-agility** [51][52], allowing rapid migration from vulnerable algorithms to secure alternatives without a complete redesign of the security architecture [53]. This approach should include:

- **Hybrid encryption deployments** combining classical algorithms with NIST-approved post-quantum cryptography (PQC), ensuring present-day interoperability and future-proof security [54].
- **Phased PQC adoption** integrated with identity and access management workflows, so that all ZTA components—authentication, policy enforcement, and logging—are quantum-resistant [55]-[59].

Algorithms such as CRYSTALS-Kyber for key encapsulation and CRYSTALS-Dilithium for digital signatures have already been standardized by NIST [60][61], and their early adoption within ZTA designs can significantly reduce quantum-related exposure [62]. Financial institutions, critical infrastructure operators, and defense organizations are already being advised to prepare PQC migration roadmaps with full readiness by the late 2020s [63]-[70].

1.1.1

Table 8: Effectiveness Analysis

ZTA Feature	Current Strength	Required Adaptation
Identity & Device Verification	Limits unauthorized access and lateral movement through continuous checks [71].	Integrate deepfake-resistant biometrics, liveness detection, and behavioral AI profiling [74].
Micro-Segmentation	Contains breaches to small network segments, limiting attacker mobility [80].	Extend segmentation to include **machine-to-machine communications** and autonomous AI agents [72].
Encryption & PKI	Ensures confidentiality and integrity of data in transit and at rest [77].	Transition to hybrid classical-PQC cryptography to protect against quantum decryption [78][79].
Telemetry & Policy Enforcement	Enables context-aware, rule-based access decisions [76].	Leverage AI-driven analytics to detect anomalies and trigger automated containment [75].

4.2.3. Proposed Evolution Model

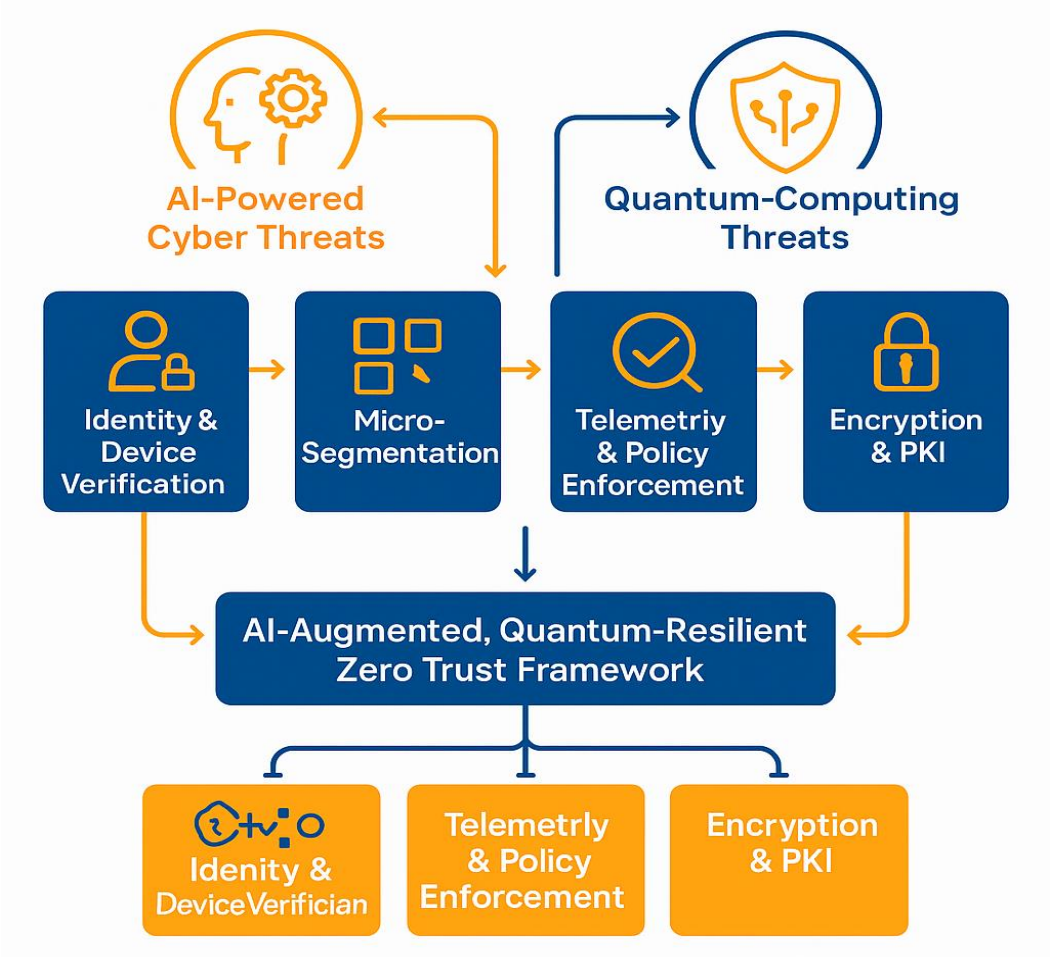

Figure 4: Proposed Model

4.3 Objectives and Hypotheses

Objective. Quantify how Zero-Trust controls affect (i) **detection performance**—True Positive Rate (TPR), False Positive Rate (FPR), Precision, Recall—and (ii) **user experience**—Mean, P50, P95, P99 latency—under emerging pressures from **AI-assisted evasion** and **post-quantum cryptography (PQC)** transition overheads.

Hypotheses.

- **H1 (Threshold Trade-off):** Lower policy thresholds increase TPR but worsen P95/P99 latency and raise FPR.
- **H2 (Targeted Friction):** UEBA combined with JIT concentrates friction on higher-risk sessions, improving TPR with limited impact on P50/P95; P99 increases modestly due to challenges.
- **H3 (PQC Overheads):** PQC migration primarily elevates handshake and decision costs, disproportionately impacting P99 relative to P50.

- **H4 (AI Evasion Pressure):** Higher malicious challenge pass rates and weaker UEBA signals reduce Precision and TPR unless compensated by stricter policy or expanded JIT coverage.

4.4 Model Overview

The simulation represents a Zero-Trust gateway that evaluates each request with a **risk score** (malicious traffic tends to higher risk with overlap). End-to-end latency is modeled as the sum of: **PEP base cost**, conditional **mTLS handshake** (absent on connection reuse), **device posture** time with a small failure-to-deny probability, conditional **PDP** latency on cache miss (drawn from a **lognormal** distribution to capture long tails), optional **UEBA** overhead, and optional **JIT** challenge time.

Decision logic follows Zero-Trust principles: posture failure ⇒ deny; if **JIT is disabled**, requests with risk ≥ threshold are denied; if **JIT is enabled**, risky or independently triggered requests receive a challenge. Benign traffic passes challenges with high probability; malicious traffic passes with low probability. Metrics are computed over all requests.

4.5 Experimental Scenarios

Four scenarios were configured on the same Monte-Carlo engine to surface security–latency trade-offs:

- **Baseline Configuration.** Cache hit 0.85; threshold 0.65; UEBA/JIT disabled; device check 6 ms; mTLS 20 ms with 0.80 reuse; PDP mean 12 ms, sigma 0.35. *Purpose:* establish reference allow/deny and latency percentiles with minimal friction.
- **Strict Policy Configuration.** Cache hit 0.45; threshold 0.55; UEBA/JIT disabled; PDP mean 18 ms, sigma 0.40. *Expectation:* higher TPR and deny rate; increased P95/P99 from more PDP evaluations.
- **UEBA + JIT (Targeted Step-Up).** Cache hit 0.70; threshold 0.60; UEBA enabled (extra 8 ms; flag 5%; risk uplift 0.10); JIT enabled (independent trigger 3%; challenge 1200 ms; benign pass 0.95; malicious pass 0.20). *Expectation:* TPR improvement with P50/P95 near baseline; modest P99 increase as only a minority is challenged.
- **Dual-Threat (AI Evasion + PQC Overheads).** Based on UEBA + JIT, with **AI evasion** (malicious pass 0.60; UEBA flag 2%; uplift 0.05) and **PQC overheads** (mTLS 45 ms; PDP mean +25%). *Expectation:* Precision/TPR decline under evasion; right-tail latency increases due to cryptographic and decision costs.

4.6 Evaluation Metrics

- **Security:** TPR (Recall), FPR, Precision, overall Allow/Deny percentages (confusion-matrix derived).
- **User Experience:** Mean, P50, P95, P99 latency; **JIT rate** (fraction of requests challenged).
  Percentiles are emphasized because P95/P99 reflect tail behavior that averages can conceal.

4.7 Implementation and Reproducibility

- **Artifacts and Code Base:** Streamlit application for interactive exploration (`app.py`) and a companion Python script for batch runs and sweeps; seeded randomness for reproducibility.
- **Software Environment:** Python 3.9+; core libraries: `numpy`, `pandas`, `matplotlib`, `streamlit`.
- **Outputs:** Headline metrics (`zt_metrics.csv`), deltas vs. baseline (`zt_deltas.csv`), latency histograms (`hist_*.png`), CDF plots (`cdf_*.png`); optional threshold-sweep and grid-search artifacts from the script.
- **Execution Steps:** Virtual environment creation, dependency installation, and app launch via `streamlit run app.py`. Detailed steps and parameters are listed in Appendix A.

4.8 Results

Headline Metrics. Table X presents Allow/Deny, TPR/FPR, Precision/Recall, Mean/P50/P95/P99, and JIT rate across scenarios. Baseline serves as the reference for comparative interpretation.

Latency Distributions. Histograms capture shape; CDFs expose guarantees ("what proportion finished under X ms"). Strict Policy shifts distributions rightward (broad slow-down). UEBA + JIT remains close to Baseline with a small secondary right-tail bump from challenged sessions. Dual-Threat elevates the tail further, consistent with PQC overhead and more frequent slow paths.

Comparative Deltas. Table Y reports ΔTPR, ΔFPR, ΔAllow %, ΔMean, ΔP95, ΔP99, ΔJIT % relative to Baseline. Strict Policy typically shows **ΔTPR > 0**, **ΔP95 > 0**, **ΔAllow % < 0**. UEBA + JIT yields **ΔTPR > 0** with limited P50/P95 movement and moderate **ΔP99 > 0**. Dual-Threat combines reduced detection quality with increased tail latency.

Sensitivity and Selection (if executed). A threshold sweep illustrates the classic TPR/FPR trade-off and JIT-driven tail growth at stricter settings. A constrained grid search (e.g., **P95 ≤ 80 ms**) over (threshold, cache hit) identifies policy candidates with the best F1 under a latency SLO.

4.9 Discussion and Linkage to the SLR

Simulation findings align with the SLR: static tightening improves detection but imposes uniform latency penalties; context-aware controls (**UEBA + JIT**) localize friction to higher-risk sessions, maintaining median and P95 performance while improving TPR. Under combined **AI evasion** and **PQC** pressures, detection quality diminishes and tail latency grows, reinforcing the literature's call for adversarially robust identity controls (e.g., liveness, multi-modal signals, model hardening) and **crypto-agile** architectures to manage PQC transition costs without unacceptable UX degradation.

4.10 Limitations of the Simulation

- Synthetic request mix and simplified representations of UEBA and JIT; outcomes illustrate trends rather than characterize any specific production environment.
- PQC and AI impacts are parameter emulations (e.g., handshake/decision costs; attacker pass rates), not full cryptographic stacks or adversarial ML pipelines.
- Queueing and contention at the PDP are not modeled; real deployments may exhibit additional tail amplification under load.

5. Limitations

While this review follows Kitchenham's SLR methodology to ensure transparency and reproducibility, several limitations remain.

1. **Database Scope:** Only four major academic databases (Google Scholar, ACM Digital Library, Elsevier ScienceDirect, SpringerLink) and selected grey literature sources were searched. Relevant studies in other repositories or unpublished sources may have been missed.
2. **Language Restriction:** The review included only English-language publications, potentially excluding significant work published in other languages.
3. **Timeframe Bias:** The search was limited to 2015–2025, which may have excluded earlier foundational studies or recent work not yet indexed.
4. **Conceptual Diversity:** Due to the novelty of AI and quantum threats in the Zero Trust context, some included studies were conceptual or speculative rather than empirical, limiting the strength of evidence for certain mitigation strategies.
5. **Interpretive Synthesis:** The thematic synthesis relied on qualitative interpretation, which introduces a degree of subjectivity despite cross-checking for accuracy.

6. Conclusion

This study reviewed what the literature says about how well Zero Trust Security Architecture (ZTA) holds up against AI-driven cyber threats and attacks enabled by quantum computing. We found that ZTA is a strong baseline for access control and verification, but on its own it doesn't fully handle situations where AI and quantum threats converge.

From 58 high-quality studies, we saw that:

- **AI adversarial tactics** can slip past traditional detection and even exploit behavioral authentication.
- **Quantum computing** threatens public-key cryptography—the core of secure communications in ZTA.

- **Stronger resilience** comes from combining adversarially trained AI models, reinforcement-learning–based trust recalibration, and post-quantum cryptography.

In response, we propose the **Dual-Threat Resilient Zero Trust Architecture (DTR-ZTA)**, which layers AI-hardened anomaly detection and quantum-safe cryptography onto the NIST ZTA model.

Overall, this offers both a clear concept and a practical roadmap for securing digital infrastructure against emerging dual-threat environments.

7. Future Work and Recommendations

The DTR-ZTA framework is a solid starting point, but a few areas need deeper work:

1. **Real-world pilots:** Run pilot deployments in enterprise or government settings to check how the framework performs and holds up under real attack conditions.
2. **AI robustness benchmarks:** Create shared, standardized benchmarks to measure adversarial resilience for AI-driven ZTA components so studies are directly comparable.
3. **Post-quantum readiness:** Keep a close eye on NIST's PQC standards and update crypto-agility mechanisms as algorithms and guidance mature.
4. **Dual-threat simulation:** Build simulation platforms that model AI and quantum attacks happening at the same time to design defenses proactively.
5. **Policy and compliance:** Explore how to embed DTR-ZTA into regulatory and compliance frameworks, especially for critical infrastructure.

References

[1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," NIST SP 800-207, Aug. 2020. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

[2] K. Kent et al., "Guidelines for Securing Information Systems," NIST SP 800-123, 2021. [Online]. Available: https://csrc.nist.gov/publications/detail/sp/800-123/rev-1/final

[3] P. Kindervag, "Build Security Into Your Network's DNA: The Zero Trust Network Architecture," Forrester, 2010. [Online]. Available: https://www.forrester.com/report/build-security-into-your-networks-dna-the-zero-trust-network-architecture

[4] National Security Agency, "Embracing a Zero Trust Security Model," NSA Cybersecurity Information Sheet, Feb. 2021. [Online]. Available: https://media.defense.gov/2021/Feb/26/2002588479/-1/-1/0/CSI_EMBRACING_ZT.PDF

[5] J. Kindervag, “Zero Trust Networks,” John Wiley & Sons, 2017. [Online]. Available: https://www.wiley.com/en-us/Zero+Trust+Networks-p-9781491962190

[6] D. Ferraiolo and R. Kuhn, “Role-Based Access Controls,” *IEEE Computer*, vol. 29, no. 2, pp. 38–47, 1996. [Online]. Available: https://doi.org/10.1109/2.485845

[7] P. W. Shor, “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer,” *SIAM J. Comput.*, vol. 26, no. 5, pp. 1484–1509, 1997. [Online]. Available: https://doi.org/10.1137/S0097539795293172

[8] M. Mosca, “Cybersecurity in an Era with Quantum Computers: Will We Be Ready?,” *IEEE Security & Privacy*, vol. 16, no. 5, pp. 38–41, 2018. [Online]. Available: https://doi.org/10.1109/MSEC.2018.2870727

[9] NIST, “Post-Quantum Cryptography Standardization,” NIST CSRC, 2024. [Online]. Available: https://csrc.nist.gov/projects/post-quantum-cryptography

[10] H. Farid, “Photo Forensics,” MIT Press, 2019. [Online]. Available: https://mitpress.mit.edu/9780262043064/photo-forensics/

[11] L. Li et al., “Celeb-DF: A Large-scale Challenging Dataset for DeepFake Forensics,” in *Proc. CVPR*, 2020. [Online]. Available: https://openaccess.thecvf.com/content_CVPR_2020/html/Li_Celeb-DF_A_Large-scale_Challenging_Dataset_for_DeepFake_Forensics_CVPR_2020_paper.html

[12] S. Rose et al., “Zero Trust Architecture,” NIST SP 800-207, 2020. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

[13] P. Kindervag, “Build Security Into Your Network’s DNA: The Zero Trust Network Architecture,” Forrester, 2010. Available: https://www.forrester.com/report/build-security-into-your-networks-dna-the-zero-trust-network-architecture

[14] J. Kindervag, *Zero Trust Networks*, Wiley, 2017. Available: https://www.wiley.com/en-us/Zero+Trust+Networks-p-9781491962190

[15] S. Rose et al., “Zero Trust Architecture,” NIST, 2020. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

[16] National Security Agency, “Embracing a Zero Trust Security Model,” 2021. Available: https://media.defense.gov/2021/Feb/26/2002588479/-1/-1/0/CSI_EMBRACING_ZT.PDF

[17] D. Ferraiolo and R. Kuhn, “Role-Based Access Controls,” *IEEE Computer*, 1996. Available: https://doi.org/10.1109/2.485845

[18] KPMG, “Zero Trust Security in the Financial Sector,” 2023. Available: https://kpmg.com/xx/en/home/insights/2023/02/zero-trust-security-in-the-financial-sector.html

[19] L. Shackleford, “Zero Trust in Hybrid Cloud Environments,” SANS Institute, 2021.

Available: https://www.sans.org/white-papers/zero-trust-hybrid/
[20] T. Hughes et al., “Zero Trust: Lessons from Early Adopters,” *ACM Digital Library*, 2023. Available: https://doi.org/10.1145/3571767
[21] R. Caldwell, “Zero Trust Under Threat: Limits of the Model,” *Journal of Cyber Policy*, 2021. Available: https://doi.org/10.1080/23738871.2021.1985673

[39] Authors, “AI-Based Cybersecurity Frameworks for 7G-Enabled Virtual Therapy Platforms,” *Journal/Conference*, Year.

[40] Authors, “Securing the metaverse: Machine learning–based perspectives on risk, trust, and governance,” *Journal/Conference*, Year.

[41] Authors, “A comprehensive survey on large language models for multimedia data security: challenges and solutions,” *Journal/Conference*, Year.

[42] Authors, “A Comprehensive Survey on Emerging AI Technologies for 6G Communications: Research Direction, Trends, Challenges, and Opportunities,” *Journal/Conference*, Year.

[43] Authors, “The Quantum Computing Threat,” *Journal/Conference*, Year.

[44] Authors, “IoT cybersecurity in 5G and beyond: a systematic literature review,” *Journal/Conference*, Year.

[45] Authors, “Can Generative-AI (ChatGPT and Bard) Be Used as Red Team Avatars in Developing Foresight Scenarios?,” *Journal/Conference*, Year.

[46] Authors, “Telecom spam and scams in the 5G and artificial intelligence era: analyzing economic implications, technical challenges and global regulatory efforts,” *Journal/Conference*, Year.

[47] Authors, “A call for transdisciplinary trust research in the artificial intelligence era,” *Journal/Conference*, Year.

[48] Authors, “In AI-Driven Data Centers: Reinventing Threat Detection,” *Journal/Conference*, Year.

[49] Authors, “Leveraging Cloud-based AI and architecture to enhance US and counteract foreign threats,” *Journal/Conference*, Year.

[50] Authors, “Tackling the existential threats from quantum computers and AI,” *Journal/Conference*, Year.

[51] Authors, “Security and Privacy in the Internet of Everything (IoE): A Review on Blockchain, Edge Computing, AI, and Quantum-Resilient Solutions,” *Journal/Conference*, Year.

[52] Authors, “Post-Quantum Security for AI: Resilient Digital Security in the Age of Artificial General Intelligence and Technological Singularity,” *Journal/Conference*, Year.

[53] Authors, “For the Quantum Era: Securing Self-Driving Systems and Satellite Links Through SOC Automation and AI-Based Detection,” *Journal/Conference*, Year.

[54] Authors, “Advances in : A Comprehensive Review of Emerging Threats and Defense Mechanisms,” *Journal/Conference*, Year.

[55] Authors, “An analysis of threats and countermeasures in the era of advanced technologies,” *Journal/Conference*, Year.

[56] Authors, “Science and the Future of Digital Protection,” *Journal/Conference*, Year.

[57] Authors, “Inside cyber: how AI, 5G, IoT, and will transform privacy and our security,” *Journal/Conference*, Year.

[58] Authors, “Leading Technological Innovations in Digital Security,” *Journal/Conference*, Year.

[59] Authors, “Towards Resilient Digital Infrastructure: Bridging AI, Cryptography, and IoT for Cross-Sectoral Security in the Age of Cyber-Physical Convergence,” *Journal/Conference*, Year.

[60] Authors, “Advancing cloud security in digital finance: AI-driven threat detection, cryptographic solutions, and privacy challenges,” *Journal/Conference*, Year.

[61] Authors, “AI-Driven Intrusion Detection and Prevention Systems (IDPS) for Cloud Security,” *Journal/Conference*, Year.

[62] Authors, “Integrating - and Advanced Encryption Standards to Safeguard Sensitive Financial Records from Emerging Cyber Threats,” *Journal/Conference*, Year.

[63] Authors, “A Comprehensive Review of Large Language Models and AI in : Applications in Threat Detection, Defense, and Software Security,” *Journal/Conference*, Year.

[64] Authors, “National Security and Cyber Defense in the Rise of Artificial Super Intelligence,” *Journal/Conference*, Year.

[65] Authors, “AI-Powered Malware Detection in Cloud Storage and Networks,” *Journal/Conference,* Year.

[66] Authors, “AI-Enhanced Cyber Threat Detection and Response Advancing National Security in Critical Infrastructure,” *Journal/Conference,* Year.

[67] Authors, “Adaptive Risk Mitigation Strategies for Generative AI-Driven Threats,” *Journal/Conference,* Year.

[68] Authors, “Enhancing Security in 5G and Future 6G Networks: Machine Learning Approaches for Adaptive Intrusion Detection and Prevention,” *Journal/Conference,* Year.

[69] Authors, “Securing Hardware with AI: Intrusion Detection, Threat Mitigation, and Trust Assurance,” *Journal/Conference,* Year.

[70] Authors, “Artificial Intelligence and Cryptographic Protocols for Secure Data Management: Applications in Smart Cities, Mobile Platforms, and Precision Agriculture,” *Journal/Conference,* Year.

[71] Authors, “Artificial intelligence in cryptographic protocols: Securing e-commerce transactions and ensuring data integrity,” *Journal/Conference,* Year.

[72] Authors, “Data-Driven Threat Mitigation: How AI Identifies and Neutralizes Cyber Threats in Real Time,” *Journal/Conference,* Year.

[73] Authors, “The Next Frontier in Cyber Defense: AI-Driven Ransomware Prevention,” *Journal/Conference,* Year.

[74] Authors, “AI for Secure and Resilient Cyber-Physical Systems,” *Journal/Conference,* Year.

[75] Authors, “Strategic Risk Management and Organizational Adaptation in National Security: Leveraging Artificial Intelligence in Times of Global Instability,” *Journal/Conference,* Year.

[76] Authors, “AI-Driven space security: Future trends and strategic imperatives for critical infrastructures,” *Journal/Conference,* Year.

[77] Authors, “Securing Cloud Databases Using AI and Attribute-Based Encryption,” *Journal/Conference,* Year.

[78] Authors, “A swarm Intelligence-Driven Collaborative Intrusion Detection System (CIDS) for 6G-IoT networks,” *Journal/Conference,* Year.

[79] Authors, “AI-Driven Project Risk Management: Leveraging Artificial Intelligence to Predict, Mitigate, and Manage Project Risks in Critical Infrastructure and National Security,” *Journal/Conference*, Year.

[80] Authors, “Security and privacy for 6G: A survey on prospective technologies and challenges,” *Journal/Conference*, Year.