www.ijarp.com ISSN 2456-9992 Page: 01-27

Research Article

International Journal Advanced Research Publications

Volume: 02
Issue: 10

BEYOND ZERO TRUST PRINCIPLES: A STRUCTURED REVIEW AND EVIDENCE-APPRAISAL FRAMEWORK FOR ZERO TRUST ARCHITECTURE RESEARCH

Anish Ganesh Bandal*[1], Durvesh Bhadgaonkar[1], Prof. Ajitkumar Khachane[1]

[1]*Department of Information Technology, Vidyalankar Institute of Technology, Mumbai, Maharashtra, India.*

Article Received: 19 September 2026, Article Revised: 09 October 2026, Published on: 29 October 2026

***Corresponding Author: Anish Ganesh Bandal**

Department of Information Technology, Vidyalankar Institute of Technology, Mumbai, Maharashtra, India.

Doi: https://doi-doi.org/101555/ijarp.9141

ABSTRACT

Zero Trust Architecture (ZTA) replaces implicit, location-based trust with explicit access decisions that are re-evaluated as conditions change. Research on the topic has spread quickly across enterprise migration, Zero Trust Network Access (ZTNA), cloud-native systems, the Internet of Things (IoT), software-defined networking (SDN), blockchain and 5G/6G networks. Existing surveys map this breadth, but they seldom ask how well the published claims are actually supported. This review examines how recent ZTA research is distributed, what evidence stands behind its security and performance claims, which suspected research gaps survive scrutiny, and how studies can be appraised on a common footing.

Following PRISMA 2020 reporting guidance, we re-verified the 15 studies of an earlier survey and screened 97 new records, which gave 35 included studies published between 2021 and 2026 (21 primary, 14 secondary). Primary studies were profiled with a proposed four-dimension Zero Trust Evidence-Appraisal Framework (ZT-EAF) covering security effectiveness, operational performance, deployment feasibility and evidence quality, deliberately without an aggregate score. Only nine primary studies reported any operational cost, and just three measured both a security outcome and its overhead; all three are non-peer-reviewed preprints. Metrics, environments and baselines differed too much to pool, public artefacts could be confirmed for a single study, and none evaluated failure of the policy engine. Where both were measured, stronger enforcement came at the price of extra latency or resource use.

The field's main shortfall is therefore comparable evidence rather than new architectures. We offer ZT-EAF and a nine-item minimum reporting checklist to make future ZTA evaluations reproducible and comparable.

KEYWORDS: Zero Trust Architecture; Systematic Review; Evidence Appraisal; Access Control; Performance Evaluation; Reproducibility.

1. INTRODUCTION

1.1 Background and motivation

For a long time, enterprise security rested on a simple assumption: a user or device inside the corporate network is more trustworthy than one outside it. That assumption has weakened with cloud services, hybrid and remote work, mobile endpoints, Internet of Things (IoT) deployments and distributed applications. Once an attacker gains a foothold inside a flat network, lateral movement is often easy. Zero Trust Architecture (ZTA) answers this by dropping implicit trust based on network location. Every request for a resource is authenticated, authorised and checked against policy before a session begins, and the decision can be revisited as conditions change [1].

The idea was popularised by Kindervag at Forrester Research [2], and Google's BeyondCorp programme showed that an enterprise could grant access on the basis of device and user credentials instead of network position [3]. NIST Special Publication 800-207 later gave the concept a vendor-neutral reference architecture built around a Policy Engine (PE), a Policy Administrator (PA) and Policy Enforcement Points (PEPs) [1]. Government guidance has since added maturity models [4], access-control patterns for cloud-native, multi-location environments [5] and example implementations assembled from commercial products [6]. Figure 1 shows the logical components that most of the reviewed studies build on.

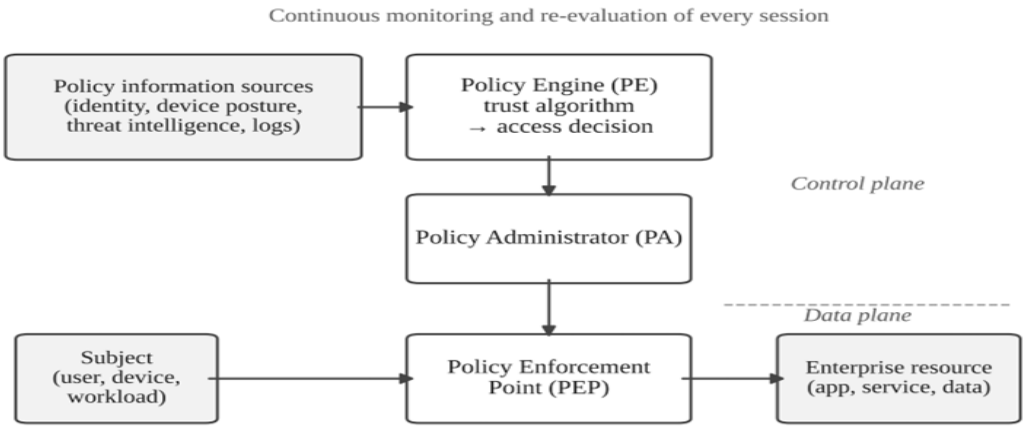


Figure 1. Simplified logical components of a Zero Trust Architecture. Adapted by the authors from NIST SP 800-207 [1].

1.2 Core principles

Five principles recur across standards and surveys [1, 4, 7]. *Verify explicitly*: decisions draw on identity, device posture, resource sensitivity and context rather than network location. *Apply least privilege*: access is limited to the resource, operation and time window actually needed. *Assume breach*: controls are designed to contain lateral movement and limit the blast radius once a compromise occurs. *Protect resources, not segments*: the objects of protection are applications, services, data and workloads. *Evaluate continuously*: trust is reassessed whenever identity, device state, behaviour or risk changes. For adoption, the Cybersecurity and Infrastructure Security Agency (CISA) arranges the work into five pillars (identity, devices, networks, applications and workloads, and data) and four maturity stages running from Traditional to Optimal [4].

1.3 Limitations of existing reviews

ZTA research has grown quickly and fragmented along the way. Some studies propose enterprise migration frameworks [8, 9]; others address adaptive trust [10], IoT and edge constraints [11, 12], hardware-rooted identity [13, 14], software-defined networking [15] or 5G/6G networks [16, 17]. Secondary literature has kept pace (Table 1). Recent systematic reviews map application domains, enabling technologies and adoption barriers [18] and classify security dimensions across domains [19], while earlier surveys supply vocabulary, taxonomies and migration challenges [7, 20–22].

What these reviews do less often is judge how strong the evidence behind a ZTA claim is. Kang et al. note that some approaches were validated only on artificially generated traffic, and that qualitative evaluation shows feasibility rather than effectiveness [21]. Dhiman et al. identify real-world deployment as an open problem [22], and Buck et al. find that academic work concentrates on architecture and performance improvement while economic and user-related questions are neglected [20]. Mushtaq et al. code each study as tested, simulated or conceptual [19] but, as far as we can tell, do not report security effectiveness and operational cost side by side. A reader therefore cannot easily tell which proposals have shown a security benefit at an acceptable cost and which remain conceptual.

Table 1. Existing ZTA surveys and reviews included in this study.

Review	Status	Scope and stated contribution	Method
Buck et al. (2021) [20]	Journal	Multivocal review of academic and grey ZT literature	Multivocal literature review
Teerakanok et al. (2021) [8]	Journal	Review of ZTA migration approaches and challenges	Narrative review

Syed et al. (2022) [7]	Journal	Comprehensive ZTA survey (authentication, access control, micro-segmentation)	Narrative survey
Phiayura & Teerakanok (2023) [9]	Journal	Process-driven migration framework synthesised from literature	Literature synthesis → framework
Kang et al. (2023) [21]	Journal	Trust-centred theory of ZT; cloud and IoT applications (WoS 2014–2022)	WoS search (2014–2022), selective analysis
Ghasemshirazi et al. (2023) [23]	Preprint	Narrative survey of ZT applications and challenges	Narrative survey
Bobelin (2023) [24]	Conf.	Industrial (vendor) literature review of ZT for IoT	Industrial (vendor) literature review
Dhiman et al. (2024) [22]	Journal	Review and comparison of ZT network model approaches	Narrative review, comparative tables
Nahar et al. (2024) [16]	Journal	ZTA survey for 6G networks	Narrative survey
Weinberg & Cohen (2024) [10]	Journal	ZT in emerging technologies (AI/ML, quantum, automation)	Narrative survey
Mavroudis (2024) [25]	Preprint	ZTNA principles, architectures and applications	Narrative analysis
Nasiruzzaman et al. (2025) [26]	Conf.	Evolution of ZTA from concept to implementation	Narrative review
Mushtaq et al. (2025) [19]	Journal	PRISMA SLR of 74 studies across domains; nine security dimensions	PRISMA SLR, six databases, 74 studies
Gambo & Almulhem (2026) [18]	Journal	PRISMA SLR of 91 studies (2016–2025); domain taxonomy, adoption barriers	PRISMA SLR, three databases, 91 studies

Source: compiled by the authors from the cited works. Methods are as stated by the authors of each review.

1.4 Research questions and contributions

This article extends an earlier 15-study survey by the authors into a structured review with an explicit appraisal step. It asks four research questions (RQs):

RQ1. How is ZTA research published between 2021 and 2026 distributed across application domains, techniques and publication types?

RQ2. What kinds of evidence (conceptual, analytical, simulated, prototype, hardware or organisational) support ZTA claims, and how are security effectiveness and operational cost reported?

RQ3. Which candidate research gaps are supported by the included studies, and to what extent have existing reviews or recent work already addressed them?

RQ4. How can ZTA studies be appraised consistently across security, performance, deployment and evidence quality?

The paper makes five contributions: (i) a re-verified and expanded set of 35 studies with corrected bibliographic details; (ii) the Zero Trust Evidence-Appraisal Framework (ZT-EAF), a four-dimension ordinal coding scheme that yields an evidence profile for each study instead of a single score; (iii) an application of ZT-EAF to all 21 primary studies, with a catalogue of the operational metrics they report; (iv) a research-gap matrix in which each candidate gap is tested against the included evidence and against existing reviews; and (v) a minimum reporting checklist for future ZTA evaluations. Section 2 describes the protocol, Section 3 reports the results, Section 4 discusses what they mean, and Section 5 concludes.

2. METHODOLOGY

2.1 Review design

We designed the work as a structured systematic literature review guided by Kitchenham and Charters [27] and reported it against the PRISMA 2020 items that apply [28]. Because ZTA practice is documented partly outside academic venues, we also admitted government technical reports, a doctoral dissertation and preprints under explicit rules adapted from guidance on multivocal reviews [29]. Publication status is recorded for every study and carried into the synthesis. No meta-analysis was attempted, since the included studies use incompatible designs and metrics (Section 3.4).

2.2 Starting set and information sources

The starting set was the 15 studies of our earlier survey. That report did not keep database hit counts or exclusion records, so we do not try to reconstruct them. Instead, every original study was re-verified against publisher, repository or arXiv records (authors, title, venue, year, identifier and publication status), and a new, documented search was run.

New records came from a general web search engine and from direct inspection of arXiv, the Directory of Open Access Journals (DOAJ), PubMed Central, publisher pages (IEEE, Elsevier, MDPI, Springer, ACM, OAE) and institutional repositories. Subscription databases (Scopus, Web of Science, IEEE Xplore search) and the Crossref and OpenAlex APIs could not be queried directly from our review environment; we treat this as a threat to validity (Section 2.9). Searches were run on 13 October 2026.

2.3 Search strategy

Each query paired a zero-trust term with a domain or evaluation term. Table 2 lists the query families used to find new records. Extra targeted queries were run only to verify the bibliographic details of individual studies; relevant records surfaced that way were screened under the same rules.

Table 2. Search query families used to identify new records.

No.	Query (as executed)	Purpose
Q1	zero trust architecture evaluation metrics benchmark testbed performance overhead latency experimental study 2023 2024	Evaluation and benchmarking
Q2	NIST SP 1800-35 Implementing a Zero Trust Architecture final 2025 NCCoE example implementations	Practice-based implementations
Q3	explainable AI zero trust access control trust score XAI paper	Explainable adaptive trust
Q4	zero trust microservices service mesh Kubernetes evaluation paper workload identity SPIFFE multi-cloud	Cloud-native and multi-cloud
Q5	zero trust IoT authentication scheme lightweight continuous trust evaluation simulation results journal 2023	IoT and constrained devices
Q6	software-defined perimeter zero trust performance evaluation experimental latency SDP	ZTNA / SDP performance
Q7	blockchain-based zero trust architecture access control evaluation 2023 2024 journal paper	Blockchain-based trust
Q8	"zero trust" survey taxonomy "evaluation" "evidence" maturity assessment framework comparing approaches security effectiveness performance deployment	Existing evaluation frameworks
Q9	zero trust dataset benchmark public dataset for zero trust access decisions identity device posture network flows	Datasets and benchmarks
Q10	zero trust trust evaluation machine learning user behavior continuous authentication experimental results dataset IEEE Transactions 2022 2023	AI/ML-driven trust
Q11	zero trust architecture healthcare OR industrial IoT implementation testbed evaluation journal 2024 2025 latency overhead measured	Sector and IIoT deployments
Q12	"zero trust" metrics measuring effectiveness quantitative assessment paper "evaluation framework" arXiv 2024 2025	Novelty check for ZT-EAF

Source: search log maintained by the authors.

2.4 Eligibility criteria

Inclusion criteria

- **I1.** Published or posted between January 2021 and October 2026.
- **I2.** Primary focus on ZTA, Zero Trust Network Access (ZTNA), Software-Defined Perimeter (SDP) or a trust-aware access-control mechanism explicitly framed as zero trust.
- **I3.** Enough technical or methodological detail to identify the contribution and how it was evaluated.
- **I4.** Bibliographic identity (authors, title, venue or repository, year) verifiable from a stable record.
- **I5.** For grey literature, issued by a standards body or archived as a full-text dissertation or preprint with a persistent identifier.

Exclusion criteria

- **E1.** Vendor marketing, blogs, solution briefs and patents.
- **E2.** Work not focused on zero trust.
- **E3.** Published before 2021 (foundational standards are cited for context but not counted as included studies).
- **E4.** Unpublished student theses without a persistent identifier.
- **E5.** Records whose full text or abstract could not be retrieved or read within the review window.
- **E6.** Dataset records rather than studies (discussed separately in Section 3.6).

2.5 Selection process and duplicate removal

Mirror links to the same work (for example a publisher page, an arXiv copy and an institutional repository copy) were collapsed into a single record before logging, and hits that duplicated a study already in the starting set were removed. The remaining records were screened on title, source type and date. We then assessed the full report against the eligibility criteria or, where only the abstract was accessible, the abstract and metadata. Every decision and its reason went into a screening log. Screening and eligibility assessment were carried out by a single reviewer.

2.6 Data extraction

For each study we extracted the following into a literature-extraction matrix: bibliographic details and publication status; study type (primary or secondary); application domain; ZTA mechanism or contribution; evaluation design (environment, tools, scale, baselines); reported security and operational metrics; stated limitations; and the material on which the coding was based (full text, partial full text, abstract or metadata). When a coding could not be confirmed from the accessed material, we recorded it as not verifiable (NV) rather than as absent.

2.7 Evidence appraisal: the ZT-EAF coding scheme

Existing maturity models assess *organisations* [4, 30], and existing reviews classify studies by domain or security dimension [18, 19]. Neither asks how strongly a single study supports its own claims, which is the role ZT-EAF is designed to fill. It codes four dimensions on short ordinal scales (Table 3, Figure 2). Results are reported as a profile and are deliberately **not** summed or weighted, because there is no empirical basis for trading a unit of security evidence against a unit of performance evidence. Secondary studies were not coded, since their evidence is the literature they review.

Table 3. ZT-EAF dimensions, codes and assessment criteria.

Dimension	Code	Assessment criterion
A. Security effectiveness	—	No security benefit, argument or outcome reported.
	Q	Security benefit asserted or described qualitatively; no structured analysis or measurement.
	F	Structured security analysis: formal proof, threat-model-based argument, or attack scenarios analysed against each ZT tenet.
	M	Measured security outcome, e.g. detection metrics on a dataset, blocked or detected attacks in a testbed, or statistically analysed incident data.
B. Operational performance	—	No operational cost reported.
	An	Analytical cost model (computation or communication complexity) without measurement.
	Sim	Latency, throughput, runtime or resource use measured in simulation or emulation.
	Meas	Latency, throughput, CPU, memory or energy measured on a prototype, testbed or real infrastructure.
C. Deployment feasibility	D	Infrastructure, interoperability, scalability or resource requirements discussed but not demonstrated.
	I	Integrated with real or standard platforms (e.g. production-grade service mesh, mobile core, blockchain platform, FPGA hardware, commercial products).
D. Evidence quality	Type	C conceptual · A analytical · S simulation/emulation · T prototype or lab testbed · H hardware testbed · E empirical organisational data.
	Status	Journal · conference · preprint · dissertation · government report.
	Artefacts	Code, configuration or data publicly released and confirmed (Y), or not confirmed (NV).
All dimensions	NV	Not verifiable from the material accessed; never interpreted as a negative finding.

Source: developed by the authors for this review.

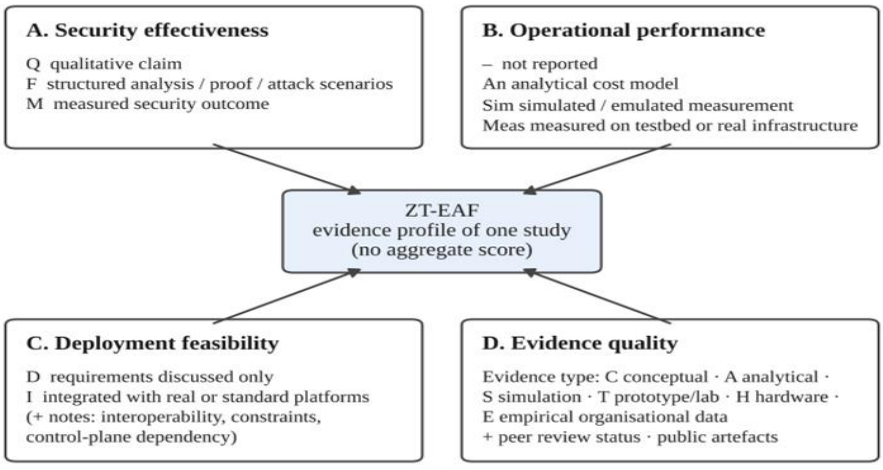


Figure 2. Conceptual structure of the proposed Zero Trust Evidence-Appraisal Framework (ZT-EAF). New figure developed by the authors.

2.8 Synthesis

We grouped studies into a domain taxonomy, tabulated the evidence profiles and counted them, and catalogued the reported metrics per study to test whether results were comparable. Each candidate research gap in the review protocol (benchmarking, real-world validation, joint reporting of security and cost, the security–overhead trade-off, lightweight trust for constrained devices, explainability, control-plane resilience, cross-domain interoperability and reproducibility) was treated as a hypothesis and checked against both the extracted evidence and the findings of existing reviews. Quantitative pooling was not attempted.

2.9 Threats to validity

- **Coverage.** Subscription databases could not be searched directly, and 49 reports sought for retrieval could not be read within the review window. Relevant studies, particularly paywalled journal articles, are therefore likely to be missing, and the review is not claimed to be exhaustive.
- **Coding basis.** Several codings rely on abstracts or partial full text. The NV code limits, but does not remove, the risk that an evaluation reported only in the full text was missed.
- **Reviewer bias.** ZT-EAF codes involve judgement. Screening and coding were carried out by a single reviewer, so reviewer bias in the coding process was not assessed.
- **Publication status.** Twelve of the 35 included studies are preprints. They are labelled throughout, and conclusions that rest mainly on preprints are flagged as such.
- **Time window.** The field moves quickly, and studies posted after the search date are not covered.

3. RESULTS

3.1 Study selection

The searches produced 101 records once mirror links were collapsed. Four duplicated studies already in the starting set and were removed, leaving 97 records for screening. Screening on title, source type and date excluded 24 of them (10 not focused on zero trust, 10 non-scholarly, vendor or patent items, 3 published before 2021 and 1 student thesis). Of the 73 reports sought for retrieval, 49 could not be retrieved or read within the review window, and the remaining 24 were assessed against the eligibility criteria. Four of these were excluded (two dataset records and two with unverifiable or incomplete metadata), so 20 new studies were included. Together with the 15 original studies, this gives 35 studies for synthesis (Figure 3).

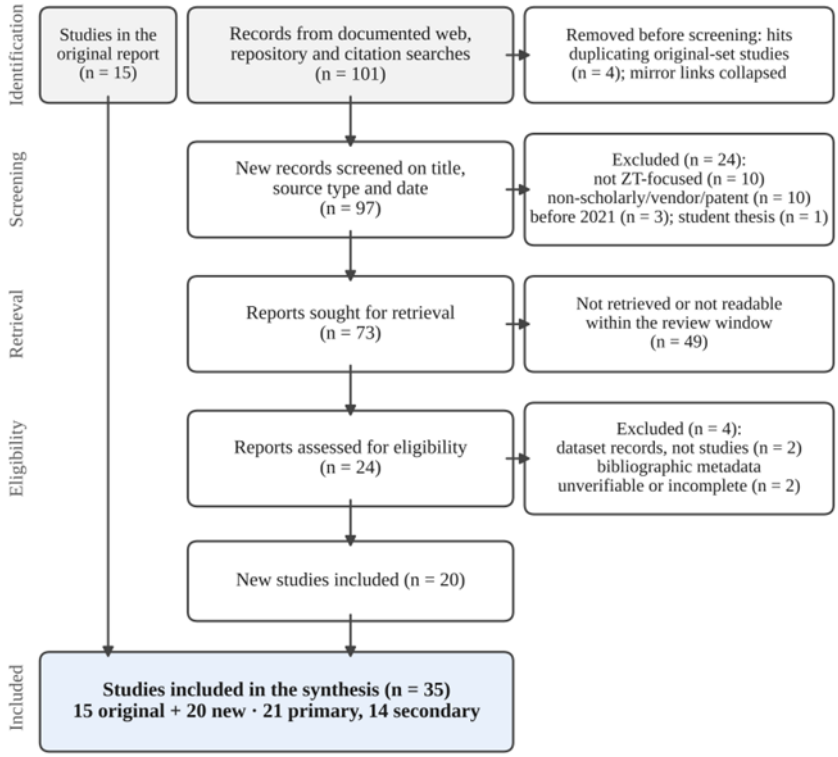


Figure 3. Identification, screening and inclusion of studies, using the counts recorded in the screening log. Layout follows the PRISMA 2020 flow diagram [28].

3.2 Characteristics of the included studies

The 35 studies comprise 21 primary studies and 14 secondary studies (surveys, reviews and literature-derived frameworks). Output grew from three studies in 2021 to eight in 2023 and peaked at ten in 2024, followed by eight in 2025 and two in 2026 up to the search date (Figure 4). Twenty-one studies appeared in peer-reviewed journals or conference proceedings, twelve were preprints, one was a doctoral dissertation and one a government technical report. The preprint share is noticeably higher among primary studies: only nine of the 21 were peer-reviewed, compared with twelve of the 14 secondary studies. Table 4 profiles each primary study.

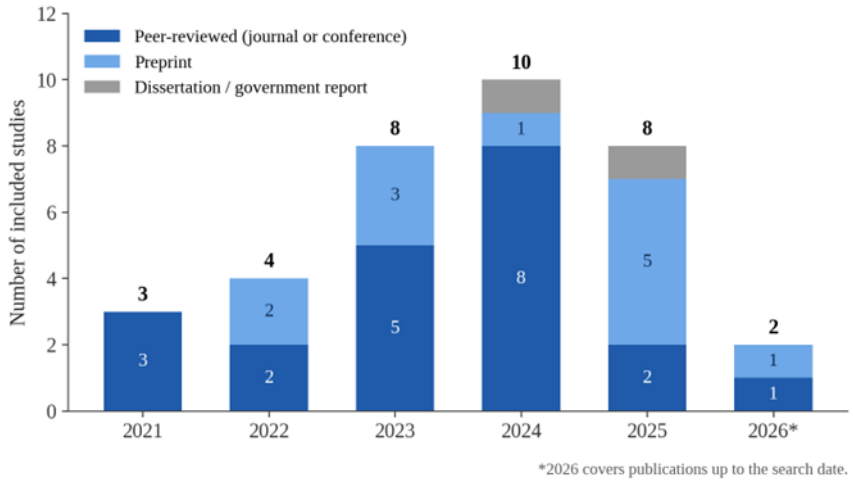


Figure 4. Included studies by year and publication status (n = 35). Chart prepared by the authors from the extraction matrix.

Table 4. Characteristics and ZT-EAF evidence profiles of the included primary studies.

ID	Study	Status	Domain	Contribution	Evidence type	A · B · C
P1	Rodigari et al. (2021) [31]	Conf.	Cloud / Multi-cloud	Performance cost of Istio-based ZT across multiple clouds	Prototype/ lab	Q · Meas · I
P2	García-Teodoro et al. (2022) [32]	Journal	Enterprise / ZTNA	SADAC: security-attribute-based dynamic access control	Prototype/ lab	Q · NV · I
P3	Chen et al. (2022) [33]	Preprint	5G / 6G	Software-defined ZTA for 6G with blockchain-based trust services	Simulation	M · NV · D
P4	Ramezanpour & Jagannath (2022) [17]	Preprint	5G / 6G	Intelligent ZTA (i-ZTA) for 5G/6G and O-RAN	Conceptual	Q · — · D
P5	Fang et al. (2023) [34]	Conf.	5G / 6G	Lightweight continuous authentication jointly designed with NOMA access	Simulation	F · Sim · D
P6	Mohseni-Ejiyeh (2023) [35]	Preprint	IoT / Edge	Lightweight ZT access-control protocol for cloud-based IoT sensors	Analytical	F · An · D
P7	Gai et al. (2023) [36]	Journal	Blockchain / Hardware	Blockchain RBAC for ZT cross-organisational data sharing	Prototype/ lab	NV · Meas · I

P8	Abdelhay et al. (2023) [37]	Preprint	5G / 6G	SDP with moving-target defence for the mobile core	Prototype/ lab	M · Meas · I
P9	Li et al. (2024) [11]	Journal	IoT / Edge	ZT architecture and blockchain device authentication for IIoT	Conceptual	Q · — · D
P10	Kulkarni et al. (2024a) [13]	Conf.	Blockchain / Hardware	Blockchain + RO-PUF ZTA for IoT	Hardware	NV · NV · I
P11	Kulkarni et al. (2024b) [14]	Journal	Blockchain / Hardware	Blockchain + RO-PUF ZT framework for FPGA supply chain	Hardware	F · NV · I
P12	Ashraf et al. (2024) [15]	Journal	SDN	ZFort: risk-aware ZT trust management and traffic engineering in SDN-IoT	Simulation	M · Sim · D
P13	Sedjelmaci et al. (2024) [38]	Journal	5G / 6G	ZTA-based 6G security framework with hierarchical AI defence agents	Conceptual	NV · NV · D
P14	Aiello (2024) [30]	Diss.	Enterprise / ZTNA	Prescriptive ZT controls and four-tier maturity model	Org. data	M · — · D
P15	Ma et al. (2025) [12]	Preprint	IoT / Edge	Cross-domain ZTA with decentralised federated learning for IoT	Simulation	F · Sim · D
P16	Shakya et al. (2025) [39]	Preprint	AI / ML	Zero-touch ZT AI/ML framework for IoT DDoS	Simulation	M · — · D
P17	Ghasemshirazi et al. (2025) [40]	Preprint	Sector-specific	ZTA for pharmaceutical supply chain	Conceptual	Q · — · D
P18	Dumitrescu & Pouwelse (2025) [41]	Preprint	Enterprise / ZTNA	TrustZero: self-sovereign trust tokens across organisations	Prototype/ lab	M · Meas · I
P19	Rajendran et al. (2025) [42]	Preprint	Cloud / Multi-cloud	Identity-federated ZT for microservices (OIDC, SPIFFE/SPIRE, Istio, OPA)	Prototype/ lab	M · Meas · I
P20	NIST NCCoE (2025) [6]	Gov. report	Enterprise / ZTNA	19 example ZTA builds with 24 collaborators (SP 1800-35)	Prototype/ lab	F · NV · I
P21	Deng & Inthiran (2026) [43]	Preprint	Enterprise / ZTNA	SME ZTA readiness pilot survey (n = 64, Asia-Pacific)	Org. data	— · — · D

Source: extraction matrix compiled by the authors. A = security effectiveness, B = operational performance, C = deployment feasibility; codes as in Table 3.

3.3 Taxonomy and comparison across domains

Table 5 groups all 35 studies into eight domains and summarises the dominant mechanisms and evidence pattern of each, and Figure 5 shows how the studies are spread. Reviews dominate the foundational cluster, enterprise and ZTNA work is the largest primary-study cluster, and 5G/6G follows closely. SDN is represented by a single study.

Table 5. Taxonomy of ZTA research domains in the included studies.

Domain	Studies	Representative mechanisms	Evidence pattern
Foundations and cross-domain reviews	[7, 18–23, 26]	Principles, taxonomies, adoption barriers, trust theory	Secondary; two PRISMA-based SLRs
Enterprise, migration and ZTNA	[6, 8, 9, 25, 30, 32, 41, 43]	Migration frameworks, security-profile-based access control, trust tokens, maturity tiers	Mixed: lab prototypes, organisational data, government example builds
Cloud-native and multi-cloud	[31, 42]	Service mesh with mutual TLS, OIDC/OAuth, SPIFFE/SPIRE workload identity, policy-as-code	Measured overhead on testbeds; preliminary or weakly documented
IoT, edge and federated learning	[11, 12, 24, 35]	Lightweight access control, cross-domain authentication, federated trust evaluation	Analytical or simulated; no constrained-device measurement
SDN	[15]	Risk-aware trust scoring, MILP-based traffic engineering	Emulation with baselines
5G/6G	[16, 17, 33, 34, 37, 38]	Software-defined ZTA, intelligent trust engines for O-RAN, continuous physical-layer authentication, SDP for the core	Mostly conceptual or simulated; one lab testbed
Blockchain and hardware-rooted trust	[13, 14, 36]	Blockchain ledgers with RO-PUF device identity, consortium-chain RBAC	Hardware and platform integration; overhead largely unreported in abstracts
AI/ML-driven trust and sector-specific use	[10, 39, 40]	ML-based DDoS detection, AI-assisted policy, supply-chain controls	Dataset-based detection metrics or qualitative cases

Source: compiled by the authors.

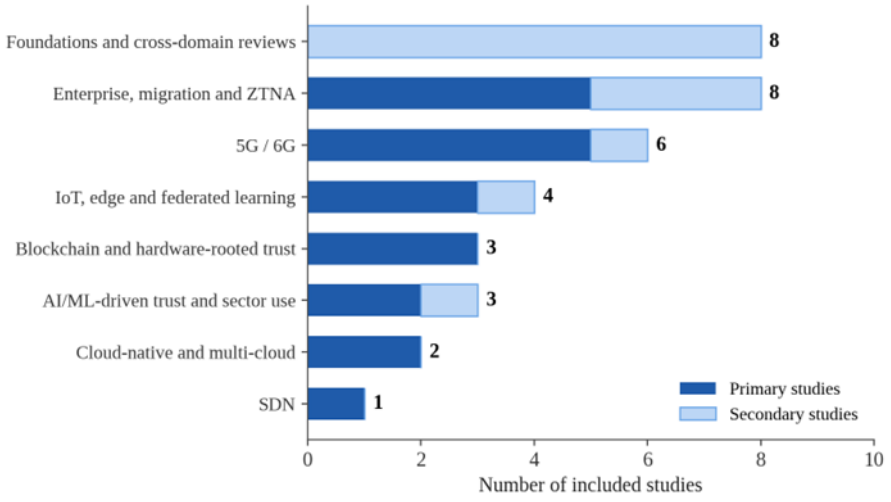


Figure 5. Distribution of the 35 included studies across the eight domains of Table 5, split into primary and secondary studies.

Three patterns cut across the domains. First, identity-centric verification, least privilege and continuous evaluation appear in nearly every study, whatever the domain [1, 7, 18]. Second, trust decisions are shifting from static rules to context- and risk-driven evaluation, using CVE/CVSS-based risk [15], federated models [12], machine-learning detectors [39] or hierarchical AI agents [38]. Third, the domains with the tightest resource budgets, IoT and 5G/6G, have the thinnest measured evidence; their proposals are mostly analytical, simulated or conceptual [12, 17, 33, 35, 38].

3.4 Evaluation techniques and metrics

Ten primary studies reported their evaluation in enough detail to catalogue (Table 6). No two used the same combination of metrics, environment and baseline, and three consequences follow. First, performance means a different quantity in each study: tail-latency variability [31], session initialisation time [37], rule-processing time [41], achievable sum rate [34] or solver runtime [15]. Second, where security effectiveness was measured, it was measured against different threats: DDoS traffic classification [39], port scanning and a known TLS vulnerability [37], signature tampering [41] or token replay [42]. Third, baselines range from a non-ZT configuration to competing ZT schemes, so improvement percentages cannot be compared across studies. Pooling the numbers would therefore be misleading, which is why our synthesis is qualitative.

Table 6. Evaluation environments and metrics reported by primary studies.

Study	Environment	Metrics reported	Baseline
Rodigari et al. [31]	Istio service mesh across multiple clouds	Data-plane latency and its variability under load; CPU and memory; control-plane impact	Mesh disabled
Fang et al. [34]	Simulation of NOMA uplink	Achievable sum rate under continuous authentication	Existing access schemes
Mohseni-Ejiyeh [35]	Analytical	Symbolic computation cost per phase and entity	Feature checklist vs four schemes
Gai et al. [36]	Hyperledger Fabric	Platform benchmarks via Caliper and BFT-SMaRT	NV
Abdelhay et al. [37]	Open5GS core on virtual machines; SDP vs OpenVPN	Initialisation time; RTT and throughput under DoS; CPU and memory; port-scan and Heartbleed outcomes	OpenVPN
Ashraf et al. [15]	Mininet with OpenDaylight; 50–200 and 500 nodes	Flow satisfaction, trust compliance, average path trust, runtime	MRTS, SECOD, greedy variants
Ma et al. [12]	Simulation	Latency, throughput; security proof	Representative existing schemes
Shakya et al. [39]	ML evaluation on IoT traffic	Accuracy, precision, recall, F1, ROC-AUC	Five ML classifiers
Dumitrescu & Pouwelse [41]	Docker containers behind ModSecurity; Android phones	Rule-processing time, latency vs signatures, load up to 2,000 users	Zero-signature case
Rajendran et al. [42]	Istio, Keycloak, SPIRE, OPA prototype	Authentication and authorisation latency, API response time, requests per second, attack counts	Non-ZT configuration

Source: compiled by the authors from the cited studies; NV = not verifiable from the material accessed.

3.5 Evidence quality and reproducibility

Figure 6 shows the ZT-EAF profile of every primary study, and Figure 7 summarises the codes. By evidence type, four studies were conceptual, one analytical, five based on simulation or emulation, seven on prototypes or lab testbeds, two on hardware testbeds and two on organisational data. Seven reported a measured security outcome and nine demonstrated integration with real or standard platforms. Only nine reported any operational cost, five of them measured on a testbed or real infrastructure, and for a further six we could not verify whether performance was evaluated at all.

	D. Evidence type	A. Security effectiveness	B. Operational performance	C. Deployment feasibility	D. Publication status
P1 Rodigari et al. (2021)	Prototype/lab	Q	Meas	I	Conference
P2 García-Teodoro et al. (2022)	Prototype/lab	Q	NV	I	Journal
P3 Chen et al. (2022)	Simulation	M	NV	D	Preprint
P4 Ramezanpour & Jagannath (2022)	Conceptual	Q	none	D	Preprint
P5 Fang et al. (2023)	Simulation	F	Sim	D	Conference
P6 Mohseni-Ejiyeh (2023)	Analytical	F	An	D	Preprint
P7 Gai et al. (2023)	Prototype/lab	NV	Meas	I	Journal
P8 Abdelhay et al. (2023)	Prototype/lab	M	Meas	I	Preprint
P9 Li et al. (2024)	Conceptual	Q	none	D	Journal
P10 Kulkarni et al. (2024a)	Hardware	NV	NV	I	Conference
P11 Kulkarni et al. (2024b)	Hardware	F	NV	I	Journal
P12 Ashraf et al. (2024)	Simulation	M	Sim	D	Journal
P13 Sedjelmaci et al. (2024)	Conceptual	NV	NV	D	Journal
P14 Aiello (2024)	Org. data	M	none	D	Dissertation
P15 Ma et al. (2025)	Simulation	F	Sim	D	Preprint
P16 Shakya et al. (2025)	Simulation	M	none	D	Preprint
P17 Ghasemshirazi et al. (2025)	Conceptual	Q	none	D	Preprint
P18 Dumitrescu & Pouwelse (2025)	Prototype/lab	M	Meas	I	Preprint
P19 Rajendran et al. (2025)	Prototype/lab	M	Meas	I	Preprint
P20 NIST NCCoE (2025)	Prototype/lab	F	NV	I	Gov. report
P21 Deng & Inthiran (2026)	Org. data	none	none	D	Preprint

Shading increases with strength of reported evidence (white = none; lightest = qualitative/discussed; darkest = measured/integrated). NV = not verifiable from the material accessed (hatched). Codes follow the ZT-EAF definitions.

Figure 6. ZT-EAF evidence profiles of the 21 included primary studies, coded by the authors from the extraction matrix.

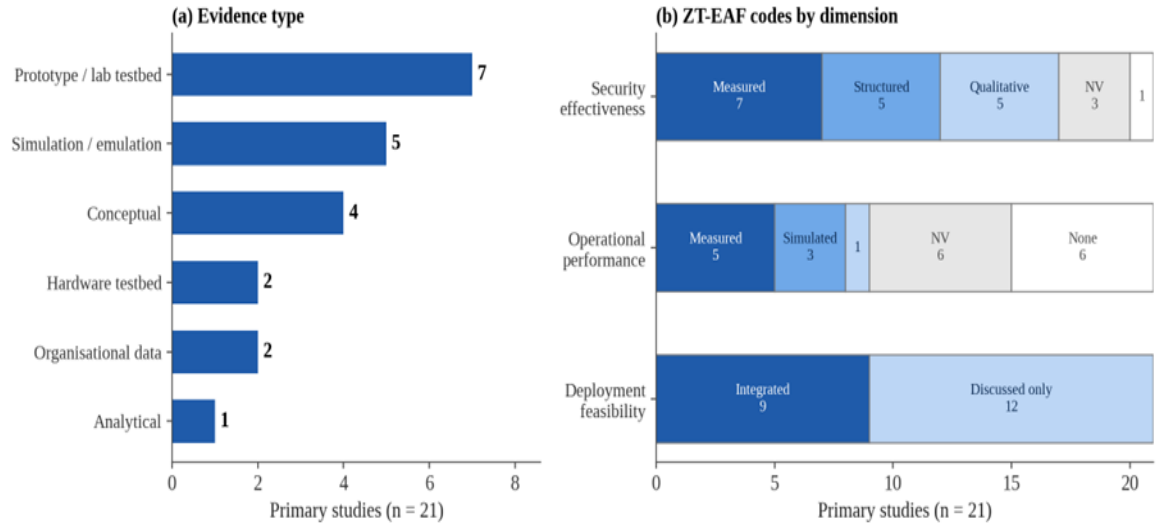


Figure 7. Distribution of evidence types (a) and ZT-EAF codes (b) across the 21 primary studies. NV = not verifiable; An = analytical cost model; None = nothing reported.

Only three studies combine a measured security outcome with a measured operational cost [37, 41, 42]. All three are preprints, and one of them documents neither its workload nor its statistical procedure and contains citation inconsistencies [42]. Public release of code, configurations or data could be confirmed for just one primary study: the NIST example implementations, which are documented well enough to be replicated [6]. Organisational evidence is thin, resting on a pilot survey of 64 professionals in the Asia-Pacific region [43] and a dissertation that analysed data from 138 organisations to propose a four-tier maturity model [30].

3.6 Research-gap mapping

We tested each candidate gap against the extracted evidence and against existing reviews (Table 7). A gap is marked *validated* when the included evidence supports it and no included study or review resolves it, and *partially addressed* when recent work has begun to close it. Of the nine candidates, six were validated, two were partially addressed and one was only partially supported.

Table 7. Research-gap matrix.

ID	Research gap (status)	Supporting evidence	Existing review coverage	Remaining problem and contribution of this review
G1	No standard ZTA benchmark or metric set *(partially addressed)*	Heterogeneous metrics in Table 6 [15, 31, 37, 41]; ZTA-specific datasets appeared only in 2025–2026 [44, 45]	Noted qualitatively [21, 22]	**Remaining:** datasets exist but are new, partly subscription-only and not yet adopted; no shared evaluation protocol. **This review:** metric catalogue; minimum reporting checklist (Table 8).
G2	Limited real-world validation *(validated)*	Four conceptual and five simulation studies; no peer-reviewed production measurement	Identified as open [21, 22]	**Remaining:** government example builds [6] are lab-based and do not report comparative performance. **This review:** evidence-type coding makes the deficit visible per study.
G3	Security and cost rarely reported together; claims exceed demonstrated outcomes *(validated)*	Three of 21 primary studies measure both, all preprints	Not reported side by side in prior reviews	**Remaining:** proposals cannot be ranked on a security–cost frontier. **This review:** ZT-EAF profile places both dimensions side by side.

G4	Security–overhead trade-off unquantified in general *(validated)*	Added latency or resource use where measured [31, 37, 42]	Overhead noted qualitatively [21, 26]	**Remaining:** no study varies enforcement strength systematically. **This review:** recommends controlled ablation of ZT controls.
G5	Lightweight adaptive trust for constrained devices *(validated)*	IoT proposals are analytical or simulated [11, 12, 35]	Identified [21, 22]	**Remaining:** no energy or memory measurement on constrained hardware in the included set. **This review:** energy and memory added as required metrics.
G6	Explainability and auditability of AI-driven trust *(partially supported)*	AI-based studies report detection metrics but not explanations [17, 38, 39]	AI integration discussed as opportunity [10, 22]	**Remaining:** absence of evidence in this sample; a candidate XAI study could not be retrieved. **This review:** flagged as priority; checklist item on decision records.
G7	Resilience of centralised policy engines and controllers *(validated)*	Controller dependency in SDN [15]; control-plane vulnerability noted [26]	Single point of failure discussed [21]	**Remaining:** no included study measures PE/PA or controller failure, fail-open/fail-closed behaviour or recovery time. **This review:** control-plane failure added as an evaluation item.
G8	Cross-domain and multi-cloud policy interoperability *(partially addressed)*	Cross-domain or cross-organisation proposals [12, 36, 41, 42]; multi-cloud measurement [31]; guidance [5]	Interoperability listed among challenges [18, 24]	**Remaining:** no evaluation of policy consistency across heterogeneous engines or administrative domains. **This review:** interoperability recorded under deployment feasibility.
G9	Limited reproducibility and comparability *(validated)*	Public artefacts confirmed for one of 21 primary studies	Not appraised in prior reviews	**Remaining:** results cannot be re-run or compared. **This review:** artefact availability coded for every study.

Source: developed by the authors from the extraction matrix and the cited reviews.

Two further observations sit outside the matrix. Buck et al. reported that economic and user-centred studies of zero trust were neglected [20], and the included set offers little to change that conclusion: organisational evidence is limited to one pilot survey [43] and one dissertation [30]. Separately, two ZTA-specific datasets released in 2025–2026 were excluded under criterion E6 because they are data records rather than studies [44, 45]. They

matter for the benchmarking gap (G1), but they are new and have not yet been adopted as a shared evaluation protocol.

4. DISCUSSION

4.1 What the evidence profiles reveal

Applying ZT-EAF separates three groups of studies that are usually presented together. The first demonstrates *integration* without a measured security benefit or cost, for example a dynamic access-control scheme deployed on a corporate wireless network [32] or blockchain-and-PUF designs implemented on FPGA hardware [14]. The second demonstrates *security or performance in simulation*, such as risk-aware routing in SDN-based IoT [15], federated cross-domain authentication [12] and lightweight continuous authentication [34]. The third, small group *measures both* on a prototype [37, 41, 42]. Conceptual frameworks for 6G and industrial IoT [11, 17, 38] sit outside all three; they set research agendas rather than demonstrate outcomes.

The profiles also show where claims outrun evidence. Several abstracts report improved security, resilience or efficiency without accompanying figures [14, 32, 40]. Such claims may well be substantiated in the full texts, but they cannot be compared, and a reader who relies on abstracts, as many surveys do, would overestimate the maturity of the field.

4.2 Security–overhead trade-offs

Where both security and cost were measured, the direction of the trade-off was consistent (Figure 8). In the mobile-core testbed, SDP left no scanned port open (the VPN server exposed SSH) and its gateway refused a Heartbleed attempt, yet it needed slightly longer session initialisation (3.26 s against 2.95 s) and more gateway CPU and memory than OpenVPN [37]. In the microservice prototype, adding identity federation, workload identity and policy-as-code raised mean authentication latency from 95 ms to 118 ms and average API response time from 210 ms to 250 ms, while throughput fell from 980 to 915 requests per second [42]. A multi-cloud service-mesh study found that CPU and memory overhead depended on configuration and cloud environment, whereas latency variability for sequential requests decreased [31].

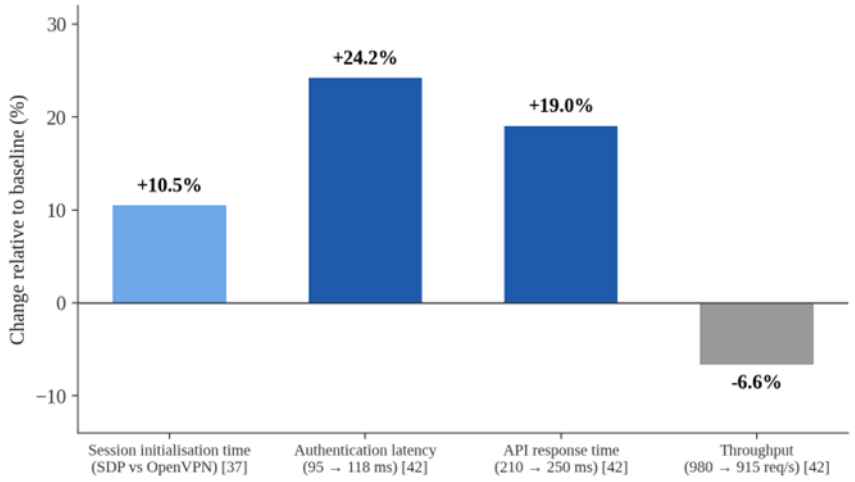


Figure 8. Relative change in measured cost between a ZT-enabled and a baseline configuration, computed by the authors from values reported in [37] and [42]. Each pair comes from a single non-peer-reviewed study and should not be generalised.

These figures come from single, non-peer-reviewed or preliminary studies, so they illustrate rather than establish the size of the cost. What they do show is that ZTA controls carry measurable overhead that practitioners need to budget for. Because no included study varied enforcement strength systematically, the shape of the curve between protection and cost remains unmapped (gap G4).

4.3 Implications for practice and research

For practitioners, three points follow. First, vendor-neutral guidance and example builds [1, 4–6] remain the best-documented basis for design, because academic evidence on comparative effectiveness is still sparse. Second, organisations should measure decision latency and enforcement overhead in their own environments before extending continuous verification to latency-sensitive or resource-constrained systems. Third, the policy engine and controllers deserve treatment as critical assets whose failure behaviour is tested, since no included study evaluated it.

For researchers, we propose the minimum reporting checklist in Table 8, derived directly from the gaps in Table 7. Studies that follow it can be placed in the same ZT-EAF profile and compared directly.

Table 8. Proposed minimum reporting checklist for ZTA evaluations.

No.	Item	What to report
R1	Threat model	Attacker capabilities, assets and the ZT tenets each control is meant to enforce.
R2	Security outcome	Detection or prevention metrics (including false-positive and false-negative rates) against a stated baseline.
R3	Decision latency and throughput	Median and tail access-decision latency and sustained requests per second under a documented load.
R4	Resource and communication cost	CPU, memory, bandwidth and, for constrained devices, energy for PE, PA and PEP components.
R5	Scale	Numbers of subjects, devices, policies and requests tested, and how cost scales with each.
R6	Control-plane failure	Behaviour when the PE, PA or controller is unavailable (fail-open or fail-closed) and recovery time.
R7	Environment and interoperability	Platforms, versions, standards and administrative domains involved.
R8	AI-based trust	Dataset provenance, leakage-aware data splits and an auditable record or explanation of each decision.
R9	Artefacts and evidence type	Public code, configuration and data, and an explicit statement of evidence type (Table 3).

Source: developed by the authors.

4.4 Comparison with existing surveys

This review overlaps with recent systematic reviews in scope but differs in purpose. Gambo and Almulhem screened 1,774 records from Scopus, Web of Science and IEEE Xplore and included 91 studies (as reported in the preprint version) to build a domain taxonomy and analyse adoption barriers [18]. Mushtaq et al. included 74 studies from six databases and coded nine security dimensions [19]. Both are broader than ours, which draws on 35 studies from a narrower search. What we add is a different lens: a per-study appraisal of evidence strength across security, performance, deployment and evidence quality, a catalogue of reported metrics, and gap statements tested against that appraisal. Narrative surveys on 6G [16], ZTNA [25], emerging technologies [10] and industrial IoT offerings [24] provide domain depth that we do not attempt to reproduce.

4.5 Limitations

The limitations in Section 2.9 bound these findings. In particular, 49 reports could not be read and several codings rest on abstracts, so the counts in Section 3.5 describe the included evidence rather than the field as a whole. The finding that few studies jointly measure security and cost holds in direction under these limits, but its magnitude may change with

fuller retrieval. ZT-EAF itself has not yet been validated by independent coders, and its ordinal levels, while explicit, involve judgement.

5. CONCLUSION

This review re-verified a 15-study survey of Zero Trust Architecture and extended it into a structured review of 35 studies published between 2021 and 2026, adding ZT-EAF to appraise the evidence behind them.

On **RQ1**, research is spread across eight domains, with the largest clusters in foundational reviews, enterprise and ZTNA work, and 5G/6G, and with a heavy reliance on preprints among primary studies. On **RQ2**, the evidence leans toward conceptual and simulated work, operational cost is rarely reported, and joint measurement of security and cost is rarer still; differences in metrics, environments and baselines rule out pooling. On **RQ3**, six of the nine candidate gaps were validated (real-world validation, joint reporting, trade-off quantification, lightweight trust, control-plane resilience and reproducibility), two were partially addressed (benchmarks and datasets, cross-domain interoperability) and one, explainability of AI-driven trust, was only partially supported. On **RQ4**, ZT-EAF offers a transparent way to profile studies on four dimensions without an arbitrary aggregate score, and the accompanying checklist spells out what future evaluations should report.

For researchers, the priority is not another architecture but comparable evidence: shared datasets and testbeds, evaluations that report security outcomes and costs together, and released artefacts. Promising directions include measuring lightweight trust evaluation on constrained hardware, controlled studies that vary enforcement strength, failure testing of policy engines and controllers, interoperability tests across administrative domains and clouds, and auditable AI-based trust decisions. For practitioners, standards-based guidance combined with in-house measurement of decision latency and overhead remains the sound route. These conclusions are bounded by the coverage limits of the review and should be revisited as the newly released ZTA datasets and evaluation studies mature.

ACKNOWLEDGEMENTS

The authors thank Prof. Ajitkumar Khachane, Department of Information Technology, Vidyalankar Institute of Technology, for guidance during the survey on which this article builds.

REFERENCES

1. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero trust architecture* (NIST Special Publication 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207
2. Kindervag, J. (2010). *No more chewy centers: Introducing the zero trust model of information security*. Forrester Research.
3. Ward, R., & Beyer, B. (2014). BeyondCorp: A new approach to enterprise security. *;login: The USENIX Magazine, 39*(6), 6–11.
4. Cybersecurity and Infrastructure Security Agency. (2023). *Zero trust maturity model* (Version 2.0). https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf
5. Chandramouli, R., & Butcher, Z. (2023). *A zero trust architecture model for access control in cloud-native applications in multi-location environments* (NIST Special Publication 800-207A). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207A
6. National Institute of Standards and Technology, National Cybersecurity Center of Excellence. (2025). *Implementing a zero trust architecture: High-level document* (NIST Special Publication 1800-35). https://doi.org/10.6028/NIST.SP.1800-35
7. Syed, N. F., Shah, S. W., Shaghaghi, A., Anwar, A., Baig, Z., & Doss, R. (2022). Zero trust architecture (ZTA): A comprehensive survey. *IEEE Access, 10*, 57143–57179. https://doi.org/10.1109/ACCESS.2022.3174679
8. Teerakanok, S., Uehara, T., & Inomata, A. (2021). Migrating to zero trust architecture: Reviews and challenges. *Security and Communication Networks, 2021*, 9947347. https://doi.org/10.1155/2021/9947347
9. Phiayura, P., & Teerakanok, S. (2023). A comprehensive framework for migrating to zero trust architecture. *IEEE Access, 11*, 19487–19511. https://doi.org/10.1109/ACCESS.2023.3248622
10. Weinberg, A. I., & Cohen, K. (2024). Zero trust implementation in the emerging technologies era: A survey. *Complex Engineering Systems, 4*(3), 16. https://doi.org/10.20517/ces.2024.41
11. Li, S., Iqbal, M., & Saxena, N. (2024). Future industry internet of things with zero-trust security. *Information Systems Frontiers, 26*(5), 1653–1666. https://doi.org/10.1007/s10796-021-10199-5

12. Ma, X., Fang, F., & Wang, X. (2025). *Dynamic authentication and granularized authorization with a cross-domain zero trust architecture for federated learning in large-scale IoT networks* (arXiv:2501.03601). arXiv. https://doi.org/10.48550/arXiv.2501.03601
13. Kulkarni, A., Niraula, A., Bhattarai, H., & Niamat, M. (2024a). A zero trust architecture employing blockchain and ring oscillator physical unclonable function for Internet-of-Things. In *2024 IEEE International Conference on Electro Information Technology (eIT)* (pp. 508–513). IEEE. https://doi.org/10.1109/eIT60633.2024.10609880
14. Kulkarni, A., Hazari, N. A., & Niamat, M. Y. (2024b). A zero trust-based framework employing blockchain technology and ring oscillator physical unclonable functions for security of field programmable gate array supply chain. *IEEE Access, 12*, 89322–89338. https://doi.org/10.1109/ACCESS.2024.3418572
15. Ashraf, U., Al-Naeem, M., Bhutta, M. N. M., & Yuen, C. (2024). ZFort: A scalable zero-trust approach for trust management and traffic engineering in SDN based IoTs. *Internet of Things, 28*, 101419. https://doi.org/10.1016/j.iot.2024.101419
16. Nahar, N., Andersson, K., Schelén, O., & Saguna, S. (2024). A survey on zero trust architecture: Applications and challenges of 6G networks. *IEEE Access, 12*, 94753–94764. https://doi.org/10.1109/ACCESS.2024.3425350
17. Ramezanpour, K., & Jagannath, J. (2022). *Intelligent zero trust architecture for 5G/6G networks: Principles, challenges, and the role of machine learning in the context of O-RAN* (arXiv:2105.01478v3). arXiv. https://doi.org/10.48550/arXiv.2105.01478
18. Gambo, M. L., & Almulhem, A. (2026). Zero trust architecture: A systematic literature review. *Journal of Network and Systems Management, 34*(1), 25. https://doi.org/10.1007/s10922-025-09998-x
19. Mushtaq, S., Mohsin, M., & Mushtaq, M. M. (2025). A systematic literature review on the implementation and challenges of zero trust architecture across domains. *Sensors, 25*(19), 6118. https://doi.org/10.3390/s25196118
20. Buck, C., Olenberger, C., Schweizer, A., Völter, F., & Eymann, T. (2021). Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust. *Computers & Security, 110*, 102436. https://doi.org/10.1016/j.cose.2021.102436
21. Kang, H., Liu, G., Wang, Q., Meng, L., & Liu, J. (2023). Theory and application of zero trust security: A brief survey. *Entropy, 25*(12), 1595. https://doi.org/10.3390/e25121595

22. Dhiman, P., Saini, N., Gulzar, Y., Turaev, S., Kaur, A., Nisa, K. U., & Hamid, Y. (2024). A review and comparative analysis of relevant approaches of zero trust network model. *Sensors, 24*(4), 1328. https://doi.org/10.3390/s24041328
23. Ghasemshirazi, S., Shirvani, G., & Alipour, M. A. (2023). *Zero trust: Applications, challenges, and opportunities* (arXiv:2309.03582). arXiv. https://doi.org/10.48550/arXiv.2309.03582
24. Bobelin, L. (2023). Zero trust in the context of IoT: Industrial literature review, trends, and challenges. In *Proceedings of the Computer & Electronics Security Application Rendezvous (C&ESAR 2023)*. CEUR Workshop Proceedings. https://arxiv.org/abs/2604.06272
25. Mavroudis, V. (2024). *Zero-trust network access (ZTNA)* (arXiv:2410.20611). arXiv. https://doi.org/10.48550/arXiv.2410.20611
26. Nasiruzzaman, M., Ali, M., Salam, I., & Miraz, M. H. (2025). The evolution of zero trust architecture (ZTA) from concept to implementation. In *2025 29th International Conference on Information Technology (IT)*. IEEE. https://doi.org/10.1109/IT64745.2025.10930254
27. Kitchenham, B., & Charters, S. (2007). *Guidelines for performing systematic literature reviews in software engineering* (Technical Report EBSE-2007-01). Keele University and Durham University.
28. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ, 372*, n71. https://doi.org/10.1136/bmj.n71
29. Garousi, V., Felderer, M., & Mäntylä, M. V. (2019). Guidelines for including grey literature and conducting multivocal literature reviews in software engineering. *Information and Software Technology, 106*, 101–121. https://doi.org/10.1016/j.infsof.2018.09.006
30. Aiello, S. T. (2024). *Prescriptive zero trust: Assessing the impact of zero trust on cyber attack prevention* [Doctoral dissertation, Dakota State University]. arXiv:2508.12953. https://doi.org/10.48550/arXiv.2508.12953
31. Rodigari, S., O'Shea, D., McCarthy, P., McCarry, M., & McSweeney, S. (2021). Performance analysis of zero-trust multi-cloud. In *2021 IEEE 14th International Conference on Cloud Computing (CLOUD)*. IEEE. https://arxiv.org/abs/2105.02334

32. García-Teodoro, P., Camacho, J., Maciá-Fernández, G., López-Marín, V. J., & Gómez-Hernández, J. A. (2022). A novel zero-trust network access control scheme based on the security profile of devices and users. *Computer Networks, 212*, 109068. https://doi.org/10.1016/j.comnet.2022.109068
33. Chen, X., Feng, W., Ge, N., & Zhang, Y. (2022). *Zero trust architecture for 6G security* (arXiv:2203.07716). arXiv. https://doi.org/10.48550/arXiv.2203.07716
34. Fang, H., Wang, X., Al-Dhahir, N., & Schober, R. (2023). Joint design of multi-dimensional multiple access and lightweight continuous authentication in zero-trust environments. In *GLOBECOM 2023 – 2023 IEEE Global Communications Conference* (pp. 3366–3371). IEEE. https://doi.org/10.1109/GLOBECOM54140.2023.10437443
35. Mohseni-Ejiyeh, A. (2023). *Zero trust real-time lightweight access control protocol for mobile cloud-based IoT sensors* (arXiv:2309.01293). arXiv. https://doi.org/10.48550/arXiv.2309.01293
36. Gai, K., She, Y., Zhu, L., Choo, K.-K. R., & Wan, Z. (2023). A blockchain-based access control scheme for zero trust cross-organizational data sharing. *ACM Transactions on Internet Technology, 23*(3), Article 38. https://doi.org/10.1145/3511899
37. Abdelhay, Z., Bello, Y., & Refaey, A. (2023). *Towards zero-trust 6GC: A software defined perimeter approach with dynamic moving target defense mechanism* (arXiv:2312.17271). arXiv. https://doi.org/10.48550/arXiv.2312.17271
38. Sedjelmaci, H., Tourki, K., & Ansari, N. (2024). Enabling 6G security: The synergy of zero trust architecture and artificial intelligence. *IEEE Network, 38*(3), 171–177. https://doi.org/10.1109/MNET.2023.3326003
39. Shakya, S., Abbas, R., & Maric, S. (2025). *A novel zero-touch, zero-trust, AI/ML enablement framework for IoT network security* (arXiv:2502.03614). arXiv. https://doi.org/10.48550/arXiv.2502.03614
40. Ghasemshirazi, S., Shirvani, G., Ranjbar Tavakoli, M., Ghaedi, B., & Langarizadeh, M. A. (2025). *Implementing zero trust architecture to enhance security and resilience in the pharmaceutical supply chain* (arXiv:2508.15776). arXiv. https://doi.org/10.48550/arXiv.2508.15776
41. Dumitrescu, A.-T., & Pouwelse, J. (2025). *TrustZero – open, verifiable and scalable zero-trust* (arXiv:2502.10281). arXiv. https://doi.org/10.48550/arXiv.2502.10281
42. Rajendran, R. N., Anumula, S. K., Rai, D. K., & Agrawal, S. (2025). *Zero trust security model implementation in microservices architectures using identity federation* (arXiv:2511.04925). arXiv. https://doi.org/10.48550/arXiv.2511.04925

43. Deng, Y., & Inthiran, A. (2026). *Towards zero trust architecture: A pilot study on information systems security readiness amongst small and medium enterprises* (arXiv:2605.18901). arXiv. https://doi.org/10.48550/arXiv.2605.18901
44. Rocha, R. S. (2025). *ZTA-RAD – Zero trust architecture – risk assessment dataset* (Version 1) [Data set]. Mendeley Data. https://doi.org/10.17632/d6hvsrxdts.1
45. Xu, S., & Qian, Y. (2026). *ZGG-RR: An enhanced multi-layer zero trust access dataset with leakage-aware benchmarks* [Data set]. IEEE DataPort. https://doi.org/10.21227/tmxr-6b67