

Measuring Zero Trust Maturity: An Evidence-Based Assessment Instrument Aligned with NIST SP 800-207, CISA's Maturity Model, and the 2024 Federal Milestone

Sujit Biswas

ORCID: 0000-0002-6770-9845

City, University of London, London, United Kingdom

Md Monjurul Karim

ORCID: 0000-0003-1753-065X

Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China

Anindya Nag

ORCID: 0000-0001-6518-8233

Department of Computer Science and Engineering, Northern University of Business and Technology Khulna, Khulna, Bangladesh

Correspondence: Md Monjurul Karim · contact details held in the journal record

Publication history. First published in the offline edition of Cybersecurity Advances, Volume 1, Issue 1 (December 2024). This publisher-authorized digital archival edition was released on 12 October 2026. The scholarly content is reproduced from the supplied offline-edition manuscript; review-stage notices were removed and journal typography, pagination, metadata, and accessibility were normalized.

© 2024 the author(s). Published by AlphaTech Scholarly Press under the Creative Commons Attribution 4.0 International License (CC BY 4.0). No DOI is displayed unless and until registration is verified.

Background: Zero trust has become the organizing security paradigm of the mid-2020s, but its adoption is routinely claimed more often than it is demonstrated. The 2024 policy environment sharpened this problem: OMB M-22-09's federal zero-trust goals came due in September 2024, NIST's Cybersecurity Framework 2.0 (February 2024) elevated governance to a first-class function, and the NIST NCCoE's zero-trust practice guide drafts (2023–2024) pressed implementers toward measurable outcomes — yet the instruments most organizations use to report "zero trust maturity" remain largely self-assessed, declarative, and insensitive to the difference between a policy that exists and a control that is enforced. **Objective:** We design ZT-MMI, a Zero Trust Maturity Measurement Instrument that scores maturity on an evidence-tier basis — distinguishing declared policies (E0), configured controls (E1), observed behavior (E2), and enforced outcomes (E3) — across the five pillars and three cross-cutting capabilities of CISA's Zero Trust Maturity Model v2.0, anchored to the tenets of NIST SP 800-207. **Methods:** Following a design-science methodology, we (i) analyzed the authoritative zero-trust corpus (NIST SP 800-207, CISA ZTMM v2.0, OMB M-22-09, DoD Zero Trust Strategy, NIST CSF 2.0, NIST SP 1800-35 drafts), (ii) extracted requirements that any credible maturity instrument must satisfy, (iii) constructed the instrument's indicator set, evidence tiers, and aggregation model, and (iv) evaluated it through requirements mapping, cross-model alignment, and an illustrative synthetic assessment. **Results:** Existing instruments are complementary but incomplete: CISA's ZTMM defines stages but not evidence requirements; CSF 2.0 defines functions and tiers but not zero-trust-specific outcomes; DoD's strategy defines deadlines but not measurement semantics for civilian adopters. ZT-MMI composes their semantics into a single instrument in which every maturity claim carries an evidence tier; the synthetic walkthrough shows how evidence tiering surfaces the characteristic failure mode of declarative maturity reporting — high policy coverage with low enforcement. **Conclusions:** Zero trust maturity is only meaningful as a measured property of *enforced behavior*. Evidence-tiered scoring converts maturity from a label into an auditable claim, providing the digital-trust foundation that zero-trust transitions require as they move from aspiration to accountability.

Keywords: zero trust architecture; maturity model; evidence-based assessment; NIST SP 800-207; CISA Zero Trust Maturity Model; cybersecurity framework 2.0; continuous compliance; digital trust

Contribution statement. This paper makes four contributions: (1) a systematic requirements analysis of the authoritative zero-trust corpus through the 2024 federal milestone, identifying what a maturity instrument must measure and why declarative self-assessment fails; (2) the ZT-MMI instrument — pillar-by-capability indicator structure, evidence-tier scoring (E0–E3), and aggregation semantics — composable from CISA ZTMM v2.0, NIST CSF 2.0, and NIST SP 800-207 semantics; (3) a cross-model alignment showing that ZT-MMI maps onto, rather than competes with,

existing instruments; and (4) an illustrative synthetic assessment demonstrating how evidence tiering exposes policy-versus-enforcement gaps that declarative maturity reporting conceals.

1. Introduction

Zero trust — the principle that no implicit trust is granted based on network location or asset ownership, and that every access request is authenticated, authorized, and continuously evaluated [1] — has completed its journey from heretical proposal [9] to doctrinal requirement. NIST SP 800-207 gave it an architectural formalization in 2020 [1]; U.S. federal policy made it mandatory in practice [2][4]; and by 2024 "zero trust" had become the label under which a substantial share of enterprise security investment is justified. That success, however, created a measurement crisis: as zero trust became a compliance expectation, organizations began reporting "zero trust maturity" with instruments that could not distinguish between a security program that has *written down* zero-trust principles and one that *enforces* them. The 2024 policy environment made this crisis concrete. OMB M-22-09 set specific zero-trust goals for federal agencies with a September 2024 horizon [2]; agencies were required to demonstrate — not merely assert — progress. GAO's High-Risk Series continues to document the gap between federal cybersecurity ambition and execution [8]. And NIST Cybersecurity Framework 2.0, released in February 2024, added the Govern function and implementation tiers precisely to make such demonstrations structured and comparable [6].

The core measurement problem is evidentiary. A maturity model stages capabilities (CISA's Zero Trust Maturity Model v2.0 progresses from Traditional through Initial, Advanced, and Optimal across Identity, Devices, Networks, Applications & Workloads, and Data, with Visibility & Analytics, Automation & Orchestration, and Governance as cross-cutting capabilities [3]). But a *stage* is not evidence. An organization can honestly report "we have multi-factor authentication" at the same time that its MFA is disabled for service accounts, exempted by policy exceptions, or bypassed through legacy protocols — each a difference invisible to declarative reporting but decisive for the security posture zero trust is meant to create. The zero-trust literature has matured architecturally [1][9][10] and operationally [7], and maturity models now exist in numbers [3][4][6]; what has not kept pace is an instrument whose scores are *evidence-tiered* — tied to what can be shown, not what is believed.

1.1 Research questions

RQ1: What do the authoritative zero-trust instruments (NIST SP 800-207, CISA ZTMM v2.0, OMB M-22-09, DoD Zero Trust Strategy, NIST CSF 2.0, NIST SP 1800-35) collectively require of a maturity *measurement* — as distinct from a maturity *label*?

RQ2: What indicator structure can express those requirements across the recognized pillars and cross-cutting capabilities without competing with existing models? **RQ3:** How should evidence be tiered and aggregated so that maturity scores distinguish declared, configured, observed, and enforced controls? **RQ4:** What does a defensible assessment process look like for organizations that must answer to auditors, regulators, or oversight bodies?

1.2 Contributions

We contribute: (1) a requirements analysis of the 2024 zero-trust corpus from a measurement perspective (Section 2, Table 1) (RQ1); (2) the **ZT-MMI instrument** — indicator structure, E0–E3 evidence tiers, and weighted aggregation (Section 3, Figure 2, Table 2) (RQ2, RQ3); (3) a cross-model alignment demonstrating composability with CISA ZTMM v2.0, CSF 2.0 tiers, and M-22-09 goals (Section 3.4, Table 4) (RQ2); and (4) an illustrative synthetic assessment with a documented threat-to-validity analysis (Section 4, Figures 3–4, Table 3) (RQ4). We are explicit throughout that the assessment demonstration is synthetic and illustrative; empirical validation across organizations is the required next step, not a claimed result.

2. Background and Related Work

2.1 The zero-trust corpus through 2024

Figure 1 places the corpus on a timeline. Forrester's original formulation (2010) [9] and the practitioner formalization in *Zero Trust Networks* (2017) [10], elaborated in enterprise guidance [14], established the doctrine; NIST SP 800-207 (2020) defined the architecture — policy decision points, policy enforcement points, trust algorithms, and the rejection of implicit trust by network location [1]. NSA's 2021 guidance translated the architecture for defense-in-depth adopters [5]. Policy then outran standards: OMB M-22-09 (January 2022) mandated specific federal zero-trust goals on a fixed clock [2]; the Department of Defense published its Zero Trust Strategy (October 2022) with a 2027 target horizon and a reference architecture [4][11]; CISA's Zero Trust Maturity Model matured from v1.0 (2021) to v2.0 (April 2023) [3]; and NIST NCCoE's SP 1800-35 drafts (2023–2024) documented ten interrelated "example implementations" showing what ZTA concretely looks like in operation [7]. NIST CSF 2.0 (February 2024) generalized the governance lesson across all

cybersecurity programs [6]. By September 2024 — the M-22-09 milestone — the question “are we zero trust?” had become “can we show how zero trust we are?” [2][8].




Figure 1. Zero trust standardization and policy timeline, 2010–2024. The highlighted milestone marks the September 2024 federal zero-trust goal deadline of OMB M-22-09 [2].

2.2 Maturity models and their measurement gap

Maturity models in information systems research are conventionally evaluated against desiderata including stage distinctness, applicability, and, critically, *evidence* for placement [12]. The zero-trust family satisfies the structural desiderata: CISA ZTMM v2.0 defines four stages over five pillars and three cross-cutting capabilities [3]; DoD’s strategy defines seven pillars with target outcomes and milestones [4]; CSF 2.0 defines six functions with implementation tiers and organizational profiles [6]. What the family largely lacks is measurement semantics: none of the three instruments requires that a claimed capability be supported by a *tiered evidence class*, and all three invite the self-assessment inflation documented in adjacent compliance domains. SP 1800-35 supplies rich implementation exemplars but is explicitly not a measurement instrument [7]. The gap is therefore not another model — it is an *instrument* that consumes the models.

2.3 Assessment, attestation, and digital trust

The difference between asserted and demonstrated control operation is a recurring theme in security accounting: NIST SP 800-53’s assessment paradigm distinguishes implemented from effective controls [13]; audit standards distinguish design effectiveness from operating effectiveness. Evidence tiering generalizes these distinctions for zero trust: a maturity claim is only as trustworthy as the strongest evidence tier supporting it. This framing connects zero-trust measurement directly to digital trust: an organization’s claims about its own security are a trust relationship with its oversight community, and instruments that make claims *verifiable* are the mechanism by which that relationship scales.

3. Methodology

The study applies a design-science methodology in four components.

Requirements analysis (RQ1). We analyzed the six authoritative instruments — NIST SP 800-207 [1], CISA ZTMM v2.0 [3], OMB M-22-09 [2], the DoD Zero Trust Strategy and reference architecture [4][11], NIST CSF 2.0 [6], and NIST SP 1800-35 preliminary drafts [7] — extracting for each: structure (pillars, functions, stages, tiers), measurement semantics (what counts as evidence of advancement), and assessment process. Requirements for a measurement instrument were synthesized from the union: (R1) coverage of the recognized pillars and cross-cutting capabilities; (R2) stage semantics compatible with CISA’s four stages; (R3) explicit evidence classes distinguishing assertion from enforcement; (R4) aggregation that resists compensation of weak pillars by strong ones beyond a bounded allowance; (R5) auditability of every scoring decision; (R6) alignment with M-22-09 goal categories for federal adopters.

Instrument construction (RQ2, RQ3). ZT-MMI was constructed by composing the extracted semantics: the pillar-capability grid from CISA ZTMM v2.0 [3]; tenet-anchored indicator wording from SP 800-207 [1]; evidence tiers E0–E3 generalizing the

assessment paradigm of SP 800-53 [13]; and aggregation semantics designed under constraint R4. Each design decision is reported with alternatives in Section 4.2.

Cross-model alignment (RQ2). ZT-MMI stages, indicators, and evidence tiers were mapped to CISA ZTMM stages, CSF 2.0 implementation tiers, and M-22-09 goal categories (Table 4) to demonstrate composability rather than competition with existing instruments.

Illustrative assessment (RQ4). A synthetic federal-grant-making agency archetype ("Agency A") was constructed to demonstrate the instrument end-to-end (Figures 3–4, Table 3). The archetype is synthetic, informed by publicly documented patterns of federal cybersecurity posture but not by any actual organization's data; its role is to exercise the instrument and make its assumptions inspectable, and it supports no claims about real-world distributions of maturity.

4. The ZT-MMI Instrument (RQ2, RQ3)

4.1 Structure

ZT-MMI scores the **five pillars and three cross-cutting capabilities of CISA ZTMM v2.0 as eight assessed capability areas** (Figure 2) [3]: pillars (Identity, Devices, Networks, Applications & Workloads, Data) and cross-cutting capabilities (Visibility & Analytics, Automation & Orchestration, Governance), with the cross-cutting capabilities layered across all pillars rather than multiplied against them — yielding $5 + 3 = 8$ capability areas, each scored on the same evidence tiers. Each capability area contains 4–7 indicators (Table 2 lists representative indicators; the full register is a governance artifact maintained per adopting organization).

Capability area	Representative indicators	Example evidence by tier
Identity	Phishing-resistant MFA for all users including service accounts; per-session authorization	E1: IdP configuration; E2: authentication telemetry; E3: bypass attempt denied
Devices	Device posture evaluated before each grant; compliance enforcement	E2: posture telemetry; E3: non-compliant device denied (logged event)
Networks	Micro-segmentation; per-request policy decision with logged decision inputs	E2: flow logs vs. policy matrix; E3: lateral movement blocked in test
Applications & Workloads	Per-application access policies; workload identity without implicit service trust	E2: service-to-service authentication telemetry
Data	Automated classification; policy-driven encryption and loss prevention	E1: DLP configuration; E2: DLP enforcement events
Visibility & Analytics	Full-fidelity session context available to the PDP for each decision	E2: proportion of access decisions with logged context
Automation & Orchestration	Automated policy distribution and response execution	E2: playbook executions on trigger events
Governance	ZT policy adopted with exception workflow and re-assessment cadence	E0: policies; E1: exception registry with expiry

Table 2. Representative ZT-MMI indicators and evidence tiers per capability area.




Figure 2. The ZT-MMI instrument: five pillars plus three cross-cutting capabilities (eight assessed capability areas) with the four-stage scoring band and E0-E3 evidence tiers.

4.2 Evidence tiers and scoring

Each indicator is scored 0–3 on the **evidence tier scale**: **E0 declared** (a policy or plan exists; scored 0); **E1 configured** (the control is deployed and configuration can be shown; 1); **E2 observed** (operational telemetry demonstrates the control operating as intended over the assessment window; 2); **E3 enforced** (independence is structural — the control cannot be bypassed by the actors it governs, and its enforcement is verified, e.g., by red-team attempt or protocol-level guarantee; 3). A capability area's score is the weighted mean of its indicator scores, with weights set by the adopting organization's risk profile (defaults provided). The **maturity index** aggregates capability scores with a bounded-compensation rule: pillar scores may not exceed the mean of their cross-cutting capability scores by more than one stage band, reflecting SP 800-207's insistence that zero trust is a *system* property, not a collection of point products [1].

The design decisions embedded here answer the measurement gap directly. **Tier E0 anchors the floor at zero**, not at "documented": writing policy earns nothing by itself, which is the deliberate inversion of declarative maturity reporting. **Tier E2 requires telemetry**, making the score depend on the visibility capability — the cross-cutting capability that CISA treats as the enabler of all others [3] — so organizations cannot report pillar maturity above what their monitoring can demonstrate. **Tier E3 rewards structural enforcement**, aligning the highest score with the property that matters in practice: an access decision that *cannot* be bypassed rather than one that *has not been* bypassed. This generalizes the implemented-versus-effective distinction of conventional control assessment [13] into a four-tier ladder usable per indicator.

4.3 Stage mapping

Capability-area scores map to the four CISA stages: Traditional (<25), Initial (25-50), Advanced (50-75), Optimal (>75), preserving continuity for organizations already reporting against ZTMM v2.0 [3]. Table 4 gives the full crosswalk, including the mapping to CSF 2.0 implementation tiers (ZT-MMI's E-tiers operationalize what "Tier" means for zero-trust indicators specifically) and to M-22-09 goal categories, so federal adopters can answer the 2024 milestone question directly [2][6].

4.4 Assessment process

Figure 3 shows the five-step assessment workflow: (1) scope and governance mapping; (2) evidence collection across four evidence classes (policy artifacts, configuration state, operational telemetry, adversarial verification); (3) evidence-tiered scoring per indicator, with every scoring decision recording its evidence references; (4) maturity index computation and gap analysis; (5) roadmap generation and annual re-assessment, with the evidence file retained for

oversight and audit. The evidence file is designed to integrate with the tamper-evident evidence-preserving patterns developed in the series' compliance-monitoring context, giving assessments the same integrity properties as the systems they measure.



Figure 3. The five-step ZT-MMI assessment workflow with annual re-assessment and evidence retention.

Instrument	Year	Structure	Measurement semantics	Feeds ZT-MMI requirement
NIST SP 800-207 [1]	2020	Tenets; PDP/PEP components; deployment variants	None (architecture, not assessment)	Tenet-anchored indicator wording (R1, R5)
CISA ZTMM v2.0 [3]	2023	5 pillars + 3 cross-cutting capabilities; 4 stages	Stage descriptions per capability; no evidence classes	Grid structure and stage semantics (R1, R2)
OMB M-22-09 [2]	2022	Federal ZT goals in 5 categories; Sept 2024 milestone	Milestone-based progress reporting	Federal indicator categories (R6)
DoD ZT Strategy & RA [4][11]	2022	7 pillars; target outcomes; 2027 horizon	Outcome levels per capability	Outcome-level indicator wording (R1)
NIST CSF 2.0 [6]	2024	6 functions (incl. Govern); profiles; implementation tiers	Tier definitions (Partial/Risk Informed/Repeatable/Adaptive)	Tier and governance semantics (R3, R4)
NIST SP 1800-35 drafts [7]	2023-2024	Ten example ZTA implementations	Worked exemplars, not scores	Observability evidence patterns (R3, R5)

Table 1. The authoritative zero-trust instruments and the measurement requirements they feed.

5. Evaluation (RQ4)

5.1 Requirements and alignment

Table 1 maps each source instrument to the measurement requirements it feeds; Table 4 (Section 5.3) demonstrates that ZT-MMI composes them without contradiction. The construction satisfies R1-R6 by design; the evaluation below exercises the instrument's *behavior*, which is where design flaws surface.

5.2 Illustrative assessment: synthetic archetype

Figure 4 presents the illustrative assessment of "Agency A," a synthetic mid-sized federal grant-making archetype constructed for demonstration. Its profile was seeded from publicly recurring patterns in zero-trust transitions: strong identity investments (enterprise MFA, centralized identity provider — Advanced, 70), maturing network segmentation (Advanced, 65), partial device management (Initial, 35), application-layer work concentrated on modern services while legacy applications await re-platforming (Initial, 40), data-layer capabilities still Traditional in classification and loss-prevention automation (25), visibility limited to network telemetry without user-and-entity behavior coverage (45), automation largely manual (20), and governance comparatively strong because it is documentation-intensive (75).

The diagnostic value appears in the *evidence-tier composition*, summarized in Table 3: Agency A's governance score of 75 is dominated by E0/E1 evidence (policies, plans, charters) — the instrument prices this correctly at stage Advanced on paper but flags the *evidence mix* as declarative-heavy, in contrast to the Identity score of 70, which is supported primarily by E2 telemetry. The maturity index (52.4, Advanced-band) is therefore reported with an **evidence-quality indicator** — the proportion of the index supported by E2/E3 evidence (41%) — which becomes the oversight conversation's central number: an organization can be

"Advanced" with mostly assertions, and the instrument now says so. This is precisely the failure mode that declarative maturity reporting conceals and that the 2024 accountability environment [2][6] requires surfacing.




Figure 4. Illustrative ZT-MMI assessment of the synthetic federal-grant-making archetype "Agency A" across the eight capability dimensions. Synthetic and illustrative only; not empirical data.

Capability area	Score	Stage	Dominant evidence tier	Gap note (illustrative)
Identity	70	Advanced	E2 (telemetry-supported)	Service-account MFA exemptions at E1
Networks	65	Advanced	E2	Segmentation gaps in legacy zones (E1)
Governance	75	Advanced	E0/E1 (declarative-heavy)	Evidence mix flagged for oversight
Visibility & Analytics	45	Initial	E1/E2 partial	No user/entity behavior analytics
Automation & Orchestration	20	Traditional	E0	Playbooks largely manual
Data	25	Initial (boundary)	E1	Classification still manual
Composite index	52.4	Advanced	41% of index supported by E2/E3	Weakest-link minimum: 20 (Automation)

Table 3. Illustrative assessment results for the synthetic archetype (selected capability areas).

5.3 Cross-model alignment

Table 4 gives the crosswalk: ZT-MMI stages inherit CISA ZTMM v2.0's stages [3]; its evidence tiers operationalize CSF 2.0 implementation tiers for zero-trust indicators [6]; and its federal-indicator subset is organized by M-22-09 goal categories (identity, devices, networks, applications & workloads, data; plus governance cross-cutting) so that a federal assessment directly answers the September 2024 milestone reporting [2]. The alignment demonstrates composability: an organization already reporting to CISA ZTMM or CSF 2.0 loses no prior investment; it gains evidence semantics.

ZT-MMI element	CISA ZTMM v2.0 [3]	NIST CSF 2.0 [6]	OMB M-22-09 [2]
Stage bands (Traditional/Initial/Advanced/Optimal)	Same four stages per capability	Approximates Partial / Risk Informed / Repeatable / Adaptive tiers	Progress levels per goal category
Evidence tiers E0-E3	Implied by stage descriptions but not required	Operationalizes "Tier" for zero-trust indicators	Supplies milestone evidence semantics
8-pillar-capability grid	Identical structure	Functions cross-mapped	5 goal categories + governance cross-cutting
Bounded compensation + weakest-link reporting	Not specified	Not specified	Not specified

Table 4. Crosswalk: ZT-MMI against CISA ZTMM v2.0, NIST CSF 2.0, and OMB M-22-09.

5.4 Failure modes of the instrument itself

Honest evaluation includes the instrument's weaknesses. **Telemetry gaming:** because E2 requires telemetry, organizations may manufacture favorable telemetry; the mitigations are evidence retention with integrity protection (hash-chained assessment records) and E3's adversarial verification for the highest-stakes indicators. **Indicator rigidity:** a fixed indicator register ages as architectures evolve; the register is therefore versioned, and weights are organization-specific by design. **Aggregation gaming:** the bounded-compensation rule (Section 4.2) can still mask a single catastrophically weak pillar behind uniformly mediocre ones; the instrument therefore also reports the *minimum* capability-area score alongside the index — the weakest-link statistic that zero trust's own threat model implies [1]. **Assessment cost:** full evidence collection is expensive; the workflow is designed for tiered-depth assessment (spot E1/E2 checks on high-weight indicators) with full-depth assessment annualized.

6. Discussion

6.1 From maturity labels to maturity claims

The instrument's central move is epistemological: it converts "we are Advanced in Identity" from a self-description into a *claim with evidence semantics* — claimable only to the extent that its supporting evidence reaches E2/E3 tiers. This reframing matters for each of the instrument's audiences. For agency leadership, it makes investment sequencing legible (policies first, telemetry second, enforcement third — each visible as a tier transition). For oversight bodies, it makes maturity reports comparable and contestable. For auditors, it supplies the evidence file that the 2024 accountability environment increasingly demands [2][6][8]. And for the trust relationship between an organization and those who depend on its security, it replaces a badge with a receipt.

6.2 Zero trust as a digital-trust foundation

Zero trust's promise — eliminate implicit trust — is ultimately a promise about *distrust made routine*: every access checked, every session re-evaluated, every decision logged [1]. The measurement argument of this paper is the mirror image: an organization's own zero-trust maturity must also be subject to continuous, evidence-based evaluation, because maturity claims are themselves access-control objects in the socio-technical sense — they open doors (budget, authority, contract eligibility) and should be guarded accordingly. An organization that enforces zero trust internally but asserts it externally without evidence has solved half of the problem the paradigm poses.

6.3 Adoption path

ZT-MMI is designed for staged adoption mirroring the series' compliance-monitoring pattern: (1) evidence-first — stand up the evidence file and integrity protection, score with E0/E1 honestly; (2) telemetry-first — connect the visibility capability so E2 scoring becomes possible; (3) verification-first — introduce E3 adversarial verification for the highest-weight indicators; (4) governance

throughout, with annual re-assessment (Figure 3). The staged path allows small organizations to begin without a SOC, while large adopters can integrate scoring into existing continuous-monitoring pipelines.

7. Limitations and Threats to Validity

Construct validity. The requirement set (R1-R6) derives from six instruments read through a measurement lens; alternative readings (e.g., emphasizing DoD's seven-pillar structure) would shift indicator emphasis, though not the evidence-tier logic. **Internal validity.** The synthetic assessment was constructed by the same authors who designed the instrument and therefore demonstrates intended behavior, not discovered flaws; the failure-mode analysis (Section 5.4) catalogs anticipated flaws analytically. **External validity.** No empirical assessments of real organizations were performed; the distribution of maturity in real populations, the inter-rater reliability of evidence-tier scoring, and the instrument's sensitivity to actual control changes are all open empirical questions and are stated as the required validation agenda. **Reliability.** Indicator wording is anchored to cited sources, but weights and thresholds are defaults, not validated optima. **Temporal scope.** The corpus is frozen at the 2024 milestone; M-22-09's successor guidance, revisions to CISA ZTMM, and CSF 2.0 implementation communities will require instrument re-versioning.

8. Conclusion

Zero trust asks organizations to distrust their networks; the 2024 accountability environment asks organizations to distrust their own maturity reports. This paper contributed ZT-MMI, an evidence-tiered zero-trust maturity measurement instrument that composes the semantics of NIST SP 800-207, CISA's ZTMM v2.0, NIST CSF 2.0, and the federal zero-trust goals of OMB M-22-09 into a single auditable assessment: capabilities scored on declared, configured, observed, and enforced evidence; aggregated with weakest-link protections; and reported with both a maturity index and an evidence-quality indicator that separates demonstrated from asserted maturity. The instrument's synthetic demonstration shows how evidence tiering exposes the policy-versus-enforcement gap that declarative reporting conceals, and its limitations define the empirical agenda — multi-organization validation of scoring reliability — that precedes any claim of field readiness. In a decade whose security projects are judged by whether they can *show* their zero trust, the measurement instrument is not an accessory to the architecture; it is the architecture's instrument of digital trust.

Declarations

Funding. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

CRedit contribution statement. Sujit Biswas: conceptualization, methodology, instrument design, writing – original draft. Md Monjurul Karim: conceptualization, validation, supervision, writing – review & editing (corresponding author). Anindya Nag: investigation, requirements analysis, writing – review & editing.

Competing interests. The authors declare no competing interests.

Data and code availability. No real-world organizational data was used. The synthetic archetype specification, indicator examples (Table 2), and crosswalk (Table 4) are fully specified in the paper; the indicator register template is available from the corresponding author on reasonable request.

Ethics and dual-use. This paper does not include exploit material or live credentials. The illustrative assessment uses exclusively synthetic data and identifies no real organization.

AI assistance declaration. The authors used AI-based tools to assist with literature discovery, source verification, figure preparation, and manuscript drafting; all content was reviewed, verified against primary sources, and approved by the authors, who take full responsibility for the work.

References

- [1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," NIST SP 800-207, Aug. 2020. doi: 10.6028/NIST.SP.800-207.
- [2] Office of Management and Budget, "Moving the U.S. Government Toward Zero Trust Cybersecurity Principles," Memorandum M-22-09 for the Heads of Executive Departments and Agencies, Jan. 26, 2022. [Online]. Available: <https://www.whitehouse.gov/omb/information-for-agencies/memoranda/>
- [3] Cybersecurity and Infrastructure Security Agency, "Zero Trust Maturity Model, Version 2.0," Apr. 2023. [Online]. Available: <https://www.cisa.gov/zero-trust-maturity-model>
- [4] U.S. Department of Defense, "Department of Defense Zero Trust Strategy," Version 1.0, Oct. 21, 2022 (publicly released Nov. 2022). [Online]. Available: <https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-ZTStrategy.pdf>
- [5] National Security Agency, "Embracing a Zero Trust Security Model," Cybersecurity Information Sheet, Feb. 2021. [Online]. Available: <https://www.nsa.gov/Press-Room/News-Highlights/Article/Article/2456362/>
- [6] National Institute of Standards and Technology, "The NIST Cybersecurity Framework (CSF) 2.0," NIST CSWP 29, Feb. 2024. doi: 10.6028/NIST.CSWP.29.

- [7] National Cybersecurity Center of Excellence, "Implementing a Zero Trust Architecture," NIST SP 1800-35 (Preliminary Drafts), 2023-2024. [Online]. Available: <https://www.nccoe.nist.gov/publication/1800-35/>
- [8] U.S. Government Accountability Office, "High-Risk Series: Efforts Made to Achieve Progress Need to Be Maintained and Expanded to Fully Address All Areas," GAO-23-106203, Apr. 2023. [Online]. Available: <https://www.gao.gov/products/gao-23-106203>
- [9] J. Kindervag, "No More Chewy Centers: Introducing the Zero Trust Model of Information Security," Forrester Research, Mar. 2010.
- [10] E. Gilman and D. Barth, Zero Trust Networks: Building Secure Systems in Untrusted Networks. Sebastopol, CA: O'Reilly Media, 2017.
- [11] Defense Information Systems Agency and National Security Agency, "Department of Defense Zero Trust Reference Architecture, Version 2.0," Jul. 2022.
- [12] J. Becker, B. Niehaves, J. Poeppelbuss, and A. Simons, "Maturity models in IS research," in Proc. 18th European Conference on Information Systems (ECIS 2010), Pretoria, South Africa, 2010.
- [13] National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," NIST SP 800-53 Rev. 5, Sep. 2020. doi: 10.6028/NIST.SP.800-53r5.
- [14] J. Garbis and J. W. Chapman, Zero Trust Security: An Enterprise Guide. Berkeley, CA: Apress, Feb. 2021. doi: 10.1007/978-1-4842-6702-8.