International Journal of Safety and Security Engineering
Vol. 16, No. 9, September, 2026, pp. 2035-2047
Journal homepage: http://iieta.org/journals/ijsse

Risk Exposure and Zero Trust Maturity Assessment Using ISO/IEC 27005, CISA ZTMM, and PPTG: A Case Study of a Hybrid-Infrastructure Bank in Indonesia

Lumban Sopian*, Nilo Legowo

Master of Information Systems Management Program, Bina Nusantara University, Jakarta 11480, Indonesia

Corresponding Author Email: lumban.sopian@binus.ac.id

Copyright: ©2026 The authors. This article is published by IIETA and is licensed under the CC BY 4.0 license (http://creativecommons.org/licenses/by/4.0/).

https://doi.org/10.18280/ijsse.160909

Received: 28 July 2026
Revised: 14 September 2026
Accepted: 20 September 2026
Available online: 30 September 2026

Keywords:
banking cybersecurity, CISA Zero Trust Maturity Model, ISO/IEC 27005, risk assessment, Zero Trust Architecture

ABSTRACT

Banks are adopting hybrid infrastructure that exposes identities, devices, applications, and data beyond the network perimeter. Although Zero Trust Architecture (ZTA) offers a 'never trust, always verify' approach, risk and maturity are often assessed separately. This study (1) evaluates the risk profile of an Indonesian hybrid-infrastructure bank using ISO/IEC 27005:2022, (2) assesses ZTA maturity using CISA Zero Trust Maturity Model (ZTMM) v2.0 through a People-Process-Technology-Governance (PPTG) analytical lens against NIST SP 800-207, and (3) examines where risk and maturity converge or diverge at the analytical-dimension and item levels. A qualitative single-case study triangulated seven data sources, including a 15-item maturity questionnaire completed by 45 respondents across three organizational levels; risk scenarios were scored using Likelihood × Consequence. Three scenarios were Critical and one was High. The overall analytical maturity mean was 2.77 of 4, with Automation & Orchestration the lowest dimension (2.42) and Identity the highest (2.90). Automation & Orchestration also contributes to Process, the lowest PPTG pillar (2.68), although only modestly below Technology (2.77) and Governance (2.83). Of the seven CISA elements implicated in the Bank's Critical or High risk scenarios, only Automation & Orchestration and Devices fall below the Bank's overall analytical dimension band; the remaining elements retain relatively high aggregate scores, while item-level analysis reveals specific weaknesses that aggregate scores can mask. The findings suggest that ZTA readiness constraints are not uniformly distributed and that joint risk–maturity analysis is most informative when interpreted at both aggregate and item levels in an Indonesian banking context.

1. INTRODUCTION

Cybersecurity risk in Indonesia's financial sector has become systemic rather than episodic. In 2024, Indonesia's National Cyber and Crypto Agency (BSSN) reported more than 330 million anomalous cyber-traffic events, 1,367 cyber-incident notifications, 241 suspected data leaks, and more than 56 million darknet-exposed records, with 47 notifications concerning the financial sector [1] (Figure 1). The 2023 LockBit 3.0 attack on Bank Syariah Indonesia, which disrupted mobile-banking and ATM services and was reportedly followed by the leak of customer and employee data, illustrates the exposure Indonesian banks face; the incident did not involve the anonymized bank studied here. Cyberattacks in digital financial services create financial, reputational, and regulatory consequences in addition to direct technical loss [2].

This exposure is compounded by an infrastructural shift already under way. Banks are migrating, gradually and unevenly, from on-premise systems toward hybrid models that combine core banking platforms with public and private cloud services. The migration expands the attack surface: identities, devices, applications, and data now routinely operate beyond the perimeter that conventional security controls were built to defend. A perimeter-based model, which trusts entities by default once inside the network, offers little resistance to credential compromise, lateral movement, or cloud-workload exploitation.

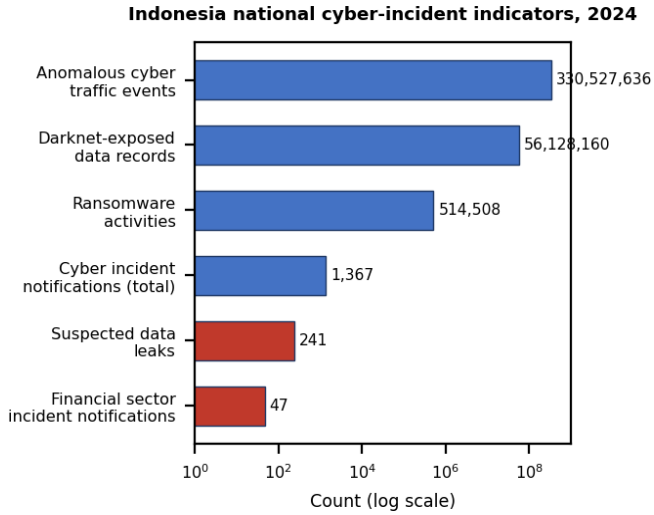


Figure 1. Indonesia national cyber-incident indicators, 2024

Zero Trust Architecture (ZTA) responds by replacing implicit trust with continuous, per-request verification: no access decision rests on network location alone. NIST SP 800-207 formalizes the target-state architecture around three components, the Policy Engine, Policy Administrator, and Policy Enforcement Point, while CISA's Zero Trust Maturity Model (ZTMM) v2.0 operationalizes readiness across five pillars, Identity, Devices, Networks, Applications & Workloads, and Data, and three cross-cutting capabilities: Visibility & Analytics, Automation & Orchestration, and Governance [3].

Despite their widespread use, ISO/IEC 27005:2022 and CISA ZTMM v2.0 are usually applied separately: risk assessment identifies urgent threat scenarios, while maturity assessment indicates the distance from a target state. Few studies test whether the two identify the same organizational weaknesses, a distinction that matters in Indonesia, where POJK No. 11/POJK.03/2022, PADG No. 24 of 2024, and Law No. 27 of 2022 impose obligations on specific risk and control domains rather than on a composite maturity score.

This study therefore integrates ISO/IEC 27005:2022 risk-assessment results with CISA ZTMM v2.0 maturity measurements, analyzed through a People-Process-Technology-Governance (PPTG) lens and benchmarked against NIST SP 800-207 and current Indonesian financial regulations. The inquiry centers on Bank XYZ, an anonymized national commercial bank operating a hybrid cloud infrastructure and holding an active ISO/IEC 27001:2022 certification.

The research proceeds through three analytical phases: mapping the Bank's threat profile via ISO/IEC 27005:2022, measuring Zero Trust readiness via CISA ZTMM v2.0 across its five core pillars and three cross-cutting capabilities, and reading the two results jointly through the PPTG analytical lens against the NIST SP 800-207 target state and prevailing national regulations. The principal contribution is case-based evidence showing how joint risk–maturity analysis can identify specific points where elevated risk and weak maturity coincide, while also revealing cases in which high-risk domains retain relatively high aggregate maturity and the underlying weakness is visible only at the item level.

2. RELATED WORK

2.1 Zero Trust Architecture: From concept to measurable maturity

NIST SP 800-207 defines Zero Trust as a set of concepts intended to minimize uncertainty in enforcing accurate, per-request access decisions regardless of the network location from which a request originates [4]. The architecture rests on three logical components, a Policy Engine that computes access decisions from policy and trust signals, a Policy Administrator that issues access tokens based on the Policy Engine's decision, and a Policy Enforcement Point that enforces the decision at every communication session [4], built on the “never trust, always verify” principle first popularized by Kindervag [5] and since established as a dominant reference point in cross-sector cybersecurity transformation.

Translating this concept into something measurable is where the literature diverges. CISA ZTMM v2.0 operationalizes Zero Trust into a maturity rubric spanning five core pillars, Identity, Devices, Networks, Applications & Workloads, and Data, and three cross-cutting capabilities, Visibility & Analytics, Automation & Orchestration, and Governance, each scored across four levels from Traditional to Optimal [3]. A comparable eight-dimension model, derived through a three-round Delphi study [6], illustrates how much deliberate work establishing content validity requires before a maturity instrument can be deployed, a concern this study addresses through the two-source triangulation. Both models are widely adopted by government and private institutions, but their application is largely descriptive: an organization reports a maturity position without reading it against an independent measure of where its actual risk sits, precisely the separation this study is designed to close.

A parallel strand of the literature documents why that separation matters. A consistent observation recurs across a targeted commentary and broader systematic reviews spanning different methods [7-9]: ZTA's promised reduction in attack surface is real, but the practical work of cross-domain implementation, technical and organizational, is consistently underestimated at the conceptual stage. That broader pattern narrows to a specific weak point, automation and orchestration capability, in a targeted analysis of ZTA implementation barriers [10], a finding this study's own lowest-scoring domain is consistent with. None of this literature, however, connects the implementation gaps it documents to a concurrent, independent risk assessment of the same organization, the specific gap that this study is positioned to close.

2.2 Risk-based approaches to information security management

ISO/IEC 27005:2022 provides an information-security risk-management methodology comprising context establishment, risk assessment (identification, analysis, and evaluation), risk treatment, and continuous monitoring [11], adopting a semi-quantitative approach that combines likelihood and impact values into an ordinal risk matrix yielding Low, Medium, High, and Critical categories [11]. Grounding this in economic terms rather than compliance checklists alone matters because the actual cost of cyber incidents frequently exceeds initial estimates when mitigation prioritization is not backed by adequate quantitative risk assessment [12].

Within Indonesian information-security research specifically, risk assessment and maturity assessment have so far been applied as parallel rather than integrated exercises. An ISO/IEC 27001-based IT security risk assessment in a banking context treats the risk register and the security control documentation as separate artifacts rather than reading them against a maturity model [13], a gap that partly motivates this study's joint reading of risk and maturity results. This line of Indonesian risk-assessment work extends in two further directions: an Indonesian manufacturing risk register anchored in three documented incidents rather than survey perception alone [14], and an SPBE-guided risk assessment of information-system audit tooling in the Indonesian public sector [15]. Together, these studies indicate that formal, standard-based risk assessment is being operationalized across Indonesian institutions well beyond banking, but none of the three reads its risk findings against an independent maturity measure.

2.3 Organizational lens and regulatory context in Indonesian banking

The PPTG lens is used in this study as an analytical lens developed by the authors to organize Zero Trust capabilities according to the organizational dimension primarily responsible for their implementation: People, Process, Technology, or Governance. It is not treated as a replacement for the CISA ZTMM structure. Rather, it provides an organizational interpretation of the CISA dimensions, allowing a technical maturity gap to be attributed to the dimension responsible for implementing or governing the relevant capability. Context- and risk-aware access control requires organizational alignment across these dimensions, not merely technical implementation at the network layer [16], a conclusion this study's Process-pillar findings reinforce empirically.

That organizational alignment plays out differently depending on regulatory context, which existing ZTA-in-banking literature engages with unevenly. A blockchain-integration framework for Zero Trust in the financial industry [17] and a report on infrastructure-readiness challenges in Nigeria's banking sector similar to those found here [18] both leave the compliance regime unspecified, which limits how directly either result transfers to institutions under binding regulatory obligations. Indonesia supplies exactly that regulatory backdrop: POJK No. 11/POJK.03/2022 requires banks to meet adequate IT risk-management principles [19], PADG No. 24 of 2024 further regulates information-system security [20], and Personal Data Protection Law No. 27 of 2022 mandates breach notification within 3 × 24 hours [21]. Because these obligations attach to specific control and risk domains rather than to an aggregate maturity figure, the joint reading attempted here speaks directly to regulatory compliance, not only to ZTA maturity in the abstract.

2.4 Self-assessment validity and this study's positioning

Self-reported maturity data carries a documented validity risk. Critical success factors identified for Zero Trust adoption come with a specific caution: organizations relying on self-assessment without independent technical validation are prone to overestimating their own readiness [6]. This caution underlies the methodological decision to triangulate seven instruments in this study, combining self-assessment questionnaires with independent vulnerability assessment and participatory observation, precisely to mitigate that overestimation risk.

Read together, the literature surveyed in this section shows that ISO/IEC 27005:2022, CISA ZTMM v2.0, NIST SP 800-207, and the PPTG analytical lens each supply a necessary element for a comprehensive picture of ZTA readiness, yet studies applying them to hybrid banking typically report risk findings and maturity findings side by side without examining whether the two describe the same underlying weaknesses. This study addresses that gap by treating the ISO/IEC 27005:2022 risk-assessment output and the CISA ZTMM v2.0 maturity results as complementary views of the same organization, organized through the PPTG lens and interpreted against the NIST SP 800-207 target state. The analysis therefore does not assume that high risk must correspond to low maturity; instead, it tests where convergence occurs and where aggregate maturity scores can mask item-level weaknesses.

3. METHODOLOGY

3.1 Research design and setting

This study uses a qualitative single-case-study design [22], triangulating quantitative and qualitative data. The research was conducted at a national commercial bank with hybrid infrastructure, supervised by the Financial Services Authority and certified to ISO/IEC 27001:2022 for its Mobile Banking and Online Onboarding scope. The evaluation boundary covers components under the Bank's direct operational control and excludes core banking systems operated by third parties.

The first author served as an insider researcher, holding the Chief Information Security Officer (CISO) role at the Bank throughout the study. This position gave direct access to internal policy documents, operational data, and informants with deep working knowledge of the Bank's security practice, but it also carries an evident risk of bias. Six control measures were applied to manage this risk: informant participation was voluntary and anonymized and was not linked to performance evaluation; informants were selected by role relevance rather than reporting relationship to the researcher; findings were triangulated across the seven data sources; the coding and evidence-mapping process was documented as an audit trail; interpretations were confirmed with informants through member checking and discussed with the academic supervisor to counter confirmation, authority, and social-desirability bias; and empirical data are reported separately from researcher interpretation. These measures partially, rather than fully, offset the independence risk: the Risk Management Committee and the Cyber Risk function sit outside the CISO's direct reporting line within the Bank's organizational structure, which gives the validation of Likelihood and Consequence ratings a degree of structural independence, whereas the corroboration by the SOC and DevOps teams and business-unit leads, who are not necessarily independent of the CISO's reporting line, is better read as technical verification of specific evidence, such as vulnerability-assessment findings and phishing-simulation results, than as an independent authority check.

3.2 Research procedure

The study proceeded through a preparatory phase followed by three analytical stages. Throughout all three stages, findings are read against a common reference point: the three-component target-state architecture of NIST SP 800-207 (Policy Engine, Policy Administrator, Policy Enforcement Point), which supplies the vocabulary connecting the risk-domain findings of Stage 1 to the PPTG-pillar findings of Stage 2, rather than being introduced only at the interpretation stage.

The preparatory phase comprised a literature review and preliminary field engagement with the Bank to scope the research problem and confirm access to informants and documentation. Stage 1 applied ISO/IEC 27005:2022 to identify, analyze, and evaluate the Bank's principal threat scenarios, producing a Risk Level classification for each scenario. Stage 2 measured Zero Trust maturity across the five CISA ZTMM v2.0 core pillars and three cross-cutting capabilities using the 45-respondent questionnaire and triangulated evidence; these eight analytical dimensions were then aggregated for case-level comparison and organized into the four PPTG pillars. Stage 3 read the two sets of results

jointly, examining where Critical or High risk scenarios were associated with weaker maturity at the domain or item level, and where risk remained elevated despite relatively strong aggregate maturity, interpreted against the NIST SP 800-207 target state and applicable national regulations.

Each stage produced a documented, traceable deliverable: the risk register, the domain and pillar score set, and the joint risk-maturity reading, rather than an undocumented intermediate judgment.

3.3 Instruments and data sources

Data were triangulated from seven instruments during the collection phase: (1) semi-structured interviews with 45 respondents comprising five strategic-level, 17 tactical-level, and 23 operational-level informants, each session exploring risk perception, control practice, and incident history relevant to the informant's role; (2) a 15-item CISA ZTMM v2.0 maturity questionnaire administered to the same 45 respondents; (3) participatory observation of infrastructure configuration and operational routines; (4) analysis of six internal policy and procedure documents; (5) an independent cloud and on-premise vulnerability assessment; (6) Key Risk Indicator (KRI) and IT security dashboard data covering January-April 2026; and (7) an internal phishing-simulation report. The first five instruments informed the risk assessment and the maturity assessment; the KRI dashboard and phishing-simulation data corroborated specific risk scenarios discussed later.

Content validity of the maturity questionnaire was reviewed by two senior information-security practitioners prior to administration. The 15-item instrument showed high pooled internal consistency (Cronbach's alpha, $\alpha = 0.932$), but this pooled coefficient does not establish reliability at the individual domain level. Domain-level alpha, computed for the six dimensions represented by two or more items, is uneven: Networks ($\alpha = 0.813$) and Data ($\alpha = 0.798$) reach levels conventionally regarded as good to acceptable, Identity ($\alpha = 0.737$) reaches a level conventionally regarded as acceptable, while Governance ($\alpha = 0.615$), Devices ($\alpha = 0.584$), and Applications & Workloads ($\alpha = 0.563$) fall below the 0.70 threshold commonly used for acceptable internal consistency.

Because Devices, Networks, Applications & Workloads, Data, and Governance are each estimated from only two items, their alpha coefficients are statistically less stable than Identity's three-item estimate and should be interpreted with caution regardless of whether they individually clear the 0.70 threshold. Visibility & Analytics and Automation & Orchestration, each measured by a single item, cannot be assessed for internal consistency. Each item was mapped to the corresponding CISA ZTMM v2.0 element based on its technical content, as shown in Table 1. Item allocation is intentionally uneven because CISA ZTMM v2.0 distinguishes five core pillars from three cross-cutting capabilities rather than defining eight equivalent constructs [3]. Accordingly, this study treats the eight elements as eight analytical dimensions for descriptive comparison, not as eight psychometrically equivalent scales. Domain scores are therefore interpreted as indicative composite measures and are triangulated with technical, documentary, observational, and organizational evidence rather than treated as uniformly reliability-validated constructs.

Table 1. Mapping of questionnaire items to CISA Zero Trust Maturity Model (ZTMM) v2.0 domains and dominant People-Process-Technology-Governance (PPTG) pillar

Item	CISA ZTMM Domain	PPTG Analytical Assignment
Q1–Q3	Identity	People
Q4–Q5	Devices	Technology
Q6–Q7	Networks	Technology
Q8–Q9	Applications & Workloads	Process
Q10–Q11	Data	Technology
Q12	Visibility & Analytics	Process (0.5); Technology (0.5)
Q13	Automation & Orchestration	Process
Q14–Q15	Governance	Governance

Source: compiled by the authors based on CISA ZTMM v2.0 [3].

Table 2. Wording of the 15-item CISA Zero Trust Maturity Model (ZTMM) v2.0 maturity questionnaire (n = 45 respondents)

Item	CISA ZTMM Domain	Item Wording (English Translation)
Q1	Identity	Does access to critical systems (Core Banking/Cloud) already use phishing-resistant MFA (e.g., FIDO2/Token)?
Q2	Identity	Are administrator accounts managed using a Privileged Access Management (PAM) solution based on Just-in-Time access?
Q3	Identity	Do access decisions incorporate real-time User Behavior Analytics (UBA)?
Q4	Devices	How accurate is the real-time inventory of devices connected to the network?
Q5	Devices	Is access automatically denied when a device does not comply with policy (e.g., outdated OS, disabled antivirus)?
Q6	Networks	How is network segmentation implemented? Has micro-segmentation been adopted?
Q7	Networks	Is internal (east-west) traffic fully encrypted?
Q8	Applications & Workloads	Is access to internal applications free of static VPN, using an Identity Proxy instead?
Q9	Applications & Workloads	Is security testing (SAST/DAST) automatically integrated into the CI/CD pipeline (DevSecOps)?
Q10	Data	Is customer data automatically classified and labeled by the system?
Q11	Data	Is an active Data Loss Prevention (DLP) mechanism in place to prevent data leakage in the cloud?
Q12	Visibility & Analytics	Are logs from the entire environment (on-premise and cloud) centralized in a single SIEM dashboard?
Q13	Automation & Orchestration	How automated is the security incident-response process? Is SOAR already in use?
Q14	Governance	Is third-party / vendor access strictly limited under least-privilege principles and fully monitored?
Q15	Governance	Can the system detect and block unauthorized SaaS applications (Shadow IT) on the network?

Source: developed by the authors based on CISA ZTMM v2.0 [3], administered as a self-assessment questionnaire scored on a four-level scale (Traditional = 1, Initial = 2, Advanced = 3, Optimal = 4).

Table 2 reports the wording of all 15 items is translated from the Indonesian-language instrument administered to respondents; the English wording was prepared by the authors for reporting purposes and was not the version shown to respondents.

Q14 and Q15 are classified under Governance rather than Visibility & Analytics because both measure the presence and enforcement of an organizational policy, vendor access governed by least-privilege (Q14) and application usage governed by an acceptable-use boundary (Q15), rather than a specific infrastructure capability. Q15's wording, "detect and block," reflects policy enforcement rather than the passive telemetry centralization that Visibility & Analytics (Q12) measures.

3.4 Risk identification and scoring procedure

Risk identification and scoring follows the ISO/IEC 27005:2022 process end to end, from establishing the asset context through to a validated Risk Level, in five sequential steps.

Step 1: Identify assets, threats, and vulnerabilities. Risk identification used an asset-based approach: the Bank's hybrid infrastructure was first partitioned into six information-asset categories within the evaluation boundary of the study. For each asset category, threats and vulnerabilities were identified by triangulating the seven data sources, vulnerability-assessment findings, interviews, participatory observation, document analysis, KRI data, and the phishing-simulation report, rather than drawn from a generic threat catalog. Table 3 summarizes the six asset categories, and their associated threats are listed; this identification step produced the four threat scenarios (R-01 to R-04) carried into Steps 2-5.

R-04 (access policy misaligned with identity) cuts across A-01 and A-04, since it concerns the policy governing how external parties reach identity and network assets rather than a single asset category.

Step 2: Rate Likelihood (L). Each of the four scenarios identified in Step 1 is rated on a five-level ordinal scale, from Rare (L = 1) to Almost Certain (L = 5), using evidence recorded in the risk register. The evidence base includes vulnerability-assessment findings, incident or simulation history, observed control effectiveness, and documented threat activity where available. Table 4 defines the observable indicators are defined to distinguish the five levels; the resulting assignments are traceable to the scenario evidence rather than to theoretical estimation alone.

Table 3. Information-asset categories and associated threats (ISO/IEC 27005:2022 identification stage)

Asset	Description	Principal Threat Identified
A-01	User, privileged, and workload identity accounts	Credential compromise via phishing (R-01)
A-02	Customer data, transaction data, security telemetry	Undetected data exfiltration / insider access
A-03	API gateway, cloud services, container workloads	Cloud vulnerability exploitation (R-03)
A-04	VLANs, DirectConnect, hybrid integration points	Undetected lateral movement (R-02)
A-05	On-premise servers, endpoints, network devices	Infrastructure vulnerability exploitation
A-06	Incident-response SOPs and access policy documents	Manual, inconsistent incident response

Table 4. Likelihood and consequence rating scales (ISO/IEC 27005:2022)

Value	Likelihood Level	Likelihood Indicator	Consequence Level	Consequence Indicator
1	Rare	Effective controls in place; no history of similar incidents	Insignificant	Minor technical finding, fixed immediately with no business impact
2	Unlikely	Controls in place but not yet fully effective	Minor	Minor incident, no service disruption, minimal customer impact
3	Possible	Vulnerability exists; partial controls applied	Moderate	Limited service disruption; incident containable within normal operations
4	Likely	Threat known to be active across the sector	Major	Service disruption exceeding 24 hours; significant regulatory fine; widespread negative publicity
5	Almost Certain	Active evidence or a recent documented incident exists	Catastrophic	Massive data breach, license-revocation-level sanction, or unrecoverable financial loss

Step 3: Rate Consequence (C). Each scenario is rated on a five-level ordinal scale, from Insignificant (C = 1) to Catastrophic (C = 5), covering financial, operational, legal, and reputational impact. Because the Bank did not maintain documented numeric thresholds for every consequence level at the time of assessment, the scale is explicitly semi-quantitative; the study does not claim that the ratings represent measured probabilities or monetary loss estimates.

These indicators remain semi-quantitative rather than fully quantitative: only the Major consequence level is anchored to a hard numeric threshold (service disruption exceeding 24 hours), since the Bank did not, at the time of assessment, maintain a documented Risk Appetite Statement or KRI dashboard specifying numeric thresholds such as loss amounts or affected-customer counts for every level. The indicators nonetheless make each Likelihood and Consequence assignment traceable to a specific evidentiary condition. To avoid implying false precision, the study reports the resulting ratings as ordinal expert-assessment categories and uses the L × C product only as a consistent classification rule.

Step 4: Compute the Risk Level. A Threat Rating (TR) score is calculated from the Likelihood (L) and Consequence (C) ratings assigned in Steps 2 and 3, as shown in Eq. (1):

$$TR = L \times C \tag{1}$$

TR is then classified against a threshold: TR ≥ 16 is Critical, TR 10-15 is High, TR 5-9 is Medium, and TR ≤ 4 is Low. This quantitative threshold supplements the qualitative two-dimensional risk matrix conventionally used under ISO/IEC 27005:2022 and resolves borderline cases consistently.

Step 5: Validate. Because the first author's dual role as

CISO and researcher placed sole initial responsibility for the Likelihood and Consequence ratings, each rating was cross-checked against an independent or corroborating source appropriate to the evidence type before being finalized. The Risk Management Committee reviewed the full risk profile in a closed session, providing structural independence from the CISO reporting line, while Strategic-level informants provided additional corroboration. The SOC and DevOps teams verified the technical basis of scenarios tied to vulnerability-assessment findings, and business-unit leads confirmed the phishing-simulation evidence underlying the credential-compromise scenario. These latter groups were not treated as fully independent because they may fall within the CISO's reporting influence. No rating was retained on the researcher's judgment alone. Only scenarios rated Critical or High after validation are carried forward as priority domains for the maturity assessment.

3.5 Maturity scoring procedure

CISA ZTMM v2.0 maturity results and PPTG analytical scores are derived through three sequential steps applied to the 15-item questionnaire (n = 45). The eight CISA elements are retained as analytical dimensions for comparison; the aggregation below does not imply that CISA defines the five core pillars and three cross-cutting capabilities as eight equivalent psychometric constructs.

Step 1: Score each item and compute item statistics. Every respondent rates each of the 15 items on the four-level CISA ZTMM scale (Traditional = 1, Initial = 2, Advanced = 3, Optimal = 4). This produces 45 raw numeric ratings per item, one per respondent, denoted xi for respondent i's rating. Two descriptive statistics are computed directly from these 45 raw ratings for each item: the Item Mean, defined in Eq. (2), and the Item Median.

$$Item\ Mean = (1/n) \times \Sigma x_i, n = 45 \quad (2)$$

The two statistics serve different purposes and are not interchangeable. The Median is reported per item as the primary statistic for labeling that item's qualitative maturity level (for example, a Median of 3.0 is labeled "Advanced"), following established guidance that the Median rather than the Mean is the appropriate central-tendency statistic for assigning a category to ordinal Likert-type data [23, 24]. The Mean, by contrast, is the statistic carried forward into Step 2 to build the domain-level composite index, because averaging several items into one domain score is an arithmetic aggregation task for which the Mean is the conventional and more distribution-sensitive choice. This dual use, Median for per-item labeling, Mean for cross-item aggregation, follows standard practice for composite indices built from ordinal survey items. Table 5 reports both statistics, alongside the standard deviation, are computed directly from the n = 45 raw responses for all 15 items.

Q13 is the only item with a Median below 3.0 (Initial rather than Advanced) and also records the lowest Mean (2.42) and the second-highest standard deviation among the 15 items (SD = 1.011, after Q3's 1.031), indicating that low ratings dominate but a minority of Operational-level respondents rated it higher, a pattern examined further below. The Mean for each item is not a single respondent's answer or an assumption; it is the arithmetic average of that item's 45 individual numeric ratings. Because the 45 respondents are not a homogeneous group, they span five Strategic, 17 Tactical, and 23 Operational informants with different vantage points on the same systems, the same item mean can be produced by very different underlying response patterns.

Table 6 decomposes each domain's score by hierarchical level, using the mean of that level's respondents on the domain's mapped items, to show how the ratings that feed the item means are distributed across the organization rather than concentrated in one group.

Table 5. Item-level descriptive statistics and maturity level (n = 45)

Item	Domain	Mean	Median	SD	Level
Q1	Identity	3.04	3.0	0.82	Advanced
Q2	Identity	3.04	3.0	0.85	Advanced
Q3	Identity	2.60	3.0	1.03	Advanced
Q4	Devices	2.76	3.0	0.74	Advanced
Q5	Devices	2.60	3.0	0.92	Advanced
Q6	Networks	2.87	3.0	0.92	Advanced
Q7	Networks	2.78	3.0	0.93	Advanced
Q8	Apps & Workloads	2.96	3.0	0.93	Advanced
Q9	Apps & Workloads	2.82	3.0	0.78	Advanced
Q10	Data	2.73	3.0	0.94	Advanced
Q11	Data	2.89	3.0	0.86	Advanced
Q12	Visibility & Analytics	2.78	3.0	0.74	Advanced
Q13	Automation & Orchestration	2.42	2.0	1.01	Initial
Q14	Governance	3.04	3.0	0.74	Advanced
Q15	Governance	2.62	3.0	0.98	Advanced

Table 6. Domain scores by hierarchical respondent level (unweighted level means)

Domain	Strategic (n = 5)	Tactical (n = 17)	Operational (n = 23)
Identity	2.87	2.69	3.06
Devices	2.60	2.53	2.80
Networks	2.80	2.38	3.15
Applications & Workloads	2.70	2.76	3.02
Data	2.80	2.53	3.02
Visibility & Analytics	2.60	2.59	2.96
Automation & Orchestration	2.00	2.00	2.83
Governance	2.70	2.62	3.02

Two patterns stand out. First, Tactical-level respondents rate every domain at or below the Strategic-level rating, and often well below the Operational-level rating; this group has direct operational visibility into implementation gaps without either the Strategic level's summary reporting or the Operational level's hands-on familiarity with day-to-day workarounds, a pattern consistent with the caution against unvalidated self-assessment raised by Yeoh et al. [6]. Second, Automation & Orchestration is the only domain where Strategic and Tactical respondents agree exactly (2.00), indicating that this weakness is recognized at the leadership level rather than being visible only to operational staff. The domain scores used in the remainder of the analysis are computed from the pooled n = 45 response-level data, not from averaging the three level columns; the level breakdown here is diagnostic context, not an alternative calculation path.

These descriptive patterns were checked against inferential testing. A Kruskal-Wallis test comparing the three hierarchical levels found a statistically significant difference in only two of the eight domains: Networks (H = 8.44, p = .015) and Automation & Orchestration (H = 7.14, p = .028); the remaining six domains showed no significant level-based difference (p > .05). Post-hoc Mann-Whitney U tests on these two domains located the significant difference specifically between the Tactical and Operational levels: Networks (U = 91.5, p = .003) and Automation & Orchestration (U = 107.5, p = .013). This result narrows the first pattern above: Tactical-level respondents rating most domains lower is a consistent descriptive trend across level breakdown, but it reaches statistical significance only for Networks and Automation & Orchestration, and specifically as a Tactical-versus-Operational gap rather than a Tactical-versus-Strategic one. This check is exploratory and complements the qualitative-quantitative triangulation central to the study's single-case design; it is not an inferential hypothesis test that the design was powered to support, and it is reported here to substantiate, rather than to independently establish, the level-based patterns visible in the level breakdown.

Step 2: Aggregate items to analytical-dimension scores. Each item is pre-mapped to one of the five CISA ZTMM v2.0 core pillars or three cross-cutting capabilities. A dimension score is the average of the Item Mean values from Step 1 for its mapped items, as shown in Eq. (3), where k is the number of items mapped to that CISA element. This equal item contribution is an analytical convention for descriptive comparison; it is not presented as a CISA-defined weighting scheme.

$$Domain\ score = (1/k) \times \Sigma Item\ Mean_{\mathrm{i}} \qquad (3)$$

Because each Item Mean is itself already an average of the same 45 individual respondents, averaging Item Means within a domain is mathematically equivalent to computing one combined mean directly from all of that domain's individual ratings, in turn equivalent to a proportional weighted average across the three hierarchical strata (5 Strategic, 17 Tactical, 23 Operational respondents) rather than an unweighted average of three separately computed level means; this prevents the five Strategic responses from being over-weighted relative to the 23 Operational responses. Table 7 shows this computation using the Item Mean values from Step 1. Domain scores are computed from Item Means carried at full arithmetic precision, not from the two-decimal values displayed in Table 5; consequently, a domain score will not always reproduce exactly from multiplying or averaging Table 5's rounded display values by hand (for example, Networks averages Q6 and Q7 at full precision, 2.8667 and 2.7778, to 2.822, which rounds to 2.82, rather than averaging the rounded 2.87 and 2.78 shown in Table 5, which would round to 2.83).

Step 3: Aggregate analytical dimensions to PPTG pillars. Each CISA dimension is assigned to the PPTG organizational dimension primarily responsible for implementing or governing it. People, Process, and Governance each receive explicitly assigned dimensions; Visibility & Analytics is treated as cross-cutting and therefore contributes 0.5 to Process and 0.5 to Technology. The 0.5/0.5 allocation is a transparent non-preferential convention adopted because the study has no empirical basis in this single case for assigning a different organizational weight. The pillar score is the weighted mean of the contributing dimension scores, as shown in Eq. (4). This PPTG aggregation is an analytical device developed for this study and should not be interpreted as a formal component structure of CISA ZTMM or NIST SP 800-207.

$$Pillar\ score = (1/m) \times \Sigma Domain\ score_{\mathrm{j}} \qquad (4)$$

Table 8 shows this aggregation explicitly. The pillar scores were computed as People = Identity (2.90); Process = (Applications & Workloads + Automation & Orchestration + 0.5 × Visibility & Analytics) / 2.5 = (2.89 + 2.42 + 1.39) / 2.5 = 2.68; Technology = (Devices + Networks + Data + 0.5 × Visibility & Analytics) / 3.5 = (2.68 + 2.82 + 2.81 + 1.39) / 3.5 = 2.77; and Governance = 2.83. The overall analytical mean of 2.77 is calculated separately as the equal mean of the eight analytical-dimension scores. It is used only as a descriptive case-level benchmark against the 3.00 Advanced reference point; neither the 2.77 index nor the PPTG pillar scores are presented as CISA-defined composite metrics.

Table 7. Domain score computation from item-level response means (n = 45)

Domain	Mapped Item(s)	Item Mean (Step 1)	Domain Score
Identity	Q1, Q2, Q3	3.04, 3.04, 2.60	2.90
Devices	Q4, Q5	2.76, 2.60	2.68
Networks	Q6, Q7	2.87, 2.78	2.82
Applications & Workloads	Q8, Q9	2.96, 2.82	2.89
Data	Q10, Q11	2.73, 2.89	2.81
Visibility & Analytics	Q12	2.78	2.78
Automation & Orchestration	Q13	2.42	2.42
Governance	Q14, Q15	3.04, 2.62	2.83

Table 8. People-Process-Technology-Governance (PPTG) pillar score computation from domain scores

PPTG Pillar	Pillar Score
People	2.90
Process	2.68
Technology	2.77
Governance	2.83

Visibility & Analytics is split evenly across the Process and Technology pillars, consistent with its dual classification in the item mapping and with CISA ZTMM v2.0's characterization of Visibility & Analytics as a cross-cutting capability not uniquely anchored to one organizational pillar.

Self-reported questionnaire scores in both tables were cross-checked against the vulnerability-assessment and document-analysis evidence, consistent with recommended practice for guarding against self-assessment overestimation [21]. The resulting values are the domain and pillar scores.

4. RESULTS AND DISCUSSION

4.1 Risk profile based on ISO/IEC 27005:2022

The risk assessment identified four principal threat scenarios in the Bank's hybrid infrastructure.

Credential compromise via phishing (R-01). An internal

phishing simulation showed 138 employees susceptible to credential-harvesting attempts, including 16.7% from the IT and Operations Directorate who hold privileged access, with the susceptibility trend rising 15% year-on-year. No user-behavior analytics capability exists to flag anomalous activity once a session begins, so a captured credential can be reused undetected after login.

Undetected lateral movement (R-02). Cloud telemetry is not centralized into the Bank's SIEM, leaving a blind spot for movement between workloads, while network segmentation still relies on perimeter-style VLANs rather than identity-verified micro-segmentation, so an intrusion that enters one workload can reach others without triggering an alert.

Cloud infrastructure vulnerability exploitation (R-03). A publicly disclosed OpenSSH remote-code-execution vulnerability (CVE-2024-6387) remained unpatched on three production cloud hosts at the time of assessment, with a public exploit already available; remediation had been deferred over downtime concerns.

Access policy misaligned with identity (R-04). Eight external connection paths used by vendors and third parties rely on static IP whitelisting rather than per-session identity verification, so a whitelisted address can grant access beyond its intended scope once a VPN tunnel is established.

Table 9 reports the Likelihood (L), Consequence (C), Threat Rating (TR = L × C), and resulting Risk Level for each scenario, applying the scoring procedure in Section 3.4. The evidence supporting each assignment is described in the corresponding risk narrative immediately above.

The first three scenarios are classified Critical and the fourth High under the ISO/IEC 27005:2022 threshold, as visualized on the risk matrix in Figure 2. This risk-first reading of the environment follows the same logic as earlier ISO/IEC 27001-based risk assessment work in Indonesian banking [13], though that study evaluated the risk register independently of a maturity model. This study closes that separation by reading the risk register against the maturity results.

What ties these four scenarios together is not their individual severity but their common root: each depends on a control that is either absent, manual, or perimeter-scoped rather than identity-scoped. That pattern is a structural observation, not a coincidence of scenario selection, and it is the reason the maturity results matter to interpreting them.

Each scenario is mapped to one or more CISA ZTMM v2.0 analytical dimensions representing the control capability implicated by the scenario. Table 10 sets out this mapping and the rationale behind it.

These seven CISA ZTMM v2.0 analytical dimensions are therefore the elements against which the risk scenarios carry practical weight, and their maturity results are compared with the corresponding risk evidence.

Table 9. Risk scoring and level classification

Risk ID	Scenario (Short)	L	C	TR (L × C)	Risk Level
R-01	Credential compromise via phishing	4	4	16	Critical
R-02	Undetected lateral movement	4	4	16	Critical
R-03	Cloud vulnerability exploitation	5	4	20	Critical
R-04	Access policy misaligned with identity	3	4	12	High

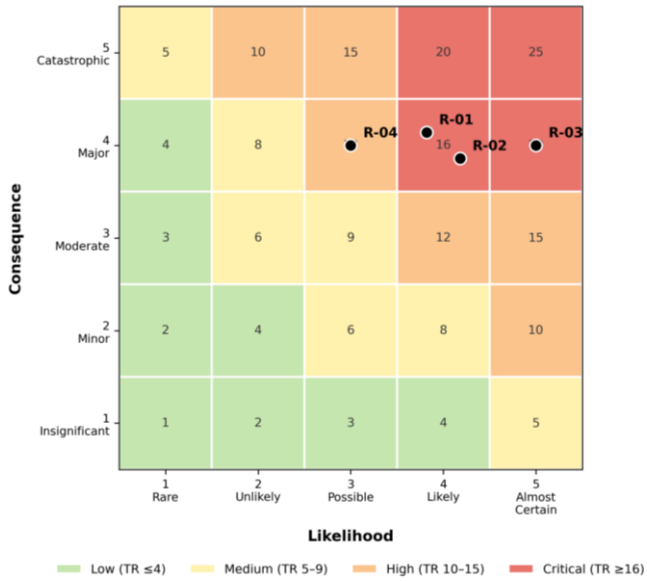


Figure 2. Mapping of risk scenarios on the ISO/IEC 27005:2022 risk matrix

Table 10. Mapping of risk scenarios to priority CISA Zero Trust Maturity Model (ZTMM) v2.0 domains

Risk ID	Risk Level	Priority CISA ZTMM Domain	Mapping Rationale (Brief)
R-01	Critical	Identity; Automation & Orchestration	Credential compromise requires adaptive identity control and automated post-authentication response.
R-02	Critical	Visibility & Analytics; Networks	Undetected lateral movement requires unified visibility and identity-based segmentation.
R-03	Critical	Applications & Workloads; Devices	Vulnerability exploitation reflects weak workload patch management and device posture.
R-04	High	Networks; Governance	Network-based access policy needs a formally mandated shift to identity-based policy.

4.2 Risk treatment, acceptance, and monitoring

Following Risk Evaluation, each of the four scenarios was assigned a treatment option from the ISO/IEC 27005:2022 set (Avoid, Reduce, Transfer, Accept), selected on the basis of feasibility within the Bank's current capacity and cost-benefit balance. Reduce was selected for all four scenarios: each risk stems from an internal control gap addressable through control strengthening rather than one requiring risk avoidance, insurance-style transfer, or unconditional acceptance at its current level. Table 11 summarizes the core control for each scenario and the projected Threat Rating and Risk Level attributable to that single core control once fully implemented. This single-control projection is conceptually distinct from the Bank's final organizational target residual level, which is defined uniformly as Low for all four scenarios; a single core control narrows exposure but does not by itself guarantee that a scenario reaches the Bank's target position. The position at

the time of assessment, before every control was complete, is reported separately below.

The projected position after a single core control differs by scenario because the proposed treatment controls do not produce the same residual Likelihood. R-01 projects to High (TR = 12), R-02 projects to Medium (TR = 8), and R-03 and R-04 project to Low (TR = 4). These projected positions describe the effect of the core control in isolation and are not themselves the Bank's treatment objective; the regulatory and standards-based sources summarized in Table 12 provide the basis for the Bank's documented risk appetite, which sets Low as the uniform final target residual level reported in Table 13, rather than a universal requirement that every single control must independently achieve that level.

For three of the four scenarios, the Bank's uniform Low target residual position could not yet be achieved at the time of assessment, and an interim position above the target was formally recorded pending completion of treatment. This interim position is recorded alongside the reason retention was necessary. The distinction between target residual risk and interim risk avoids treating a planned control outcome as if it were an already achieved control state.

Interim retention above target is permitted only where documented with a commitment and timeline; acceptance of the residual risk was formally validated by the Risk Management Committee as risk owner, consistent with its validation role in the risk-rating process.

Monitoring and review are operationalized through six mechanisms summarized in Table 14, three already in place and three recommended to sustain the cycle going forward.

Semi-annual and quarterly cadences are adequate for tracking structural maturity trends but are not responsive to fast-moving exploitation risk such as R-03, whose publication-to-exploitation window is measured in days rather than months. This is a specific, checkable claim rather than a general caution: a vulnerability disclosed and weaponized within days, as CVE-2024-6387 was, will not be caught by a semi-annual review cycle regardless of how mature the surrounding governance appears on paper. It is this mismatch, not a general critique of periodic review, that underscores why real-time detection through SIEM-SOAR integration is a priority rather than a completed control.

Table 11. Risk treatment and projected residual risk

Risk ID	Treatment	Core Control	L	TR	Projected Level
R-01	Reduce	Phishing-resistant FIDO2 MFA	3	12	High
R-02	Reduce	Centralize cloud telemetry into SIEM	2	8	Medium
R-03	Reduce	Critical-vulnerability patch SLA ≤ 30 days	1	4	Low
R-04	Reduce	Identity-based, per-session access control	1	4	Low

Table 12. Risk acceptance basis and residual risk position

Basis	Requirement Applied
ISO/IEC 27005:2022 §6.4.2	Acceptance criteria set against risk appetite defined by top management
POJK No. 11/POJK.03/2022	Board-approved Risk Appetite Statement; no open High/Critical IT risk without a formal action plan
NIST SP 800-207 §6.1	Acknowledges an acceptable, non-zero residual risk aligned to business mission
PADG No. 24 of 2024	Organization-specific risk-tolerance thresholds and SLAs for critical cyber findings

Table 13. Interim risk position during treatment implementation

Risk ID	Target Level	Interim Level	Interim Justification
R-01	Low	High	Post-authentication behavioral detection not yet mature
R-02	Low	Medium	Residual zero-day exposure window
R-03	Low	High	Vulnerability publicly exploitable and unpatched at time of assessment; 30-day SLA in progress, not yet completed
R-04	Low	Low	No interim retention required

Table 14. Monitoring and review mechanisms

Mechanism	Cadence	Corresponds to
KRI-2: dark-web data-leakage indicator	Monthly (existing)	R-01
KRI-3: Core Banking downtime	Monthly (existing)	R-02, R-03
KRI-4: disruption frequency	Monthly (existing)	R-04
Vulnerability-assessment findings into risk register	Semi-annual (recommended)	R-03
CISA ZTMM score as maturity KRI vs. 3.00 target	Annual (recommended)	All domains
MTTD/MTTR vs. 72-hour UU PDP notification deadline	Quarterly (recommended)	R-01, R-02

4.3 Maturity measurement using CISA Zero Trust Maturity Model v2.0

Maturity measurement across the eight CISA ZTMM v2.0 analytical dimensions produced an overall analytical mean of 2.77 of 4 (Traditional = 1, Initial = 2, Advanced = 3, Optimal = 4), below the Advanced-level milestone of 3.00 used as a reference point in this study. This 2.77 value is a case-level descriptive index calculated by equally averaging the eight analytical dimension scores; it is not a composite score defined by CISA ZTMM v2.0. Identity recorded the highest score (2.90), followed by Applications & Workloads (2.89), Governance (2.83), Networks (2.82), Data (2.81), Visibility & Analytics (2.78), Devices (2.68), and Automation & Orchestration, which recorded the lowest score by a clear margin (2.42). Seven of the eight dimensions cluster within a comparatively narrow band (2.68 to 2.90), while Automation & Orchestration stands out as a distinct outlier, 0.26 points

below the next-lowest dimension. The dimension-score distribution is shown in Figure 3.

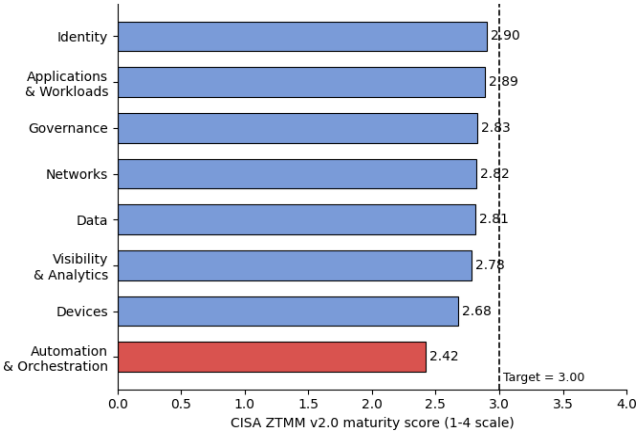


Figure 3. CISA Zero Trust Maturity Model (ZTMM) v2.0 maturity scores across the eight analytical dimensions

Analytical-dimension scores were aggregated into the four PPTG pillars using the organizational assignment and weighting convention of the PPTG lens.

The same four pillar scores are plotted in Figure 4 against the 3.00 target milestone; the NIST SP 800-207 relationships shown in Table 15 are interpretive rather than one-to-one structural correspondences.

Among the eight analytical dimensions, Automation & Orchestration is the clear outlier: at 2.42, it sits 0.26 points below the next-lowest dimension (Devices, 2.68) and 0.48 points below the overall analytical mean, whereas the remaining seven dimensions cluster within a comparatively narrow 0.22-point band (2.68 to 2.90). At the pillar level, this concentrates in Process (2.68), which remains the lowest of the four PPTG pillars but by a modest margin, 0.09 points below Technology (2.77), rather than the wide separation a dimension-only view might suggest. The pillar-level gap is therefore better read as corroborating evidence for a domain-specific weakness than as evidence of a broad Process-versus-Technology divide.

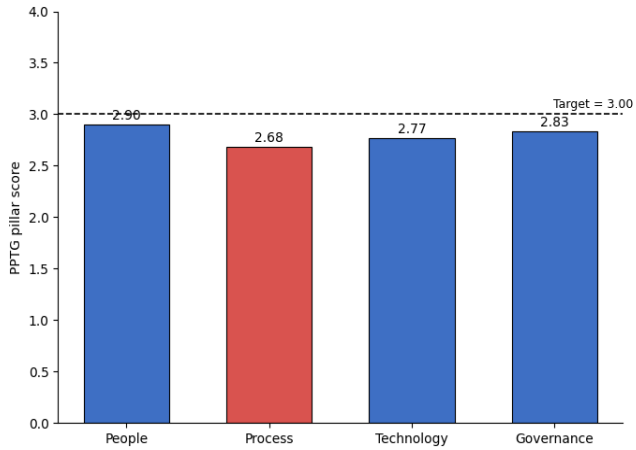


Figure 4. People-Process-Technology-Governance (PPTG) pillar scores against the 3.00 Advanced-level target

Table 15. Aggregation of CISA Zero Trust Maturity Model (ZTMM) analytical dimensions to the People-Process-Technology-Governance (PPTG) pillars and their interpretive relationship to NIST SP 800-207

PPTG Pillar	Mapped Domain(s)	Pillar Score	Related NIST SP 800-207 Relevance (Interpretive)
People	Identity (2.90)	2.90	Provides identity/subject and contextual inputs considered by Policy Engine decisions
Process	Apps & Workloads (2.89), Automation & Orch. (2.42), Visibility & Analytics (0.5 weight)	2.68	Supports Policy Administrator functions through policy translation, orchestration, and response processes
Technology	Devices (2.68), Networks (2.82), Data (2.81), Visibility & Analytics (0.5 weight)	2.77	Supports Policy Enforcement Point functions through device/network/workload enforcement controls
Governance	Governance (2.83)	2.83	Shapes policy context, risk criteria, accountability, and regulatory constraints considered in policy decisions

Read against NIST SP 800-207, the practical picture is narrower than a pillar-level reading alone would indicate. The Policy Administrator function, responsible for translating policy decisions into enforcement actions, remains manual and reactive because incident response still depends on human intervention rather than SOAR-driven automation (Automation & Orchestration: 2.42). The Policy Engine's supporting telemetry is also incomplete, since cloud activity is not yet centralized into a SIEM (Visibility & Analytics: 2.78), though this domain sits within the Bank's normal range rather than at a comparably weak level. The Policy Enforcement Point, by contrast, is comparatively closer to target: Networks (2.82) and Devices (2.68) both sit within or near the Advanced band. The composite mean of 2.77 therefore masks a specific, addressable gap in automated response rather than reflecting a broad shortfall across the Bank's ZTA implementation.

This finding sits in an interesting relationship to Cao et al. [10], who identify automation and orchestration as a persistent weak point across ZTA implementations generally, not specific to any one sector or region. The consistency between that broader finding and this single-case result is worth noting carefully rather than treated as confirmation: one case cannot validate a cross-sector pattern, but it does suggest the Automation & Orchestration gap identified here is unlikely to be an artifact specific to this Bank's circumstances, since independent literature identifies the same weak point through different methods. The practical implication is that automation deficits may be structurally difficult to close through policy alone; Cao et al.'s [10] finding and this study's finding both point toward implementation-level barriers, such as tooling integration and SOAR adoption cost, rather than governance awareness, since Automation & Orchestration is, notably, the one domain where this study's Strategic and Tactical respondents agree exactly, meaning the weakness is already recognized at the leadership level without yet being resolved operationally.

Bertino [7] separately cautions that ZTA's practical implementation challenges are often underestimated at the

conceptual level even when the promised reduction in attack surface is well understood. The Automation & Orchestration finding here is a concrete instance of that caution: understanding that Zero Trust requires continuous, automated policy enforcement (the concept) is evidently not the same as having built the SOAR and telemetry-centralization infrastructure that enforcement requires (the implementation), and the gap between the two is exactly what the 0.26-point separation between Automation & Orchestration (2.42) and the next-lowest domain, Devices (2.68), makes visible.

Both single-item domains benefit from corroboration that does not depend on their questionnaire score alone. Automation & Orchestration's 2.42 is reinforced by the cross-level agreement noted earlier, where Strategic and Tactical respondents converge exactly at 2.00, and by the independent literature pattern Cao et al. [10] report; Visibility & Analytics' 2.78 is reinforced by the Risk Management Committee's review of the full risk profile underlying R-02, including its telemetry gap, and by the participatory observation of infrastructure configuration among the study's seven data sources, rather than by the Q12 questionnaire item alone. Neither domain's score, in other words, rests on its single questionnaire item in isolation. The item-level reliability limitation nonetheless remains a valid methodological constraint.

4.4 Reading risk and maturity together

The scenario-to-domain mapping links the Bank's four risk scenarios to seven CISA ZTMM v2.0 domains: Identity and Automation & Orchestration (R-01), Visibility & Analytics and Networks (R-02), Applications & Workloads and Devices (R-03), and Networks and Governance (R-04). Reading these against the maturity results shows a pattern that is more specific than a uniform risk-maturity convergence. Only two of the seven domains, Automation & Orchestration (2.42) and Devices (2.68), score below the Bank's overall domain band; the remaining five, Identity (2.90), Applications & Workloads (2.89), Governance (2.83), Networks (2.82), and Visibility & Analytics (2.78), all sit at or above the composite mean despite being implicated in Critical or High risk scenarios. A flat convergence claim, that high-risk domains are uniformly low-maturity domains, is therefore not supported by the data as a whole.

The convergence that does hold is narrower and, on inspection, more informative. Automation & Orchestration anchors R-01's post-authentication response gap specifically, and its domain score is also the Bank's single lowest, making it the one point where elevated risk and weak maturity clearly coincide. R-02's lateral-movement detection gap maps instead to Visibility & Analytics and Networks; both domains sit within the Bank's comparatively strong band (2.78 and 2.82 respectively), consistent with the pattern discussed later in this section. A second, item-level pattern reinforces the Automation & Orchestration finding in Identity: although the domain's aggregate score (2.90) sits among the Bank's highest, driven by strong ratings on phishing-resistant MFA (Q1: 3.04) and just-in-time privileged access (Q2: 3.04), the item most directly tied to R-01's threat mechanism, user-behavior analytics for detecting anomalous post-login activity (Q3), scores 2.60, appreciably below the domain average. The Identity domain's high aggregate score therefore does not indicate the specific capability that R-01 depends on; it reflects strength in access-granting controls alongside a comparative weakness in post-access monitoring that the domain-level figure alone does not surface.

For the remaining domains implicated in Critical or High risk, Applications & Workloads, Devices, Networks, Governance, and Visibility & Analytics, maturity scores are comparatively strong, indicating that these risks originate less from a general capability gap and more from specific unresolved conditions: an unpatched vulnerability awaiting a 30-day SLA (R-03), static IP whitelisting not yet replaced by per-session identity verification (R-04), and the absence of centralized cloud telemetry despite an otherwise Advanced-rated visibility posture (R-02). Aggregate domain and pillar scores, in other words, can mask specific unresolved conditions that a risk assessment surfaces directly; the practical value of triangulating both assessments lies specifically in surfacing this kind of masked gap, rather than in confirming a uniform pattern across every domain.

These findings are also significant from a regulatory perspective. POJK No. 11/POJK.03/2022 and PADG No. 24 of 2024 evaluate compliance across specific control domains rather than through a single maturity indicator, which is consistent with what the domain-level analysis above shows: an aggregate ZTMM score of 2.77, or even a pillar-level reading, would not on its own have flagged the R-03 patch-SLA gap or the R-01 behavioral-detection gap, both of which are visible only once the risk register and the item-level maturity data are read together.

The findings extend previous concerns regarding maturity self-assessment raised by Yeoh et al. [6]. While earlier studies emphasized the risk of overestimating readiness through self-reported maturity at the aggregate level, this study shows that the same overestimation risk can operate at the domain level: a domain's aggregate score can obscure a weak constituent item that a risk assessment, read independently, is able to surface. This highlights the value of triangulating risk and maturity assessments down to the item level, not only the domain or pillar level, when evaluating Zero Trust readiness in hybrid banking environments.

5. CONCLUSIONS

This study integrated ISO/IEC 27005:2022 risk assessment with CISA ZTMM v2.0 maturity assessment to evaluate Zero Trust readiness in a national commercial bank operating a hybrid infrastructure. The assessment identified three Critical and one High risk scenarios, while the overall analytical maturity mean reached 2.77 out of 4. Automation & Orchestration recorded the Bank's lowest analytical dimension score (2.42), contributing to Process being the lowest of the four PPTG pillars (2.68), although only modestly below Technology (2.77) and Governance (2.83). Reading risk and maturity together shows that convergence is selective rather than uniform: Automation & Orchestration clearly combines elevated risk relevance with the weakest maturity, while several other risk-implicated dimensions retain relatively high aggregate scores. Item-level analysis nevertheless exposes specific weaknesses, such as Q3 user-behavior analytics, that aggregate dimension scores can mask. The principal implication is therefore methodological: risk and maturity assessments are most informative when read jointly at both aggregate and item levels.

This study shows that risk assessment and maturity assessment should be interpreted together because each

provides different insights into an organization's security posture. A maturity score alone may overlook Critical exposure, whereas a risk assessment identifies urgent issues without explaining the procedural factors behind them. Integrating both assessments through the People, Process, Technology, and Governance perspective provides a more effective basis for prioritizing Zero Trust implementation. Although this assessment approach can be applied in other organizations, the results reported in this study are specific to the investigated bank.

This study is limited by its single-case design, the use of partially self-reported maturity data, and the first author's dual role as CISO and researcher. A further instrument limitation concerns reliability: domain-level Cronbach's alpha, computed for the six dimensions with two or more items, ranges from 0.563 (Applications & Workloads) to 0.813 (Networks), with three dimensions, Devices, Applications & Workloads, and Governance, falling below the 0.70 threshold conventionally used for acceptable internal consistency; Visibility & Analytics and Automation & Orchestration, each measured by a single item, cannot be assessed for internal consistency. The pooled α=.932 should therefore not be interpreted as evidence that all eight analytical dimensions are equally reliable. Dimension and pillar scores are interpreted as indicative composite measures rather than uniformly reliability-validated scales, with corroborating evidence used to strengthen interpretation of the single-item dimensions. Validation drawing on the Risk Management Committee and the Cyber Risk function, both outside the CISO's direct reporting line, offers a degree of structural independence, but corroboration from the SOC, DevOps, and business-unit informants, who are not necessarily independent of that reporting line, cannot be treated as equivalent to an external check. The PPTG mapping and 0.5/0.5 allocation for Visibility & Analytics are analytical conventions developed for this case rather than formal CISA or NIST structures. Future studies should engage an assessor entirely external to the organization, expand and psychometrically refine the questionnaire to strengthen item-level reliability across all eight analytical dimensions, validate the integrated approach across multiple banking institutions, and develop empirically justified weighting or quantitative prioritization models for Zero Trust implementation planning.

ETHICS STATEMENT

Bank management authorized the study and approved publication of anonymized findings. Participation in interviews and the perception survey was voluntary. No personally identifiable customer data were accessed or analyzed; only de-identified organizational and respondent-level results are reported. The study followed the research-ethics expectations of the Master of Information Systems Management Program at Bina Nusantara University.

DATA AVAILABILITY STATEMENT

The item-, domain-, and pillar-level results are available from the corresponding author on reasonable request. Interview transcripts, internal incident and audit documents, and individual questionnaire responses remain confidential under the Bank's data-sharing terms and are not publicly deposited.

CONFLICTS OF INTEREST

The corresponding author served as CISO of the participating bank during the study. This role enabled access to the relevant systems, personnel, and records and also created a potential conflict of interest. The author's initial risk ratings were independently validated by the Risk Management Committee. No other conflicts of interest are declared.

ACKNOWLEDGMENT

The authors thank the Bank's management for approving the research and the informants and survey participants for their contribution. The first author also thanks Prof. Nilo Legowo for supervision and feedback and the Master of Information Systems Management Program at Bina Nusantara University for institutional support.

REFERENCES

[1] Badan Siber dan Sandi Negara (BSSN). (2024). Lanskap Keamanan Siber Indonesia 2024. Republic of Indonesia.

[2] Adekoya, O.A., Atlam, H.F., Lallie, H.S. (2025). Quantifying the multidimensional impact of cyber attacks in digital financial services: A systematic literature review. Sensors, 25(14): 4345. https://doi.org/10.3390/s25144345

[3] Cybersecurity and Infrastructure Security Agency. (2023). Zero trust maturity model. Department of Homeland Security.

[4] Rose, S., Borchert, O., Mitchell, S., Connelly, S. (2020). Zero trust architecture. NIST Special Publication, 800(207): 1-52. https://doi.org/10.6028/NIST.SP.800-207

[5] Kindervag, J. (2016). No more chewy centers: Introducing the zero trust model of information security. Forrester Research. https://www.forrester.com/report/No-More-Chewy-Centers-The-Zero-Trust-Model-Of-Information-Security/RES56682.

[6] Yeoh, W., Liu, M., Shore, M., Jiang, F. (2023). Zero trust cybersecurity: Critical success factors and a maturity assessment framework. Computers & Security, 133: 103412. https://doi.org/10.1016/j.cose.2023.103412

[7] Bertino, E. (2021). Zero trust architecture: does it help? IEEE Security & Privacy, 19(5): 95-96. https://doi.org/10.1109/MSEC.2021.3091195

[8] 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

[9] 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

[10] Cao, Y., Pokhrel, S.R., Zhu, Y., Doss, R., Li, G. (2024). Automation and orchestration of zero trust architecture: Potential solutions and challenges. Machine Intelligence Research, 21(2): 294-317. https://doi.org/10.1007/s11633-023-1456-2

[11] International Organization for Standardization. (2022).

ISO/IEC 27005:2022: Information security, cybersecurity and privacy protection: Guidance on managing information security risks (4th ed.). ISO, Geneva.
[12] Eling, M., Wirfs, J. (2019). What are the actual costs of cyber risk events? European Journal of Operational Research, 272(3): 1109-1119. https://doi.org/10.1016/j.ejor.2018.07.021
[13] Legowo, N., Juhartoyo, Y. (2022). Risk management; risk assessment of information technology security system at bank using ISO 27001. Journal of System and Management Sciences, 12(3): 181-199. https://doi.org/10.33168/JSMS.2022.0310
[14] Lai, R., Legowo, N. (2026). Incident-driven information system risk management using ISO 31000:2018, STRIDE, and CIA triad: A case study of an Indonesian manufacturer. International Journal of Safety and Security Engineering, 16(6): 1287-1296. https://doi.org/10.18280/ijsse.160611
[15] Prafitia, H.A., Legowo, N. (2025). Risk assessment of information system audit tools based on SPBE risk management guidelines. International Journal of Safety and Security Engineering, 15(7): 1471-1480. https://doi.org/10.18280/ijsse.150714
[16] Xiao, S., Ye, Y., Kanwal, N., Newe, T., Lee, B. (2022). SoK: Context and risk aware access control for zero trust systems. Security and Communication Networks, 2022(1): 7026779. https://doi.org/10.1155/2022/7026779
[17] Daah, C., Qureshi, A., Awan, I., Konur, S. (2024). Enhancing zero trust models in the financial industry through blockchain integration: A proposed framework. Electronics, 13(5): 865. https://doi.org/10.3390/electronics13050865
[18] Ashogbon, A., Ukpere, O. (2025). Evaluating the application of zero trust architecture (ZTA) implementation in Nigeria's banking industry. Journal of Electrical and Electronic Engineering, 13(3): 131-142. https://doi.org/10.11648/j.jeee.20251303.12
[19] Otoritas Jasa Keuangan. (2022). POJK No. 11/POJK.03/2022 on the implementation of information technology by commercial banks [in Indonesian]. https://ojk.go.id/id/regulasi/Pages/Penyelenggaraan-Teknologi-Informasi-Oleh-Bank-Umum.aspx.
[20] Bank Indonesia. (2024). Peraturan Anggota Dewan Gubernur (PADG) No. 24 Tahun 2024 tentang Keamanan Sistem Informasi dan Ketahanan Siber bagi Penyelenggara Sistem Pembayaran, Pelaku Pasar Uang dan Pasar Valuta Asing, serta Pihak Lain yang Diatur dan Diawasi Bank Indonesia. https://www.bi.go.id/id/publikasi/peraturan/Pages/PADG_242024.aspx.
[21] Republic of Indonesia. (2022). Law No. 27 of 2022 on personal data protection [in Indonesian]. https://peraturan.bpk.go.id/Details/229798/uu-no-27-tahun-2022.
[22] Yin, R.K. (2018). Case Study Research and Applications. Sage Publication.
[23] Jamieson, S. (2004). Likert scales: How to (ab) use them? Medical Education, 38(12): 1217-1218. https://doi.org/10.1111/j.1365-2929.2004.02012.x
[24] Sullivan, G.M., Artino Jr, A.R. (2013). Analyzing and interpreting data from Likert-type scales. Journal of Graduate Medical Education, 5(4): 541-542. https://doi.org/10.4300/JGME-5-4-18