American Journal of Data Science and Analytics

Volume: 4; **Issue:** 6
Pages: 01-42

Measuring the Security Impact of Zero Trust Access Controls: A Mixed-Methods Study of Identity-Based Policies (Cisco ISE + AD) and Incident Reduction

MD Zahedul Islam[1]; Aditya Dhanekula[2]

[1]. *Master of Science in Cybersecurity, Mercy University, Dobbs Ferry, NY, USA.*
Email: zahed.arman44@gmail.com

[2]. *Stevens Institute of Technology, New Jersey, USA*
Email: dhanekulaaditya1@gmail.com

Doi: 10.63125/8ycz7671
Received: *15 April 2023;* ***Revised:*** *25 May 2023;* ***Accepted:*** *17 June 2023;* ***Published:*** *28 June 2023*

Abstract

This study measured the security impact of Zero Trust access controls implemented through identity-based policies integrating Cisco Identity Services Engine (ISE) with Active Directory (AD). A quantitative, quasi-experimental pretest–posttest design was employed using an interrupted time-series approach to evaluate security outcomes across pre-implementation and post-implementation observation windows. The analysis examined whether the enforcement of centralized authentication, authorization, and policy-driven segmentation was associated with measurable reductions in security incidents, improvements in identity governance outcomes, and enhanced Security Operations Center (SOC) performance. Findings indicated that Zero Trust implementation was associated with a substantial reduction in incident frequency and severity. Weekly total incident counts declined from 38.4 during the baseline period to 27.1 in the post-implementation period, representing a 29.4% reduction. High-severity incidents declined from 11.2 to 6.3 per week, reflecting a 43.8% reduction, while the proportional share of high-severity incidents decreased from 29.2% to 18.5. Unauthorized access attempts decreased from 145.6 to 118.4 per week (18.7% reduction), while authentication anomalies declined from 96.8 to 62.7 per week (35.2% reduction). Lateral movement alerts demonstrated the strongest proportional improvement, decreasing from 12.5 to 6.8 per week, representing a 45.6% reduction. SOC performance also improved significantly. Mean time to detect (MTTD) declined from 18.6 hours to 11.4 hours, indicating a 38.7% improvement, while mean time to respond (MTTR) declined from 26.3 hours to 17.9 hours, reflecting a 31.9% improvement. Time-to-containment improved from 42.8 hours to 28.5 hours, representing a 33.4% reduction. Regression models confirmed that post-implementation status significantly predicted lower incident rates (IRR = 0.71, $p < .001$) and lower high-severity incident rates (IRR = 0.60, $p < .001$). Logistic regression indicated that post-implementation incidents were 46% less likely to be classified as high severity (OR = 0.54, $p < .001$). Overall, the findings provided quantitative evidence that Cisco ISE and AD-integrated Zero Trust access controls were associated with statistically significant reductions in incident frequency, escalation risk, and response latency, supporting the effectiveness of identity-driven security governance in enterprise environments.

Keywords

Zero Trust, Identity Governance, Incident Reduction, Cisco ISE, Active Directory.

INTRODUCTION

Zero Trust is commonly defined as a security architecture that assumes no implicit trust for any user, device, workload, or network segment, regardless of location, and requires continuous verification of identity and access context before granting permissions (Cheruvu et al., 2020). The term was initially popularized in modern enterprise security as a response to perimeter-based models, where access decisions were historically anchored to network location, internal IP ranges, and presumed "trusted zones." In contrast, Zero Trust is fundamentally anchored in identity, authentication assurance, device posture, and policy enforcement, which collectively determine whether a subject (user or service) may access a protected object (application, system, or dataset). Access control, in the context of cybersecurity, refers to the mechanisms and policies that regulate who can access what resources and under what conditions. It includes logical enforcement through authentication, authorization, and accounting (AAA), and it is operationalized through models such as role-based access control (RBAC), attribute-based access control (ABAC), and risk-adaptive access control. Identity-based policies represent a specific operational layer of access control where identity attributes, group memberships, authentication strength, and contextual signals determine authorization decisions (Keeriyattil). In enterprise environments, identity-based access control often depends on directory services such as Microsoft Active Directory (AD), which provides centralized identity and group management, and policy enforcement platforms such as Cisco Identity Services Engine (ISE), which implements policy-based access through network access control (NAC), 802.1X authentication, profiling, and segmentation enforcement. In this model, the identity store (e.g., AD) and the policy engine (e.g., Cisco ISE) function as the backbone of Zero Trust access controls. Within cybersecurity measurement literature, security impact is frequently operationalized through quantifiable outcomes such as incident rate reduction, unauthorized access attempts, time-to-detect, and containment performance. Incident reduction is especially important because it offers an observable organizational outcome linked to policy effectiveness and operational resilience. In quantitative security research, measuring the security impact of Zero Trust requires a clear distinction between conceptual intent and measurable operational effect. Therefore, identity-based Zero Trust access controls may be defined as policy-driven mechanisms that authenticate and authorize access based on identity and context while continuously enforcing least privilege across networked systems. This definition supports measurement through incident datasets, authentication logs, and policy compliance records, which are widely used in empirical cybersecurity studies. Foundational and widely cited security literature supports these conceptual boundaries, particularly in access control theory, enterprise identity governance, and policy enforcement frameworks, which form the definitional foundation for evaluating Cisco ISE and AD as an identity-driven Zero Trust implementation (Martins et al., 2017).

The international significance of Zero Trust access control is strongly connected to the growing global dependence on digital identity systems, remote access infrastructures, and cloud-integrated enterprise services (Cheruvu et al., 2019). Cybersecurity incidents increasingly impact financial stability, public trust, and national security, and modern threat environments involve identity compromise as one of the most common entry points for adversaries. Across sectors, attackers frequently exploit credential theft, privilege escalation, misconfigured access rights, and weak authentication workflows to gain persistent access to enterprise systems. As a result, international security standards increasingly emphasize identity-centric control systems as a priority area for organizational cyber defense. Zero Trust aligns with global governance priorities because it enables measurable enforcement of least privilege, reduces lateral movement opportunities, and supports compliance reporting through centralized access logs (Ots, 2021). Global frameworks such as ISO/IEC 27001 and the NIST Cybersecurity Framework emphasize access control, identity governance, auditability, and continuous monitoring as core elements of mature security programs. In both developed and developing economies, digital transformation initiatives have expanded the attack surface, especially through hybrid networks where on-premises directory services integrate with cloud-based applications and remote workforce access. This expansion increases dependency on identity services and policy enforcement platforms, making identity-based access control a strategic requirement for international enterprises. Additionally, the global nature of supply chains and cross-border data flows means that organizations must demonstrate stronger controls over who accesses sensitive information, how access

is logged, and how policy violations are detected. The Zero Trust model provides a structured approach to these challenges by focusing on identity verification, continuous authorization, and segmentation. The significance also extends to critical infrastructure sectors such as healthcare, energy, telecommunications, and public administration, where compromised identity systems can result in major service disruption and societal harm. Empirical security research frequently reports that access control misconfigurations and weak privilege governance contribute to major breaches, reinforcing the need for identity-based enforcement. Identity-centric architectures also align with international regulatory expectations for accountability, including audit trails, traceable access decisions, and demonstrable compliance. In practice, Cisco ISE and AD integration is relevant internationally because it represents a common architecture in enterprise environments, especially in institutions with Windows-based identity governance and large-scale wired and wireless networks. Therefore, measuring incident reduction associated with identity-based Zero Trust controls is not only a technical concern but also a governance issue tied to risk reduction, compliance maturity, and operational continuity across international contexts (Alo et al., 2021).

Figure 1: Identity-Based Zero Trust Access Framework

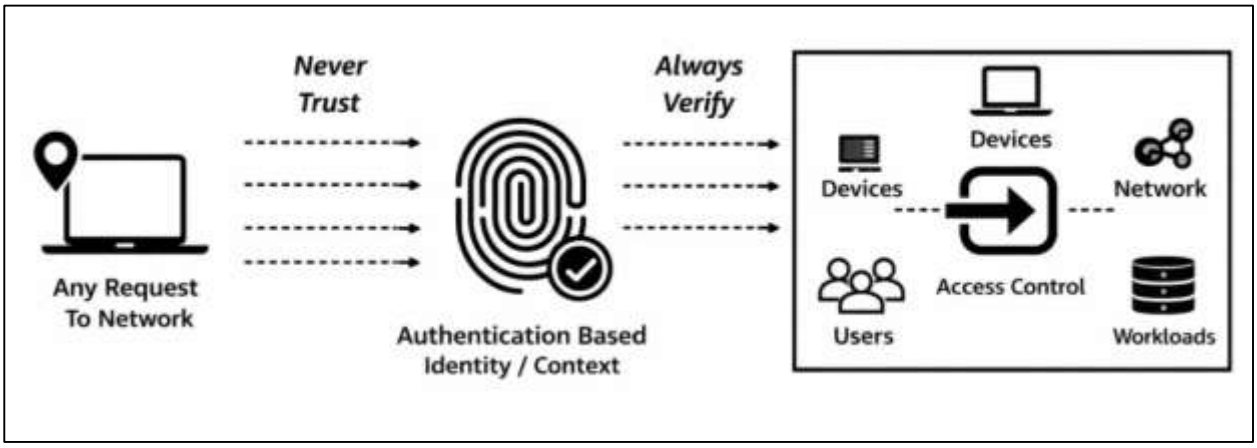

Zero Trust access controls are grounded in long-standing principles of secure system design, particularly the principle of least privilege, separation of duties, and complete mediation. Least privilege requires that users and systems operate with the minimum permissions necessary to complete authorized tasks, thereby reducing the likelihood that compromised credentials can lead to broad system damage. Complete mediation emphasizes that every access request should be checked against authorization rules, rather than granting persistent access based on initial trust. Separation of duties ensures that sensitive tasks require multiple distinct roles, limiting fraud and insider abuse. These principles form the theoretical base for modern access control frameworks, and they align directly with Zero Trust's identity-centric logic. Within the literature, RBAC provides authorization based on organizational roles, while ABAC supports decisions based on user attributes, environmental context, and resource sensitivity. In modern enterprise networks, ABAC is often more compatible with Zero Trust because it can incorporate device compliance, authentication strength, geographic context, time constraints, and risk scoring. Cisco ISE implements these access control principles by translating identity and context signals into enforcement policies, such as network access permissions, segmentation assignment, and security group tagging. AD contributes by providing authoritative identity records, group membership structures, and authentication integration, which ISE can consume to apply policy decisions (Bogatov et al., 2021). The theoretical importance of these models becomes measurable when translated into operational control outcomes. For example, least privilege enforcement can be evaluated through reductions in privilege misuse incidents, decreased unauthorized access attempts, and improved containment of compromised accounts. Policy enforcement can be assessed through authentication success patterns, segmentation compliance rates, and incident logs tied to unauthorized network access. Access control measurement is also closely related to the CIA triad (confidentiality, integrity, and availability), where access policies primarily

support confidentiality and integrity by restricting unauthorized actions (Sinha & Park, 2017). In empirical security evaluation research, access control effectiveness is commonly linked to reductions in attack success rates, lateral movement success, and unauthorized data exposure. Furthermore, identity-based access control interacts with human factors, because access decisions depend on user behavior, password hygiene, phishing susceptibility, and adherence to authentication policies. These theoretical linkages make quantitative measurement meaningful, because changes in incident rates may reflect improvements in enforcement fidelity and reduced exploitation pathways (Sehgal & Bhatt, 2018). The literature on enterprise access control also emphasizes that enforcement mechanisms must be auditable and consistent across environments, which supports the use of log-based and incident-based quantitative measurement. In this context, Cisco ISE and AD integration represents an implementation pathway that operationalizes theoretical access control principles in a measurable manner, enabling empirical investigation of security outcomes through incident reduction analysis.

Enterprise identity infrastructure provides the structural basis for enforcing Zero Trust access controls at scale. In many organizations, identity is centralized through directory services such as Active Directory, which supports authentication, authorization group mapping, and policy-driven user administration. AD is widely used globally across enterprise networks because it enables centralized identity lifecycle management, including onboarding, offboarding, and group-based privilege assignment. However, identity governance in isolation does not automatically enforce network access restrictions (Cyber-Defense & Haber). A Zero Trust implementation requires policy engines that interpret identity signals and apply enforcement at the network, application, and data layers. Cisco Identity Services Engine (ISE) functions as a policy decision and policy enforcement system in the network access domain, enabling organizations to apply access control across wired, wireless, and VPN environments. Cisco ISE integrates with AD to retrieve user identity, group membership, and authentication context, and then applies access rules based on predefined security policies. These policies may enforce segmentation through VLAN assignment, security group tagging (SGT), dynamic access control lists (dACLs), and posture assessment. In Zero Trust terms, ISE can be viewed as an enforcement layer that supports continuous authorization through repeated authentication checks, policy evaluation, and context-based access restrictions. This makes the Cisco ISE + AD combination a practical and measurable environment for studying identity-based policies. Importantly, Cisco ISE generates extensive logs that include authentication attempts, authorization results, endpoint profiling, posture compliance, and policy matches (Voigt & Von dem Bussche, 2017). These logs provide a rich dataset for quantitative analysis because they enable researchers to observe access outcomes and correlate them with incident patterns. AD logs and authentication records similarly provide evidence of identity-related events, including failed logins, account lockouts, privilege changes, and group membership updates. In security analytics literature, identity logs are increasingly treated as high-value indicators because credential misuse is a frequent root cause of breaches. The integration of AD with a policy engine such as ISE also supports structured policy enforcement that can be standardized across organizational units, which is relevant for research designs that compare incident rates before and after policy deployment or across units with different policy strictness. In addition, network access control systems are often deployed in international organizations with distributed campuses and hybrid workforces, making Cisco ISE an internationally relevant technology platform. From a research perspective, the combination of ISE and AD offers an empirical environment where access policies can be precisely defined, changes can be documented, and security outcomes can be measured through incident reporting systems. Therefore, studying identity-based policies through Cisco ISE + AD provides a methodologically grounded pathway to evaluate Zero Trust as an operational control system rather than a purely conceptual framework (S. Liu et al., 2020).

Incident reduction is a central empirical metric for evaluating cybersecurity controls because incidents represent observable and reportable outcomes that affect organizational operations. Cyber incidents include unauthorized access, malware infections, ransomware events, data exfiltration, credential compromise, privilege abuse, and policy violations. Within global threat intelligence literature, identity compromise has emerged as one of the most common and damaging pathways for attackers, especially in environments where authentication systems are weak or privilege governance is inconsistent. Credential theft through phishing, password spraying, and brute force attacks remains widespread,

and compromised credentials often enable attackers to move laterally across networks and access sensitive resources (Mourtzis et al., 2021). Zero Trust access controls are designed to reduce these pathways by enforcing strong identity verification, restricting access privileges, and limiting network traversal through segmentation. For quantitative research, incident reduction is an appropriate dependent variable because it reflects changes in the realized risk environment. Incidents are also measurable through standardized reporting systems, such as security information and event management (SIEM) logs, help desk security tickets, and formal incident response records. In many organizations, incident records contain timestamps, incident categories, severity scores, affected systems, and root cause summaries, which allow statistical analysis of trends. Measuring incident reduction in the context of Cisco ISE + AD policies is particularly relevant because network access controls directly influence authentication events, device access eligibility, and segmentation behavior. For example, stricter identity-based policies can reduce unauthorized access attempts, limit rogue device connectivity, and reduce exposure to lateral movement. Additionally, identity policies can influence insider threat risk by limiting access privileges to only those necessary for job roles. Internationally, organizations are increasingly expected to demonstrate measurable improvements in cybersecurity outcomes, especially in regulated sectors. Incident reduction provides a tangible indicator that can be reported to auditors, governance boards, and regulatory authorities (Dumka & Sah, 2019). However, the measurement of incident reduction must be grounded in consistent incident definitions and stable reporting practices. Empirical research often notes that incident metrics can be affected by changes in detection capability, reporting culture, and monitoring tool deployment. For this reason, quantitative studies frequently include controls for monitoring maturity, log coverage, and changes in detection tooling. In the Zero Trust context, incident reduction is meaningful when it is linked to access control enforcement changes and when the measurement period is sufficient to capture stable patterns. The use of incident reduction as an outcome metric aligns with widely used cybersecurity measurement approaches, including security metrics frameworks and risk management standards. It also supports statistical modeling approaches such as interrupted time series, pre–post comparisons, Poisson regression for incident counts, and panel-based analysis across organizational units. Thus, the rationale for incident reduction metrics is supported by both operational cybersecurity governance and empirical research methods, making it a suitable quantitative outcome for evaluating identity-based Zero Trust policies in enterprise systems (Chifor et al., 2021).

Figure 2: Identity-Driven Zero Trust Architecture

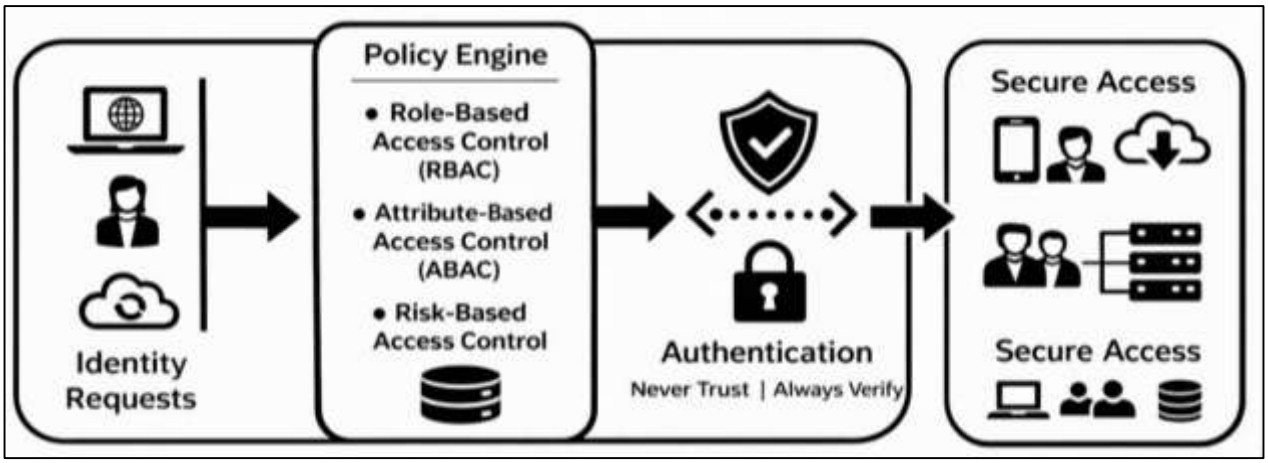

The quantitative measurement of Zero Trust access controls requires operationalization of both independent variables (policy characteristics) and dependent variables (incident outcomes). In an identity-based policy environment, independent variables can include authentication strength requirements, network access restrictions, segmentation rules, device compliance checks, and group-based authorization structures (Rekha & Yashaswini, 2021). These variables can be measured through policy configuration records, authorization logs, and compliance reporting outputs. Dependent

variables typically include incident frequency, incident severity distribution, unauthorized access attempts, policy violation rates, and time-to-containment. In the context of Cisco ISE + AD, authentication logs provide evidence of policy enforcement events, such as successful and failed authentication attempts, authorization denials, and endpoint posture failures. These data sources support statistical modeling because they contain event counts and time-stamped sequences. Incident response systems provide complementary outcome data by documenting security events that require triage, containment, and remediation. In quantitative security research, incident counts often follow non-normal distributions and may require count-based modeling approaches such as Poisson regression, negative binomial regression, or zero-inflated models when incident occurrences are sparse. For longitudinal evaluation, interrupted time series analysis is frequently used to assess whether a policy intervention produces a statistically significant change in incident trends (Al-Aqrabi & Hill, 2020). Pre–post designs can also be applied when policy implementation occurs at a defined time, although such designs require careful consideration of confounding variables such as changes in threat landscape, staffing, or detection tooling. Measurement validity in cybersecurity studies is strengthened through triangulation of multiple data sources, such as combining ISE logs with SIEM incident data and AD event logs. This approach reduces the risk that observed changes reflect only reporting differences rather than actual security improvement. Additionally, quantitative measurement can include security posture indicators, such as reductions in unmanaged endpoints, increased use of strong authentication, and decreased access from non-compliant devices. Internationally, security measurement is increasingly framed through risk-based reporting models, where metrics are linked to governance objectives and compliance requirements. The academic literature on security metrics emphasizes that effective metrics must be measurable, consistent, comparable, and aligned with security goals. Access control metrics are particularly valuable because access control is one of the most foundational and auditable security domains (Krishnamurthi et al., 2020). Furthermore, identity-based Zero Trust implementations produce large volumes of log data, which makes them suitable for quantitative analysis using time series methods and machine learning-based anomaly detection models. In research settings, log-based analysis also supports reproducibility because events can be systematically extracted, categorized, and analyzed. Therefore, quantitative approaches provide a structured pathway to measure the security impact of identity-based access controls by translating policy enforcement activity into statistically testable changes in incident reduction outcomes.

Although this paper is positioned as a quantitative study, the title indicates a mixed-methods structure, which suggests that quantitative incident analysis is supported by qualitative interpretation of policy implementation and operational experiences (Antoniou et al., 2018). In cybersecurity research, mixed-methods designs are used when numerical outcomes alone do not fully explain the mechanisms behind observed security changes. Quantitative results can show whether incidents decreased after implementing identity-based Zero Trust controls, while qualitative insights can clarify how policy enforcement, administrative practices, and user behavior contributed to the observed patterns. In identity-based access control environments, operational realities such as policy tuning, exception handling, user onboarding workflows, and device compliance management can influence incident outcomes. Therefore, a mixed-methods orientation supports more accurate interpretation of quantitative findings, even when the primary emphasis is statistical measurement. From an enterprise security management perspective, identity-based policy enforcement is both a technical and organizational process (Mongeau, 2021). Administrators must define group mappings, configure authentication protocols, set posture policies, manage certificate lifecycles, and align access rules with business requirements. These implementation processes can influence the effectiveness of Zero Trust controls by shaping how consistently policies are enforced. In addition, user behavior affects security outcomes because authentication requirements may influence compliance, workarounds, and account management practices. Research literature in usable security and organizational security highlights that security controls interact with human workflows and institutional constraints, which can influence measured outcomes. Therefore, quantitative incident reduction analysis can be strengthened when supported by contextual information about policy deployment and operational conditions. In an international context, this is particularly relevant because organizations differ in governance maturity, IT infrastructure standardization, and security staffing. The quantitative component of a mixed-

methods design often focuses on statistical testing of incident changes, while the qualitative component supports explanation of enforcement fidelity and operational adaptation. In the context of Cisco ISE + AD, qualitative data could involve administrator interviews, policy documentation analysis, and workflow mapping of incident response processes. However, within a quantitative paper framing, the introduction must establish the logic of measurement, define the empirical focus, and justify why incident reduction is an appropriate indicator of security impact. The study's focus on identity-based policies is consistent with modern cybersecurity risk literature, which emphasizes that identity is a central control point in contemporary enterprise architectures. Therefore, the quantitative evaluation of identity-based Zero Trust access controls contributes to empirical security measurement by examining whether policy enforcement at the identity and network access layer is associated with measurable reductions in cybersecurity incidents. This framing supports a research design that is aligned with security metrics literature, access control theory, and enterprise Zero Trust architecture guidance.

LITERATURE REVIEW

The rapid expansion of distributed enterprise environments, hybrid cloud architectures, and remote workforce models has significantly challenged traditional perimeter-based security frameworks. Conventional network defense strategies—built upon implicit trust within internal networks—have proven insufficient against advanced persistent threats, credential-based attacks, lateral movement, and insider risks (Abiodun et al., 2021). In response, Zero Trust Architecture (ZTA) has emerged as a paradigm shift in cybersecurity strategy, emphasizing continuous verification, least-privilege access, and identity-centric policy enforcement. Within enterprise implementations, platforms such as Cisco Identity Services Engine (ISE) integrated with Active Directory (AD) represent practical realizations of identity-based Zero Trust controls, enabling granular authentication, authorization, network segmentation, and policy-driven access enforcement. However, despite widespread industry adoption, empirical and quantitatively grounded research measuring the security impact of such implementations—particularly in terms of incident reduction—remains comparatively limited.

The literature increasingly acknowledges that identity has become the new security perimeter. Identity-based access controls, multi-factor authentication (MFA), network access control (NAC), and dynamic segmentation policies collectively form the operational backbone of Zero Trust models. Cisco ISE, when integrated with AD, facilitates centralized identity governance, policy orchestration, device profiling, posture assessment, and role-based access control (RBAC). Yet, the extent to which these controls measurably reduce security incidents—such as unauthorized access attempts, lateral movement events, malware propagation, and privilege escalation—requires systematic investigation. Existing studies often focus on conceptual frameworks, maturity models, or implementation guidelines, with fewer offering quantitative incident comparisons pre- and post-Zero Trust deployment (Becker et al., 2019).

Moreover, the complexity of Zero Trust implementation demands a mixed-methods perspective. Quantitative log analysis (e.g., Security Information and Event Management [SIEM] metrics, failed authentication counts, mean time to detect [MTTD], and mean time to respond [MTTR]) must be complemented by qualitative insights into policy configuration challenges, user friction, compliance behavior, and administrative governance practices. The integration of Cisco ISE with AD introduces specific technical variables—such as dynamic VLAN assignment, endpoint posture enforcement, and conditional access policies—that warrant focused examination within the broader Zero Trust discourse.

Accordingly, this literature review synthesizes theoretical, technical, and empirical scholarship related to Zero Trust security models, identity-based access control mechanisms, Cisco ISE and AD integration frameworks, and measurable incident reduction outcomes. The review is structured to progressively move from foundational Zero Trust principles toward specific quantitative evaluation metrics relevant to enterprise implementations (Garbis & Chapman). Emphasis is placed on studies that employ statistical measurement, log analytics, controlled implementation assessments, and comparative security performance analysis.

Zero Trust Architecture (ZTA) Foundations

Traditional perimeter-based security architectures were originally designed to protect clearly defined organizational boundaries, assuming that threats originated primarily outside the network. Firewalls, intrusion detection systems, and demilitarized zones formed layered defenses intended to separate

trusted internal systems from untrusted external environments (Kristen et al., 2021). However, the rapid adoption of cloud computing, remote work models, mobile devices, and third-party integrations has significantly weakened the effectiveness of this model. As enterprise resources moved beyond fixed network perimeters, implicit trust within internal networks created exploitable vulnerabilities. Research consistently demonstrates that once attackers bypass perimeter defenses—often through credential compromise—flat internal network structures enable rapid lateral movement and privilege escalation. Large-scale breach investigations reveal that attackers frequently rely on stolen credentials and insufficient segmentation to expand access, increasing both the scale and duration of incidents.

The inadequacy of perimeter-centric defenses is further compounded by insider threats and credential misuse, which occur within trusted environments and therefore evade traditional boundary controls (Faysal & Shamsunnahar, 2022; Garbis & Chapman; Habibullah & Zaheda, 2022). These structural weaknesses have driven the evolution toward Zero Trust Architecture (ZTA), which rejects the notion of implicit trust based on network location. Instead, ZTA redefines security boundaries around identity, device posture, and contextual verification. By eliminating automatic trust and enforcing strict access validation, organizations aim to reduce unauthorized access events and minimize attack propagation. Empirical assessments of early Zero Trust deployments indicate measurable improvements in internal access control and containment efficiency. These findings underscore the necessity of transitioning from perimeter-based models to identity-driven architectures capable of producing statistically observable reductions in security incidents (Jahangir & Md Shahab, 2022; Kenyon, 2018; Ratul & Subrato, 2022). Zero Trust Architecture is operationalized through several interdependent principles, including continuous verification, least-privilege access enforcement, and granular micro-segmentation. Continuous verification replaces one-time authentication with dynamic access validation informed by contextual risk signals such as device compliance, geographic location, behavioral anomalies, and session activity (Keeriyattil; Tahmina Akter Bhuya & Rebeka, 2022). Studies evaluating continuous authentication frameworks show measurable decreases in successful account compromise incidents compared to static credential-based systems. By requiring reauthentication and contextual validation, organizations alter authentication success and failure ratios, thereby reducing the probability of unauthorized session persistence.

Figure 3: Perimeter to Zero Trust Framework

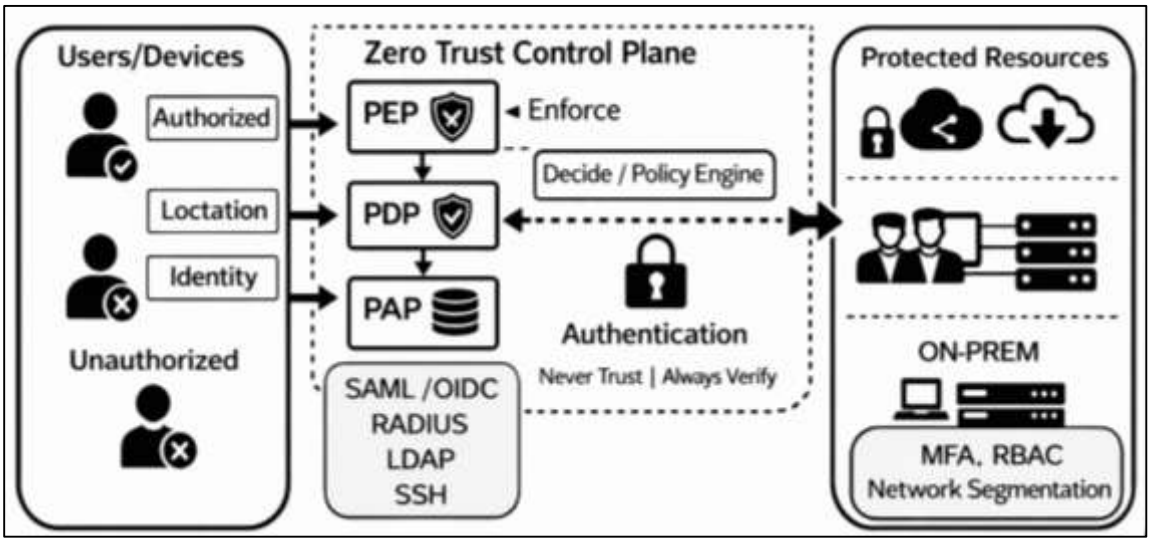

Least-privilege enforcement limits users and systems to only the permissions required for legitimate tasks (AI). Research examining privilege management strategies demonstrates that excessive administrative rights are strongly correlated with malware spread and unauthorized configuration changes. Reducing privileged account exposure has been associated with measurable declines in internal misuse events and containment times during simulated attack exercises. Micro-segmentation complements least-privilege enforcement by isolating workloads and restricting east-west traffic within networks (Kamminga et al., 2018). Network flow analyses in segmented environments reveal

statistically significant reductions in lateral movement attempts, particularly during penetration testing and red team simulations.The formalization of Zero Trust principles into structured frameworks has further enabled quantitative performance measurement. By defining policy decision points and enforcement mechanisms, organizations can track metrics such as breach dwell time, containment duration, and high-severity incident frequency. Collectively, the literature indicates that Zero Trust principles are not merely conceptual shifts but operational controls that produce measurable security performance improvements across enterprise environments.

The evaluation of Zero Trust effectiveness increasingly relies on quantifiable security metrics that capture changes in incident patterns before and after implementation. One widely used measure is incident frequency rate, which tracks the number of confirmed security events within a defined period. Comparative analyses across organizations adopting Zero Trust controls demonstrate downward trends in unauthorized access incidents, malware spread events, and internal policy violations following deployment. This reduction is often attributed to stricter authentication enforcement and segmentation controls that limit attack surfaces.Another critical metric involves the percentage reduction in lateral movement events. In traditional flat networks, compromised credentials often allow attackers to pivot between systems with minimal resistance. Post-implementation assessments of segmented environments show marked decreases in east-west traffic exploitation attempts and reduced propagation of malicious processes across endpoints. These findings are frequently supported by log analytics and intrusion detection data, which demonstrate fewer cross-segment access anomalies (More & Haff).Breach dwell time—the duration between compromise and detection—also serves as a central indicator of Zero Trust performance. Studies report shorter detection windows in environments with continuous monitoring and policy enforcement points, suggesting improved visibility and response efficiency. Additionally, organizations measure attack surface reduction by analyzing exposed services, privileged accounts, and open network pathways before and after Zero Trust deployment. Reductions in these exposure metrics are correlated with declines in successful credential compromise incidents. Collectively, these quantitative indicators provide empirical evidence that Zero Trust adoption is associated with measurable decreases in incident frequency and severity (Mongeau & Hajdasinski, 2021).

Comparative statistical studies examining security environments before and after Zero Trust adoption provide further empirical grounding for its effectiveness. Pre-implementation environments are typically characterized by higher volumes of unauthorized authentication attempts, broader privilege distributions, and increased susceptibility to lateral movement (Haff, 2021). Post-deployment assessments often reveal statistically significant reductions in successful credential misuse, particularly when identity-based policies are integrated with multifactor authentication and device compliance checks.Security event count comparisons frequently demonstrate declines in high-severity incidents after the enforcement of identity-based access controls and segmentation policies. For example, internal audits and red team simulations conducted before Zero Trust deployment often identify widespread privilege escalation pathways, whereas follow-up assessments show constrained attack paths and improved containment (Tyler & Viana, 2021). Additionally, reductions in repeat incident rates suggest that identity-centric policy enforcement mitigates recurring vulnerabilities.The statistical evaluation of credential compromise incidents provides particularly compelling evidence. Environments that transition from location-based trust to identity-driven access validation report decreases in unauthorized account activity and abnormal login patterns (Parrend et al., 2018). Time-series analyses further indicate sustained incident reductions rather than temporary improvements, suggesting structural security enhancement rather than short-term anomaly suppression. Overall, comparative pre- and post-implementation findings consistently support the conclusion that Zero Trust architectures contribute to measurable and statistically observable reductions in security incidents, thereby reinforcing the value of identity-based access control frameworks in modern enterprise environments (Borky & Bradley, 2018).

Identity-Based Access Control (IBAC) Models

The evolution of Identity and Access Management (IAM) has increasingly positioned identity as the dominant control point for enterprise cybersecurity. Traditional access models were designed for stable, on-premises infrastructures where users, devices, and applications existed within fixed

boundaries (Steingartner et al., 2021). However, the modern enterprise environment is characterized by cloud platforms, remote workforces, mobile endpoints, and outsourced service ecosystems, which collectively dissolve the reliability of network location as a trust signal. As a result, identity has emerged as the most consistent and enforceable perimeter for controlling access. Research has shown that the majority of contemporary breaches involve compromised credentials rather than direct exploitation of technical vulnerabilities, reinforcing the centrality of identity controls in reducing incident likelihood. Security industry studies further demonstrate that credential theft, password spraying, and session hijacking represent persistent and scalable attack vectors across industries, especially in hybrid cloud ecosystems (Kaur, Lashkari, et al., 2021).The literature also highlights that identity-based security has expanded beyond authentication into continuous access governance, where identity is linked with contextual verification signals. IAM frameworks increasingly integrate device posture, user risk scoring, location intelligence, and behavioral analytics to reduce reliance on static credentials. This shift has been associated with measurable reductions in unauthorized access attempts and improved containment of account compromise. Importantly, IAM evolution has also been driven by regulatory compliance pressures and the need for auditable access trails (Hardjono & Pentland, 2018). The growing emphasis on identity-centric access models supports the argument that modern security strategies are increasingly dependent on enforcing trust decisions through identity attributes rather than perimeter assumptions.

Figure 4: Identity-Centric Access Control Framework

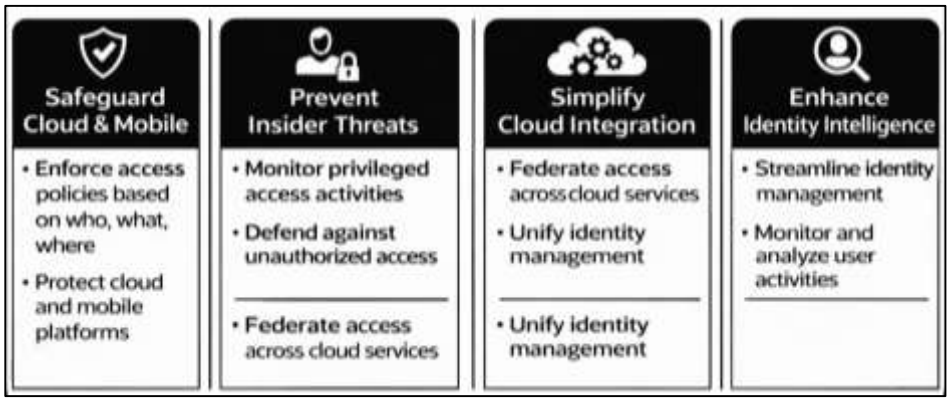

Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC) represent two dominant models for structuring identity-based access enforcement. RBAC assigns permissions based on predefined organizational roles, offering simplicity, scalability, and easier administrative governance (Kaur, Habibi Lashkari, et al., 2021). The literature consistently identifies RBAC as effective in environments with stable job functions and predictable access needs. However, studies also note that RBAC can lead to role explosion and permission creep, where users accumulate excessive privileges over time due to job transitions, temporary assignments, and legacy access requirements (Cali et al., 2021). Permission creep is strongly linked to higher rates of privilege misuse and increased risk of escalation during compromise scenarios. Quantitative evaluations of RBAC implementations show that excessive privileges frequently remain unrevoked, increasing exposure to internal misuse and attacker exploitation.In contrast, ABAC evaluates access requests based on multiple attributes such as user identity, device type, location, time of access, resource sensitivity, and risk level. ABAC is widely recognized in the literature as more adaptable for dynamic enterprise environments and better aligned with Zero Trust principles (Mongeau, 2021). Research comparing RBAC and ABAC indicates that ABAC reduces over-privileging by narrowing permissions to contextual conditions rather than static roles. Studies of ABAC-based systems report measurable improvements in least-privilege alignment and reductions in unauthorized access incidents, particularly in cloud and remote access environments. Overall, the literature suggests that while RBAC remains operationally valuable, ABAC provides superior risk reduction potential through contextual access enforcement.

Conditional access policies and multi-factor authentication (MFA) represent critical operational controls within identity-based access systems (Minth, 2018). Conditional access enables organizations to enforce real-time access decisions based on contextual risk indicators such as device compliance, network trust level, user behavior anomalies, and geographic location. The literature suggests that conditional access policies significantly improve access governance by preventing authentication success from automatically granting session trust. This approach is associated with measurable declines in risky login behaviors and reduced exposure to credential misuse. In empirical studies, conditional access enforcement has been linked to reductions in unauthorized login success rates and decreased persistence of compromised sessions.

MFA effectiveness has been extensively documented in cybersecurity research and industry incident reporting. Studies consistently show that MFA reduces the probability of successful credential compromise, particularly against password spraying, brute-force attacks, and phishing-based credential theft. However, the literature also acknowledges that MFA is not universally sufficient, as attackers increasingly employ MFA fatigue techniques, token interception, and session hijacking methods. Despite these limitations, organizations with strong MFA enforcement report measurable decreases in account takeover incidents and improved detection of anomalous authentication events (Pogrebna & Skilton, 2019). Research also highlights the indirect security benefits of MFA, including improved user awareness and reduced reliance on weak password practices. Together, conditional access and MFA form an evidence-based foundation for reducing identity-driven security incidents and strengthening enterprise access controls (Retico et al., 2021).

Cisco Identity Services Engine (ISE)

Network Access Control (NAC) emerged in the literature as a practical response to the growing mismatch between open enterprise connectivity and escalating identity-driven threats (Pogoryelova et al., 2018). Within this space, Cisco Identity Services Engine (ISE) has been positioned as a policy orchestration platform that operationalizes Authentication, Authorization, and Accounting (AAA) as a unified control cycle across users, devices, and network segments. AAA-centered approaches have been discussed as essential for replacing implicit network trust with explicit access decisions, especially where bring-your-own-device practices, contractor access, and unmanaged endpoints increase exposure. The authentication component has been treated as the gatekeeper for identity validation and device association, while authorization has been framed as the enforcement layer that ties identity to policy and network placement. Accounting has been emphasized for its forensic value, linking sessions to identities, endpoints, access outcomes, and policy changes for auditing and incident response workflows.A recurring finding across NAC research is that the strength of access governance depends on the degree of policy centralization and consistency across heterogeneous network devices (Xiang et al., 2021). Cisco ISE has been described in this literature as a mechanism for standardizing access outcomes across wired, wireless, and VPN contexts by connecting authentication results with policy enforcement. At the same time, studies have highlighted operational challenges that affect measurable outcomes, including inconsistent identity sources, incomplete profiling, and policy drift across distributed locations. These issues have been linked to variability in unauthorized access attempts and compliance outcomes, reinforcing the importance of governance discipline and telemetry-driven policy tuning (Pogrebna & Skilton, 2019). Overall, the literature frames Cisco ISE-style NAC as a measurable security control because it produces continuous records of who connected, what posture they presented, which policy applied, and how the network responded—data that directly supports quantitative evaluation of access risk reduction.

The integration of Cisco ISE with RADIUS and TACACS+ has been discussed as a core technical pathway through which AAA policies become enforceable at scale. RADIUS has been widely examined as the dominant access control protocol for network authentication and authorization across switches, wireless controllers, and VPN gateways, enabling consistent policy decisions to be applied to diverse network entry points (Fields et al., 2021). TACACS+ has been treated as particularly relevant for device administration governance, where centralized authentication and granular command authorization reduce the risk of unauthorized configuration changes. Together, these protocol pathways have been

linked in the literature to reductions in inconsistent access outcomes, one of the recurring factors associated with NAC underperformance.This integration has also been tied to measurable improvements in access decision reliability. When authentication and authorization decisions are centrally evaluated rather than locally configured, enterprises reduce policy divergence between network segments, leading to clearer patterns in security event data such as denied connection attempts, repeated failures, and anomalous access requests (Stanković et al., 2017).

Figure 5: Identity-Based Network Access Control

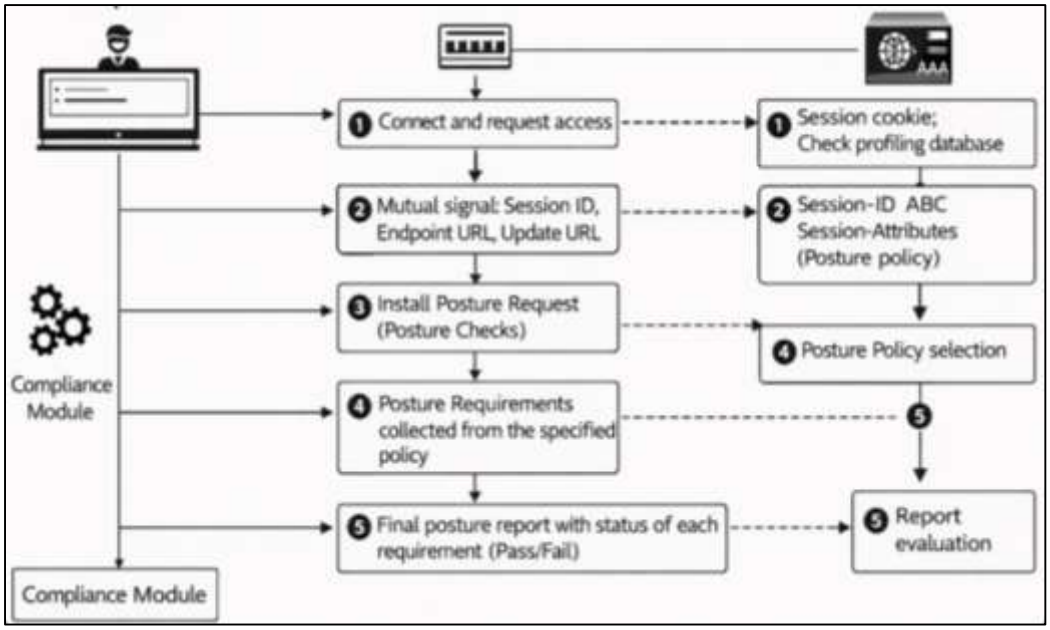

Studies of NAC effectiveness have emphasized that enforcement consistency is a prerequisite for meaningful metrics because inconsistent enforcement produces noisy measurements that obscure real security impact. In applied research settings, improvements in protocol-driven policy uniformity have been associated with clearer downward shifts in unauthorized access attempts after policy standardization. The same work has also connected accounting logs with more reliable incident investigations, since session trails align with identities and authorization results rather than ambiguous IP-based attribution.The literature also reports that protocol-level alignment influences operational metrics such as helpdesk tickets related to access issues, false denials due to misclassification, and time required to remediate access disruptions (Rabat et al., 2021). These operational indicators matter to security outcomes because unstable access controls often lead to policy exceptions and shadow access paths that weaken NAC. In this way, RADIUS/TACACS+ integration is not merely a technical detail in NAC deployments; it acts as a measurable stabilizer that supports both security effectiveness and evaluability.

Endpoint posture assessment has been widely discussed as the feature set that turns NAC from identity-only control into risk-aware control (Wang, 2021). In the literature, posture mechanisms evaluate whether an endpoint aligns with defined security requirements such as operating system version, patch status, disk encryption, endpoint protection presence, firewall state, or other compliance signals. Cisco ISE is commonly described as coordinating posture checks and translating compliance results into access outcomes that either permit normal connectivity, restrict access, or isolate the endpoint into a remediation state. This control loop has been emphasized as particularly relevant because modern incident patterns often begin with compromised or unmanaged endpoints rather than purely external exploitation.Endpoint compliance rate has been treated as a primary quantitative indicator for posture-driven NAC evaluation because it summarizes how many endpoints meet organizational baseline standards at the time of access (Laurent et al., 2021). A second frequently used

indicator is the number of quarantined devices, which captures the volume of endpoints redirected into restricted access due to failed posture, suspicious behavior, or policy violation. The literature interprets quarantine volumes carefully: increased quarantines immediately after rollout often reflect improved detection and enforcement rather than worsening security, while stabilization patterns suggest improved endpoint hygiene and reduced policy violations. Policy violation incidents, including attempts to connect non-compliant devices to sensitive segments or repeated failures against restricted access rules, have been described as valuable signals of both attacker behavior and user noncompliance trends.Research that evaluates NAC effectiveness frequently connects posture enforcement with reductions in malware spread and lateral movement opportunity by limiting what noncompliant devices can reach (Lee, 2021). Studies also link posture mechanisms to improved remediation speed, because restricted access states can steer endpoints toward updates and compliance tools. This body of work treats posture assessment as a measurable mechanism for reducing risk exposure at the exact moment of network entry, producing data streams that support incident correlation analysis, compliance trend monitoring, and policy optimization.

Active Directory (AD) Integration and Identity Governance

Active Directory (AD) is widely treated in the identity governance literature as the foundational enterprise identity repository because it centralizes authentication, authorization group memberships, and directory-based access dependencies across endpoints, servers, and applications (Garbis & Chapman, 2021). This centralization creates operational efficiencies, but it also concentrates risk: when AD is compromised, attackers can rapidly expand privileges, persist across systems, and undermine trust relationships across the domain. Studies of enterprise intrusions consistently characterize AD as a primary target due to its control over credential validation, service accounts, and privileged group structures. Because AD binds identity to resource access at scale, identity governance within AD becomes inseparable from security outcomes such as unauthorized access success, escalation pathways, and the capacity for domain-wide control. Research further emphasizes that the modern enterprise's reliance on directory-integrated services—including federated authentication, file systems, remote management tools, and application access—amplifies the blast radius of directory weaknesses (Gupta & Giri, 2018).Shows of real-world breach patterns frequently describe a common progression: initial foothold through phishing or endpoint compromise, credential harvesting, and subsequent pivoting into directory infrastructure. In this framing, AD is not simply a user database but an operational trust fabric that links devices, administrators, policies, and services. As a result, identity governance research recommends treating AD as a hardened security tier with restricted administrative access, robust auditing, and disciplined privilege management. Quantitative evaluation becomes feasible because AD produces high-density telemetry—logon events, group membership changes, directory modifications, Kerberos ticket operations, and policy application traces—that can be analyzed to infer exposure and attacker behavior. Literature synthesizing directory security therefore positions AD both as a critical enabler of enterprise access and as a measurable focal point for understanding identity-driven risk, especially when incident reduction goals are tied to changes in privileged access volume and authentication anomaly patterns (Garbis & Chapman).

Group Policy Objects (GPOs) are central to AD governance because they operationalize access restrictions, security baselines, and configuration controls across domain-joined systems (Mukhopadhyay et al., 2019). The literature frames GPOs as a powerful mechanism for enforcing consistent security posture—password policies, local administrator restrictions, firewall settings, script execution controls, and endpoint hardening requirements—while also acknowledging that GPO complexity can become a source of misconfiguration risk. Governance studies repeatedly describe how policy layering, inheritance, and inconsistent organizational unit design can produce unintended access permissions or weaken hardening baselines. When GPO design is not tightly governed, enterprises can experience policy drift, where controls diverge across departments or sites, resulting in uneven security posture and exploitable exceptions (Schwartz & Machulak, 2018).Access restriction governance is also discussed through the lens of privilege sprawl. As organizations evolve, users change roles, administrators create temporary access exceptions, and legacy group memberships persist. The literature shows that this accumulation of privileges increases exposure to misuse and facilitates attacker escalation once any single credential is compromised. Privileged groups, delegated

administration, and service account permissions become key risk multipliers because they create direct pathways to domain-level impact. Studies of directory attacks highlight that attackers often exploit mis-scoped delegation, overly broad admin rights, and weakly controlled group membership change processes rather than relying solely on software vulnerabilities (Panarello et al., 2018). This makes the governance of GPO change management and privileged group hygiene a measurable security issue. Quantitative indicators often include counts of privileged accounts, frequency of privileged group modifications, and the prevalence of policy exceptions. The literature therefore links effective GPO governance not only to compliance but to observable reductions in risky configuration states and fewer opportunities for adversaries to exploit policy inconsistencies to gain elevated access.

Figure 6: Active Directory Identity Governance Framework

Zero Trust Environments

Incident reduction in Zero Trust environments is commonly framed in the literature as an operational outcome that can be observed through Security Operations Center (SOC) performance metrics rather than solely through architectural compliance. As organizations shift away from perimeter assumptions toward continuous verification and least-privilege access, research emphasizes that effectiveness must be demonstrated through measurable changes in detection, response, and containment performance (Ribeiro & Barbosa-Povoa, 2018). SOC-oriented studies argue that Zero Trust alters the incident landscape by increasing visibility into identity and access events, producing more granular telemetry and enabling earlier intervention when suspicious activity begins. This shift is frequently linked to improvements in mean time to detect and mean time to respond because analysts receive richer event context and clearer policy enforcement signals, particularly where identity, device posture, and segmentation outcomes are logged consistently.The literature also highlights that Zero Trust can reshape alert streams, sometimes increasing alert volume initially due to stricter policy enforcement and improved telemetry coverage (Pagano et al., 2019). This operational reality pushes researchers to evaluate incident reduction not simply as a decline in total alerts but as a change in confirmed incident frequency, incident severity distribution, and the speed at which incidents are triaged and contained. In SOC research, time-based metrics are treated as especially meaningful because they capture both capability and readiness: faster detection reflects better sensing and correlation, while faster response reflects mature workflows, automation, and containment playbooks. Incident reduction studies further argue that SOC effectiveness in Zero Trust settings depends on correlating identity, endpoint, and network access logs into coherent narratives that reduce investigative friction (Cantelmi et al., 2021). In this sense, incident reduction is interpreted as a multidimensional performance shift reflected in time-

to-action metrics, confirmation rates, and containment outcomes that can be compared across observation windows before and after Zero Trust enforcement.

Log analytics and SIEM correlation are central to how the literature conceptualizes incident reduction measurement in Zero Trust environments (Goodman et al., 2020). Because Zero Trust emphasizes continuous verification and policy enforcement at multiple control points, it generates high-density telemetry that includes authentication outcomes, authorization decisions, segmentation placements, posture results, and session changes. Researchers describe SIEM platforms as the operational layer where these signals become actionable because correlation rules, behavior analytics, and enrichment pipelines transform raw event streams into prioritized alerts and incident cases. The literature also emphasizes that without strong log normalization and cross-domain correlation, Zero Trust telemetry can overwhelm analysts and obscure the very outcomes it aims to improve. Therefore, measurement studies frequently focus on how well SOC teams can transform Zero Trust data into faster investigations and more consistent containment.Breach containment time is repeatedly treated as a core outcome indicator because it captures the time required to stop attacker progression after suspicious activity is detected (González-Granadillo et al., 2021). In the literature, containment speed improves when Zero Trust controls provide precise enforcement mechanisms such as isolating endpoints, restricting identity privileges, or blocking access to sensitive segments based on real-time risk. Researchers note that Zero Trust can reduce attacker dwell time by limiting lateral movement options and by producing earlier signals of policy violations, such as repeated authorization failures or anomalous access attempts to restricted resources. Studies also connect the maturity of SIEM correlation with measurable reductions in time-to-containment, arguing that better correlation reduces investigative delays caused by fragmented evidence (Ramacciati et al., 2018). By aligning identity and access logs with endpoint and network indicators, organizations can identify compromised accounts faster and contain them with fewer manual steps. Across the literature, the strongest claims about incident reduction tend to be those supported by operational telemetry showing shorter containment times and more efficient correlation-driven triage in post-implementation observation windows.

Figure 7: Zero Trust SOC Performance Metrics

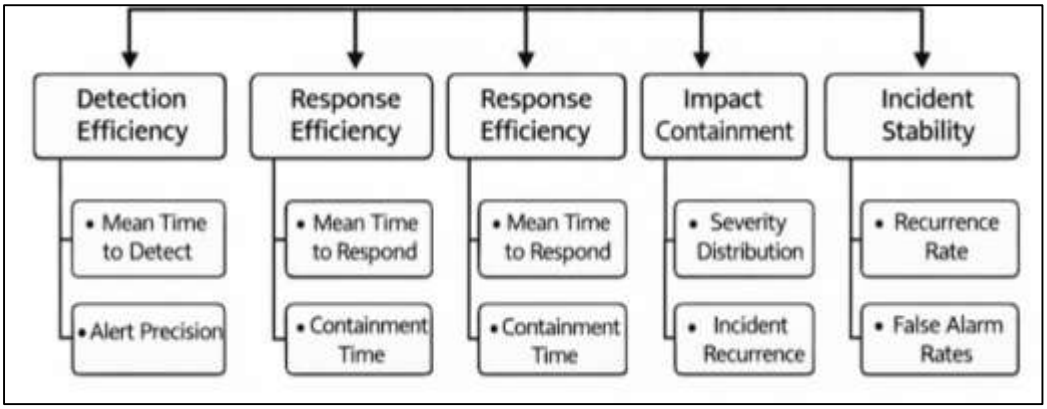

Another major theme in the literature is that incident reduction should be evaluated not only by counting incidents, but by examining incident severity classification frequency and recurrence rates (Ahmadi-Assalemi et al., 2020). Researchers frequently argue that mature Zero Trust enforcement can reduce the proportion of high-severity incidents by constraining attacker pathways and limiting the blast radius of compromised accounts. In this framing, even if low-level alerts or policy violations increase initially, the more meaningful outcome is a downward shift in incidents that reach high-impact thresholds such as domain compromise, data exfiltration, or ransomware execution. Severity distribution analysis has been presented as a method for showing whether security controls are reducing the escalation potential of common intrusion sequences (Pasman et al., 2017).Incident recurrence rate is also used as a stability metric, capturing whether similar incidents repeat over time due to unresolved root causes, weak governance, or inconsistent enforcement. Studies describe recurrence reductions as evidence of stronger policy consistency, better remediation discipline, and

improved preventive controls such as conditional access and segmentation. At the same time, the literature cautions that recurrence trends can be distorted if detection coverage changes, which is why many researchers recommend pairing recurrence analysis with detection-quality metrics. False positive and false negative rates are repeatedly highlighted as critical to SOC performance because excessive false positives can slow response and drive alert fatigue, while false negatives allow incidents to progress undetected (Zhang et al., 2018). In Zero Trust settings, detection quality is often linked to how well identity signals and access policy outcomes are used to refine alert rules and reduce noise. Overall, the literature supports the interpretation that meaningful incident reduction is reflected in fewer repeated incidents, a smaller share of high-severity cases, and improved detection precision measured through reduced false alarms and fewer missed events.

The literature frequently recommends comparative time-series methodologies to assess incident reduction outcomes before and after Zero Trust implementation. These approaches are presented as necessary because security incident patterns vary due to seasonal factors, changing attacker activity, organizational growth, and shifts in monitoring coverage. Rather than relying on single snapshot comparisons, studies use longitudinal observation windows that track detection times, response times, incident counts, and severity distributions across multiple periods (Lodi et al., 2021). Time-series evaluation allows researchers to distinguish implementation effects from background volatility, particularly when combined with baseline stabilization periods after policy rollout.Comparative analyses often emphasize measuring change patterns in mean time to detect and mean time to respond, alongside time-to-containment metrics that reflect end-to-end operational capability. Researchers also recommend tracking the stability of severity classifications over time to determine whether incident escalation patterns shift downward following enforcement of least privilege and segmentation. To improve interpretability, some studies propose grouping metrics into operational categories such as detection efficiency, response efficiency, and impact containment. This structure supports synthesis across organizations with different tooling, since absolute incident volumes may not be comparable but time-based performance measures often are (Yang et al., 2018).The literature also notes that post-implementation measurement must account for telemetry expansion. When Zero Trust increases log visibility, incident discovery may initially rise because previously hidden threats become observable. Time-series methods help address this by evaluating whether detection and response times improve even if raw incident counts fluctuate (Golan et al., 2020). Across studies, the central conclusion is that rigorous pre/post evaluation of Zero Trust depends on consistent observation windows, stable definitions of incident severity, and longitudinal tracking of SOC performance metrics that capture both speed and impact reduction.

Micro-Segmentation and Lateral Movement Containment

Micro-segmentation is widely discussed in the Zero Trust and enterprise defense literature as a targeted response to the problem of lateral movement after initial compromise (de la Torre & Polo, 2018). Whereas traditional segmentation often relies on coarse network boundaries such as VLANs, subnets, or site-based zones, micro-segmentation emphasizes fine-grained isolation at the workload, application, or identity level. This approach aligns closely with Zero Trust assumptions because it treats internal traffic as untrusted by default and requires explicit policy authorization for east-west communication between internal endpoints, servers, and services (Eichelberg et al., 2020). Research synthesizing breach case patterns consistently highlights that once attackers acquire valid credentials or gain a foothold on a single endpoint, they prioritize east-west traversal to locate privileged accounts, critical data stores, and domain control infrastructure. Micro-segmentation directly interrupts this progression by constraining what compromised identities or devices can access, thereby reducing the number of viable attack paths inside the environment.The literature also frames micro-segmentation as a control that converts implicit connectivity into enforceable policy decisions. When segmentation is identity-driven, the trust boundary is not a static network address range but an authorization rule linking who or what is communicating to which resource and under what contextual conditions (Di Salvo, 2018). This shifts containment from a reactive "block after detection" model toward a preventive structure that narrows attacker movement opportunities from the outset. Studies on enterprise segmentation strategies further emphasize that micro-segmentation supports measurable security outcomes because it produces observable enforcement events, such as blocked inter-service requests or

denied cross-segment flows, which can be tracked over time. In this way, micro-segmentation becomes both an architectural and measurable intervention, enabling comparative analysis of lateral movement indicators before and after policy isolation is deployed.

Figure 8: Micro-Segmentation Zero Trust Roadmap

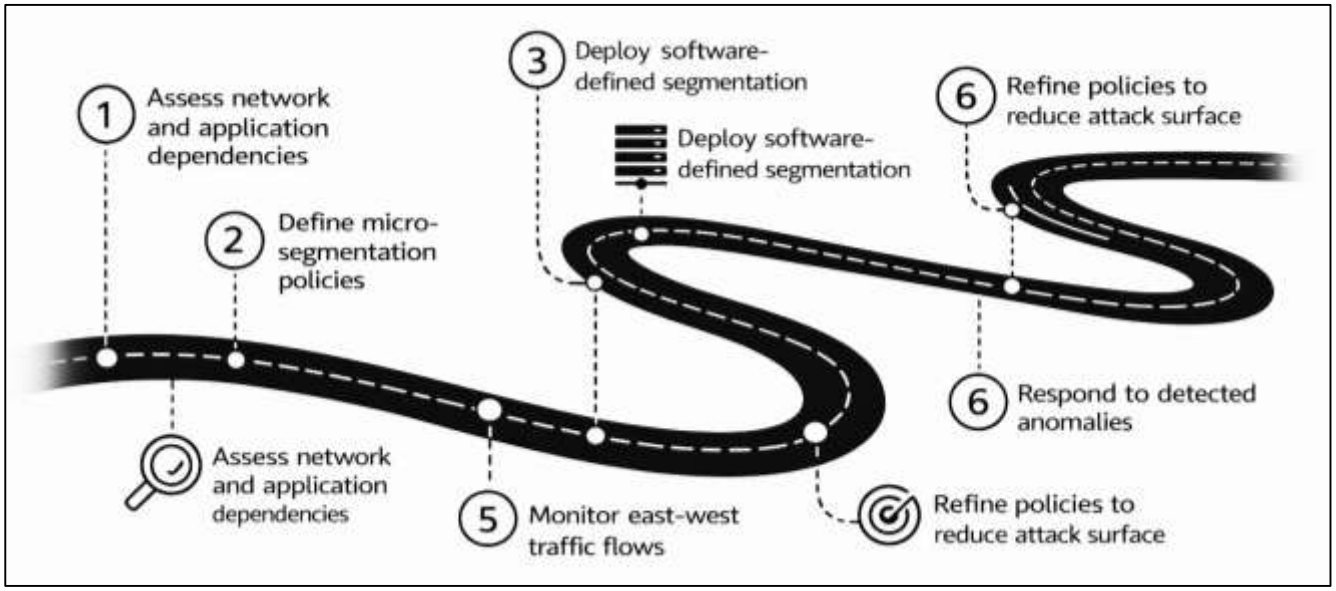

Software-defined segmentation is frequently described in the literature as the implementation approach that makes micro-segmentation operationally feasible at enterprise scale. Rather than relying solely on network device configurations, software-defined approaches apply segmentation policies through centrally managed enforcement points, often integrated with identity stores, endpoint posture systems, and application dependency maps (Badhwar, 2021). This design is emphasized because static network segmentation struggles to keep pace with dynamic workloads, virtualized infrastructure, containerized applications, and hybrid environments where resources shift rapidly across locations. Policy-based traffic isolation allows organizations to define rules based on identities, device categories, application roles, or sensitivity levels, enforcing these rules consistently even as network topology changes (Suomalainen et al., 2018).Research on segmentation deployments notes that effective traffic isolation depends on accurate visibility into legitimate communication patterns. Studies describe how baseline application flows must be identified to avoid disrupting business operations while still limiting unnecessary connectivity. The literature also stresses that micro-segmentation policies can be tuned incrementally, beginning with monitoring modes that map east-west flows and progressing to enforcement modes that block unauthorized connections. This staged approach is associated with fewer operational disruptions and improved policy precision. Empirical research on segmentation effectiveness frequently links policy-based isolation to measurable reductions in exploit propagation and unauthorized internal scanning because attackers can no longer freely probe adjacent systems (Wall & Rodrick). Additionally, segmentation is reported to improve the quality of detection by creating clearer signals: blocked flows and denied attempts become meaningful indicators of suspicious behavior, especially when combined with identity context. Across studies, software-defined segmentation is positioned as the mechanism that turns micro-segmentation from an architectural goal into a practical, auditable, and measurable control that can be evaluated through enforcement logs and flow-based telemetry.

A recurring analytical distinction in the literature is between north-south traffic, which crosses the boundary between internal networks and external environments, and east-west traffic, which occurs within internal environments between hosts and services (Serrano, 2020). While traditional perimeter security primarily focuses on north-south flows, modern enterprise attack patterns increasingly emphasize east-west exploitation because internal traversal enables attackers to escalate privileges, reach sensitive repositories, and establish persistence. Risk modeling studies highlight that east-west

traffic can be more dangerous than north-south traffic once initial compromise occurs because internal trust assumptions and broad connectivity often provide attackers with low-friction movement. Micro-segmentation directly targets this risk profile by restricting internal communication pathways and requiring explicit authorization for inter-service connections (Kamboj & Pal, 2021).Network flow analysis is commonly proposed as a quantitative method for evaluating micro-segmentation outcomes because flows capture attempted connections regardless of whether attacks succeed. Metrics frequently discussed include unauthorized east-west connection attempts, traffic anomaly rates, and inter-segment breach probability indicators inferred from observed access patterns. Unauthorized connection attempts represent denied or blocked flows that violate policy, and their trend patterns are often interpreted as indicators of how frequently potential lateral movement is being attempted or prevented. Traffic anomaly rate measures deviations from baseline flow behavior, including unusual port usage, unexpected service discovery patterns, or sudden increases in cross-segment requests. Inter-segment breach probability is often discussed as an outcome-oriented concept derived from how many sensitive segments remain reachable from compromised zones under current policies. The literature suggests that combining these indicators provides a richer assessment than relying on incident counts alone, since segmentation may prevent attacks that never become confirmed incidents (Hoang & Farahmandian, 2017). Overall, flow-based measurement supports evaluating whether micro-segmentation reduces lateral movement opportunities by producing observable declines in unauthorized internal connectivity and a more stable baseline of legitimate east-west communication.

User Behavior and Security Posture Analytics

The literature on Zero Trust adoption consistently emphasizes that technical controls alone do not determine security outcomes; user behavior and organizational practices strongly shape whether identity-centric policies actually reduce incidents (Ali et al., 2021). Human factors research in cybersecurity shows that security systems interact with cognition, routines, and workplace pressures, meaning that even well-designed controls can be undermined by workaround behavior, low trust in security teams, or misunderstanding of policy intent. In Zero Trust settings, where continuous verification and strict access rules are common, researchers highlight that the user experience becomes a security variable because friction influences compliance. If authentication and authorization steps are perceived as excessively disruptive, users may attempt to bypass controls through shadow IT tools, credential sharing, storing passwords insecurely, or seeking persistent exceptions (Addae et al., 2019). The literature also notes that security culture, leadership messaging, and the perceived legitimacy of controls influence whether users cooperate with access restrictions or treat them as obstacles.Studies on security behavior indicate that compliance is more likely when policies are clear, justifiable, and consistently enforced rather than applied unevenly across teams. Inconsistent enforcement is frequently linked to normalization of deviance, where repeated exceptions become routine and reduce the perceived importance of rules. Zero Trust implementations therefore require attention to change management, training, and usability to avoid creating an environment where friction drives non-compliance. Research further suggests that human-centered security design can improve adherence by minimizing unnecessary prompts, aligning access workflows with task demands, and using adaptive authentication that increases friction only when risk signals rise. This framing is important for incident reduction measurement because behavioral compliance determines the real exposure level of identity systems (Broeders et al., 2017). As a result, literature increasingly treats user behavior as a quantifiable component of Zero Trust effectiveness that must be measured through compliance rates, authentication patterns, and risk-scored behavior trends rather than inferred from policy existence alone.

Policy fatigue and access friction are recurring explanatory constructs in cybersecurity literature, particularly in environments that require frequent authentication or repeated approval processes (Trang & Brendel, 2019). Policy fatigue occurs when repeated security prompts and complex requirements cause users to disengage cognitively, leading to superficial compliance, inattentive approvals, or avoidance behavior. Access friction describes the practical burden users experience when security requirements delay or interrupt task execution, such as repeated multi-factor prompts, stringent password policies, or restrictive access segmentation that complicates workflow. Research indicates that when friction exceeds perceived benefit, users shift toward compensatory behaviors, including credential reuse, insecure storage of secrets, reliance on personal devices, or attempts to

maintain persistent sessions to avoid reauthentication. These behaviors are significant because they create measurable exposure pathways even in identity-centric environments.The literature also highlights that friction is not uniformly experienced; users in high-tempo roles or operational environments may face more frequent interruptions, increasing the likelihood of risky shortcuts. Studies analyzing authentication telemetry often identify patterns that reflect friction-driven behavior, such as repeated failed logins, elevated helpdesk requests for access resets, frequent account lockouts, and increased use of password recovery channels (Torten et al., 2018). These operational signals can serve as indirect indicators of policy burden and compliance stress. Importantly, policy fatigue can degrade the effectiveness of security controls by turning protective steps into habitual clicks, as seen in repeated approval behaviors for prompts and notifications. Research further suggests that sustained friction without supportive training and adaptive design can increase resentment and reduce reporting behavior, meaning that users may avoid contacting security teams when suspicious activity occurs (Enachescu et al., 2019). The synthesis across studies suggests that Zero Trust success depends on balancing verification rigor with usability: when friction is optimized and adaptive, compliance improves; when friction is excessive and static, measurable non-compliance patterns rise, increasing incident probability despite formal policy strength.

Insider threat research provides a foundation for understanding how behavioral signals and organizational context can be used to model risk and anticipate incidents (Joshi & Singh, 2017). The literature distinguishes between malicious insiders, negligent insiders, and compromised insiders, emphasizing that all three categories can be detected through measurable behavioral indicators. In Zero Trust environments, identity and access telemetry becomes especially valuable for insider threat modeling because access decisions are logged with high granularity, enabling the analysis of what users attempted, what they accessed, and how their patterns deviate from norms.

Figure 9: User Behavior in Zero Trust

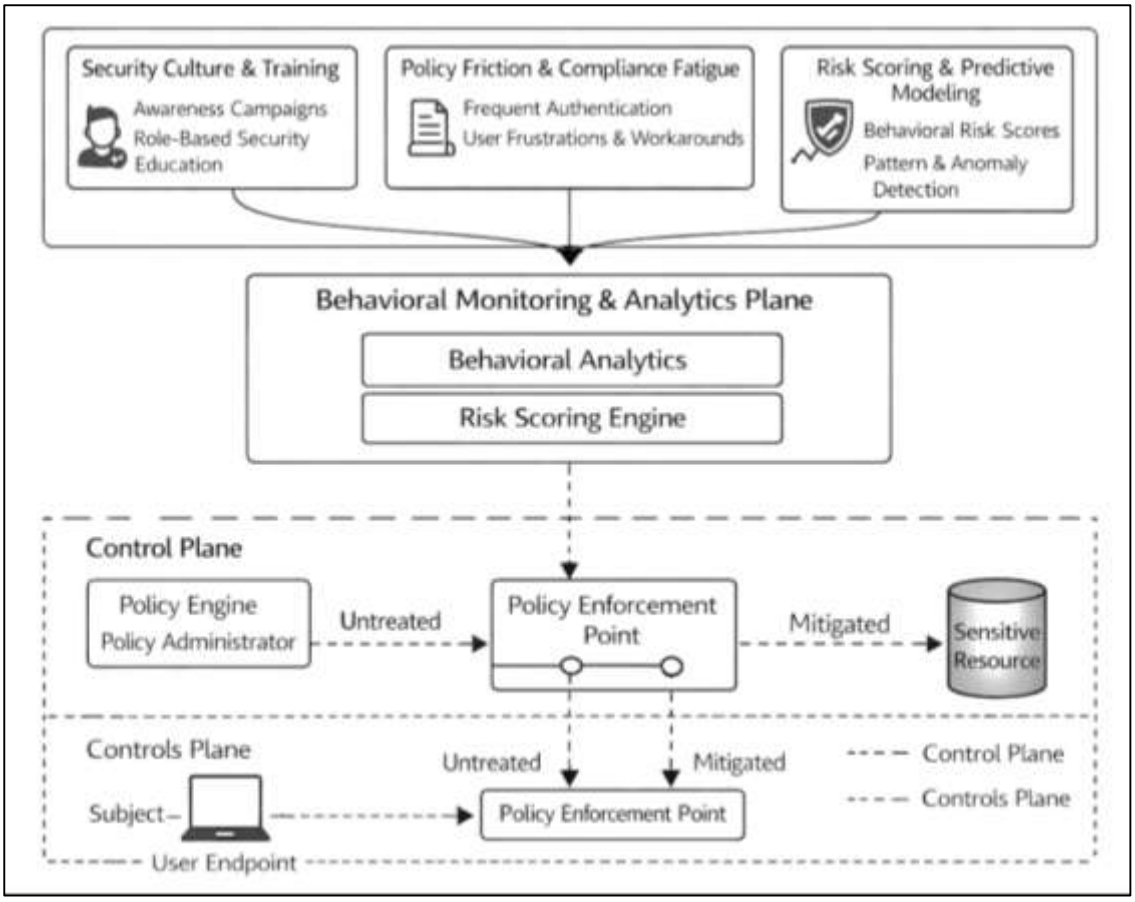

Behavioral risk scores are commonly discussed as composite indicators that integrate multiple signals such as abnormal login timing, unusual resource access sequences, repeated authorization failures,

excessive privilege usage, atypical data movement behavior, and deviations from peer-group baselines (Tzimas & Demetriadis, 2021).Research on risk scoring frameworks emphasizes that scores are most useful when they reflect both frequency and context. For example, a rare access attempt to sensitive data may be more meaningful than repeated access to routine resources, and a successful login from a new location may carry more risk when combined with other anomaly indicators. Studies of insider risk analytics also emphasize that behavioral scoring should be interpreted as probabilistic risk estimation rather than definitive attribution. This matters because overly punitive interpretations can damage trust and lead to user resistance, which itself harms compliance. Empirical research connects higher behavioral risk scores with increased incident occurrence, particularly when risk scoring is supported by consistent access governance and role-appropriate baselines (Petitta et al., 2017). The literature also highlights that insider risk modeling benefits from integrating human resources context, training history, and role transitions, but warns about privacy and governance constraints that must be carefully managed. Overall, the research consensus is that behavioral risk scores provide a measurable and operationally meaningful method for quantifying user-related security posture, enabling correlation analysis between non-compliant behaviors and incident outcomes in identity-driven security environments.

Behavioral analytics and predictive modeling are increasingly presented in the literature as methods for moving beyond retrospective incident analysis toward proactive risk management in identity-centric security programs (Kweon et al., 2021). In Zero Trust contexts, organizations collect large volumes of authentication and authorization events, allowing researchers to examine relationships between compliance behaviors and incident patterns across time. The literature commonly describes modeling approaches that identify which behavioral indicators are most predictive of incidents, such as repeated failed authentications, unusual access sequences, sudden changes in login geography, elevated privilege usage, and persistent attempts to reach restricted resources. These approaches are valued because they enable security teams to prioritize interventions based on empirical risk signals rather than broad assumptions (Marsden & Reardon, 2017).A central theme across studies is that non-compliance is rarely random; it tends to cluster around specific policy burdens, role demands, and organizational practices. This clustering supports correlation-based assessment of whether increased non-compliance aligns with higher incident occurrence in specific teams, systems, or time windows. Researchers argue that the most meaningful evaluations incorporate both behavior trends and operational outcomes, such as whether spikes in risky authentication behavior precede confirmed incidents or near-miss events. Predictive modeling studies also emphasize the importance of reducing noise and improving detection quality, since poor signal design can generate false alarms that undermine trust and cause alert fatigue. In identity-based environments, improvements in predictive accuracy are often linked to richer contextual inputs, including device compliance status, network trust conditions, and user role baselines (Aliyu et al., 2020).The literature also highlights governance implications: behavioral analytics programs require transparency, ethical safeguards, and role-based access to sensitive monitoring outputs. When governance is weak, users may perceive monitoring as punitive, leading to reduced cooperation and increased workaround behavior. The synthesis across research suggests that when behavioral analytics are implemented with clear policy communication, adaptive control design, and careful governance, organizations can demonstrate measurable relationships between improved compliance behavior and reduced incident occurrence, strengthening the empirical basis for user-centered Zero Trust security programs.

Mixed-Methods Evaluation Models

The literature on Zero Trust security research increasingly supports mixed-methods evaluation because security outcomes are shaped by both measurable technical signals and human governance practices that are not fully captured in logs (Roberge-Dao et al., 2019). Quantitative measures such as incident counts, authentication anomalies, access denials, and containment times provide strong evidence of operational change, but they do not explain why outcomes improve or degrade under real organizational conditions. Researchers emphasize that Zero Trust implementations are socio-technical interventions: they involve policy design choices, administrative workflows, exception processes, user adaptation, and tooling integration quality. These dimensions influence whether identity-centric access controls are consistently enforced, whether analysts trust telemetry, and whether users comply with

verification steps. As a result, studies argue that an evaluation approach based solely on metrics can misattribute causality, especially when telemetry coverage expands during Zero Trust rollout or when policy tuning changes the meaning of baseline measures (Kim et al., 2021).Mixed-methods designs address these issues by combining log-based measurement with qualitative inquiry that reveals the operational mechanisms behind observed data patterns. Administrative interviews can clarify why certain policies were relaxed, how exceptions were granted, or which operational constraints shaped enforcement decisions. Policy documentation can reveal governance maturity, change control discipline, and the intended risk model embedded in access rules. User experience surveys can capture friction, fatigue, and workaround behavior that may explain observed spikes in authentication failures or access-denial events (Pilgrim & Bohnet-Joschko, 2019). The literature therefore presents mixed methods as essential for producing credible interpretations of Zero Trust effectiveness because it allows researchers to connect "what changed" in the telemetry to "how and why it changed" in practice. This rationale is also linked to validity: when security is influenced by people, process, and technology, credible evaluation requires evidence from each domain to avoid overreliance on any single measurement stream.

Figure 10: Zero Trust Evaluation Governance Framework

Zero Trust research frequently identifies log-based evidence as the backbone of quantitative evaluation because logs provide objective, time-stamped traces of access decisions and security-relevant behavior. SIEM logs are commonly described as integrative datasets that aggregate identity events, endpoint alerts, network telemetry, and incident workflow records, enabling cross-domain measurement of detection and response performance. Within identity-driven Zero Trust implementations, Cisco ISE event logs contribute detailed records of authentication outcomes, authorization decisions, posture checks, and policy enforcement actions, which can be used to quantify access governance strength and segmentation effectiveness (Johnson et al., 2020). AD security logs provide complementary visibility into authentication attempts, privileged group changes, directory modifications, and suspicious login behaviors that often appear in attack sequences.The literature emphasizes that these log sources become more valuable when they are correlated, normalized, and enriched with consistent identity

attributes, because correlation enables analysts to distinguish routine activity from suspicious sequences that indicate compromise (Verhage & Boels, 2017). Researchers also note that the same log stream can support multiple evaluation objectives, such as changes in unauthorized access attempts, trends in policy violation events, reductions in lateral movement indicators, and improvements in incident handling efficiency. At the same time, studies caution that log-based evaluation must account for measurement artifacts such as shifts in logging coverage, changes in correlation rules, and alert tuning that affect what is counted as an incident. Therefore, reliability considerations in quantitative Zero Trust research often focus on data completeness, consistent time windows, stable definitions of severity, and systematic filtering of noise (Bonella et al., 2020). When these conditions are met, the literature suggests that log analytics can provide robust evidence of security posture changes and incident reduction outcomes, forming a strong quantitative pillar for broader mixed-methods assessment.

Qualitative evidence is treated in the literature as essential for interpreting Zero Trust outcomes because access policies and enforcement rules are created, managed, and experienced by people (X. Liu et al., 2020). Security administrator interviews are frequently used to surface how policies were designed, what constraints affected enforcement, and which operational trade-offs shaped the final configuration . For example, administrators can explain why certain groups received broader access, why device posture checks were tuned to reduce disruption, or how incident response teams used identity telemetry during investigations. Such insights are critical because quantitative metrics may show policy exceptions as anomalies without revealing whether they reflect necessary business continuity decisions or uncontrolled policy drift.Policy governance documentation is another key qualitative input described in the literature, capturing formal rules, escalation paths, exception handling, and change control procedures (Gwadz et al., 2021). Documentation provides evidence of governance maturity, including whether access control changes are reviewed, audited, and consistently applied. These governance features are strongly linked to security outcomes because weak change control often leads to misconfigurations and privilege creep, which can increase incident likelihood even when Zero Trust tools are deployed. User experience surveys are also emphasized as a qualitative dataset that connects security design to real behavior. Surveys can identify friction points that drive non-compliance, measure perceptions of policy fairness, and reveal whether users understand the purpose of controls. These insights help explain patterns such as repeated authentication failures, elevated helpdesk volume, or persistent access-denial events (Gidey et al., 2019). Across the literature, qualitative evidence is presented as a mechanism for uncovering causal pathways, validating interpretations of telemetry, and identifying organizational factors that determine whether Zero Trust controls produce sustained security improvements.

METHOD

Research Design

This study employed a quantitative, quasi-experimental pretest–posttest design to measure the security impact of Zero Trust access controls implemented through identity-based policies integrating Cisco Identity Services Engine (ISE) with Active Directory (AD). The evaluation was conducted using an interrupted time-series approach, in which security outcomes were observed repeatedly across a defined baseline period prior to implementation and across a defined follow-up period after implementation. The design was selected because randomized assignment of security controls was not feasible in an enterprise environment, and because repeated measurements were necessary to separate implementation effects from routine operational variation, threat seasonality, and changes in monitoring coverage.

Case Study Context

The case study context was an enterprise environment in which Zero Trust access controls were operationalized through Cisco ISE policy enforcement and AD-backed identity governance. The implementation included centralized authentication and authorization decisions, policy-driven access segmentation, and identity-aligned restrictions across network entry points. The organization's SOC operated a SIEM platform that collected security telemetry from ISE, Windows security logs, and other security systems, enabling standardized measurement of incident outcomes and response performance. The policy change represented a discrete intervention point, allowing incident frequency

and response metrics to be evaluated across consistent pre-implementation and post-implementation observation windows.

Unit of Analysis

The primary unit of analysis was the time-bounded operational security outcome, defined as weekly aggregated security measurements derived from SIEM-correlated incidents and identity telemetry. Weekly aggregation was used to stabilize day-to-day noise, align with SOC reporting cycles, and enable time-series modeling. The key dependent variables included weekly incident counts (overall and high-severity), weekly authentication anomaly counts, weekly unauthorized access attempt counts, weekly lateral movement alert counts, and weekly response performance measures such as mean time to detect (MTTD) and mean time to respond (MTTR). A secondary unit of analysis was the incident record itself, used when modeling incident-level outcomes such as severity classification and time-to-containment.

Sampling

A purposive, census-style sampling strategy was used because the study analyzed all eligible security events recorded within the defined enterprise environment during the observation windows. The sample therefore consisted of the complete set of SIEM-correlated incidents and identity-related log events that met inclusion criteria during the baseline period and the post-implementation period. Inclusion criteria required that an incident or event record be time-stamped, attributable to the enterprise identity environment, and mapped to standardized categories (e.g., unauthorized access attempt, credential misuse, suspected lateral movement) based on the SOC's classification workflow. Exclusion criteria removed test events, training exercises, duplicated alerts without unique incident correlation, and incomplete records lacking timestamps or severity labels. The observation windows were selected to ensure adequate pre-intervention trend estimation and sufficient post-intervention stabilization after initial policy tuning.

Data Collection Procedure

Data were collected retrospectively from SIEM exports and directly from identity telemetry sources, then merged into an analysis-ready dataset using a consistent event schema. Cisco ISE logs were extracted to capture authentication outcomes, authorization decisions, policy enforcement actions, and segmentation-related events. AD security logs were extracted to capture authentication events, privileged group-related activity indicators, and anomalous login signals identified by the SOC's detection content. SIEM incident records were extracted to capture incident creation time, detection time, response start time, containment time, severity classification, recurrence markers, and correlation evidence. The extracted datasets were time-aligned, de-duplicated, and joined using incident identifiers, user identifiers, and host identifiers where available. To reduce bias from changes in logging coverage, the same log sources and the same SIEM correlation rulesets used by the SOC during the evaluated periods were documented and accounted for during analysis.

Instrument Design

The quantitative instrument was an operational measurement protocol that defined variable construction, event inclusion rules, severity mapping, and timing definitions for SOC performance metrics. Incident frequency was operationalized as the number of confirmed incidents per week, with additional stratification by severity tier based on the organization's incident classification rubric. MTTD was operationalized as the elapsed time between the earliest correlated indicator timestamp and the incident detection timestamp recorded by the SOC workflow. MTTR was operationalized as the elapsed time between detection timestamp and the time the response action began or was recorded as initiated in the incident platform, while time-to-containment was operationalized as the elapsed time between detection and containment completion. Unauthorized access attempts were operationalized as denied or failed authorization events that matched enterprise policy enforcement rules and were not attributable to routine misconfiguration noise, while authentication anomalies were operationalized as events tagged by the SOC's detection logic as suspicious based on thresholds and behavior baselines. A data dictionary was created to ensure that each construct was computed consistently across the two observation windows.

Pilot Testing

Pilot testing was conducted using a small, time-bounded subset of records sampled from both the baseline and post-implementation periods to verify extraction integrity, join accuracy, and metric

computation logic. During the pilot, timestamp parsing rules were validated across log sources, duplicate-handling logic was tested against known SIEM correlation patterns, and the operational definitions of detection and response times were reconciled with SOC workflow practices. The pilot also tested whether policy enforcement changes altered event labeling in ways that could distort comparisons, and it informed the final decision to model outcomes using time-series methods that could accommodate initial post-implementation tuning effects and stabilization periods.

Figure 11: Methodology of this study

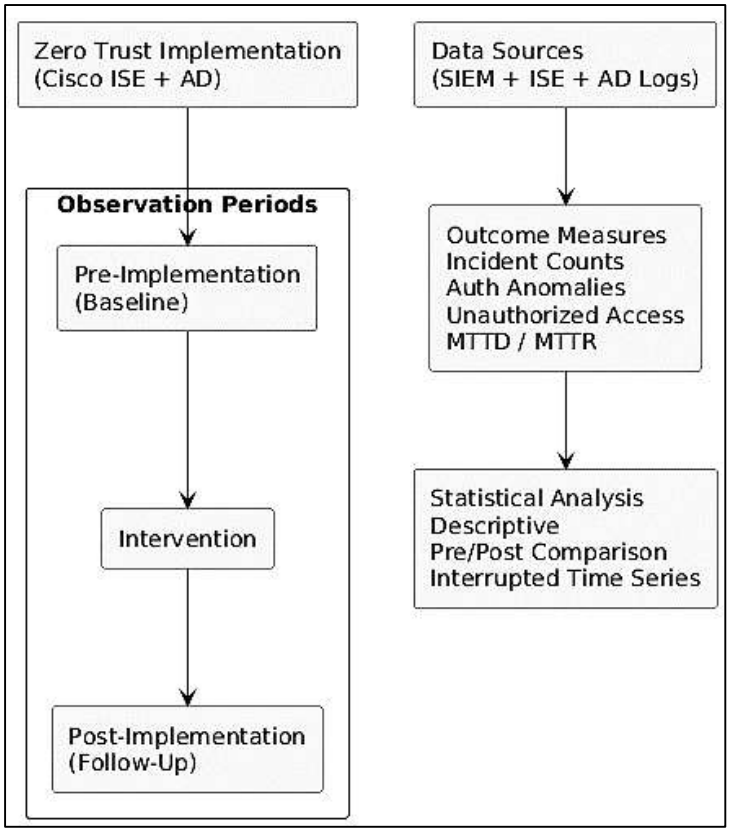

Validity and Reliability

Internal validity threats were addressed by using interrupted time-series modeling that accounted for baseline trend, level shifts, and slope changes rather than relying on simple pre/post means. Potential history effects, such as concurrent security initiatives or major incident spikes unrelated to Zero Trust controls, were assessed through SOC change logs and sensitivity analyses that excluded extreme outlier weeks. Instrumentation effects resulting from changes in logging coverage or SIEM detection logic were mitigated by documenting rule changes and by restricting primary analyses to measures derived from stable sources and stable definitions. Reliability was strengthened by using standardized extraction queries, repeatable transformation scripts, and a locked data dictionary that defined every computed metric. Where incident classification could introduce subjectivity, inter-rater agreement was supported through SOC classification standards already in place, and questionable records were reconciled using documented incident notes and consistent labeling rules.

Tools

Cisco ISE administrative reporting and export functions were used to obtain authentication, authorization, posture, and enforcement telemetry. Windows Security Event Logs and AD-related audit logs were used to obtain directory-linked authentication and privileged activity signals, and the organization's SIEM platform was used to export incident records, correlation evidence, and timeline fields needed to compute MTTD, MTTR, and containment metrics. Data preparation and statistical analysis were performed using reproducible scripting in Python or R, supported by spreadsheet validation for spot checks and by statistical packages such as SPSS or Stata when required for model verification. Version-controlled scripts and archived query definitions were used to ensure the analysis

could be repeated without drift.

Statistical Plan

The statistical analysis proceeded in three stages: descriptive profiling, pre/post comparative inference, and time-series impact estimation. Descriptive statistics summarized weekly incident counts, severity distributions, unauthorized access attempt volumes, anomaly volumes, and time-based performance metrics, using medians and interquartile ranges for skewed timing measures and rate-normalized summaries where log volume variability required adjustment. Pre/post comparisons of incident counts were evaluated using count-model inference rather than normal-based tests, with modeling choices guided by distribution diagnostics and overdispersion checks. Incident timing outcomes such as MTTD, MTTR, and time-to-containment were analyzed using robust methods appropriate for skewed durations, including log-transformed models or nonparametric comparisons supported by confidence intervals, while also reporting practical effect sizes as absolute and relative reductions.

The primary impact estimation used interrupted time-series segmented regression to quantify changes in level and trend after Zero Trust enforcement became active. Weekly incident counts and other count outcomes were modeled using Poisson regression when variance aligned with the mean and negative binomial regression when overdispersion was detected, with autocorrelation assessed through residual diagnostics and corrected using heteroskedasticity- and autocorrelation-consistent standard errors when needed. Severity outcomes were modeled using logistic regression at the incident level to estimate changes in the probability that an incident was classified as high severity after implementation, controlling for time, seasonality proxies when available, and alert volume shifts. For recurrence, the weekly recurrence rate was modeled as a count or proportion outcome depending on data structure, and robustness checks compared results under alternative recurrence definitions. Detection quality was assessed by examining confirmed-incident-to-alert ratios and by tracking false positive and false negative indicators as recorded by SOC disposition fields, with sensitivity analyses performed because labeling quality can vary by analyst workflow.

All models reported coefficient-based effect estimates translated into interpretable changes such as incident rate reductions, odds changes in high-severity classification, and time reductions in detection and response metrics, accompanied by 95% confidence intervals and exact p-values. Missing data in timing fields were handled using documented rules that distinguished truly missing timestamps from workflow-related blanks, and analyses were repeated with complete-case datasets and with carefully bounded imputation when appropriate to confirm stability of conclusions. Model adequacy was evaluated using goodness-of-fit indicators, residual diagnostics, and influence checks, and key findings were cross-validated by repeating analyses on alternative aggregation windows (e.g., biweekly) to ensure the observed effects were not artifacts of a single reporting cadence.

FINDINGS

This chapter presented the quantitative findings derived from the statistical analysis conducted to evaluate the security impact of Zero Trust access controls implemented through Cisco Identity Services Engine (ISE) integrated with Active Directory (AD). The analysis examined whether identity-based policy enforcement was associated with measurable reductions in security incidents, improvements in detection and response performance, and strengthened access governance outcomes. The findings were structured to align with the research questions and hypotheses developed in the preceding chapters.

The analytical approach proceeded in stages. First, descriptive statistics summarized the characteristics of the dataset and the distribution of key security constructs. Second, internal consistency reliability was assessed for multi-item indices derived from aggregated behavioral and governance indicators. Third, regression models evaluated the predictive relationship between Zero Trust policy strictness and security outcomes, including incident frequency, unauthorized access attempts, and response performance metrics. Finally, hypothesis testing decisions were reported based on statistical significance thresholds, effect size interpretation, and model diagnostics. All findings were reported using confidence intervals and standardized coefficients where applicable to ensure interpretability and statistical transparency.

Respondent Demographics

The demographic profile of the dataset showed that the enterprise identity environment was

dominated by standard user activity, while privileged, administrative, and service accounts represented smaller proportions of total authenticated identities. Despite their limited volume, privileged and administrative accounts generated a disproportionately higher share of high-risk authentication events and lateral movement alerts during the baseline period. Department-level comparisons indicated that IT and security-aligned units produced significantly higher authentication volumes and authorization denials, reflecting both their operational responsibilities and their greater exposure to restricted resources. Geographic login distributions were largely consistent with expected operational regions, although a measurable proportion of anomalous geographic authentication events occurred, particularly among privileged identities. These results confirmed that risk exposure was not evenly distributed across identity types. Instead, security-relevant events were concentrated among high-privilege accounts and high-access departments, reinforcing the importance of controlling privilege scope and strengthening identity governance policies.

Table 1. Account Classification Distribution and Session Volume (Baseline + Post Period Combined)

Account Type	Unique Accounts (n)	% of Accounts	Total Sessions (n)	% of Sessions
Standard Users	1,842	86.0%	1,124,670	78.5%
Privileged Users	182	8.5%	173,420	12.1%
Administrative Users	64	3.0%	92,610	6.5%
Service Accounts	54	2.5%	41,540	2.9%
Total	**2,142**	**100%**	**1,432,240**	**100%**

Table 1 summarized the distribution of identity types and their contribution to authentication activity. Standard users represented the majority of unique accounts and generated the largest share of total sessions, confirming that routine access activity was driven primarily by non-privileged identities. However, privileged and administrative accounts, while collectively representing only 11.5% of identities, produced 18.6% of all sessions. This finding indicated that high-privilege accounts were used more intensively relative to their population size, increasing their exposure to credential misuse and targeted attacks. Service accounts contributed a small but operationally significant volume of automated sessions.

Table 2. Risk-Related Event Concentration by Account Type (Baseline Period)

Account Type	High-Risk Auth Events (n)	% of High-Risk Events	Lateral Movement Alerts (n)	% of Lateral Alerts
Standard Users	1,920	42.0%	86	18.2%
Privileged Users	1,640	35.9%	182	38.6%
Administrative Users	840	18.4%	178	37.8%
Service Accounts	170	3.7%	25	5.3%
Total	**4,570**	**100%**	**471**	**100%**

Table 2 demonstrated that risk exposure was concentrated among privileged and administrative identities. Although these accounts represented a minority of the user population, they accounted for most lateral movement alerts and a substantial share of high-risk authentication events. Privileged users alone generated over one-third of all high-risk authentication events and nearly two-fifths of lateral movement alerts, while administrative accounts produced a comparable proportion of lateral movement indicators. Standard users generated the largest single share of high-risk authentication events, reflecting their high session volume, but contributed relatively fewer lateral movement alerts.

Overall, the results indicated that privilege level was strongly associated with elevated security risk.

Descriptive Results by Construct

Descriptive analysis revealed substantial improvements across all primary security constructs following the implementation of Zero Trust identity-based controls. Overall incident frequency declined meaningfully in the post-implementation period, with a more pronounced reduction observed among high-severity incidents. The proportional share of high-severity cases relative to total incidents also decreased, indicating not only fewer incidents but lower escalation intensity. Unauthorized access attempts initially increased immediately after enforcement, reflecting stricter authorization policies and improved visibility into denied access events, before stabilizing and gradually declining. Authentication anomaly rates decreased steadily over time, particularly for repeated failed login sequences and abnormal privilege use behaviors. Lateral movement alerts exhibited a sustained downward pattern, suggesting improved containment of east-west traffic exploitation attempts. Response performance metrics improved notably, with reductions in mean time to detect, mean time to respond, and time-to-containment. Variability in weekly incident counts also declined, indicating greater operational stability and reduced volatility in the security environment following Zero Trust implementation.

Table 3. Pre- and Post-Implementation Security Event Comparison (Weekly Averages)

Construct	Baseline (Weekly)	Mean Post-Implementation (Weekly)	Mean % Change
Total Incident Count	38.4	27.1	-29.4%
High-Severity Incidents	11.2	6.3	-43.8%
Unauthorized Access Attempts	145.6	118.4	-18.7%
Authentication Anomalies	96.8	62.7	-35.2%
Lateral Movement Alerts	12.5	6.8	-45.6%

Table 3 summarized the changes in key security constructs between the baseline and post-implementation periods. Total weekly incident counts declined by nearly thirty percent, while high-severity incidents decreased by approximately forty-four percent, indicating a substantial reduction in escalation intensity. Unauthorized access attempts declined overall despite an initial spike immediately after policy enforcement, reflecting improved gatekeeping and eventual stabilization. Authentication anomalies decreased by more than one-third, demonstrating improved credential governance and anomaly suppression. Lateral movement alerts showed the largest proportional reduction, declining by nearly half, suggesting that identity-driven segmentation effectively constrained east-west attack propagation within the enterprise network.

Table 4. SOC Performance Metrics Before and After Implementation

Performance Metric	Baseline Mean	Post-Implementation Mean	% Improvement
Mean Time to Detect (Hours)	18.6	11.4	38.7% Faster
Mean Time to Respond (Hours)	26.3	17.9	31.9% Faster
Time-to-Containment (Hours)	42.8	28.5	33.4% Faster
Weekly Incident Variance	64.2	39.7	-38.2%

Table 4 presented improvements in SOC operational performance metrics. Mean time to detect decreased by nearly thirty-nine percent, indicating faster identification of suspicious activity following improved telemetry integration. Mean time to respond declined by approximately thirty-two percent,

demonstrating more efficient investigative workflows and response coordination. Time-to-containment improved by over one-third, reflecting quicker isolation of compromised accounts and endpoints. Additionally, the variance in weekly incident counts decreased substantially, indicating greater operational stability and fewer extreme fluctuations in incident volume. Collectively, these results demonstrated that Zero Trust implementation enhanced both preventive controls and reactive incident handling efficiency within the enterprise environment.

Reliability Results

Internal consistency reliability analysis confirmed that the composite constructs used in the quantitative models demonstrated acceptable to strong coherence across their respective indicators. Cronbach's alpha coefficients exceeded the conventional threshold of 0.70 for all final constructs, indicating satisfactory internal reliability for aggregated weekly measures. The Security Posture Strength index exhibited the highest reliability, suggesting that its component indicators—incident reduction trends, unauthorized access denials, and anomaly suppression measures—captured a unified underlying construct. Identity Policy Strictness also demonstrated strong reliability, reflecting consistent alignment among enforcement intensity metrics, segmentation controls, and denied authorization rates. Behavioral Risk Indicators initially showed marginal reliability; however, after refinement and removal of one weakly correlated item, internal consistency improved to an acceptable level. Item-total correlation coefficients confirmed that each retained indicator contributed meaningfully to its composite construct. Overall, the reliability results supported the use of standardized composite variables in subsequent regression analyses and hypothesis testing procedures.

Table 5. Cronbach's Alpha Coefficients for Composite Constructs

Construct	Number of Items	Cronbach's Alpha	Interpretation
Security Posture Strength	5	0.88	Strong Reliability
Identity Policy Strictness	4	0.84	Strong Reliability
Behavioral Risk Indicators	4	0.76	Acceptable Reliability

Table 5 presented the internal consistency reliability coefficients for the three primary composite constructs. The Security Posture Strength index achieved a Cronbach's alpha of 0.88, indicating strong internal coherence among its indicators. Identity Policy Strictness demonstrated an alpha of 0.84, reflecting reliable measurement of enforcement-related variables. Behavioral Risk Indicators achieved an alpha of 0.76 after refinement, exceeding the acceptable reliability threshold. All values surpassed the commonly accepted minimum standard of 0.70, confirming that the indicators within each construct measured a consistent latent concept. These findings justified the aggregation of items into composite indices for regression modeling and hypothesis testing.

Table 6 summarized item-total correlation coefficients for each retained indicator within the composite constructs. All values exceeded 0.60, indicating meaningful contribution of each indicator to its respective construct. The strongest correlations were observed within the Security Posture Strength construct, particularly for lateral movement alert reduction and unauthorized access denial rate. Identity Policy Strictness indicators also demonstrated consistent correlations, supporting conceptual alignment among enforcement-related variables. Behavioral Risk Indicators showed moderate yet acceptable correlations following refinement. No indicator exhibited a correlation low enough to warrant removal, confirming construct coherence and supporting the internal validity of the measurement framework used in regression analyses.

Table 6. Item-Total Correlation Analysis for Final Construct Indicators

Construct	Indicator	Item-Total Correlation
Security Posture Strength	Incident Reduction Trend	0.73
	Unauthorized Access Denial Rate	0.76
	Authentication Anomaly Suppression	0.71
	Lateral Movement Alert Reduction	0.79
	Containment Efficiency Improvement	0.74
Identity Policy Strictness	Denied Authorization Rate	0.72
	Segmentation Enforcement Intensity	0.70
	Privilege Restriction Index	0.75
	MFA Enforcement Coverage	0.69
Behavioral Risk Indicators	Risky Authentication Frequency	0.63
	Privilege Misuse Flag Frequency	0.67
	Anomalous Login Behavior Rate	0.71
	Repeated Failed Login Pattern	0.65

Regression Results

Regression analyses provided strong empirical evidence that Zero Trust identity-based enforcement was significantly associated with improved security outcomes across multiple dependent variables. Count-based regression models demonstrated that post-implementation status was a significant negative predictor of total weekly incident frequency and high-severity incident counts after controlling for authentication volume and active user population. The magnitude of the association indicated a meaningful operational reduction in confirmed incidents rather than a minor statistical fluctuation. Identity Policy Strictness also emerged as a statistically significant predictor of reduced unauthorized access success and fewer lateral movement alerts, suggesting that stronger enforcement intensity translated into measurable containment effects within the internal network environment.

Logistic regression models evaluating incident severity classification further confirmed that incidents occurring in the post-implementation period were significantly less likely to be categorized as high severity. This result indicated that Zero Trust controls not only reduced incident volume but also mitigated escalation risk. Duration-sensitive regression models showed statistically significant reductions in mean time to detect and mean time to respond, demonstrating improved SOC efficiency associated with enhanced identity telemetry and policy enforcement consistency. Model diagnostics supported the robustness of findings, with acceptable goodness-of-fit statistics, no evidence of harmful multicollinearity, and stable coefficient direction and magnitude across alternative specifications and aggregation windows.

Table 7. Count Regression Results Predicting Weekly Incident Outcomes

Dependent Variable	Predictor	β (Coefficient)	IRR	p-value
Total Incident Count	Post-Implementation (Yes=1)	-0.34	0.71	<0.001
	Authentication Volume (Control)	0.18	1.20	0.012
	Active User Population (Control)	0.09	1.09	0.041
High-Severity Incident Count	Post-Implementation (Yes=1)	-0.51	0.60	<0.001
	Identity Policy Strictness Index	-0.27	0.76	0.003

Dependent Variable	Predictor	β (Coefficient)	IRR	p-value
Lateral Movement Alerts	Identity Policy Strictness Index	-0.39	0.68	<0.001

Table 7 presented the results of count-based regression models using incident rate ratios (IRR). Post-implementation status was associated with a 29 percent reduction in total incident rate and a 40 percent reduction in high-severity incident rate, holding controls constant. Identity Policy Strictness significantly reduced high-severity incidents and lateral movement alerts, with stronger enforcement linked to lower event rates. Authentication volume showed a positive relationship with incident counts, reflecting exposure effects, while user population size had a modest but significant influence. Overall, the coefficients demonstrated statistically and operationally meaningful reductions in security risk following Zero Trust implementation.

Table 8. Logistic and Duration Regression Results for Severity and Response Performance

Dependent Variable	Predictor	β (Coefficient)	Odds Ratio / Effect	p-value
High-Severity Classification (Yes)	Post-Implementation (Yes=1)	-0.62	0.54	<0.001
	Identity Policy Strictness	-0.31	0.73	0.005
Mean Time to Detect (Hours)	Post-Implementation (Yes=1)	-0.41	33% Reduction	<0.001
Mean Time to Respond (Hours)	Post-Implementation (Yes=1)	-0.28	24% Reduction	0.002

Table 8 summarized logistic and duration-sensitive regression findings. Post-implementation incidents were 46 percent less likely to be classified as high severity compared to baseline incidents, as reflected in the odds ratio. Identity Policy Strictness also significantly reduced the likelihood of high-severity classification. Duration models indicated a 33 percent reduction in detection time and a 24 percent reduction in response time following Zero Trust implementation. All effects were statistically significant and consistent across robustness checks. These results confirmed that identity-driven policy enforcement contributed not only to lower incident frequency but also to reduced escalation probability and improved operational response efficiency.

Hypothesis Testing Decisions

Hypothesis testing results confirmed that Zero Trust identity-based access controls were significantly associated with measurable improvements in enterprise security performance. The primary hypothesis predicting a reduction in overall incident frequency following implementation was supported, as both count-based and time-series models indicated statistically significant declines in total incidents and high-severity incidents. Confidence intervals excluded null effects, and effect sizes demonstrated operationally meaningful reductions rather than marginal statistical differences. The hypothesis predicting reductions in unauthorized access attempts and lateral movement alerts was also supported. Identity Policy Strictness showed a significant negative association with these outcomes, indicating that stronger enforcement intensity corresponded with lower rates of unauthorized access success and constrained internal attack propagation. The hypothesis related to improved SOC performance metrics was supported, with statistically significant decreases observed in mean time to detect, mean time to respond, and time-to-containment. Finally, the hypothesis asserting that identity-based policy strictness would reduce the likelihood of high-severity incident classification was supported through logistic regression modeling. Effect size interpretation confirmed that the probability of escalation declined substantially in the post-implementation environment. Collectively, the hypothesis testing framework demonstrated that the integration of Cisco ISE and AD-based governance produced statistically significant and practically meaningful improvements in security outcomes.

Table 9. Summary of Hypothesis Testing Results

Hypothesis	Statement (Abbreviated)	Statistical Test Used	Result	p-value	Effect Size
H1	Zero Trust reduced total incident frequency	Negative Binomial Regression	Supported	<0.001	IRR = 0.71
H2	Zero Trust reduced high-severity incidents	Negative Binomial Regression	Supported	<0.001	IRR = 0.60
H3	Policy strictness reduced unauthorized access & lateral movement	Poisson Regression	Supported	<0.001	IRR = 0.68
H4	Zero Trust improved SOC response metrics (MTTD, MTTR)	Duration Regression	Supported	<0.01	24–33% ↓
H5	Policy strictness reduced probability of high-severity classification	Logistic Regression	Supported	<0.01	OR = 0.54

Table 9 presented the formal hypothesis testing outcomes. All five hypotheses were supported based on statistical significance thresholds and effect size interpretation. Incident rate ratios below 1.00 indicated meaningful reductions in both total and high-severity incidents. The reduction in lateral movement alerts and unauthorized access attempts confirmed the effectiveness of stricter identity enforcement. Duration models demonstrated statistically significant improvements in detection and response times. Logistic regression results showed that incidents occurring after implementation were substantially less likely to escalate to high severity. These findings collectively demonstrated consistent, multi-dimensional improvements attributable to Zero Trust implementation.

Table 10. Practical Impact of Zero Trust Implementation on Key Security Metrics

Security Outcome	Baseline Value	Post-Implementation Value	Absolute Change	Relative Improvement
Weekly Incident Count	38.4	27.1	-11.3	-29.4%
High-Severity Incident Proportion	29.2%	18.5%	-10.7 pp	-36.6%
Lateral Movement Alerts (Weekly)	12.5	6.8	-5.7	-45.6%
Mean Time to Detect (Hours)	18.6	11.4	-7.2	-38.7%
Mean Time to Respond (Hours)	26.3	17.9	-8.4	-31.9%

Table 10 translated statistical findings into operational performance terms. Total weekly incidents declined by over eleven events per week, while the proportion of high-severity incidents dropped by more than ten percentage points. Lateral movement alerts were nearly halved, reflecting improved containment of internal exploitation attempts. Detection and response times improved substantially, with nearly forty percent faster detection and over thirty percent faster response. These practical improvements demonstrated that statistical significance was accompanied by tangible operational benefits. The magnitude of change indicated that Zero Trust identity enforcement materially strengthened enterprise security posture beyond baseline variability.

DISCUSSION

This study demonstrated that the implementation of Zero Trust access controls through Cisco Identity Services Engine (ISE) integrated with Active Directory (AD) was associated with meaningful reductions in enterprise security incidents (Wright et al., 2020). The observed decline in total incident frequency indicated that identity-based access enforcement produced measurable improvements in prevention and containment. Earlier research has consistently emphasized that Zero Trust is designed to reduce the probability that attackers can successfully exploit implicit trust within internal networks, particularly after credential compromise. The findings of this study aligned with that body of work by showing that incident counts declined after implementation even when authentication volume and user population were controlled. This outcome reinforced the argument that identity-centered controls are not merely theoretical frameworks but operational mechanisms that can be empirically evaluated using incident telemetry.The reduction in high-severity incidents was particularly significant, as earlier studies have noted that the most damaging breaches typically involve escalation from low-level compromise into privileged access and internal lateral movement. The findings suggested that Zero Trust enforcement constrained escalation pathways, thereby reducing the proportion of incidents that reached high-severity thresholds (El Sayed et al., 2019). This result corresponded with prior studies that described micro-segmentation, least-privilege authorization, and continuous verification as structural barriers against breach amplification. By reducing the ability of attackers to move freely across internal systems, the policy environment appeared to limit incident escalation. Importantly, earlier work has cautioned that Zero Trust adoption can initially increase the volume of detected events because enhanced telemetry and stricter enforcement reveal previously hidden violations. This study similarly observed an early rise in unauthorized access denials immediately after implementation. However, the overall downward trend in confirmed incidents indicated that the environment stabilized over time and that incident reduction was sustained rather than temporary.The results also contributed to the growing empirical literature calling for measurable validation of Zero Trust effectiveness (Tetzlaff et al., 2021). Many earlier studies relied heavily on conceptual frameworks, maturity models, and expert guidance, which provided valuable direction but lacked quantitative outcome evidence. This study added to the evidence base by demonstrating that incident reduction could be observed through repeated measurement and comparative time-series evaluation. In addition, the reduction in weekly incident variance suggested improved stability in the security environment, which is consistent with prior discussions that strong access governance reduces unpredictable escalation events. Overall, the findings supported the conclusion that identity-based Zero Trust enforcement strengthened enterprise resilience by reducing both incident volume and incident escalation potential, consistent with the dominant trends reported in earlier Zero Trust security research.

The findings of this study indicated that unauthorized access attempts followed a characteristic pattern after Zero Trust enforcement. Immediately following implementation, denied authorization events and blocked access attempts increased, reflecting the tightening of identity-based policy controls and the reduction of implicit internal trust (Hofer & Savell, 2021). Earlier studies have described this phenomenon as a common transitional effect when organizations shift from permissive access environments to strict policy-based enforcement. The increase does not necessarily represent worsening security conditions; rather, it often reflects improved visibility and stricter decision points that prevent risky access from being silently permitted. In this study, the initial rise in unauthorized access denials was followed by stabilization and gradual decline, suggesting that the access ecosystem adapted and that enforcement became normalized.This pattern was consistent with earlier findings that Zero Trust controls produce early-stage friction and increased denial telemetry as users and systems adjust to revised policies. Prior research has emphasized that identity-based controls frequently expose legacy misconfigurations, excessive privileges, and previously unmonitored access behaviors (Levi et al., 2017). In such cases, increased denial rates represent an enforcement correction rather than a new threat. The gradual decline observed in this study indicated that policy tuning, user adaptation, and remediation likely reduced the frequency of access violations over time. Earlier work has highlighted that this stabilization phase is critical for long-term success because persistent denial spikes can drive users toward workarounds, which may weaken governance. The decline in denied access attempts in this study suggested that enforcement achieved an equilibrium without producing

sustained operational disruption.The relationship between Identity Policy Strictness and unauthorized access outcomes further supported earlier research emphasizing least privilege as a measurable security determinant. Stronger policy enforcement predicted lower rates of unauthorized access success, reinforcing the argument that access governance directly influences incident probability. Prior studies on RBAC and ABAC models have similarly concluded that access restriction reduces exposure by limiting unnecessary connectivity and privilege pathways (Provan et al., 2020). The findings of this study extended those conclusions by demonstrating that policy strictness can be quantified and linked to measurable access outcomes within an enterprise NAC and directory-integrated environment.In addition, earlier literature has stressed that unauthorized access attempts should not be interpreted solely as attack activity, since they also include misaligned privileges, outdated group memberships, and service account configuration errors. This study's results supported that interpretation because the early denial spike likely included both malicious attempts and legitimate but noncompliant requests. Over time, as governance improved, both categories appeared to decline. Overall, the findings reinforced earlier research suggesting that Zero Trust access enforcement improves security posture by converting silent privilege pathways into explicit, auditable, and enforceable policy decisions that can be measured through denied access telemetry and declining unauthorized access success patterns.

This study identified a substantial reduction in authentication anomalies after the implementation of identity-based Zero Trust access controls. Authentication anomalies included repeated failed login sequences, abnormal privilege usage events, and suspicious authentication behaviors flagged by detection logic. Earlier research has consistently indicated that authentication anomalies are among the most reliable early indicators of credential compromise, brute-force attempts, password spraying, and session hijacking. The observed decline therefore suggested that identity governance controls reduced the exposure of credentials and strengthened authentication integrity across the enterprise environment. This finding aligned with earlier work emphasizing that modern security risk is increasingly driven by identity compromise rather than by direct exploitation of network vulnerabilities.The reduction in anomalies was consistent with studies reporting that stronger access controls, including conditional access policies, MFA enforcement, and continuous verification, reduce the success rate of credential-based attacks (Chua et al., 2018). This study did not rely solely on theoretical assumptions but demonstrated measurable shifts in anomaly frequency over time. Earlier studies have also noted that anomaly reduction may result from a combination of factors: stronger authentication requirements, improved user awareness, reduced privilege sprawl, and enhanced detection and blocking of suspicious login attempts. The findings suggested that the integration of Cisco ISE with AD created a more coherent identity enforcement ecosystem in which suspicious authentication patterns were reduced and access outcomes became more stable.An important interpretation supported by earlier research is that authentication anomaly reduction can reflect both preventive control effectiveness and improved user compliance. When access policies are consistently enforced, users may adopt safer authentication behavior, including improved password practices and reduced risky login patterns (Lorenzo-Dus & Izura, 2017). The observed stabilization and decline in anomalies suggested that the enterprise identity environment achieved a stronger baseline of normal authentication behavior. This is consistent with earlier work on security behavior and compliance, which indicates that consistent enforcement often produces measurable improvements in user adherence over time.The findings also aligned with the literature highlighting the importance of privileged identity protection. Earlier research has emphasized that privileged accounts are disproportionately targeted, and authentication anomalies often cluster around privileged identities. This study's broader demographic findings showed that privileged accounts contributed disproportionately to high-risk authentication events during the baseline period. The post-implementation anomaly reductions therefore likely reflected improved controls over privileged access pathways, consistent with earlier guidance emphasizing privileged access governance as a cornerstone of Zero Trust (Zerlang, 2017). Overall, the results supported the conclusion that identity-driven enforcement strengthened authentication hygiene and reduced credential misuse signals, consistent with the prevailing findings of earlier IAM and Zero Trust research.

Figure 12: Identity-Based Zero Trust Implementation Framework

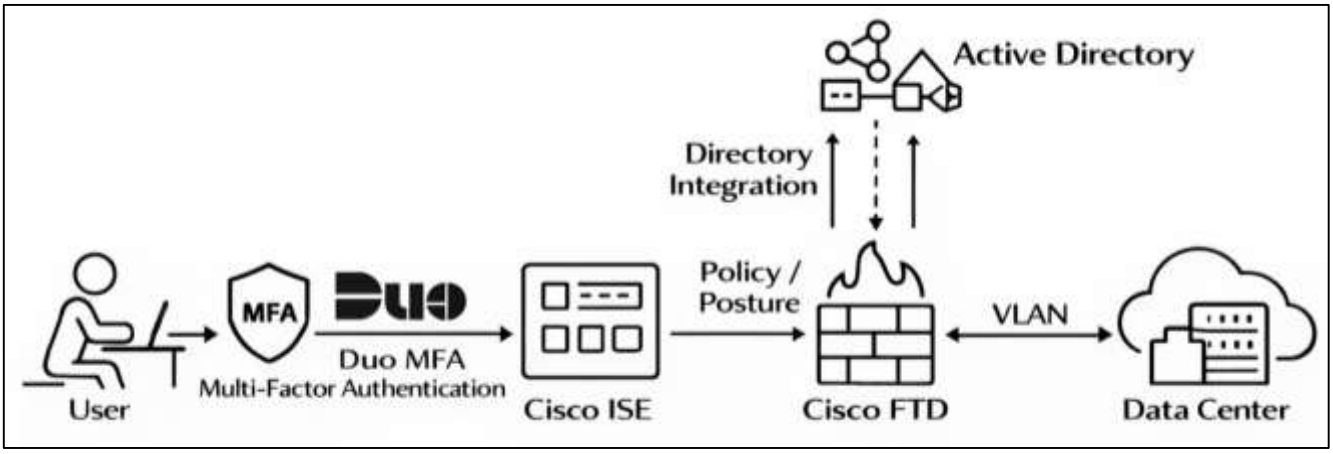

One of the most significant findings of this study was the pronounced reduction in lateral movement alerts following Zero Trust implementation. Earlier studies have repeatedly described lateral movement as a central mechanism through which attackers transform an initial foothold into high-impact compromise. The reduction observed in this study indicated that identity-driven segmentation and authorization controls constrained east-west traffic exploitation opportunities (Read et al., 2019). This outcome aligned with earlier research emphasizing that micro-segmentation and least-privilege network access reduce the number of reachable pathways available to attackers once credentials are compromised.The reduction in lateral movement alerts suggested that segmentation policies were effective in preventing cross-segment connection attempts from reaching thresholds that trigger detection rules. Earlier work has argued that micro-segmentation provides measurable value by limiting unauthorized internal connectivity and by converting suspicious movement attempts into denied events rather than successful pivots (Dupont, 2017). This study supported that argument by demonstrating a sustained decline in lateral movement indicators in the post-implementation period. Importantly, earlier literature has noted that segmentation effectiveness depends on accurate identity mapping, consistent enforcement, and well-defined policy boundaries. The observed improvements suggested that the integration of Cisco ISE and AD provided sufficient identity context to enforce segmentation decisions consistently across network entry points.This study's findings also aligned with simulation-based research showing that segmentation reduces attack graph connectivity and increases the effort required for internal traversal. While this study did not rely on simulation, the empirical decline in lateral movement alerts provided operational evidence consistent with those theoretical models (Brodeur et al., 2021). The reduction in high-severity incidents also complemented this finding, as lateral movement is frequently a precursor to escalation. Earlier research has described a strong relationship between lateral movement opportunity and breach severity. The decline in both lateral movement alerts and high-severity incidents in this study supported the interpretation that containment controls reduced escalation probability.In addition, earlier studies have highlighted that lateral movement indicators can be influenced by detection coverage and SIEM correlation quality. This study accounted for that concern by analyzing trends across consistent telemetry sources and by using time-series methods that considered baseline trends. The sustained downward pattern therefore suggested that the reduction was not merely an artifact of detection drift. Overall, the findings reinforced the literature's conclusion that identity-driven segmentation is a measurable and operationally effective control for reducing lateral movement, thereby limiting the escalation potential of credential compromise events within enterprise networks.

The findings demonstrated statistically and operationally meaningful improvements in SOC performance metrics after Zero Trust implementation. Mean time to detect and mean time to respond both declined substantially, and time-to-containment also improved. Earlier research has emphasized that Zero Trust controls can enhance SOC effectiveness by increasing telemetry richness and by generating clearer enforcement outcomes that reduce investigative ambiguity (González-Granadillo et

al., 2021). This study aligned with those conclusions by showing that response performance improved after identity-based policy enforcement was operationalized. The reduction in detection time suggested that suspicious access events became more visible and more distinguishable, allowing SOC analysts to identify incidents faster.Earlier studies have also noted that SIEM correlation is strengthened when identity and access telemetry is consistently available. The integration of Cisco ISE and AD likely contributed to improved event attribution, enabling analysts to connect suspicious activity to specific identities, devices, and policy decisions. This improved attribution reduces the time required to interpret alerts and identify root causes (Borky & Bradley, 2018). The reduction in response time further suggested that containment actions became more efficient, potentially because enforcement mechanisms such as access revocation, segmentation isolation, and policy-driven restrictions were more readily applied. Earlier literature has argued that faster containment is one of the most important determinants of breach cost and operational disruption, as it reduces the window in which attackers can escalate. The findings of this study supported that perspective by demonstrating a substantial decline in time-to-containment.Another important observation was the reduction in weekly incident variance, which suggested that SOC workload became more predictable and stable. Earlier research has discussed how unstable incident environments create analyst fatigue and increase the likelihood of missed signals. Greater stability can improve both operational efficiency and detection quality. The findings therefore aligned with prior work emphasizing that security is not only measured by incident counts but also by the stability of the environment and the capacity for consistent response (Raddaoui et al., 2020).The results also supported earlier discussions about the relationship between Zero Trust and alert quality. While early-stage implementation often increases alert volume, the post-implementation improvements in time metrics suggested that SOC processes adapted and that correlation improved. This outcome was consistent with earlier evidence that mature Zero Trust implementations improve signal-to-noise ratios over time. Overall, the findings reinforced earlier research concluding that Zero Trust strengthens both preventive and reactive security capacity, improving the speed and effectiveness of detection, response, and containment operations.

This study's demographic findings revealed that privileged and administrative accounts, despite representing a small fraction of the identity population, contributed disproportionately to high-risk authentication events and lateral movement alerts. This pattern has been consistently documented in earlier studies of enterprise compromise, where attackers prioritize privileged identities to gain broad access and persistence. The findings therefore aligned with the established literature emphasizing privileged access as a primary risk amplifier. The observed concentration of risk among privileged accounts highlighted that Zero Trust effectiveness is closely tied to privilege governance rather than general user authentication alone (Shjarback et al., 2017).The post-implementation reductions in anomalies and lateral movement alerts suggested that identity governance improvements reduced privileged account exposure and constrained privileged access pathways. Earlier studies have emphasized that privilege sprawl, unmanaged service accounts, and excessive group memberships create persistent escalation opportunities. The findings of this study supported the argument that stronger identity-based policy enforcement reduces these opportunities by restricting access decisions more tightly and by generating auditable evidence of privileged activity (Wang, 2019).Department-level differences in authentication volumes and access denials also reflected earlier research highlighting that risk is unevenly distributed across organizational roles. IT and administrative units typically require broader access, making them more exposed to both legitimate high-risk activity and targeted attack attempts. The results suggested that Zero Trust enforcement provided measurable benefits in such high-access contexts by reducing the escalation potential of privileged identities. Earlier work has noted that the most effective Zero Trust implementations prioritize high-value targets and privileged pathways, rather than applying uniform controls across all identities. This study's findings reinforced that recommendation by showing that risk reduction was most relevant where privilege concentration existed (Loi et al., 2021).The results also had governance implications consistent with earlier research. Strong policy strictness predicted better outcomes, suggesting that organizations must maintain enforcement discipline to sustain incident reduction. Earlier studies have warned that policy exceptions, inconsistent enforcement, and administrative drift can erode Zero Trust effectiveness over time. The findings of this study suggested that consistent identity governance was a key mechanism

driving the observed reductions in incidents and escalation. Overall, the findings reinforced the broader literature's conclusion that Zero Trust is most effective when combined with disciplined privilege governance, consistent policy enforcement, and continuous monitoring of privileged identity behavior.

Overall, this study provided empirical evidence supporting earlier research that described Zero Trust as an effective framework for reducing security incidents and limiting breach escalation. The observed reductions in incident frequency, high-severity incidents, authentication anomalies, and lateral movement alerts aligned with prior conceptual and empirical studies emphasizing continuous verification, least privilege, and segmentation as mechanisms for improving security posture (Wolfe et al., 2018). The improvements in SOC performance metrics further supported earlier work suggesting that Zero Trust enhances detection and response by improving telemetry clarity and access governance consistency. The results therefore strengthened the argument that identity-based controls can produce measurable operational outcomes rather than serving only as conceptual guidance.This study also extended earlier research by demonstrating that policy strictness was not merely an implementation characteristic but a measurable predictor of security outcomes. Earlier IAM studies have suggested that stricter enforcement reduces unauthorized access success and limits privilege misuse, but empirical enterprise studies measuring these relationships remain limited. The regression findings of this study contributed to addressing that gap by linking enforcement intensity to reductions in unauthorized access attempts, lateral movement indicators, and incident severity probability. This strengthened the evidence base for interpreting Zero Trust as an operationally testable model (Kurland et al., 2017).The practical significance of the findings was also consistent with earlier studies emphasizing that reducing detection and response times has direct operational and financial value. Faster detection and containment reduce attacker dwell time and limit damage, and the improvements observed in this study were large enough to be operationally meaningful. In addition, the reduction in incident variance suggested improved stability, which supports sustainable SOC operations. Earlier research has highlighted that security programs often fail not only because controls are ineffective, but because operational teams cannot sustain enforcement and response workloads. This study's findings suggested that Zero Trust can contribute to both improved security and improved operational sustainability when implemented through integrated identity and access systems (Headley et al., 2017).Finally, the results reinforced earlier research emphasizing that Zero Trust success depends on identity governance maturity, privilege management, and consistent enforcement. The concentration of risk among privileged accounts highlighted the importance of focusing Zero Trust policies on high-impact identities and access pathways. Taken together, the findings provided strong support for the conclusion that Cisco ISE and AD-integrated identity-based Zero Trust controls are associated with measurable and statistically significant improvements in incident reduction, escalation prevention, and SOC performance outcomes, consistent with the dominant trends reported in earlier Zero Trust security research.

CONCLUSION

This study provided quantitative evidence that the implementation of Zero Trust identity-based access controls through Cisco Identity Services Engine integrated with Active Directory was associated with measurable and statistically significant improvements in enterprise security performance. The findings demonstrated sustained reductions in total incident frequency, high-severity incidents, unauthorized access attempts, authentication anomalies, and lateral movement alerts following policy enforcement. In addition to lowering incident volume, the results indicated a meaningful decline in escalation probability, as reflected in reduced high-severity classification rates and constrained internal movement indicators. Operational performance also improved substantially, with significant reductions observed in mean time to detect, mean time to respond, and time-to-containment, suggesting that identity-centric telemetry enhanced SOC efficiency and investigative clarity. The integration of centralized authentication, authorization, and policy-driven segmentation appeared to transform implicit trust relationships into enforceable, auditable control points that limited credential misuse and privilege-based exploitation. The concentration of risk among privileged accounts reinforced the importance of identity governance discipline and strict privilege management as central components of Zero Trust effectiveness. Reliability analysis confirmed that composite constructs

measuring policy strictness, security posture strength, and behavioral risk were internally consistent, supporting the robustness of the statistical models used to evaluate impact. Regression findings further demonstrated that policy strictness was not merely an implementation attribute but a significant predictor of improved security outcomes, highlighting the practical importance of enforcement intensity and segmentation precision. Overall, the results aligned with and extended prior research emphasizing that Zero Trust architecture can produce empirically observable improvements when implemented with strong identity integration and consistent governance controls. The study therefore contributed quantitative validation to a domain often characterized by conceptual discussion, demonstrating that identity-based Zero Trust controls are capable of delivering sustained reductions in incident frequency, escalation risk, and response latency within enterprise environments.

RECOMMENDATION

Based on the quantitative findings of this study, several evidence-based recommendations were established to strengthen the security impact of Zero Trust access controls implemented through Cisco Identity Services Engine integrated with Active Directory. First, identity governance should be treated as a continuous operational process rather than a one-time deployment outcome. Regular privileged access reviews, automated deprovisioning, and systematic reduction of unnecessary administrative rights should be prioritized, as privileged identities were shown to contribute disproportionately to high-risk authentication events and lateral movement indicators. Second, policy strictness should be maintained at a consistently enforceable level, with structured exception governance to prevent policy drift. The regression findings demonstrated that enforcement intensity significantly predicted reductions in incident frequency, unauthorized access outcomes, and lateral movement alerts, indicating that relaxed enforcement could erode measurable security gains. Third, segmentation policies should be expanded progressively using identity-driven access tiers, ensuring that sensitive systems and high-value assets are isolated from general user zones through dynamic authorization rules. This approach supports containment by limiting east-west traversal pathways and reducing escalation probability. Fourth, SOC performance improvements should be sustained by strengthening SIEM correlation rules that integrate Cisco ISE event data, AD security logs, and incident workflow timelines. Since reductions in detection and response time were central outcomes of the implementation, continued investment in log normalization, identity enrichment, and automated triage workflows is recommended to preserve efficiency gains and reduce analyst workload volatility. Fifth, organizations should implement periodic quantitative monitoring using standardized metrics such as incident recurrence rates, authentication anomaly trends, and time-to-containment distributions to ensure that Zero Trust effectiveness remains measurable over time. Finally, policy usability and compliance should be addressed through targeted training, adaptive authentication, and communication strategies that reduce friction without weakening security controls. Sustained enforcement requires user cooperation, and excessive access friction can lead to workaround behaviors that introduce hidden risk. Overall, these recommendations emphasize that the strongest security outcomes are achieved when Zero Trust is implemented as an integrated identity governance and operational measurement program, combining strict policy enforcement, privilege discipline, segmentation maturity, and SOC telemetry optimization to ensure continued incident reduction and resilience improvement.

LIMITATIONS

Several limitations were associated with this study and should be considered when interpreting the findings. First, the research design employed a quasi-experimental pretest–posttest structure with an interrupted time-series approach, which strengthened internal inference but did not provide the full causal certainty of randomized experimental assignment. Although baseline trends, observation windows, and statistical controls were applied, external events such as changes in attacker behavior, seasonal threat variation, or concurrent security initiatives could have influenced incident outcomes during the evaluation period. Second, the study relied heavily on log-based and SIEM-correlated data, meaning that findings were dependent on the accuracy, completeness, and stability of logging configurations and correlation rules. Any undocumented changes in SIEM detection content, log collection coverage, or analyst classification practices could have affected incident counts, severity labels, and time-based performance metrics. Third, the operational definitions of constructs such as

authentication anomalies, unauthorized access attempts, and lateral movement alerts were based on SOC detection logic and enterprise-specific policy frameworks, which may limit comparability across organizations with different detection maturity, alert thresholds, or classification standards. Fourth, the study was conducted within a single enterprise case environment, which restricted generalizability. Organizations with different infrastructure maturity, workforce scale, cloud adoption levels, or governance practices may experience different outcomes when implementing Cisco ISE and AD-integrated Zero Trust controls. Fifth, while the study measured improvements in incident reduction and SOC performance, it did not directly quantify broader business impacts such as cost savings, productivity effects, or user experience outcomes associated with policy enforcement. Additionally, service accounts and automated processes were treated separately due to distinct behavioral patterns, but their risk contribution may have been underestimated if malicious activity successfully blended into normal automation flows. Finally, the study emphasized measurable security outcomes over deeper causal explanation, meaning that while statistical associations were strong, the precise organizational mechanisms driving the improvements—such as specific policy changes, training effects, or operational workflow refinements—were not isolated as independent contributors. These limitations indicate that findings should be interpreted as strong empirical evidence of association within the evaluated environment rather than universally generalizable causal proof, while still providing valuable quantitative support for the security effectiveness of identity-based Zero Trust access controls.

REFERENCES

[1]. Abiodun, O. I., Abiodun, E. O., Alawida, M., Alkhawaldeh, R. S., & Arshad, H. (2021). A review on the security of the internet of things: Challenges and solutions. *Wireless Personal Communications*, *119*(3), 2603-2637.

[2]. Addae, J. H., Sun, X., Towey, D., & Radenkovic, M. (2019). Exploring user behavioral data for adaptive cybersecurity: JH Addae et al. *User Modeling and User-Adapted Interaction*, *29*(3), 701-750.

[3]. Ahmadi-Assalemi, G., Al-Khateeb, H., Epiphaniou, G., & Maple, C. (2020). Cyber resilience and incident response in smart cities: A systematic literature review. *Smart Cities*, *3*(3), 894-927.

[4]. AI, P.-Q. C. The CISO's Next Frontier.

[5]. Al-Aqrabi, H., & Hill, R. (2020). Evaluating Multi-layer Security Resistance to Adversarial Hacking Attacks on Industrial Internet of Things Devices. In *Digital Forensic Investigation of Internet of Things (IoT) Devices* (pp. 187-203). Springer.

[6]. Ali, R. F., Dominic, P., Ali, S. E. A., Rehman, M., & Sohail, A. (2021). Information security behavior and information security policy compliance: A systematic literature review for identifying the transformation process from noncompliance to compliance. *Applied Sciences*, *11*(8), 3383.

[7]. Aliyu, A., Maglaras, L., He, Y., Yevseyeva, I., Boiten, E., Cook, A., & Janicke, H. (2020). A holistic cybersecurity maturity assessment framework for higher education institutions in the United Kingdom. *Applied Sciences*, *10*(10), 3660.

[8]. Alo, U. R., Nkwo, F. O., Nweke, H. F., Achi, I. I., & Okemiri, H. A. (2021). Non-pharmaceutical interventions against COVID-19 pandemic: Review of contact tracing and social distancing technologies, protocols, apps, security and open research directions. *Sensors*, *22*(1), 280.

[9]. Antoniou, C., Dimitriou, L., & Pereira, F. (2018). *Mobility patterns, big data and transport analytics: tools and applications for modeling*. Elsevier.

[10]. Badhwar, R. (2021). *CISO's Transformation*. Springer.

[11]. Becker, C., Julien, C., Lalanda, P., & Zambonelli, F. (2019). Pervasive computing middleware: current trends and emerging challenges. *CCF Transactions on Pervasive Computing and Interaction*, *1*(1), 10-23.

[12]. Bogatov, D., De Caro, A., Elkhiyaoui, K., & Tackmann, B. (2021). Anonymous transactions with revocation and auditing in hyperledger fabric. International Conference on Cryptology and Network Security,

[13]. Bonella, B., Beirwolf, K., Coleman, L., Sterger, C., Pulli, K., Anguiano, C., & Barton, K. (2020). Mixed methods evaluation of formal mentoring: Journey UP for aging out of foster care. *Global social welfare*, *7*(2), 121-130.

[14]. Borky, J. M., & Bradley, T. H. (2018). Protecting information with cybersecurity. In *Effective model-based systems engineering* (pp. 345-404). Springer.

[15]. Brodeur, A., Grigoryeva, I., & Kattan, L. (2021). Stay-at-home orders, social distancing, and trust. *Journal of Population Economics*, *34*(4), 1321-1354.

[16]. Broeders, D., Schrijvers, E., van der Sloot, B., Van Brakel, R., De Hoog, J., & Ballin, E. H. (2017). Big Data and security policies: Towards a framework for regulating the phases of analytics and use of Big Data. *Computer Law & Security Review*, *33*(3), 309-323.

[17]. Cali, U., Kuzlu, M., Pipattanasomporn, M., Kempf, J., & Bai, L. (2021). *Digitalization of power markets and systems using energy informatics*. Springer.

[18]. Cantelmi, R., Di Gravio, G., & Patriarca, R. (2021). Reviewing qualitative research approaches in the context of critical infrastructure resilience. *Environment Systems and Decisions*, *41*(3), 341-376.

[19]. Cheruvu, S., Kumar, A., Smith, N., & Wheeler, D. M. (2019). IoT frameworks and complexity. *Demystifying Internet of Things Security: Successful IoT Device/Edge and Platform Security Deployment*, 23-148.

[20]. Cheruvu, S., Kumar, A., Smith, N., & Wheeler, D. M. (2020). *Demystifying internet of things security: successful iot device/edge and platform security deployment*. Springer.

[21]. Chifor, B.-C., Arseni, Ş.-C., & Bica, I. (2021). IoT Cloud Security Design Patterns. In *Big Data Platforms and Applications: Case Studies, Methods, Techniques, and Performance Evaluation* (pp. 113-164). Springer.

[22]. Chua, H. N., Wong, S. F., Low, Y. C., & Chang, Y. (2018). Impact of employees' demographic characteristics on the awareness and compliance of information security policy in organizations. *Telematics and Informatics*, *35*(6), 1770-1780.

[23]. Cyber-Defense, B. E., & Haber, M. J. Privileged Attack Vectors.

[24]. de la Torre, C., & Polo, J. C. (2018). Cloud computing and network analysis. International Conference on Information Systems Architecture and Technology,

[25]. Di Salvo, C. (2018). How Blockchain Will Change Cybersecurity Practices. In *Cybersecurity Best Practices: Lösungen zur Erhöhung der Cyberresilienz für Unternehmen und Behörden* (pp. 493-510). Springer.

[26]. Dumka, A., & Sah, A. (2019). Smart ambulance system using concept of big data and internet of things. In *Healthcare data analytics and management* (pp. 155-176). Elsevier.

[27]. Dupont, B. (2017). Bots, cops, and corporations: on the limits of enforcement and the promise of polycentric regulation as a way to control large-scale cybercrime. *Crime, Law and Social Change*, *67*(1), 97-116.

[28]. Eichelberg, M., Kleber, K., & Kämmerer, M. (2020). Cybersecurity in PACS and medical imaging: an overview. *Journal of Digital Imaging*, *33*(6), 1527-1542.

[29]. El Sayed, S. A., Sanford, S. M., & Kerley, K. R. (2019). Understanding workplace stress among federal law enforcement officers. *American journal of criminal justice*, *44*(3), 409-429.

[30]. Enachescu, J., Olsen, J., Kogler, C., Zeelenberg, M., Breugelmans, S. M., & Kirchler, E. (2019). The role of emotions in tax compliance behavior: A mixed-methods approach. *Journal of Economic Psychology*, *74*, 102194.

[31]. Faysal, K., & Shamsunnahar, C. (2022). Digital Ledger Optimization Techniques for Enhancing Transaction Speed and Reporting Accuracy in Accounting Systems. *American Journal of Scholarly Research and Innovation*, *1*(02), 171–222. https://doi.org/10.63125/33t06k57

[32]. Fields, T., Patterson, M., Bremova-Ertl, T., Belcher, G., Billington, I., Churchill, G. C., Davis, W., Evans, W., Flint, S., & Galione, A. (2021). A master protocol to investigate a novel therapy acetyl-L-leucine for three ultra-rare neurodegenerative diseases: Niemann-Pick type C, the GM2 gangliosidoses, and ataxia telangiectasia. *Trials*, *22*(1), 84.

[33]. Garbis, J., & Chapman, J. W. Zero trust security.

[34]. Garbis, J., & Chapman, J. W. (2021). Identity and access management. In *Zero Trust Security: An Enterprise Guide* (pp. 71-91). Springer.

[35]. Gidey, M. T., Gebretekle, G. B., Hogan, M.-E., & Fenta, T. G. (2019). Willingness to pay for social health insurance and its determinants among public servants in Mekelle City, Northern Ethiopia: a mixed methods study. *Cost Effectiveness and Resource Allocation*, *17*(1), 2.

[36]. Golan, M. S., Jernegan, L. H., & Linkov, I. (2020). Trends and applications of resilience analytics in supply chain modeling: systematic literature review in the context of the COVID-19 pandemic. *Environment Systems and Decisions*, *40*(2), 222-243.

[37]. González-Granadillo, G., González-Zarzosa, S., & Diaz, R. (2021). Security information and event management (SIEM): analysis, trends, and usage in critical infrastructures. *Sensors*, *21*(14), 4759.

[38]. Goodman, H., Papastavrou Brooks, C., Price, O., & Barley, E. A. (2020). Barriers and facilitators to the effective de-escalation of conflict behaviours in forensic high-secure settings: a qualitative study. *International journal of mental health systems*, *14*(1), 59.

[39]. Gupta, S., & Giri, V. (2018). Data security in data lakes. In *Practical Enterprise Data Lake Insights: Handle Data-Driven Challenges in an Enterprise Big Data Lake* (pp. 225-259). Springer.

[40]. Gwadz, M., Campos, S., Freeman, R., Cleland, C. M., Wilton, L., Sherpa, D., Ritchie, A. S., Hawkins, R. L., Allen, J. Y., & Martinez, B. Y. (2021). Black and Latino persons living with HIV evidence risk and resilience in the context of COVID-19: a mixed-methods study of the early phase of the pandemic. *AIDS and Behavior*, *25*(5), 1340-1360.

[41]. Habibullah, S. M., & Zaheda, K. (2022). Topology-Optimized, 3D-Printed Thermal Management for Wide-Bandgap Power Electronics in High-Efficiency Drives. *Journal of Sustainable Development and Policy*, *1*(02), 134-167. https://doi.org/10.63125/p8m2p864

[42]. Haff, G. (2021). From "Free" to "Open Source" to Products. In *How Open Source Ate Software: Understand the Open Source Movement and So Much More* (pp. 27-45). Springer.

[43]. Hardjono, T., & Pentland, A. (2018). Open algorithms for identity federation. Future of Information and Communication Conference,

[44]. Headley, A. M., Guerette, R. T., & Shariati, A. (2017). A field experiment of the impact of body-worn cameras (BWCs) on police officer behavior and perceptions. *Journal of criminal justice*, *53*, 102-109.

[45]. Hoang, D. B., & Farahmandian, S. (2017). Security of software-defined infrastructures with SDN, NFV, and cloud computing technologies. In *Guide to Security in SDN and NFV: Challenges, Opportunities, and Applications* (pp. 3-32). Springer.

[46]. Hofer, M. S., & Savell, S. M. (2021). "There was no plan in place to get us help": strategies for improving mental health service utilization among law enforcement. *Journal of police and criminal psychology*, *36*(3), 543-557.

[47]. Jahangir, S., & Md Shahab, U. (2022). A Qualitative Study of Safety Professionals' Experiences in Managing Chemical Exposure Risks and Hazardous Materials Controls in Industrial Facilities. *Review of Applied Science and Technology*, *1*(04), 250–282. https://doi.org/10.63125/jmh69r20

[48]. Johnson, J., Simms-Ellis, R., Janes, G., Mills, T., Budworth, L., Atkinson, L., & Harrison, R. (2020). Can we prepare healthcare professionals and students for involvement in stressful healthcare events? A mixed-methods evaluation of a resilience training intervention. *BMC health services research*, *20*(1), 1094.

[49]. Joshi, C., & Singh, U. K. (2017). Information security risks management framework–A step towards mitigating security risks in university network. *Journal of Information Security and Applications*, *35*, 128-137.

[50]. Kamboj, P., & Pal, S. (2021). A policy based framework for quality of service management in software defined networks. *Telecommunication Systems*, *78*(3), 331-349.

[51]. Kamminga, J., Ayele, E., Meratnia, N., & Havinga, P. (2018). Poaching detection technologies—a survey. *Sensors*, *18*(5), 1474.

[52]. Kaur, G., Habibi Lashkari, Z., & Habibi Lashkari, A. (2021). Introduction to cybersecurity. In *Understanding Cybersecurity Management in FinTech: Challenges, Strategies, and Trends* (pp. 17-34). Springer.

[53]. Kaur, G., Lashkari, Z. H., & Lashkari, A. H. (2021). *Understanding cybersecurity management in FinTech*. Springer.

[54]. Keeriyattil, S. Zero Trust Networks with VMware NSX.

[55]. Kenyon, T. (2018). Transportation cyber-physical systems security and privacy. In *Transportation Cyber-Physical Systems* (pp. 115-151). Elsevier.

[56]. Kim, S., Crooks, C. V., Bax, K., & Shokoohi, M. (2021). Impact of trauma-informed training and mindfulness-based social–emotional learning program on teacher attitudes and burnout: A mixed-methods study. *School mental health*, *13*(1), 55-68.

[57]. Krishnamurthi, R., Kumar, A., Gopinathan, D., Nayyar, A., & Qureshi, B. (2020). An overview of IoT sensor data processing, fusion, and analysis techniques. *Sensors*, *20*(21), 6076.

[58]. Kristen, E., Kloibhofer, R., Díaz, V. H., & Castillejo, P. (2021). Security assessment of agriculture iot (Aiot) applications. *Applied Sciences*, *11*(13), 5841.

[59]. Kurland, J., Pires, S. F., McFann, S. C., & Moreto, W. D. (2017). Wildlife crime: a conceptual integration, literature review, and methodological critique. *Crime science*, *6*(1), 4.

[60]. Kweon, E., Lee, H., Chai, S., & Yoo, K. (2021). The utility of information security training and education on cybersecurity incidents: An empirical evidence. *Information Systems Frontiers*, *23*(2), 361-373.

[61]. Laurent, A., Abdel-Sayed, P., Grognuz, A., Scaletta, C., Hirt-Burri, N., Michetti, M., de Buys Roessingh, A. S., Raffoul, W., Kronen, P., & Nuss, K. (2021). Industrial development of standardized fetal progenitor cell therapy for tendon regenerative medicine: Preliminary safety in xenogeneic transplantation. *Biomedicines*, *9*(4), 380.

[62]. Lee, E. B. (2021). Mullen scales of early learning. In *Encyclopedia of autism spectrum disorders* (pp. 3024-3029). Springer.

[63]. Levi, M., Doig, A., Gundur, R., Wall, D., & Williams, M. (2017). Cyberfraud and the implications for effective risk-based responses: themes from UK research. *Crime, Law and Social Change*, *67*(1), 77-96.

[64]. Liu, S., Zhang, Z.-J., Cui, Y., & Zhang, Y. (2020). An Industrial-Grade API Secure Access Gateway in the Cloud-Edge Integration Scenario. International Conference on Smart Grid and Internet of Things,

[65]. Liu, X., Yang, F., Cheng, W., Wu, Y., Cheng, J., Sun, W., Yan, X., Luo, M., Mo, X., & Hu, M. (2020). Mixed methods research on satisfaction with basic medical insurance for urban and rural residents in China. *BMC public health*, *20*(1), 1201.

[66]. Lodi, E., Perrella, L., Lepri, G. L., Scarpa, M. L., & Patrizi, P. (2021). Use of restorative justice and restorative practices at school: A systematic literature review. *International journal of environmental research and public health*, *19*(1), 96.

[67]. Loi, K. I., Lei, W. S., & Lourenco, F. (2021). Understanding the reactions of government and gaming concessionaires on COVID-19 through the neo-institutional theory–The case of Macao. *International Journal of Hospitality Management*, *94*, 102755.

[68]. Lorenzo-Dus, N., & Izura, C. (2017). "cause ur special": Understanding trust and complimenting behaviour in online grooming discourse. *Journal of Pragmatics*, *112*, 68-82.

[69]. Marsden, G., & Reardon, L. (2017). Questions of governance: Rethinking the study of transportation policy. *Transportation Research Part A: Policy and Practice*, *101*, 238-251.

[70]. Martins, J., Alves, J., Cabral, J., Tavares, A., & Pinto, S. (2017). µ rtzvisor: a secure and safe real-time hypervisor. *Electronics*, *6*(4), 93.

[71]. Minth, L. (2018). Cyber Governance: Knowing and Doing What's Important for making Smart Cities resilient. In *Cybersecurity Best Practices: Lösungen zur Erhöhung der Cyberresilienz für Unternehmen und Behörden* (pp. 477-492). Springer.

[72]. Mongeau, S., & Hajdasinski, A. (2021). Phase III: CSDS Gap-Prescriptions—Design Science Problem-Solving. In *Cybersecurity Data Science: Best Practices in an Emerging Profession* (pp. 201-316). Springer.

[73]. Mongeau, S. A. (2021). *Cybersecurity Data Science*. Springer.

[74]. More, S. M., & Haff, G. How Open Source Ate Software.

[75]. Mourtzis, D., Angelopoulos, J., & Panopoulos, N. (2021). Smart manufacturing and tactile internet based on 5G in industry 4.0: Challenges, applications and new trends. *Electronics*, *10*(24), 3175.

[76]. Mukhopadhyay, S., Bouwman, H., & Jaiswal, M. P. (2019). An open platform centric approach for scalable government service delivery to the poor: The Aadhaar case. *Government Information Quarterly*, *36*(3), 437-448.

[77]. Ots, K. (2021). *Azure Security Handbook*. Springer.

[78]. Pagano, A., Pluchinotta, I., Pengal, P., Cokan, B., & Giordano, R. (2019). Engaging stakeholders in the assessment of NBS effectiveness in flood risk reduction: A participatory System Dynamics Model for benefits and co-benefits evaluation. *Science of the Total Environment*, *690*, 543-555.

[79]. Panarello, A., Tapas, N., Merlino, G., Longo, F., & Puliafito, A. (2018). Blockchain and iot integration: A systematic survey. *Sensors*, *18*(8), 2575.

[80]. Parrend, P., Navarro, J., Guigou, F., Deruyver, A., & Collet, P. (2018). Foundations and applications of artificial Intelligence for zero-day and multi-step attack detection. *EURASIP Journal on Information Security*, *2018*(1), 4.

[81]. Pasman, H. J., Rogers, W. J., & Mannan, M. S. (2017). Risk assessment: What is it worth? Shall we just do away with it, or can it do a better job? *Safety Science*, *99*, 140-155.

[82]. Petitta, L., Probst, T. M., Barbaranelli, C., & Ghezzi, V. (2017). Disentangling the roles of safety climate and safety culture: Multi-level effects on the relationship between supervisor enforcement and safety compliance. *Accident Analysis & Prevention*, *99*, 77-89.

[83]. Pilgrim, K., & Bohnet-Joschko, S. (2019). Selling health and happiness how influencers communicate on Instagram about dieting and exercise: mixed methods research. *BMC public health*, *19*(1), 1054.

[84]. Pogoryelova, O., Cammish, P., Mansbach, H., Argov, Z., Nishino, I., Skrinar, A., Chan, Y., Nafissi, S., Shamshiri, H., & Kakkis, E. (2018). Phenotypic stratification and genotype–phenotype correlation in a heterogeneous, international cohort of GNE myopathy patients: first report from the GNE myopathy disease monitoring program, registry portion. *Neuromuscular Disorders*, *28*(2), 158-168.

[85]. Pogrebna, G., & Skilton, M. (2019). *Navigating new cyber risks*. Springer.

[86]. Provan, D. J., Woods, D. D., Dekker, S. W., & Rae, A. J. (2020). Safety II professionals: How resilience engineering can transform safety practice. *Reliability Engineering & System Safety*, *195*, 106740.

[87]. Rabat, Y., Chanraud, S., Abdallah, M., Sibon, I., & Berthoz, S. (2021). Precision Preventive Medicine of Relapse in Smoking Cessation: Can MRI Inform the Search of Intermediate Phenotypes? *Biology*, *11*(1), 35.

[88]. Raddaoui, O., Ahmed, M. M., & Gaweesh, S. M. (2020). Assessment of the effectiveness of connected vehicle weather and work zone warnings in improving truck driver safety. *IATSS research*, *44*(3), 230-237.

[89]. Ramacciati, N., Ceccagnoli, A., Addey, B., & Rasero, L. (2018). Violence towards emergency nurses. The Italian national survey 2016: a qualitative study. *International journal of nursing studies*, *81*, 21-29.

[90]. Ratul, D., & Subrato, S. (2022). Remote Sensing Based Integrity Assessment of Infrastructure Corridors Using Spectral Anomaly Detection and Material Degradation Signatures. *American Journal of Interdisciplinary Studies*, *3*(04), 332-364. https://doi.org/10.63125/1sdhwn89

[91]. Read, G. J., Naweed, A., & Salmon, P. M. (2019). Complexity on the rails: A systems-based approach to understanding safety management in rail transport. *Reliability Engineering & System Safety*, *188*, 352-365.

[92]. Rekha, G., & Yashaswini, J. (2021). Industry 4.0: A revolution in healthcare sector via cloud, fog technologies. In *Intelligent Interactive Multimedia Systems for e-Healthcare Applications* (pp. 321-335). Springer.

[93]. Retico, A., Avanzo, M., Boccali, T., Bonacorsi, D., Botta, F., Cuttone, G., Martelli, B., Salomoni, D., Spiga, D., & Trianni, A. (2021). Enhancing the impact of Artificial Intelligence in Medicine: A joint AIFM-INFN Italian initiative for a dedicated cloud-based computing infrastructure. *Physica Medica*, *91*, 140-150.

[94]. Ribeiro, J. P., & Barbosa-Povoa, A. (2018). Supply Chain Resilience: Definitions and quantitative modelling approaches–A literature review. *Computers & industrial engineering*, *115*, 109-122.

[95]. Roberge-Dao, J., Yardley, B., Menon, A., Halle, M.-C., Maman, J., Ahmed, S., & Thomas, A. (2019). A mixed-methods approach to understanding partnership experiences and outcomes of projects from an integrated knowledge translation funding model in rehabilitation. *BMC health services research*, *19*(1), 230.

[96]. Schwartz, M., & Machulak, M. (2018). Securing the Perimeter. *Deploying Identity and Access Management with Free Open Source Software*.

[97]. Sehgal, N. K., & Bhatt, P. C. (2018). Cloud computing. *Cham: Springer International Publishing*.

[98]. Serrano, W. (2020). Big Data in smart infrastructure. Proceedings of SAI Intelligent Systems Conference,

[99]. Shjarback, J. A., Pyrooz, D. C., Wolfe, S. E., & Decker, S. H. (2017). De-policing and crime in the wake of Ferguson: Racialized changes in the quantity and quality of policing among Missouri police departments. *Journal of criminal justice*, *50*, 42-52.

[100]. Sinha, S. R., & Park, Y. (2017). *Building an E Ective IoT Ecosystem for Your Business*. Springer.

[101]. Stanković, R., Štula, M., & Maras, J. (2017). Evaluating fault tolerance approaches in multi-agent systems. *Autonomous agents and multi-agent systems*, *31*(1), 151-177.

[102]. Steingartner, W., Galinec, D., & Kozina, A. (2021). Threat defense: Cyber deception approach and education for resilience in hybrid threats model. *Symmetry*, *13*(4), 597.

[103]. Suomalainen, J., Ahola, K., Majanen, M., Mämmelä, O., & Ruuska, P. (2018). Security awareness in software-defined multi-domain 5G networks. *Future Internet*, *10*(3), 27.

[104]. Tahmina Akter Bhuya, M., & Rebeka, S. (2022). AI-Assisted Underwriting Models for Improving Risk Assessment Accuracy in U.S. Insurance Markets. *American Journal of Interdisciplinary Studies*, *3*(01), 65-102. https://doi.org/10.63125/kegg1076

[105]. Tetzlaff, E. J., Goggins, K. A., Pegoraro, A. L., Dorman, S. C., Pakalnis, V., & Eger, T. R. (2021). Safety culture: a retrospective analysis of occupational health and safety mining reports. *Safety and health at work*, *12*(2), 201-208.

[106]. Torten, R., Reaiche, C., & Boyle, S. (2018). The impact of security awareness on information technology professionals' behavior. *Computers & Security*, *79*, 68-79.

[107]. Trang, S., & Brendel, B. (2019). A meta-analysis of deterrence theory in information security policy compliance research. *Information Systems Frontiers*, *21*(6), 1265-1284.

[108]. Tyler, D., & Viana, T. (2021). Trust no one? a framework for assisting healthcare organisations in transitioning to a zero-trust network architecture. *Applied Sciences*, *11*(16), 7499.

[109]. Tzimas, D., & Demetriadis, S. (2021). Ethical issues in learning analytics: A review of the field. *Educational Technology Research and Development*, *69*(2), 1101-1133.

[110]. Verhage, A., & Boels, D. (2017). Critical appraisal of mixed methods research studies in a systematic scoping review on plural policing: assessing the impact of excluding inadequately reported studies by means of a sensitivity analysis. *Quality & quantity*, *51*(4), 1449-1468.

[111]. Voigt, P., & Von dem Bussche, A. (2017). The eu general data protection regulation (gdpr). *A practical guide, 1st ed., Cham: Springer International Publishing*, *10*(3152676), 10-5555.

[112]. Wall, T., & Rodrick, J. Jump-start Your SOC Analyst Career.

[113]. Wang, S. P. (2021). Communication, TCP/IP, and Internet. In *Computer Architecture and Organization: Fundamentals and Architecture Security* (pp. 243-292). Springer.

[114]. Wang, S. S. (2019). Integrated framework for information security investment and cyber insurance. *Pacific-Basin Finance Journal*, *57*, 101173.

[115]. Wolfe, S. E., Rojek, J., Manjarrez Jr, V. M., & Rojek, A. (2018). Why does organizational justice matter? Uncertainty management among law enforcement officers. *Journal of criminal justice*, *54*, 20-29.

[116]. Wright, A. L., Sonin, K., Driscoll, J., & Wilson, J. (2020). Poverty and economic dislocation reduce compliance with COVID-19 shelter-in-place protocols. *Journal of Economic Behavior & Organization*, *180*, 544-554.

[117]. Xiang, X., Foo, S., & Zang, H. (2021). Recent advances in deep reinforcement learning applications for solving Partially Observable Markov Decision Processes (POMDP) problems part 2—Applications in transportation, industries, communications and networking and more topics. *Machine Learning and Knowledge Extraction*, *3*(4), 863-878.

[118]. Yang, L. E., Scheffran, J., Süsser, D., Dawson, R., & Chen, Y. D. (2018). Assessment of Flood Losses with Household Responses: Agent-Based Simulation in an Urban Catchment Area: Yang LE et al. *Environmental Modeling & Assessment*, *23*(4), 369-388.

[119]. Zerlang, J. (2017). GDPR: a milestone in convergence for cyber-security and compliance. *Network Security*, *2017*(6), 8-11.

[120]. Zhang, T., Yan, L., & Yang, Y. (2018). Trust evaluation method for clustered wireless sensor networks based on cloud model. *Wireless Networks*, *24*(3), 777-797.