Computer Fraud and Security
ISSN (online): 1873-7056

Zero-Trust Identity and Access Architecture for Securing Large-Scale Distributed Enterprise Systems

Ishu Anand Jaiswal

Affiliation: Apple Inc., Cupertino, California, USA,

ijaiswal@apple.com

Abstract

The rapid adoption of cloud computing, hybrid infrastructures, Software-as-a-Service platforms, microservices, mobile computing, remote working, application programming interfaces, Internet of Things devices, and geographically distributed enterprise applications has significantly altered the security requirements of modern information systems. Traditional enterprise-security architectures generally rely on network boundaries in which internal users, systems, and devices receive a relatively higher level of implicit trust than entities operating outside the organizational perimeter. However, distributed enterprise environments increasingly place applications, identities, workloads, and sensitive data across multiple cloud platforms, remote locations, partner networks, and heterogeneous computing infrastructures. Under such conditions, network location alone cannot provide sufficient assurance regarding the legitimacy of an access request. This study proposes a Zero-Trust Identity and Access Architecture for Large-Scale Distributed Enterprise Systems (ZTIAA-DES) that places verified identity, device posture, contextual trust, least-privilege authorization, continuous access evaluation, and policy-driven enforcement at the center of enterprise security. The architecture integrates Identity and Access Management, Multi-Factor Authentication, federated identity, Single Sign-On, Role-Based Access Control, Attribute-Based Access Control, device verification, dynamic trust assessment, behavioral monitoring, resource-sensitivity analysis, Policy Decision Points, Policy Enforcement Points, session-risk monitoring, and identity lifecycle governance. Every enterprise access request is evaluated independently according to who or what is requesting access, the current security state of the device or workload, requested resource, business context, behavioral characteristics, and current risk level.

Keywords: Zero Trust Architecture, Identity and Access Management, Distributed Enterprise Systems, Continuous Authentication, Adaptive Access Control.

Introduction

Modern enterprises increasingly depend on large-scale distributed information systems to support business operations, communication, data processing, collaboration, customer services, analytics, cloud applications, and digital service delivery. Enterprise resources that were once maintained primarily within centralized organizational data centers are now distributed across public clouds, private clouds, hybrid infrastructures, Software-as-a-Service platforms, regional data centers, edge environments, mobile devices, remote offices, and third-party service providers. This transformation has improved organizational flexibility, scalability, availability, and accessibility, but it has also created a significantly more complex security environment. Users, applications, devices, APIs, workloads, and services continuously interact across different administrative and network boundaries, making conventional assumptions about trusted internal environments increasingly difficult to maintain. Traditional enterprise-security architectures are generally based on perimeter-oriented security. In this approach, organizational networks are protected using firewalls, virtual private networks, intrusion-detection mechanisms, network segmentation, authentication gateways, and other preventive security controls. Users and devices located within a protected internal network may receive a higher degree of implicit trust than entities connecting from external networks. Such an approach was practical when most enterprise applications, employees, devices, and data were located within clearly defined organizational boundaries. However, modern distributed enterprise systems no longer operate within a single trusted perimeter. Employees may work remotely, business partners may access selected enterprise resources, cloud applications may be hosted outside organizational infrastructure, and users may connect through mobile networks or unmanaged environments. Consequently, physical or network location alone can no longer provide sufficient evidence that a user, device, or application should be trusted.

The growing adoption of cloud computing has further accelerated this change. Organizations frequently distribute enterprise workloads among several cloud providers while continuing to maintain important applications within private infrastructure. Sensitive information may therefore move among internal systems, external cloud platforms, SaaS applications, mobile endpoints, and partner organizations. At the same time, enterprises increasingly rely on APIs and microservices that communicate automatically without direct human interaction. This creates a security environment in which both human and machine identities require reliable authentication, authorization, monitoring, and lifecycle

management. A compromised account, stolen access token, malicious service identity, or insecure endpoint can potentially provide an attacker with access to distributed organizational resources even when traditional perimeter defenses remain operational. Zero-trust security has emerged as an important architectural response to these challenges. The fundamental principle of zero trust is that no user, device, application, or service should receive implicit trust solely because of its network location or previous successful authentication. Every access request should instead be explicitly evaluated before access is granted. Identity, authentication strength, device security posture, requested resource, user behavior, environmental context, organizational policy, and current risk conditions can all contribute to the access decision. Trust is therefore treated as a dynamic condition rather than a permanent characteristic assigned to an internal user or system. This approach is particularly suitable for large-scale distributed enterprises because resources and identities frequently exist outside conventional organizational network boundaries.

Identity becomes a central security element within such an architecture. Modern enterprise identities extend beyond employees and administrators. They also include contractors, customers, external partners, service accounts, applications, APIs, virtual machines, microservices, cloud workloads, automated processes, and connected devices. Each identity may require different privileges and security controls depending on its purpose and organizational relationship. Therefore, effective zero-trust implementation requires a comprehensive Identity and Access Management framework capable of creating, maintaining, authenticating, authorizing, monitoring, modifying, and revoking identities throughout their lifecycle. Identity management must also operate consistently across heterogeneous enterprise systems so that access decisions remain coordinated even when applications are distributed across multiple platforms. Authentication is the first important component of identity-centered zero-trust security. Conventional password-based authentication alone provides limited protection because passwords can be stolen through phishing, malware, credential reuse, brute-force attacks, social engineering, or data breaches. Multi-Factor Authentication strengthens identity verification by requiring multiple independent forms of evidence. These may include passwords, security tokens, mobile authentication applications, hardware keys, certificates, biometric characteristics, or other authentication mechanisms. Strong authentication significantly reduces the probability that possession of a single compromised credential will provide immediate access to sensitive enterprise resources. However, successful authentication at the beginning of a session should not automatically establish permanent trust for the complete duration of the user's activity.

Continuous verification therefore represents an important characteristic of zero-trust architecture. The security condition of an authenticated session may change after access has been established. A previously trusted device may become compromised, unusual resource-access behavior may emerge, a user may suddenly connect from an unexpected location, or a session token may be stolen. Large-scale enterprise platforms must therefore continuously or periodically reassess identity confidence, device posture, behavior, and contextual risk. If suspicious conditions are detected, the system may request additional authentication, reduce the user's privileges, block access to sensitive resources, revoke authorization tokens, or terminate the session. This continuous evaluation transforms access control from a one-time login decision into an adaptive security process. Device trust is also essential because valid user credentials do not guarantee that the device being used is secure. An employee may authenticate from a fully managed enterprise laptop in one situation and from an unknown or vulnerable endpoint in another. Although the human identity remains unchanged, the associated security risks are different. Device-security information such as operating-system integrity, patch status, endpoint protection, encryption, ownership, configuration compliance, and security health can therefore contribute to trust evaluation. Sensitive enterprise applications may require both a strongly authenticated identity and a compliant enterprise device before access is permitted.

Literature Review

Ward and Beyer (2014) introduced Google's BeyondCorp approach as an alternative to conventional perimeter-oriented enterprise security. The authors argued that traditional firewall-based architectures create significant risk because an attacker who successfully penetrates the organizational perimeter may subsequently obtain comparatively easy access to privileged internal resources. BeyondCorp therefore removed the assumption that an internal corporate network should automatically be trusted and shifted application access toward user- and device-oriented controls. The approach represented an important early practical foundation for zero-trust enterprise security by demonstrating that corporate applications can be protected according to identity and device information rather than the originating network. Hu et al. (2014) formalized Attribute-Based Access Control through NIST SP 800-162 and established a framework in which access decisions can be derived from attributes associated with subjects, resources, requested operations, and environmental conditions. ABAC provides greater fine-grained authorization than purely role-oriented mechanisms because access can depend on identity characteristics, resource classification, location, time, device information, and operational context. This principle is particularly important for zero-trust enterprise systems because successful authentication alone should not automatically grant broad access. Instead, every request can be evaluated against attributes and organizational policy before access is permitted.

Sakimura et al. (2014) developed OpenID Connect as an interoperable identity layer over OAuth 2.0. The specification allows applications to verify an end user's identity through a trusted authorization server and obtain standardized claims describing the authenticated subject. OpenID Connect has important implications for large distributed enterprise systems

because multiple applications can rely on a common identity provider rather than maintaining separate authentication infrastructures. Federated identity also supports Single Sign-On and centralized identity governance, although strong token protection, lifecycle management, and authorization remain necessary when federation is used across sensitive enterprise applications. Schaffer (2015) examined continuous authentication using mobile devices and highlighted a fundamental limitation of conventional security mechanisms: traditional authentication verifies an identity primarily at a particular point in time. The increasing availability of device sensors and behavioral characteristics creates opportunities for continuously reassessing whether the current user remains consistent with the initially authenticated identity. This concept directly supports zero-trust architecture because trust should not remain permanently fixed after successful login. Continuous verification provides a mechanism for detecting changes in identity confidence during an active enterprise session.

Dos Santos et al. (2016) investigated risk-based access control for cloud computing environments. The authors observed that conventional access-control models may be too rigid for dynamic and collaborative cloud systems and proposed an architecture that extends XACML-based enforcement with risk assessment. Their framework combined conventional access-control mechanisms with dynamically calculated risk and experimentally evaluated the resulting implementation. The work demonstrated that authorization can incorporate current security risk rather than relying entirely on predetermined static permissions, providing an important foundation for adaptive zero-trust access decisions. Osborn et al. (2016) described the practical design and deployment of BeyondCorp within Google. Unlike conventional enterprise-security systems that determine trust partly according to physical location or network origin, BeyondCorp evaluated information concerning the user, device, and current device state. It dynamically assigned access tiers and treated both internal and external networks as untrusted. The deployment demonstrated that identity- and device-based security can operate at enterprise scale and that policy decisions can be separated from traditional network-location assumptions. This study is highly relevant to large distributed enterprise platforms because it provides practical evidence that zero-trust principles can be incorporated into real organizational infrastructure.

Shahzad and Singh (2017) investigated continuous authentication and authorization in Internet of Things environments. The researchers argued that one-time authentication is often unsuitable in highly dynamic environments because users may not maintain continuous control over connected devices after initial authentication. Their work emphasized the importance of repeatedly evaluating both authentication and authorization while a session is active. Although the study focused on IoT systems, its central principle applies directly to distributed enterprise platforms containing users, endpoints, cloud workloads, services, and connected devices whose security states may change during operation. Grassi et al. (2017) developed the NIST Digital Identity Guidelines, which established comprehensive guidance for digital identity proofing, enrollment, authentication, authenticator lifecycle management, federation, and identity assertions. The guidelines demonstrated that secure identity management must extend beyond password validation and encompass the complete identity lifecycle. For zero-trust enterprise architecture, these principles are particularly important because every access decision ultimately depends on the reliability of identity information, authentication assurance, and the mechanisms used to maintain or revoke credentials over time.

Sartoli and Siami Namin (2019) investigated adaptive access-control policy modelling using Answer Set Programming. The authors highlighted limitations of conventional monotonic policy systems when contextual conditions change and exceptional rules need to be introduced dynamically. Their approach supported modification of policy decisions according to newly available contextual information and runtime exceptions. The research demonstrates that large distributed enterprise environments require authorization policies capable of adapting to operational changes rather than relying exclusively on permanently fixed rules. Such adaptive policy behaviour complements zero-trust principles by enabling access decisions to respond to changing identity, environment, and risk conditions. Nespoli et al. (2019) proposed PALOT, a context-aware framework for continuous and non-intrusive authentication and authorization. The system incorporated ontology-based contextual modelling together with behavioral-pattern analysis and a confidence-management component. Experimental evaluation indicated the feasibility and scalability of the proposed mechanism. The study demonstrated that user behavior and surrounding contextual information can contribute to continuous confidence assessment rather than restricting authentication to login-time credentials. This provides direct support for the proposed architecture's continuous trust and session-monitoring functions.

Rose et al. (2020) formalized Zero Trust Architecture through NIST SP 800-207. The publication represents a major development in zero-trust research because it explicitly states that implicit trust should not be granted to users or assets solely according to physical or network location or enterprise ownership. Authentication and authorization of subjects and devices are performed before a session to an enterprise resource is established. NIST further shifted the security focus from network segments toward individual users, assets, services, and enterprise resources. This model provides the principal architectural foundation for the present research. Rose et al. (2020) also defined major logical components involved in zero-trust implementation, including policy decision and policy enforcement functions. This separation is particularly useful for large distributed enterprises because security governance can remain logically coordinated while access decisions are enforced near individual applications and resources. NIST additionally recognized remote users, BYOD, and cloud-hosted resources as important drivers for the transition from traditional network-perimeter protection toward resource-

centric security. The work therefore provides a direct connection between zero-trust architecture and contemporary distributed enterprise computing.

Teerakanok et al. (2021) examined the migration of traditional enterprise infrastructure toward Zero Trust Architecture. The authors reviewed the zero-trust concept using NIST SP 800-207 and emphasized that organizations cannot normally migrate from perimeter security to zero trust through a single technology replacement. Migration requires consideration of legacy systems, implementation stages, existing policies, organizational requirements, and operational challenges. Their study is particularly significant for practical enterprise adoption because many organizations must introduce zero-trust security gradually while continuing to operate existing applications and IAM infrastructure. He et al. (2022) conducted a comprehensive survey of Zero Trust Architecture and identified identity authentication, access control, and trust assessment as three major technical foundations. Their analysis emphasized the zero-trust principle of requiring verification regardless of whether a requester is located inside or outside the organizational network. The authors also discussed existing methods, advantages, limitations, implementation challenges, and future research requirements. They observed that practical zero-trust adoption remained relatively immature and highlighted trust assessment and dynamic access control as important areas requiring further development.

Nahar and Gill (2022) investigated Identity and Access Management from an enterprise-architecture perspective and developed an integrated IAM metamodel and pattern system. The study treated IAM as a fundamental component of secure enterprise architecture and proposed a technology- and business-domain-independent metamodel based on ontology and graph modelling. Eight IAM patterns were subsequently developed for recurring identity-management problems. Their research is particularly relevant to zero-trust enterprise architecture because it demonstrates the importance of integrating identity governance and access management with broader enterprise-system architecture rather than implementing authentication as an isolated security function. Arshad et al. (2022) reviewed Semantic Attribute-Based Access Control and identified interoperability limitations associated with conventional ABAC in distributed and heterogeneous environments. Attributes representing equivalent concepts may use different names or structures across enterprise systems, potentially preventing policies from interpreting them consistently. Semantic ABAC combines semantic technologies with attribute-based authorization to improve policy expressiveness and attribute interoperability. This problem is important for large-scale distributed enterprise platforms because identity and resource information may originate from multiple directories, applications, cloud providers, and administrative domains.

The development of zero-trust architecture continued significantly in 2023 as enterprise systems increasingly adopted cloud-native and multi-cloud application models. Chandramouli and Butcher (2023), through NIST SP 800-207A, extended zero-trust principles specifically to cloud-native applications operating across multiple locations and cloud environments. The publication emphasized a shift from network-parameter-based security toward user, application, service, and workload identities. It recommended identity-tier and network-tier policies together with API gateways, sidecar proxies, service-identity infrastructure, and dynamic authentication and authorization mechanisms. This work is highly relevant to the present study because large distributed enterprise platforms increasingly depend on microservices and machine-to-machine communication rather than only human-user access. Chandramouli and Butcher (2023) further emphasized that zero-trust policy enforcement in cloud-native systems requires continuous monitoring and telemetry regarding users, services, enterprise directories, resources, and access events. Such telemetry can be used to adjust access privileges or require step-up authentication when operational conditions change. Their model therefore extends zero-trust identity concepts from user-level access toward comprehensive service-identity and workload-oriented authorization in distributed multi-cloud environments.

Methodology

This study adopts a Systematic Literature Review (SLR) combined with a Conceptual Experimental Framework to investigate the design and operation of a Zero-Trust Identity and Access Architecture for securing large-scale distributed enterprise systems. This methodological organization follows the uploaded reference paper, where a systematic review is combined with a conceptual framework, mathematical functions, an algorithmic strategy, and subsequent performance evaluation.

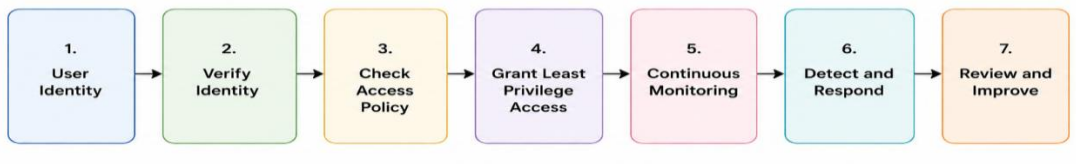

Figure 1. Methodology for Zero-Trust Identity and Access Architecture in Large-Scale Distributed Enterprise Systems

Figure 1 presents the proposed methodology for implementing a Zero-Trust Identity and Access Architecture in large-scale distributed enterprise systems. The process begins with User Identity, where users or devices request access to enterprise resources. The Verify Identity stage authenticates users using secure identity verification mechanisms. Next, Check Access Policy evaluates authorization rules based on organizational security policies. If the request satisfies the required conditions, Grant Least Privilege Access provides only the minimum permissions necessary to perform the requested task. The Continuous Monitoring stage observes user activities, device behavior, and session integrity throughout the access period. Any suspicious or unauthorized activity is handled in the Detect and Respond phase through automated security responses and threat mitigation mechanisms. Finally, Review and Improve analyzes security logs and policy effectiveness to continuously refine access control strategies and enhance the resilience of the enterprise security architecture.

Proposed Algorithm

Step 1: Identity and Access Request Acquisition

The algorithm begins when a human user, application, API, service account, device, microservice, or cloud workload requests access to an enterprise resource.

The access request is represented as:

$$AR_i = \{I, D_i, D, V_i, W, I_i, R, ES_i ACT_i LOC_i T_i CTX_i\}$$

where:

ID_i= User or service identity
DV_i= Device identity
WI_i= Workload or application identity
RES_i= Requested resource
ACT_i= Requested operation
LOC_i= Location or network context
T_i= Request time
CTX_i= Additional contextual information

The framework retrieves the identity profile:

$$IP_i = \{R, o, l, e_i, Attribute_i Privilege_i Status_i History_i\}$$

If the identity cannot be resolved, is suspended, expired, or revoked:

$$Status_i \neq Active \Rightarrow Deny$$

Thus, only valid and currently authorized enterprise identities proceed to the next stage.

Step 2: Strong Authentication Verification

The identity is authenticated using available authentication mechanisms.

Authentication may involve:

Password validation
One-Time Password
Security token
Hardware security key
Digital certificate
Biometric verification
Federated identity
Single Sign-On
Adaptive Multi-Factor Authentication

The authentication confidence is calculated conceptually as:

$$AC_i = f(K, F_i, P, F_i, BF_i CF_i FF_i)$$

where:

KF_i= Knowledge-factor confidence
PF_i= Possession-factor confidence
BF_i= Biometric confidence
CF_i= Certificate confidence
FF_i= Federated-authentication confidence

If authentication fails:

$$AC_i < T_{auth} \Rightarrow Deny$$

If authentication succeeds but the requested resource is highly sensitive, the algorithm may require step-up MFA.

$$ResourceRisk_i = High \wedge AC_i < T_{strong} \Rightarrow Challenge$$

This prevents low-assurance authentication from being used for highly sensitive operations.

Step 3: Device and Workload Trust Verification

After identity authentication, the security state of the device or workload is evaluated.

For user endpoints, the Device Trust Score is:

$$DTS_i = w_1 OS_i + w_2 PATCH_i + w_3 ENC_i + w_4 EDR_i + w_5 COMP_i$$

where:

OS_i= Operating-system integrity
$PATCH_i$= Security-patch status
ENC_i= Encryption state
EDR_i= Endpoint-security status
$COMP_i$= Enterprise compliance

For machine identities, a Workload Trust Score is calculated:

$$WTS_i = f(S, I, D_i, C, ERT_i IMG_i ATTEST_i RUNTIME_i)$$

where:

SID_i= Service identity
$CERT_i$= Certificate validity
IMG_i= Image integrity
$ATTEST_i$= Attestation result
$RUNTIME_i$= Runtime security state

The resulting endpoint state is classified as:

$$TrustState \in \{T, r, ustedConditionalUntrusted\}$$

An untrusted endpoint requesting a critical enterprise resource may be immediately rejected.

Step 4: Contextual and Behavioral Risk Assessment

The algorithm analyzes current access conditions.

The relevant characteristics include:

Current location
Network type
Login time
Authentication history
Device changes
Historical access patterns
Requested resource sensitivity
Data-transfer behavior
Threat intelligence
Privilege-sensitive activity

Behavioral deviation is represented as:

$$BD_i = d(B_{current}, B_{historical})$$

A larger deviation indicates a greater difference between current and expected enterprise behavior.

The contextual risk score is:

$$CRS_i = v_1 LR_i + v_2 DR_i + v_3 BD_i + v_4 RR_i + v_5 TI_i$$

where:

LR_i= Location/network risk
DR_i= Device risk
BD_i= Behavioral deviation

RR_i= Resource risk
TI_i= Threat-intelligence risk

The risk is conceptually categorized as:

$$Risk_i \in \{L, o, w, ModerateHighCritical\}$$

A critical-risk request can be denied before resource access occurs.

Step 5: Dynamic Trust Calculation

The system calculates a combined trust score instead of treating successful authentication as permanent trust.

$$TS_i = \alpha AC_i + \beta DTS_i + \gamma WTS_i + \delta BC_i - \epsilon CRS_i$$

where:

TS_i= Dynamic Trust Score
AC_i= Authentication confidence
DTS_i= Device trust
WTS_i= Workload trust
BC_i= Behavioral confidence
CRS_i= Contextual risk

The trust level can be classified as:

Trust Level	Interpretation
High	Strong identity and contextual confidence
Moderate	Acceptable with additional monitoring
Low	Additional verification required
Critical Risk	Access should be denied

The trust requirement is also linked to resource sensitivity:

$$RequiredTrust_j = f(ResourceSensitivity_j)$$

Therefore, highly sensitive resources require stronger trust than low-risk enterprise information.

Results and Findings

Authentication Security Performance

Table 1. Comparative Authentication Characteristics

Authentication Method	Identity Assurance	Credential-Theft Resistance	Adaptive Capability	Enterprise Suitability
Password Only	Low	Low	Very Low	Moderate
Password + OTP	High	High	Moderate	High
Certificate Authentication	Very High	Very High	Moderate	Very High
Federated SSO	High	High	High	Very High
MFA	Very High	Very High	High	Very High
Adaptive MFA	Very High	Very High	Very High	Very High

Analysis

Table 1 demonstrates that password-only authentication provides insufficient assurance for sensitive distributed enterprise environments. Multi-Factor Authentication improves resistance to compromised credentials by requiring independent verification factors. Federated SSO improves enterprise usability and centralizes authentication, while certificate-based authentication provides strong device and workload identity verification. Adaptive MFA provides the greatest flexibility

because additional authentication can be requested only when resource sensitivity or contextual risk requires stronger assurance.

Device and Workload Trust

Table 2. Device and Workload Verification Capability

Trust Parameter	Security Importance
Device Identity	Very High
Patch Status	High
Endpoint Security	Very High
Disk/Device Encryption	High
Compliance Status	Very High
Service Identity	Very High
Workload Certificate	Very High
Runtime Integrity	Very High
Attestation	Very High

Analysis

Table 2 indicates that zero-trust security must evaluate both user endpoints and machine workloads. A valid human identity does not guarantee that the associated device is safe. Similarly, trusted user authentication does not protect service-to-service communication when workload identities are compromised. Therefore, the proposed architecture incorporates endpoint compliance, service certificates, workload integrity, and runtime status as independent trust inputs.

Access Control Performance

Table 3. Comparative Authorization Models

Authorization Model	Role Awareness	Attribute Awareness	Context Awareness	Dynamic Adaptability
ACL	Low	Low	Very Low	Very Low
RBAC	Very High	Low	Low	Low
ABAC	High	Very High	Very High	High
Risk-Based Access	Moderate	High	Very High	Very High
Proposed Hybrid Zero-Trust Model	Very High	Very High	Very High	Very High

Analysis

Table 3 demonstrates that Role-Based Access Control remains valuable for organizational permission management but provides limited adaptability when security conditions change. ABAC introduces fine-grained subject, resource, action, and environmental attributes. Risk-aware authorization adds current contextual conditions. The proposed hybrid architecture combines these capabilities so that business roles establish baseline eligibility while attributes, risk, trust, and resource sensitivity determine the final decision.

Conclusion and Discussion

The present study investigated Zero-Trust Identity and Access Architecture for Securing Large-Scale Distributed Enterprise Systems through a systematic examination of research and standards published between 2014 and 2023 and through the development of the proposed Zero-Trust Identity and Access Architecture for Distributed Enterprise Systems (ZTIAA-DES). The study focused on the growing limitations of perimeter-oriented enterprise security and examined how identity, authentication, device and workload verification, contextual trust, fine-grained authorization, continuous monitoring, and lifecycle governance can be integrated into a unified security architecture. The overall approach is consistent with the

detailed discussion structure of the reference paper, which relates the proposed framework back to the major problems, capabilities, and limitations identified throughout the study. The central conclusion is that network location can no longer provide an adequate basis for enterprise trust. Modern organizational resources are distributed across internal data centers, public and private clouds, SaaS environments, mobile endpoints, remote offices, APIs, microservices, partner platforms, and multi-cloud infrastructures. Users and applications may legitimately access sensitive enterprise resources without ever entering a conventional organizational network. At the same time, attackers possessing valid credentials or compromising internal endpoints may operate from locations traditionally considered trusted. Consequently, the conventional distinction between trusted internal networks and untrusted external networks becomes progressively less meaningful. Early practical work on BeyondCorp demonstrated this transition by shifting enterprise application access from network location toward user and device characteristics. Ward and Beyer described the removal of privileged-network assumptions, while later deployment work demonstrated that user identity and device state could be used to control enterprise access at organizational scale. This established an important practical foundation for identity-centered zero-trust architecture. NIST SP 800-207 subsequently formalized zero-trust principles by removing implicit trust based solely on network location, affiliation, or ownership and by requiring authentication and authorization before access to enterprise resources is established. Zero trust therefore represents more than stronger authentication. It changes the underlying trust model of enterprise security by requiring each access request to be explicitly evaluated against current identity, resource, device, and policy conditions.

References

1. Ward, R., & Beyer, B. (2014). BeyondCorp: A new approach to enterprise security. *;login: The USENIX Magazine, 39*(6), 6–11.
2. Hu, V. C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., & Scarfone, K. (2014). *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*. NIST Special Publication 800-162. https://doi.org/10.6028/NIST.SP.800-162
3. Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., & Mortimore, C. (2014). *OpenID Connect Core 1.0*. OpenID Foundation.
4. Schaffer, K. B. (2015). Expanding continuous authentication with mobile devices. *Computer, 48*(11), 92–95. https://doi.org/10.1109/MC.2015.333
5. Dos Santos, D. R., Marinho, R., Schmitt, G. R., Westphall, C. M., & Westphall, C. B. (2016). A framework and risk assessment approaches for risk-based access control in the cloud. *Journal of Network and Computer Applications, 74*, 86–97. https://doi.org/10.1016/j.jnca.2016.08.013
6. Gummadi, V. P. K. (2023). MuleSoft batch processing: High-volume streaming architecture. Computer Fraud & Security, 2023(12), 50–57. https://doi.org/10.52710/cfs.886.
7. Grassi, P. A., Garcia, M. E., & Fenton, J. L. (2017). *Digital Identity Guidelines*. NIST Special Publication 800-63-3. https://doi.org/10.6028/NIST.SP.800-63-3
8. Shahzad, M., & Singh, M. P. (2017). Continuous authentication and authorization for the Internet of Things. *IEEE Internet Computing, 21*(2), 86–90. https://doi.org/10.1109/MIC.2017.33
9. Sartoli, S., & Siami Namin, A. (2019). Modeling adaptive access control policies using answer set programming. *Journal of Information Security and Applications, 44*, 49–63. https://doi.org/10.1016/j.jisa.2018.10.007
10. Nespoli, P., Zago, M., Huertas Celdrán, A., Gil Pérez, M., Gómez Mármol, F., & García Clemente, F. J. (2019). PALOT: Profiling and authenticating users leveraging Internet of Things. *Sensors, 19*(12), 2832. https://doi.org/10.3390/s19122832
11. Zaheer, Z., Chang, H., Mukherjee, S., & Van der Merwe, J. (2019). eZTrust: Network-independent zero-trust perimeterization for microservices. *Proceedings of the ACM Symposium on SDN Research*, 49–61. https://doi.org/10.1145/3314148.3314349
12. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture*. NIST Special Publication 800-207. https://doi.org/10.6028/NIST.SP.800-207
13. Campbell, M. (2020). Beyond zero trust: Trust is a vulnerability. *Computer, 53*(10), 110–113. https://doi.org/10.1109/MC.2020.3011081
14. Teerakanok, S., Uehara, T., & Inomata, A. (2021). Migrating to Zero Trust Architecture: Reviews and challenges. *Security and Communication Networks, 2021*, 9947347. https://doi.org/10.1155/2021/9947347

15. Bertino, E., & Brancik, K. (2021). Services for zero trust architectures—A research roadmap. *2021 IEEE International Conference on Web Services (ICWS)*, 14–20. https://doi.org/10.1109/ICWS53863.2021.00016
16. Jaiswal, I. A. (2023). Zero-trust identity and access architecture for securing large-scale distributed enterprise systems. Computer Fraud & Security.
17. Nahar, K., & Gill, A. Q. (2022). Integrated identity and access management metamodel and pattern system for secure enterprise architecture. *Data & Knowledge Engineering, 140*, 102038. https://doi.org/10.1016/j.datak.2022.102038
18. Arshad, H., Johansen, C., & Owe, O. (2022). Semantic Attribute-Based Access Control: A review on current status and future perspectives. *Journal of Systems Architecture, 129*, 102625. https://doi.org/10.1016/j.sysarc.2022.102625
19. Chandramouli, R., & Butcher, Z. (2023). *A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Cloud Environments*. NIST Special Publication 800-207A. https://doi.org/10.6028/NIST.SP.800-207A
20. Kang, H., Liu, G., Wang, Q., Meng, L., & Liu, J. (2023). Theory and application of zero trust security: A brief survey. *Entropy, 25*(12), 1595. https://doi.org/10.3390/e25121595