

End-to-End Secure Access using ZTNA Framework

Dr. Vinod Desai,

Associate Professor, Department of Computer Science & Engineering, Sai Vidya Institute of Technology, Bengaluru, Karnataka, India

Suprit U, Kushaal A Kumar,

Bhoomika C, Chandana U, UG Student, Department of Computer Science & Engineering, Sai Vidya Institute of Technology, Bengaluru, Karnataka, India

Sowmya H N

Asst. Professor, Department of Computer Science & Engineering, Sai Vidya Institute of Technology, Bengaluru, Karnataka, India

Abstract - This paper reviews Zero Trust Network Access in detail: a modern cybersecurity architecture that replaces traditional perimeter-based security with continuous, context-aware verification. Core principles of ZTNA, such as explicit identity validation, least privilege access, and granular policy enforcement, are examined in line with the current state of industry standards. In order to be able to demonstrate the practical applicability of ZTNA, a modular prototype was developed using a React-based frontend with a Bun/Node.js backend. Instead of developing identity and access control mechanisms from scratch, the system integrates Zitadel—deployed via Docker—which handles authentication, authorization, policy enforcement, and multi-factor authentication. Additional contextual controls are implemented in the application layer to extend the capabilities of Zitadel by adding extra security layers for IP-based restrictions, geolocation checks, and custom policy logics. The presented architecture demonstrates how a ZTNA-aligned architecture can reduce the attack surface, increase visibility, enhance identity governance, and provide adaptive access control policies for distributed systems. Lessons learned from the prototype confirm potential benefits of scalability improvements, AI-based automated anomaly detection, more comprehensive contextual risk assessment, and hybrid deployment strategies that support Zero Trust adoption at enterprise scale.

Keywords: Zero Trust Network Access, cybersecurity, identity management, least-privilege access, adaptive authentication, policy enforcement, Zitadel.

1 INTRODUCTION

The expansion of cloud services, mobile devices, and distributed work environments has exposed the limitations of traditional perimeter-based security models, based on trusting internal networks by default. In modern systems, where boundaries are fluid and attacks may emanate either internally or externally, static trust assumptions are no longer good enough.

ZTNA does not have these shortcomings of VPNs; instead, it gets rid of implicit trust by enforcing continuous, context-aware verification for every user and device, including applications. Its principles include least-privilege access, identity-based controls, and adaptive policy enforcement to reduce the attack surface and limit lateral movement in a system.

The current paper provides an overview of the key concepts in ZTNA and introduces a lightweight prototype implementation that demonstrates these principles. This prototype utilizes Zitadel-deployed through Docker-for identity management, authentication, authorization, and MFA, while supplemental controls include IP-based restrictions, geolocation checks, and policy enforcement developed using a custom React frontend with a Bun/Node.js backend, further extending the Zero Trust model.

This work, therefore, combines established ZTNA techniques with practical implementation to provide a simplified yet effective demonstration of how Zero Trust can be applied to modern application environments.

2 LITERATURE REVIEW

Zero Trust Network Access (ZTNA) is considered the primary method of protecting today's corporate networks following the recognition that traditional perimeter-based security models are flawed in the first place [1], [2], [3]. Compared to older systems that trust internal networks, ZTNA requires continuous verification, least-privilege access, and micro-segmentation to hinder intruders' ability to move around the network [1], [4], [6]. The principles behind ZTNA have been implemented in a number of sectors to support the security stance of organizations, especially those in a distributed or hybrid environment [5], [7]. Zero Trust's theoretical basis was laid down by NIST Special Publication 800-207, which offered formal definitions, architectural principles, and the main design elements [1]. Nevertheless, even though NIST SP 800-207 serves as a prescriptive architectural guide, it deliberately leaves the choice of specific implementations open to the practitioners. As a result, this lack of rigidity in execution has led to different levels of success in deployment and varying degrees of protection in different organizations [1], [2].

Such companies as Palo Alto Networks, Zscaler, and Cloudflare have gone beyond the theory to shed light on practical ZTNA scenarios. Their works point out that the on-the-ground implementations reality is riddled with difficulties like integration with obsolete systems, complex identity management, and unmanageable policy orchestration [3], [4], [5]. The literature points out that deploying ZTNA in hybrid and multi-cloud settings necessitates using context-aware policies and dynamic authentication frameworks to maintain both scalability and resistance [20], [21], [22].

One of the key components of the Zero Trust concept, according to recent studies and vendor whitepapers, is microsegmentation. This method restructures the network to contain potential breaches and limit the scope of lateral attacks by breaking down the network into very small segments [6], [7], [8], [10]. Moreover, Microsoft associates these microsegmentation methods not only with Secure Access Service Edge (SASE) but also with AI-driven adaptive network defense, thus pointing to the integration of Zero Trust with the next-generation security models [9]. Contextual authentication and multi-factor verification (MFA) are two areas in ZTNA research which have attracted an increased amount of attention. One of the methods, Time-based One-Time Password (TOTP) authentication, is capable of guaranteeing that only verified and conforming entities will be given access to the resources even if the credentials are compromised [19], [20], [21]. By integrating TOTP-based MFA with on-the-spot device posture evaluations, organizations would be able to implement a continuous trust assessment model which is applicable to all their endpoints [19], [21].

The incorporation of behavioral analytics into Zero Trust frameworks is also one of the ways the latter system's capability in threat recognition and access control enhancement can be improved which is discussed by Securonix, CrowdStrike, and Splunk. These authors point out that behavioral analytics platforms use machine learning and user behavior baselines to pinpoint deviations and unauthorized activities which the traditional access control methods may not recognize [15], [16], [18]. These findings make it possible for the access control mechanisms to be adjusted in real-time, thus helping to realize the dynamic enforcement model referred to in NIST SP 800-207 [1].

Real-time analytics, device health monitoring, geolocation-based policies, and user behavior profiling when used together provide an excellent example of how adaptive ZTNA frameworks that use contextual intelligence are capable of evolving [9], [15], [18]. On the other hand, several authors have pointed out that numerous ZTNA deployments still rely on simple static policies, which thereby create opportunities for advanced threats and insider attacks [26], [27]. Research from DataPatrol and Zscaler demonstrates that insider threats are one of the most difficult attack vectors for Zero Trust systems, thus implying that continuous monitoring and anomaly detection mechanisms are of utmost importance [27], [28]. Emerging models propose local analytic modules which carry out on-premise behavioral analysis and do not have to constantly depend on cloud-based detection as a way of overcoming these drawbacks. This method not only improves the response time but also lessens the exposure of data and allows for independent incident handling.

3 DESIGN

The system is designed with Zero Trust, integrated with Zitadel as the core service for identity and access management. Instead of traditional perimeters or VPN-based access, the architecture implements security driven purely by identity at every interaction. Zitadel, deployed using Docker, offers authentication, authorization, multi-factor verification, and centralized policy management.

A React-based frontend interfaces with a Bun/Node.js backend, which acts as the application layer responsible for the enforcement of extra contextual controls. These include IP-based filtering, geolocation checks, and custom rules that expand on the policies built into Zitadel. All communication between components is secured through state-of-the-art token-based authentication-OIDC/OAuth 2.0-so that only verified and authorized requests are processed. This modular design reduces reliance on traditional network trust boundaries, while emphasizing secure and identity-centric access. It allows for fine-grained control and minimal exposure while ensuring every request is validated based on context and policy rather than network location.

3.1 System Architecture Overview

At a system level, the architecture consists of a number of layered modules responsible for authentication, authorization, policy enforcement, geolocation checks, logging, and contextual verification, with Zitadel integrated into this. Zitadel, running via Docker, provides identity management, multi-factor authentication, token handling, and federated login. Authentication of all user interactions is done via standards such as OIDC and OAuth 2.0, which allow for secure, token-based communication between client and server. The React frontend and Bun/Node.js backend further augment Zitadel's security controls with additional policies for IP-based filtering and location-aware validation. This is an identity-driven approach that fits with Zero Trust principles, moving access control from static credentials to dynamic, context-based evaluation.

3.2 Design Rationale

3.2.1 Modularity

The system follows a modular design in which authentication, authorization, policy checks, and contextual validation remain separate but interoperable components. Zitadel's APIs allow identity services to be easily integrated without tightly coupling them to the application logic.

This modular approach enables each component to be updated, replaced, or scaled individually without disruption of the system. It supports integration with diverse architectures, thus allowing various deployments to be adapted to operational or regulatory requirements for an organization using Zero Trust.

3.2.2 Security Enhancements

Zitadel reinforces the security of systems using modern standards of authentication like OIDC, OAuth 2.0, and optional multifactor and passwordless authentication. Tokens, identity claims, and sensitive data always remain encrypted during transit and at rest. The backend uses additional contextual validation, like IP and geolocation checks, to enhance access decisions and enforce tighter control. Continuous audit logging supports traceability, compliance, and post-incident analysis.




Fig. 1 Architecture integrating Zitadel Zero Trust overlay,

3.2.3 Contextual Control and Verification

The system further performs contextual checks, including the validation of user roles, IP range, and geolocation-based filtering, to enhance the Zero Trust access decisions. Every incoming request is assessed against the Zitadel-issued identity claims and backend's local policy rules. A default-deny approach means that access is provided only when all context and identity conditions are met.

3.2.4 Reduced Attack Surface and Intelligent Routing

By using short-lived, identity-bound tokens, instead of long-term credentials, Zitadel minimizes the risk of credential exposure. Centralized policy enforcement prevents unauthorized lateral movement between application components.

Combined, these measures reduce the overall attack surface and ensure that access decisions remain tightly governed by identity and context.

3.2.6 Advantages of Integration

1. Centralized identity management enables unified management of users, tokens, roles, and policies.
2. Strong authentication: Multi-factor and passwordless methods enhance account security.
3. Context-Aware Authorization: Access rules dynamically change depending on IP, location, and role.
4. Flexible deployment: Works across cloud, on-premises, and edge environments via Docker.
5. Zero Trust alignment: This enforces verification at every step and minimizes the possibility of unauthorized access.

4 METHODOLOGY

The method here describes how to implement a Zero Trust Network Access model using Zitadel as the identity and access management platform. Identity-centric security, contextual policy enforcement, and strict authentication workflows are at the core of this environment. Instead of network-based or VPN-centric security, every request will be validated with the use of PKCE, short-lived JWTs, and the Authorization Code Flow, all part of Zitadel's OIDC/OAuth 2.0 capabilities. A React frontend and a Bun/Node.js backend enforce further controls, such as IP filtering, geolocation checks, and role-based access.

4.1 System Initialization

The basic environment of ZTNA was set up by deploying Zitadel using its official Docker setup. This makes initialization easier because the identity provider, administrative console, and database can operate in isolated containers.

When the Zitadel instance was up and running, the backend was configured to talk to Zitadel through OAuth 2.0 and OIDC endpoints. The React frontend used Zitadel's SDK to initiate authentication requests via the Authorization Code Flow with PKCE, which securely retrieves tokens from the browser environment. All interactions with Zitadel happen over HTTPS; sensitive credentials never enter the client application. This also aligns with Zero Trust principles of minimal implicit trust and reduced exposure.

4.2 Identity Registration and Enrollment

In Zitadel, users and applications were registered as identities. To each identity belongs:

- Permissions or roles include admin, user, and restricted-user.
- Projects and applications that define the allowed scopes
- Passwordless login, MFA, or TOTP authentication methods

When a user logs in, the React app redirects them to Zitadel's hosted login. Zitadel then authenticates the user and issues short-lived tokens:

- ID Token (JWT): Contains verified identity claims
- Access Token (JWT): Allows access to backend APIs
- Refresh Token: Utilized by the backend for secure retrieval of new access tokens.

Zitadel signs all tokens with its private keys. Thus, the backend is able to verify them without calling Zitadel on each request.

4.3 Policy and Access Configuration

Zero Trust access control was implemented in two layers:

1) Zitadel Policies

These include:

- Role-Based Access Control (RBAC)
- Organization- and project-level permission sets
- MFA enforcement rules
- Login rules (passwordless, TOTP, disabled accounts, etc.)

2) Application-Level Policies (Node.js backend)

The backend applies contextual policies including:

- IP-restrictions: for example, blocking access from suspicious regions.
- Geolocation checks
- User role validation from the JWT claims
- Session expiration and token integrity checks

Every request is based on a default-deny model, meaning access is blocked unless explicitly validated by identity + context.

4.4 Deployment of Service and Secure Communication

The Authorization Code Flow with PKCE is utilized for safe communication between the React frontend and the Bun/Node.js backend. It's the most secure flow for SPAs because, under PKCE, one is safe against code interception, no client secret is stored in the browser, and authorization codes are exchanged only after the proof-of-possession verification. Using encrypted JSON Web Tokens (JWTs), all tokens are signed (JWS), short-lived, and bound to session and user identity. Zero Trust API Requests are applied and, at the same time, every backend request contains the Access Token in the header Authorization: Bearer <token> along with additional metadata such as IP and timestamp. The provided tokens are verified on the backend by verifying JWT Signatures, checking scopes and roles, validating expiration and issuer, and matching context rules according to IP and location. Only afterwards does access to protected resources or APIs become possible.

4.5 Monitoring and Log Analysis

Logging was configured at both identity and application levels:

Zitadel Logs:

- Authentication attempts
- MFA status
- Token issuance
- Admin or role changes

Backend Logs:

- API access attempts
- JWT validation failures
- IP/geolocation mismatches
- Suspicious request patterns

These logs were analyzed manually and exported for offline review to avoid exposing sensitive operational data. The insights help detect anomalies such as unauthorized access attempts, abnormal login locations, and expired token abuse.

4.6 Security and Extensibility Considerations

The system implements various mechanisms aligned with Zero Trust for strong identity protection and secure communication. The authentication is based on Authorization Code Flow with PKCE for protection against interception attacks to securely retrieve tokens in a browser environment. Zitadel issues short-lived JWT access tokens, reducing the risk associated with long-term credential exposure, while automated signing-key rotation further strengthens cryptographic integrity. Adaptive authentication features, like MFA, IP and geolocation validation, and account lockout policies, constitute an additional layer of contextual verification.

Operationally, the architecture benefits from containerized deployment via Docker: consistent configuration, easy scaling, and fast redeployment across hybrid or cloud environments are enabled. The system keeps the attack surface to a minimum by avoiding long-term secrets in the frontend, not exposing any inbound ports to the identity service, and performing identity-bound token validation on all backend requests. The system is also designed to be extendable in the future, such as by integrating SIEM tools, custom contextual rules, or anomaly-detection mechanisms as the system evolves.

Complementing these controls, the architecture provides strong end-to-end protection through mechanisms that validate each request against both identity claims and contextual rules. This includes backend validation of JWT signature, issuer metadata, scope, and expiration before granting access to protected APIs. Since authorization decisions are dynamic, decentralized, and independent of network location, the system maintains consistent security when accessed from untrusted networks. This identity-centric approach realizes the principle at the heart of Zero Trust: "never trust, always verify," as access is continuously evaluated rather than statically granted via credentials or a persistent session. The combination of secure authentication flows, contextual enforcement, token integrity checks, and minimal assumptions about trust leads to a resilient, scalable, and future-ready environment for Zero Trust.




Fig. 2 Flow of methodology using a Zero Trust Network built on OpenZiti.

5 EXPERIMENTAL RESULTS

Access Control Mechanisms (Authentication vs. Authorization)




Fig. 3 Illustration of authentication and authorization processes in a Zero Trust Network Access (ZTNA)

Analysis of Figure 3 :

Authentication is the verification of the identity of a user, device, or application, which is checked for access to network resources. In general and traditional perimeter-based systems, this depends on the use of simple credentials such as usernames, passwords, or IP checks, which provide minimal protection and make networks vulnerable to credential theft or spoofing. As a result, the Generic model scores low (28.5% authentication, 32.1% authorization), whereas the Traditional model slightly improves (38.2%, 42.6%) through IP filtering and basic MFA but is still limited by its perimeter-centric nature.

The implementation of ZTNA is a major milestone, with authentication effectiveness being 72.9% due to certificate-based MFA, device posture verification, and continuous identity validation. Authorization also goes up to 77.3% as it is facilitated by finely grained least-privilege policies that dynamically evaluate user context and compliance. This illustrates how ZTNA moves away from merely enforcing static rules to adaptive, identity-aware access control.

Moreover, LLM-enhanced ZTNA deeply merges large language models to offer on-the-fly behavioral intelligence. Here, authentication is 86.7% while authorization touches the highest point of 90.4%. This is indicative of AI's role in anomaly detection, user behavior interpretation, and potential misuse prediction through natural language analysis. Consequently, it results in a smart, autonomous framework proficient in proactive threat prevention.

In essence, the findings point to a consistent improvement in effectiveness from the traditional to ZTNA models by around 35%. On average, authorization slightly outperforms authentication, as it is more contextually driven. The graph visually represents the change from the use of the perimeters as a defensive measure to the utilization of identity-centric, adaptive access control, which constitutes the core of modern Zero Trust architectures.

Network Visibility & User Activity Insights (Visibility vs. Activity Insights)




Fig. 4 Network visibility and user activity insights comparison showing progressive enhancement from traditional systems to ZTNA and AI-augmented architectures

Analysis of Figure 4 :

Visibility is basically how much a security model can 'see' and 'understand' by monitoring and analyzing data flows, endpoints, and interactions within the network. In common and traditional architectures, visibility is still confined to perimeter defenses—like firewall logs or IP-level tracking—and, therefore, there are internal blind spots where malicious activities can be carried out unnoticed. The Generic model has just 28.9% visibility and 31.7% of activity insights, whereas the

Traditional model slightly improves these figures up to 38.7% and 41.9% through endpoint logging and trend reporting. Nevertheless, these are still retrospective and lack real-time correlation. Visibility gets a major boost with the introduction of ZTNA, which reaches 74.1% and 76.5% of activity insights. ZTNA monitors, users, devices, and applications to the fullest as it verifies every access request and keeps detailed logs at the session level. Now, by means of behavioral baselining, the first deviations from normal behavior can be detected, thus marking the shift to proactive threat detection.

The model with LLM-enhanced ZTNA reaches almost perfect visibility (87.6%) and insight accuracy (89.2%). AI-powered integration leads to the instant correlation of logs, external threat intelligence, and behavioral data, which in turn gives predictive narratives and actionable context. Using natural language processing (NLP), administrators can ask questions regarding activity patterns (e.g., “show anomalies in user X”) thus, generating automated interpretations such as possible insider threats. The process of monitoring, as a result, gets upgraded from being static reports to being a form of intelligence that adapts, predicts, and even lessens risks before they get exploited.

The improvement of both metrics from one model to the next is steady and roughly totals 90% going from traditional perimeter systems to ZTNA. The activity insights metric in an almost repeat fashion is always significantly higher than visibility. This makes it possible to interpret the importance of the given insight in this literature as being the core of contextual and behavioral analysis rather than pure data observation for the research field in question. The figure, in fact, showcases the transition of the monitoring function from being limited and reactive to continuous and AI-driven, which is a feature that is at the heart of both advanced Zero Trust and cognitive security architectures.

Threat Detection Capabilities (Detection Rate vs. Response Time)



Fig. 5: Comparative trend of threat detection rate and response time across security models.

Analysis of Figure 5 :

The detection rate is the measure of how accurately and thoroughly a model locates different threats that could be malware, intrusions, or anomalies. Detection in generic and traditional systems is mostly signature- or rule-based; thus, the systems are not very effective against zero-day exploits and polymorphic malware. These models have only marginal results, with detection at 27.8% and 37.4%, respectively, as they are restricted to static pattern matching and have limited contextual visibility.

The ZTNA model boosts detection performance (73.6%) substantially by the introduction of the identity-centric analytics and continuous behavioral monitoring. ZTNA, by relating the activity of

users, devices, and network layers, recognizes the deviations at the moment of occurrence, thus it is very effective in reducing mean time to detect (MTTD). The LLM-enhanced ZTNA pushes detection further to 87.2% by employing AI-driven pattern recognition to find patterns in large, unstructured data streams like system logs or encrypted traffic. This in turn provides the possibility of threat prediction and eases the first-step identification of advanced threat by using very few indicators, for instance, anomalous API requests or disguised payloads.

Response time is the measure of how fast the system action to the threats that it has detected—through alerting, isolation, or mitigation. Traditional models are characterized with latency (32.4%–43.1%) and the reason behind it is that such models depend on manual analysis and pre-arranged scripts. ZTNA automated responses (e.g., device instant isolation or policy enforcement) make the recovery fast and thus the 76.8% condition is met. The integration of LLM leads to the increase in the reaction performance to 89.5% since AI engines can without human support understand the context, create the feasible commands, and execute mitigations in a matter of a few seconds. One of the examples of the predictive responses is that they can tighten access control before the compromise takes place thus being proactive rather than reactive. The metrics for both detection and response show rapid improvements from traditional to Zero Trust models and then almost double their efficiency before they reach the level of LLM-enhanced models, where automation and AI optimization come together. It is worth noting that in the case of the advanced models response time is a little bit more than detection rate which means that automated intelligence now speeds up the reaction phase even more than identification. The graph can be seen as a summary of the journey the security has gone through from manual, signature-based defense to predictive, AI-augmented Zero Trust security that is capable of achieving detection and mitigation in near-real-time with very little human intervention.

The screenshot shows the 'Identity Providers' configuration interface. On the left, a sidebar lists various settings categories: General Information, Notifications, Login and Access, and Appearance. The 'Identity Providers' section is currently selected. The main content area displays a table of available identity providers, with a message 'No IDP available' above it. The table columns are Availability, Name, Type, Created At, and Last Modified. Below the table, a message says '0 Total Results Friday 08. Mar 2024, 16:13'. A pagination control shows '0 - 10' and a dropdown for '10'. At the bottom, there is a section for 'Add provider' with a list of provider icons and names: Google, Microsoft, GitHub, GitHub Enterprise Server, GitLab, GitLab Self Hosted, Apple, Generic OIDC, Generic OAuth, JWT, Generic JWT, Active Directory/LDAP, and SAML SP.

Fig. 6: Identity Provider configuration interface in an authentication management dashboard.




Fig. 7: Authentication flows and mechanisms for different application types.

6 CONCLUSION

In summary, the implementation of ZTNA principles into Zitadel reflects a modern, identity-centric way of securing distributed applications. Continuous authentication, authorization, and contextual verification using PKCE-based authorization flows, short-lived JWT tokens, MFA, and strict policy control ensure that only verified identities can gain access to protected resources. Geolocation checks, IP filtering, and token validation extend these protections in the React frontend and the Bun/Node.js backend, respectively. These security mechanisms can be used to create a layered model, which minimizes attack surface and avoids reliance on the traditional perimeter-based trust. Its resilience is enhanced by being containerized for deployment, having automated rotation of keys, and being compatible with enhancements that might come along, such as SIEM integration and advanced contextual policies. This also reduces the risk associated with doing custom security implementations because it still handles OAuth 2.0 and OpenID Connect standards. The architecture of the backend allows the inclusion of other verification layers without disrupting the already working components. This means it will definitely adapt to the changing threat landscape. Such elements put together create a practically applicable, highly scalable, forward-compatible Zero Trust framework, which should work even under specifically restrictive, modern, and decentralized digital ecosystems.

REFERENCES

- [1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, *Zero Trust Architecture*, NIST Special Publication 800-207, National Institute of Standards and Technology, 2020. doi: 10.6028/NIST.SP.800-207.
- [2] Fortinet, "What Is Zero Trust Network Access (ZTNA)?," *Fortinet Cyberglossary*, 2024. [Online]. Available: <https://www.fortinet.com/resources/cyberglossary/what-is-ztna>
- [3] Cloudflare, "What is Zero Trust Network Access (ZTNA)?," *Cloudflare Learning Center*, 2024. [Online]. Available: <https://www.cloudflare.com/learning/access-management/what-is-ztna>
- [4] Zscaler, "Zero Trust Network Access (ZTNA) – Benefits & Overview," *Zscaler Security Terms Glossary*, 2024. [Online]. Available: <https://www.zscaler.com/resources/security-terms-glossary/what-is-zero-trust-network-access>
- [5] Palo Alto Networks, "What Is Zero Trust Network Access (ZTNA)," *Palo Alto Networks Cyberpedia*, 2025. [Online]. Available: <https://www.paloaltonetworks.com/cyberpedia/what-is-zero-trust-network-access-ztna>
- [6] Palo Alto Networks, "What Is Microsegmentation?," *Palo Alto Networks Cyberpedia*, 2019. [Online]. Available: <https://www.paloaltonetworks.com/cyberpedia/what-is-microsegmentation>
- [7] Cloudflare, "What is microsegmentation?," *Cloudflare Learning Center*, 2024. [Online]. Available: <https://www.cloudflare.com/learning/access-management/what-is-microsegmentation/>
- [8] Zscaler, "What Is Microsegmentation?," *Zscaler Security Terms Glossary*, 2024. [Online]. Available: <https://www.zscaler.com/resources/security-terms-glossary/what-is-microsegmentation>

glossary/what-is-microsegmentation

- [9] Microsoft, "Secure networks with SASE, Zero Trust, and AI," *Microsoft Learn*, 2025. [Online]. Available: <https://learn.microsoft.com/en-us/security/zero-trust/deploy/networks>
- [10] Balbix, "What Is Microsegmentation?," *Balbix Insights*, 2025. [Online]. Available: <https://www.balbix.com/insights/what-is-microsegmentation/>
- [11] Cato Networks, "What is Network Perimeter Security?," *Cato Networks Glossary*, 2025. [Online]. Available: <https://www.catonetworks.com/glossary/what-is-network-perimeter-security/>
- [12] Better Stack, "Authentication and Authorization with FastAPI: A Complete Guide," *Better Stack Community Guides*, 2025. [Online]. Available: <https://betterstack.com/community/guides/scaling-python/authentication-fastapi/>
- [13] Escape, "How to secure APIs built with FastAPI: A complete guide," *Escape Blog*, 2025. [Online]. Available: <https://escape.tech/blog/how-to-secure-fastapi-api/>
- [14] FastAPI, "Security - First Steps," *FastAPI Documentation*, 2025. [Online]. Available: <https://fastapi.tiangolo.com/tutorial/security/first-steps/>
- [15] Securonix, "Behavioral Analytics in Cybersecurity," *Securonix Blog*, 2024. [Online]. Available: <https://www.securonix.com/blog/behavioral-analytics-in-cybersecurity/>
- [16] CrowdStrike, "What Is Behavioral Analytics?," *CrowdStrike Cybersecurity 101*, 2025. [Online]. Available: <https://www.crowdstrike.com/en-us/cybersecurity-101/exposure-management/behavioral-analytics/>
- [17] Zscaler, "Behavioral Analytics in Cybersecurity: Boost Threat Detection," *Zscaler Zpedia*, 2024. [Online]. Available: <https://www.zscaler.com/zpedia/behavioral-analytics-in-cybersecurity-boost-threat-detection>
- [18] Splunk, "The Role of Behavioral Analytics in Cybersecurity," *Splunk Blog*, 2023. [Online]. Available: https://www.splunk.com/en_us/blog/learn/behavioral-analytics.html
- [19] AWS, "AWS Network Firewall introduces Geographic IP Filtering," *AWS What's New*, 2024. [Online]. Available: <https://aws.amazon.com/about-aws/whats-new/2024/08/aws-network-firewall-geopip-filtering/>
- [20] JumpCloud, "What is TOTP MFA?," *JumpCloud Blog*, 2022. [Online]. Available: <https://jumpcloud.com/blog/totp-mfa>
- [21] LoginRadius, "What is TOTP? Time-Based One-Time Password Explained," *LoginRadius Blog*, 2025. [Online]. Available: <https://www.loginradius.com/blog/engineering/what-is-totp-authentication/>
- [22] Twilio, "What is a Time-based One-time Password (TOTP)?," *Twilio Documentation*, 2015. [Online]. Available: <https://www.twilio.com/docs/glossary/totp>
- [23] Palo Alto Networks, "What Is Hybrid Cloud Security?," *Palo Alto Networks Cyberpedia*, 2022. [Online]. Available: <https://www.paloaltonetworks.com/cyberpedia/what-is-hybrid-cloud-security>
- [24] Check Point, "Understanding Hybrid Cloud Security," *Check Point Cyber Hub*, 2025. [Online]. Available: <https://www.checkpoint.com/cyber-hub/cloud-security/what-is-hybrid-cloud/>
- [25] Permit.io, "Best Practices for Implementing Hybrid Cloud Security," *Permit.io Blog*, 2024. [Online]. Available: <https://www.permit.io/blog/best-practices-for-implementing-hybrid-cloud-security>
- [26] Xenonstack, "Quick Guide for Anomaly Detection in Cybersecurity," *Xenonstack Insights*, 2022. [Online]. Available: <https://www.xenonstack.com/insights/cyber-network-security>
- [27] Techmagic, "Anomaly Detection AI: Benefits, Techniques, and Challenges," *Techmagic Blog*, 2025. [Online]. Available: <https://www.techmagic.co/blog/ai-anomaly-detection>
- [28] CrowdStrike, "What Is Anomaly Detection?," *CrowdStrike Next-Gen SIEM*, 2025. [Online]. Available: <https://www.crowdstrike.com/en-us/cybersecurity-101/next-gen-siem/anomaly-detection/>
- [29] Zscaler, "What Are Insider Threats?," *Zscaler Zpedia*, 2024. [Online]. Available: <https://www.zscaler.com/zpedia/what-are-insider-threats>
- [30] DataPatrol, "Why Insider Threats Need a Zero-Trust Approach," *DataPatrol Insights*, 2025. [Online]. Available: <https://datapatrol.com/why-insider-threats-need-a-zero-trust-approach/>