Zero Trust Architecture and Ransomware Mitigation

Ely Johnson
joh21256@morris.umn.edu
Division of Science and Mathematics
University of Minnesota, Morris
Morris, Minnesota, USA

Abstract

Ransomware has become a critical threat to modern enterprises, exploiting excessive privileges and flat network architectures to spread rapidly. Traditional perimeter-based security models are insufficient, as they rely on implicit trust within internal networks. This paper examines how Zero Trust Architecture (ZTA) mitigates ransomware through least privilege access, continuous monitoring, and micro-segmentation. Experimental results show that ZTA can significantly reduce impact, limiting encryption to about 20% of targeted files while preserving most data. Continuous monitoring enables rapid detection (5.3 seconds) with high accuracy (up to 97.2%) and a 78% reduction in false positives. Micro-segmentation further restricts lateral movement, reducing attack paths and network exposure by over 99%. Despite implementation challenges, ZTA strengthens enterprise defenses against ransomware attacks.

Keywords: Zero Trust Architecture, ransomware, Policy Enforcement Point, Policy Decision Point, micro-segmentation, least privilege, continuous monitoring

1 Introduction

Ransomware has emerged as one of the most significant cybersecurity threats confronting modern enterprises. According to the 2025 Sophos State of Ransomware report, 3,400 organizations across 17 countries experienced attacks in the past year, with 50% of incidents resulting in data encryption and 28% involving data exfiltration (the unauthorized transfer of sensitive data out of an organization's systems) [7]. Median ransom demands reached $1.32 million, and organizations spent an average of $1.53 million on recovery costs, excluding ransom payments [7]. High-profile incidents affecting healthcare systems, governments, educational institutions, and critical infrastructure demonstrate the disruptive and financially devastating consequences of large-scale encryption attacks. Beyond operational downtime and ransom payments, organizations often face regulatory penalties, reputational damage, and long-term data exposure. As ransomware campaigns have grown more sophisticated, attackers increasingly combine credential theft, privilege escalation, and lateral movement (where adversaries move between systems within a network after the initial compromise) to maximize impact across interconnected systems [5].

Traditional perimeter-based security architectures are ill-suited to defend against these strategies. They assume threats originate outside the network and that internal entities can be implicitly trusted. Ransomware operators bypass these defenses via phishing, compromised credentials, remote service exploitation, or supply chain vulnerabilities [5]. Once inside, flat networks and excessive privileges enable attackers to move laterally with minimal resistance. As enterprises adopt cloud services, remote endpoints, and hybrid environments, the notion of a clear network perimeter has become obsolete.

Zero Trust Architecture has become a strategic response to these structural weaknesses. Rather than relying on network location as a proxy for trust, Zero Trust operates on the principle that no user, device, application, or workload should be trusted by default. Access to resources must be continuously verified, explicitly authorized, and limited according to least privilege principles. Formalized in NIST Special Publication 800-207 [6], Zero Trust redefines enterprise security around identity-based controls, contextual policy evaluation, and continuous monitoring.

This paper examines how Zero Trust Architecture can mitigate ransomware risk by addressing the internal trust relationships that enable large-scale compromise. Specifically, it analyzes how least privilege enforcement (Section 4.1), continuous authentication and monitoring (Section 4.2), and micro-segmentation (Section 4.3) restrict attacker mobility and reduce the blast radius of successful intrusions. The paper first outlines the evolving ransomware threat model and the limitations of perimeter-based security, then explores the logical components of Zero Trust systems. It subsequently evaluates how Zero Trust strategies are applied in practice and discusses challenges with enterprise adoption.

2 Background

2.1 Ransomware as a Multi-Stage Enterprise Threat

Ransomware is a form of malicious software that encrypts an organization's data or systems and demands payment in exchange for restoring access. While early ransomware campaigns primarily targeted individual users through automated malware distribution, modern ransomware operations have evolved into highly organized attacks directed at enterprise infrastructure [4, 5].

Modern ransomware campaigns typically begin with an initial compromise, frequently achieved through phishing, credential theft, exploitation of software vulnerabilities, or

abuse of remote access services. Following initial access, adversaries conduct reconnaissance to identify high-value systems, escalate privileges, and move laterally across interconnected assets before deploying encryption payloads.

This lateral movement phase, the process of moving within a network to gain wider access, is central to the effectiveness of contemporary ransomware. Enterprise environments often contain shared administrative credentials, broad access permissions, and flat network architectures that allow attackers to traverse systems with limited resistance. Many ransomware groups also perform data exfiltration prior to encryption, increasing financial leverage through double-extortion tactics, in which attackers both encrypt systems and threaten to release stolen data unless a ransom is paid. As a result, effective defense mechanisms must address not only perimeter protection but also internal trust relationships and privilege structures that enable widespread compromise [5].

2.2 Limitations of Traditional Perimeter-Based Security

Traditional enterprise security architectures are designed around the concept of a trusted internal network protected by a hardened external perimeter. Firewalls, intrusion detection systems, and virtual private networks (VPNs) serve as gatekeepers that prevent unauthorized external access. Once authenticated and granted entry, however, users and devices often operate with relatively broad internal trust.

This implicit trust assumption presents significant weaknesses in modern enterprise environments. If an attacker successfully compromises a legitimate credential or bypasses perimeter defenses, the internal network frequently offers limited barriers to lateral movement. Additionally, the expansion of cloud services, remote work, and virtualized infrastructure has blurred traditional network boundaries, rendering the concept of a clearly defined perimeter increasingly obsolete. In this context, security models that rely primarily on network location as a proxy for trust are insufficient to contain ransomware and other advanced threats. [4]

2.3 Zero Trust Architecture (ZTA): Foundational Tenets

Zero Trust Architecture (ZTA) represents a fundamental shift from perimeter-based security toward a model that eliminates implicit trust within enterprise systems. Rather than assuming that entities inside a network are trustworthy, Zero Trust requires continuous verification of users, devices, applications, and workloads (individual applications, services, or computing resources such as CPUs and GPUs) regardless of their network location.

NIST Special Publication 800-207 formalized Zero Trust through a set of foundational tenets [6]. These tenets establish that all data sources and computing services are treated as protected resources, all communication must be secured regardless of location, and access to resources is granted on a per-session basis using least privilege principles. Access decisions are determined dynamically through policy evaluation that considers identity attributes, device posture (the device's current security state, including patch level, antivirus status, and configuration compliance), application state (the current security posture and behavior of the application), behavioral indicators, and environmental context.

Furthermore, Zero Trust requires continuous monitoring of asset integrity and security posture. Authentication and authorization are not one-time events but are dynamically enforced and reevaluated throughout ongoing interactions. Enterprises are expected to collect telemetry across systems and network interactions to inform policy decisions and improve security posture over time. Collectively, these principles redefine trust as conditional, context-driven, and continuously assessed rather than implicitly granted. [6]

3 Core Zero Trust Logical Components

Zero Trust Architecture is composed of several interrelated logical components that collectively enforce dynamic, policy-driven access control across enterprise resources. These components may be deployed as on-premises services, cloud-based services, or hybrid implementations.

Figure 1 presents the conceptual logical model defined in NIST SP 800-207. The figure represents an idealized architectural abstraction rather than a specific product implementation. It illustrates how initially untrusted access requests are intercepted and evaluated through distinct logical components before a subject is considered trusted and permitted to communicate with an enterprise resource.

Access is never implicitly granted based on network location. Instead, every request is evaluated against defined policy and contextual inputs, and communication with the protected resource is permitted only after explicit authorization. The core logical components enabling this enforcement are described below.

3.1 Policy Enforcement Point (PEP)

The Policy Enforcement Point (PEP) enables, monitors, and ultimately terminates communication between a subject and an enterprise resource. It functions as the security checkpoint that enforces policy decisions in real time.

When a subject attempts to access a resource, the PEP intercepts the request and forwards it to the Policy Decision Point (PDP) for evaluation. Based on the authorization decision returned by the PDP, the PEP either permits the session to proceed or blocks the connection.

Although represented as a single logical component in ZTA, the PEP may be implemented in multiple forms, including a client-side agent on a user device, a gateway positioned in front of a protected resource, or a centralized portal component controlling communication paths. Beyond the PEP lies the trust zone hosting the enterprise resource. Access to

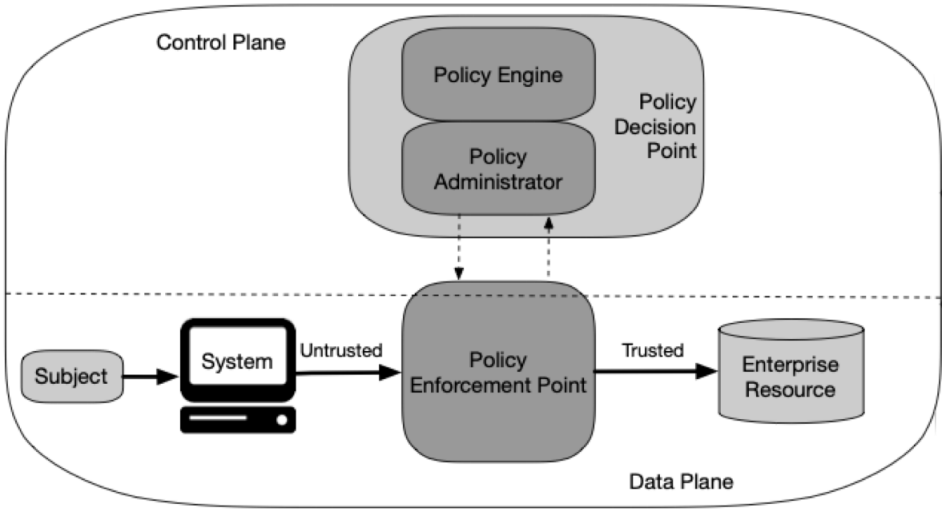

Figure 1. Core logical components of Zero Trust Architecture, illustrating the interaction between the Subject, Policy Enforcement Point, Policy Decision Point, and enterprise resource, as defined in NIST SP 800-207 [6].

this zone is permitted only after successful policy evaluation and enforcement [6].

3.2 Policy Decision Point (PDP)

The Policy Decision Point is responsible for determining whether a subject should be granted access to an enterprise resource. It evaluates access requests and produces authorization decisions based on enterprise-defined policies and contextual inputs. Enterprise-defined policies specify which users, devices, or workloads may access which resources and under what conditions, while contextual inputs encompass identity attributes, device posture, application state, behavioral indicators, and environmental factors such as location or network.

In a Zero Trust model, the PDP is divided into two components: the Policy Engine (PE) and the Policy Administrator (PA). While some implementations may combine these functions into a single service, they are separated here to distinguish decision-making from execution [6].

3.2.1 Policy Engine (PE). The Policy Engine makes the ultimate decision to grant, deny, or revoke access to enterprise resources. It evaluates each access request using enterprise-defined security policies alongside contextual data from external systems, such as Continuous Diagnostics and Mitigation platforms and threat intelligence feeds.

This evaluation is performed through a trust algorithm, which considers inputs including the access request itself, the subject's attributes and historical behavior, the security posture of the device or workload, and resource-specific requirements. Based on these factors, the PE produces an authorization decision, which it records and forwards to the Policy Administrator for enforcement through the appropriate Policy Enforcement Points. Trust algorithms may be implemented in different ways: they can be criteria-based, granting access only when all required conditions are met, or score-based, computing a confidence level from weighted inputs. Additionally, they can be singular (evaluating each request independently) or contextual (taking into account recent activity and behavior patterns to detect anomalous or suspicious access attempts). [6]

3.2.2 Policy Administrator (PA). The Policy Administrator executes the decision made by the Policy Engine. Based on the PE's determination, the PA establishes or terminates the communication path between the subject and the enterprise resource.

If access is granted, the PA generates session-specific authentication tokens or credentials and configures the Policy Enforcement Point to permit the session, with each subsequent request remaining subject to policy enforcement. If access is denied, or if a previously approved session must be revoked, the PA instructs the PEP to terminate the connection. Communication between the PA and the PEP occurs via the control plane [6].

3.3 Control Plane and Data Plane Separation

In addition to defining logical components, the ZTA model distinguishes between two communication planes: the control plane and the data plane.

The *control plane* carries policy-related communications, including access requests, authorization decisions, session establishment commands, and policy updates exchanged between the PEP and PDP. It is responsible for decision-making, policy evaluation, and session management.

The *data plane*, in contrast, carries the actual application traffic exchanged between the subject and the enterprise

resource after access has been authorized. Once the PA configures the PEP to permit a session, user data flows through the PEP to the protected resource via the data plane.

Separating the control plane from the data plane ensures that policy evaluation and enforcement mechanisms remain logically distinct from application data transfer. This architectural separation enhances security, scalability, and manageability within a Zero Trust deployment [6].

4 Key Zero Trust Strategies for Ransomware Mitigation

Zero Trust Architecture mitigates ransomware not through a single control, but through the coordinated application of multiple reinforcing security principles. Rather than relying on perimeter-based defenses, ZTA restructures trust relationships across users, devices, and network resources [6]. This section examines three core strategies that are particularly effective against ransomware threats: least privilege access with identity-based controls, continuous authentication and monitoring, and micro-segmentation for containment. Together, these mechanisms limit initial compromise, restrict lateral movement, and reduce the operational impact of successful intrusions.

4.1 Least Privilege Access and Identity-Based Controls

A fundamental principle of Zero Trust Architecture is the enforcement of least privilege access, which grants users, devices, and applications only the minimal permissions required to perform their tasks [6]. Unlike traditional perimeter-based security models that assume trust based on network location, Zero Trust requires explicit verification and authorization for every access request, including internal requests. This principle is particularly important for mitigating ransomware attacks, which often rely on excessive permissions and unrestricted internal movement to spread across a network after an initial compromise [5].

In many ransomware incidents, compromised accounts with administrative privileges allow attackers to escalate privileges and encrypt large portions of enterprise infrastructure. Modern environments often integrate on-premises systems, cloud platforms, and distributed services, meaning a single credential may grant access to multiple resources [3]. Least privilege mitigates this risk by granting users, devices, and applications only the permissions necessary for their role, enforced through access control policies [6]. By limiting permissions in this way, ransomware propagation beyond the initial compromise is significantly reduced.

Experimental studies demonstrate the effectiveness of this approach. In a controlled ransomware simulation, a Zero Trust framework combining least privilege access, continuous monitoring, and host-based micro-segmentation was implemented [3]. The study was conducted in a lab environment replicating an enterprise network, including virtual machines, file servers, and endpoints. Ransomware was introduced under controlled conditions, and its impact was measured by monitoring the percentage of files encrypted, the spread of the attack across systems, and the ability of containment mechanisms to limit lateral movement. The results showed that, on average, only about 20% of files in the targeted directory were encrypted, leaving roughly 80% of the data intact [3]. These findings highlight how combining least privilege, continuous monitoring, and micro-segmentation can significantly reduce the impact of ransomware attacks.

Identity-based controls further strengthen least privilege enforcement by tying authorization decisions to verified identities and contextual attributes rather than static network assumptions. According to NIST, Zero Trust systems evaluate identity, device posture, and environmental context before granting access to enterprise resources [6]. For example, an endpoint attempting to access a sensitive database must not only present valid credentials but must also demonstrate compliance with security requirements such as up-to-date patches, trusted device configurations, and appropriate access roles. This evaluation ensures that privileges remain tightly scoped and dynamically enforced.

By combining least privilege policies with identity-based verification, Zero Trust architectures significantly reduce opportunities for ransomware operators to escalate privileges, traverse networks, or access sensitive resources.

4.2 Continuous Monitoring and Behavioral Detection

Continuous monitoring is a fundamental component of Zero Trust Architecture. While traditional perimeter-based security models often include monitoring capabilities, this monitoring is typically focused on network boundaries and external traffic. In contrast, Zero Trust environments require constant observation of internal system activity to detect anomalies and malicious behavior in real time. According to NIST, Zero Trust security systems continuously evaluate telemetry from endpoints, networks, and applications in order to dynamically reassess trust decisions throughout an active session [6]. This emphasis on internal behavioral monitoring is particularly important in defending against ransomware attacks, which frequently spread through lateral movement and privilege abuse after an initial compromise. Without visibility into internal activity, ransomware can encrypt large volumes of data within minutes before traditional boundary-focused defenses detect the threat.

Recent studies demonstrate the effectiveness of continuous internal monitoring systems integrated with modern security platforms. One experimental framework combined machine learning models with the Wazuh Security Information and Event Management platform, an open-source system for intrusion detection, log analysis, and endpoint

monitoring, to observe internal activity across multiple enterprise endpoints [2]. The system achieved real-time detection with inference times between 35ms and 62ms while maintaining high classification accuracy, correctly distinguishing between malicious and benign events. The machine learning models achieved up to 97.2% accuracy with a false-positive rate of only 0.03%, demonstrating the potential for these techniques to significantly improve threat detection.

In addition to improving detection accuracy, continuous monitoring of internal system activity significantly reduces alert fatigue for security analysts. When the machine learning–enhanced system was evaluated on a labeled dataset containing 15,427 security events, the approach achieved a 78% reduction in false positives, decreasing the false alarm rate from 23% to 5% compared with a traditional rule-based monitoring approach. For more details on the specific models and configurations used, see *"Improving Threat Detection in Wazuh Using Machine Learning Techniques"* [2].

These findings are consistent with other experimental Zero Trust ransomware defense studies. In a controlled simulation, internal continuous monitoring using Linux auditing (`auditd`) and the Wazuh security platform detected ransomware activity within approximately 5.3 seconds of the start of file encryption [3]. This early detection allowed automated containment measures to limit the attack's impact, resulting in only 12% of files in the targeted directory being encrypted during that period. The results demonstrate that real-time telemetry analysis can significantly reduce ransomware damage by enabling rapid response actions.

These results demonstrate that internal continuous monitoring enables rapid detection of ransomware activity. When combined with automated response mechanisms and access control policies, such monitoring allows organizations to identify and contain ransomware attacks before they spread throughout the enterprise environment.

4.3 Micro-Segmentation and Containment

Micro-segmentation is a key architectural component of Zero Trust environments that limits lateral movement within enterprise networks. Instead of allowing unrestricted internal communication, networks are divided into smaller logical segments governed by fine-grained access policies [1]. Communication between segments is permitted only through explicitly defined policy enforcement points.

As illustrated in Figure 2, micro-segmentation creates clearly defined trust boundaries between workloads and network zones (groups of systems or devices with similar functions). In traditional flat network architectures, systems may freely communicate with each other, and firewalls are typically deployed only at the perimeter to control north-south traffic (communication between internal systems and external networks such as the internet). Micro-segmentation, in contrast, enforces strict policies not only on north-south traffic but also on east-west traffic (communication between

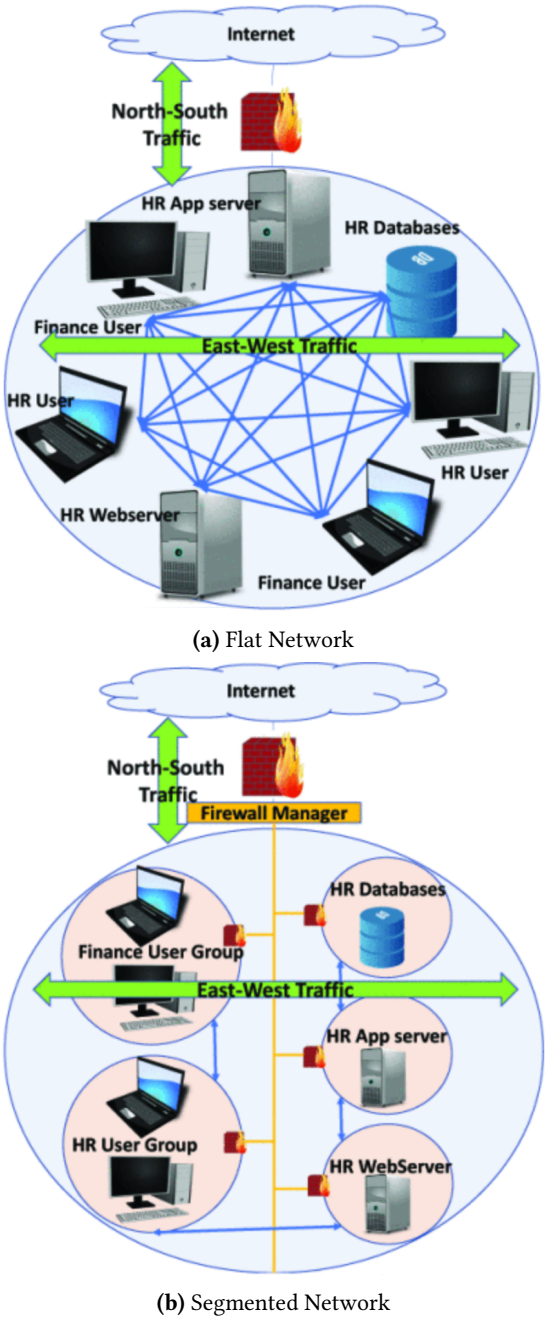

(a) Flat Network

(b) Segmented Network

Figure 2. Comparison of network architectures: (a) a flat network where all systems are interconnected and (b) a micro-segmented network where the system is divided into isolated segments [1].

systems within the enterprise network or data center). Firewalls or policy enforcement points are applied to both directions, preventing lateral movement and limiting attackers' ability to move across segments. [1]

Experimental evaluations demonstrate that micro-segmentation can greatly improve network security. In one study,

researchers analyzed two enterprise networks, a university and a life-care organization, by mapping how devices and servers communicate [1]. In flat networks, all internal hosts could communicate freely, so an attacker who compromised a single system could reach any other system in one step. After applying micro-segmentation, allowed connections between systems dropped, reducing overall network exposure by over 99%. Metrics such as average node connectivity (the number of direct connections per system) and clustering coefficient (how tightly groups of systems are interconnected) also declined substantially, falling by 95% and 70–80%, respectively. As a result, micro-segmentation greatly increases the effort required to reach high-value targets. In flat networks, the shortest path between any two hosts typically has a length of one, meaning that a compromised system can directly communicate with any other asset. After micro-segmentation, the shortest path increased to two or even three in some cases. Overall, the effort needed to reach a high-value target doubled, even in the worst case. [1]

Most importantly, micro-segmentation dramatically reduces the number of potential attack paths available to adversaries by restricting communication between internal systems to only explicitly authorized connections. Attack graph analysis revealed that deploying micro-segmentation decreased the total number of viable attack paths by approximately 99%. The average number of successful attack opportunities following a compromised privilege, representing a typical account, was reduced by roughly 93%, while the worst case attack opportunities, representing the most powerful privilege that could reach the largest number of systems, decreased by nearly 69%. Micro-segmentation also strengthens network robustness by limiting access to vulnerable assets. Even though software vulnerabilities still exist on devices and applications, micro-segmentation reduces the number of attack graph root nodes, the initial points where an attacker could gain access, by over 65%, effectively making these entry points unreachable due to strict communication controls between network segments [1].

These results highlight the critical role of micro-segmentation in ransomware defense. Once ransomware gains an initial foothold, its ability to spread depends heavily on unrestricted lateral communication between systems. By enforcing strict isolation between workloads and network zones, Zero Trust micro-segmentation limits the blast radius of a compromise. Even if a single system becomes infected, the malware is confined to a restricted segment of the network, preventing widespread encryption of enterprise resources and enabling faster containment during incident response.

5 Practical Challenges and Limitations of Zero Trust

Despite its security advantages, implementing Zero Trust Architecture introduces several practical challenges. One major difficulty is the scale and complexity of modern enterprise networks. ZTA requires fine-grained access control policies that specify which users, devices, and services can access particular resources. While these policies strengthen security by limiting unnecessary privileges, they can also become difficult to manage in large and highly distributed environments. Organizations operating cloud services, mobile devices, and Internet of Things systems must maintain consistent policies across heterogeneous infrastructure, which can significantly increase administrative overhead. [4]

ZTA also relies heavily on continuous authentication and dynamic authorization. Although these mechanisms improve security by ensuring that every access request is verified, they may introduce performance overhead and increased latency, particularly in large networks with high volumes of authentication requests. Additionally, designing an effective ZTA requires integrating multiple technologies such as identity management systems, network micro-segmentation mechanisms, and policy enforcement components, which can increase implementation complexity. [4]

Beyond technical challenges, organizations must also address operational and organizational barriers. Migrating from traditional perimeter-based security models to Zero Trust often requires integrating legacy systems that were not designed for identity-based access control. This transition may require infrastructure upgrades, policy redesign, and workforce training, all of which can increase costs and slow adoption. Consequently, while ZTA provides strong security benefits, careful planning and resource allocation are necessary for successful implementation. [4]

6 Conclusion

Ransomware continues to pose a significant threat to modern organizations, particularly as enterprise environments become increasingly distributed and interconnected. Traditional perimeter-based security models are insufficient for defending against attackers who can bypass external defenses and exploit implicit trust within internal networks. Zero Trust Architecture addresses these weaknesses by removing assumptions of trust and enforcing continuous verification of users, devices, and applications.

This paper examined how key Zero Trust principles, including least privilege access, continuous monitoring, and micro-segmentation, can significantly reduce the effectiveness of ransomware attacks. By restricting access privileges, monitoring system activity in real time, and isolating network segments, organizations can limit lateral movement and contain the impact of a successful compromise. Although implementing Zero Trust introduces technical and organizational challenges, its ability to reduce attack surface and contain breaches makes it a critical strategy for improving enterprise resilience against ransomware and other advanced cyber threats.

Acknowledgments

I would like to thank my advisor, Kristin Lamberty, Cassie Bonte, and Peter Dolan for their guidance and support throughout the research and writing process of this paper.

References

[1] Nardine Basta, Muhammad Ikram, Mohamed Ali Kaafar, and Andy Walker. 2022. Towards a Zero-Trust Micro-segmentation Network Security Strategy: An Evaluation Framework. In *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. 1–7. doi:10.1109/NOMS54207.2022.9789888

[2] Samir Achraf Chamkar, Mounia Zaydi, Yassine Maleh, and Noreddine Gherabi. 2025. Improving Threat Detection in Wazuh Using Machine Learning Techniques. *Journal of Cybersecurity and Privacy* 5, 2 (2025). doi:10.3390/jcp5020034

[3] Atharva Dhumal, Mustafa Ghaleb, Samah Abdelsalam, Arghir-Nicolae Moldovan, and Mosab Hamdan. 2025. Zero Trust Architecture for Ransomware Defense in Virtualized Environment. In *Proceedings of the IEEE/ACM 12th International Conference on Big Data Computing, Applications and Technologies (BDCAT '25)*. Association for Computing Machinery, New York, NY, USA, Article 24, 7 pages. doi:10.1145/3773276.3774876

[4] Muhammad Liman Gambo and Ahmad Almulhem. 2025. Zero Trust Architecture: A Systematic Literature Review. *Journal of Network and Systems Management* 34, 1 (Nov. 2025). doi:10.1007/s10922-025-09998-x

[5] Microsoft Threat Intelligence. 2022. Ransomware as a service: Understanding the cybercrime gig economy and how to protect yourself. https://www.microsoft.com/en-us/security/blog/2022/05/09/ransomware-as-a-service-understanding-the-cybercrime-gig-economy-and-how-to-protect-yourself/

[6] Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly. 2020. *Zero Trust Architecture*. Technical Report Special Publication (SP) 800-207. National Institute of Standards and Technology (NIST), Gaithersburg, MD. doi:10.6028/NIST.SP.800-207

[7] Sophos. 2026. The State of Ransomware 2025. https://www.sophos.com/en-us/content/state-of-ransomware