Micro-Segmentation for Zero Trust Architecture

A Framework for Legacy Systems Integration

EPA Master Thesis

Sofia Sakhi

Delft University of Technology

TU Delft

EY Building a better working world

Micro-Segmentation for Zero Trust Architecture

A Framework for Legacy Systems Integration

by

Sofia Sakhi

Student Name	Student Number
Sofia Sakhi	5417015

Thesis Chair:	Professor Martijn Warnier
First Supervisor:	Dr. Yury Zhauniarovich
Second Supervisor:	Professor Martijn Warnier
EY Supervisor:	M. Meesters
EY Supervisor:	J. Gros
Project Duration:	February, 2025 - July, 2025
Faculty:	Faculty of Technology, Policy and Management, Delft

Cover:	Cybersecurity, lock, technology
Style:	TU Delft Report Style, modified by Sofia Sakhi

TU Delft

EY Building a better working world

Preface

This thesis, , stems from my work during the ICT Architecture course, where I and my team were tasked with redesigning the network of a real-world client. While exploring ways to enhance the network's security, I encountered the Zero Trust model; a concept that immediately stood out for its relevance in today's threat landscape. However, it quickly became apparent that implementing Zero Trust in a legacy-heavy environment posed unique and often overlooked challenges. This experience sparked the idea for this research.

The process of working on this thesis has deepened my understanding not only of Zero Trust Architecture, but also of the real-world tensions between ideal security models and the operational constraints that organizations face. I learned how technical feasibility, organizational priorities, and architectural decisions all intersect when security strategies are implemented. Developing a framework that bridges these gaps required not only academic analysis but also learning to extract meaningful insights from real-world cases and practitioner interviews.

The study is intended to support organizations facing similar challenges by offering a framework to select and evaluate micro-segmentation strategies that align with their infrastructure realities. It combines technical considerations with practical implementation criteria to bridge the gap between modern security models and legacy constraints.

I would like to express my sincere gratitude to my university supervisors and the company mentors who guided me throughout this research, as well as the professionals I interviewed, whose insights were essential to shaping the practical aspects of this work. The journey involved both technical deep dives and real-world learning, particularly around balancing ideal architectures with operational feasibility; a lesson that I believe is central to any meaningful security design.

A special thanks goes to Haiko van der Voort, for his commitment to academic integrity and for fostering a critical, systems-thinking mindset that shaped how I approached this thesis. It's possible that I may have forgotten to mention someone who supported me in one way or another, if so, please know that your help is genuinely appreciated.

This thesis is submitted as part of my Master's degree and reflects not only academic research, but also personal growth, professional development, and a deeper appreciation for the complexity of securing digital systems in imperfect conditions. I hope it offers practical value to others facing similar challenges.

Sofia Sakhi
Delft, July 2025

Abstract

With the rise of cloud computing, remote work, and interconnected systems, the traditional model of trusting everything inside the network perimeter has become increasingly risky. In response, Zero Trust Architecture (ZTA) has emerged as a model where no user or device is inherently trusted, regardless of their location. Micro-segmentation, one of ZTA's key mechanisms, aims to limit lateral movement and enforce fine-grained access control. However, implementing micro-segmentation in environments with legacy systems remains particularly challenging due to outdated technologies, rigid configurations, and limited visibility.

Although the literature offers many advanced segmentation models, these are often designed for cloud-native or modern infrastructures and fail to address the constraints inherent in legacy-heavy environments. This research investigates how organizations with legacy systems can identify and evaluate appropriate micro-segmentation strategies based on their network characteristics and technical constraints. The central research question is:

> *How can organizations choose and implement micro-segmentation strategies that align with their network architectures and the technical constraints posed by legacy systems?*

To address this question, the study began with a structured literature review, which revealed a lack of detailed and context-aware guidance for applying micro-segmentation in constrained environments. To fill this gap, a series of semi-structured interviews were conducted with cybersecurity professionals experienced in industrial and legacy-heavy systems. The interview data were analysed using thematic coding to uncover recurring patterns, practical challenges, and the decision-making logic used in real-world settings.

The findings reveal three principal categories of micro-segmentation strategies: network-based, agent-based, and hybrid. The choice between these depends on technical factors such as system compatibility, network architecture, and the level of administrative control available. In addition to strategy selection, the study identifies a phased implementation approach typically followed by organizations. The study also defines key criteria for evaluating segmentation outcomes, including enforcement of access control, operational continuity, system performance, network visibility, and many more. These insights form the basis of a structured, step-by-step implementation guide and a decision-support framework. Together, they enable organizations to assess their technical constraints and select appropriate strategies. The main contribution of this research is the development of a practice-oriented framework that supports the secure integration of legacy systems into Zero Trust environments.

Executive Summary

The cybersecurity landscape is evolving rapidly, with increasingly complex threats challenging traditional defence models. Perimeter-based approaches, which once relied on securing the boundaries of a network, are now proving inadequate against sophisticated adversaries capable of moving laterally within systems after breaching a single point. As a response, Zero Trust Architecture (ZTA) has gained traction in both public and private sectors. Built on the principle of "never trust, always verify," ZTA requires continuous verification of identities, strict access controls, and segmentation of systems regardless of their location within a network. At the core of ZTA is micro-segmentation, a method for breaking down networks into smaller, more secure zones to prevent attackers from easily moving through internal systems.

While micro-segmentation is widely endorsed in theory, its practical application, especially in environments with legacy systems, remains underexplored. Legacy systems are still heavily relied upon in many critical sectors, such as healthcare, energy, transport, and finance. These systems often lack the flexibility, compatibility, or performance required for modern security solutions. They are also deeply embedded in daily operations, making replacement costly and disruptive. As a result, organizations face a difficult challenge: how to apply advanced security strategies like micro-segmentation in infrastructure that was not designed for it.

This research addresses that challenge directly. Using a qualitative research method based on semi-structured interviews with domain experts, this study explores how micro-segmentation can be effectively applied in legacy-heavy environments by identifying practical strategies, trade-offs, and evaluation criteria involved in the implementation process. The study contributes a structured decision-making framework (see Figure 6.1)and a phased implementation roadmap, designed to help organizations make informed choices based on their network characteristics and technical constraints. By grounding its insights in real-world practices rather than theoretical ideals, the research aims to bridge the gap between academic models and practical needs.

The main research question guiding this work is:

> *'How can organizations choose and implement micro-segmentation strategies that align with their network architectures and the technical constraints posed by legacy systems?*

This question was broken down into four sub-questions that structure the study and its findings. First, the research asked what strategies are used to implement micro-segmentation and how network characteristics influence the choice of these strategies. A thorough review of existing literature revealed that most academic assume the presence of modern infrastructure compatible with strategies such as software-defined networking (SDN), hypervisors, or context-aware dynamic access control. However, interviews with 16 experienced cybersecurity professionals showed that real-world environments often require simpler, more adaptable methods.

One common approach in legacy-heavy systems is to implement micro-segmentation by creating security zones using virtual LANs (VLANs), with access between these zones managed through firewall rules based on network attributes such as IP addresses. For more granular control within those VLAN-defined zones, some organizations use Private VLANs (PVLANs) to further restrict communication between individual systems. This method provide logical isolation without requiring endpoint changes, making them suitable for outdated systems. In more modern parts of the network, organizations use agent-based controls to enforce fine-grained policies. This method works by installing software agents on endpoints that leverage the host firewall to monitor and control traffic at a granular level between systems. Often, organizations adopt a hybrid approach that blends different segmentation methods based on technical feasibility. Agent-based controls are used where infrastructure supports them, while network-based methods are applied in areas where host-level enforcement is not possible.

Third, the study examined the step-by-step process organizations follow to implement micro-segmentation and how these steps can be organized into a repeatable roadmap. Implementation typically begins with network separation, where critical assets are isolated from the rest of the environment using basic network controls. From this foundation, the approach diverges based on the technical capabilities of the environment.

Depending on the level of agent support, organizations adopt different strategies. If the majority of systems are agent-compatible, micro-segmentation is implemented using an agent-based approach, following the dedicated roadmap outlined in Figure 5.4. This includes validating endpoint readiness, deploying agents in observation mode, collecting traffic data, simulating policies in monitor-only mode, and gradually shifting to full enforcement using tagged workloads. If agent support is limited or inconsistent, a network-based micro-segmentation strategy is adopted, following the corresponding network segmentation roadmap (see Figure 5.3). This involves conducting a detailed asset inventory, defining the target segmentation architecture, creating high- and low-level design documents, and implementing segmentation using VLANs, routing, and firewall policies. Enforcement is applied gradually, with extensive validation and testing. In mixed environments, a hybrid approach is used. Agent-based enforcement is applied where feasible, and network-based controls are used elsewhere. In particularly sensitive or high-risk areas, some organizations adopt a layered approach, where broad segmentation is first implemented at the network level to define zones, and agent-based controls are added within those zones to apply more granular policies.

Fourth, the research investigated how the success of micro-segmentation implementations is evaluated. Practitioners reported using a combination of qualitative and technical measures. The primary test is whether access control behaves as expected, systems should only communicate with authorized peers, and violations should be blocked or logged. Operational continuity is another key benchmark. If segmentation disrupts normal activities or introduces performance bottlenecks, it may be rolled back. Participants also evaluate segmentation based on security improvements, such as reduced attack surface and visibility into traffic patterns.Coverage, or in other words the percentage of the network that is effectively segmented, is another important criterion mentioned by interviewees. If the segmentation design is too complex to manage over time, it is considered a failure, even if it works in the short term.

By addressing these questions, the research makes several key contributions. It provides a realistic view of micro-segmentation implementation grounded in practitioner experience. This study critically builds on a phased Zero Trust Architecture model that emphasizes system upgrades or replacements for legacy environments but lacks implementation depth. While some existing research offers limited solutions for securing legacy systems through isolation or encapsulation, they are often narrow in scope or not broadly applicable. In response, this work provides a technically detailed micro-segmentation guide and introduces a generalized decision-making framework to bridge the gap between high-level theory and real-world implementation.

While this study offers a practical framework for micro-segmentation in Zero Trust environments, it has several limitations. The framework has not been validated through large-scale or cross-sector testing, and its generalizability remains limited. It focuses only on micro-segmentation, excluding other core Zero Trust components like identity governance or continuous monitoring. Human factors such as resistance to change and limited organizational readiness also affect implementation but were not fully explored. Additionally, the role of AI in supporting micro-segmentation is not well understood, and future research should examine how effectively it is applied in real-world settings.

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
ACL	Access Control List
CCTV	Closed-Circuit Television
CPU	Central Processing Unit
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
HMI	Human-Machine Interfaces
IP	Internet Protocol
OS	Operating System
OT	Operational Technology
PLC	Programmable Logic Controllers
PVLAN	Private Virtual Local Area Network
SDN	Software-Defined Networking
VLAN	Virtual Local Area Network
VRF	Virtual Routing and Forwarding
ZTA	Zero Trust Architecture

1 Introduction

1.1. Research Background

In traditional perimeter-based security models, trust is granted to users and devices once they are inside the network. However, this assumption no longer holds in an environment where internal networks are just as exposed as external ones. Over the past decade, the enterprise IT landscape has undergone a fundamental transformation. The increasing reliance on cloud computing, mobile devices, and third-party integrations has significantly changed how organizations structure and operate their digital infrastructure (Chang et al., 2010). Corporate networks are no longer confined to centralized offices and data centres, they now extend across geographically distributed users, virtual environments, and externally hosted services (Masuda et al., 2017). Additionally a growing number of employees, contractors, and partners access internal systems from personal devices and unmanaged networks. As remote and hybrid work models continue to rise (see Figure 1.1), the number of potential entry points into organizational systems has expanded substantially. This shift has dissolved the traditional notion of a secure network perimeter (Anjum et al., 2022; Gadkari, 2025; Kindervag, 2010).

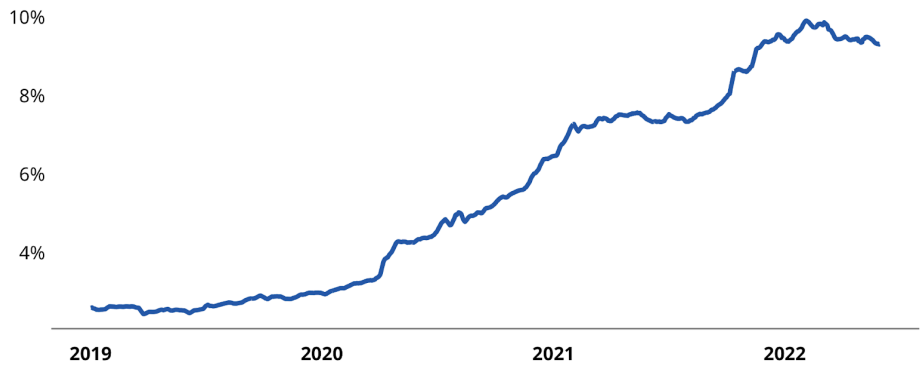

Figure 1.1: Increase in the percentage of remote job postings between 2019 and 2022 (Bunker, 2022)

At the same time, attackers have adapted to this complexity. Once inside a network, often through phishing, compromised credentials or vulnerable endpoints, they rely on lateral movement to escalate privileges and exfiltrate data (H. Kang et al., 2020; Siadati and Memon, 2017).

The lack of internal segmentation or access controls enables small breaches to turn into full-scale incidents (Herranz-Oliveros et al., 2024). A well-known example is the 2013 Target data breach, where attackers gained access through stolen credentials from a third-party HVAC vendor and moved laterally within the network to steal payment card data from point-of-sale systems (U.S. Senate Committee on Commerce, Science, and Transportation, 2014). This breach exposed 40 million credit card numbers and the personal data of over 70 million customers (Jones, 2025). Recent data shows a continued rise in phishing activity and credential compromise (see Figure 1.2), underlining the importance of strong containment mechanisms within the network.

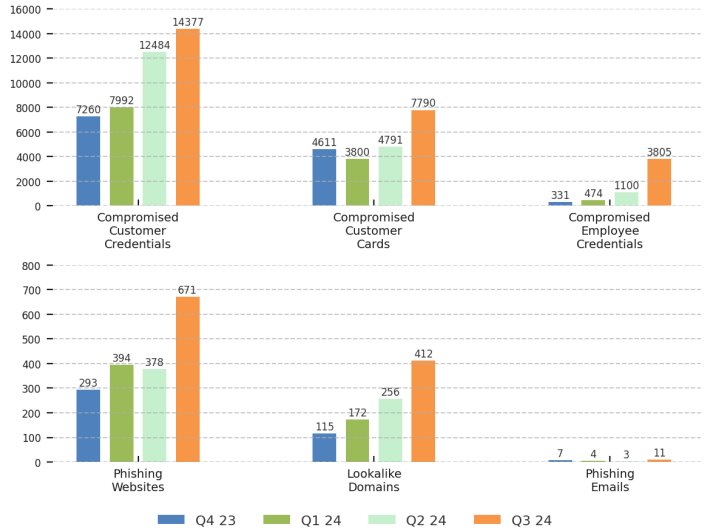

Figure 1.2: Quarterly distribution of key cyber threats, including phishing and compromised credentials (Cyberint, 2024)

In response to this increasingly fragmented and exposed landscape, Zero Trust Architecture (ZTA) has emerged as a new security paradigm. Rather than assuming that systems or users within the network are trustworthy, ZTA operates on the principle of "never trust, always verify" (Ofili and Obasuyi, 2025). It enforces continuous authentication, strict access control, and real-time monitoring of traffic between users, devices, and applications, regardless of their location (Stafford, 2020). A visual comparison between the traditional perimeter-based security model and the Zero Trust model is illustrated in Figure 1.3. One of the core techniques enabling this model is micro-segmentation, which involves dividing the network into smaller, logically isolated segments and applying granular policies to control communication between them (Markus, 2024). This containment strategy is especially effective in limiting lateral movement, allowing organizations to reduce the blast radius of potential breaches (Mämmelä et al., 2016).

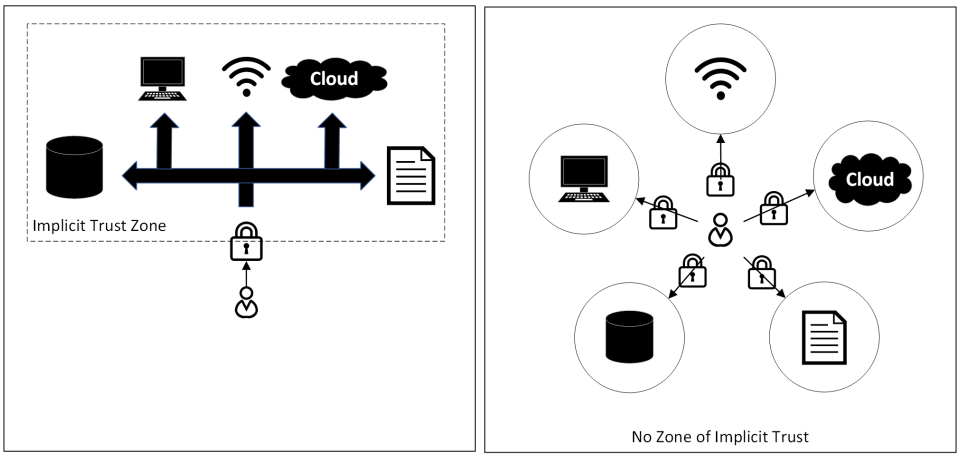

Figure 1.3: Classic Perimeter-Based Security (Left) VS Zero Trust Security Models (Right)

1.1.1. Problem Statement

Adopting Zero Trust and implementing micro-segmentation is not a simple technical upgrade, it is an architectural transformation (Oladimeji, 2024). Most existing implementation models assume that organizations already possess modern infrastructure compatible with Zero Trust concepts, such as centralized identity governance and dynamic policy enforcement. However, integrating micro-segmentation into existing infrastructures presents significant challenges, particularly for organizations with legacy systems that lack the flexibility to support such granular controls (Kolawole, 2025).

A legacy system refers to outdated computing software or hardware that remains in use despite the availability of more modern alternatives. These systems frequently present integration challenges, primarily because they lack compatibility with current cybersecurity standards (Binbeshr et al., 2025). Many organizations continue to rely on legacy systems due to their critical role in supporting core business functions and replacing them is often costly, technically complex, and carries operational risk (Binbeshr et al., 2025). The extensive network reconfiguration required for Zero Trust adoption, combined with high implementation costs, often poses a significant barrier for many enterprises, particularly small and medium-sized businesses (SMEs) with limited resources (Aslam and Steltzer, 2025). These constraints highlight the difficulty of aligning Zero Trust principles with operational realities in legacy-heavy environments.

As a result, there is a lack of practical, context-aware guidance to help organizations apply Zero Trust principles, particularly micro-segmentation, in environments dominated by legacy systems. Without such guidance, these organizations face increased security risk and limited ability to contain breaches.

1.1.2. Research Structure

This thesis is structured into six chapters, each building progressively toward answering the main research question. Chapter 2 provides the theoretical background, covering the concept of Zero Trust Architecture, micro-segmentation techniques, and related technologies Chapter 3 presents the literature review, identifying the current state of the art, outlining the research gap, and formulating the research questions. Chapter 4 details the research methodology, including data collection through expert interviews, ethical considerations, and the analytical approach used to develop the framework.

Chapter 5 presents the research results. It defines how micro-segmentation is understood in practice, explores the strategies used in legacy-heavy environments, and outlines the challenges identified through interviews. It also describes the concrete steps toward micro-segmentation and the criteria used to evaluate different approaches. Chapter 6 focuses on framework development. It synthesizes the findings into a structured decision-support framework, including requirements, strategy selection guidance, and potential implications for organizations. Chapter 7 is dedicated to validation. It describes the validation protocol and discusses the outcomes of the validation process. Chapter 8 concludes the thesis by directly answering each sub-question and the main research question. It also offers a broader discussion, highlights the study's limitations, and proposes directions for future research.

1.1.3. Research Scope

This research focuses on micro-segmentation as a technical and strategic component of Zero Trust Architecture, specifically within organizations that rely heavily on legacy systems. The study does not cover the full range of Zero Trust components instead, it concentrates on understanding how segmentation is implemented in real-world environments where infrastructure limitations, visibility gaps, and operational constraints are common.

The analysis is based on expert interviews from a single large multinational organization with substantial legacy infrastructure. While this provides deep insights into practical challenges and solutions, the findings may not represent the full variety of experiences across different sectors, company sizes, or regions. The goal is not to create a one-size-fits-all model, but to develop a flexible and constraint-aware framework that organizations can adapt to their specific context. The study prioritizes strategies that can be realistically applied without requiring full infrastructure replacement. It focuses on the selection, application, and evaluation of micro-segmentation methods in constrained settings. Broader issues such as regulatory influences, economic cost analysis, and end-user adoption are acknowledged but remain outside the direct scope of this research.

1.2. Relevance

This section outlines the broader significance of the research by examining its relevance across societal, scientific, and disciplinary contexts. It highlights how the findings contribute to addressing real-world cybersecurity challenges in legacy-heavy environments, advance academic understanding of micro-segmentation strategy selection, and align with the interdisciplinary focus of the Engineering and Policy Analysis (EPA) program.

1.2.1. Societal Relevance

Legacy systems form the backbone of many critical sectors, including energy, water, healthcare, transportation, and finance, making them high-value targets for cyberattacks with severe societal consequences (Aljohani, 2022; Krause et al., 2021). These interconnected, often aging infrastructures are particularly vulnerable to malware outbreaks or ransomware campaigns that can disrupt essential services, cause economic losses, and threaten public safety (Aljohani, 2022; Krause et al., 2021).

As cyber threats, particularly those linked to nation state actors, become increasingly sophisticated, the need for advanced security models such as micro-segmentation has become more urgent. Protecting critical services now requires infrastructure designs capable of containing and surviving cyber incidents without resulting in widespread disruption (Durojaye and Raji, 2022; George, 2025).

Micro-segmentation serves as an effective mitigation strategy in these contexts by isolating network segments and enforcing least-privilege principles between zones. This approach significantly restricts lateral movement, thereby limiting the potential impact of a breach originating in a legacy-dependent system (Basta et al., 2022; George, 2025). The critical importance of segmentation is underscored by real-world incidents such as the 2015 Ukrainian power grid attack. Sophisticated malware penetrated a substation network and propagated laterally across flat infrastructure, disrupting power supply to approximately 230,000 consumers in western Ukraine (Smith and Cardenas, 2024; Wikipedia contributors, 2025). Had micro-segmentation been in place, fundamental isolation measures could have significantly limited the breach's spread and impact.

In essence, this research advances societal relevance by delivering actionable insights and a structured framework for micro-segmentation, thereby supporting stable public services, enabling rapid breach containment, and strengthening community trust in essential systems.

1.2.2. Scientific Relevance

This study builds on the phased methodology proposed by Paul and Sheriffdeen (2022), which outlines a structured approach to transitioning from perimeter-based security to Zero Trust Architecture. While their work provides a high-level roadmap, it lacks technical depth on how micro-segmentation should be implemented, particularly in environments constrained by legacy systems. This research extends their contribution by introducing a detailed, step-by-step implementation guide specifically focused on micro-segmentation, grounded in the practical realities of mixed-infrastructure networks.

Further, Paul and Sheriffdeen (2022) suggest that legacy systems must be significantly upgraded or replaced to align with Zero Trust principles. This study challenges that assumption by building on the more inclusive approaches of Rose et al. (2020), Tyler and Viana (2021), and Køien (2021), who propose ways to secure legacy systems through isolation or encapsulation. However, their solutions remain limited to specific use cases or lack general applicability. In response, this study introduces a generalized decision-making framework that helps organizations evaluate their existing network characteristics and select suitable micro-segmentation strategies. Combined with the implementation roadmap and evaluation criteria, this work contributes a comprehensive and adaptable guide for integrating legacy systems into Zero Trust environments, bridging the gap between conceptual proposals and real-world application.

1.2.3. EPA Relevance

The Engineering and Policy Analysis (EPA) program is designed to prepare students to tackle complex, multi-layered problems that combine technical systems, organizational behaviour, and strategic decision-making under uncertainty. The problem addressed in this thesis fits well within that scope. Implementing micro-segmentation in legacy-heavy environments is not simply a technical exercise, it is a systems challenge that involves navigating outdated infrastructure, organizational constraints, competing priorities, and evolving security risks.

What makes this problem especially relevant to EPA is its inherently socio-technical nature. Decisions about segmentation strategies must consider not only what is technically feasible but also what is operationally sustainable, politically acceptable, and aligned with broader organizational goals. These decisions often involve trade-offs between security, performance, cost, and maintainability. In addition, uncertainty, such as unpredictable cyber threats, forces organizations to make strategic choices without full visibility, which is a core concern in EPA-related decision-making.

This thesis contributes to the EPA field by offering a structured, constraint-aware decision framework that helps organizations evaluate and select appropriate micro-segmentation strategies. It applies systems thinking to a real-world cybersecurity challenge and highlights how technical choices intersect with organizational and human factors.

2

Theoretical Background

2.1. Zero Trust Architecture

The concept of Zero Trust emerged as a response to the limitations of traditional perimeter-based security models. In 2004, the Jericho Forum introduced the idea of de-perimeterization, emphasizing the need for security beyond network boundaries (Stafford, 2020). Forrester Research formally coined the term "Zero Trust" in 2009, advocating for strict identity verification and least-privilege access control (Cunningham, 2020). In 2014, Google introduced "BeyondCorp," a framework that implemented Zero Trust principles to eliminate reliance on internal network security (Sarkar et al., 2022). By 2017, Forrester released the Zero Trust eXtended (ZTX) framework (Cunningham et al., 2018), while Gartner introduced the Continuous Adaptive Risk and Trust Assessment (CARTA) model, reinforcing Zero Trust as a necessary security paradigm. The U.S. National Institute of Standards and Technology (NIST) further standardized Zero Trust in 2019 with the publication of Special Publication (SP) 800-207, which provided organizations with structured guidelines for Zero Trust implementation (Stafford, 2020). Today, Zero Trust has become a cornerstone of modern cybersecurity, addressing the evolving threats posed by cloud computing, remote work, and increasingly sophisticated cyberattacks.

Zero Trust Architecture (ZTA) relies on a set of core logical components to enforce security policies and control access to enterprise resources (Stafford, 2020). It divides the network into two parts: the **Control Plane** and the **Data Plane** as shown in Figure 2.1. The control plane is responsible for making decisions about access, while the data plane handles the actual movement of information between the subject (user or device requesting access) and enterprise resources (the protected data or services).

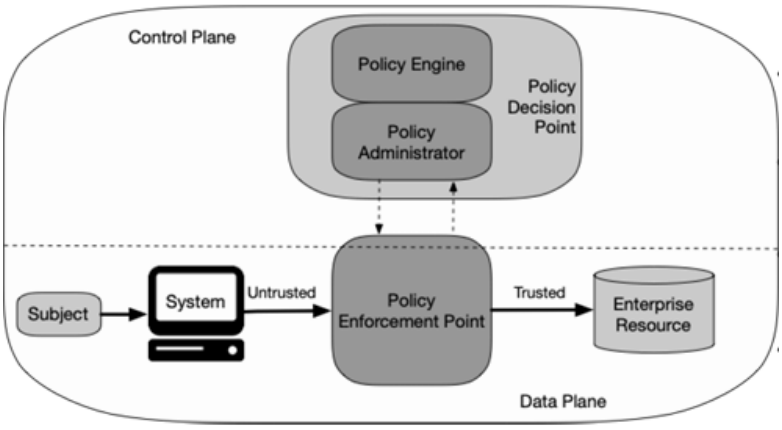

Figure 2.1: Core Zero Trust Logical Components

1. **Policy Decision Point (PDP):** Consists of the Policy Engine (PE), which decides if access should be granted based on security policies, and the Policy Administrator (PA), which enforces these decisions by instructing the Policy Enforcement Point (PEP). The PDP operates on the control plane to manage who gets access to which resources.
2. **Policy Enforcement Point (PEP):** Acts as a gatekeeper on the data plane, enforcing access controls by allowing or blocking data traffic between the untrusted side (the subject) and the trusted side (enterprise resources) based on the decisions made by the PDP.
3. **Access Policy:** Access is granted based on fine-grained access policies that are created based on either predefined rules or contextual factors such as device security posture, threat intelligence, compliance status, and activity logs.

Building on these core components, the Zero Trust Architecture (ZTA) can be implemented using three main methods, as outlined in by NIST (Stafford, 2020):

- **Enhanced Identity Governance:** This approach focuses on using the identity of subjects as the primary factor for creating access policies. Access policies are created and enforced based on identity and assigned attributes. This method is particularly effective in environments with open networks or frequent non-enterprise devices, as well as for cloud-based services where enterprise-owned ZT security components might not be applicable.
- **Micro-Segmentation:** This approach involves placing individual or groups of resources on unique network segments protected by gateway security components acting as PEPs.
- **Network Infrastructure and Software Defined Perimeters (SDP):** In this approach, the network infrastructure is used to implement ZTA, often leveraging overlay networks and SDP concepts. This method is applicable to both on-premises and cloud environments, providing flexibility for managing diverse network infrastructures.

Among the methods mentioned above, micro-segmentation is widely recognized in the literature as a common approach to implementing Zero Trust Architecture (ZTA) (Fernandez and Brazhuk, 2024). Identity-based Zero Trust models are impractical in such contexts, as they require robust identity and access management systems, typically absent in legacy environments (Omar and Abdelaziz, 2020). Similarly, Software-Defined Networking (SDN) approaches necessitate extensive network redesign and specialized expertise, posing integration challenges when layered over older infrastructure (Jammal et al., 2014; Kreutz et al., 2015).

2.1.1. Micro-segmentation

Micro-segmentation is a network security technique that divides a network into smaller, isolated zones or segments, each with its own access policies (Sheikh et al., 2021). A key concept behind micro-segmentation is the control of east-west traffic, which refers to the internal flow of data between systems within the same network, with the goal of limiting attacker movement in the event of a perimeter breach (Sheikh et al., 2021). This control is enforced through strict communication policies between components, only allowing the minimum necessary interactions (Liu et al., 2024). For example, a web server may be permitted to communicate only with the specific application server it requires, and that application server may be allowed to access only the necessary database server (Sheikh et al., 2021). This granularity of access control is distinct from traditional network segmentation, which typically divides the network into broader zones or VLANs and often assumes some level of trust within each zone (Foltz, 2022). Approaches to micro-segmentation can be broadly categorized into following categories (ColorTokens, 2025):

- Network-Based Micro-segmentation
- Hypervisor-Based Micro-segmentation
- Host-Based Micro-segmentation

Network-Based Micro-segmentation

Network-based micro-segmentation operates at the network layer, using technologies such as VLANs, firewalls, and routing to divide a larger network into smaller, logically isolated segments (Li et al., 2024). This approach VLANs (Virtual LANs) are commonly used to group devices within the same broadcast domain regardless of their physical location (Jeuk et al., 2015). Traffic between VLANs is routed through firewalls or Layer 3 devices, where access policies are applied to regulate communication based on predefined rules (Jeuk et al., 2015). This model is relatively simple to implement without requiring changes at the endpoint level, but it offers less granularity than modern alternatives due to its reliance on network topology and static configurations. (Li et al., 2024).

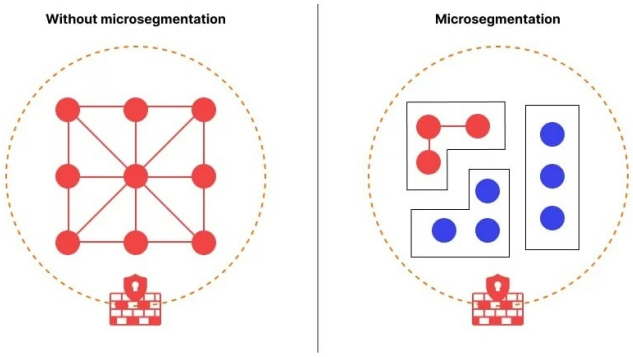

Figure 2.2: Micro-Segmentation Architecture using VLAN to create security zones (Cyber Kendra, 2023)

In more advanced deployments, Software-Defined Networking (SDN) enhances this approach by introducing a centralized controller that manages traffic rules dynamically across the entire network (Kreutz et al., 2015). SDN separates the control plane (which makes decisions) from the data plane (which forwards traffic), allowing network administrators to define and update segmentation policies centrally (Kreutz et al., 2015). Unlike traditional network-based segmentation, which requires manual configuration of each device, SDN enables policies to be deployed automatically and adjusted in real time as systems or applications change (Nunes et al., 2014). This ensures consistent enforcement even in large-scale or frequently changing environments, significantly improving agility and security visibility (Zhao et al., 2017). While SDN offers greater flexibility and control, it is more complex to implement and often incompatible with legacy infrastructure, making it more suitable for modern or greenfield environments.

Hypervisor Level Micro-segmentation

A hypervisor is a software layer that creates and manages virtual machines (VMs) on a physical server, enabling multiple VMs to run isolated operating systems and applications (Bushouse and Reeves, 2018). A virtual machine (VM) is a software-based emulation of a physical computer that runs its own operating system and applications (Carbone et al., 2008). In hypervisor-level segmentation, each VM is connected to a virtual switch managed by the hypervisor, where traffic between VMs is filtered based on defined security policies (Zaenchkovski et al., 2023). This method typically leverages distributed firewalls deployed at the hypervisor layer to inspect and control traffic between virtual machines within the same host or across hosts (Amoroso, 2017). Hypervisor-level segmentation offers a virtualization-centric approach to micro-segmentation by embedding security controls directly into the virtual infrastructure (Amoroso, 2017; Z. Ma et al., 2025). Among the most widely used implementations of this approach is VMware NSX, a network virtualization and security platform that integrates tightly with hypervisor environments to deliver micro-segmentation capabilities (VMware, 2020).

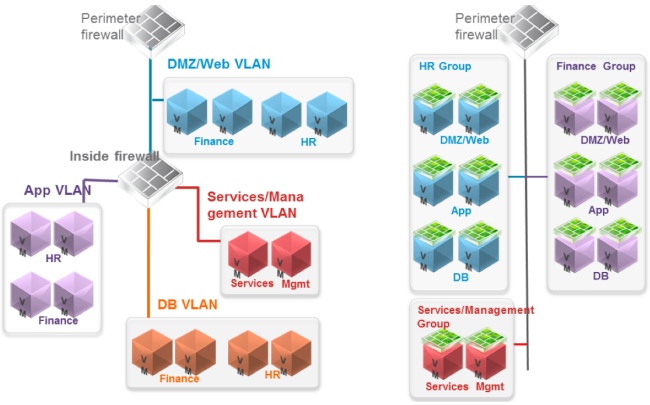

Figure 2.3: Hypervisor-Based Environment Without (Left) vs. With Micro-Segmentation (Right)(Suresh, 2016)

Host-Based or Agent-based Micro-segmentation

Agent-based micro-segmentation involves installing software agents directly on servers, virtual machines, or endpoints to monitor traffic and enforce security rules locally (T V et al., 2023). These agents collect information about network behaviour and communicate with a central controller or policy engine, which uses this data to create and update security policies in real time (Suri et al., 2003). The agents typically rely on the host's built-in firewall to enforce these policies at the operating system level (Tuglular, 2008).

Figure 2.4: No Micro-segmentation (Left) vs Host-Based Micro-Segmentation (Right) (Mastering Nutanix, 2020)

3

Literature Review

3.1. Literature Search Strategy

Despite extensive research on micro-segmentation, significant challenges continue to hinder its adoption in real-world environments (Xie et al., 2021).This section aims to find the main challenges associated with the implementation of Zero Trust Architecture (ZTA), with a particular focus on micro-segmentation as a foundational component. The overall study selection process is summarized in Figure 3.1, which follows the PRISMA flow diagram approach (Page et al., 2021).

As illustrated in Figure 3.1, the Identification phase began with a systematic search strategy using the following Boolean query:

> (("Zero Trust Architecture" OR "Zero Trust model" OR "ZTA") AND
> ("microsegmentation" OR "micro-segmentation" OR "micro segmentation") AND
> ("challenge" OR "implementation challenges" OR "barrier" OR "technical barriers" OR "limitation"))

This query was designed to capture a diverse range of studies addressing technical, operational, organizational, and policy-related challenges in ZTA implementation. The combination of OR operators ensures the inclusion of studies using different terminologies for the same concepts, while the AND operators refine the search to focus on research that explicitly discusses both Zero Trust and micro-segmentation.

Applying this search query resulted in the retrieval of approximately 554 relevant papers from multiple academic databases. After removing 15 duplicate and non-English records, a total of 539 articles remained for screening. To refine the dataset to the most relevant studies, a structured screening process was applied.

In the first stage, the titles and abstracts of these articles were reviewed to check if they were relevant to the topic of micro-segmentation within Zero Trust Architecture (ZTA). At this point, the inclusion and exclusion criteria (see Table 3.1) were used to remove articles that clearly did not fit the purpose of the review. As a result, 300 articles were excluded.

Table 3.1: Inclusion and Exclusion Criteria for Literature Review

Inclusion Criteria
▪ Articles written in English. ▪ Literature published after 2020. This criterion is chosen due to the rapidly evolving nature of technology, to focus on the most recent developments and challenges in micro-segmentation within Zero Trust Architecture (ZTA). ▪ Industry reports, whitepapers, and case studies that provide empirical evidence or practical insights into ZTA adoption.
Exclusion Criteria
▪ Articles that focus exclusively on ZTA with minimal or no attention to micro-segmentation. ▪ Literature that provides only introductory overviews of Zero Trust or micro-segmentation and does not offer technical depth, such as analysis, discussion of implementation challenges, or practical architectural guidance.

This step left 239 articles that were considered potentially relevant and were selected for full-text reading. In the second stage, each of these articles was read in full to assess their eligibility in greater depth. During this step, the inclusion and exclusion criteria (see Table 3.1) were applied more strictly. Articles were excluded if, upon detailed examination, they did not sufficiently address micro-segmentation within the context of Zero Trust Architecture, or if they lacked substantive insights, empirical data, or practical implementation details. Some were removed due to being overly theoretical, lacking methodological transparency, or repeating well-established introductory information without new contributions. In total, 194 articles were excluded during the full-text review and 45 articles were included in the final review. An overview of the study selection process is presented in Figure 3.1.

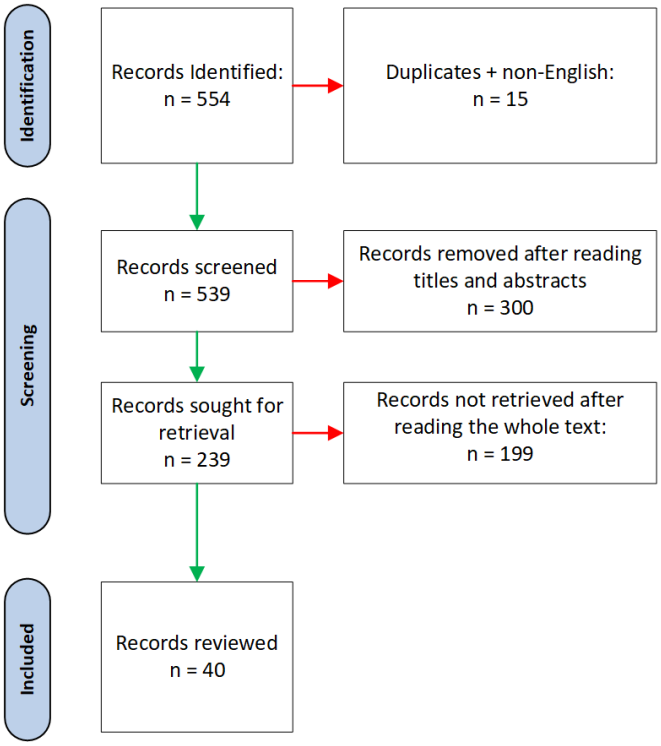

Figure 3.1: Flowchart of the study selection process

3.2. State-of-the-Art

This section reviews the current academic and industry understanding of Zero Trust Architecture and micro-segmentation. It explores key concepts, technical approaches, and existing challenges, with a focus on how these strategies are applied in environments that include legacy systems. The aim is to identify knowledge gaps and limitations in existing models to better position the need for this study.

Micro-Segmentation with Software-Defined Networking (SDN)

Several studies explore SDN-based micro-segmentation as a foundation for implementing Zero Trust security in modern networked environments. Arora and Hastings (2024) propose a layered architecture using open-source tools Istio and Calico within a Kubernetes environment. Their system enforces segmentation at the network, resource, and application levels, with Calico managing network-layer policies through SDN and Istio handling service-level authentication using mutual TLS. The architecture supports identity-based policy enforcement across multi-cloud containerized deployments and illustrates how segmentation can operate across communication layers.

Desai and Patil (2020) conduct a literature-based investigation into how SDN can enable fine-grained, automated micro-segmentation in cloud-native setups. Their study emphasizes the role of centralized SDN controllers that dynamically apply policies based on workload identity, behaviour, and contextual factors. They argue that this approach improves scalability, automation, and policy uniformity, especially in multi-tenant and dynamic environments.

Zanasi et al. (2024) introduce a modular SDN-based framework grounded in Zero Trust principles, using the concept of Security Domains to create independently governed network segments. Each domain uses metadata-rich digital certificates to enforce access control, and multi-domain certificates are designed to allow secure interactions without compromising individual domain policies. Their system formally defines rules to prevent policy override, ensuring containment even when one domain is compromised. Kim et al. (2024) address the rigidity of traditional network architectures in defence systems and propose SDN-based logical segmentation to replace static, manually configured designs. Their approach is grounded in a conceptual analysis of current technologies and Zero Trust requirements. They suggest that SDN can enable more adaptive and centrally controlled security postures suited to dynamic threat environments.

While their approaches differ, these studies share common limitations. Most rely on specific infrastructures such as Kubernetes or SDN-compatible platforms, which may hinder general applicability. Several proposals remain at the conceptual or prototype stage without real-world deployment.

Hypervisor-Level Micro-Segmentation and Virtualization Platforms

Koskinen (2020) explores the application of VMware NSX for micro-segmentation within virtualized data centers, focusing on its suitability for military computing environments. The study examines NSX's distributed firewall architecture, which enforces security policies at the hypervisor level and eliminates reliance on centralized inspection points. Through a controlled lab setup, traditional firewall-based segmentation is compared with NSX-enabled segmentation in terms of security, performance, and application throughput. Findings show that NSX improves segmentation granularity without introducing measurable overhead, and its integration with existing infrastructure supports gradual deployment in legacy systems.

Al-Ofeishat and Alshorman (2023) extend this direction by integrating VMware NSX-T with Sky ATP and a policy enforcer to build a more advanced segmentation environment suited for large, complex infrastructures. Their evaluation compares VLAN-based segmentation, standard NSX, and their enhanced NSX-T model across criteria such as security, cost, complexity, workload mobility, and threat detection. Results show that the enhanced setup offers superior performance across nearly all metrics, demonstrating the potential of SDN-enabled micro-segmentation in hybrid and multi-cloud scenarios.

Despite their promising results, both studies are limited to lab-based evaluations and lack validation in real-world deployments. Additionally, they assume the presence of a fully virtualized network environment and rely on compatibility with VMware-based infrastructure, which restricts the applicability of their findings in non-virtualized or heterogeneous systems.

Agent-Based Micro-segmentation

Agent-based mechanisms have been widely adopted to enforce micro-segmentation policies in cloud and distributed environments. Sheikh et al. (2021) implement a zero trust security model in a Microsoft Azure cloud setting using the Illumio Adaptive Security Platform. Their approach relies on Virtual Enforcement Node agents deployed on hosts to collect telemetry data and coordinate with a central Policy Compute Engine, which enforces whitelist-based rules based on port and protocol. The system effectively isolates workloads and maintains stable traffic flows under test conditions. However, it does not address the creation of policies themselves, and challenges remain in mapping security rules to business logic and scaling the system in high-traffic or heterogeneous infrastructures.

In contrast, several studies aim to tackle the policy creation challenge identified by Ma et al. (2023), who argue that solutions like Illumio, while effective at enforcement, lack mechanisms for recommending or generating segmentation policies. Ma proposes a system that combines static code and configuration analysis with dynamic traffic monitoring and machine learning to automatically generate behaviour-aware access control policies. The architecture includes a translation engine that converts abstract policy models into low-level rules, addressing the policy-definition gap.

Arifeen et al. (2021) also contribute to automated policy creation with a model designed for Industrial IoT environments. Their method groups devices based on traffic behavior using clustering and employs a decision tree classifier to detect malicious traffic. From this analysis, security rules are generated to define micro-segment boundaries, thereby limiting lateral movement. The system demonstrates strong classification performance and containment potential. Although not explicitly framed as agent-based, the model implies host-level data collection and rule enforcement, aligning functionally with agent-driven designs.

Xie et al. (2021) present a conceptual Zero Trust micro-segmentation framework that addresses both enforcement and policy generation. The system features a centralized policy control centre that collects traffic data from agent plug-ins, performs intelligent traffic analysis, and automatically generates segmentation policies. These policies are then reviewed by administrators before being enforced by the agents, which are responsible for adapting internal traffic isolation based on current endpoint conditions. The architecture includes modules for north-south authentication, east-west segmentation, and environment-aware policy adaptation, though it remains untested in operational settings.

Liu et al. (2024) propose a hierarchical micro-segmentation framework for next-generation networks using large language model-empowered agents. These agents operate within a graph-based segmentation model that integrates both trust-level and physical infrastructure attributes. The LEGD algorithm generates optimized segmentation structures, while the LEGD-AM mechanism adapts them dynamically based on service changes and endpoint trust updates. The agent's role involves applying policy filters and guiding adaptive reconfiguration using a reward-based control loop. Experimental results show notable improvements in efficiency and resilience, though deployment remains limited to a testbed environment.
As also been identified in the previously discussed methods, these approaches share a common limitation, which is the lack of real-world deployment. All are evaluated in simulated or conceptual environments, leaving questions about their practical applicability in diverse infrastructures unresolved.

Micro-Segmentation through Traditional Segmentation Methods

Traditional segmentation techniques rely on mechanisms such as VLANs to create security segments, and firewalls to control traffic between defined network zones (2023). These methods typically operate by grouping devices based on static attributes like IP address or physical location, and enforcing access restrictions through manually configured rules (Shaik, 2024). While it offers simple initial implementation and effective basic isolation, it lacks the flexibility and granularity required to respond to dynamic workloads, identity-based access, or real-time threat conditions. (Álvarez et al., 2023).

Some studies have explored how traditional segmentation mechanisms like VLANs and firewalls can be adapted or combined with additional technologies to implement micro-segmentation with improved isolation and control. Rocha et al. (2021) introduce a Zero Trust-based security model that incorporates micro-segmentation and Next-Generation Firewalls to prevent lateral movement and advanced persistent threats in LAN environments populated with IoT devices. The model segments the network into VLANs for users

and devices, each protected by a firewall configured as a Policy Enforcement Point. Threat modelling is performed using attack trees, comparing the system before and after segmentation. The evaluation, conducted in a simulated LAN using Cisco Packet Tracer, shows that micro-segmentation significantly improves access control and limits unauthorized access to critical assets such as IP cameras.

Li et al. (2024) propose a hybrid segmentation approach that combines VLAN and VxLAN technologies to enhance east-west traffic isolation in large-scale data centres. Their method introduces a dual-tagging mechanism that maps each VLAN ID to a VxLAN Network Identifier, allowing traffic to remain logically segmented even when routed across Layer 3 infrastructure. The system enables scalable segmentation without changing IP addressing schemes and supports a greater number of segments than VLANs alone. Experimental simulations using GNS3 demonstrate that this technique significantly reduces the risk of lateral movement compared to VLAN-only configurations.

Granular and Context-Aware Access Control Policies

As explained in Chapter 2, Zero Trust Architecture defines access control through distinct logical components and emphasizes the use of fine-grained policies that consider contextual and identity-based factors. Building on this foundation, recent research has proposed various approaches to designing and enforcing advanced access policies that enhance micro-segmentation effectiveness. Chandramouli (2022) emphasizes that assigning unique identities to systems and devices is fundamental to enabling scalable and secure micro-segmentation. The author argues that identity-driven policies provide the necessary abstraction to manage complex and dynamic environments, allowing segmentation decisions to adapt based on contextual and trust-related factors.

Zaheer et al.(2019) propose ezTrust, a network-independent micro-segmentation framework for cloud-native microservice environments. Unlike traditional enforcement methods, ezTrust uses metadata such as application version and runtime context to define access control identities. Packets are tagged at the source with these identity tokens, and receiving agents apply access decisions based solely on this verified identity context. Evaluated in a Kubernetes testbed, ezTrust demonstrates low performance overhead and strong enforcement accuracy. However, the system's complexity, along with its reliance on containerized infrastructures, limits its applicability in broader enterprise settings.

Chen et al. (2020) present a dynamic access control model within a Zero Trust framework tailored for 5G-enabled smart healthcare. Their approach defines micro-segmentation policies across four dimensions: subject, object, behaviour, and environment. This model allows real-time evaluation of user behaviour and system context to continuously adjust access control decisions. Tested in an industrial-grade simulation environment, the model proves technically feasible but requires high computational resources.

Although differing in scope and technical design, these approaches share a common emphasis on adaptive, identity-aware access control as a foundation for effective micro-segmentation. While they demonstrate strong potential, their reliance on modern infrastructure raises concerns about deployment feasibility in real-world and heterogenous environments.

Legacy Infrastructure as a Barrier to Modern Micro-Segmentation Methods

Across all implementation methods, whether SDN based, hypervisor level, agent driven, based on enhanced traditional techniques, or focused on advanced access policy models, studies consistently share a core set of limitations. The most common one is the lack of real-world deployment. Most proposals are evaluated only in conceptual models, testbeds, or emulated environments, leaving open questions about their practical feasibility under real operational conditions. Another recurring limitation is the assumption that organizations already operate modern and compatible infrastructures. These include systems capable of supporting context aware and identity based access control, SDN controllers, advanced hypervisors, or host level agents. In practice, this assumption rarely holds.

Many organizations still rely heavily on legacy network systems, which were not originally designed to meet modern cybersecurity standards (SnapLogic, 2025). Multiple studies emphasize that legacy infrastructure poses a significant barrier to the implementation of Zero Trust Architectures. Bellamkonda (2022) highlights that legacy systems are often incompatible with modern Zero Trust technologies, requiring intermediary solutions or upgrades for effective integration. Phiayura and Teerakanok (2023) reinforce this view, identifying legacy systems as major obstacles due to outdated authentication mechanisms, incompatibility with modern controls, and elevated resource demands during migration. Ojo (2025) finds that the technological incompatibility and high costs associated with legacy infrastructure significantly hinder zero trust architecture implementation in critical infrastructure contexts. Bell et al. (2024) similarly note that legacy systems introduce integration complexity, financial strain, and increased operational overhead when transitioning to Zero Trust security models.

Beyond hindering Zero Trust adoption broadly, these legacy system limitations directly impact the practical implementation of micro-segmentation strategies discussed earlier. Syed et al. (2022) highlight SDN as a promising enabler for micro-segmentation in Zero Trust Architectures, particularly when combined with technologies like network function virtualization (NFV) and software-defined perimeters (SDP). The authors acknowledge open challenges such as policy scalability and the difficulty of integrating SDN-based approaches into environments that rely on legacy infrastructure. This is largely because legacy systems lack the programmability and centralized control needed for SDN integration, which depends on dynamic policy enforcement and real-time network reconfiguration (Lin and Lin, 2014). For example, many legacy environments use static, hardcoded network configurations that are deeply embedded in proprietary hardware and software stacks, making it difficult to adapt them to SDN's flexible architecture without significant modification (Sokappadu and Mungur, 2021).

Similarly, modern micro-segmentation strategies that depend on virtualization platforms such as hypervisors, encounter serious limitations in legacy environments. Modernization of legacy systems, for example through virtualization, frequently requires substantial refactoring before they can effectively be virtualized or containerized (2020). This transformation is rarely immediate, it typically demands an extensive, phased approach involving careful planning, system decomposition, and gradual migration of services (2022). For instance, the transformation process involves parsing the source code of the legacy system and converting it into a language-independent representation (Hunold et al., 2009). A common issue with legacy systems, however, is that the original developer is often no longer available, leaving current developers to spend significant effort deciphering the code's intent and functionality (Hunold et al., 2009). Ponnusamy and Eswararaj (2023) noted that this inherent inefficiencies of legacy systems can strain resources and inflate costs. As a result, even when retrofitting is attempted, the performance and security trade-offs often outweigh the benefits (Gaska).

Agent-based micro-segmentation approaches also face serious limitations in legacy environments. Sundareswaran et al. (2023) found that installing enforcement agents on legacy operating systems like older versions of Linux and Windows was cumbersome, with low success rates and limited scalability, making concurrent deployment impractical without automation. Kang et al. (C. Kang et al., 2022) noted that traditional host-agent-based micro-segmentation can strain legacy systems by consuming excessive resources, potentially disrupting normal program execution, especially when complex functions are installed on older hardware.

Following the deployment challenges associated with different micro-segmentation approaches, advanced access policy mechanisms, such as context-aware, identity-driven, and continuously adaptive controls, face their own limitations in legacy settings. Ahmadi (2024) emphasizes that while these policies enhance Zero Trust implementations through real-time analytics and contextual verification, they often remain impractical due to technical complexity, high operational demands, and incompatibility with legacy systems. Rajasekharan (2025) further notes that integrating attribute-based access control (ABAC) in legacy environments is hindered by policy complexity, performance constraints, and the absence of contextual data sources typically required for fine-grained enforcement.

Efforts Toward Legacy System Integration into Micro-segmentation

Several efforts have been made to address the integration challenges posed by legacy systems and to incorporate them into micro-segmentation strategies within Zero Trust Architecture (ZTA).
Rose et al. (2020) introduce the Zero Trust Architecture (ZTA) model as a strategic response to evolving cybersecurity threats, emphasizing the principle of "never trust, always verify" to protect modern enterprise environments. To address legacy systems, the authors suggest incremental integration through overlays or gateways that act as policy enforcement points, allowing these systems to participate in a ZTA without being re-architected. However, the guidance provides only a high-level mention of placing legacy systems into enclaves without detailing how trust boundaries within those enclaves are maintained or how policy enforcement is applied internally. No specific tools, technologies, or operational mechanisms are proposed to support this integration, nor is there discussion of how the size, scope, or composition of such enclaves should be determined.

In their, Køien (Køien, 2021)proposes a structured framework for integrating legacy devices (LDs) into modern Industrial Control Systems (ICS) using principles rooted in Zero Trust (ZT) security. The work introduces a 12-rule model tailored to mitigate the vulnerabilities of legacy devices typically unprotected, unpatchable, and irreplaceable components such as sensors or actuators within increasingly digitized and cloud-integrated environments. The author lays out a model architecture involving a Legacy Interface Function (LIF) and a Legacy Encapsulating Gateway (LEG), which together encapsulate legacy components and implement ZT principles such as entity authentication, access control, channel separation, mandatory logging, and data validation. Findings are cantered on the feasibility of isolating and securing LDs through layered defences without replacing them, using encapsulation and rigorous policy enforcement. This method addresses real-world challenges where full replacement of legacy infrastructure is technically or economically infeasible.

Tyler and Viana (2021) present a comprehensive framework to support healthcare organisations in transitioning from traditional perimeter-based security models to a zero-trust architecture, with particular attention to the limitations posed by legacy systems and unmanageable medical devices. Micro-segmentation is addressed in depth, proposing techniques for isolating vulnerable legacy systems. The authors test two methods, proxy servers and hardware firewalls. Proxy servers are used at the application layer to mediate communication between legacy systems and the rest of the network. This method is ultimately rejected due to high latency and reduced security, making proxy use impractical in time-sensitive medical environments. In the second method, legacy devices are placed behind firewall clusters with policy enforcement based on access control lists. Although clustered firewalls are found to offer better performance and stronger security, they introduce trade-offs in packet loss during failover. The findings show that the proposed architecture successfully limits the spread of simulated attacks and achieves significant isolation of assets. However, the framework is constrained by its reliance on simulated environments and assumptions about network configurations. It also acknowledges that the cost and complexity of implementing firewall clusters or behavioural monitoring may not be feasible for all healthcare providers. Furthermore, the framework is specifically designed for healthcare environments and may not transfer effectively to other sectors with different security priorities, network structures, or regulatory requirements.

Paul and Sheriffdeen (2022)present a structured approach for organizations to transition from traditional perimeter-based security models to a zero-trust framework. The authors outline a phased methodology to facilitate this transition. The process begins with assessing the current security posture to identify critical assets, vulnerabilities, and the limitations of legacy systems. This is followed by defining granular access control policies based on least privilege, utilizing role-based or attribute-based access control and enforcing multi-factor authentication. The study concludes that this phased approach enables organizations to overcome the limitations of legacy systems by gradually transitioning toward a Zero Trust framework. However, the framework assumes that legacy systems are fundamentally incompatible with the technologies required for Zero Trust, such as advanced identity management, micro segmentation, and continuous monitoring, and therefore must be significantly upgraded or replaced. In practice, many legacy systems lack the technical capacity to support such upgrades, and full replacement is often not a viable option due to high investment costs or the risk of disrupting critical services.

3.2.1. Research Gap

Existing micro-segmentation strategies are largely built around advanced technologies such as SDN or hypervisors and emphasize context-aware or identity-based access controls. While these methods offer fine-grained control and strong security enforcement, they typically rely on modern, flexible infrastructure being already in place. As a result, they are difficult to apply in environments that rely on legacy systems, which lack the compatibility and flexibility these methods require. Most studies are limited to controlled or simulated settings and fail to address the practical constraints of real-world legacy infrastructures. Some efforts have been made to integrate legacy systems into micro-segmentation strategies, typically by isolating outdated components using firewalls, gateways, or proxy servers. Most of them often present isolated technical solutions without considering the limitations and characteristics of the existing network. These models provide high-level recommendations or sector-specific strategies but do not deliver generalized, repeatable frameworks that organizations can follow across different environments.

Despite various proposed solutions, there remains no comprehensive framework to help organizations with legacy systems select appropriate micro-segmentation strategies or take the first steps toward implementation based on their network characteristics. What is also needed is a detailed, step-by-step implementation roadmap along with methods and criteria for evaluating the implementation of micro-segmentation. Moreover, the absence of clear criteria and evaluation methods to determine when micro-segmentation is complete and fully functional highlights a substantial gap in current research, as noted by Fernandez and Brazhuk (2024).

Addressing this gap, the present study aims to guide organizations with legacy systems in selecting appropriate micro-segmentation strategies based on their network characteristics, while also addressing the technical and operational challenges of implementation. Specifically, it seeks to define criteria for successful implementation, develop methods to evaluate whether these criteria are met, and provide a detailed, step-by-step implementation guide to support a structured transition to Zero Trust Architecture.

3.2.2. Research Objectives

This study aims to develop a framework that guide organizations with legacy systems in selecting appropriate micro-segmentation strategies based on their network characteristics, addressing the technical and operational challenges of implementing these strategies. It also identifies criteria for successful implementation and explores methods to assess whether these criteria are being met.

3.2.3. Research Questions

This study investigates how organizations can make structured decisions about micro-segmentation in environments that include legacy systems. The goal is to understand how different strategies are selected, what constraints shape those decisions, and how implementations can be evaluated in a practical context. The research is guided by one main question and three sub-questions that define the scope of inquiry.

Main Research Question

> *"How can organizations choose and implement micro-segmentation strategies that align with their network architectures and the technical constraints posed by legacy systems?"*

Sub-questions
To support the investigation of the main research question, the following sub-questions are considered:

1. *What are the common strategies for implementing micro-segmentation and how do network characteristics influence the choice of these strategies?*
2. *What technical and operational challenges do organizations face with legacy systems when applying these micro-segmentation strategies, and how can these challenges be addressed?*
3. *What steps are involved in the implementation of micro-segmentation and how can these steps be structured into a repeatable roadmap?*
4. *What criteria define a successful micro-segmentation implementation, and what methods can be used to evaluate whether these criteria are met?*

4

Research Methodology

4.1. Data Collection

This study adopts semi-structured interviews as the primary method for collecting qualitative data from cybersecurity professionals. This approach was selected because the specific insights required were not sufficiently documented in the existing literature. In particular, there is a lack of information on how practitioners implement and adapt micro-segmentation in the presence of legacy systems. As such, a qualitative, exploratory method was necessary to access expert knowledge embedded in practice. Semi-structured interviews provide a flexible yet guided format, allowing the researcher to explore key topics while also accommodating unanticipated but relevant details (DiCicco-Bloom and Crabtree, 2006; Kallio et al., 2016). This method is particularly effective when addressing complex, under-researched subjects, as it enables both comparability across participants and in-depth, context-specific narratives (Adams, 2015; Gill et al., 2008). It also facilitates rapport and adaptability during the interview process, improving the depth and relevance of responses.

4.1.1. Semi-Structured Interviews

The data collection process followed a structured sequence of six steps, as shown in Figure 4.1. In Step 1, approval was obtained from the Human Research Ethics Committee (HREC). Step 2 involved identifying potential participants based on predefined criteria. Although the identification process took place prior to HREC approval, no recruitment or contact was initiated until approval was granted.

Participants were required to have a minimum of five years of experience in security architecture or a closely related field such as network security, cybersecurity engineering, or infrastructure security. In addition, they needed to have direct experience in implementing or contributing to the implementation of Zero Trust Architecture (ZTA), with micro-segmentation as a core component. All experts were identified from a single organization, using the organization's internal website, which provided access to their professional profiles, including areas of expertise, years of experience, CVs, and completed projects relevant to the research focus.

The aim was to identify at least 15 participants to ensure sufficient coverage across roles and responsibilities within the organization. This number was selected to balance depth and diversity of insights while keeping the data manageable for qualitative analysis. Participants were chosen to reflect a variety of functions and experience levels relevant to the implementation of micro-segmentation in environments with legacy systems. The final list of selected participants, their expertise and the years of experience is presented in Table 4.1.

Table 4.1: Overview of Participants, Their Expertise, and Years of Experience (Y-of-E)

Identifier	Expertise	Y-of-E
P1	Security Network operations and design	26
P2	Security Network Architecture and implementation	13
P3	Security Network Architecture and implementation	10
P4	OT Security Network Architecture and implementation	15
P5	OT and IT Security Network Architecture and implementation	15
P6	Security Network Architecture and implementation	17
P7	Security Network Architecture and implementation	12
P8	Cybersecurity management and strategic planning	10
P9	Network security architecture + IAM	7
P10	Network engineering and Architecture	8
P11	Security Network Architecture and implementation	13
P12	Network security, Identity and access management	20
P13	Security Transformation Strategy and management	7
P14	Cybersecurity advisory and implementation	9
P15	Cybersecurity transformation leadership	25
P16	Cybersecurity transformation leadership	21

Step 3 involved scheduling the interviews. Participants were contacted via email, provided with an information sheet and consent form, and invited to take part in a one-on-one interview. Interviews were then scheduled and conducted using the Microsoft Teams video conferencing platform. In Step 4, the interviews were conducted and audio-recorded. Each interview followed a semi-structured format based on an interview guide covering three main themes: current segmentation strategies, challenges associated with legacy systems, and evaluation or control mechanisms. Additionally, brief notes were taken during the interviews to document contextual observations and initial impressions.

In Step 5, interviews were transcribed verbatim, and in Step 6, all transcribed data were anonymized. Identifying information was removed and replaced with coded identifiers (P1, P2, ...). All data, including audio files, transcripts, and notes, are securely stored in encrypted folders on TU Delft's OneDrive environment, with access restricted exclusively to the research team.

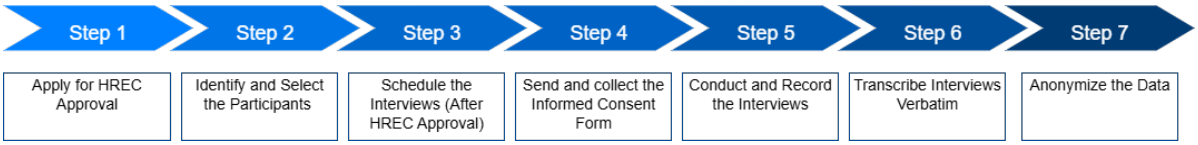

Figure 4.1: Data Collection Steps

4.1.2. Ethical Considerations

This study was conducted in accordance with ethical research standards and received approval from the Human Research Ethics Committee (HREC) at TU Delft. Participants were informed of the research aims, procedures, and data handling practices via a digital information sheet sent prior to the interviews. Informed consent was obtained through digitally signed forms, and no data collection commenced without this consent. The study involved collecting personal data, including job titles, professional experience, interview recordings, and interview transcripts. The transcripts were anonymized after transcription and prior to analysis to protect participant identities. Data was securely stored within TU Delft's OneDrive environment, with access restricted to the researcher and the academic supervisor. The processing of personal data was based on informed consent and did not involve any high-risk or sensitive categories, as confirmed by the university's data management guidelines and supported by the Data Management Plan (DMP) reviewed by the faculty's data steward. The research output includes only anonymized segments from the transcripts, with no raw data or personally identifiable information disclosed.

4.2. Data Analysis

This study employs thematic analysis as the primary method for analysing data from semi-structured interviews with cybersecurity professionals. Thematic analysis is a widely used and methodologically flexible approach in qualitative research, designed to identify, analyse, and interpret recurring patterns of meaning across a dataset (Braun and Clarke, 2006; Clarke and Braun, 2016). It is particularly well-suited to semi-structured interviews, which allow for both comparability across participants and the exploration of context-specific detail (Nowell et al., 2017). The approach accommodates inductive and deductive reasoning, facilitating the integration of new insights with existing theoretical frameworks (Braun and Clarke, 2006). Furthermore, thematic analysis aligns with qualitative research best practices, as it promotes transparency, reliability, and clear documentation throughout the analytical process (Campbell et al., 2013; Nowell et al., 2017).

4.2.1. Data Analysis Process

As a continuation of the thematic analysis framework, this study adopts the eleven-step process proposed by Duruk (2019) to analyse the interview data. The process begins with the preparation of the starting documents, including verbatim transcripts and written notes taken during the interviews (Step 1). Each transcript is then summarized qualitatively to identify initial impressions and key points (Step 2). From these summaries, preliminary strands (recurring ideas or patterns) are identified in each interview (Step 3). In Step 4, these strands are reviewed and regrouped to eliminate overlaps and improve thematic clarity. The next step (Step 5) involves classifying coded data segments (units of meaning) under the refined strands. These coded units are then combined across interviews to form a shared thematic structure (Step 6).

A reassessment of all strands is performed (Step 7) to ensure consistency, accuracy, and alignment with the data. Following this, more detailed levels of analysis are introduced: Level 1 dimensions (Step 8) are identified within each strand, representing broader thematic categories. These are further broken down into Level 2 dimensions (Step 9), providing deeper thematic resolution. The final stages involve refining the theme structure through splitting and splicing, merging overlapping dimensions and removing unnecessary ones (Step 10). This leads to the final delimitation (Step 11), where a clear, coherent set of themes or codes and sub-codes (see Figure A)is established for interpretation and discussion. Figure 4.2 gives an overview of the data analysis steps discussed in this section).

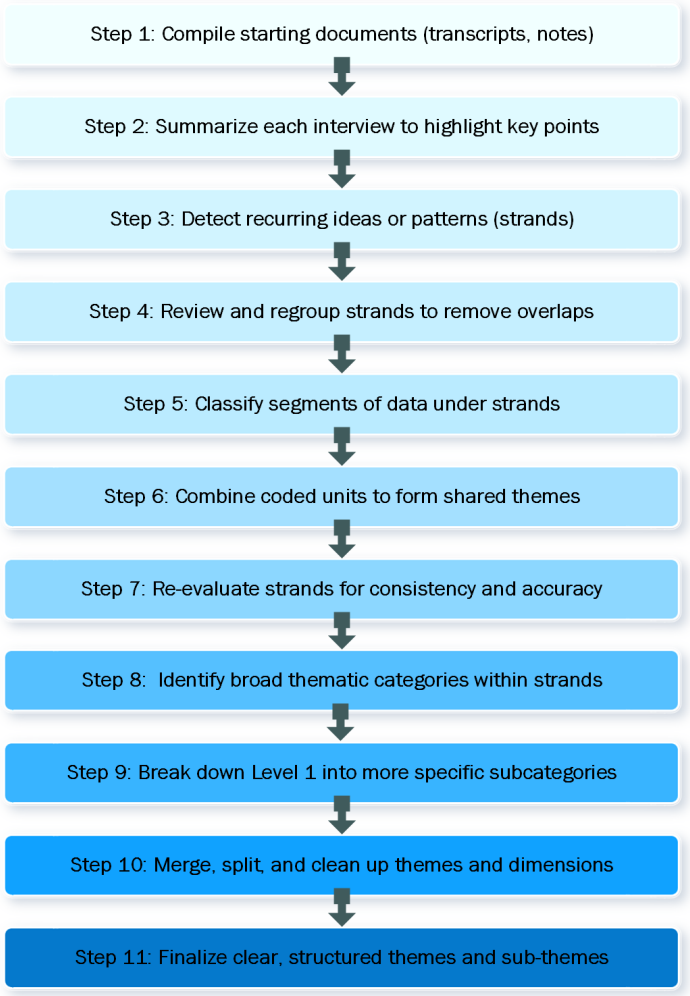

Figure 4.2: Data Analysis Process in Steps

5 Results

This chapter presents the key findings from the expert interviews conducted as part of this study. The results provide insight into how micro-segmentation is understood in practice and how it is adapted to the constraints of legacy-heavy environments.The chapter begins by outlining how practitioners define and approach micro-segmentation in practice. It then explores the challenges and constraints that influence decision-making, followed by a discussion of the criteria used to evaluate segmentation options. Finally, the chapter explains how these findings were synthesized into a practical framework to support micro-segmentation strategy selection in legacy-heavy environments.

5.1. Definition of Micro-segmentation In Practice

Although the primary aim of the interviews was to understand implementation practices, one of the recurring theme emerged early in the conversations: each expert began by describing what micro-segmentation means in their specific context. Experts described micro-segmentation in diverse ways, possibly shaped by their roles, technical environments, and organizational maturity. Below is a synthesis of how the interviewees defined micro-segmentation. Similar views have been grouped together to reflect shared themes.

Granular Control Within the Internal Network

Many experts (P1, P2, P4, P5, P6, P7, P8 and P16) described micro-segmentation as the process of dividing a network into distinct, tightly controlled segments using fine-grained access policies. These controls are typically implemented through firewalls, filtering rules, or access management systems. Micro-segmentation goes beyond traditional perimeter defences by applying traffic control at the level of individual devices, virtual machines, or workloads. This enables organizations to manage communication even between systems that reside on the same subnet or broadcast domain. As P7 explained:

> *"The goal is to create multiple zones within the environment and enforce access controls between them, even if systems are technically close or running on the same host. That's what sets micro-segmentation apart from traditional segmentation." (P7, 2025)*

This perspective highlights one of the the core purpose of micro-segmentation which is isolating communication paths and only allowing explicitly authorized interactions between systems.

Controlling East-West Traffic to Prevent Lateral Movement

P3, P6, P7, and P10 defined micro-segmentation primarily through its security function, focusing on its ability to control east-west traffic within internal networks. Rather than focusing on where controls are applied, they described it in terms of its purpose. For them, micro-segmentation is about preventing lateral movement after an attacker gains access. East-west traffic refers to the flow of data between systems within the same data centre or subnet, unlike north-south traffic, which typically flows between external users and internal services. According to these experts, micro-segmentation is about placing fine-grained controls on this internal communication to limit how far an attacker can move inside the network. One expert explained:

> *"We need to put a policy or a control in place such that the restriction applies even though the servers are in the same zone. This kind of control has to be there to protect server-to-server*

> *communication. So this level of granular control that you apply for the traffic that is happening within the application, within the servers, is called East-West traffic. The protection that you apply to control East-West traffic is called micro-segmentation." (P6, 2025)*

Micro-Segmentation at the Application or Service Level

Interviewees P11 and P13 pushed the definition further down the stack, describing micro-segmentation as control applied at the application or service layer. For them, it's not just about separating network segments, it's about isolating specific services or components, such as APIs, microservices, or even individual files. This view sees micro-segmentation as a way to enforce policy between the smallest possible units of functionality, reflecting a very advanced or cloud-native environment. One of the intreviewee mentioned:

> *"With micro-segmentation, it's more specific, that could be done on application level, virtual machine level... it's more detailed, more specific than traditional network segmentation" (P14, 2025).*

Variable Definitions Based on Organizational Maturity

Some interviewees (P9, P12 and P15) cautioned against assuming a fixed definition. They pointed out that what counts as micro-segmentation depends heavily on the maturity of the organization. In some environments, simply restricting traffic between servers or subnets could be seen as a form of micro-segmentation. In others, that might be considered a basic step, far from a truly granular model. According to these participants, definitions should be seen on a sliding scale rather than a fixed standard.

> *"Some people may consider using a host firewall to block something as micro-segmentation, and I think that's true. I think it's a maturity scale. As you progress, you get better at it, and you do more complex things." (P9, 2025)*

Tiered, Zone-Based Segmentation Approach:

One interviewee (P13) offered a more conceptual critique of the term "micro-segmentation" itself.

> *"For me, micro-segmentation is actually not the best terminology. I would look at zone-based segmentation, meaning that I will create a tiered approach to defining the taxonomy of a zone. A zone can be very large and you can start from a high level and then go smaller to a common denominator like one application or one asset." (P13, 2025)*

He argued that the "micro" in the name can be misleading, as the real focus should be on creating meaningful security zones that align with business needs and risk profiles. His approach begins with broad segmentation and gradually adds more detailed controls in higher-risk areas. He described this as a tiered model of segmentation, where depth and precision grow over time, rather than starting with extremely fine granularity everywhere.

5.2. Micro-segmentation Strategies in Legacy-heavy Environment

Organizations adopt various strategies to implement micro-segmentation depending on their technical landscape, risk profile, and operational constraints. This diversity is especially evident in environments with legacy systems, where flexibility and adaptability are essential.

Based on the data gathered via expert interviews, multiple approaches to micro-segmentation were identified. These approaches vary in terms of how segmentation is achieved, how access policies are enforced, and what tools or technologies are used to support them. Figure 5.1 summarizes the most commonly referenced strategies across the interviews.

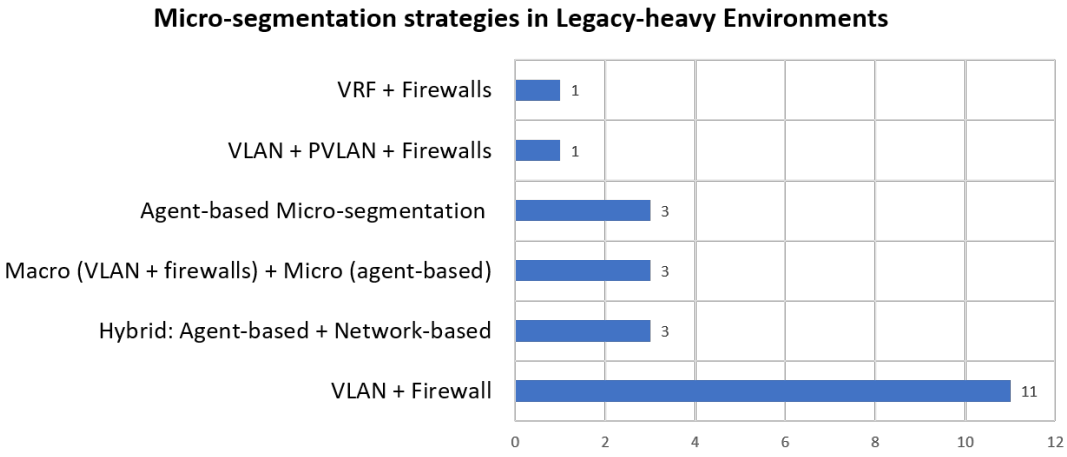

Figure 5.1: Micro-segmentation Strategies

In the sections that follow, each of these strategies will be examined in detail. The discussion will cover how segments are formed, the logic used to group assets, how access policies are designed and enforced, the technologies that enable them, and the practical considerations for applying each method in environments where legacy systems are present.

Network-Based Micro-Segmentation Using VLANs and Firewalls

As shown in the chart, a prevalent micro-segmentation strategy adopted by organizations operating in legacy-heavy environments is the use of VLANs combined with firewalls. This method was highlighted by numerous interviewees (P1, P3, P4, P7, P8, P11, P12, P13, P14, P15, P16) as both practical and compatible with existing network architectures. According to many interviewees, the networks they worked with had evolved from a flat topology. This was often the result of organizations starting small and gradually expanding, leading to environments where systems with varying levels of criticality such as development servers, printers, and production databases were all connected without internal separation. As P1 explained:

> *"The company grows, started with this single production environment [...] and they placed in a very flat network, all their servers, equipment, printer, database in only one segment. That it's really often the case" (P1, 2025)*

While this approach does not fully meet the definition of Zero Trust micro-segmentation, it represents a meaningful step in that direction. It provides logical separation and traffic control between zones based on function or sensitivity, even though it lacks the full granularity, identity-awareness, and dynamic policy enforcement typical of agent-based or software-defined solutions. As P8 put it in simple terms:

> *"This kind of setup is more like Zero Trust version 0.6. It's not perfect, but it's much better than having nothing at all. For many legacy environments, it's the only practical option." (P8, 2025)*

In this method, VLANs are used to logically divide the network into segments based on function, risk level, device type, and other criteria, effectively containing broadcast traffic within each segment. Firewalls then

enforce Layer 3 access policies between these segments, allowing only explicitly permitted IP- and port-based communication. Interviewees like P1 and P14 emphasized that this approach leverages existing infrastructure and aligns well with operational constraints. It provides a manageable entry point into micro-segmentation without requiring radical changes to the environment, making it suitable even in organizations with legacy systems that cannot support modern enforcement mechanisms. P14 addressed this by noting:

> *"A lot of clients are still using just the traditional network segmentation with different zones... not from a Zero Trust principle. But it's a manageable way to get started, especially where legacy systems make deeper segmentation difficult." (P14, 2025)*

Participants such as P7 and P11 noted that the VLAN and Firewall model becomes especially useful in environments where host-based segmentation is not feasible. This typically includes embedded systems or legacy devices with outdated operating systems, limited CPU and memory resources, or architectural constraints that make upgrades impractical or impossible. As P7 mentioned, in these cases segmentation had to occur entirely at the network layer:

> *"When you're working with equipment that simply can't support modern agents, network-layer controls like VLANs and firewalls are often the only realistic option."" (P7, 2025)*

Some interviewees, including P3, P12, and P15, discussed the varying levels of granularity achievable with the VLAN and Firewall approach. These range from coarse-grained segmentation where multiple devices or applications with similar functions or risk profiles share a VLAN to highly granular implementations in which each critical system or device is placed in its own dedicated VLAN. For example, P7 described the grouping as follows:

> *"Often, devices that serve the same function, have similar security requirements, or operate within the same zone like engineering workstations or lab equipment, are grouped together to simplify segmentation and policy enforcement. It's a compromise that keeps things manageable." (P7, 2025)*

In operational technology (OT) environments, as P5 and P4 noted, grouping is frequently guided by structured models such as the Purdue Enterprise Reference Architecture and ISA-95, which segment assets by industrial control system layers and operational roles. P4 highlighted the grouping strategy in OT environment saying:

> *"We often start with models like Purdue or ISA-95. They help us define zones based on process cells or operational areas like packaging, utilities, or production lines—so systems that support the same function end up grouped together." (P4, 2025)*

When grouping OT assets, system function and communication dependencies are key considerations. For example, devices such as PLCs and HMIs, typically found at Levels 1 and 2 of the ISA-95 model, often require direct and time-sensitive communication with each other. As P4 noted:

> *"Communication between Level 1 and Level 2 requires high-speed performance, typically in the range of milliseconds, and in some cases, even microseconds. This level of responsiveness is often impossible to maintain if additional controls introduce latency." (P4, 2025)*

these dependencies must be respected during segmentation, meaning that assets tightly coupled in operation such as a PLC and the HMI are usually grouped within the same security zone. This minimizes latency and avoids breaking essential process interactions. In contrast, systems located at higher levels, such as MES (Manufacturing Execution Systems) or ERP platforms at Levels 3 and 4, can be segmented more aggressively, with communication tightly controlled through firewalls and access policies.

For organizations that pursue finer granularity, the trade-off is a noticeable rise in operational complexity. High levels of granularity demand not only scalable tooling, but also highly skilled personnel and organizational commitment. As P13 explained:

> *"You can have the best tools, but if your team doesn't fully understand the environment or can't manage the operational overhead, fine-grained segmentation quickly becomes unsustainable." (P13, 2025)*

Given these constraints, some organizations find that extreme granularity is not operationally feasible, especially in environments with limited resources or high legacy burden. Recognizing these limitations, many opt for a trade-off, implementing segmentation at a moderately granular level that reduces risk and limits lateral movement, while keeping the architecture manageable within existing resource and infrastructure constraints. P12 highlighted this trade-off and mentioned:

> *"There's always a trade-off, you want tighter control, but the more granular you go, the harder it gets to maintain. At some point, you have to balance security with what your team and infrastructure can realistically handle." (P13, 2025)*

This network-based micro-segmentation method (VLAN + Firewall) supports multiple configuration models, each offering different trade-offs in visibility, control, and operational overhead. One configuration involves creating VLAN subnets that include only devices which need to communicate with one another, with each initial VLAN assigned a dedicated switch and firewall. Within these VLANs, further segmentation is achieved by creating additional nested VLANs. This layered approach provides strong fault isolation, high network visibility, and granular, agentless access control, making it well-suited to legacy-heavy environments. However, it is also highly resource-intensive, requiring extensive cabling and hardware, and introduces substantial configuration complexity and an elevated risk of misconfiguration. In addition, this method can lead to inefficient IP address utilization, as the large number of subnets consumes significant address space and complicates IP planning, particularly in IPv4-based environments. P12 highlighted these issues saying:

> *"You can definitely isolate every small group with its own VLAN and firewall, it's doable and gives you strong control. But it quickly becomes a maintenance headache. Lots of hardware, too much IP address consumption, not a smart choice if you don't have the staff or tools to manage it well." (P12, 2025)*

A second configuration places each individual device in its own VLAN, with access control enforced using a central firewall (or multiple firewalls) capable of handling traffic between these isolated VLANs. While assigning a dedicated VLAN to each critical system can provide stronger isolation, high network Visibility, and better alignment with Zero Trust principles, it also introduces significant overhead in terms of configuration, firewall rule management, and long-term maintenance. P3 highlighted this challenge by stating:

> *"In some cases, we had to go as far as putting every critical device into its own VLAN just to make sure we could isolate them properly, especially in regulated environments. But of course, it adds a lot of complexity." (P3, 2025)*

Despite their theoretical advantages, these highly granular configurations are rarely implemented at scale in real-world environments. The operational burden, infrastructure requirements, and IP address consumption make them impractical for most organizations, especially those working with constrained resources or legacy infrastructure. As a result, many practitioners opt for a more pragmatic and commonly adopted method which based on grouping multiple devices/ systems based on shared functions, communication needs, or operational zones, and applying access controls at the segment boundaries. P7 mentioned this method and said:

> *"In most environments I've seen, grouping by function or zone is just more realistic. You might lose some granularity, but it's manageable, especially when you're dealing with legacy systems or don't have a large team to maintain complex firewall rules." (P7, 2025)*

Although this method does not achieve fine-grained access control, interviewees consistently described it as a foundational control that can reduce lateral movement and isolate high-risk systems effectively. It also serves as a baseline for organizations pursuing Zero Trust principles, offering a scalable and relatively low-disruption way to begin restructuring legacy networks. Despite its limitations in granularity, its compatibility with legacy systems, alignment with existing operational models, and capacity for phased implementation make it a cornerstone method in legacy-heavy environments.

Network-Based Micro-Segmentation Using VLANs, PVLANs and Firewalls

This approach, mentioned by P5, extends traditional VLAN-based segmentation by incorporating Private VLANs (PVLANs) to further restrict lateral communication between devices within the same broadcast domain. Segments are initially formed using VLANs to logically separate traffic based on function, criticality, or operational role. PVLANs are then applied within those VLANs to enforce more granular isolation, for example, preventing communication between peer devices or endpoints that must reside in the same subnet due to hardcoded configurations or application dependencies.

> *"PVLANs gave us a way to isolate systems that couldn't be readdressed or modified, especially when agents weren't an option. It let us tighten control without touching the legacy setup." (P5, 2025)*

Access policies are enforced through a combination of Layer 2 PVLAN rules and Layer 3 firewall policies. While firewalls regulate communication between VLANs, PVLANs restrict peer-to-peer traffic within a VLAN itself. For instance, allowing devices to communicate with a central server but not with each other.

> *"We used PVLANs to make sure devices could reach the central server, but couldn't talk to each other directly. Then we layered firewall rules on top to control any cross-VLAN traffic. It gave us tight control without touching the endpoints." (P5, 2025)*

Technologically, this method depends on switches and firewalls that support PVLAN configurations, typically available in enterprise-grade network hardware. In legacy-heavy environments, the approach is often used to introduce micro-segmentation without disrupting existing IP schemas or application behaviour. It is particularly valuable as an alternative in cases where a granular access control is required, but agent-based micro-segmentation cannot be implemented.

Network-Based Micro-Segmentation Using VRF and Firewalls

In router-centric environments, especially common in Operational Technology (OT) or legacy-heavy networks, Virtual Routing and Forwarding (VRF) can be used to create logically isolated Layer 3 routing domains on the same physical router. This method is particularly relevant when the network is already built around routers instead of switches. As P4 pointed out:

> *"VRF is rather used when the network relies on routers. It can be effective, but in most environments, VLANs are more common because they're easier to implement." (P4, 2025)*

This method enables logical separation without needing to redesign the switching infrastructure, and can scale efficiently for large, distributed networks. Firewalls are typically deployed at the core layer, enforcing policy across VRF boundaries and ensuring tight control over east-west and north-south traffic. As P4 explained:

> *"VRF works well in networks that are built around routers. It allows for clean separation between domains and can scale effectively, but it's not commonly used unless the infrastructure is already designed for it." (P4, 2025)*

However, VRF-based segmentation introduces its own challenges. It requires advanced configuration, not all routers support VRF, and the operational complexity can be high, especially in environments without consistent network documentation or routing expertise. P4 highlighted these challenges and said:

> *"VRF depends on having the right routing infrastructure in place. It's not just about turning it on, you need the expertise to configure it properly, and if the environment isn't well documented, it can become very difficult to manage." (P4, 2025)*

If the routers in place use hard-coded static routes and cannot be replaced, organizations may fall back on isolating the routers and any dependent systems, applying strict boundary controls via perimeter firewalls. While this alternative is easier to implement and minimally disruptive, it only provides coarse-grained segmentation and offers limited visibility into intra-zone activity:

> *"When routers are fixed or hard-coded, we don't try to change them we just isolate the entire block and control what goes in or out with firewalls. It's simple, but you lose visibility inside that segment." (P4, 2025)*

Agent-based Micro-segmentation

Agent-based micro-segmentation is a method that enforces access control policies at the host level through the use of software agents deployed directly on endpoints. According to interviewees P5, P6, and P9, this strategy is seen as one of the common strategies used to implement micro-segmentation and control east-west traffic within data centres and legacy-heavy environments, particularly when visibility and granular enforcement are top priorities.

> *"Agent-based micro-segmentation is a common approach, especially when you need deep visibility and precise control at the host level [...] It allows us to define access policies with much more granularity than traditional network methods." (P6, 2025)*

In this method, Systems are grouped logically using labels or tags that reflect their application context or operational role, and communication is then permitted only within these defined groups.

P5 and P6 mentioned that in environments where a high proportion of systems can support software agents, this method preferred and can be applied as a stand-alone solution. It leverages the host's native firewall capabilities to enforce fine-grained control, thereby eliminating the need to restructure the existing network topology. P6 articulated this point saying:

> *"These agent-based tools typically use the underlying host firewall of the OS... These agents will leverage that inbuilt capability to implement the rule... to control the policy... . So these kinds of solutions are good because I don't need to restructure my network. I don't need to go and deploy firewall, I don't need to touch my VLAN. I can just deploy it on the systems." (P6, 2025)*

While agent-based micro-segmentation is widely applicable, interviewees noted that certain legacy systems such as outdated Windows servers, industrial control devices, or embedded platforms may pose challenges due to OS limitations, resource constraints, or architectural incompatibilities. In these situations, organizations may fall back on network-based controls to maintain segmentation for these systems while still applying agent-based enforcement elsewhere. P2 also discussed about this and mentioned:

> *"When they [legacy systems] are not compatible with agents, it's better to isolate them... to make sure that we are limiting the blast radius... What we are following... is to isolate the legacy systems in a dedicated VLAN at Layer 2 and then limiting the traffic filtering for those legacy systems at firewall level... allowing only the specific source and destinations with specific ports." (P2, 2025)*

When agent support is only partial, a hybrid enforcement model is often considered. The proportion of systems that can support agents depends on several factors, including cost constraints, licensing models, organizational priorities, and the availability of skilled personnel to manage the infrastructure. This leads many organizations to implement segmentation using both host-based agents and network-level controls simultaneously, depending on what is technically and operationally feasible for each system. P5 discussed this issue and explained:

> *"When you can't get full agent coverage due to technical limitations on some systems, you end up needing a hybrid model. But that comes with its own challenges. Now you have to manage both agent-based and network-based segmentation, which means more tools, more policies, and more people to keep it all running smoothly." (P5, 2025)*

Agent-compatible systems are brought under host-level policy enforcement, while agent-less systems are isolated using traditional methods such as VLANs, dedicated firewall rules, or proxy-based inspection. However, when fewer than 40–50% of systems can support agents, a hybrid approach may not be justified. The limited coverage often fails to offset the added complexity, licensing costs, and operational overhead of maintaining an agent-based solution. In such scenarios, organizations typically opt for a network-based strategy. This consideration was highlighted by P2:

> *"There's no strict standard, but generally if fewer than 40 or 50 percent of your systems can support agents, it's hard to justify a hybrid setup. The overhead just doesn't pay off" (P5, 2025)*

As mentioned by P2, P5 and P10, some organizations go beyond only agent-based or hybrid approach and implement a layered enforcement approach, using both network- and host-based controls within the same environment.

> *"A lot of organizations start by creating basic segmentation and boundaries using network elements, then they bring in agents to tighten control within those segments" (P10, 2025)*

Typically, this starts with macro-segmentation using VLANs, subnets, or firewalls to create broad containment zones, and then adds micro-segmentation using agents for finer control within those zones. The network layer offers foundational containment, while the agent-based layer enforces granular host-level access policies. The interviewees also discussed the advantages and limitation of these methods. The layered model offers enhanced the security by leveraging the strength of network-based isolation and the precision of host-level policy enforcement. It enhances security visibility and control while requiring fewer specialized networking components compared to purely network-based solutions. However, this method can also be demanding in terms of time and resources, increased system load, and elevated operational complexity.

5.3. Legacy System Challenges and Solutions

The implementation of micro-segmentation in networks containing legacy systems is often hindered by a distinct set of technical and operational limitations. These include constraints such as outdated operating systems, lack of visibility, flat network dependencies due to hard-coded settings, and compatibility issues with modern security tools. To highlight which legacy-related challenges are most prevalent in real-world environments, Figure 5.2 shows how frequently each issue was mentioned during the expert interview. This section expands on the challenges shown in the diagram and discusses how they were addressed in practice by professionals interviewed during the research.

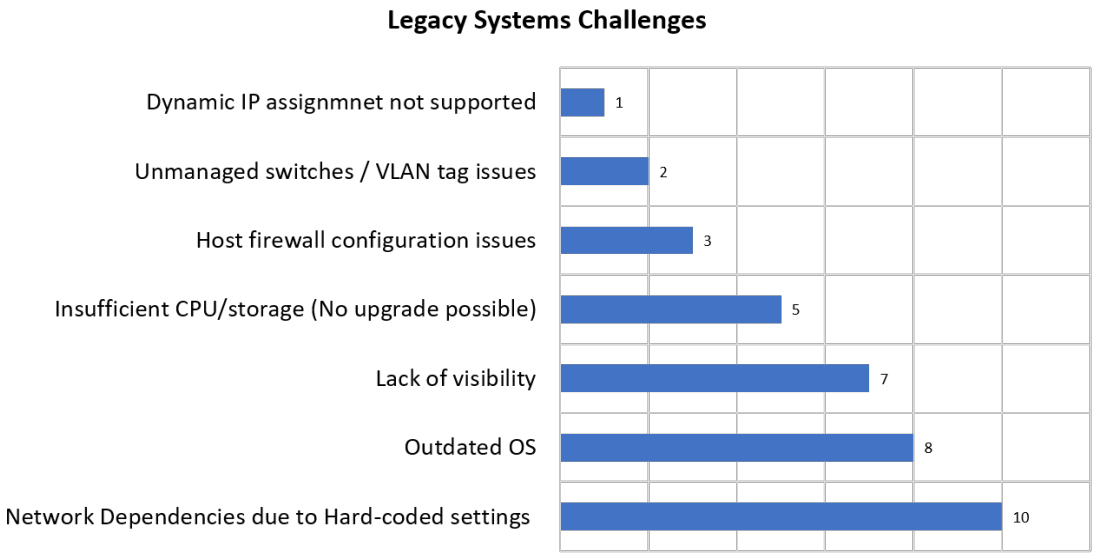

Figure 5.2: Frequency of legacy system challenges mentioned by interviewees

Network Dependencies due to Hard-Coded Settings

One of the most pervasive and disruptive challenges in implementing micro-segmentation in legacy environments is the presence of hard-coded settings, with hard-coded IP addresses being the most frequently cited example among interviewees. Legacy applications are often designed with static configurations that bind them to specific IPs, making them fragile in the face of network changes and creates network dependencies. This becomes particularly problematic during segmentation, where systems are reassigned to new network zones or VLANs and are given new addressing schemes. In the case of P1, this limitation was described as a direct barrier to segmentation. Many systems in their environment used hardcoded IPs, and due to the absence of source code or the unavailability of original developers, updates were impossible. P1 mentioned:

> *"If we move the printer to a new network segment, it will receive a new IP address. However, the application that communicates with the printer was specially developed, and the old IP address is hard-coded into its source code. As a result, once the printer is moved, the application can no longer reach it, because it still tries to use the old IP address." (P1, 2025)*

Without the ability to update IP addresses and settings, changes due to segmentation can break interdependencies and lead to critical service outages. According to P2, this constraint force the organization to cluster dependent systems into the same security zone, essentially limiting the granularity of segmentation.

> *"When you talk about micro-segmentation, if one server communicates with another using an IP address and that IP or the server's location changes, then the communication can break, especially if the settings are hard-coded. To avoid this, you need to understand system dependencies and group those dependent servers within the same security zone. In practice, we often work backward: identifying which systems communicate closely and clustering them together to ensure connectivity is preserved." (P2, 2025)*

This issue is not only relevant to network-based micro-segmentation, but also agent-based micro-segmentation. For agent-based segmentation, although enforcement occurs at the host level, communication rules are often still defined using IP-based logic. P2 reported this obstacles and the workaround involved detailed asset discovery to identify tightly coupled systems and keep them within the same firewall ruleset, accepting reduced segmentation as a trade-off for stability. P2 explained:

> *"You should be aware that agent-based also uses IP communication. It's not like agent to agent talks to each other directly via some proprietary mechanism. Ultimately, traffic still goes through the network, and enforcement happens at the IP level... So even with agents, you are defining rules based on IP, port, protocol, etc." (P2, 2025)*

The presence of hardcoded IP addresses forces organizations to compromise between segmentation depth and operational continuity. The most sustainable mitigation involves identifying such systems early during asset discovery and applying segmentation at the perimeter using VLANs or firewalls. This challenge makes agentless network segmentation with strict firewalls or VLAN isolation the default strategy.

Outdated Operating Systems (OS)

Another recurring challenge to implementing micro-segmentation in legacy-heavy environments is the prevalence of outdated operating systems, which limits the applicability of agent-based and hybrid micro-segmentation strategies. Modern agent-based solutions rely on deploying host-level software that can monitor and enforce fine-grained access control. However, many legacy systems, especially those running unsupported versions of Windows, Unix, or proprietary OT platforms, cannot accommodate such agents due to OS-level compatibility issues. P6 also discussed this obstacle saying:

> *"Let's say you're running Windows 2003, which is like a 20-year-old operating system... some of the modern systems may not be able to support that. So... we might have to put them in a separate VLAN... We may not be able to install agent-based segmentation tools on those servers." (P6, 2025)*

This challenge makes it impossible to install the security components needed to enforce segmentation policies at the host level. To overcome these challenges, organizations typically resort to network-based isolation strategies.

> *"If we cannot deploy the enforcement component on the endpoint because of some legacy operating system or because of any restriction, then in that case we on network-based isolation." (P5, 2025)*

Outdated operating systems force organizations to prioritize containment over precision. While this limits the granularity of control, it also underscores the role of network-level segmentation as a practical and often necessary strategy in environments where modernization is not immediately feasible.

Lack of Visibility

A foundational barrier to implementing micro-segmentation in legacy environments is the lack of visibility into existing assets, communication flows, and system dependencies. In environments with legacy infrastructure, visibility is often undermined by missing documentation, decommissioned asset tracking tools, or the absence of institutional knowledge. As noted in P4 and P9, many organizations do not have an up-to-date inventory of their systems. Teams are unaware of which systems are running, where they're located, who owns them, or how they interconnect. P9 highlighted this gap by saying:

> *"In many cases you don't even have an inventory. You don't know what systems are working, where they are working, who is using them. You just know there is something because there is blinking or producing data." (P9, 2025)*

This issue is particularly problematic in flat networks where everything can technically communicate with everything else, making it almost impossible to draw safe segmentation boundaries without accidentally cutting off critical services. During the discussion about this issue, P4 mentioned:

> *"We are entering environments which are quite flat, everything talks to everything else... and we are afraid to touch anything because we may break something. So first step is even before segmentation, just understanding what is there." (P4, 2025)*

P13 described how the lack of visibility became apparent during the initial steps of segmentation. An attempt to simply separate production and non-production zones revealed thousands of undocumented firewall dependencies. Legacy applications had long-standing but informal data flows that were neither documented nor owned by any identifiable team.

> *"We took a flat network and... did a first iteration by separating the production environment versus the non-production environment... The challenge was significant because we had almost 1000 application individual assets... they have no clue about what configuration the firewall policy should have... the challenge is always the same. It's the lack of understanding, lack of awareness or lack of documentation, a lack of understanding how my IT landscape is composed of and how is it operating." (P13, 2025)*

The lack of visibility affects all segmentation strategies but is especially critical for agent-based and hybrid approaches, where granular policy enforcement depends on precise knowledge of which systems communicate, how they interact, and what normal behaviour looks like. When the existing state of communication is unknown or unpredictable, the risk of enforcing incorrect rules increases significantly. P6 pointed out that tagging servers accurately requires confidence in their role and application context, something that is difficult to do without clear visibility into system function and traffic patterns:

> *"When we talk about tagging a server, the challenge is always: what tag should I give? Is it a DB server? Is it an app server? Is it a middleware? Unless you have that information very clearly identified, you cannot apply tagging properly." (P6, 2025)*

From a network-based segmentation perspective, visibility is equally essential. P15 emphasized that grouping systems into VLANs or firewall zones without knowing their communication requirements can lead to broken dependencies, outages, or overly permissive rules that defeat the purpose of segmentation:

> *"If you don't understand the actual communication pattern between systems, and you just go and put them into separate VLANs or segments, you might break something. Or worse, you overcompensate by allowing too much, and now your segmentation doesn't really mean anything." (P15, 2025)*

These findings highlighted how deeply this challenge can delay or complicate Zero Trust segmentation initiatives. To address this issue, most interviewees stressed the need for a dedicated discovery phase. Tools like passive network monitoring (NetFlow, TAP/SPAN), vulnerability scanners, or CMDBs can be used to build an initial asset and traffic inventory. As P9 and P1 noted, automated tools alone are not enough. Manual verification, cross-functional workshops, and close coordination with application owners are required to translate raw traffic data into meaningful segmentation rules. P9 clarified this by saying:

> *"You can use something like Gigamon or whatever to extract data, but that's just giving you what the system is doing in that moment in time. Then you need to go back to the application owners and validate that the data flows are correct. That's why we do workshops to make sense of the data." (P9, 2025)*

P4 mentioned several methods to establish network visibility in legacy environments, including existing documentation, conducting manual walkthroughs, and using passive monitoring tools such as IDS systems. When asked which method works best, he recommended a hybrid approach, using all of these techniques in combination to compensate for the limitations of any single method and build a more complete picture of the environment:

> *"So probably combination. As not all all asset can be covered by passive monitoring, so I always recommend to start what we have documented already. Even if not much sure if it's next little spreadsheet. Let's let's start with it. Maybe have a manual walkthrough and then have a complementary information from the network scans. Yeah. So let's say the mix hybrid approach." (P4, 2025)*

Active polling, which involves directly querying devices to collect real-time data, can be added to the hybrid approach to improve accuracy. However, P4 emphasized that organizations must be aware that it introduces additional network traffic and may disrupt sensitive or latency-critical systems:

> *"Active polling, of course, is better, but can be dangerous in some networks, especially OT... Active polling may create traffic that is not expected by devices, which may crash or misbehave. So, passive is safer, although it's less precise." (P4, 2025)*

To conclude, the documentation is often outdated and manual methods alone can overlook hidden interdependencies, using these techniques in combination enables a more reliable foundation for safe and effective segmentation planning.

Insufficient CPU and Storage

A major barrier to deploying agent-based or hybrid micro-segmentation in legacy environments is the insufficient hardware capacity of legacy systems specially when upgrading them is not possible. Many legacy servers, appliances, and OT devices were designed with minimal CPU, memory, or storage headroom, making them unable to support the overhead introduced by modern security agents. These agents perform real-time traffic monitoring, enforce host-level rules, and sometimes integrate with identity systems or external controllers. As noted by P2, P5 and P6, legacy systems often cannot accommodate the additional processing demand, particularly when they are already operating near capacity to handle business-critical workloads. In many cases, just installing the agent causes system instability or application timeouts:

> *"We can't just plug in our XP Windows Desktop and we can't expect that Windows machine to respond to the micro segmentation solution... There is a requirement at physical level as well like... the resources... should have at least these many processes or memory or storage, which most legacy systems simply don't meet." (P2, 2025)*

To mitigate this, organizations turn to network-based segmentation techniques that do not rely on host-level enforcement. P6 highlighted this issue and suggested:

> *"some systems may not support agents. For those, we have to go with more traditional segmentation like VLANs, firewall segmentation... because we cannot install the agent. Some of them are so old or resource-constrained that just installing the agent causes the application to crash or timeout." (P6, 2025)*

In short, insufficient CPU and storage is not just a technical limitation, it fundamentally shapes what forms of micro-segmentation are feasible, often forcing organizations to compromise on granularity, enforcement.

Host Firewall Configuration Issues

In agent-based and hybrid segmentation strategies, host firewalls play a central role in enforcing access policies at the system level. However, many legacy devices either lack built-in firewall functionality or have it disabled due to performance concerns, misconfiguration, or lack of administrative access. P10 highlighted this issue and said:

> *"Some of the legacy systems don't even have host firewalls, or they're turned off because they caused performance issues or nobody knew how to configure them properly. In those cases, you really can't rely on the endpoint itself" (P10, 2025)*

In some cases, even when the capability exists, teams may be reluctant to modify settings on those legacy systems due to fear of causing service disruption. To compensate, organizations typically rely on network-based controls such as VLANs and firewalls:

> *"Even if the system technically supports it, no one wants to touch it. It's like, if it's working, don't mess with it [...]. So instead, we just rely on network level segmentation." (P5, 2025)*

While this approach lacks the granularity of host-level enforcement, it still provides containment when applied appropriately.

Unmanaged switches and VLAN-tagging Issues

One of the key challenges highlighted by P15 and P3 is the presence of unmanaged switches in legacy and OT environments. These switches do not support VLAN tagging, making Layer 2 segmentation impossible. As a result, micro-segmentation strategies based on VLANs cannot be fully applied. All systems connected to the unmanaged switch must remain in the same broadcast domain, which limits the ability to isolate them into smaller zones.

> *"In many OT environments, we see unmanaged switches. These are not capable of VLAN tagging, so we can't apply any segmentation there. Everything connected stays in the same Layer 2 segment." (P3, 2025)*

This forces organizations to adopt a coarser segmentation model, reducing the overall granularity of the network-based strategy. P3 explained that the ideal solution is to replace unmanaged switches with managed ones, but this is often resisted due to cost and the operational risk of disrupting critical systems. In many OT environments, organizations have already made significant investments in industrial-grade switches, which are costly due to their need to withstand extreme conditions like high temperatures, dust, and vibrations. As a result, replacing these expensive switches is difficult to accept:

> *"Companies are usually not very happy when someone tells them, 'You need to replace 200 switches to support VLANs.' Then they ask, 'How much would that cost?' and I say, 'It's going to be €800,000,' and they respond, 'No, we're not doing that." (P3, 2025)*

He also noted that many legacy devices, such as PLCs and older Windows systems, do not support VLAN tagging. However, this does not block VLAN-based strategies entirely. The workaround is to assign these devices to native VLANs so that the switch handles tagging on their behalf, allowing them to participate in segmented networks without requiring changes to their configuration.

Dynamic IP Assignment not Supported

In some legacy environments, certain devices do not support dynamic IP assignment via DHCP and must be configured with static IP addresses. As noted by P2, this limitation was encountered during segmentation efforts, particularly with older systems like surveillance cameras or legacy infrastructure equipment:

> *"Some of the legacy devices, they do not support DHCP at all. So they have to be configured with a static IP address and those devices were mainly surveillance cameras or industrial equipment which are there in the data centre for years." (P2, 2025)*

While this does not directly affect the choice of micro-segmentation strategy, it introduces extra operational overhead during network restructuring. Static addressing requires more careful planning and documentation to ensure consistent device identification and policy enforcement, especially when moving systems between zones. An overview of how these challenges influence the applicability of different micro-segmentation strategies is presented in Appendix F, which includes visual summaries linking each challenge to affected strategies, along with a comparison of their respective limitations and advantages.

5.4. Steps Toward Micro-segmentation

Many interviewees described the phases and practices involved in implementing micro-segmentation in detail. One interviewee (P2) provided a particularly comprehensive and structured breakdown, which is used here as a foundation. Other participants contributed complementary insights that enrich and validate this multi-step process. Together, these accounts reveal a clear and **phased** progression from planning to deployment and continuous evaluation, with necessary adaptations for legacy-heavy environments.

5.4.1. Network Separation

Before segmentation can begin, several experts, particularly those working in operational technology (OT) environments, emphasised the importance of an initial network separation phase. This step involves isolating critical systems that are often still embedded within flat enterprise networks, exposing them to unnecessary risk. The goal here is not to define detailed segments, but rather to relocate high-value or vulnerable systems behind firewalls or other boundary controls to establish basic isolation from general-purpose IT infrastructure. Particularly in fragile or safety-critical environments, this gradual, containment-focused approach is critical for avoiding operational risk.

> *"First step is even before we touch network segmentation, we actually start with what we call network separation. So very often when we visit our clients, in almost every case, we can meet critical systems located in the corporate environment. So we are first doing the clean-up, identifying these critical OT systems that should be moved behind a firewall." (P4, 2025)*

In many OT environments, systems like programmable logic controllers (PLCs), human-machine interfaces (HMIs), or vendor-specific control units are deployed on shared networks with minimal protection. Interviewees explained that isolating these systems by moving them into controlled zones or placing them behind dedicated firewalls creates a buffer that reduces exposure and prepares the environment for more structured segmentation. This step also serves as an opportunity to begin assessing what is on the network, identifying undocumented assets, and observing system behaviours without making disruptive changes. At this stage, the asset inventory is usually not yet detailed; rather, it provides a broad overview of what exists and where immediate risks may lie.

5.4.2. Network Segmentation

Once critical systems have been isolated, organizations typically move into the network segmentation phase, where logical groupings are defined and enforced. This is where the architecture of segmentation begins to take shape, usually through the deployment of VLANs at Layer 2 and the implementation of firewalls or access control lists (ACLs) at Layer 3. Interviewees described this as the foundational layer of any micro-segmentation effort, a macro-structuring of the network into zones based on functional roles, risk profiles, operational domains or system/ device types. As P2 explained:

> *"We are dividing all of our web application servers into one segment or in one zone. Then we are creating another zone for application servers and then similarly for the databases. So now we have three zones created in our environment." (P2, 2025)*

These large-scale groupings serve to reduce unnecessary communication between system classes and prepare the environment for more granular controls. P3 added a complementary perspective, emphasizing segmentation based on access requirements and operational needs:

> *"You might have different network zones, such as management, internal, and public-facing segments. And then you can have something we call dedicated network conduct. So that is, for example, if a vendor needs to patch a system or update a system, then you use a dedicated network path that is very restricted and isolated." (P3, 2025)*

A clear and structured breakdown of network segmentation was provided by P2. The process begins with a high-level design where assets are categorized by device types, such as laptops, printers, surveillance systems, or production machines and grouped into conceptual zones. This grouping is primarily based on type, but the rationale is tied to the functionality and communication behaviour of the devices. As P2 stated:

> *"Combining similar device types in a particular zone. Let's say out of those eight device categories, there are chunks which you can made like OK these two device categories can fall in one zone.*

> *Remaining three can fall in another zone and then remaining remaining three can fall in another zone. So we have 3 zones defined." (P2, 2025)*

Many interviewees such as P2, P4, P5, P6, P8, P9 and P12, emphasized that a comprehensive asset inventory is required to support the design and enforcement of effective segmentation policies. P2 highlighted the importance of a detailed asset inventory for both high-level and low-level design, stating:

> *"First of all, we need to have a comprehensive inventory with us... I should have the exact number of laptops, notebooks, printers, scanners, then my CCTV cameras,..." (P2, 2025)*

Other interviewees also reinforced the importance of accurate asset inventories and traffic flow analysis at this stage. Without understanding what devices exist, how they communicate, and what services they depend on, segmentation efforts can easily disrupt critical operations. Some professionals noted that even at this level, legacy systems can pose challenges such as hardcoded IP addresses, protocol limitations, or undocumented dependencies which must be accounted for in the design. Such constraints make a comprehensive inventory critical, allowing segmentation to be planned with full awareness of potential legacy-related risks. As emphasized by P5:

> *"When you have legacy systems, a clear asset inventory is even more important. Some of these devices don't fully support segmentation, so you need to know exactly what's there and how it behaves before you start." (P5, 2025)*

The subsequent low-level design mentioned by P2, introduces technical specifics such as VLAN allocation, IP subnet sizing, firewall interface assignments, and access policies. To guide the high level firewall setup, heatmaps are created to map communication flows between zones. Within each zone, traffic is generally unrestricted, but inter-zone communication is tightly controlled to reduce the blast radius.

Once the low-level design is complete, the focus shifts to preparation and implementation. This involves configuring switches and firewalls, deploying VLANs, updating routing tables, and ensuring that DNS, DHCP, and other infrastructure services are aligned with the new segmentation model. Change management becomes critical at this stage to avoid disruptions during cutover. P2 described this as the transition to the migration phase, mentioning that all communication flows and access paths must be validated before devices are moved into their new zones. Heatmaps and traffic baselines help ensure that critical services remain accessible during and after migration.

The final steps include cutover execution, where devices are gradually migrated into their assigned segments, followed by post-migration testing and validation. This includes checking connectivity, verifying that security policies are enforced correctly, and ensuring that no essential services are inadvertently blocked. Interviewees emphasized the need for rollback options in case issues arise, particularly in environments with sensitive or legacy systems. A summary of these steps are illustrated in Figure 5.3.

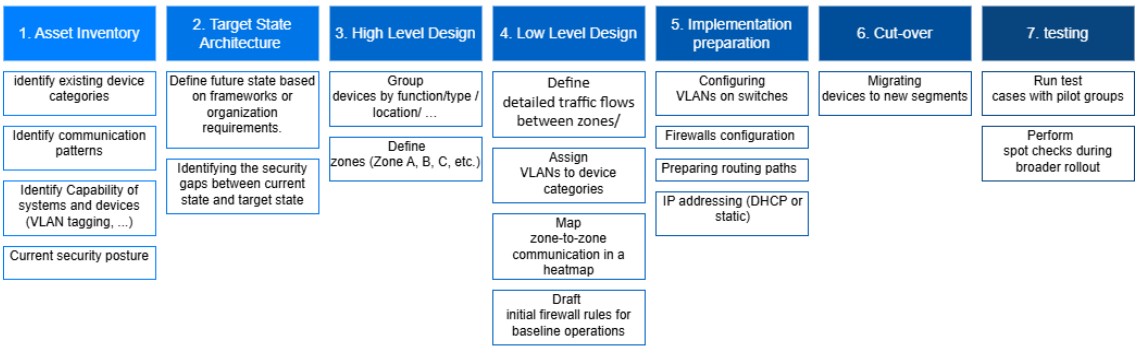

Figure 5.3: Network Segmentation Steps described by Interviewees

5.4.3. Micro-Segmentation

With the network now divided into zones, the next phase focuses on controlling communication within those zones, what interviewees described as true micro-segmentation. This stage introduces finer-grained access controls, either through host-based agents or by subdividing existing zones into smaller segments using the organization's network infrastructure. The choice between these approaches depends on factors such as device compatibility, required control granularity, available infrastructure, business requirements, and more. One of the interviewees referred to one of these influencing factors, saying:

> *"For some of the legacy systems, deploying agents wasn't feasible, so we had to rely on the network controls we already had in place to enforce segmentation." (P5, 2025)*

In cases where network-based micro-segmentation is selected, the same set of steps described earlier for network segmentation can be followed (see Figure 5.3). However, the focus shifts from broad segmentation across the environment to creating additional isolation zones within already segmented areas, allowing for more controlled lateral movement and layered defence.

When the chosen approach is *agent-based micro-segmentation*, the overall structure remains similar, but the process introduces additional steps and greater detail at the host level. The strategy selection itself typically occurs after the asset inventory phase, once systems are classified as agent-capable or not. From this point, organizations proceed with agent-specific phases. P2 describes this path in detail, beginning with early preparation activities such as validating endpoint readiness, reviewing existing policies, and identifying which systems can support agents. The interviewee mentioned:

> *"You can't just drop in micro-segmentation and expect it to work, you have to first assess which systems can support agents, what the fallback options are, and make sure your policies don't break essential communication paths." (P2, 2025)*

Following the inventory, the organization defines its target state architecture, choosing whether to adopt an agent-based, network-based, or hybrid segmentation strategy. This decision is informed by the earlier classification of systems and the intended enforcement goals, such as implementing a deny-by-default posture between application zones or functional roles.

The next phase, high-level design, involves mapping the segmentation intent. Using available documentation such as application diagrams and known traffic dependencies, interviewees described how teams outline conceptual grouping logic and preliminary trust boundaries. These assumptions guide the upcoming observation process but remain subject to validation.

In the low-level design phase, agents are deployed on eligible systems in observation mode. Over a defined period, they passively capture real-world traffic flows, recording port usage, frequency, and communication dependencies between systems. This telemetry is analysed to identify actual behaviour and compare it to existing documentation. P6 described this saying:

> *"So typically what we do is we install the agents in all the servers that we want to protect. We let it run for sometime like 2-3 weeks or four weeks or something like that, so that we observe the traffic... They'll have then application level architecture. They'll also have network architecture, so we'll have to kind of reconcile everything between, let's say, what the application says, what the network architecture says versus what the actual traffic is... We will try to reconcile and then agree what should be allowed and what should not be allowed." (P6, 2025)*

Based on the findings, systems are then grouped using tags that reflect functional or application-level relationships. For example, servers involved in supporting a single application such as a web server, app server, and database, might share a common tag. From here, policies are drafted to allow traffic within these groups and block unnecessary traffic between them. P2 emphasized the need to fine-tune policies to accommodate valid exceptions uncovered during observation:

> *"There are always some edge cases. Like, maybe a middle-tier server needs to talk to a backend DB directly, even though normally that goes through an app layer. Those kinds of exceptions based on how the application is actually built or specific business needs, we account for them during the fine-tuning." (P2, 2025)*

Once policy logic is established, the implementation preparation phase begins. Tags are assigned in the agent management platform, and the drafted rules are applied in monitor-only mode to simulate enforcement. This enables teams to validate expected system behaviour and refine policies before going live. P6 highlighted this approach mentioning:

> *"We'll create the labels, define the rules, but we won't enforce right away. We just let it run in simulation, so we can tweak and avoid outages." (P6, 2025)*

During the cutover phase, policies transition from monitor to enforce mode in controlled phases. Systems are gradually brought under live segmentation controls, with teams monitoring closely for unintended service disruptions or blocked flows. The testing phase includes both structured and exploratory efforts to confirm that essential communications are preserved. Policies are adjusted as needed based on real-world outcomes and user feedback. Importantly, the work does not end at enforcement. Interviewees emphasized that micro-segmentation requires ongoing evaluation and refinement. Organizations rely on log data, monitoring tools, and support team input to assess performance, detect issues such as latency or blocked processes, and iteratively improve policy configurations. A detailed summary of the steps involved in the implementation of agent-based micro-segmentation is illustrated in Figure 5.4.

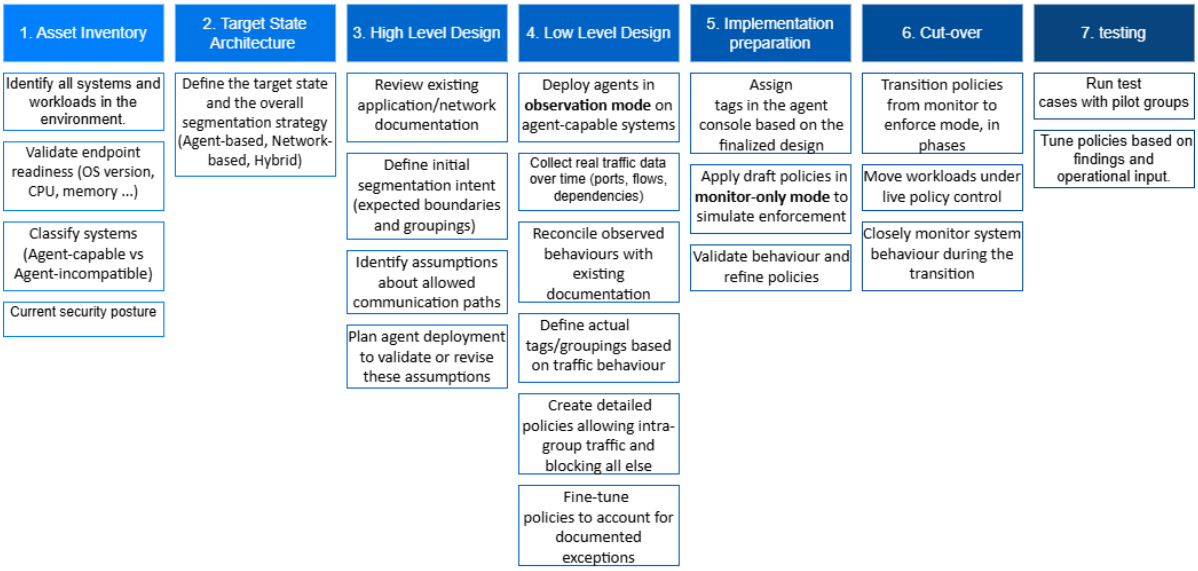

Figure 5.4: Agent-based Micro-Segmentation Steps described by Interviewees

5.5. Evaluation Criteria and Methods

Evaluation of micro-segmentation strategies in legacy-heavy environments focuses on how well the implementation meets practical goals without disrupting business operations. The criteria most frequently mentioned by interviewees provide a clear picture of what matters most in real-world deployments. The insights from interviews are summarized in the frequency chart below, which highlights the most commonly cited evaluation criteria. In the remainder of this section, each criterion is discussed in order of frequency, along with the reasoning behind its importance and the typical methods used to evaluate it in practice.

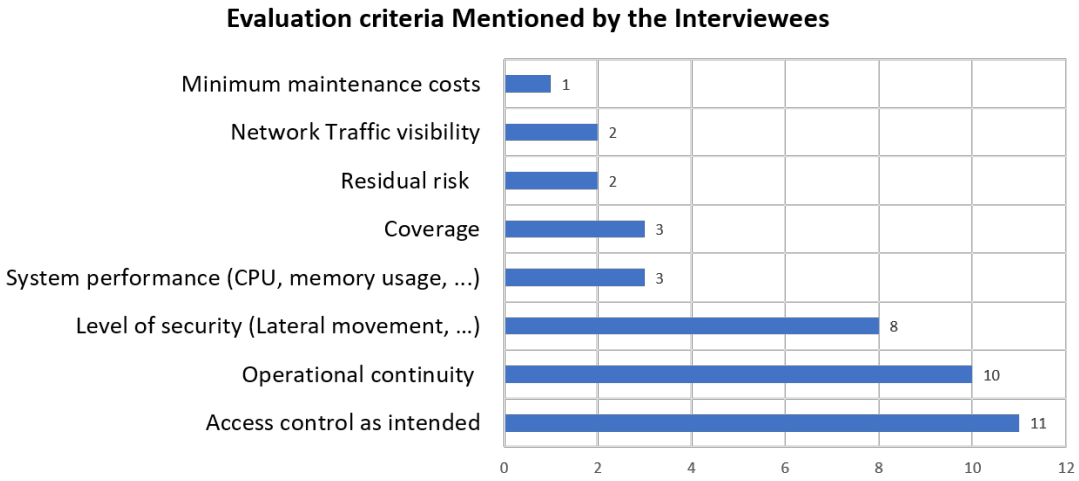

Figure 5.5: Evaluation Criteria and Their Mention Frequency During the Interviews

Access Control as Intended

The most important criterion, mentioned by almost all participants, is Access Control as Intended. This means ensuring that only the right systems can communicate with each other and that unauthorized access is fully blocked. P9 mentioned this criteria and said:

> *"I mean, the way we validated is by checking that from that subnet I cannot reach other subnet. That's it. So if this machine is supposed to talk only with that machine, I should not be able to ping other ones or connect to other ports, that's the only way to be sure it works as expected." (P11, 2025)*

To evaluate this, organizations typically use both predefined test cases and exploratory testing. Predefined tests check whether necessary connections (like between a database and its front-end) still function. Exploratory tests, on the other hand, aim to simulate unauthorized access attempts to see if segmentation rules are effective:

> *"We also test if the database can talk to the application server, or if this backend service can still be reached. But we also try the opposite, try to access something that should not be reachable, like from a user zone to a server zone. That way we know both what works and what is blocked" (P9, 2025)*

Logs from applications and firewall consoles, along with monitoring tools, help confirm whether segmentation is working correctly.

Operational Continuity

Next is Operational Continuity, which was highlighted by nearly as many participants. It refers to making sure that segmentation does not interrupt day-to-day business activities. This is checked by performing functional tests after deployment, making sure that users can still access services, applications respond normally, and performance is not degraded. P2 gave some example of functional testing:

> *"After implementing the segmentation rules, we tested the connectivity to ensure all business-critical services are still accessible. We checked if users could log in, applications were reachable, and there was no performance degradation." (P12, 2025)*

Some organizations also monitor helpdesk tickets or system alerts to catch any disruptions early.

Level of Security

The third most cited factor is the Level of Security, particularly the reduction of lateral movement within the network. The goal is to stop attackers from moving between systems once they get in. This is usually assessed using penetration testing and by analysing traffic flows after segmentation.

> *"Another approach is penetration test. After network segmentation, I always recommend to perform the network penetration testing to confirm if our segmentation is efficient and this properly implemented and supports security." (P4, 2025)*

If unauthorized movements are blocked and the attack surface is reduced, the implementation is considered successful from a security standpoint.

System Performance

System Performance, especially resource usage like CPU and memory, is also an important consideration. Adding firewalls or agents can put extra load on systems. Organizations measure performance before and after segmentation to ensure the added controls don't slow things down or require hardware upgrades. P2 highlighted this point and said:

> *"The only point is when you're doing something at the host level—putting an agent on a PLC or Windows server, you have to monitor CPU, memory, and see if it's stable. Some of these systems are already running near capacity, and any additional load can cause issues." (P2, 2025)*

While no specific tools were named, they referred to the use of monitoring and analytical tools to track system resource usage, such as CPU load, memory consumption, and latency, particularly in legacy environments where hardware capacity is limited.

Coverage

Coverage refers to how many of the intended systems were actually segmented. A segmentation effort that leaves out important systems can create gaps as mentioned by P14:

> *"The one thing I always want to understand is: what is included? Did we include everything that should be segmented, or are there blind spots? That's how you assess whether your coverage is complete" (P14, 2025)*

To check coverage, organizations use asset inventories, network discovery tools, and traffic mapping to verify that all relevant devices and applications are included in the segmentation design.

Residual Risk

Residual Risk is about what's left unprotected, either because of technical limitations or deliberate exceptions. Even if perfect segmentation isn't possible, isolating high-risk systems and documenting any exceptions helps reduce exposure. This is typically reviewed through audits or manual inspections of the network and segmentation policies. P4 mentioned:

> *"If we can address most of the risks identified at the beginning, then what remains after segmentation, that's what we consider residual risk." (P4, 2025)*

Network Traffic Visibility

Improved network visibility is also considered an important evaluation criterion. After implementing micro-segmentation, organizations expect to gain clearer insights into how systems communicate across the network. In legacy-heavy environments where undocumented dependencies and unknown traffic flows are common, segmentation should help make these interactions more transparent. P14 mentiond this criteria and adde:

> *"After implementing segmentation, we had a much better view on what was actually communicating. Some flows we thought were inactive turned out to be still in use, and we discovered a few things nobody had documented, and that is the kind of visibility we were aiming for." (P14, 2025)*

Interviewees explained that enhanced visibility enables better detection of unauthorized connections, easier troubleshooting, and more informed policy adjustments.

Minimum Maintenance Costs

Lastly, Minimum Maintenance Costs and efforts were mentioned less often, but they are still relevant, especially over the long term. If segmentation policies are too complex or require frequent updates, they become hard to manage. P14 explained:

> *"We realized some policies were too detailed, and every little application change forced us to revisit the rules. That's not sustainable long-term." (P14, 2025)*

Evaluation in this case is more informal, often based on feedback from administrators and how much manual effort is needed to keep the setup running smoothly.
These criteria define an approach to evaluating micro-segmentation. A strong implementation not only enforces access restrictions and improves security but also keeps systems running smoothly and remains manageable over time.

6

Framework Development

6.1. Framework Development Requirements

The requirements of the proposed framework were derived from the insights that emerged across the 16 expert interviews. This section explains how those insights were translated into actionable design requirements for the framework.

The interviews began with an overview of current network-based micro-segmentation practices, particularly in environments dominated by legacy systems. As discussed by interviewees most of the interviewees, VLANs combined with firewalls represent a widely used strategy in such environments due to their compatibility with existing infrastructure. Several participants emphasized that VLAN tagging must be supported by switches, and that segmentation can only be safely applied if network dependencies due to hard-coded settings are understood and accounted for. These findings lead to the requirement to verify VLAN tagging capability and dependency constraints before recommending VLAN-based segmentation.

In cases where VLAN tagging is not supported or network dependencies are present in the network that prevent system separation, interviewees like P2 and P4 described fallback approaches that involve grouping dependent systems into a single, isolated zone. Although this reduces segmentation granularity, it preserves critical communication paths and avoids service disruption. These insights informed the requirement that the framework should recommend isolating connected or dependent systems into a common zone when the necessary technical criteria for segmentation are not met.

Participants also described a spectrum of VLAN-based configurations, ranging from coarse zones to highly granular VLAN-per-device architectures. When discussing Private VLANs (PVLANs), P5 explained how these enable intra-subnet isolation in scenarios where devices cannot be readdressed due to application or configuration constraints. As a fine-grained network-based micro-segmentation strategy, PVLANs allow tighter control without modifying endpoint settings. Other configurations mentioned during the interviews included VLAN-per-zone segmentation, placing each device in its own VLAN with centralized firewall enforcement, and grouping systems into VLANs based on function, location, or communication requirements. The framework should therefore include all of these configurations as final strategy options, reflecting the range of approaches discussed. Moreover, the framework must clearly distinguish between high- and low-granularity strategies.

The topic of router-centric environments and VRF-based segmentation was raised by P4, who emphasized that VRF is effective only when the infrastructure is built around routers and supports flexible routing configurations. He noted that issues such as hard-coded static routes can make VRF implementation infeasible. This supports the idea that routers must support VRF and allow routing flexibility; otherwise, the framework should propose isolating dependent systems.

Interviewees consistently discussed the importance of agent-based micro-segmentation as a strategy that offers the highest granularity and visibility. According to P5, P6, and P9, this method is preferred when systems can support agents, meaning that the framework must begin by evaluating the feasibility of agent-based segmentation, and only proceed to alternative strategies if full agent compatibility cannot be achieved across the environment. Some interviewees such as P2 and P6 stressed that many legacy systems cannot support agents due to OS limitations, CPU/memory constraints, or host firewall configuration issues. Therefore, agent-based segmentation should only be proposed as a stand-alone solution if full compatibility exists across the environment.

When full agent deployment is infeasible but partial support exists, P5 and P2 described using hybrid approaches, combining network and host-based controls. However, they warned that hybrid enforcement should only be used when agents can be deployed to at least 40-50% of systems, as anything less results in disproportionate complexity. This means that hybrid strategies should only be proposed when agent support is available for the majority of systems.

Another pattern mentioned by several interviewees, such as P2, P5, and P10, is the adoption of layered enforcement models. In these models, VLAN-based macro-segmentation is combined with host-level controls for critical systems. Given these insights, the framework should allow the option to apply layered approaches when agents are supported. However, P2 and P3 also noted that the layered approach only works if the network infrastructure supports it. This led to an additional condition which is the layered segmentation should only be considered if both agent and network prerequisites are met.

Finally, throughout the interviews, participants emphasized the importance of understanding the trade-offs, constraints, and operational overhead associated with each strategy. P12, P13, and others highlighted how fine-grained segmentation can become unsustainable in resource-limited environments, and that any strategy must be balanced against organizational capacity. This formed the basis of the requirement to include the advantages and limitations of all end strategies in the framework to enable informed and context-sensitive decision-making.

All these insights have been translated into 15 requirements for the development of the framework, as summarized in Table 6.1.

Table 6.1: Development Requirements for Micro-Segmentation Strategy Selection Framework

Requirements	
1	The framework must confirm if switches support VLAN tagging and if there are network dependencies before suggesting VLAN-based segmentation.
2	If VLAN tagging is not supported and network dependencies are present, the framework should propose isolating connected/dependent systems into a separate zone.
3	If VLAN-based segmentation is supported, the framework should allow for choosing granular and less granular approaches.
4	All possible configurations mentioned during the interviews should be included as end strategies.
5	The framework must provide network-based micro-segmentation strategies based on network architecture (router-centric vs. switch-centric).
6	The framework must check the compatibility of routers with VRF (no hard-coded routes, etc.). If not compatible, isolation of connected systems must be suggested as the final solution.
7	The framework must begin by evaluating the feasibility of agent-based segmentation and only proceed to alternative strategies if full agent compatibility cannot be achieved across the environment.
8	The framework should only propose agent-based micro-segmentation if agent support is available for all systems.
9	If systems can't support agents, the framework should fall back to network-based segmentation.

10	The framework should only propose hybrid enforcement if agent support is available for a majority of systems (approximately 50%). If not, a network-based approach should be recommended.
11	If agents are supported, the framework must allow the option to implement a layered approach combining network- and host-based controls.
12	If a layered approach is not preferred, the framework should recommend agent-based strategy as a stand-alone solutions.
13	In the case of a layered approach, the framework must check if VLAN tagging is not supported by the switches and if there are network dependencies due to hard-coded settings. If so, only the agent-based approach should be suggested as the final strategy.
14	If VLAN tagging is supported and no network dependencies exist, the framework should recommend creating VLAN subnets based on communication needs and deploying host-level agents within those subnets.
15	The framework must include the advantages and limitations of all end strategies (mentioned during the interviews) to allow for informed decision-making.

6.2. Framework for Micros-egmentation Strategy Selection

This section explains how the micro-segmentation framework was developed based on the findings from expert interviews. The framework is intended to be applied after the asset inventory phase (Figure 5.4), once organizations have determined which systems are agent-compatible and which are not. By this stage, the environment's visibility, legacy system constraints, and technical limitations should be sufficiently understood to inform a strategic choice. The Framework is shown in Figure 6.1.

During the interviews, experts shared practical insights into the segmentation strategies they had implemented. These included agent-based, network-based, and hybrid approaches. Each method had its own benefits and drawbacks depending on the organization's infrastructure, operational model, and the proportion of legacy systems. Interviewees were also asked about the specific challenges encountered during implementation. The most common issues included outdated operating systems, flat networks due to hard-coded settings, unmanaged switches, resource constraints (CPU, memory), and host-based firewall configuration issues. These challenges strongly influenced the design of the framework's logic and the placement of its decision points. To reflect these operational realities, the framework was explicitly designed to account for such practical and technical constraints, fulfilling Requirement 15 that were discussed in previous section.

As shown in Figure 6.1, the first decision point in the framework asks whether there are agent-incapable systems in the network, or whether resource constraints make host-level enforcement infeasible. This question was chosen as the initial branch not only because nearly all interviewees emphasized the importance of agent compatibility, but also because agent-based micro-segmentation offers the most granular, scalable, and architecture-independent form of control. It is typically the preferred approach when technically feasible, and evaluating its viability early avoids unnecessary complexity or compromise in the segmentation design. This reflects Requirement 7, which states that the framework must begin by evaluating the feasibility of agent-based segmentation before considering alternatives.

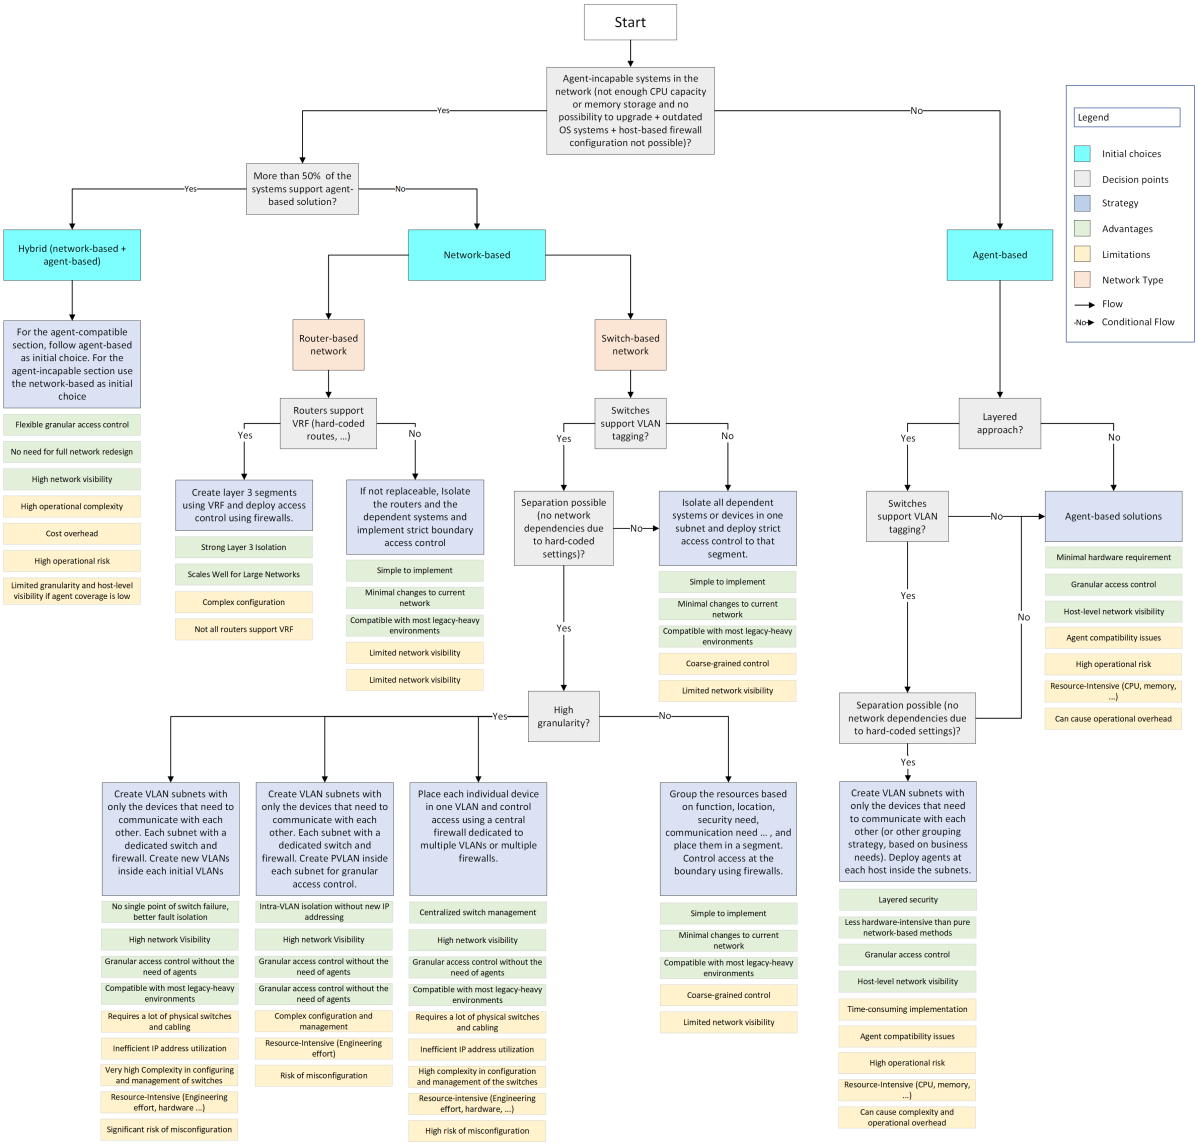

Figure 6.1: Micro-segmentation Strategy Selection Framework for Legacy-Heavy Environments

Agent-based solutions offer fine-grained control and visibility, but they are only viable when the underlying systems can support the necessary software. Multiple participants noted that many legacy systems lacked the hardware resources or modern OS versions required to run agents, and that attempting to do so often caused instability or failure. Therefore, distinguishing between agent-capable and agent-incapable systems is a foundational step that determines the broader direction of the strategy. As such, the framework complies with Requirement 8 by recommending agent-based micro-segmentation only when full compatibility exists across all systems.

If the environment includes a significant number of agent-incapable systems or systems that cannot be upgraded, the framework proceeds down the network-based or hybrid path. This approach aligns with Requirement 9, which calls for a fallback to network-based segmentation if agents are not feasible. If not, it recommends an agent-based strategy. However, rather than using a binary split, the next decision point refines the logic further by asking whether more than 50% of the systems are agent-compatible. This threshold was selected based on the insights shared by participants, who noted that hybrid strategies only become manageable when a majority of systems can be brought under agent-based enforcement. If fewer than 40–50% of systems support agents, the cost and complexity of maintaining a hybrid environment often outweigh the benefits. Therefore, when the majority of systems are agent-compatible but not all, the framework recommends a hybrid (network-based + agent-based) approach. For environments where agent coverage is

low, the framework recommends a network-based approach to avoid operational inefficiencies. Therefore, the framework supports Requirement 10 by proposing hybrid enforcement only when agent support is available for a majority of systems. Otherwise, it defaults to a network-based approach.

Once a main strategy is selected (agent-based, network-based, or hybrid), the framework leads users through a series of decision points that determine the most suitable implementation approach. For the network-based path, the framework distinguishes between router-based and switch-based architectures. This happens in accordance with Requirement 5 which states that the framework should explicitly distinguishes between router-centric and switch-centric architectures to ensure relevant strategy selection. This branching was informed by P4's observation that some OT environments are router-centric and some are switch-centric. The router/switch distinction is necessary because it directly affects which segmentation methods are feasible.

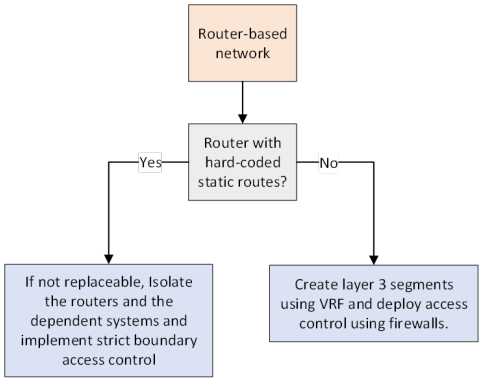

Figure 6.2: Micro-segmentation strategy selection framework:Router-centric network sub-strategies

For router-based networks, the framework introduces Virtual Routing and Forwarding (VRF) as a preferred option where supported. Interviewees acknowledged that VRF provides strong isolation but also mentioned that not all routers support it and configuration is complex. If VRF is not feasible, the framework defaults to isolating legacy devices using perimeter firewalls and static controls. This logic directly supports Requirement 6, which states that the framework must check router compatibility with VRF and, if not feasible, recommend isolation of connected systems as the fallback strategy.

For switch-based networks, the framework checks whether VLAN tagging is supported, since unmanaged switches or legacy devices without tagging support can restrict segmentation. If VLAN tagging is not possible, typically due to the unmanaged switches, the framework recommends grouping all dependent systems or devices into a single subnet and applying strict access control at the boundary. This step partially fulfils Requirement 2.

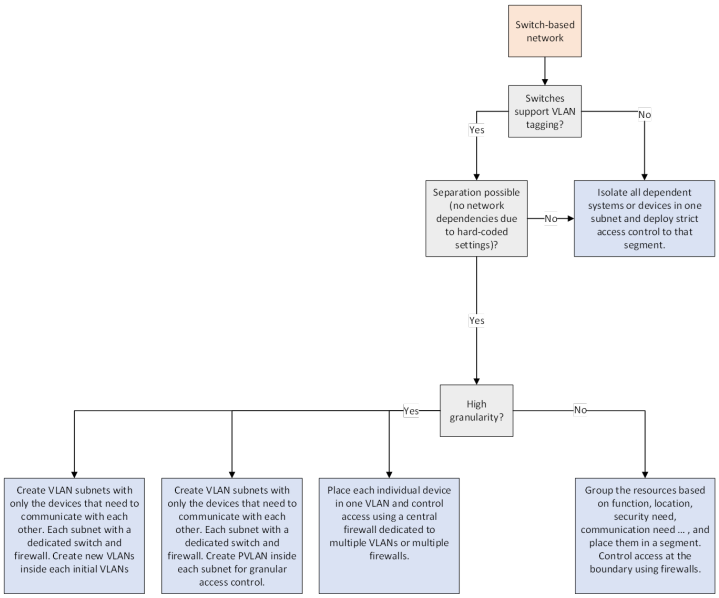

Figure 6.3: Micro-segmentation strategy selection framework:Switch-centric network sub-strategies

If VLAN tagging is possible, the framework introduces an additional decision point: whether separation is feasible meaning there is no network dependencies due to hard-coded configurations such as hard-coded IP addresses. If separation is not feasible, the framework again suggests grouping the dependent systems into one subnet and enforcing access control at that segment's boundary. With this step, Requirement 2 is fully addressed. If separation is feasible, the framework allows organizations to choose between high- and low-granular segmentation, fulfilling Requirement 3. It also incorporates all configurations mentioned during the interviews, thereby satisfying Requirement 4. The strategies include:

- Creating nested VLANs (PVLANs) for finer isolation within broader segments.
- Assigning each device to a separate VLAN and managing access via centralized or dedicated firewalls.
- Grouping devices based on business or operational function and placing them in segmented zones with firewall-based control.

On the agent-based side of the framework, a separate decision branch checks whether a layered approach is preferred, fulfilling Requirement 11. This reflects practices described by participants who combined macro-segmentation at the network level with micro-segmentation enforced by agents. This layered strategy provides foundational control through VLANs or firewalls, while applying finer access restrictions within those zones through host-based policies.

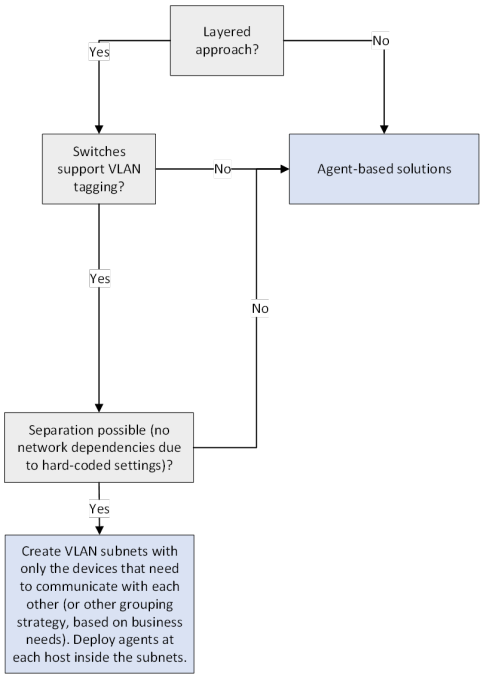

Figure 6.4: Micro-segmentation strategy selection framework:agent-based sub-strategies

If a layered approach is not preferred, the framework defaults to recommending pure agent-based solutions, which enforce policies directly on each host. This supports Requirement 12 by offering stand-alone agent-based enforcement when layering is not selected. If a layered approach is chosen, the framework further evaluates whether the supporting switches allow VLAN tagging . In case tagging is unsupported, granular segmentation at the network level becomes unfeasible, and the framework reverts to host-level enforcement without VLAN-based separation. When VLAN tagging is supported, a final decision point checks whether segmentation is possible without disrupting system functionality, particularly in environments with network dependencies due to hard-coded configurations. If such constraints is present, the framework advises implementing agent-based solution without macro-segmenting the network. This logic fulfils Requirement 13, which mandates fallback to agent-only enforcement in layered models when VLAN tagging or separation is not viable.

In case separation is feasible, the framework recommends creating VLAN subnets that group only those devices that need to communicate (based on business or operational needs), and deploying agents at the host level within those subnets, satisfying Requirement 14. This approach offers the benefits of layered security and granular control, but also introduces increased configuration complexity, resource consumption (CPU, memory), and operational overhead.

To support informed decision-making, the framework attaches the advantages and limitations to each final strategy option. This presentation allows users to weigh trade-offs such as complexity, compatibility, and resource requirements before selection, thereby fulfilling Requirement 15.

The framework is structured to reflect both technical feasibility and operational constraints, organizing decision points according to conditions commonly encountered in legacy-heavy environments. Each decision path is informed by patterns identified across expert interviews, with the aim of representing the range of factors that influence segmentation planning. Strategies that involve higher complexity or resource demands are included only where interview data indicated they are realistically applicable. In cases where visibility is limited, hardware is constrained, or system modifications are not feasible, the framework incorporates alternatives that require fewer changes.

6.3. Framework Implication

This framework serves as a structured representation of how micro-segmentation strategies are selected and applied in environments where legacy systems are present. It synthesizes interview-based insights into a decision model that accounts for a variety of technical and operational factors influencing real-world implementation.

Its primary function is to organize available micro-segmentation approaches according to the conditions in which they are typically deployed. These conditions include, but are not limited to, the architectural nature of the network, the capabilities of deployed components, the compatibility of systems with enforcement mechanisms, and the level of segmentation granularity that can be operationally sustained. By mapping these factors into a logical sequence of decisions, the framework makes visible how strategy selection tends to follow from environmental constraints rather than abstract ideals.

The framework does not prescribe a universal method or prioritize a particular approach. Instead, it reflects how different strategies are adopted depending on what is technically and operationally feasible. This includes recognizing when host-level enforcement is achievable, when segmentation must rely on network infrastructure, or when a combination of both becomes necessary. It also accounts for different network types, such as switch-based and router-based environments which influence the segmentation options that are realistically available. The framework enables practitioners to evaluate trade-offs between granularity, control, complexity, and resource demands. For example, some strategies may offer stronger isolation but introduce significant configuration overhead or scalability issues.

Finally, the framework offers a structured way to assess and compare the advantages and limitations of different segmentation strategies. It supports informed decision-making, facilitates coordination across technical and operational teams, and helps make segmentation choices more transparent to non-technical stakeholders. Although based on current practitioner input, its logic is adaptable and can evolve as system capabilities change or new implementation methods become feasible.

While the framework is grounded in practitioner experience, it has certain limitations. It assumes that organizations have completed an initial asset inventory and have sufficient visibility into their network environment to make informed decisions. In contexts where such visibility is lacking, or where system ownership and dependencies are unclear, the framework's decision paths may be harder to apply effectively. Additionally, the framework is based on conditions and technologies common at the time of the interviews; it does not account for future changes in segmentation tools, network architectures, or organizational practices. It is also focused on technical feasibility and does not model broader factors such as budget constraints, staffing limitations, or policy decisions that may influence implementation. As such, the framework should be viewed as a practical guide shaped by observed patterns in legacy-heavy environments, rather than as a universally applicable model.

7

Framework Validation

Validating the proposed framework is essential to ensure its relevance, clarity, and practical applicability in real-world settings. As Hevner et al. (2004) emphasize in design science research, evaluation is a critical step to confirm that the designed artefact addresses the identified problem and meets user needs. In the context of cybersecurity and architectural frameworks, expert validation is a common and accepted approach to assess the quality and usefulness of conceptual models (Peffers et al., 2007).

7.1. Validation Protocol

The validation of the proposed framework was conducted through expert review using a structured validation questionnaire. Validation Participants (VPs) were selected through the same criteria and process used to identify experts during framework development, specifically targeting professionals with expertise in network security and practical experience in implementing Zero Trust Architecture or micro-segmentation strategies. Table 7.1 provides an overview of the participants, highlighting their areas of expertise and years of experience.

Table 7.1: Overview of Validation Participants, Their Expertise, and Years of Experience (Y-of-E)

Identifier	Expertise	Y-of-E
VP1	Security Network operations and design	12
VP2	Security Network Architect	14
VP3	Security Network Architect	19

These experts were not involved in the development of the framework to allow for an independent assessment. Initially, the intention was to conduct in-depth semi-structured interviews with at least five experts. However, due to their busy schedules during the validation period, only three experts were able to participate, providing their feedback through completed questionnaires rather than formal interviews.

Each participant was provided with the final version of the framework and asked to review the material. They were then requested to answer a focused set of questions designed to validate key aspects of the framework. Specifically, participants assessed whether the logic for choosing between agent-based, network-based, or hybrid segmentation strategies was clear, realistic, and practically applicable. Additionally, they evaluated whether the criteria guiding the selection of segmentation strategies were complete and applicable in practice. They were also asked to suggest any improvements or additions that could enhance decision accuracy. Finally, participants identified which types of users or teams would benefit most from the framework and noted any parts of the material they found unclear or difficult to interpret.

All procedures adhered to the ethical principles described earlier (Section 4.1.2). The goal was not to measure quantitative outcomes but to capture practical insights from experienced practitioners that could confirm or refine the framework's design. The full validation interview protocol can be found in Appendix D.

7.2. Results and Discussion

This section presents the results of the expert validation process, organized into three thematic areas: the completeness and clarity of the decision logic, suggestions for improving the framework, and the type of users or teams who would benefit most from its application. The responses are based on the structured questionnaires completed by three validation participants (VP1, VP2, and VP3), as described in Section 7.1.

Completeness and Clarity of the Decision Logic

All three experts agreed that the framework presents a clear and practical structure for deciding on segmentation strategies. VP1 said the first step in the decision process, which checks if systems can support agent-based segmentation, made sense and matched real-world limitations. If systems cannot support endpoint agents, leaving out agent-based approaches is the obvious choice. This first decision point was seen as a good way to start. VP1 also mentioned that the framework focuses more on network-level segmentation, which is reasonable in situations where legacy systems are common and more advanced solutions are not practical. They considered the framework a useful orientation tool, particularly for organizations in the early planning stages or those without established segmentation policies.

VP2 also mentioned that the framework is a strong starting point for understanding different segmentation options. They liked the flowchart and felt it was easy to follow. They believed the framework had enough detail and that it would be supported by clear documentation explaining the different conditions. The criteria used in the decision-making process were seen as complete and useful in real situations. VP2 suggested the framework could be especially helpful after a company has reviewed its current network setup and is getting ready to plan its segmentation.

VP3 had a similar view and said the decision process between agent-based, network-based, and hybrid strategies made sense. It followed the kind of choices that people usually make in real-world situations. The framework was regarded as particularly helpful for organizations operating legacy systems and still in the initial stages of segmentation planning. VP3 stated that the framework can help users understand the different options and trade-offs clearly. They also said the presentation was easy to understand and the logic behind the decisions was clear.

Areas for Improvement

The experts provided a range of suggestions to improve the framework while generally agreeing that it offers useful guidance. VP1 pointed out that the second decision point in the framework, which asks whether agent coverage is high enough to consider a hybrid approach, is more subjective and depends on more than just technical coverage. Factors like available resources, organizational preferences, and how much complexity a team is willing to handle all play a role. Some organizations might move forward with hybrid segmentation even with low agent coverage, while others may avoid it despite having the technical capability. For this reason, the expert appreciated that the framework allows for flexibility and doesn't force a single path.

VP1 also liked that the framework includes some less common but technically possible options, such as assigning each system to its own VLAN. While this shows the full range of choices, they noted that such detailed strategies are rarely used in real environments because they are complex to manage, especially when legacy systems are involved. They suggested that the diagram could be easier to follow for less experienced users if it came with a short explanation of the key terms or an example use case.

VP2 emphasized that in real-world situations, micro-segmentation is rarely achieved using only one method. They believed a hybrid approach, combining agent-based and network-based techniques, is necessary. Agent-based tools offer more detailed protection and real-time monitoring, while agentless solutions provide a broader view of the system's health and structure. Using both approaches together helps improve protection and gives better insight into network and cloud environments.

For improvements, VP2 said that assigning one switch or firewall to each VLAN is not realistic. In practice, it's more common to create several VLANs within a single switch or across multiple devices, and then use firewall zones to control communication between them. To manage communication within a VLAN, techniques like Private VLANs, VLAN Access Control Lists, or switchport protections should be used, depending on what equipment is available. They added that it's rare to find switches that do not support VLAN tagging, but if such a device exists, it should be replaced before starting any segmentation project. Otherwise, the effort would be based on weak and outdated infrastructure.

VP3 also discussed the agent coverage threshold and said that the 50 percent figure could vary depending on the organization's needs, risk level, and resources. They explained that in many environments, a mix of routers and switches influences how segmentation is done, which might lead to network-based hybrid strategies. Although the framework includes VRF-based segmentation, VP3 suggested adding more details to help users understand when it is a suitable option. However, they did not consider it a problem that this strategy was included, since it is not commonly used and can simply be skipped if not relevant.

Intended Users of the Framework

All three experts agreed on who would benefit most from using the framework. VP1 and VP2 both stated that the framework is especially useful for teams responsible for reviewing the current network setup before starting any segmentation work. These are the people who need to understand the technical conditions of the environment and make early decisions about the direction of the segmentation strategy. VP3 gave a similar answer, adding that the framework would also be helpful for network architecture and security teams involved in both the assessment and planning stages. These teams are typically in charge of shaping the overall strategy and making key decisions, so having a clear and structured guide like this framework can support their work and help ensure a more effective segmentation plan.

8

Conclusion and Discussion

8.1. Conclusion

This section brings together the key findings of the study, addressing the main research question and its sub-questions, and reflecting on how organizations can implement effective micro-segmentation strategies within legacy-heavy environments.

8.1.1. Sub-question 1: The Main Strategy and Steps

In response to the question of what common strategies are used to implement micro-segmentation and how network characteristics influence their selection, this study finds that organizations do not follow a single model. Instead, they adopt strategies that align with their infrastructure, system capabilities, and operational limitations. Although micro-segmentation is a key pillar of Zero Trust, its implementation in legacy-heavy environments is often shaped by what is technically and operationally feasible.

The most frequently adopted strategies include network-based segmentation using VLANs and firewalls, Agent-based micro-segmentation solutions; and hybrid models that combine both. Network-based methods are especially common in legacy-heavy environments due to their compatibility with older systems and minimal endpoint requirements. Factors such as unmanaged switches, hard-coded IPs, and flat topologies often steer organizations toward simpler, coarse-grained network-based segmentation. Agent-based approaches offer greater granularity but are limited by OS compatibility, system resources, and administrative access. In environments with mixed capabilities, hybrid models are used only when agent coverage is high enough to justify the added complexity. As such, the choice of strategy depends heavily on network architecture, device compatibility, and the availability of skilled personnel.

8.1.2. Sub-Question 2: Legacy System Challenges

In response to the question of what technical and operational challenges organizations face when implementing micro-segmentation in legacy-heavy environments. The most critical challenges include hard-coded dependencies, unsupported operating systems, resource-constrained hardware, and unmanaged network infrastructure. These factors often prevent the use of modern, agent-based approaches and require organizations to fall back on more compatible network-based methods such as VLANs or PVLANs. Visibility gaps stemming from the presence of legacy systems whose functions, dependencies, and communication patterns are often poorly understood, further complicate segmentation efforts, making a structured discovery phase even more difficult.

8.1.3. Sub-Question 3: Implementation steps

In response to the question of what steps organizations follow to implement micro-segmentation in legacy-heavy environments, this study finds that the process typically unfolds include some general defined phases. While the exact order and detail of steps may vary, most organizations follow a consistent pattern that begins with network separation, continues to network-segmentation and ends with micro-segmentation.

The first step, often referred to as network separation, focuses on isolating high-value or vulnerable systems from flat, enterprise-wide networks. This is especially common in operational technology environments, where critical systems are often exposed to unnecessary risks. Network separation involves moving these systems behind firewalls or other boundary controls to establish a basic level of isolation before deeper segmentation begins.

Following this, organizations move into network segmentation, where broader zones are defined based on system types, risk profiles, or operational roles. This stage uses VLANs, firewalls, and access control lists to establish high-level groupings. In legacy-heavy environments, network-based segmentation follows a phased and infrastructure-driven approach. It begins with identifying and grouping systems through an asset inventory, followed by defining the target architecture based on zones such as user devices, servers, or operational domains.

In the high-level design, teams outline expected communication paths between zones. The low-level design phase includes assigning VLANs, configuring IP subnets, and drafting firewall rules. Implementation preparation involves applying these configurations and validating key services like DNS and routing. During cut-over, zones are activated gradually, and communication is restricted based on predefined rules. The final testing phase ensures that services remain available and security policies are effective. This approach supports stronger segmentation while accommodating the limitations of legacy infrastructure.

Once high-level network segmentation is established, organizations can progress to micro-segmentation to control traffic within those broader zones. At this point, the asset inventory plays a critical role in determining the appropriate path forward. Systems are classified as either agent-compatible or not. If systems are agent-compatible, organizations can proceed with agent-based micro-segmentation, which follows a structured, phased approach designed to reduce risk. It begins with defining the target architecture and mapping communication assumptions. Agents are then deployed in observation mode to collect real traffic data, which informs the creation and fine-tuning of policies. During implementation, tags and rules are applied in monitor-only mode. Enforcement is introduced gradually in the cut-over phase, followed by testing and ongoing policy adjustments. This method provides deep visibility and fine-grained control, especially valuable in complex or sensitive environments.

For systems that are not compatible with agents, organizations continue with network-based micro-segmentation, extending the same roadmap used for initial segmentation. However, at this stage, the zones are made more granular, subdividing existing segments based on more specific roles, communication needs, or risk levels. This allows for improved lateral movement control using the existing network infrastructure, even in legacy-heavy environments.

8.1.4. Sub-Question 4: The Evaluation of Implementation

In response to the question of what defines a successful micro-segmentation implementation and how it is evaluated, this study finds that success depends not on technical complexity alone, but on achieving a balanced outcome across security, stability, visibility, and manageability.

The core criterion is enforcing access control as intended allowing only authorized communication while blocking all other traffic. This is verified through connectivity tests, exploratory attempts to simulate unauthorized access, and log analysis. Equally important is operational continuity meaning that segmentation must not disrupt critical services. This criteria can be confirmed through post-implementation functional testing and monitoring.

Security validation focuses on reducing lateral movement, often measured through penetration testing and traffic analysis. Performance is another practical concern, particularly in legacy environments where additional load from agents or firewalls must be carefully monitored. Successful implementations also ensure broad coverage, leaving no critical systems unsegmented, and minimize residual risk by documenting and controlling any necessary exceptions. Improved visibility is another outcome, enabling better understanding of system interactions and policy refinement. Finally, maintainability is crucial meaning that policies must be sustainable without excessive administrative effort.

8.1.5. Main Research Questions: Micro-segmentation in Legacy-heavy Environment

The implementation of micro-segmentation should follow a phased approach to manage complexity and minimize operational risk. This process consists of three main steps which are network separation, network segmentation and micro-segmentation. At each step, asset inventory is conducted with increasing depth. It begins with a coarse inventory to identify major systems and their dependencies and progresses toward a detailed mapping of communication flows, system configurations and endpoint capabilities. The appropriate segmentation strategy is selected during the micro-segmentation step when the most comprehensive inventory is available.

At this stage, systems can be distinguished based on their compatibility with host-level controls. Agent-based strategy is most appropriate when more than half of the systems in the environment are agent-compatible. This includes having sufficient CPU, memory and modern operating systems that support host-level firewall control. It is ideal when minimal hardware constraints exist and when the organization seeks fine-grained control over east-west traffic with high visibility and scalability. However, this strategy may introduce operational overhead due to agent deployment on each system and device in the network.

Network-based strategy is recommended when agent deployment is not feasible. This approach uses switches and routers to group resources into VLANs or routing domains, and access is controlled at the boundary using firewalls. Depending on the level of control needed, segments can be coarse or more granular using techniques like Private VLANs (PVLANs). This method is compatible with legacy environments and avoids endpoint changes, though it provides less visibility and precision than agent-based solutions.

Hybrid strategy is used in mixed environments where agent deployment is only feasible in part of the network. In such cases, agent-based controls are applied to systems that support them while network-based segmentation secures the rest. It offers flexible enforcement and improved coverage but comes with higher implementation complexity and requires coordinated management between host-level and network-level controls.

Evaluation of micro-segmentation success requires both technical and operational metrics. Access control must function as intended, unauthorized connections should be blocked and permitted paths should align with business requirements. Operational continuity is also critical, as segmentation should not degrade system performance or interrupt essential workflows. Additionally, organizations assess improvements in network visibility, threat containment, and overall security posture. Coverage, defined as the proportion of the network effectively segmented, serves as a key indicator of maturity. A successful implementation balances control granularity with maintainability.

8.2. Discussion

This study set out to understand how organizations implement micro-segmentation in environments with legacy systems to develop a practical framework that reflects real-world constraints. The findings show that while micro-segmentation is central to Zero Trust security, its meaning and application vary widely depending on infrastructure, maturity, and available tools. These variations are not well captured in the existing literature, which often assumes ideal conditions and modern environments.

A key insight is the difference in how micro-segmentation is defined. Experts in this study described it in varied ways. Some referred to technical controls such as host-based enforcement, others focused on limiting lateral movement, and a few emphasized control at the application level. In practice, the definition depends on what is achievable. For instance, organizations with limited visibility or rigid infrastructure may view VLAN-based separation as micro-segmentation. This contrasts with academic literature, which typically treats micro-segmentation as a fixed concept with identity-aware policies and deep granularity. The practical definitions shared by experts reflect a more flexible view. They suggest that segmentation progress should be judged based on feasibility and operational context rather than strict technical benchmarks. This is especially relevant for legacy-heavy environments where full compliance with theoretical models may not be possible.

The study also revealed a variety of strategies that organizations use to implement micro-segmentation. These include VLANs, firewalls, PVLANs, VRFs, and agent-based enforcement. While these methods are known in the literature, this research explains how and why organizations choose among them. The main factor is technical constraint. Many legacy systems cannot support modern agents or virtualization tools. As a result, network-based segmentation remains the most common and practical approach. VLANs are widely used because they require minimal changes to endpoints and integrate with existing hardware. Although this approach may not meet all Zero Trust goals, it provides an achievable way to improve segmentation and limit lateral movement.

Agent-based methods offer fine-grained control and improved visibility but are only used when systems are compatible. Often, organizations apply a hybrid model that combines agent-based and network-level controls. However, this approach adds complexity and is avoided when only a small portion of systems support agents. In most cases, organizations choose strategies based on a careful balance between security benefit, technical capability, and operational effort. Legacy systems introduce several practical challenges that complicate segmentation. Instead of aiming for complete replacement, organizations focus on isolating them using tools that are already available.

Visibility was highlighted as a fundamental prerequisite. Many organizations lack accurate asset inventories or do not fully understand system communication patterns. This makes segmentation risky, especially in flat networks where all devices can interact. Before implementing controls, most teams invest in a discovery phase using passive monitoring, manual checks, and input from system owners. While automation helps, expert judgement remains critical. This aspect is discussed in literature but this study provides more detailed insight into how visibility is achieved in practice.

The implementation process followed by organizations reflects a phased approach. Instead of deploying full segmentation at once, teams start by isolating critical systems or functions that are unnecessarily exposed. After initial isolation, they move to broader segmentation based on business roles or risk levels. Granular controls are added last, and only when systems and staff capacity permit. This step-by-step structure aligns with some academic models but includes adjustments that account for real-world complexity. In legacy-heavy environments, sudden changes can break key systems. Phased segmentation helps reduce this risk and allows for continuous learning and adaptation. The evaluation criteria used by participants show a practical perspective on success. Organizations measure outcomes based on access control accuracy, service stability, and maintainability. Other important factors include coverage, residual risk, and operational effort. If segmentation disrupts normal activities or requires constant manual adjustments, it is seen as unsustainable. The focus is not on achieving perfection but on reducing the most significant risks while maintaining a manageable and stable environment. This differs from much of the literature which tends to emphasize complete policy enforcement and ideal technical outcomes.

This study also shows that success depends on organizational capacity. Teams must match their security goals to available tools, skills, and infrastructure. Attempting advanced controls in unsupported environments can lead to failure. The decision-support framework developed here reflects that reality. It helps organizations identify what is feasible and prioritize actions accordingly. This makes it especially useful in constrained settings where careful planning and incremental improvement are necessary.

Taken together, the findings confirm that micro-segmentation in legacy-heavy environments is shaped by practical constraints rather than ideal designs. While modern technologies offer powerful options, they are not always compatible with existing systems. VLANs, firewalls, and zone-based separation remain essential tools because they align with operational needs and infrastructure limits. This study addresses the gap between academic models and real-world practices by offering a flexible, structured, and validated framework for implementation.

8.3. Research limitation and future direction

While this study offers valuable insights into how organizations can approach micro-segmentation in legacy-heavy environments, it also faces several limitations that should be acknowledged. First, the research was limited to a selected group of cybersecurity experts with experience in industrial and legacy-heavy environments. All interviewees were from the same global organization, a large, multinational firm with operations across various regions. While this provided access to deep expertise and real-world practices, the findings may reflect a specific organizational context and may not fully capture the diversity of approaches used in other sectors, company sizes, or geographical settings.

Second, the study focuses on technical and operational dimensions but does not extensively address the economic, regulatory, or organizational governance aspects that may also influence micro-segmentation strategies. Cost constraints, compliance requirements, and internal policy dynamics could all affect implementation feasibility, especially in critical infrastructure sectors. Third, the study relied on semi-structured interviews for data collection, which carry known methodological limitations. While suitable for exploring expert perspectives, this format can lead to uneven data across participants and may pose challenges in maintaining thematic consistency (DeJonckheere and Vaughn, 2019). The flexibility of the method can also introduce interviewer influence, where subtle biases in tone or phrasing affect participant responses (Jamshed, 2020). Additionally, the open-ended nature of interviews often results in large, unstructured datasets that are time-consuming to analyse, which can limit the feasible number of participants (Morse and Coulehan, 2015).

Fourth, thematic analysis, used to analyse the interview data, also presents limitations. Its flexibility, while useful, can introduce subjectivity, particularly during coding and theme development where researcher interpretation plays a central role (Nowell et al., 2017). Unlike more structured methods such as grounded theory, thematic analysis lacks a standard procedure, which can impact replicability and credibility if not carefully managed (Holloway and Todres, 2003).

Another key limitation concerns the validation of the proposed decision-support framework. Although expert validation was conducted through a structured questionnaire, responses were received from only three experts, all from the same organization as the interviewees involved in the framework's development, which limits both the diversity and quantity of validation input. As noted by Venable et al. (2016), expert evaluations may lack generalizability, particularly when the sample size is small or not diverse across sectors.

While this study contributes a practical and adaptable framework for micro-segmentation in Zero Trust environments, several limitations and opportunities for future research remain. First, the framework has not yet been validated through large-scale simulations, structured surveys, or field studies, and would benefit from input across a broader range of sectors, organizational sizes, and geographic regions. Second, its applicability could be enhanced through sector-specific studies, particularly in healthcare, manufacturing, and energy, where legacy infrastructure poses unique challenges. Third, the analysis focuses solely on micro-segmentation and does not address other essential Zero Trust components such as identity governance, continuous monitoring, or data-centric policy enforcement. Expanding the framework to include these elements would improve its overall strategic relevance.

Furthermore, human factors present a significant challenge in Zero Trust adoption. These include resistance to organizational change, gaps in cybersecurity awareness and training, and the lack of internal readiness or leadership support. Addressing these issues is crucial, as technical solutions alone are insufficient without cultural and procedural alignment. Lastly, it should be investigated how and to what extent artificial intelligence is used to support micro-segmentation in practice. While existing literature indicates that AI can assist in tasks such as creating security zones, generating access control policies, and analysing traffic patterns, the actual depth and consistency of its use across different operational contexts remain unclear. Understanding the maturity and limitations of AI-driven approaches is essential, particularly as organizations seek scalable and adaptive solutions for managing complex network environments.

References

A Al-Ofeishat, H., & Alshorman, R. (2023). Build a secure network using segmentation and micro-segmentation techniques. *International Journal of Computing and Digital Systems*, *16*(1), 1499–1508. https://doi.org/10.12785/ijcds/1601111

Adams, W. C. (2015). Conducting semi-structured interviews. In *Handbook of practical program evaluation* (pp. 492–505). John Wiley & Sons, Ltd. https://doi.org/https://doi.org/10.1002/9781119171386.ch19

Ahmadi, S. (2024). Zero trust architecture in cloud networks: Application, challenges and future opportunities. *Ahmadi, S.(2024). Zero Trust Architecture in Cloud Networks: Application, Challenges and Future Opportunities. Journal of Engineering Research and Reports*. https://doi.org/10.9734/jerr/2024/v26i21083

Aljohani, T. M. (2022). Cyberattacks on energy infrastructures: Modern war weapons. *arXiv preprint arXiv:2208.14225*. https://doi.org/10.48550/arXiv.2208.14225

Álvarez, D., Nuño, P., González, C. T., Bulnes, F. G., Granda, J. C., & García-Carrillo, D. (2023). Performance analysis of software-defined networks to mitigate private vlan attacks. *Sensors*, *23*(4). https://doi.org/10.3390/s23041747

Amoroso, E. G. (2017). Chapter 67 - protecting virtual infrastructure. In J. R. Vacca (Ed.), *Computer and information security handbook (third edition)* (Third Edition, pp. 945–951). Morgan Kaufmann. https://doi.org/10.1016/B978-0-12-803843-7.00067-3

Anjum, I., Kostecki, D., Leba, E., Sokal, J., Bharambe, R., Enck, W., Nita-Rotaru, C., & Reaves, B. (2022). Removing the reliance on perimeters for security using network views. *Proceedings of the 27th ACM on Symposium on Access Control Models and Technologies*, 151–162. https://doi.org/10.1145/3532105.3535029

Arifeen, M., Petrovski, A., & Petrovski, S. (2021). Automated microsegmentation for lateral movement prevention in industrial internet of things (iiot). *2021 14th international conference on security of information and networks (SIN)*, *1*, 1–6. https://doi.org/10.1109/SIN54109.2021.9699232

Arora, S., & Hastings, J. (2024). Microsegmented cloud network architecture using open-source tools for a zero trust foundation. *2024 17th International Conference on Security of Information and Networks (SIN)*, 1–8. https://doi.org/10.1109/SIN63213.2024.10871361

Aslam, N., & Steltzer, H. (2025). Cybersecurity and network security: Strengthening defenses against emerging threats. https://doi.org/10.13140/RG.2.2.16350.96321

Basta, N., Ikram, M., Kaafar, M. A., & Walker, A. (2022). Towards a zero-trust micro-segmentation network security strategy: An evaluation framework. *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*, 1–7. https://doi.org/10.1109/NOMS54207.2022.9789888

Bell, C., Broklyn, P., & Egon, A. (2024). Zero-trust security model for enhanced cloud security and data privacy. *Available at SSRN 4904958*.

Bellamkonda, S. (2022). Zero trust architecture implementation: Strategies, challenges, and best practices. *International Journal of Communication Networks and Information Security*.

Bentaleb, O., Belloum, A. S., Sebaa, A., & El-Maouhab, A. (2022). Containerization technologies: Taxonomies, applications and challenges. *The Journal of Supercomputing*, *78*(1), 1144–1181. https://doi.org/10.1007/s11227-021-03914-1

Binbeshr, F., Imam, M., & Ghaleb, M. (2025). The rise of cognitive socs: A systematic literature review on ai approaches. *IEEE Open Journal*. https://doi.org/10.1109/OJCS.2025.3536800

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, *3*(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

Bunker, N. (2022). Employer appetite for remote work remains elevated for some jobs [accessed June 25, 2025]. https://www.hiringlab.org/2022/06/16/remote-work-update-2022/

Bushouse, M., & Reeves, D. (2018). Goalkeeper: Comprehensive process enforcement from the hypervisor. *Computers & Security*, *73*, 459–473. https://doi.org/https://doi.org/10.1016/j.cose.2017.11.020

Campbell, J. L., Quincy, C., Osserman, J., & Pedersen, O. K. (2013). Coding in-depth semistructured interviews: Problems of unitization and intercoder reliability and agreement. *Sociological Methods & Research*, *42*(3), 294–320. https://doi.org/10.1177/0049124113500475

Carbone, M., Zamboni, D., & Lee, W. (2008). Taming virtualization. *IEEE Security & Privacy*, *6*(1), 65–67. https://doi.org/10.1109/MSP.2008.24

Chandramouli, R., & Chandramouli, R. (2022). *Guide to a secure enterprise network landscape*. US Department of Commerce, National Institute of Standards; Technology. https://doi.org/10.6028/NIST.SP.800-215

Chang, W. Y., Abu-Amara, H., & Sanford, J. F. (2010). *Transforming enterprise cloud services*. Springer Netherlands. https://doi.org/10.1007/978-90-481-9846-7

Chen, B., Qiao, S., Zhao, J., Liu, D., Shi, X., Lyu, M., Chen, H., Lu, H., & Zhai, Y. (2020). A security awareness and protection system for 5g smart healthcare based on zero-trust architecture. *IEEE internet of things journal*, *8*(13), 10248–10263. https://doi.org/10.1109/JIOT.2020.3041042

Clarke, V., & Braun, V. (2016). Thematic analysis. *The Journal of Positive Psychology*, *12*, 1–2. https://doi.org/10.1080/17439760.2016.1262613

ColorTokens. (2025). Approaches to microsegmentation [Accessed: 2025-03-03]. https://colortokens.com/microsegmentation/

Cunningham, C. (2020). A look back at zero trust: Never trust, always verify. *Forrester*.

Cunningham, C., Blankenship, J., Balaouras, S., Murphy, R., & Cyr, M. (2018). The zero trust extended (ztx) ecosystem. *Forrester, Cambridge, MA*.

Cyber Kendra. (2023, May). Strengthening cloud security through micro-segmentation and zero trust architecture [Accessed: June 23, 2025]. https://www.cyberkendra.com/2023/05/strengthening-cloud-security-through.html

Cyberint. (2024). Europe threat landscape report [Accessed: July 5, 2025]. https://cyberint.com/blog/research/europe-threat-landscape-report/

da Rocha, B. C., de Melo, L. P., & de Sousa, R. T. (2021). Preventing apt attacks on lan networks with connected iot devices using a zero trust based security model. *2021 Workshop on Communication Networks and Power Systems (WCNPS)*, 1–6. https://doi.org/10.1109/WCNPS53648.2021.9626270

DeJonckheere, M., & Vaughn, L. M. (2019). Semistructured interviewing in primary care research: A balance of relationship and rigour. *Family Medicine and Community Health*, *7*(2), e000057. https://doi.org/10.1136/fmch-2018-000057

Desai, B., & Patil, A. (2020). Zero trust with micro-segmentation: A software-defined approach to securing cloud-native applications. *Annals of Applied Sciences*, *1*(1).

DiCicco-Bloom, B., & Crabtree, B. F. (2006). The qualitative research interview. *Medical Education*, *40*(4), 314–321. https://doi.org/10.1111/j.1365-2929.2006.02418.x

Durojaye, H., & Raji, O. (2022). Impact of state and state sponsored actors on the cyber environment and the future of critical infrastructure. *arXiv preprint arXiv:2212.08036*. https://doi.org/10.48550/arXiv.2212.08036

Duruk, Ü., Akgün, A., & Tokur, F. (2019). Prospective early childhood teachers' understandings on the nature of science in terms of scientific knowledge and scientific method. *Universal Journal of Educational Research*, *7*, 675–690. https://doi.org/10.13189/ujer.2019.070306

Fernandez, E. B., & Brazhuk, A. (2024). A critical analysis of zero trust architecture (zta). *Computer Standards & Interfaces*. https://doi.org/10.1016/j.csi.2024.103832

Foltz, S. (2022). Zero trust technology integration issues. https://www.jstor.org/stable/resrep34846

Gadkari, B. (2025). Ai integration in zero trust security architecture: A technical overview, 7. https://doi.org/10.56726/IRJMETS67329

George, A. S. (2025). The critical role of micro-segmentation in modern cybersecurity architectures: A comprehensive review. https://doi.org/10.5281/zenodo.15063176

Gill, P., Stewart, K., Treasure, E., & Chadwick, B. (2008). Methods of data collection in qualitative research: Interviews and focus groups. *British Dental Journal*, *204*(6), 291–295. https://doi.org/10.1038/bdj.2008.192

Herranz-Oliveros, D., Marsa-Maestre, I., Gimenez-Guzman, J. M., Tejedor-Romero, M., & de la Hoz, E. (2024). Surgical immunization strategies against lateral movement in active directory environments. *Journal of Network and Computer Applications*, *222*, 103810. https://doi.org/https://doi.org/10.1016/j.jnca.2023.103810

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, *28*(1), 75–105. https://doi.org/10.2307/25148625

Holloway, I., & Todres, L. (2003). The status of method: Flexibility, consistency and coherence. *Qualitative Research*, *3*(3), 345–357. https://doi.org/10.1177/1468794103033004

Hunold, S., Krellner, B., Rauber, T., Reichel, T., & Rünger, G. (2009). Pattern-based refactoring of legacy software systems. In J. Filipe & J. Cordeiro (Eds.), *Enterprise information systems* (pp. 78–89). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-01347-8_7

Jammal, M., Singh, T., Shami, A., Asal, R., & Li, Y. (2014). Software-defined networking: State of the art and research challenges. https://doi.org/10.48550/arXiv.1406.0124

Jamshed, S. (2020). Qualitative research method-interviewing and observation. j basic clin pharm [online] 5 (4): 87–88. https://doi.org/10.4103/0976-0105.141942

Jeuk, S., Salgueiro, G., & ZHou, S. (2015). A novel approach to classify cloud entities: Universal cloud classification (ucc). *2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, 801–804. https://doi.org/10.1109/CCGrid.2015.127

Jones, C. (2025). Warnings & lessons of the 2013 target data breach [Accessed July 2, 2025]. https://redriver.com/security/target-data-breach

Kallio, H., Pietilä, A.-M., Johnson, M., & Kangasniemi, M. (2016). Systematic methodological review: Developing a framework for a qualitative semi-structured interview guide. *Journal of Advanced Nursing*, *72*(12), 2954–2965. https://doi.org/10.1111/jan.13031

Kang, C., Li, Y., Gao, L., Li, X., Wang, L., & Gao, Y. (2022). Automatic generation model of host micro-segmentation in distribution master station based on distributed intrusion detection. *2022 2nd International Conference on Computer Science, Electronic Information Engineering and Intelligent Control Technology (CEI)*, 54–59. https://doi.org/10.1109/CEI57409.2022.9948463

Kang, H., Liu, B., Mišić, J., Mišić, V. B., & Chang, X. (2020). Assessing security and dependability of a network system susceptible to lateral movement attacks. *2020 International Conference on Computing, Networking and Communications (ICNC)*, 513–517. https://doi.org/10.1109/ICNC47757.2020.9049748

Kim, Y., Sohn, S.-G., Jeon, H. S., Lee, S.-M., Lee, Y., & Kim, J. (2024). Exploring effective zero trust architecture for defense cybersecurity: A study. *KSII Transactions on Internet and Information Systems (TIIS)*, *18*(9), 2665–2691. https://doi.org/10.3837/tiis.2024.09.011

Kindervag, J. (2010). *Build security into your network's DNA: The zero trust network architecture* (tech. rep.). Forrester Research. https://www.virtualstarmedia.com/downloads/Forrester_zero_trust_DNA.pdf

Køien, G. M. (2021). Zero-trust principles for legacy components: 12 rules for legacy devices: An antidote to chaos. *Wireless Personal Communications*, *121*(2), 1169–1186.

Kolawole, I. (2025). Leveraging cloud-based ai and zero trust architecture to enhance u. s. cybersecurity and counteract foreign threats. *World Journal of Advanced Research and Reviews*. https://doi.org/10.30574/wjarr.2025.25.3.0635

Koskinen, J. (2020). Microsegmentation as part of organization's network architecture: Investigating vmware nsx for vsphere.

Krause, T., Ernst, R., Klaer, B., Hacker, I., & Henze, M. (2021). Cybersecurity in power grids: Challenges and opportunities. *Sensors*, *21*(18), 6225. https://doi.org/10.3390/s21186225

Kreutz, D., Ramos, F. M. V., Veríssimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. *Proceedings of the IEEE*, *103*(1), 14–76. https://doi.org/10.1109/JPROC.2014.2371999

Lenarduzzi, V., Lomio, F., Saarimäki, N., & Taibi, D. (2020). Does migrating a monolithic system to microservices decrease the technical debt? *Journal of Systems and Software*, *169*, 110710. https://doi.org/https://doi.org/10.1016/j.jss.2020.110710

Li, D., Yang, Z., Yu, S., Duan, M., & Yang, S. (2024). A micro-segmentation method based on vlan-vxlan mapping technology. *Future Internet*, *16*(9), 320.

Lin, L., & Lin, P. (2014). Software-defined networking (sdn) for cloud applications. In Z. Mahmood (Ed.), *Cloud computing: Challenges, limitations and r&d solutions* (pp. 209–233). Springer International Publishing. https://doi.org/10.1007/978-3-319-10530-7_9

Liu, Y., Liu, G., Du, H., Niyato, D., Kang, J., Xiong, Z., Kim, D. I., & Shen, X. (2024). Hierarchical micro-segmentations for zero-trust services via large language model (llm)-enhanced graph diffusion. *arXiv preprint arXiv:2406.13964*. https://doi.org/10.48550/arXiv.2406.13964

Ma, M., Yu, Z., & Liu, B. (2023). Automatic generation of network micro-segmentation policies for cloud environments. *2023 4th International Seminar on Artificial Intelligence, Networking and Information Technology (AINIT)*, 1–5. https://doi.org/10.1109/AINIT59027.2023.10212857

Ma, Z., Li, C., Zhang, Y., You, R., & Tu, B. (2025). Ims: Towards computability and dynamicity for intent-driven micro-segmentation. *IEEE Transactions on Dependable and Secure Computing*, *22*(1), 677–694. https://doi.org/10.1109/TDSC.2024.3413752

Mämmelä, O., Hiltunen, J., Suomalainen, J., Ahola, K., Mannersalo, P., & Vehkaperä, J. (2016). Towards micro-segmentation in 5g network security.

Markus, K. (2024). Strategies for network segmentation : A systematic literature review [Accessed: 2025-03-3]. https://jyx.jyu.fi/jyx/Record/jyx_123456789_92952

Mastering Nutanix. (2020, November). What is micro-segmentation? [Accessed: June 23, 2025]. https://masteringnutanix.com/2020/11/28/what-is-micro-segmentation/

Masuda, Y., Shirasaka, S., Yamamoto, S., & Hardjono, T. (2017). An adaptive enterprise architecture framework and implementation: Towards global enterprises in the era of cloud/mobile it/digital it. *International Journal of Enterprise Information Systems*, *13*, 1–22. https://doi.org/10.4018/ijeis.2017070101

Morse, J. M., & Coulehan, J. (2015). Maintaining confidentiality in qualitative publications. *Qualitative Health Research*, *25*(2), 151–152. https://doi.org/10.1177/1049732314563489

Nowell, L. S., Norris, J. M., White, D. E., & Moules, N. J. (2017). Thematic analysis: Striving to meet the trustworthiness criteria. *International Journal of Qualitative Methods*, *16*(1), 1609406917733847. https://doi.org/10.1177/1609406917733847

Nunes, B. A. A., Mendonca, M., Nguyen, X.-N., Obraczka, K., & Turletti, T. (2014). A survey of software-defined networking: Past, present, and future of programmable networks. *IEEE Communications Surveys & Tutorials*, *16*(3), 1617–1634. https://doi.org/10.1109/SURV.2014.012214.00180

Ofili, B., & Obasuyi, O. (2025). Enhancing federal cloud security with ai: Zero trust, threat intelligence and cisa compliance. *World Journal of Advanced Research and Reviews*, *25*, 2377–2400. https://doi.org/10.30574/wjarr.2025.25.2.0620

Ojo, A. O. (2025). Adoption of zero trust architecture (zta) in the protection of critical infrastructure. *Path of Science*, *11*(1), 5001–5009. https://doi.org/10.22178/pos.113-2

Oladimeji, G. (2024). A critical analysis of foundations, challenges and directions for zero trust security in cloud environments. https://doi.org/10.48550/arXiv.2411.06139

Omar, R. R., & Abdelaziz, T. M. (2020). A comparative study of network access control and software-defined perimeter. *Proceedings of the 6th International Conference on Engineering & MIS 2020*, 1–5. https://doi.org/10.1145/3410352.3410754

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, M., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., … Moher, D. (2021). The prisma 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, *372*, n71. https://doi.org/10.1136/bmj.n71

Paul, F. (2022). From legacy systems to zero trust: Transitioning your organization's security model. https://www.researchgate.net/publication/385782309

Peffers, K., Tuunanen, T., Rothenberger, M. A., & and, S. C. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Phiayura, P., & Teerakanok, S. (2023). A comprehensive framework for migrating to zero trust architecture. *Ieee Access*. https://doi.org/10.1109/ACCESS.2023.3248622

Ponnusamy, S., & Eswararaj, D. (2023). Navigating the modernization of legacy applications and data: Effective strategies and best practices. *Asian Journal of Research in Computer Science*, *16*(4), 239–256. https://doi.org/10.9734/AJRCOS/2023/v16i4386

Rajasekharan, D. (2025). Simplifying attribute-based access control (abac) for modern enterprises. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, *11*(1), 3140–3149. https://doi.org/10.32628/cseit251112332

Sarkar, S., Choudhary, G., Shandilya, S. K., Hussain, A., & Kim, H. (2022). Security of zero trust networks in cloud computing: A comparative review. *Sustainability*. https://doi.org/https://doi.org/10.3390/su141811213

Shaik, A. (2024). A survey of emerging techniques for large networks of virtual local area networks (vlans) with benefits and limitations.

Sheikh, N., Pawar, M., & Lawrence, V. (2021). Zero trust using network micro segmentation. *IEEE INFOCOM 2021-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 1–6. https://doi.org/10.1109/INFOCOMWKSHPS51825.2021.9484645

Siadati, H., & Memon, N. (2017). Detecting structurally anomalous logins within enterprise networks. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1273–1284. https://doi.org/10.1145/3133956.3134003

Smith, J., & Cardenas, A. (2024). *Ukraine blackouts caused by malware attacks warn against evolving cybersecurity threats to the physical world* [UCSC News]. https://news.ucsc.edu/2024/05/ukraine-cybersecurity/

SnapLogic. (2025). The potential of generative ai to overcome legacy tech and data challenges [Accessed: 2025-03-04]. https://www.snaplogic.com/blog/genai-potential-legacy-tech-data-challenges

Sokappadu, B., & Mungur, A. (2021). A middleware for integrating legacy network devices into software-defined networking (sdn). In R. Zitouni, A. Phokeer, J. Chavula, A. Elmokashfi, A. Gueye, & N. Benamar (Eds.), *Towards new e-infrastructure and e-services for developing countries* (pp. 121–139). Springer International Publishing. https://doi.org/10.1007/978-3-030-70572-5_8

Stafford, V. (2020). Zero trust architecture. *NIST special publication*, *800*, 207.

Suresh, J. (2016, April). How to choose a micro-segmentation technology [Accessed: June 23, 2025]. https://itmt535-janarthansuresh.blogspot.com/2016/04/how-to-choose-micro-segmentation.html

Suri, N., Carvalho, M., Bradshaw, J., Breedy, M., Cowin, T., Groth, P., Saavedra, R., & Uszok, A. (2003). Enforcement of communications policies in software agent systems through mobile code. *Proceedings POLICY 2003. IEEE 4th International Workshop on Policies for Distributed Systems and Networks*, 247–250. https://doi.org/10.1109/POLICY.2003.1206981

Syed, N. F., Shah, S. W., Shaghaghi, A., Anwar, A., Baig, Z., & Doss, R. (2022). Zero trust architecture (zta): A comprehensive survey. *IEEE access*. https://doi.org/10.1109/ACCESS.2022.3174679

T V, S., Karunakaran, K., Raju, K., J, B. J., & Jebaraj Solomon, E. (2023). Automation of enforcer agent installation. *2023 International Conference on Innovative Computing, Intelligent Communication and Smart Electrical Systems (ICSES)*, 1–4. https://doi.org/10.1109/ICSES60034.2023.10465458

Tuglular, T. (2008). Automatic enforcement of location aware user based network access control policies.

Tyler, D., & Viana, T. (2021). Trust no one? a framework for assisting healthcare organisations in transitioning to a zero-trust network architecture. *Applied Sciences*. https://doi.org/10.3390/app11167499

U.S. Senate Committee on Commerce, Science, and Transportation. (2014). A "kill chain" analysis of the 2013 target data breach [Accessed June 25, 2025]. https://www.commerce.senate.gov/services/files/24d3c229-4f2f-405d-b8db-a3a67f183883

Venable, J., Pries-Heje, J., & and, R. B. (2016). Feds: A framework for evaluation in design science research. *European Journal of Information Systems*, *25*(1), 77–89. https://doi.org/10.1057/ejis.2014.36

VMware. (2020). Vmware nsx micro-segmentation [Accessed: 2025-06-28]. https://www.vmware.com/docs/vmware-nsx-microsegmentation

Wikipedia contributors. (2025). 2015 ukraine power grid hack [Accessed: May 2025].

Xie, L., Hang, F., Guo, W., Lv, Y., & Chen, H. (2021). A micro-segmentation protection scheme based on zero trust architecture. *ISCTT 2021; 6th International Conference on Information Science, Computer Technology and Transportation*, 1–4.

Zaenchkovski, A., Lazarev, A., & Shtempel, A. (2023). Intelligent interface for secure routing of tcp/ip connections to the kvm hypervisor. *2023 International Russian Automation Conference (RusAutoCon)*, 871–875. https://doi.org/10.1109/RusAutoCon58002.2023.10272865

Zaheer, Z., Chang, H., Mukherjee, S., & Van der Merwe, J. (2019). Eztrust: Network-independent zero-trust perimeterization for microservices. *Proceedings of the 2019 ACM Symposium on SDN Research*. https://doi.org/10.1145/3314148.3314349

Zanasi, C., Marchetti, M., & Colajanni, M. (2024). Cybersecurity domains: A design pattern for creating zero trust architectures through microsegmentation. *2024 IEEE Conference on Dependable, Autonomic and Secure Computing (DASC)*, 15–22. https://doi.org/10.1109/DASC64200.2024.00009

Zhao, Z., Hong, F., & Li, R. (2017). Sdn based vxlan optimization in cloud computing networks. *IEEE Access*, *5*, 23312–23319. https://doi.org/10.1109/ACCESS.2017.2762362

A

Interview Coding Results

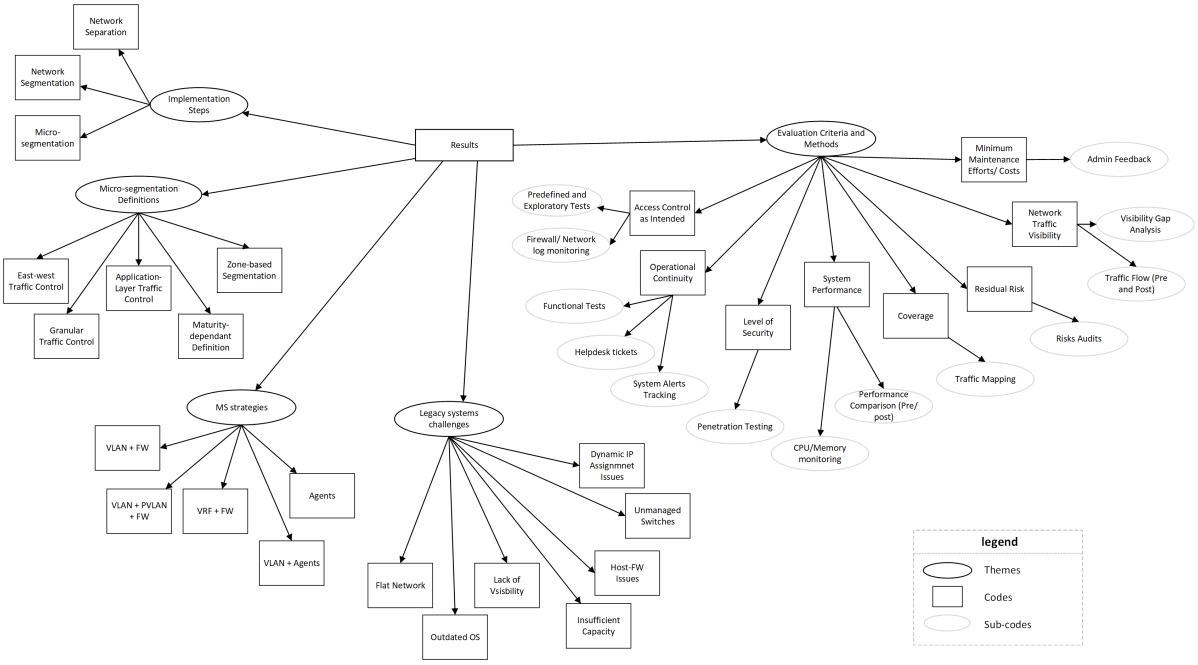

Figure A.1: Interview Coding Results

B

Informed Consent

Dear Participant,

You are being invited to participate in a research study on the selection of micro-segmentation strategies for legacy systems within a Zero Trust Architecture (ZTA). The purpose of this study is to understand how organizations approach the integration of legacy systems into Zero Trust environments, which factors influence the selection of micro-segmentation strategies, and what practical challenges and trade-offs professionals encounter during this process. The interview will take approximately 40–50 minutes to complete.

The collected data will be used for academic research, including publication in a Master's thesis, and may also provide insights relevant to industry practice. You will be asked to share your professional experiences, decision-making approaches, and any challenges you've faced related to micro-segmentation strategy selection specifically for legacy environments. If the interview is conducted in person, it will be audio recorded. If conducted remotely via Microsoft Teams, both audio and video may be recorded, with your consent.

As with any online activity, the risk of a data breach is always possible. To the best of our ability, your responses in this study will remain confidential and anonymised. No personally identifiable information will be included in any research outputs, and no personal or sensitive data will be shared with EY. All personal data collected during the project will be stored at TU Delft, accessible only to the TU Delft research team. The personal data will be handled according the European General Data Protection Regulation (GDPR).

After the interview, we will produce an anonymous technical summary of the conversation. That summary will be sent to you for review. If you have any concerns regarding the content, you are welcome to reach out to us. The summary will be included in the publicly available MSc thesis.

Your participation in this study is entirely voluntary, and you may withdraw at any time without providing a reason. You are also free to omit any questions. If you wish to withdraw after participation, you may contact the project members at any time. All personal data collected will be deleted ± 2 years after graduation. The data may be reused for further scientific or educational activities on the topic of secure ICT infrastructure. Should the material be used, you will be anonymous in any and all outputs.

For any questions regarding this study, you may contact Sofia Sakhi at [School email] or [Company email]. If you have any concerns or complaints or questions about this research, please contact Michiel Meesters, the EY supervisor at [email] or Professor M. Jansen the TU Delft supervisor at [email]. By signing below, you confirm that you have read and understood the study information above or that it has been read to you. You have had the opportunity to ask questions, and they have been answered to your satisfaction.

I have read and understood the study information above and I consent to participate to the study and to the data processing described above.

Full Name:
Date:
Signature:

C

Framework Interview Protocol

Introduction

- Introduction of the interviewer (affiliation, purpose of the interview ...)
- Structure of the interview
- Context and importance of the research
- Clarify confidentiality, anonymity, and how data will be used
- Collecting informed consent forms

Background Questions

- Could you briefly describe your current role and responsibilities?
- Could you briefly describe your professional background, particularly your experience with micro-segmentation as a critical component of Zero Trust Architecture?

Main Questions

Micro-segmentation Strategies:

- What strategies have you used or observed being used for micro-segmenting the network?
- How are segments defined and how is the access policy created?
- What tools and technologies are used to create the segments and enforce access policies?
- Which factors led you to select this strategy in those cases?
- Could you provide the main advantages and limitations of these strategies?
- What are the key steps involved in implementing these micro-segmentation strategy, from initial planning to full deployment?

Legacy System Challenges:

- What specific challenges have you encountered with legacy systems when implementing micro-segmentation? Could you share examples?
- How did you adjust your segmentation strategy to address those challenges?

Evaluation of Implementation:

- What defines a successful implementation of micro-segmentation in your organization?
- What criteria do you use to assess whether the strategy has achieved its goals?
- What tools do you use to measure or validate these criteria?

Closing the Interview

- Summarize the main points discussed and ask if any aspects were missed.
- Ask if the interviewee can recommend others with relevant experience.
- Request permission to follow up if clarification is needed.
- Mention a technical summary will be shared for their review.
- Thank the interviewee and formally close the interview.

D

Validation Interview protocol

Introduction

- Brief introduction of the interviewer
- Explain the objective: validating a practical framework based on earlier findings
- Clarify what kind of feedback is being sought (clarity, usefulness, completeness, etc.)
- Remind participant of confidentiality and ethical practices (as discussed in the main study)
- Confirm informed consent and that participation is voluntary
- Asking if they had the chance to review the framework materials provided?

Validation Questions

Decision Flow:

- Is the decision flow for selecting between agent-based, network-based, or hybrid segmentation strategies clear and realistic?
- Do you think the criteria used to guide the choice are complete and applicable in practice?
- Would you suggest any changes or additions to improve decision accuracy?

General Usefulness:

- Do you believe this framework could be useful for organizations with legacy systems?
- What types of users or teams would benefit most from using this framework?
- Is anything unclear or difficult to interpret in the materials?

Closing the Interview

- Ask if the participant has any general comments or concerns
- Thank them for their time and input
- Offer to share final results or acknowledgments, if interested
- Ask if follow-up is okay in case of clarification needs

E

Informed Consent Validation

Dear Participant,

You are being invited to participate in a research study on the selection of micro-segmentation strategies for legacy systems within a Zero Trust Architecture (ZTA). The purpose of this study is to understand how organizations approach the integration of legacy systems into Zero Trust environments, which factors influence the selection of micro-segmentation strategies, and what practical challenges and trade-offs professionals encounter during this process. The interview will take approximately 40–50 minutes to complete.

The collected data will be used for academic research, including publication in a Master's thesis, and may also provide insights relevant to industry practice. You will be asked to share your professional experiences, decision-making approaches, and any challenges you've faced related to micro-segmentation strategy selection specifically for legacy environments. If the interview is conducted in person, it will be audio recorded. If conducted remotely via Microsoft Teams, both audio and video may be recorded, with your consent.

As with any online activity, the risk of a data breach is always possible. To the best of our ability, your responses in this study will remain confidential and anonymised. No personally identifiable information will be included in any research outputs, and no personal or sensitive data will be shared with EY. All personal data collected during the project will be stored at TU Delft, accessible only to the TU Delft research team. The personal data will be handled according the European General Data Protection Regulation (GDPR).

After the interview, we will produce an anonymous technical summary of the conversation. That summary will be sent to you for review. If you have any concerns regarding the content, you are welcome to reach out to us. The summary will be included in the publicly available MSc thesis.

Your participation in this study is entirely voluntary, and you may withdraw at any time without providing a reason. You are also free to omit any questions. If you wish to withdraw after participation, you may contact the project members at any time. All personal data collected will be deleted ± 2 years after graduation. The data may be reused for further scientific or educational activities on the topic of secure ICT infrastructure. Should the material be used, you will be anonymous in any and all outputs.

For any questions regarding this study, you may contact Sofia Sakhi at [School email] or [Company email]. If you have any concerns or complaints or questions about this research, please contact Michiel Meesters, the EY supervisor at [email] or Professor M. Jansen the TU Delft supervisor at [email]. By signing below, you confirm that you have read and understood the study information above or that it has been read to you. You have had the opportunity to ask questions, and they have been answered to your satisfaction.

I have read and understood the study information above and I consent to participate to the study and to the data processing described above.

Full Name:

Date:

Signature:

F

Challenges, Solutions, advantages and limitations of each Strategy

F.1. Network-based (VLAN and Firewall)

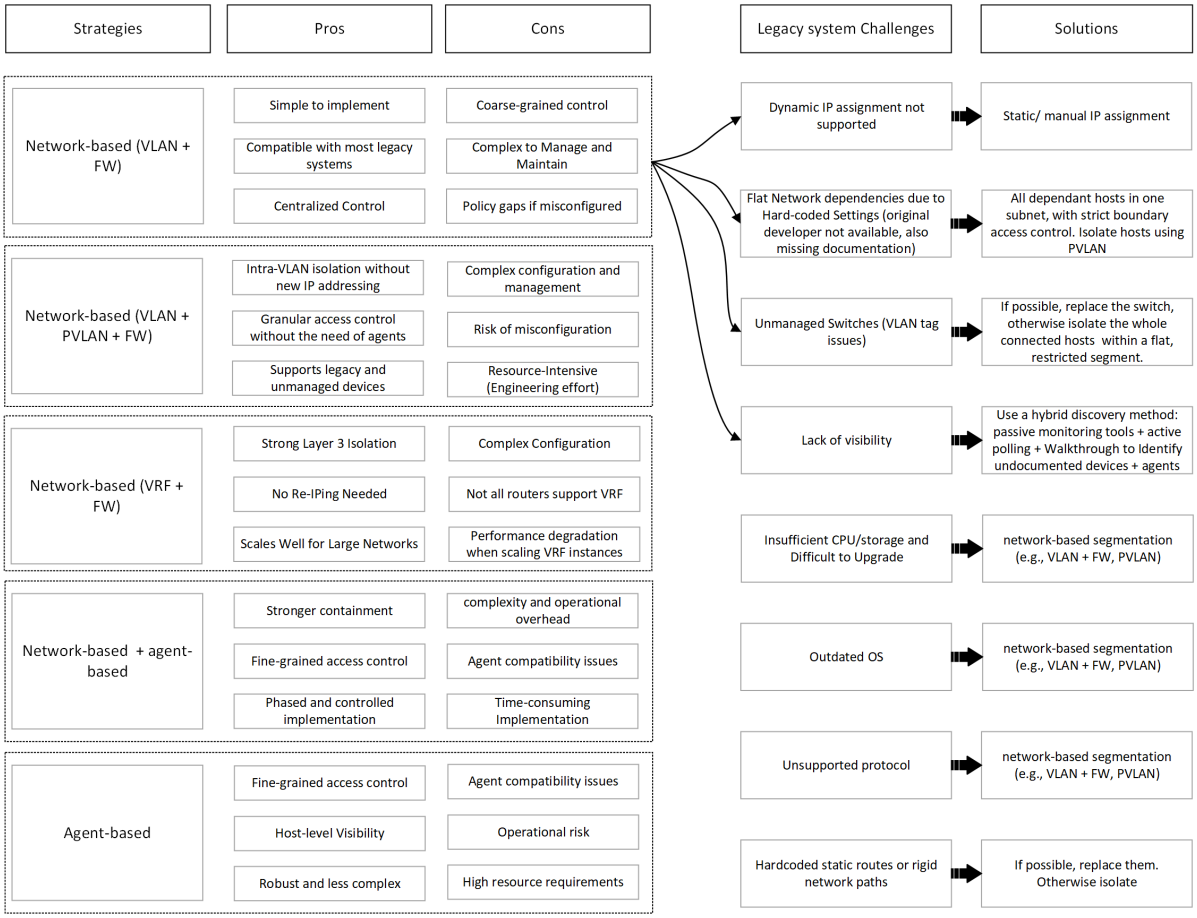

Figure F.1: Legacy system challenges when implementing VLAN and firewall micro-segmentation

F.2. Network-based (VLAN + PVLAN and Firewall)

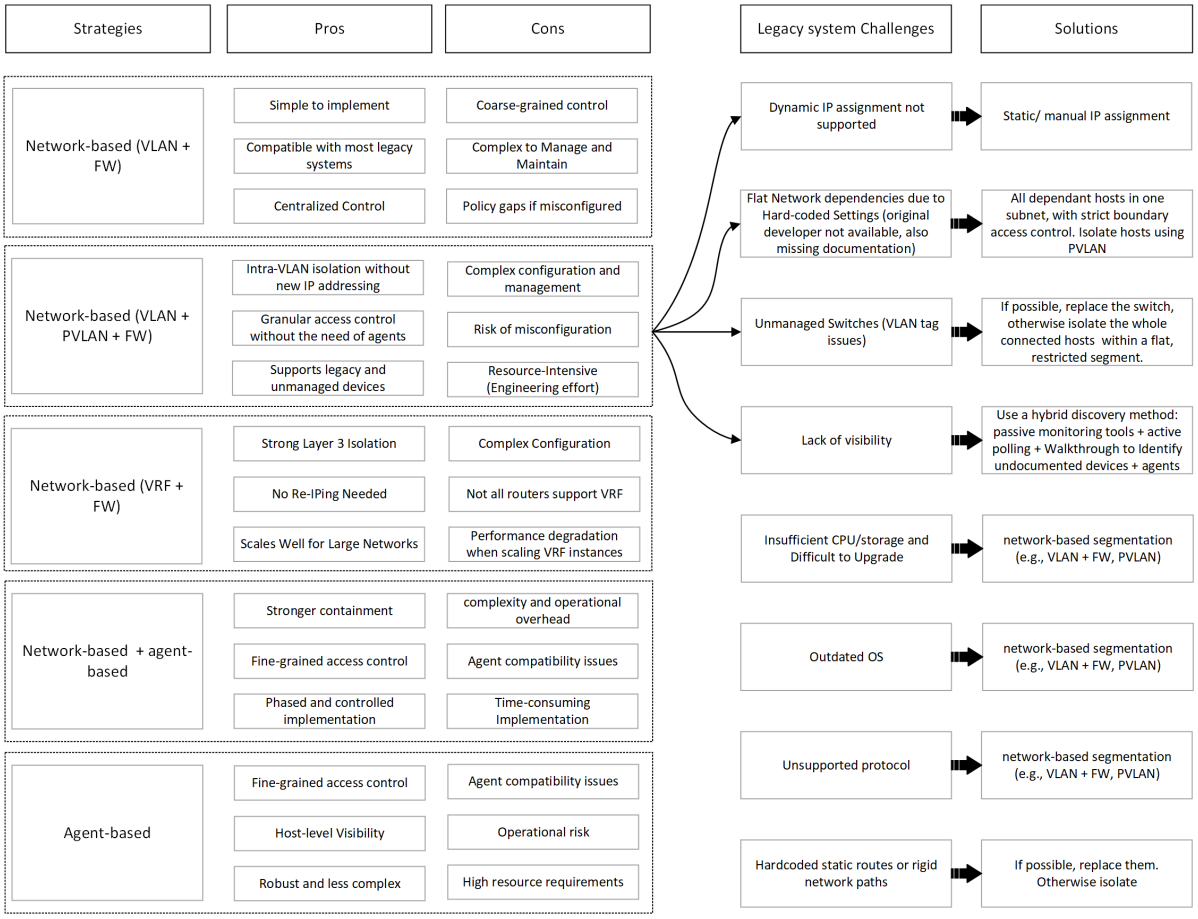

Figure F.2: Legacy system challenges when implementing VLAN + PVLAN and firewall micro-segmentation

F.3. Network-based (VRF and Firewall)

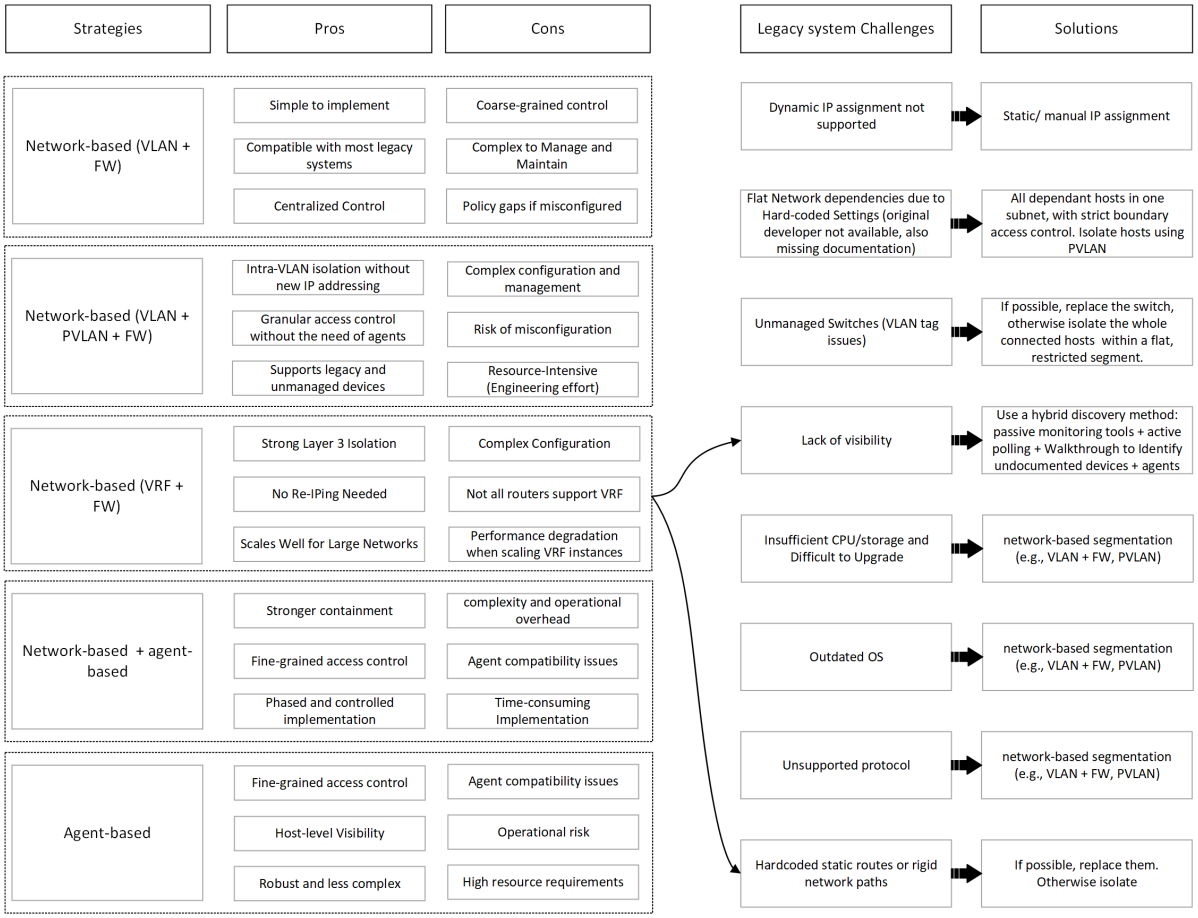

Figure F.3: Legacy system challenges when implementing VRF and firewall micro-segmentation

F.4. Agent-based

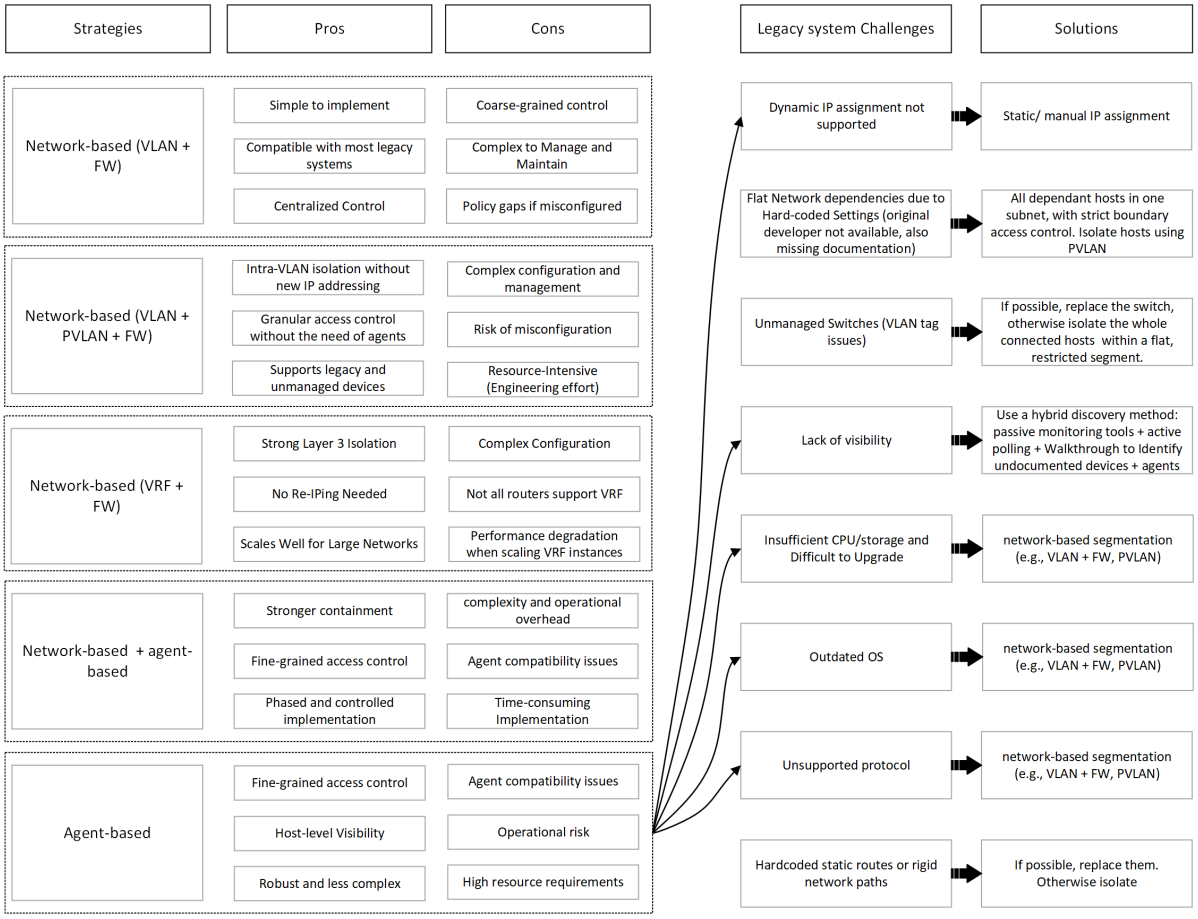

Figure F.4: Legacy system challenges when implementing Agent-based micro-segmnetation

F.5. Hybrid approach

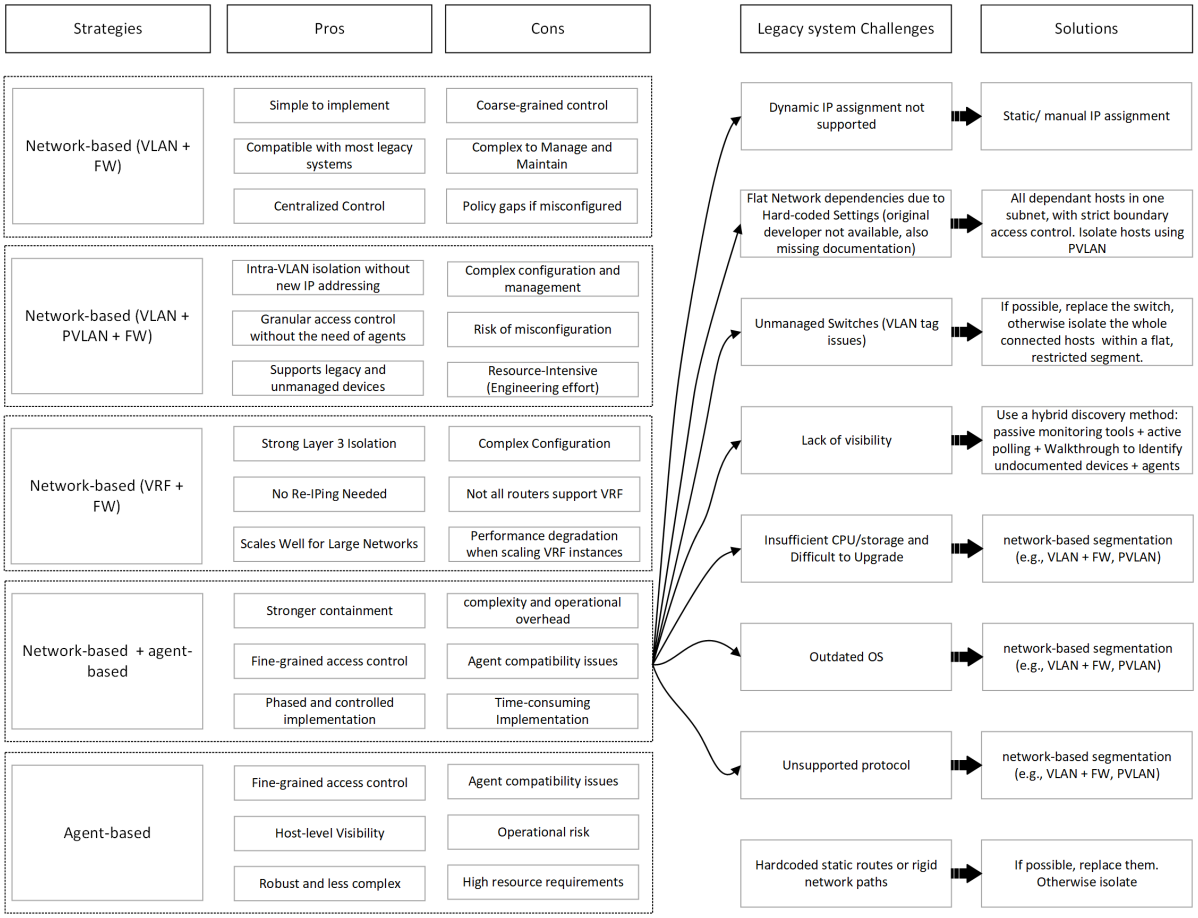

Figure F.5: Legacy system challenges when implementing hybrid approach (network + agent-based micro-segmentation)