

A Zero Trust Architecture for Secure Connected and Autonomous Vehicles

Final Report

by

Long Cheng
Clemson University

Zhenkai Zhang, Clemson University
Gurcan Comert, Benedict College

July 2025



**NATIONAL CENTER FOR TRANSPORTATION
CYBERSECURITY AND RESILIENCY (TraCR)**





DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by the National Center for Transportation Cybersecurity and Resiliency (TraCR) under Grant No. 69A3552344812 and 69A3552348317 which is headquartered at Clemson University, South Carolina, USA, from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

Non-exclusive rights are retained by the U.S. DOT.

CONTACTS

For more information:

Long Cheng
210 McAdams Hall
Clemson University
Phone: 8642380125
Email: lcheng2@clemson.edu

TraCR
Clemson University
One Research Dr
Greenville, SC 29607
tracr@clemson.edu



ACKNOWLEDGMENT

This work is based upon the work supported by the National Center for Transportation Cybersecurity and Resiliency (TraCR) (a U.S. Department of Transportation National University Transportation Center) headquartered at Clemson University, Clemson, South Carolina, USA. Any opinions, findings, conclusions, and recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of TraCR, and the U.S. Government assumes no liability for the contents or use thereof.



National Center for Transportation Cybersecurity and Resiliency (TraCR)

Technical Report Documentation Page

1. Report No. 11	2. Government Accession No. N/A	3. Recipient's Catalog No. N/A	
4. Title and Subtitle A Zero Trust Architecture for Secure Connected and Autonomous Vehicles		5. Report Date: July 2025	
		6. Performing Organization Code: N/A	
7. Author(s) Long Cheng, Ph.D.; https://orcid.org/0000-0003-1736-0873 Zhenkai Zhang, Ph.D.; https://orcid.org/0000-0002-9025-3460 Gurcan Comert, Ph.D.; https://orcid.org/0000-0002-2373-5013		8. Performing Organization Report No. 11	
9. Performing Organization Name and Address National Center for Transportation Cybersecurity and Resiliency (TraCR) Clemson University 414 A One Research Dr, Greenville, SC 29607 School of Computing, Clemson University McAdams Hall, 821 McMillan Rd, Clemson, SC 29634		10. Work Unit No. N/A	
		11. Contract or Grant No. 69A3552344812 and 69A3552348317	
12. Sponsoring Agency Name and Address U.S. Department of Transportation Office of the Assistant Secretary for Research and Technology 1200 New Jersey Avenue, SE, Washington, DC 20590		13. Type of Report and Period Covered Final Report, 01/01/2024 - 05/31/2025	
		14. Sponsoring Agency Code OST-R	
15. Supplementary Notes Conducted under the U.S. DOT Office of the Assistant Secretary for Research and Technology's (OST-R) University Transportation Centers (UTC) program.			
16. Abstract Connected and Autonomous Vehicles (CAVs) are the future of personal and public transportation. As CAVs increasingly rely on cyber-based control, navigation, and communication, security has become a pressing concern in future transportation systems. The complexity and inter-connectedness of CAVs offer myriad opportunities for security compromise, potentially resulting in unsafe operation or leakage of confidential information about the user. Zero Trust Architectures (ZTA) for networks have emerged as a fundamentally new way of approaching network security. The zero-trust security model does not automatically trust any user or device inside or outside the network perimeter. Instead, it enforces a set of policies (i.e., rules that are dynamically maintained and enforced) to verify and ensure the security of resources. In this project, we mainly designed a Zero Trust Architecture for CAV (named ZTA-CAV) to provide fundamental protection for individual components within CAV systems and their supporting infrastructure. In addition, as Artificial Intelligence (AI) continues to become an integral part of cybersecurity, we discuss the potential opportunities and challenges of integrating AI into a Zero Trust framework to enhance the effectiveness of Zero Trust security. We also developed education materials designed to equip students with a comprehensive understanding of the application of Zero Trust security in protecting CAVs. This project provides stakeholders with effective Zero Trust strategies for securing CAVs and their supporting infrastructure in an increasingly connected and threat-prone environment.			
17. Keywords Network security, Zero Trust Architecture		18. Distribution Statement No restrictions.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 25	22. Price N/A



TABLE OF CONTENTS

DISCLAIMER	2
CONTACTS	2
ACKNOWLEDGMENT.....	3
EXECUTIVE SUMMARY	6
CHAPTER 1	7
Introduction.....	7
CHAPTER 2	9
Literature Review.....	9
2.1 Components of Zero Trust Architectures	9
2.2 ZTA Enabling Technologies.....	11
2.3 Open Challenges and Future Trends.....	12
CHAPTER 3	14
Methods -- Zero Trust Architecture for CAVs	14
CHAPTER 4	17
Case Study and Evaluation	17
4.1 Excessive Speeding.....	17
4.2 Harsh Acceleration.....	17
4.3 Abrupt Lane Changes	18
4.4 Collision Scenarios	18
4.5 Evaluation	18
CHAPTER 5	20
Case Study and Evaluation	20
CHAPTER 6	22
Discussion.....	22
CHAPTER 7	23
Conclusions.....	23
REFERENCES	24



EXECUTIVE SUMMARY

Connected and autonomous vehicles (CAVs) are highly complex and interconnected systems that often involve many sensors, such as GPS, LiDAR (Light Detection and Ranging Technology), cameras, IMU (Inertial Measurement Unit), radars and ultrasonic sensors, to capture the environmental circumstances for improving vehicle safety, efficiency, and mobility on roadways. This reliance on multiple sensors makes CAV vulnerable to a variety of threats. As CAVs develop in scale, capability, and complexity, so do the vulnerabilities and the corresponding opportunities for malicious actors. These potential threats might target the physical systems or software of the vehicles themselves or the enabling hardware and software in the surrounding support infrastructure. Additionally, the data generated by the operation of the self-driving systems may expose personal and confidential information that could be misused by unauthorized parties, creating a threat to user privacy. Security and privacy considerations must be considered a foundational layer in CAVs.

Researchers have shown that various vehicle systems can be hacked to enable a remote takeover of the vehicle. Miller et al. used a Wi-Fi open port to hack the multimedia system of Jeep Cherokee and reprogram the firmware of ECUs. They were able to control a wide range of automotive functions (e.g., cutting brakes, shutting down engines, and driving off roads), triggering a recall of 1.4 million hackable vehicles. On the other hand, deep learning algorithms (such as perception modules for obstacle detection and traffic light detection) are increasingly integrated into CAVs that perform complex control functions under highly dynamic operating environments. Despite benefits from adopting deep learning techniques, it also extends the attack surface (e.g., adversarial machine learning) of the already vulnerable vehicle systems. Learning-based modules rely on the training data to make control decisions. However, the training data is usually incomplete and covers only a subset of possible scenarios that could be encountered operationally in CAVs, rendering design-time safety assessment and verification incomplete. As a result, runtime continuous verification and authentication are critically important for safety assurance in CAVs.

Zero Trust Architectures (ZTA) for networks have emerged as a fundamentally new way of approaching network security. It offers new paradigms for defining and enforcing policy through various means rooted in modeling trust relationships. The zero-trust security model does not automatically trust any user or device inside or outside the network perimeter. Instead, it enforces a set of policies (i.e., rules that are dynamically maintained and enforced) to verify and ensure the security of resources. It is a multi-layered approach with the idea of never trusting and always verifying every access to a resource. ZTA is gaining popularity in automotive industry. It can aid in reducing potential risks in CAVs by guaranteeing that only approved users and devices can access sensitive systems and data.

This research project investigates how ZTA can be adapted to CAVs to provide fundamental protection for individual components within CAV systems and their supporting infrastructure. This research provides stakeholders with effective Zero Trust strategies for securing CAVs and their supporting infrastructure in an increasingly connected and threat-prone environment.



CHAPTER 1

Introduction

As Connected and Autonomous Vehicle (CAV) research and adoption continue to grow in both academia and industry and CAVs are increasingly rely on cyber-based control, navigation, and communication, many security-related concerns arise [1]. The complexity and inter-connectedness of CAVs offer myriad opportunities for security compromise, potentially resulting in unsafe operation or leakage of confidential information about the user.

Researchers have shown that various vehicle systems can be hacked to enable a remote takeover of the vehicle [2, 3]. Miller et al. used a Wi-Fi open port to hack the multimedia system of Jeep Cherokee and reprogram the firmware of ECUs [4]. They were able be control a wide range of automotive functions (e.g., cutting brakes, shutting down engines, and driving off roads), triggering a recall of 1.4 million hackable vehicles. On the other hand, deep learning algorithms (such as perception modules for obstacle detection and traffic light detection) are increasingly integrated into CAVs that perform complex control functions under highly dynamic operating environments. Despite benefits from adopting deep learning techniques, it also extends the attack surface (e.g., adversarial machine learning) of the already vulnerable vehicle systems. Learning based modules rely on the training data to makes control decisions. However, the training data is usually incomplete and covers only a subset of possible scenarios that could be encountered operationally in CAVs, rendering design-time safety assessment and verification incomplete. As a result, runtime continuous verification and authentication are critically important for safety assurance in CAVs.

Zero Trust Architectures (ZTA) for networks have emerged as a fundamentally new way of approaching network security [5]. It offers new paradigms for defining and enforcing policy through various means rooted in modeling trust relationships. The zero-trust security model does not automatically trust any user or device inside or outside the network perimeter. Instead, it employs a set of access rules to verify user identity and ensure the security of resources. In 2020, The National Institute of Standards and Technology (NIST) outlines seven tenets of zero trust, which focus on resource perspective, communication, access control, and the enterprise perspective [6]. In the zero-trust view of a network, neither the private network, the devices on the network, nor resources can be inherently trusted or treated as the trust zone. The zero trust policy, in this case, refers to a set of access rules that are dynamically maintained and enforced using the Policy Decision Point (PDP) and Policy Enforcement Point (PEP). The resources, which include data sources and computing services, can only be trusted after processing with the PEP. The Department of Defense (DoD) further refined the zero-trust reference architecture in 2022 [7]. Buck et al. [8] summarized and structured the current literature on zero-trust research. The zero trust security model has been discussed in several fields and applied in various systems [9]. Chen et al. [10] presented a software-defined ZTA for 6G networks for establishing an elastic and scalable security regime. Sedjelmaci et al. [11] proposed a ZTA-empowered intrusion detection framework to protect edge computing from network attacks. Researchers in [12] applied the zero-trust model to smart medical information systems, where the environment pertains to situational security during access requests and processing.



National Center for Transportation Cybersecurity and Resiliency (TraCR)

The overarching objective of this project is to investigate how ZTA can be adapted to CAVs to provide fundamental protection for individual components within CAV systems and their supporting infrastructure. As Artificial Intelligence (AI) continues to become an integral part of cybersecurity, this project also explores integrating AI into a Zero Trust framework to enhance the effectiveness of Zero Trust security. This research provides stakeholders with effective Zero Trust strategies for securing CAVs and their supporting infrastructure in an increasingly connected and threat-prone environment. This report presents the outcomes of an extensive investigation how ZTA can be utilized to ensure the safe operation of CAVs.

This report is organized into five chapters. Chapter 1 introduces the research context, outlines the problem statement and defines the project objectives. Chapter 2 presents a literature survey of the ZTA. Chapter 3 introduces an AI-enhanced ZTA for CAVs. Chapter 4 presents a proof-of-concept prototype of the designed ZTA for CAV systems, and case studies that demonstrate the application of the zero-trust policy and the proposed enforcement architecture.



CHAPTER 2

Literature Review

This literature review is organized into three main sub-sections: (1) Key Components of Zero Trust Architectures (ZTA); (2) ZTA Enabling Technologies; and (3) Open Challenges and Future Trends.

2.1 Components of Zero Trust Architectures

In this section, we will explore the key components from the agreed ZTA model and use those as a foundation in describing the interesting technologies and challenges that have arisen through industry use of ZTA. Figure 1 presents the basic ZTA components and their interactions.

Threat Model. In ZTA, no component within or outside the vehicle is inherently trusted. The threat model assumes that adversaries may compromise Electronic Control Units (ECUs), inject malicious messages into in-vehicle networks such as the CAN bus, manipulate sensor data from cameras, LiDARs, radars, or GPS receivers, and interfere with actuator control pathways. However, we assume that the core ZTA security infrastructure itself is trustworthy and protected against compromise. For example, the ZTA components may be protected using hardware-assisted security technologies such as Trusted Execution Environments (TEEs) and secure enclaves. Under these assumptions, the ZTA infrastructure serves as a trusted foundation for continuously verifying the security posture of all other vehicle subsystems while minimizing the trusted computing base and reducing the attack surface available to adversaries.




Figure 1: High-level Zero Trust Architecture model illustrating the interaction between the control and data plane components, and external inputs, which together evaluate and translate policy rules for enforcement within the network.



At the core of the ZTA framework is the Policy Decision Point (PDP), which is tasked with evaluating access requests in real-time. The PDP is a composite function that comprises the Policy Engine (PE) and the Policy Administrator (PA), both of which play critical roles in its operation. The Policy Engine serves as the decision-making core of the PDP, leveraging predefined and dynamically updated policies to determine whether access should be granted or denied [6]. This component evaluates access requests against a set of rules and contextual factors, including user identity, device posture, location, and behavior analytics [13]. The outcomes are derived through rigorous application of access policies and may also incorporate risk scoring mechanisms to assess the trustworthiness of the request [14]. Complementing the Policy Engine is the Policy Administrator, which acts as the operational arm of the PDP. It is responsible for translating the decisions of the Policy Engine into actionable commands and ensuring their proper execution across the system. The Policy Administrator communicates directly with the Policy Enforcement Point (PEP) to apply the determined access permissions, revoke existing access where necessary, or implement additional security measures. Together, these two subcomponents ensure that the PDP operates effectively as the central authority for access control decisions in a Zero Trust framework.

Supporting the PDP is the Policy Enforcement Point (PEP), which acts as a gatekeeper by enforcing access decisions. The PEP ensures that only authorized entities can interact with the system's resources, thereby maintaining stringent access control. Another essential component is the Policy Information Point (PIP), which serves as a repository and aggregator of contextual data. This data includes attributes related to users, devices, and operational environments, which are critical for informed decision-making by the PDP. The PIP is the ideal place with the ZTA Control Plane to implement AI models to facilitate the aggregation and processing of vital data. Within this conceptual ZTA framework, resources are identified as the protected digital assets, applications, and services that constitute the system's operational scope. These resources are accessed by endpoints, which are devices or applications interfacing with the ZTA. Endpoint security and monitoring are integral to preventing unauthorized access and mitigating potential vulnerabilities.

To facilitate secure interactions, ZTA employs communication channels that ensure the confidentiality and integrity of data exchanged between components. These channels are underpinned by encryption protocols and continuous monitoring to detect and address potential breaches. Additionally, access policies represent a dynamic and granular set of rules that dictate permissible actions within the system. These policies adapt to changing contexts, ensuring that access decisions remain relevant and effective.

By integrating these components, the ZTA framework establishes a security model that minimizes implicit trust and proactively mitigates threats. NIST emphasizes the importance of leveraging advanced contextual analysis, such as behavioral analytics and device health assessments, to fortify the decision-making process. This approach not only enhances security but also ensures operational resilience in the face of evolving cyber-threats.

As-is evident throughout the architecture, the use of continuous checks vs. point-in-time during connection establishment is a vital principal of ZTA, and thus authentication and authorization are no different. Continuous authentication leverages dynamic and context-aware factors, such as



behavioral biometrics, device posture, geolocation, and patterns of user interaction, to continually assess the legitimacy of access requests [15]. Similarly, continuous authorization evaluates ongoing access conditions and adjusts permissions dynamically based on real-time contextual information [16]. This approach significantly reduces the risk of unauthorized access by addressing potential threats as they emerge, maintaining a robust and adaptive security posture.

2.2 ZTA Enabling Technologies

In practice, ZTA relies on a diverse set of enabling technologies to support its core principles, including continuous verification, least privilege access, and dynamic security policies. These technologies facilitate robust identity management, access control, network segmentation, and security analytics, ensuring that Zero Trust frameworks remain effective in mitigating modern cyber threats. This section explores key enabling technologies that underpin the practical deployment of ZTA across various environments.

2.2.1 Identity and Access Management

Identity and Access Management (IAM) is fundamental to ZTA, as it ensures that only authenticated and authorized users and devices gain access to protected resources. Modern IAM solutions incorporate Multi-Factor Authentication (MFA), Single Sign-On (SSO), and federated identity management to enhance authentication security [17]. Attribute-Based Access Control (ABAC) and Role-Based Access Control (RBAC) mechanisms dynamically enforce access policies based on user roles, device attributes, and real-time contextual information. Advanced IAM frameworks leverage continuous authentication, ensuring that access remains restricted even after initial verification [18].

2.2.2 Endpoint Protection

Endpoint Detection and Response (EDR) and Extended Detection and Response (XDR) solutions play a crucial role in continuously monitoring and securing endpoints within a Zero Trust framework [8, 9, 15]. EDR systems detect and respond to anomalous behavior at the device level, mitigating risks associated with malware, unauthorized access, and data exfiltration [16]. By continuously assessing device health and enforcing compliance policies, EDR solutions reinforce ZTA's zero trust posture. XDR extends this capability by integrating telemetry from multiple security sources, e.g., network traffic, cloud services, and user behavior analytics, to provide a comprehensive threat detection and incident response mechanism [19].

2.2.3 Micro-Segmentation

Micro-segmentation enhances security by dividing networks into smaller, isolated segments, ensuring that lateral movement is restricted even in the event of a breach. This approach can be implemented using Software Defined Networking (SDN), which dynamically enforces security policies based on real-time network activity [20]. SDN-based micro-segmentation allows organizations to apply fine-grained access controls, reducing attack surfaces and containing potential threats within defined security zones [15].



2.2.4 Edge Network Access

Secure Access Service Edge (SASE) integrates networking and security services into a unified cloud-based model, enabling organizations to enforce consistent security policies across distributed environments [21]. SASE represents a convergence of wide-area networking (WAN) and network security services—such as firewall-as-a-service (FWaaS), secure web gateways (SWG), cloud access security brokers (CASB), and Zero Trust Network Access (ZTNA)—into a single, cloud-managed platform [22]. This convergence reduces complexity, improves scalability, and supports mobile and remote users by bringing security enforcement closer to the point of access. A core component of SASE, ZTNA replaces traditional VPNs by verifying user and device identities at each access request, ensuring that trust is never implicitly granted [23]. SASE combines capabilities such as secure web gateways, ZTNA, and data loss prevention (DLP) to ensure consistent policies across access points. ZTNA solutions provide encrypted, least-privilege access to applications, minimizing exposure to potential attackers [24].

2.2.5 Encryption and Secure Communication Protocols

Encryption is a cornerstone of ZTA, ensuring the confidentiality and integrity of data both in transit and at rest [25]. Transport Layer Security (TLS), Secure Shell (SSH), and end-to-end encryption mechanisms prevent unauthorized interception of sensitive information [26]. Secure communication protocols, such as IPsec and DNS-over-HTTPS (DoH), further strengthen Zero Trust implementations by securing data exchange across networks [27].

2.2.6 Threat Intelligence, AI, and Automated Response

Threat intelligence and artificial intelligence (AI) technologies play a critical role in enhancing ZTA by enabling real-time detection, prediction, and response to evolving threats. Threat intelligence platforms leverage machine learning (ML) and AI to aggregate and analyze security event data from various sources, supporting proactive threat hunting and dynamic risk assessment. Behavioral analytics and anomaly detection continuously evaluate user activities, flagging deviations from established baselines to inform access decisions and threat responses [22]. AI-driven security solutions further extend these capabilities by automating threat detection and response processes [26]. Algorithms analyze behavioral patterns, device telemetry, and network traffic to identify and contain threats in real-time [19]. Automated mechanisms, such as Security Orchestration, Automation, and Response (SOAR) platforms, facilitate rapid and coordinated mitigation of incidents without the need for human intervention [27]. AI also enables predictive analytics, improving the accuracy and contextual relevance of security policies while enhancing the overall resilience of ZTA systems.

2.3 Open Challenges and Future Trends

2.3.1 Usability and Privacy

Despite growing adoption, the practical realization of ZTA continues to face several unresolved challenges. One of the primary difficulties lies in balancing stringent security policies with usability. Overly aggressive enforcement of least privilege or continuous authentication can



degrade user experience and operational efficiency. This trade-off necessitates research into adaptive trust mechanisms that can dynamically calibrate access controls based on risk context. ZTA continuously monitors users' requests and analyzes their behaviors, posing an increased risk to users' privacy and data protection [28, 29]. Zero-trust solutions may collect users' sensitive personal information for user authentication, e.g., emails and passwords, or for further tasks like anomaly detection. However, limited work has been conducted on the privacy compliance of zero-trust systems. Researchers have identified privacy concerns in remote work and Bring Your Own Device (BYOD) scenarios, where participants reported feeling their privacy was being invaded [30, 31]. In a previous study [30], 14.8% of participants responded "yes" and 37% answered "sometimes" feeling their privacy is invaded because of active BYOD security measures. The perception of privacy invasion may stem from uncertainty about the extent of the monitoring policies or a sense that the company does not trust them. Therefore, it is important to address employees' privacy concerns and improve the trust and transparency regarding the data disclosure between employees and the company.

2.3.2 Scalability and Performance

Scalability and performance constraints also pose significant concerns. Zero Trust demands pervasive telemetry collection, real-time policy evaluation, and frequent identity verification, which can strain network and processing resources, particularly in high-throughput environments [8]. Efficient architectures and lightweight cryptographic primitives are essential to maintain performance while preserving security guarantees. The advent of edge computing and 6G networks introduces further complexities, as decentralized and ephemeral endpoints require continuous verification in resource-constrained environments [32]. Emerging solutions such as blockchain, decentralized identifiers, and digital twins offer promise, but their integration into Zero Trust ecosystems remains an open area of exploration [33], [34].

CHAPTER 3

Methods -- Zero Trust Architecture for CAVs

3.1.1 Introduction

In this chapter, we explain the concept of our ZTA for CAVs through a description of the high-level policy constructs that dictate the communication channels and the event processing components of CAVs. We then propose a reference architecture that is used to enforce these policy constructs by implementing and adhering to ZTA principles. Our ZTA is designed to be applied to different automotive electronic/electrical architectures, from the central gateway and domain architecture to newer zonal architectures but has a focus on CAN bus implementations.

3.1.2 Abstract View of the Enforcement Architecture

Our ZTA for CAV architecture builds upon NIST's ZTA concepts and serves to clarify specific roles and how they are to be implemented in internal autonomous vehicle environments. To this end, we make use of standard components, including Policy Enforcement Points (PEP), Policy Decision Points (PDP), Policy Administrator (PA), and Policy Engine (PE), and the NIST ZTA constructs of subject and resource. In our architecture, we implement the four central components--Policy Engine, Policy Administrator, Policy Decision Point, and Context Manager--on a single system and collectively refer to this as the Zero Trust Controller or just the Controller, as illustrated in Figure 2.




Figure 2: High-level design showing major architectural components in our ZTA for CAV



The first significant component of our architecture is the Context Manager (CM), tasked with monitoring the vehicle's CAN bus network and providing what amounts to flow information to the controller. Because the CAN bus communication protocol relies on broadcast messages across the entire bus medium, the CM can easily access all communication. Still, it also means the CM must constantly filter out the flows that are not of interest. This is a function of parsing the policy to make these decisions. The CM's role is to process the communication flow information to maintain each resource's current and accurate state (sensor, ECU, etc.). We previously discussed that continuous authentication & authorization is a critical aspect of our solution, and the CM implements this functionality. The CM is responsible for polling resources regularly to update open connections' authentication state. This state is fed into the PE, which may act upon changes in authorization in real time. We refer to these regular polling intervals as the continuous context check. These context checks are also an integral part of the overall ZTA within the CAN bus architecture.

The PA's role is to interpret the policy intent and dictate actions in the system through the PE. While it is still true that the nature of the policy in the automotive world remains primarily stagnant, the future will undoubtedly hold a greater frequency of software updates. For this, the PA remains crucial in adapting to changing policy requirements. On the other hand, the PDP has the job of taking the intent from the PA, plus current environmental information from the CM, to make policy decisions that are to be enforced in the system by way of the PE.

The PDP is responsible for inspecting all current state information provided by the CM. In a given flow inspection, the PE will generate the enforcement intent for the PEPs, and then these PEPs will act upon that intent. The CAN bus nodes contain the PEPs in our architecture and receive dynamic policy updates from the PDP through the CAN network. This enforcement action is for the purpose of establishing or closing a flow between resources. From an architectural standpoint, the PEP is an entity between the resource (sensor) and its connection to the CAN bus. Thereby each resource effectively implements its own PEP, but it is vital to consider the separation of duties in that the PEP is a part of the ZTA, which exists to protect the resource.

When a sensor resource wants to access another sensor or ECU in the network, it must request permission from the PDP, for example, a proximity sensor wishing to interface with a parking-assist ECU. Therefore, a request will be made from the PEP and sent to the PDP, depicted by ① in Figure 2. The PDP will attempt a policy match with input from the most recent subject and device context from the CM. Unlike the routine device continuous context check, this flow initiation request will include the resource to access. This is depicted as ② in Figure 2.

When a policy match occurs that results in an allowed flow, the PE will compile the policy specification down to PEP-specific syntax, which is pushed to the PEPs involved in the flow, as shown by ③ in Figure 2. Figure 3 shows this overall workflow. A valid flow remains active in the CAN network for as long as it is needed or dictated by policy. The PDP does, however, retain the critical right to revoke a flow at any point in time if the dynamic context of the subject or device changes in a negative way. The highlighted portion of Figure 3 shows how the PDP takes this action upon a continuous context check.




Figure 3: Sequence diagram showing the steps and interactions required for the ECU/sensor to open a communication flow with another resource on the CAN bus

All the communications within the network among the resources, CM, PDP, PE, and PEP, are encrypted to ensure the integrity of the message. The encryption and decryption keys are assigned and stored before the resource connects to the CAN bus network. The keys can be generated and configured when registering the resource using standard TLS during the manufacturing stage of the vehicle.

3.1.3 Continuous Authentication and Authorization

Our design conforms to the abstract architecture specified in [6]. Each time a resource wants to access the network, it must communicate with the controller that implements the PDP to verify its identity and authorize its permission to acquire access. This is achieved through a control protocol over the CAN bus. The access control workflow in our system is shown in Figure 3. After the sensor/resource is registered and connected to the network, its embedded PEP component will continuously report profile metadata about the resource to the CM. As stated previously, the PEP agent is also responsible for sending a flow request to the CM for evaluation when the resource wishes to access the network. The PDP will evaluate this request and either allow or deny the flow. In the case of an allowed flow, beyond the PEP involvement already covered, our system implements continuous authentication and authorization for the duration of the flow. Because the agent will periodically report the state, the CM and PDP can reevaluate the eligibility of the resource to access its requested peer. Indeed, the CM may also take an external context, for example, from a Security Information Event Manager (SIEM), which may inform the CM that a resource is compromised in some meaningful way. The CM aggregates all this information and makes it available to the PDP for evaluation. At each interval, the PDP may decide that the resource now violates the previously allowed policy. In such an event, the PDP would instruct the PE to revoke the flow policy from the PEPs, effectively disconnecting the resource from the network. By understanding this context, our architecture would be able to react to this dynamic authorization change and act, accordingly, revoking any active flows.



CHAPTER 4

Case Study and Evaluation

We developed and implemented a proof-of-concept prototype of the designed ZTA for CAV systems. We conducted case studies that demonstrate the application of the zero-trust policy and the proposed enforcement architecture. This chapter presents both the case study of implementing our ZTA in CAV and the evaluation of our prototype.

4.1 Excessive Speeding

Excessive speeding is operating a vehicle well above designated limits, posing risks to the driver and others. Addressing it includes enforcing speed limits, using monitoring technologies, and raising awareness. To detect and address the risks associated with excessive speeding, the CM continuously monitors the speed data of the vehicle. This data is retrieved from the CAN through the ECU. To effectively manage speeding risks, we have defined a ZT policy that establishes a speed threshold. The policy checks whether subsequent CAN data indicates an acceleration-induced increase in speed, potentially leading to an unsafe condition. If the CM detects this pattern in the profiles of the upcoming time windows, it triggers an event transmitted to the PDP. Subsequently, the PDP handles the event by determining the appropriate action to be taken and dispatches the decision to the PEP. The PEP then modifies the functionality related to accelerator input, preventing any further acceleration and mitigating the risk of continued speeding. Implementing this mechanism will effectively identify and address excessive speeding instances in real-time, thereby enhancing safety and reducing potential hazards on the road.

4.2 Harsh Acceleration

Harsh acceleration refers to abrupt and forceful changes in a vehicle's speed, often associated with aggressive driving. Detecting harsh acceleration provides insights into driving habits and promotes safer practices, reducing accident risks and enhancing road safety. Our PDP executes the ZT policy to address this concern, specifically designed to monitor acceleration patterns. The PDP utilizes the CM profiles, including information about the vehicle's acceleration behavior over a specific time window. By examining these profiles, the PDP checks for instances of frequent full-speed acceleration signals. If the frequency of such acceleration signals surpasses a predefined threshold, the PDP initiates an action decision. This decision is then transmitted to the PEP, which is responsible for enforcing the necessary measures to mitigate excessive acceleration. The PEP can implement actions such as limiting certain acceleration actions or applying braking actions to counterbalance the excessive speed gained through acceleration. Through the coordinated effort of the PDP and PEP, it proactively identifies and addresses instances of excessive acceleration, promoting safer driving conditions within city environments. It aims to balance vehicle performance and compliance with recommended driving practices by enforcing appropriate actions.



4.3 Abrupt Lane Changes

Abrupt lane changes involve sudden, unanticipated maneuvers where a vehicle shifts lanes without warning or signaling. Understanding the risks and promoting responsible driving practices contribute to a safer traffic environment. To mitigate the risks of lane changes, we incorporate a ZT policy specifically designed to detect and respond to these unsafe lane change behaviors. The ZT policy examines the profiles obtained from the CM within a short time window, focusing on instances where the vehicle undergoes frequent steering changes without activating the corresponding turn indicators at certain speeds. If the ZT policy identifies such patterns within the returned profiles, the PDP triggers an action decision. This decision is then communicated to the PEP, which takes necessary actions to address the issue. The PEP's actions will include slowing down the vehicle and ensuring that appropriate turn indicators are activated during left or right steering operations. Additionally, the system generates an alert to remind the operator about the importance of avoiding frequent abrupt lane changes without signaling. Implementing these measures actively promotes safer driving practices, reducing the risk of accidents caused by sudden lane changes and encouraging the use of turn indicators as a crucial aspect of responsible driving.

4.4 Collision Scenarios

The collision scenarios replicate real-world conditions, incorporating various elements like speeds, weather, road layouts, traffic patterns, and vehicle types. By observing vehicle behavior, including braking, acceleration, and steering, we can gain insights into their impact on occupant safety, vehicle damage, and surrounding traffic conditions. Another straightforward detection will be focusing on the critical data from the CAN bus that serves as indicators of collisions within a vehicle. One prominent indicator is the activation of the airbag system. When the CM identifies these collision indicators within the CAN bus data, it promptly triggers an alert to the PDP. Upon receiving the alert, the PDP thoroughly assesses the vehicle's current speed and whether the operator is still applying acceleration. These evaluations are crucial in identifying potentially unsafe conditions. If the PDP detects either of these unsafe conditions, it initiates an action decision and transmits it to the PEP. Subsequently, the PEP executes the action decision by engaging the braking system, gradually reducing the speed of the vehicle until it comes to a complete stop. This intervention aims to mitigate the consequences of the collision and ensure the safety of the vehicle's occupants. This proactive collision detection and response mechanism enables timely intervention in critical situations, effectively minimizing collisions' potential impact and severity.

4.5 Evaluation

Our prototype implementation is based on the design illustrated in Figure 2. To ensure secure and reliable operations, we utilized an open-source project to simulate the CAN bus data [37]. Moreover, we adopted the ZT policy format proposed in [38]. With our prototype, we could conduct realistic simulations using simulated vehicle CAN data. This allows us to assess the effectiveness and performance of our system within a controlled environment. In our prototype, the CM retrieves CAN bus data at five-second intervals. The PDP applies predefined ZT policies to examine whether these CAN bus data contain any unsafe violations.



```
(1686285873.624312) vcan0 100#5A7436616877C53A
(1686285873.824552) vcan0 100#98AA085185718221
(1686285874.024673) vcan0 19B#000000E000000
(1686285874.224915) vcan0 100#DA25545FA9D09C
(1686285874.425125) vcan0 100#
(1686285874.625360) vcan0 100#15CC373ADE42A53C
(1686285874.825600) vcan0 100#9C
(1686285875.025867) vcan0 100#
(1686285875.225998) vcan0 100#7A
(1686285875.426256) vcan0 100#4E48
(1686285875.626533) vcan0 244#0000003800
(1686285875.826812) vcan0 100#34336C30EB
(1686285876.027047) vcan0 100#EEE40E16FD16307E
```

Figure 4: Security Violation Example.

For example, in Figure 3, apart from the timestamp and CAN interface, the red-boxed lines “19B#000000E000000” and “244#0000003800” represent CAN data indicating unlocking the left front door and immediate acceleration. The sequence of unlocking the door, followed by stepping on the accelerator, may pose a potential safety hazard for the passengers in the vehicle. Consequently, it violates the defined ZT policy, triggering the activation of the PEP. The PEP then takes appropriate action, either removing the acceleration command or adding an action to lock the door. We conducted the evaluation of our ZT prototype in CAV systems. Considering the diverse range of real-world CAV systems, each with its unique hardware capabilities, we focused our evaluation on quantifying the additional time overhead incurred by implementing our ZT solution within a simulated environment.

To assess our prototype’s effectiveness, the CAN bus simulator ran on a desktop equipped with 12th Gen Intel® Core™ i5-12400 CPU with the Ubuntu operating system 22.04. We utilized the simulator in [37] and generated CAN bus data using the open-source project [39]. Within our evaluation process, we examined the CAN bus data every 5 seconds, which contained a total of 455 frames. During this 5-second time window, the PDP executed 50 conditional checks and then let PEP do the correction action. The evaluation results indicated an additional delay of an average of 0.155 milliseconds for executing the ZTA interactions when incorporating the ZT framework into CAV systems. Given the 5-second time window, the additional time overhead incurred can be considered negligible.



CHAPTER 5

Case Study and Evaluation

With the focus of our work being on internal vehicle networks (IVN) and specifically CAN bus networks, the threat vectors of interest mainly revolved around access to that bus. In the contemporary model, CAN buses are insecure because the protocol itself is a broadcast medium, meaning any participating node on the bus can communicate with and listen to any other node. This simplicity offers practical applications in automotive use cases but the security implications are clear and further exacerbated by the total lack of authentication or authorization mechanisms or even integrity checks in the fundamental CAN protocol. The authors in [40] defined three types of attacks in this category: Fabrication, Suspension, and Masquerade attacks. Fabrication attacks involve forged communication on the CAN bus by manipulating CAN protocol header messages. On their own, Fabrication attacks seek to confuse the CAN network by broadcasting messages at higher frequencies and in some cases, simply amount to Denial of Service (DoS) style activity. Fundamentally though, fabrication means simply that an actor is able to broadcast a valid message on the CAN bus. Suspension attacks focus on taking an individual ECU node offline from the network. This usually involves an actor manipulating the ECU directly. Finally, Masquerade attacks combine the aspects of the other two types by taking a legitimate ECU node offline and replacing it with a nefarious actor.

Our ZTA for CAV design is able to address these types of IVN attacks in several specific ways. First, because each ECU node requires registration to the ZTA controller, it is very easy for the CM to detect unauthorized usage of the CAN bus, through its continuous context checks. The flow request model between the PEPs and the controller also means that each valid flow on the CAN bus is authorized at all times. This means from the CM it is feasible to flag unauthorized flows based on its understanding of the current network state and the visibility of the bus due to its broadcast nature. Communication is encrypted in our architecture using TLS, and because of the ECU registration, the controller is able to identify and discriminate nodes that have attached to the bus but were not a part of the original network. This means CAN protocol messages cannot be easily spoofed and the encryption serves to provide confidentiality and help with the integrity of message delivery.

One potential major open question in our architecture is addressing DoS attack vectors. The broadcast nature of the CAN bus means that DoS attacks have large outcomes for minimal effort. In such cases, an attacker needs only clog the CAN bus with garbage communication, consuming the entire bus media bandwidth. Because the bus is the only communication channel, the inability to process control or response messages to counteract the DoS attack means the network is left entirely unusable. CAN bus networks typically support between 1Mbps and 5Mbps throughput so it is relatively easy for attackers to saturate the available bandwidth. Our architecture would allow the individual ECUs to continue their independent operation, but this means the CAV would likely need to cease operation to ensure the safety of its occupants.

Next, remote attacks appeared, mostly focused on the connected nature of autonomous vehicles through their interaction with cellular and other networks. Notably, these attack vectors are generally covered in our architecture through existing means because of our reliance on pre-



registration and defined policy flows. Even if an attacker is able to compromise the individual Telematic Control Units (TCU), which communicates with the outside world, its embedded PEP will still enforce the policy intent of the controller. The current generation of attack vectors sees a focus on ancillary applications and connectivity software such as phone and vehicle entertainment system applications. This is where our architecture's dynamic policy adaptation aspects come into play. As new and emerging threats come about in this area, the policy intent of the IVN can be kept up to date to keep a clamp on authorized flows. This could be as simple as enforcing that there is no need for a passenger's tablet to connect to an ECU that controls the engine's fuel flow.

There are also, of course, many attack vectors that are specific to different types of ECU or sensors. These can include Location (GPS), LiDAR, Perception (Computer Vision), or Ultrasonic attacks, just to name a few. While our architecture strives to solve the general sorts of security issues that might arise with any given sensor it is important to highlight the need to understand specific attack scenarios. Returning to the point about CAN bus segmentation and operability of the vehicle, if a GPS sensor were to be compromised, it is theoretically possible for instance that the CAV may be able to continue operation for some period of time, whereas if a perception ECU were compromised, perhaps that means to doom the CAV to an immediate halt. In both cases, our ZTA would be able to isolate the compromised node to prevent further degradation of the CAN bus. Perhaps with enough redundancy in the sensor architecture, the CAV would be able to survive such one-off attacks with the advent of our design.



CHAPTER 6

Discussion

The integration of Artificial Intelligence (AI) into ZTA represents a transformative advancement in cybersecurity. AI's ability to support real-time threat detection, behavioral analytics, and adaptive access control enhances the foundational principles of ZTA, including continuous verification and least privilege. By processing high-volume telemetry data across identity, device, network, and application layers, AI systems enable organizations to proactively identify anomalous behavior and automate policy enforcement. These capabilities are especially valuable in dynamic environments such as multi-cloud infrastructures, where traditional perimeter defenses are ineffective. Machine learning models embedded within ZTA frameworks can analyze historical and contextual access patterns to make predictive inferences, thereby improving the speed and precision of threat mitigation.

Despite these benefits, the incorporation of AI into ZTA frameworks introduces significant challenges. One of the foremost concerns is the vulnerability of AI models to adversarial inputs. Attackers can craft malicious data to deceive classification models, bypassing security controls or generating false positives, thus undermining the trustworthiness of automated decisions. Moreover, the extensive data collection and monitoring required to train and operate AI systems can raise privacy and compliance concerns, especially when sensitive personal or behavioral information is involved. Ensuring robust data anonymization, encryption, and governance is critical to aligning AI-enhanced ZTA with privacy frameworks and regulatory mandates.

Another major obstacle is the lack of transparency and explainability in many AI models, particularly those based on deep learning. These "black box" systems complicate the task of auditing decisions, understanding incident causality, and ensuring accountability within automated access control systems. This challenge has prompted growing interest in explainable AI (XAI) approaches that provide interpretable outputs suitable for operational security environments. Furthermore, the integration of AI capabilities into existing ZTA infrastructure presents non-trivial technical hurdles. Legacy systems may lack the compatibility or processing overhead to support real-time AI inference, necessitating costly upgrades or rearchitecting efforts.

To fully realize the potential of AI in Zero Trust contexts, organizations must balance innovation with caution. Strategic AI governance, responsible data stewardship, and interdisciplinary collaboration among cybersecurity professionals, data scientists, and policy stakeholders are essential to operationalizing these technologies effectively. While AI substantially augments the resilience and responsiveness of ZTA systems, its success depends on thoughtful design and sustained oversight.



CHAPTER 7

Conclusions

In this project, we show how a ZTA can be adapted for use in CAV security relative to local communication on the CAN bus network. Through an analysis of existing and pertinent attacks on CAN networks we proposed solutions that can be implemented as a part of our overall architecture. These include attacks that focus on the CAN bus' communication paths directly, and the nodes participating in the network. Despite its promise, ZTA adoption remains hindered by practical challenges such as integration with legacy systems, scalability constraints, user experience trade-offs, and the complexity of continuous verification in decentralized environments. These limitations underscore the need for further research into hybrid deployment models, lightweight enforcement mechanisms, and resilient architectures tailored to domain-specific requirements.



REFERENCES

- [1] C. Young, J. Zambreno, H. Olufowobi and G. Bloom, "Survey of Automotive Controller Area Network Intrusion Detection Systems," in *IEEE Design & Test*, vol. 36, no. 6, pp. 48-55, 2019.
- [2] Stephen Checkoway, Damon McCoy, Brian Kantor, Danny Anderson, Hovav Shacham, Stefan Savage, Karl Koscher, Alexei Czeskis, Franziska Roesner, and Tadayoshi Kohno. Comprehensive experimental analyses of automotive attack surfaces. In *Proceedings of the 20th USENIX Conference on Security, SEC'11*, page 6, USA, 2011. USENIX Association.
- [3] J. Petit and S. E. Shladover. Potential cyberattacks on automated vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 16(2):546–556, 2015.
- [4] Raul Quinonez, Jairo Giraldo, Luis Salazar, Erick Bauman, Alvaro Cardenas, and Zhiqiang Lin. SAVIOR: Securing autonomous vehicles with robust physical invariants. In *29th USENIX Security Symposium*, 2020.
- [5] Warner, Jeannie, "What is Zero Trust Security?", <https://www.crowdstrike.com/cybersecurity-101/zero-trust-security/>, 2021.
- [6] Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly. Zero trust architecture. Technical report, National Institute of Standards and Technology, 2020.
- [7] Defense Information Systems Agency (DISA) and National Security Agency (NSA) Zero Trust Engineering Team. Department of defense (DoD) zero trust reference architecture. 2022.
- [8] Christoph Buck, Christian Olenberger, Andr  Le Schweizer, Fabiane V  Nolter, and Torsten Eymann. Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust. *Computers & Security*, 110:102436, 2021.
- [9] Songpon Teerakanok, Tetsutaro Uehara, and Atsuo Inomata. Migrating to zero trust architecture: Reviews and challenges. *Security and Communication Networks*, 2021.
- [10] Xu Chen, Wei Feng, Ning Ge, and Yan Zhang. Zero trust architecture for 6g security, in *IEEE Network*, vol. 38, no. 4, pp. 224-232, July 2024.
- [11] Hichem Sedjelmaci and Nirwan Ansari. Zero framework to secure 6g edge computing. *IEEE Network*, pages 1–13, 2023.
- [12] Baozhan Chen, Siyuan Qiao, Jie Zhao, Dongqing Liu, Xiaobing Shi, Minzhao Lyu, Haotian Chen, Huimin Lu, and Yunkai Zhai. A security awareness and protection system for 5g smart healthcare based on zerotrust architecture. *IEEE Internet of Things Journal*, 2020.
- [13] R. Chandramouli and Z. Butcher. A zero-trust architecture model for access control in cloud-native applications in multi-location environments. 2023.
- [14] U.H. Park, J.H. Hong, A. Kim, and K.H. Son. Endpoint device risk-scoring algorithm proposal for zero trust. *Electronics*, 12(8):1906, 2023.
- [15] N.F. Syed, S.W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss. Zero trust architecture (ZTA): A comprehensive survey. *IEEE Access*, 10:57143–57179, 2022.
- [16] A. Alagappan, S.K. Venkatachary, and L.J.B. Andrews. Augmenting zero trust network architecture to enhance security in virtual power plants. *Energy Reports*, 8:1309–1320, 2022.
- [17] John Kindervag et al. Build security into your network's DNA: The zero trust network architecture. Forrester Research Inc, 27, 2010.
- [18] R. Chandramouli and Z. Butcher. A zero-trust architecture model for access control in cloud-native applications in multi-location environments, 2023.
- [19] H. Kang, G. Liu, Q. Wang, L. Meng, and J. Liu. Theory and application of zero trust security: A brief survey. *Entropy*, 25(12):1595, 2023.
- [20] Y. He, D. Huang, L. Chen, Y. Ni, and X. Ma. A survey on zero trust architecture: Challenges and future trends. *Wireless Communications and Mobile Computing*, 2022.
- [21] Gartner Secure access service edge (sase),2023.Accessed:2025-05-12.



- [22] P. Phiayura and S. Teerakanok. A comprehensive framework for migrating to zero trust architecture. *IEEE Access*, 11:19487–19511, 2023.
- [23] S. Li, M. Iqbal, and N. Saxena. Future industry internet of things with zero-trust security. *Information Systems Frontiers*, pages 1–14, 2022.
- [24] L. Bobelin. Zero trust in the context of iot: Industrial literature review, trends, and challenges. In *Proceedings of the 30th C&ESAR*, page 37, 2023.
- [25] D. Marikyan, S. Papagiannidis, and E. Alamanos. A systematic review of the smart home literature: A user perspective. *Technological Forecasting and Social Change*, 138:139–154, 2019.
- [26] B. Karabacak and T. Whittaker. Zero trust and advanced persistent threats: Who will win the war? In *International Conference on Cyber Warfare and Security*, volume 17, pages 92–101, March 2022.
- [27] C. Shepherd. Zero trust architecture: Framework and case study. In *International Conference on Cyber Warfare and Security*, volume 17, pages 92–101, March 2022.
- [28] Aizhan Tursunbayeva, Claudia Pagliari, Stefano Di Lauro, and Gilda Antonelli. The ethics of people analytics: risks, opportunities and recommendations. *Personnel Review*, 51(3):900–921, 2022.
- [29] Jason RC Nurse, Nikki Williams, Emily Collins, Niki Panteli, John Blythe, and Ben Koppelman. Remote working pre-and post-covid-19: an analysis of new threats and risks to security and privacy. In *HCI International 2021-Posters: 23rd HCI International Conference, HCII 2021, Virtual Event, July 24–29, 2021, Proceedings, Part III* 23, pages 583–590. Springer, 2021.
- [30] Kathleen Downer and Maumita Bhattacharya. Byod security: A study of human dimensions. In *Informatics*, volume 9, page 16. MDPI, 2022.
- [31] Aljuaid Turkea Ayedh M, Ainuddin Wahid Abdul Wahab, and Mohd Yamani Idna Idris. Systematic literature review on security access control policies and techniques based on privacy requirements in a byod environment: State of the art and future directions. *Applied Sciences*, 13(14):8048, 2023.
- [32] K. Ramezanpour and J. Jagannath. Intelligent zero trust architecture for 5g/6g networks: Principles, challenges, and the role of machine learning in the context of o-ran. *Computer Networks*, 217:109358, 2022.
- [33] K. Wu, R. Cheng, H. Xu, and J. Tong. Design and implementation of the zero trust model in the power internet of things. *International Transactions on Electrical Energy Systems*, 2023.
- [34] I.A. Ridhawi, S. Otoum, and M. Aloqaily. Decentralized zero-trust framework for digital twin-based 6g. *arXiv preprint arXiv:2302.03107*, 2023.
- [35] Liangkai Liu, Sidi Lu, Ren Zhong, Baofu Wu, Yongtao Yao, Qingyang Zhang, and Weisong Shi. Computing systems for autonomous driving: State of the art and challenges. *IEEE Internet of Things Journal*, 8(8):6469–6486, 2020.
- [36] Cloud Security Alliance. How is ai strengthening zero trust?, 2025. Accessed: 2025-05-13.
- [37] Instrument Cluster Simulator for SocketCAN. <https://github.com/zombieCraig/ICSim>, 2020.
- [38] John Anderson, Qiqing Huang, Long Cheng, and Hongxin Hu. Byoz: Protecting Boyd through zero trust network security. In *2022 IEEE International Conference on Networking, Architecture and Storage (NAS)*, pages 1–8. IEEE, 2022.
- [39] CAN-utils. <https://github.com/linux-can/can-utils>, 2022.
- [40] Mert D. Pese, Jay W. Schauer, Junhui Li, and Kang G. Shin. S2-can: Sufficiently secure controller area network. In *Annual Computer Security Applications Conference*, page 425–438, 2021.