



Article

ZSM Framework for Autonomous Security Service Level Agreement Life-Cycle Management in B5G Networks

Rodrigo Asensio-Garriga ¹, Alejandro Molina Zarca ^{2,*}, Jordi Ortiz ², Ana Hermosilla ¹, Hugo Ramón Pascual ³, Antonio Pastor ³ and Antonio Skarmeta ¹

¹ Department of Information and Communications Engineering, University of Murcia, 30100 Murcia, Spain; rodrigo.asensio@um.es (R.A.-G.); ana.hermosilla@um.es (A.H.); skarmeta@um.es (A.S.)

² Spanish Air Force Academy, University Center of Defense, 30729 San Javier, Spain; jordi.ortiz@cud.upct.es

³ Telefónica I + D, 28013 Madrid, Spain; hugo.ramonpascual@telefonica.com (H.R.P.); antonio.pastorperales@telefonica.com (A.P.)

* Correspondence: alejandro.mzarca@cud.upct.es

Abstract: In the rapidly evolving landscape of telecommunications, the integration of commercial 5G solutions and the rise of edge computing have reshaped service delivery, emphasizing the customization of requirements through network slices. However, the heterogeneity of devices and technologies in 5G and beyond networks poses significant challenges, particularly in terms of security management. Addressing this complexity, our work adopts the Zero-touch network and Service Management (ZSM) reference architecture to enable end-to-end automation of security and service management in Beyond 5G networks. This paper introduces the ZSM-based framework, which harnesses software-defined networking, network function virtualization, end-to-end slicing, and orchestration paradigms to autonomously enforce and preserve security service level agreements (SSLAs) across multiple domains that make up a 5G network. The framework autonomously manages end-to-end security slices through intent-driven closed loops at various logical levels, ensuring compliance with ETSI end-to-end network slice management standards for 5G communication services. The paper elaborates with an SLA-triggered use case comprising two phases: proactive, wherein the framework deploys and configures an end-to-end security slice tailored to the security service level agreement specifications, and reactive, where machine learning-trained security mechanisms autonomously detect and mitigate novel beyond 5G attacks exploiting open-sourced 5G core threat vectors. Finally, the results of the implementation and validation are presented, demonstrating the practical application of this research. Interestingly, these research results have been integrated into the ETSI ZSM Proof of Concept #6: ‘Security SLA Assurance in 5G Network Slices’, highlighting the relevance and impact of the study in the real world.

Keywords: ZSM; B5G; security; automation; E2E network slicing; multi-domain; MANO



Academic Editor: Carlo Blundo

Received: 10 January 2025

Revised: 7 February 2025

Accepted: 8 February 2025

Published: 12 February 2025

Citation: Asensio-Garriga, R.; Molina Zarca, A.; Ortiz, J.; Hermosilla, A.; Pascual, H.R.; Pastor, A.; Skarmeta, A. ZSM Framework for Autonomous Security Service Level Agreement Life-Cycle Management in B5G Networks. *Future Internet* **2025**, *17*, 86. <https://doi.org/10.3390/fi17020086>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Fifth generation networks and beyond come with the ability to integrate a very wide range of technologies and devices of all kinds, forming a heterogeneous network whose complexity continues to increase [1]. The variety of requirements brought by these architectures is tightly coupled with complexity and hardware-specific constraints, which make non-automated End-to-End (E2E) service deployment and management an extremely costly activity in terms of time, security, and expenses [2]. In addition, E2E services logically interconnect different functional domains and fall into three categories: Ultra Reliable Low

Latency Communications (URLLC), massive Machine Type Communications (mMTC), and enhanced Mobile Broadband (eMMB). Thus, connectivity requirements differ significantly depending on the service type and device capabilities. In this context, ensuring system security and robustness in the presence of network heterogeneity and variability becomes a particularly challenging task [3]. To address this task, research on novel technologies is mainly focused on enabling the softwarization of the network through the virtualization and programmability of its components, as well as its integration with AI and ML engines that learn and manage the system in an optimal way [4].

The envisioned next generation of cellular networks will strongly rely on a multi-stakeholder infrastructure. Seamless resource sharing based on AI automation will enable service interconnection across diverse domains, in the so-called Cloud-in-Continuum. Cooperation between stakeholders strengthens E2E service delivery by addressing competition and trust models in service provisioning. To this aim, 5G establishes well-known technological pillars to fulfill these needs. E2E network slicing, leveraging NFV and SDN paradigms, logically groups and isolates resources and services, while NFV-MANO enables the automation of the NFV life-cycle management. Lastly, DLT-based solutions enhance trustworthiness in information-sharing processes and establish these as three key pillars [5]. However, of the three aforementioned, slicing management through orchestration often relies on human intervention for decision making and responding to context evolution which reduces its effectiveness. To mitigate the limitations of manual interventions, an abstraction layer [6] is proposed as a solution. As a consequence, ETSI defined the Zero-touch network and Service Management (ZSM) standardization working group [7], oriented toward providing autonomous governability of cross-domain service networks, taking humans out of the loop. ZSM is driven by closed loops that coordinate logical steps. The closed loops defined in ZSM enable interaction between system components via intents, aiming to deliver concrete tasks such as slice deployment. ZSM divides the 5G architecture into different Security Management Domains (SMDs), such as RAN, Edge, Transport, Cloud, or Core. Intra-domain management is addressed by the ZSM architecture [8], which defines functional modules that rationally split functionalities (e.g., orchestration, Policy Management, Analytics). These modules and the domain's infrastructure are connected by a closed loop, driven by an intent. Furthermore, ZSM defines an E2E domain that ensures cross-domain integration. A global perspective of the system allows the E2E domain to manage the life-cycle of E2E slices via intents. In this way, ZSM transforms network management into an autonomous, efficient, and agile operation, simplifying cross-domain service composition.

This work relies on the theoretical concepts from ETSI's ZSM architecture [8] and E2E Management and Orchestration of Network Slicing [9] proposals, extending the concept with inherent security as part of the closed loops [10], avoiding the typically too-late amendments to services when security has already been undermined. This extension is achieved through the adoption of a Security Service Level Agreement (SLA) agreement (SSLA) as a security intent definition, which, in turn, is orchestrated, scheduled, and enforced via security policies. This paper describes the proactive E2E security slice enforcement as a consequence of an SSLA high-level description, as well as the reactive security enforcement triggered by a security state alteration. This work also presents, implements, and evaluates a use case chosen by ETSI as the sixth Proof of Concept [11] for ZSM, titled ETSI ZSM: Security SLA Assurance in 5G Network Slices. In this use case, the analysis of encrypted traffic is addressed as a consequence of an SSLA, and a VNF infection caused by a malicious attacker is mitigated without human intervention. The contributions of this work are as follows:

- Security is integrated as a native component in the process of creating and managing E2E slices.
- An automated closed-loop security operation is achieved.
- An extensible intent-based approach with different levels of abstraction, security models, and services actions is proposed.
- Continuous and autonomous Security Service Level Agreement (SSLA) compliance through tailored monitoring and cross-domain reactive capabilities provided by the framework is achieved.
- Real-world validation on a 5G testbed encompassing multiple administrative domains and different security properties is achieved.
- As result of the above, we demonstrate the feasibility of ZSM architectures for managing B5G networks.

The rest of this paper is organized as follows: Section 2 provides the state of the art. Section 3 provides the security framework overview. Section 4 shows a detailed explanation regarding how the framework provides dynamic 5G core security slice management. Sections 5 and 6 provide the implementation and performance results, respectively. Finally, Section 7 provides the conclusions.

2. Related Work

This section provides the state of the art, focused on the four principal topics of our research efforts: orchestration, network slicing, Zero-touch network and Service Management, and Security and AI.

2.1. Zero-Touch Network and Service Management

The ZSM architectural framework has been designed to automate and simplify the deployment, management, and assurance of services in Beyond 5G (B5G) networks. Extensive research has been conducted on this subject, as presented in works like [12,13], where the authors provide a comprehensive study of ZSM's background, highlighting its major advantages, architectural choices, and the most relevant standards. An in-depth study analyzing the threat surfaces introduced by the technologies constituting ZSM is documented in [14]. Taking into account various standards related to end-to-end (E2E) slicing, such as the ZSM-defined framework and the Transport Slicing solution proposed by the IETF, an abstraction model and corresponding interfaces for Transport Slicing have been introduced in [15]. Although ZSM specifies how closed loops should operate within the system, it does not define how different closed loops should communicate to ensure SLA assurance coordination. Gomes et al. [16] propose an intent-driven model based on delegation, describing an escalation of intents across different closed loops as a hierarchical coordination solution. Their work offers a detailed explanation of the ZSM architecture, emphasizing the definition of the intent model. However, their study focuses only on the model's definition and does not include any evaluation.

Furthermore, in [17], a monitoring model integrated into ZSM-based service life-cycle management is defined. However, despite describing their approach as end to end, it effectively represents a single domain, as they treat each entity within the infrastructure as distinct domains. Finally, in [18], a ZSM-based architectural framework for managing slicing in 3GPP networks is proposed. However, their approach introduces a nested structure of management domains for the same logical part of the infrastructure. This deviates from the ZSM standard definition and significantly increases the complexity of coordination between the various levels of closed loops.

2.2. Orchestration

The orchestration layer offers great potential to seamlessly manage diverse domains and, due to the wide range of possibilities it provides, it has attracted the attention of the scientific community. In Cassaza et al. [19], an orchestration model is proposed, focusing on providing high availability for service function chains, where multiple VNFs that perform the same function but with different roles (Master or Slaves) are instantiated on distinct servers, generating multiple paths for delivering the same service.

Recently, lightweight forms of orchestration have gained significant attention due to their shorter deployment time and ease of portability and scalability. In [20], a survey on container orchestration is presented, where among the identified gaps are container monitoring capabilities, tools, and autonomous features in orchestration frameworks that enable self-healing and self-optimization. In the study by Salhab et al. [21], an open-source Core Network (CN) was orchestrated using Docker Swarm as the manager. Their approach involved deploying multiple nodes with a load balancer and a floating IP to establish a highly available 5G core (5GC). However, this work primarily leveraged Docker's existing capabilities to demonstrate their applicability in the context of 5G. The orchestration lacked dynamism, and the security measures implemented were rudimentary.

Nevertheless, these efforts in cloud-based VNF orchestration are critical, as they lay the foundation for NFV orchestration within 5G and B5G cellular networks. Many of these works tend to focus on specific, isolated aspects, often overlooking the involvement of multiple actors across different domains within the overarching management of 5G networks. In contrast, our approach bridges these gaps by introducing autonomous, policy-intent-based orchestration. This enables us to enforce security measures and countermeasures seamlessly, whether applied to standalone VNFs or to E2E network slices covering the entire 5G system.

2.3. Network Slicing

Network slicing presents itself as the centerpiece for splitting the physical infrastructure into multiple logical, isolated, and customized instances of network resources [22]. When the slices span across multiple domains, encompassing different network segments and technologies, they are referred to as E2E network slices. The scientific community has focused on developing the intricacies of E2E network slicing as it is one of the main pillars of B5G networks. For instance, in [23], the authors present a reference architecture for network slicing end-to-end orchestration, using an ML approach to improve performance during the decision phases of the slicing life-cycle. However, it does not consider the use of ZSM during this process. Additionally in this line, Ref. [24] offers a comprehensive study of network slicing security concerns, exploring the threats, challenges, and issues, highlighting the existing gaps in current research, with the aim to ease the development of a secure network slicing framework.

Furthermore, in [25], a work with open-source tools such as OAI is presented, proposing a four-building-blocks classification that characterizes resource allocation, whereas various AI techniques are used in the orchestration to optimize each of these four blocks. Although the paper presents multiple useful insights, the presented network slicing orchestration only considers RAN slicing. In addition, it uses 5G Non-Standalone (NSA), which means that the slice has no logical management at the core, and security is not considered at all. In [26], Liu et al. propose an architectural solution for E2E slice resource management composed of three different layers: network, orchestration, and intelligence based on Deep Reinforcement Learning (DRL). This work shows promising results in leveraging DRL as an enabling technology to enhance decision making in resource sharing. Additionally,

its architectural proposal could be seamlessly integrated into the well-established and standardized ZSM framework.

Regarding solutions that could be integrated into ZSM, [27] introduces a large-scale E2E network slice facility offering network slices as a service. This initiative effectively coordinates several domains via a federation approach. However, despite mentioning that the challenges associated with the project are being addressed by the ZSM standard, they do not adopt the reference architecture to integrate their solution. In our study, we expanded the concept of E2E slices, customizing them to address security requirements, resulting in the creation of E2E security slices. In particular, we have devised a ZSM-compliant approach to manage these E2E security slices within the context of B5G Networks.

2.4. Security and ML/AI

Utilizing AI to address security challenges is not just a current trend but also a fundamental aspect shaping the future telecommunications landscape. When integrated with NFV and SDN, AI becomes a powerful tool, enabling early detection and rapid responses to a wide set of attacks. The authors in [28] present the potential use of ML/AI-enabled network slicing in current and future infrastructures, offering a comprehensive survey of the use of ML for the life-cycle management of network slices, starting from the slice preparation phase and covering all network domains, but without considering ZSM.

In [29], the authors present a comprehensive survey of the security challenges associated with NFV, proposing the use of anomaly detection techniques to mitigate potential risks of cyberattacks. However, they do not consider network slicing in their work. Additionally in this line, and as mentioned before, [23] presents an architecture that uses an ML approach to enhance the process and increase security. In [30,31], different types of ML techniques used for identifying cybersecurity risks, including phishing, malware, and spam, are reviewed. These techniques fall into two primary categories: unsupervised and supervised learning. In the first category, clustering methods are utilized to detect anomalies based on data similarities, thereby identifying potential security risks. The second category focuses on classifying different types of attacks into predefined categories. Additionally, in recent years, deep neural networks have proven to be decisive for multiple cybersecurity applications [32], such as malware detection, intrusion detection, or malicious traffic classification, among others.

A systematic review of recent publications investigating the application of ML for security within the context of NFV and SDN [33] highlights important shortcomings. Our study aims to address these deficiencies through a practical demonstration. The shortcomings addressed include the lack of a specific ML integration architecture tailored for SDN/NFV, the limited availability of ML solutions beyond well-established attack categories such as DDoS, and inadequate provision of coordinated mitigation and response mechanisms. In particular, we integrated a cryptomining detection ML model to expand this scope [34], autonomously managed by leveraging SDN/NFV paradigms and enabling seamless mitigation with an E2E perspective.

2.5. Related Frameworks

Concerning existing approaches that aim to integrate AI/ML mechanisms into frameworks to improve security, the authors in [35] propose a security framework to automatically mitigate different IoT-related threats. This work employs AI-based reactive agents, powered by ML models, to analyze monitored data and utilize this information to raise intrusion alerts. However, it does not contemplate network slicing or E2E mechanisms for the management or orchestration of infrastructure security. In parallel to that idea, the authors in [36] present a model architecture that combines a validation framework

for applications in conjunction with a reactive AI-driven system to enhance the security of the applications and the infrastructure in real time. That model leverages ML/AI to dynamically monitor and respond to security threats; however, it solely presents a general framework and does not incorporate network slicing or Zero-touch Service Management (ZSM) techniques.

An AI-based network slice management framework is proposed in [37], but the introduced AI mechanisms are exclusively oriented to enhance radio link performance. A similar situation occurs in [38], where AI mechanisms are integrated into 5G networks to improve their security, among other aspects; however, they are mainly focused on the radio link rather than improving the security of the infrastructure itself. In [39], the authors propose a Federated Network Intelligence Orchestration approach for anomaly detection in B5G networks using Federated Learning (FL). To achieve this, the work orchestrates network intelligence to detect and prevent cyberattacks, integrating the system within a Zero-touch Service Management (ZSM) security framework, including multi-domain and multi-tenant orchestration. However, they neither present nor contemplate the inclusion of network slices in their orchestration. Lastly, in [40], the incorporation of certain security controls at the network layer is proposed. These controls are classified by efficiency and suitability, aiming to allocate resources dynamically with enhanced security assurance and isolation capabilities, which are tested to mitigate DDoS attacks. Nevertheless, although the work suggests the suitability of zero-touch mechanisms for resource and reactive orchestration, the theoretical solution outlined in the paper has not been implemented.

3. Security Framework Overview

Traditional ZSM frameworks have made significant strides in specifying automation for service life-cycle management, even though they are still grappling with challenges in security. Generally, security is often considered an overarching aspect, but its definition is either omitted or left ambiguous. Our approach [10] relies on ETSI's ZSM reference architecture [8] to integrate security assets in autonomous and intelligent intent-driven E2E service management (Figure 1). We adapt the standard's closed loops to enhance the framework's autonomy through self-healing, self-repairing, self-configuring, and self-protecting capabilities. We therefore define the closed loops' interactions covering single domain security management and cross-domain coordination. To drive them, we follow an intent-driven approach, proposing three distinct security abstraction levels: SSLAs, MSPL-OP, and final configuration.

The decisions taken within the closed loops are tailored to the requirements specified in the intent. As an innovation, we integrate DLT-based trust management to ponder decision making at its various steps; more details can be found in [41]. Trust is evaluated through continuous SSLA monitoring and analysis. Trust metrics represent the historical assets' and domains' behavior for SSLA compliance. Genuine behavior is rewarded, while vulnerabilities are punished. In ETSI's ZSM definition, single management domains are able to self-manage their own services and infrastructure, but for cross-domain actions, an E2E domain is used. As a novelty, our E2E domain vision presents a centralized data analysis that correlates security information from different domains. This view expands the security coverage across the entire infrastructure, including Federated Cyberthreat Intelligence (CTI) sharing, cross-domain threat mitigation, and the redeployment of mistrusted assets among other E2E domain functionalities.

The remainder of this section is structured as follows: first, we present the various intent models and outline the procedures for transforming them. Next, we detail the key steps involved in both intra-domain and cross-domain closed-loop processes. Finally, we

provide a precise definition of the ZSM framework interactions, focusing on the proactive and reactive components of the closed loops.




Figure 1. ZSM security framework architecture.

3.1. Intent-Driven Security Model

Automated systems require motivations that lead to decisions, ultimately flowing into actions. Intents provide a concrete specification of the desired goals to drive these decisions. The level of accuracy in intent definition can vary. On the one hand, a client with no knowledge of the underlying system only needs to specify the wanted system state. In contrast, a security manager must define specific actions to take in a precise and standardized format.

In this context, intents simplify the security management of the underlying hardware and technologies by abstracting their complexity, providing a flexible and comprehensible approach. For our work, we define three distinct levels of intent specification, ranging from high-level abstraction to detailed configurations: SSLA, MSPL-OP, and final configuration (as illustrated in Figure 2).




Figure 2. Security abstraction levels.

The SSLA serves as the entry point to the evolved ZSM architecture as explained in [42] inspired by the work in [43]. The SSLA is a formal agreement between a service requester, such as an over-the-top (OTT) or other operator, and the provider. The SSLA is the first intent: it defines security and service level objectives (SSLOs), spanning from the QoS requirements (throughput, latency, resource usage, etc.) to the security requirements (confidentiality, privacy, authentication, specific attacks protection, etc.). SSLOs define the genuine system's status for a concrete service. The ZSM framework uses SSLOs to drive decisions toward their fulfillment. During this process, monitoring probes and analytic engines are deployed to predict and detect SSLA violations, and finally prevent or mitigate the threats. Therefore, the final step for an SSLA enforcement process is the appliance of specific actions to continuously maintain the system's status specified by the SSLOs. The SSLA enforcement process usually covers cross-domain coordination. Final actions cannot be taken directly; however, the SSLOs require proper refinement to map their objectives with a domain-related intent. In this regard, the refinement transforms the SSLA into an orchestration policy, named MSPL-OP (see Figure 3 and Figure 4 step 1).

To bridge the gap between the SSLA and actionable deployments, the MSPL-OP serves as an intermediate abstraction. It establishes the means to achieve every SSLO but does not specify concrete actions. In fact, the refinement process maps SSLOs to security capabilities. Every MSPL contains one or more capabilities, while an MSPL-OP contains one or more MSPLs. The MSPL-OP model is specified in Figure 3. SSLOs are categorized to be covered by one or more capabilities. This is precisely what the refinement process entails: transforming SSLO categories into a set of defined security capabilities. For this work, the MSPL-OP model has been extended in several ways. Relevant to this section is the incorporation of *5GSecuritySlice*, allowing the specification of a 5G slice. A 5G slice is composed of a set of intertwined 5G services with security capabilities associated with those services. The 5G slice also differs from the regular MSPL-OP, as it requires multi-domain coordination. Intertwined services and security must be deployed to cover every 5G domain. In summary, the MSPL-OP model has been extended to address the requirements of a 5G slice deployment.

To deploy a 5G slice, this work also considers domain selection and customization. This means that once the MSPL-OP has been created, security capabilities and services must be assigned to their corresponding domains. These domains are where the respective components and services will be deployed. We address this challenge with a straightforward and logical approach. If only one solution is possible, the decision is straightforward. However, if multiple domains can implement the same capability, then *orchestration requirements* come into play. Orchestration requirements define the driving factors for selecting one domain over another. For this work, we adopted trust-based orchestration. In this approach, we select the domain with the most trusted assets to comply with the required capabilities [41]. Various algorithms, such as trust ranking or load balancing, could be applied or combined based on system requirements. Once the capabilities and services are

distributed across domains, the framework generates an appropriate MSPL-OP for each domain. Each MSPL-OP contains only the MSPLs relevant to that specific domain. For example, one MSPL-OP might include 5G core as a service and protection requirements for the control plane, while another MSPL-OP focuses exclusively on protection schemes for the communication channel.




Figure 3. MSPL-OP diagram.

Each SMD, upon receiving an MSPL-OP, before translating it into a final asset deployment and configuration action, analyzes feasibility and checks for conflicts with active policies. To ensure consistency, SMDs extract individual MSPL policies from the MSPL-OP and evaluate them for conflicts using rule-based engines. This step ensures that the translation of MSPLs into assets and configurations maintains system consistency. At this stage, each MSPL might imply actions across various devices, such as deploying VNFs, reconfiguring existing services, or other possibilities dependent on the SMD's infrastructure, geographical location, or any other physical or administrative property (e.g., support for Kubernetes). This means the translation process is infrastructure-, topology-, and environment-aware. To achieve this, the translation process (see Figure 4, step 4) selects assets that handle the capability specified in the MSPL and are available for deployment and configuration within the current SMD premises. For a 5G slice, all translations are packaged as a network slice template (NST), which, in an overview, represents a sub-slice of the E2E slice. The NST then provides the means to enforce every capability and service

requested in the policy model, as well as to ensure the interconnection between components of the sub-slice and across domains.




Figure 4. Intent-driven life-cycle security orchestration.

This paper presents a complementary extension to our previous work [10], where we used a ZSM-based framework to protect a vehicular infrastructure from DDoS attacks. For this work, we provide a deeper view of the intent model and how we have extended it to integrate new security assets such as Smart Traffic Analyzer (STA) (see Section 4.1.1), Proof of Transit (PoT) (see Section 4.1.3), I2NSF agent (I2NSFA) (see Section 4.1.2), and 5G core agent (see Section 4.1.4), which will be used by the PoC in the next section. This work aligns with ETSI’s vision about 5G network slice provisioning from an SSLA definition.

3.2. Intent-Driven Security Closed-Loop Workflow

Before delving into the PoC definition, it is essential to understand the relationship between the entities that compose the ZSM framework. This section presents the concept of a closed loop as a set of recurrent steps that interconnect entities in a protocol-less manner. The closed loop was also defined in our previous work [10], but with a high-level approach. Here, we present a more detailed view of the closed-loop phases, both proactive and reactive. These processes are illustrated in Figure 5 and Figure 6, respectively. While the proactive process focuses on initial SSLA enforcement, the reactive process ensures the slice remains operational and secure through continuous monitoring and adjustment. To maintain clarity, we provide a brief description of each process.

The process begins upon SSLA reception or E2E-level MSPL-OP generation when a slice is identified as necessary. This proactive scenario involves the Plan and Execute steps of a MAPE-K (Monitor, Analyze, Plan, Execute, and Knowledge) framework. This proactive scenario is detailed in Figure 5, where steps 3, 8, and 15 involve MSPL-OP refinement, while steps 1 and 9 highlight conflict detection with existing services or infrastructure capabilities. From this point, the Monitor, Analyze, and Knowledge-related steps of the closed loop run continuously until the service’s end of life. A vulnerability might be detected during the Analyze step of the closed loop, as shown in steps 1 and 2 in Figure 6.

If an available countermeasure is identified, a new MSPL-OP message is generated and sent to the SMD's Security Orchestrator. This orchestrator then checks the resources for availability and adequacy (e.g., in terms of trust), which may lead to a reconfiguration of the already deployed slice. This reconfiguration might involve the removal of a compromised asset whose trust has been reduced as a result of step 2.




Figure 5. Proactive ZSM-aligned security slice orchestration solution.




Figure 6. ZSM 5G component anomalous behaviour reaction diagram.

4.1. Security Assets

Security assets are the core components that enable protection measures to be enforced within our ZSM architecture. These assets include instantiable entities, such as Virtual Security Functions (VSFs), as well as configurable physical and virtual components, like switches or SDN controllers. Both types are responsible for implementing advanced security functions. Assets are managed using a modular approach, allowing seamless integration into the framework. This modularity not only improves system scalability but also enhances security flexibility by enabling the on-demand addition of new security functions as they become available. This approach ensures that the system can evolve rapidly and adapt to new security threats without disrupting existing operations.

The asset integration process starts by expanding the policy model to define the capabilities and configuration parameters of the new security components. This step ensures that the framework understands the asset's functionality. Next, a plugin is developed to translate high-level policies into technology-specific configurations, such as YANG. Finally, a driver is created to directly operate and configure the asset, for example, using NETCONF, thus completing the life-cycle management process.

In this work, several security assets have been integrated into the autonomous management of the ZSM framework. These include as follows:

- STA: Analyzes 5G core control traffic and detects anomalous encrypted traffic.
- I2NSF: Establishes a secure communication channel.
- PoT: Verifies that traffic follows the secure communication channel.
- Fifth generation core agent: Acts as a middleware controller to manage network functions (NFs) and retrieve contextual UE data from the 5GC.

For each of these assets, the aforementioned integration steps have been applied. This section introduces these security modules and explains how they have been integrated into the proposed architecture as part of this work.

4.1.1. Smart Traffic Analyzer (STA)

The STA is an ML-trained security asset designed to retrieve communication flows and classify their behavior. The STA is built with flexibility in mind, allowing different ML models to be loaded and configured based on TSAT [44] values, as outlined in [34]. The STA's configuration defines the type of traffic to be analyzed, the ML model for pattern detection to be loaded, and the possible classification outputs. In the PoC, the STA is utilized for 5G core protection by monitoring and analyzing 5G signaling traffic. The ML model loaded into the STA is trained to analyze monitored flows and classify them as either crypto-mining or genuine traffic. When anomalous traffic is detected, the STA reports an alert, providing a malicious flow classification along with a confidence level.

The STA plays an important role in the presented architecture by handling the Security Data Collection (monitoring) and Security Analytics Engine. STA life-cycle management follows the three steps presented in the introduction of the section. The policy model is extended to include new security analysis and monitoring capabilities, along with additional configuration actions. These extensions allow the inclusion of additional monitoring parameters, such as specifying Service Based Interface (SBI) traffic and defining analysis parameters. These analysis parameters indicate to the trained ML model how to classify abnormal behavior in SBI traffic. Based on the 5G core context and the type of protection specified in the policy, the translation plugin selects the ML model to be used, the TSAT values required by the model to classify traffic types, and the potential classification outputs for the traffic. This ensures alignment between the policy and the STA's configuration. Finally, the configuration driver interacts with the STA API, which operates alongside the 5G core NFs and retrieves its URI from the Data Services. The driver identifies the interface

that requires analysis and sends the appropriate configuration requests to the STA API. Upon receiving these requests, the STA applies the configuration and automatically begins analyzing the traffic.

This process enables the STA to efficiently monitor and analyze traffic in real time, providing a dynamic and robust layer of protection adapted for the 5G core.

4.1.2. I2NSF Agent

The I2NSF IPsec solution [45] enables dynamic VPN deployments by employing a centralized controller to handle requests for securing the traffic between private networks. In the presented scenario, two I2NSF agents are autonomously deployed establishing an IPsec tunnel between the RAN and the computing nodes hosting 5G core components, encrypting and securing all traffic cross-domain. Agents are managed via a controller and utilize the ikeless model eliminating the dependency of agents on IKE protocols [46] for key material establishment, Security Associations, and rekey processes since all the configuration parameters are provided by the controller.

Concerning the policy extension, channel protection capability has been enhanced with the IPSec configuration action, which extends channel protection with new fields and available values for confidentiality and integrity to be compliant with cypher suites needed for private 5G networks' interconnection to public base stations. The plugin takes this policy to build a configuration containing: (i) the networks to be interconnected via the IPsec tunnel, (ii) the addresses of the agents, and also, (iii) security parameters for IPsec, such as encryption and integrity algorithms, the lifetime of the SA for rekeying, and the lifetime of the SA to be removed in case the rekeying process fails. The configuration driver uses the I2NSF controller's REST API to pass these parameters. The I2NSF controller then generates the necessary key material (e.g., initialization vectors and keys), Security Policies (SPs), and SAs. The configuration is securely transmitted to the agents using NETCONF, employing the ietf-ipsec-ikeless YANG model [45]. During the rekeying process, the controller must be subscribed to the agents' NETCONF notifications. These notifications inform the controller when an SA is about to expire, enabling it to generate and provide the key material required for the new SA. This ensures seamless security operations without interruption.

4.1.3. Proof of Transit (PoT)

PoT is a security asset designed to provide trust in the routing path followed by a packet. Based on the IETF draft [47] and the mathematical operations outlined in [48], PoT adds a small portion of operational data to each packet traversing the network devices that form the infrastructure. These data enable the determination of the path's properties and the detection of any missing conditions. In the PoC, PoT ensures trust for traffic expected to traverse a protected channel. It verifies that encrypted traffic arriving at one of the endpoints—in our PoC, the RAN and the core—has followed the desired hops. PoT denies traffic that fails to meet this condition and identifies where the failure occurred. Nodes with PoT enabled are managed by a centralized controller.

The policy model expansion involves extending the monitoring capability to include a configuration action definition for Proof of Transit. This allows specification of the allowable traffic profile for specific interfaces. For example, it can be used to detect plain traffic appearing on an interface where all traffic is expected to be encrypted. The plugin for configuration creation is responsible for supplying the controller with the endpoints of the nodes involved. The driver utilizes the REST API defined by the controller to process the configuration requests. The controller then generates Shamir's Secret Sharing Scheme (SSSS) keys and symmetric keys for each pair of nodes [48]. These configurations are forwarded to the nodes using the IETF PoT profile YANG model via NETCONF.

Once the final agents receive and validate the configuration, they initiate probes to verify the trustworthiness of the path. This process ensures the integrity of the traffic and enables detection of any deviations from the specified routing path.

4.1.4. Fifth Generation Core Agent

The 5G core agent developed in this work functions as a middleware component between the ZSM-based framework and the 5G core. Its primary role is the internal management of various network functions (NFs) within the 5G core network. Specifically, the agent dynamically adapts the configurations and versions of these NFs based on the policy intent defined by the Security Orchestrator (SO). This dynamic adaptation ensures that the 5G core network operates in compliance with the specified security requirements and service parameters dictated by the SO. By intelligently managing NF configurations, the system maintains the desired security and trust levels in response to evolving network conditions and security threats. A key feature of this security asset is its repository of NF images. When a vulnerability is identified in any of the images, a more reliable image is on-boarded, with configurations tailored to the image version and the current status of the 5G core. This proactive image management enhances the overall resilience and security of the network. In addition, the 5G core agent can initiate a request toward the management plane of the RAN, for instance, requesting the incorporation of a PLMN to be broadcast, and register/delist serving AMFs.

The policy model has been extended to introduce the Secured Service MANO capability for the 5G core. This capability ensures that the 5G core provides a mechanism (e.g., an agent) for managing its internal components and features with security in mind. In this context, the 5G core agent becomes a necessary dependency for secure 5G core deployment. The plugin translates the policy into configurations required by the 5G core agent driver. This driver equips the framework with a REST API that facilitates a wide range of actions to implement procedures, such as registering a new AMF in the gNB. The plugin models configurations using a lightweight JSON format, enabling the specification of a diverse array of actions—including redeployment, NF registration, and UE identifier retrieval—along with the associated action-specific parameters such as endpoints, protocols, identifiers, and versions. The Data Services provide additional support for these configurations.

This concise JSON-based approach simplifies the configuration process while maintaining flexibility and precision in defining management and security actions. By enabling seamless interaction between the ZSM-based framework and the 5G core, the 5G core agent ensures that trusted NF versions are effectively deployed in the 5G core.

4.2. 5G Core E2E Security Slice Deployment and Configuration

In this section, we focus on how intent-driven security is utilized in the PoC, as described in Section 3.1, and aligned with the closed-loop operations presented in Section 3.2. Intent-driven security is employed to instantiate a cross-domain slice. In the PoC scenario depicted in Figure 7, a customer requests a private 5G network to interconnect employees and devices across different geographic areas with internal services. To achieve this, the deployment leverages a public gNB, dynamically establishing synchronization with 5G core NFs to enable on-demand cross-domain connectivity. Regarding security, the customer and the infrastructure provider establish an SLA for the private 5G network. The SLA specifies the following security objectives: data confidentiality and integrity for external connectivity with the RAN, protection against inappropriate use of resources in the 5G core, and some contextual information such as the geographical area where employees require coverage.

Figure 8 maps the PoC scenario into the ZSM architecture. Two distinct SMDs are identified: one managed by TID and the other by UMU. Each SMD possesses unique properties and varying security assets. While some capabilities are shared, others are specific to each domain. SSLA deployment involves the creation of a cross-domain slice that spans both SMDs, coordinated via the E2E SMD. The defined SSLA is refined into channel protection with minimum key length and encryption algorithms, Security Traffic Analysis capabilities, and monitoring capabilities to autonomously enable SSLA compliance analysis. The 5G core, combined with these security capabilities, forms a refined MSPL-OP. This MSPL-OP is subsequently passed through the SMD selection and customization phase, where the appropriate domains are determined. For this, the geographical location of users is considered to select the serving RAN SMD. During this step, additional adaptations are made, such as selecting the available channel protection protocol for the selected domains to prevent conflicts caused by protocol incompatibilities.




Figure 8. Proactive ZSM-aligned E2E B5G security slice orchestration solution.

The refined MSPL-OPs are dispatched to the SMD orchestrators, triggering policy translation, sub-slice deployment, and configuration enforcement. For TID, the MSPL-OP is translated into a sub-slice that includes Free5GC as the fully containerized 5G core, STA for traffic analysis and monitoring to protect against abnormal resource use in cryptomining situations, I2NSF for IPsec channel protection, and PoT for monitoring the IPsec tunnel.

Meanwhile, in the UMU SMD, the MSPL-OP is translated into a sub-slice providing the IPsec tunnel endpoint and its associated PoT monitoring capabilities, specifically within the RAN domain. Table 1 outlines the capabilities included in each MSPL-OP sent to the respective domains.

Table 1. Available domain’s capabilities.

Capability	UMU	TID
5G core	✓	✓
Channel protection	✓	✓
Monitoring channel protection	✓	✓
Monitoring and analysis 5G core		✓
Radio Access Network	✓	

Once the slice components are prepared, the Security Orchestrator enforces the configurations as explained in Security Assets (Section 4.1). This enforcement process prioritizes dependencies and ensures that the security capabilities are consistently applied across the entire infrastructure. Briefly summarizing, once the IPsec tunnel is established, the 5G core can securely communicate with the RAN. Subsequently, the 5G core agent configures the 5G core NFs and sends a PLMN broadcast request to the gNB. Upon acceptance, the gNB starts irradiating the requested PLMN and registers the serving AMF, thereby establishing the control plane for 5G communication through the protected channel. Monitoring and analysis probes are activated only after their dependency on operational components are satisfied.

4.3. B5G E2E Security Slice Reaction

The reactive aspect of the use case is illustrated in Figure 9, showing the local and E2E reactions within the system. The reactive mitigation process follows the steps outlined in Figure 6. The STA, which monitors and analyzes the 5G core SBI, detects abnormal traffic in the AMF. After analyzing the traffic anomaly, a cryptomining alert is sent to the Decision Engine (DE), accompanied by a confidence level in the detection result. The DE subsequently updates the trust values for the STA, the affected NFs, and the domain in the Trust and Reputation Management (TRM) system. Specifically, the trust value for the STA increases, while the values for the affected version of the NF, Free5GC, and the domain decrease.

Upon receiving the alert, two countermeasures are generated. First, a local reaction is triggered, which redeploys a more reliable version of the vulnerable NF within the local domain. Second, an E2E reaction is initiated, which involves scanning other domains for similar vulnerabilities and applying appropriate countermeasures. This ensures end-to-end security compliance and addresses inter-slice security correlation.

These reactive mechanisms enable the system to adapt swiftly to security threats without human intervention, maintaining compliance with SLA requirements and ensuring continuous protection across the infrastructure.




Figure 9. Reactive ZSM-aligned guarantee of SLA compliance orchestration solution.

5. Testbed Implementation and Integration

To validate the proposed architecture and processes, a realistic multi-vendor, multi-site setup was employed, featuring a real 5G network including UEs and virtualized backend infrastructure managed by our proposed ZSM framework. This section outlines the testbed specifics and the implementation efforts undertaken to achieve the most realistic evaluation, apart from production environments.

The testbed is supported through the coordinated effort of two geographically separated locations with three independent management domains. The University of Murcia operates two separate Security Management Domains (SMDs): UMU SMD, which has a Release 16 compliant 5G network that delivers high coverage capabilities through two cells connected to a single gNB, forming the RAN domain, and E2E SMD, which independently manages the global state of the entire infrastructure. Additionally, Telefónica I + D contributes by hosting the 5G core and managing the transport domain within the TID SMD.

The key properties of the testbed include as follows:

1. Multi-domain integration, ensuring seamless coordination across various domains.
2. Virtualized management, NFV and SDN.
3. Real-world 5G NR communications, with dedicated hardware providing real 5G NR connectivity.
4. Interoperability, offering mechanisms to integrate different technology solutions.
5. Dynamic scalability, facilitating modular expansion of capabilities.

6. Intent-driven architecture, precisely modeling system requirements and enabling E2E enforcement.
7. Continuous monitoring, ensuring compliance with system policies through persistent monitoring.
8. Security instrumentation, incorporating tools to simulate security threats and test system responsiveness.
9. Machine learning integration, where trained ML models monitor and analyze security aspects.
10. RESTful APIs, providing communication between ZSM entities across domains.

To meet these objectives, a broad range of components was implemented, primarily in Python 3.x. ZSM architectural components are developed using REST APIs for simplified communication and are containerized for deployment in a Kubernetes ecosystem. The Security Orchestrator was extended to utilize OSM for automated deployment and configuration of 5G security slices. Dynamic and reactive agents were deployed as virtual machines (VMs) through OpenStack. Several drivers and plugins were developed to handle the translation of policies and the configuration of security assets: I2NSFControllerDriver, PoTControllerDriver, STADriver, and 5GCoreAgentDriver.

Efforts were also focused on developing the Integration Fabric, a critical component of the architecture, based on Kubernetes, Istio, Kafka, and metallB. These technologies were selected for their ability to meet ETSI's requirements [8]. Kubernetes provides a high-availability environment for containerized service deployment and life-cycle management, while Istio simplifies traffic management, observability, and security. Kafka enhances event streaming within the service mesh, ensuring secure and balanced routing of network traffic.

Regarding the technological and computational resources used for performance evaluation, the E2E SMD and 5G RAN SMD were deployed at UMU. The control plane was hosted in a Kubernetes cluster virtualized on an AMD EPYC 730P 16-core processor with 32 threads and 128 GB of RAM, while the data plane was hosted on an OpenStack Rocky infrastructure with two identical compute nodes running Intel Xeon Gold 6138 processors, each with 40 cores, 80 threads, and 256 GB of RAM. The RAN was deployed using AW2S, running the Amarisoft Radio stack on an Intel Core i7-9700K CPU @3.60 GHz with eight cores, eight threads, and 8 GB of DDR4 2666 MHz RAM. Meanwhile, the 5G core/transport SMD was deployed at TID. The control and data planes were hosted on a server equipped with an Intel Xeon Gold 5318N CPU, offering 24 cores, 48 threads, and 120 GB of RAM. For the UE, we used a Nemo Handy Handheld Measurement solution, a 5G Standalone (SA) mobile phone with monitoring capabilities. Despite the use of high-performance machines for realistic deployments, the core of the framework's control plane, including the fabric, requires 16 GB of RAM and eight cores, making it feasible for deployment in less powerful environments.

6. Performance Evaluation

To determine the feasibility of our ZSM-based framework, we extracted various metrics for both proactive and reactive stages as part of a common PoC.

For validating the proactive phase, we measured the time taken by the framework to deploy and configure an SSLA as an E2E 5G security slice, as outlined in Section 4.2. We provided: (i) time measurements from an E2E perspective, and (ii) detailed time measurements per phase (orchestration, deployment, and configuration enforcement) within each SMD. For the reactive phase, we triggered a cryptomining attack as described in Section 4.3. We considered (iii) the number of false positives and false negatives in the STA, (iv) the E2E and SMD Mean Time To Resolve (MTTR) the incident, (v) the Time to Detect, and finally (vi) the AMF Downtime.

6.1. Proactive Experimental Results

In this section, we explain the different KPIs used to evaluate the proactive stage of our ZSM-based framework for managing the security of a real B5G system, as well as provide insights gained during the process. The total initial time is calculated by summing the per-domain times. Specifically, the E2E SMD initial time is measured from the moment the E2E SO receives the slice and security requirements until these policies are distributed across each domain. For the UMU and TID SMDs, the initial time refers to the time elapsed from when each domain's SO receives the B5G security slice policy to its full deployment and configuration across the infrastructure.

Figure 10a illustrates the initial time required to deploy and configure the security slice across the E2E, UMU, and TID SMDs. The results show that the E2E initial time is negligible, less than 2 s, as there are no VNF instantiation or configuration tasks at the E2E domain. In contrast, the TID and UMU domains have significantly longer times, averaging 95 s and 180 s, respectively. These times are attributed to the more extensive tasks handled in these domains, such as 5G core setup, channel protection, STA monitoring, and Proof of Transit (PoT). Figure 10b further breaks down these initial times by key processes in proactive slice deployment for each domain (UMU and TID). Key stages include as follows:

- **Orchestration time:** The time taken by the SO to create and enforce the orchestration plan, including policy translation. This stage is minimal in both domains, taking less than 3 s.
- **Deployment time:** The time required to dynamically deploy security enablers, constituting a significant portion of the total time. This is primarily influenced by the type of virtualization used and the dependencies and coordination required between the assets.
- **Enforcement time:** The time required for security enablers to be configured and ready. This varies based on dependencies. For instance, PoT becomes operational within 7.2 s, whereas other processes, like the I2NSF agent, require longer wait times due to dependency sequencing.

The I2NSF agent's enforcement time is the largest component of UMU's total time, requiring approximately 100 s to be fully operational. In contrast, STA in TID takes about 25 s to reach readiness, with 15 s dedicated to its final configuration. Full details are presented in the following table.

Table 2 shows the distribution of times required by each domain to apply each policy. It presents both the total elapsed time and a breakdown into specific operational phases, such as startup and configuration. These measurements were obtained using the technological and computational resources described at the end of Section 5. For I2NSF and PoT, measurements appear only in the UMU SMD domain because the timing is recorded by controllers at the UMU premises. As a result, the reported startup time encompasses the period required for both the TID and UMU VNFs to become ready for configuration, while the configuration time refers to the additional time needed for full operational readiness. PoT requires the E2E channel protection to be operational before its configuration can take place. Additionally, note that the STA VNF is deployed exclusively in TID SMD alongside the 5G core VNF. These results are further compared with the state-of-the-art solutions in Section 6.3. Among the resources used, the 5G core with the STA enabler required a flavor powered by four cores, 4 GB of RAM, and 32 GB of disk space. The same flavor was used for the I2NSF agents and PoT VMs.




Figure 10. E2E B5G slice deployment and configuration initial time. (a) Use case initial time; (b) time consumed per procedure per SMD.

Table 2. Percentage distribution of time for UMU and TID SMDs.

Domain	Process	Percentage of Time
UMU SMD	Total time	180 s
	Deployment time	36%
	I2NSF agent (startup)	55%
	I2NSF agent (configuration)	3.32%
	Enforce PoT	3.97%
TID SMD	Total time	95 s
	Deployment time	55%
	STA (startup)	26%
	STA (configuration)	15%

Analyzing the extracted data, the experimental results highlight the effectiveness and scalability of the ZSM-based framework for managing B5G system security. The results show a minimal initial deployment time in the E2E domain, with more significant setup requirements in the UMU and TID domains due to their complex, resource-intensive tasks, such as 5G core deployment, channel protection, and active monitoring, achieving a full B5G network deployment in under 300 s. It is important to note that the enforcement time for critical components like the I2NSF agent and STA varies due to dependency constraints, impacting overall readiness. These findings demonstrate the framework's potential to streamline multi-domain security deployment and suggest opportunities to further reduce deployment times through optimized resource allocation, parallel task execution, and container-based deployments.

6.2. Reactive Experimental Results

Regarding the reactive stage performance, we evaluate the system using multiple Key Performance Indicators (KPIs). These include Detection Time (time taken by the monitoring system to identify malicious activity), False Positive Rate (the percentage of benign traffic misclassified as malicious), False Negative Rate (the percentage of malicious traffic misclassified as benign), and Mitigation Time (time taken to detect and neutralize the threat).

We start with the monitoring probes. Figure 11a shows the time it took for the STA detector to identify cryptomining traffic. The results are promising, with an average detection time of just 87.08 milliseconds, demonstrating the efficiency of the detection mechanism. Additionally, Figure 12a highlights the occurrences of false positives and false negatives when analyzing a pcap file captured from the 5G core, containing both cryptomining and 5G signaling traffic. With prior knowledge of the expected outcomes, the model achieved an average false negative rate of 0.022% and a false positive rate of 0.013%, reflecting outstanding accuracy. Over 3200 classifications per iteration were performed, and the system consistently exhibited robust performance throughout the experiments. It is important to highlight that in cybersecurity contexts, false positives are particularly significant because they can trigger unnecessary actions within the framework. Our model demonstrates better performance in this regard than other models as explained in [34]. Furthermore, the STA in this scenario is confined to a network slice, resulting in only a minimal number of erroneous values. Nevertheless, it is important to consider the error deviation associated with false positives and apply appropriate threshold quotas to avoid misguided actions that could negatively impact overall system performance. Regarding false negatives, although the results show a minimal detection rate of this kind, in a large-scale system, even a small fraction of undetected cryptomining traffic can translate into significant illicit resource usage. Employing more sophisticated models that leverage multi-layer inspection and incorporating enhanced data quality and variety can further improve detection performance and mitigate the impact of false negatives.

Finally, Figure 12b illustrates the time taken by the system to detect and mitigate the cryptomining attack at both the SMD and E2E levels. Impressively, the attack was mitigated within 20 s at the SMD level and in under 10 s at the E2E level, involving cross-domain interaction. Interestingly, despite requiring more interactions, the E2E mitigation process was quicker. This is because E2E mitigation benefits from a preventive approach, which is triggered at the same time as the alert. In contrast, during the attack, the UMU SMD exhibited a faster response, whereas TID SMD's resources were strained, with nodes overwhelmed by the ongoing attack.




Figure 11. Time to Detect and AMF Service Downtime. (a) Time to Detect; (b) AMF Service Downtime.




Figure 12. False positives/negatives and Mean Time to Resolve. (a) False positives/negatives; (b) Mean Time To Resolve.

6.3. Comparison with Other Works

According to the results we obtained, the research outcomes are promising. In Table 3, we compare our findings with those from other works in the scientific community. For instance, in [49], the OSM slice deployment time was measured using vCache and vDPI, where a 4G Evolved Packet Core (EPC) with Network Services was deployed in a single domain, averaging 68 s for full slice deployment. This result falls within the standard deviation range of our single-domain deployment times.

Table 3. Time results comparison with related scientific works.

Work	Deployment Time	Configuration Time	Technology
[49]	68 s	no data	OSM
[50]	58 s	no data	OSM + OpenStack
[51]	40 s	125 s	OSM + OpenStack
UMU SMD	65 s	106 s (I2NSF) + 7 PoT	OSM + OpenStack
TID SMD	53 s	40 s	OSM + OpenStack

Similarly, in [50], OSM-based slice deployment in a single domain was evaluated, where the network slice included a VNF with two Virtual Deployment Units (VDUs) deployed on OpenStack as virtual machines. The deployment time was approximately 58 s, which also aligns with our single-domain standard deviation. Furthermore, Ref. [51] validated network slice deployments using OSM and OpenStack in a single domain. Their work, which included the configuration time for a network slice composed of a mobile core and RAN, reported slice deployment times of around 40 s in the worst-case scenario, with configuration times extending up to 125 s—figures consistent with our configuration times. Regarding STA, further insights into the ML model’s performance compared to other approaches can be found in [34].

From these comparisons, we can conclude that, with current technology, approximately 52 s per domain are typically required for each sub-slice instantiation via OSM. The remaining time is accounted for by the E2E orchestration process, the time needed for services to become operational, and configuration tasks. As noted, these times can vary significantly depending on the software and hardware used (e.g., external controllers) and the load on the services. Additionally, real E2E coordination across multiple domains has been scarcely tested, making our work particularly innovative in this area.

7. Discussions

As we move into the era of 6G networks, the expectations for always-on, uninterrupted service are higher than ever. In this fast-paced environment, relying on humans to manage and intervene in every aspect of network operations is not sustainable. Instead, we are aiming for “zero-touch management”—an ideal where networks can manage and secure themselves with minimal human input. However, even with automation, we must remember that networks are ultimately built and used by people. This means that while machines can handle much of the work, human guidance in the form of clear policies and security standards will remain critical.

In this paper, we explored a new way of managing and securing these next-generation networks, focusing on integrating security into every part of the process. We aligned our work with ETSI standards, but what makes our approach stand out is how it works across different domains. Essentially, it can manage multiple sections of the network at the same time, even when those sections are run by different entities. The framework we proposed ensures that security is not something you add later; it is built into the very core of how the network operates. Whether we are preventing potential attacks or reacting to ongoing threats, the system responds quickly and automatically to maintain security.

One of our key contributions is the definition and development of an intent-driven security management, with different models for intents within the Zero-touch Service Management (ZSM) framework. This model management allows network administrators to define high-level security goals in the form of intents, which are then translated into specific policies and eventually into configurations by the system. By leveraging this model, we enable a clear, human-readable way to guide the autonomous management of networks while minimizing manual intervention. This intent-driven approach ensures that complex, multi-domain networks remain aligned with overarching security and operational objectives, even as they dynamically adapt to evolving conditions.

Another significant contribution is the creation and management of E2E security slices upon the intent-driven security management. This technology brings together several innovative security measures, such as tools to analyze threats, ensure data integrity, and monitor network traffic to prevent attacks and maintain security requirements. These tools work in real time to detect and respond to threats, maintaining trust and security across the network, even when different domains are involved. We also made sure our approach is flexible, meaning it can adapt to various types of networks and infrastructures, and scale up as needed.

In simple terms, our framework empowers networks to defend themselves in a more dynamic, human-like way—anticipating potential threats and responding instantly when something goes wrong. It also makes it easier for security administrators to manage, as they can set broad policies and let the network handle the details, without needing to micromanage every step.

In conclusion, the results of our research demonstrate the potential for autonomous security management in beyond 5G networks. By integrating advanced technologies and real-time monitoring, we have addressed some of the biggest challenges in securing these complex, multi-domain systems. As networks continue to evolve, we believe that our approach will help lay the foundation for even more innovative solutions, bringing together the strengths of automation and human oversight to ensure safer, smarter networks for the future.

Author Contributions: Conceptualization, R.A.-G., A.M.Z., A.P. and A.S.; methodology, R.A.-G., A.M.Z., A.P.; software, R.A.-G., A.M.Z., H.R.P.; validation, R.A.-G., H.R.P., A.M.Z.; investigation, R.A.-G., A.M.Z., J.O., A.H., H.R.P., A.P. and A.S.; resources, A.H., H.R.P.; data curation, R.A.-G.; writing—original draft preparation, R.A.-G., A.H., A.M.Z., J.O., A.P., H.R.P. and A.S.; writing—review and editing, R.A.-G., A.H., A.M.Z., J.O.; visualization, R.A.-G.; supervision; project administration, R.A.-G., A.M.Z.; funding acquisition, A.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by ONOFRE 3 Grant number PID2020-112675RB-C44. The APC was funded by MCIN/AEI/10.13039/5011000011033.

Data Availability Statement: No data set was saved during this research.

Acknowledgments: This research is supported by ONOFRE 3 Grant number PID2020-112675RB-C44. The APC was funded by MCIN/AEI/10.13039/5011000011033. It was also supported by the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union—nextGeneration EU [TSI-063000-2021-36, TSI-063000-2021-44, TSI-063000-2021-45, TSI-063000-2021-62], the Horizon project RIGOUROUS funded by the European Commission, GA: 101095933, and by the Spanish Ministry of Science and Innovation under the DIN2019-010827 Industrial PhD Grant, co-funded by Odin Solutions S.L.

Conflicts of Interest: Hugo Ramón Pascual, and Antonio Pastor were employed by the company Telefónica. The authors declare no conflicts of interest.

References

1. Simsek, M.; Zhang, D.; Ohmann, D.; Matthe, M.; Fettweis, G. On the Flexibility and Autonomy of 5G Wireless Networks. *IEEE Access* **2017**, *5*, 22823–22835. [\[CrossRef\]](#)
2. Saraiva de Sousa, N.F.; Lachos Perez, D.A.; Rosa, R.V.; Santos, M.A.; Esteve Rothenberg, C. Network Service Orchestration: A survey. *Comput. Commun.* **2019**, *142–143*, 69–94. [\[CrossRef\]](#)
3. Ji, X.; Huang, K.; Jin, L.; Tang, H.; Liu, C.; Zhong, Z.; You, W.; Xu, X.; Zhao, H.; Wu, J.; et al. Overview of 5G security technology. *Sci. China Inf. Sci.* **2018**, *61*, 081301. [\[CrossRef\]](#)
4. Siriwardhana, Y.; Porambage, P.; Liyanage, M.; Ylianttila, M. AI and 6G Security: Opportunities and Challenges. In Proceedings of the 2021 Joint European Conference on Networks and Communications & 6G Summit EuCNC6G Summit, Porto, Portugal, 8–11 June 2021; pp. 616–621. [\[CrossRef\]](#)
5. Gkonis, P.; Giannopoulos, A.; Trakadas, P.; Masip-Bruin, X.; D’Andria, F. A Survey on IoT-Edge-Cloud Continuum Systems: Status, Challenges, Use Cases, and Open Issues. *Future Internet* **2023**, *15*, 383. [\[CrossRef\]](#)
6. Gutierrez-Estevez, D.M.; Wang, Y.; Gramaglia, M.; Domenico, A.D.; Dandachi, G.; Khatibi, S.; Tsolkas, D.; Balan, I.; Garcia-Saavedra, A.; Elzur, U. Artificial Intelligence for Elastic Management and Orchestration of 5G Networks. *IEEE Wirel. Commun.* **2019**, *26*, 134–141. [\[CrossRef\]](#)
7. ETSI GS ZSM 006 Version 1.1.1. Zero Touch Network and Service Management (ZSM); Proof of Concept Framework. 2018. Available online: https://www.etsi.org/deliver/etsi_gs/ZSM/001_099/006/01.01.01_60/gs_ZSM006v010101p.pdf (accessed on 1 March 2024).
8. ETSI. Zero Touch Network and Service Management (ZSM); Reference Architecture. 2019. Available online: https://www.etsi.org/deliver/etsi_gs/ZSM/001_099/002/01.01.01_60/gs_ZSM002v010101p.pdf (accessed on 15 March 2024).
9. ETSI. ETSI GS ZSM 003 v1.1-cdn.standards.iteh.ai. Available online: <https://cdn.standards.iteh.ai/samples/54284/23c11fe05cfb4e018436ac371db62dac/ETSI-GS-ZSM-003-V1-1-1-2021-06-.pdf> (accessed on 15 March 2024).
10. Asensio-Garriga, R.; Alemany, P.; Zarca, A.M.; Sedar, R.; Kalalas, C.; Ortiz, J.; Vilalta, R.; Muñoz, R.; Skarmeta, A. ZSM-Based E2E Security Slice Management for DDoS Attack Protection in MEC-Enabled V2X Environments. *IEEE Open J. Veh. Technol.* **2024**, *5*, 485–495. [\[CrossRef\]](#)
11. ETSI. ETSI ZSM PoC 6: Intent-Based Assurance for 5G and Beyond. 2022. Available online: https://zsmwiki.etsi.org/images/e/e1/ZSM_POC_6.pdf (accessed on 16 December 2024).
12. Coronado, E.; Behraves, R.; Subramanya, T.; Fernández-Fernández, A.; Siddiqui, M.S.; Costa-Pérez, X.; Riggio, R. Zero Touch Management: A Survey of Network Automation Solutions for 5G and 6G Networks. *IEEE Commun. Surv. Tutorials* **2022**, *24*, 2535–2578. [\[CrossRef\]](#)
13. Liyanage, M.; Pham, Q.V.; Dev, K.; Bhattacharya, S.; Maddikunta, P.K.R.; Gadekallu, T.R.; Yenduri, G. A survey on Zero touch network and Service Management (ZSM) for 5G and beyond networks. *J. Netw. Comput. Appl.* **2022**, *203*, 103362. [\[CrossRef\]](#)

14. Benzaid, C.; Taleb, T. ZSM Security: Threat Surface and Best Practices. *IEEE Netw.* **2020**, *34*, 124–133. [\[CrossRef\]](#)
15. Rokui, R.; Yu, H.; Deng, L.; Allabaugh, D.; Hemmati, M.; Janz, C. A Standards-Based, Model-Driven Solution for 5G Transport Slice Automation and Assurance. In Proceedings of the 2020 6th IEEE Conference on Network Softwarization (NetSoft), Ghent, Belgium, 29 June–3 July 2020; pp. 106–113. [\[CrossRef\]](#)
16. Gomes, P.H.; Buhrgard, M.; Harmatos, J.; Mohalik, S.K.; Roeland, D.; Niemöller, J. Intent-driven Closed Loops for Autonomous Networks. *J. ICT Stand.* **2021**, *9*, 257–290. [\[CrossRef\]](#)
17. Saraiva de Sousa, N.F.; Lachos Perez, D.; Esteve Rothenberg, C.; Henrique Gomes, P. End-to-End Service Monitoring for Zero-Touch Networks. *J. ICT Stand.* **2021**, *9*, 91–112. [\[CrossRef\]](#)
18. Sajjad, M.M.; Jayalath, D.; Tian, Y.C.; Bernardos, C.J. ZSM-based Management and Orchestration of 3GPP Network Slicing: An Architectural Framework and Deployment Options. *arXiv* **2022**, arXiv:2203.12775.
19. Casazza, M.; Bouet, M.; Secci, S. Availability-driven NFV orchestration. *Comput. Netw.* **2019**, *155*, 47–61. [\[CrossRef\]](#)
20. Casalicchio, E. Container Orchestration: A Survey. In *Systems Modeling: Methodologies and Tools*; Puliafito, A., Trivedi, K.S., Eds.; Springer International Publishing: Cham, Switzerland, 2019; pp. 221–235. [\[CrossRef\]](#)
21. Salhab, N.; Rahim, R.; Langar, R. NFV Orchestration Platform for 5G over On-the-fly Provisioned Infrastructure. In Proceedings of the IEEE INFOCOM 2019—IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), Paris, France, 29 April–2 May 2019; pp. 971–972. [\[CrossRef\]](#)
22. Afolabi, I.; Taleb, T.; Samdanis, K.; Ksentini, A.; Flinck, H. Network Slicing and Softwarization: A Survey on Principles, Enabling Technologies, and Solutions. *IEEE Commun. Surv. Tutorials* **2018**, *20*, 2429–2453. [\[CrossRef\]](#)
23. Martins, J.S.B.; Carvalho, T.C.; Moreira, R.; Both, C.B.; Donatti, A.; Corrêa, J.H.; Suruagy, J.A.; Corrêa, S.L.; Abelem, A.J.G.; Ribeiro, M.R.N.; et al. Enhancing Network Slicing Architectures With Machine Learning, Security, Sustainability and Experimental Networks Integration. *IEEE Access* **2023**, *11*, 69144–69163. [\[CrossRef\]](#)
24. De Alwis, C.; Porambage, P.; Dev, K.; Gadekallu, T.R.; Liyanage, M. A Survey on Network Slicing Security: Attacks, Challenges, Solutions and Research Directions. *IEEE Commun. Surv. Tutorials* **2024**, *26*, 534–570. [\[CrossRef\]](#)
25. Salhab, N.; Langar, R.; Rahim, R. 5G network slices resource orchestration using Machine Learning techniques. *Comput. Netw.* **2021**, *188*, 107829. [\[CrossRef\]](#)
26. Liu, Q.; Choi, N.; Han, T. Deep Reinforcement Learning for End-to-End Network Slicing: Challenges and Solutions. *IEEE Netw.* **2022**, *37*, 222–228. [\[CrossRef\]](#)
27. Ordóñez-Lucena, J.; Tranoris, C.; Rodrigues, J.; Contreras, L.M. Cross-domain Slice Orchestration for Advanced Vertical Trials in a Multi-Vendor 5G Facility. In Proceedings of the 2020 European Conference on Networks and Communications (EuCNC), Dubrovnik, Croatia, 15–18 June 2020; pp. 40–45. [\[CrossRef\]](#)
28. Donatti, A.; Corrêa, S.L.; Martins, J.S.B.; Abelem, A.J.G.; Both, C.B.; de Oliveira Silva, F.; Suruagy, J.A.; Pasquini, R.; Moreira, R.; Cardoso, K.V.; et al. Survey on Machine Learning-Enabled Network Slicing: Covering the Entire Life Cycle. *IEEE Trans. Netw. Serv. Manag.* **2024**, *21*, 994–1011. [\[CrossRef\]](#)
29. Zehra, S.; Faseeha, U.; Syed, H.J.; Samad, F.; Ibrahim, A.O.; Abulfaraj, A.W.; Nagmeldin, W. Machine Learning-Based Anomaly Detection in NFV: A Comprehensive Survey. *Sensors* **2023**, *23*, 5340. [\[CrossRef\]](#)
30. Buczak, A.L.; Guven, E. A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Commun. Surv. Tutorials* **2015**, *18*, 1153–1176. [\[CrossRef\]](#)
31. Torres, J.M.; Comesaña, C.I.; Garcia-Nieto, P.J. Machine learning techniques applied to cybersecurity. *Int. J. Mach. Learn. Cybern.* **2019**, *10*, 2823–2836. [\[CrossRef\]](#)
32. Berman, D.S.; Buczak, A.L.; Chavis, J.S.; Corbett, C.L. A survey of deep learning methods for cyber security. *Information* **2019**, *10*, 122. [\[CrossRef\]](#)
33. Herrera, J.A.; Camargo, J.E. A survey on machine learning applications for software defined network security. In Proceedings of the International Conference on Applied Cryptography and Network Security, Bogota, Colombia, 5–7 June 2019; pp. 70–93.
34. Pastor, A.; Mozo, A.; Vakaruk, S.; Canavese, D.; Lopez, D.R.; Regano, L.; Gomez-Canaval, S.; Lioy, A. Detection of encrypted cryptomining malware connections with machine and deep learning. *IEEE Access* **2020**, *8*, 158036–158055. [\[CrossRef\]](#)
35. Bagaa, M.; Taleb, T.; Bernabe, J.B.; Skarmeta, A. A Machine Learning Security Framework for Iot Systems. *IEEE Access* **2020**, *8*, 114066–114077. [\[CrossRef\]](#)
36. Hermosilla, A.; Gallego-Madrid, J.; Martinez-Julia, P.; Ortiz, J.; Kafle, V.P.; Skarmeta, A. Advancing 5G Network Applications Lifecycle Security: An ML-Driven Approach. *Comput. Model. Eng. Sci.* **2024**, *141*, 1447–1471. [\[CrossRef\]](#)
37. Bega, D.; Gramaglia, M.; Garcia-Saavedra, A.; Fiore, M.; Banchs, A.; Costa-Perez, X. Network Slicing Meets Artificial Intelligence: An AI-Based Framework for Slice Management. *IEEE Commun. Mag.* **2020**, *58*, 32–38. [\[CrossRef\]](#)
38. Jiang, W.; Anton, S.D.; Dieter Schotten, H. Intelligence Slicing: A Unified Framework to Integrate Artificial Intelligence into 5G Networks. In Proceedings of the 2019 12th IFIP Wireless and Mobile Networking Conference (WMNC), Paris, France, 11–13 September 2019; pp. 227–232. [\[CrossRef\]](#)

39. Fernández Saura, P.; Bernabé Murcia, J.M.; García de la Calera Molina, E.; Molina Zarca, A.; Bernal Bernabé, J.; Skarmeta, A. Federated Network Intelligence Orchestration for Scalable and Automated FL-Based Anomaly Detection in B5G Networks. *Comput. Mater. Contin.* **2024**, *80*, 163–193. [\[CrossRef\]](#)
40. Wichary, T.; Mongay Batalla, J.; Mavromoustakis, C.X.; Żurek, J.; Mastorakis, G. Network Slicing Security Controls and Assurance for Verticals. *Electronics* **2022**, *11*, 222. [\[CrossRef\]](#)
41. Perez Palma, N.; Matheu Garcia, S.N.; Zarca, A.; Ortiz, J.; Skarmeta, A. Enhancing trust and liability assisted mechanisms for ZSM 5G architectures. In Proceedings of the 2021 IEEE 4th 5G World Forum (5GWF), Montreal, QC, Canada, 13–15 October 2021; pp. 362–367. [\[CrossRef\]](#)
42. Vilalta, R.; Alemany, P.; Sedar, R.; Kalalas, C.; Casellas, R.; Martínez, R.; Vázquez-Gallego, F.; Ortiz, J.; Skarmeta, A.; Alonso-Zarate, J.; et al. Applying Security Service Level Agreements in V2X Network Slices. In Proceedings of the 2020 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), Leganes, Spain, 10–12 November 2020; pp. 114–115. [\[CrossRef\]](#)
43. Casola, V.; De Benedictis, A.; Rak, M.; Villano, U. SLA-Based Secure Cloud Application Development: The SPECS Framework. In Proceedings of the 2015 17th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC), Timisoara, Romania, 21–24 September 2015; pp. 337–344. [\[CrossRef\]](#)
44. Mellia, M.; Carpani, A.; Lo Cigno, R. TStat: TCP STatistic and Analysis Tool. In *Quality of Service in Multiservice IP Networks*; Marsan, M.A., Corazza, G., Listanti, M., Roveri, A., Eds.; Springer: Berlin/Heidelberg, Germany, 2003; pp. 145–157.
45. Marin-Lopez, R.; Lopez-Millan, G.; Pereniguez-Garcia, F. A YANG Data Model for IPsec Flow Protection Based on Software-Defined Networking (SDN). 2021. Available online: <https://www.rfc-editor.org/info/rfc9061> (accessed on 2 June 2024).
46. Frankel, S.; Krishnan, S. IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap. 2011. Available online: <https://www.rfc-editor.org/info/rfc6071> (accessed on 3 May 2024).
47. Brockners, F.; Bhandari, S.; Mizrahi, T.; Dara, S.; Youell, S. Proof of Transit. Internet-Draft draft-ietf-sfc-proof-of-transit-08. *Internet Eng. Task Force* 2020, *work in progress*.
48. Martinez, J.V.; Forcada, M.A.E.; Pastor, A.; Lopez, D.; Alonso-López, J.A. Implementation of a traffic flow path verification system in a data network. In Proceedings of the 2024 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit), Antwerp, Belgium, 3–6 June 2024; pp. 943–948. [\[CrossRef\]](#)
49. Kourtis, M.A.; Anagnostopoulos, T.; Kuklilski, S.; Wierzbicki, M.; Oikonomakis, A.; Xilouris, G.; Chochliouros, I.P.; Yi, N.; Kostopoulos, A.; Tomaszewski, L.; et al. 5G Network Slicing Enabling Edge Services. In Proceedings of the 2020 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), Leganes, Spain, 10–12 November 2020; pp. 155–160. [\[CrossRef\]](#)
50. Meneses, F.; Fernandes, M.; Corujo, D.; Aguiar, R.L. SliMANO: An Expandable Framework for the Management and Orchestration of End-to-end Network Slices. In Proceedings of the 2019 IEEE 8th International Conference on Cloud Networking (CloudNet), Coimbra, Portugal, 4–6 November 2019; pp. 1–6. [\[CrossRef\]](#)
51. Fernández-Fernández, A.; Colman-Meixner, C.; Ochoa-Aday, L.; Betzler, A.; Khalili, H.; Siddiqui, M.S.; Carrozzo, G.; Figuerola, S.; Nejabati, R.; Simeonidou, D. Validating a 5G-Enabled Neutral Host Framework in City-Wide Deployments. *Sensors* **2021**, *21*, 8103. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.