SUCCESS-6G: DEVISE

WP5 Deliverable E13

Project Title:	SUCCESS-6G: DEVISE
Title of Deliverable:	Initial testing and preliminary validation of service KPIs
Status-Version:	v1.0
Delivery Date:	18/02/2025
Contributors:	Allen Abishek, Ricard Vilalta, Raul Muñoz (CTTC), Maria A. Serrano (NBC), Miguel Fornell, Francisco Paredes (IDNEO)
Lead editor:	CTTC
Reviewers:	Charalampos Kalalas (CTTC)
Keywords:	Security as a Service; secure software update management; SLA compliance

Document revision history

Version	Date	Description of change
v0.1	30/09/24	Table of Contents (ToC)
v0.2	15/10/24	Content added
v0.3	30/11/24	Additional inputs
v0.4	31/01/25	Final inputs and revision
v1.0	18/02/25	Final version uploaded to the website

Disclaimer

Acknowledgment

---

[1] http://creativecommons.org/licenses/by-nc-nd/3.0/deed.en_US

Executive Summary

Security remains a paramount concern in the deployment of over-the-air (OTA) software updates, as vulnerabilities in the vehicle software can expose systems to cyber threats. SUCCESS-6G-DEVISE introduces a security-first approach by integrating Security as a Service (SECaaS), AI-enhanced threat detection, and blockchain-based integrity verification. These mechanisms ensure that software updates are authenticated, encrypted, and delivered securely, preventing unauthorized modifications and cyberattacks. Experimental results on the ADRENALINE Testbed demonstrate enhanced security enforcement, reduced threat detection times, and improved compliance with cybersecurity standards. The adoption of AI-driven security policies and network slicing for isolated update distribution solidifies SUCCESS-6G-DEVISE as a robust framework for secure vehicular software updates.

Table of Contents

List of Figures

1 Introduction

The increasing reliance on software-driven functionalities in modern vehicles necessitates a robust and secure method for delivering over-the-air (OTA) software updates. As connected and autonomous vehicle ecosystems continue to evolve, manufacturers and service providers must ensure that software updates are not only timely but also protected from cyber threats. Vehicle-to-Everything (V2X) communication is a key enabler of this transformation, facilitating seamless and reliable OTA updates while incorporating advanced security measures. However, ensuring data integrity, preventing unauthorized access, and safeguarding against cyberattacks remain critical challenges. SUCCESS-6G-DEVISE aims to address these security concerns through the integration of Software-Defined Networking (SDN), Security as a Service (SECaaS), and AI-driven threat detection.

A fundamental requirement for secure OTA software updates is ensuring end-to-end encryption and authentication mechanisms to prevent unauthorized modifications. Cellular V2X (C-V2X) technology, enabled by 5G and edge computing, enhances secure communication channels by enabling encrypted, tamper-proof data exchanges between vehicles and update servers. The SUCCESS-6G-DEVISE framework incorporates threat detection to continuously monitor OTA traffic for anomalies and potential cyber threats. Additionally, the use of network slicing ensures that OTA updates are transmitted over dedicated, isolated channels to prevent unauthorized interception and manipulation.

Beyond connectivity, cybersecurity remains a primary concern for OTA software updates. Outdated vehicle software is a prime target for cyberattacks, necessitating stringent security protocols to authenticate update sources and verify software integrity. SUCCESS-6G-DEVISE integrates Security as a Service (SECaaS) mechanisms, incorporating AI-enhanced threat detection, real-time anomaly detection, and blockchain-backed update verification. These security measures ensure that only authorized updates are deployed, mitigating risks such as firmware tampering, data breaches, and ransomware attacks.

Efficient security management is another key factor in optimizing OTA software updates. MEC-based security processing reduces the computational burden on centralized cloud infrastructure by distributing security monitoring and threat mitigation to edge nodes. This allows for real-time security assessments, rapid threat response, and proactive risk mitigation. SUCCESS-6G-DEVISE employs threat prediction models to dynamically adjust security policies based on evolving cyber threats, vehicle density, and update criticality. By leveraging federated learning and distributed intelligence, the system ensures proactive security enforcement while minimizing processing delays.

This deliverable presents primary results on the implementation and validation of secure OTA software updates within a V2X connectivity framework. Experimental evaluations conducted on the ADRENALINE Testbed demonstrate significant improvements in security enforcement, threat mitigation efficiency, and update integrity verification. The integration of SDN, SECaaS, and AI-driven security mechanisms has resulted in a scalable and adaptive solution capable of addressing the evolving cybersecurity demands of connected vehicle ecosystems.

The preliminary findings from this research highlight the potential of SUCCESS-6G in revolutionizing vehicular software update security methodologies. By leveraging cutting-edge networking and security technologies, the proposed framework ensures that vehicles remain protected against cyber threats while maintaining seamless software update deployment. As the automotive industry continues to transition towards fully connected and autonomous systems, the implementation of secure and efficient OTA update mechanisms will be instrumental in enhancing vehicle safety, data protection, and regulatory compliance.

The subsequent sections present the specific methodologies employed, experimental setup, and detailed performance evaluations of the proposed secure OTA update system, providing a comprehensive analysis of its benefits and potential industry applications.

2 Use case 2: Automated software updates for vehicles

2.1 General description and overall objectives

Over-the-air software updates are delivered remotely from a cloud-based server, through a cellular connection, to the connected vehicle with the aim of providing new features and updates to the vehicle's software systems. Such software updates may include changes to any software that controls the vehicle's physical parts or electronic signal processing system. In practice, the updates often tend to apply more to user interfaces like infotainment screens and navigation (i.e., vehicle maps). The update procedure, when performed over-the-air, enables a vehicle's performance and features to be continuously up-to-date and improved. The integration of advanced data analytics, automated and remote service delivery eliminates the need for visiting repair/service centres, while technological advancements in these updates give vehicle manufacturers the freedom to constantly "freshen up" finished products remotely. C-V2X technology plays a crucial role in the update process, enabling efficient, scalable, and seamless wireless communication between vehicles and software management platforms. Figure 1 illustrates the implementation phases for this use case.

Figure 1: Implementation phases for the automated software updates

The overall **objectives** of this use case can be summarized as follows:

- Safer and more entertaining driving experience.
- Hardware and software components maintained and updated regularly during a vehicle's lifespan, implying a slower rate of depreciation.
- Prevention of cyberattacks targeting outdated software.
- Compliance to new rules and standards.
- Lower repair costs and elimination of labour charges.
- Lower warranty costs for manufacturers and lower downtime for customers

The key **stakeholders** involved in the use case are:

- The *Mobile Network Operator (MNO)*, providing wireless connectivity between the vehicle, the edge computing infrastructure, and the vehicular software management system. The MNO is interested in optimizing the network operation by enhancing its energy efficiency and coverage, while offering novel services to accommodate more users.
- The *edge infrastructure provider*, offering and managing computational resources at the edge and supporting real-time services as well as virtualized network functions and AI-empowered algorithms for advanced computational tasks.
- The *equipment provider,* providing in-vehicle embedded devices, e.g., hardware components and sensor devices, that can be remotely reconfigured and updated.
- The *vehicular software management system,* operated by the equipment provider or vehicle manufacturer, is responsible for issuing periodically new software updates.
- The *software developers*, devising and applying data-processing modules for automated update of vehicular components' software.

- The *cloud providers* can optionally be involved, offering additional computational resources to host the service.

Note that, without loss of generality, some stakeholders may assume multiple roles or, equally, some roles may be assumed by multiple stakeholders. For instance, the MNO could also be the owner of the edge infrastructure, or an equipment provider may also be responsible for the operation of the vehicular software management system or outsource it to a third party.

2.2   User story 2.2: Over-the-air vehicular software updates with security guarantees

Over-the-air software updates deliver critical information to on-board vehicular devices. As vehicles introduce new functionalities (such as advanced driver-assist features like self-parking) and the number of connected vehicles keeps growing, automakers need to handle the regular software updates required in a secure and trustworthy way. Thus, the integration of intelligent security enforcement solutions and effective prediction/mitigation of security threats are deemed essential for the secure operation of over-the-air update service and to preserve trustworthiness. Additionally, by instantiating virtual security functions and by exploiting secure edge provisioning empowered by AI-driven capabilities, the threat risk for software updates can be further minimized.

2.3   Overall UC2 architecture and network deployments

The elaboration of Figure 2 details a system architecture specifically designed for Over-the-air (OTA) software updates, integral to the SUCCESS-6G framework. This architecture addresses the complex requirements of the use case 2. In particular, Figure 2 Proposed overall UC2 system architecture provides a high-level system architecture for OTA vehicular software updates within a robust Vehicle-to-Everything (V2X) connectivity framework, leveraging ETSI TeraFlowSDN for network automation and control. The figure illustrates the key components enabling software update dissemination to connected vehicles via 5G mobile edge computing (MEC) nodes. At the core of this system is the ETSI TeraFlowSDN Controller, which manages the network infrastructure, including the gNBs (5G base stations) and Transport Network. The NFV Orchestrator (NFV-O) enables dynamic deployment and scaling of virtualized network functions, such as Distributed User Plane Functions (D-UPF) within MEC nodes.

Each edge node (Edge Node 1 & Edge Node 2) hosts a Software Update Server, responsible for caching and distributing updates to C-V2X On-Board Units (OBU) in connected vehicles. These updates are delivered via the 5G network, passing through the transport network, controlled by the TeraFlowSDN controller. To ensure security and integrity, the system integrates a Security-as-a-Service module, providing firewall protection and secure communications for software updates. The updates originate from local cloud infrastructure, which includes 5G Core Control Plane components such as SMF (Session Management Function), AMF (Access and Mobility Management Function), and UPF (User Plane Function). The software update client within the vehicle's C-V2X OBU interacts with the Software Update Servers over the network, ensuring efficient and timely delivery of critical updates for vehicle applications.

Figure 2 Proposed overall UC2 system architecture

This architecture highlights the interplay between 5G, MEC, SDN, and V2X technologies to facilitate secure and efficient OTA software updates, enabling reliable vehicle connectivity and automation.

2.4 Facilities for use case 2

2.4.1 ADRENALINE Testbed

The ADRENALINE testbed® is an open and disaggregated SDN/NFV-enabled packet/optical transport network and edge/core cloud infrastructure for 6G, IoT/V2X and AI/ML services, constantly evolving since its creation in 2002, and reproducing operators' networks from an End to End (E2E) perspective and Data Centre Interconnect (DCI). The figure below summarizes the networking scenario of the ADRENALINE testbed, to be used for the execution of SUCCESS-6G-DEVISE.

Figure 3 ADRENALINE Testbed to be used for Use Case 2

ADRENALINE spans the access, aggregation-metro, and core segments, and includes distributed Data Centres (DCs) geographically dispersed and located at the edge or in central locations. As depicted in Figure 3, the key elements are: (1) an SDN-controlled optical network (flexi-grid DWDM photonic mesh), with 4 ROADM nodes and over 600km of amplified DWDM links. Currently, all the links of the mesh are based on amplified C-band transmission, but one of them also supports amplified flexible L-band transmission; (2) packet-optical nodes with optical pluggable transceivers, providing aggregated 400G data rates (muxponders) for transporting traffic flows between the access networks and the core central offices or data centers; (3) programmable SDN-enabled S-BVTs able to transmit multiple flows at variable dat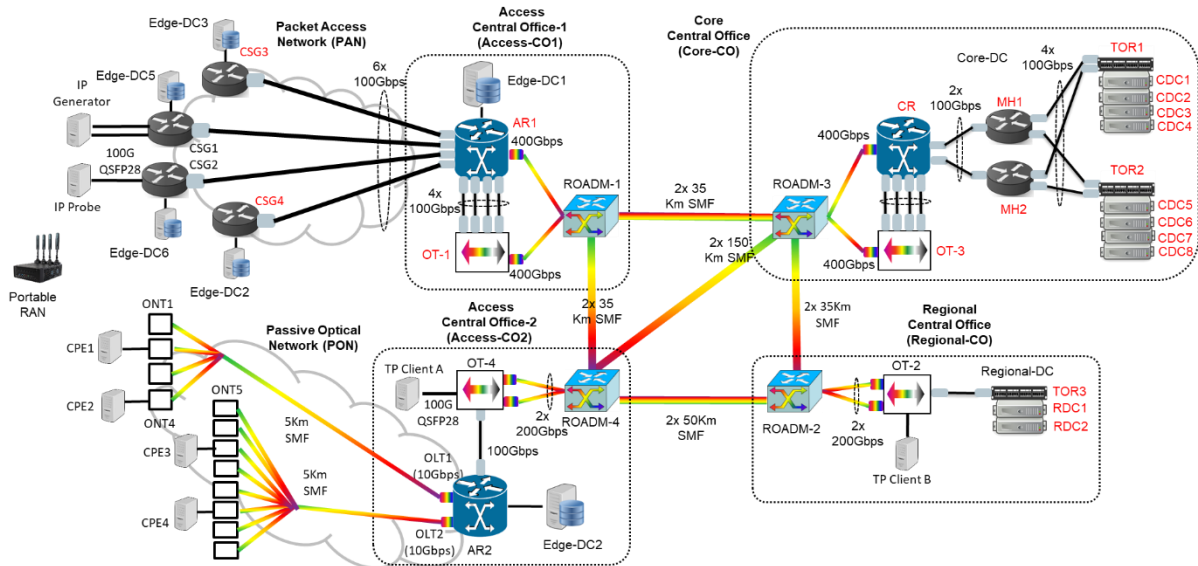a rate/reach up to 1 Tb/s; (4) a Packet Access Network (PAN) connected to the metro infrastructure with IP Cell Site Gateways (CSGs); (5) a PON tree formed by disaggregated Optical Network Terminals (ONTs), offering connectivity to several Customer Premises Equipment (CPEs). ADRENALINE also includes a Portable 5G RAN platform for testing and validation of 5G and beyond use cases.

The different access networks (i.e., PON) and the photonic mesh are managed by dedicated orchestrators and controllers (e.g., CTTC FlexOpt Optical Controller) to automatically handle the connectivity services entailing the de-/allocation of heterogeneous network resources (i.e., packet and optical devices).

The *domain-specific* controllers and orchestrators are coordinated hierarchically by the ETSI TeraFlowSDN controller, which exposes a North Bound Interface to allow interaction of resources to request network connectivity services. This service platform orchestrates the transport (optical/packet) and computing:

i)   Multi-VIM (virtualized infrastructure managers) combining OpenStack and K8s controllers for virtual machines and containers;

ii)   TeraFlowSDN controller for E2E connectivity among virtual machines, containers, and endpoints. The service platform is also in charge of managing the life-cycle of network services and network slices: i) a network service is composed of chained NFs; ii) a network slice is composed of one or several concatenated network services that deploy a set of NFs.

2.4.2   Nextcloud platform

NextCloud is an open-source private cloud platform based on a dedicated server with extensive resources, orchestrated by a Debian-based distribution called Proxmox to manage virtualization

environments, offering a simple web interface to manage both virtual machines (VMs) and containers. These machines can be of various types, but we ideally use two of them:

- KVM (Kernel-based Virtual Machine): It is used for full virtualization, allowing to run complete operating systems as VMs.

- LXC (Linux Containers): It also supports LXC containers, which are lighter and more efficient environments to run applications or services.

These virtual machines can be scaled on demand, whether in storage, memory, or network resources. NextCloud bundles core technologies such as PHP, JavaScript, HTML, CSS, MySQL, SQLite, Redis and Nginx, and provides full control over the data and privacy, offering self-host option to deploy custom applications. The following capabilities ensure that NextCloud is ideal for delivering SUCCESS-6G-DEVISE catalogue services such as vehicle data storage and Firmware Over-The-Air (FOTA). NextCloud will be fully customized to fit the needs of the project.

3 Over-the-air vehicular software updates with security guarantees: Implementation at the ADRENALINE testbed

3.1   Security as a Service

Figure 4 illustrates a sophisticated network architecture designed for service management through Software Defined Networking (SDN) and Network Function Virtualization (NFV) technologies. At the apex of this architecture is the API (NBI), the Northbound Interface, which functions as a pivotal point of access for service management, facilitating communication between the service management layers and the underlying network infrastructure.

Central to this architecture is the Security as a Service component. It includes a Slice Manager, a component that manages network slices—distinct segments of the network tailored for specific operational requirements, each with its own set of performance parameters and policies. This management is intricately connected to the SSLA & Policies Mapper, which translates service level agreements and policies into practical, enforceable rules for each network slice.

These slices are cataloged in the Slice DB (Abstracted Domain Resources), a comprehensive repository that maintains details about network slices and their corresponding resources. The Orchestrator operates in concert with the Slice DB, orchestrating the deployment and lifecycle of network services across various network segments. The Provider Mapper (SBI), or the Southbound Interface, communicates with physical and virtual network functions, translating the orchestrated service management directives into actionable tasks within the network fabric. This includes the NFV Orchestrator (NFV-O), which is responsible for the overall management of virtualized network functions, ensuring their proper instantiation, scaling, and termination.

At the foundation of the network control plane is the Transport SDN controller, which governs data plane devices such as routers and switches, enabling efficient and dynamic routing of traffic within the network. The infrastructure components, including vehicles, cellular towers, and data centers, are depicted at the bottom of the figure, highlighting the endpoints of this network architecture. These components represent the tangible elements where data and services are consumed and delivered, completing the ecosystem of this advanced network architecture.
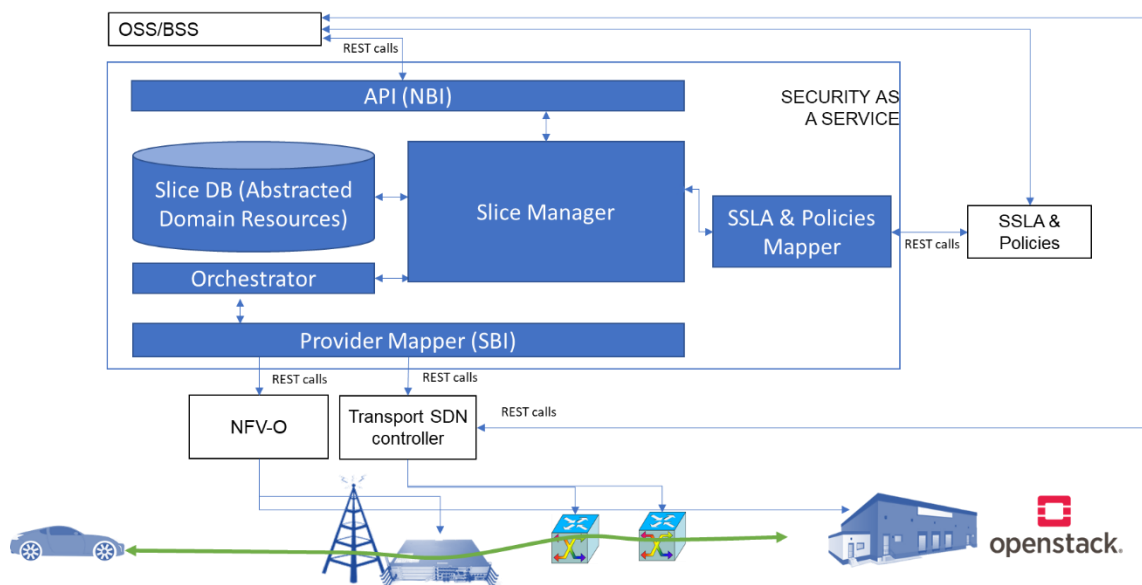

Figure 4 Proposed architecture for OTA vehicular software updates with security guarantees

3.2   Requirements and KPIs

The rapid evolution of connected vehicle technology has heightened the need for stringent security, reliability, and performance standards. As modern vehicles increasingly depend on software to manage critical functionalities—ranging from engine control and braking systems to autonomous navigation—ensuring robust mechanisms for software updates and network management is paramount. To address these concerns, this section outlines the key requirements that safeguard automotive software ecosystems against a broad spectrum of cyber threats while maintaining operational efficiency.

Each requirement detailed in the following subsections underscores a unique aspect of secure and reliable software lifecycle management. From Secure Software Update Management that focuses on authentication, authorization, and integrity checks, to Security as a Service (SECaaS) offering scalable and automated policy enforcement, these requirements collectively establish a comprehensive defense strategy. Further, Software-Defined Networking (SDN) & Network Function Virtualization (NFV) Integration and Slice-Based Service Management enable agile, resource-efficient network operations tailored to the varying needs of different vehicles and update types. Finally, Northbound and Southbound Interface Support ensures seamless interoperability across multiple platforms and vendors, facilitating smooth update distribution and real-time monitoring.

Together, these requirements form a holistic framework designed to protect and enhance the software-driven functionalities of connected vehicles. By emphasizing security, scalability, and interoperability, they serve as the foundational building blocks for maintaining trust, minimizing risks, and supporting the advanced features of next-generation automotive systems.

3.2.1   Requirements

3.2.1.1 Secure Software Update Management

Ensuring that software updates are securely distributed is critical to preventing unauthorized modifications, malware insertion, or other cybersecurity threats. The system must authenticate and authorize update sources before distribution, verifying the integrity and authenticity of software packages before deployment to vehicles. This process helps mitigate security risks associated with malicious attacks or accidental corruption of updates, ensuring that only trusted software is installed on connected vehicles.

The rationale behind this requirement is to enhance the safety and reliability of vehicular software systems. Modern vehicles rely heavily on software to control essential functionalities, including braking, navigation, and autonomous driving. Any compromise in the update process could lead to severe security breaches or operational failures. By implementing stringent authentication mechanisms and encryption methods, the system can maintain trust in over-the-air updates while complying with automotive cybersecurity standards.

3.2.1.2 Security as a Service (SECaaS)

The security framework should be designed as a scalable and adaptable service that automates security policy enforcement across different network slices. This ensures a consistent security posture while accommodating various levels of security requirements based on the specific needs of each vehicle or update type. Security policies should be dynamically adjustable to respond to evolving cyber threats and regulatory requirements without requiring significant infrastructure changes.

The rationale for integrating Security as a Service is to provide a centralized and efficient approach to managing security controls across multiple network domains. Instead of applying security configurations manually to individual network components, SECaaS enables automation, monitoring, and policy enforcement in a streamlined manner. This approach reduces operational overhead, improves compliance with cybersecurity regulations, and enhances the ability to respond quickly to emerging threats in the automotive ecosystem.

3.2.1.3 Software-Defined Networking (SDN) & Network Function Virtualization (NFV) Integration

The system must leverage SDN and NFV to optimize network resource allocation, enhance scalability, and ensure efficient update distribution. SDN enables centralized network control, while NFV provides flexible network functions that can be deployed dynamically as needed. Together, these technologies support real-time monitoring and adaptation of network conditions, ensuring that software updates reach their intended destinations with minimal latency and high reliability.

The rationale behind integrating SDN and NFV is to address the growing complexity of vehicular networks, where software updates must be transmitted efficiently over heterogeneous network infrastructures. By decoupling network control and function deployment from traditional hardware dependencies, SDN and NFV enable more agile, responsive, and cost-effective network management. This improves update delivery speeds and enhances the resilience of the network infrastructure, preventing bottlenecks and service disruptions.

3.2.1.4 Slice-Based Service Management

Network slicing must be implemented to create isolated virtual networks tailored to different operational needs. Each slice should be assigned specific security policies and Quality of Service (QoS) parameters to ensure that critical updates receive the highest priority. The slice database should store metadata about network slices, facilitating efficient resource allocation and policy enforcement.

The rationale for using network slicing is to ensure that essential updates, such as security patches or firmware upgrades, are delivered with guaranteed performance and minimal risk of interference. Vehicles operating in different regions or under varying network conditions may have different connectivity requirements. Slicing enables customized resource allocation, allowing for differentiated services while maintaining strong security and performance guarantees.

3.2.1.5 Northbound and Southbound Interface Support

The system must expose well-defined interfaces for seamless integration with external platforms. The Northbound Interface (NBI) should facilitate communication with service management applications, while the Southbound Interface (SBI) should interact with physical and virtual network functions. These interfaces should support standardized protocols to ensure compatibility across different vendors and infrastructure providers.

The rationale for exposing standardized interfaces is to promote interoperability and simplify integration with third-party applications, including automotive manufacturers, cybersecurity monitoring tools, and cloud-based orchestration platforms. By enabling flexible API interactions, the system can support automated update workflows, real-time security monitoring, and policy-driven update deployments, enhancing the overall efficiency of over-the-air update management.

3.2.2 Key Performance Indicators (KPIs)

3.2.2.1 End-to-end encryption success rate (%)

This KPI measures the effectiveness of encryption mechanisms in securing data transmission by assessing the percentage of successful encrypted transmissions without breaches.

Encryption is crucial for ensuring the confidentiality and integrity of software updates during transmission. A high success rate indicates a robust security framework, reducing the risk of data tampering or interception by malicious actors.

Target Value: At least 99.9% of transmitted updates should be successfully encrypted without any detected breaches.

3.2.2.2 Software update delivery latency (ms)

This KPI measures the time taken for an update to reach its destination from the source.

Low latency is essential for delivering time-sensitive security patches and feature updates to vehicles. Delays in updates could leave vehicles vulnerable to exploits or operational inefficiencies.

Target Value: The system should deliver critical updates within 500 ms in optimal network conditions.

3.2.2.3 SLA compliance rate (% adherence to defined policies)

This KPI evaluates how well the system adheres to predefined service level agreements (SLAs) regarding security, reliability, and performance.

Ensuring SLA compliance helps maintain trust with stakeholders and guarantees that vehicles receive updates as promised within set parameters.

Target Value: At least 98% adherence to SLAs across all update transactions.

3.2.2.4 Mean Time to Detect (MTTD)

MTTD is the time elapsed from the moment a security threat or attack occurs to the point where it is successfully detected by the security system. It represents the efficiency of threat detection mechanisms, including AI-driven inference models, log monitoring, and real-time anomaly detection.

Importance: A lower MTTD is critical for early threat identification, reducing the window of exposure and limiting potential damage. Faster detection enables security teams to respond proactively before an attacker can exploit vulnerabilities. Optimizing MTTD is particularly important in cloud-native environments, 5G networks, and real-time cybersecurity frameworks, where high-speed attack detection is required to protect critical infrastructure and services.

Target Values/Thresholds: *Ideal*: Less than 1 minute for AI-enhanced cybersecurity systems with real-time monitoring. *Acceptable*: Between 5–10 minutes in standard SOC (Security Operations Center) environments with human-assisted analysis. *Critical Risk*: Above 30 minutes, as this increases the likelihood of widespread damage, extended data breaches, and regulatory non-compliance.

3.2.2.5 Mean Time to Respond (MTTR)

MTTR is the time elapsed from the moment a security threat is detected to the complete mitigation or resolution of the issue. It includes incident triage, threat analysis, containment, mitigation, system restoration, and validation to ensure the threat has been neutralized.

Importance: A lower MTTR reduces the impact of cyberattacks by minimizing system downtime, data breaches, and business disruptions. In environments like telco networks, IoT ecosystems, and cloud-native microservices, rapid response ensures business continuity and compliance with cybersecurity policies. Automated orchestration and AI-driven remediation significantly improve MTTR by enabling real-time threat containment and recovery.

Target Values/Thresholds: *Ideal*: Under 5 minutes for automated threat response mechanisms (e.g., SDN-controlled ACL updates, automated patching). *Acceptable*: 30–60 minutes in cases where manual intervention or forensic analysis is required. *Critical Risk*: Above 2 hours, as prolonged response times increase financial losses, system disruptions, and potential data exfiltration risks.

3.3 UC2 Architecture and network deployment

Figure 5 illustrates a comprehensive architecture for enabling software updates in a Cellular Vehicle-to-Everything (C-V2X) environment using 5G, Multi-access Edge Computing (MEC), Network Function Virtualization (NFV), and Software-Defined Networking (SDN). The system is structured into multiple layers to ensure efficient data processing, low-latency communication, and secure software updates for in-vehicle systems. These layers include the Orchestration Layer, 5G Core Control Plane (Local Cloud), MEC Layer (Edge Nodes), Network Infrastructure Layer, and In-Vehicle Layer. Among these, the dark blue elements play a crucial role in orchestrating network functions, ensuring security, and enabling seamless data transmission.
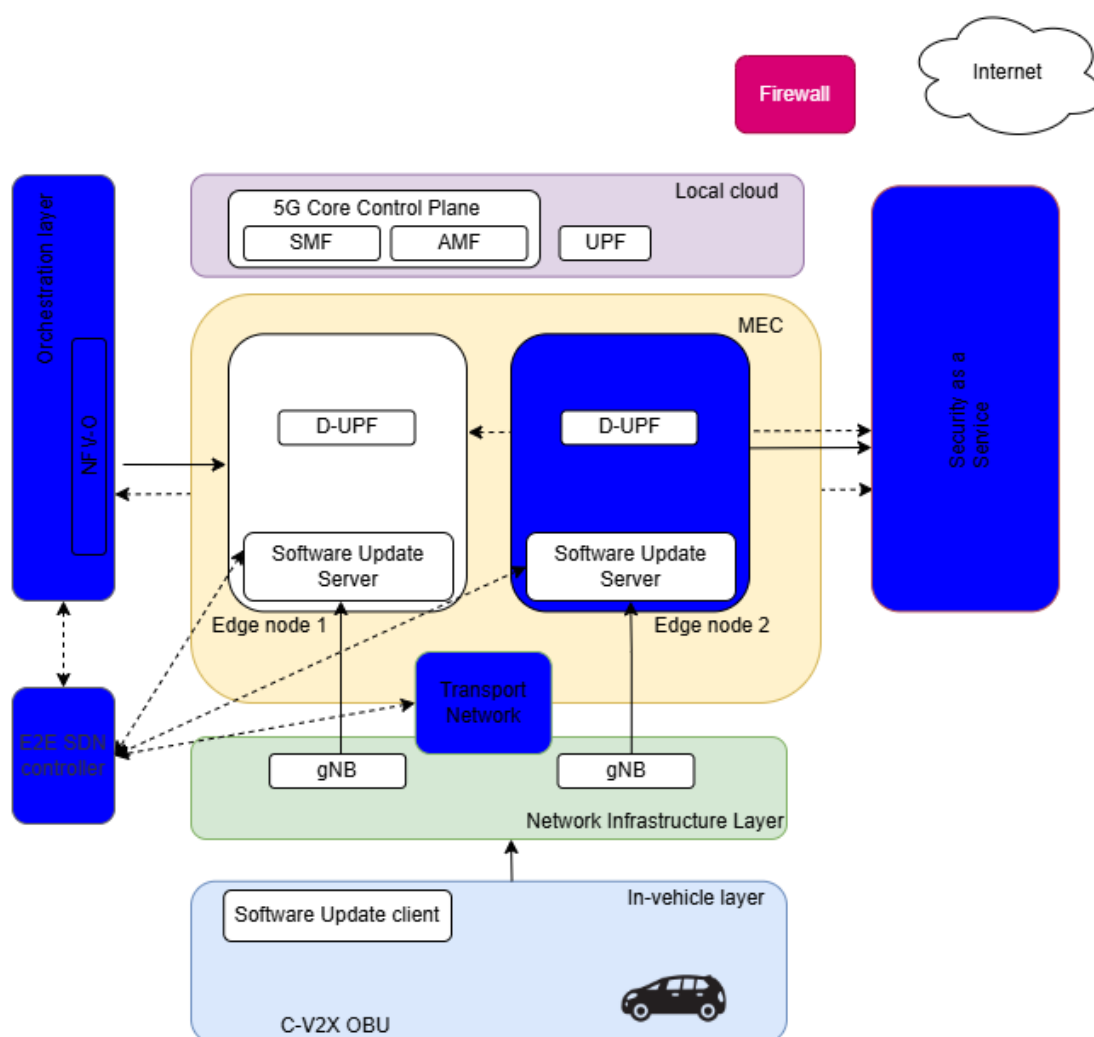
Figure 5 Instantiation of SUCCESS-6G-DEVISE architecture for OTA vehicular software updates with security guarantees

At the top of the architecture, the Orchestration Layer is responsible for managing the deployment and lifecycle of network functions using NFV and SDN technologies. The NFV Orchestrator (NFV-O) oversees the instantiation, scaling, and management of Distributed User Plane Functions (D-UPFs), which are deployed closer to the edge to reduce latency. Alongside NFV-O, the End-to-End SDN Controller provides centralized control over the transport network, ensuring efficient data routing and dynamic policy enforcement. This integration of NFV and SDN allows for adaptive network resource management, dynamic traffic steering, and enhanced service reliability, which are essential for delivering real-time software updates to connected vehicles.

The Multi-access Edge Computing (MEC) layer is a critical part of the architecture, hosting D-UPFs and Software Update Servers at edge nodes. The D-UPF (Distributed User Plane Function) is responsible for offloading data traffic from the 5G core, allowing localized data processing and reducing dependency on central cloud infrastructure. This offloading mechanism enhances the responsiveness of V2X applications, including OTA software updates. The Software Update Servers deployed at the edge nodes ensure that vehicles receive software patches and security updates efficiently, minimizing downtime and optimizing network resource utilization. By leveraging MEC, the system reduces latency, improves update delivery speed, and enhances overall network performance.

To protect the integrity and confidentiality of software updates and vehicular communications, the architecture incorporates a Security as a Service framework. This module ensures that all software updates are encrypted, verified, and transmitted securely to prevent cybersecurity threats such as man-in-the-middle attacks, data tampering, and unauthorized access. The security framework is tightly

integrated with the MEC layer and orchestration systems, ensuring continuous threat monitoring, anomaly detection, and compliance with cybersecurity policies. By implementing security measures at both the orchestration and edge levels, the system guarantees the safety and reliability of C-V2X applications.

At the transport network level, gNBs (Next-Generation Base Stations) facilitate communication between vehicles and the edge computing infrastructure. These base stations provide high-speed, low-latency connectivity to ensure that software update clients in vehicles (C-V2X OBUs) receive real-time updates. The transport network is dynamically managed by the E2E SDN controller, which optimizes routing paths based on network congestion, traffic demand, and security policies. Within the In-Vehicle Layer, the Software Update Client ensures that received updates are properly installed, validated, and synchronized with the vehicle's onboard systems. This integration ensures that connected vehicles remain updated with the latest firmware, security patches, and feature enhancements without requiring physical intervention.

3.4   Exposed interfaces

We have developed two different interfaces, dedicated to attack detector and attack mitigator.

3.4.1   Attack Detector

The provided Protobuf definition outlines an AttackDetector service within the "attack_detector" package, leveraging proto3, the latest version of the Protobuf syntax. This service is responsible for configuring and managing an attack detection system that integrates with Elasticsearch, Kafka, and machine learning-based threat detection.

The AttackDetector service is defined as a gRPC-based service that enables clients to configure and retrieve attack detection settings remotely. It consists of six Remote Procedure Call (RPC) methods, each of which facilitates specific operations. The ConfigureDetector method allows clients to configure the attack detector using a DetectorConfig message, which includes Elasticsearch and Kafka settings along with a minimum machine learning confidence level for attack detection. The GetDetectorConfiguration method enables clients to retrieve the current configuration settings, ensuring they can verify or modify existing parameters when necessary. Additionally, the ConfigureAttack method allows users to define attack specifications, while ListConfiguredAttacks provides a list of all active attack configurations. To retrieve details of a specific attack, clients can use the GetConfiguredAttack method, which takes an AttackId as input and returns the corresponding AttackSpecs. Lastly, the DeconfigureAttack method allows for the removal of an attack configuration, ensuring that outdated or irrelevant attack profiles do not persist in the system.

Each RPC function follows a request-response model, meaning a client sends a request message and expects a structured response. This approach is particularly useful for gRPC-based microservices, where real-time communication is essential. The service relies on various Protobuf messages to facilitate structured data exchange. The Empty message serves as a placeholder for functions that do not require parameters or return values, similar to an empty JSON object ({}). The ElasticSearch message defines the configuration for an Elasticsearch database, which the attack detector uses for storing and querying attack data. It includes a list of hosts, authentication details, and an index name, ensuring secure and efficient interactions with the database. Similarly, the Kafka message encapsulates settings for Kafka-based event streaming, including broker hosts, authentication credentials, consumer group identifiers, and topic names for sending and receiving messages. This configuration enables the attack detector to process real-time threat intelligence streams efficiently.

The DetectorConfig message acts as a central configuration entity, combining Elasticsearch and Kafka settings with a minimum machine learning confidence level. This confidence level determines the threshold at which the system considers a potential attack to be valid, balancing false positives and false negatives effectively. The attack identification process relies on the AttackId message, which contains a unique attack UUID and a probability threshold. This allows the system to manage and track

individual attack instances with a predefined level of certainty. Furthermore, the AttackSpecs message extends AttackId by introducing an additional parameter: the minimum confidence level for detection. This ensures that the attack detection model can be fine-tuned to only trigger alerts for high-confidence threats, reducing noise in the system. To facilitate bulk retrieval of attack configurations, the ListAttackSpecs message aggregates multiple AttackSpecs entries into a structured list.

The use of Protobuf in this architecture provides several key advantages. Firstly, Protobuf's binary format ensures efficient serialization, reducing both the size and transmission time of messages compared to traditional text-based formats. This efficiency is particularly beneficial in network security and attack detection, where large volumes of streaming data must be processed in real time. Secondly, Protobuf allows for schema evolution without breaking compatibility. Fields can be added or deprecated without disrupting existing services, making it easier to introduce new attack detection parameters as threats evolve. Thirdly, strongly typed structures in Protobuf reduce parsing errors and improve data integrity, ensuring that attack configurations remain accurate and consistent across different components of the system.

Furthermore, Protobuf is optimized for gRPC, which enables efficient and secure remote procedure calls over HTTP/2. This allows the AttackDetector service to scale effectively, leveraging features such as bidirectional streaming, load balancing, and authentication. The combination of Protobuf, Kafka, and Elasticsearch in this architecture creates a robust framework for real-time attack detection and response. By enabling efficient configuration, retrieval, and deconfiguration of attack profiles, this system ensures that security measures remain adaptive and responsive to emerging threats.

```
syntax = "proto3";
package attack_detector;

service AttackDetector {
    rpc ConfigureDetector      (DetectorConfig) returns (    Empty        ) {}
    rpc GetDetectorConfiguration(Empty        ) returns (    DetectorConfig ) {}
    rpc ConfigureAttack        (AttackSpecs  ) returns (    Empty        ) {}
    rpc ListConfiguredAttacks  (Empty        ) returns (    ListAttackSpecs) {}
    rpc GetConfiguredAttack    (AttackId     ) returns (    AttackSpecs    ) {}
    rpc DeconfigureAttack      (AttackId     ) returns (    Empty        ) {}
}

message Empty {}

message ElasticSearch {
    repeated string hosts = 1;
    string authentication = 2; // example: "api_key:<API_KEY>"
    string index         = 3;
}

message Kafka {
    repeated string hosts    = 1;
    string authentication    = 2; // example: "api_key:<API_KEY>"
    string consumer_group_id = 3;
    string consumer_topic    = 4;
    string producer_topic    = 5;
}

message DetectorConfig {
    ElasticSearch elasticsearch = 1;
    Kafka         kafka         = 2;
```

```
            float min_ml_confidence_level = 3;
        }

        message AttackId {
            string attack_uuid = 1;
            float probability_threshold = 2;
        }

        message AttackSpecs {
            AttackId attack_id       = 1;
            float    min_confidence_level = 2;
        }

        message ListAttackSpecs {
            repeated AttackSpecs attack_specs_list = 1;
        }
```

3.4.2 Attack Mitigator

RFC 8519, titled "YANG Data Model for Network Access Control Lists (ACLs)," defines a standardized YANG 1.1 data model for configuring and managing ACLs on network devices. An Access Control List (ACL) is an ordered set of rules, known as Access Control Entries (ACEs), that determine how packets are processed based on specific match criteria and corresponding actions. The YANG model provides a structured and vendor-neutral way to configure ACLs, ensuring interoperability across different network platforms.

The ACL model is structured as a hierarchical YANG schema that allows network administrators to define ACLs in a standardized manner. At the top level, the acls container holds multiple ACL configurations. Each ACL is uniquely identified by a name and a type, which specifies whether it applies to IPv4, IPv6, or Ethernet traffic.

Within each ACL, the aces container holds a list of Access Control Entries (ACEs). Each ACE contains:

- A unique name for identification.
- Match conditions that define criteria based on packet headers, organized into:
- Layer 2 (l2): Matches Ethernet fields, such as source and destination MAC addresses.
- Layer 3 (l3): Matches IP packet fields, such as source and destination IP addresses.
- Layer 4 (l4): Matches transport-layer fields, such as TCP/UDP ports.
- Actions specifying how matching packets should be handled, such as permit, deny, or log.

One of the most critical aspects of the YANG ACL model is its ability to bind ACLs to network interfaces. The model defines an acl-interfaces structure that enables administrators to attach ACLs to specific network interfaces, including physical and logical interfaces. This ensures that ACLs can be applied in a structured and policy-driven manner to control both ingress (incoming) and egress (outgoing) traffic. The interface attachments can apply to various network devices, such as routers, switches, and firewalls.

The YANG model supports different types of interfaces where ACLs can be enforced:

- Physical Interfaces – Ethernet, fiber, or any other hardware network interface.
- Logical Interfaces – VLAN interfaces, virtual tunnel interfaces, or sub-interfaces.
- Software-Defined Interfaces – Interfaces used in SDN-based environments where ACL policies are dynamically enforced.
- By defining ACLs within the YANG model, administrators can programmatically attach them to these interfaces, ensuring centralized and automated security policies across network environments.

The YANG ACL model also includes mechanisms to collect and report statistics. Each ACE can have associated counters tracking packet matches, drops, and forwarding actions. These statistics allow administrators to monitor ACL effectiveness, optimize rule sets, and troubleshoot network security policies efficiently.

The YANG-based ACL model is designed to be extensible, allowing vendors to augment the base model with additional capabilities. For instance, vendors can introduce custom match conditions, logging mechanisms, or policy-based ACL configurations while still maintaining compatibility with the standard model. Additionally, the model includes feature statements that enable devices to advertise the specific ACL capabilities they support.

3.5   Workflow

The entire architecture of the CLA platform can be seen in Figure 6. The functionality of the CLA platform begins with the presumption that a connected car has already established a Protocol Data Unit (PDU) session between itself and the OTA server via the Telecom Network. Then, the user plane packets from the vertical service (e.g: IP audio calls, music, video-streaming, OTA updates, etc.) are transmitted from the connected car through the gNB to the UPF located in the Mobile Edge Compute (MEC). Then, the UPF forwards the data plane packets to the CLA platform. The components in the CLA platform determine if incoming packets are an attack or not, if the CLA platform determines the Network is being exposed to a PoD attack, the CLA platform creates the necessary ACL which is sent to the TFS and finally the ACL is pushed into the Network Element (switch or router or firewall or server) that will block the malicious packets.
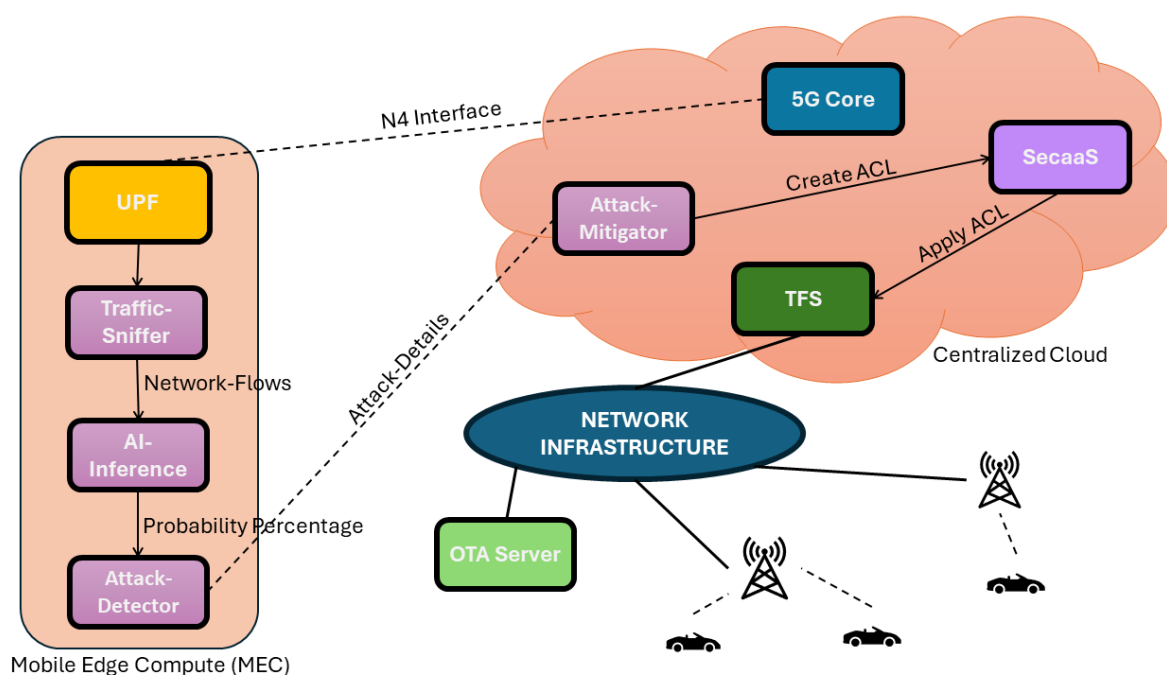

Figure 6 The specific architecture of SECaaS platform

Figure 7 illustrates a sequence diagram representing the process of network traffic analysis, attack detection, and mitigation using Elasticsearch, Kafka, AI inference, and multiple security components. The process begins with Traffic Sniffer capturing network flows and writing them into an Elasticsearch database. This step ensures that network data is logged and stored for further processing. Once the network flow is created and written, the Traffic Sniffer publishes the document ID of the network flow to Kafka, a real-time event streaming platform. This enables subsequent components in the pipeline to access and analyze the network data efficiently.

Following this, the AI Inference module extracts the document ID of the network flow from Kafka and processes it to determine the probability of an attack. Once the inference process is completed, the AI system publishes the attack probability back to Kafka, making it available for Attack Detector to retrieve. The Attack Detector listens to Kafka for incoming probabilities and assesses the likelihood of an attack occurring. If a potential attack is detected, the Attack Mitigator takes action by publishing the necessary information to create an Access Control List (ACL). This ACL is intended to restrict or control access to affected network entities based on the detected attack patterns.

The Attack Mitigator listens to Kafka for further information required to generate the ACL in JSON format. Once all required data is gathered, the TFS (presumably a security enforcement or traffic filtering system) creates and pushes the ACL. The final step involves TFS pushing the ACL into the Device, ensuring that security policies are enforced at the endpoint level. This process establishes a closed-loop security framework, where threats are detected, evaluated, and mitigated in an automated manner.
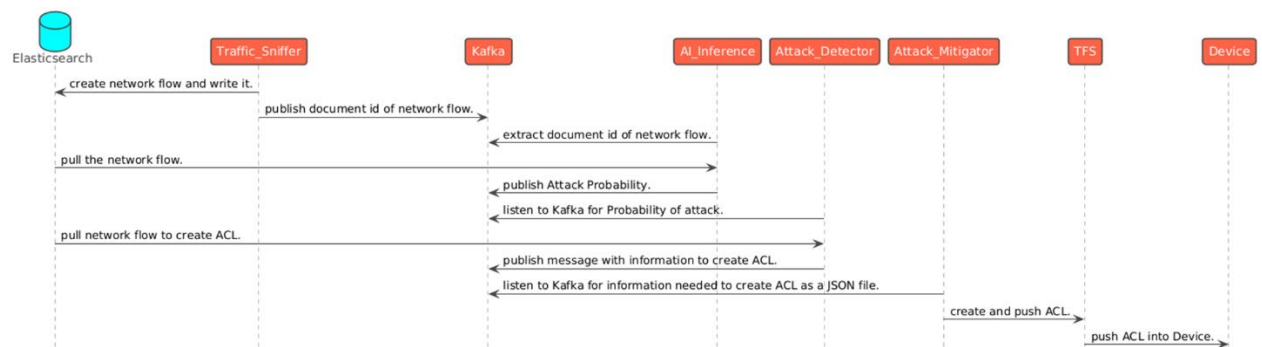

Figure 7 Sequence diagram

3.6 Preliminary experimental validation of the functionalities

Upon running the evaluation of the CLA platform, the results are shown below, The Traffic Sniffer publishes into Kafka the document id of the network flow it creates as can be seen in Figure 8.
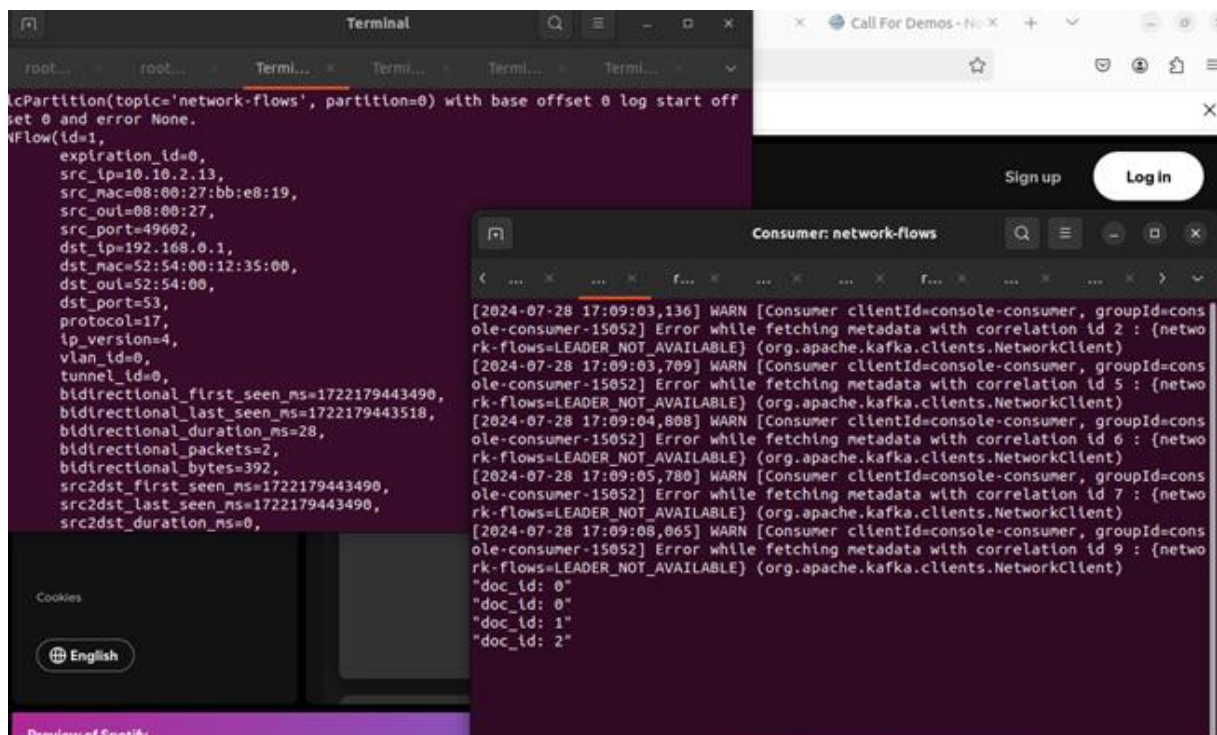

Figure 8 Creating network flows and publishing into Kafka

The NBI is used to extract network flows from the Elasticsearch database as shown in Figure 9.

```
cache
Available document IDs: [23, 24, 25, 26, 32, 33, 34, 35, 2, 3, 4, 5, 19, 20, 21, 28, 29, 30, 31, 6, 7, 0, 1, 22, 27, 8, 9, 10, 11, 1
2, 13, 14, 15, 16, 17, 18, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48]
Enter the document ID you want to retrieve: 25
Network Flow Details:
data {
  features {
    doc_id: 25
    source_oui: "08:00:27"
    destination_oui: "52:54:00"
    ip_version: 4
    bidirectional_first_seen_ms: 1722182462876
    bidirectional_last_seen_ms: 1722182463466
    bidirectional_duration_ms: 590
    src2dst_first_seen_ms: 1722182462876
    src2dst_last_seen_ms: 1722182463466
    src2dst_duration_ms: 590
    src2dst_packets: 13
    src2dst_bytes: 2321
    dst2src_first_seen_ms: 1722182462893
    dst2src_last_seen_ms: 1722182463466
    dst2src_duration_ms: 573
    dst2src_packets: 9
    dst2src_bytes: 4256
    application_name: "TLS"
    application_category_name: "Web"
    application_confidence: 6
    requested_server_name: "ab.chatgpt.com"
    client_fingerprint: "b5001237acdf006056b409cc433726b0"
    server_fingerprint: "eb1d94daa7e0344597e756a1fb6e7054"
  }
  metadata {
    timestamp: "2024-07-28T18:03:57.467344"
    monitored_device: "Linux"
    monitored_device: "e2ecla"
    monitored_device: "6.5.0-44-generic"
    monitored_device: "#44-22.04.1-Ubuntu SMP PREEMPT_DYNAMIC Tue Jun 18 14:36:16 UTC 2"
    monitored_device: "x86_64"
```

Figure 9 Extract Network Flow from Traffic Sniffer

The next step in the CLA process is the AI inference which extracts network flow from elasticsearch and calculates the probability of it being a PoD attack using the Random Forest Classifier machine learning algorithm, this demonstration can be seen in Figure 10. The calculated probability is published as a message is published in Kafka.

```
of_death probability of: 0.04, has a mail spam probability of 0"}
{"message": "The flow with ID: 7 has a DoS attack probability of: 0, has a Ping_
of_death probability of: 0.04, has a mail spam probability of 0"}
{"message": "The flow with ID: 8 has a DoS attack probability of: 0, has a Ping_
of_death probability of: 0.18, has a mail spam probability of 0"}
{"message": "The flow with ID: 9 has a DoS attack probability of: 0, has a Ping_
of_death probability of: 0.04, has a mail spam probability of 0"}
{"message": "The flow with ID: 10 has a DoS attack probability of: 0, has a Ping
of_death probability of: 0.04, has a mail spam probability of 0"}
```

Figure 10 AI Inference message

The Attack-Detector steps in and is configured remotely if the probability published by the AI-Inference is an attack or not, this is seen in Figure 11.

```
Consumer: attack-details

{"Attack Type": "Ping_of_Death", "Timestamp": "2024-07-28T15:13:19.631319", "det
ails": {"doc_id": 8, "source_port": 44790, "destination_ip": "104.16.100.215", "
destination_port": 443, "protocol": 6, "document_timestamp": "2024-07-28T17:13:1
2.866307"}}
```

Figure 11 Attack Details Message

The user can set a threshold limit dynamically such that any probability value above the threshold is classified as an attack and published as an attack in Kafka, this can be seen in Figure 12.

```
root@sniffer-poc:~/complete-sniffer-poc/Attack-Detector-Architecture/itr-1# pytl
on3 client.py
Enter min_ml_confidence_level (0 to 1): 0.8
ConfigureDetector response:
GetDetectorConfiguration response: elasticsearch {
  hosts: "localhost:9200"
  authentication: "api_key:dummy"
  index: "network-flows"
}
kafka {
  hosts: "localhost:9092"
  authentication: "api_key:dummy"
  consumer_group_id: "group1"
  consumer_topic: "probability-flows"
  producer_topic: "attack-details"
}
min_ml_confidence_level: 0.8
```

Figure 12 Attack Detector Configuration

The next step is for the Attack-Mitigator to create the ACL and post it into the TFS controller, which can be seen below in Figure 13 and Figure 14.
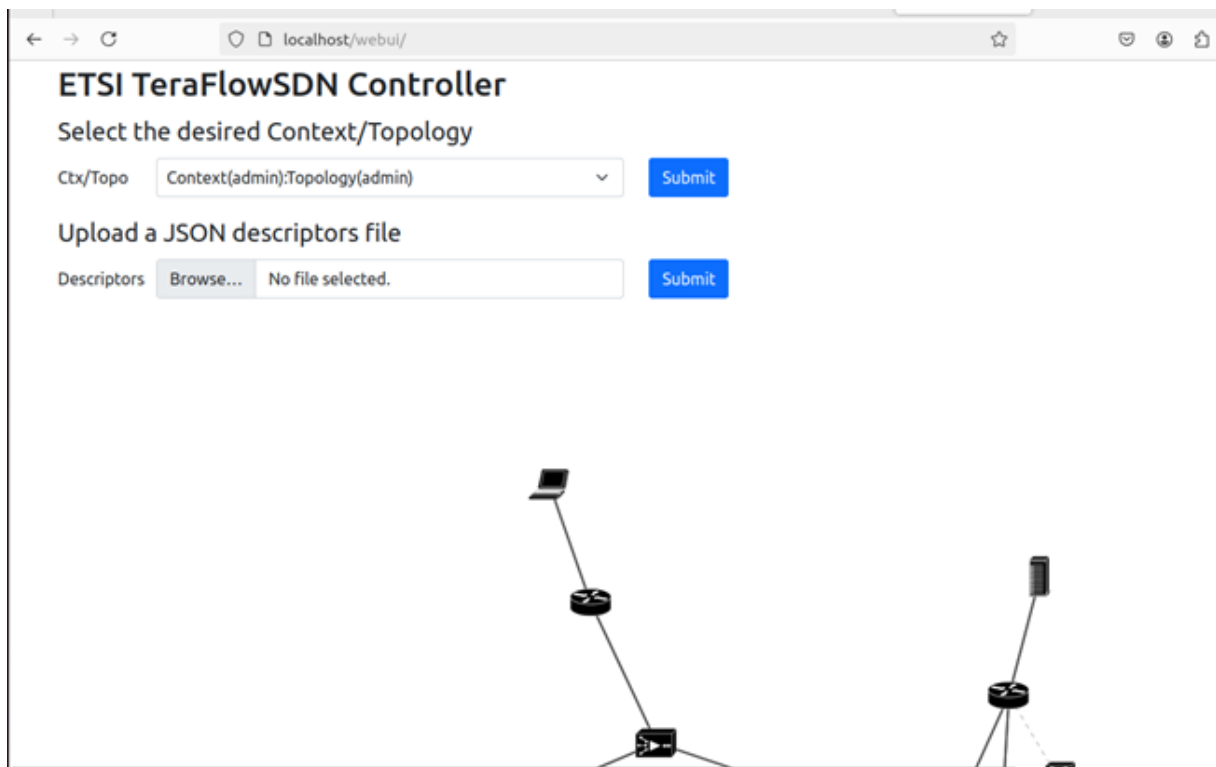

Figure 13 TFS Web UI

Received message: {'Attack Type': 'Ping_of_Death', 'Timestamp': '2024-07-28T16:03:06.267690', 'details': {'doc_id' : 3, 'source_port': 49474, 'destination_ip': '104.18.13.201', 'destination_port': 443, 'protocol': 6, 'document_ti mestamp': '2024-07-28T18:03:03.990214'}}
ACL posted at: 2024-07-28 18:03:08
Post response: {}

Document timestamp: 2024-07-28T18:03:03.990214, Time difference: 4.133747 seconds
Received message: {'Attack Type': 'Ping_of_Death', 'Timestamp': '2024-07-28T16:04:02.481628', 'details': {'doc_id' : 23, 'source_port': 47614, 'destination_ip': '104.18.13.201', 'destination_port': 443, 'protocol': 17, 'document_ timestamp': '2024-07-28T18:03:57.204019'}}
ACL posted at: 2024-07-28 18:04:04
Post response: {}

Document timestamp: 2024-07-28T18:03:57.204019, Time difference: 6.996824 seconds
Received message: {'Attack Type': 'Ping_of_Death', 'Timestamp': '2024-07-28T16:04:12.593039', 'details': {'doc_id' : 28, 'source_port': 42466, 'destination_ip': '104.18.31.2', 'destination_port': 443, 'protocol': 17, 'document_ti mestamp': '2024-07-28T18:04:08.476605'}}

Figure 14 POST ACL into TFS

4 Use case 2 proof-of-concept (PoC)

4.1 Phase 1: PoC at Nextcloud

To carry out phase 1 testing, a test setup has been provided where the entire team can work on the same devices. This has the following advantages.

The acquisition of vehicular data requires tools to facilitate development work. For this, logs were created from the data provided by the sensors, which we called Datasets. These corresponded to the acquired analog values along with their respective timestamps. Figure 15 shows the distinction between the actual implemented architecture and the development architecture.

Remote Collaboration: Developers can access hardware from anywhere in the world, allowing them to work on the project regardless of their physical location.

Reduced operational costs: Only one set of hardware is needed, which can be shared by all developers. This reduces the costs of purchasing multiple devices for each team member.

More efficient testing: Hardware can be accessed 24/7, allowing for continuous testing and debugging, even when developers are in different time zones.

Ease of maintenance and updating: Updates to the operating system or applications built for the project can be performed centrally and instantly available to all developers.

Scalability: If more developers need to work on a project, access to hardware can be easily scaled without the need to purchase more physical equipment.

Security and Access Control: Security controls can be implemented to restrict access to hardware to only those who need it, ensuring that sensitive or confidential data is protected.
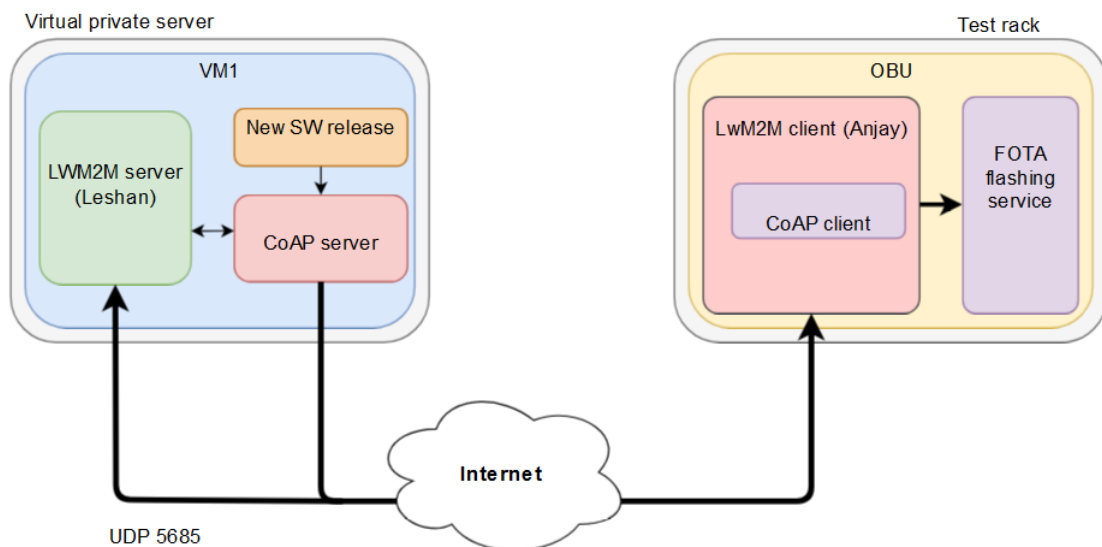

Figure 15: UC2 phase 1: development architecture

4.1.1 On-board unit test setup

The on-board unit (OBU) developed, called VMAX, is based on an aarch64 architecture, as opposed to the PCs used for development, which is based on the X86-64 architecture.

```
vmaxsetup@vmax:~ $ ssh -oHostKeyAlgorithms=+ssh-rsa root@192.168.2.4
root@192.168.2.4's password:
~ # uname -a
Linux ag215scnaa 4.14.206-perf #1 SMP PREEMPT Tue Dec 12 09:58:25 UTC 2023 aarch64 GNU/Linux
```

This means our code must undergo a process known as cross-compilation, which involves compiling a software project on a machine with one architecture, in this case, X86-64, so that it can run on a machine with a different architecture. For this reason, a test setup is necessary for the cross-compiled code. This setup consists of a Raspberry Pi connected to two device units equipped with a C-V2X PC5 interface for radio communications, linked via an Ethernet switch (see Figure 16). The Raspberry Pi is accessible over the Internet through a TLS reverse tunnel to a dedicated server managed by IDNEO, and from there, access to the VMAX units is provided via serial and/or local network interfaces.

Alternatively, the setup also allows access to the VMAX units through a Uu cellular connection, either via access permissions granted to the public IP of the IDNEO dedicated server or through other permanent TLS reverse tunnels.
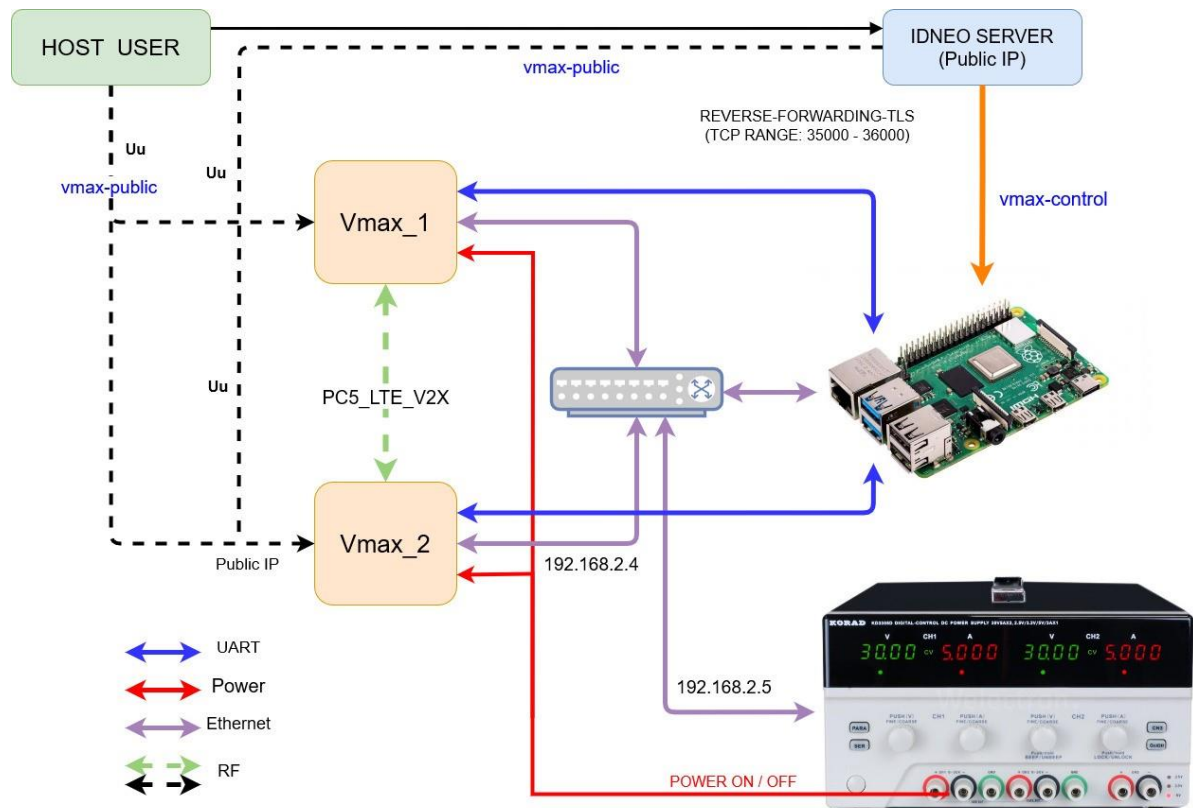

Figure 16: Test setup

In Figure 16, it is possible to identify the access interfaces that allow remote control of the Vmax. The control that this setup allows with respect to each Vmax, goes from turning it on or off, through a network-managed power source or serial interfaces that allow debugging from the Vmax boot start-up from the Raspberry PI.

4.1.2 TLS Reverse Tunnel

This tunnel allows to create a permanent connection between the Raspberry Pi, which is located behind a local network, and the IDNEO server. To run this tunnel, the openssl TLS client is used in reverse forwarding mode, which establishes a port on the remote host and redirects it locally to the Raspberry Pi TLS server. To do this, a systemd service is created that executes the corresponding command and manages the recovering process.

```
* backdoor.service - Backdoor Service
     Loaded: loaded (/etc/systemd/system/backdoor.service; enabled; preset: enabled)
     Active: active (running) since Thu 2024-10-10 15:34:38 CEST; 6min ago
      Until: Thu 2024-10-10 18:34:38 CEST; 2h 53min left
   Main PID: 28106 (bash)
      Tasks: 2 (limit: 8734)
        CPU: 136ms
     CGroup: /system.slice/backdoor.service
             |-28106 /bin/bash /usr/bin/backdoor.sh
             `-28108 /usr/bin/ssh -N -R 35990:localhost:22 ficosa-root@193.70.33.60 -p 27022

oct 10 15:34:38 vmax systemd[1]: Started backdoor.service - Backdoor Service.
oct 10 15:34:38 vmax bash[28106]: New port:  35990
```

On the IDNEO server side, this TLS client has automatic access because the public key of the client host, Raspberry Pi, is in the list of authorized keys on the TLS server, so no user and password authentication is needed.

4.1.3   Firewall on the on-board unit

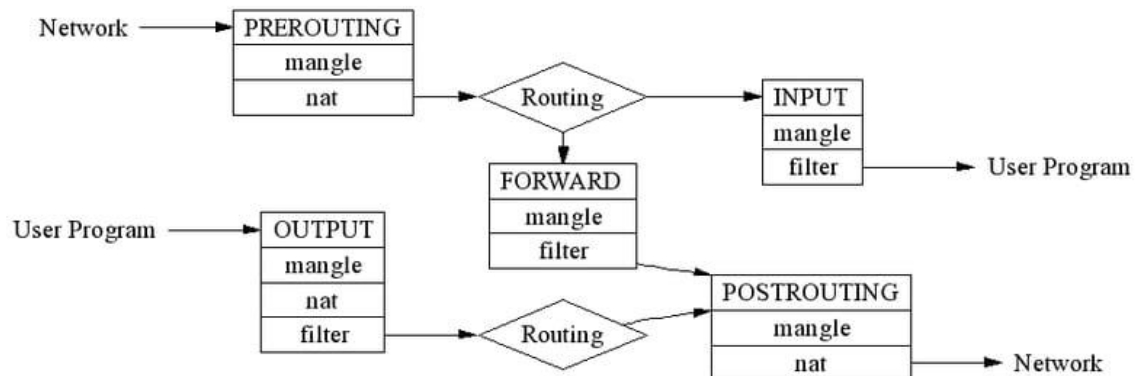The iptables architecture relies on these three components working together:

- **Tables:** Organize rules by function (e.g., filtering, mangling).
- **Chains:** Define the path packets take and the order in which rules are applied.
- **Rules:** Specify the criteria for matching packets and the actions to take (e.g., ACCEPT, DROP, REJECT).

The filter table is the default table used to filter packets. Its rules determine whether a packet should be accepted or rejected. This table contains the core chains used for traffic management: INPUT, OUTPUT, and FORWARD.

INPUT: This chain handles incoming traffic destined for the local system itself. When a packet arrives at a network interface and is intended for the device, it is processed by the rules within the INPUT chain.

OUTPUT: This chain handles outgoing traffic originating from the local system. Any packets generated by the device and sent out through a network interface are processed by the rules in the OUTPUT chain.

FORWARD: This chain handles traffic that is neither destined for nor originating from the local system. It is used for packets that are being routed or forwarded to other networks or machines. This chain is essential for devices acting as routers or gateways.

In the OBU, restrictive input rules have been configured, which are automatically configured at system startup. OBUs have access to the public network through the Quectel AG550 modem. In this modem, incoming and outgoing traffic rules are configured for network management. By configuring the INPUT and OUTPUT chains to discard unauthorized traffic and the FORWARDING chain to only allow access to port 22 (TLS server of the AG215 AP) to certain configured public IP addresses, a filter is achieved for all traffic that is not allowed access.

```
/ # iptables -L -t filter
Chain INPUT (policy DROP)
target      prot opt source              destination
ACCEPT      icmp --  anywhere            anywhere
ACCEPT      all  --  anywhere            anywhere
ACCEPT      all  --  192.168.225.0/30    192.168.225.0/30

Chain FORWARD (policy ACCEPT)
target      prot opt source              destination
ACCEPT      tcp  --  193.70.33.60        anywhere            tcp dpt:ssh
ACCEPT      tcp  --  88.98.97.234        anywhere            tcp dpt:ssh
ACCEPT      tcp  --  46.136.172.101      anywhere            tcp dpt:ssh
DROP        tcp  --  anywhere            anywhere            tcp dpt:ssh

Chain OUTPUT (policy DROP)
target      prot opt source              destination
ACCEPT      icmp --  anywhere            anywhere
ACCEPT      all  --  anywhere            anywhere
ACCEPT      all  --  192.168.225.0/30    192.168.225.0/30
```

This configuration allows to avoid brute force attacks since the TLS server will only serve the clients that the iptables rules allow.

4.2   Phase 2: PoC at Castellolí

To deploy phase 2 of the tests, all software, both server-side and OBU software, has been migrated to their final locations: the Castellolí server and the IDNEO-developed OBU installed in the test vehicle. Figure 17 details the phase 2 architecture.
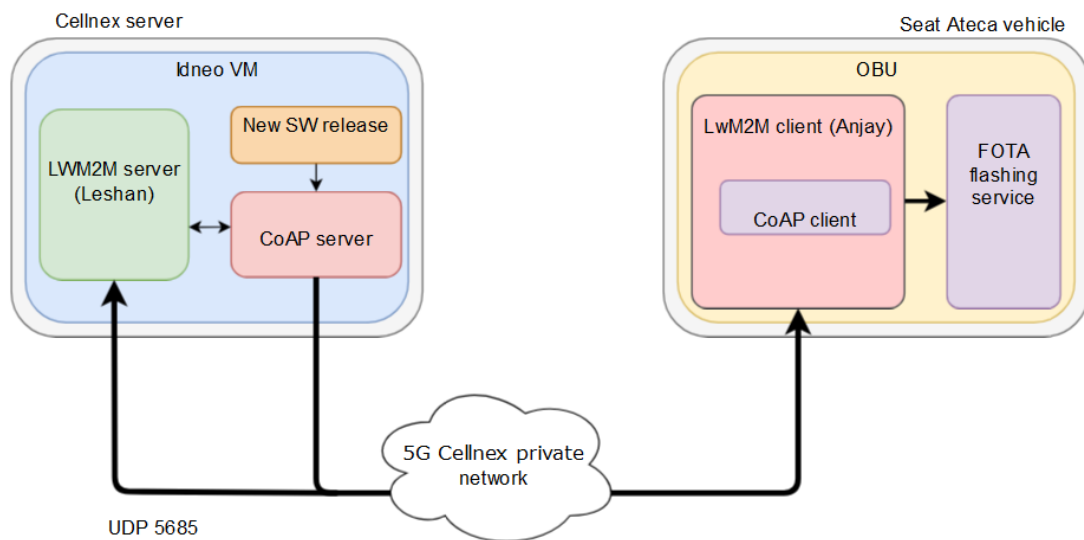
Figure 17: UC2 phase 2: real architecture (Castelloli)

This tool simplifies development to the extent that accessing the vehicle for algorithm testing was no longer necessary, as real values were used for testing, albeit with different timestamps.

5 Conclusions

Ensuring the security of OTA vehicular software updates is essential to protect vehicles from cyber threats and unauthorized modifications. The SUCCESS-6G-DEVISE framework demonstrates the effectiveness of AI-driven security policies, network slicing, and SECaaS in mitigating risks associated with OTA updates. The validation results show a significant reduction in security vulnerabilities, enhancing overall trust in the software update process. As cybersecurity threats continue to evolve, ongoing research into AI-driven threat detection, blockchain-based verification, and real-time anomaly monitoring will be critical in maintaining the integrity and resilience of vehicular software updates.

One of the key challenges in secure OTA updates is maintaining end-to-end encryption while minimizing performance trade-offs. The use of blockchain-based verification techniques ensures that software packages remain tamper-proof throughout the update process. By leveraging distributed ledger technology, updates can be validated across multiple nodes, reducing the risk of data breaches and unauthorized alterations.

Additionally, integrating AI-driven threat detection mechanisms enhances the real-time security posture of the OTA ecosystem. By continuously monitoring network activity and analyzing behavioral anomalies, the system can proactively identify and mitigate potential cyber threats. Future advancements should focus on expanding adaptive security frameworks that dynamically respond to emerging attack vectors, ensuring robust protection against evolving cyber risks.

6 References

[1] Abishek A, Vilalta R, Gifre L, Alemany P, Manso C, Casellas R, Martínez R, Muñoz R. Network Extensions to Support Robust Secured and Efficient Connectivity Services for V2X Scenario. In2024 24th International Conference on Transparent Optical Networks (ICTON) 2024 Jul 14 (pp. 1-4). IEEE.

[2] Vilalta R, Vía S, Mira F, Casellas R, Muñoz R, Alonso-Zarate J, Kousaridas A, Dillinger M. Control and management of a connected car using sdn/nfv, fog computing and yang data models. In2018 4th IEEE Conference on Network Softwarization and Workshops (NetSoft) 2018 Jun 25 (pp. 378-383). IEEE.

[3] Muñoz R, Vilalta R, Yoshikane N, Casellas R, Martínez R, Tsuritani T, Morita I. Integration of IoT, transport SDN, and edge/cloud computing for dynamic distribution of IoT analytics and efficient use of network resources. Journal of Lightwave Technology. 2018 Apr 1;36(7):1420-8.

[4] Fallgren M, Dillinger M, Alonso-Zarate J, Boban M, Abbas T, Manolakis K, Mahmoodi T, Svensson T, Laya A, Vilalta R. Fifth-generation technologies for the connected car: Capable systems for vehicle-to-anything communications. IEEE vehicular technology magazine. 2018 Jul 17;13(3):28-38.

[5] Garg S, Kaur K, Kaddoum G, Ahmed SH, Jayakody DN. SDN-based secure and privacy-preserving scheme for vehicular networks: A 5G perspective. IEEE Transactions on Vehicular Technology. 2019 May 20;68(9):8421-34.

[6] Varma IM, Kumar N. A comprehensive survey on SDN and blockchain-based secure vehicular networks. Vehicular Communications. 2023 Aug 22:100663.

[7] Meyer P, Hackel T, Langer F, Stahlbock L, Decker J, Eckhardt SA, Korf F, Schmidt TC, Schüppel F. A security infrastructure for vehicular information using sdn, intrusion detection, and a defense center in the cloud. In2020 IEEE Vehicular Networking Conference (VNC) 2020 Dec 16 (pp. 1-2). IEEE.