CONNECT

CCAM TRUST & RESILIENCE

D2.1

Operational Landscape, Requirements and Reference Architecture – Initial Version

Project number	101069688
Project acronym	CONNECT
Project title	Continuous and Efficient Cooperative Trust Management for Resilient CCAM
Start date of the project	1st September 2022
Duration	36 months
Call	HORIZON-CL5-2021-D6-01-04

Deliverable type	Report
Deliverable reference number	D6-01-04/ D2.1/ V1.1
Work package contributing to the deliverable	WP2
Due date	31st August, 2023 - M12
Actual submission date	07th December 2023

Responsible organisation	UBITECH
Editor	Anna Angelogianni (UBITECH), Thanassis Giannetsos (UBITECH)
Dissemination level	PU-Public
Revision	1.1 (disclaimer updated)

Abstract	Deliverable D2.1 defines the technical requirements of CONNECT, alongside the specification of the conceptual reference architecture, the functional components, and interfaces between them. It also provides an analysis and point of reference for CONNECT in relation to the use cases and reference scenarios including an analysis of the trust models, that need to be considered for capturing the security and privacy requirements of each of the target CCAM services, and the trusted computing anchors to be further investigated and deployed as part of all actors – both at the MEC and Edge. Its purpose is to define the parameters for the rest of the CONNECT project and provide the necessary input for the design and implementation of all security enablers and models towards the dynamic trust assessment of complex CCAM ecosystems.
Keywords	Architecture Specification, Functional Components, Interfaces & APIs, Requirements Analysis, Trust Assessment, Use Cases, User Stories, MVP

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

Editor
Anna Angelogianni, Thanassis Giannetsos (UBITECH)

Contributors (ordered according to beneficiary numbers)
Anna Angelogianni, Dimitris Karras, Thanassis Giannetsos (UBITECH)
Ana Petrovska, Ioannis Krontiris (HUAWEI)
Panagiotis Pantazopoulos, Pavlos Basaras (ICCS)
Nataša Trkulja, Artur Hermann, Frank Kargl (UULM)
Antonio Kung, Guillaume Mockly (TRIALOG)
Francesca Bassi, Ines Ben Jemma (IRTSX)
Chris Newton, Liqun Chen (SURREY)
Alexander Kiening, Anderson Ramon Ferraz de Luce (DENSO)
Matthias Schunter, Dmitrii Kuvaiskii (INTEL)
Konstantinos Latanis, Sotiris Kousouris (SUITE5)
Marco Zanzola (CRF)
Chirag Arora, Adam Henschke (UTWENTE)
Marco Rapelli, Claudio Casetti (POLITO)

Executive Summary

The emergence of Connected, Cooperative, and Automated Mobility (CCAM) systems is envisioned to transform to great extent the future of transportation, providing innovative solutions to challenges and improving commuter experiences. These technologies aim to improve road safety, optimizing traffic control, and reduce transportation times and fuel expenses. However, ensuring the trustworthiness, security, and privacy of these systems becomes paramount, given their integration into various essential aspects of people's daily lives. The task is particularly challenging due to the multitude of devices and communication interfaces that expand the attack surface.

Especially considering the current trend of decomposition of systems across the entire CCAM continuum, including Multi-access Edge Computing (MEC). Despite the benefits it introduces, it further opens up to new vulnerabilities which affect the entire CCAM topology. The dynamic nature of interactions between vehicles, infrastructure, and services, coupled with the time-sensitive nature of decision-making for safety-critical functions like collision avoidance, adds another layer of complexity to this imperative task. Continuous assessment and evaluation of trust levels for both actors/nodes and data are essential enablers of the cooperative (i.e., Day 3) vision of CCAM. Interoperability among diverse entities and cooperative decision-making within the CCAM ecosystem rely on mutual trust in data and communications.

The CONNECT project aims to facilitate the development and further adoption of CCAM ecosystems by being the first to introduce a continuous trust assessment mechanism enabling the operation of the entire CCAM continuum in a zero-trust manner. Its vision is to allow participants to define trust models and assess dynamic trust relationships, establishing trust for cooperative safety-critical decisions. CONNECT will provide secure data exchange between data sources in the CCAM ecosystem, federated trust assessment and also provides the blueprint alongside and reference implementation of a Digital Twin architecture that can allow the secure and authenticated offloading of resource extensive tasks (i.e., trust calculations or application related tasks).

The project extends beyond traditional security requirements that mainly focus on integrity of comprised CCAM actor to also include other trust properties of CCAM ecosystems such as reliability, resilience and robustness, by enabling an overarching trust management. This can ensure that participating nodes can be trusted against generic trust models and the data they exchange is secure and trustworthy, adhering to the strictest security requirements. Furthermore, the broad adoption of autonomous vehicles and connected services is heavily dependent on the trustworthiness of the technology, which directly impacts user acceptance. Trust is crucial for protecting user privacy, establishing strong protection against cyber-attacks, and conforming to the regulatory norms.

The present deliverable, D2.1, sets the foundation for the entirety of the CONNECT project in terms of technical innovations and contributions. Towards this direction the CONNECT architecture is presented, focusing on the modes of operation (i.e., cloud, Multi-access Edge Computing and far-edge Vehicle), their position and interconnection. The architectural framework revolves around two principal pillars: i) the node and data centric *trust assessment framework,* and ii) the novel *Trusted Execution Environment (TEE) extensions* for the runtime verification of trustworthiness evidence. this document provides an initial, yet comprehensive, high-level view of their functionalities and message exchanges. A more detailed analysis regarding the architecture, including thorough descriptions of interfaces, message types, and information exchange, will be furnished in the subsequent deliverables of WP3, WP4, and WP5.

The deliverable highlights both the functional and non-functional requirements integral to the proposed CONNECT framework. Special attention is given to key aspects such as trust assessment, security and operational requirements as well as privacy requirements. The document goes beyond requirements by providing a comprehensive definition of use cases and detailed user stories. This use-case-driven approach, in conjunction with the identified requirements, serves as a foundation for eliciting the timeline for the development of the core CONNECT integrated framework. The analysis of use cases and requirements has resulted in a two-phased approach to evaluation plans, with the first round of experimentation scheduled for M21. This structured approach ensures that the

development and evaluation of the CONNECT framework align with the identified requirements and use cases, contributing to its effectiveness and reliability in real-world applications. The second round of experimentation will take place during M30.

Table of Content

List of Figures

List of Tables

1 Introduction

1.1 Towards Dynamic Trust Assessment in Future CCAM Services

The future of transportation is on the verge of a significant change, led by the emergence of Connected, Cooperative, and Automated Mobility (CCAM) systems. These cutting-edge technologies are ready to introduce a new era, providing innovative solutions to critical challenges and improving the daily experience of commuters with autonomous driving through perception sharing, path planning, real-time local updates, and coordinated driving. The core of this shift is the focus on improving road safety and traffic control, while reducing transportation times, as well as fuel expenses. By adopting CCAM technologies, society enters a future where mobility surpasses beyond merely a means of transportation, becoming a domain of enhanced effectiveness and ease. This future holds the potential to revolutionise passenger's mobility but also their interaction with the environment. More details on the specific applications enabled by the CCAM services can be found in chapter 2 of the present deliverable.

Nevertheless, as CCAM systems are gradually becoming more integrated into many aspects of our everyday lives, it is crucial to prioritise their reliability, security, and privacy. Considering the diverse range of stakeholders that participate in this ecosystem like sensors, vehicles, and others that provide valuable input, such as traffic conditions, it is evident that the attack surface is exceptionally broad and multifaceted. Given the significance of CCAM systems, and their potential influence on passenger safety in the event of a failure, security and trustworthiness are key properties; hence, it is crucial to develop effective protection and trust evaluation methods, for these systems, before their broad usage. The security factor is not solely addressed to a single vehicle itself, but furthermore on the ability of this vehicle to assess the trustworthiness of the incoming information from other vehicles, sensors, the cloud, the Multi-access Edge Computing (MEC), etc.

Towards this direction, the primary objective of the CONNECT project is to facilitate the development and adoption of CCAM ecosystems, by introducing novel security and trustworthiness assessment mechanisms for ensuring advanced protection levels. Through CONNECT, the participating entities of the system will be able to define a trust model and assess dynamically the trust relationships, based on which trust is established for cooperatively executing safety-critical decisions. CONNECT will provide secure exchange of data between data sources in the CCAM ecosystem that lack or have inadequate pre-existing trust relationships, as well as enabling the reliable outsourcing of tasks to the MEC and the cloud.

It is important to highlight that CONNECT extends its scope beyond traditional safety requirements that ensure that the system operates without causing harm to its users or surroundings to prevent accidents. CONNECT further covers the overall security and trustworthiness of CCAM ecosystems by extending its focus to the domain of trust management. CONNECT enhances the overall level of trust in the system, ensuring that participating nodes in this ecosystem (i.e., vehicles, sensors, etc.) can be trusted, and the data they exchange is also secure and trustworthy.

1.2 Scope and Purpose

The present deliverable plays a foundational role within CONNECT, setting forth the essential technical requirements that have shaped the architectural framework of the project. This architectural design doesn't solely prioritise security; it also emphasises safety in the context of CCAM ecosystems. It comprehensively addresses the mobility and connectivity prerequisites, including factors like latency and performance. Beyond functional requirements, CONNECT incorporates non-functional requirements, particularly concerning privacy aspects across the Edge-Cloud CCAM Continuum.

This deliverable dives deep into the reference architecture, functional components, and their vital interconnections, offering an extensive exploration of the CONNECT framework's design. The research carried out in this document not only provides in-depth insights into CONNECT but also highlights its significance in various use cases and reference scenarios.

The focal point examined in this deliverable, that further drove the choice of the use case scenarios is trust, which serves as the crucial element in any CCAM ecosystem. Hence, we examine trust models that play a critical role in capturing the security and privacy needs that are intrinsic to each CCAM service. Moreover, we delve into the world of trusted computing anchors, investigating their function and deployment across various actors within the Multi-access Edge Computing (MEC) and Edge domains.

The insights and characteristics presented in serve as a point of reference for the CONNECT project, leading the way towards the dynamic assessment of trust within complex and ever-evolving CCAM ecosystems. With the fulfilment of this deliverable, CONNECT establishes the ground for the development of a complete framework which incorporates security, privacy, and, most crucially, trustworthiness.

1.3 Relation To Other WPs and Deliverables

The CONNECT framework aims to address the gap of trust management in the CCAM ecosystem to enable a reliable assessment of the involved actors in the service graph chain, regardless of their layer (i.e., local. MEC, cloud), as well as the data exchanged. The present deliverable (D2.1), covers the entire operation of CCAM ecosystems both from a security and a safety standpoint. The definition of the Security, Trust, and Operational Assurance core requirements, per layer, along with the documentation of the CONNECT reference architecture and the further decomposition of the basic system components and the intercommunication among them, is crucial for the establishment of such a framework.

WP2, further provides the definition of the use cases and reference scenarios that correspond to CCAM functions, along with technical requirements for CONNECT and the respective Key Performance Indicators (KPIs). This serves as the core information to extract more detailed Security, Trust and Operational Assurance and translate them into concrete architectural specifications in order to meet the desired capabilities of the CONNECT framework.

Figure 1 depicts the direct and indirect relationship of the D2.1 to the other Work Packages (WPs). The definition of the system-wide reference architecture is cornerstone, in order to drive the technical work of WP3-WP6 and steer the validation processes in WP6.

The outcome of D2.1 is intended to support the definition of later activities in the project. In relation with the rest of the WPs of the project, D2.1 serves as a point of reference for the technical developments of the project as it offers a set of direction to each WP. More specifically, D2.1 though the definition of the use cases, paves the way to the clear definition of the trust properties that will be employed for the evaluation of the Required Trust Level (RTL) and the Actual Trust Level (ATL) per case for the trust assessment. Similarly, D2.1 provides a conceptual architectural definition of the Trusted Execution Environment (TEE) Guard which will be used to execute the security critical functions in a trusted environment (WP4). It is worth noting that the CONNECT framework is agnostic to the specific Root of Trust (RoT) employed and opts to evaluate the enablers without the constraints of this environment. WP5 inherits the defined data models regarding the data sharing tasks and the calculations that are executed at the MEC-level. Last but not least WP6 aims to the validation of the CONNECT framework in the context of the D2.1 defined pilots i.e., *“Intersection Movement Assistance”, “Vulnerable Road-User (VRU) Protection,* the *Cooperative Adaptive Cruise Control (C-ACC)”,* and *“Slow Traffic Movement Detection”.*

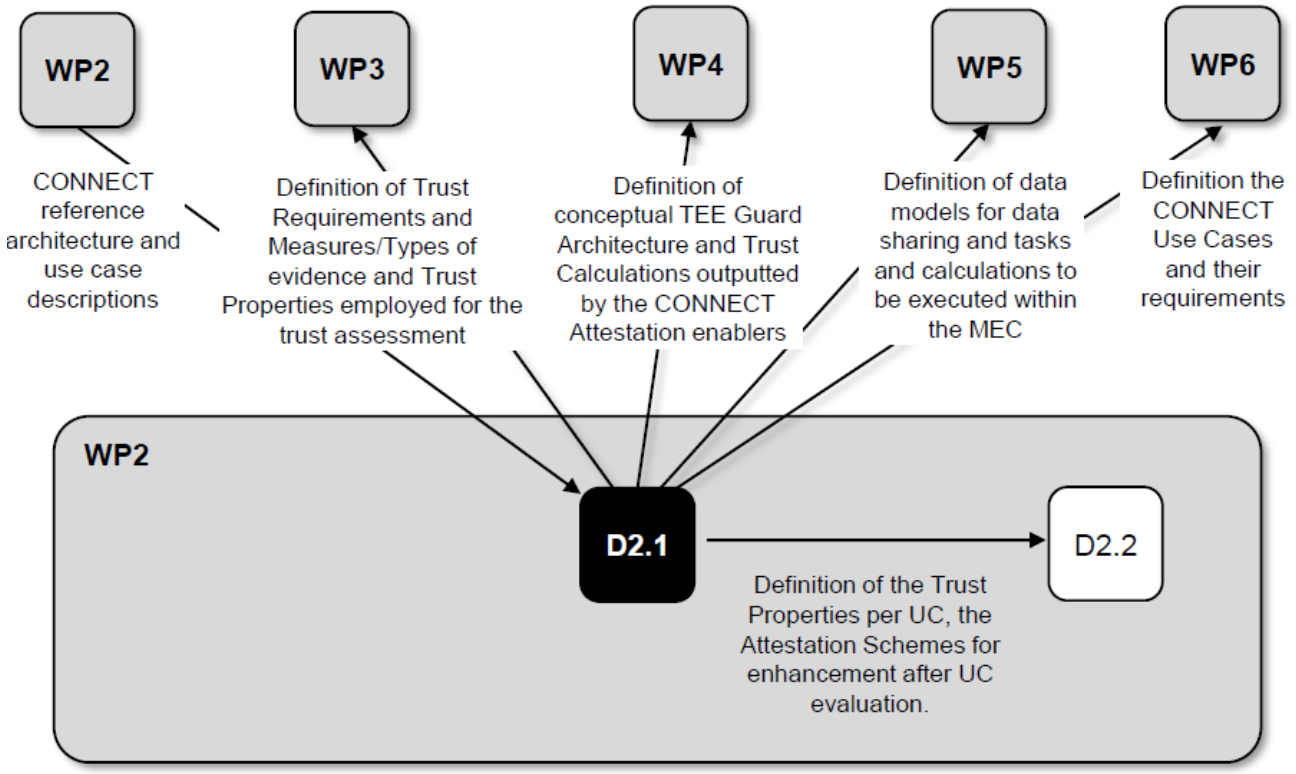

Figure 1 - WPs Relation

1.4 Deliverable Structure

Chapter 2 begins by outlining the vision and background of the CONNECT project, offering insights into the current state of the art within the CCAM landscape. Additionally, it emphasizes the encountered challenges and sets the ground for introducing the novelties presented by the CONNECT framework. **Chapter 3** places emphasis on the domain of Trust to define the chosen approach and considerations as it pertains to the trust assessment framework. **Chapter 4** explores the CCAM environment, offering in-depth analysis of the stakeholders involved and the types of exchanged messages and information shared between these complex ecosystems. **Chapter 5** outlines the methodology employed for eliciting functional requirements, shedding light on the process that ensures the framework's effectiveness. **Chapter 6** provides a comprehensive overview of the functional components and architecture of CONNECT, including a detailed examination of the project's architectural design. **Chapter 7** provides an in-depth analysis of the use cases, specifically examining the realistic scenarios in which CONNECT's capabilities are applied. **Chapter 8** provides a comprehensive and detailed description of the functional and non-functional requirements of CONNECT. It offers a thorough overview of the project's technological prerequisites. In **Chapter 9**, the mapping of the use case requirements is defined, providing a practical viewpoint on how CONNECT effectively caters to the distinct requirements of each use case scenario. Lastly **chapter 10** draws the conclusions.

2 CONNECT Vision and Background

2.1 Getting CCAM On the Road: Future Automation Needs to be Built on Zero Trust

Connected vehicles, as part of the emerging Cooperative Intelligent Transportation Systems (C-ITS) are positioned to transform the future of mobility. This change is enabled by the vehicle's communication with other entities formulating the Vehicle-to-everything (V2X) landscape. V2X communication systems are expected to greatly improve road safety and traffic control efficiency. Today, the introduction of C-ITS has already started. Vehicles and infrastructure transmit information regarding their status as well as information describing unexpected events. Exchanging this type of information allows realising warning applications (Day 1 applications [1]), such as collision avoidance warning. The communication services adopted for the generation and management of this information are based on the exchange of Cooperative Awareness Messages (CAMs), and Decentralised Environmental Notification Messages (DENMs).

Public Key Infrastructure (PKI) serves as the foundation of security, privacy and trust in C-ITS by providing a secure framework for authentication, confidentiality, and integrity of the communications between different entities within the system. Through the use of digital certificates, the system makes sure that only authorised and legitimate participants can interact with each other. The Certificate CAs provide the digital certificates that establish a connection between public keys and identities, enabling the process of authentication. A vehicle is also able to trust the authenticity of the message by verifying the signature of the sender vehicle against a trusted certificate chain. The foundation of trust in PKI lies in the concept of trusted root CAs, which are pre-installed and are inherently trusted, acting therefore as trust anchors in the C-ITS ecosystem. It shall be noted that past PKI implementations were subjected to identity and location leakage-related attacks. Novel implementations strengthen the privacy guarantees by employing short term pseudonyms, that conceal the actual identity of the vehicle.

The Car2Car consortium has developed a comprehensive deployment roadmap consisting of three distinct phases: **Day 1, Day 2, and Day 3**, as depicted in Figure 2. The development of this roadmap is based on empirical observations of automation trends in the commercial vehicle industry and the growing prevalence of V2X-equipped systems. The evolution of V2X technology can be observed, starting with its initial function of exchanging status information, and gradually progressing towards the development of cooperative automated driving capabilities. The aforementioned vision not only presents a trajectory for the seamless development of intelligent mobility, but also emphasises the pivotal significance of V2X technologies in shaping the safety and interaction of vehicles.

Several of the Day2 and beyond applications are (mainly) intended for automated vehicles [1]. These connected and automated vehicles will benefit from increased connectivity with other vehicles, the infrastructure and other road users. This heightened level of connectivity allows them to exchange planned trajectories/routes and coordinate manoeuvres with other traffic participants as well as the infrastructure. Such information sharing paves the way for the implementation of cooperative automated driving scenarios where automated vehicles can collaborate implicitly or explicitly to execute manoeuvres while avoiding conflicts and ensuring overall safety. In that way, C-ITS (mostly referring to connected and cooperative vehicles) converges into CCAM (Connected Cooperative and Automated Mobility).

The 5G Automotive Association also provides a comprehensive definition of communication protocols and virtualization technologies that can be used for deploying a service and safety-critical functionalities. Their primary focus is directed toward the utilization of 3GPP standards, emphasizing V2X communication, encompassing both direct C-V2X and the 5G set of protocols (5GAA) [2]. In addition to use cases that depend only on short-range communication, it also specifies use cases that are only possible with long-range communication. A lot of emphasis is placed on sensor sharing use cases with different variations (e.g., data collection and sharing for HD maps, data sharing of dynamic objects, non-analysed sensor signal sharing). Sensor sharing is the cornerstone building

block required to enable Advanced Driver Assistance Systems (ADAS) ranging from Level 2 (AD L2+) to Level 3 (AD L3), as well as connected ADAS assistance.

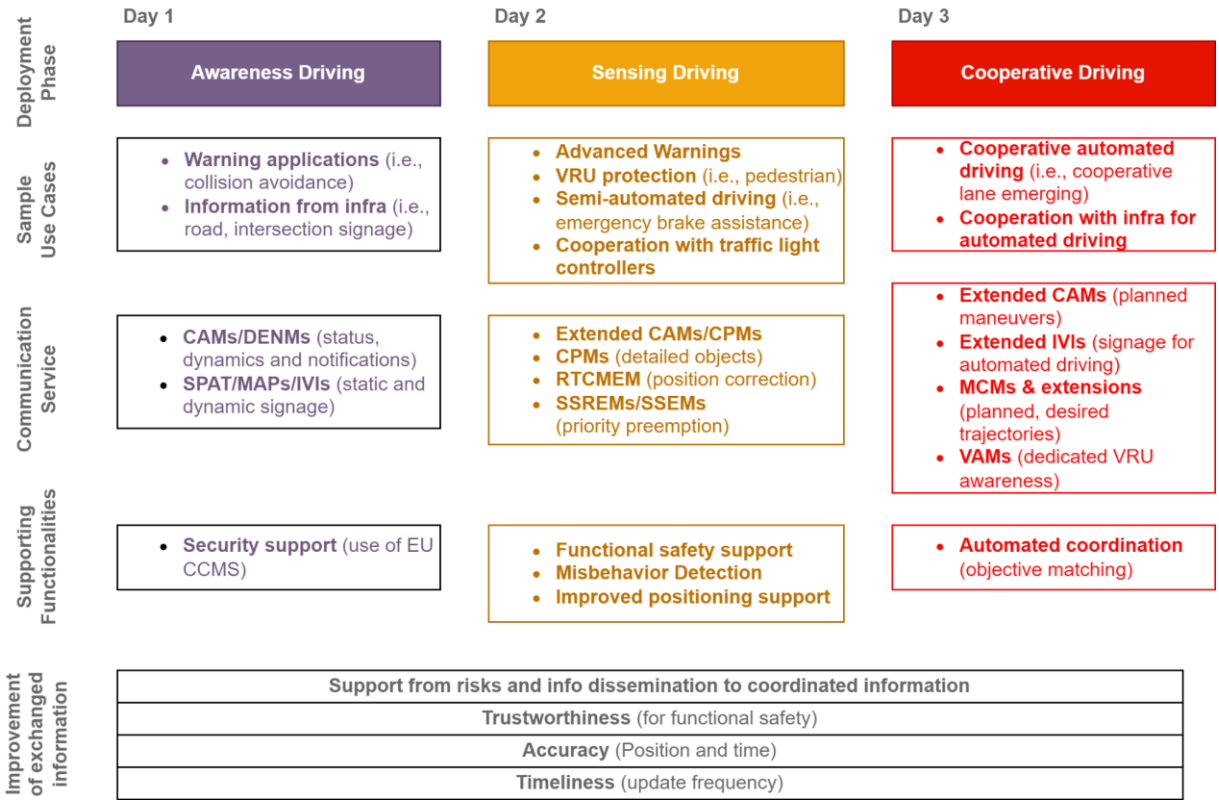

Figure 2 - Roadmap according to C2C consortium

However, the shift towards higher levels of automation poses a significant challenge - *the need for external data to facilitate partially automated or fully automated driving functions*. In this context, the integrity and trustworthiness of external data sources, such as external sensor information, maps, and positioning data, becomes paramount. If the integrity of this data is compromised or not provided with the expected quality, the building blocks of the automated operation functions will use incorrect data to control the vehicle. There is a broad set of security attacks that have consequences on the trustworthiness of the data and data sources. The dependability and resilience of CCAM systems can be seriously affected by these attacks at run-time. Furthermore, there are many sources and reasons that can negatively impact dependability and safety that are not related to security. Mechanical defects, failure of ECUs or decreased accuracy of sensors are just some examples of events from this category.

The need to solve this problem becomes increasingly pressing as we move towards the Day2 and Day3+ phase. As we advance, entities increasingly depend on this information to make safety-critical decisions. Consequently, for all forthcoming use cases in the realm of C-ITS and CCAM to effectively utilize external information, it becomes imperative **to explicitly define and quantify the trustworthiness of exchanged data**, which is used as evidence. The integrity of any evidence, particularly when it is used in safety-critical decision-making, should be trustworthy; hence verifiable.

Even though the security through integrity of V2X communication, as part of C-ITS systems is somehow established (regardless of the challenges pertaining to the scalability of existing PKI solutions), the problem of assessing trust on the exchanged information in such a highly dynamic, distributed, and ubiquitous environment, remains open. That is because we lack tools to reason about trust relationships between data sources that were previously unknown to each other. In CCAM emerging scenarios, it might be the case that the sources of evidence offered by others are untrusted, or the evidence is indirect and obtained through a referral chain.

Therefore, we need tools to measure and manage levels of trust under uncertainty, based on incomplete and/or subjective information provided by potentially untrustworthy sources. Furthermore, these tools should accommodate dynamically changing trust relationships due to the high level of mobility exhibited during the operational time of the systems at run-time. This is one of the core challenges to be resolved for unlocking the full potential of CCAM ecosystems. **The core disruptive technology created by CONNECT, Dynamic Trust Assessment, achieves exactly that: probabilistic assessment of data and data source trustworthiness.**

The issue of trust in CCAM extends beyond the realm of data and data sources. ETSI introduced the term Multi-access Edge Computing (MEC) [3] with the goal to bring processing power near the vehicle to meet ultra-low-latency requirements, and to reduce network traffic towards a centralised datacentre. This has two important advantages: First, with the help of MEC, massive computation and storage tasks need not be handled in the vehicle with its limited power and resources. Instead, these functionalities can be offloaded to the MEC, which can handle it in a more cost-effective way in real-time. Second, MEC can act as a coordinating anchor for various V2X services and enable access in critical safety and real-time processing of sensor signals from various vehicles and RSUs, as described above. Ertico's investigation (see Figure 3) into the progression of CCAM connectivity from 3G to 5G, along with the integration of MEC specifically for the automotive industry, has demonstrated its significance in facilitating the proximity of vital services to vehicles, thereby enhancing system performance and accuracy.

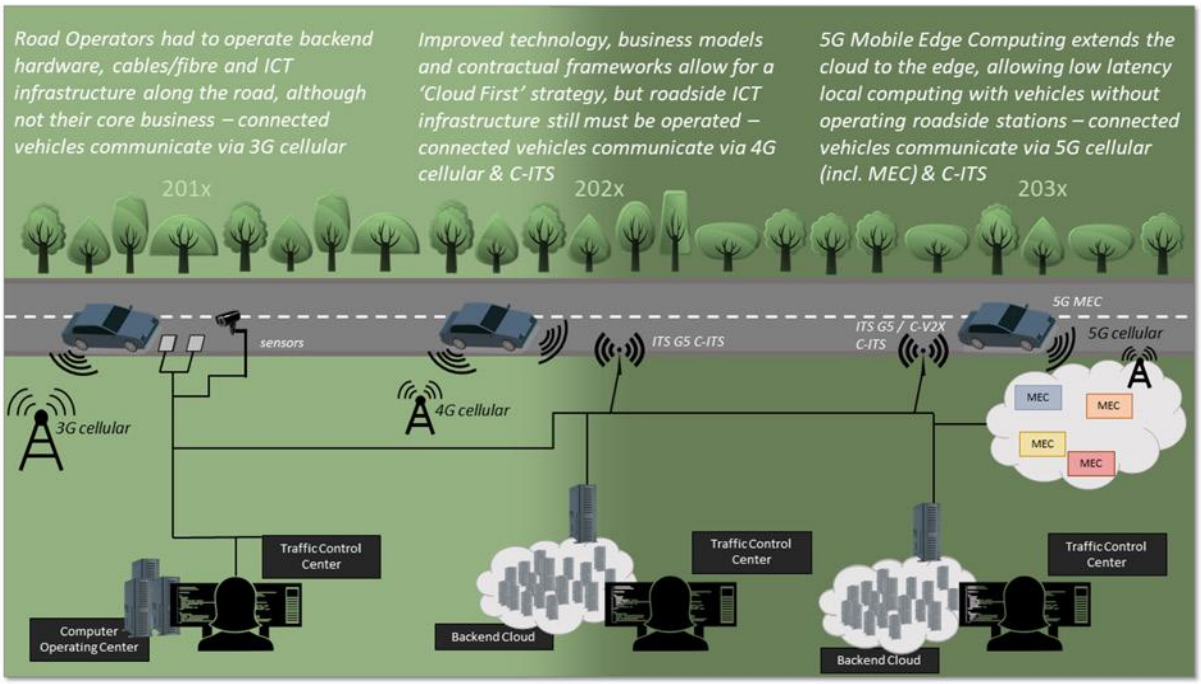

Figure 3 - CCAM Connectivity evolution (based on Ertico)

However, it is essential to acknowledge that such Edge Computing environments possess inherent characteristics of a **complex and highly heterogeneous ecosystem** due to the involvement of multiple vendors, suppliers, Original Equipment Manufacturers (OEMs) and stakeholders. Additionally, in the context of distributed systems, it is not feasible to presume the presence of a

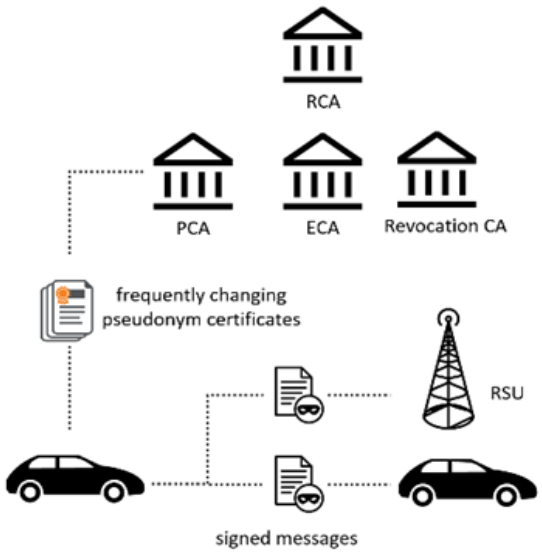

Figure 4 - A V2X security solution based on PKI

central entity responsible for implementing universal security measures (and updates) across the entirety of the system. Hence, **it becomes apparent that in such highly complex environments, trust levels vary**. Towards this end, the existence of various MEC hosts, that correspond to different trust domains, and may require seamless information exchange, necessitates the implementation of mechanisms for evaluating the level of trust for each party involved. This evaluation should take into account the dynamic nature of the environment along with its heterogeneity, particularly in relation to activities involving lifecycle management (i.e., secure enrolment or deployment). Therefore, **to prevent the formation of deceptive or impractical trust assumptions, it would be advantageous to implement a system founded on the concept of zero trust**, which entails a mindset of "*never trust, always verify*". This approach mandates that every entity involved, whether physical or virtual, must provide substantiating evidence to establish its trustworthiness, irrespective of its location within the system.

CONNECT adopts the "*never trust, always verify*" approach of zero trust architecture, since it implements security measures at all network and infrastructure levels, regardless of the user or resource's location. It treats any user, device, or application attempting to access resources as untrusted and trust is continuously evaluated based on evidence.

2.2 SOTA Analysis

This section delves into the various facets of CCAM technology, offering valuable perspectives on the most recent advancements and groundbreaking ideas that are influencing the trajectory of transportation in the coming years. The topics encompassed in this discussion include Public Key Infrastructures (PKIs) and their extensions, cryptographic security, the identification and detection of misbehaviour, the assessment of trust, the concept of edge computing, the utilisation of digital twins, the establishment of a Zero Trust environment, and the implementation of confidential and trusted computing mechanisms. The aforementioned concepts will have a crucial impact on facilitating advancements in day 2 and day 3+ applications, while also providing significant levels of security and operational safety.

2.2.1 Public key Infrastructures (PKIs) and Beyond

The correctness and reliability of V2X messages is a fundamental step and a key enabler for the envisioned C-ITS applications, as it ensures a certain level of security and trust through authenticity among the actors. The use of digital certificates was suggested very early on in order to authenticate messages in vehicular communications and prevent an outsider attacker from injecting false messages. However, since V2X applications rely on broadcasting continuous and detailed location information, as for example, through CAM, it raises the need to protect privacy as well.

Addressing this challenge, current approaches are based on PKI-based solutions with privacy-friendly authentication services through the use of short-term pseudonyms (for a survey, see [4]). The common denominator in such architectures is the existence of trusted (centralised) infrastructure entities (as demonstrated in Figure 4) for the support of services such as authenticated vehicle registration, pseudonym provision, revocation, etc. Hence, the location privacy is protected by requiring that each vehicle uses multiple pseudonyms that are frequently updated.

Prominent solutions include the Security Credential Management System (SCMS) [5], which is a product of vehicle OEM consortia and the US Department of Transport (USDOT), the Cooperative-ITS Certificate Management System (CCMS) developed by the European Committee for Standardization (CEN) and European Telecommunications Standards Institute (ETSI), with support from the European Commission [6] and the Chinese C-SCMS developed by CCSA [7]. 5GAA has evaluated the SCMS and the CCMS system designs, and it has concluded that they can be improved to take advantage of cellular connectivity. The effort to identify potential design simplifications to increase efficiency and harmonise technologies across regions has resulted in an updated system design for large-scale deployment and cross-regional interoperability called “Efficient Security Provisioning System” (ESPS) [8].

Broadly speaking, in all of the above systems, Privacy and Cybersecurity features have been realised by-design by defining the Certificate and Security Policy based on PKI management and pseudonymizing of the signatures messages.

The question then becomes*, how the vehicles are provided with the set of pseudonyms*. In the PKI approach, a set of certification authorities (CAs) provide credentials to the vehicles. In the general case, there is a set of different authorities with distinct roles:

- ✓ Root Certificate Authority (RCA): This entity is the trust anchor of the PKI that is responsible for issuing certificates to sub-CAs. The certificate of the RCA is signed by itself.
- ✓ Enrolment Certification Authority (ECA): This entity is responsible for registering vehicles and issuing long-term certificates. Entities with enrolment certificates can then apply to other CAs, like for example to the pseudonym CA for issuing pseudonym certificates.
- ✓ Pseudonym Certification Authority (PCA): This entity is responsible for issuing certificates that do not contain any identifying information.
- ✓ Certificate Revocation CA: responsible for issuing certificate revocation list for all kinds of certificates.

Private key material associated with pseudonym credentials should be stored securely inside vehicle OBUs and should not be extracted or transferred outside the vehicle. For this reason, the integration of hardware security modules (HSM) or tamper-proof devices (TPD) in OBUs have been proposed for secure key storage and management [9].

At the same time, there is an increasing effort by researchers to explore decentralised solutions, capable of shifting trust from the back-end infrastructure to the edge (i.e., vehicles), in order to reduce the vector of entities for which we want to make sound statements in terms of their configuration, security settings and trustworthiness. For example, more recent work started investigating the use of trusted computing to transfer the root of trust inside the vehicle and also leverage the power of anonymous credentials through the use of advanced cryptographic primitives such as Direct Anonymous Attestation (DAA) [10]. This offers the following advantages to us:

- ✓ Create an immutable root of trust in the vehicle. A trusted computing represents a commonly trusted and immutable initial processing step, to which trust in other processing steps can be bootstrapped.
- ✓ Policies to access the secure memory, which can be used to store certificates that are then used to attest to security of system components. This “secure memory” is immutable, except via the TC component itself.

An in-vehicle trusted computing component and the use of DAA enables vehicles to create an **unlimited amount of trusted pseudonym certificates** in the vehicle itself. Most notably, one of the biggest advantages of applying the DAA protocol is the redundancy (and removal) of a number of authorities, such as the Pseudonym Certification Authority; vehicles can now create their own pseudonyms, and DAA signatures are used to self-certify each such credential. This allows vehicles to have better control over their privacy, since no trusted third-party is involved in the pseudonym

creation phase. The combined use of trusted computing and crypto such as DAA can also provide the based for empowering the vehicles to control what information they disclose and the level of privacy.

Moreover, DAA proves to be an efficient solution for **pseudonym revocation** in large-scale C-ITS systems, as it does not necessitate the use of Certificate Revocation Lists (CRLs), avoiding delays associated with computational and communication overhead. Typically, CRLs include a list of revoked certificates, published by the Certification Authority, but in scenarios with numerous certificates per vehicle due to pseudonym usage for privacy, pseudonym resolution is further required, and this may add to the complexity of the revocation task.

Whitefield et al. [11] first applied this DAA-based solution for generating and managing pseudonym certificates in vehicle scenarios. Larsen et al. further enhanced DAA for V2X in terms of security, privacy, scalability, and revocation capabilities [12]. More recently, Angelogianni et al. [13] implemented the DAA-enabled solution and put forth the first complete analysis between DAA-enabled and PKI-enabled security configurations for V2X, in order to showcase how the former overcomes a series of shortcomings of the conventional centralised solutions in terms of scalability and computational footprint. It is worth noting that while PKI-based solutions provide robust support for security and privacy, trust had not been the primary focus until now. The study's results don't explicitly favour one solution over the other but instead underscore the need for a holistic strategy that combines elements from both approaches to better address the evolving requirements of future Cooperative Intelligent Transport Systems (C-ITS).
Towards this direction, CONNECT will provide support to the traditional PKI schemes, as standardised by ETSI, but will further investigate the DAA scheme for the authentication of the trustworthiness claims.

2.2.2 Trust Assessment in CCAM

This chapter does not aim to present a comprehensive overview of the State of the Art (SOTA) in the field of CCAM Trust Assessment, as such an analysis has already been covered in D3.1. Instead, our focus here is to provide a deeper understanding of the necessity for dynamic trust assessment. We specifically address situations where the intricacy arises from two key factors: i) the sheer volume of incoming data and ii) the potential conflicts and inconsistencies within this data. We delve into the unique requirements and challenges presented by these complex CCAM systems, shedding light on the critical role of dynamic trust assessment in managing the flow of data and ensuring the reliability and safety of these advanced ecosystems.

Despite the robust cryptographic measures in place, exemplified by the use of PKI as elaborated in the previous section, the need to incorporate trust assessment mechanisms persist within the complex and heterogeneous ecosystem that encompasses an extensive array of devices and data sources. This inclusivity extends to sources that inherently lack trustworthiness, including vehicles with sensors transmitting CAM and CPM messages. In this landscape, these trust assessment mechanisms assume a critical role in not only establishing and quantifying the extent of trust existing between entities, based on both their own opinions as well as the opinion of others without a central authority, but further enabling the provision of trust evaluations even for devices that lack inherent trustworthiness. Achieving this requires a multifaceted approach, leveraging a diverse set of properties and criteria to form trust relationships primarily based on integrity-related attributes. As the ecosystem thrives on interconnectedness, these mechanisms pave the way for a secure and resilient environment where interactions among diverse entities can transpire with confidence.

Understanding, classifying, measuring, and assessing trust have been fundamental research challenges in trust management in the last twenty years. Trust assessment has a focal point in sociology, technology, and computing, including e-commerce, access control, and security risk analysis. For example, Kurdi et al. [14] proposed a trust assessment approach for a cloud ecosystem where multiple Cloud Service Providers (CSPs) cooperate. Here, a CSP forms a subjective trust opinion on other CSPs based on Service Level Agreements and the reputation it has with these CSPs. Based on this, a CSP can select which other CSP to cooperate with. Furthermore, Garlichs et al. [15] describe an approach for trust assessment in the context of vehicle platooning. Here, a host vehicle determines a trust opinion to its predecessor in the platoon by comparing the actual

position of the predecessor (known due to the on-board sensors) to the contents of the received V2X messages over an extended period. Based on the determined trust opinion, the safety distance is adjusted. Lastly, Cheng et al. [16] use epistemic logic to quantify the trustworthiness of agents in multi-agent systems. The quantification is done by an observer monitoring the behaviour of the agents, based on which the trust opinion of the agents is adjusted. In all of these works, trust is assessed between two peers who communicate directly with each other. In other words, *these works assess the level of trust in the scope of a specific relationship between (two) actors*. However, this does not suffice in the context of CCAM, where trust needs to be assessed on complex trust networks comprising multiple such trust relationships with a high degree of dependencies. This is also depicted in the Figure 9, which highlights different data sharing profiles (i.e., trust relationships) encountered in CCAM ecosystems.

Namely, different emerging scenarios and applications from different domains, including CCAM systems, introduce novel challenges for the research field of trust assessment. These emerging systems and the newly introduced use cases for trust assessment mainly originate from the recent technological and engineering trends for increased system autonomy as envisioned in the Day 3 applications in the Car2Car consortium (see Sections 2.1 and 7.1). The current state of the art focuses on highly dynamic, distributed, and ubiquitous systems that operate with a high degree of autonomy in unpredictable, uncertain and continuously changing environments, which inherently introduces a high degree of uncertainty in data. For example, a vehicle in an Intersection Movement Assistance (IMA) use case can be malicious, or a vehicle's sensor could be faulty, reporting wrong information to other vehicles it communicates with. Also, these modern systems become more connected and collaborative, forming more complex systems-of-systems that enable them to attain goals that one system in isolation cannot (e.g., multiple cars communicate and collaborate to autonomously traverse an intersection. In the vision of Day 2 Car2Car consortium, intent sharing is as additional information that can be exchanged as part of the communication and the collaboration between different vehicles). Furthermore, the systems themselves consist of a growing number of components or subsystems (e.g., it is estimated that modern cars contain more than 100 ECUs).

In response, there are some new requirements for trust assessment in the CCAM domain that need to be considered if we are to be able to capture all the aforementioned requirements. Additionally, we exemplify each of the requirements using the IMA use case (cf. Section 7.2). Please note that some of the requirements have already been identified in existing works; however, many of the requirements are unique for CCAM domains and the CONNECT use cases, and to the best of our knowledge have not been considered before in the literature.

First, due to the interconnected, collaborative, and complex nature of CCAM systems, instead of trust assessment of a single trust relationship (i.e., assessing trust among two entities, a trustor and a trustee), there is a need to assess trust on complex networks consisting of multiple trust relationships. For example, as part of the IMA use case, we do not want to assess trust only between two vehicles (e.g., vehicle A and B, forming one trust relationship) or between a vehicle and a MEC (forming another trust relationship), but we want to assess trust on the trust network consisting of different trust relationships of this type that might have dependencies to each other. Please note that when we talk about trust relationships, we consider the construction of atomic propositions related to the fulfilment of those trust properties of interest (e.g., integrity, safety, resilience, etc.) for these specific types of trust relationships (cf. Section 4.1 in D3.1). Additionally, assessing trust on trust networks allows establishing trust in an entity based on input from multiple cooperative entities. Namely, there should be higher confidence in the trustworthiness of an entity (e.g., vehicle E in Figure 5), when the trust is built based on opinions or inputs from multiple entities (e.g., vehicles A, B, C and D in Figure 5), versus only one entity (e.g., vehicle A in Figure 5).

Second, trust levels need to be measured and assessed based on incomplete and/or subjective information provided by potentially untrustworthy entities, and as a result the evidence they provide cannot be trusted. Additionally, there might be cases where evidence is not available in a direct trust relationship, or in general, the direct trust relationship between two entities (e.g., between two vehicles or between vehicle and MEC) might not exist. As a result, the trust assessment shall allow transitive trust, obtained through a single or a chain of referral trust relationships as part of a trust

network. The authors in [17][18] also differentiate between direct and indirect trust, as well as transitive trust.

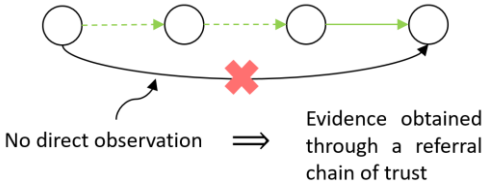

In transitive trust relationships, entity A does not have a (direct) relationship with entity X, for which it needs to build trust. Instead, entity A has a relationship with entity B, and entity B has a relationship with entity X, as shown in the figure.

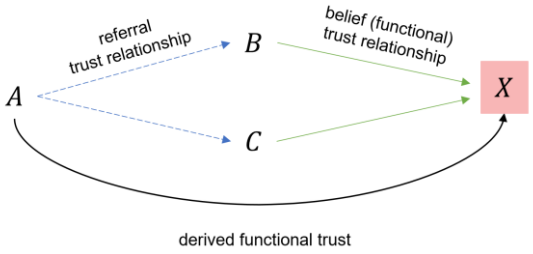

Therefore, entity B can give a recommendation to entity A about entity X, which enables entity A to build an indirect trust relationship (also called a referral trust relationship) with entity X, through B's recommendation about entity X, and A's trust on B. To summarise, as part of a trust network, the problem of lack of direct trust between trust objects can be easily tackled by utilising the notion of transitive trust relationships. In the IMA use case, the entities A, B, C can be different vehicles or MEC, and X could be data that is created, replayed, or processed by these entities. For example, X can be data generated by vehicle B. Although vehicle A does not have a direct observation on the data generated by vehicle X, it can still build trust on X through trust discount over vehicle B.

Third, the systems in the CCAM domain are highly dynamic and new trust relationships are formed ad-hoc at run-time, which is directly related to the mobility feature in the CCAM. In other words, as part of a trust network, there is an emerging requirement for trust relationships between trust objects to be formed and decomposed continuously during run-time. For example, the graphs below, show how the trust network can change in a single time step (from $t = x$ to $t = x + 1$). Concretely, at time $t = x + 1$, the trust object E has been removed/has left the trust network, which has triggered the trust relationships A→E, B→E, D→E and E→C to be broken out. However, trust objects M and N have become part of the trust network, forming new relationships A→M, B→M, D→N and B→N. As part of the IMA use case, this translates to vehicles joining and leaving the intersection. In order for this dynamicity to be reflected in the trust model, we allow the trust model to make new and break existing trust relationships during run-time. There are a couple of other works and white papers in the literature that have identified the need for dynamic trust and have incorporated the temporal aspects [17] [65] [19].

Closely related to the previous requirement, due to the dynamic aspects of the systems, the trust assessment shall be done under strict time requirements. The entire trust assessment process should be fast and robust to be used at run-time for real-time and safety-critical applications. Concretely, in the IMA use case, the trust assessment shall be done in real-time, in order to be seamlessly integrated and not to affect the normal operation of the vehicles.

As a fifth requirement, the trust assessment shall allow assessing trust not only on entities, i.e., nodes, for example, vehicles, MEC, etc., but also shall enable assessing the data produced by these nodes. In other words, the trust assessment should support evaluating trust based on node-centric and data-centric trust relationships. As mentioned above, in the IMA use case, we do not only assess trust on vehicles and the MEC, but also, we assess trust on data created, replayed, or processed by those entities.

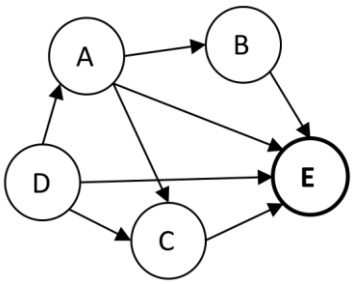

$t = x$

Figure 5 - Trust Relationship Example in time t=x

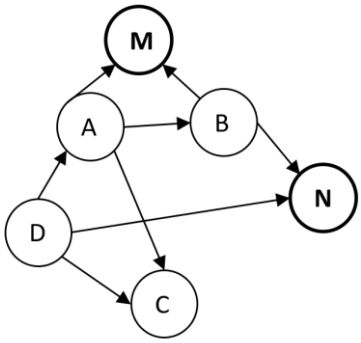

$t = x + 1$

Figure 6 - Trust Relationship Example in time t=x+1

Complex trust networks contain different heterogeneous nodes. As a result, the trustworthiness and trust are assessed based on multiple, varying sources of trust for each trust object in the trust network. Hence, as a sixth requirement, the trust assessment shall build trust based on evidence by many, different trust sources that could also potentially change over time. For example, in the IMA use case we use the misbehaviour detection as a trust source (explained further in Section 2.2.3). In this use case, to determine whether a node is misbehaving, various detectors can be used to analyse the behaviour of the node or the data it sends. Misbehaviour in this context refers to a node sending incorrect data, such as incorrect position data, so we focus on the veracity of the data. Based on the results of these detectors, the trustworthiness of the node can be increased or decreased. As part of Section 6.3 in Deliverable D3.1, we explain different trust sources that we use for other use cases in CONNECT, related to: i) communication, ii) system integrity, iii) applications and iv) entity behaviour.

Finally, due to the distributed and ubiquitous nature of these systems, in combination with the Zero Trust concept that we further explain in Section 2.2.6, centralised solutions for trust assessment in CCAM are insufficient, and there is a need for decentralised and distributed solutions. This inherently requires the notion of subjectivity when it comes to assessing trust. The subjectivity in trust has also been previously considered in other existing works [20][18]. Therefore, the trust assessment shall include belief ownership in order to reflect the subjective nature of beliefs that, once merged, reflect the objective world more precisely than all the opinions in independence. To some degree, this is in relation to the first requirement, where we explain that multiple entities can build trust cooperatively. In the IMA use case, trust assessment will be done in a decentralised manner in each of the vehicles and the MEC. In other words, the Trust Assessment Framework will be instantiated for each vehicle and MEC in our overall architecture, and this will be the focal point of research in the upcoming WP3 deliverables that will focus on TAF federation.

There are various mathematical theories that have been proposed in the literature to ascertain information in uncertain and unpredictable conditions that could be potentially used for trust assessment: Probabilistic Logic, Fuzzy Logic, Bayesian Probability, Dempster-Shafer Theory and Subjective Logic. As part of Deliverable D3.1, we provide an in-depth analysis of all of these mathematical theories, and we explain the decision to choose **Subjective Logic** for our Trust Assessment Framework (TAF) as part of CONNECT. In short, Subjective Logic is the only theoretical framework that enables i) **considering subjective beliefs from multiple agents on**

the same proposition, ii) **merging conflicting sources of evidence** and iii) **puts trust transitivity at the forefront**, i.e., enables trust assessment on complex trust networks instead of trust relationships. For more information refer to Chapter 5 in D3.1

2.2.3 Misbehaviour Detection

As presented in Section 2.1, CCAM services are based on collaborative data sharing between communicating entities and as highlighted in the previous section, the **assessment of the trustworthiness of these shared data and their generators is a fundamental aspect that needs to be ensured since the design phase of the CCAM service**. **Misbehaviour detection** acts at an intermediate level between V2X shared data and CCAM services. It is a set of intermediate functionalities that are activated upon the reception of the V2X shared data that will be considered for the CCAM operations. By continuously verifying the semantics of the received V2X shared data, misbehaviour detection is an additional trust source that contributes among other sources to the establishment of a dynamic and flexible node and data trustworthiness evaluation.

Misbehaviour detection has captured a lot of attention in the past few years in cooperative vehicular systems. **It focuses on developing techniques to detect anomalous behaviours revealed from analysing the semantics of the transmitted V2X messages.** These messages contain specifically kinematic data characterising either the message transmitters or some other elements of the environment, such as the perceived objects. Several techniques are proposed in the literature.

Van der Heijden et al. [21] provided a taxonomy of **data-centric detection techniques and node-centric detection techniques**. Data-centric mechanisms rely on the contents of the message to verify its plausibility and consistency. Node-centric mechanisms rely on the behaviour of a node and generally assign a trust value to every message transmitter. Both categories perform several verifications in order to test if the received data meets some predefined rules or models. If the data does not respect these rules, the sender is potentially misbehaving, and a misbehaviour report is sent to the backend security system for more extensive analysis and verifications [22]. The process of misbehaviour detection at the backend security system is called global misbehaviour detection. Even if there are some works which address the global misbehaviour detection, most of the existing efforts concentrate on the local misbehaviour level.

Generally, the detection of potential misbehaviours relies on the combination of several verifications of the plausibility and the consistency of the kinematic attributes. Using predefined kinematic and communication thresholds [23] or some signal propagation characteristics [24] were the first introduced techniques that enable the detection of non-semantically correct V2X data. In addition to this, some approaches use the cooperative aspects of the vehicular networks to consolidate misbehaviour detection among communicating entities, either by exchanging common kinematic information about neighbours or by relying on entities such as RSUs. Recent approaches rely mainly on non-cooperative detection to avoid the honest majority problems [25].

A second category of work considers the use of Bayesian estimation. [26] was one of the first works that uses particle filters to detect sybil attacks through CAM message transmission and on-board sensor data. [27] reuses Bayesian filters to estimate the trust of the received collective perception data in the context of the manipulation of the perceived object attacks.

Another category of promising approaches uses **fusion techniques to assess the trust in the received data**. [28] is one of the first works that propose to attribute trust to data rather than agents in the context of ephemeral networks such as the vehicular network. They propose a generic framework template where they attribute trust for each individual piece of information called evidence reporting a specific event. The ultimate trust establishment phase is based on different data fusion techniques. The authors show that using the Dempster Shafer theory in the fusion process works well when there is a high uncertainty about the event while the Bayesian inference performs better when the a priori knowledge is provided. Recently, [29] uses the subjective logic concept to detect misbehaving nodes in intersection scenarios. The same concept is used by authors of [30] to assess communicating entities' trustworthiness and thus detect misbehaving ones when cooperative perception data are exchanged. They target especially conflictual situations at the intersections when contradictory perception data are received. They show that, depending on some parameters such

as CPM frequency and on-board sensor configurations, their approach allows them to detect misbehaving entities at acceptable times before entering critical zones in the intersection.

Machine learning techniques are another category of used misbehaviour detection mechanisms. Lately, several works have experimented with machine learning [31][32] to detect kinematic data falsification attacks and show performant results. However, computation cost and time are still major constraints for on-board equipment.

Many challenges are still open in misbehaviour detection for V2X communications. One important challenge is the ability to provide reliable results in the context of vehicular networks being ephemeral networks with high dynamics and changing topologies. Short communication lifetimes, and sometimes unreliable communication channels, do not allow to share a considerable amount of V2X data and thus to build strong proofs of misbehaviour. In addition to this, the sensor perception of the environment is a very challenging task often characterised by imperfection [33] (such as uncertainty and impreciseness). This leads to the assumption that V2X data generators may be unreliable and thus could not be assumed to be fully trustworthy. For this, detecting misbehaviour on V2X data is one key element toward assessing trustworthiness on V2X data and data generators but it is not sufficient to build a reliable and robust trustworthiness estimation of the whole system. For this, there is an urgent need to combine the output of misbehaviour detection with the output of other sources of trust. Finally, further studies are needed to confirm the benefits of **deploying misbehaviour detection modules, as a source of trustworthiness**, at the CCAM services level. To the best of our knowledge, this kind of studies, which are one of the concerns in CONNECT, are not yet published in the literature.

2.2.4 Introducing the Edge Computing Concept

The present chapter does not aim to provide an exhaustive survey of the state of the art in Mobile Edge Computing (MEC) research, given the extensive and continually evolving nature of this field. Instead, it offers an introductory overview by providing an indicative breakdown of research directions. This approach allows us to touch upon key areas of interest within MEC without delving into the comprehensive landscape of ongoing research, setting the stage for a more focused exploration of MEC in the subsequent sections of the deliverable.

Edge Computing represents a transformative distributed computing paradigm with the primary objective of relocating computing and storage resources from the distant and centralized cloud infrastructure to the proximity of the data's source. In other words, it aims to position these capabilities closer to the user, the user equipment and the processes that generate and analyse data. This fundamental shift in computational infrastructure enables more efficient and responsive data handling, as it takes place at the very edge where data is being collected and initially processed [34]. Although the edge computing paradigm is naturally adaptable and can support different communication technologies, there are some situations where its adoption becomes more of a necessity rather than an option. Traditional cloud-centric application deployment is inadequate for achieving demanding criteria such as low roundtrip latency, particularly in the context of Vehicle-to-Everything (V2X) communications and Internet of Things (IoT) contexts. In these contexts, data volumes can escalate rapidly, resulting in excessive bandwidth consumption when transmitting data to remote cloud locations. Moreover, misbehaviour instances (see paragraph 7.2) are frequently limited to a specific geographic area. Consequently, the process of sending data to a remote cloud infrastructure becomes inefficient and unproductive. As a result, a growing range of distributed applications, including autonomous driving, distributed robotics, video analytics, and augmented and extended reality applications, are increasingly adopting edge computing to leverage its benefits.

The aforementioned notion, specifically in the context of mobile communications, was standardized and defined by ETSI under the title "Multi-access Edge Computing (MEC)"[1]. The 3GPP, a global consortium of standards organizations, has also contributed to the development of an edge computing architecture that bears significant similarities to the ETSI MEC framework. In this context, the MEC Platform aligns with the 3GPP Edge Enabler Server, while MEC Applications correspond to the Edge Application Servers defined by 3GPP. Notably, the 3GPP service environment is

[1] https://www.etsi.org/technologies/multi-access-edge-computing

positioned at the edge of the Next Generation Radio Access Network (NG-RAN), and its computing capabilities are inherently integrated into the 5G Core[2]. It's worth mentioning that CONNECT aligns more closely with European developments, as evidenced by its collaboration with ETSI in advancing the MEC[3] framework. The MEC has gained substantial recognition and attention because of its profound scientific, technological, and commercial implications [35]. As a result, it has become a central focus in the study and development of forthcoming 5G and beyond-5G (B5G) technologies. Indeed, it enables rapid and flexible deployment of 5G applications having the 5G core (data plane) functions orchestrated close to the data sources, i.e., cellular local breakout. Overall, service deployment at the edge of the network can lower operating costs, reduce energy expenditure, save bandwidth, reduce application latency, and enhance the user (and service) quality of experience (QoE).

Along these lines, the research background on the MEC has so-far addressed a broad[4] set of involved challenges. One can roughly identify (at least) *three* focus areas of MEC research:

A) MEC resource allocation and scheduling, which mainly tackles the management of MEC resources to achieve latency reduction, optimise processing power (compute), and energy (monetary) savings. A number of proposed schemes, optimise slicing resources of MEC servers sharing slices among multiple MEC operators [36], address the data redundancy problem through collaborative task computing that utilises idle MEC resources [37] or even manage (according to user requirements) virtualized resources for virtual reality streaming on MECs [38]. The area also includes the **computation-task offloading problem,** which is typically explored along the axes of decision-making, resource allocation, and mobility management between different MNOs and MEC SPs (details appear in the CONNECT D5.1). The utilisation of seamless *task offloading* is observed in various contexts, including the internal structure of vehicles and external infrastructures such as MEC or Cloud systems. One significant advancement is the advent of the Digital Twin, a novel virtual representation of the tangible vehicle, which signifies a substantial expansion in the domain of offloading capabilities. Additional information regarding Digital Twins is elaborated upon in section 2.2.5 of the current deliverable. The act of offloading serves to enhance the decision-making capabilities of individual vehicles, and also has wide-ranging implications within the broader domain. The result is an increased level of efficiency and performance in various system functionalities and services [39]. The dynamic approach, specifically in the context of Digital Twins, has the ability to allocate additional resources to individual vehicles and greatly enhance the speed of computation, precision, and accuracy of the provided services [40]. The cumulative outcome is an enhanced level of decision-making, spanning at a region-level. Moreover, within the realm of mobility management, MEC assumes a pivotal role in its facilitation. However, the integration of MEC and its associated building blocks introduces a substantial augmentation in the complexity of establishing trust relationships across the diverse actors and within the CCAM ecosystem.

B) Support of MEC-enabled distributed ML pipelines. Efficient utilisation of resources has been studied in relation to the ML convergence rate [41] while a stochastic minimization of the number of federated learning iterations between the learners and the MEC is proposed in [42], subject to ML accuracy guarantees. Along this thread, a large body of works has appeared lately in light of the ever-increasing interest in ML studies.

C) MEC relevance to network security and data integrity. MEC networks exhibit significant relevance in the context of network security and data integrity, encompassing essential security aspects such as confidentiality, integrity, and availability of MEC-stored information. These networks also incorporate mechanisms for user authentication and authorization [43]. Relevant research has identified a wide array of threat vectors spanning from the access network to the mobile edge and mobile core, with a particular focus on MEC virtualization technologies (VNFs). These concerns are generally addressed through the utilization of Trusted Platform Modules (TPMs) [44]. Furthermore, security measures can be integrated into the virtual infrastructure when creating virtual entities, such as containers or virtual machines (VMs), as discussed in paragraph 6.4.1.2. It's important to note

[2] https://www.3gpp.org/technologies/edge-computing#_ftn1

[3] In this document we persistently use the term MEC for the sake of a clear -to the reader- reference. What we mean is in-general any edge computing infrastructure in the user vicinity hosting trust/security/CCAM applications, even tolerating reasonable technical diversion from the typical MEC specification proposed by ETSI.

that lightweight virtualization technologies may offer fewer security mechanisms, indicating a need for further research and experimentation in this area. Additionally, MEC networks face privacy challenges related to exchanged and stored data, user identity, positioning, and mobility. Various privacy-preserving techniques are employed, including mutual agreements for anonymous authentication keys [45] and the exploration of blockchain solutions [46].

Towards the direction of **MEC infrastructures tailored for V2X communications**, the deployment of MEC hosts across various Mobile Network Operators (MNOs) has become a common occurrence within the expansive V2X ecosystem. This is primarily driven by the wide geographical coverage of V2X communications and the necessity for continuous connectivity. In specific cases, there is a need for MEC hosts from different MNOs to cooperate in sharing information for ensuring smooth and reliable delivery of V2X services.

One example of such a scenario pertains to **cross-border communication**, wherein vehicles originating from different countries interact while traversing international borders. This particular situation requires the seamless transfer of MEC hosts associated with various MNOs, as vehicles move between different geographical areas. In order to enhance this transnational collaboration, it is imperative to have effective data traffic steering mechanisms in place to guarantee the provision of uninterrupted service with minimal latency. Another paradigm is found in the intersections within densely populated urban areas. Real-time communication and synchronisation between vehicles and infrastructure are imperative for the effective management of traffic. Considering the unlikelihood of a single MNO having the capacity to supervise all city intersections, it becomes necessary for multiple MEC hosts, belonging to different MNOs, to participate in the supervision and exchange of crucial traffic-related data.

The MEC4AUTO working group, a part of the 5G Automotive Association (5GAA) [47], has conducted a comprehensive analysis to examine various scenarios that demonstrate the **operational complexities of MEC**. These scenarios clearly highlight the possibility of deploying MEC in various geographical areas. This specific dynamic serves as an essential example of the **complex web of trust relationships** that require careful evaluation. The framework presented delineates three primary MNO scenarios, wherein two MEC platforms are involved alongside a shared MEC application denoted as X in Figure 7. The primary objective of these scenarios is to delineate the architectural framework and collaboration strategies between two MNOs, whether they possess distinct or similar infrastructures. This aims to streamline the management of data traffic effectively, ultimately enabling the smooth transitions of V2X services as vehicles traverse diverse operational environments. The scenarios are designed to ensure a harmonious and uninterrupted flow of V2X services, even when vehicles move across different MNO domains, by defining the necessary infrastructure and cooperation mechanisms.

The aforementioned scenarios can be further investigated to obtain an initial comprehensive understanding of the intricacy of **trust relationships and the corresponding trust requirements in scenarios involving MEC**. Trust requirements encompass the need for a MEC service, which can be implemented in three ways: i) The MEC service can be distributed across multiple infrastructures, each protected by distinct security mechanisms that offer varying levels of assurance.

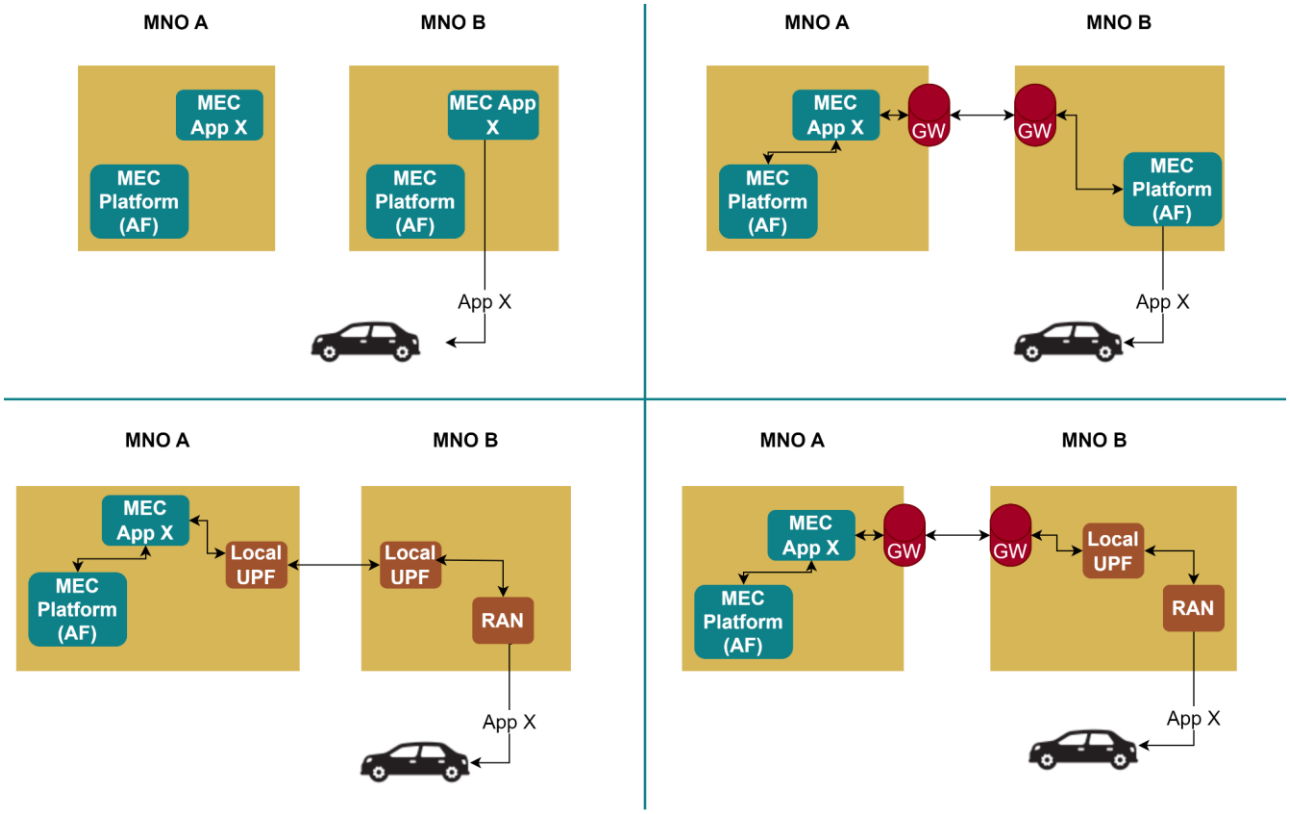

Figure 7 - MNOs with/without MEC platform and MEC application

Alternatively, ii) the MEC service can be deployed in infrastructure A, but accessed by MEC platforms from infrastructure B. In this scenario, stricter security measures are applied to establish a secure and authenticated channel. Another possibility is iii) the MEC service switching between operator A and operator B. This introduces concerns regarding service consumption and privacy.

In the CONNECT context, edge computing is utilised as part of the networking infrastructure, offering a prominent location for hosting instances of both trust computation services (see Chapter 6) and CCAM applications. The introduction of MEC and the corresponding stakeholders and building blocks increases the number of trust relationships that need to be established between the actors and the CCAM ecosystem. There is a plethora of initiatives, for instance PODIUM project [48], where focus is placed on showcasing the exact functional benefits of introducing edge computing to host day3 CCAM applications.
Obviously, the utilisation of the MEC increases the complexity of the system; hence, the trust relationship. Nevertheless, as mentioned in Trust Assessment in CCAM2.2.2, trust assessment is highly crucial in the automotive industry which is composed of multiple and heterogeneous entities and data sources. Towards this direction, even though MEC offers many advantages in the V2X ecosystem, its effects on trust shall be specifically scrutinised. A region-based Trust Assessment Framework, responsible to provide a trust evaluation, based on operational and security assurances (i.e., through verifiable evidence), requires further resources as well as information and updates from various sources within a common region. Contrary to the vast majority of the so-far research works, CONNECT invests effort on a combination of trust modelling principles (see paragraph 6.3), efficient attestation mechanisms (see paragraph 6.4) and advanced trusted environments for virtualized (MEC) resources (see paragraph 6.5.2).

2.2.5 Digital Twins

Similarly to Edge Computing (i.e., as discussed in section 2.2.4), the notion of Digital Twins (DT) has emerged as a disruptive technological paradigm with extensive applications that encompass a wide range of industries. The ISO/IEC 30173 standard defines digital twins as follows: “A digital

representation of a target entity with data connections that enable convergence between the physical and digital states at an appropriate rate of synchronisation".

According to the StandICT landscape report [49], digital twins are recognized as highly innovative and promising concepts within the broader context of digital transformation. The term "Digital Twins" was introduced by Professor Michael Grieves from the University of Michigan in 2002 and gained significant recognition, particularly from NASA, around 2010. The fundamental idea behind the Digital Twins concept involves the creation of a virtual replica or simulation of physical objects in the digital domain. The virtual counterpart exhibits an impressive ability to integrate real-time data and advanced analytics features, accurately replicating the behaviour and attributes of its physical counterpart. Industries, including manufacturing and product development, have shown a strong inclination towards harnessing the capabilities of Digital Twins. These industries utilise Digital Twins to proactively identify defects or malfunctions, optimise maintenance approaches, and reduce operational interruptions. Three notes are further added to the definition:

- ✓ Digital twins have some or all of the capabilities of connection, integration, analysis, simulation, visualisation, optimization, collaboration, etc.
- ✓ Digital twins may provide an integrated view throughout the life cycle of the target entity.
- ✓ Target entity, which provides some functional purpose in reality, can be either a physical entity or a digital entity under consideration.

The DTs landscape encompasses four key research directions [50]. The first focuses on coordinating **cloud-edge-end** DTs, leveraging cloud resources for computational power and edge capabilities for real-time synchronisation, and enhancing intelligent application systems. The second emphasises **space-air-ground** DTs, addressing cross-domain authentication, integrated sensing, communication, and computation, with the incorporation of blockchain technology for efficient interactions. The third highlights **interoperable and regulatory** DTs, emphasising exchange capabilities and the necessity of creating new standards and processes to ensure seamless operation across diverse blockchain networks. Regulatory aspects can be streamlined through misbehaviour detection, AI-driven decision-making, automated enforcement, and decentralised governance, involving elements such as soft law, **explainable AI**, smart contract security, consensus mechanisms, and controlled blockchains. The fourth research direction delves into enhancing DTs through explainable AI, focusing on understanding AI model components' semantics, and generating explanations to enhance transparency and understanding.

It is important to note that the utilisation of Digital Twins is not without its inherent difficulties. Prior to investing in such a solution, it is crucial to have clearly identified the advantages of this expensive strategy as well as the requirements across the entirety of the life cycle. To this extent, it should be considered that support for a DT should be provided across its complete product lifespan. However, currently not all stages are supported (i.e., especially the early and the later stages) [51]. In addition, achieving accurate and two-way synchronisation is not an easy task. A significant number of resources is needed to support the wide range of software tools used in any industry that wants to benefit from such a solution. Achieving smooth interoperability between these outdated systems and digital technologies is a challenging undertaking that might result in delays during the installation process. Hence it becomes apparent that the complex characteristics of DTs need the incorporation of diverse elements, real-time instruments, collaborative optimization approaches, and extensive data resources, which may lead to lengthy deployment procedures that demand meticulous handling by companies.

Furthermore, cybersecurity is of utmost importance, as it involves ensuring the security of digital ecosystems, defending against cyber-attacks, and preventing data breaches in situations when DTs extend across several industrial partners and inventory locations. In light of the considerable aggregation of sensitive information within these digital constructs, it becomes crucial to prioritise the protection of these assets against cyber threats and unauthorised entry. The implementation of strong security measures and encryption protocols is essential in order to maintain the integrity and confidentiality of the data associated with Digital Twins. The security and privacy of digital twins are currently being extensively researched, even at the level of standardisation. An example of this is the ISO/IEC PWI 27568 initiative, which focuses on addressing these issues.

To fully harness the capabilities of Digital Twins in many fields, it is crucial to directly confront and tackle the aforementioned issues. The integration of DT within the automotive sector has sparked notable exploration. An example of a usage of a DT for an Internet of Vehicles (IoV) implementation can be found in [52]. In this study, a vehicle, with cloud-based DT functionality, is proposed for establishing an in-car autonomous traffic control ecosystem. The overall goal of this DT is to promote safety at all times. Towards this direction, the traffic system model is based on inter-vehicle networks and takes into account a multitude of factors, including vehicles, roads, pedestrians, climate, and geographic factors. The system operates in real-time to monitor traffic conditions. Following this approach, the server may improve situational awareness, forecast the viability of certain plans, and accelerate decision-making. These functionalities may significantly reduce the number of road accidents. Figure 8 illustrates an example of such a scenario according to [52]. It should be underlined that the architecture is based on a backend cloud infrastructure for supporting these tasks.

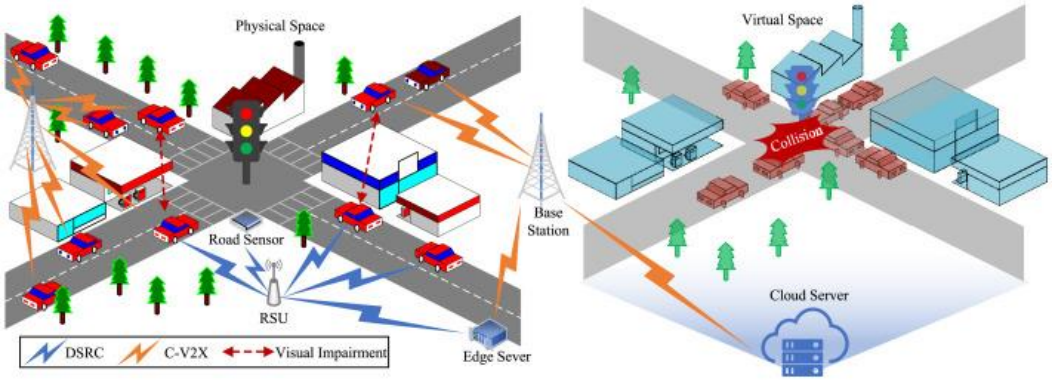

Figure 8 - Example of a traffic system with a DT [52]

Apart from cloud-based DT architectures, research has been conducted in the integration of DT in the context of MEC. The MEC provides further capabilities in this sense, promoting the efficiency of the service execution, by bringing it closer to the vehicle, but overcoming in parallel the local challenges caused by limited resources. In [53], a novel architectural framework is presented that combines the principles of Vehicle Edge Computing (VEC) with the capabilities of DTs. In this particular model, DTs are utilised to replicate the virtual models of vehicles and Roadside Units (RSUs). The integration of VEC and DTs presents three significant benefits: firstly, the capacity to monitor vehicles in real-time; secondly, the establishment of a virtual twin layer that connects physical VEC entities with vehicular applications; and thirdly, the provision of insights into the dynamic network topology. The utilisation of this dynamic framework enables VEC networks to flexibly adjust schedules, incorporating tasks offloading, resource allocation, and network management. The architectural framework comprises a tangible VEC network, a parallel network layer, and a layer dedicated to vehicular applications. The twin network layer is responsible for tasks such as model construction, item mapping, and strategy optimization. The primary aim of this replication endeavour is to effectively incorporate and evaluate reinforcement learning (DRL) techniques in order to improve adaptive and intelligent decision-making with regards to offloading strategies. Leveraging digital twins, offloading processing and enhancing decision-making capabilities in wireless edge networks can be achieved, as further demonstrated in various studies such as [54], [55], [56], and [57].

The introduction of DTs in the automotive sector brings forward additional complexities related to the V2X communications. V2X technology encompasses various forms of communication, including inter-vehicle, vehicle-infrastructure, vehicle-pedestrian, and vehicle-road user interactions. This technology facilitates the exchange of data in a seamless manner, thereby enabling the implementation of advanced safety applications. Through the adoption of DTs, automotive manufacturers and service providers can significantly augment safety and operational efficiency. By creating digital replicas of vehicles and their immediate environments, they gain the ability to simulate

and analyse diverse driving scenarios, evaluate the performance of autonomous systems, and optimise overall vehicle functionality.

Relevant DTs examples can be found in the Tesla digital twin [58], a digital twin of networking capabilities exemplified by the 5G-CARMEN project [59], which replicates the 5G physical network to support the study and optimization of real-world use cases, or a digital twin of the traffic system as demonstrated in the DIGEST project [60].

As highlighted in [61], the field of vehicular DTs specifically emphasises two aspects of the communication design related to: i) the communication within each twin (i.e., **intra-twin**) and ii) the communication between twins (i.e., **inter-twin**). The first covers the essential exchanges between a digital representation and its corresponding physical vehicle, hence a reliable and synchronised channel is essential. The second case refers to the communication between the digital entities on the edge or cloud infrastructure. This necessitates the development of capabilities for effective peer-to-peer communication, support for multi-agent interactions, and the secure protection of private data assets.

Regarding security challenges, it's essential to highlight the critical issue of privacy leakage attacks, which can pose a severe threat to DTs. These attacks may enable malicious actors to intercept and collect information exchanged between DTs, which might possibly lead to a breach of user privacy. For instance, an attacker could gather information regarding the current location of the vehicle, thus deduce its movements. This information could then be used to make educated guesses about the user's personal interests, hobbies, or behavioural patterns. Given the sensitive nature of this data, implementing robust security mechanisms becomes absolutely imperative to safeguard against such privacy breaches.

In the context of CONNECT, the notion of Digital Twins (DT) contributes to the exploration of the MEC capabilities, as introduced in section 2.2.4. The MEC instantiates the DT not only for task offloading and migration, particularly for the TAF, but also to facilitate the efficiency and scalability of the CCAM ecosystem (as more elaborately discussed in D5.1). Consequently, some of the core DT-related challenges, such as synchronisation, become less pivotal in this context. The concept of the Digital Twin showcases an inherent ability to accurately replicate not just the capabilities, but also the functionalities of the vehicles or infrastructures it represents.
Various architectures have been proposed to support DT capabilities in the automotive sector, forming the basis for CONNECT. However, the proposed architectures will be expanded to address the specific needs of CONNECT. Notably, there still exists a substantial lack of convergence within these architectural approaches, presenting an opportunity for CONNECT's exploration and integration. Although there are existing standards and frameworks discussing the diverse applications of Digital Twins across various domains, a critical need persists for the establishment of a comprehensive and unified architectural perspective.
In line with CONNECT's vision, the use of DTs extends to support a variant of the Trust Assessment Framework, known as the TAF-DT, as defined in D3.1. The TAF-DT offers a virtual representation of a Vehicle as a Digital Twin. This DT serves as a supporting element, allowing resource-intensive trust calculations to be offloaded from the Vehicle TAF. By doing so, collective trust calculations and trust-aware decision-making strategies can be formulated to improve the overall system's performance. Further details on the TAF-DT will be described in D3.3. Details regarding the architecture of CONNECT's DT are described in chapter 6 of the present deliverable.

2.2.6 Towards Zero Trust Environment

As indicated in sections 2.2.4 and 2.2.5, the incorporation of novel technologies such as the MEC and the DT into V2X communications has introduced numerous opportunities (i.e., performance, efficiency, etc.). Nevertheless, it has further posed various challenges, specifically in terms of trust establishment due to the inclusion of additional entities. As a result, the traditional trust assumptions, relying on centralised infrastructures or presuming the invulnerability of certain components, are no longer applicable in the dynamic and decentralised nature of modern MEC-enabled V2X systems, considering the multitude of actors as well as the numerous stakeholders (i.e., OEMs, manufacturers, MNOs, service providers, etc.). Previously, trust assumptions extended to concepts like Public Key Infrastructure (PKI) solutions used for V2X communications, relying on central

authorities, and assuming the main On-Board Unit (OBU) within vehicles could not be compromised, leaving CAN messages unprotected. **However, the evolution of day 3 operations has complicated the landscape, necessitating a paradigm shift in trust assumptions.**

In response to these challenges, the notion of zero trust has gained popularity, even in the context of V2X communications. Adopting a "*never trust, always verify*" approach, zero trust architecture implements security measures at all network and infrastructure levels, regardless of the user or resource's location. It treats any user, device, or application attempting to access resources as untrusted until verified through continuous authentication and data/resource attestation. This approach strengthens security guarantees. Applying the zero-trust concept to V2X presents certain difficulties. The heterogeneity of the V2X landscape, involving multiple services, MNOs, and vehicles, results in varying data sharing profiles and patterns between infrastructures and vehicles, further complicating trust considerations (see Figure 9). Prior to the emergence of MEC, the cloud served a crucial role in enabling the exchange of data profiles among collaborating vehicles and between vehicles and the infrastructure. Nevertheless, the incorporation of MEC brings forth a novel aspect, characterised by the inclusion of additional layers. This primarily encompasses the dynamics between vehicles and MEC, as well as the complex interactions among multiple instances of MEC. Currently, the predominant focus of trust assessment mechanisms concentrates on individual vehicles or standalone systems. Although reputation-based systems have been proposed as potential solutions for evaluating trust, they are still susceptible to facing similar limitations.

It becomes apparent that in order to create a comprehensive framework that includes both security and safety, it is necessary to integrate the trust paradigm throughout every level of the communication and networking stack, starting from the vehicle and expanding all the way up to the MEC and the cloud. To embrace the zero-trust notion holistically, **dynamic and flexible trust assessment mechanisms are needed to accommodate the needs of modern V2X systems.** These mechanisms must provide a comprehensive assessment of trust across the entire ecosystem while handling data with high uncertainty or contradictory information effectively. Furthermore, it is essential to recognize that **different services within the V2X ecosystem may have diverse data sharing profiles when using MEC.** These services range from misbehaviour detection to congestion control management, each with specific data sharing needs and security requirements. Therefore, adherence to specific security profiles becomes imperative to ensure the integrity and confidentiality of the exchanged information.

For services dealing with sensitive data or safety-critical functions, stringent security profiles with strong encryption methods, access control mechanisms, and secure authentication procedures should be enforced to prevent unauthorised access and data breaches. Continuous monitoring of the data sharing process is essential to detect and mitigate potential vulnerabilities or threats promptly. Additionally, services involving the exchange of personally identifiable information (PII) or location data require special attention to privacy requirements. Compliance with data protection regulations and privacy policies is crucial to protect the privacy and anonymity of individuals involved in V2X interactions. Incorporating anonymization techniques, data minimization, and consent management into security profiles ensures the appropriate handling of sensitive data.

As V2X services increasingly rely on the MEC infrastructure, it is crucial to establish standardised security and privacy profiles tailored to the diverse needs of different applications. These profiles should accommodate various use cases and specific data sharing requirements associated with each service. By implementing robust security and privacy measures aligned with defined profiles, the V2X ecosystem can maintain a trustworthy and secure environment, fostering innovation and the advancement of cutting-edge V2X services.

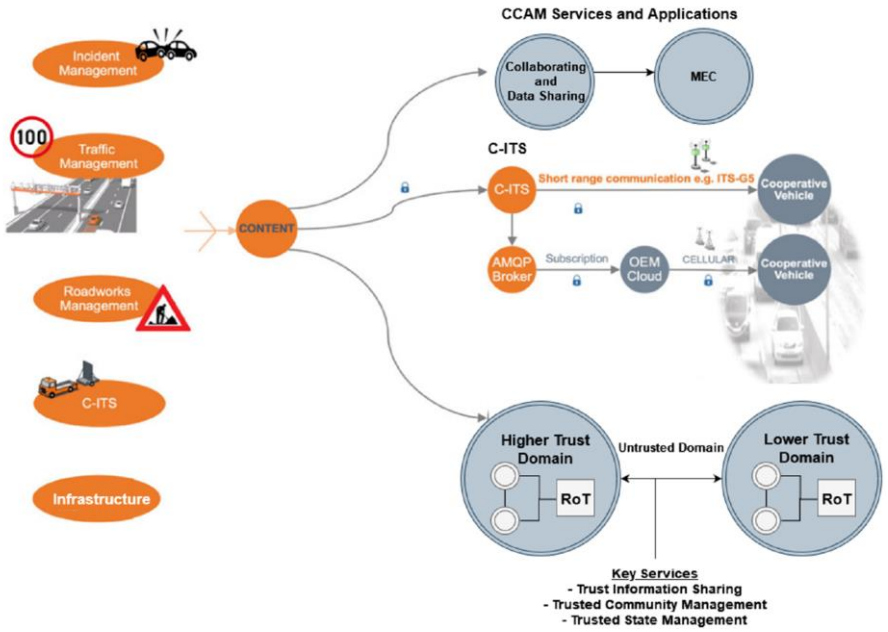

Figure 9 - CCAM Services in different trust domains

2.2.7 Confidential & Trusted Computing

Towards the direction of achieving a zero-trust network architecture for the V2X ecosystem, the confidential and trusted computing notion is pivotal. As the system moves away from the centralised paradigm, the utilisation of decentralised entities as trust anchors, along with the provision of verifiable evidence, is a prerequisite for the Trust Assessment Framework. The concept of confidential and trusted computing emerges as a catalyst in this endeavour, providing the critical enabling capabilities. To achieve a decentralised and dynamic trust assessment, security mechanisms such as cryptographic protection, continuous authentication, validation of system components, hardware security, and misbehaviour detection all work in tandem. In addition to security, privacy guarantees are equally crucial. These guarantees should not just focus on safeguarding sensitive information and handling individuals' personal data appropriately within the system, but also ensuring that data sent outside the vehicle cannot be used to track or otherwise monitor the vehicle itself. However, the dynamic trust assessment framework may create a privacy-security trade-off, as more information sharing may be required for entities to be considered trustworthy, affecting privacy. Striking the right balance between trust and privacy is essential.

Trusted Computing is employed to support the zero-trust concept, increasing the trust per entity while maintaining a decentralised environment. Trusted Computing is used to guarantee secure storage, to oversee the execution of computations and to provide the trust evidence and anchors necessary for trust assessment. The secure storage will be used for system measurements and data, and for the management and control of keys used for ensuring the protection of systems and communications and reporting on the assessment of the system's state correctness. Trusted computing makes use of secure enclaves, Trusted Execution Environments, or Trusted Platform Modules to act as Roots of Trust (RoT) that offer strong encryption mechanisms and ensure the protection of data and functionalities even in untrusted settings. These solutions also provide evidence that can be trusted by other entities, asserting that no unauthorized modifications have occurred, and the device operates in the correct state.

The Root of Trust (RoT) is an essential element in both vehicle and cloud infrastructures, whether they are software or hardware based. It plays a critical role in providing important features such as secure storage and establishing the foundation for a secure system. Expanding upon the concept of the RoT, there exist software stacks that make use of these inherent capabilities. The current research endeavours revolve around investigating the interoperability and communication mechanisms between different Root of Trust (RoT) components, including the Trusted Platform Module (TPM) and Trusted Execution Environments (TEE). The primary objective is to enhance the

integration of diverse RoTs, particularly in heterogeneous environments. This is not an approach that has been taken before, but it will be important, for enabling state migration from one RoT to another in the case of failure, or to support interactions between the vehicle and the Mobile Edge Computing (MEC) or cloud and, in particular, the setup of Digital Twins (DT).

In this context, it is essential to prioritise the safeguarding of communication by implementing and enforcing mechanisms that guarantee the integrity and security of data at various endpoints. The transmission of data from the RoT of the vehicle to the RoT of the cloud, commonly known as Chip to Cloud assurance, remains an area of active investigation and standardisation endeavours. This is due to the requirement for strong security protocols and dependable communication channels in order to uphold a continuous chain of trust. How the protection of communication between diverse RoTs, implemented across distinct endpoints, can be efficiently achieved holds significant importance in the research endeavours of CONNECT. The protocols used for the establishment of resilient and reliable communication channels must be flexible enough to work with the range of roots of trust in the system.

On the infrastructure side, Confidential Computing offers much of the same benefits as trusted computing, process isolation and remote attestation, but on a much larger scale. It plays a crucial role in safeguarding the isolation and proper execution of virtualized services and containers. To address the security challenges posed by traditional isolation and virtualization techniques used in Cloud Computing, Confidential Computing leverages various technologies. Some solutions employ cryptographic primitives, while others rely on formal methods and verification techniques to ensure robust security and isolation. Additionally, hardware-based mechanisms offer further layers of protection [62]. In the context of the V2X ecosystem, Confidential Computing becomes particularly vital as it facilitates the attestation of the environment. This is of utmost importance in scenarios where sensitive data such as vehicle location, driving habits, and identifiers must be shielded from unauthorised access, interception, or tampering. By employing Confidential Computing, the V2X infrastructure can effectively protect this sensitive information and maintain the integrity and confidentiality of the exchanged data.

In V2X communications, diverse hosts, or services, especially those deployed on multiple trust domains, possess distinct security and safety requirements. As a result, they should be assigned varying levels of trust based upon the attestation mechanisms available to them and their cryptographic capabilities. The establishment of trust relationships between different trust domains is crucial to achieve a holistic view of security within the V2X ecosystem. A trust domain denotes a distinct and isolated boundary within the system where a level of trust can be established among the entities operating within that boundary. By following common security policies to protect shared resources, entities within the trust domain can collaborate securely and effectively.

An illustrative example of trust between different domains is the one in the context of MEC communication for efficient traffic management and coordination at intersections. Different MEC hosts may perform different tasks such as analyse real-time traffic data and exchange critical information with vehicles to optimise traffic flow, minimise congestion, and enhance overall road safety. To establish trust in this collaboration, between the vehicle and the MEC must prove their integrity, authenticity, and trustworthiness to the vehicles, and vice versa. Vehicles need to trust that the MEC hosts are genuine and authorised to provide the services. Similarly, MEC hosts must ensure that the vehicles communicating with them are legitimate and authorised to receive traffic-related instructions. This trust relationship is crucial to ensure that vehicles can confidently follow the guidance provided by the MEC hosts, leading to safer and more efficient intersection management.

The implementation of confidential and trusted computing aligns seamlessly with the principles of zero trust, actively working to continuously assess and verify the trustworthiness of both vehicles and MEC hosts. Within the MEC ecosystem, dynamic trust assessment mechanisms are employed to continually evaluate the behaviour and trustworthiness levels of vehicles and MEC hosts, enabling the swift identification of any suspicious or malicious activity. By strictly adhering to zero trust principles, this V2X collaboration gains the confidence to rely on real-time data sharing between vehicles and MEC hosts, resulting in seamless intersection management, heightened traffic safety, and efficient decision-making, all while upholding the highest standards of security and privacy.

2.3 Consortium's Shared Vision for CONNECT

CONNECT project envisages a future CCAM ecosystem in which trust plays a central role in ensuring the system's correct and trustworthy operation, which further aligns with the user and stakeholder expectations. The notion of trust refers to the degree of confidence that the users and stakeholders can have that the system operates as expected. Trust covers the whole application stack, including configuration and operational behaviour. It serves as the basis for user confidence in the performance and security of the system. In the past, security, safety, and privacy requirements were treated as separate entities, but as modern systems encounter mixed criticality issues and services, the need for their convergence becomes evident. This convergence is especially noticeable in V2X contexts, where services may display variable levels of trust depending on their alignment with specific requirements. For example, certain services may prioritize operational availability as their primary priority, while others may place a greater emphasis on strengthening security. The integration of diverse criteria is a crucial element in contemporary system design, addressing the complex demands of a rapidly changing technological environment.

CONNECT aims to overcome the challenges related to the overall trust to the system, by firstly addressing security, safety, and privacy requirements with minimal performance impact. The project adopts a holistic zero-trust approach, employing a dynamic and flexible trust assessment framework to evaluate trust across the ecosystem, especially when handling uncertain or contradictory data. Towards this direction, Trusted Computing, utilizing secure enclaves and Roots of Trust (RoT), becomes a pivotal component for ensuring data integrity and operational correctness. This approach introduces trust in real-time data sharing, leading to improved traffic safety, efficient decision-making based on the received information from other entities of the CCAM ecosystem, and the maintenance of high security and privacy standards. Within the CCAM ecosystem, the MEC is also considered in the cases where outsourcing tasks is required for effective resource management.

CONNECT aims to enhance the integration of security and safety in CCAM through the establishment of dynamic trust connections and a strong trust model. This trust reasoning framework enables entities to establish trust for executing safety-critical functions based on information received from other entities, that can be verified. The goal is to facilitate secure data sharing among data sources within the CCAM ecosystem, fostering trust where pre-existing relationships may be absent or insufficient. Such a framework will gradually introduce the new era of Autonomous Vehicles (Avs), paving the way to the establishment of day 3 secure and trustworthy operations and decision making.

3 Trust and Trustworthiness

3.1 Definition of Trustworthiness, Trust & CONNECT High-Level Trust Goals

The ability to perform dynamic and continuous trust assessment for V2X communications is a major pillar for the CONNECT project. This trust assessment extends beyond the limitations of existing systems, aiming to provide elevated levels of trust assurance. Such a system is particularly crucial in the context of safety-critical decisions made by connected vehicles. Overall, this trust assessment feature holds immense significance, especially considering the future adoption of autonomous vehicles (AVs) and their associated socio-technology systems. These technologies attempt to improve the user's road experience by providing optimised traffic management and intersection assistance, ultimately leading to better driving behaviour in urban areas.

Evidently, and as aforementioned in Chapter 2, as we are progressing to higher levels of automation in autonomous vehicles, there is a growing realisation that trust, and trustworthiness will play a pivotal role in enabling this path towards higher levels of automation which in turn enable better services and functionalities. Recent times have witnessed a notable emphasis on **incorporating trustworthiness and trust as integral concepts within technological design**. These two are not merely enabling better and more accurate decision-making in such complex safety-critical systems, but also contribute to the user acceptance. However, realising this vision necessitates the establishment of precise and unambiguous definitions for trustworthiness and trust. In line with this, CONNECT is actively exploring and definition of both trust and trustworthiness, which are the two main notions, aligned with the scope of the project. The initial explanation of these principles can be found in D3.1, marking a significant step forward in the area of future mobility.

These definitions must not only be theoretically sound but also highly practical within the complicated framework of CCAM systems. It is crucial to acknowledge that the dynamic and complex nature, hence the challenges as it pertains to CCAM systems (i.e., as presented in Chapter 2), demand a holistic approach towards trust and trustworthiness. This approach should accommodate a **wide spectrum of relationships along with a high volume of data**, spanning from interactions within a vehicle; within the CCAM environment in general and with the backend systems. Moreover, apart from the high volume of data; that should be assessed in terms of trust and trustworthiness, the assessment should produce results even when the data are contradictory. In order to reach its goals, CONNECT will use a systematic approach. The first step is a thorough examination at the most important questions that need to be answered when making a trust assessment framework. Also, it's important to remember that these **questions and challenges**, must be addressed in the context of **specific use-case cases**. This part is especially important as the **properties of trust** are use-case specific. The above sentence demonstrates how challenging and complex the trust evaluation problem is.

This chapter serves as a roadmap constructed around these pivotal questions and the inherent challenges that need to be overcome to develop a comprehensive trust assessment framework. The proposed framework shall align with the regulations for safety, security and privacy in autonomous vehicles. Thus, this chapter not only presents the current landscape but also underscores the critical aspects that require thorough examination in response to these questions and challenges. The following sections aim to provide the fundamental concepts of trustworthiness and trust definitions in the automotive sector. Table 1 serves as a comprehensive resource, systematically documenting the range of challenges pertaining to the area of trust within the framework of CONNECT. It begins with a fundamental examination of trust properties that emerge as most prominent within the scope of CONNECT, then proceeds to establish mechanisms to assure the correct operation of the trust assessment. More specifically, the process involves identifying the specific challenges that need further clarification, determining the questions that need to be resolved, and identifying the components that must be coordinated in order to provide a thorough evaluation of trust.

It should be mentioned that building trust in CCAM systems is a complicated and interdependent process with many steps, such as characterization, quantification of trust, capability operation and verifiability. Ensuring runtime trustworthiness is an ongoing process that requires **real-time monitoring and adjustment**, which can be influenced by the measurements and characterizations made earlier. The process is further hardened by the **high degree of interdependence and cross-over** between the steps. For example, in the context of challenges, verifiability is tightly connected to both characterization and quantification, necessitating the establishment of precise definitions and metrics. This interdependence further explains why the process of characterising, measuring, and verifying properties of trustworthiness in CCAM systems is not linear in terms of chronology. On the contrary, it is an **iterative and interconnected** endeavour. Insights gained from one challenge may prompt revisions or refinements in others, reflecting the dynamic nature of trustworthiness in CCAM systems. For instance, in the characterisation phase, defining what trustworthiness means for a specific system can impact how we subsequently measure and quantify those characteristics. As we attempt to quantify trustworthiness, we may discover that certain capabilities or trust relationships within the system need to be redefined or adapted to meet the desired trustworthiness criteria. Since trust and trustworthiness are dynamic in such complex scenarios as the ones introduced in CCAM set ups, there is a need for a holistic and adaptable approach to ensure their safe and reliable operation.

Table 1 –Trust-related challenges

Challenges	Important Aspect	Description
Characterization based also in the given context that trust is to be assessed	**What are the trustworthiness properties?** **Possible sub-questions to consider:** **1.** Which properties are the most essential to evaluate a system's/component's or relationship's trustworthiness? **2.** What factors determine the importance of a property or of a particular trust relationship (i.e., ethical principles, performance aspects, user acceptance criteria, regulatory compliance, etc)? **3.** What is the context that trust is to be assessed?	First, it is essential to establish a **clear definition** of both trust and trustworthiness, which will serve as the basis for the subsequent identification of **properties-to-be-used** for assessing the trust. These properties are enumerated in D3.1. They are the essential attributes that may crucially affect the system's overall trustworthiness level. However, it should be noted that not all properties influence trustworthiness equally. For instance, particularly in the context of CONNECT's envisioned Use Cases, integrity stands as a fundamental attribute that is common across all scenarios. It's noteworthy though that certain attribute, such as the privacy protection and its implications for adhering to regulatory requirements, may influence how **users perceive trust** within the system. Hence it becomes evident that the user perception is not always aligned with the technical attributes. In the previous example, a device that offers privacy guarantees ultimately enhances the perceived trustworthiness for users by addressing their concerns that are related to the potential misuse or unauthorised access of their personal data. To ensure a holistic approach to trust in CCAM systems, it's imperative to consider the user's perspective while defining properties of trustworthiness. In doing so, it becomes paramount that the collection of technical properties used to define trust levels does not cover all aspects. For example, the system performance or the data privacy are neglected. Nevertheless, any negative impact on performance, such as latency, due to the addition of security measures, could quickly affect how users perceive the system's trustworthiness. Similarly, if privacy requirements are not adequately met, users may lose confidence in the system's

		ability to protect their personal data, further undermining trust. This delicate balance between technical and non-technical aspects regarding system's trust and user's perception is pivotal for ensuring the successful adoption of such technologies. A system that excels in both is more likely to gain user acceptance and foster trust in this rapidly evolving landscape. All in all, it is vital to have a comprehensive understanding of the complex relationship between trustworthiness properties, their varying levels of impact in order to foster a thorough comprehension of how trust is established and maintained within the vehicles and as a result within the CCAM ecosystem**.**
Quantification of Trust	**What are the quantification means for trustworthiness; hence, how the acquired evidence/properties can be leveraged to quantify trust?** **Possible sub-questions:** **1.** What factors determine calculation of actual and required trustworthiness levels? **2.** Are such quantified trustworthiness levels dynamic (or are they affected by/during runtime or post initial deployment considerations)?	After determining the critical properties of trustworthiness and establishing mechanisms to capture properties in a manner that maintains security, privacy, while offering verifiability (i.e., characterization), the next crucial step is to leverage these properties to define the distinct trustworthiness levels. Whether for the device with its software components or for the hardware modules, a clear expression of Levels of Trust (LoT) is required. As previously discussed, it is crucial to prioritise the avoidance of potential implications that may affect user's perception (i.e., privacy), and arise from the pursuit of increased trust (i.e., from a technical perspective). Such implications have the potential to undermine the basic principles of the trust we aim to foster. Therefore, the quantification shall consider both technical and non-technical aspects, in order to offer a holistic view. In parallel due to the complexity of the CCAM ecosystem and the criticality of the trust-related decisions the assessment shall be performed in real-time to capture all possible modifications and updates. In parallel, the assessment is performed iteratively to further (re)consider the interdependencies between the different attributes of trust.
Capability	**How does the system/ component provide safety?** **Possible sub-questions:** **1.** How does the system provide for authentication?	In addition to the conventional principles (i.e., confidentiality, integrity, and availability), it is imperative for the device to effectively present the collected evidence regarding trust properties in a manner that ensures both **security and authenticity.** Towards this direction, and in order to provide the fundamental security guarantees, the system must adhere to protection mechanisms, and thus implement mechanisms for authentication and authorisation, as well as privacy preservation and operational failure notification. The inclusion of these crucial elements is imperative in mitigating potential failures and preventing

	2. How does the system provide for authorization? **3.** How does the system provide for privacy? **4.** How does the system provide for operational failure, including notifications? **5.** How does the system provide a means for detecting component and system operational state, failures and alerting?	attacks, not within the evidence per se, but rather within the process of **evidence gathering**. Authentication and authorization in particular, are of utmost importance, specifically the domain of **referral trust**. This is because referral trust necessitates adopting the trust opinions of other nodes. Consequently, these nodes must be authenticated and authorised. Within this particular framework, it is imperative that the device has the capacity to verify and validate the individuals or entities who are accountable for disseminating their opinions of trust. Furthermore, an equally significant aspect pertains to the **preservation of privacy assurances**. The operations of the device must be carefully coordinated to ensure the establishment of trust and evidentiary support, while simultaneously preserving the privacy of all individuals involved.
Runtime Operation	**What are the means to ensure trustworthiness during operation?** **Possible sub-questions:** **1.** How does the system provide secure runtime operation? **2.** How are all components updated to ensure trust before and after the update? **3.** How would you detect a potential compromise of a component at runtime? **4.** How to ensure trust after a component is no longer supported or at the end of its service life? **5.** How is a component destroyed when not used or requires replacement? **6.** How are users/owners notified of a potential component breath? Including 3rd party provide components?	The evaluation of trustworthiness and trust should further ensure **operational correctness and precision**. This implies that the framework should be able to respond to any change, during runtime. In parallel, for compliance with regulatory requirements, trust decisions are required to be conducted within a specific timeframe. In this regard, the system's architectural design and security protocols must be in accordance with established regulations, ensuring compliance with **regulations and standards** and further increasing the user's confidence. Equally important is that the trustworthiness and trust assessment process **does not interfere with the operational efficacy** of CCAM systems. A critical endeavour is to strike a careful balance between security measures and system efficiency. The system's model should be designed to achieve not only peak performance but also the highest levels of security and reliability, ensuring that both aspects work in tandem.
Verifiability	**What are the means to verify trustworthiness?** Possible sub-questions:	The incorporation of **verifiable mechanisms** within the framework for assessing trustworthiness and trust should be considered in order to enable the validation of trust levels. This practically underlines the need to have the appropriate guarantees in the collected trustworthiness evidence.

	1. What controls should be applied that will provide auditable means for validation for trust level? **2.** Is there an independent 3rd party certification process (or self-certification regime)?	Consideration should be given to the potential implementation of either a **self-certification or an independent third-party certification** process, in conjunction with the aforementioned measures. In the latter case, **privacy implications** are important to be considered, so that personal information is not leaked.

In general, **trust** can be conceived as a three-place relation involving a trustor (one who trusts), a trustee (one who is trusted), and the entrusted task or domain [63]. The general idea here is that trust is related to an expectation by the trustor that the trustee will achieve some entrusted tasks on behalf of, or for the trustor. **Trustworthiness** can be broadly conceived as a measure of the trustee's ability to achieve the entrusted task and respond to the trust placed in it by the trustor. Further, in some cases, the trust relationship may depend on the "entrusted task" to be conceived more broadly than just a performance outcome. Lee and See [64], for example, advance a 3P (performance, process, and purpose) model of trust in automation, signifying that the trustor's expectations from the trustee are a function of not just the performance (outcome) of the entrusted task, but also the process (through which the entrusted task was carried out), as well as the purpose for which the entrusted task was chosen and fits into the overall scheme of the technological system in consideration. It should not be noted though that for trust relationships to work successfully, trustor's expectations need to be appropriate or reasonable, otherwise there may be a threat for misuse and disuse [64]. It is critical to avoid, for example, overtrust, where the trustor's expectations exceed the trustee's capabilities.

There are at least two main aspects associated with trustworthiness of a given trustee: its ability to deliver the expected performance and secondly, the extent to which it is aligned with the goal of the trustor. For example, a given CCAM system needs to have the required technical ability to exhibit the relevant properties of safety, robustness, usability, etc. and this ability needs to be aligned with the expectations of the stakeholders, for example, the users of this system or the policymakers regulating the design and use of such systems, such as regarding what the appropriate level of safety is, or what criteria need to be satisfied in order to deem a system trustworthy.

Further, trustworthiness here is to be defined within a specific "context". Here, context refers to a restriction on a set of circumstances under which the trustee is expected to perform or achieve the given tasks [65]. That is, the trustee is not expected to fulfil expected tasks under all circumstances, but under a limited set of defined circumstances. For example, a CCAM system may be expected to conform to relevant safety standards under *proper conditions of use*.

Given this discussion, Trustworthiness can be defined as **the likelihood of the trustee to fulfil trustor's expectations in a given context**, where such expectations can be a function of the entrusted task, the process through which it was achieved, and the purpose for which the task was chosen.

In this context, expectations may encompass, for example, the correctness of data provided by a sensor (a trustee). However, expectations can also extend to the behaviour of the trustee, the process employed for entrusted tasks, and the purpose behind task selection.

Formally, given a trustor A and a trustee B, one can denote Trustworthiness of B for A's reasonable expectations regarding B's behaviour (Rx) in a context C as:

$$\mathbf{Tw_{B,A}(C) = The\ likelihood\ that\ B\ will\ exhibit\ behaviour\ R(x)\ in\ Context\ C} \quad (1)$$

Further, Trustworthiness can be a matter of degree or levels. That is, a trustee B may have the likelihood to fulfil the trustor's expectations to some degree or level L between 0 to 1, where 0 denotes no such likelihood and 1 denotes a maximal likelihood to fulfil trustor's expectations.

(1) can then be re-written as:

$Tw_{B,A}$ (C,L) – The likelihood that B will exhibit behaviour R(x) in Context C to a level L (2)

Moreover, assuming the principle of Zero Trust described earlier in Chapter 2, which means no initial implicit trust between nodes (e.g., vehicle A and vehicle B) is presumed, **verifiability becomes essential in establishing trustworthiness**. Verifiability relies on the trustee's (e.g., vehicle B) ability to offer evidence of meeting the trustor's (e.g., vehicle A) expectations in a measurable or demonstrable manner. The evaluation of trustworthiness depends on the trustee's capacity to provide verifiable evidence of fulfilling those expectations.

For the ideal/maximal evidence E, which would warrant appropriate trust in the trustee, we can write:

$Tw_{B,A}$ (C,L,E) – The likelihood that B will exhibit behaviour R(x) in Context C to a level L established by evidence E (3)

E can potentially have many sources. Some examples include:

- Evidence (direct or indirect) of B's past behaviour (ideally in context C, or a similar context) that is available to A
- An assessment made by an independent agent Z about B's ability (and willingness) to exhibit R(x) in context C made available to A (referral or transitive trust)
- Information about compliance with, for example, legal regulations that incentivize B to exhibit R(x) or disincentivize/prohibit B to deviate from exhibiting R(x))

3.1.1 Trust

A widespread view on the distinction between trust and trustworthiness among philosophers of *Trust* is implicit in the following statement by [66]: "*Trust is an attitude we have towards people whom we hope will be trustworthy, where trustworthiness is a property not an attitude*." Yet, statements such as these also carry the problem of generalizability and application to technical systems as they are too anthropocentric. It is hard to say in what way a technical component can have an "*attitude*" or "*hope*" for some desirable outcome.

The general point in the above quote of distinguishing trust and trustworthiness by deeming trust as an act (such as having an attitude) done by the trustor (that is, the one who trusts) versus trustworthiness being a property of the trustee (that is, one who is trusted) is, however, still worthwhile and reflected in other, less anthropocentric conceptions of trust and trustworthiness. Consider, for example, [67] (in p.31) comment on the matter: "*In a sense, trusting someone in some context is simply to be explained as merely the expectation that the person will most likely be trustworthy*". While Hardin is still talking about trust relationships between humans here, he invokes the same distinction as McLeod: trust is an *act*, of having expectations, towards someone who has the *property* of being trustworthy.

With this distinction in mind, one can then propose a non-anthropocentric conception of trust as follows:
Given two entities A and B, where A is the trustor (one who trusts) and B is the trustee (one who is trusted),
A Trusts B implies that A has expectations that B will have the property of being Trustworthy.

In other words, when A trusts B, A deems that the likelihood that B will meet A's expectation is very high, or higher than what may be required given A's expectations and risks taken by A. Again, in trusting B, it is critical that A's expectations and evaluation of B's trustworthiness is reasonable, appropriate and calibrated to B's actual trustworthiness. In our approach, we employ a Trustworthiness Assessment framework (TAF), which assesses the decision to trust an entity based on whether or not the entity meets the required level of trustworthiness. We explain this further in Section 3.2.2, and a more detailed explanation can be found in D3.1.

3.2 Methodological considerations towards assessment of trustworthiness and trust

One of the main concerns and challenges regarding evaluation of trust in the context of automation and automated systems is that **measuring trust** can be relatively difficult as trust is conceptualised as a "latent variable" which is hard to observe directly and rather needs to be inferred [68]. Kohn et. al. [68] discusses, for example, three broad types of indirect methods to measure and evaluate trust – namely, self-reports, behavioural measures, and psychological measures. Such measures, however, are better suited for trust relationships involving humans and automated components, rather than the diverse set of relationships, including relationships between automated components themselves, which is the primary focus of CONNECT. In this section, we explain three-steps methodological considerations that need to be exploited as part of the design of a holistic Trust Assessment Framework, corresponding to the challenges that are required to be resolved in order to derive trust assessment, as mentioned earlier in this chapter: Characterization (including context consideration), Quantification of Trust, Capability, Runtime Operation and Verifiability.

Please note that as part of this section we use the terms trust evaluation and trust assessment interchangeably.

3.2.1 Properties for Evaluating Trustworthiness

To evaluate trust, we first need to identify the specific properties for which we are interested in assessing trust. For example, assume that we need to assess trust between a zonal controller within a vehicle and a camera during the Cooperative Adaptive Cruise Control (CACC) function, where the zonal controller relies on the camera to deliver non-compromised camera data to it. Here, the camera needs to exhibit, among others, *the property of integrity* which implies that the transferred data has not been altered in an unauthorised manner.

In D3.1, we describe an indicative set of properties that are relevant for evaluation of trustworthiness of CCAM systems and their components. The list of properties described there has been extracted from sources such as documentation on standards (such as ISO 5723 [69], ISO/IEC 22624 [70] and ITU-T [71]), existing literature on autonomous vehicle systems and trustworthiness (such as [72]), and existing documentation on CCAM systems (*Cooperative, Connected and Automated Mobility)*. This indicative list of properties would be further refined as the project progresses based on different purposes for CONNECT.

3.2.1.1 Properties characterization

We categorise the proposed list of properties to evaluate trustworthiness in three broad categories:

3. **Performance-based** – These properties are linked to the performance. For example, reliability and robustness. Such properties are vital in CCAM to ensure the safe and efficient operation of vehicles. Properties like reliability, accuracy, and robustness are critical for providing consistent and dependable performance, while a property like resilience is essential for adaptability to various real-world scenarios, fostering user trust.
4. **Based on Ethical aspects** – These properties are linked to the ethical aspects and implications in the given context. For example, privacy protection and safety. Ethics-based properties play a crucial role in CCAM as they define the moral framework governing the behaviour of vehicles and other key components. Ethics-based properties are of paramount importance when considering trustworthiness due to their direct impact on public perception and societal implications. Properties such as accountability and transparency are essential for holding the system, and manufacturers, responsible and providing insights into decision-making, promoting accountability and regulatory compliance. Explainability ensures that system actions are interpretable to users and regulators, addressing concerns about the "black-box" nature of AI (or components based on AI-based technologies). Usability and authenticity reinforce the system's commitment to user objectives and protect against malicious actors, enhancing public trust in CCAM. These properties are essential to address concerns related to liability, unintended consequences, and the potential for unethical behaviour, which can significantly influence public trust and acceptance of automated vehicles. By upholding strong ethical principles, CCAM systems can build a foundation of trust with users and society,

promoting widespread adoption and contributing to the safe and responsible advancement of autonomous mobility technologies.

5. **Based on User acceptance** – These properties are linked to properties that have implications for acceptance of the overall system by the users. For example, transparency and usability. Such properties are paramount in CCAM to gain public confidence and adoption of automated vehicles. Privacy protection ensures the safeguarding of personal data, alleviating privacy concerns and respecting users' rights. Usability addresses the system's ease of use and interaction, making the technology accessible and user-friendly for a broader audience. Safety and security instil confidence in passengers by prioritising their well-being and mitigating cybersecurity risks. Relevance and consistency provide accurate and relevant information, bolstering user confidence in the system's capabilities. Recency and credibility emphasise the importance of up-to-date and trustworthy data, enhancing user trust in the information provided. Equitable access ensures fair market opportunities for various CCAM providers, fostering a competitive and diverse landscape.

These properties can be overlapping in the sense certain properties may belong to more than one category, or even in all three categories. For example, security-related aspects such as integrity and confidentiality are properties linked to the performance of the CCAM system, but they are also related to required ethical values for the system to exhibit, especially when personal and important private user data may be involved. In fact, the integrity of communication involved in CCAM systems has critical significance in terms of performance, protection of key ethical values such as safety and privacy protection, as well as for gaining user trust and acceptance.

3.2.1.2 Building trust relationships based on the concrete properties.

Trust relationship is a directional relationship between two trust objects that can be called trustor and a trustee. The trust relationship is always defined in relation to a concrete property. For example, even if we have the same trustor and trustee, the trust relationships would be different if we want to assess different properties.

The trustor is the “source” trust object as part of a trust relationship for which trust is assessed (one who trusts, the “thinking entity”, the assessor). The trustee is a “sink” trust object as part of a trust relationship for which trust is assessed (one who is trusted). If various trust relationships are combined, they form a trust network (also referred to as trust model throughout the project).

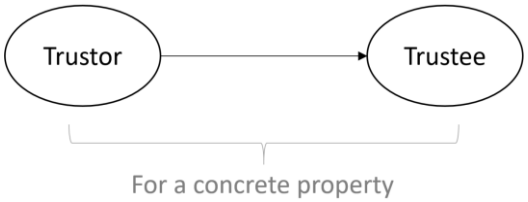

Please note that there is an implicit hierarchy between different properties. Concretely, we cannot create trust relationships based on all the properties that we have listed in D3.1. Namely, we can assess integrity as trust property and respectively build a trust relationship for the property of integrity. However, on the other side, there are properties like functional safety or reliability that can be measured and are causally related to (increased or reduced) trustworthiness. Increased trustworthiness in the system, can in result increase the property of safety. The concrete property hierarchy will be covered in more detail in the upcoming deliverable D3.2.

CCAM systems require evaluating trust on complex trust networks. However, the emerging scenarios in highly dynamic and distributed systems introduce novel challenges for trust assessment, such as the need to measure and evaluate trust and trustworthiness based on incomplete and subjective information from potentially untrustworthy sources. As part of complex trust networks transitive trust relationships are also introduced as a means to build trust indirectly through referral trust relationships within the trust network. Additionally, trust should be assessed for continuously changing trust relationships within the trust network, and the trust assessment shall support **node-based** and **data-based trust**, meaning that trust shall be assessed for both the entities that produce, process or relay data, as well as the data itself. Due to the distributed nature

of these systems, decentralised and distributed solutions are essential, and we described our choice of **Subjective Logic** as the mathematical theory for trust assessment due to its ability to handle subjective beliefs from multiple agents, merge conflicting evidence, and support trust transitivity on complex trust networks.

3.2.2 Defining and identifying evidence for evaluating a concrete property

As explained in the previous section, a key part of the assessment of trust and trustworthiness is defining the relevant property for a concrete trust relationship. However, to assess a specific trust relationship, it is important to include the **evidence** that is necessary to assess the particular property, i.e., the concrete trust relationship. This includes how the evidence is chosen, what is the best evidence or the best set of evidence to represent and assess a specific property, how evidence is managed, etc. Namely, depending on the property, appropriate trust sources need to be defined that provide the evidence for the fulfilment of the corresponding property. Decisions on trust are rarely made on a single parameter, and trust is always contextual. Thus, depending on the trust properties of interest, different trust sources are selected to do the trustworthiness assessment and quantify the trust opinion of the trust relationship.

In D3.1, we divide the trust sources into four categories: (I) trust sources related to communication, (II) trust sources related to system integrity, (III) trust sources related to applications and (IV) trust sources related to entity behaviour.

However, depending on the trustee, not all categories may be relevant. The trust sources of the first three categories are predominantly security mechanisms. In addition to security mechanisms, a fourth category was added that considers the behaviour of the node to get further evidence about the trustworthiness of the node. Some trust sources might require regular evaluations, while others only require one-time assessments at system startup.

When it comes to the evidence and the trust sources there are various research challenges that originate from different use cases in the project. For example, how are different trust sources chosen to calculate the trustworthiness of a trust relationship? Based on which trust sources and evidence can we quantify the fulfilment of a certain property within the trust relationship? How are the trust sources chosen in an automated manner in use cases where the trust relationships within the trust model change dynamically at run-time?

To summarise, evaluating trustworthiness and trust is the exercise of going through the various measures of trust applicable for a trust relationship, evaluating the levels of assurance, and if they meet the criteria set, validating the trust relationship. This is done based on evidence, received from the trustee as sources of trust, conveying information on the status of those properties of interest that can be used in a verifiable manner to calculate and quantify the trustworthiness. We detail more on verifiability of evidence in the following section.

3.2.3 Quantification of trust

In Section 3.1.1, we define trust as an act (such as having an attitude) done by the trustor (that is, the one who trusts) versus trustworthiness being a property of the trustee (that is, one who is trusted) is. Based on this definition, as part of the CONNECT project, we define *the act of trust* when the Actual Trust Level (ATL) is greater or equal to the Required Trust Level (RTL) and the act of mistrust otherwise. We explain this more in detail in the following.

3.2.3.1 Impact Assessment on Risks and Threats

In assessing the trustworthiness of an autonomous entity, a comprehensive approach involves both impact assessment during run-time operation and risk assessment during design time. Impact assessment focuses on evaluating the real-time consequences of the entity's properties in practical scenarios. It seeks to answer critical questions such as: What if the property of integrity, where data should not be altered without authorization, is not exhibited during the entity's operation? How would it affect the system's performance and overall reliability? Additionally, impact assessment addresses concerns about properties that might have been overlooked or not adequately taken into account

during system development. For instance, what if robustness or resilience is insufficiently addressed in the decision-making process, and how would it impact user trust and acceptance? On the other hand, risk assessment is performed during the design phase and aims to identify potential vulnerabilities, weaknesses, or uncertainties related to the entity's properties. It involves examining the likelihood of certain properties not being met and the potential consequences. Key questions to address may, for example, include: *What are the risks associated with insufficient privacy protection or security measures in the design that could compromise data integrity*? *How might goal alignment issues impact user experience and trust in the entity's data management*? Risk assessment also considers the availability and quality of evidence to support the entity's claim of meeting specific properties. For instance, what if there is insufficient evidence regarding the entity's authenticity or accountability related to data integrity, and how would it impact stakeholder confidence?

By combining impact assessment during run-time operation and risk assessment during design time, stakeholders can gain a comprehensive understanding of the entity's trustworthiness and identify potential areas of improvement. These evaluations allow for proactive measures to address issues and ensure the entity operates with the highest level of trustworthiness in managing data integrity, promoting user confidence and widespread acceptance of autonomous systems. More importantly, the risk assessment serves as a foundation for calculating the Required Trust Assessment (RTL) that is later compared to the ATL to do the trust assessment. The RTL quantifies, how much we **need** to trust, or in other words, what level of trustworthiness is required so we can proceed with the act of trust.

3.2.3.2 Trust Assessment

As previously explained in Section 3.1.1, Trust, the Trustworthiness Assessment Framework (TAF) assesses the decision to trust an entity based on whether or not the entity meets the required level of trustworthiness. In other words, if the Actual Trust Level (ATL) is greater than the Required Trust Level (RTL), then we proceed with the decision to trust. As part of the TAF, the ATL is calculated at run-time by the Trustworthiness Level Expression Engine (TLEE) which calculates the trustworthiness level, i.e., the trust opinion using different evidence and trust sources. For more details on the TLEE and the architecture of the TAF, in general, please refer to D3.1. On the other hand, the RTL is calculated based on different risk assessment methods, as explained in the section above.

3.2.4 Verifiability of evidence for evaluation of trustworthiness

The process of evaluating trust and trustworthiness involves assessing various measures applicable to a trust relationship. This evaluation relies on evidence from the trustee as sources of trust, providing verifiable information about the relevant property. For example, in the context of "Vulnerable User Protection through Cooperative Adaptive Cruise Control" (C-ACC), the integrity of data used for maintaining safe distances between vehicles becomes critical. Reactive security mechanisms like Misbehaviour Detection (MD) are in place to detect data integrity compromises, but they may be bypassed by clever attackers. To address this, a combination of trust sources is used to assess trustworthiness, i.e., to create a subjective trust opinion for the trust relationship.

With this, we highlighted the last essential methodological consideration towards assessment of trustworthiness and trust, which is **Verifiability**. Verifiability involves the trustee providing evidence justifying the trustor's decision to trust them. The element of verifiability is a key part of the approach that aims to define precise conceptions of applicable and relevant properties for evaluating a trustee and how they can demonstrate these properties. In response, investigation to attestation and trusted computing mechanisms is needed (see the models in Chapter 6).

Additionally, it specifies the evidence required by trustors for positive evaluation of trustworthiness, ensuring a robust trustworthiness assessment process. Verifiability, thus, is crucial for addressing key questions such as how does a given trustee in a given trust relationship, for example, exhibit the property of integrity? *What evidence is needed to demonstrate that this trustee indeed delivers data with integrity*? *How can this evidence be made available so it can be assessed by the trustor*?

Evaluating trustworthiness, therefore, involves verifying trust relationships through a comprehensive assessment of relevant properties and trust sources. The methodology emphasises verifiability, enabling trustors to make informed decisions about the trustworthiness of trustees based on concrete evidence.

3.3 Interplay between Trust and Privacy

According to [72], the issue of **trustworthiness in CCAM systems spans in three dimensions: i) the *technical*, ii) the *policy-making*, and iii) the *societal***. When it comes to trust assessment, the examination of the trust-related evidence, from a technical perspective, holds significant importance for enabling informed decision-making. Nevertheless, it is equally important to broaden the scope and consider additional factors beyond purely technical interpretations. More specifically, when discussing the notion of trust in CCAM, we cannot ignore the dimension of human trust from the side of the passenger that will eventually make use of the AV. In that respect, trust of people to the technology is a factor directly affecting the acceptance and adoption of AVs. Research has already demonstrated that the level of trust influences the acceptance of AVs [73] [74] [75] [76]

One compelling interpretation of trust revolves around the sense of vulnerability experienced by individuals inside a vehicle due to the loss of control. In that sense, trust is defined as "*the extent to which drivers willingly become vulnerable when using an AV"* [77]. Another interpretation of trust is closer to how we have defined trust in the previous sections, i.e., the degree of confidence drivers and passengers have in the predictability and functionality of the vehicle [78].

In order to better understand the human aspect of trust, [76] break down trust into three categories as follows: i) trust in the performance of the AV, ii) trust in the manufacturers of the AV, and iii) trust in the institutions responsible for regulating AVs. These dimensions of trust have been elaborated in previous research as well. Eiser et al. [79] point out that people might reject an innovation even if the technology is trustworthy, simply because the organisations behind the technology are not themselves considered as trustworthy. Liu et al. [80] adds another aspect to this dimension, raising the aspect of trust in jurisdiction. Hence, the concept of competence goes beyond mere ability, as it also includes the element of trust in governmental bodies tasked with formulating and implementing laws and regulations that assess the proficiency of these companies. These regulatory authorities grant certificates to brands that exhibit consistent adherence to the specified regulations. For instance, individuals can readily determine the extent to which different businesses adhere to the GDPR, highlighting the importance of proficiency and adherence to regulations as crucial factors in establishing confidence in the domain of data protection and adoption of technology.

Kenesei et al. [76] also explore the intricate interplay between trust and **perceived risk**. Indeed, when using an AV, the user should have sufficient trust that reduces the perceived risk of potential failure and misuse. More specifically, the authors examine two dimensions of risk: i) the perceived risk of the performance and hence security of the AV, and ii) the risk of misuse of the personal data that is exposed during use, which intersects with privacy protection considerations. Interestingly, their results indicate that privacy risk is influenced by trust in OEMs: trust in the manufacturer decreases the perceived risk of incorrect data handling.

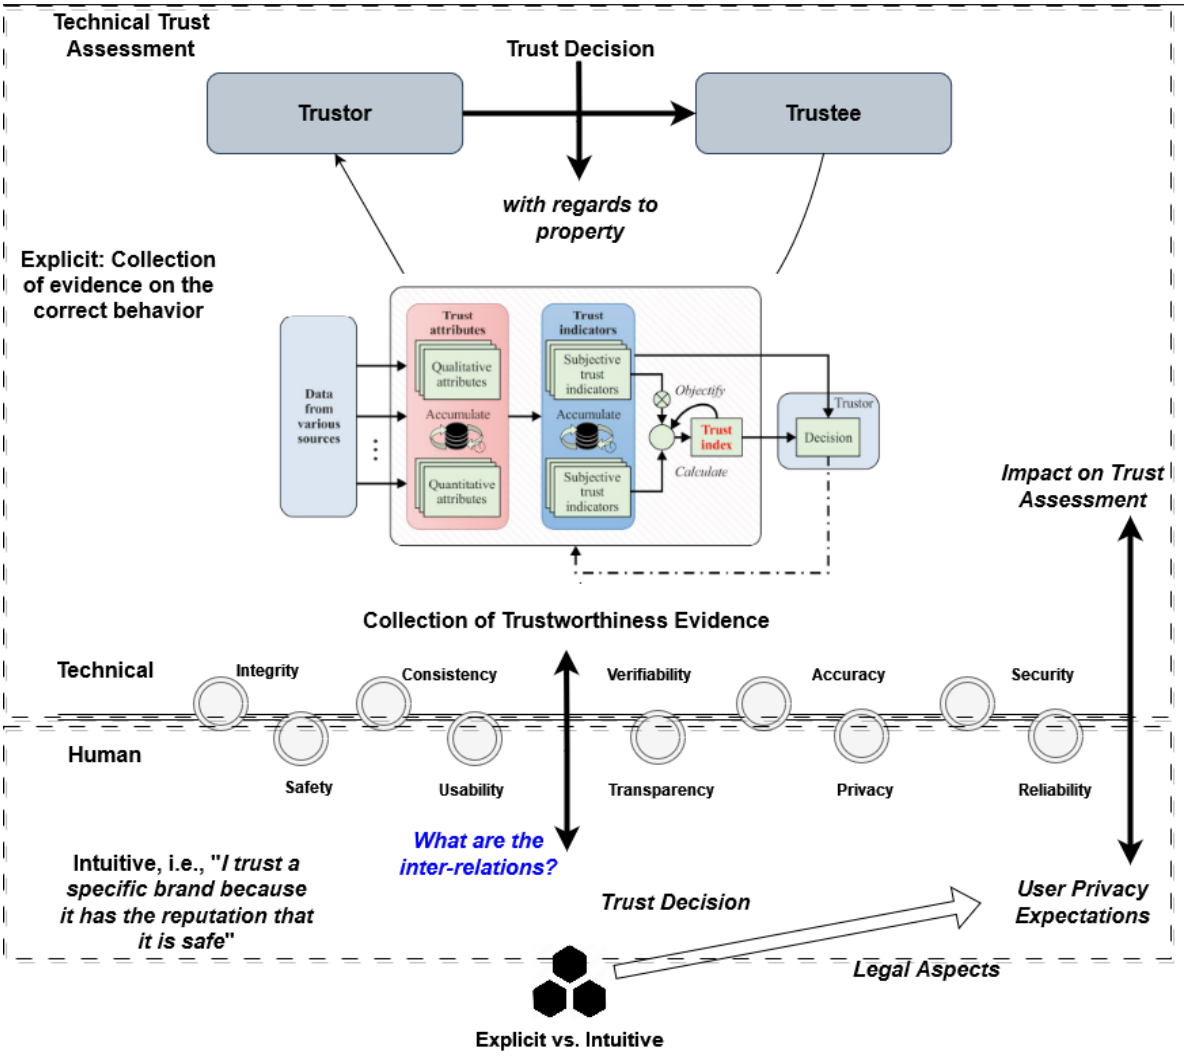

Figure 10 - Technical and Human aspects in Trust Assessment

In this light, it becomes evident that the ***policies*** governing how OEMs manage and process the collected data, should be considered. Concurrently, the ***societal*** dimension, intricately linked to user acceptance, assumes a pivotal role in shaping trust perceptions. It becomes apparent though that the **policy-making and the societal factors are intricately intertwined,** characterised by strong interdependencies that warrant thorough investigation. It is imperative to comprehend how these intertwined factors can influence users' perceived trust, and consequently, their acceptance of emerging technologies.

The complexity of this task is heightened by the current knowledge and understanding of users with respect to various regulations. The **acceptability** of a technology in contemporary society is heavily influenced by the demonstrated ***competency*** **and** ***willingness*** of the parties concerned, to use AVs in the future [81]**.** Essentially, consumers critically examine the extent to which OEMs demonstrate their willingness to comply with rules and their level of competence in doing so.

This can be further linked with the system's ability not only to adhere to relevant regulations but also to ***transparently*** articulate its compliance mechanisms by possessing certain qualifications/certifications [72]. This is a crucial step so that users can have the explainable guarantees that the system is compliant to regulations, offering transparency which has a positive effect on the perceived trust [75]. When a service is deemed “GDPR compliant”, for example, users are more likely to perceive it as trustworthy, reinforcing the interplay between policy-making, societal expectations, and the establishment of trust in CCAM systems. In addition, the need for transparency becomes evident in circumstances characterised by malfunctions, such as the event with the Tesla Model Y. The implementation of independent third-party investigations, as suggested in similar instances, has a significant role in fostering the building of trust [82].

Towards this direction, it becomes imperative to not only consider but also address the **potential conflict between the necessity of safeguarding data privacy and the overarching objective of enhancing road safety.** It is essential to acknowledge that in such situations, the unequivocal prioritisation and assurance of road traffic safety must take precedence over addressing data privacy concerns [72].

CONNECT will investigate the usage of crypto primitives for harmonising and minimising the type of evidence needed for the trust decision, maintaining a certain level of privacy. Figure 10 essentially illustrates the previously mentioned interdependencies based on a trust relationship between a Trustor and a Trustee. In this context, the trust level is determined not only by the amalgamation of collected evidence, which encompasses various trust attributes, both qualitative and quantitative, along with subjective trust indicators, but also by the human perceptual dimension that significantly influences the trust decision-making process. A comprehensive trust assessment should, therefore, strive for an equilibrium, taking into account both explicit, quantifiable metrics and more intuitive, subjective aspects. Consequently, it becomes imperative to investigate the repercussions of human-perceived trust and user privacy expectations on the overall implications of the trust assessment. This means that the inclusion of additional systematic qualities as evidence of trustworthiness has the potential to enhance the accuracy of trust decisions; however, this approach might raise concerns over the potential violation of vehicle and driver privacy profiles.

In the domain of vehicle data exchange, the sharing of information plays a crucial role, especially in trust assessment processes. This shared information can include specific parameters like the ECU type and software versions, aimed at enhancing the overall functionality of the system. However, the unintended consequences of this data exchange become apparent when viewed through the lens of privacy concerns within the automotive context.
The exchange of such data presents a double-edged sword. On one hand, it's intended to improve system performance and enhance user experiences. On the other hand, it raises significant privacy concerns. One notable risk is the potential for vehicle fingerprinting, a process through which malicious entities could deduce the exact make and model of a vehicle based on this shared information. This intrusion into the specific details of a user's vehicle poses a clear threat to user privacy, as it exposes potentially sensitive information about their vehicles.
Consequently, this privacy-trust trade-off is at the heart of a critical question that requires careful consideration and resolution. In addressing this complex issue, it's paramount to strike a balance between enhancing trust and ensuring the privacy and security of users' personal information and data. The resolution of this dilemma forms a key focus area for investigation in the context of CONNECT, as discussed in chapter 6 of the present deliverable and detailed upon in D5.1. Within CONNECT, the application of harmonisation techniques to build data models for trustworthiness evidence without compromising user privacy is a crucial aspect of addressing the privacy vs trust challenge.

3.4 Towards Trustworthiness Profiles for CCAM Ecosystems

As aforementioned, trustworthiness can be defined as **the likelihood of the trustee to fulfil trustor's expectations in a given context**. The expectations vary depending on the entrusted task, the process through which it was achieved, and the purpose for which the task was chosen. Four notes should be considered together with this definition:

- ✓ **Context dependency**: Trustworthiness depends on the context, sector, product, service, data, technology and process used. Different characteristics apply and need verification to ensure stakeholders' expectations are met.
- ✓ **Characteristics** can include but are not limited to accountability, accuracy, authenticity, availability, controllability, integrity, privacy, quality, reliability, resilience, robustness, safety, security, transparency, or usability.
- ✓ **Application to different entities of interest**: Trustworthiness is an attribute that can be applied to services, products, technology, data and information as well as to organisations.
- ✓ **Verifiability:** Trustworthiness includes measurability and demonstrability by means of objective evidence.

Practically achieving trustworthiness involves establishing a common understanding of expectations and employing an assurance approach to reach a level of confidence.

Trustworthiness in standardization is ensured by a layered approach that incorporates reference architectures (RAs) and implementation architectures, as shown in Figure 11. Architectural descriptions in RAs are governed by the ISO/IEC/IEE 42010 standards (Architecture description) [83]. These reference architectures should also conform to a guideline document established at the ISO/IEC JTC1 level [84]. Some examples of these RAs are the IoT reference architecture [85] and the digital twin reference architecture [86].

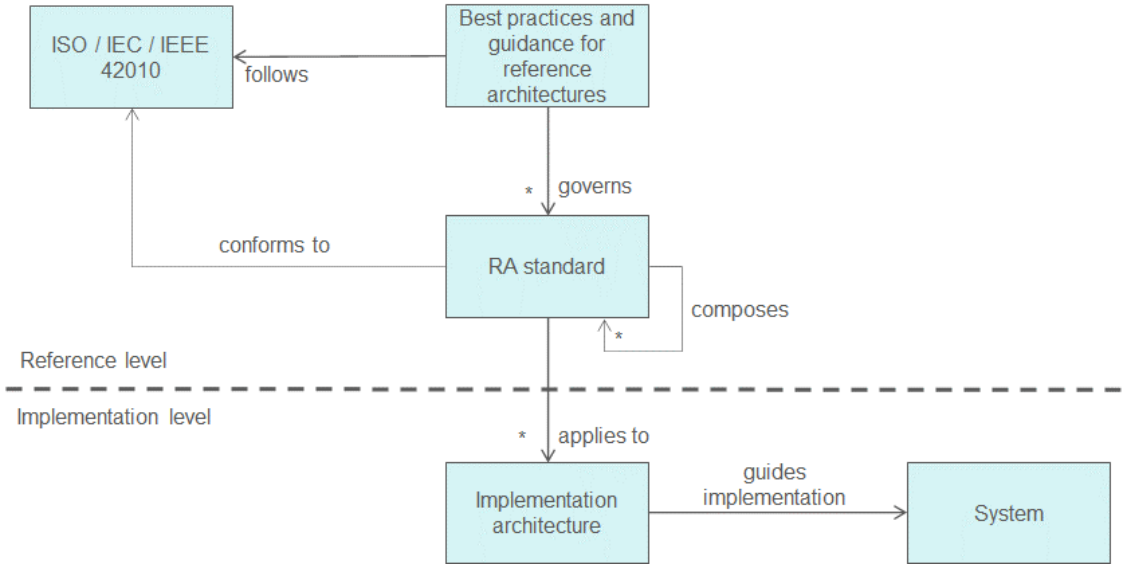

Figure 11 - Using reference architectures

The implementation level can be further structured into two tiers: i) a **CCAM reference level** where CCAM architecture profiles and CCAM trustworthiness profiles can be defined and ii) a **CCAM implementation level** which can apply a CCAM architecture profile and a CCAM trustworthiness profile. Figure 11 show how the implementation level can be structured further into two levels:

The term profile is widely used in standardisation. It consists of a single specification which groups other references. A profile is also a reference, so profiles can contain other profiles. In particular an architecture profile can include a trustworthiness profile. The difference between the two can be summarized as follows:

1. **CCAM Architecture Profiles**: These profiles pertain to the architectural description of specific CCAM services or applications, such as Intersection Movement Assist (IMA) or other services within the CCAM domain. They provide a structured representation of the architecture, components, and interactions required to implement a particular CCAM service. Essentially, CCAM architecture profiles describe how a specific service is designed and structured within the CCAM system.
2. **CCAM Trustworthiness Profiles**: In contrast, CCAM trustworthiness profiles are concerned with characterizing the trust model for a particular type of CCAM service, such as IMA. These profiles focus on defining those trustworthiness attributes and requirements specific to the trustworthiness level of a service. They provide a detailed understanding of the reliability, security, privacy, and other trust-related aspects associated with this service of interest and can serve as a guideline to Service Providers to also associate the appropriate security controls together with their deployed service. Overall, CCAM trustworthiness profiles guide how trust is established, maintained, and verified for a given CCAM service.

The objective of CONNECT is therefore to investigate the definition of such **trustworthiness profiles** mapping varying levels of trust to the type of trustworthiness evidence that need to accompany a service. This, essentially, provides a guideline (as aforementioned) on the type of security controls that need to be deployed with a service (or as part of the virtualized infrastructure – MEC – hosting a service). Whether or not we are able to define a profile that can also be advertised in the standards specifications also depends on the existence of an **approved methodology** to

create a profile. CONNECT will follow the work on this topic with the endmost goal of defining harmonized trustworthiness profiles to best capture the trust requirement of CCAM services operating in a zero-trust environment.

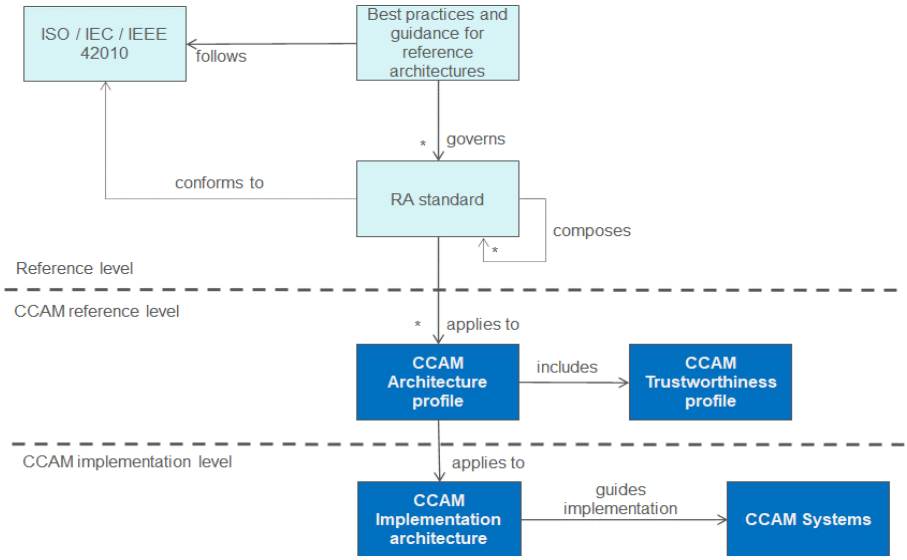

Figure 12 - Building trustworthiness profiles

In pursuit of this objective, CONNECT recognizes the significance of establishing **a common harmonization and obfuscation approach that all stakeholders can uniformly adopt, to exchange information without breaking privacy**. This common approach will facilitate **the calculation of a consistent trust score**, ensuring that trust assessment is coherent and meaningful across the entire CCAM ecosystem. Achieving this uniformity is essential to support interoperability and trustworthiness within complex CCAM systems, particularly when there are multiple entities with diverse responsibilities and characteristics involved in the data sharing and decision-making processes.

3.5 Threat Model

In the context of CONNECT, our primary focus centres on **enhancing the security and trustworthiness of CCAM** ecosystems. Our objective is to support the resilience of such environments by enabling them to effectively assess the level of trust of a communicating entity by providing verifiable evidence on their equipped security controls that enhance their security posture against a wide attack vector. This entails not only the detection of attacks but also the **analysis of their consequences on trustworthiness, empowering trust decisions**. Ultimately, these efforts contribute to the enhanced dependability and safety of CCAM systems.

However, it's essential to acknowledge that there are numerous factors that can influence dependability and safety, and these factors are not always rooted in security concerns. Mechanical defects, ECU failures, or sensor accuracy degradation are examples of events falling into this category. It is important to note that CONNECT is not addressing such scenarios explicitly, although trust assessment and our trust assessment framework might be very well suited to also model and analyse them and their consequences on CCAM systems. Even more so, some of our mechanisms to detect attacks like misbehaviour detection or attestation might even be triggered by such events and our trust assessment framework might react to them, as they are not distinguishable from some attacks like Denial-of-Service or position spoofing.

In this deliverable and throughout the project, our primary focus revolves around events that **negatively affect trust** and are associated with security-related incidents. CONNECT's threat model and consideration of attacker capabilities encompass both **outsider** and **in-vehicle** attackers. This approach aligns with the **zero-trust** principle, reflecting CONNECT's fundamental posture of not inherently trusting any entity, including those within the vehicle. This zero-trust based approach assumes an "**honest but curious**" model for all entities. Outsider attackers aim to disrupt the

system's availability, while insider attackers, who already possess authorised access, seek to disrupt normal system operation or extract information.

The primary motivation for adopting such a robust adversarial model is rooted in the evolving CCAM landscape's security challenges. Recent developments, especially **the introduction of new components** and actors like the MEC in the context of AVs, **introduce potential vulnerabilities**. Given this dynamic environment, embracing the zero-trust concept provides a safer approach. The introduction of new actors and components underscores the necessity for sophisticated trust mechanisms, such as those defined in CONNECT, capable of addressing these emerging attack vectors. **This strong adversarial model, encompassing security, privacy, and trust considerations, guides the requirements defined in Chapter 8 of this deliverable**.

The methodology employed for establishing the requirements, encompassing security, privacy, and trust considerations, is firmly grounded in the real-world security requirements articulated by key stakeholders in the context of real case CCAM scenarios. For instance, in the case of the MEC related requirements, the 5GAA [87] and ETSI standards [88] provided the basis for the CONNECT requirements. For the formulation of use-case-specific requirements, the CONNECT consortium harnessed the expertise and input of relevant stakeholders, ensuring that the needs and specifications were comprehensively addressed. Furthermore, privacy considerations are also important within the CONNECT context (although not in the forefront) as was also evident in Section 3.3 and the interplay between privacy and trust – not only how privacy enhancing technologies can affect the accuracy of trust-related decisions (as it may require the obfuscation of some of the trust-centric information/evidence exchanged) but also how this double-edge sword impacts the user acceptance of AV technologies. Besides also traditional privacy requirements (described in Section 8.2.1), CONNECT is also interested how the decomposition of a service across the entire CCAM continuum (spanning across far edge-edge-cloud) can affect the privacy profile of all actors. Especially, the consideration of MEC-enabled service deployments poses questions on the anonymity and unlinkability of users as they are moving along a road and they need to be "*handed over*" from one Mobile Network Operator to another due to distinct service provision boundaries. Such newly emerged questions are also discussed in CONNECT and form a subset of the envisioned requirements of the overall framework.

A detailed threat analysis as it pertains to specific threats and vulnerabilities which can be exploited to launch attacks against confidentiality, integrity, availability, etc., will be elaborated on D2.2 [109]. Together with a detailed mapping of the requirements, as defined in Chapter 8, against specific attacks.

Furthermore, it must be clarified that once the architecture is finalized, COVNECT will conduct a LINDDUN Data Protection Impact Assessment (DPIA) on the Trust Assessment Framework (TAF), for evaluating the trustworthiness of its proposed solutions. This assessment will provide a comprehensive and systematic evaluation of data protection considerations, ensuring that our solution aligns with the highest privacy and security standards.

4 CCAM Services Landscape, Actor Definitions and their Roles

4.1 Stakeholders and their Goals

Thus far, the CONNECT's primary objectives have been introduced, in establishing trust within CCAM services. However, it's imperative to delve into this intricate landscape by identifying the diverse actors and components actively participating in this ecosystem and elucidating the nature of the information and messages they exchange. This comprehensive analysis serves a dual purpose: firstly, it provides a deeper understanding of the expansive attack surface to reach to Day 3+ automations, as well as the myriad challenges stemming from the extensive threat landscape. Secondly, it unveils the dynamic nature of trust relationships and elucidates how these dynamics impact the overarching trust assessment framework, underlining, in parallel, the significance of fostering common trustworthiness profiles to promote uniformity across the multifaceted CCAM ecosystem. The next chapters will shed light on the aspects of the CCAM landscape, to define the relevant CCAM stakeholders as well as the type of information exchanged within this ecosystem.

In general, the stakeholders within the CCAM functional areas can be summarized in the table below (i.e., Table 2)[5]. CCAM encompasses a diverse range of stakeholders contributing to the advancement of intelligent transportation systems. At its core, **industry players can be found such as Original Equipment Manufacturers (OEMs), automotive manufacturers and suppliers as well as CCAM Service Providers that are involved in the vehicle design and advanced technology/service integration**. The **telecommunications sector and Mobile Network Operators (MNOs)** play a vital role in providing essential communication infrastructure, while cloud providers assist in the management of data. Government bodies, such as **transportation authorities and road operators**, are responsible for the regulation and management of infrastructure. **Service providers**, including several sectors such as public transportation and insurance businesses, fulfil crucial functions in facilitating mobility and ensuring safety. Non-profit groups, such as **car associations** are promoting innovation in the sector. Standardization groups uphold and regulate industrial norms, whereas academics engages in research. Collectively, the aforementioned categories summarize the landscape as it pertains to the CCAM interested parties.

Table 2 - CCAM stakeholders

Type	Actor Definition
Industry	✓ Original Equipment Manufacturer (OEM) ✓ Automobile Manufacturers ✓ Automotive supplier ✓ ITS solution providers ✓ Telecom industry ✓ Mobile network operators ✓ Cloud providers ✓ Fleet operators, including operators of automated vehicles ✓ Tele-operation centres
Government	✓ Transport authorities ✓ Road authorities ✓ Road operators ✓ Emergency responders ✓ Member states ✓ Road and transport authorities ✓ National mapping agencies ✓ City government
Service Providers	✓ Roadside assistance ✓ Public transport

[5] https://www.mobilityits.eu/ccam-connected-vehicles

	✓ Weather services ✓ Mobility and logistics providers ✓ Insurance companies ✓ Toll road operators ✓ Vehicle repair and maintenance providers ✓ Road maintenance providers ✓ Parking services ✓ Mobility as a Service (MaaS) providers ✓ Roadside assistance ✓ International corporations ✓ Navigation system providers ✓ Traffic management centres and in general CCAM Services (as the ones also envisioned in the context of the Intersection Movement Assist (Section 7.2) and Slow Moving Traffic Detection (Section 7.4) use cases)
Non-profit Organizations	✓ Automobile associations ✓ Trade associations ✓ Technology clusters ✓ Road safety associations ✓ Environmental organizations
Standardization Bodies	✓ National, European, and international
Academia	✓ Universities ✓ Public research institutes ✓ Private research institutes

As aforementioned, in the context of CONNECT, the following stakeholder are considered among the most crucial ones:

Original Equipment Manufacturers (OEMs): These are the vehicle manufacturers responsible for designing and integrating vehicle architectures. They further decide which CCAM services their vehicles will support and are accountable for the cybersecurity of their vehicles throughout the entire product lifecycle (certified to be compliant with cybersecurity regulation during type approval).

Component Manufacturers: These are manufacturers of components which will be integrated into vehicles, road-side infrastructure, or cloud. Such a component can be a software component (i.e., operating system, application, or library), a pure hardware component, or a combination of both (i.e., an ECU). Component Manufacturers are responsible for designing, integrating sub-components, and producing their components. They support OEMs in integrating these components and ensuring compliance with cybersecurity regulations. Usually, Component Manufacturers are separate, independent companies serving as a supplier to the OEMs, but in some cases, the OEM can also develop their components in-house, resulting in a combination of both roles in one company.

CCAM Service Providers: These are organizations that are responsible for offering the CCAM services, to be deployed as part of the vehicle's software stack, but also the auxiliary processes (deployed on a virtualized infrastructure such as the MEC) for supporting the better and more scalable execution of a specific CCAM service. The role of an SP can also be taken by an OEM for equipping the vehicle with automated driving functionalities such as cruise control, etc.

Mobile Network Operators (MNOs): These are the organizations responsible for the orchestration and management of the virtualized infrastructure (CONNECT MEC) where various services will be deployed. They further decide on the type of security controls and built-in security capabilities of their infrastructure and also employ orchestration techniques for the optimal deployment strategy of all services so as to not violate any requirements as described in the Service level Agreements (SLAs).

4.1.1 In-Vehicle Actors

At the core of CCAM services, vehicles play a pivotal role by serving as Intelligent Transport System Stations (ITS-Ss). In this context, a typical 'legacy' vehicle is transformed into an ITS-S with the

inclusion of various essential components, which are typically found in the in-vehicle environment. A high-level overview of the following components is further depicted in Figure 13.

1. **Vehicle Computer / On-Board Unit (OBU):** The OBU is a sophisticated hardware unit that incorporates dedicated software to facilitate the seamless communication between the vehicle and other ITS-Ss via V2X communication protocols, as specified by ETSI standards. The OBU provides, in essence, provides the connection between the Vehicle Internal Network (i.e., the CAN bus) and the components outside the CAN bus (e.g., V2X OBU). This technology also aligns with IEEE and SAE standards in the United States, ensuring uniformity.
   For CONNECT, this component is of specific importance, as described in Chapter 6.5, due to its linkage with the security related functionalities that CONNECT will require and enable through the newly developed TEE Extensions. Essentially, the Vehicle Computer will be equipped with the CONNECT Trusted Computing Base (TCB) leveraging the underlying built-in HW-based Root-of-Trust capabilities for supporting the new breed of attestation capabilities towards the provision (in a verifiable manner) of trustworthiness evidence. Apart though from the security related functionalities, the In-Vehicle computer for CONNECT supports all CCAM applications (i.e., including the IMA, the C-ACC and the SMTD discussed in the use cases in Chapter 7), the Misbehaviour Detection (MD), the proposed Trust Assessment Framework (TAF), the Attestation Integrity Verification (AIV), the Trustworthiness Claim Handler (TCH) and the Identity and Authentication Management (IAM) service.
2. **Electronic Control Unit (ECU)** is a vital component in modern vehicles and other complex systems, such as industrial machines and consumer electronics. ECU is essentially a dedicated microcontroller or computer that is responsible for monitoring, controlling, and optimizing specific subsystems within the system. In the automotive context, different types of ECUs manage critical functions like the engine, transmission, anti-lock brakes, airbags, and emissions control. These ECUs receive data from various sensors and make real-time decisions to ensure the safe and efficient operation of the vehicle. ECUs play a pivotal role in enhancing performance, safety, and reducing emissions in modern vehicles by constantly analysing and adjusting various parameters to meet performance and efficiency goals.
   a. **A-ECU (Advanced ECU)**: the ECU is powerful enough to execute all kinds of cryptographic algorithms, including asymmetric (e.g., ECC, RSA) and symmetric (e.g., AES), and hashing (e.g., SHA-2/3) algorithms.
   b. **S-ECU (Symmetric ECU)**: the ECU's capabilities are limited to only execute symmetric and hashing algorithms.
   c. **N-ECU (No-Crypto ECU)**: the ECU has no cryptographic capabilities.

3. **Supplemental ECUs**: These ECUs support the in-vehicle communication architecture.
   a. **Zonal Controller:** It provides a communication interface between ECUs (in a sub-network) and the central In-Vehicle Computers (i.e., the OBU). As will be seen later on there can be different type of such interfaces on-boarded depending on the requirements of the specific application, i.e., CAN[6], Ethernet Switch, etc. As illustrated in Figure 13, a Zonal Controller may be responsible for relaying information received from more than one ECUs.
   b. **Smart Antenna**: provide an interface between the in-vehicle network and the outside world including the MEC, cloud-based services and other vehicles. When communicating with the MEC and other vehicles, the Smart Antenna acts as mediator for forwarding the V2X messages that have been constructed by the CAM/CPM Encoder/Decoder running on an ECU.
   c. **Global Navigation Satellite System (GNSS) receiver:** This component proves invaluable for capturing kinematic data related to the vehicle's movements and location. It can be integrated within the Vehicle Internal Network or exist as an external unit. GNSS receivers come in various forms, including conventional ones, those equipped with Real Time Kinematic (RTK) systems for high-precision accuracy, and models that incorporate an Inertial Measurement Unit (IMU) featuring accelerometers and gyroscopes to support Dead Reckoning navigation. It's worth

[6] ISO 11898 Road vehicles — Controller area network (CAN)

noting that most "legacy" vehicles typically include a GNSS receiver as a standard component.

4. **In-vehicle sensors**: These sensors produce data about the vehicle's environment to support its decision making. These sensors include:
    a. **GNSS** (Global Navigation Satellite System): produces data on the vehicle's position. This sensor is implemented in pure hardware as an ASIC (Application-Specific Integrated Circuit) system. It is connected to the Camera ECU via a LIN[7] bus.
    b. **Camera**: This sensor captures the visible vehicle surroundings in a video. It also forwards the vehicle position data from the GNSS sensor.
    c. **Lidar**: a laser-based imaging system to detect objects on the road.
    d. **Radar**: a radio-based imaging system to detect objects on the road.
5. **Vehicle Internal Network:** It manages the data generated by various hardware elements within the vehicle. Depending on the architecture, there are different sub networks, the most common of which today are the CAN bus and Ethernet. Within these networks, Electronic Control Units (ECUs) play a pivotal role by implementing one or more vehicle functions, such as engine management, wheel control, transmission, and the operation of sensors like cameras, radar, LIDAR, and proximity sensors. All these ECUs interact and share data with one another through this network, ensuring coordinated functioning and responsiveness.

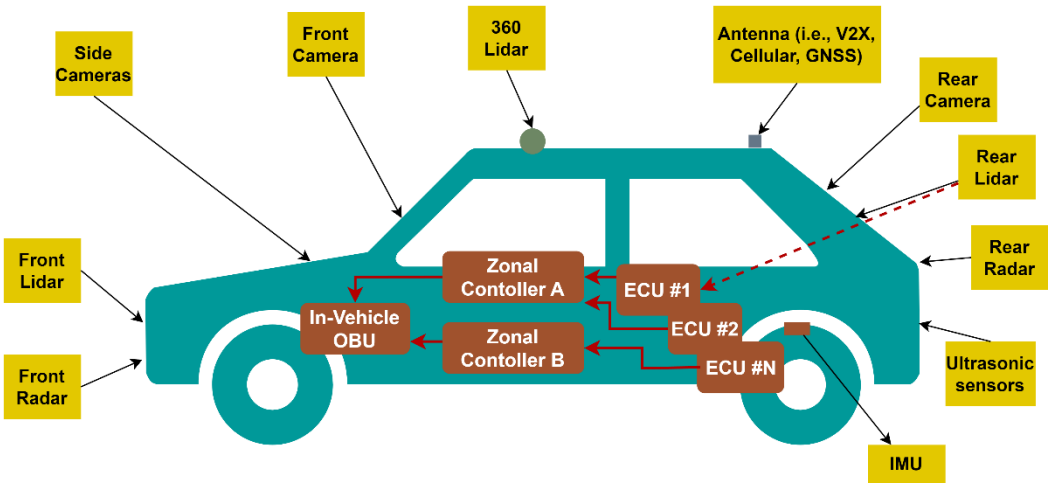

Figure 13 - In-Vehicle CCAM actors

4.1.2 External to the Vehicle Actors

In this context, a typical 'legacy' vehicle is transformed into an ITS-S with the inclusion of various essential components, which are typically found in the outside the vehicle environment (i.e., meaning other entities of the CCAM ecosystem) or processes that may take place at a central infrastructure such as the MEC or the Cloud.

1. **Road Side Units (RSUs):** These are the basic infrastructure elements for ITS-Ss, consisting of hardware boards usually fixed on a dedicated pole and elevated from the ground, with dedicated software running on. RSUs are equipped with dedicated software to communicate using various vehicular communication technologies and may include advanced sensors like cameras, radar, and LIDAR.
2. **Vulnerable Road Users (VRUs):** VRUs encompass all ITS-Ss whose involvement in car accidents may lead to severe damage to the user. This category includes pedestrians (communicating via smartphones), bicyclists, motorcyclists, and other non-vehicular road users.
3. **Multi-access Edge Computing (MEC) Server:** A key element in the network architecture introduced by ETSI, the MEC Server enables cloud computing capabilities and services at the network's edge. More details on the MEC are discussed in section 2.2.4 of the present

[7] ISO 17987 Road vehicles — Local Interconnect Network (LIN)

deliverable. These servers are located close to end users (i.e., at the EnodeB, that is the cellular antenna), enhancing their role in managing V2X message flows.

   a. **Local Misbehaviour Detection Service (Local MD Service):** This service may operate either within the vehicle or at the MEC level. Its primary responsibility is to inspect incoming V2X messages, such as CAM or CPM, to identify semantic inconsistencies that could indicate data manipulation attacks. The outputs generated by this service may be used directly by the vehicle's on-board system to identify untrustworthy data. Additionally, the service may include the generation of Misbehaviour Reports sent to the Misbehaviour Authority for further action. [89].

4. **Cloud Server:** These servers, often deployed in a hierarchical manner, contribute to data management and services on a broader urban or regional scale, complementing the capabilities of MEC Servers, which are found closer to the user side (i.e., MEC server for every block and two or three of them for every district).

   a. **Traffic Control Centre:** Preceding up in a hierarchical architecture, the Traffic Control Centre is not a simple message flow manager. It further analyses the high-level traffic situation to emit some alerts or infotainment messages to be dispatched towards ITS-Ss of specific regions.
   b. **Vehicular PKI (VPKI):** The VPKI plays a crucial role in ensuring secure communication within the CCAM ecosystem by managing public key infrastructure for vehicles, as introduced in Section 2.2.1 and further elaborated in Section 4.2.1.
   c. **Certification Body:** The Certification Body plays a pivotal role in granting approval to vehicle manufacturers, component suppliers, and service providers, attesting that their offerings comply with the applicable security and safety regulations. This certification process includes rigorous testing, validation, and assessment, with a focus on ensuring the integrity, confidentiality, and availability of CCAM services. In addition to certification, the CB is also responsible for ongoing monitoring and auditing, ensuring that certified components and services continue to meet the established trustworthiness requirements throughout their operational life cycle. It collaborates with other CCAM stakeholders to develop and maintain industry-specific standards and best practices, contributing to the continuous improvement of trustworthiness in the CCAM ecosystem. In the context of CONNECT, such Certification Bodies might be interested in auditing the safety or integrity of deployed vehicles by getting access to the history of trust information that is stored and securely shared through the CONNECT Blockchain infrastructure (Section 6.7)

5. **Vulnerability and Threat Knowledge:** This component identifies, monitors, and manages vulnerabilities and threats that may pose risks to the system's integrity and security. It helps to maintain a comprehensive database of known vulnerabilities and threats specific to the CCAM environment, keeping abreast of emerging risks and attack vectors, providing real-time threat intelligence and risk assessments (part of the CONNECT Risk Assessment engine – Section 6.2). Through its continuous evaluation, it improves the overall security posture of CCAM systems.

4.2 CCAM Services and Communications

The continuous growth of the discussions revolving around CCAM services, has triggered the research on the field of communications aiming to provide all the necessary enablers to cover the demanding needs as it pertains to network coverage, capacity, latency, and security, among others. As mentioned in Chapter 2 of the present deliverable, a wide range of services that are anticipated with CCAM, including Cooperative Manoeuvring, Situation Awareness, Autonomous Driving, Platooning, etc. To this end, the European Telecommunications Standards Institute (ETSI) has defined a reference architecture designed to support a multitude of use cases spanning road safety, traffic efficiency, infotainment, and business applications [90].

Figure 14 presents the layers defined in the ETSI Intelligent Transport Systems (ITS) architecture of communications [91][92]. In terms of this communications architecture, our focus in this chapter is

basically placed on the **Facilities** layer, as illustrated. This layer consists of the messages (i.e., CAM, DENM) that are, in essence, **enabling the future collaborative aspect** of the ITS Applications (i.e., Autonomous Driving, Traffic Management, etc.), paving the way to day 3+ automations. It is worth noting that the access technologies may vary, according to the ETSI definition (i.e., V2X, GNSS or IEEE 802.11p).

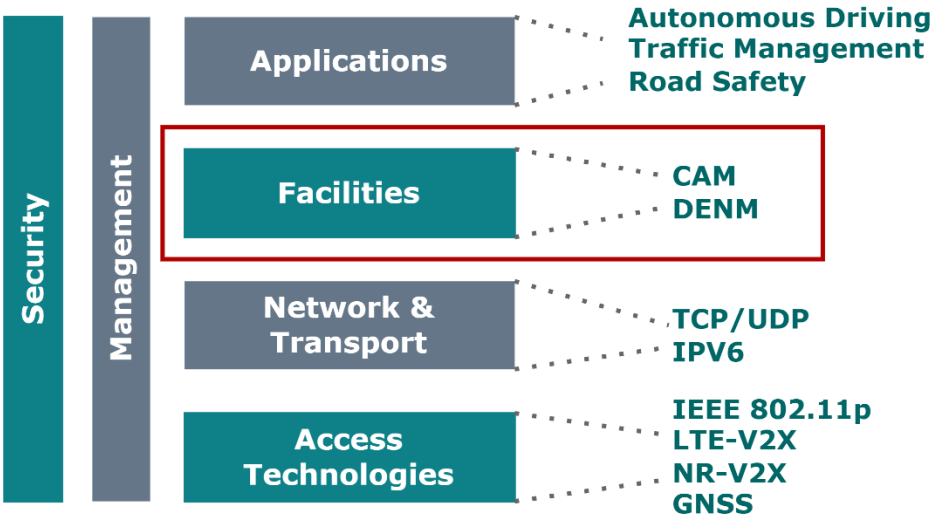

Figure 14 - ETSI ITS architecture

In terms of access technologies, the growth of C-V2X has progressed in three distinct stages. Stage 1, which began in 2015, involved the use of LTE technology for basic V2X applications such as CAM, DENM, and BSM. Stage 2, introduced in Rel-15, focused on supporting more advanced V2X scenarios such as remote driving, vehicle platooning, extended sensors, and advanced driving. Stage 3, implemented in 2016, aimed to encompass road safety solutions and connected cars. The necessary services for common V2X applications encompass a message transfer frequency ranging from 10 to 50 Hz, a broad communication range, support for V2X communication both within and outside of network coverage, and a transmission delay of less than 100 ms [93].

In terms of the CONNECT project, more specifically, the V2V and V2N radio interfaces will be employed.

- **V2V radio interface:** The V2V radio interface is used by the vehicles to broadcast messages which are received by neighbouring stations within radio range. This is also called short-range communication or direct-communication radio interface. The V2V radio interface may be based on the ITS-G5 radio access technology or on the cellular sidelink. In the IMA use case, the V2V radio interface is used for the exchange of CAM and CPM messages between vehicles.
- **V2N radio interface:** The V2N radio interface allows the vehicle to connect with the MEC. It exploits the mobile network radio access (e.g., 4G or 5G) to reach the mobile base station, to which the MEC is connected. The V2N radio interface allows connection both in uplink and downlink. The communication in the uplink is unicast, while the downlink may be both unicast and broadcast. In the IMA use case, the V2N uplink is used to upload MRs from the vehicle to the MEC; to upload vehicle-TCs from the vehicle to the MEC; to upload T-CAMs and T-CPMs from the vehicle to the MEC. The V2N downlink is used to disseminate NTMs in broadcast; to disseminate geo-CPM in broadcast, as defined in the following sections.

4.2.1 V2X Security Services

V2X systems incorporate some basic security mechanisms to safeguard communication and ensure trustworthiness within the n Cooperative Intelligent Transport Systems (C–ITS) ecosystem. Two vital components that support these security measures are the Local Misbehaviour Detection Service and

the Vehicular Public Key Infrastructure (PKI). In the subsequent subsections, we delve into these integral elements, shedding light on their functions and significance. The Local Misbehaviour Detection Service plays a pivotal role in scrutinizing incoming messages for integrity and authenticity, aiming to detect any potential misbehaving entities. These checks are then analysed by a local detection mechanism to estimate the overall plausibility of a message. Concurrently, the Vehicular PKI provides a robust framework for secure communication among ITS–Ss entities, establishing authenticity, integrity, non-repudiation and optionally privacy within the V2X environment. The Local Misbehaviour Detection Service is complementary to the PKI as the digital signature alone cannot ensure the accuracy and validity of a message. For instance, a malicious vehicle with a valid certificate could send inaccurate or false data on the C–ITS network. This is where the necessity for a Misbehaviour Detection (MD) system becomes apparent, to protect and mitigate the effects of these malicious or faulty ITS–Ss [94]. The combination of these two components contributes significantly to enhancing the overall security and resilience of the ecosystem.

4.2.1.1 Local Misbehaviour Detection Service

It is assumed that the ego vehicle is equipped with a **Local Misbehaviour Detection (MD) service**, which is a crucial component responsible for processing incoming CAMs in order to determine the presence of misbehaviour events. This process entails subjecting the contents of the CAMs (i.e., the observation) to a series of predetermined checks, as exemplified in the following examples. The results of these checks constitute the activation pattern. Depending on the activation pattern, the Local Dynamic Map (LDM) update function records the observation in the LDM, in accordance with predetermined policies. A commonly observed policy entails that data which does not exhibit any activated misbehaviour detectors, is deemed **authentic**, and subsequently incorporated into the LDM. Conversely, data linked to any instances of activation is disregarded.

The set of misbehaviour checks relevant to the problem of making sure that correct data are added to the LDM are defined in the literature as data-based [139] and rely on the inspection for inconsistency of the data semantics in the messages. Examples of **data-based misbehaviour checks** on CAMs are given on Table 3 [95]. They can be classified in **plausibility checks**, which are performed on a single message and aim to reveal non-plausible data; and **consistency checks**, which are performed comparing the contents of successive messages to reveal incompatible data.

ETSI's view of the Misbehaviour Detection Management system is described in [96]. This is composed by the Local Misbehaviour Detection system, which runs on the vehicle. The aim of the Local Misbehaviour Detection system is to detect and log misbehaviour events by running checks on the received data. The Local Misbehaviour Detection system includes the Misbehaviour Reporting Service [89], which allows the node to report the observation of misbehaviour events to the Misbehaviour Authority (MA).

A **Misbehaviour Report** (MR) contains the message that activated the observed misbehaviour detectors, along with evidence for the MA to independently verify the misbehaviour event (see Section 7.2.2). The role of the MA is to pool MRs and use them to identify nodes for which misbehaviour events have been reported repeatedly over time. These nodes, whose identity at the MA is provided by the pseudonym certificates from the Vehicular PKI, are considered suspicious. Using the available information, the MA tries to determine whether the suspicious node may be executing an attack. The MA, which does not have the ability to link the pseudonym certificate of a node to its permanent credentials in the system, may then decide to interact with the PKI for, e.g., preventing the provision of further pseudonym credentials of the node. Notice that this process does not happen in real-time and may conclude several days after the initial Misbehaviour Report.

- **MR message:** The **Misbehaviour Reporting Service (MRS**) is standardised by ETSI in [89]. The vehicle supporting the MRS encodes **Misbehaviour Reports (MR)**, whose objective is to make the receiver aware that the Local MD Service on the vehicle has observed the activation of a set of MD checks. A single MR may report the activation of one or several detectors, all triggered by the reception of the same V2X message (e.g., a CAM or a CPM). The MR contains the triggering V2X message, the identity of the activated detectors, and enough related information to allow the receiver to verify the activation. For instance, if the MR reports the activation of a consistency detector upon reception of the most recent CAM

from a station, the MR shall include the most recent CAM along with the previous one, so that the receiver may verify the activation of the consistency detector. According to [89] the vehicle supporting the MRS is not in the obligation to issue an MR for each observation of misbehaviour detector activation by its Local MD Service. In CONNECT, MRs are extended to be able to accommodate V-TCs in the message.

Local Dynamic Map (LDM): The LDM stores all observations whose reference time is more recent than the last n seconds. In this scenario, an **observation** may be extracted from a CAM, from a CPM or from the on-board perception of the ego. A CAM always contains a single observation (the self-description of the node); the Originating Station container of a CPM contains a single observation (the self-description of the node); the Perceived Object container of a CPM may contain multiple observations (the descriptions of the perceived objects).

An observation is always logically linked to the CAM or CPM message it derives from, so that it is always possible to identify its source (thanks to the PKI certificate). An observation must at least contain:

- the **object identifier**, expressed in the name domain of the transmitting V2X node. If the observation is extracted from a CAM or from the Originating Station container of a CPM, the object identifier is the PKI certificate of the emitting V2X node, which is the identifier of the object also in the absolute name domain; if the observation is extracted from the Perceived Object container of a CPM, the object identifier is the same used in the CPM.
- the **position** of the object and the reference time.

The observation may contain optional attributes as the *velocity*, *acceleration*, *type*, *size* of the object; it may contain the vehicular PKI certificate of the object, if the observation comes from a self-description of the node (e.g., if the observation comes from a CAM).

Each observation in the LDM refers to a physical object. The Object Association Algorithm, which runs periodically, has the objective of clustering observations of the same physical object. The clusterization is performed in part based on the object identifiers (e.g., in the case of consecutive CPMs from the same node, where the same physical object is always attributed the same identifier); in part based on the kinematic information contained in the observation (e.g., two observations of the same object contained in CPMs from different nodes are recognized as referring to the same physical object for their similarity in position).

Kinematic data fusion (Vehicle): The kinematic data fusion is the module responsible for building the consolidated view, from the data contained in the LDM. At any moment, the LDM contains several observations referring to the same physical object, with the respective TL. Each observation constitutes a measurement of the kinematic state of a physical object. The Kinematic data fusion module processes all entries in the LDM by grouping those associated to the same physical object (classification task); and then, for each group, processing the observations (measurements) to produce the *estimate* of the current kinematic state of the physical object (estimation task). Since the measurements are all associated with a TL, the estimation is able to associate a TL to the resulting estimate. At any moment, the collection of the estimates of the kinematic states of all physical objects appearing in at least one observation in the LDM constitute the consolidated view, which is consumed by the IMA application.

Local Misbehaviour Detection Service (Vehicle): The Local Misbehaviour Detection service processes each incoming CAM, CPM and the local perception, extracting their observations. Each observation is then subjected to the battery of misbehaviour checks. The presence of data coming from sensor data and the possibility of comparing multiple observations of the same physical object makes the panel of possible misbehaviour checks way richer than what is employed when CAM only are considered. The misbehaviour checks may be classified as described in Table 3.

Table 3 - Misbehaviour checks classification

Plausibility checks. They are applied to the observation, individually. They aim to check that the observation agrees with the basic physical rules that describe the world.	1. **Position plausibility**: checks that the position in the observation is compatible with the map of the scene (e.g., a vehicle may not be over a building). 2. **Speed plausibility**: check if the speed in the observation is inferior to a predefined threshold (maximum physical speed). 3. **Range plausibility**: if the observation is a self-description of a source, check if the position in the observation is compatible with the radio range of the ego (e.g., a vehicle declaring to occupy a position out of the estimated radio range of the ego)
Consistency checks. They aim to check that the current observation is in agreement with previous observations of the same object, by the same source. For example, this may apply to the self-description of a source obtained from a CPM and from the immediately preceding CAM; or to the description of the same object in two successive CPMs from the same source.	1. **Position consistency**: this applies only to observations which have an identifier in the absolute namespace, i.e., to observations corresponding to self-description of sources (they are identified by their associated PKI certificate). It checks if the position in the most recent observation is compatible with the Kalman prediction of the position of the node, based on the previous observation [97]. 2. **Speed consistency**: check if the speed in the most recent observation is compatible with the speed and acceleration in the previous observation. 3. **Size/type consistency**: check that the size/type of the object in the most recent observation matches the values of the previous observation.
Redundancy checks. They aim to check that observations that the Object Association Algorithm associates to the same physical object are in agreement. Observations are considered in a pairwise fashion.	1. **Speed redundancy**: check if the speed in the tested observation is compatible with the speed in the other observation. 2. **Acceleration redundancy**: check if the acceleration in the tested observation is compatible with the acceleration in the other observation. 3. **Size/type redundancy**: check if the size/type of the tested observation is compatible with the other observation.
Source spatial checks. This is the class of most complex checks. They aim to verify that the declared position of a V2X node is in agreement with the positions of the objects it declares to perceive.	4. **Source perception region**: checks if the observations extracted from the Perceived Object container of a CPM are within the perception region declared in the Sensor container. 5. **Pairwise perception regions**: for each pair of known CPM sources, evaluates the intersection of the relative perception regions, and compares the declared lists of objects. This allows it to activate detectors for ghost objects (when a V2X node produces an observation of a non-existent object) or for omitted objects (when a V2X node does not produce an observation for an object which would be in its perception region). These checks are especially complex because they require the evaluation of the occlusion zones in the perception regions caused by the presence of known objects.

4.2.1.2 Vehicular PKI

Nodes wishing to legitimately participate in V2X communications in C-ITS need to be enrolled by the Vehicular PKI [98], which is the trusted provider of credential and cryptographic material used for authenticating the vehicle in the system. At the time of commissioning, the node receives a permanent certificate from the Vehicular PKI through secure enrolment; this certificate is proof of the legitimacy of the node in the system and of its long-term identity [96].

V2X messages such as CAM and CPM are broadcasted on the control channel in plaintext, in order to guarantee that they may be timely consumed by the receivers. The authentication of the sender and the integrity of the message is guaranteed through the cryptographic signature [99]. However, in order to protect the privacy of the sender and prevent tracking, the node does not sign V2X

messages using its permanent certificate. The Vehicular PKI provides, instead, legitimate nodes with ephemeral, pseudonym certificates [13][121]. Hence, ***legitimate nodes broadcast V2X messages in plaintext and sign them using a pseudonym certificate***. Nodes use the same pseudonym certificate to sign all their V2X messages during a random period of time, and then change pseudonym. This strategy, mandatory in Europe, efficiently prevents a receiver to trace back to the same source all the messages that have been emitted by the same vehicle during an extended operating period, thus protecting its privacy.

This solution will be used within CONNECT to ensure privacy-preserving broadcast of V2X messages, also for the exchange of V2X messages introduced by CONNECT to support the trust assessment functions. Moreover, CONNECT will make use of the Vehicular PKI pseudonymous certificates as a means to identify vehicles in information exchanged as a part of the trust assessment process, as not to increase the surface for privacy breaches.

4.2.2 CONNECT Data Lifecycle

Within CCAM ecosystems, data play a crucial role towards facilitating communication, decision-making and coordination among different stakeholders, such as vehicles, infrastructure, and control units. The data lifecycle in CCAM ecosystems is dynamic and continuous, supporting the operation of connected and autonomous vehicles, traffic management systems and other components of smart mobility. This data lifecycle, targeting at the transformation of road safety and mitigation of the occurrence of accidents, may involve various stages, as follows:

1. **Data Collection**: Sensors and devices in vehicles, roadside infrastructure and nearby meteorological stations collect data continuously. Such data may pertain to information like vehicle speed, location, weather conditions, road conditions and traffic status.
2. **Data Transmission**: Upon collection of data, their transmission to the relevant stakeholders is necessary. This involves wireless communication technologies, such as Dedicated Short-Range Communication (DSRC), Cellular Vehicle-to-Everything (C-V2X) or other IoT communication options.
3. **Data Processing**: The data that are collected from various sources are processed in real-time towards becoming available in a format that will facilitate the extraction of valuable information and insights. This may involve mapping to standard-based data models, data cleaning for mitigation of any errors or inconsistencies, data anonymization for safeguarding any personal or sensitive information or data fusion for combination of data coming from multiple sensors and devices.
4. **Data Quality Management**: In liaison with data processing activities, maintaining data quality ensures that the information used for decision-making is accurate and trustful. To this end, data validation and curation mechanisms, such as leveraging misbehaviour detection for checking the semantics of the data, may be employed to facilitate quality of collected data.
5. **Data Storage**: After their processing, data can be stored in databases or cloud-based systems in a secure manner towards being available for various CCAM purposes. In some cases, there may be a need to store also data inside vehicles for a specific period of time so as to have them available in case of extreme events, such as accidents. For example, in the SMTD use case of CONNECT, the CAM/CPM messages are stored in the vehicle for 60 seconds to facilitate insurance services.
6. **Data Analysis**: Data analytics are crucial for gaining insights from the collected and processed data, through the design and execution of appropriate analysis mechanisms, such as Subjective Logic or AI/ML modelling. Such analytics can provide information regarding identification of traffic patterns, congestion prediction, assessment of road conditions, misbehaviour detection and trustworthiness level assessment.
7. **Decision-Making**: The insights derived from data analysis are exploited for decision-making activities. Control units, autonomous vehicles and other stakeholders use such information to make real-time decisions, such as adjusting traffic signals, rerouting vehicles or sending alerts to drivers, among others.
8. **Data Sharing**: Data sharing is important for cooperative and connected mobility, as data are shared among vehicles, infrastructure and control units to enhance road safety and traffic

efficiency. However, sensitive and personal data should be always treated in a secure manner by applying the necessary privacy and security measures.

9. **Data Privacy and Security**: As much of the data involved in a CCAM ecosystem are sensitive, privacy and security mechanisms for the protection of such data are essential. Among these mechanisms can be encryption, anonymization and access control policies that prohibit any unauthorised access and protect sensitive and personal information.
10. **Data Retention and Archiving**: In accordance with data storage, data retention and archiving is essential for maintaining data for extended periods for activities, such as incident investigation or legal compliance. In this direction, the blockchain technology may be exploited for keeping information safely, such as attestation results that can be used in the future for trust assessment purposes.

As envisioned through the above phases of the data lifecycle in CCAM ecosystems, there is an imperative need for enabling access to trustworthy data derived from trustful sources. Eventually, **decentralised/distributed trust is inevitably necessary**, while until now only centralised trust had been employed through the definition of PKIs. Following, the types of such necessary trustworthy data that are exchanged as part of the applications running in the envisioned CCAM services, are presented.

4.2.3 Data Models & Design Space

To comprehend the data landscape of CCAM services, it is necessary to examine the many forms of data that are exchanged within the ecosystem. This encompasses a range of data, such as V2X communications, kinematic data, ITS-S events, misbehaviour reports, and application-specific data. Examining these data types offers valuable understanding of the dynamic interactions that form the basis of CCAM services and establishes a framework for evaluating security, trust, and privacy. This explanation will offer a thorough comprehension of the massive data exchanges that are crucial for CCAM services, establishing the foundation for a full grasp of their data-driven operations. Table 4 offers a detailed analysis of the messages.

Table 4 –Data Models

Data & Message Type	Definition/Description
Application Data	**Kinematic data** of the ego-vehicle **and estimated kinematic data** of other sensed vehicles**.**
ITS-S Operational Data/Messages	**Cooperative Awareness Messages (CAMs)**: The CAM is standardised by ETSI in TS 103 900 [137]. It is a beaconing message transmitted by the connected vehicle to inform the receivers of its position at the indicated reference time and other kinematic attributes, such as the velocity, the acceleration, the size and type. The kinematic attributes are associated with an indication of their uncertainty (e.g., a confidence region). The generation algorithm for CAMs, detailed in [137], establishes that a vehicle needs to transmit CAMs respecting a minimum frequency constraint, to ensure that neighbouring stations may receive timely information. The frequency constraint depends on the context, e.g., on the speed of the vehicle. Vehicles supporting the Cooperative Awareness Service (CAS) [137] are able to emit and receive CAMs.
	Collective Perception Messages (CPMs): The CPM is standardised by ETSI in TS 103 324 [100]. It is a message transmitted by the connected vehicle to inform the receivers of the kinematic state of the dynamic objects that the connected vehicle perceives thanks to its on-board system. If the same physical object is described in successive CPMs from the same vehicle, it is labelled with the same identifier in the perceived object lists. The CPM is composed of the Originating Station container, which specifies the reference time and the position of the originating vehicle; the Sensor Information and the Perception Region containers, which provide information about the perception region of the vehicle, i.e. define the region where the on-board sensors are able to perceive the presence of dynamic objects; the Perceived Objects

	container, which specifies the list of perceived objects and their kinematic attributes (position, orientation, velocity, acceleration), along with their relative uncertainties. If the same physical object is described in successive CPMs from the same vehicle, it is labelled with the same identifier in the perceived object lists. The generation algorithm for CPMs is detailed in [100] and establishes that a vehicle needs to transmit CPMs respecting a minimum frequency constraint. The frequency of generation of CPMs is considerably lower than the frequency of generation of CAMs. Vehicles supporting the Collective Perception Service (CPS) [100] are able to emit and receive CPMs.
	Decentralised Environmental Notification Messages (DENMs): The DENM is standardised in ETSI EN 302 637-3 V1.3.1 [101]. It is mainly used in vehicular applications in order to alert road users of a detected hazardous event. Traffic incidents, roadwork, weather conditions, and other data that may impact road safety and traffic management are a few of the alerts that might be covered by DENM messages. These messages are typically sent by vehicles or roadside infrastructure units to notify other vehicles and traffic management systems about particular circumstances on the road.
	Misbehaviour Reports (MR): Report providing evidence about a misbehaviour event. The Misbehaviour Reporting Service (MRS) is a standardised system that encodes Misbehaviour Reports (MR) to notify receivers of the activation of a set of MD checks. These reports contain the triggering V2X message, activated detectors' identities, and related information for verification. The vehicle supporting the MRS is not required to issue an MR for each observation of misbehaviour detector activation by its Local MD Service. More info on the MR message is available on Section 4.2.1.1.
	Local Perception: The Local Perception (LP) comprises the list of dynamic objects perceived by the vehicle's on-board system, as encoded for inclusion in outgoing CPMs. It can hence be assimilated to a self-issued CPM. Whenever the vehicle supports the CPS, the LP is available to the Local MD Service and is included in the LDM.
	Observation: An observation refers to the kinematic description of a single dynamic object. The observation must include the identifier of the object, the reference time, and the position. It may also include other kinematic attributes such as the velocity, the acceleration, the type, the size of the object. The indication of the relative uncertainty (e.g., confidence region) is associated with each attribute. A CAM contains a single observation, consisting of the self-description of the emitting station; the identifier of the object is in this case the PKI pseudonym certificate used for signing the CAM message. A CPM contains several observations. The first observation comes from the Originating Station container, it corresponds to the self-description of the emitting station and the associated object identifier is, as for the CAM, the pseudonym PKI certificate; the other observations come from the elements in the Perceived Objects container, one per perceived object. The object identifiers are the ones used in the Perceived Object container of the CPM. The observations available in the LP are defined as those in the CPM.
CONNECT defined messages	**Vehicle Trustworthiness Claims (V-TCs)** V-TCs are produced by the on-board Trustworthiness Claims Handler (TCH) based on the output of the Attestation Integrity Verification component which is responsible for the runtime collection and verification of trustworthiness evidence (Section 6.4.3) from the in-vehicle sensors to be forwarded to the TAF for appraising the trustworthiness level of the entire service graph chain (Section 6.3). For privacy protection, when the aim is to report such evidence outside the vehicle, the TCH generates harmonised attributes (which is equivalent to adding a level of abstraction to the monitored evidence "hiding" detailed information that might lead to the identification/fingerprinting of the vehicle – Section 6.1) that ensure that V-TCs do not contain any granular information that may be used to identify/fingerprint a vehicle. V-TCs provide evidence about attributes pertinent to the ability of the vehicle to construct and

	transmit correct V2X messages, e.g., system integrity, communication integrity, and system safety (only certified applications are running in the system and are generating/processing the kinematic data). For attestation, V-TCs are signed using the internal Attestation Key of the underlying trust anchor; this is then wrapped under the signature based on the PKI pseudonym certificate. Notice that this allows linking the V2X messages emitted by a node with the V-TCs. V-TCs may be transmitted as standalone messages or may be included in another message (e.g., in MRs).	
	NTM messages: The **V2X Node Trustworthiness Message (NTM)**, defined in CONNECT, is a message that contains a list of pseudonym PKI certificates, each corresponding to an active V2X node, along with their most recent trustworthiness levels, as attributed by the **V2X Node Trustworthiness Assessment Service (NTS)** running at the MEC. An NTM also includes MEC-TCs.	
	MEC Trustworthiness Claims (MEC-TCs) Trustworthiness Claims, defined in CONNECT, provide evidence that a specific entity has (or exhibits **ownership** and **correctness** of) a set of attributes. Intuitively, the MEC-TCs provide the necessary guarantees that the virtualized environment where the service is running has integrity and has not been compromised/altered, thus, resulting in the service having the same configuration as when deployed. Thus MEC-TCs provide evidence on the integrity of the computing support used to provide the MEC service.	
	T-CAM: In CONNECT T-CAMs are regular CAM messages, extended such that they can accommodate V-TCs for the purposes of CONNECT. V-TCs are included in a periodic fashion in T-CAMs produced by a given station, so that not all T-CAMs by that station are expected to include V-TCs.	
	T-CPM: In CONNECT T-CPMs are regular CPM messages, extended such that they can accommodate V-TCs. V-TCs are included in a periodic fashion in T-CPMs produced by a given station, so that not all T-CPMs by that station are expected to include V-TCs.	
	geo-CPM: This message is defined for the purposes of CONNECT and is disseminated by the geo-CPS service at the MEC. As a regular CPM, it contains the description of objects on the road. More specifically, it contains a Detected Objects container, corresponding to the Perceived Objects container of the regular CPM, which specifies a list of objects and their kinematic attributes (position, velocity, acceleration, type, etc.) along with their relative uncertainties and the trustworthiness level attributed by the MEC to the observation. The geo-CPM may contain MEC-TCs.	
	MR message The MR contains the triggering V2X message, the identity of the activated detectors, and enough related information to allow the receiver to verify the activation. In CONNECT, MRs are extended to be able to accommodate V-TCs in the message	
Trust Identification Data	**Models:**	**Trust Model:** Graph-based model which represents all components/data to perform a certain function. **Trust Model Template:** Template based on which a Trust Model is instantiated.
	Requests:	**Setup Request:** Request to setup items for the TAF **Trustworthiness Assessment Request (TAR):** Request to conduct a trustworthiness assessment for a specific data item
	Evidence:	**Trustworthiness Claims (TC):** Evidence that the hardware/software of an external node is in a correct configuration state
	Output of the TAF:	**Trust Decision**: Decision if a trust object is trustworthy or not. **Actual Trustworthiness Level (ATL):** Extent to which a trust object can be considered trustworthy
	Others:	**Trust Policies:** Contain several configuration parameters for the TAF, such as the RTL.

		Trust Value List (TVL): List of trustworthiness levels of active V2X nodes
Trustworthiness Evidence	**Attestation Evidence**: This is an instance of Trustworthiness Evidence that serve as a trust source to be leveraged by the CONNECT trust assessment framework when conducting a fresh trust evaluation of a specific node or data item. It focuses on specific system attributes that can be monitored during runtime and can serve as evidence on its correct configurational and executional state. Attestation evidence are the result of the newly developed CONNECT attestation capabilities (Section 6.4.3) that are capable of providing such data in a verifiable manner.	
	Harmonised Attributes: Attributes of the vehicle which have been obfuscated so that they can provide the necessary evidence on the vehicle's trustworthiness level but in a privacy-preserving manner, that ensures that a Trustworthiness Claim vector will not leak any information on the assets of the vehicle. Note that harmonisation of the evidence entails the abstraction or removal of any personally identifiable information and is meant to ensure that a TC vector reported outside of the vehicle will not lead to vehicle fingerprinting or implementation disclosure attacks.	
	Trustworthiness Claims: A Trustworthiness Claim (TC) is a data structure created by the TC Handler (TCH), for holding the harmonised version of the trustworthiness evidence (originating from an entity of CONNECT and produced by a security control). The TC is part of the presentation that is transmitted from the vehicle to external parties and is described in detail in D5.1 [107].	

4.2.4 Data Structures in the context of the Use Cases

The previous section delved into the data types considered by CONNECT, as a framework. Here, the scope is to further analyse the type and description of data as expected per each of the three CONNECT use cases (see Table 5). Those use cases in specific are the i) Intersection Movement Assistance, the ii) Cooperative Adaptive Cruise Control, and iii) Slow-moving Traffic Detection. More information regarding the use cases is available on chapter 7.

Table 5 – Data type and description according to Use Case

Item	Use case	Type	Definition/Description
1.	**Intersection Movement Assistance**	Object Data	Observation (description of the kinematic state of an object)
		Numeric Data	V2X CAM sent/received from/by vehicles/MEC used for the use case
			V2X CPM sent/received from/by vehicles/MEC used for the use case
			Data about objects on the road from the in-vehicle sensors (Local Perception)
			V2X MR sent/received from/by vehicles/MEC used for the use case
			Result of Assessment of trustworthiness
			Public key certificate from vehicular PKI
2.	**Cooperative Adaptive Cruise Control**	Object Data	Data about objects on the road from the in-vehicle sensors Radar and Lidar
		Stream Data	Video feed to augment precision of object data from the in-vehicle Camera
		Numeric Data	Ego Vehicle position from the in-vehicle GNSS
			V2X CAMs sent/received from/by vehicles, used for the use case
			In-vehicle acceleration command to control the ego vehicle's speed

			Evidence on an in-vehicle ECU's integrity at design time
			Evidence on an in-vehicle ECU's integrity at runtime
			Evidence on another vehicle's integrity at design time
			Evidence on another in-vehicle's integrity at runtime
			Request for Assessment of Trustworthiness
			Result of Assessment of Trustworthiness
3.	**Slow-moving Traffic Detection**	Object Data	V2X CAM sent/received from/by vehicles used for the use case
			V2X CPM sent/received from/by vehicles used for the use case
			V2X DENM sent/received from/by vehicles used for the use case
			Ego Vehicle position from the in-vehicle GNSS
			Public key certificate to be periodically added on V2X messages
			Digest of the public key certificate to be periodically added on V2X messages
			ECDSA signature to certify all the content of the V2X message
		Numeric Data	Data about objects on the road from the in-vehicle camera sensor

5 Methodology

The methodology that has been accommodated for the creation of this deliverable is presented in this chapter, commencing from the design of the CONNECT **MVP (Minimum Viable Product)** for showcasing the common consortium vision and continuing with the elicitation of the technical and use case requirements (Section 8) for the elaboration of the core MVP features. In this context, using the methodology described here, the full landscape of CONNECT has been designed in terms of research and specifications. The correlation between technical and functional requirements and various use cases and use situations is explained in Section 7. This systematic approach guarantees a unified and organized basis for the following stages of the CONNECT project.

5.1 Methodology for MVP Design

As a term introduced by Frank Robinson in 2001, the “Minimum Viable Product” (MVP) pertains to an agile and lean approach for the planning and design of a product. An MVP constitutes a product version with an adequate number of features that can satisfy the needs of initial users, while targeting at offering input for further product improvements in the future. By providing early feedback during product development, an MVP has a lower cost compared to the implementation of a wider stack of features and the subsequent collection of feedback at the end of the implementation process of the product. In addition, an MVP can help newly established companies to discover business opportunities by experimenting on the reactions of customers while trying different business models.

As CONNECT is a Research and Innovation project aiming at reaching a Technology Readiness Level (TRL) equal to or higher than 5 at its finalisation, there is a different need for an MVP compared to a startup company, as in this case the CONNECT consortium desires to share its common vision. In this sense, even from the proposal phase, the consortium has set the goals of the project and has agreed on the core features of the CONNECT framework. Thus, a gradual MVP approach is important for the smooth development of the project and its adoption from the customers.

The three use case demonstrators of the CONNECT project will contribute to the design of an effective MVP. Even from the initial phase of WP2, they have given input about their activities that should be described in the CONNECT methodology. The demonstrators' engagement contributed to establishing a clear definition of the project's scope and purpose. This was achieved by delving into the actual benefits that the proposed solution would offer to the end users. The consortium engaged in a collaborative process, where the demonstrators shared their insights and the technical, research, and academic partners analysed the information. This process allowed the consortium to answer the fundamental question of “What Problem You're Solving.” Additionally, by describing the “as-is” and “to-be” scenarios for each demonstrator, the consortium identified areas where the project vision could be refined and demonstrated how the existing workflows could be enhanced with the intervention of CONNECT. Moreover, by acknowledging existing gaps, CONNECT also aims to develop additional valuable services that provide competitive advantages to end-user organisations.

5.1.1 CONNECT Requirements Definition Process

After discussing the initial vision during the project's DoA preparation, the consortium worked together to refine the value propositions of CONNECT. The consortium partners agreed on the core services outlined in Section 2.1 and conducted an analysis of the State-of-the-Art for each of the identified project innovations, which is further detailed in Section 2.2. This analysis was particularly important given the research nature of CONNECT as it allowed the consortium to gain insights into the competition and understand the current state of the market. The findings from this analysis will be utilised in WP7, which focuses on Dissemination, Standardization, Exploitation, and Impact Creation. This information will drive the definition of future pathways for exploiting the project's outputs and defining the go-to-market strategy for both the framework itself and its individual assets.

Moving forward with the MVP definition, the consortium proceeded to identify the stakeholders within the supply chain context. They also determined the specific actors and entities involved in the envisioned use cases, along with their respective goals (Section 7). This exercise aimed to gain a deeper understanding of the intended beneficiaries for whom the consortium is developing the solution.

To achieve this goal, it is necessary to gather and prioritise requirements. Within the context of CONNECT, requirements are gathered in two ways. Firstly, technical requirements are derived from the technical partners within the consortium who are responsible for designing and developing the overall solution. Secondly, requirements are defined based on the use cases or user stories provided by the project's demonstrators. Once all the requirements have been collected, the partner in charge of requirements elicitation, with the consensus of the entire consortium, will prioritise them.

As the project progresses, an essential component of the MVP definition is the implementation of the "Build, Test, and Learn" cycle. This cycle is integrated into the development of the core technical Work Packages 3, 4, and 5, which involve the creation of both early and advanced versions of all project modules, including the CONNECT framework itself. Through an iterative process, all technical partners will develop their respective modules and conduct testing at early and mature trial sites (demonstrators). This approach aims to maximise the value added both to each module and the whole CONNECT framework. A figure depicting the overall methodology for the CONNECT MVP definition is presented below.

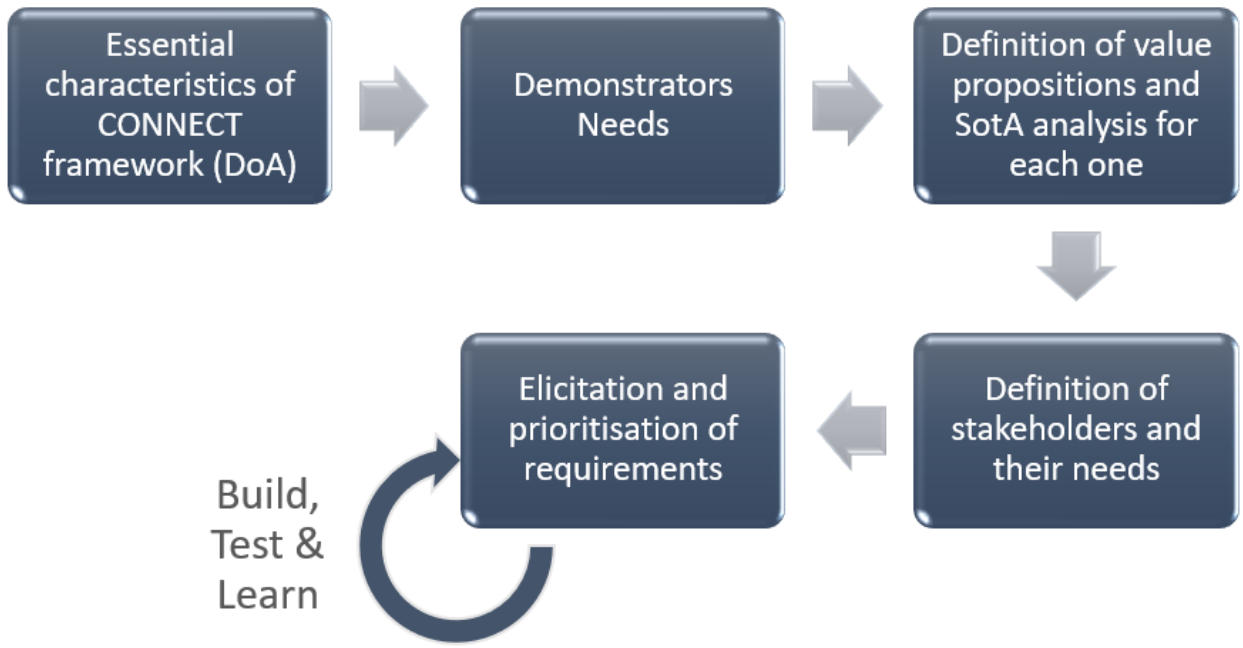

Figure 15 - Methodology for definition of CONNECT MVP

The MVP definition process outlined in this section aims to provide a specific and comprehensive depiction of the CONNECT platform. This includes a detailed examination and analysis of the technical, functional, and security requirements of the various architectural components involved.

At its core, the MVP aligns with the overall vision of the consortium and the adopted product development process, taking into consideration the expectations of the users. Its purpose is to provide tangible value, validate methodological ideas and hypotheses, and serve as a guide for design and development activities throughout the project implementation. The CONNECT MVP is expected to play a crucial role in directing and shaping the ongoing work within the project.

5.2 Requirements Elicitation Methodology

A requirement is a statement which translates or expresses a need and its associated constraints and conditions with the purpose to transform through their analysis the stakeholder, requirement-driven view of desired services into a technical view of a required product that could deliver those services [102].

The process of eliciting requirements serves as a significant foundation for developing the business value of the CONNECT framework. Since these requirements form the basis for addressing identified needs, it is vital that they are comprehensive, transparent, and accurate. To ensure successful elicitation, it is essential for the individual or organisation responsible to actively involve all relevant stakeholders and engage them in this process. The elicitation of requirements is typically not a one-time event. For example, during the elaboration phase of a project, requirements may be collected through interviews or requirements workshops. As these requirements are utilised to define and validate models or products, gaps in the requirements may become apparent, necessitating the elicitation of additional details for those newly identified requirements.

In the next subsections, an overview of the methods employed within the CONNECT project for defining stakeholder requirements is provided.

5.2.1 Extracting Technical Requirements

The CONNECT project has employed an agile methodology, characterised by its iterative nature. This methodology emphasises clear communication and understanding between the business, technical, and scientific aspects of the project. It establishes transparent expectations at the project's commencement and at each milestone (software release), fostering collaboration and progress [103].

The CONNECT approach to system requirement specification initiates with individual interviews conducted with the consortium technical partners. These interviews were carried out using online tools, resulting in the collection of raw requirements. Raw requirements refer to requirements that have not undergone analysis or been formally documented in a well-structured requirement notation [104]. The collected raw requirements are then subjected to an iterative internal process led by the CONNECT system architects to yield more refined results.

During this phase, brainstorming techniques are employed through ad-hoc calls to further enhance the requirements. The outcomes of these brainstorming sessions are subsequently compared with the technical aspects of the system and linked to one or more of the specified value propositions of the project. This iterative process ultimately leads to the identification and formulation of the necessary technical requirements, which are detailed in Section 8.

Nevertheless, it is important to acknowledge that requirements derived from interviews can occasionally be unclear. This ambiguity may arise due to potential misinterpretation of stakeholder needs by the system architects, or when stakeholders struggle to comprehend certain questions or lack technical expertise to provide accurate answers. As a result, alongside the process of requirements collection, the CONNECT academic partners have conducted a comprehensive document analysis. This analysis takes the form of a State-of-the-Art review, including an extensive list of industry practices and relevant literature. The purpose of this analysis is twofold: to validate the specified requirements and to identify applicable standards and constraints. By conducting this analysis, the aim is to enhance the quality of the requirements generated within the project.

The process of requirements collection underwent continuous review and examination by the WP2 focus group. The WP2 focus group conducted dedicated teleconferences, involving specific partners or the whole CONNECT consortium, to discuss and assess the gathered requirements. Additionally, one of their responsibilities was to offer valuable feedback on the interview participants. The success of eliciting user requirements relies significantly on the knowledge and experience of the stakeholders, making the input from the WP2 focus group crucial in ensuring the efficacy and accuracy of the requirements gathering process.

5.2.2 Extracting Use Case Requirements

In addition to the overall input on technical requirements, the use case partners provided a detailed description of their user stories and a technical explanation of how they intended to utilize the CONNECT framework. They also specified the core functionalities they planned to leverage. This allowed for a more precise alignment of requirements with the demonstrators, providing opportunities for further research and enabling the demonstrators to explore multiple implementations, if time and resources permit.

The use case partners collaborated closely with the research partners to refine the requirements and elaborate on the technical details derived from the narratives. This followed an approach commonly observed in agile projects. The refinement process involved translating the narratives into specific user stories. User stories are concise units of ideas used to provide high-level descriptions of requirements. They depict a particular feature from the perspective of an end user, typically a system user or customer. User stories are designed to be easily understandable and expressible by non-technical partners. Despite being short, typically consisting of a single sentence, user stories possess the unique ability to be self-explanatory and contain sufficient information to describe the requirement. This allows developers to provide a reasonable estimate of the effort required for implementation.

For each of the three distinct CONNECT Use Cases, the demonstrator partners were initially requested to provide user stories that describe the "to-be" reference scenario. To elicit requirements from these user stories, the perspective that a user story serves as a well-defined requirement is embraced, in the sense that:

- Emphasises the perspective of a specific role that will utilise or be affected by the solution.
- Defines the requirement using language that is meaningful to that role.
- Clarifies the reasoning behind the requirement.
- Facilitates the definition of high-level requirements without delving into detailed specifics very early.
- Takes into account the user's goals and the business value associated with each requirement.

These aspects are captured using a straightforward textual template to create a comprehensive sentence. While various templates are available, the Connextra template is employed by 70% of practitioners [105].

As a < ***type of user*** >*, I want to* < ***some goal*** > *so that* < ***some reason*** >

In an agile project, new or updated user stories may arise at any stage of implementation, leading to changes in the backlog. This behaviour is highly desirable as it ensures a continual focus on aspects that are meaningful to users, while potentially excluding features that may have less importance in terms of the value they offer to both the system and its surrounding environment.

To ensure that the user stories developed by the project's demonstrators in this step meet the necessary quality and criteria, a user story validation process is incorporated into the methodology. The objective of validating each user story is to assess the extent to which it fulfils the INVEST characteristics. The acronym INVEST serves as a mnemonic for a well-established set of criteria or checklist used to evaluate the quality of a user story. If a user story fails to satisfy any of these criteria, the team may consider rephrasing it or even rewriting it. This technique has been recommended by Mike Cohn [105].

An optimal user story, based on INVEST characteristics, should be:

- **Independent**: The user story should possess self-contained characteristics, meaning that it does not rely on another user story for its completion or functionality.
- **Negotiable**: User stories are subject to change and rewriting until they become part of an iteration.
- **Valuable**: A user story should provide value to the end user.

- **Estimable**: The size of a user story should always be able to be estimated by the team.
- **Small**: User stories should be sized appropriately to ensure they can be effectively planned, tasked and prioritised with a reasonable level of certainty.
- **Testable**: The user story or its related description should include the necessary information to facilitate the development of tests.

The user stories described for each of the three CONNECT Use Cases are presented in Section 7 as part of the corresponding “to-be” reference scenarios.

6 CONNECT Functional Components

CONNECT is envisioned as a technical solution advocating the integration of **advanced (HW-based) trusted computing primitives** for enabling the conversion of CCAM ecosystems into **trustable (heterogeneous) communication environments** allowing for the continuous trust assessment of all resources and workloads deployed over this complex landscape. As aforementioned in Section 2.1, one core foundation behind the design of CONNECT is to facilitate the realization of Day 3 CCAM services envisioning the exchange of rich information between vehicles and the backend infrastructure towards the development of cooperative automated driving capabilities. The composition of infrastructure entities extends from the traditional view of cloud-based services (including applications such as traffic control centres and/or intersection movement assistance services) and centralized security solutions (such as PKIs) to also align with the vision of **disaggregating the services over the entire compute continuum** so as to benefit from resource availability and optimal latency capabilities closer to the edge – features provided from emerging networking technologies and schemes like (B)5G and Multi-Access Edge Computing (MEC) architectures. In this context, CONNECT considers the integration of ETSI MEC[8] as a driving factor for bringing processing power near the vehicle to meet ultra-low-latency requirements, and to reduce network traffic towards a centralized data-centre. Whereas the collective consideration (treatment) of resources deployed over the entire (far edge-edge-cloud) continuum presents opportunities of increased networking performance, the **individual Software (SW) and Hardware (HW) infrastructure may exhibit diverse yet dynamic trust states**; and when those infrastructures are brought together to make-up the CCAM service ecosystem, the complexity of the established relationships increases, thus, also requiring for an **overarching trust characterization framework capable of coping with all the intrinsic attributes of the actors exchanging (kinematic) data towards safety-critical services**; i.e., enhanced mobility, low-latency requirements, zero-trust entities, etc. Thus, questions are raised on the trustworthiness level of all actors and infrastructure entities that need to behave in a secure (trustworthy) manner providing high assurances for the CCAM workloads operation over (end-to-end) communication paths, stressing the efficient (life cycle) security management need under a zero-trust approach.

This is the main goal of CONNECT: *Providing all the necessary means and security controls that can allow for the bootstrapping and dynamic assessment of the trust level of all CCAM HW- and SW-based resources; from the application to the execution environment and device hardware from the vehicle up to the MEC and cloud environments*. Towards this direction, this chapter aims at giving more details, about the phases and components, to serve as a guide for the further development of all core technical artefacts throughout the duration of the project.

More specifically, we present the conceptual architecture of the CONNECT framework (Figure 16) by documenting the interactions among its building blocks, while also elaborating on the modes of operation of the framework, namely the **Setup and Runtime phases**. The **Setup phase encompasses all operations needed for the correct establishment and deployment of CONNECT-related components** needed to support the continuous trust assessment throughout the entire lifecycle of a (data and/or entity) resource (cf. Section 6.3): From the **definition of the appropriate trust model templates considering the most prominent types of risks and attacks** against the entire CCAM ecosystem (cf. Section 6.2), dictating the type of trustworthiness evidence that need to be continuously monitored, from each resource, so as to quantify its trust level (Actual Trust Level-ATL) and compare it to the Required Trust Level (RTL), to the **deployment of all trust-related information and components as part of CONNECT'S Trusted Computing Base** (TCB – cf. Section 8.1.4) towards enforcing the circulated trust policies. The **Runtime phase** encompasses the **(runtime) operation of all CONNECT attestation schemes** (cf. Section 6.4) **and security controls** (i.e., Misbehaviour Detection service – cf. 4.2.1.1), protected through their instantiation in a Trusted Execution Environment (cf. Section 6.4.1), for allowing the secure monitoring and exchange (in a verifiable manner, as defined in D3.1 [106]) of a resource's trustworthiness evidence based on which the trust assessment/quantification will occur. Such trustworthiness evidence will be

[8] ETSI MEC ISG, "Mobile Edge Computing (MEC); Framework and reference architecture," ETSI, DGS MEC 003, April 2016

further processed, before been transmitted outside the vehicle, so as to be adequately **abstracted** (**harmonized trustworthiness evidence** as introduced in Section 6.4.3 and further elaborated in D5.1 [107]) in order to avoid privacy implications: Sharing of such sensitive information on the software stack been instantiated in a vehicle can lead to the identification of the vehicle per se (e.g., manufacturing brand) which in turn can lead to implementation disclosure attacks [108]. This evidence will be selectively disclosed, based on the trust information defined as part of the trust model per (CCAM) service, through the **construction of formally defined Trustworthiness Claims (TCs)** that need to be shared between entities that want to establish a trust relationship. Such claims (that follow the Yang Data Model[9], as has been defined by the IETF standardization body for capturing trust state evidence of network elements) are generated by the CONNECT TCB and the provided attestation enablers and comprise evidence on the system properties of a resource. Validation properties may range from static properties such as integrity measurements of the host CCAM actor (e.g., vehicle), enabling the generation of static evidence of the vehicle's components correct configuration, to dynamic properties for verifying that RTL calculations are performed by a trustworthy software. Additionally, such evidence might also originate from additional security controls such as Misbehavior Detection and cover also the **operational assurances needed for the MEC infrastructure** hosting both application- and security-related workloads (cf. Section 6.4.1.2, Section 6.5.2.1 and Section 6.6).

The rest of the sections provide an indicative workflow of interactions among the components so that the reader is able to understand the functional objective of the overall framework, as depicted in Figure 16. We have to highlight that for better conveying the flow of actions across the entire CCAM continuum, Figure 16 is broken down into two conceptual parts: On the **upper part depicting the CONNECT-related components operating on the cloud (and are mainly focused on supporting the runtime node- and data-centric trust assessment operations)** and the **lower part** capturing the positioning and interactions of the **CONNECT security enablers deployed over the MEC and the (far-edge) vehicle component**. In this context, the left MEC and vehicle configuration (MEC-A and Vehicle i) provide a more nuanced view of all CONNECT operations (that will be further described in Section 6.6) while the right configuration (MEC-B and Vehicle ii) convey a more abstract view of the CONNECT components and how they support the security lifecycle management of all CCAM actors.

6.1 CONNECT Conceptual Architecture

Overall, the CONNECT framework follows a **three-tiered architecture** for providing the envisioned trust assessment functionalities with components been deployed over: i) *the cloud*, ii) *the MEC* and iii) *the (far-edge) Vehicle*, thus, capturing all layers and actors of the CCAM continuum. The endmost goal, as already described, is to provide the tools for the **trustworthy communication between entities/resources with no inherent trust**. Thus, the central element of the entire architecture is the **CONNECT Trust Assessment Framework (TAF)** (cf. Section 6.3) that is capable to efficiently assess and quantify the trustworthiness level of any CCAM actors that want to establish a relationship for further exchanging information in a secure/trustworthy manner. In this context, **trust relationships pertain both to node-centric trust but also data-centric trust**. For instance, in the context of the envisioned *Intersection Movement Assist* use case (cf. Section 7.2), where intersection manoeuvre decisions are made per vehicle based on its own perception (of the environment) but also on the perception information exchanged by neighbouring vehicles, trust calculations mainly target each exchanged CPM message. Nonetheless, for a more accurate data-centric trust quantification, this also needs to be based on the trust level of the data originator, i.e., integration and deployment of different types of security controls can provide varying levels of assurance on the correct operation of the target device, thus, resulting in different RTL requirements.

Therefore, each type of trust relationship to be assessed (i.e., Vehicle-to-Vehicle, Vehicle-to-Infrastructure, and vice versa) needs to be modelled through different **trust model templates defining both the RTL that needs to be exhibited by each communicating party but also the type of evidence that need to be monitored for allowing the calculation of the (runtime) ATL**. This is performed by the **Trust Management Framework (TMF)**, deployed over the cloud, that

[9] Internet Engineering Task Force (IETF) RATS Working Group. Attestation Results for Secure Interactions, Sept. 2021.

works in tandem with the **CONNECT Risk Assessment Engine** (cf. Section 6.2). The CONNECT RA is responsible for (statically) generating assessments reflecting the baseline cyber-security posture for the target CCAM deployment. This assessment contains possible **vulnerabilities and threats that may compromise the deployment** based on the threat landscape and adversarial model (detailed definition of the considered threat landscape will be documented in D2.2 [109]), as provided by the respective Original Equipment Manufacturer (OEM). Based on this, an initial **risk graph** is generated capturing all this type of most prominent threats and attack vectors as well as their impact on the operation of the target actor. The reasoning is to allow the TMF to identify the RTL for this specific type of trust relationship: *Depending on the assets involved and the possible set of security controls deployed (spanning from mechanisms such as secure boot, integrity verification to intrusion detection), a different trustworthiness level will need to be exhibited for appraising an entity as trustworthy*. For instance, a Vehicle been equipped with secure boot capabilities (as part of its in-vehicle set of Electronic Control Units (ECUs)) will need to exhibit a higher RTL from a Vehicle equipped with runtime integrity monitors capable of monitoring its host integrity beyond the boot-up procedure and into its runtime operation. All this information that essentially guide the (runtime) trust assessment process is performed in tandem between the CONNECT TMF and RA components in a supervised manned under the OEM or a Security Administrator. It needs to be noted that risk assessments will not actually be based on the cloud since they will run in the premises of each Original Equipment Manufacturer (OEM) and will then be communicated to the CONNECT cloud, i.e., each OEM will perform its own risk assessments.

There are **two types of trust model templates considered for capturing the different characteristics of the relationships to be established across the different layers of the CCAM continuum, namely static and dynamic trust models**. The differentiating factor is on the frequency with which the topology modelled as part of the trust template can be subject to changes. For instance, a Vehicle-to-Everything relationship encompasses the communication across different layers; e.g., between Vehicles, or between the Vehicle and the MEC where a service has been deployed, or even the Vehicle and the backend cloud infrastructure where either application services or security-supporting services (e.g., Public Key Infrastructures) might have been deployed. Such communication topologies need to be able to cope with this high degree of mobility exhibited by moving Vehicles. In addition, different brands of Vehicles might be equipped with different sets of security controls, thus, requiring the definition of separate trust models. Such requirements are due to the highly dynamic, distributed and ubiquitous nature of CCAM services that operate with a high degree of autonomy in unpredictable, uncertain and continuous changing environments, which inherently introduces a high degree of uncertainty in the exchanged data. They are captured by the **dynamic trust model templates** that are defined by the CONNECT TMF, for this exact V2E type of relationships, and they are deployed to the CONNECT TAF agents through the policy-compliant Blockchain infrastructure (cf. Section 6.7).

On the other hand, all these types of V2E relationships share the communication of ETSI V2X messages (such as CAM and CPM messages, as detailed in Section 4.2.3) based on which safety-critical decisions are made. This type of kinematic (or collective perception) messages originate from the in-vehicle network of sensors and ECUs such as cameras, LIDAR, etc. Thus, it is also vital to enable a **vehicle-wide trust assessment that can then be elevated to a CCAM continuum-wide trust quantification capturing this highly distributed operational environment**, as aforementioned. This in-vehicle E/E topology, however, is rather static since ECUs or other in-vehicle equipment can be changed mainly due to maintenance purposes but not in a dynamic manner – but rather in a controlled environment (e.g., garage). This type of **in-vehicle trust relationships are captured through the CONNECT static trust model templates** that are directly pushed by the OEM (who is aware of the in-vehicle E/E topology) to the CONNECT TAF agent instantiated in the target vehicle.

It is also important to note that the RTL per trust model (or per type of relationship modelled through a specific trust model template) can change during the lifecycle of the target system. As aforementioned, the definition of the RTL is based on a combination of the type of threats that can impact the operation of the system and the set of security controls deployed for safeguarding against identified attacks. Thus, in the case of a **zero-day vulnerability been identified (during runtime)**, this will result in the creation of an updated risk graph which, in turn, might lead to the deployment

of new RTL values and security policies to be enforced throughout the entire CCAM continuum. Such information on possible indications of risk in the operation a system come from the execution of the attestation schemes: CONNECT enables the provision of both **Configuration Integrity Verification (CIV) but also more advanced Control-Flow Attestation (CFA) capabilities** that can verify the configurational and executional correctness of a system (cf. Section 6.4.2). This is also enhanced with allowing for a **local attestation/verification to occur based on newly defined key restriction usage policies**: As will be described in D4.2 [110], in contrast to the current state-of-the-art attestation mechanisms that require from a device (acting as Prover) to share its attestation evidence to a remote Verifier that will perform the verification process, CONNECT converts the underlying CONNECT TEE Guard as the local Verifier. **A HW-based key is established and binded to the correct state of the device**, thus, the CONNECT trust anchor will not allow the use of the key for signing any outgoing trustworthiness evidence unless the device is deemed to be at a correct state. This transforms the **verification process to a simple signature verification without any assumptions on the trustworthiness of the Verifier** (facilitating the vision of zero-trust) and without requiring the disclose of the low-level attestation evidence that contain sensitive information about the system's operation.

The **(auditable) security and trust policy enforcement** is facilitated then through the **CONNECT Blockchain infrastructure** (cf. Section 6.7) that allows for the secure execution of data sharing agreements. Essentially, CONNECT DLT has a two-fold role: It allows for the secure exchange and deployment of updated security and trust policies only to the intended CCAM actors and their instantiated TAF agents. Furthermore, this infrastructure hosts a "*hygiene pool*" by recording the history of trust levels (per CCAM actor) and the low-level system evidence (traces) in case of a failed attestation result. The former allows the built-up of some form of "*reputation*" which allows for a higher accuracy in the trust quantification process for this specific system while the latter allows for the OEM (or Security Administrator) to get access to detailed logs and traces so as to pinpoint any new types of vulnerabilities that were exploited. This, in turn, triggers the CONNECT RA process for re-calculating the overall risk graph based on which the new trust policies and RTLs are calculated.

This enforcement of dynamically calculated security and trust policies, is governed through the CONNECT security controls and trust enablers that are deployed both at the Vehicle and MEC layers (Figure 16) and they constitute the core building blocks of the previously described Runtime Phase. The **Runtime Phase** is an ongoing and critical stage that spans the entire lifetime of a CCAM actor (mainly the Vehicle and the MEC where safety-critical services are deployed subject to trust assessment), encompassing all major software processes within the CONNECT framework. This phase is characterized by the **active execution of trust assessment procedures (following the trust model templates defined)**, where an entity performs its local trust quantification by *fusing* evidence originating from its local sensors/actuators but also from neighbouring entities and the MEC. For instance, a Vehicle calculates a local trust opinion on its perception of an area based on the kinematic data it produces leveraging the trustworthiness evidence monitored from all ECUs comprising this specific service graph chain; i.e., from the LIDAR, to the Zonal Controller and subsequently to the vehicle's facility layer where CAM/CPM message blocks are collected to be forwarded to the CAM/CPM Encoding/Decoding component (see Figure 53 for a reference scenario in the context of the Collaborative Adaptive Cruise Control use case). This local trust opinion is then enhanced/compared/amended with trust calculations on the CPM messages received by neighbouring vehicles but also from the MEC on the trust level of the vehicle transmitted this message (combination of data- and node-centric trust quantification). This advanced trust assessment/fusion is performed for every message communicated based on the defined trust models, building on top of Subjective Logic as introduced in Section 6.3 and further described in D3.1: Overall, it computes trust based on both direct evidence (as produced by the local CONNECT security controls and attestation enablers) but also indirect evidence (obtained via referral paths of the trust topology defined) and can combine this with referrals about reputation of sources (as extracted from the CONNECT DLT).

The runtime phase involves constant monitoring, evaluation, and adaptation, empowering the trust assessment framework to respond dynamically to changes in the operational environment. Through continuous data collection and analysis, the framework remains agile and adaptive, establishing a secure and trustworthy environment for seamless CCAM operations. This phase stands as a

cornerstone in the overall effectiveness of the trust assessment framework, providing a responsive mechanism to uphold trust in the intricate and ever-evolving landscape of CCAM.

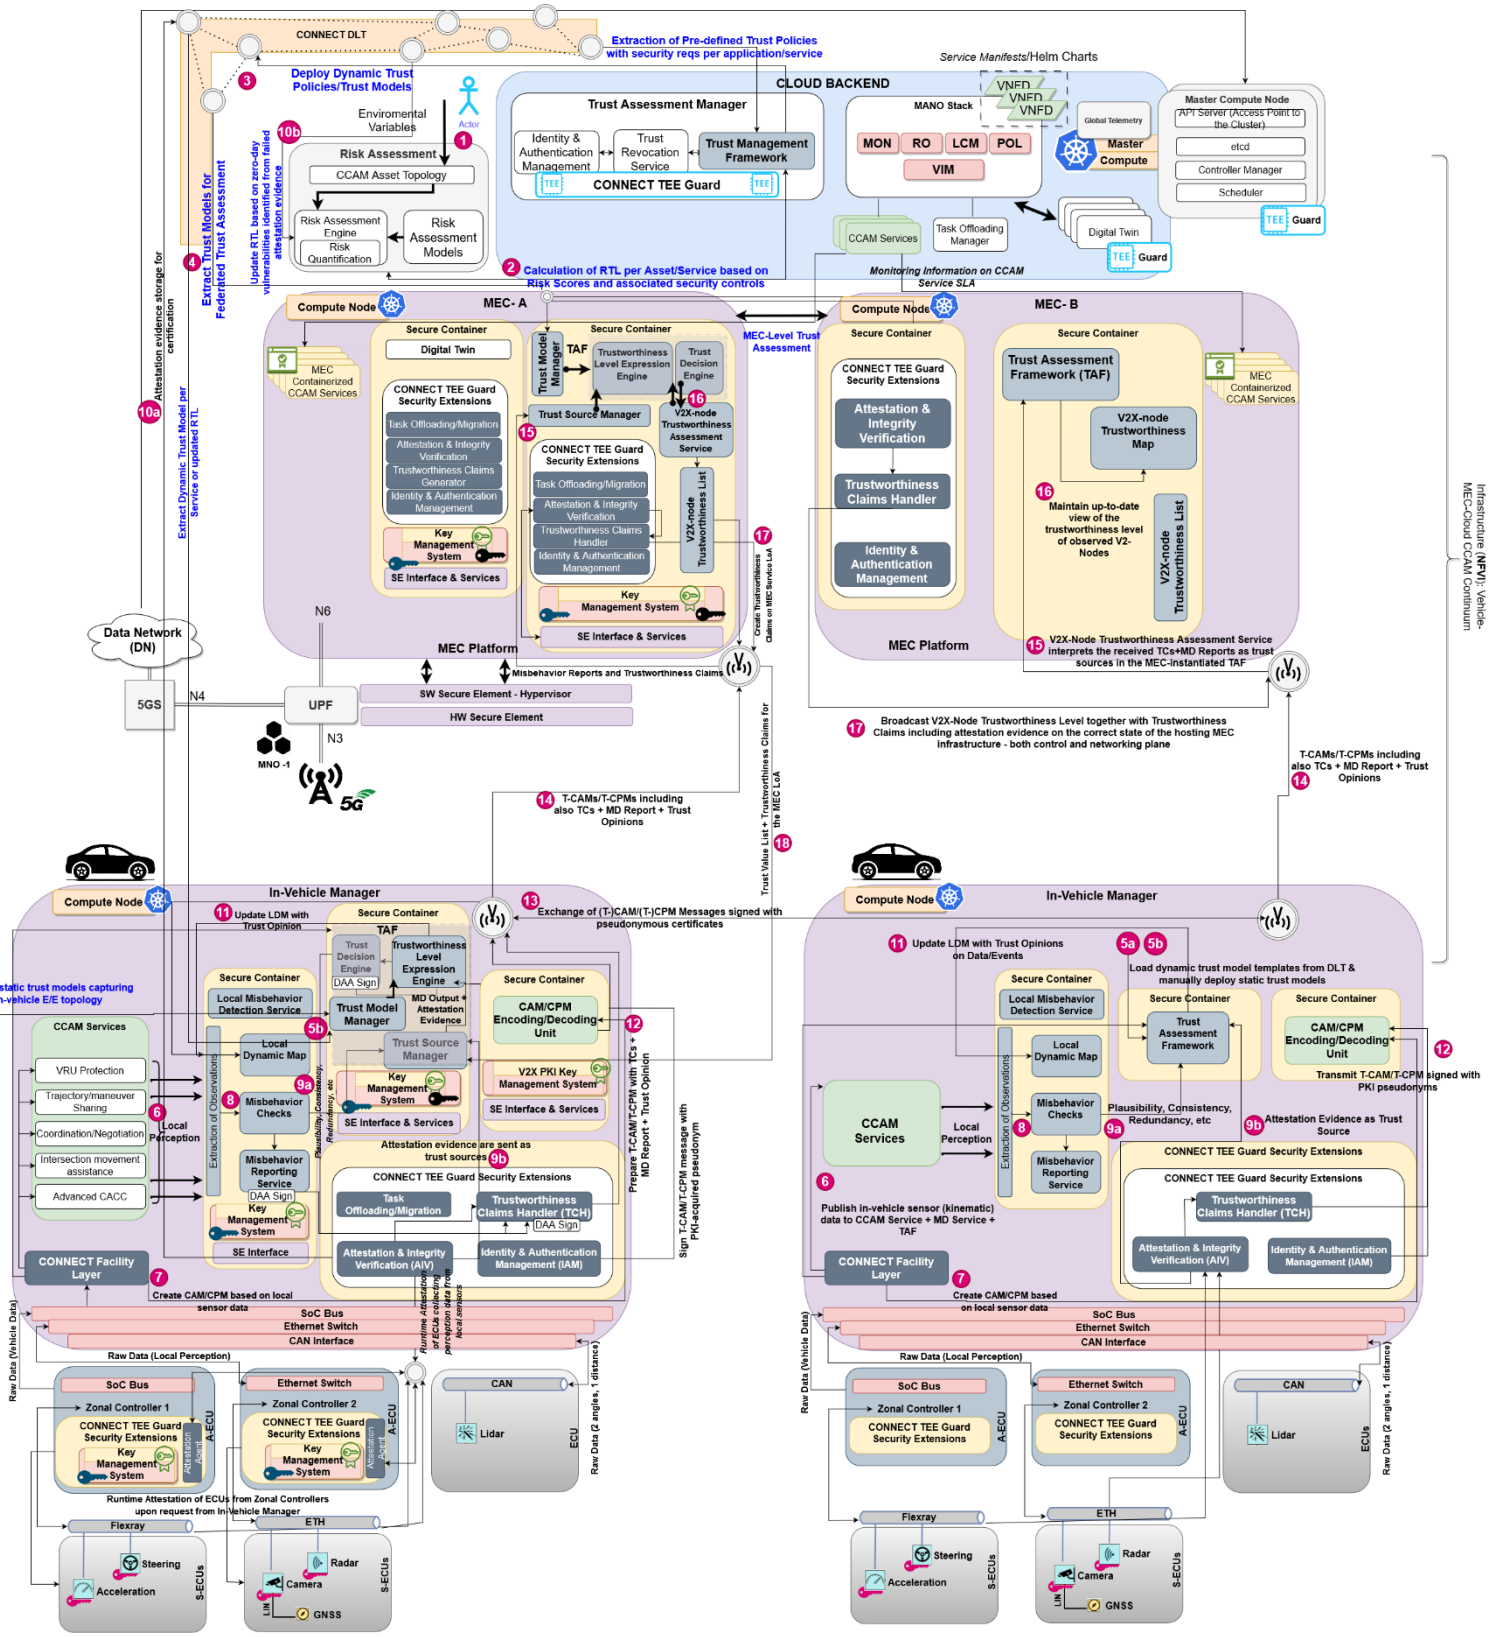

Figure 16 - CONNECT Reference Architecture

6.1.1 High-Level Sequence of Actions

Building on top of this high-level description of CONNECT's conceptual architecture and building blocks, in what follows, we provide a brief overview of the sequence of actions that take place for both setting up and deploying the appropriate trust model templates (Setup Phase) but also the runtime trust quantification process (Runtime Phase) that takes place across all layers of the CCAM continuum. Figure 16 provides a detailed view of the positioning of all CONNECT components across the Cloud, MEC and the Vehicle for supporting this continuous trust assessment process. Recall that a more nuanced view of the architecture is depicted on the left part of the figure whereas on the right part the same flow of actions is depicted with the internal CONNECT building blocks (as specified in Section 6.1.2) been abstracted.

The first step pertains to the **definition of the appropriate (static and dynamic) trust model templates** for capturing all type of V2E relationships to be assessed. Towards this direction, the Security Administrator, the OEM and the Mobile Network Operator (MNO - that manages the MEC infrastructure where any CCAM services and CONNECT security services will be deployed) provide an extensive list of assets, covering both software and hardware infrastructure elements, along with their interdependencies. The goal is for the CONNECT Risk Assessment engine to use this information so as to create risk graphs quantifying the risk score of each prominent type of threats so as to enable the CONNECT TMF to then calculate the necessary RTL per trust model. This information contributes to the creation of a comprehensive in-vehicle topology (**Step 1**) for which a Threat Analysis and Risk Assessment (TARA) process will be employed (following the ISO SAE 21434) as well as a risk graph depicting the threats of the MEC virtualized infrastructure. For the latter, the CVSS v3.0 risk quantification methodology will be employed (D3.2 [112]) alongside the security controls identified by ETSI[10] for enabling the provision of different Levels of Assurance both for the SW but also HW infrastructure elements of the (virtualized) MEC infrastructure (**Step 2**). Based on such evidence, CONNECT also assess the level of trust of MEC-deployed services.

This numerical score reflects the assessed impact of identified threats, further considering their feasibility. The CONNECT Trust Management Framework receives the risk score and, together with information also provided by the OEM and MNO on the security controls deployed in the vehicles and the MEC, calculates the trust model template comprising the RTL to be deployed in the CONNET TAF Agents. As aforementioned, this trust model template can be deployed and enforced either through the DLT (in case of dynamic trust models, as depicted in **Step 4** and **Step 5b**) or directly (by the OEM during the manufacturing process of an ECU, as depicted in **Step 5a**) to the Vehicle Computer capturing the in-vehicle (static) E/E topology and the internal trust relationships to be assessed. The accessibility of this information extends to the Trust Model Manager of the Trust Assessment Framework (TAF) at both the Multi-Access Edge Computing (MEC) and In-Vehicle levels **(Steps 4, 5a and 5b),** complemented by the deployed Trust Policies and Trust Models.

Each trust model is linked to a different CCAM service deployed over both the vehicle, MEC and (possibly) the backed cloud infrastructure. It shall be highlighted that even though the in-vehicle E/E topology, that is reflected by a static trust model, cannot be dynamically altered (i.e., unless the vehicle itself is modified in a garage), the RTL may need to be updated since new threats and vulnerabilities may be identified during the runtime operation of the vehicle (detected through the CONNECT attestation enablers as part of the overall provided trusted computing base – **Step 10a**). Hence, after the static information is inserted, the TMM may further receive dynamic updates regarding the Trust Policies (i.e., including RTL and TMs) from the DLT (**Step 10b**).

The aforementioned steps mark the conclusion of the Setup Phase, triggering the real-time, evidence-based trust assessment which is initiated by a CCAM service through the communication of a Trust Assessment Request (TAR) (**Step 6**). Given the predefined nature of the RTL, it is used as an enabler for the runtime trust decisions, serving as the benchmark against which the Actual Trust Level (ATL) is compared.

[10] ETSI GR NFV-SEC 007, "Networks Function Visualization (NFV) Trust; Report on Attestation Technologies and Practices for Secure Deployments", 2017, [Available Online: https://www.etsi.org/deliver/etsi_gr/NFV-SEC/001_099/007/01.01.01_60/gr_nfv-sec007v010101p.pdf]

During the runtime phase, it is essential to highlight the simultaneous occurrence of two distinct flows: one pertaining to the operation of the *CCAM application/service* and the other dedicated to the *CONNECT-enabled continuous trust assessment*. These parallel flows denote the dynamic and interconnected nature of all these processes during runtime, ensuring the seamless operation of both the underlying CONNECT infrastructure and the specific functionalities offered by the CCAM application or service.

Regarding the CCAM application/services, observations collected by the sensors are sent to the Zonal Controller and forwarded to the CONNECT Facility Layer for further communication to all vehicle's internal components; i.e., to the CCAM Services portfolio (i.e., VRU Protection, Trajectory/manoeuvre sharing, coordination/negotiation, Intersection Movement Assistance, Advanced CACC), the CAM/CPM Encoding/Decoding Component (**Step 7**) for starting to create the appropriate V2X message structures (including the observed kinematic data), and also the TAF and Misbehaviour Detection Services for analysing any inconsistencies. Among these observations data related to the vehicle's position, speed, direction, and other relevant kinematic parameters are included. These raw data are utilised by the CCAM services to accommodate specific tasks, while further serving a role in creating the vehicle's Local Perception (LP). This involves processing the raw sensor data to form a comprehensive and localized understanding of the vehicle's immediate surroundings. The LP essentially represents the vehicle's perception of its environment. The raw data will further be integrated into the CAM/CPM payload by the CAM/CPM Encoder/Decoder (**Step 7**) ready to be then shared with other vehicles in the vicinity. This can either be done immediately, without the inclusion of the local trust opinion as calculated by the CONNECT TAF Agent, or it will also include the Trustworthiness Claims (capturing the trustworthiness evidence of all in-vehicle sensors that produce or process the kinematic data, "*harmonized/obfuscated*" by the CONNECT Trustworthiness Claims Handler (TCH) so as to avoid any vehicle privacy implications – see below **Step 12**). This depends on the defined trust models and whether they have considered the trust assessment based also on indirect evidence that, as aforementioned, can be extracted via referral paths of the trust network (i.e., TAF Agents of the MEC and/or neighbouring vehicles). In case, such Trustworthiness Claims together with the trust opinion (as calculated by the local TAF) is included in the CAM/CPM messages then this forms a **CONNECT T-CM/T-CPM message** that is broadcasted to all vehicles in the vicinity.

The Local Perception (LP) of a vehicle, as established by the CCAM applications and further amended with the trust opinions (outputted by the CONNECT TAF – **Step 11**) and processed based on the misbehavior detection reports (observations that fail the BD plausibility checks are discarded from the LP – **Step 9a**), is thoroughly scrutinised by the Local Misbehaviour Detection (MD) Service to identify any discrepancies or anomalies in the processed data. This step is crucial in identifying potential misbehaviours or malicious activities in the interpreted sensor data. The MD Report, resulting from this analysis, provides a tangible indication of any irregularities or abnormalities in the vehicle's perception or sensor data, thereby ensuring the trustworthiness and reliability of the connected CCAM ecosystem.

As it pertains to the CONNECT Trust Assessment framework, the output of the Local MD Service, that is the MD Report, is shared with the Trust Source Manager (TSM) of the TAF **(Step 9a)**. This, essentially, constitutes one of the trust sources based on which the runtime trust assessment procedure will be performed. The other trust source is based on any other security controls deployed; e.g., attestation enablers and/or intrusion detection systems. In the context of CONNECT, the TAF triggers the Attestation and Integrity Verification (AIV) component (through the transmission of a Request for Evidence (RFE) – **Step 9b**) to collect and verify the attestation evidence, stemming from all in-vehicle sensors (i.e., A-ECUs, S-ECUs, and N-ECUs and including both sensors and zonal controllers as defined in Section 6.5), confirming that the entities providing the evidence are at the expected state which is equivalent to not been compromised. This is enabled through CONNECT's new breed of **runtime behavioural attestation enablers** (D4.1 [111] and D4.2 [110]), targeting both the software and hardware layers and covering all phases of a device's execution. From the trusted boot and integrity measurement of an ECU, enabling the generation of static, boot-time or load-time evidence of the system's components correct configuration (Configuration Integrity Verification (CIV)), to the runtime behavioural attestation of those safety-critical components of a system providing strong guarantees on the correctness of the control- and information-flow properties. All

these capabilities compose the main CONNECT TEE Guard extensions that are built on top of the underlying Root-of-Trust (RoT). As explained in Section 8.1.4, CONNECT is agnostic to the type of RoT employed as long as it can provide specific properties; i.e., Root-of-Trust for Measurements, Storage and Reporting. However, for supporting the implementation activities of CONNECT, the entire Trusted Computing Base has been built on top of the Grammine Trusted Execution Environment (cf. Section 6.4.1).

The result of this attestation process is summarised in an attestation report that is shared with the Trustworthiness Claim Handler (TCH) and the TAF **(Step 9b)**. In the event of an attestation failure, the AIV also securely transmits encrypted raw evidence to the DLT, providing a transparent repository accessible to OEMs for in-depth analysis and investigation into the underlying reasons behind this failure, thus, possibly identifying zero-day vulnerabilities **(Steps 10a and 10b)**. This secure recording of the low-level system traces is protected through the employment of Attribute-based Encryption (ABE) scheme so as to allow for only authenticated and authorized users and entities (that can exhibit ownership of the appropriate attributes) to get access to the attestation evidence. In this context, CONNECT adopts the Self-Sovereign identity (SSI) concept for allowing the vehicles data sovereignty – they define the usage control policies that are associated with each data recorded on the Blockchain including also the type of attributes that an entity needs to present for getting access. This decentralized identity management is then based on the construction and provision of Verifiable Credentials (VCs) and Verifiable Presentations (VPs) for allowing an entity to selectively disclose only the required attributes in a verifiable manner (constructed and verified by the underlying CONNECT TCB).

The TCH plays a pivotal role in the post-AIV phase, where it not only harmonizes the information within the AIV report but also constructs the necessary Trustworthiness Claims (TCs). It ensures that the trustworthiness of the vehicle can be verified, while ensuring that vehicle fingerprinting is not feasible when this information is shared beyond the boundaries of the vehicle.

The TAF, and more specifically the Trust Source Manager (TSM) harnesses the collected information, encompassing i) the attestation report, ii) the MD report, and optionally iii) the Intrusion Detection (IDS) report, and provides them (as trust sources) to the Trust Assessment (TA) which will leverage the help of Trustworthiness Level Expression Engine (TLEE) to derive an ATL **(Step 11)**. In conjunction with this threefold information, the TLEE further needs information regarding the Trust Models (TMs). The TM, as defined earlier on, leverages the static information as provided by the security administrator, while it can further receive dynamic updates regarding the Trust Policies (i.e., including RTL and TMs) from the DLT. To define the TM for the specific vehicle setup, the Trust Model Manager leverages the templates available at the DLT. The TLEE leverages all information available to it (threefold information which is used as evidence and TM) to execute the complex calculations necessary for creating subjective trust expressions that are then converted to trust quantifiable values, by the TA, determining the Actual Trust Level (ATL). The calculated ATL is then compared against the pre-defined RTL, making a decision on whether this node or data is deemed trustworthy (Trust Decision (TD)). In parallel, information goes into the Local Dynamic Map (LDM), but is annotated with Trust Opinions (TOs), as calculated by the TAF **(Step 11)**.

The TCH, which is responsible for the harmonisation of attributes (more information is to be provided in D5.1 [107] where the data model to be adopted for the Trustworthiness Claims will be defined) receives the following information: i) the attestation report, ii) the MD report as a signed Verifiable Credential (VC), iii) the TAF Report as a signed VC and optionally iv) the IDS report also signed as a VC. Results from the attestation report are obfuscated (attestation results per in-vehicle sensors is hidden in a group-based attestation result depicting whether the entire service has (for instance) integrity or not) and used to generate a VC to protect the privacy of the vehicle when this information is broadcast. These three (or four) VCs are combined by the TCH in a Verifiable Presentation so as to disclose only this information necessary for the receiving entity to assess as part of its own trust calculations. The resulting VP is then forwarded to the Identity and Authentication Management (IAM) component that is responsible for managing all the necessary keys (both CONNECT attestation-related but also CCAM application-related as well as traditional ETSI PKI-enabled pseudonymous certificates – cf. Section 6.5) for signing the VP with one of its pseudonyms. This VP is finally forwarded to the CAM/CPM Encoder/Decoder for been added as part of the outgoing CAM/CPM messages resulting in the construction of the T-CAM/T-CPM message (**Step 12**).

This VP can then be evaluated from standalone TAFs that reside in other vehicles (**Step 13**) or at the TAF instantiated on the MEC **(Step 14)**. The collection of TAF agents, running across different layers of the CCAM continuum, composes the CONNECT Federated Trust Assessment capabilities allowing for a more accurate trust quantification process since the TAF, running on the MEC, is capable of constructing trust opinions by considering multiple trust sources as collected by the entire CCAM ecosystem whereas (vehicle) local TAF agents are based only on their local data and those received from the neighbouring nodes (**Steps 15-17**). Further information on the detailed analysis of all CONNECT TAF modes of operation can be found in D3.1.

Within the MEC-instantiated TAF, the node-based trust assessment result **(Step 15 and 16)** is examined, according to the pre-defined Trust Model. The MEC-instantiated TAF collects information from multiple vehicles in the vicinity, hence it may possess a more consolidated view. This analysis gives a trustworthiness list of nodes **(Step 17)** based on the VPs (i.e., which include the TCs) received from multiple vehicles in the area. The list, together with the TCs are signed by the MEC-based IAM and transmitted back to the in-vehicle TAF in the form of an IAM VP **(Steps 17 and 18)**. The in-vehicle TAF can further utilise this input to take trust-related decisions. For example if the result of the Federated TAF gives another indication for a node, then the in-vehicle TAF may decide to place more emphasis on the opinion of the Federated TAF; hence update its own trust opinion regarding this specific node.

6.1.2 Building Blocks

After the provided overview of the CONNECT conceptual architecture, this section describes the components that work in synergy in order to provide the envisioned functionalities for supporting this continuous CCAM-wide trust assessment process. Our aim is to define the high-level framework that provides the necessary services to enable dynamic trust assessment when composing a CCAM ecosystem, alongside with the new CONNECT Trusted Computing Base and security control services required for allowing enhanced monitoring of all trustworthiness evidence needed as sources in the trust evaluation.

The core architectural components set the scene of implementation in the context of WP3-WP5 and need to support the overall framework by providing scalability and efficiency. In light of the above, we proceed to a component-oriented analysis of the framework. *Recall that in the previous section we have provided a description on the indicative workflow of actions that can help the reader understand the operational behaviour of the framework as a whole.*

1. **Master Compute Node:** It acts as the orchestrator of this distributed CCAM service architecture. As described in Section 6.1, the CONNECT MEC infrastructure hosts both CCAM services but also CONNECT security-related services. The orchestration and deployment of these services is facilitated through the Master Compute Node following the Kubernetes virtualization and orchestration technology. The main role of this component is to manage the resource allocation to the services, deployed on the MEC, based on the requirements defined by the Service provider in the form of Service Level Agreements (SLAs).
2. **Management and Orchestration (MANO) stack:** It is a framework designed to manage and coordinate virtualized network resources. The system consists of many elements, including the Virtualized Infrastructure Manager (VIM), the Virtual Network Function Descriptor (VNFD), the Network Function Virtualization Orchestrator (NFVO), Lifecycle Management (LCM), Policy (POL), Resource Orchestrator (RO), and Monitoring (MON). The VIM, or Virtual Infrastructure Manager, is responsible for managing physical resources. The VNFD, or Virtual Network Function Descriptor, provides detailed information on the structure and needs of a Virtual Network Function (VNF). The NFVO, or Network Function Virtualization Orchestrator, manages the deployment and interconnection of VNFs. The LCM, or Lifecycle Management, guarantees continuous monitoring of VNFs. Lastly, the Policy sets rules for resource allocation. The Resource Optimizer (RO) improves the allocation of resources, hence increasing the efficiency of the network. Meanwhile, the Monitoring (MON) module observes network performance to make informed decisions. The MANO stack offers a comprehensive and flexible framework for dynamically controlling virtualized network operations. This is considered an auxiliary service to support the CONNECT trust-related operations deployed on the MEC. Both the Master Compute

Node and the MANO Software Stack constitute the services responsible for setting up and managing the entire lifecycle of MEC-deployed services as well as the CONNECT-related services deployed as part of the Vehicle's main on-board unit. All these services are packaged through secure containers as described in Section 6.5.

3. **Risk Assessment**: Within the CONNECT project, the Risk Assessment component plays a pivotal role in evaluating and quantifying potential security risks regarding all the CCAM actors and entities, as reported by OEMs or other relevant institutions and MNOs. This component is initiated by the OEM/MNO/Security Administrator, who inputs a comprehensive list of assets, encompassing both hardware and software elements, along with their interdependencies. The Asset Modelling and Visualization component then utilizes this information to create a detailed asset graph, providing a visual representation of the asset relationships. Subsequently, the Risk Assessment engine automatically calculates risk scores for each asset, quantifying potential vulnerabilities. In essence, the Risk Assessment component provides a critical foundation for the calculation of the Required Trustworthiness Level (RTL) understanding and addressing potential security threats, contributing valuable insights for subsequent phases in the CONNECT architecture.
4. **Trust Assessment Manager (TAF):** Within the CONNECT project, the Trust Management Framework serves as a central component responsible for defining the RTLs through the risk scores as identified by the Risk Assessment module as well as the dynamic trust models (per service), and uploading them at the DLT side so as to be automatically deployed to the TAF agents instantiated across the different layers of the CCAM continuum.

   - ✓ **Trust Revocation Service:** It oversees events that may require the revocation of trust, such as security breaches or misbehaviour detection. The Trust Revocation Service implements trust revocation decisions according to pre-established policies, guaranteeing the security and trustworthiness of the system. The system acts upon modifications that affect the trust levels, conveying revocation decisions, and meticulously recording occurrences.
   - ✓ **Trust Management Framework:** The TMF leverages the information outputted by the Risk Assessment Engine to enable the supervised calculation of the initial Required Trustworthiness Levels (RTLs) needed for bootstrapping the trust level of all actors considered in a trust model. This process is essentially a methodology that will be followed by the (for instance) Security Administrator so as to define the appropriate RTLs based on the considered threat model. As highlighted previously, the RTL can be dynamically updated during runtime in case of any new vulnerabilities been identified. This RTL, in essence, represents a baseline for the accepted Trustworthiness Level while it further identifies the attributes that need to be attested during runtime for the Actual Trustworthiness Level (ATL), and will be used to form the security claims. Note that the centralised instance of the Trust Assessment does not encapsulate the same components as the In-Vehicle and MEC based Trust Assessment Framework, since the scope is different. The cloud-based Trust Assessment calculates the RTL and sends it to the Distributed Ledger Technology (DLT) for further dissemination. In addition to the RTL, Trust Policies and Trust Models are further incorporated. By residing on the cloud, the TAM facilitates centralised control and coordination, enabling the efficient management and distribution of trust-related information across the entire CONNECT CCAM ecosystem, through the DLT (for the dynamic TMs).

5. **CONNECT Trust Assessment Framework (TAF):** It provides one of the core functionalities of the CONNECT framework. The TAF can be employed either on the In-vehicle side or at the MEC-level depending on several factors, such as resource availability. The following subcomponents are integral parts of the overall TAF architecture.

   - ✓ **Trust Model Manager (TMM):** It is an integral component of the TAF. Its primary role involves the derivation, storage and provisioning of Trust Models (TM), crucial for the Trust Assessment (TA). Specifically, the TMM stores and disseminates to the Trustworthiness Assessment (TA), trust models tailored for distinct functions operating within a CCAM system – considering both CCAM services/applications deployed in a vehicle (recall that there is a different TM per brand of vehicle considering the type of equipped security controls), and processes supporting a CCAM application deployed on the MEC. It takes into account

diverse scopes that these trust models may encompass, ensuring a comprehensive coverage of trust-related considerations. The TMM plays a vital role in maintaining a repository of trust models that are crucial for evaluating and determining the trustworthiness of entities and processes within the CONNECT ecosystem. Trust Models can further derive from the DLT of CONNECT, comprising of an additional source of information for the TMM.

- ✓ **Trust Source Manager (TSM):** It is a component of the Trust Assessment Framework (TAF), that manages the diverse Trust Sources (TS), representing the multiple sources of trustworthiness evidence (i.e., the AIV, the MD and the IDS).
- ✓ **Trust Assessment (TA):** The Trust Assessment (TA) is the component inside the TAF responsible for the derivation of the Actual Trustworthiness Level (ATL). For this, it compiles all necessary information (TMs, TSs, and specific proposition to be evaluated) and initiatives the evaluation of the specific proposition by the TLEE.
- ✓ **Trustworthiness Level Expression Engine (TLEE)**: The Trustworthiness Level Expression Engine (TLEE) implements the actual analysis and reasoning over the Trust Models (TM) by substituting the concrete values of Atomic Trust Opinions (ATO) on the trust sources gathered by the TSM and conducting the quantifiable assessment of trustworthiness tailored to the unique context of the Trust Model (TM) and proposition under consideration.
- ✓ **Trust Decision Engine (TDE):** It is the final component within the internal TAF, executing the conclusive step before delivering a Trust Decision (TD) to the requesting application. The decision-making process is based on the comparison of the ATL with the RTL, in a predetermined manner.

6. **Trust Assessment Framework on Digital Twin (TAF-DT)**: In the context of CONNECT it acts as an advanced virtual replica of specific tasks. More specifically the Digital Twin replicates the TAF, including its TMs and TSs in a MEC. This allows a vehicle to outsource trust assessment process to a MEC where the TAF-DT is expected to run inside a TEE so that its data and state can be kept confidential. The CONNECT Digital Twin facilitates in-depth analysis, scenario testing, and informed decision-making, contributing to the optimization and advancement of CCAM capabilities. It may reside on the MEC or at the cloud side.
7. **Task Offloading Service:** This service enables the distribution and delegation of computational tasks from in-vehicle systems to the MEC environment, ensuring efficient utilization of resources. By facilitating seamless and dynamic task distribution, this service contributes to the efficient utilisation of the available resources between in-vehicle processing and MEC-based computing.
8. **Distributed Ledger (DLT):** The DLT in the context of CONNECT acts as a decentralized, secure and permanent (auditable) record-keeping system that maintains a transparent and immutable history of transactions. This ledger employs the Blockchain technology to ensure data integrity, accountability, transparency and trustworthiness of information exchange across various participants, including vehicle manufacturers, OEMs, and other stakeholders. It is used to store Trust Models (TMs) and Trust Policies (TP) that can be accessed by the In-Vehicle Manager in order to use them as input to the Trust Assessment process.
9. **TEE Guard Extensions:** The TEE Guard Extensions in the context of the CONNECT project refer to additional components or functionalities that run within Trusted Execution Environments (TEEs). A TEE serves as an isolated enclave within the computer system, whether it is the In-Vehicle computer or the Multi-Access Edge Computing (MEC) infrastructure. The TEE enhances security by providing both secure storage capabilities and a protected environment for executing cryptographic functions and managing cryptographic keys. The TEE Guard Extensions within CONNECT extend the capabilities of the TEE, encompassing features that contribute to the secure execution of specific cryptographic operations and the overall protection of data within the CCAM ecosystem. Hence, the following components are considered as the CONNECT TEE Guard Extensions:

   - ✓ **Attestation and Integrity Verification (AIV**): It serves as a critical element for ensuring the security and trustworthiness of the CCAM ecosystem. AIV is responsible for verifying the attestation evidence generated by different components, such as the in-vehicle sensors and actuators (i.e., Asymmetric-capable ECUs and/or Symmetric-capable ECUs, as defined in Section 6.5). It plays a crucial role in confirming the integrity of these components and their compliance with an expected output. By examining attestation reports, AIV contributes to the

overall trust assessment process, providing a layer of assurance regarding the authenticity and security posture of the involved entities within the CCAM architecture. This verification process enhances the system's ability to detect and respond to potential security threats, reinforcing the overall trustworthiness of communication and collaboration in connected and automated mobility scenarios. When attestation failures occur, the AIV sends the encrypted raw evidence to the Distributed Ledger Technology (DLT), in order to be subjected to further examination by the interested stakeholders (e.g., OEMs)

- ✓ **Identity and Authentication Management (IAM)**: It plays a pivotal role in securely onboarding devices to the In-Vehicle. This process involves integrating various components, including sensors, Zonal Controllers, and Electronic Control Units (ECUs), into the vehicle. The IAM establishes cryptographic keys and key restriction policies, ensuring a secure and controlled environment for device interactions. Furthermore, the IAM is responsible for communicating with an ETSI-enabled Public Key Infrastructure (PKI) for acquiring the necessary authentication tokens (and pseudonymous certificates) for supporting the vehicle's lifecycle as part of the CCAM ecosystem operation.
- ✓ **Trustworthiness Claim Handler (TCH):** It serves a critical function in the CONNECT architecture, acting as a key component responsible for harmonizing and processing information related to the AIV. In collaboration with the Trust Assessment Framework (TAF), the TCH incorporates attestation reports from the AIV, Misbehaviour Detection (MD) reports, and Trust Opinions (TOs) into a unified Verifiable Presentation (VP). This comprehensive presentation, containing harmonized attributes (expressed through Trustworthiness Claims), is a pivotal output that contributes to updating the Local Dynamic Map (LDM) based on Node Trust Level (NTL). Additionally, the TCH plays a crucial role in facilitating privacy-preserving exchanges, ensuring the secure that possible fingerprinting of the vehicle is not possible while ensuring the trustworthiness of the exchanged message.
- ✓ **Live Migration Management:** This component essentially enables the enforcement of a reaction strategy in the case of an observed decrease in the trustworthiness level of a vehicle resource. For instance, if a CCAM service (deployed in the Vehicle) is notified (by the TAF) of a decrease in the ATL of an ECU, then this serves as an indication of risk for the specific ECU. In this case, CONNECT facilitates the enforcement of flexible policies (overcoming the current hurdle of reaction strategies that disrupt the operation of a system/device) capable of migrating (in a secure and verifiable manner) the operational state of this ECU (definition of what constitutes state and the specific migratable parts of an application are service-dependent and defined by the OEM or the Service Provider) to a neighbouring ECU (or its Digital Twin on the MEC) with an ATL that characterizes a higher trustworthiness level from what is required. All these processes are managed by the Live migration Manager which governs this live migration between the secure enclaves (isolated environments hosting the secure and trusted execution of parts of the application to be migrated) as provided by the underlying CONNECT Trusted Computing Base.

10. **Key Management:** It includes the generation, distribution, storage, and revocation of cryptographic keys and the compilation of cryptographic algorithms during attestation tasks. The Key Management can support the interaction with two types of systems, depending on the device's in-vehicle topology capabilities. For devices supporting asymmetric cryptography, the Key Manager binds the keys to the corresponding Root of Trust. For devices lacking computational power, the CONNECT TCB uses a Hardware Security Module (HSM) with a pre-installed root-ID key.
11. **Attestation Agent:** It verifies the integrity and authenticity of entities within the network, specifically focusing on the attestation process for the A-ECUs (Application Execution Control Units). This component plays a crucial role in assessing the attestation evidence generated by various entities, inspecting for any indicators of compromise that affect the operational integrity. To enhance this (usually static attestation), a dedicated tracer is configured, tasked with gathering configuration traces essential for verifying the integrity of the system or component under evaluation. Further insights into the functionality and specifics of the tracer will be documented in the context of Deliverable D4.1 [111].

12. **Misbehaviour Detection**: The components included in the Misbehaviour Detection within CONNECT, following the notions as described in Section 4.2.1.1.

    - ✓ **Local Misbehaviour Detection Service (Local MDS):** It examines CAMs, CPMs, and other in-vehicle observations to identify any anomalies or inconsistencies that may identify potential security threats or system errors. Employing a variety of checks, including plausibility, consistency, and redundancy checks, it assesses the reliability of incoming data. The Local Misbehaviour Detection Service, inspects sensor data, comparing multiple observations of the same object, and detecting any aberrations. Any detected misbehaviour is reported to the TAF by representation of the MDS as a trust source and quantifying positive or negative trustworthiness evidence as trust opinions, and also sent as a VC to the TCH for inclusion in the trustworthiness VP that is sent outside of the vehicle with the data to which it refers. This contributes crucial information for the comprehensive evaluation and decision-making processes within the CCAM environment, ultimately enhancing the overall security and reliability of the system.
    - ✓ **Local Dynamic Map (LDM):** It stores and manages observations of physical objects with reference times over the last n seconds, extracted from a CAM, CPM messages received, or on-board perception of the ego. These observations are logically linked to the CAM or CPM message, and must contain an object identifier, position, and reference time. The Object Association Algorithm groups observations of the same physical object based on object identifiers and kinematic information. This algorithm runs periodically. The LDM ensures the source of the observation is identifiable through the PKI certificate.
    - ✓ **Misbehaviour Checks:** They are essential in the Local Misbehaviour Detection Service, scrutinizing incoming CAMs, CPMs, and Local Perceptions. These checks ensure adherence to physical rules, consistency and redundancy validations, as well as spatial alignment, enhancing the consistency and security of communication in the dynamic vehicular environment.
    - ✓ **Misbehaviour Reporting Service:** This service is responsible for generating and disseminating Misbehaviour Reports (MR) when suspicious behaviour is detected by the Local Misbehaviour Detection Service. The MRs serve as crucial notifications, promptly informing relevant entities within the CCAM ecosystem about potential anomalies.

13. **CCAM/Local Perception (LP):** It refers to the process of collecting and interpreting real-time data from various sensors, such as Lidar, Camera, and Acceleration sensors, within the vehicle's environment. The CONNECT Facility Layer aggregates data from these sensors and forwards it to the CAM/CPM Serializer. This raw data is utilized to generate the LP, a representation of the vehicle's surroundings. The Local Perception may include information about the vehicle's position, objects in proximity, and other relevant environmental details, gaining situational awareness. Subsequently, this Local Perception data can be incorporated into CAM/CPM to be shared with other connected entities in the CCAM ecosystem.

    - ✓ **VRU Protection:** It provides the service that ensures the safety of Vulnerable Road Users (VRUs) in CCAM ecosystem. The service requires observations from the sensors such as cameras, radar, and LIDAR. If there are any possible concerns, the component will activate responding activities such as notifications and safety systems.
    - ✓ **Trajectory/manoeuvre sharing:** It provides the service that facilitates communication and coordination among connected vehicles within the CCAM ecosystem. It allows vehicles to share their trajectories and manoeuvres with neighbouring vehicles and infrastructure elements, improving situational awareness and traffic flow.
    - ✓ **Coordination/Negotiation:** It enables connected vehicles to take decisions that optimise traffic flow, reduce accidents, and improve mobility efficiency. V2X communication enables vehicles to exchange data regarding their intentions, such as projected lane changes or route choices. Advanced algorithms and decision-making logic enable negotiation, guaranteeing fair and efficient cooperation in intricate traffic scenarios. This fosters a cohesive and collaborative transportation ecosystem.
    - ✓ **Intersection Movement Assistance (IMA):** It improves safety and efficiency at intersections IMA leverages V2X communication to facilitate vehicles in exchanging information that will

help them take well-informed decisions and carry out coordinated movements. IMA leverages information such as velocity, trajectory, and intentions of neighbouring vehicles, to offer immediate support. Advanced algorithms and predictive models enhance traffic management by mitigating congestion, minimizing collision hazards, and optimizing traffic flow, therefore promoting safer and more efficient junction interactions.

- ✓ **Advanced Cooperative Adaptive Cruise Control (CACC):** It optimizes vehicle platooning for improved traffic flow and fuel efficiency. It builds on traditional adaptive cruise control and uses real-time data exchange among platoon members to maintain closer following distances. This technology synchronizes vehicle movements, resulting in smoother driving patterns, reduced aerodynamic drag, and improved fuel efficiency. It also enhances safety by allowing rapid reaction to change in the lead vehicle's behaviour.

14. **V2X Communication Interface**: For the purposes of CONNECT, it facilitates seamless communication between vehicles (V2V) and between vehicles and the infrastructure (V2I). It uses V2X technology for safety, traffic conditions, and other relevant information. Leveraging Dedicated Short-Range Communication (DSRC) or Cellular Vehicle-to-Everything (C-V2X) technologies, this interface ensures low-latency and high-reliability communication. It plays a central role in enabling cooperative functionalities, such as collision avoidance, traffic optimization, and platooning.
15. **CAM/CPM Construction:** It enables vehicles to share information about their current state, surroundings, and perceptions with other connected entities. The process is dynamic, aggregating data from sensors like Lidar, cameras, and accelerometers. The collected data, accompanied also with the Trustworthiness Claims (TCs) of CONNECT, is processed into CAMs and CPMs, which are shared through the V2X Communication Interface. This process facilitates cooperative functionalities like real-time traffic awareness, safety enhancement, and advanced driver assistance systems.

6.2 Security & Trust Policies Enforcement: Runtime Risk Assessment

A stripped-down version of the CONNECT Risk Assessment (RA) architecture is depicted on Figure 17 highlighting the core activities of constructing and maintaining a risk quantification graph that further allows the CONNECT TMF to calculate the RTL per defined trust model: Per CCAM service deployed in the vehicle, and per process (deployed over the MEC virtualized infrastructure) for supporting the correct execution of a specific CCAM application. Furthermore, as previously described, the Risk Assessment process spans both during the Setup and Runtime Phases:

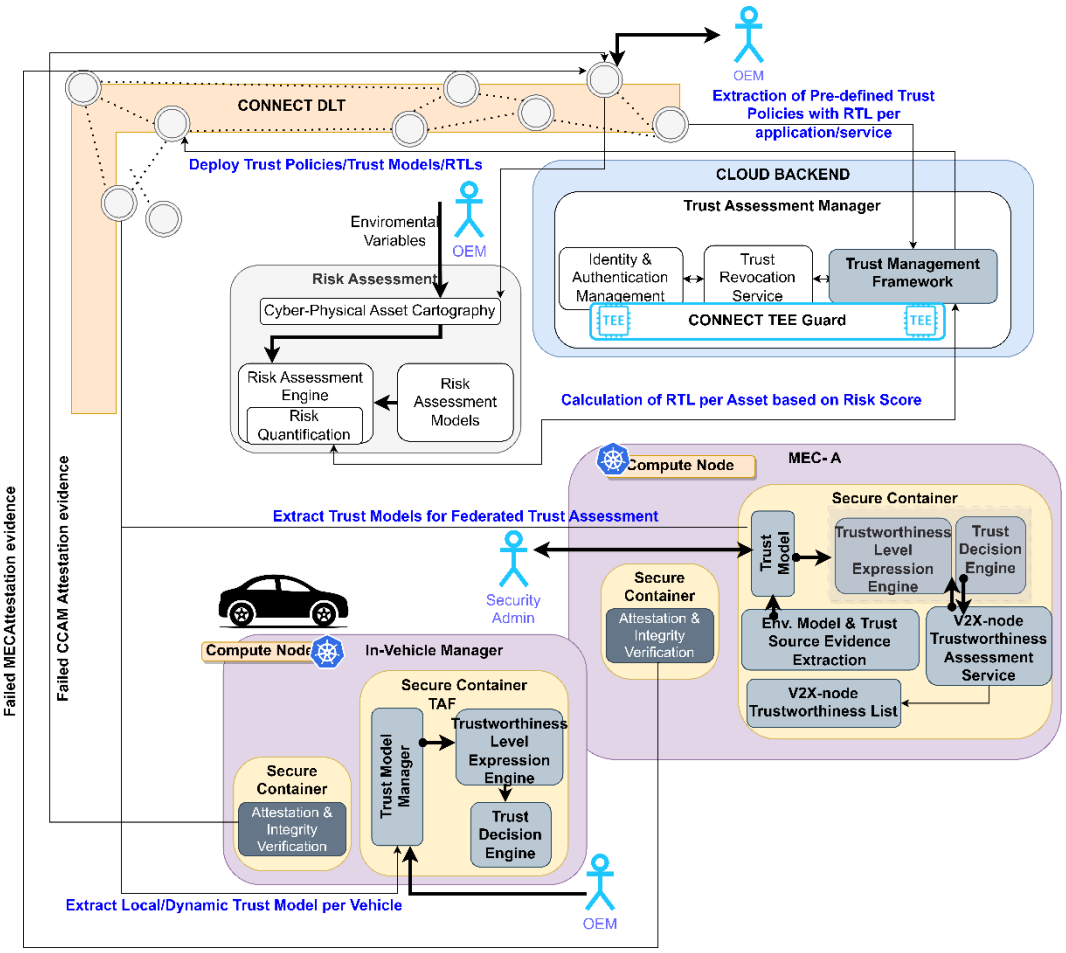

Figure 17 - Zoom-in Risk Assessment architecture

Especially for the latter, the runtime variant of the risk assessment framework will capitalize on an innovative methodology for performing risk calculations considering the relationships of the CCAM ecosystem and the comprising actors and SW and HW infrastructure assets based on well-defined risk models (converging on the benefits of multiple risk quantification models including the TARA for the quantification of automotive environment vulnerabilities and CVSS for enabling the risk analysis of MEC virtualized environments, both defined by ISO). This continuous risk quantification will also take into consideration any newly identified vulnerabilities (zero-day exploits) that may be the result of a detailed (supervised) learning of a system's behaviour based on the real-time extracted low-level system traces from the CONNECT Attestation Enablers (Section 6.4.3). Essentially, in the case of a failed attestation result (serving as an indication of risk), the monitored attestation evidence are shared (through the Blockchain infrastructure) with authorized users, administrators and OEMs that can leverage them for further identifying the exact point of intrusion.

After analysing such possible safety-critical events, optimal security and trust policies/strategies can be recommended. This, for instance, can lead to the dynamic update of the RTL for a specific trust model resulting into a more strict trust assessment process taking place for appraising the trustworthiness profile of a system. Or it could lead to the deployment and enforcement of optimal mitigation actions through the compilation of attestation security policies and/or task migration policies. Overall, through this setup, this continuous risk assessment will enable CONNECT framework to always maintain an up-to-date view of all identified threats and better evaluate the identified mitigation actions and consider their impact on the defined RTL.

As depicted, the RA component may reside anywhere, including locally to the OEM or MNO side side, while it has a direct communication with the CONNECT cloud where the Trust Assessment Manager (TAM) component resides. These two components enable the estimation of the RTL per asset based on the risk score. The trust policies, models and RTLs are uploaded to the Distributed

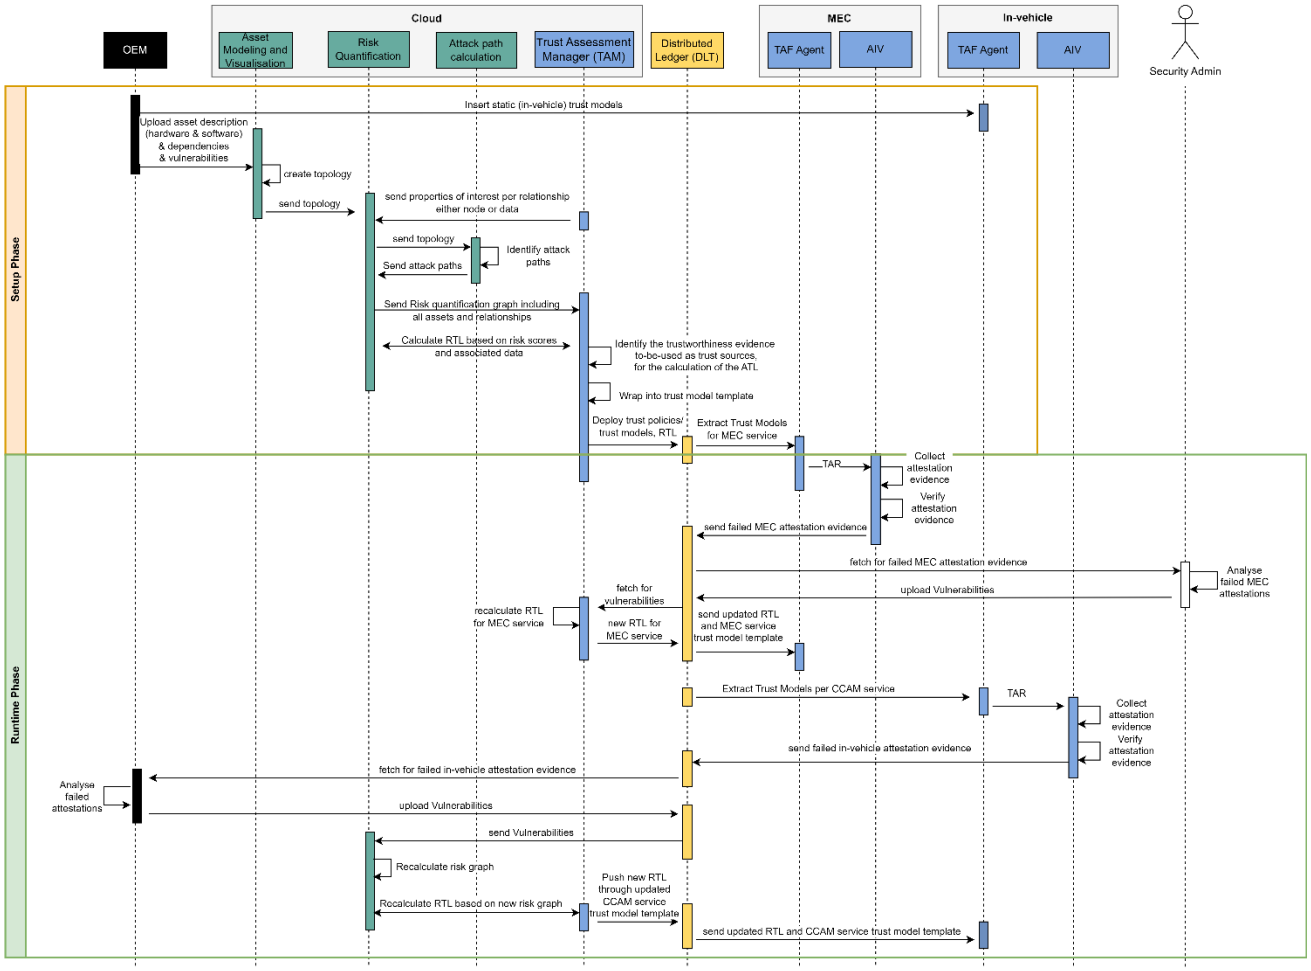

Figure 18 - Risk Assessment flow of actions

Ledger (DLT) which is accessible both be the vehicle and the MEC. These comprise of the dynamic trust models of the CONNECT platform. More details on the Risk Assessment component and its internal architecture will become available at Deliverable D3.2 [112].

Despite the in-vehicle topology remaining static post the setup phase, the evolving threat landscape may lead to the discovery of new vulnerabilities. Consequently, the RTL may require periodic updates over time to ensure its ongoing relevance and effectiveness in adapting to emerging security challenges. Hence, two phases are relevant here: the setup and the runtime, which are elaborated in the following paragraphs.

Setup Phase: The setup phase is initiated by the Original Equipment Manufacturer (OEM), which is the firm constructing the vehicle. The OEM inserts the static (in-vehicle) trust models (TMs) and the RTL, to the TAF Agent residing at the vehicle. Similarly, the OEM uploads details about all the parts of the car—both hardware and software, their interdependencies as well as vulnerabilities. This information is sent to the Asset Modelling and Visualization tool that creates a detailed topology considering both in-vehicle assets but also vehicle to vehicle. This graph is then shared with the Risk Quantification component.

Specific details about the properties of interest per relationship (i.e., either node or data) are sent by the Trust Assessment Manager to the Risk Quantification. Furthermore, to cover the possible paths that an attacker may follow in order to launch an attack, attack path calculation is also a part of the assessment. To achieve that, the Risk Quantification tool shares the topology with the Attack Path Calculation tool, which figures out potential ways an attacker could target the vehicle (attack paths). This information is sent back to the Risk Quantification.

The latter now has all the information needed to perform the risk assessment for the given topology, considering also the possible attack paths. Towards this direction, the Risk Quantification component leverages the TARA (Threat and Risk Assessment) methodology, to calculate the risks based on specific properties of interest, as identified by the Trust Assessment Manager (TAM). The result is a

graph that demonstrates the risks for all the components (thus nodes) along with their interconnections. This graph, along with the calculated risk scores are shared with the TAM, to proceed with the calculation of the Required Trust Level (RTL).

The TAM leverages the RTL to deduce the information needed to evaluate the trustworthiness of vehicle against potential attacks. This information refers to the trustworthiness evidence, that are used as trust sources for the calculation of the ATL. This evidence is leveraged prove their integrity and authenticity of specific nodes of the graph. The identified trustworthiness evidence is placed into a trust model template, that is sent to the Distributed Ledger Technology (DLT), along with the RTLs. The DLT can be accessed by the Trust Assessment Framework (TAF) agent (i.e., both at the MEC and the vehicle side). These agents may extract the trust models either referring to the MEC services (i.e., MEC-based TAF agent) or to the CCAM services (i.e., In-vehicle based TAF agent).

Runtime Phase: After receiving the TMs both the MEC and the vehicle TAF agent may send the Trust Assessment Request (TAR) to the Attestation and Integrity Verification (AIV) component, to initiate the collection and verification of evidence.

Since the steps for the failed attestation evidence in the case of the MEC and the vehicle differentiate, we will start our description with the MEC and continue with the flow of actions for the vehicle. Whenever a failed attestation evidence for the MEC is received, it is being uploaded to the DLT. The DLT is accessible by a security administrator, which analysed the failed MEC attestation evidence, identified vulnerabilities and sends them back to the DLT, where they can be accessed by the TAM. The TAM re-evaluates the RTL and sends it back to the DLT, to be consumed by the MEC-based TAF agent.

Similarly, the in-vehicle AIV, receives the TAR from the in-vehicle TAF Agent and starts the collection and verification of attestation evidence. Whenever the attestation fails, the relevant evidence is being uploaded to the DLT so that the OEM may further investigate it and report back to the DLT new vulnerabilities. The Risk Quantification tool may access information regarding the identified vulnerabilities through the DLT. Hence, it may recalculate the risk graph and the updated RTL(s), considering the discovered vulnerabilities. The new RTL is being uploaded to the DLT, through the updated CCAM service TM template, to be consumed by the In-vehicle TAF agent. The flow described above is illustrated on Figure 18.

6.3 Trust Assessment Framework (TAF)

Continuing the description of CONNECT's core building blocks, Figure 19 illustrates a simplified version of the main technical artefact responsible for providing the continuous trust assessment functionalities of the overall framework, namely the Trust Assessment Framework (TAF). For showcasing the underpinnings and internal mode of operation, we have opted for focusing on the TAF variant that is deployed as part of the Vehicle's main computing unit. This is subject to secure execution safeguarded by CONNECT Trusted Execution Environment for providing the necessary guarantees and assurances on the verifiable output of the TAF when calculating Trust Opinions (TOs). However, as aforementioned in Section 6.1.1 (and further elaborated in D3.1 [106]), CONNECT employs a federated trust assessment architecture where multiple TAF agents – deployed in all Vehicles but also the MEC – collaborate for providing higher accuracy in the CCAM-wide trust evaluation and quantification. The internal architecture of the MEC-instantiated TAF follows a similar structure, hence, it is omitted from the following figure.

The TAF process is integral to determining the Actual Trust Level (ATL) and subsequently making trust-related decisions. It operates with both static trust models (predefined within the vehicle capturing topologies of in-vehicle sensors and actuators for providing kinematic and other V2X-related data as needed by various CCAM services) and dynamic trust models accessed directly the Distributed Ledger. This dynamic approach, characteristic of CONNECT TAF, incorporates inputs from multiple sources for ATL calculation. Internal sources include i) the output of the Misbehaviour Detection Service and ii) the result of the attestation process. Notably, the framework is extensible, allowing for the incorporation of additional sources as needed. More details on the TAF component and its internal architecture are available at D3.1 [106], while more elaborated descriptions will be provided in D3.2 [112].

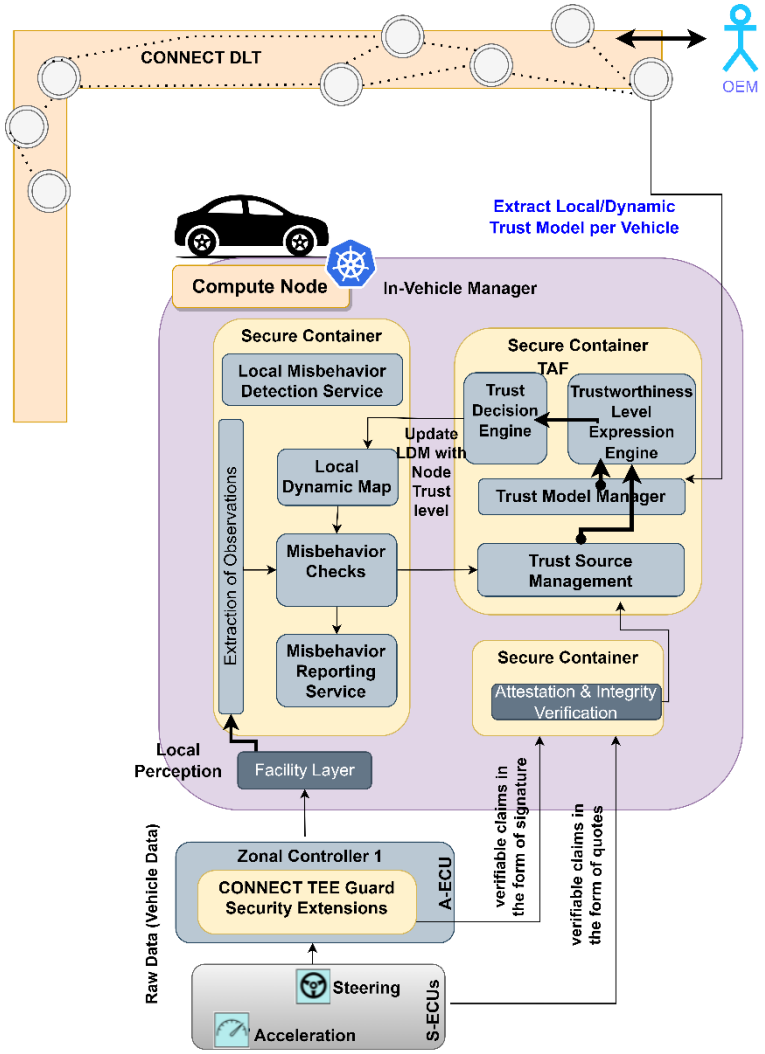

Figure 19 - Zoom-in (in-vehicle) Trust Assessment Framework architecture

In the Figure 20, the interactions between the in-vehicle TAF and the rest of the system during two phases of the vehicle lifetime: Setup Phase, and Runtime Phase, are depicted through a sequence diagram. The definition of the aforementioned phases is elaborated in Section 6.1.

Setup Phase: There are two types of trust models that the TAF uses: **static and dynamic trust models**. Both types are application-specific; however, static trust models do not change during runtime, and dynamic trust models do. Static trust models are created at design time for in-vehicle applications based on in-vehicle component diagrams and application-specific data flows. Dynamic trust models are created during run-time and based on trust model templates fetched from the CONNECT DLT. The only event happening during the deployment phase is loading the in-vehicle static trust models by the OEM onto the TAF. Each model has an ID associated with it which allows the TAF to distinguish between them and later call upon the appropriate trust model.

The first time an application is run, the application sends a Setup Request to the TAF. The goal of the Setup Request is to set up the necessary items for the TAF to start the trustworthiness

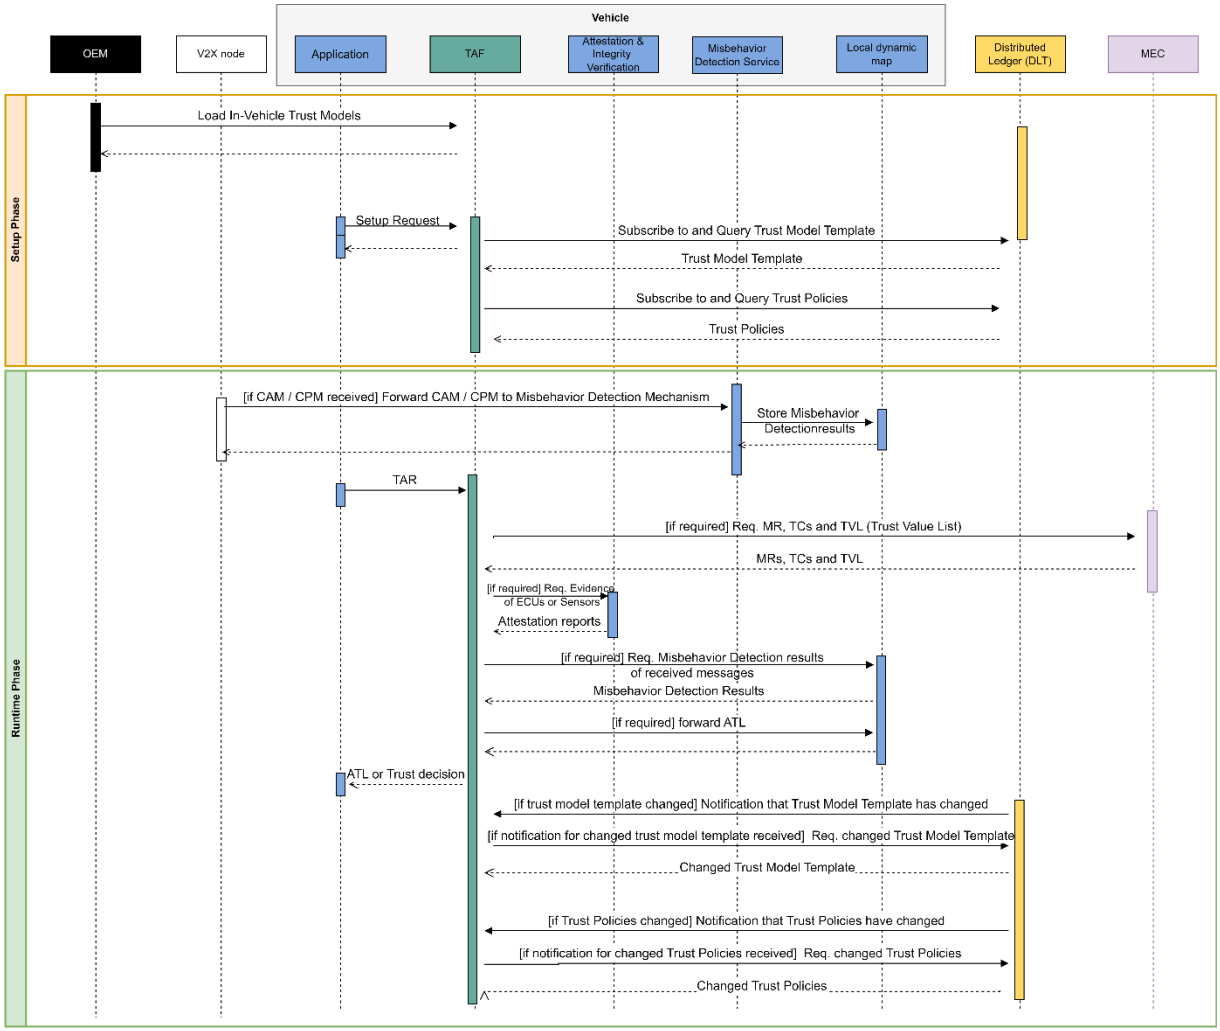

Figure 20 - Trust Assessment flow of actions

assessment for the application. The Setup Request contains the application ID, which allows the TAF to either locate a pre-stored static trust model or to send a Trust Model Template Subscription Request to the CONNECT DLT, where the application-specific templates are stored. A query to get the latest trust model template is automatically sent along with the Subscription Request. As a result of the query, a Trust Model Template is sent by the CONNECT DLT to the in-vehicle TAF. The CONNECT DLT has also registered the vehicle's subscription request and will inform the vehicle of any changes to the template during the Runtime Phase. Finally, the TAF sends an application-specific Trust Policy Subscription Request along with a query for the latest Trust Policy to the CONNECT DLT. Similarly, to the Trust Model Template Request, as a result of the query, the current Trust Policy is sent to the TAF. The CONNECT DLT also registers the vehicle's Trust Policy Subscription Request and will notify the vehicle during the Runtime Phase if the Trust Policy has changed. The Trust Policy contains the Required Trustworthiness Level (RTL) for a specific application, which will be used during runtime to compare to the Actual Trustworthiness Level (ATL).

Runtime Phase: Several critical processes are involved in the trustworthiness assessment during the vehicle Runtime Phase. However, the Misbehaviour Detection process occurs independently of whether the trustworthiness assessment is running. Misbehaviour Detection takes place every time a new CAM or CPM is received. A new CAM/CPM is forwarded to the Misbehaviour Detection Service, and the results of the Misbehaviour Detection Checks performed are then stored inside the Local Dynamic Map. An application sends a Trustworthiness Assessment Request (TAR) to the TAF when it wants to run a safety-critical function that uses specific data as input. The TAR triggers the process of trustworthiness assessment of this data. It contains a parameter which indicates if the **ATL should only be sent once (synchronised TAR), sent periodically (periodic TAR), or sent whenever the ATL changes (event-based TAR)**. A periodic TAR could be helpful in situations where the trustworthiness of the trust object changes frequently because new evidence is constantly

arriving. For example, this could be helpful for the trustworthiness of an observation in the Intersection Movement Assist use case. Since messages are constantly arriving here, the trustworthiness of past observations could change frequently. On the other hand, an event-based TAR could be helpful in situations where the trustworthiness of a trust object is not expected to change very often, but a fast reaction is needed when it does. An example of this could be the trustworthiness of an ECU in the C-ACC use case. If an ECU is no longer trustworthy, the application needs to be moved quickly to another ECU. Once the TAF receives a Trustworthiness Assessment Request (TAR) sent by the application, the trustworthiness assessment begins. The last two TAR types, periodic and event-based, result in TAF continually providing an ATL to the application during application run-time, essentially running a monitoring service.

The TAF already has access to the application-specific trust model (template). Based on the trust model (template), the TAF decides which evidence needs to be collected. If needed, the TAF requests the Misbehaviour Reports, Trustworthiness Claims, and the Trust Value List from the MEC, which delivers them to the TAF. The Misbehaviour Reports contain observations and results of the Misbehaviour Detection Service that indicate vehicle misbehaviour. The Trustworthiness Claims attest that the hardware and software of the MEC are in a correct configuration state. Finally, the Trust Value List contains the active V2X nodes and their most recent trustworthiness levels determined by the MEC. Moreover, the TAF can also request evidence about the internal ECUs and sensors from the AIV, which then provides the Attestation Reports on the relevant ECUs and sensors to the TAF. Finally, the TAF can also request results of the Misbehaviour Detection on the received CAMs/CPMs from the Misbehaviour Detection Mechanism, which forwards these to the TAF upon request.

Once the TAF collects all the evidence needed, it can perform a trustworthiness assessment on the relevant data and produce a set of ATLs. The ATLs are then either forwarded to the application which sent the TAR or compared to the RTLs to produce a set of Trust Decisions to be sent to the application. The ATLs can also be stored inside the Local Dynamic Map if needed.

A trust model template and trust policies update process can also occur if the TAF is notified by the CONNECT DLT that there have been changes. In this case, the TAF queries the CONNECT DLT for the updated version of the template and the policies, which then forwards the updates to the TAF. The aforementioned flow is illustrated in Figure 34.

6.4 Trustworthy Platform Configuration & Attestation

Figure 20 presents a simplified illustration of the interactions with the CONNECT Trusted Computing Base (TCB) offering the new set of (runtime) attestation capabilities for providing the necessary trustworthiness evidence based on which a recent trustworthiness appraisal can occur from the CONNECT TAF. Such evidence (produced in the form of **verifiable security claims**) are dictated by the deployed trust models and are mapped to the trust properties of interest based on which an evaluation (for a node or piece of data) will take place. For instance, such marketed security claims might focus on the verification of: (i) the secure boot capabilities of a device with known hardware and firmware as an indicator of **design time integrity**, (ii) the vendor of a system/device (e.g., originating from a specific OEM) and are running known software containing the latest patches as an indicator of **configuration correctness**, (iii) the runtime configuration integrity of the loaded software stack (and Operating System) as an indicator of **runtime integrity**, (iv) the software stack loaded in the device so as to make sure that only certified application is been running (with the appropriate certificates checked and attested by the OEM) as an indicator of the **safety level of the system**, etc. CONNECT TCB is equipped with the Attestation and Integrity Verification (AIV) component (Section 6.4.3) that governs all the TEE Device Interfaces[11] capable of monitoring such type of runtime evidence.

[11] IETF RATS, "TEE Device Interface Security Protocol", 2022, [Available Online: https://members.pcisig.com/wg/PCI-SIG/document/18268?uploaded=1]

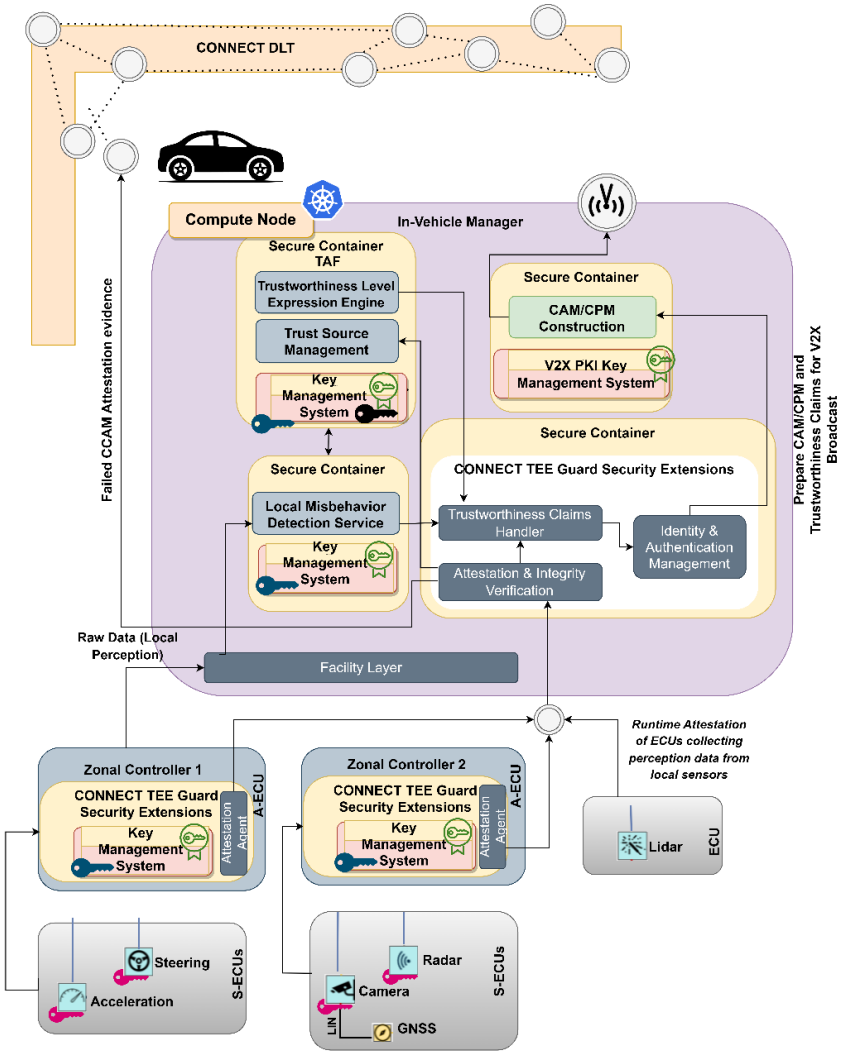

Figure 21 - Zoom-in CONNECT Trustworthy Platform Configuration and Attestation architecture

The AIV resides within the boundaries of a secure enclave, deployed by a secure container. More information on the exact technologies is available at the following paragraphs. Trust-related information regarding the vehicle's trustworthiness can be shared with the external CCAM environment through advanced encryption techniques. These techniques ensure a time-efficient and cost-effective protection of trust-related information, preserving the privacy of the vehicle by concealing its identity. Another point in this overall approach is that failed attestation evidence is being uploaded to the Distributed Ledger (DLT) so that it can by further scrutinized by the OEMs, with the intent of making the entirety of the CCAM landscape more resilient against attacks. This facilitates further scrutiny by Original Equipment Manufacturers (OEMs), enhancing the overall resilience of the CCAM landscape against potential attacks.

The runtime attestation capabilities of CONNECT are crucial for achieving trustworthy vehicular communication and data exchanges. This means that ensuring the continuous verification of the integrity and security state of various components within the system is a priority. This verification is conducted during operation, with regular checks on the integrity of software, hardware, and communication channels. The AIV component, as described in section 6.1.2 carries out this verification. It receives evidence from the nodes, utilising pre-configured Tracer, which it verifies in terms of integrity. The result of the attestation is included in an Attestation Report. More information regarding the tracer and the attestation capabilities will be reported in D4.1.

The security of any platform relies on the effective operation of both its software and hardware components. This chapter explores the Root of Trust (RoT) concept for CONNECT, which acts as

the foundation for collecting evidence from In-Vehicle and MEC environments, thus enabling the execution of critical operations. A RoT offers vital security services that are considered resilient against compromise by malicious software. Examples services include:

16. **Cryptographic operations** utilizing keys securely stored in hardware
17. **Attestation services** that provide verifiable evidence of executed software
18. **Trusted execution** that protects specific software components from unauthorized modification by other software (usually including the operating system)

The Trusted Execution Environment (TEE) is a well-established component in contemporary system architectures, acting as a RoT and offering a safe setting for important operations such as cryptographic calculations. Using TEE provides separate environments for "trusted" and "untrusted" applications, protecting “trusted “applications from unauthorised modifications or leakage of sensitive data. TEEs collaborate with hardware components such as the CPU and Hardware Security Module (HSM) to guarantee the protection of applications running within the TEE.

It is important to mention that CONNECT, as a framework, remains agnostic towards the particular RoT that a platform provides, which allows for various implementations and flexibility. However, as a proof of concept, a hardware RoT is leveraged as an example for the demonstrators. The prototype development primarily focuses on using Intel's Software Guard Extensions (SGX). SGX provides a flexible trusted execution environment for user-space programs, isolating the “trusted” and “untrusted” applications and data. It allows user-level code to allocate exclusive memory locations, known as enclaves, to protect against programs operating at higher privilege levels. Enclaves can autonomously produce and store their own exclusive signing/attestation keys, preventing external access. Data can be signed using keys linked to specific instructions operating within each enclave.

The following paragraphs elaborate on the In-Vehicle and MEC-based TEEs, utilising the example of the SGX as a RoT technology. As explained in Section 2.2.7, the overall goal of CONNECT is to establish Chip-to-Cloud assurance. Therefore, the provision of secure, TEE-based, enclaves in both the In-Vehicle side as well as the MEC is critical to achieve this goal.

6.4.1 Secure Elements (SE) in CONNECT

The Secure Element (SE) in CONNECT is in essence the TEE, which is a well-protected enclave designed to execute applications with an elevated level of security. Since the vision of CONNECT is to be agnostic to the exact technology employed, the investigation is focused on SE for two distinct settings: i) the in-vehicle environment and ii) the MEC infrastructure. In both contexts, for our example of an implementation for the employed SE Intel's SGX hardware is leveraged for the sake of the implementation, that acts as a TEE. The ability of TEE to be used in both the in-vehicle and MEC domains highlights its versatility and flexibility. The inclusion of a SE in both the vehicles and the MEC further enables the Confidential Computing paradigm that CONNECT adheres to.

In addition, delving more into the in-vehicle side, ECUs may have diverse capabilities in terms of cryptographic support. More specifically, an A-ECU offers additional functionalities than an S-ECU, which supports solely symmetric cryptographic capabilities. Hence this is further considered in the CONNECT architecture, where different schemes are envisaged for the two options, as well as hardware (i.e., TEE for A-ECU and HSM for S-ECU). The following paragraphs further elaborate on the TEE leveraged as a means of protection in the in-vehicle and the MEC domain.

6.4.1.1 In-Vehicle Secure Elements using Trusted Execution Environments (TEEs)

For in-vehicle SEs, the focus is on ease of deployment and ease of use. While the underlying hardware extension (i.e., Intel SGX) provides protection for user-space processes, these processes must be self-contained, since the operating system might be compromised, thereby necessitating a reliance on self-contained execution units. This self-contained execution units are referred to the secure enclaves. Figure 22 illustrates this isolation provided by the enclaves in terms of the OS. On the implementation side, for ease of deployment, a library OS called “Gramine” is leveraged, that executes a linux-style application inside an Intel SGX enclave - with no security-critical dependency on the operating system.

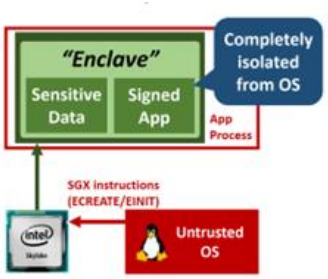

Figure 22 - Isolated from untrusted OS Enclaves (i.e., Intel SGX)

The deployment of service in an Intel SGX-protected environment includes multiple components:

1. The **application** (i.e., the TAF) to be protected inside a TEE enclave (i.e., Intel SGX enclave).
2. The **Gramine library-OS**, that provides operating-system services that are also protected by a TEE (i.e., an Intel SGX enclave).
3. A **signed manifest** where the TEE developer (i.e., Intel SGX enclave) specifies the correct software and configuration of the enclave, including the original application, its dependencies, the Gramine library OS, and any required configuration. When such a package is then installed and executed within an Intel SGX enclave, the hardware verifies this signature and installs the software within an Intel SGX enclave if and only if the signature in the manifest corresponds with the package to be installed.

Regarding the use cases, practically the execution of most applications within a TEE-based (i.e., SGX) environment is feasible. However, this would affect the cost in terms of performance as well as development complexity. Hence, in reality this layer of protection is reserved only for critical components that adhere to the specific criteria. For instance, i) components that a possible compromise could lead to unrecoverable or undetectable system failure, or ii) components that external stakeholders need to validate their integrity, or iii) components that need to enforce security policies defined by external stakeholders, are all included in the critical category hence they should be executed within a TEE.

Please note that critical services may diverge from safety-critical services. While the first category would benefit from advanced security, the second could be negatively affected. For example, in safety-critical scenarios where decisions must be taken instantly to protect passengers, the latency and overhead introduced by running the application in a TEE could seriously delay the process. Thus safety-critical services are excluded from this definition.

Applying these criteria to the Cooperative Adaptive Cruise Control (C-ACC) use case that is defined in Section 7.3 of the present deliverable, then the following components can profit from enhanced hardware-backed security:

- ✓ **Central Security Services**: Since these services protect cryptographic keys, external stakeholders often depend on the requirement that these keys securely identify a given device. As a consequence, hardware-protection can ensure that keys cannot be copied even if the platform has been compromised.
- ✓ **Secure variant coding:** The manufacturer needs to be sure that only licensed features can be used. Hardware protection ensures that this is protected even if some parts of the software have been hacked.
- ✓ **Secure diagnostics:** Diagnostics include logging of potential attack traces (e.g., failed signature verification counter for firmware updates). Therefore, to reliably detect intrusions, it is important that attackers cannot overcome diagnostics.

Basic Gramine Functionality

The basic functionality of Gramine is to allow hardware-protected execution of most Linux-style binaries. Utilizing Intel SGX, the binary is executed within the enclave, which guarantees a compromised software outside the enclave cannot attack/tamper/compromise the binary, even if an attacker has managed to get root access to the host operating system. However, a constraint of this approach is that the operating system is outside of the "trust boundary" of the binary. Consequently, the binary should not rely on services provided by the OS. Instead, Gramine implements system

calls inside the enclave, augmenting them with additional security controls. For instance, when the protected binary performs file read and write operations, Gramine seamlessly introduces encryption and integrity protection to mitigate potential compromise by the untrusted operating system.

Cybersecurity regulation mandates maintaining cybersecurity maintenance of vehicles during their full product lifecycle. This includes responding to new threats resulting from discovering relevant vulnerabilities or unforeseen technical advancements by attackers. For a mid- and long-term response, components need to be updated or even upgraded, for a short-term or immediate response, systems should be able to maintain a secure and safe operational state. These requirements also include components running in secure enclaves. Gramine can help implementing appropriate response strategies with the following extensions:

Extension 1: Secure Upgrade

Gramine, in conjunction with Intel SGX, protects the integrity of the software residing within the enclave, leveraging a manifest file that contains a signature on the specific enclave executable. This mechanism ensures that any modifications to the software/executable prevent the enclave from initiating, thereby preserving restrict access to the data residing within the enclave.

In the context of a vehicular ECU, this requirement would mean that key software packages that are to be deployed inside a hardware-protected enclave, must be signed by the manufacturer. At start time (i.e., deployment phase), only those packages that successfully validate against the manufacturer signature are allowed to commence execution. If the package has been modified, the startup of this component will fail. In addition, such components can only access their storage (e.g., the secret keys of the Authentication and Access Management) once the signature has been verified and the integrity of the component has been ascertained.

CONNECT aims to introduce a secure software stack *upgrade* and *mitigation* feature in Gramine, ensuring a smooth transition of state from the old to the new version with minimal downtime. This process adheres to the defined Service Level Agreement (SLA) established during service deployment to prevent any violations. The feature allows i) the introduction of new features and ii) the mitigation of potential security vulnerabilities while preserving the component's state. Currently Intel SGX currently supports enclave dumping and restoration to maintain state over reboots, similar to the hibernation and restoration process of a PC. To this day, hardware support for secure upgrades and mitigation is not provided as this would depend on the trust decisions of stakeholders and must adhere to any Service Level Agreements (SLAs) to avoid disrupting the normal execution of the target system. Implementing a secure upgrade involves several phases:

1. **Build a New Updated Enclave**
    a. Build a new enclave with the desired software stack. Create a manifest for this new stack and start the updated enclave.
2. **State Migration from Outdated to Updated Enclave**
    a. The outdated enclave verifies that the updated enclave is indeed authorized to receive its state (e.g., it can verify that it has a higher version number and was signed by an authorized party).
    b. Migrate the state from the outdated to the new enclave.
    c. The migration protocol needs to ensure that (a) the state can only be migrated to an updated enclave that is an authorized update of the outdated enclave and (b) that the state is protected during migration (e.g., using end-to-end encryption and authentication).
3. **Disable Outdated enclave and Enable Updated Enclave**
    a. Atomically disable the outdated enclave and enable the updated enclave.
    b. Perform any required cleanup operations (e.g., some keys may be replaced by new keys).

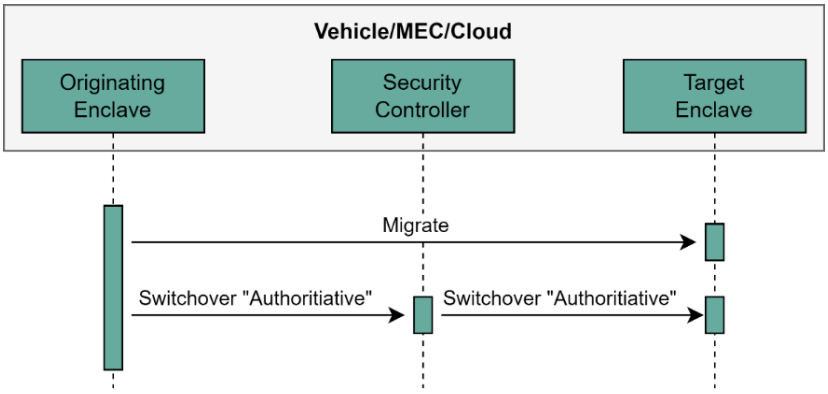

Figure 23 - Authoritative State Migration

The reliability and integrity of ECUs is key for maintaining security in autonomous driving scenarios, where vulnerabilities can significantly impact safety. Therefore, a mechanism for securely upgrading ECUs is essential to promptly address any identified issues, improving the overall functionality of the system. To implement this update feature, multiple extensions need to be designed:

- ✓ **Ability to disable old enclave**: The capability of enclaves to be snapshotted at any time introduces a potential security vulnerability, as prior snapshots can be restored without rollback protection. In the absence of mechanisms to prevent this, even if an enclave has marked itself as "disabled," the untrusted operating system can revert to a previous version that hasn't been flagged as "disabled." This scenario poses a risk of a "version downgrade" attack, where an adversary could intentionally revert to an older software version with known vulnerabilities, creating opportunities for subsequent attacks. Mitigating such risks is crucial to ensure the security and integrity of enclave-based systems.
- ✓ **Ability to securely export/import/migrate application-level state of Gramine:** Intel SGX only supports a dump and unmodified restore which poses a challenge when it comes to exporting the application-level state within Gramine. The key issue is that exporting the state might inadvertently expose critical long-term keys or sensitive information to the untrusted filesystem of the operating system. For example, a disclosure attack against an ECU could severely affect the security of all services executed. To prevent such leakage during migration, an end-to-end protection scheme is required for ensuring that only authorised enclaves can load migrated/exported state.
- ✓ **Downgrade protection:** To prevent downgrade attacks, it is important that upgrade can only be done towards authorised higher versions. Without such protection, attackers could downgrade software to versions with known vulnerabilities that can then be exploited to leak state/secrets. One way to achieve this is to use a security controller as depicted in the following picture (i.e., Figure 23). The security controller that maintains keys and also tracks what enclave is the current authoritative version (i.e., this security controller can basically blacklist the outdated enclave).

Figure 24 - Enclave-CC framework flow

Extension 2: Secure Migration

Upgrades are provided locally and aim to enforce the security requirements listed above without any dependency on additional services. Unlike upgrades, migration aims to move an enclave service from one machine to another. The goal is e.g., to offload some security-critical services or (other application-related) calculations from the vehicle to the MEC to simplify communication or increase performance. Examples include collaborative intersection management or the TAF. In principle, we plan to enable two flavours - depending on the services provided by an enclave:

- **Remote Clone:** An enclave is cloned from one machine to another. Therefore, after remote cloning, two identical enclaves then co-exist on two machines and continue to execute in parallel. One example application are multiple replicas of a digital twin.
- **Migration:** Which is similar to remote cloning, except that the protocols guarantee that the originating enclave is no longer active.

These services may require an additional trusted third party that we named “Security Controller”. The services required are similar to upgrades and needs to implement the following phases:

1. **Target Enclave**
    a. Build a new enclave with the desired software stack. Create a manifest for this new stack and start the enclave.
    b. Register the intent to import state from a source enclave.
2. **Source Enclave:**
    a. Register the intent to migrate the enclave to the specific updated enclave.
3. **Both: State Migration and (if migrate) Switchover**
    a. Migrate the state from the outdated to the new enclave.
    b. If migration: Atomically disable the outdated enclave and enable the updated enclave.
    c. Perform any required cleanup operations (e.g., some keys may be replaced by new keys).

6.4.1.2 MEC Secure Elements using Secure Containers

As mentioned in Section 2.3, CONNECT aims at reaching Chip-to-Cloud assurance, following the notions of Confidential Computing. To achieve this, apart from the secure and isolated enclave for critical operations offered by the TEE (i.e., Intel SGX), ensuring the security of the containers is also a requirement for the MEC and the cloud-based infrastructure. In the context of the MEC, and more specifically the secure containers that are envisaged to be instantiated there, the RoT is extending its domain of applicability to establish trust in the entirety of the computing stack. Hence, this MEC-based RoT paradigm incorporates the protection of both the hardware and software that run in isolated enclaves, supported by the TEE (i.e., Intel SGX) but it further extends this protection to cover secure containers. The scope is to protect the contents of the container(s), including applications, files, and memory, from potential threats originating outside the container boundaries. By combining the isolated enclaves (i.e., of Intel SGX) with the secure containers, CONNECT achieves a RoT that manages to guarantee an elevated security posture for virtualized platforms and consequently, contribute to the vision of CONNECT for Chip-to-Cloud assurance.

Leveraging Gramine, which enables the execution of most Linux applications inside an Intel SGX enclave, CONNECT embraces containers to package larger multi-component applications; thus, simplify deployment and management. The core idea for using containers is to allow standard packaging of composed applications into a Docker container. The Enclave-CC framework

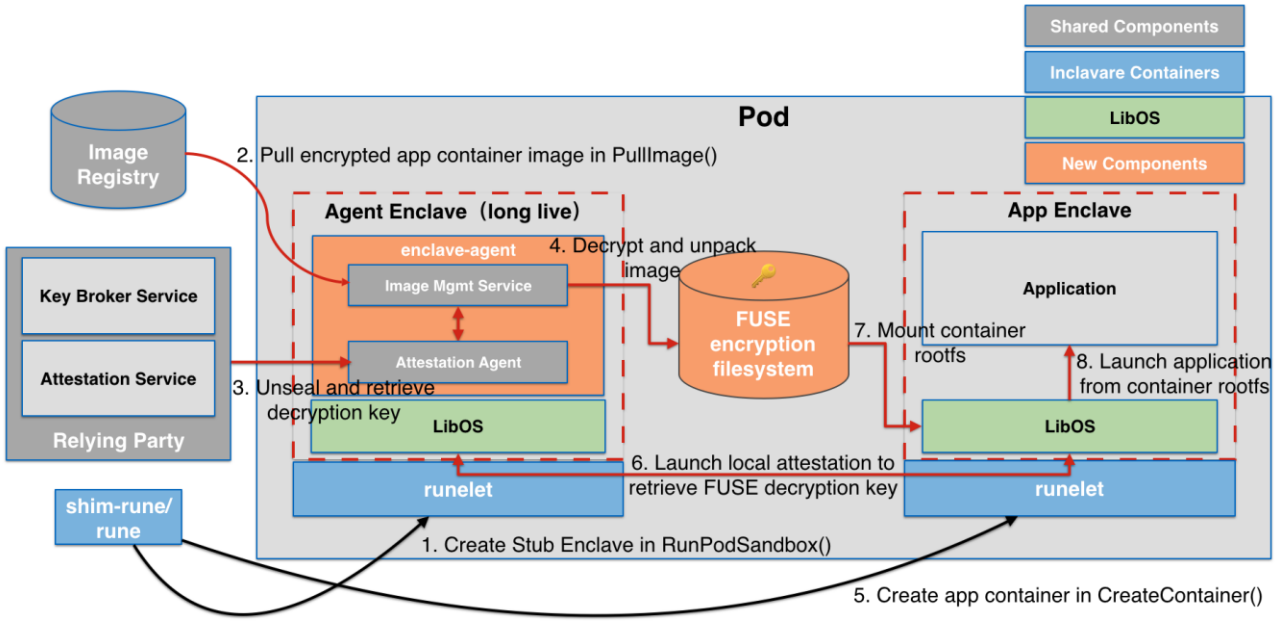

Figure 25 - Architecture of the Enclave-cc Confidential Container [113]

seamlessly orchestrates the execution of Docker containers within an Intel SGX enclave, thus offers protection against adversaries who may have compromised elements outside the container confines. This approach ensures that security measures are extended cohesively from the chip-level enclave to the encapsulation provided by containers, reinforcing the goal of achieving chip-to-cloud security in the CONNECT architecture. Since the Enclave-CC offers an open-source solution regarding Confidential Containers, it is eligible for consideration and extension by CONNECT [114]. It starts with a default image that is encrypted and signed to protect integrity and confidentiality. The image is then pushed to the container registry where it can be pulled and provisioned onto an Intel-SGX-enabled platform. The aforementioned flow is illustrated on Figure 24.

The current architecture of enclave-cc is illustrated on Figure 25. The services are provided by a set of secure containers (i.e., Pod) that jointly ensure protection of the application container:

- ✓ A long running “agent enclave” provides protected container services such as key management, attestation, and image management.
- ✓ An App Enclave then executes a given container while transparently adding protection by e.g., encrypting/authenticating the file system.

The goal of this architecture is to give a stakeholder (e.g., a vehicle) full control over a container that runs in a cloud (e.g., or MEC). The establishment of a running secure container that contains a Gramine Library OS and an application comprises the following depicted steps:

1. The first step is to create the long-lived Agent enclave out of a stub enclave.
2. The next step is to pull an encrypted container image from the registry.
3. The stakeholder “unlocks” the enclave by unsealing a decryption key that can decrypt the encrypted container.
4. With the permission of the stakeholder, the image is decrypted and unpacked into a FUSE encrypted file system.
5. A new application container is created.
6. The agent enclave verifies the integrity of the new container using local attestation and releases the FUSE decryption key if any only if the integrity verification was successful.
7. Using this key, the FUSE file system is mounted and auto decrypted.
8. The application inside the container is launched - accessing files on the transparently decrypted filesystem.

Note that the “agent enclave” provides a generic infrastructure for creating, starting, stopping, storing, restoring, restarting containers. As a consequence, any services required by CONNECT will be implemented as specific containers that are then executed in one or more “app enclaves”.

6.4.2 CONNECT Attestation Technologies & Practice for Secure Deployment

CONNECT is designed as a zero-trust framework, as mentioned in section 2.2.6, emphasizing the verification of every component's integrity within the in-vehicle environment. This adherence to a hierarchical architecture involves a step-by-step approach to achieving a fully attested system (see Figure 26). Starting from components like sensors, which are part of S-ECUs (i.e., camera, lidar, acceleration, radar, etc.), constraints on resources and cryptographic capabilities limit them to supporting symmetric encryption. They may possess an HSM module for the secure storage of keys, nevertheless cryptographic operations are not considered protected; hence, the level of protection provided by HSM is limited to storage. In cases where even an HSM is not available to the S-ECU, cryptographic operations may still take place, nevertheless this complete lack of any type of protection is reflected by the TAF. Asymmetric encryption capabilities and the use of a TEE, elevate the security posture for A-ECUs.

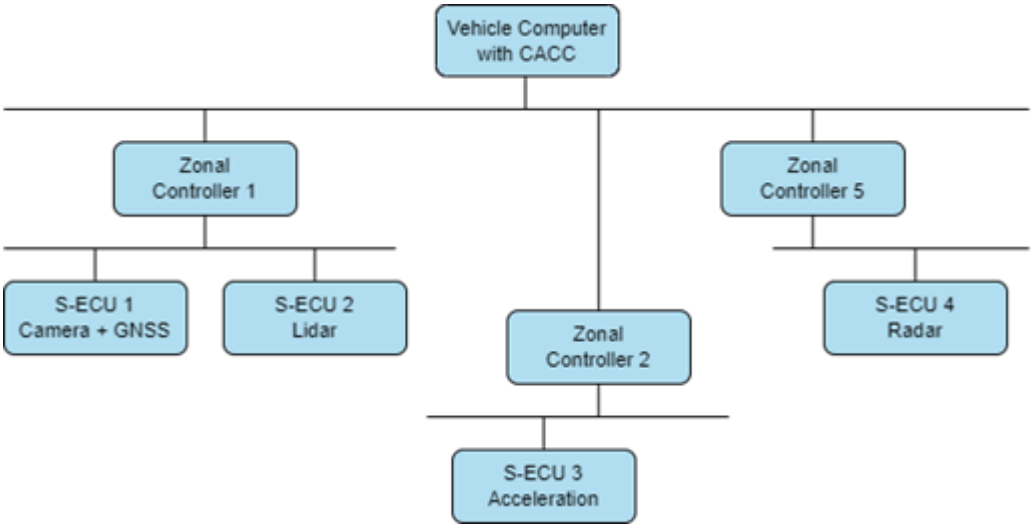

Figure 26 - Key Hierarchy between S-ECUs and Zonal Controllers

In scenarios where an ECU connects to a Zonal Controller, the latter acts as the verifier, and depending on the vehicle topology, these components (i.e., ECU or Zonal Controller) may undergo verification by the In-vehicle computer. For example, in Figure 26, since the S-ECUs are connected to a Zonal Controller, the Zonal Controller verifies the S-ECUs and itself and then it provides the evidence to the In-Vehicle computer. The In latter will attest the provided evidence to verify that they have not been tampered in the meantime. Notably, the In-vehicle computer is mandated by the standards to have a TEE for enhanced security controls. In contrast, the Zonal Controller can opt for either a TEE or an HSM, which offer distinct security guarantees. The remaining in-vehicle components have the option to leverage an HSM but are not expected to inherently provide robust security guarantees. Figure 27 illustrates the ECU classes and profiles, outlining individual integrity protection features. White boxes indicate mandatory cybersecurity controls, while the grey ones are included based on security risk assessments (TARA). This standards-driven approach primarily emphasizes securing the In-vehicle computer, with coverage extending to the Zonal Controller, while other in-vehicle components are not considered only as optional.

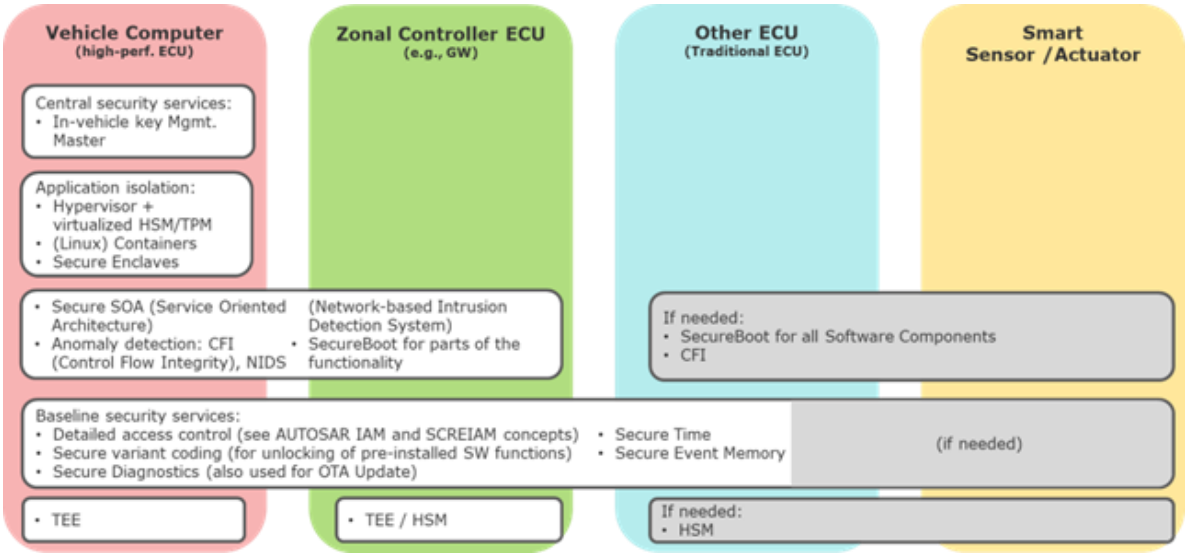

Figure 27 - ECU classes and their platform integrity cybersecurity controls.

6.4.3 Attestation and Integrity Verification and the Trustworthiness Claims Handler

Trustworthiness Claims (TCs) are signed items of evidence about the state of a given component, such as its integrity or its configuration. These claims are generated in each of the different parts of the CONNECT system and take several forms:

- ✓ **Direct evidence from devices or components in the system**. For example, attestation evidence from an ECU in a vehicle, or from a container somewhere in the system.
- ✓ **Indirect Evidence provided indirectly by other components across the CCAM continuum**. For example, trustworthiness claims about a vehicle wishing to report on its state while maintaining its privacy, or the assessment of V2X messages by the MEC.

TCs are transmitted outside the vehicle to supply evidence about the trustworthiness of the devices. Providing this type of data could be sent as part of the CAM/CPM messages but, if included with each message, this may introduce too large of an overhead for the over-the-air system. In this case, CONNECT will investigate caching strategies similar to those used for caching pseudonym certificates [115].

Inside the vehicle, the Trustworthiness Claims Handler (TCH) generates **harmonised attributes** used for reporting outside the vehicle. It does this using verifiable attestation reports from the Attestation and Integrity Verification (AIV) component of the vehicle and configuration data; for example, software versions, confirmation that keys were installed from the Identity and Authentication Management (IAM) component. The TCH will store the harmonised attributes and use them as required to provide information for the other parts of the system, such as that for CPM/CPM construction. The stored harmonised attributes will be updated each time that there is a new attestation. Note that together with the **Live Migration Management (LMM), IAM, AIV and TCH, form the CONNECT TEE Guard Security Extensions (TEE-GSE)**. These extensions will be further elaborated in D4.1 [111].

The devices in the vehicle form a hierarchy with the ECUs which control the sensors and actuators connected to a set of Zonal Controllers which are themselves connected to the Vehicle Computer, as shown in Figure 26 (which illustrates the devices relevant to the CACC use case, described in Section 7.3). Working on the zero-trust principle, all devices provide attestation results or evidence (i.e., as dictated by the deployed trust models and depending on their capabilities) to the AIV component of the Vehicle Computer. For instance, the A-ECUs and Zonal Controllers may be equipped with the CONNECT TEE extensions capable of supporting the advanced attestation capabilities based on the use of key restriction usage policies. One of its roles is to enforce the policies on the use of its keys, provided by the IAM component of the Vehicle Computer when the device is installed or updated. The policy for the attestation key will ensure that it can only be used to sign the attestation evidence, if it matches the known reference value. These devices are therefore able to perform local attestation/verification, based on the key restriction usage policies and send

back to the AIV the attestation result. On the other hand, devices such as the S-ECUs, that due to their limited resources cannot support local attestation leveraging the key restriction usage policies, provide their signed attestation evidence (as qutoes) to be remotely verified by the AIV. An outline of the attestation process and how these results are used is shown in Figure 28 - CONNECT Attestation Process Protocol where the AIV is acting as verifier for the S-ECU attestation evidence. This may change as we move to a federated TAF and, if necessary, this will be detailed in D3.2.

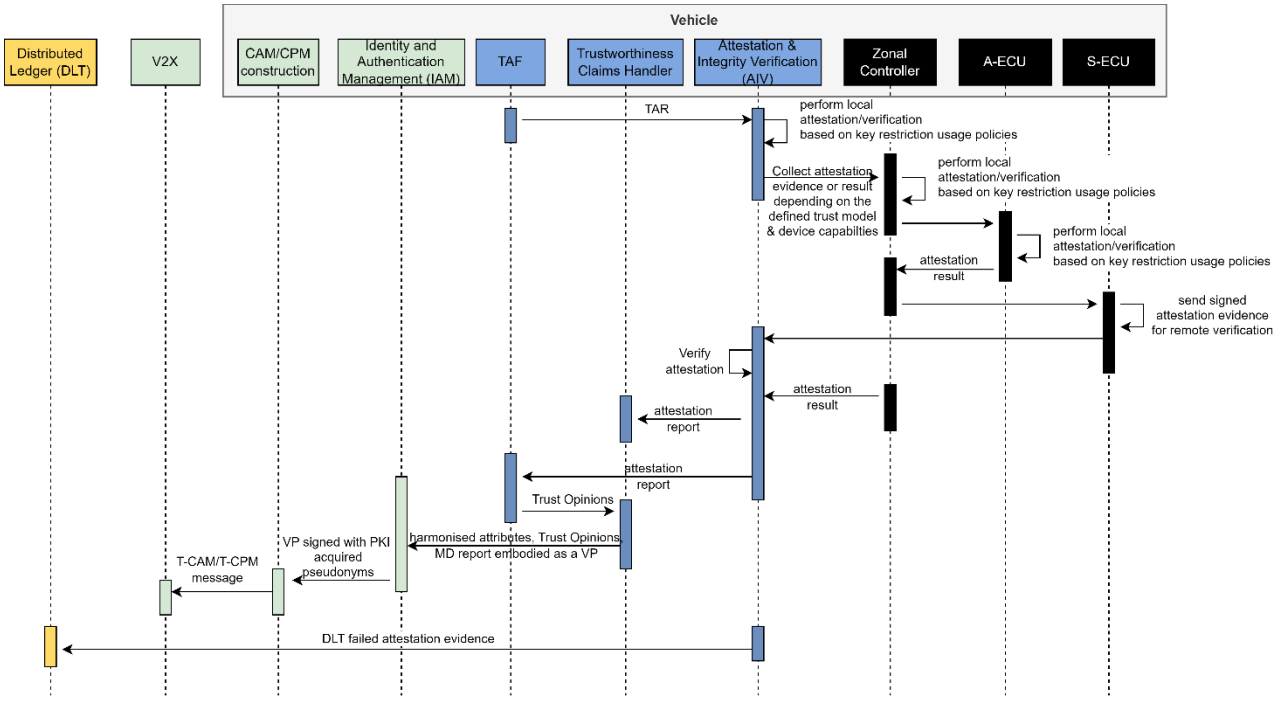

Figure 28 - CONNECT Attestation Process Protocol

The attestation evidence collected by the AIV will be signed and passed as an attestation report to the TCH and to the Trust Source Manager (TSM) component of the TAF. It shall be underlined though that regardless of the attestation approach (i.e., local or remote), both the A-ECUs/Zonal Controllers and the S-ECUs will send the evidence in the case of a failed attestation so that the AIV may forward them to the Distributed Ledger (DLT) for further investigation by OEMs as elaborated on Section 6.7.The failed attestation reports may be immediately sent to be stored on the CONNECT DLT, or stored locally and sent when the system is idle. More specifically, as the size of the evidence may be large the evidence itself will be stored by the Blockchain Peer on off-chain storage with an (encrypted) pointer stored on the DLT (as thoroughly explained in section 6.7).

Although running in different secure containers the AIV, TCH and TAF are running on the same vehicle and the attestation reports can therefore be protected with a traditional signing key. The harmonised attributes since they are sent outside the vehicle, they will be signed using a privacy preserving key mechanism. Hence, outside the vehicle, the Trustworthiness Claims for a given data item, or data stream, are represented by harmonised attributes. Harmonised attributes for attestation will depend upon the attestation results of all the devices involved in providing that data.

As an example and referring to Figure 26, suppose that the vehicle is sending a CAM message to report on its position. The GNSS data is read by S-ECU 1 and it is then passed to Zonal Controller 1 and finally to the AIV. The TCH will use the attestation reports from the ECU, its Zonal Controller and the AIV's container when generating the attestation harmonised attributes for the position data. Only if all devices attest correctly will the TCH allow the attestation harmonised attributes to be included with the position data in the CAM message. Note that, harmonised attributes will not just refer to the attestation results for devices in the system but will also refer to the wider range of trust

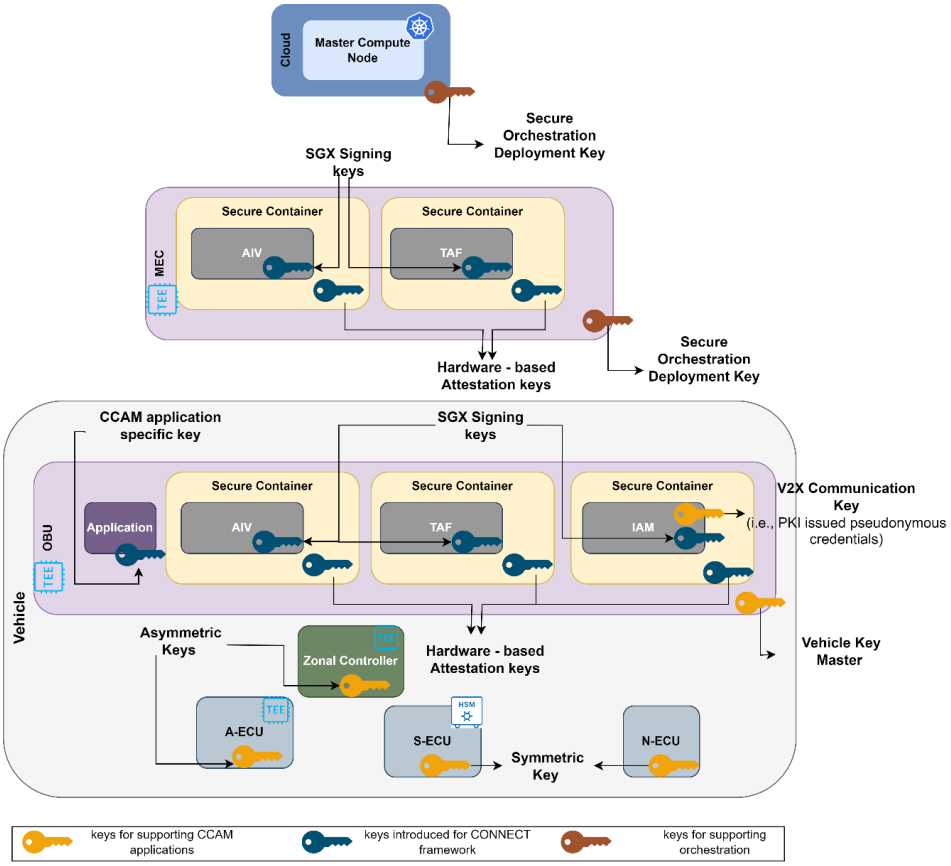

Figure 29 – CCAM application and CONNECT keys across all layers of the CCAM continuum (i.e., Vehicle, MEC, Cloud)

properties described in D3.1, such as Trust Opinions or Misbehaviour Detection Reports. The TCH component sends the harmonised attributes along with trust-related information in the form of a Verifiable Presentation (VP) and the Identity and Authentication Management (IAM) component, signs it with the PKI acquired pseudonyms, in order for the information to be sent outside the vehicle without revealing any information that could identify the vehicle.

The TEE-SGE are also implemented in the secure containers of the MEC and similarly used to attest to and provide TCs for the services supplied by the MEC. In particular, the MEC will provide evidence about the trustworthiness of V2X messages that it has received, allowing vehicles to judge how to handle the data that these messages provide.

6.5 CONNECT Key Management System

Keys in the CONNECT system are used for a number of different purposes: **for ensuring data integrity, for authentication, for confidentiality, for controlling access to data and for confirming device integrity (attestation)**. Protecting (as much as possible) the privacy of the vehicle puts extra requirements on some of these keys, in particular those used for reporting outside the vehicle on the harmonised attributes the trustworthiness claims and any misbehaviour reports. To do this, we will investigate the use of a number of anonymous signature schemes, such as BBS+ [116] and Attribute-based Direct Anonymous Attestation (DAA-A) [117].

For V2X communication, we will use the public key infrastructure (PKI) and pseudonyms as standardised by ETSI and 5GAA for safeguarding the identity and location privacy of communicating vehicles. From time to time, as dictated by the vehicle's Trust Model (i.e., defined during the Setup phase of the TAF as explained in Section 6.3), the messages will need to contain Trustworthiness Claims (TC) comprising evidence on the state of the CCAM actors based on which a fresh trust appraisal can occur. These TCs will need to be independently signed to confirm their authenticity and these signatures will also need to protect the privacy of the vehicle.

Inside the vehicle, a **number of keys will be required to secure communications between the car's components and for signing the attestation reports from those components**. Other keys will also be used to support different applications running inside the vehicle. How these in-vehicle keys are assigned and used will vary from one OEM to another, the **CONNECT key management system must be flexible enough** to allow for this. How the keys will be installed and used will be detailed in D4.1 [111]. As these keys are only used inside the vehicle, privacy is not a concern and straightforward signing mechanisms can be used. Keys will also be used to control access to the data and ensure that only 'certified' applications can process it. Actuators, like those for acceleration and braking, need to be sure that the data has come from the vehicle computer, while applications like those processing data for CAM/CPM messages, or building the LDM, need to know the provenance of this data.

The in-vehicle computer will use several secure containers, and these will have their own keys used to secure communications between the containers and for their attestation.

The MEC will have a number of secure containers, and these will also have their own key managers. The MEC does not have the same privacy requirements for its own data and messages and can therefore use a standard PKI for messages sent outside the MEC. These are broadcast messages, although in some circumstances they will be meant for a specific vehicle, for example when using a digital twin. Any TCs generated by the MEC can also use the standard PKI.

There are a number of different requirements for the keys, and these are discussed in the following sections and more specifically in section 8.1. Details of how the keys are derived and installed in the different components are not included here, they will be described in deliverable D4.1. The following paragraphs outline the fundamental principles using a top-down approach.

6.5.1 Vehicle Component Keys

As shown in Figure 27, different components in the vehicle have different cryptographic capabilities; hence, use different keys to prove their authenticity. A few will have no built-in HW capabilities for supporting **secure and efficient crypto agility** (i.e., N-ECU where crypto capabilities are offered through the “untrusted” host of the device) whereas some **can only perform symmetric crypto** (S-ECU) leveraging underlying secure elements that can only through provide Root-of-Trust for Secure Storage property (Section 8.1.4). This differentiation is a critical aspect that dictates the verifiability dimension of the trustworthiness evidence that need to be provided to the TAF for appraising the trustworthiness level of the target device. For instance, N-ECUs will be able to provide *attestation quotes* (holding their runtime system state evidence) signed under a cryptographic primitive with no assurances on its integrity; i.e., since the crypto software stack will be part of the overall operating system of the device not supported with any level of isolation. This means that the trustworthiness evidence received by the TAF will have a lesser weight in the trust assessment process as it cannot verify that it is not a result of an attacker’s attempt to manipulate the trust evaluation process.

On the other hand, the A-ECUs are more powerful in-vehicle sensors (including also the Zonal Controllers) that have the resources to instantiate a fully-fledged Trusted Execution Environment (TEE), thus, they will be equipped with the CONNECT TEE Extensions. In this case, not only real-time (verifiable) security claims will be able to be produced through the CONNECT TEE extensions, enabling the continuous monitoring of different types of system measurements depicting information needed for varying trust properties (as dictated by the underlying trust model) to be collected by the TAF, but will also be able to perform this operation in an efficient manner through the **newly developed key restriction usage policies**: The underlying HW-based attestation (signing) key will be binded to the expected state of the device, thus, the underlying trust anchor will constraint its usage if and only the device still remains at the correct state. This allows for the verification process to be performed by simply verifying a traditional signature, thus, also avoiding disclosing the details low-level system traces that can expose to possible implementation disclosure attacks (zero-knowledge configuration integrity verification [118]).

In what follows, we elaborate on the management of those keys and their differences, focusing on the components that reside on the vehicle side. Furthermore, the keys introduced in this section are illustrated in Figure 29, starting from the ECUs (i.e., A-ECUs, S-ECUs, and N-ECUs) and reaching all the way to the cloud. As explained in the following paragraphs, the vehicle side is comprised of i)

the component keys and ii) the secure container keys. While the first type (i.e., keys coloured in yellow in Figure 29) is present in any CCAM application, the second type is introduced specifically for the context of the CONNECT framework (i.e., keys coloured in blue in Figure 29). The V2X communication key (i.e., PKI-issued pseudonymous credentials) is protected and managed by the IAM component within a secure enclave. This key is the one that enables communication with other vehicles as well as the MEC without revealing any further details regarding the identity of the vehicle.

6.5.1.1 S-ECU Keys

The S-ECUs may have an HSM, often integrated into the System on a Chip (SoC). The HSM manages the secure storage of keys and it also preserves results of the secure boot process, which at their simplest form are an 'extended' hash. Nevertheless, the presence of an HSM in S-ECUs is not guaranteed. Therefore, in the cases of S-ECUs that lack an HSM, keys may still be used for communication and for signing attestation data, however it is not considered as trustworthy, and this will be reflected in the calculations of the TAF. The zonal controller, depending on OEM-specific configurations, may directly verify the attestation data or forward it to the AIV component for verification.

During the manufacturing phase, the S-ECU is provided with a **unique identity** and a **symmetric key** that is used for the initial integration of the ECU into the vehicle. Other keys will be installed to enable the communications to be integrity checked and, where necessary, to control access to the data being sent or prove the provenance of data being received.

6.5.1.2 A-ECU and Zonal Controller Keys

These devices are equipped with a TEE for securely storing their keys, attestation results and the attestation reference values. The inclusion of TEE enables these devices to independently verify their attestation and enforce key restriction usage policies for ensuring that attestation data can only be signed if it is correct. This allows the attestation process to be short-circuited, that means that if the OBU receives a signed and verified attestation report it will know that the attestation succeeded.

During the manufacturing phase, the A-ECUs and zonal controllers are provided with **asymmetric signing keys** and **certificates**, these are used when integrating the devices into the vehicle. As with the S-ECU, other keys will be installed and used to protect communications, to control access to data being sent and allow the provenance of data being received to be confirmed.

6.5.1.3 Vehicle Master Key (OBU)

In addition to the keys used by the secure containers (see Section 6.5.2), the main Vehicle Computer will have its own identity and associated with this an integration key. The Vehicle Computer is equipped with a TEE and uses asymmetric cryptography for all the instantiated secure containers. Recall that this component will also have instantiated a network compute node for allowing its secure interactions with the MEC deployed services (Section 6.6). Thus, all these containers will be equipped with HW-based keys, as identity keys (leveraging the secure element and confidential computing technology described in Section 6.4.1.2), used to sign the Verifiable Credentials needed for supporting the continuous authentication of the vehicle with the other CCAM actors. Whenever its attestation status needs to be confirmed externally a privacy-preserving key mechanism is needed to sign its attestation report without revealing any information that could fingerprint the vehicle.

6.5.2 Secure Container Keys

As outlined in Section 6.4.1, the enclave-cc infrastructure allows a stakeholder to create a container, check its integrity, and then allow the container to use keys if and only if the integrity verification has successfully completed.

There are multiple types of keys related to the container security infrastructure:

1. **Hardware Attestation Keys** that allow stakeholders to verify attestation tokens generated by the Intel SGX hardware The stakeholder can use SGX attestation to verify the integrity of the Container Agent enclave that manages other “payload” enclaves.
2. **SGX Signing keys:** Keys that allow stakeholders to sign an executable to be executed within an SGX enclave. The corresponding signatures are then maintained inside the so-called manifest. It is worth noting that the SGX Signing keys are based on the hardware attestation key, since the latter provides the underlying Root of Trust (RoT). The SGX key provides, in essence, the signing key of each enclave that is built on top of the hardware.
3. Each “payload” application container contains **application-specific keys** – including a key for encrypting/decrypting the FUSE encrypted file system. The application inside the container may need additional keys that can then be stored inside the encrypted file system.\
4. Lastly, **RoT – Protected communication key** for the governance of the Root of Trust (RoT) is needed, to secure this part of the communication between the secure elements (SEs) of the MEC and the OBU (i.e., irrespectively of the underlying SE). Since CONNECT aims at avoiding a fragmented solution caused by the support in different SEs in different endpoints, the construction of an additional layer named “mediation layer” is envisioned, for the exchange of trust related information irrespectively of the SE.

6.5.2.1 MEC Containerized Services Key Management

The MEC layer follows a similar approach to the vehicle, adopting the notion of Hardware-based keys, SGX signing keys, as well as the RoT- Protected communication key, according to the ones described in the previous paragraph for the vehicle. These keys are based on TEE, to protect the CONNECT-introduced components (i.e., AIV, TAF, etc.). Furthermore, the RoT - Protected communication key introduces an added layer of security for the communication of the two secure containers: the one residing at the MEC side and the one residing at the vehicle side.

A distinctive key in this context is the Secure Orchestration and Deployment key, serving a critical function in guaranteeing the authenticity and integrity of orchestration and service deployment processes. Unlike the keys mentioned earlier, the Secure Orchestration and Deployment key uniquely contributes to securing the processes involved in orchestrating and deploying services, emphasizing its role in fortifying the overall security posture of the MEC layer within the CONNECT architecture.

6.5.3 Attestation Keys

The CONNECT architecture ensures a secure and authentic exchange of Trustworthiness Claims (TCs), which contain information crucial for trust assessment and may be received from various sources, either internal or external to the vehicle. These TCs include Misbehavior Detection (MD), Attestation and IDS-related reports. To preserve privacy outside the vehicle, the system employs direct anonymous attestation, as well as harmonization techniques that ensure that fingerprinting of the vehicle is not possible.

In the verification process, CONNECT deviates from the conventional State of the Art (SOTA) remote attestation approach, where the ECU typically sends quotes to the verifier using a pre-established key. In contrast, CONNECT introduces an innovative local/remote codesign approach for attestation. This flexible approach adapts to the device's capabilities, proposing key restriction usage policies that bind the key with the authentication result. This verification confirms the correctness of the device/component state. Simultaneously, the key restriction usage policy ensures the binding of the key used for signing the attestation evidence to the entity's identity. In cases where asymmetric key cryptography is unsupported (e.g., in S-ECUs or N-ECUs), CONNECT leverages the SOTA approach, relying on pre-established keys for remote attestation, where a quote is transmitted to the Verifier. Detailed information on this process will be provided in document D4.1.

6.6 Mobile Edge Computing (MEC) and Secure Service Deployments

In this section, we discuss the way CONNECT exploits **containerization technologies** to deploy services in the considered automotive setting. A number of locations in the 'vehicle-infrastructure domain' (as part of the overall CCAM continuum) (see the proposed architecture in Section 6.1, including Cloud, edge and in-vehicle deployments) can be exploited to leverage vehicle, sensor, network, service, and cluster telemetry data, potentially of relevance to misbehaviour detection. The MEC (together with the in-vehicle platform) are two such prominent locations, explored in CONNECT, where the selection criterion of deployment, (i.e., where to deploy a service), can depend on various factors such as processing the data close to the area of interest (thus, preserving local relevance) instead of sending them to remote locations, or ensuring some service key performance indicator such as latency. Regardless of the location, any deployed container should at the same time be enabled to cope with involved security challenges.

The following subsections provide a brief overview of a Kubernetes system, extended to support the security features relevant to CONNECT. *Note that the proposed CONNECT system adheres to the cloud native principles regarding virtualization and orchestration technologies and is not solely restricted to a Kubernetes system*. Hence, other orchestration platforms (e.g., Docker Swarm) can be exploited. In CONNECT we utilise Kubernetes, the de facto standard for container orchestration, for the efficient management of container applications at scale, taking advantage of its advanced features on resource management, security, self-healing and multi-tenancy, supported by a large and active community, and supervised by the Cloud Native Computing Foundation (CNCF) ensuring its ongoing development and standardization. More technical details of all relevant components and implementation specifications will be delivered on Deliverable D5.2.

6.6.1 Automation of CONNECT Service Deployment

We now briefly introduce the Kubernetes technology and present the corresponding CONNECT service deployment. The Kubernetes (abbreviated as k8s) is an open-source, production-ready technology which controls distributed compute nodes where containerized[12] applications are deployed to work as a single entity composed of potentially many micro-services. **Virtualization allows the removal of dependencies on individual machines going beyond the legacy deployment model whereby applications, installed as packages, were closely coupled/integrated with the host** [119]. Particularly, CONNECT exploits three main nodes (see Figure 30) to facilitate a highly automated service deployment and lifecycle management spanning across the automotive setting: **the master node residing at the cloud back-end; the MEC nodes and vehicle nodes (see the proposed architecture in paragraph 6.1) which run containerized applications and handle networking between CCAM/CONNECT applications in the cluster and across clusters**.

In a k8s system, the **Master Compute Node** (see Figure 30) is responsible for control plane operations, managing and monitoring the cluster's state. This involves observing the cluster resource usage/availability (CPU, Storage, etc.) in real time, which is used for guiding the service orchestration decisions; self-healing operations for maintaining the desired state of a service (e.g., always run on at least 10 CPU cores) by monitoring running/failed containers and relevant lifecycle management (LCM) operations; scheduling the workload of containerized services in k8s worker/compute nodes (whether VMs or physical computers) for running the applications; as well as adding/removing worker nodes to/from the cluster. Kubelet (see Figure 30), an agent installed in all cluster nodes (master and worker) is responsible for these mechanisms.

Additionally, kubelet interfaces with the container run-time (e.g., containerized, docker-shm, cri-o, etc.) to start, stop, and monitor containers (i.e., perform LCM operations), as well as with kube-proxy that facilitates that network traffic within the k8s cluster is correctly routed and load-balanced to the appropriate containers and services. Table 6 lists a subset of control plane agents that facilitate the control plane functionalities of a kubernetes cluster. For more details, please refer to [13].

[12] Containers can be briefly described as software packages that include all necessary elements to virtualise the underlying operating system and therefore, capable to run in any environment.

[13] Kubernetes Documentation Home. Retrieved from https://kubernetes.io/docs/home/

Table 6 - Agents/utilities in K8s nodes

K8s agents/utilities	Functionality
Kube API-server	Performs all administrative tasks on the master node
ETCD	A distributed key-value store that is used to store the cluster state
Kube-scheduler	Used to schedule the work to different worker nodes
Kube-Controller task	Obtains the desired state from the API Server (desired state VS current state)
Kube-proxy	Used to communicate between multiple worker nodes (network plane)

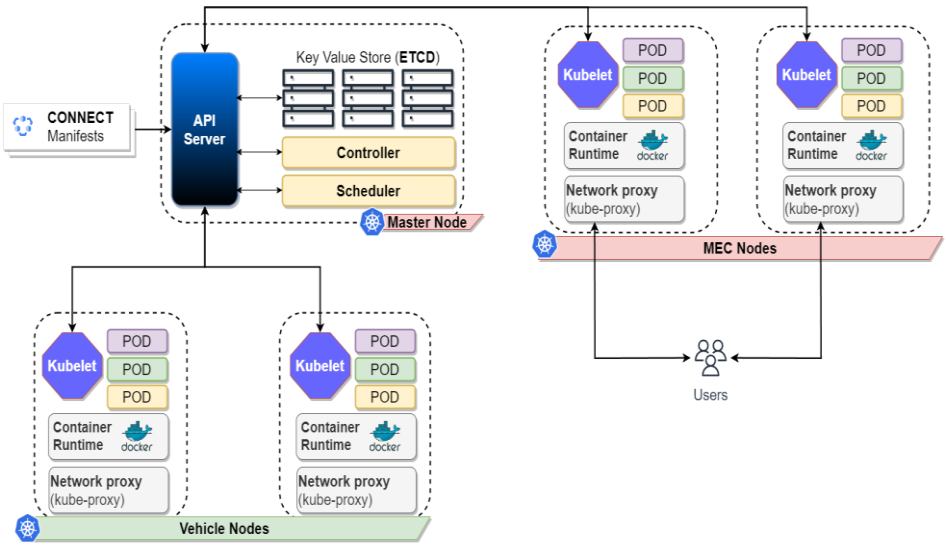

Figure 30 - A high-level notion of the CONNECT K8s deployment approach (right)

A K8s deployment involves the establishment of a pod hosting an application, rather than directly setting-up a single container. A pod is a k8s abstraction that manages a group of application containers and the shared resources (such as storage, networking, IP addresses and metadata) that those applications require. All containers in a pod share the same IP address, and typically a Pod is an abstraction introduced by kubernetes to be agnostic of the container technology used for building the micro-service (e.g., docker containers or linux containers, etc.).

A simple yet indicative example of a CONNECT k8s deployment is depicted in Figure 31 whereby a master node deploys through control plane directives (see k8s API) and the scheduling capabilities of the kube-scheduler ; the latter collaborates with etcd and kube-api for finding the appropriate (worker) node for the desired k8s Pod to run on (i.e., for nodes that can support the services requirements in terms of e.g., compute or network capacity).

6.6.2 Securing the Deployed Containers

Containers technology needs definitely to come with security capabilities to address a number of (NFV) risks[14]. Numerous (security) dimensions are involved and need to be covered by the enabled controls/mechanisms of the involved virtual functions. Along this line, the so-called container security seeks to implement security tools and policies to ensure that **container code is running as intended, including protection of the container application (and containerized workloads), the deployment environment and (to some extent) the infrastructure**. In CONNECT, all such containers hosting applications to be deployed over the MEC are encoded as secure containers allowing for the secure operation of the application to be instantiated as part of an enclave. In this

[14] Those risks essentially relate to the (VNF-related) threats identified in Section 3.5. Clearly, the virtualization risks span across a broad range of software engineering/secure-coding challenges that go-beyond the scope of CONNECT.

context, confidential computing technologies are adopted (Section 6.4.1.2) been able to operate on top of any type of underlying secure element – capable of providing the properties defined in Section 8.1.4. All the CONNECT-enabled confidential containers are equipped with the appropriate attestation capabilities for providing different Levels of Assurance as defined by ETSI[15].

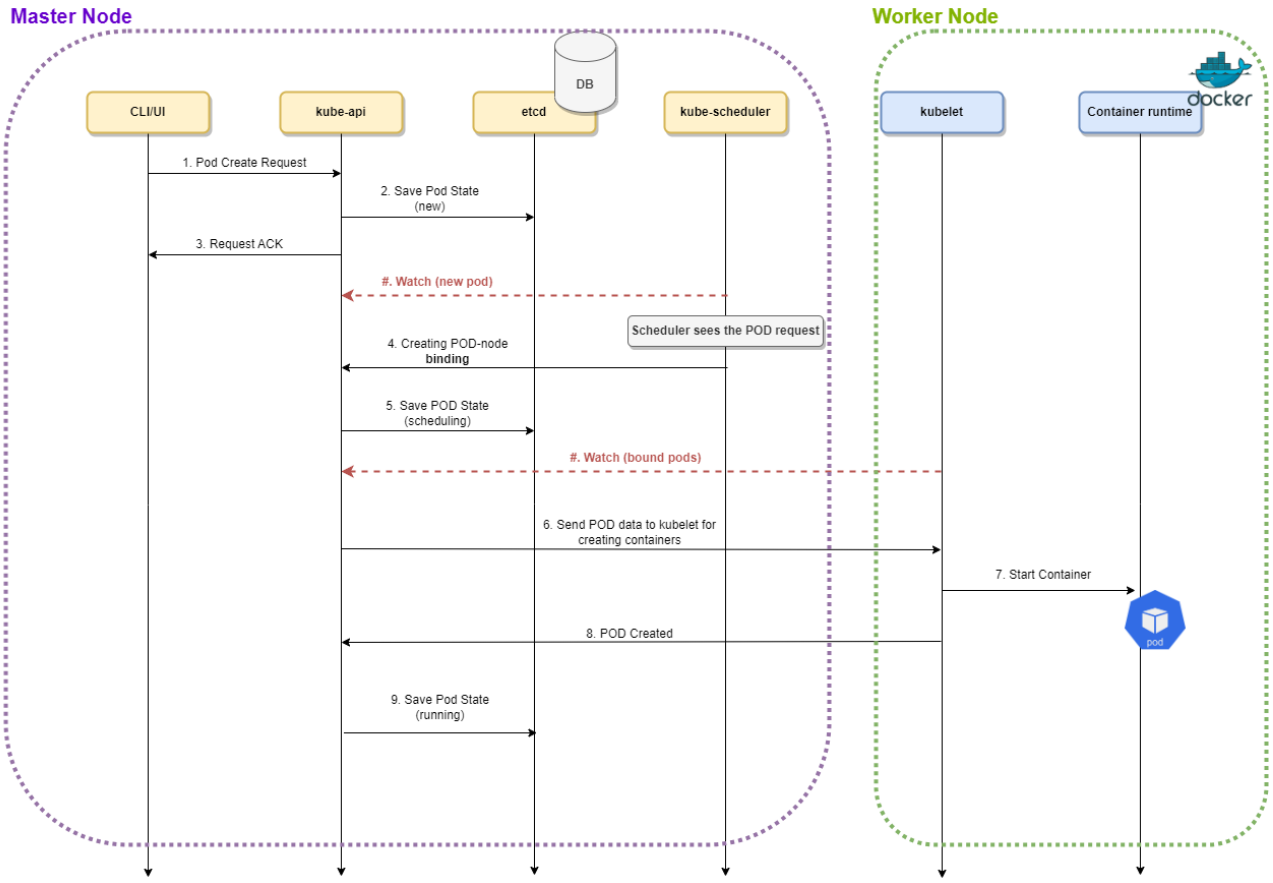

Figure 31 - The CONNECT k8s pod creation (driven by the Master node)

Towards that end, CONNECT employs a TEE in the form of an enclave (see Section 6.4.1). The main concept involves running applications in a well-protected environment, potentially backed-up by hardware support (i.e., SGX) acting as a root of trust. The corresponding framework used for the (execution of binaries of the) CONNECT containers is named Enclave-CC [120] and allows the execution of a container inside the CONNECT enclave where no other malicious software can have access-to. In a typical scenario, access to a (secure/certified) repository of VNFs (e.g., services or applications) is controlled and managed by an organisation or service provider that offers VNFs for deployment within the network infrastructure. Access to the VNF repository is usually controlled and restricted with authentication and authorization policies and access control mechanisms. For CONNECT, the certified applications VNF registry can be collocated with the Kubernetes Maste Compute Node following the CONNECT principles (in terms of privacy, security, TEE guard) offering CONNECT protected containerized applications.

In more technical terms, the enclave-cc framework [120] exploits components from the cloud-native stack (i.e., orchestration and containerization), Intel SGX technology and existing key management approaches to offer a process-based confidential container solution to the Kubernetes platform. A reference architecture (see a short description of the modules in Table 7) is depicted in the following figure.

[15] ETSI GR NFV-SEC 007, "Networks Function Visualization (NFV) Trust; Report on Attestation Technologies and Practices for Secure Deployments", 2017, [Available Online: https://www.etsi.org/deliver/etsi_gr/NFV-SEC/001_099/007/01.01.01_60/gr_nfv-sec007v010101p.pdf]

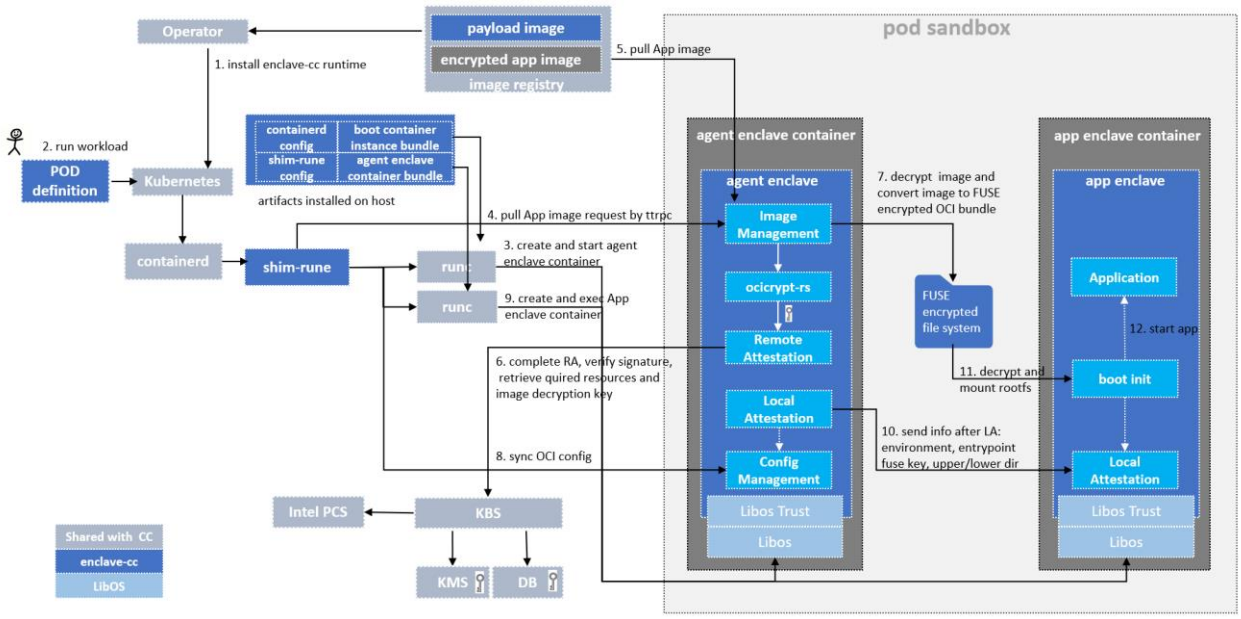

Figure 32 - Main functionalities of Enclave-CC framework

In terms of the depicted workflow, we very briefly recognise two phases: First, the enclave-CC is installed (Step 1 in Figure 32) in the k8s environment by securely transferring binaries and configuration files to the host directory. Then, a Pod configuration is defined to describe the workload and runtime requirements with shim-rune getting the container creation request, creating an agent enclave container (step 3) and asking the enclave-agent to pull the encrypted workload image. Enclave-agent verifies (step 7) image signatures and takes care of key management (step 10) and communicates with shim-rune towards the creation of the app enclave container. More details will be provided in the coming WP4 and WP5 deliverables.

Table 7 - Main modules/utilities in the Enclave-CC framework

K8s ag	Functionality
enclave-agent	a process inside LibOS inside the enclave. It accepts request from shim-rune and handles image management, attestation, config management, and fuse encrypted file system management
shim-rune	a standard shim component sits between containerd and runc. It accepts request from containerd, starts pause and agent enclave container, asks enclave-agent to perform image management actions, and starts agent enclave container. All the containers started by shim-rune are all instantiated by runc
operator related tools	help to build the enclave-cc payload image for CoCo operator. The payload image includes binaries of shim-rune and modified containerd, the image bundles of agent enclave container and boot instance, the configurations shim-rune, and the enclave-cc runtime deployment script.
operator	a Kubernetes operator to help install/uninstall enclave-cc runtime in a cloud native way
LibOS	Occlum is supported and Gramine in future
FUSE	encrypted file system: enclave-agent puts decrypted app image into an encrypted fs and later referenced by LibOS as roots to start app enclave container

6.6.3 Leveraging Edge Computing in the CONNECT Data Exchange

Data exchanges in CONNECT may involve a diverse set of endpoints (i.e., from ECU to ECU over in-vehicle buses to V2I communications over cellular networks – see Chapter 7) depending on the considered use-case. When focusing on its task-offloading concept (mainly showcased in the Slow-Moving Traffic Detection use-case, as described in Section 7.4), CONNECT seeks to heavily rely on

the convergence of edge computing with secure-containers distributed orchestration technology, both highlighted in the previous paragraphs. As such, the points of interest in the considered setting are mainly the vehicle platform/OBU and the MEC (infrastructure) while the back-end cloud may also host services of relevance.

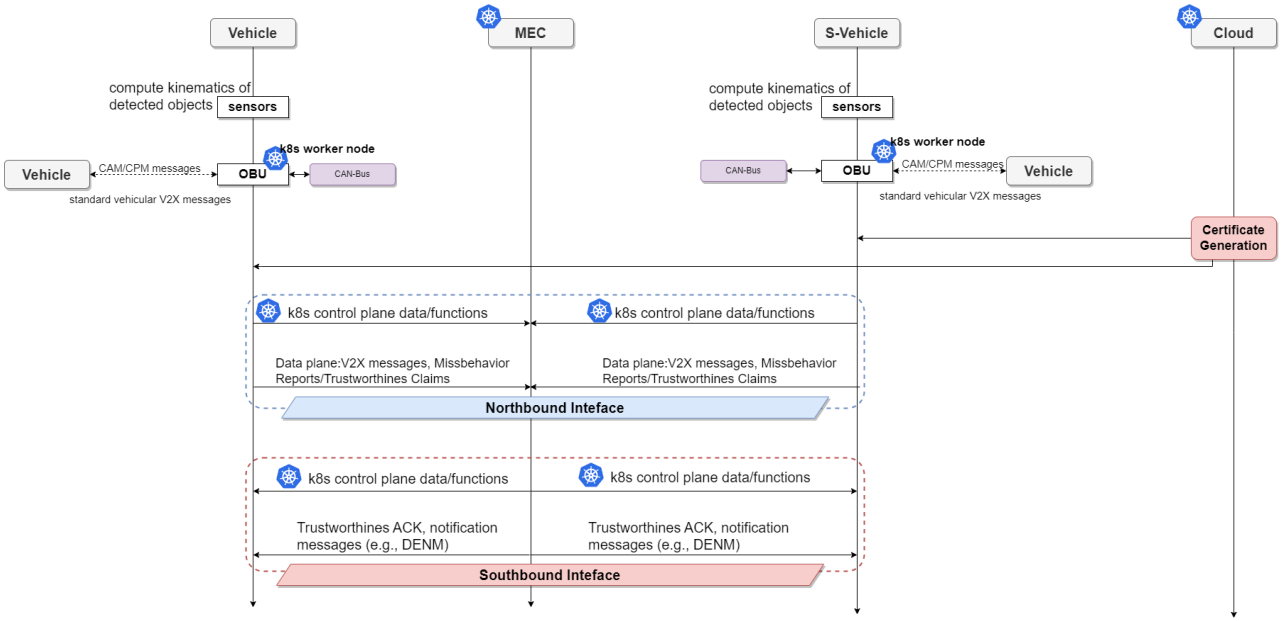

Figure 33 - Data-plane view of the MEC functionality: an example based on the slow moving traffic use-case (see Section 7.4)

In that sense, the MEC constitutes a prominent location to offload[16] demanding tasks (i.e., safety-critical services) from the vehicle towards the infrastructure, especially when latency concerns are involved. In Figure 33 we briefly showcase the interactions between the above points abstracting the hosted applications/services with the corresponding k8s pods which reside in each one of the aforementioned locations. Any offloading algorithm (each one designed to fulfil different objectives - see D5.1) requires information both from operation of the vehicles (e.g., kinematic information, computation resources) and the infrastructure (e.g., storage, memory, computation resources) while in parallel ensuring that all relevant information is accompanied by the CONNECT trustworthy evidence (see D3.1). Furthermore, as part of the process, CONNECT-related data and (potentially) certain C-ITS messages need to be communicated to the vehicles.

Having the CONNECT k8s Master Compute Node placed at the cloud node (cf. Section 6.1) coordinating also the edge/MEC infrastructure, we preliminary identify two ways to gain access to the above pieces of information, namely through a northbound and a southbound interface (see In Figure 33): the former accommodates data that are generated at the vehicle platform/OBU and the hosted CONNECT/CCAM applications and communicated (whether multicasted or unicated) to the edge; the latter involved the information that is the outcome of edge-processing, whether CONNECT-of typical CCAM-related) and is to be communicated to the participating vehicles/OBUs.

Finally, apart from the MEC, cloud locations can serve as other candidate locations for offloading (as is for instance the case of digital twins software running at the cloud back-end). K8s pods deployed in the cloud infrastructure will again help manage the involved services. In parallel, cloud servers may host third-party services (i.e., external to the CONNECT services ecosystem) which provide information of relevance in the automotive setting, (see Figure 33). A prominent example is the identity management services which (being outside of CONNECT work) may be implemented as a collection of Certification Authorities and Trust Service Providers (e.g., as in [121]). They are to provide the required certificates which ensure that all involved C-ITS entities are legitimate (certified) parties in the relevant communications. CONNECT services may consume such kind of information, considered trusted.

[16] Although such offloading scenarios have already attracted the interest of industrial innovation (i.e., 5GAA), their implementation over the cellular network which requires the careful integration of the legacy C-ITS/automotive messages has rarely been shown in real-world demos before.

6.7 DLT for Establishing a Chain of Trust

The CONNECT DLT (Distributed Ledger Technology) acts as a core element for storing critical data within the CONNECT system, which depict the state of the vehicle's components. This architectural decision in terms of storage not only ensures auditability but also facilitates the monitoring of trustworthiness evidence, particularly in attestation processes, resulting in the establishment of a historical trust record. This signifies that any authenticated and authorized entity can verify, at any given moment, the current state of a component. Moreover, the TAF (Trust Assessment Framework) can harness this detailed information as a foundation for a reputation system, leveraging the historical trust data stored on the DLT for more sophisticated trust assessments.

This data includes: i) *failed attestation evidence* that are leveraged by the TAF to calculate a trust level for the the ECUs (i.e., CCAM service) or the MEC (i.e., MEC service) that failed the runtime service or component verification task; hence the related evidence (i.e., low level system traces or logs) is inserted onto the DLT to be checked by OEMs, and ii) *trust model templates* including *RTLs*, derived from the Trust Assessment Manager (TAM) with the support of the Risk Assessment (RA). This type of data is crucial for the relevant stakeholders (i.e., OEMs, MEC security administrators, etc.) to analyse the reasons behind the failed attestation, with the intent to propose appropriate mitigation measures (i.e., security controls), thus update accordingly the RTL, which lays the foundation of the trust calculation process, as thoughtfully explained in section 6.3. Upon accessing this information, the OEM (i.e., for the CCAM) or the security administrator (i.e., for the MEC) may delve into the exact trust property (i.e., integrity, communication integrity, safety, etc) that failed the attestation for the given component and consequently be able to identify and pinpoint the point of intrusion, thus recognize possible zero-day exploits. This information is translated into vulnerabilities and uploaded to the DLT, for the consumption of the TAM. The latter, in tandem with the RA, may re-evaluate the risk graph and update the RTLs accordingly, which plays a pivotal role in the extraction of a Trust Decision.

It has to be noted that within the CONNECT DLT, no personally identifiable information (PII) will be stored, such as location or other sensitive data. Only low-level system trace will be maintained, as derived from the attestation processes. CONNECT considers that among these traces, sensitive information may be included, therefore, through the use of Attribute-based encryption (ABE) and Attribute-based Access Control (ABAC), it ensures that only authenticated and authorised users may access such information, while it further enables data sovereignty leveraging concepts from Self Sovereign Identity (SSI). CONNECT Impact Assessment will be performed and documented in the context of D5.2.

The following paragraphs summarise the flow of actions for the DLT based on the two phases, as introduced in section 6.1, the *setup* and the *runtime*. During the setup phase, we consider the deployment of the trust models, RTLs and trust policies onto the DLT, as well as the authentication of the TAF Agents (i.e., both the one residing at the MEC and the one residing at the vehicle). The aforementioned trust models, policies and RTLs are strictly accessible to authorized entities, within the Private Ledger, leveraging the Blockchain Peer. During the runtime phase, we consider the recording of the failed attestation evidence, either for the MEC or the CCAM services and the identification of vulnerabilities by security administrators and OEMs respectively. Whenever new vulnerabilities are recognised, the RTLs are updated.

Setup Phase: In Figure 34, the main operations of the CONNECT DLT are depicted. The first step is for the TAM (Trust Assessment Manager) running on the cloud, to send to the DLT, and more specifically to the Blockchain Peer, the trust policies, including the trust model templates and the RTLs, as defined in section 6.3. The aforementioned information is recorded to the Private Ledger; in a form of a smart contract, hence it is accessible only by authenticated and authorised entities, leveraging Verifiable Credentials (VCs), adopted from the Self Sovereign Identity (SSI) instance. We consider that the interested parties have been previously registered to the DLT, therefore already possess valid VCs. Whenever a new smart contract is created, all registered parties will be notified. In order though to access and extract the new trust policies, trust models and RTL(s), the TAF Agent (i.e., either MEC-based or in-vehicle) shall be authenticated first. For this a novel Attribute Based Access Control (ABAC) scheme is supported by CONNECT, which enables access to specific information based on attributes (i.e., entity characteristics).

After their authentication, the TAF Agent(s) send a request to the Blockchain Peer to extract the trust policies, the trust model templates and the RTLs. These are available at the Private Ledger, hence the Blockchain Peer sends a query with the block address to the Private Ledger, extracts the requested information and sends it back to the TAF Agent(s). The MEC-based TAF Agent receives information regarding the MEC service, while the in-vehicle TAF Agent receives information regarding the CCAM service. Upon receiving the trust model template from the Ledger, the TAF agent may initiate a trust assessment process.

Runtime Phase: Regarding the storage of the failed attestation evidence in the CONNECT DLT, the TAF Agent upon extracting the trust model template, it may initiate the collection of trustworthiness evidence, including the attestation evidence collected by the AIV (Attestation & Integrity Verification), by sending a Trustworthiness Assessment Request (TAR). The DLT may store failed attestation evidence both for the MEC service as well as for the CCAM service. Our description starts with the MEC and then the CCAM follows.

The AIV executes the attestation process and receives a failed attestation evidence. In this case, the AIV leverages the underlying CONNECT Trusted Computing Base (TCB) to construct the appropriate authentication tokens and credentials (its characteristics are elaborated on Section 6.4.3), to authenticate to the Blockchain Peer. In particular, the AIV discloses the entirety of the attestation evidence as part of the VCs signed with the underlying hardware-based key (e.g., DAA key) to the Blockchain Peer.

The confidentiality of this evidence is protected by using CONNECT newly designed Attributed Based Encryption (ABE). Essentially the AIV will instruct the underlying CONNECT TCB to generate encryption keys linked to the ownership of specific system attributes. Such attributes can span from identity related attributes to device-related characteristics (e.g., version of OS, type of stack, etc.).

This approach ensures that only the entities possessing the appropriate attributes have the capability to create corresponding decryption keys while the overall process is governed in a secure and verifiable manner through the underlying TCB, as will be further documented in D4.2. This ABE-encrypted failed MEC attestation evidence is uploaded by the Blockchain Peer to the Off-chain storage, primarily due to the substantial size of the information. Simultaneously, pointers to this encrypted evidence are relayed back to the Blockchain Peer, prompting an update to a smart contract on the Private Ledger. This smart contract is amended with the encrypted pointers, ensuring a secure linkage between the evidence and its storage location, thereby reinforcing the auditability and confidentiality of the failed attestation records.

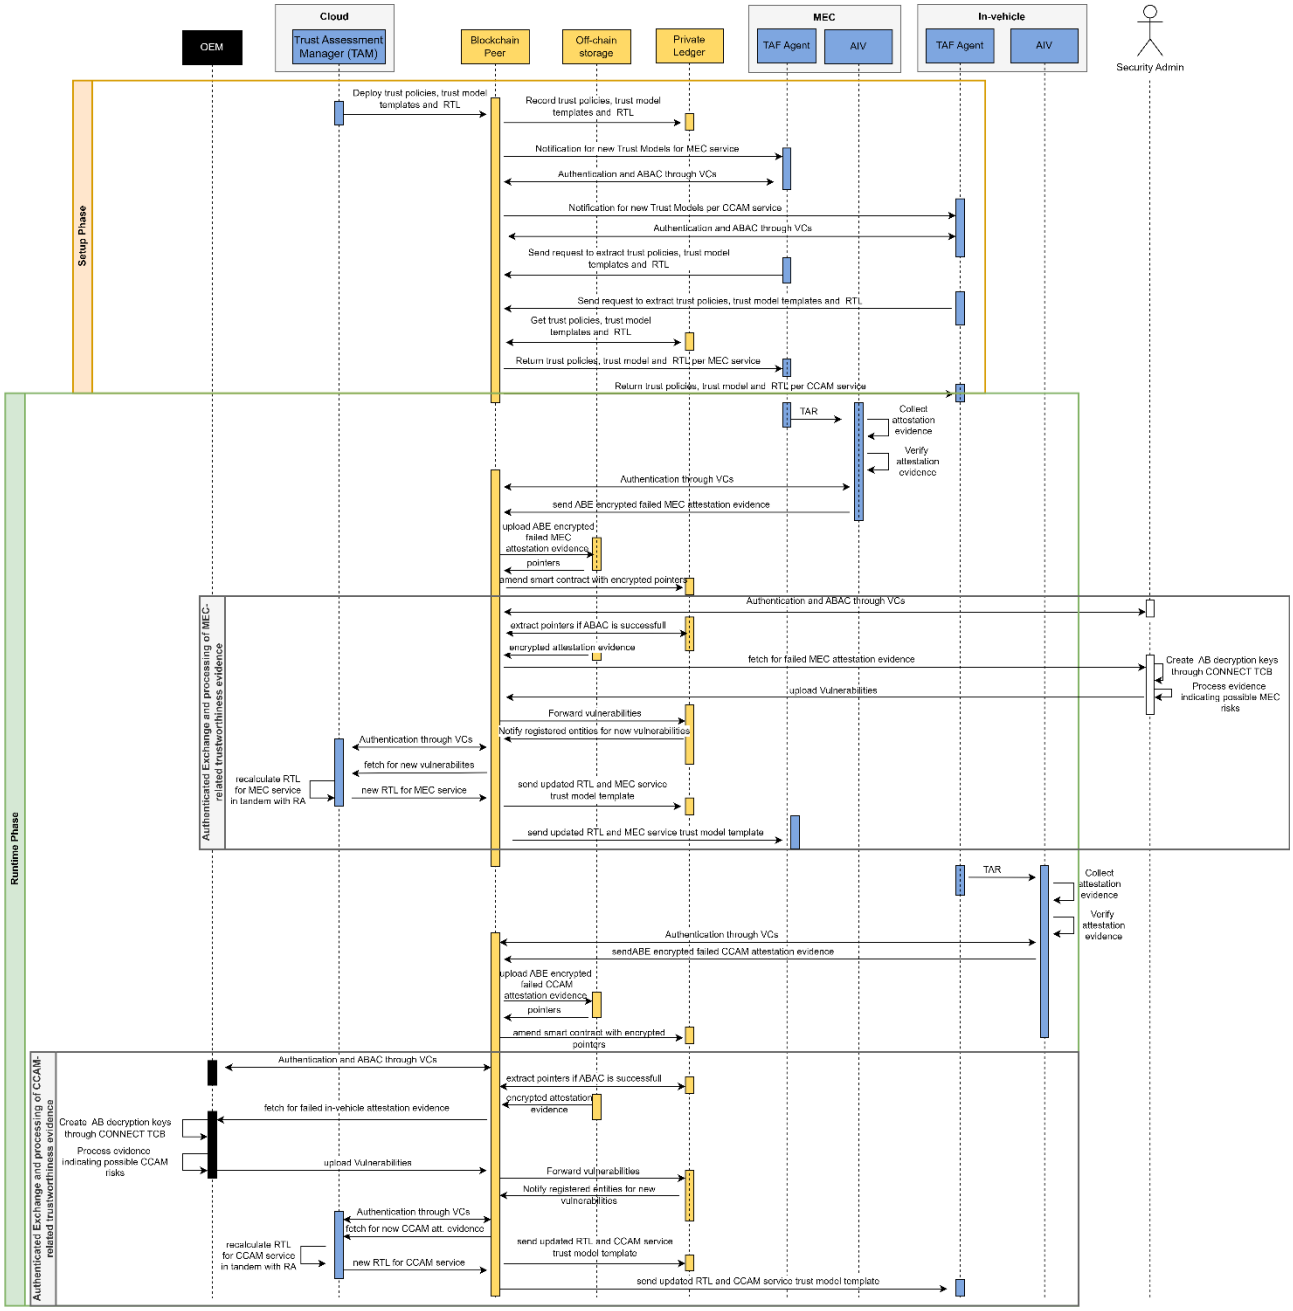

Figure 34 - CONNECT Blockchain Foundational View

This failed MEC attestation evidence can be further accessed by a Security Administrator, with the intent of analysing them to possibly discover and report new vulnerabilities. Towards this direction, the Security Administrator is authenticated granted access, based to its attributes (i.e., ABAC) to the Blockchain Peer, utilizing VCs. If the ABAC is successful, the Blockchain Peer is able to extract the pointers from the Private Ledger and consequently access the encrypted evidence stored at the Off-chain storage. This encrypted evidence is sent to the Security Administrator by the Blockchain Peer. The Security Administrator may now decrypt the evidence, by leveraging the CONNECT Trusted Computing Base (TCB). Decryption keys are generated by the CONNECT TCB only if the system attributes required for access exist and are verified. The ABE decryption key is calculated based on the VCs and specific attributes that the admin possess. After the decryption of the evidence takes place, Security Administrator is able to analyse the evidence and identify possible MEC-related risks. The next step is for the Security Administrator to upload the identified vulnerabilities to the Blockchain Peer, which will forward them to the Private Ledger.

This specific piece of information is particularly useful for the TAM, which is notified on the inclusion of new vulnerabilities. The TAM, after being authenticated to the Blockchain Peer (i.e., through VCs), receives the discovered vulnerabilities which uses in order to recalculate the RTL, in tandem with

the Risk Assessment (RA) component, for the MEC service, as described in section 6.2. The new RTL for the MEC service is sent to the Blockchain Peer and then to the Private Ledger. The TAF Agent will be notified on the updated RTL.

Similarly, the In-vehicle TAF Agent extracts trust model templates and initiates collection of trustworthiness evidence, from the AIV (among other sources), through the TAR.

The in-vehicle AIV executes the attestation process and receives a failed attestation evidence for a CCAM service. Like the MEC-based AIV, the in-vehicle AIV leverages the underlying CONNECT TCB to construct the appropriate authentication tokens and credentials, to authenticate to the Blockchain Peer. In particular, the AIV discloses the entirety of the attestation evidence as part of the VCs signed with the underlying hardware-based key (e.g., DAA key) to the Blockchain Peer.

The confidentiality of this evidence is protected by using CONNECT ABE scheme which, as defined for the MEC case, it constructs keys linked to specific attributes; thus, only the entities that possess the defined attributes may create the decryption keys. The process is further protected by the underlying TCB. This ABE-encrypted failed CCAM attestation evidence is uploaded by the Blockchain Peer to the Off-chain storage, primarily due to the size of the information. This encrypted failed attestation evidence is being uploaded the by the Blockchain Peer to the Off-chain storage, where it can be retrieved at any given time utilizing pointers. The pointers are sent from the Off-chain storage to the Blockchain Peer which amends the smart contract with the encrypted pointer and upload the updated information onto the Private Ledger.

The failed attestation CCAM evidence is a useful information for OEMs. OEMs may be granted access, based to their attributes (i.e., ABAC) to the Blockchain Peer, utilizing VCs. If the ABAC is successful, the Blockchain Peer is able to extract the pointers from the Private Ledger and consequently access the encrypted evidence stored at the Off-chain storage. This encrypted evidence is sent to the OEM by the Blockchain Peer. The OEM may now decrypt the evidence, based on Attribute-based decryption keys through the CONNECT TCB. As defined for the MEC case, similarly for the vehicle the decryption keys are generated by the CONNECT TCB only when specific system attributes are present and verified. The ABE decryption key is calculated based on these OEMs attributes as well as the VCs. After the decryption of the evidence takes place, OEM is able to analyse the evidence and identify possible CCAM service-related risks. The next step is for the OEM to upload the identified vulnerabilities to the Blockchain Peer, which will forward them to the Private Ledger.

The TAM will be notified on the inclusion of new vulnerabilities, and after being authenticated to the Blockchain Peer (i.e., through VCs), it will receive the discovered vulnerabilities. These vulnerabilities are leveraged to recalculate the RTL, in tandem with the Risk Assessment (RA) component, for the CCAM service. The new RTL for the CCAM service is sent to the Blockchain Peer and then to the Private Ledger. The TAF Agent will be notified on the updated RTL.

6.8 Digital Twin for Functional Offloading

As discussed in Section 2.2.5 the introduction of **Digital Twins (DTs) within CONNECT can support the offloading processing for both the trust assessment process but also for resource-intensive operations**. The endmost goal is to be able to benefit from the abundance of resources currently available at the edge so as to create a virtual replica of the vehicle that can support its lifetime activities through the execution of tasks and the communication (through the ultra-low-latency edge networking capabilities) of the results that can allow for the efficient and more accurate decision making (specifically in the CCAM context). DTs are gaining significant attention across diverse industries [122], [123], [124], [125]. According to [49], DTs are powerful enough to enable the integration of various technologies such as AI, simulation and metaverse, while addressing transversal concerns like security and real-time considerations within domains such as smart cities and energy. DTs are defined by ISO/IEC 30173 as the "*digital representation of a target entity with data connections that enable convergence between the physical and digital states at an appropriate rate of synchronisation".* In the field of automotive E/E architectures and CCAM, the concept of the Digital Twin showcases an inherent ability to accurately replicate not just the capabilities (vehicle configuration), but also the functionalities of the vehicles or infrastructures it represents. However, it is important to note that there is still a significant absence of convergence in

these architectural approaches. It is within this realm of exploration that the efforts of CONNECT become relevant. Although there are currently standards and frameworks that discuss the various applications of Digital Twins in different domains, there is still a need for a comprehensive and unified architectural perspective to be firmly established. Furthermore, **the concept of a digital twin has not been applied in the domain of trust modelling and trust assessment which is what CONNECT focuses on**.

Figure 35 recapitulates terms used to describe an architecture in ISO/IEC/IEEE 42010 [126]. The first term is **entity of interest**, defined as the subject of an architecture description. A list of examples of entity of interest is provided by the standard: *enterprise, organization, solution, system (including software systems), subsystem, process, business, data (as a data item or data structure), application, information technology (as a collection), mission, product, service, software item, hardware item, product line, family of systems, system of systems, collection of systems, collection of applications*. The standard also points out that interest in an entity is intended to encompass interest in that entity's environment, life cycle, architecture, requirements, design, implementation, and operation. Such interests are captured via aspects, concerns, and stakeholder perspectives. The second term is **environment**, defined as the **context** of surrounding things, conditions, or influences upon an entity. This includes external entities that can have various influences upon an entity, such as developmental, technological, business, operational, organizational, political, economic, legal, regulatory, ecological, and social influences as well as external physical effects such as electromagnetic radiation, charged particles, gravitational effects, and electric and magnetic fields. The standard further points out that a label can be attached as a qualifier to the term environment to identify a particular context within another context, such as development environment, test environment, and operational environment.

Figure 35 - Relation between entity of interest and environment

In terms of CONNECT, it would be translated as presented in Figure 36. The digital twin entity of interest combines the **edge server vehicle twin** (in charge of assistance functions) and the **vehicle** (in charge of target functions) while the target function is the vehicle's trust assessment capability provided by the TAF or any other calculation or application task to be offloaded to the Digital Twin. For instance, CONNECT also envisions to leverage the edge server vehicle twin for supporting the execution of more resource-intensive attestation processes, and especially Control-Flow Attestation (CFA). CFA is the most prominent type of verification of a device's runtime executional behaviour. It operates by monitoring the control-flow graph of a device's execution and then comparing it with a reference value depicting the nominal device behaviour. CONNECT envisions to employ AI-assisted CFA mechanisms where the training of the classification process can be offloaded to the DT while the runtime inference and comparison of the runtime control-flow graph with the reference value/model can be done locally at the vehicle. This will overcome the current limitations and assumptions on the performance overhead of such advanced mechanisms that hinder their applicability in environments as the ones envisioned in CONNECT.

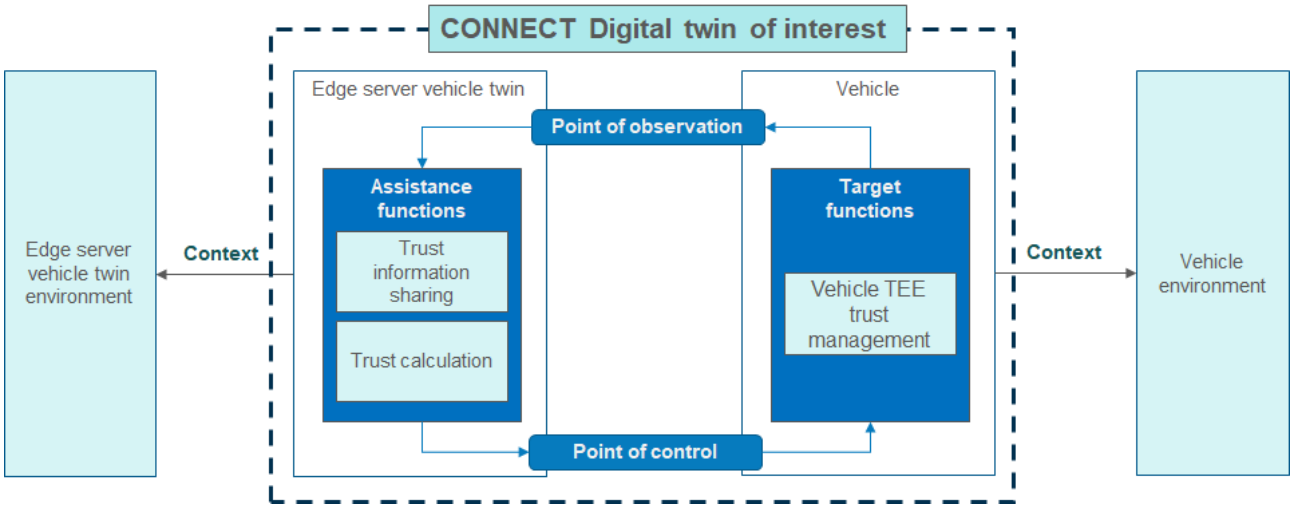

Figure 36 - Digital Twin for TEE trust management offloading

The assistance function consists of the trust calculation (or other security calculation) capability as well as the distributed trust information sharing capability. The environment of the CONNECT digital twin includes the vehicle environment and the edge server vehicle twin environment as well as their secure and authenticated communication. The interactions between the edge server vehicle twin and the vehicle are the following: the edge server vehicle twin is offloading the system by carrying trust (or other security) calculation capabilities. It does so through two interfaces, a point of observation (PO) interface and a point of control (PC) interface. The PO is used to provide real-time information that allows the edge server vehicle twin of calculate trust. **The scope of this twin is to facilitate the trust assessment process by outsourcing this task to a MEC where the TAF-DT is expected to run inside a TEE, keeping the state confidential**.

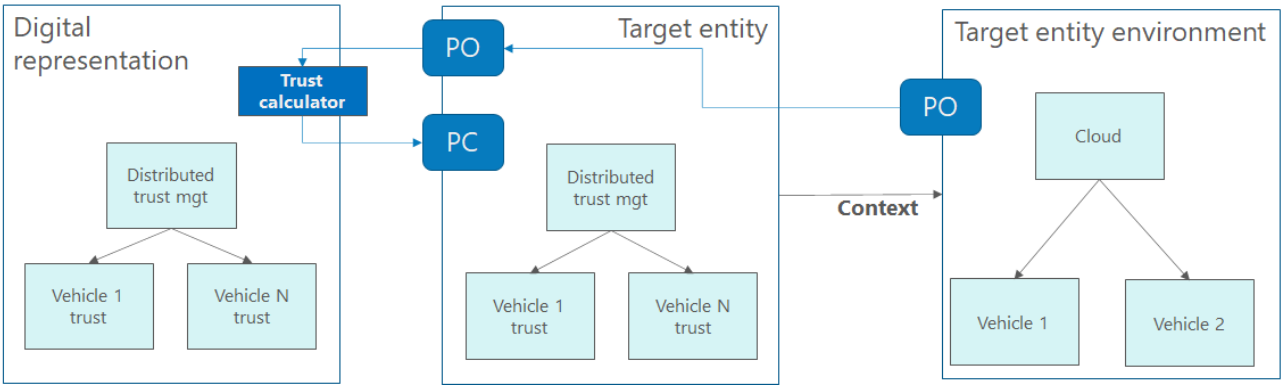

Figure 37 - Digital Twin for Functional Offloading

A blueprint of the architectural framework is depicted in Figure 38. It comprises a physical VEC network, managing distributed vehicles and RSUs, a digital twin network layer, and a layer dedicated to vehicular applications. The digital twin network layer is responsible for tasks such as model construction (vehicle model, RSU model, network model), item mapping, and strategy optimization. As aforementioned, the primary aim of this replication endeavour is to enable for the vehicles to better facilitate the resources that are currently available closer to the edge and, thus, offload resource-intensive security and trust calculations in order to not impact the operational (safety) profile of the physical vehicle entity.

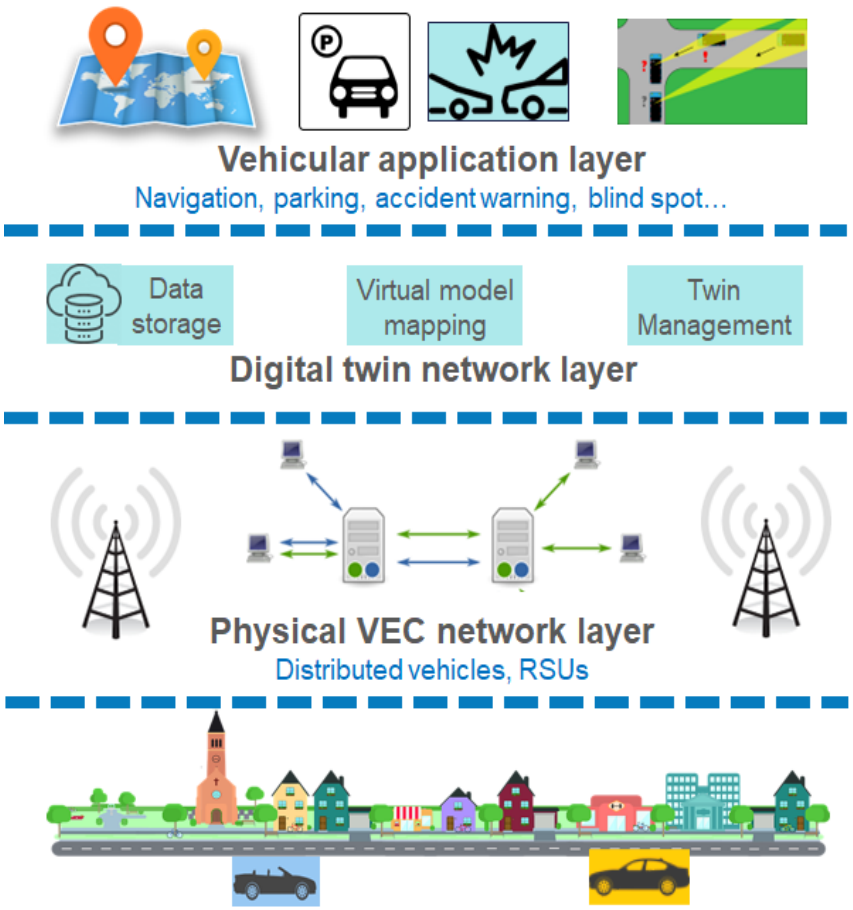

Figure 38 - Digital twin based on vehicle edge computing

The CONNECT research project will demonstrate distributed trust management in a configuration that does not involve handover, as this capability is not yet available in a mainstream infrastructure. The demonstration will include the management of digital twin registration and activation within a given base station cell. In addition, it should be noted that CONNECT will also consider security and privacy implications when offloading tasks to the DT so as to make sure that there are no additional privacy implications posed to the physical vehicle when sharing such information. With its Trust Assessment Framework – Digital Twin (TAF-DT) as developed and described in WP3 and its deliverables, CONNECT pioneers the concept of a digital twin in the context of trust assessment.

7 CONNECT Use Cases

7.1 High-Level Introduction of the CONNECT Use Cases Towards Cooperative Automated Driving

Road traffic crashes pose a significant global challenge, being classified as the sixth most prominent cause of death on a global scale [127]. The aforementioned alarming statistic underscores the pressing necessity for novel approaches aimed at bolstering road safety and mitigating the tragic loss of innumerable lives. In light of this urgent matter, the advent of V2X technologies has presented novel opportunities for directly addressing this concern. V2X technologies facilitate the exchange of information between vehicles and their surroundings, encompassing other vehicles, infrastructure, pedestrians, and the overall road environment. The potential of enabling instantaneous **communication and collaboration between vehicles** and their environment has significant implications for transforming road safety and mitigating the occurrence of accidents. V2X technologies have the capability to enhance the awareness of vehicles beyond their immediate surroundings, **enabling both drivers and automated systems to make well-informed decisions, identify potential risks, and implement precautionary measures**. V2X technologies offer a significant opportunity to transform the road safety landscape and mitigate the impact of road traffic accidents worldwide.

The progression towards attaining higher levels of vehicle automation is primarily driven by the ongoing advancements in perception, positioning, control, and communication technologies [128]. The evolutionary process of developing and seamlessly integrating automated driving functions is categorised into distinct deployment phases known as "days", as mentioned in Section 2 of the present deliverable. The progression moves from initial awareness (day 1) to enhanced environmental-sensing capabilities (day 2), ultimately aiming to achieve cooperative automation (day 3) [129].

Despite advances in enabling technologies, communications have the greatest potential to improve vehicle automation efficiency and road safety. The collaborative framework of ITS allows all participants to share valuable information [130]. The fifth generation (5G and beyond) cellular networks have unmatched capabilities in connecting countless devices and vehicles. These networks set a new standard for reliability and scalability, operating at gigabyte data rates. The rise of MEC [131] has followed these advances. MEC's ability to host services in close proximity to users, especially vehicles, could make it faster than cloud-based services. This proximity ensures quick access to locally relevant data, reducing latency and transportation infrastructure load. This setup also promises significant energy conservation, supporting modern transportation systems' sustainability goals.

Along these lines all involved (vehicle-to-infrastructure, vehicle-to-network) communications, utilising the 5G channels and/or MEC interactions, are expected to cope with/assist critical (automated) driving decisions that directly relate to mobility efficiency and road safety. However, that would require the involved communication technologies together with the accommodated data to meet demanding security (see Section 8.1.2) requirements. Even more challenging requirements (see Section 8.1.1) arise for (the quantification of) trust and trustworthiness of exchanged automotive messages that may be subject to node (vehicle) misbehaviours; the various network operators (MNOs), the different vehicle platforms (OEMs), the plethora of edge computing and automotive-application providers stress the need for interoperability but also call for trust considerations (of potentially increased complexity) across diverse domains.

CONNECT invests effort to address those challenges focusing on three carefully selected use-cases. Their selection serves the purpose of covering functionalities that may appear up-to Day 3 applications and at the same time exhibit a certain level of complementarity (with respect to a number of dimensions):

- ✓ inter- and intra-vehicle trust management and secure communication challenges are explored.

- ✓ both means of simulation studies and real-world experiments (demos) are utilised.
- ✓ a varying level of information and (edge computing) infrastructure support is employed.

Table 8 provides a synopsis of the characteristics of each CONNECT use-case.

Table 8 - CONNECT Use Cases Characteristics

ID	Name	Description	Category	Included Communication Patterns	CONNECT Security & Privacy Enablers
1.	**Intersection Movement Assistance (IMA) & Misbehaviour Detection**	Combine heterogeneous sources of evidence to assess the trust indicator of incoming kinematic data. The UC has two sub-scenarios: **#1** V-TCs are sent to the MEC due to inability of the ego to process it. **#2** V-TCs are evaluated within the vehicle.	Safety and Traffic Efficiency [Day 1, Day 2]	V2V, V2N	TAF, TCH, TEE Guard, MEC for #1,
2.	**Cooperative Adaptive Cruise Control (CACC)**	For keeping a safe distance to the vehicle in front, the C-ACC bases its decisions on various data items received from in-vehicle sensors, as well as from other vehicles.	Safety and Traffic Efficiency [day 2, conditionally day 3]	V2V	TAF, AIV, TEE Guard, TCH
3.	**Slow-moving Traffic Detection (SMTD)**	Sensor based object info about non-connected road users is sent by an ITS-S to a central ITS-S.	Safety and Traffic Efficiency [day 1]	V2N	TAF, TCH, TEE Guard, MEC

7.1.1 Increasing Trust in CCAM

In the highly inter-connected and perplexed environments like those envisioned in CCAM, the importance of a framework capable of giving a trust opinion on data collected from the various sources, both internal and external to the vehicle, is crucial. This requirement emerges for two reasons. Firstly, it permits the vehicle to harness the whole spectrum of the acquired data, leveraging all available sources. This inclusion is crucial for full situational awareness and a holistic grasp of the environment. Secondly, and more critically, such a framework becomes the foundation for well-informed decision-making that is based on trust. Without such a trust assessment framework, the risk receiving and basing decisions on false or manipulated kinematic data increases, potentially leading to unsafe or inefficient outcomes. Establishing trust in the exchanged kinematic data establishes a foundation for collaborative and secure operations, fostering a trustworthy environment where entities can share critical information, ultimately contributing to the safety and efficiency of complex systems. By analyzing the trustworthiness and authenticity of incoming data, the system may make decisions with an increased level of confidence, leading to improved safety, efficiency, and overall effectiveness in navigating the obstacles offered by the dynamic and linked landscapes of CCAM.

In what follows all exchanged info based on which TA is performed, we name this TCs. The explicit definition of TCs and its structure is defined on D5.1.

7.2 Intersection Movement Assistance & Misbehaviour Detection

The first Use Case of CONNECT concerns the **Intersection Management Assist (IMA) application**, a critical component in Intelligent Transportation Systems. An ego vehicle approaching an intersection, utilizes CAM messages to predict the trajectories of other vehicles, identifies potential collision zones, and issues timely warnings to the driver as collision probabilities reach predefined thresholds. More specifically, the vehicles in this Use Case employ a Local Dynamic Map (LDM) to predict the positions of dynamic objects in the intersection, primarily relying on information from CAM messages. The LDM, serves to store comprehensive information about the environment, encompassing both static and dynamic data. Addressing challenges related to legacy vehicles without communication capabilities, the introduction of the Collective Perception Service (CPS) complements CAM messages by periodically broadcasting Collective Perception Messages (CPMs). These CPMs enhance the LDM by providing information about dynamic objects perceived by the ego vehicle's on-board sensors, ensuring awareness of legacy vehicles.

The goal of this Use Case is to use the CONNECT's trust assessment framework to evaluate the **trust in the incoming V2X data before employing them in the IMA's decision process**. This requires the development of a framework capable of dynamically combining heterogeneous information relating to different properties of interest. To meet these challenges, the IMA application necessitates a dynamic trust assessment system. CONNECT explores two reference scenarios for such IMA trust assessment. In the first one, **vehicles exchange standard V2X messages and create their own view of the intersection based on the kinematic data they receive**. At the same time, they send their trustworthiness claims in separate messages to the MEC. The MEC takes over the role of interpreting these claims, assess the trustworthiness of each node and disseminate the results back to the vehicles. In the second scenario, **vehicles integrate the trustworthiness claims inside their CAM and CPM messages, so each vehicle can interpret the trustworthiness of received V2X messages locally**. The MEC has a different role here, which is closer to the Intersection Manager: it collects the kinematic data from all the CAM and CPM messages and builds a consolidated perception of the whole intersection, which then is broadcast back to the vehicles by the MEC. In that way, vehicles can complement their own local view with the consolidated view of the MEC. In both of these scenarios, the MEC

The V2X landscape is gradually evolving. From day 1 to day 3+ applications, the paradigm shift is beyond noteworthy. The era of AVs introduces many novel possibilities, but it is further prominent to security attacks. Malfunctions are not unknown to modern vehicles, even to those that support V2X applications. For example, the authors in [132] mention an incident on Tesla Model Y, that took place on November the 13th of 2022 where video footage demonstrates the vehicle malfunctioning by speeding through the streets of a Chinese city and eventually killing two people.

Within the complex network of modern cities, where cars, smart sensors, road infrastructure, and various entities interact, the situation becomes intricate. This network facilitates the exchange of data from multiple sources, providing vehicles with valuable information to prevent accidents. However, it also exposes vehicles to potentially malicious or false data that could lead to misinformed decisions. Therefore, ensuring the protection and trustworthiness assessment of the information used for decision-making is paramount.

The primary objective of misbehaviour detection in V2X systems is to **identify and address instances of misbehaviour, which can span a broad spectrum of behaviours, including data injection, impersonation, message manipulation, and denial of service**. The objective is to evaluate the reliability of data, entities, and interactions in the V2X ecosystem. Misbehaviour detection in C-ITS can be classified into three main categories: node-centric detection, data-centric detection, and hybrid, while the techniques for the detection of an anomaly vary from Machine Learning (ML) based, to probabilistic and statistical [133]. Misbehaviour detection leverages the notion of PKI for authenticity of the exchanged messages while it also enables the revocation of certificates in the case of a malicious event.

Among the various safety applications enabled by V2X communications, the Intersection Movement Assist (IMA) application [134] offers assistance to vehicles wanting to cross an intersection. An application similar to IMA and called Intersection Collision Risk Warning (ICRW)

is described in ETSI TS 101 539-2 v1.1.1 [135] and referred to as Intersection Collision Warning (ICW) in [136]. **The primary objective of both the IMA and ICRW applications is to provide drivers with prompt notifications regarding potential collisions with other vehicles in the intersection, enabling them to take appropriate corrective actions**. The main difference is that the ICRW application assumes, in order to operate, the presence of a connected traffic light or roadside unit, whereas the IMA application can operate at any kind of intersection.
Within the CONNECT Trust Assessment Framework (TAF), message exchange is crucial for evaluating incoming kinematic data, especially from sources that do not have inherent trust. This framework acts as a crucial element, allowing the ego vehicle to carefully evaluate and determine the reliability of the data it receives. The IMA scenario in CONNECT is examined in two variations; the first leverages the MEC while the second does not employ the MEC. In parallel, while in the first case the Trustworthiness Claims (TCs) are sent directly to the MEC, in the second case, the TCs are part of the CAM/CPM messages that are exchanged between the vehicles (i.e., hence we refer to them as T-CAM and T-CPM messages).

7.2.1 "As-Is" Scenario

An example of how the IMA application works is provided in Figure 39 [134]. The ego vehicle, in blue, is approaching the intersection. The ego knows the geometry of the intersection and knows the position and kinematic information of the red vehicle thanks to its CAM messages. The ego predicts the possible trajectories of the red vehicle and identifies the possible crash zones. As the red vehicle continues broadcasting CAMs, the ego learns which trajectory is taken, and continuously estimates the probability of collision in the identified crash zones. As the collision probability reaches a threshold the application issues a warning to the driver.

This simple example is already indicative of the important fact that the efficacy of the IMA application, and ultimately its safety, depends on the ability of the vehicle to assess, dynamically and in real time, the level of trust in the data contained in the incoming V2X messages.

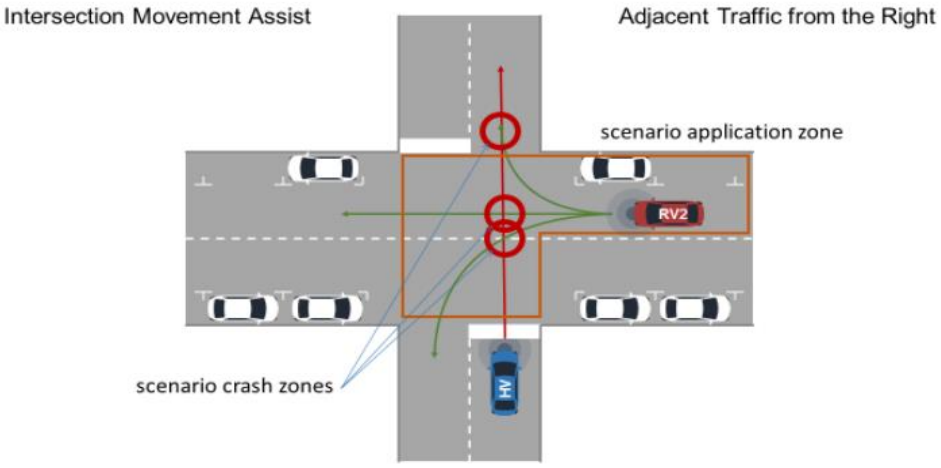

Figure 39 - Example of IMA scenario [134].

Figure 40 describes the radio interfaces involved in the "As-Is" scenario as detailed in Section 7.2.2. The connected vehicles exchange messages on the V2V interface via direct (short-range) communication e.g., the ITS G5 radio interface or the C-V2X sidelink. They are provided pseudonym certificates by the vehicular PKI, which they use to assure authentication and communication integrity in privacy-preserving way. They support the Cooperative Awareness Service [137] and they

periodically broadcast Collective Awareness Messages (CAMs) containing their position and kinematic state.

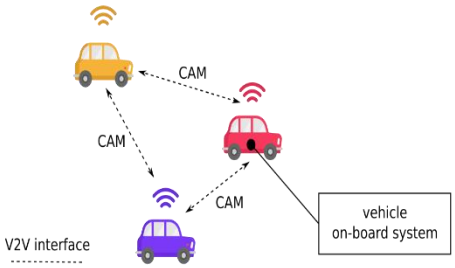

Figure 40 - Radio interfaces in the "As-is" scenario.

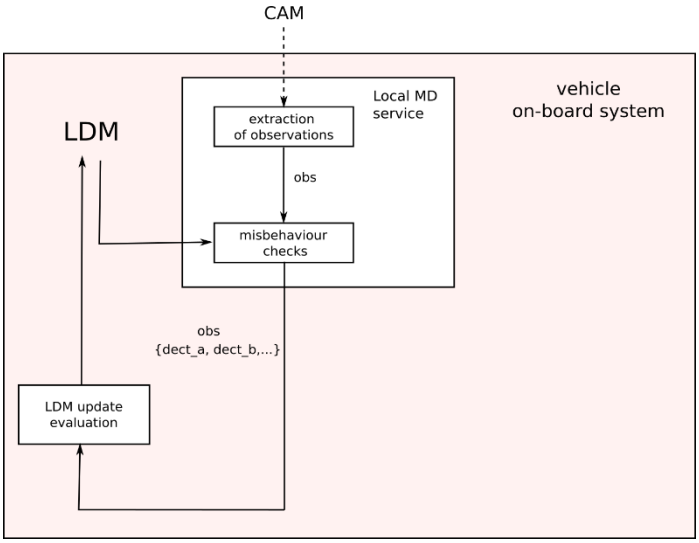

Figure 41 - Functional description of the vehicle

on-board system for the "As-is" scenario.

Figure 41 presents the functional description of the facilities of the ego on-board system that supports the IMA application. The IMA predicts the position of the dynamic objects in the intersection leveraging the information stored in the Local Dynamic Map (LDM), which records the kinematic states of the other connected vehicles. These are known due to the received CAM messages. The correctness of the data in the LDM, hence, depends on the correctness of the data in the received CAMs which, in turn, highlights the need for integrity guarantees on the source of this data (i.e., neighbouring vehicles).

The fact that a transmitting node is authorised within the system since it utilises a valid vehicular PKI certificate to sign their messages does not provide an assurance regarding the correctness of the kinematic data in the message [138], [139]. Consider the following example from Figure 39: Assume that the blue vehicle broadcasts a CAM in which it falsely declares to be at the intersection's centre; since the blue vehicle correctly signs their V2X messages, the incorrect position of the blue vehicle is recorded in the LDM of the red vehicle, and the IMA application triggers a sudden braking. A **misbehaviour event** is operationally defined as the occurrence of erroneous kinematic data transmission within a V2X message. In other words, it refers to the transmission of data that significantly deviates from the actual physical reality, surpassing the predetermined acceptable margin of error established by the industry [138]. A misbehaviour event can arise from either intentional action, such as the data manipulation attack mentioned earlier, or unintentional occurrences, such as a malfunctioning sensor.

The efficacy of the IMA application depends on the ability of the on-board system to process correct kinematic data. For this reason, the vehicle is equipped with a Local Misbehaviour Detection service, which serves the purpose of detecting incoherence or discrepancy in the data that may be the sign of a misbehaviour event. This allows to mitigate the occurrence of the storage of incorrect data in the LDM.

It should be underlined that in this specific scenario, the **traditional** vehicle internal topology is considered, to cover the trust related needs of current services (i.e., TAF been executed at the MEC side).

7.2.2 Communication Interfaces and Messages in the context of Intersection Movement Assistance

Building upon the main stakeholders of the CCAM landscape, as identified in Chapter 4, in what follows we mention specific components and roles that we consider in the specific UC. The fundamental enabler of the IMA application on the vehicle is V2X communications, i.e., the ability to

exchange messages with other vehicles and the infrastructure. In this section, we briefly introduce the messages and radio interfaces supporting the IMA application.

- **Vehicular PKI:** To participate in V2X communications in C-ITS, nodes must enrol in the Vehicular PKI (ETSI TS 102 940), which provides credential and cryptographic material for authenticating vehicles. A permanent certificate is issued at commissioning, proving node legitimacy and long-term identity. V2X messages are broadcasted in plaintext, with cryptographic signatures for authentication. To protect sender privacy and prevent tracking, nodes use ephemeral pseudonym certificates. This strategy is mandatory in Europe and will be used within CONNECT to ensure privacy-preserving broadcasts and support trust assessment functions.
- **V2X Messages:**
  - ✓ **Local Perception:** The Local Perception (LP) is a list of dynamic objects perceived by a vehicle's on-board system, encoded for inclusion in outgoing CPMs. It can be assimilated to a self-issued CPM and is available to the Local MD Service when the vehicle supports the CPS.
  - ✓ **Observation:** An observation is a kinematic description of a single dynamic object, including its identifier, reference time, and position. A CAM contains a single observation, while a CPM contains several observations.
  - ✓ **MR message:** The Misbehaviour Reporting Service (MRS) is a standardised system that encodes Misbehaviour Reports (MR) to notify the receiver of the activation of a set of MD checks. The MR contains the triggering V2X message, the identity of the activated detectors, and related information to verify the activation.
- **Radio Interface**
  - ✓ **V2V**: The V2V radio interface is used by vehicles to broadcast messages to neighbouring stations within radio range, based on ITS-G5 or cellular sidelink technology. It is used for CAM and CPM messages exchange in the IMA use case.
  - ✓ **V2N**: The V2N radio interface connects vehicles to mobile base stations using mobile network radio access, allowing both uplink and downlink communication. In IMA use cases, the V2N uplink uploads MRs, vehicle-TCs, T-CAMs, and T-CPMs, while the V2N downlink disseminates Node Trustworthiness Messages (NTMs) and geo-CPM.

7.2.3 IMA Scenario Needs from CONNECT

The efficacy of the IMA application in enhancing driving safety, significantly relies on the correctness of its decision process mechanism, whose foundation is in the requirement to use dependable and consistent data. However, this presents a challenge, owing to the fact that the V2X data originates from sources, including also other neighbouring vehicles, lacking inherent trust.

As a result, it becomes **essential to introduce mechanisms to evaluate the trust in the incoming V2X data before employing them in the IMA's decision process**. This requires the development of a framework capable of **dynamically combining heterogeneous information relating to different properties of interest**. This may include, for instance, the verification of the message integrity, the appraisal of the plausibility of the kinematic data, as well as the verification of the operational integrity of the node originating the data.

Most importantly, the ability of combining heterogeneous trust evidence is essential to enable decisions on the trustworthiness in kinematic data (included in the exchanged CAM/CPM messages), **even when data confirmation through redundancy observations, of the data itself (i.e., repeated measurements of the same physical phenomena from different sources/vehicles), is not available**. In what follows, this inherent challenge is further elaborated, in the context of CONNECT, as one of the main hurdles to overcome that can unlock the full potential of the Misbehaviour Detection service in such complex ecosystems.

In the “As-Is” scenario, the IMA application is based only on CAM messages from other connected vehicles. One major problem is the potential presence, in the intersection, of dynamic objects without communication capabilities (legacy vehicles), who remain invisible in the LDM and, hence, to the IMA application. This issue is mitigated by introducing the Collective Perception Service (CPS) on connected vehicles. In addition to CAM messages, a connected vehicle periodically broadcasts

Collective Perception Messages (CPMs) containing information about the dynamic objects that it perceives thanks to its on-board sensors (see Section 7.2.2). The received CPMs contribute towards enhancing the scene description in the LDM and make it possible for the ego to be aware of the presence of legacy vehicles. As a consequence of supporting the CPS, the ego has access to its Local Perception structure, i.e., the scene description produced by its own on-board perception system, which contributes to the LDM as well. **Within CONNECT we assume that the IMA application is supported by both the Cooperative Awareness and Collective Perception Service, and that the LDM is populated by data coming from CAMs, CPMs and Local Perception.**

On the ego, CPS support has an impact on the Local Misbehaviour Detection service and on the LDM. The ego may now receive independent observations of the same object from distinct V2X nodes (e.g., two vehicles send CPMs containing information about the same perceived object; a vehicle sends a CAM and is simultaneously described in the CPM of another vehicle). This allows to define several new misbehaviour checks, and in particular the class of **redundancy checks which are activated in presence of discrepancies between the descriptions of the same scene by two distinct V2X nodes**. A more detailed description of the evolution of the Local Misbehaviour Detection service and of the LDM is provided in Section 4.2.3.

The data in the Collective Perception Service comes from the on-board perception system of the vehicles. The performance of the on-board sensors in general depends on the quality of the equipment but may also be affected by the environmental conditions (think for instance of the different detection capabilities of a camera sensor in daylight or at night; or a Lidar's perception range reduction in case of bad weather). The occasional activation of some misbehaviour detector due to an unintentional error may become a likely event. In the current instantiation of an IMA service, these messages are implicitly considered as untrustworthy and excluded from the LDM. **This policy is, however, too restrictive and risks erasing too much information, too prematurely, compromising the ability of the on-board system to provide a consistent reconstruction of the scene for the IMA application.** As an example, think about a sensor whose reading is temporarily perturbed because of the weather, providing a series of inconsistent consecutive speed measurements of the same physical object in successive CPMs: *These would activate consistency misbehaviour checks at the receiver, triggering exclusion of the perceived object from the receiver's LDM, and eventually making the IMA application temporarily blind with respect to the presence of the physical object*. Providing the IMA application with the indication of the existence of the object in the area, albeit with limited trust on the precision of its kinematic state, would much better serve the purpose of the application.

Thus, the first need towards enhancing the knowledge based on which the IMA can rely its decision-making process is the definition of a trust indicator and the development of a framework for the assessment of the trust in incoming V2X kinematic data, able to function in real-time and to account for the high dynamicity of the scene. **This is particularly challenging, because the same trust object in the scene will typically be part of multiple trust relationships, and each relationship may have a different evolution in time**. This will allow the ego to record observations in the LDM, along with their individual trust indicators. Finally, the fusion of LDM observations referring to the same physical object will allow the ego to obtain the consolidated view of the scene, where the inclusion of a physical object will be decided as a function of the trust indicators on all its observations in the LDM. The IMA application will consume the consolidated view of the scene.

The second need of the IMA application **is the ability to combine heterogeneous sources of evidence to assess the trust indicator of incoming kinematic data.** The misbehaviour checks activation pattern on the target observation is the first, natural evidence (notice that it is already considered in the "As-Is" use case, albeit less explicitly, as evidence to determine whether data were trustworthy enough to be recorded in the LDM). The activation of a check on the target observation may decrease the observation's trust; on the other hand, if it corroborates another observation already in the LDM, it may increase the trust in both. This works well when redundant observations of the same physical object are available. However, in highly dynamic scenes, where the appearance of new nodes and objects is frequent; or in low-density scenarios, or scenarios with prevalence of legacy vehicles, the desired level of redundancy in the observations may not

be attainable, thus, **resulting in less rich information available for making a trust decision**. The consequence is that the misbehaviour detection checks alone do not provide enough evidence for a robust assessment. For this reason, when trying to evaluate the trust indicator of a specific kinematic observation, it is essential to be able to integrate in the assessment process evidence pointing to the satisfaction of properties distinct from the veracity of the kinematic observation itself. These properties may refer, e.g., to the integrity of the communication and integrity of the operational state of the device originating the transmission of the message.
To resolve these two pressing needs, within CONNECT, we are going to consider two reference scenarios for the IMA application. In the first scenario we wish to illustrate how such a dynamic trust assessment system may be integrated, with minimum level of intrusion/disruption to the current standardised architectures. The trust assessment framework of this scenario is deployed while keeping communications over the V2V radio interface compliant with the current standardisation landscape, as specified in Europe by ETSI. In the second scenario we will explore solutions leveraging extensions of currently standardised V2V messages tailored to provide better support for trust assessment.

7.2.4 "To-Be" Reference Scenario #1: MEC-based V2X Node Trustworthiness Assessment Service

The aim of the ego vehicle is to assess the trustworthiness levels associated with the observations contained in received CAMs and CPMs, and in the LP, and to record them in the LDM, so that they can be used to produce a trustworthy consolidated view of the scene, to be consumed by the IMA application. The rationale behind Scenario #1 is to keep the communication over the V2V radio interface compliant with the current standardisation landscape, as specified by ETSI. On the V2N interface, vehicles can communicate with the MEC, which provides services to assist the vehicle in its task.

This reference scenario is enabled by the following CONNECT components: The **Trustworthiness Assessment Framework (TAF)** is the CONNECT component whose functionality is to provide the assessment of the trustworthiness level of an item, by processing heterogeneous evidence belonging to possibly distinct semantic levels. The TAF may be deployed in the vehicle as well as at the MEC. Another essential element is **CONNECT's Trustworthiness Claims (TCs)** used as one of the trust sources based on which the dynamic trust assessment is performed. More specifically, and as described in Section 6.1.1, TCs are the means for an entity to provide evidence that can be used by the TAF for assessing their level of trust, in a verifiable and privacy-protecting manner. More specifically, TCs are constituted by evidence (signed by a trust anchor instantiated in the target node) about the state of a given (node and data) component, thus, providing the necessary proofs on the runtime status of those properties of interest based on which a detailed trust assessment needs to be performed, such as, e.g., device integrity.

The TAF is able to exploit **Vehicular-TCs (V-TCs)**, i.e., the TCs generated by a vehicle, as a trust source in order to assess the trustworthiness level of the V2X-node with respect to its ability of transmitting correct kinematic data in V2X messages. Details regarding these messages are available in section 4.2.3 of the present deliverable. The trustworthiness level of the emitting V2X-node can in turn be used by the TAF, when it assesses the trustworthiness level of the kinematic data contained in a received V2X message.

In this Reference Scenario #1, the communication over the V2V radio interface is compliant with the current standardisation landscape, so vehicles cannot directly exchange V-TCs. However, V-TCs, which are generated on the vehicle in a periodic fashion, are periodically uploaded to the MEC. The MEC hosts V2X-Node Trustworthiness Assessment Service (NTS). This service interprets the received MRs and V-TCs as trust sources to assess the trustworthiness levels of known V2X-nodes. V2X nodes, which are identified by their pseudonym PKI certificate, are known if they uploaded V-TCs or if they are the object of a MR. The MEC service disseminates the trustworthiness levels on known V2X-nodes in bulk, encoding them in V2X-Node Trustworthiness Messages (NTMs), periodically broadcast on the V2N downlink. NTMs also contain MEC-TCs produced by the MEC, concerning its ability to correctly deliver the service, i.e., to include reliable data in the NTMs.

The TAF on the ego vehicle uses the MEC-TCs to assess the level of trust of the virtualized infrastructure where the service is running, so as to attest that received NTM messages from the MEC have been calculated in a correct manner. The TAF on the ego then combines the trustworthiness value of the NTS service by the MEC with the trustworthiness values in the received NTMs, in order to assess the local trustworthiness values for the V2X-nodes. Finally, when an incoming V2X message is received, the kinematic data it includes is assessed by the TAF on the ego, which can use both the activation pattern of the misbehaviour detectors and the local trustworthiness level on the V2X message emitter as trust sources. The fact that the NTS service is provided by the MEC, which serves all vehicles in the same geographic region, ensures that with high probability at the moment of the reception of a message the ego TAF will be able to rely on the trustworthiness value for the emitter V2X-node. This ensures that a trustworthiness assessment can be made also in case of scarce evidence from the misbehaviour detection service (because, e.g., no previous data had been already received by the same V2X-node).

The defining feature of this reference scenario is the **inability of the ego to observe the V-TCs**, produced by the senders of the V2X messages, which does not enable the ego TAF to directly exploit them as a trust source in the process of assessing the trustworthiness levels of the kinematic contained in the exchanged V2X messages (due to aligning with the existing V2X messages structures as defined by the ETSI standards where they limit the size of trust related information that can be added as part of the security header). Hence, the need for the ego TAF to collaborate and rely upon the assessment made by the TAF at the MEC on the trustworthiness levels of V2X message sources. This type of a distributed trust assessment process (referred to as Federated TAF in CONNECT [D3.1]) enables the two entities to collaborate to exploit all available evidence towards the assessment of the trustworthiness level of the kinematic data.

In this reference scenario the following messages are considered, in addition to what described in Section 7.2.2:

- **Vehicle Trustworthiness Claims (V-TCs)**, as defined in 4.2.4, V-TCs are generated by the Trustworthiness Claims Handler (TCH) using Attestation Integrity Verification (AIV) output. They provide harmonised evidence about a vehicle's ability to construct and transmit correct V2X messages, such as system integrity, communication integrity, and system safety. V-TCs are signed using the internal Attestation Key of the trust anchor, allowing linking V2X messages with them.
- **NTM messages:** The **V2X Node Trustworthiness Message (NTM)**, as defined in 4.2.4, it contains a list of pseudonym PKI certificates, and their trustworthiness levels, as attributed by the V2X Node Trustworthiness Assessment Service (NTS) running at the MEC.
- **MEC Trustworthiness Claims (MEC-TCs),** as defined in 4.2.4, they guarantee the integrity of the virtualized environment, ensuring the service's configuration remains unchanged, thus confirming the computing support used.

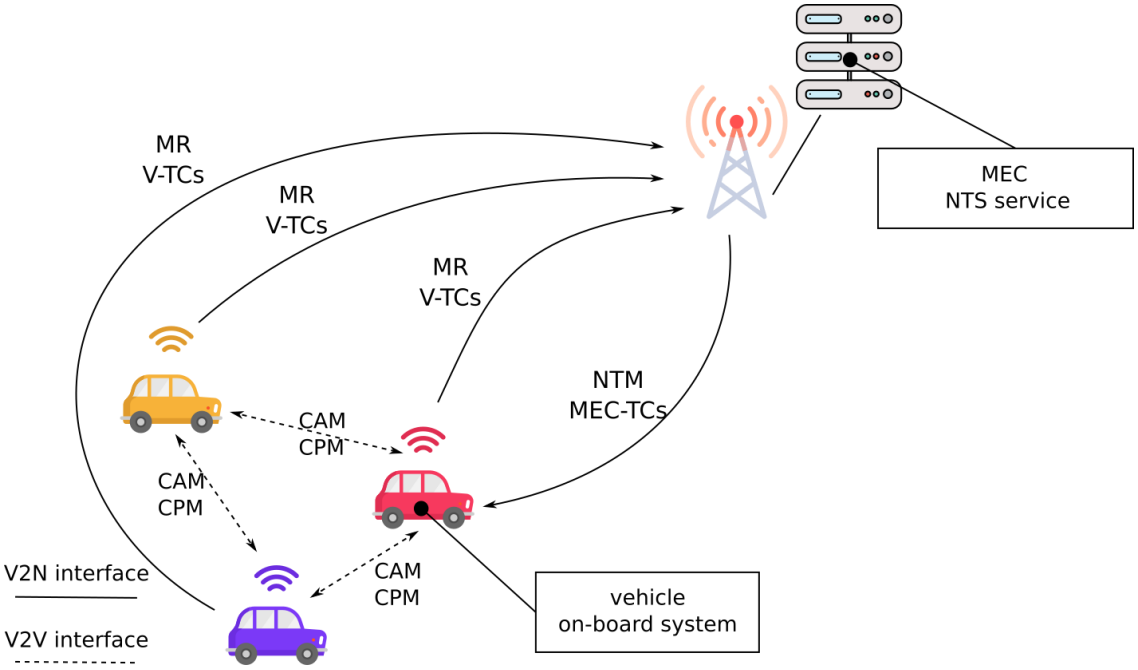

Figure 42 - Radio Interfaces in the MEC-based V2X Node Trustworthiness Assessment (#1)

Figure 42 represents the radio interfaces considered in this scenario. The exchange of messages on the **V2V radio interface** is as in the "As-Is" scenario: vehicle exchange only CAM and CPM messages, in compliance with ETSI TS 103 900 [137] and ETSI TS 103 324 [100], respectively (see Section 7.2.2 and Chapter 4). We assume that connected vehicles possess credentials, in the form of pseudonym certificates, delivered by the Vehicular PKI. The connection to the MEC is made through to the **V2N radio interface**:

- **Uplink**: V2X nodes share with the MEC Misbehaviour Reports (MR) according to ETSI [89] (see Section 7.2.2.) and V-TCs. V-TCs are uploaded in conjunction with MRs in a periodic fashion based on predefined policies dictating the “freshness” of the evidence needed for the successful execution of the trust assessment process.
- **Downlink**: the MEC disseminates to the vehicles NTMs and MEC-TCs (see Section 7.2.2) relative to the Source Trustworthiness Service at the MEC. The MEC-TCs are embedded in the NTMs messages. The NTMs have geographical relevance, and they are broadcasted to V2X nodes targeting their localization. This may be achieved using beamforming capabilities of the base station.

Figure 43 presents the functional description of the on-board system and of the system at the MEC providing the MEC-based Node Trustworthiness assessment Service (NTS). We start by describing the functionality at the MEC Service Provider and we describe the functionality of the vehicle's on-board system. We conclude detailing the LDM and the Local Misbehaviour Detection Service at the vehicle.

MEC NTS Provider: The CONNECT Trustworthiness Assessment Framework (TAF) is the core component used to deliver the MEC-based NTS. At the MEC the **Active V2X Node Directory (AND)** is a database maintaining the list of active known V2X nodes. V2X nodes are identified in the AND by their vehicular PKI pseudonym certificates. A V2X node is known to the MEC either because it is a sender of a MR (and hence includes V-TCs by itself); either because it has uploaded a V-TC message to the MEC; or because one of its CAMs or CPMs are included as evidence in a MR received by the MEC. The AND stores a trustworthiness level associated with each entry. The TAF is responsible for the evaluation of the trustworthiness levels.

MEC NTS Provider: The CONNECT Trustworthiness Assessment Framework (TAF) is the core component used to deliver the MEC-based NTS. At the MEC the **Active V2X Node Directory (AND)** is a database maintaining the list of active known V2X nodes. V2X nodes are identified in the AND by their vehicular PKI pseudonym certificates. A V2X node is known to the MEC either because it is a sender of a MR (and hence includes V-TCs by itself); either because it has uploaded a V-TC message to the MEC; or because one of its CAMs or CPMs are included as evidence in a MR

received by the MEC. The AND stores a trustworthiness level associated with each entry. The TAF is responsible for the evaluation of the trustworthiness levels.

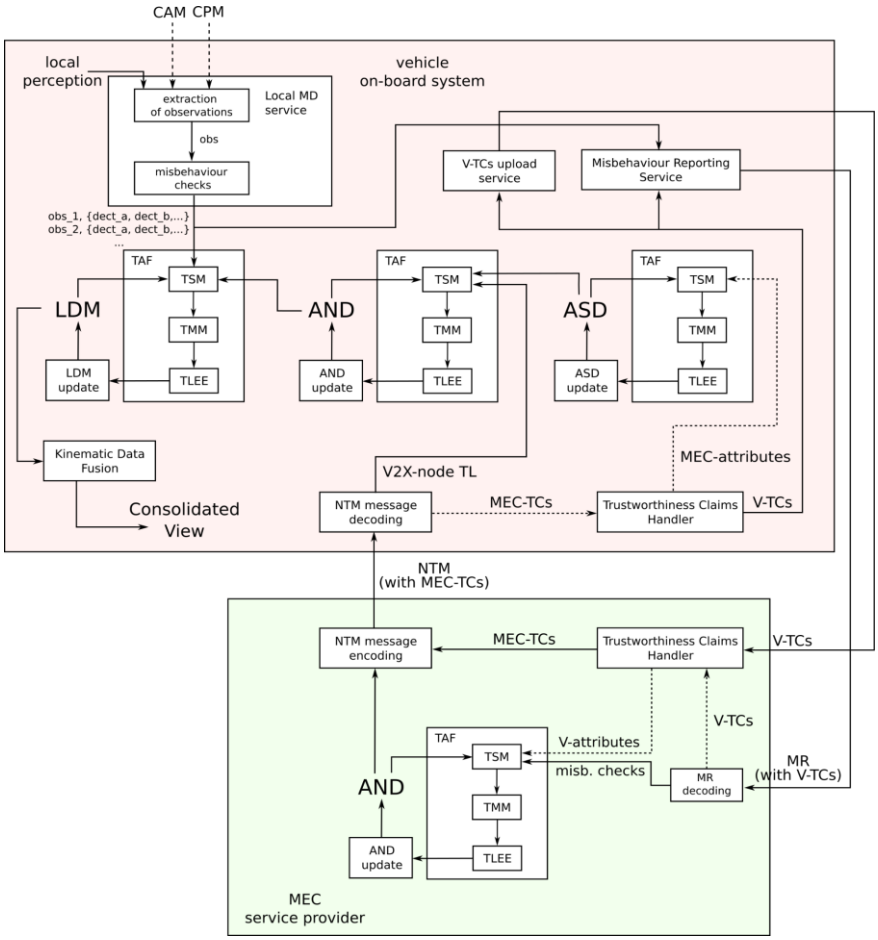

Figure 43 - Functional description of the vehicle on-board system and of the MEC service provider for the service for the MEC-based V2X Node Trustworthiness Assessment Service

The TAF is represented in Figure 43. The Trust Source Manager (TSM), the Trust Model Manager (TMM) and the Trustworthiness Level Evaluation Engine (TLEE) inside the TAF (as described in D3.1) have been drawn for clarity. The aim of the TAF is to output the trustworthiness level of a specific V2X node, by processing new evidence as soon as it becomes available. The relevant trust sources are of two types:

1. V-TCs signed by the target node, attesting the integrity of the system that generates and sends messages; and
2. MRs involving CAMs or CPMs transmitted by the target V2X node, concerning the accuracy of data transmitted in the past by the V2X node.

If the target node already has an entry in the AND at the time when new evidence becomes available to the TAF, the associated trustworthiness level (already been calculated based on previously collected evidence as part of the node reputation) is considered as an additional trust source and processed as an input of the TAF, as well. The output of the TAF then overwrites the trustworthiness value in the AND. Notice that for this reason the content of the AND dynamically evolves in time. In order to avoid an excessive growth of the AND, entries are removed if the associated trustworthiness level has not been updated since a predefined time window. This operation has the effect of periodically removing old pseudonymous identities no longer in use.

The NTS service at the MEC encodes NTMs. Each NTM comprises a list of V2X node identifiers and associated trustworthiness levels, as recorded in the AND at the time of the NTM encoding. The NTM also contains the TCs used by the MEC-based NTS to prove to V2X nodes that it possesses attributes relevant for its correct operation as specified by the ETSI standards for showcasing the level of assurance of the MEC infrastructure where the aforementioned services are instantiated.

Vehicle On-board System: The core CONNECT component in the vehicle on-board system is again the TAF, which is used to assess the trustworthiness levels of several entities.

The on-board system maintains the **Active MEC Service Directory (ASD)**, which is a database of the identities of known and active services delivered by the MEC, with the current associated trustworthiness level. The trustworthiness levels for the MEC-based NTS are evaluated by the local TAF by processing the available evidence, which comes in the form of MEC-TCs transmitted in NTMs.

The on-board system also maintains a local AND. At any given time, a V2X node has a record in the on-board AND only if it has been listed in a received NTM. Notice that this implies the possibility that the ego vehicle receives a CAM or CPM message by a V2X node which does not appear in the on-board AND. The trustworthiness level associated with each entry in the on-board AND is evaluated by the TAF, by considering the trustworthiness level for the target V2X node expressed in the received NTM, and the trustworthiness level associated with the MEC service in the ASD.

The TAF is finally responsible for the assessment of the trustworthiness levels of the observations contained in incoming CAMs, CPMs, and LP. The received observations are stored along with the associated trustworthiness level in the LDM. Notice that this requires extending the concept of LDM as considered in the "As-Is" use case, to accommodate the Actual Trustworthiness Levels (ATL) of each observation. To attribute a trustworthiness level to an observation, the TAF considers as a trust source the current trustworthiness level of the V2X node which encoded the message the target observation was included in, if this is available in the AND. The second trust source is the output of the Local Misbehaviour Detection service, consisting of the activation pattern of the misbehaviour checks on the target observation. The Local Misbehaviour Detection Service with CPS support is described below. In order to perform the appropriate checks (e.g., consistency or redundancy checks), the target observation may be compared with observations that already are in the LDM. The current trustworthiness values of the previous observations are considered as input to the TAF, as well. The output of the TAF may update the trustworthiness levels of all the observations involved in the misbehaviour checks, i.e., the target observation as well as previous observations.

In absence of the connectivity to the MEC, or when the Actual Trust Level of the MEC-bases NTS is not sufficient, the on-board system works in a degraded state, and the TAF only consumes the output of the Local MD service as a trust source, to assess the trustworthiness levels of incoming kinematic data.

7.2.5 "To-Be" Reference Scenario #2: geo-Collective Perception Service

In this scenario we make the assumption that **vehicles are allowed to exchange V-TCs on the V2V radio interface**. We assume that these may be incorporated in CAM and CPM messages and exchanged in a periodic fashion as T-CAMs and T-CPMs. Using V-TCs from neighbouring V2X nodes and the results of the Local MD Service as trust sources to the TAF, the vehicle is able to maintain a local AND containing the current trustworthiness levels on the active V2X nodes it has observed, and to use it in the assessment of the trustworthiness of the incoming kinematic data. With respect to Reference Scenario #1, this allows the on-board system to be more precise in the evaluation of the trustworthiness of the incoming kinematic data, when the MEC is not available; this, however, comes with an increased computational cost for the on-board system.

Moreover, when the MEC service is present, V2X nodes share with the MEC their T-CAM and T-CPM messages over the uplink of the V2N radio interface. The MEC is now able to process kinematic data and to assess their trustworthiness levels, running a similar system to the one deployed on-board. The MEC, whose sources of kinematic data extend well beyond the radio range of the vehicle, may significatively improve the extended perception in comparison to the one available at the ego, as well as the quality of the trustworthiness assessment of kinematic data. The MEC hosts the geo-Collective Perception Service (geo-CPS): its encodes geo-CPM messages, which are then disseminated to the V2X nodes over the V2N radio interface downlink. A geo-CPM contains the MEC view of the environment in the form of a collection of observations, as in a standard CPM, complete with associated trustworthiness levels. The geo-CPM also embeds MEC-TCs relevant with respect to its ability to produce accurate geo-CPMs. The V2X nodes may hence locally assess the

trustworthiness of the MEC geo-CPM service and decide to exploit geo-CPMs to form the local view of the scene which is exploited by the IMA application.

The defining feature of this Reference Scenario is the role of the MEC, which is the enabler for bolstering the extended perception capabilities of the vehicles. In this reference scenario the following messages are considered, in addition to what described in Section 7.2.2:

- **geo-CPM message** as defined in Section 4.2.2, the geo-CPM message is provided by the MEC-based geo-CPS service, which contains a Detected Objects container with kinematic attributes, uncertainties, and trustworthiness level, potentially including MEC-TCs.
- **T-CAM** as defined in Section 4.2.2 CONNECT T-CAMs are regular CAM messages that can accommodate V-TCs, which are included in a station's T-CAMs periodically, ensuring not all stations' T-CAMs include V-TCs.
- **T-CPM** as defined in Section 4.2.2 CONNECT T-CPMs are regular CPM messages that can accommodate V-TCs, which are included periodically in T-CPMs produced by a station, ensuring not all are expected to include V-TCs.

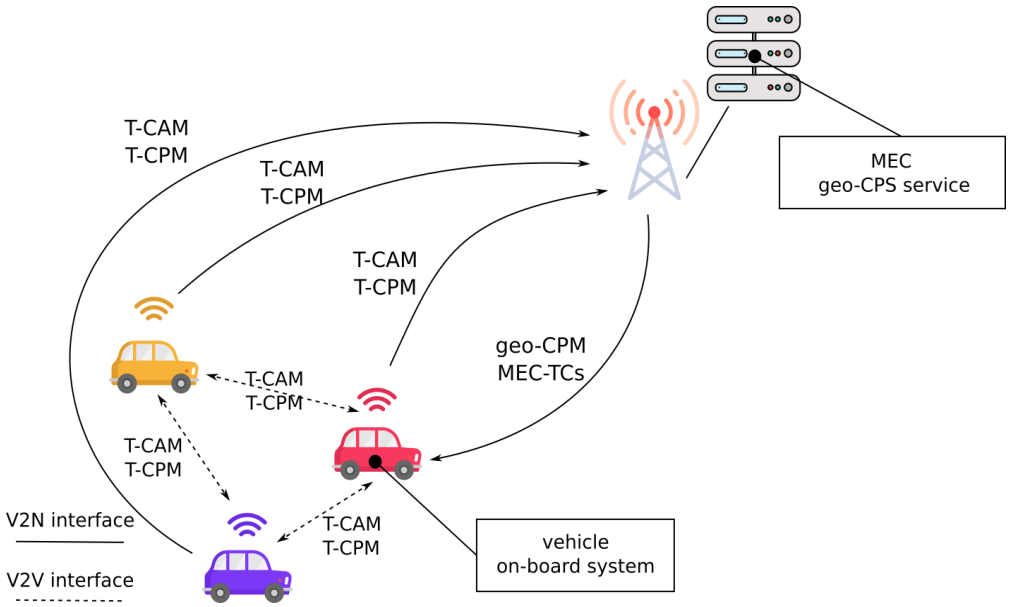

Figure 44 - Communication interfaces in the geo-Collective Perception service scenario (#2)

Figure 44 represents the radio interfaces considered in this scenario. The exchange of messages on the **V2V radio interface** consists of T-CAM and T-CPM messages (see Section 7.2.2 and Chapter 4). Recall that a T-CAM or T-CPM is a regular CAM or CPM message that may also carry V-TCs from the vehicle, and that T-CPs will be included in T-CAMs and T-CPMs in a periodic fashion. We assume that connected vehicles possess credentials, in the form of pseudonym certificates, delivered by the vehicular PKI.

The connection to the MEC is made through to the **V2N radio interface**:

- ✓ **Uplink**: V2X nodes share with the MEC the T-CAMs and T-CPMs that are also sharing in the V2V radio interface.
- ✓ **Downlink**: the MEC disseminates to the vehicles geo-CPMs, which periodically include MEC-TCs relative to the geo-CPS.

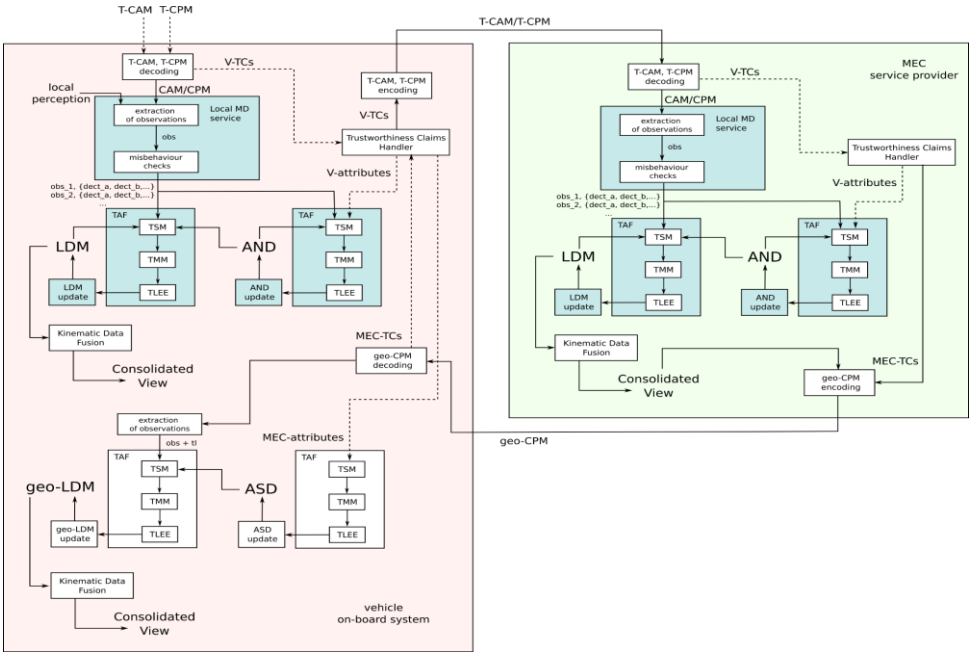

Figure 45 - Functional description of the vehicle on-board system and of the MEC system in the MEC-based Collective Perception service scenario (#2)

Figure 45 presents the functional description of the on-board system and of the system at the MEC providing the MEC-based geo-CPS. We start describing the functionality at the MEC and then we describe the functionality of the vehicle's on-board system.

MEC geo-CPS: The TAF is the core component used to deliver the MEC-based geo-CPS. As in Scenario #1 the MEC maintains the **Active V2X Node Directory (AND).** A V2X node is known to the MEC because it is a sender of T-CAMs or T-CPMs. The trust sources exploited by the TAF to maintain the trustworthiness levels in the AND are the V-TCs contained in T-CAMs and T-CPMs as well as the activation pattern of misbehaviour detectors on incoming messages.

The MEC also observes the kinematic data coming from the T-CAMs and T-CPMs uploaded from the V2X nodes. It runs a Local MD service which produces a misbehaviour detectors activation pattern for each incoming observation, which is stored in the local LDM. The trustworthiness levels of the observations are assessed by the TAF, which uses the output of the Local MD Service and the contents of the AND as trust sources, as done by the on-board system in Scenario #1.

The **geo-CPM encoding service** uses the contents of the LDM to build the consolidated view of the scene, by fusing the observations in the LDM which are attributed to the same physical object by the Object Association Algorithm. The consolidated view of the scene consists of one observation per physical object, associated with a trustworthiness level. The contents of the consolidated view are used in the Detected Objects contained in the geo-CPMs which are broadcast to the V2X nodes.

Vehicle On-board System:

On-board the vehicle maintains the LDM, where converge kinematic data contained in received T-CAMs and T-CPMs, and in the LP; and the geo-LDM, where converge kinematic data contained in received geo-CPMs. It is important to notice that the two are separate entities. The LDM and the geo-LDM are used to generate two separate consolidated views of the scene. The IMA application may choose to use one view or the other. It is important to notice that since the sets of kinematic data that contribute to form the two views partly overlap, they cannot be safely fused, and represent alternatives. Whenever it is available, the IMA will consume the consolidated view of the scene obtained from the geo-LDM.

The on-board subsystem maintaining the local LDM works in the same way as its counterpart at the MEC. This is put into evidence in Figure 45 by shading the two subsystems with the same colour. The subsystem maintaining the geo-LDM is very simple. After extraction of the observations

contained in the geo-CPM, the local TAF assesses their trustworthiness levels using as trust sources the trustworthiness levels attributed by the MEC (and contained in the geo-CPM, see Section 7.2.2 and Chapter 4) and the current trustworthiness level in the geo-CPS service, contained in the ASD. The ASD is maintained with trustworthiness levels assessed by the TAF using the MEC-TCs contained in geo-CPMs.

7.2.6 Reference Scenario User Stories

[MB.US1a] As the IMA application on the Vehicle, I want to be able to consume a consolidated view of the scene containing trustworthy data.

[MB.US1b] As the geo-CPS Service Provider at the MEC, I want to be able to build my service exploiting a consolidated view of the scene containing trustworthy data.

User story confirmation:

The considered system (vehicle or MEC Service Provider) receives V2X messages containing kinematic data. These are processed and used to populate the local LDM with observations and the relative TLs. We can distinguish the following cases:

- **At the vehicle, the incoming messages are CAM/CPM or T-CAM/T-CPM**. The kinematic data processing that allows to record observations and TLs in the local LDM is described in [MB.US2]
- **At the vehicle, the incoming messages are geo-CPMs.** The kinematic data processing that allows recording observations and TLs in the local geo-LDM is described in [MB.US3].
- **At the geo-CPS Service Provider (instantiated at the MEC), the incoming messages are T-CAM/T-CPM**. The kinematic data processing that allows recording observations and TLs in the local LDM is described in [MB.US2].

In all cases, the local LDM (or geo-LDM) records the observations along with the respective TLs. The Kinematic Data Fusion module processes the entries in the LDM, clustering the observations relative to the same physical object (decision task); and estimating the kinematic state of the physical object from the observations in each cluster (estimation task). Each estimation, or description, is associated with a consolidated TL; the collection of the descriptions and associated TLs form the consolidated view, which is consumed by the IMA application of the vehicle, and by the geo-CPM encoding module at the geo-CPS service provider at the MEC.

User Story Workflow: *(see Figure 46)*

Sequence diagram description: In this user story we have the following flows:

- The CAM/CPM message or T-CAM/T-CPM message is received and processed such that the observations and the respective TLs are recorded in the LDM (see [MB.US2] for the details of this dataflow)
- The observations and their relative TLs recorded in the LDM are accessed by the Kinematic Data Fusion module, which performs the detection and the estimation tasks.
- For each detected physical object, the Kinematic Data Fusion module builds a single description, associated with a consolidated TL, which is recorded in the Consolidated view.
- The IMA application / the geo-CPM encoding service consume the descriptions and respective TLs recorded in the Consolidated view.

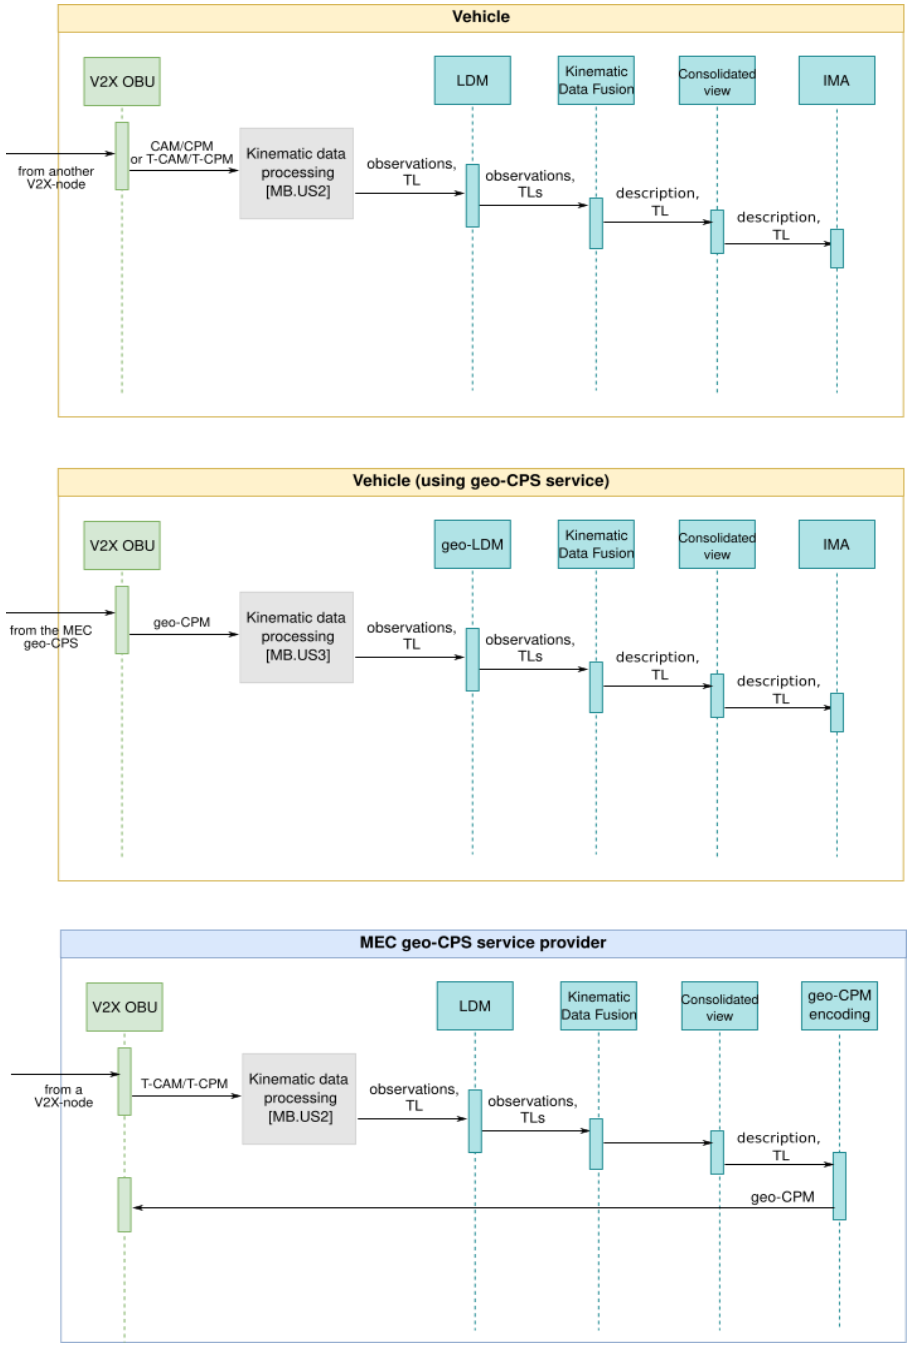

Figure 46 - Representation of the sequence diagram of user story [MB.US1]

CONNECT KPIs:

CONNECT provides the IMA use case with the availability of TLs on the incoming kinematic data, which is used by the ego to build a more accurate view of the environment. For this reason, the CONNECT impact may be measured on the Consolidated view as outputted by the Kinematic Data Fusion module. Since the IMA use case will be validated using simulation, the ground truth, defined as the accurate and truthful kinematic description of the scene, will be available. We are then interested in measuring the following KPIs, which focus on the comparison of the Consolidated view against the ground truth:

- **Object detection ratio**. The proportion of physical objects in the ground truth scene that are in the consolidated view of the scene.
- **Kinematic description accuracy**. Covariance matrix of the error vector on the kinematic object description in the consolidated view of the scene, with respect to ground truth.

For both KPIs, the performance depends on several properties of the considered driving scenario, and especially on the vehicle density, which impacts the amount of V2X data received by the ego; and the attack type and density, which impacts their quality. For this reason, several benchmark driving scenarios will be defined and used to assess the achievement of such a KPI.

However, we have to note that this evaluation property does not aim at assessing the performance of the specific Kinematic Data Fusion module, but rather in assessing the improvement granted by the availability of TLs on the observations recorded in the LDM. At the vehicle, depending on the considered Reference Scenario, a different mix of trust sources is available to the TAF assessing the TLs. In particular, the following possible configurations are of interest:

1. **(Standalone-based TAF) Vehicle LDM #1.** Vehicles broadcast regular CAM and CPM messages, hence do not broadcast V-TCs. The ego vehicle has no access to the NTS service provided by the MEC. In this configuration, the only available trust source is the misbehaviour detection activation pattern observed by the vehicle.
2. **(Federated-based TAF) Vehicle LDM #2.** Vehicles broadcast regular CAM and CPM messages, and upload V-TCs to the NTS service at the MEC. The ego receives NTMs from the MEC. In this configuration, the available trust sources at the vehicle TAF are the misbehaviour detection activation pattern and the trustworthiness levels on V2X-nodes received in NTMs. The available trust sources at the NTS Service Provider at the MEC are the harmonised attributes of the V2X-nodes and the misbehaviour detection patterns included in MRs. This configuration, where the TAF (at the NTS service running on the MEC) is in charge of evaluating TLs on the V2X-node broadcasting kinematic data, and the TAF of the ego vehicle exploits them to assess the TLs on the kinematic data itself, corresponds to the **federated TAF**.
3. **(Distributed-based) Vehicle LDM #3.** Vehicles broadcast T-CAM and T-CPM messages. In this configuration, the available trust sources are the misbehaviour detection activation pattern and the harmonised attributes of the V2X-nodes. All the evidence used in the assessment of the TL of the emitting V2X-node and in the assessment of the TL of the observation is available at the vehicle.

For each benchmark driving scenario the objective is the comparison of the Object detection ratio and of the Kinematic description accuracy in order to be able to appreciate how the increasing diversity of the available trust sources for assessing TLs impacts the ability of the system to correctly recognize the environment. When no attack is present and the kinematic data broadcast in CAMs and CPMs is **correct**, the performance is expected not to degrade with respect to the "As-Is" use case; in case of attacks triggering misbehaviour detectors, the target is the improvement of **at least 10%** in the expected average Object detection ratio with respect to the "As-Is" use case, for all the three considered configurations.

A second objective is to consider Reference Scenario #2, where the ego vehicle simultaneously maintains a Consolidated view based on the contents of the LDM (i.e., data coming from neighbouring vehicles only) and a consolidated view based on the contents of the geo-LDM (i.e., data received from the MEC service). The comparison of the Object detection ratio and Kinematic description accuracy obtained in these two cases will allow measuring, as function of parameters such as V2X-node density, the impact of the availability of the MEC service to bolster the ability of the system to correctly understand the driving environment.

Table 9 - MB.US1a and MB.US1b KPIs

KPIs	Description	Value
	Object detection ratio	It should be noted that the value is impacted by the amount of V2X data received by the ego. **Scenario #1:** In the case of attacks triggering misbehaviour detectors, the target is the improvement of at **>= 10%** in the expected average Object detection ratio with respect to the "As-Is" use case, for all the three considered configurations (i.e., Standalone, Federated and Distributed TAF). **Scenario #2:**

		The comparison of the Object detection ratio and Kinematic description accuracy obtained in the two cases (i.e., vehicle only and MEC only) will allow measuring, as function of parameters such as V2X-node density **(TRUE).**
	Kinematic description accuracy	**Scenario #1:** When no attack is present and the kinematic data broadcast in CAMs and CPMs is **correct** **Scenario #2:** The comparison of the Object detection ratio and Kinematic description accuracy obtained in the two cases (i.e., vehicle only and MEC only) will allow measuring, as function of parameters such as V2X-node density (**TRUE**).

[MB.US2a] As the Vehicle I want to be able to extract an observation contained in a CAM or CPM, attribute it a Trustworthiness Level and record it in the LDM, so that it can be used to produce the consolidated view of the scene.

[MB.US2b] As the geo-CPS Service Provider at the MEC I want to be able to extract an observation contained in a T-CAM or T-CPM, attribute it a Trustworthiness Level and record it in the LDM, so that it can be used to produce the consolidated view of the scene.

User Story Confirmation:

The system extracts kinematic data (observations) from the received CAM and CPM messages or T-CAM and T-CPM messages. The TL of each observation needs to be assessed by the TAF and to be stored in the LDM with the observation itself. This is accomplished using the output of the Local MD service, which is exploited as a trust source by the TAF. Moreover, if an entry for the emitter V2X-node is already recorded in the AND, the TAF may use its available TL. T-CAM and T-CPM messages, which are transmitted when by a predefined periodical policy fresh TCs are generated at the encoder, contain V-TCs on the emitter V2X-node. If a T-CAM or T-CPM is received, before the TL of the observation is assessed, the V-TCs are verified, and the attributes are used as trust sources by the TAF to update the TL of the emitter V2X-node in the AND. This ensures that the fresh evidence on the emitter V2X-node available in the received V-TCs is accounted for by the TAF when it assesses the TL of the observation.

User Story Workflow: (*see Figure 47*)

Sequence Diagram Description: In this user story we have the following flows:

- The CAM/CPM or T-CAM/T-CPM received by the V2X OBU is forwarded to the IAM, which checks the signature and identifies the emitter V2X-node ID (via the PKI pseudonym certificate). The V2X-node ID is forwarded to the TAF.
- If T-CAM/T-CPM is available: the T-CAM/T-CPM is decoded, and the V-TCs are forwarded to the Trustworthiness Claims Handler:
    - The TCH verifies the V-attributes and sends them to the TAF, to be used as a trust source.
    - The TAF queries the AND for the TL of the V2X-node id.
    - If present, the AND forwards the TL of the V2X-node id to the TAF.
    - The TAF uses the V-attributes and, if present, the TL of the V2X-node id to evaluate the updated TL of the V2X-node id
    - The updated value of the TL is forwarded to the AND and overwrites the preceding one.
- The CAM or CPM is forwarded to the Observation Extraction, which provides each observation to the LDM and to the Misbehaviour Checks
- The misbehaviour checks activation pattern on the observation is forwarded to the TAF, to be used as a trust source.
- The TAF queries the AND for the TL for the emitting V2X-node id.
- If present, the AND forwards the TL for the emitting V2X-node id to the TAF.

- The TAF uses the misbehaviour checks activation pattern and, if present, the TL of the V2X-node id to assess the TL of the observation and records the observation and the TL in the LDM

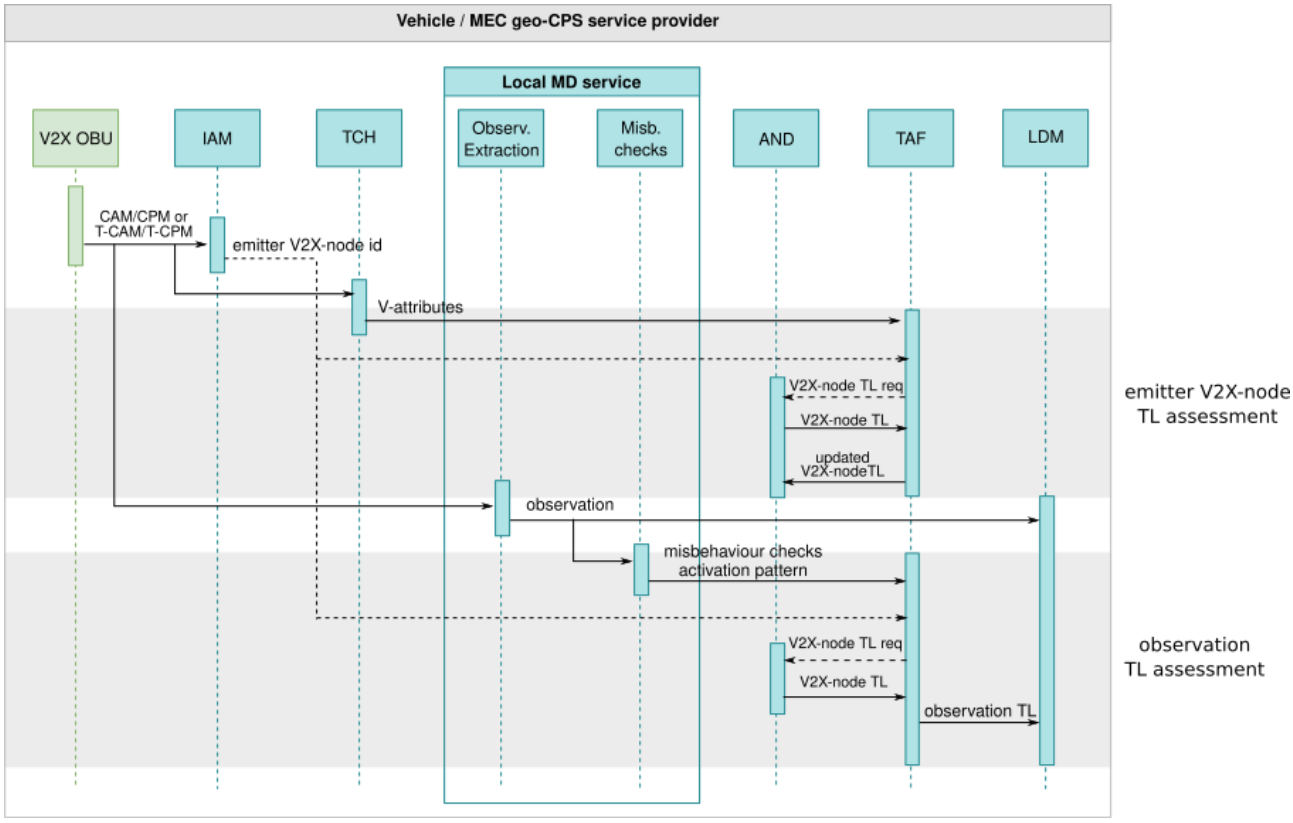

Figure 47 - Representation of the sequence diagram of user story [MB.US2].

CONNECT KPIs:

Processing complexity until LDM update: complexity of the operations between message reception and update of the observations' trustworthiness levels in the LDM. Evaluation of the complexity in the simulated setup is considered as a proxy for delay in a real deployment. This includes:

Table 10 - MB.US.2 KPIs

KPIs	Description	Value
	Processing complexity until LDM update: If V-TCs present, verification of the TCs and attributes extraction	The focus is on identifying the overhead pattern imposed based on the number of identified objects. Once this is calculated, detailed benchmarking will follow considering different vehicle neighbourhood densities
	Processing complexity until LDM update: If V-TCs present, emitter V2X-node's TL assessment by the TAF	
	Processing complexity until LDM update: Local MD on the observations. Notice that the complexity of this step may be influenced by factors such as traffic density (in denser scenarios, characterised by richer V2X information, more misbehaviour checks become available)	
	Processing complexity until LDM update: Observation's TL assessment by the TAF	

[MB.US3] As the Vehicle, I want to be able to extract an observation contained in a geo-CPM received by the MEC, attribute it a Trustworthiness Level and record it in the geo-LDM, which will be used to produce the consolidated view of the scene.

User Story Confirmation:

- The geo-CPM received by the Vehicle contains observations and their relative TLs attributed by the geo-CPS at the MEC. The geo-CPM also contains the TCs constructed by the geo-CPS Service Provider including evidence on the level of assurance provided by the MEC virtualized infrastructure where the "*geo*-enabled" service is executed. At first, the vehicle uses the TCs to update the TL in the geo-CPS service, stored in the local ASD.
- Then, to record each observation of the geo-CPM in the geo-LDM, the Vehicle needs to attribute its own observation TL. To do so, it uses the local TAF, which combines the TL of the observation contained in the received geo-CPM with the TL of the geo-CPS service, which it queries from the ASD.

User Story Workflow: *(see Figure 48)*

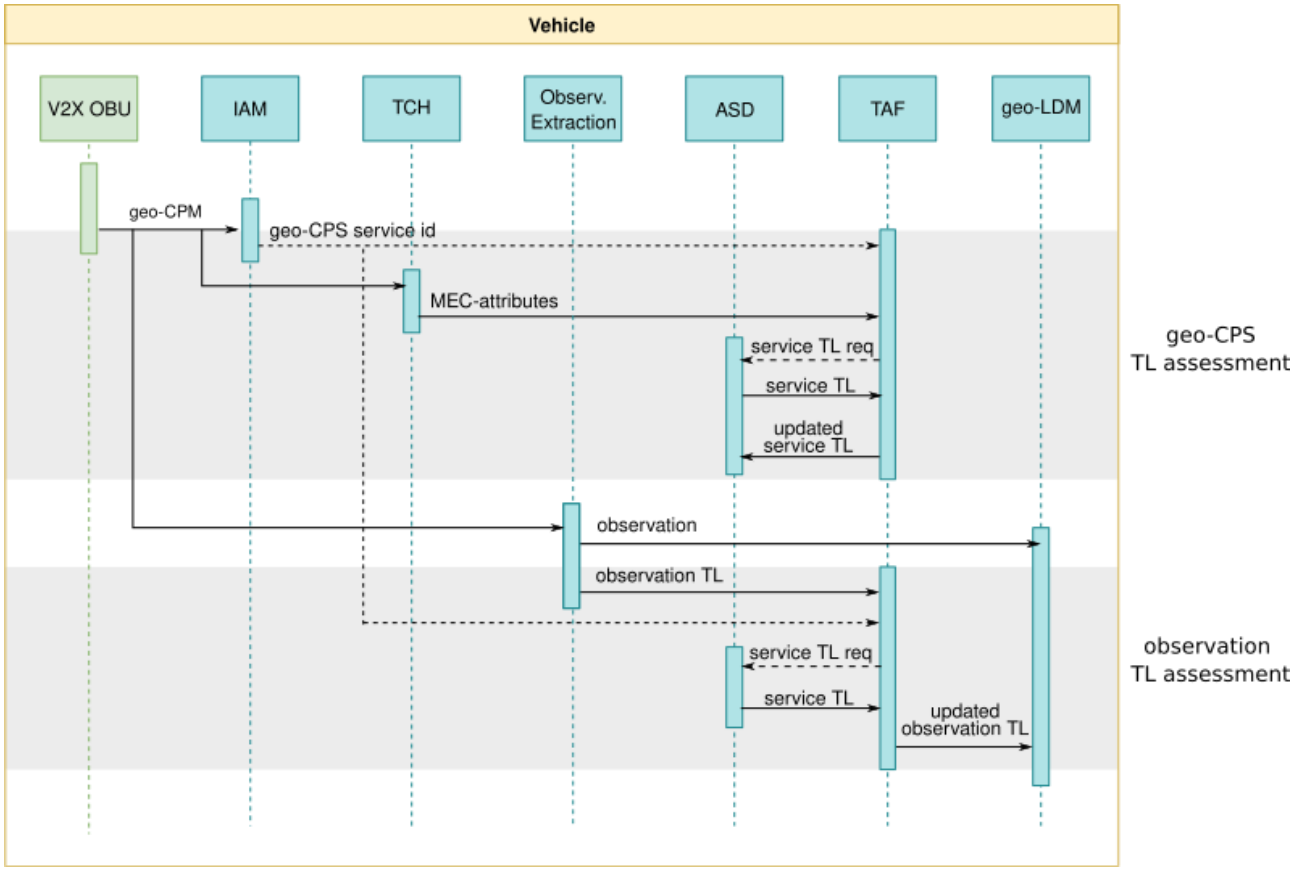

Figure 48 - Representation of the sequence diagram of user story [MB.US3]

Sequence Diagram Description: In this user story we have the following flows:

- The geo-CPM received by the V2X OBU is forwarded to the AIM, which checks the signature and identifies the geo-CPS service id. The geo-CPS service id is m\de available to the TAF.
- The geo-CPS is forwarded to the TCH, which verifies the MEC-attributes and sends them to the TAF, to be used as a trust source.
- The TAF queries the ASD for the TL of the geo-CPS service.
- If an entry for the geo-CPS service is present in the ASD, the ASD forwards the relative TL to the TAF
- The TAF uses the MEC-attributes and, if present, the TL of the geo-CPS service to assess the updated TL of the geo-CPS service
- The ASD updates the TL of the geo-CPS service or creates a new entry for it if it was not present.

- The Observation Extraction extracts the observation from the geo-CPM and forwards it to the geo-LDM.
- The TAF queries the ASD for the TL of the geo-CPS service.
- The ASD forwards the TL of the geo-CPS service to the TAF.
- The TAF uses the observation TL contained in the geo-CPM and the TL in the geo-CPS service to assess the updated TL of the observation and forwards it to the geo-LDM.

CONNECT KPIs:

- Processing complexity until geo-LDM update: time between message reception and update of the observation trustworthiness level in the geo-LDM. This includes:

Table 11 - MB.US3 KPIs

KPIs	Description	Value
	Processing complexity until geo-LDM update: Verification of the TCs and attributes extraction	The focus is on identifying the overhead pattern imposed based on the number of identified objects. Once this is calculated, detailed benchmarking will follow considering different vehicle neighbourhood densities
	Processing complexity until geo-LDM update: geo-CPS's TL assessment by the TAF	
	Processing complexity until geo-LDM update: Observation's TL assessment by the TAF	

[MB.US4a] As the Vehicle I want to update the Trustworthiness Level of the emitter V2X-node upon the reception of a CAM/CPM or T-CAM/T-CPM message.

[MB.US4b] As the geo-CPS Service Provider at the MEC I want to update the trustworthiness level of the emitter V2X-node upon the reception of a T-CAM or T-CPM message.

User Story Confirmation:

- Whenever a CAM/CPM or T-CAM/T-CPM is received, the emitter V2X-node is added to the AND or its TL is updated. In case of T-CAM or T-CPM, the contained V-TCs are used as a trust source by the TAF to assess or update the TL of the V2X-node. This, as highlighted in [MB.US2], takes place before the trust level assessment of the observations contained in the T-CAM/T-CPM. The second update of the TL of the emitter V2X-node in the AND takes place after the trustworthiness level on the kinematic data has been assessed, using the misbehaviour checks activation pattern as a trust source by the TAF.
- The misbehaviour checks activation pattern is hence twice used as a trust source: first, to assess the TL on the current observation; secondly, to update the TL on the emitter V2X-node as a result of the behaviour observed in the latest transmission.
- In this user story the TL of the known V2X-nodes are evaluated using trust sources which have been directly observed by the ego vehicle.

User Story Workflow: *(see Figure 49)*

Sequence Diagram Description: The flows belonging to this user story are represented in the shaded blue boxes in Figure 49. If V-TCs are present in the T-CAM or T-CPM, the update of the AND happens in two phases, before and after the assessment of the TL of the observations contained in the message (the complete flow is illustrated in Figure 49 for clarity). The first phase has already been described in [MB.US2]. The flow in the second phase is as follows:

- The Local MD service forwards the misbehaviour checks activation pattern to the TAF.
- The IAM module forwards the emitter V2X-node id to the TAF.
- The TAF queries the AND for the TL of the V2X-node id.
- If an entry for the V2X-node id is present in the AND, the AND forwards the relative TL to the TAF
- The TAF uses the misbehaviour detection activation pattern and, if present, the TL of V2X-node id to assess the updated TL of the V2X-node id, and forwards it to the AND

CONNECT KPIs:

Table 12 - MB.US4a and MB.US4b KPIs

KPIs	Description	Value
	The TLs of the active V2X-node evolves correctly: it increases as more evidence of correct behaviour is gathered; it degrades as malicious behaviour is injected in the scenario.	**TRUE**

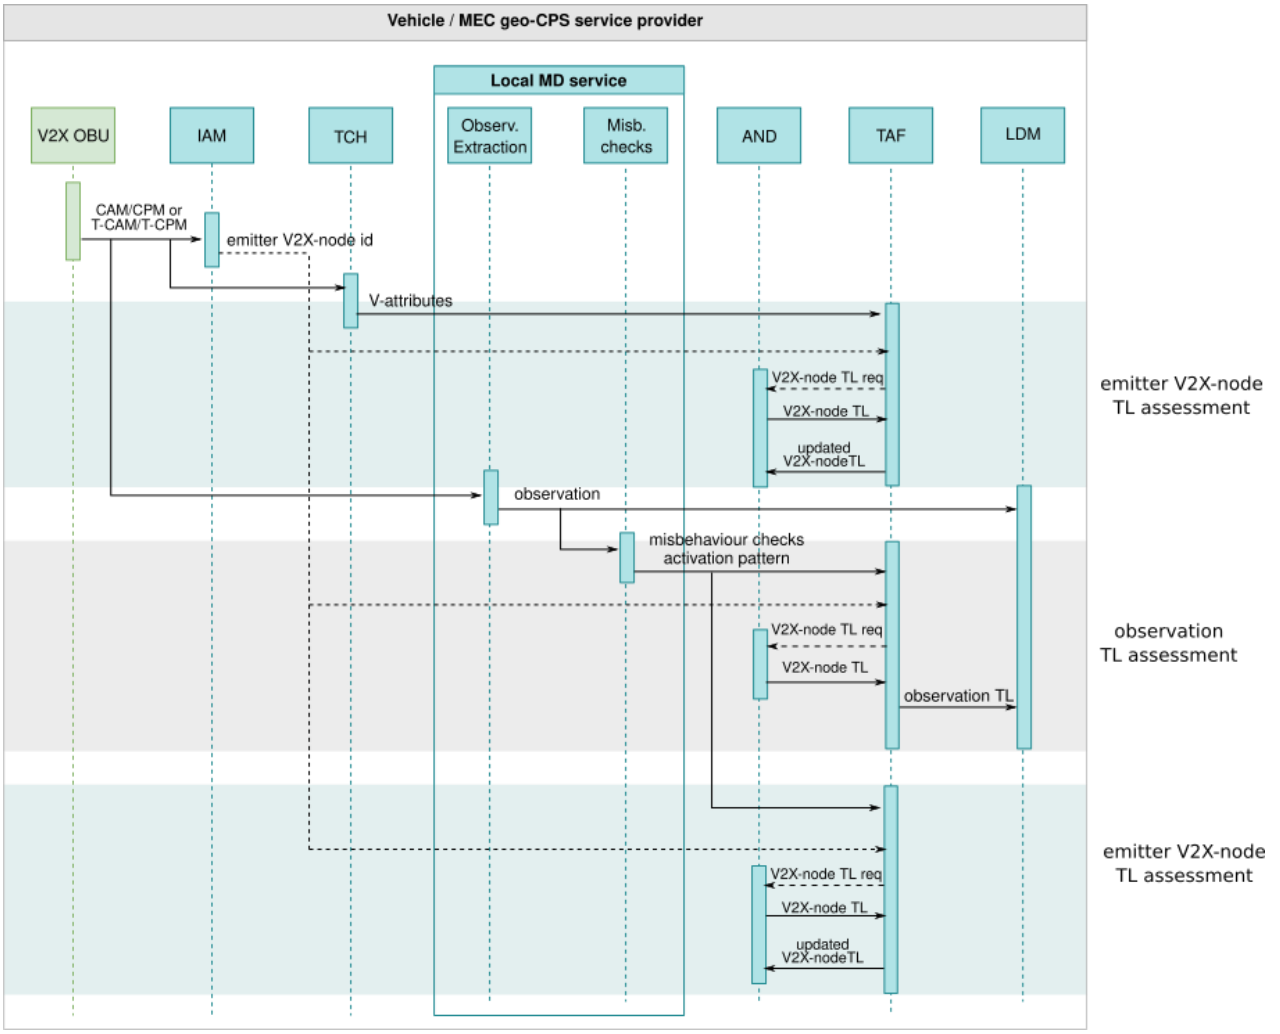

Figure 49 - Representation of the sequence diagram of user story [MB.US4]

[MB.US5] As the vehicle, I want to update the trustworthiness level of the emitter V2X-node in the AND whenever I receive a NTM message from the MEC.

User Story Confirmation:

- In S1, the vehicle receives NTMs from the MEC, containing lists of V2X-nodes and their respective TLs, as attributed by the NTS service. The NTM also contains the MEC-TCs made by the NTS Service Provider at the MEC. At first, the vehicle uses the MEC-TCs to update the TL in the NTS service, stored in the local ASD.
- Then, the Vehicle needs to attribute its own TL in the V2X-nodes. To do so, it uses the TAF, which combines the TL of the V2X-node contained in the received NTM with the TL of the NTS service, which it queries from the ASD. Finally, the AND is updated.
- The TLs in the emitter V2X-nodes in the AND of the vehicle have hence been assessed combining trust evidence observed by the NTS service at the MEC and trust evidence observed by the vehicle itself. These TLs in turn will be used by the vehicle TAF, along with other evidence observed by the vehicle, to assess TLs on the kinematic data. This constitutes a **federated TAF**. The federated TAF hence enables to combine heterogeneous evidence, collected in different places, to the purpose of making a more informed trust assessment.

User Story Workflow: *(see Figure 50)*

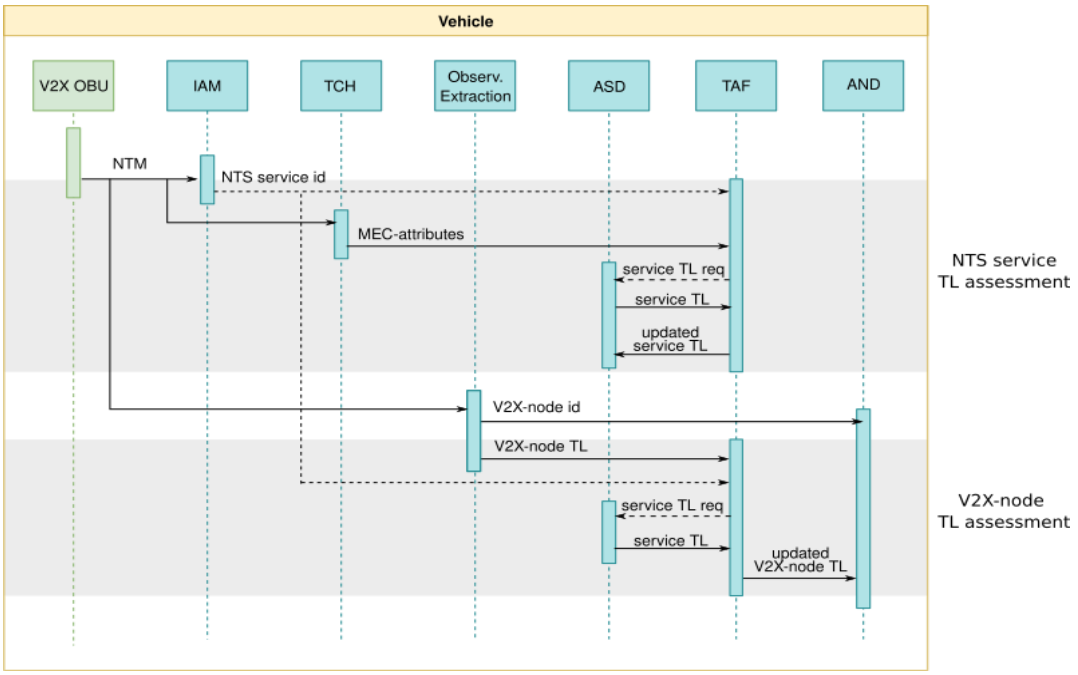

Figure 50 - Representation of the sequence diagram of user story [MB.US5]

Sequence Diagram Description: In this user story we have the following flows:

- The NTM is received by the V2X OBU and it is forwarded to the IAM, which verifies the signature and extracts the NTS service id.
- The TCH module extracts and verifies the attributes from the MEC-TCs and forwards them to the TAF.
- The TAF queries the ASD for the TL of the NTS service.
- If an entry for the NTS service is present in the ASD, the ASD forwards the relative TL to the TAF
- The TAF uses the attributes and, if present, the TL of the NTS service to assess the updated TL of the NTS service and forwards it to the ASD.
- The ASD updates the TL of the NTS service or creates a new entry for it if it was not present.
- The Observation Extraction extracts the contents from the NTM: the V2X-node id and the V2X-node TL.
- The TAF queries the ASD for the TL of the NTS service.
- The ASD forwards the TL of the NTS service to the TAF.
- The TAF uses the V2X-node id TL contained in the NTM and the TL in the NTS service to assess the updated TL of the V2X-node id, and forwards it to the AND

CONNECT KPIs:

Table 13 - MB.US5 KPIs

KPIs	Description	Value
	The TLs of the active V2X-node evolves correctly: it increases as more evidence of correct behaviour is gathered; it degrades as malicious behaviour is injected in the scenario.	**TRUE**

[MB.US6] As the NTS Service Provider at the MEC I want to update the Trustworthiness Level of the V2X-node in the AND whenever I receive a MR or a TCs. Keeping the AND updated with fresh information allows it to deliver a beneficial NTS service.

User Story Confirmation:

- The NTS Service Provider at the MEC may receive data referring to a specific V2X-node in various forms. It may receive a MR transmitted by the V2X-node, which contains TCs from the V2X-node itself; or it may receive a MR reporting a message sent by the V2X-node. The TAF uses the TCs, and the misbehaviour checks activation pattern contained in the MR as evidence allowing to update the TL of the V2X-node in the AND.

User Story Workflow: *(see Figure 51)*

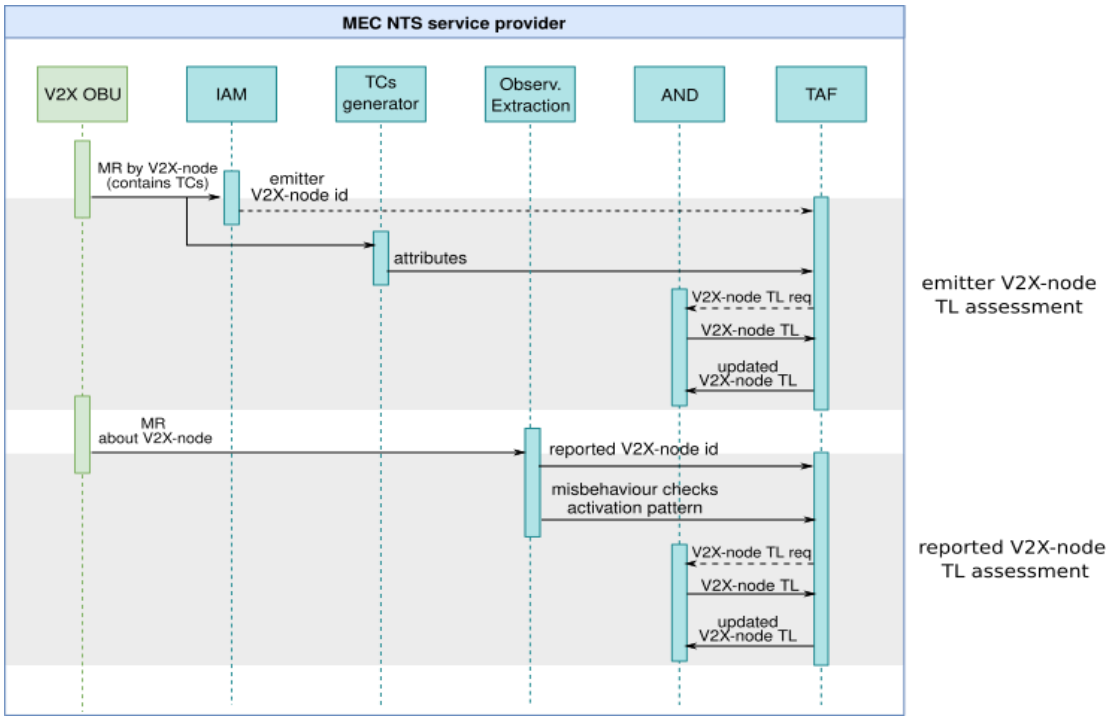

Figure 51 - Representation of the sequence diagram of user story [MB.US6]

Sequence diagram description: In this user story we have the following flows:

- The MR emitted by the V2X-node and containing the TCs is received by the V2X OBU and it is forwarded to the Observation Extraction module.
- The Observation Extraction extracts the attributes and the emitter V2X-node id.
- The Observation Extraction module forwards the attributes to the TAF.
- The Observation Extraction module forwards the V2X-node id to the TAF.
- The TAF queries the AND for the TL of the V2X-node id.
- If an entry for the V2X-node is present in the AND, the AND forwards the relative TL to the TAF
- The TAF uses the attributes and, if present, the TL of the V2X-node id to assess the updated TL of theV2X-node id, and forwards it to the AND
- The MR reporting a message emitted by the V2X-node is received by the V2X OBU and it is forwarded to the Observation Extraction module.
- The Observation Extraction extracts the misbehaviour checks activation pattern of the reported message and the reported V2X-node id.
- The Observation Extraction module forwards the misbehaviour checks activation pattern to the TAF.
- The Observation Extraction module forwards the V2X-node id to the TAF.
- The TAF queries the AND for the TL of the V2X-node id.

CONNECT KPIs:

Table 14 - MB.US6 KPIs

KPIs	Description	Value

	The TLs of the active V2X-node evolves correctly: it increases as more evidence of correct behaviour is gathered; it degrades as malicious behaviour is injected in the scenario.	**TRUE**

7.3 Cooperative Adaptive Cruise Control

Cooperative Adaptive Cruise Control (C-ACC) represents a cutting-edge application in the realm of intelligent transportation systems, enhancing traditional Adaptive Cruise Control through vehicle-to-vehicle communication. In this innovative use case, vehicles equipped with C-ACC exchange real-time messages, sharing crucial information such as current speed, acceleration, and position. This continuous communication allows the vehicles to operate in a coordinated manner, optimizing traffic flow and safety.

We evaluate our use case within the framework of a **service-oriented zonal architecture**. This architectural design features a zonal controller strategically positioned between the Electronic Control Unit (ECU) and the sensors and actuators, thereby facilitating the possibility of also enabling the delivery of specialized software updates. Our in-vehicle architecture consists of a smart antenna, enabling the communication with other vehicles, the Vehicle Computer, the zonal controllers and the sensors or actuators (i.e., lidar, etc.). For keeping a safe distance to the vehicle in front, the C-ACC bases its decisions on various data items received from in-vehicle sensors, as well as from other vehicles. Upon receiving the data, the C-ACC proceeds with making crucial driving decisions and generates two essential types of messages to ensure safe and coordinated driving: Acceleration command to the Acceleration ECU for regulating the vehicle's speed (i.e., whether it should slow down or speed up) and CAM message to inform other vehicles about driving parameters, including speed, heading, etc.

The trustworthiness of the involved technical components and data items is of paramount importance to ensure the safe and effective operation of C-ACC. **This use case focuses on how to assess the trustworthiness at runtime and respond to changes, taking appropriate action**. The C-ACC main component executed on the vehicle's computer improves the safety and reliability of its functionality through dynamic trustworthiness assessment of 1) input data (whether from local sensors or other vehicles) and how they are transferred to it, as well as 2) the execution platform of the C-ACC itself, i.e., a Vehicle Computer ECU). The C-ACC Main Component is able to respond to changes in the trustworthiness of data items it receives from in-vehicle and vehicle-external sensors, by enacting pre-defined safe responses, such as adapting function-internal calculations or increasing the distance to other vehicles. Also, in case the current Vehicle Computer is not considered trustworthy enough based on the evidence and trust assessment provided by CONNECT's TAF, it performs a migration process to move its C-ACC Main Component from one Vehicle Computer ECU to another, to ensure operational continuity.

Based on the analysis presented in the European Union's transport White Paper [140], the increasing road transport activity in the European Union is the primary cause of growing congestion and rising energy consumption, as well as a source of environmental and societal problems. In August 2019, a report from the Texas A&M Transportation Institute estimated the additional time spent by the average American commuter driving through heavy traffic [141]. The report found that American commuters spent an extra 54 hours per year in traffic delays. An "extra hour" refers to the additional time spent travelling at congested speeds rather than free-flow speeds. It also mentions that while some U.S. cities may have more congested traffic conditions compared to European cities, the average American experience is similar to that of many big cities worldwide.

The integration of dependable communication systems in intelligent vehicle cooperation not only aids in the mitigation of traffic accidents but also enhances traffic efficiency [142]. Cooperative adaptive cruise control (CACC) systems in specific, have the capacity to enhance traffic flow by permitting shorter distances between cars and enabling them to travel in a coordinated manner at a synchronised pace (i.e., platoon). CACC leverages V2V communications to enable Connected and Automated Vehicles (CAVs) to create platoons and travel at synchronised speeds with reduced time

intervals between them. In order to enable platoons, the vehicles disseminate data such as acceleration, speed, and location within a specific communication range in order to collaborate with one another. The results of such cooperation can be summarised in the following points: 1) the time it takes for a vehicle to react is reduced while the driving safety is increased, compared to manual driving, 2) the amount of time and distance between vehicles is reduced, leading to an increase in the capacity of the roadway and 3) unnecessary changes in velocity and aerodynamic drag on following vehicles are reduced, resulting in a decrease in energy consumption and pollutant emissions [143].

Given the nature of the communication between vehicles, intentionally or unintentionally malicious or falsified information may arrive. Such information though could have serious implications on the passenger safety since it could result in erroneous speed (i.e., acceleration or braking) or position information, which can disrupt the coordination of the vehicle. With the continuous advancement of C-ACC technology, it is imperative to enhance its security measures in order to guarantee the ongoing safety and dependability of this revolutionary driving system.

7.3.1 "As-Is" Scenario

Modern vehicles currently on the road can run a driving assistance system called Adaptive Cruise Control (ACC). This system empowers vehicles to autonomously maintain a safe following distance from other vehicles ahead, significantly enhancing road safety. ACC relies exclusively on sensor data captured by the vehicle itself. The ACC functionality is executed on a central high-performance OBU often referred to as a Vehicle Computer, which seamlessly integrates inputs from a suite of sensors, including cameras, radar, lidar, and GNSS (global navigation satellite service). The security and safety controls specific to each sensor are meticulously defined and certified during the system's design phase. In the event of a partial system malfunction, triggered by factors such as system errors or detected security breaches, the ACC function promptly responds by notifying the driver and suspending ACC operations.

To further elevate the reliability and effectiveness of the ACC function, the concept Cooperative Adaptive Cruise Control (C-ACC) has been created as described in the C2C-CC's Guideline for Day 2 and beyond [1] and ETSI TR 103 299 [144]. With C-ACC, the established ACC function is extended by sharing vehicle steering data with other vehicles and roadside infrastructure (data sent by pedestrians or other road users is not integrated). This may increase the vehicle's field of perception over its environment, but also introduces risks as vehicles' attack surface is also extended and **safety-critical driving decisions are influenced by external, potentially non-trustworthy sources**. Figure 52 shows two vehicles exchanging information to facilitate such a C-ACC scenario.

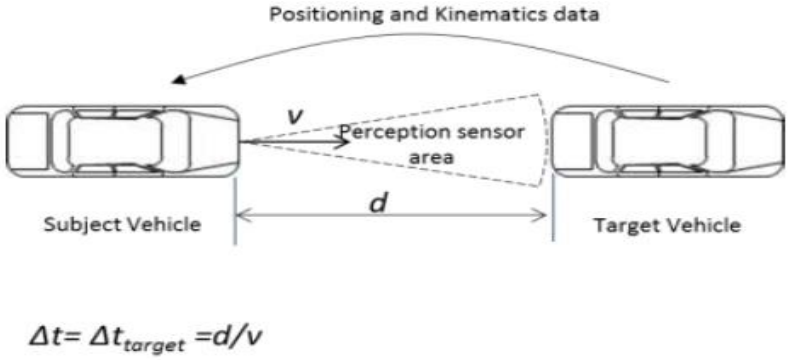

Figure 52 - Example use case scenario for C-ACC based on ETSI TR 103 299 [144]

In the context of CONNECT, we focus on the **in-vehicle point of view of the overall C-ACC functionality**. This means, we investigate the data flows and computation steps within a vehicle to produce the necessary C-ACC related information, to be exchanged, but also how to securely convert this data into in-vehicle steering commands.

It should be underlined that in this specific scenario, a service-oriented zonal architecture is considered. In the traditional architectures, the ECU connects with the rest of the in-vehicle

components through the CAN bus. However, in recent years, the need for seamlessly delivering software updates to in-vehicle components is gradually becoming more critical. Real-time software updates on vehicles provide manufacturers with cost savings, enabling the prompt correction of critical software defects, and facilitating the incorporation of new features.

In response to this challenge, zonal controllers have surfaced as a prospective solution aimed at enabling these updates, all the while preserving a level of control and maintaining a well-defined structure regarding the vehicle's internal topology (i.e., by supporting different zones). CONNECT considers both the current in-vehicle design and the expected service-oriented design, in its use cases, guaranteeing flexibility to fulfil both sets of needs.

7.3.2 In-Vehicle Components, Communication Interfaces and Messages in the context of Cooperative Adaptive Cruise Control

We evaluate our use case within the framework of a service-oriented zonal architecture, as elaborated in Maul et al's work [145]. This (forward-looking) architectural design features a zonal controller strategically positioned between the Electronic Control Unit (ECU) and the sensors and actuators, thereby facilitating the possibility of also enabling the delivery of specialised software updates. In the automotive industry, the ability to perform in-the-field software updates to vehicles is of paramount importance, as it enables the prompt resolution of critical issues and the introduction of new functionalities. Currently, only a limited number of vehicles benefit from such updates, and those that do primarily target infotainment or telematics systems. In this context, we delve into the requirements of systems that aspire to provide runtime updates to their components, enabling the vision of AVs.

Our in-vehicle architecture, which is comprised of a smart antenna, enabling the communication with other vehicles, the Vehicle Computer, the zonal controllers and the sensors or actuators (i.e., lidar, etc.), is depicted in Figure 53:

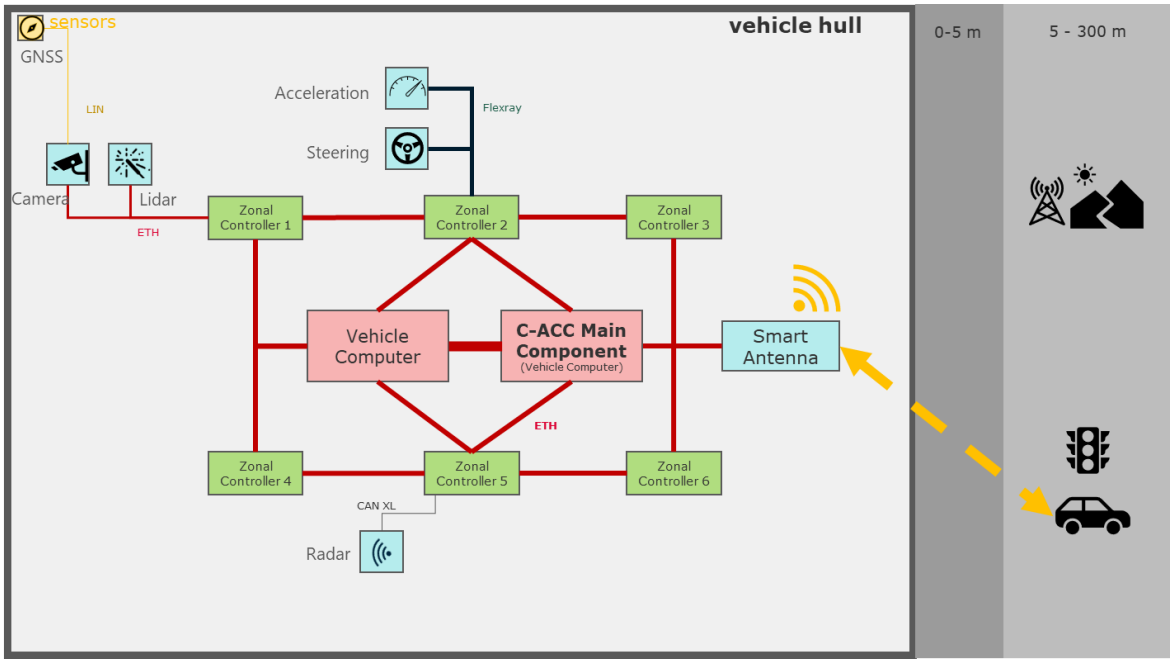

Figure 53 - CONNECT in-vehicle architecture for C-ACC

7.3.1.1 In-Vehicle Electronic Control Units (ECUs)

Building upon the main stakeholders of the CCAM landscape, as identified in chapter 4, in what follows we mention specific components and roles that we consider in the specific UC.

This architecture comprises the following stakeholders:

- **C-ACC Main Component**: This is a software component which processes all C-ACC-relevant sensor data (including data from the in-vehicle sensors, as well as data from other vehicles), produces commands to in-vehicle actors, and shares information with other vehicles. The C-ACC Main Component is executed on a Vehicle Computer type ECU.
- **CAM/CPM Encoder/Decoder Component:** Like the C-ACC Main Component, this is a software component which is executed on a Vehicle Computer type ECU. It receives incoming messages from the V2X communication network, verifies their security and decodes them. The message data is disseminated in the vehicles to all interested components. The same principle is applied for sending messages: ECUs send data to the CAM/CPM Encoder to create the corresponding V2X message and send it to the V2X communication network.
- **In-vehicle sensors**: These sensors produce data about the vehicle's environment to support the C-ACC Main Component in its decision making. These sensors include:
    - **GNSS** (Global Navigation Satellite System): produces data on the vehicle's position. This sensor is implemented in pure hardware as an ASIC (Application-Specific Integrated Circuit) system. It is connected to the Camera ECU via a LIN [146] bus.
    - **Camera**: This sensor captures the visible vehicle surroundings in a video feed sent to the C-ACC Main Component for further analysis. In addition, it also forwards the vehicle position data from the GNSS sensor to the C-ACC Main Component
    - **Lidar**: a laser-based imaging system to detect objects on the road and sends an object map to the C-ACC Main Component
    - **Radar**: a radio-based imaging system to detect objects on the road and sends an object map to the C-ACC Main Component
- **In-vehicle actors**: These actors receive commands (also) from the C-ACC component to navigate the vehicle. These actors include.
    - **Acceleration**: controls the vehicle's speed. This actor can speed up the vehicle or slow it down.
- **Supplemental ECUs**: These ECUs support the in-vehicle communication architecture.
    - **Zonal Controller:** It provides a communication interface between ECUs (in a sub-network) and the central Vehicle Computers. As will be seen later on there can be different type of such interfaces on-boarded depending on the requirements of the specific application, e.g., CAN [147], Ethernet Switch, etc.
    - **Smart Antenna**: provide an interface between the in-vehicle network and the outside world including the MEC, cloud-based services and other vehicles. When communicating with the MEC and other vehicles, the Smart Antenna acts as mediator for forwarding the V2X messages that have been constructed by the CAM/CPM Encoder/Decoder running on a Vehicle Computer ECU.

7.3.1.2 C-ACC Messages

The communication between these in-vehicle stakeholders is shown in the data flows in Figure 54 and Figure 55.

For keeping a safe distance to the vehicle in front, the C-ACC bases its decisions on various data items received from in-vehicle sensors, as well as from other vehicles. These types of data items are explained in the following paragraphs, based also on the illustration, as provided in Figure 54.

- The **Lidar** sensor offers data regarding the "**objects on the road**". These data items are sent from the Lidar ECU to Zonal Controller 1, in the given topology (see Figure 54). However, for security and network isolation reasons, Zonal Controller 1 not only forwards the data but also inspects it for potential security threats. Following this inspection, a new network connection is established to transmit this data to the Vehicle Computer.
- The **GNSS** unit locates the **vehicle's position.** This data item is sent from the GNSS ECU via the LIN (Local Interconnect Network) to the **Camera ECU**. Since the LIN bus operates as a polling system, the Camera ECU is continually querying the GNSS for updates regarding the data item. The Camera ECU translates the data item from the LIN bus to Ethernet format

and forwards it to Zonal Controller 1. The latter, inspects the data before forwarding it to the Vehicle Computer (see Figure 54).

- The **Camera ECU**, in addition to the LIN polling for GNSS data, produces a **video feed.** This video feed is transmitted directly to the Vehicle Computer without any modifications as it passes through Zonal Controller 1.
- In the meantime, the **Radar ECU** is also collecting data regarding the **objects on the road,** similarly to the Lidar. These data are transmitted initially to Zonal Controller 5. Before sending the data to the
- Vehicle Computer, Zonal Controller 5 converts the data format from CAN XL[17] (Controller Area Network XL) to Ethernet.
- Additionally, the **Smart Antenna** unit plays a role in C-ACC by receiving CAM from other vehicles. These CAM messages, as specified in ETSI EN 302 637-2 V1.3.1 [101], are transmitted by the Smart Antenna to the Vehicle Computer.

Upon receiving the data, the C-ACC proceeds with making crucial driving decisions and generates two essential types of messages to ensure safe and coordinated driving, as illustrated on Figure 55:

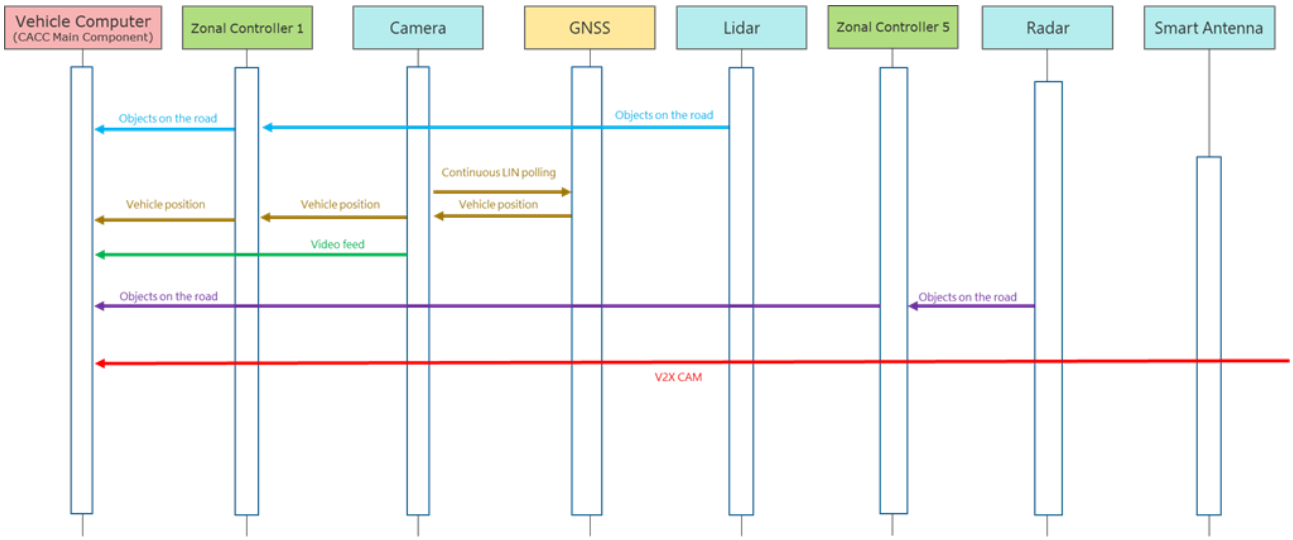

Figure 54 - Incoming data flows to the C-ACC Main Component

- **Acceleration command** to the Acceleration ECU: This data item plays a pivotal role in regulating the vehicle's speed (i.e., whether it should slow down or speed up). The Vehicle Computer is responsible for generating this command, which is subsequently sent to the Zonal Controller 2. The latter translates the command to Flexray [148] bus format, guaranteeing compliance with the vehicle's communication system, and transmits it to the Acceleration ECU. This acceleration instruction efficiently regulates the speed of the vehicle in order to sustain a secure distance from the preceding vehicle.
- **CAM message** to other vehicles: This data item informs other vehicles in the vicinity about the C-ACC-equipped vehicle's driving parameters, including its speed, heading, and other relevant information. It is important to mention that this data item is not directly produced by the C-ACC module itself, but rather by another function that operates on the same Vehicle Computer. After its creation, this message is sent to the Smart Antenna, which broadcasts it to the surrounding vehicles and infrastructure.

7.3.1.3 C-ACC Security Features

[17] A new version of the CAN bus which supports significantly larger packet sizes (> 1KB). Currently in standardization for the next edition of ISO 11898-1.

The C-ACC use case utilises the cryptographic capabilities and key management concepts (that will be enhanced by the integration of the CONNECT TEE Guard), described in Chapter 6, establishing a strong basis for incorporating other security features, especially as it pertains to the communication integrity and confidentiality (when needed by the application) of the exchanged messages, kinematic data and other in-vehicle system information. Towards this direction there is the already defined security communication architecture, breaking down the different protocols as illustrated in Figure 56.

By sequentially examining Figure 56 in a top-down and left-to-right manner, it should be noted that the LIN bus between the GNSS and Camera is **not secured.** This is due to the fact that this information should be sent in a fast and efficient manner.

On the other hand, all Ethernet sub-networks are **integrity-protected using MACsec** [149] employing the cipher suite PSK_WITH_NULL_SHA256. **Confidentiality** security measures are **purposefully excluded** in most scenarios since their implementation might unnecessarily complicate the debugging operations. For a more comprehensive assessment of the enhanced cryptographic primitives incorporated within CONNECT, a portion of the C-ACC application's communication will be encrypted. This encryption serves a critical role in ensuring the confidentiality of sensitive data, such as video feeds, as will be elaborated upon later on, in this section.

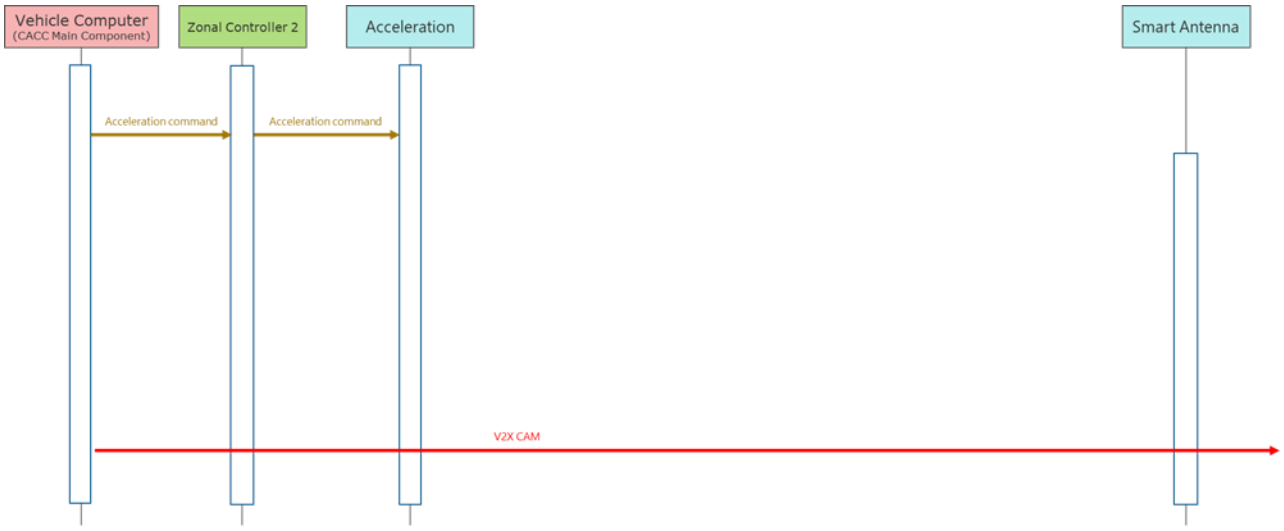

Figure 55 - Outgoing data flows from the C-ACC Main Component

Following the breakdown of the ECU capabilities, that comprise the in-vehicle topology as described in chapter 6, it becomes clear that ECUs are divided into symmetric and asymmetric capable. The asymmetric capable ECUs **leverage their own keypairs** (i.e., public, private) to establish authentic relationships. In the context of symmetric ECUs, to establish the **authenticity** of communication relationships, **pre-shared keys (PSK)** are utilised. This approach supports ECUs that may not meet the performance criteria for asymmetric cryptography methods like ECC or RSA. Given the relatively static nature of communication relationships within the vehicle, the deployment of PSK-based authentication remains suitable. As both of these set of keys (both for asymmetric and symmetric capable ECUS) are primarily intended for authenticating communication between the in-vehicle ECUs, they are not included in the migration scenario. In the event of migrating from one ECU to another, the receiving ECU possesses its own set of keys, distinct from those used by the migrating ECU. This separation ensures that the integrity and security of the communication between the ECUs remain intact during migration.

It is crucial to recognize that in our specific scenario, MACsec primarily provides security for the sub-networks. This means that **end-to-end (E2E) security is not established when data crosses the boundaries of these sub-networks**. For instance, the

The communication between the Camera and the C-ACC Main Component benefits from an additional layer of protection, offered by **TLS** with the cipher suite

TLS_PSK_WITH_AES_128_CBC_SHA256. TLS is used in addition to the underlying MACsec, to introduce **E2E security**; thus, data confidentiality. Confidentiality is particularly needed here, due to privacy reasons since the video feed can include personally identifiable information (e.g., pedestrians' faces). Notably, the integrity of the TLS protection remains intact as data traverses Zonal Controller 1, as the Zonal Controller 1 acts as a mere pass-through for TLS-secured communications.

Security measures are further implemented in the **Flexray** sub-network that includes the Acceleration, Steering, and Zonal Controller 2 components. These measures involve the use of **SecOC**[18] **with PSK authentication**. Furthermore, the CAN XL sub-network, consisting of the Radar and Zonal Controller 5, is protected by SecOC with PSK authentication. Furthermore, the **CAN XL sub-network**, consisting of the Radar and Zonal Controller 5, is protected by **SecOC with PSK** authentication.

For external vehicle communication between the Smart Antenna and other vehicles or the MEC, the **V2X security protocols** are pivotal. These protocols not only assure the security of the data, but also include procedures to preserve privacy. These mechanisms are designed to safeguard sensitive information that is transferred between the vehicle and the external environment. This all-encompassing security policy ensures the protection of all forms of communication within the C-ACC system.

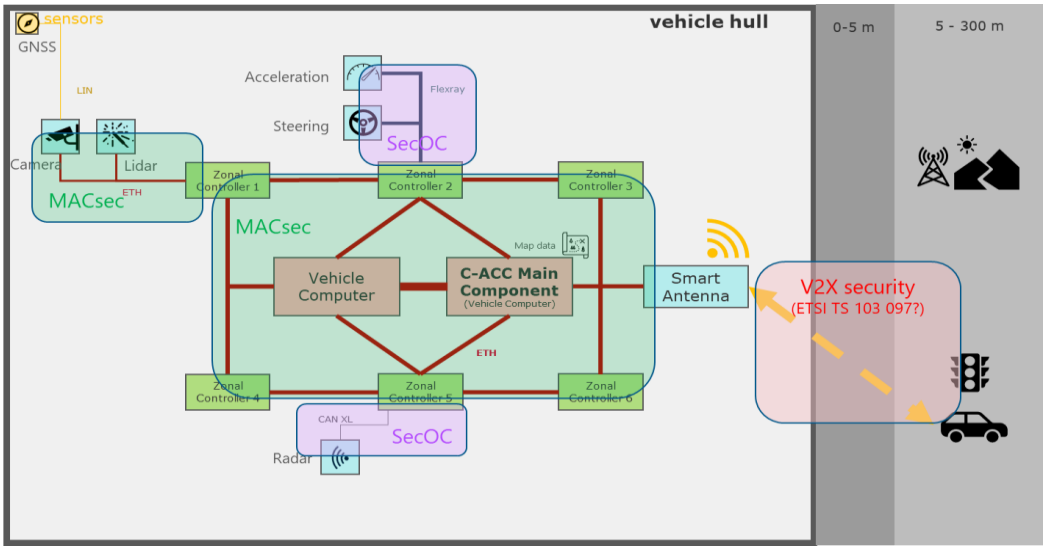

Figure 56 - Secure communication protocols in the C-ACC use case

In addition to the integrity assurances provided by MACsec for communication, the integrity of the involved ECUs is further protected by cybersecurity controls specific to each ECU class. All ECUs in this use case are categorised according to their computational and cybersecurity capabilities. This resulted in the following ECU classes and profiles of their integrity protection features as shown in Figure 56. More specifically, the white boxes include mandatory cybersecurity controls, while the grey boxes are referring to the cybersecurity controls that are only included, if needed, based on a security risk assessment (TARA).

The categorization of ECUs' cryptographic capabilities, as introduced in Section 6.5.1.3, can be succinctly summarised into three main classes:

1. **A-ECU (Advanced ECU)**: the ECU is powerful enough to execute all kinds of cryptographic algorithms, including asymmetric (e.g., ECC, RSA) and symmetric (e.g., AES), and hashing (e.g., SHA-2/3) algorithms.
2. **S-ECU (Symmetric ECU)**: the ECU's capabilities are limited to only execute symmetric and hashing algorithms.
3. **N-ECU (No-Crypto ECU)**: the ECU has no cryptographic capabilities.

[18] AUTOSAR: Specification of Secure Onboard Communication Protocol (https://www.autosar.org/fileadmin/standards/R22-11/FO/AUTOSAR_PRS_SecOcProtocol.pdf)

In Figure 57, the distribution of these ECU categories can be observed, within the context of the C-ACC use case. The colour-coding differentiates these categories, with red representing A-ECUs, and green and blue signifying S-ECUs and N-ECUs, respectively.

The above introduced in-vehicle secure communication protocols MACsec, TLS, and SecOC use Pre-Shared Keys (PSKs) for authenticating communication relationships These PSKs are generated and distributed in the vehicle by a so-called Vehicle Key Master (VKM) function.

To facilitate the secure distribution of these PSKs from the VKM to the respective ECUs, an initial secure connection is required to be established between the ECUs and the VKM. This initial connection is typically established during vehicle assembly at the manufacturer's plant or, in certain cases, in a service garage when an ECU replacement is necessary. The process involves authentication using ECU Identification keys for A-ECUs, while S-ECUs use a specialised key known as the vehicle integration key for this purpose.

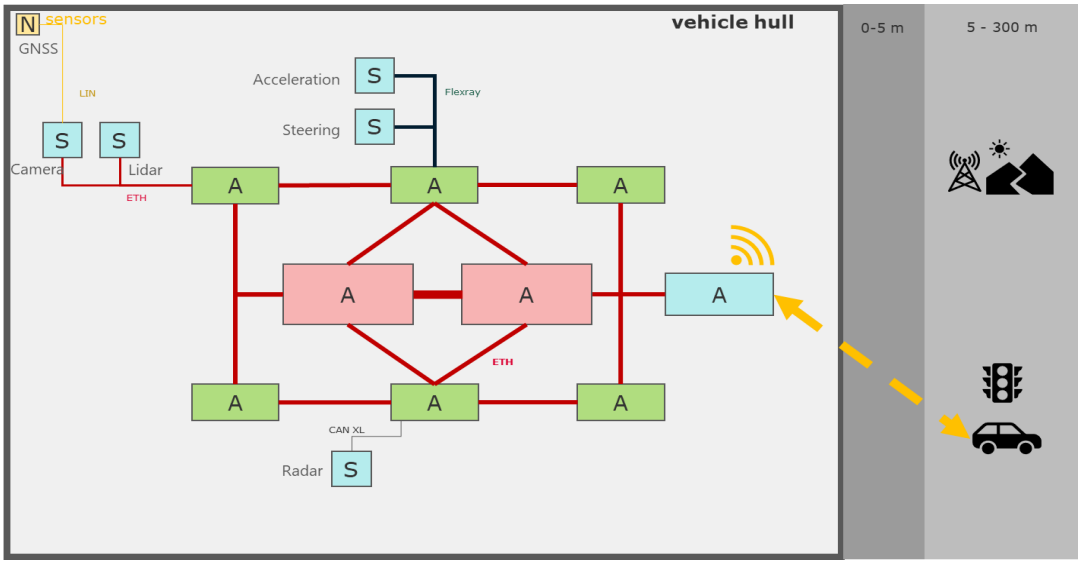

Figure 57 - Distribution of cryptographic capabilities of ECUs in the C-ACC use case

After this initial authentication, the VKM creates a symmetric key, called ECU-VKM, specific for each ECU in the given vehicle. This ECU-VKM key is then used to distribute or update function-specific keys (e.g., PSKs for Secure Communication).

Delving into more details, the following cryptographic keys are being used in the context of A-ECUs. It's important to clarify that this classification does not encompass the keys described in Chapter 6 of CONNECT's security and trust mechanisms, including attestation keys and other cryptographic primitives derived from the hardware-based key of the underlying RoT (i.e., CONNECT TEE guard). More information regarding the specific type and utilisation of these keys in enhancing the security of all communication will be provided in D4.1.

Table 15 - Keys in A-ECUs of C-ACC use case

Key	Type of crypto	Creation	Usage
ECU Identification	Asymmetric	In ECU at production	• Identify the ECU when integrating in the vehicle (production / maintenance garage) • Identify the ECU in the vehicle network at runtime
ECU Encryption	Asymmetric	In ECU at production	Encrypt data sent to the ECU and decrypt it in the ECU
Secure Communication	Symmetric	Vehicle Key Master	PSK to authenticate secure in-vehicle communication (MACsec, TLS or SecOC)

		(VKM) during vehicle integration	
ECU-VKM (Vehicle Key Master)	Symmetric	VKM during vehicle integration	PSK to authenticate confidential communication between the ECU and VKM to distribute/update in-vehicle keys (e.g., Secure Communication)
V2X	Asymmetric	During production / on the road	Special case for the Smart Antenna ECU. These keys include the V2X Enrolment keys and Authentication Tickets (Pseudonyms)

Additionally, in the context of S-ECUs, the following cryptographic keys are being used:

Table 16 - Keys in S-ECUs of C-ACC use case

Key	Type of crypto	Creation	Usage
ECU Integration key	Symmetric	In ECU at production	Identify the ECU when integrating in the vehicle (production / maintenance garage)
Secure Communication	Symmetric	VKM during vehicle integration	PSK to authenticate secure in-vehicle communication (MACsec, TLS or SecOC)
ECU-VKM (Vehicle Key Master)	Symmetric	VKM during vehicle integration	PSK to authenticate confidential communication between the ECU and VKM to distribute/update in-vehicle keys (e.g., Secure Communication)

7.3.1.4 In-Vehicle Key Management System

Many cybersecurity controls rely on cryptography and the security of the used cryptographic keys. Each of these keys have their own lifecycle which needs to be managed. This lifecycle includes the stages creation, distribution, operation/usage, and deletion.

Key creation:

- Most keys with a long lifetime - especially based on asymmetric cryptography – are created during production. These keys can serve the vehicle network (e.g., for component/ECU identification) but can also be use case specific to authenticate the vehicle against outside entities (e.g., cloud or road-side infrastructure).
- There is also the possibility to create keys in the vehicle. This can take place during production when the ECUs are integrated into the vehicle, but also in the field e.g., at maintenance workshops, during Over-The-Air (OTA) Updates, or in some use case-specific scenario. The creation of these keys can be managed by a vehicle-central component such as a Vehicle Key Manager (VKM), but also by distinct ECUs with a Trusted Execution Environment (TEE) serving specific use cases.

7.3.3 C-ACC Scenario Needs from CONNECT

The above introduced architecture and cybersecurity controls are decided in the C-ACC Item's design time where the service software stack and internal building blocks and functions are constructed. Within such a complex system**, the trustworthiness of the involved technical components and data items is of paramount importance to ensure the safe and effective operation of C-ACC.** This is particularly crucial for controlling critical functions like lane changes and speed adjustments based on proximity to preceding vehicles, with the overarching goal of preventing accidents and protecting vehicle passengers.

This trustworthiness needs to be assessed at the level of each relevant distinct artefacts. If a set of artefacts is homogenous with regard to trustworthiness, they can be treated the same and assessed

as a set. An artefact's trustworthiness depends on all aspects with potential influence on the artefact. Therefore, the granularity of the trustworthiness assessment needs to be per artefact (data item or execution platform. e.g., ECU).

For a data item, **the item's production, its communication, and all other relevant fringe elements (e.g., components with access to the same bus systems) need to be assessed**. As an example, for C-ACC's sensor data item "Vehicle position", the trustworthiness of the actively involved components (see Figure 7.3.4) GNSS, Camera, and Zonal Controller 1 as well as the passively involved components such as Lidar need to be assessed. The passively involved components' trustworthiness is relevant since they can influence the exchanged data flow by attacking the communication network and, hence, the shared data items.

To allow such detailed assessment, C-ACC needs to be prepared at design time. This includes appropriate documentation such as a Cybersecurity Threat Assessment and Risk Analysis (TARA) as well as policies to respond at runtime to changes in trustworthiness. The latter encompasses potential actions such as initiating a migration operation if the trustworthiness levels fall below a certain threshold. Additionally, it involves the pre-definition of the components of the C-ACC function that can be migrated, referred to as the "migratable function state". These definitions are crucial to be clarified at the design time, for ensuring a clear approach for managing trust-related events.

The engineers developing such a function need appropriate guidance on how to disassemble the function into an appropriately detailed set of technical artefacts; a methodology to decide on the required level of trustworthiness for each artefact; as well as a way to specify the trust relationships necessary to assess an artefact's trustworthiness in form of a trust model. One example for defining such a trust model can be found in Trkulja et al [150].

In order to introduce these assessments, the C-ACC needs the following work products from CONNECT:

- A clear definition of trust and level of trustworthiness that need to be depicted by the various artefacts before their produced data and/or outcome can be considered as part of the overall C-ACC functionality.
- Methodology on how to specify the Required Levels of Trustworthiness (RTL) during the system's design time. This is elaborated in section 3.2.3 of the present deliverable.
- A Trustworthiness Assessment Framework (TAF) with the capabilities to assess the system's actual trustworthiness at runtime based on sources of trustworthiness relevant to C-ACC's execution.
- A way for the C-ACC (i.e., via a trust model) to inform the TAF which components of the system are relevant to C-ACC, i.e., types of trust sources to be considered by the TAF. These relevant parts include the involved sensors, actors, and execution platforms that need to be continuously (or periodically) assessed prior to taking part in the overall establishment of the C-ACC workflow. Additionally, other system components might also be relevant such as other ECUs not involved in C-ACC but connected to the bus systems used by C-ACC that might target to alter the C-ACC operation (if compromised) by manipulating the content of the shared messages exchanged over the common communication bus (or disrupting the communication).
- Notifications to C-ACC when the actual trustworthiness changes.

In addition to this assessment, C-ACC needs guidance on how to specify appropriate response policies in case of the ATL not fulfilling the RTL anymore. This also includes policies for the scenario that C-ACC's execution platform (a Vehicle Computer ECU) is not trustworthy anymore and needs **to securely migrate relevant parts of C-ACC to other execution platforms to ensure operational continuity**. The application informs CONNECT beforehand about which parts are migratable. Finally, the migration needs to be verified whether it has been performed correctly.

7.3.4 "To-Be" Reference Scenario #1: C-ACC Collaboration and In-Vehicle Data Sharing Environment

The C-ACC item is supported by a dynamic Trust Assessment Framework (TAF). During design time, the Required Trust Level (RTL) for each respective artefact (e.g., components, data items in transit or in storage) for the relevant components of C-ACC have been specified, along with the necessary trust models are specified and circulated through the CONNECT Blockchain as described in section 6.7 of the present deliverable.

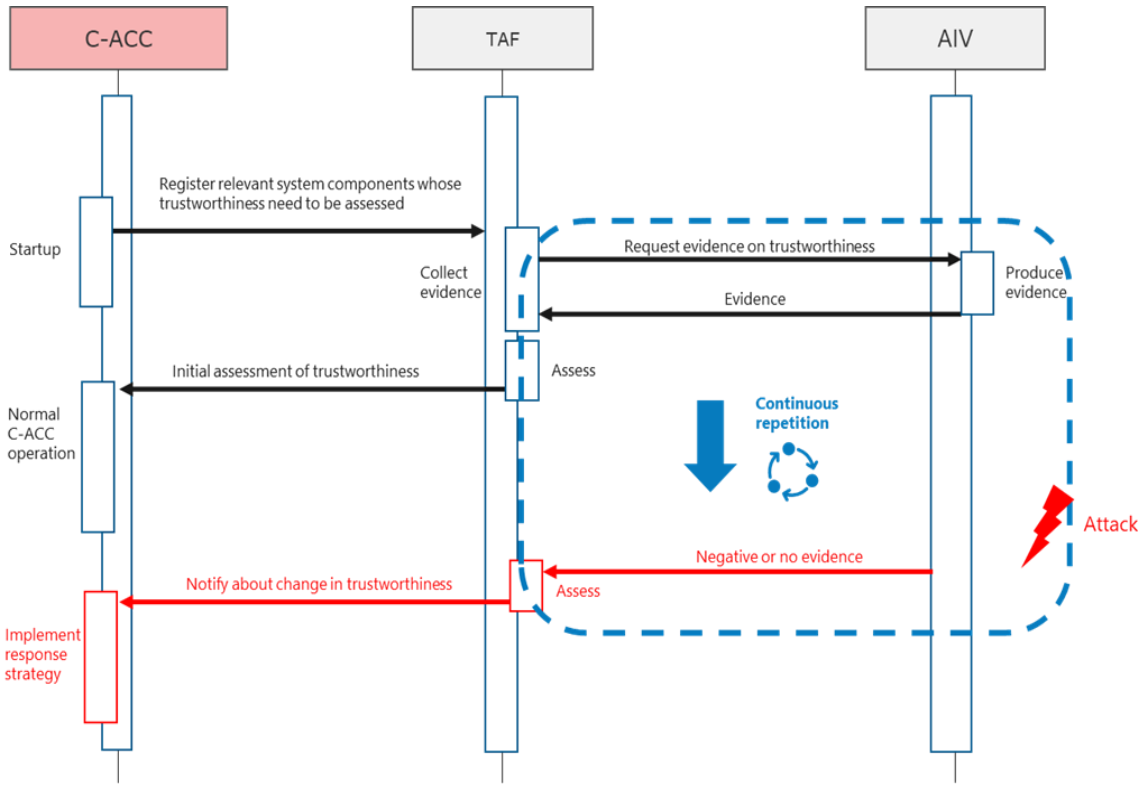

Figure 58 - Flow for assessment of trustworthiness in C-ACC use case

As shown in Figure 58, during start up, the TAF has three options to load or update the trust models:

- Option 1: the TAF loads the trust model from C-ACC directly which has extracted it from the Blockchain as was pushed by the Trust Management Component
- Option 2: the trust models are loaded from the vehicle architecture or updated based on vehicle-internal changes (e.g., components have been migrated/relocated between ECUs or the vehicle network is only partially in operation to save energy)
- Option 3: the TAF is notified for any updates regarding the trust models and the RTL through the Blockchain

With the Trust Models loaded, the C-ACC can send requests for Assessment of Trustworthiness (TAR) to the TAF which are then going to be forwarded to the Attestation & Integrity Verification (AIV) component for initiating the trustworthiness evidence collection (for the target ECUs) through the deployed security controls - in this context, the CONNECT runtime attestation mechanisms. Next to requesting a single assessment, such a TAR can also include a request for continuously monitoring the trustworthiness level of a device by asynchronously receiving the produced trustworthiness (attestation) evidence when there is a change in the stat of the respective device. After the notification of a TAR, the TAF then collects relevant evidence from the relevant C-ACC items, as part of the overall service graph chain, from the Trust Source which receives them from the AIV. Trust-related information (i.e., through exchanged Trustworthiness Claims) from neighbouring vehicles can also be considered. With this evidence, the TAF assesses the actual trustworthiness and forwards it to C-ACC. C-ACC compares the actual with the required trustworthiness level. If the requirements are fulfilled, C-ACC will start its operation (or react to a specific change in the level of trust based on predefined policies).

The TAF keeps monitoring the trust sources (e.g., by periodically requesting up-to-date proofs). If one of the requested evidence is not provided anymore or is negative (e.g., indication that secure

boot failed due to a manipulation), the TAF will reassess the trustworthiness and notify the C-ACC about the change.

C-ACC will then use this runtime assessment of trustworthiness and compare it with its requirements. **If the reduction in trustworthiness is relevant to C-ACC, it will decide on executing pre-defined (at design time) response strategies.** Such strategies can include:

- treating input data from less trustworthy sources more cautiously.
- introduce additional plausibility checks on data (e.g., checking with correlating other sensor data if there is a conflict).
- increase logging and reporting to the OEM.
- increase safety buffer (e.g., distance between vehicles).
- relocating C-ACC components between execution platforms, (e.g., another ECU or to the MEC) based on the predefined policies, and
- reduce the C-ACC's functionality (including turning it off as a last resort).

7.3.5 Reference Scenario User Stories

In the following user story workflows, the blue arrow depicts preceding actions and flows performed before the respective story's flow starts.

[CACC.US.1]: As C-ACC, I want to improve the safety and reliability of my function through dynamic trustworthiness assessments of how its input data items are generated (whether from local sensors or other vehicles) and how they are transferred to me. I want to be able to assess the trustworthiness to serve two different driving situations:

- **Imminent driving situation: perform a quick assessment.**
- **Upcoming driving situation: perform a very detailed assessment with high confidence in its results to carefully prepare.**

User Story Confirmation:

In this user story the C-ACC receives an assessment of the trustworthiness of its input data items based on the assessment of **static (representing in-vehicle objects such as ECUs) and dynamic (representing other objects in the affinity such as detected moving vehicles) trust objects** by the TAF. This assessment includes an evaluation of relevant evidence (e.g., about the data items' integrity) within the required time limit depending on the driving situation.

In some driving situations, there is not enough time to perform an assessment based on full freshly collected evidence (e.g., when there is only 1 second to decide to avoid a collision). In such situations, the assessment is done on a best effort basis using mostly cached (previously collected) evidence data. This assessment might not be as accurate as a full assessment but can help with driving decisions in time-critical situations. For upcoming driving situations, there is enough time to properly plan driving decisions. In such a situation the assessment can be detailed based on freshly collected evidence. In such situations, the available time limit is larger, but not unlimited. In some situations, the collection of fresh evidence can take even longer than permitted by the available time, leading to the assessment being based on a mixture of cached and fresh evidence.

CONNECT Functionalities:

✓ **Trust Assessment, Trustworthy Platform Configuration and Attestation**

User Story Workflow: (*see Figure 59)*

The workflow starts when C-ACC receives a Data Item A from a sensor or another vehicle and needs information on this data item's trustworthiness. Depending on the driving situation, C-ACC requests a detailed or a quick trustworthiness assessment from the TAF. In case of a detailed assessment,

the TAF requests up-to-date evidence of trustworthiness regarding how the data item was produced from the respective sensor and evidence about the transfer of the data item from its source to C-ACC from Supplemental ECUs. Based on this collected evidence, the TAF assesses the Actual Trustworthiness Level (ATL) of all in-vehicle components which provide input to the C-ACC application. Moreover, during C-ACC run-time, other vehicles send C-ACC relevant data to the ego-vehicle along with the ATLs they assessed on their data. The TAF requests trustworthiness evidence from every vehicle it receives C-ACC-relevant data from. It uses this evidence to create opinions on the trustworthiness of vehicles not to compromise the C-ACC data sent to the ego-vehicle. The ego-vehicle's TAF then uses trust discounting to produce its ATLs on the C-ACC data received from other vehicles. Once all the necessary ATLs have been assessed, the TAF informs the C-ACC application.

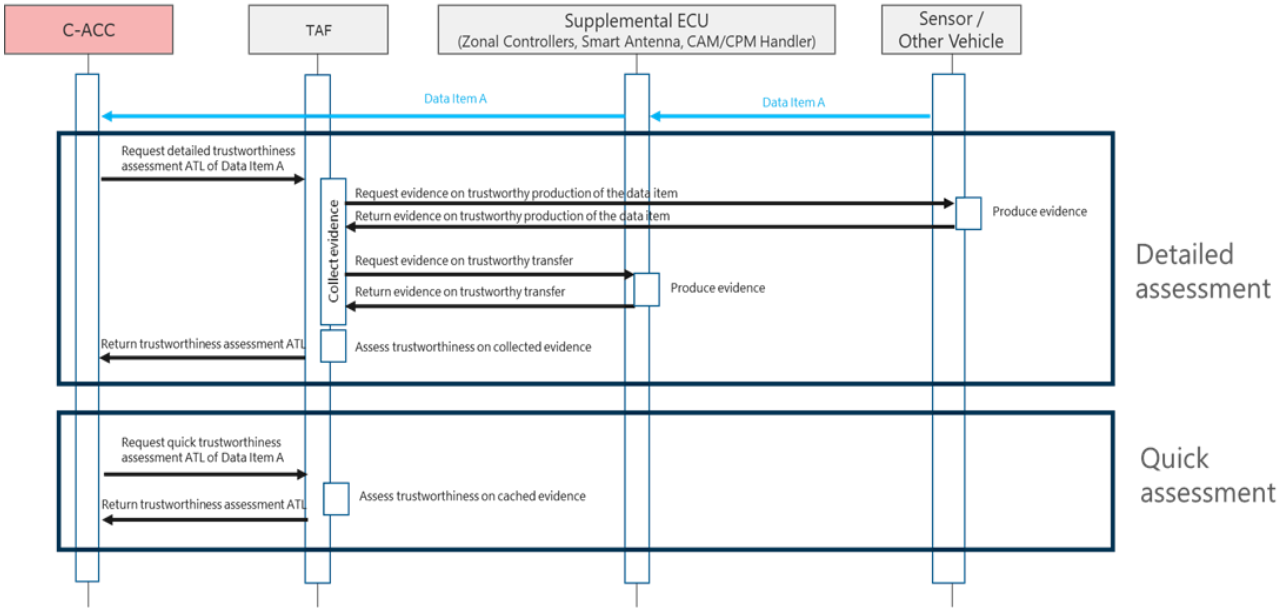

Figure 59 - Representation of the sequence diagram of user story [CACC.US.1]

CONNECT KPIs:

Table 17 - [CACC.US.1] KPIs

KPIs	Description	Value
	ATL Correctness (i.e., C-ACC Main Component receives correct ATL for requested component or data item)	**TRUE**
	Latency for imminent driving situation	**<=100 ms** delay for the TAF to respond to C-ACC request when the TAF is instantiated and executed as part of the application software stack in the target system (outside the CONNECT TEE) **<=200 ms** delay for the TAF to respond to C-ACC request when the TAF is instantiated and executed within the CONNECT TEE.
	Latency for upcoming driving situation	TAF responds to C-ACC request with assessment based on freshly collected evidence with latency **< 2sec**

[CACC.US.2]: [3] As C-ACC, I want to be able to compare the Actual with the Required Level of Trustworthiness for a specific node or data item as part of the overall service graph chain.

User Story Confirmation:

The C-ACC function has knowledge about the required trustworthiness level as defined at the function's design time and deployed together with the trust model template. **This required and the actual level of trustworthiness produced by the TAF at runtime need to be comparable at a semantic level.** With such a comparison, the C-ACC function can make a decision on the vehicle's behaviour based on whether the actual trustworthiness level matches or fulfils the required trustworthiness.

User Story Workflow: *(see Figure 60)*

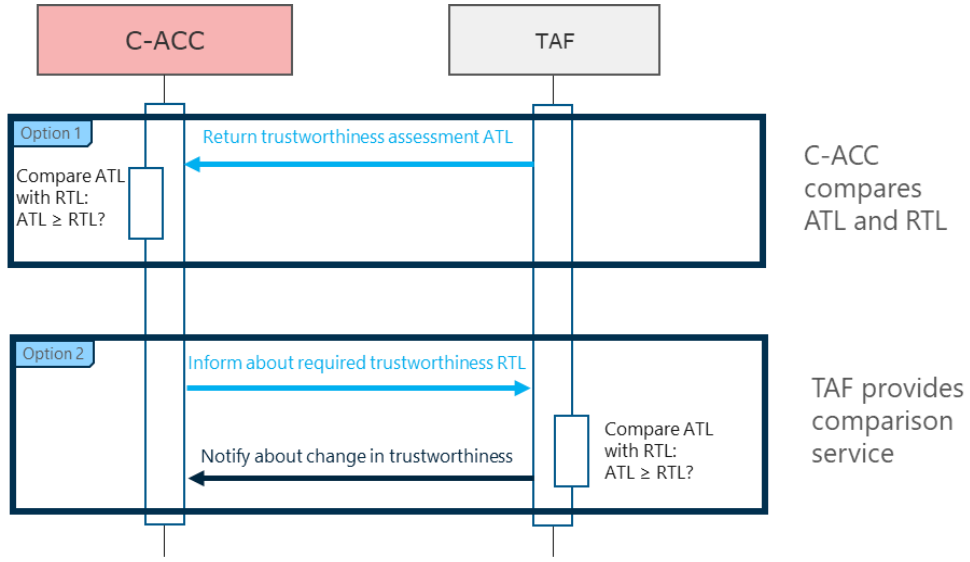

Figure 60 - Representation of the sequence diagram of user story [CACC.US.2]

For this workflow, **the CCAM function C-ACC compares the RTL and ATL directly:** the workflow starts with C-ACC receiving the trustworthiness assessment (ATL) from the TAF (see user story CACC.US.1). C-ACC compares the received ATL with its pre-defined required trustworthiness (RTL). After this comparison, C-ACC knows whether the required trustworthiness is actually fulfilled or not.

CONNECT KPIs:

Table 18 - [CACC.US.2] KPIs

KPIs	Description	Value
	Comparable ATL and RTL	ATL and RTL can be compared by the TAF within a timeslot of **50 ms** (time needed for enabling a context switch between trusted and untrusted world)
	Ability to calculate	C-ACC Main Component can calculate whether the ATL is below the RTL
	Calculation time	**< 10 ms**

[CACC.US.3]: As the C-ACC Item, based on my policies, I want my C-ACC Main Component to be able to respond to changes in the trustworthiness of my execution platform (Vehicle Computer) by migrating to a more trustworthy execution platform to ensure the correctness and safety of my functionality.

User Story Confirmation:

In this user story, C-ACC can perform a migration process to move its C-ACC Main Component from one Vehicle Computer ECU to another Vehicle Computer ECU because the current Vehicle computer is not considered trustworthy enough anymore based on the evidence and trust assessment provided by the ATL. Migrating to another component outside the vehicle (e.g., MEC) is not possible since the new communication relationships between the C-ACC Main Component and the vehicle's sensors and actors would have too much delay that would affect the safety profile of the overall C-ACC service and, thus, the vehicle itself. The migration process can start preparing a new instance of the C-ACC Main Component to the target ECU while the function continues its operation. When the new instance is ready, C-ACC switches over to the new instance.

User Story Workflow: *(see Figure 61)*

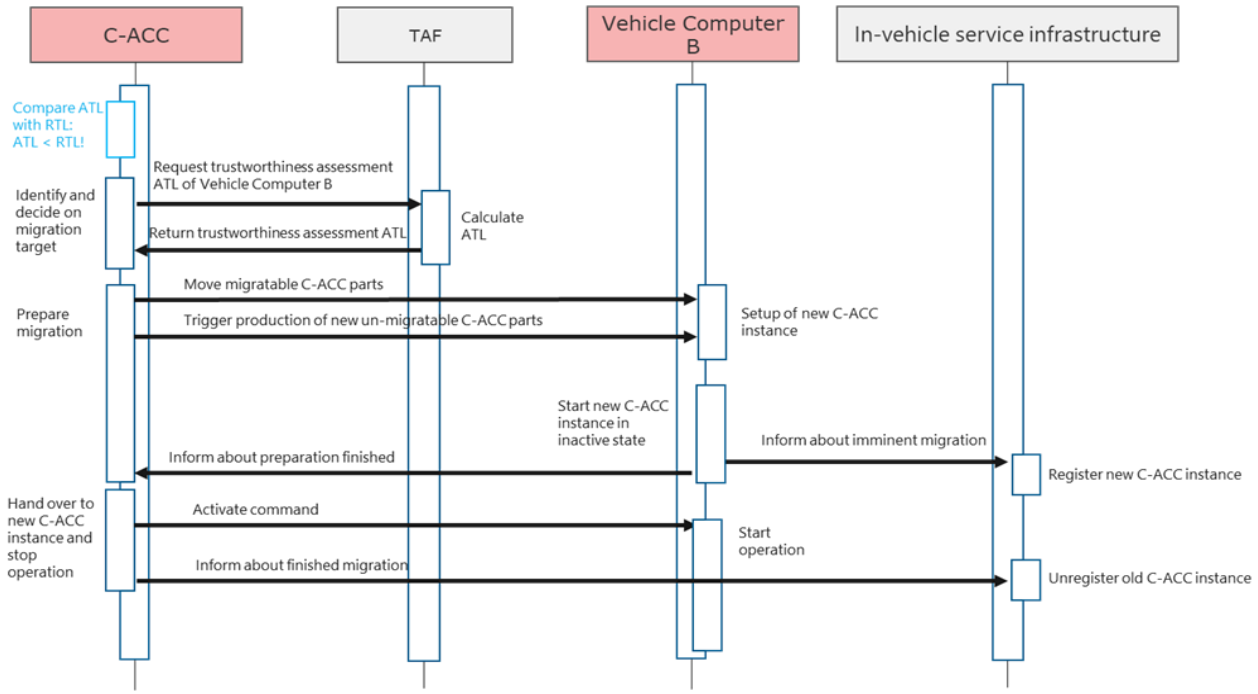

Figure 61 - Representation of the sequence diagram of user story [CACC.US.3]

The workflow starts with C-ACC noticing that the actual trustworthiness in its execution platform, the Vehicle Computer is not fulfilled anymore and decides to migrate its function to another Vehicle Computer B. First, C-ACC requests a trustworthiness assessment about Vehicle Computer B from the TAF to find out if Vehicle Computer B fulfils C-ACC's trustworthiness requirements. If the requirements are fulfilled, C-ACC starts the migration by preparing a new instance of C-ACC on Vehicle Computer B. This preparation includes moving the migratable parts of C-ACC to Vehicle Computer B and triggering the production of new un-migratable parts on Vehicle Computer B. Vehicle Computer B uses these parts to set up a new instance of C-ACC on its platform. When this new instance is set up, Vehicle Computer B starts it, but in an inactive state, and inform the In-vehicle service infrastructure about the new instance. With these steps done, Vehicle Computer B informs the old instance of C-ACC of the migration preparation having finished. C-ACC then hands over the operation to the new instance by sending an activate command to Vehicle Computer B to start the new instance's operation. Finally, the old instance stops its operation and informs the In-vehicle service infrastructure that the migration has finished.

CONNECT KPIs:

Table 19 - [CACC.US.3] KPIs

KPIs	Description	Value

	C-ACC Execution platform migration (i.e., C-ACC can select another execution platform (Vehicle Computer) fulfilling the required trustworthiness as migration target;)	For this selection, the target platform's ATL needs to fulfil C-ACC's RTL;
	C-ACC Function Migration C-ACC can prepare the function migration (e.g., migrating necessary keys to the target platform or create new keys on the target platform, starting up the new C-ACC Main Component in an inactive state and informing the in-vehicle service infrastructure of the upcoming migration)	**TRUE**
	Disabling Old Version C-ACC (i.e., C-ACC can disable the old C-ACC Main Component on the leaving platform and activate the new C-ACC Main Component on the target platform)	**TRUE**
	C-ACC function downtime during migration	**< 1sec**; During this time, the C-ACC is not operational;
	C-ACC successful migration evidence generation (i.e., C-ACC can produce evidence of successfully completing the migration process)	process as shown in the workflow (considering also the notification on the In-Vehicle Infrastructure) **< 10sec.** Note that until the evidence of successful migration is produced, the resulting uncertainty is reflected in the C-ACC Main Components trustworthiness.

[CACC.US.4]: As the C-ACC Item, based on my policies, I want my C-ACC Main Component to be able to respond to changes in the trustworthiness in data items I receive from in-vehicle and vehicle-external sensors. I want to compensate for the change of trustworthiness by enacting pre-defined safe responses such as adapting my function-internal calculations or increasing the distance to other vehicles.

User Story Confirmation:

With knowledge about the ATL being below the RTL, C-ACC is able to select between pre-defined fallback plans to enact and continue its functionality.

User Story Workflow: *(see Figure 62)*

Similar to user story CACC.US.3, the workflow starts with C-ACC noticing that the actual trustworthiness in received data items is not fulfilled anymore. With this knowledge, C-ACC decides on implementing pre-defined, alternative fallback plans (e.g., adapting functional calculations to increase safety distance between vehicles).

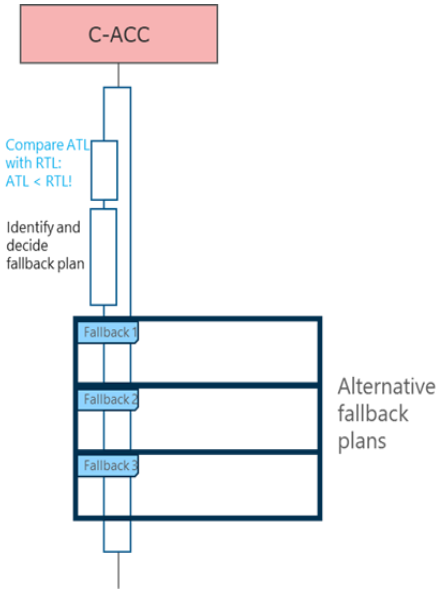

Figure 62 - Representation of the sequence diagram of user story [CACC.US.4]

CONNECT KPIs:

Table 20 - [CACC.US.4] KPIs

KPIs	Description	Value
	C-ACC Main Component changes the treatment of sensor data in its internal calculations	**TRUE**
	Update of calculation method (e.g., updated weights in how different types of data items affect the decision of the C-ACC function)	**< 50ms**

7.4 Slow Moving Traffic Detection (SMTD) Use Case

The Slow-Moving Traffic Detection (SMTD) system addresses traffic congestion by utilizing V2X communication technologies. It enables equipped vehicles to share real-time information about slow-moving vehicles on the road, enhancing road safety, reducing congestion, optimizing transport efficiency, and minimizing environmental impacts. Cooperative Perception Messages (CPMs) and Cooperative Awareness Messages (CAMs) play a crucial role, providing real-time environmental and kinematic data securely signed with short-term anonymous credentials. These messages could then be processed by a Mobile Edge Computing (MEC) server, acting as a central hub on the edge of the network for data analysis. The MEC server decodes the incoming V2X messages and checks the correctness of the received data or for possible contradictions between multiple sources. The SMTD system focuses on real-time detection of slow-moving traffic, crucial for road safety and traffic management. Legacy vehicles lacking V2X capabilities rely on ADAS sensors for real-time detection, but without V2X, they cannot share this information.

The security challenge lies in the unverified correctness of messages from V2X-equipped vehicles. In this context, the Trust Assessment Framework (TAF) becomes pivotal. It is employed at the MEC server to assess the trustworthiness of data received from V2X-equipped vehicles. The TAF evaluates input validation, identifying potential misbehaviors or contradictions among data sources. The data, along with the trust opinion, is then forwarded to the Traffic Control Centre (TCC). The TCC utilizes this information to create a high-definition Local Dynamic Map (LDM),

displaying trustworthy real-time events for traffic monitoring purposes. In case the TCC detects a slow-moving traffic situation, it issues an alert disseminated through a DENM message to all the vehicles approaching the hazardous area and a warning is displayed to the drivers.

False alert messages can cause not only misleading, but also dangerous situations for all road users. This can be caused not only by a malicious sender, but also due to faulty or inaccurate sensors. By adding attestation and security claims on messages sent, CONNECT makes it possible to exploit the data veracity created by ITS stations sending trusted messages. Thanks to this, we can recreate a Misbehaviour Detector instance able to assign a trust level to ITS stations by analysing their messages sent. The TAF plays a crucial role in assigning trust levels to vehicles based on their messages, allowing the system to identify and mitigate misbehaving entities. This enhances the overall security of the SMTD system, ensuring that alerts and notifications are based on trustworthy information.

To address the pervasive issue of traffic congestion within the EU, as previously explained in section 7.3, a collaborative effort involving various stakeholders has led to the development of advanced technologies aimed at improving the efficiency, safety, and environmental performance of road transport systems. One such cutting-edge solution is the Slow-Moving Traffic Detection (SMTD) system. SMTD leverages V2X communication technologies, allowing equipped vehicles to communicate essential information. This system encompasses both V2X-equipped and non-V2X vehicles, enabling a comprehensive network for information sharing. SMTD that can enable the timely dissemination of information about slow moving vehicles on the road, thus, improving road safety, reducing congestion, optimising transport efficiency, enhancing mobility, reducing energy use and environmental impacts.

In the context of SMTD, Cooperative Perception Messages (CPMs) [151] and Cooperative Awareness Messages (CAMs) [152] play a pivotal role. Both messages are ETSI standardised. On the one hand, CPMs are used for sharing real-time environmental data and the perception of road conditions by the equipped vehicles. CAMs on the other hand, are responsible for kinematic data. Both CAM and CPM messages, generated within the SMTD system, are securely signed, leveraging the short-term anonymous credentials from a PKI-like entity. This adds an additional layer of trust and security. Afterwards the two messages are sent to a MEC server, in order to be processed, via Vehicle-to-Network (V2N) communication.

The MEC server acts as a central hub for processing and analysing the data. It decodes the incoming V2X messages, which are typically encoded in the standard ASN.1 format [153] and extracts the essential information. The MEC is further responsible for checking the correctness of the received data. This pertains to the fact that there are potentially multiple data sources (i.e., vehicles); hence, the incoming information might be contradictory. In such a case, misbehaviour may be detected. This information is crucial to be checked prior to being sent to the Traffic Control Centre (TCC). To perform this input validation, a Trust Assessment Framework (TAF) must be used, to export a trust opinion regarding the received information.

After evaluating the trustworthiness of the received information, the data along with the trust opinion is forwarded to a Traffic Control Centre (TCC), where the information that is deemed as “trustworthy” is displayed on a Local Dynamic Map (LDM). The LDM serves as a high-definition map environment, showcasing real-time events for traffic monitoring purposes. Both CAMs and CPMs provide up-to-date, real-time information regarding the precise location of vehicles; hence the map visualisation is precise.

Following this innovative approach, equipped vehicles may automatically detect slow-moving or halted traffic conditions on their own. Therefore, real-time notifications about such situations are promptly forwarded to the Traffic Control Centre and visualised on the map environment. These notifications trigger the generation of the ETSI-standardised Decentralised Environmental Notification Messages (DENMs) [101]. Each DENM is tagged with a specific area, alerting approaching V2X-equipped vehicles about the approaching slow-moving traffic ahead. In the long run, this preventative system helps drivers slow down or speed up as needed, which improves road safety, lessens congestion, maximises transportation efficiency, expands accessibility, reduces energy use, and lessens environmental consequences.

The following figure (i.e., Figure 63) demonstrated a Slow-Moving Traffic Detection (SMTD) scenario where the slow moving, non-V2X vehicle (i.e., the blue one) is detected by a following V2X vehicle (i.e., the white one) leveraging its camera sensor. The first, the conventional 'legacy' vehicle, is lacking V2X connectivity, while the second vehicle is equipped with both a front camera sensor and a telematic box designed for sending and receiving specialised ITS messages. This telematic box is referred to as the Vehicle-to-Everything On-Board Unit, or V2X OBU for short. The V2X-connected vehicle leverages its front camera sensor to perceive the presence and speed of the vehicle(s) ahead. This capability enables it to detect potential instances of slow-moving or stationary traffic conditions.

In the given scenario, as the vehicle ahead is detected, the V2X-enabled vehicle creates and sends ETSI-standard CPMs (alongside its own CAMs) towards a MECC server, via 5G connection. The MECC server detects the slow-moving traffic situation, and it generates a DENM alert message sent via 5G to all vehicles in a dedicated approaching area (i.e., the black vehicle). In this way, all the V2X vehicles approaching can be informed of the slow-moving traffic situation ahead of them.

We have to note that in the context of the experimentation CONNECT will deploy a MEC infrastructure. This MEC deployment follows the standards, as defined by ETSI, functionally wise; nevertheless, it does not provide support over the communication medium (i.e., 5G network plane).

Hence, whenever we use the term CONNECT MEC or MECC we refer to our instance. This element is introduced specifically for the SMTD use case.

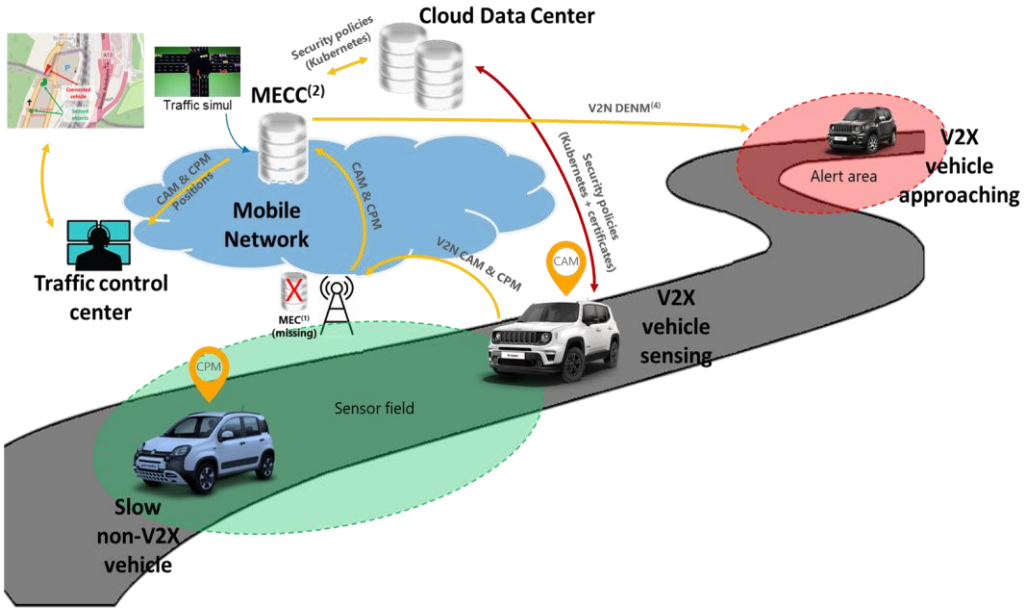

Figure 63 - Representation of the Slow-Moving Traffic Detection (SMTD) use case - Misbehaviour Case

7.4.1 "As-Is" scenario

This use case focuses on the real-time detection and identification of a slow-moving traffic condition on the road, commonly referred to as the Slow-Moving Traffic Detection (SMTD) use case. The foundation of this use case stems from the "Detection of Non-Connected Road Users" use case, proposed in the standard ETSI TR 103 562 V2.1.1 (2019-12) [151]. This use case is particularly important in the context of road safety and traffic management.

Imagine a scenario on a highway where a slow-moving traffic situation may occur, due to roadworks or other unforeseen events. When such a situation happens, the approaching vehicles have to abruptly decelerate from high travelling speeds. This sudden change in speed can lead to potential hazards. In addition, if the approaching vehicle has the Automated Cruise Control (ACC) activated, the driver may not be immediately ready to take over the control of the vehicle and avert a collision.

In this context, the value of the SMTD use case becomes evident. This proactive communication not only enhances road safety but also contributes significantly to effective traffic management.

In the current road landscape, we encounter an “As-Is” scenario, where a limited number of vehicles are equipped with V2X communication technologies. The majority of vehicles on the road are “legacy” vehicles without V2X capabilities, relying solely on GNSS receivers for position tracking and LTE cellular connections for connecting a Human-Machine Interface (HMI). Furthermore, depending on the automation level, these legacy vehicles could be equipped with ADAS sensors such as proximity sensors, radar, LiDAR, cameras, and more.

In many large cities, traffic management is overseen by a TCC. This centre not only monitors real-time traffic situations providing alerts for potential risks in case of hazardous events, deploying traffic sensors throughout the urban landscape, and conducting in-depth traffic analysis based on the collected data.

In the 'As-Is' scenario, when a slow-moving traffic situation occurs, the traffic control centre typically receives information through manual notifications from drivers using dedicated apps like Waze. Alternatively, they may become aware of the situation after it persists for a certain duration, at which point it is automatically detected. Some comprehensive traffic control applications, like Google Maps, can identify slow-moving traffic, but only after it has been ongoing for some time.

Vehicles equipped with ADAS sensors (e.g., radar, LiDAR, cameras, etc.) are able to detect a slow-moving traffic situation in real-time, but, if not equipped with a V2X OBU, they are not able to share this info with other road entities or a traffic control centre.

From a security perspective, the limited number of vehicles currently transmitting V2X messages have a specific data transmission protocol in place. These vehicles transmit data in plaintext within a 'to-be-signed-data' field, and they augment this data packet with a digest of a certificate and an ECDSA NIST P-256 signature encompassing the entire packet. Additionally, approximately every two seconds, the entire certificate is appended to the message, rather than just the digest. Notably, while the signature varies with each message, the certificate (and thus the digest) remains consistent throughout its validity period, which spans 7 days. In terms of data size, the observed certificate occupies 151 bytes (with its digest a mere 9 bytes), and the signature extends to 66 bytes. It's essential to highlight that these practices align with the standard [99], which outlines the security header and certificate formats for ITS messages.

Another important aspect to consider regarding the overall security of the system, in the as-is scenario, is that the messages sent by the vehicles are not verified for correctness. As a result, the TCC is unable to determine if the information is indeed trustworthy or whether it contains fraudulent data.

As discussed earlier in this chapter, the collection of CAMs and CPMs originates from vehicles equipped with V2X capabilities and is transmitted to the MECC infrastructure where the SMTD system is deployed. Within this MEC environment, the received data undergoes meticulous processing and evaluation to determine its trustworthiness. Subsequently, this valuable information, accompanied by the corresponding trust opinion is forwarded to the TCC which is responsible for generating the DENM designed for dissemination within the affected areas.

Imagine a scenario where the slow-moving vehicle (i.e., the blue one as depicted in Figure 64) is correctly detected by two V2X vehicles, which leverage their camera sensor (i.e., the white one on the right and the white one on the bottom). However, a third simulated vehicle (i.e., the white one on the top) fails to detect the slow-moving vehicle, either due to inaccurate data of its sensor or due to a misbehaviour. As a result, only two out of the three vehicles manage to successfully detect the slow-moving situation, and thus generate CPMs towards the MECC. The misbehaving vehicle will only send CAMs to the MECC. Nonetheless, this behaviour enables the MECC to promptly identify the misbehaviour and subsequently assign a lower level of trust to the corresponding vehicle.

As aforementioned, the Slow-moving Traffic Detection use case stems from the “Detection of Non-Connected Road Users” use case proposed in the standard ETSI TR 103 562 V2.1.1 (2019-12) [151]. However, in the standard ETSI TR 103 562 V2.1.1 (2019-12) there is no consideration of any aspects about the trustworthiness of the received data, which is the core factor of CONNECT. Indeed, a malicious sender may transmit a fictitious message, signalling a slow-moving traffic event that is not actually there. In this case, if the whole system is lacking any integrity and security checks,

there may be an alert issued for a false hazardous event. False alert messages can cause not only misleading, but also dangerous situations for all road users. Note that this misbehaviour can be caused not only by a malicious sender, but also due to faulty or inaccurate sensors. By adding attestation and security claims on messages sent, it is possible to exploit the data veracity created by ITS stations sending trusted messages. Thanks to this, we can recreate a Misbehaviour Detector instance able to assign a trust level to ITS stations by analysing their messages sent.

An innovative aspect of this scenario, which extends beyond the current state of traffic management, is the critical evaluation of the data's Level of Trust before it reaches the TCC. Towards this direction CONNECT will leverage its Trust Assessment Framework.

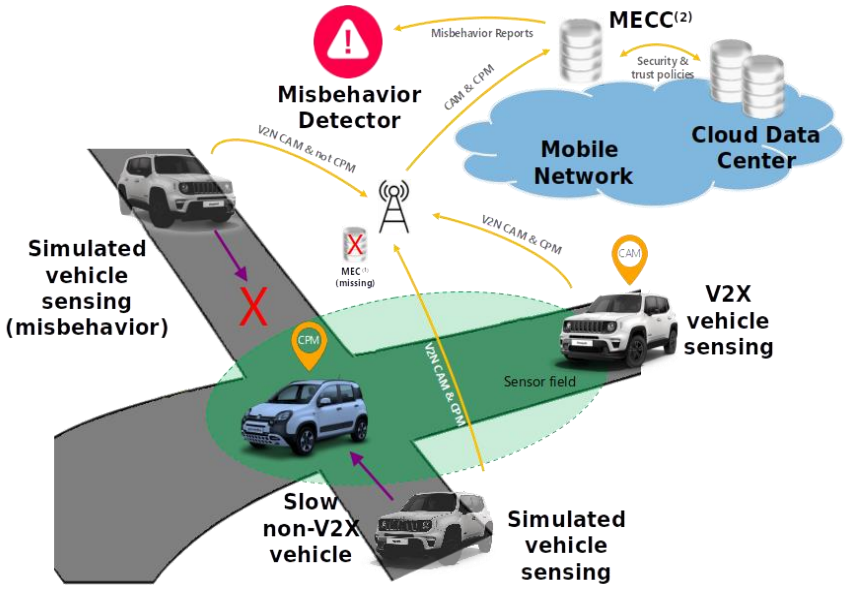

Figure 64 - Representation of a misbehaviour detection.

Figure 65 presents the high-level architecture diagram of the Misbehaviour Detection for SMTD. The vehicle is equipped with a GNSS receiver with IMU and a V2X OBU. The latter includes the V2X Module as well as the local instance of the TAF. Radar information is further sent to the V2X Module, leveraging the CAN bus. On the MECC side, the SMTD is being deployed, to be in close proximity with the vehicles, thus minimising the effects of latency. The SMTD analyses the information coming from the CAM and CPM. as sent by the vehicles. This information is being assessed by the TAF to export a trust opinion. The information (including information regarding potential misbehaviour) along with the trust score is sent to the TCC to create the DENM messages and inform other vehicles.

7.4.2 Entities, Actors, In-Vehicle Components, Communication Interfaces and Messages in the context of SMTD

Building upon the main stakeholders of the CCAM landscape, as identified in chapter 4, in what follows we mention specific components and roles that we consider in the specific UC. The actors involved in this use case will be:

- **Non-V2X vehicle:** a normal “legacy” vehicle that is not equipped with V2X capabilities to communicate with other road entities.

  Note that it may be connected in the broad sense with a GNSS and/or a classic LTE connection. In fact, most of the cars today on the road are provided with a classic LTE connection in order to exchange data only between the vehicle and its corresponding car maker. The type of data exchanged are not ETSI-standard type of messages but only sparse

data regarding the infotainment or the map update of the navigation system. It is important to underline that those kinds of vehicles are not able to communicate their information with other entities using the V2X standard vehicular communication technologies.

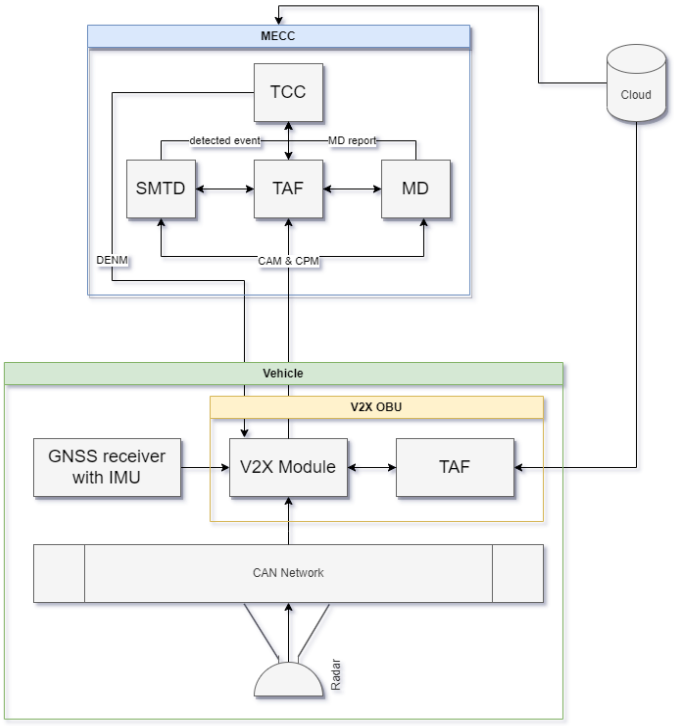

Figure 65 - Representation of the CAN bus data acquisition, the V2X OBU architecture and the interactions with the MECC.

- **V2X vehicles with front camera sensor:** a couple of vehicles equipped with a front camera system and a V2X OBU. Leveraging the camera system, they are capable of detecting objects in front of themselves and compute the distance and speed of the detected elements. They are also equipped with V2X OBUs connected to the CAN bus network of the vehicle, through which they are capable of sending/receiving standard vehicular V2X messages and to perform any kind of V2X communication (Vehicle-to-Vehicle, Vehicle-to-Infrastructure, Vehicle-to-Network, etc.). It has to be noted that the CAN bus and the camera sensor used in this use case are the real one of the vehicles according to the existing in-vehicle topology, as depicted in Figure 64.
- **Multi-access Edge Cloud Computing (MECC) server:** an ETSI-standard Multi-access Edge Computing (MEC) server is a server directly connected to the cellular antenna in order to stay close to end users (at the edge of the network) and therefore significantly reduce latency. Since in this project there is not a Mobile Network Operator participating, we cannot have direct access to a cellular antenna. For this reason, the MEC server instance we are using is not an ETSI-standard MEC but is a public IP server with an AMQP Broker instance on it and we will call it Multi-access Edge Cloud Computing (MECC) server. This element is introduced specifically for the SMTD use case.
- **Cloud Data centre:** this hardware and software element is entitled to manage the security and trust functions in the whole scenario. It is also the top-level element of a Kubernetes network which also includes the MECC server. The Cloud Data centre will be entitled to manage a "Trust Assessment" container that will be running on the MECC server. This element is introduced specifically for the SMTD use case.
- **Traffic Control centre:** it is an entity capable of monitoring the situation on the road in real-time, thanks to the incoming V2X messages. It may also act to some hazardous events by issuing an alert to a specific area. It is therefore capable of producing real-time large-scale

maps visualisation of an area of interest with all the real-time instances of monitored road events, so called Local Dynamic Maps (LDMs).

- **Local Dynamic Map (LDM):** it is a real-time map visualisation of road events, as standardised in the standard ETSI EN 302 895 V1.1.1 (2014-09) [154]. All the V2X messages (CAMs, CPMs, DENMs, etc.) received from the Traffic Control centre will be processed and sent to the LDM for visualisation thanks to the fields contained in the messages (station ID, station type, timestamp, position, speed, etc.). In this specific use case, we will visualise the real-time positions of both the V2X equipped vehicle and of the non-equipped vehicle, thanks to CPMs received.
- **Misbehaviour Detector (MD):** it is the entity capable of processing each incoming CAM and CPM extracting their kinematic data. Each data is then subjected to misbehaviour checks in order to assess whether the ITS message received is coming from a trusted vehicle or a misbehaviour one. This element is introduced specifically for the SMTD use case.

7.4.3 SMTD Scenario's Needs from CONNECT

In the SMTD use case, all ITS messages, including CAMs and CPMs, are transmitted with dedicated security features. These security features are implemented to ensure the trustworthiness and authenticity of the messages, aligning with the standards as outlined in the ETSI TS 103 097 [99]. This standard in specific, refers to the security headers and certificates that should be included in ITS messages. According to this, the SMTD scenario's needs from CONNECT are the following:

- **a dedicated public key certificate to be periodically added on ITS messages**. The sensitive info in the certificate, e.g., the issuer entity, may be hashed. Moreover, the signature of the certificate needs to be compressed, in order to reduce the total certificate size over the ITS messages. Ideally the certificate should be created by a dedicated public key infrastructure (PKI);
- **a digest of the public key certificate described above.** Since the public key certificate may be too large to be added on every ITS message to be sent, the full certificate will be attached to ITS messages just once in a while (e.g., once every five seconds or similar). All the other ITS messages will be certified just attaching the digest of the certificate previously described, since the digest it is usually few bytes long.
- **an ECDSA signature,** as described in the IEEE Std 1609.2-2022 (Revision of IEEE Std 1609.2-2016) [155] standard clauses 6.3.38 and 6.3.39, in order to certify all the content of the ITS message sent. A new ECDSA signature has to be generated for every new ITS message in order to sign the single content of the message itself. As for the signature of the certificate, the ECDSA signature has to be compressed, in order to reduce the total certificate size over the ITS messages.

All the security features listed above, will be provided by the Cloud Node directly to the OBU on the vehicle, thanks to a dedicated Kubernetes channel.

In the context of the CONNECT project we reproduce the "Detection of Non-Connected Road Users" use case adding data security claims on every ITS message in order to guarantee their data trustworthiness and authenticity. The process of faulty or misbehaviour data sent to the MECC and detected by the Misbehaviour Detector is depicted in Figure 64.

7.4.4 "To-Be" Reference Scenario

The complete flow of actions and messages exchanged in the SMTD use case is described in Figure 66. The camera sensor of a vehicle is able to detect an object (e.g., another vehicle) in front and create a message on the CAN bus with the corresponding info (i.e., the angles of the detected object and its distance from the ego position). Concurrently, the V2X OBU is able to read the kinematic data and the position accuracy level from the GNSS device and, consequently, construct ETSI CAM messages, including information of the ego vehicle. By reading the info of the sensed object from the CAN bus, the V2X OBU is also capable of constructing ETSI CPM messages.

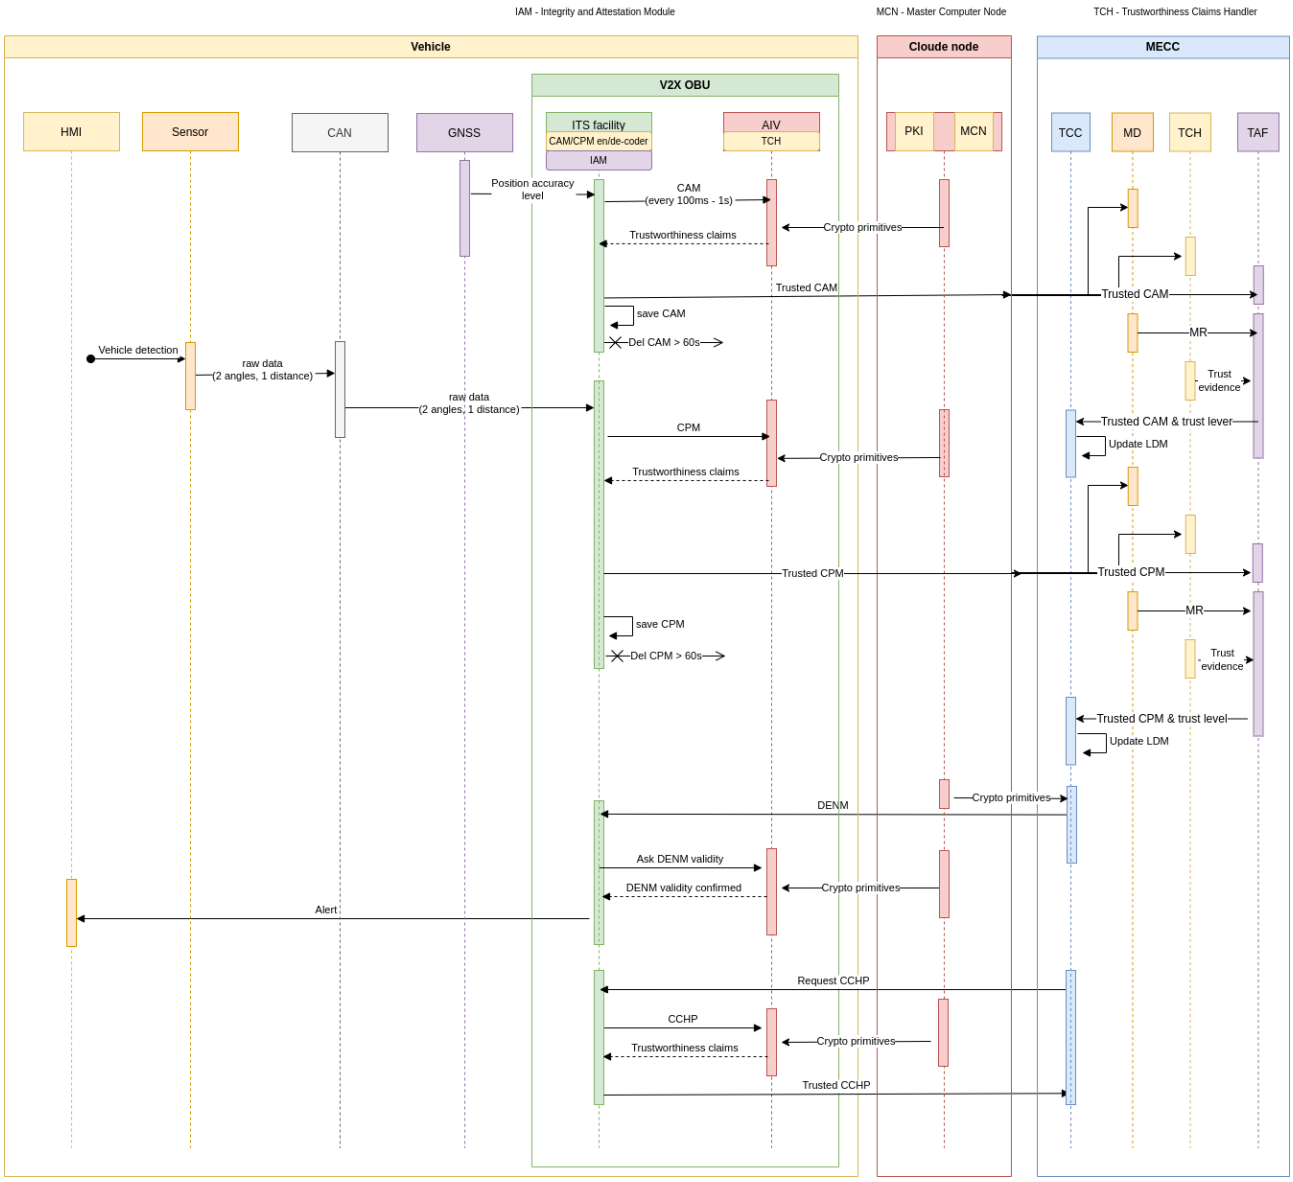

Figure 66 - Representation of a detailed sequence diagram of the Slow-Moving Traffic Detection use case

The generated CAM and CPM messages are then protected leveraging the crypto primitives (i.e., a public key certificate and short-term anonymous credentials (pseudonyms) for generating anonymous signatures) provided by the PKI module (instantiated on the Cloud Node). This PKI module essentially enables the alignment to the (ETSI standardised) PKI architecture. Then, the Identity & Authentication (IAM) module on the V2X OBU receives these short-term (pseudonymous) credentials which uses them for constructing an anonymous signature to be associated to the ready-to-transmit CAm and CPM messages (and their included Trustworthiness Claims). The IAM module is part of a Kubernetes network, and, thanks to this, it is able to communicate with the Master Computer Node (i.e., the master node of the Kubernetes network) on the Cloud Node from which the crypto primitives are transmitted. CAM and CPM messages with the crypto primitives attached become in this way verifiable CAM and CPM messages to be sent to the MECC.

Once received on the MECC's side, the SMTD processes the data and if any potential misbehaviour is detected, it is being reported by the MD module. In parallel, the trust levels of the ITS messages are assessed by the TAF. Both the information as well as the trust levels are used by the Traffic Control Center (TCC) on the MECC to recreate a Local Dynamic Map (LDM) instance. The LDM is a map instance on the MECC server which is constructed based on messages received from all vehicles (it is an equivalent instance of the so-called Global Dynamic Map of previous use cases).

By analysing the ITS messages received, the TCC may detect a slow-moving traffic situation. In this case, an alert should be issued for all the vehicles approaching the hazardous area. In a similar way with respect to the ITS messages' generation, the TCC receives crypto primitives by the Cloud Node (i.e., a public key certificate and a signature) to be inserted in the ETSI DENM message. The trusted DENM message is sent to all vehicles in a specific area.

Once a vehicle receives a DENM message, it checks its integrity and validity with respect to the expected crypto primitives sent by the Cloud Node and, once the DENM is correctly verified, an alert message is visualised on the HMI of the vehicle (i.e., a tablet) in order to alert the driver of the hazard ahead.

Simultaneously with this ongoing process, the V2X OBU collects and stores vehicle motion information from the last minute. Having this foresight is extremely useful in the case of a car accident. When this occurs, the collected motion data is promptly sent to the MECC server. The transmission of these CAMs and CPMs is executed over the same channel utilised for conveying trust-related information between the cloud node and the V2X OBU, ensuring data security and integrity.

The information received within the cloud node is considered a significant resource, specifically the data included in the CAM and CPM History Packet (CCHP). Insurance firms have the potential to utilise past data in order to recreate the dynamics and conditions that led to the occurrence of an accident. These insights play a crucial role in expediting the process of claims processing and assessing responsibility.

7.4.5 SMTD Security Features

In addition to the fundamental scenario outlined in Section 7.4.1, our objective is to further establish a "network of trust" among all relationships, whether they pertain to data exchange or node interactions. Towards this direction, CONNECT trust anchors will be employed.

The central element in this trust ecosystem is the Cloud Node, which assumes the role of the master node within a Kubernetes network. This Cloud Node is responsible for managing the security and trust policies among nodes and providing cryptographic primitives to both the MECC and the equipped vehicles.

Regarding the state-of-the-art on misbehaviour detection systems, the ETSI TR 103 460 V2.1.1 (2020-10) [156] collects several previous studies on misbehaviour detection and misbehaviour reports systems in the context of vehicular networks. It divides the detection approaches in three main categories: i) false beacon information detection, ii) false warning detection and iii) node trust evaluation. For each of them different misbehaviour detection methods are inspected from the literature. In addition, the technical study describes in detail four scenarios in which such systems may be useful and includes an appendix that addresses the difficulties of such systems when applied to CPMs.

In the latter, it is reported that for CPMs the previously listed misbehaviour detections for CAMs and DENMs may not be sufficient for providing accurate results and this is reported as a currently open issue. More recently, the ETSI TS 103 759 V2.1.1 [89] standardised the Misbehaviour Reporting Service, which allows the node to report the observation of misbehaviour events to the Misbehaviour Authority (MA). A Misbehaviour Report (MR) contains the message that activated the observed misbehaviour detectors, along with evidence for the MA to independently verify the misbehaviour event. The role of the MA is to pool reports from the nodes and use them to identify misbehaving nodes, whose identity is provided by the pseudonym certificates. The MA may then interact with the PKI to ask for the revocation of the permanent credentials of the misbehaving node. Notice that this process does not happen in real-time, and the revocation of a node may happen several days after first reports by other nodes.

7.4.6 SMTD Reference Scenario User Stories

[SMTD.US.1]: As a Traffic Control Centre, I want to be able to receive events (from the MECC) on the position of both equipped and non-equipped vehicles based on trustworthy data.

User Story Confirmation:

- In this user story the position of the equipped vehicle is identified leveraging the CAM messages, while the position of the non-equipped vehicle is provided by CPM messages.
- The CAM message is encoded at the vehicle side by the encoder module of the ITS (ETSI) facility, based on the information coming from the GNSS sensor having a predefined position accuracy level.
- The CPM message is encoded at the vehicle side by the encoder module present in the ITS facility, leveraging the data coming from the camera (distance and relative angles) along with the information coming from the GNSS sensor, the position of the non-equipped vehicle is calculated, based on the relative distance of the equipped vehicle.
- The CAM and CPM messages are sent to the TCH module to be signed. The now trusted CAM and CPM messages are sent to the MECC by the V2X OBU.
- The CAM and CPM messages after reaching on the MECC are being forwarded to i) the Misbehaviour Detector, in order to create the Misbehaviour Detector Reports, ii) to the TCH, in order to extract the trust evidence and iii) to the TAF that, leveraging the previous outputs, is able to define the trust level of the received messages.
- The Traffic Control Centre is using the trust level of the received messages, provided by the TAF, to understand the veracity of the received data (the accuracy of the data that encompasses the following three traits of data: integrity, precision and reliability).

User Story Workflow: *(see Figure 67)*

Sequence Diagram Description: In this user story we have the following flows:

- Vehicle detection by the vehicle sensor.
- Raw data from the sensor published on the CAN bus.
- Raw data read from the CAN bus by the V2X OBU.
- Position level accuracy, provided by the GNSS sensor to the V2X OBU
- CAM and CPMs are constructed by the CAM/CPM Encoder of the ITS facility module. They are also sent to the Trustworthiness Claims Handler (TCH) module on the OBU in order to have the trustworthiness claims.
- The AIV & TCH modules contains crypto primitives transmitted from the Cloud node towards the V2X OBU through a dedicated Kubernetes channel.
- The trusted CAMs & CPMs are sent to the MECC from the V2X OBU
- The CAMs & CPMs reaching on the MECC are forwarded to:
    - the Misbehaviour Detector in order to create the Misbehaviour Detector Reports.
    - the TCH in order to extract the trust evidence.
    - the TAF that, thanks to the previous outputs, is able to define the trust level of the received messages.

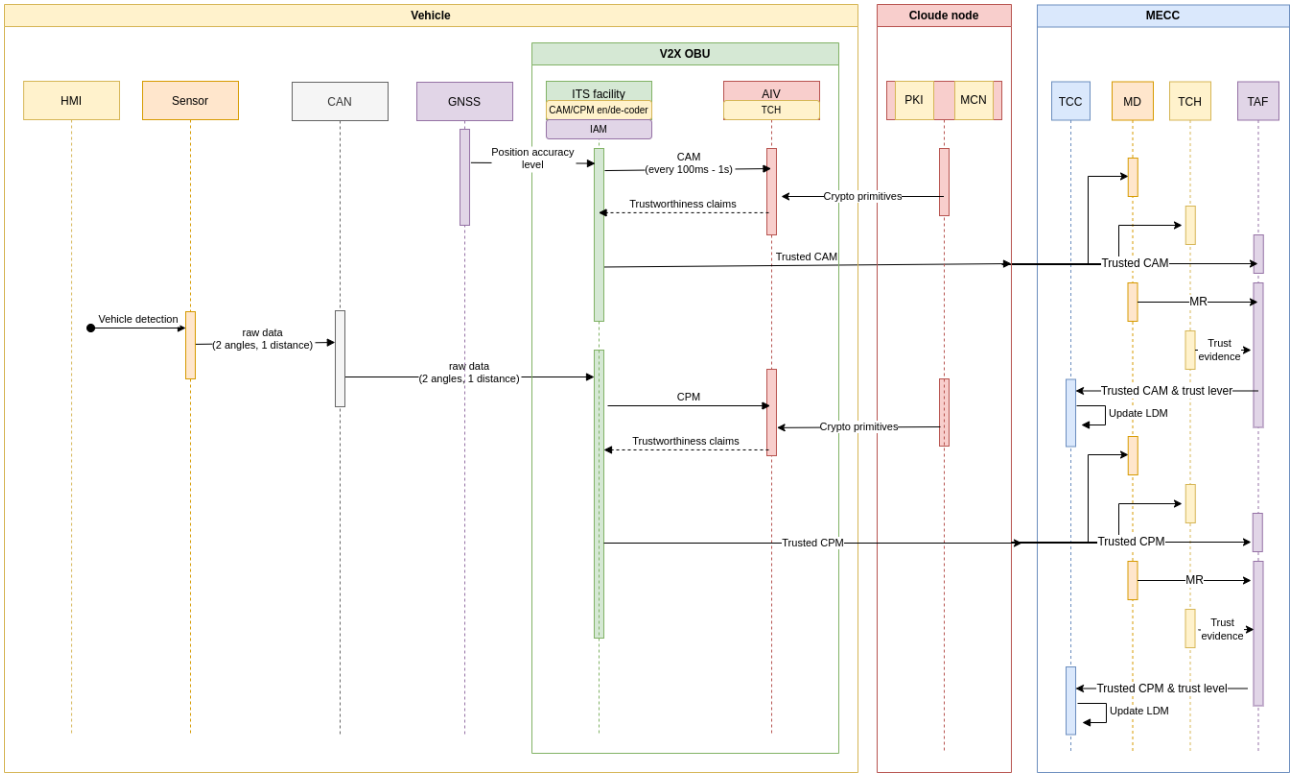

Figure 67 - Representation of the sequence diagram of user story [SMTD.US.1]

CONNECT KPIs:

The main KPI focus is on the latency introduced by the inter-generation times of the CAM and CPM messages that should be minimised in order to respect the actual standard, assuming that this trustworthiness mechanism doesn't introduce too much delay affecting the readiness of the vehicle perception.

The processing frequency needed for CAM and CPM secure construction to support a 10 Hz transmission rate KPI is needed in order to be compliant with the standard ETSI EN 302 637-2 V1.4.1 for CAM transmission and with the standard ETSI TS 103 324 V0.0.54 [100] for CPM transmission.

Table 21 - [SMTD.US.1] KPIs

KPIs	Description	Value
	Processing frequency for CAM and CPM secure construction considering the time needed for the computation of TCs (either calculated inside TEE or outside a TCB)	**>= 10 Hz** so as to be able to support ETSI requirements
	Processing frequency needed for CAM and CPM secure construction considering the time needed to perform the necessary crypto operations for ensuring integrity and confidentiality of CAM/CPM messages based on the use of PKI-based certificates	
	Vehicle to MECC latency < 1.5s in order to let the driver know in advance the existence of a traffic congestion ahead: Verification of attestation evidence (Vehicle)	**< 50 ms**
	Vehicle to MECC latency < 1.5s in order to let the driver know in advance the existence of a traffic congestion ahead: Construction of TCs (Vehicle)	**< 400 ms**
	Vehicle to MECC latency < 1.5s in order to let the driver know in advance the existence of a traffic congestion ahead: Verification of TCs (MECC)	**< 20 ms**

[SMTD.US.2]: As a Driver, I want to be notified in real-time of any traffic congestion and/or blocking points that may affect my journey. In particular, I want to be able to receive correct

and trustworthy traffic congestion events based on the LDMs calculated at the MECC Level.

User Story Confirmation:

- In this user story an alert is issued by the MECC due to a traffic congestion and/or blocking point on the road.
- The TCC encapsulates the alert inside a DENM message that is signed with the PKI-acquired keys that also this MEC-instantiated service is equipped with. The DENM message is sent to all vehicles in a given area (approaching the traffic congestion event).
- The receiving ITS facility inside the vehicle checks the veracity of the DENM message thanks to the crypto primitives inside the THC module. If the veracity is confirmed, the alert is sent to the vehicle HMI. In this way the driver of an equipped vehicle can be informed of the hazard ahead.

User Story Workflow: *(see Figure 68)*

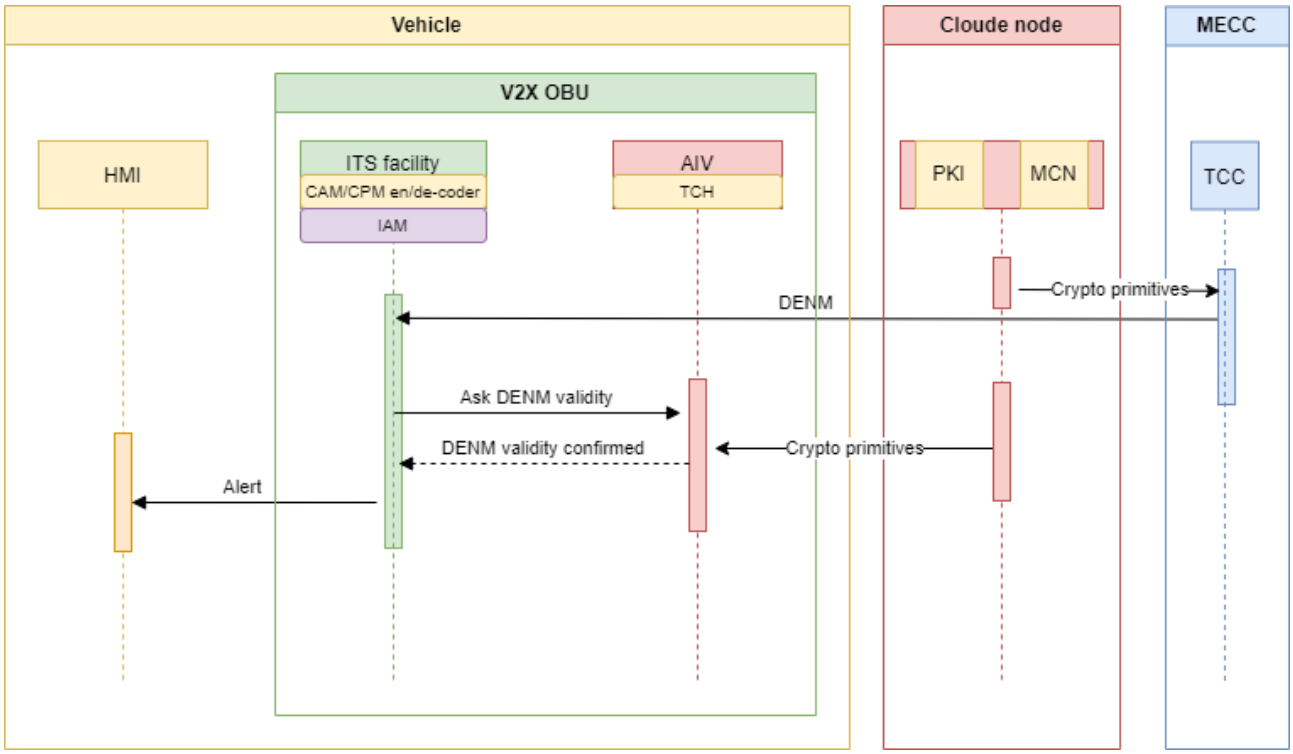

Figure 68 - Representation of the sequence diagram of user story [SMTD.US.2]

Sequence Diagram Description: In this user story we have the following flows:

- The TCC is using the crypto primitives loaded from the Cloud node to generate a trusted DENM.
- DENM sent by the TCC after a slow-moving traffic detected.
- Validity check request from the ITS facility module to theTCH module
- Validity confirmation received by TCH module on the V2X OBU witch crypto primitives was loaded from the Cloud note through a dedicate Kubernetes channel.
- Alert issued on the HMI of the vehicle for the driver.

CONNECT KPIs:

The KPI focus is to prove that the alert reached the driver with sufficient readiness to let the driver react to the highlighted dangerous situation.

- (Qualitative).
- MECC to vehicle HMI latency < 1.5sec in order to let the driver manage in time the alert.
    - DENM Construction Time (needed on the MECC)
    - Verification of trustworthiness level of DENM (on the Vehicle)

- This whole KPI will be evaluated with and without the TEE in order to assess the impact of TEE on message frequency.

Table 22 - [SMTD.US.2] KPIs

KPIs	Description	Value
	Alert visualisation on the HMI	**< 50 ms**
	MECC to vehicle HMI latency in order to allow for the driver's reaction time to a received alert. DENM Construction Time (needed on the MECC)	**<= 1.5sec**
	MECC to vehicle HMI latency in order to allow for the driver's reaction time to a received alert. Verification of trustworthiness level of DENM (on the Vehicle)	

[SMTD.US.3]: As a Vehicle I want to be able to establish a secure and authenticated (application-layer) link with the MECC Provider offering the SMTD Service.

User Story Confirmation:

- As part of the cellular connection and the data accommodated, there are already 'standard' security mechanisms to ensure a UE authorization and authentication[19]. In this user story we focus on the authenticity and integrity of the information accommodated by the CAM and CPM messages, when directed to a MECC application.
- Leveraging the TCH module that contains crypto primitives downloaded from the Cloud node on the V2X OBU through a dedicated Kubernetes channel, the ITS facility on the vehicle is able to enrich the CAM and CPM messages with the appropriate trustworthiness claims before sending them.

User Story Workflow: *(see Figure 69)*

Sequence Diagram Description: In this user story we have the following flows:

- Vehicle detection by the vehicle sensor
- Raw data from the sensor published on the CAN bus.
- Raw data read from the CAN bus by the V2X OBU
- Position level accuracy provided by the GNSS sensor to the V2X OBU
- CAMs & CPMs created by the ITS facility module sent to the TCH module to be provided with the trustworthiness claims.
- Trustworthiness claims provided by the TCH module whose primitives are loaded from the Cloud Node through a dedicated Kubernetes channel.
- The so trusted CAMs & CPMs are sent to the MECC.

[19] A user (UE) of 5G networks needs to successfully go through a primary authentication procedure whereby the UE is authorised to subsequently access additional network services. 3GPP has specified two relevant protocols i.e., 5G-AKA and EAP-AKA (https://www.3gpp.org/technologies/akma). A secondary authentication phase enables the UE to access the user plane data.

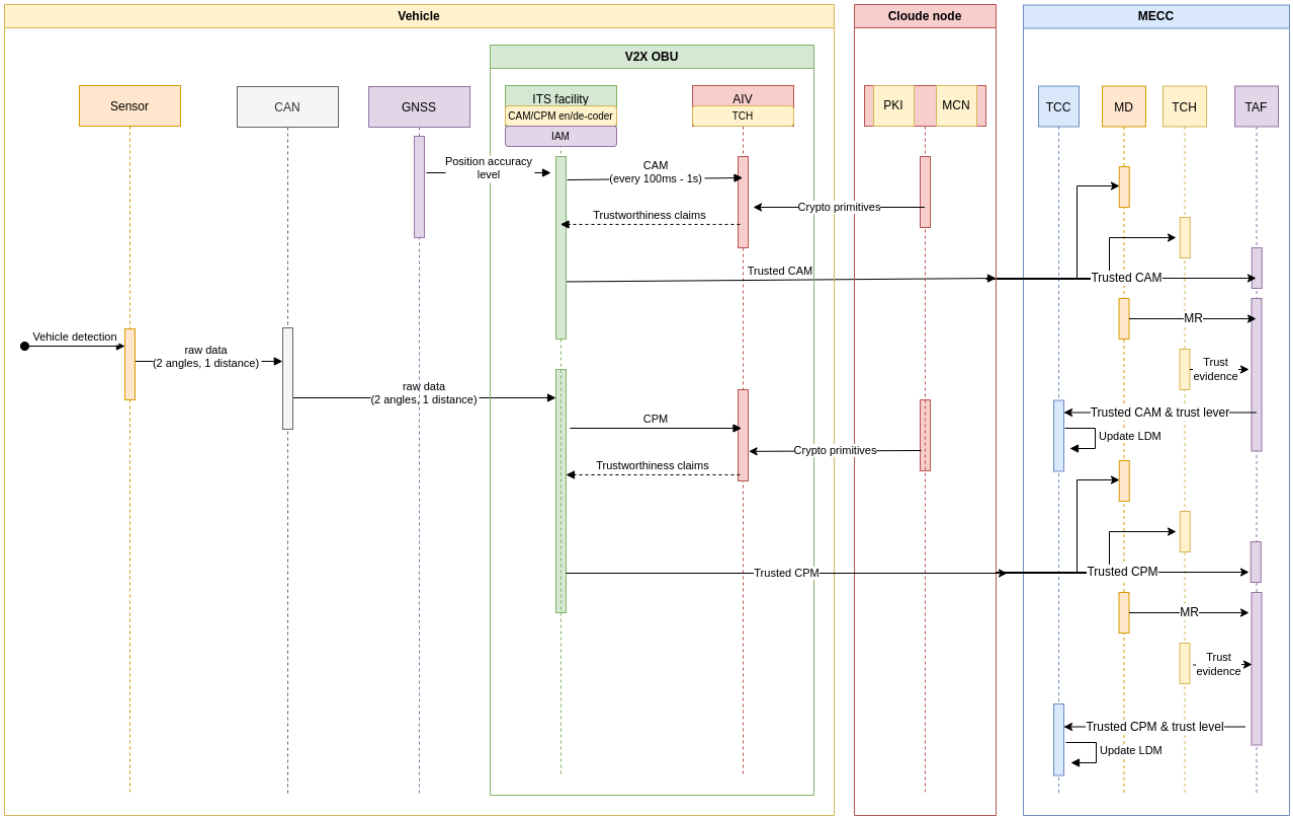

Figure 69 - Representation of the sequence diagram of user story [SMTD.US.3]

CONNECT KPIs:

The first type of KPI focuses on the time needed to enrich the CAM and CPM messages with the appropriate trustworthiness claims that should be minimised in order to respect the actual standard in order to not introduce too much delay affecting the readiness of the vehicle perception. The second type of KPI (the qualitative one) focuses on the cyber security of the trustworthiness mechanism that must be proven.

Table 23 - [SMTD.US.3] KPIs

KPIs	Description	Value
	Processing frequency needed for CAM and CPM secure construction to support a 10 Hz transmission rate. Time needed for TCs creation inside and outside the TEE	**>= 10 Hz** so as to be able to support ETSI requirements
	Processing frequency needed for CAM and CPM secure construction to support a 10 Hz transmission rate. Time needed for ensuring integrity and confidentiality of CAM/CPM messages based on the use of PKI-based certificates	
	A compromised OBU should not be able to inject fake (but cryptographically valid) CAM/CPM messages into the network - (Confirmation) Identity crypto primitives to be part of CONNECT's Trusted Computing Base	**TRUE**

[SMTD.US.4] As a Traffic Control Center, I want to be able to enhance my LDM with trustworthiness levels of all vehicles based on plausibility checks performed by the MD service.

User Story Confirmation:

- In this user story the TAF is able to define the trust level (TL) of the received messages leveraging the support of the Misbehaviour Detector that generates the Misbehaviour Detector Reports and the TCH that is able to extract the trust evidence.

- If the trustworthiness level of an ITS message is below a threshold, the message is considered misbehaviour and it is discarded, otherwise the message can be passed to the Traffic Control Center (TCC).

User Story Workflow: (see *Figure 70)*

Sequence Diagram Description: In this user story we have the following flows:

- The CAMs & CPMs reaching on the MECC are forwarded to:
  - the Misbehaviour Detector in order to create the Misbehaviour Detector reports.
  - the TCH in order to extract the trust evidence.
  - the TAF that, thanks to the previous output, is able to define the trust level of the received messages.
- The so trusted CAM & CPMs are used to update the LDM.

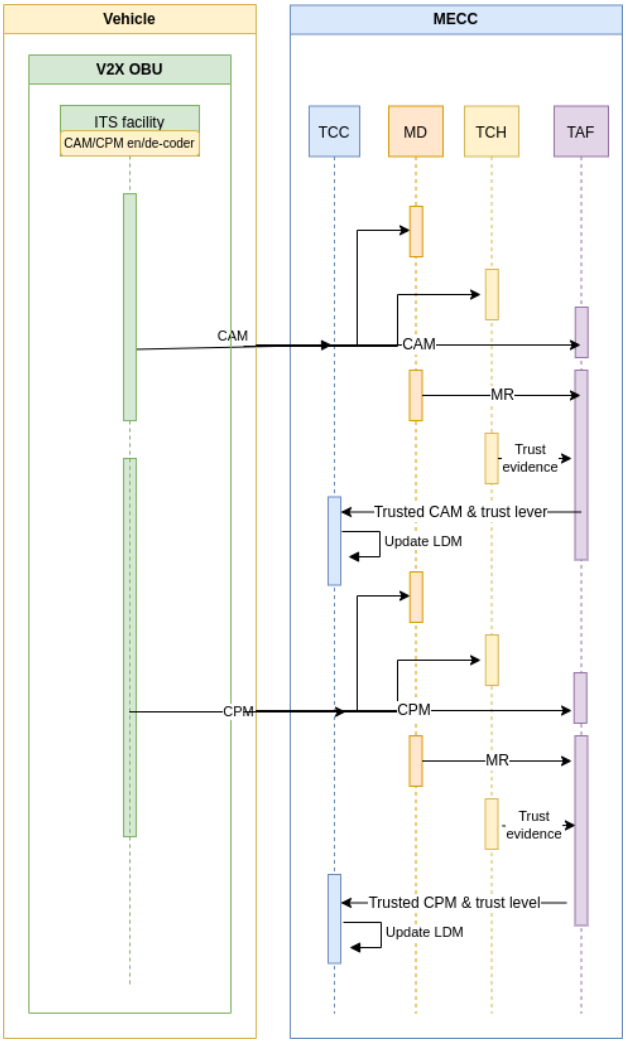

Figure 70 - Representation of the sequence diagram of user story [SMTD.US.4]

CONNECT KPIs:

The KPI focus is similar to IMA use cases (Section 7.2.6) regarding the Service Providers at the MEC that takes in consideration the quality of the generated TL and the time to generate it.

It should be clarified that the processing complexity until LDM update, refers to the complexity of the operations between message reception and update of the observations' trustworthiness levels in the LDM.

Table 24 - [SMTD.US.4] KPIs

KPIs	Description	Value

	The TLs of the active V2X-vehicle evolves correctly: it increases as more evidence of correct behaviour is gathered; it degrades as malicious behaviour is injected in the scenario.	**TRUE**
	Processing complexity until LDM update: If V-TCs present, verification of the TCs and attributes extraction	The focus is on identifying the overhead pattern imposed based on the number of identified objects. Once this is calculated, detailed benchmarking will follow considering different vehicle neighbourhood densities.
	Processing complexity until LDM update: If V-TCs present, emitter V2X-node's TL assessment by the TAF	
	Processing complexity until LDM update: Local MD on the observations. Notice that the complexity of this step may be influenced by factors such as traffic density (in denser scenarios, characterised by richer V2X information, more misbehaviour checks become available).	
	Processing complexity until LDM update: Observation's TL assessment by the TAF	

[SMTD.US.5] As a vehicle, I want to be able to offload resource-demanding tasks to the MECC.

User Story Confirmation:

- In this user story a resource-demanding task requested from a sensor on the vehicle is offloaded towards the MECC. It begins with an initialization and security check phase between a dedicated Task Offloading Agent on the vehicle and a Digital Twin on the MECC. The considered offloading process may include: i) trust calculations per-se; there may be a need (see Figure 71) to carry-out demanding trust calculations to make sure that the sensory data are trustworthy, ii) a sensor input (that relates to CCAM function) and its processing requires more resources than the currently available ones in the vehicle. In this case, the Task Offloading Agent requests the offloading procedure to the Digital Twin that elaborates the request, and it sends back the reply.

User Story Workflow: *(see Figure 71)*

Sequence Diagram Description: In this user story we have the following flows:

- Initialization and security checks between the Task Offloading Agent on the vehicle and the Digital Twin module on the MECC
- High power consuming input from a sensor on the vehicle
- Offloading task request from the Task Offloading Agent on the vehicle towards the Digital Twin module on the MECC
- Reply from the Digital Twin module on the MECC.

CONNECT KPIs: - Processing time of Task Offloading Agent in Vehicle and DT on MEC (prepare task to offload task calculation)

Table 25 - [SMTD.US.5] KPIs

KPIs	Description	Value
	Offloading trust calculations that result need to be given as input to the TAF so as to not affect safety	**< 1 sec** excluding network latency
	Offloading application-related tasks. Those tasks are shaped by the processing time needed (by the CONNECT facility layer) to orchestrate the offloading process along with the time required to have the result back (to the vehicle). A relevant KPI may be expressed as the need to offload a given task under a certain time limit 'T', measured on the application level. This limit 'T' is subject to the involved	Our experimentation will identify the time interval within the task offloading (for a set of tasks) is feasible and thus, will provide insights for each potential CCAM function.

	CCAM function (e.g., a platooning driving function may pose more stringent requirements than an infotainment one) and may attain a broad set of values.	

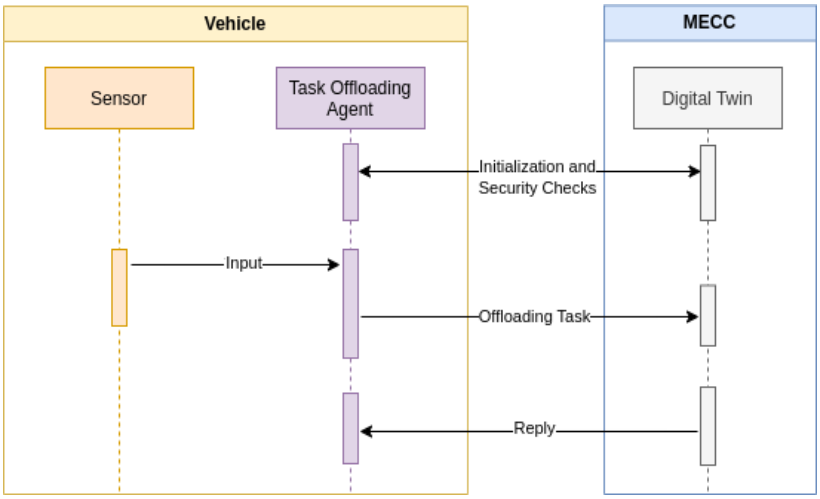

Figure 71 - Representation of the sequence diagram of user story [SMTD.US.5]

[SMTD.US.6] As a Traffic Control Centre, I want to receive from the vehicles the CAM and CPM messages generated in the last sixty seconds (CCHP - CAM and CPM History Package) in a trustworthy, conflict-free and resource-efficient manner.

User Story Confirmation:

- In this user story the motion info of the ego vehicle and of the sensed vehicles are stored on board. In particular, this process will store the last minute of motion info, while the data related to info older than one minute will be deleted.
- The above stored motion info is sent to the MECC upon a hazardous event (i.e., a car accident) or by an explicit request of the TCC (when the TCC wants to rebuild the accident scene adding information from the nearby vehicles not involved in the crash, the “witness” vehicles). In such circumstances, an ad-hoc CAM and CPM History Package (CCHP) is created, and it is transmitted to the MECC server following the same trust mechanism described in the previous use cases for the single CAM and CPM messages. The CCHP contains all the motion info of the past minute of the ego vehicle and of the sensed vehicles. The “veracity” (integrity, precision, and reliability of data) of the CCHP must be guaranteed.
- The aim of this use case is to strongly certify an historical set of data that could be used in a legal debate with the insurance company and the vehicle owners involved in the analysed situation (data veracity) or for a fine or a service payment of the Municipality (data monetization).

User Story Workflow: *(see Figure 72)*

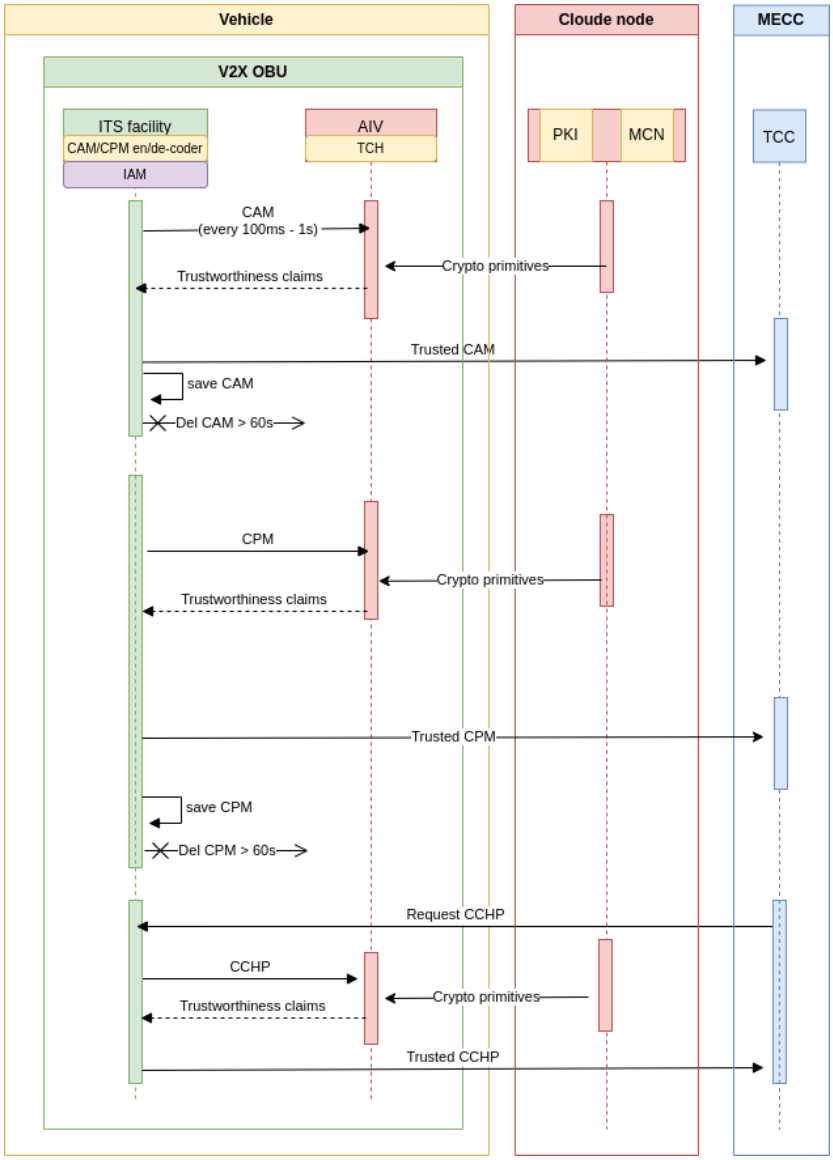

Figure 72 - Representation of the sequence diagram of user story [SMTD.US.6]

Sequence Diagram Description: In this user story we have the following flows:

- CAM and CPMs are constructed by the CAM/CPM Encoder of the ITS facility module. They are also sent to the TCH module in order to have the trustworthiness claims.
- The AIV & TCH modules contains crypto primitives downloaded from the Cloud node on the V2X OBU through a dedicated Kubernetes channel.
- The so trusted CAMs & CPMs are sent to the MECC from the V2X OBU
- The sent CAMs & CPMs are saved on the V2X OBU for a shifting time window of 60s.
- Following the request of CAM and CPM History Package (CCHP) from the TCC on the MECC or a triggering event in the vehicle, the CCHP package is created by the ITS facility module and sent to the TCH module in order to have the trustworthiness claims.
- The so trusted CCHP package is sent to the MECC from the V2X OBU

CONNECT KPIs:

The KPI focus is on the veracity of the provided CCHP package: the data stored must be complete, not corrupted and trustworthy. This last point, the trustworthy, should be guaranteed by the results of the MD, TCH and TAF that can provide the trust level on which the TCC can make a decision on the veracity of the data.

Table 26 - [SMTD.US.6] KPIs

KPIs	Description	Value

	Duration that the information is saved **(i.e., this is not to overload the OBU device with too much useless information)**	**60sec**
	Completeness of information, all needed signals are present	**TRUE**
	Integrity of data, data are not corrupted	**TRUE**
	Signature is valid	**TRUE**
	CAM/CPM Messages associated with local Trust Level to be sent to the MECC for making a decision based on trustworthy data - based on the TAF running on the MECC (to also cover the scenario when vehicles sending contradicting evidence)	Associated TCs, CONNECT security controls and all C-ITS functionalities **<= 90%** of storage and computation resources

8 CONNECT Technical Requirements

In this chapter, we describe the technical requirements of CONNECT, which have been clustered in **mandatory and desirable ones**. This classification differentiates between those technical requirements that are needed for the successful operation of the CONNECT trust assessment functionalities (and are also needed for the successful instantiation of the demonstrators, and the possible requirements of Connected Cars operational assurance as a standard in general (and as a topic of research) and those that are good to have but do not directly impact the CONNET mode of operation. *Thus, it is the mandatory requirements that will form the basis of the core technical requirements of this project.*

Towards this direction, CONNECT has separated **functional** and **non-functional** requirements. These requirements are further crucial for the long-term success and sustainability of Day 3+ operations in the V2X landscape. Within the functional domain, three crucial subcategories are investigated: i) the Trust Assessment Framework, which defines the mechanisms for assessing and establishing trustworthiness among V2X entities; ii) MEC (Multi-access Edge Computing), which focuses on the seamless integration of edge computing resources to improve security and latency-sensitive V2X applications; and iii) the Trust Computing Base, which comprises the foundational hardware and software elements responsible for supporting the execution of trust-related operations through the provision (in a verifiable manner) of the necessary trustworthiness evidence.

In addition, we evaluate non-functional requirements, with a focus on privacy considerations, as protecting user data and preserving anonymity are of the utmost importance in V2X environments. This exhaustive analysis lays the groundwork for the creation of a secure and trustworthy V2X ecosystem, supported by the CONNECT platform.

8.1 CONNECT Technical Requirements

In the enumerated listings that follow, we make a concrete mapping between “CONNECT’s Value Propositions” and the functional requirements that they correlate to each other. In parallel, a brief description of each functional requirement is provided.

It should be noted that for the performance KPIs, as it pertains to the time overhead introduced by the trust assessment and trusted computing enablers of CONNECT, the same approach is followed across all requirements. This approach revolves around the performance evaluation and the resource utilisation of all CONNECT enablers that participate in the system, without the consideration of the overhead posed by the underlying Root of Trust (RoT). This is explained by the fact that CONNECT remains agnostic to the RoT employed, while the overhead is a result specific to the selected RoT; hence, CONNECT opts to evaluate the enablers without the constraints of this environment.

To cover the overhead posed by the RoT, the percentage that the TEE execution introduces in the instantiation of the security enablers, will be estimated separately and for a specific instance (i.e., Gramine). Based on this evaluation CONNECT will be able to provide recommendations regarding the feasibility of the instantiation of real-time trust assessment enablers, in the context of TEE.

8.1.1 Trust Assessment Requirements

This subsection provides the requirements that the Trust Assessment Framework (TAF) must satisfy. The TAF is one of the core products of the work to be done as part of the CONNECT project. The TAF's task will be to assess the trustworthiness of nodes and data involved in safety critical CCAM applications from the perspective of many CCAM actors. TAF could be instantiated on different components inside a single vehicle, a Multi-access Edge Computer (MEC), the cloud, and any CCAM actor running a safety-critical application that uses data from other nodes as input. The following requirements need to be satisfied by any TAF, irrespective of whether it is running on a vehicle, a MEC, a cloud, or any other actor.

The requirements we have chosen to highlight the most important characteristics that the TAF ought to have to satisfy the needs of our use cases and the needs of any safety-critical applications running in CCAM in general. Moreover, we have structured these requirements in a way that allows us to measure how successful the TAF is in satisfying any one of them.

It should be noted that the CONNECTs TAF extends beyond purely technical aspects, considering the establishment of trust relationships between entities and addressing user trust in a system's reliability, transparency, and compliance with privacy regulations.

ID	FR.TR.1 (Mandatory)
Title	**Generalizability**
Actors/Components Involved	Trust Assessment Framework
Description	**Background:** The assessment of trust in the automotive sector is a particularly difficult task due to the complexity and the heterogeneity of the V2X landscape. Literature is rather poor regarding the topic of assessing trust in CCAM scenarios that comprise multiple stakeholders and components both in terms of software and hardware. Apart from the number and heterogeneity of the involved parties, it shall be noted that changes in the trustworthiness of the aforementioned entities and systems over time, that are taking place dynamically; variables such as software updates should be taken under account. For example, in-vehicle and V2X networks, consider different software updates over time, different types of sensors and sensor data, etc. In addition, the trust assessment should operate in a zero-trust environment, where trust is never assumed and must be continually verified. This further adds to the complexity of the overall framework. It is imperative that the Trust Assessment Framework (TAF) is generic enough to be applicable in the multitude of the different CCAM scenarios that need to operate under the zero-trust assumption and also capture the changes in the trustworthiness level that might occur over time. A generalizable TAF should be widely applicable to different use cases and would reduce or eliminate the customization effort for each use-case scenario. Ideally, it could even be updated and extended for use in completely novel use-cases or scenarios. **Description:** CONNECT, envisages a TAF capable of assessing the level of trust and trustworthiness *for any given scenario; thus, any arbitrary trust model that includes different types of trust relationships, created among heterogeneous trust objects for different properties.* For example, such a trust relationship can be between two entities inside a single vehicle (e.g., between two ECUs), between two vehicles themselves, and also between a vehicle and another entity in the V2X network, e.g., a Multi-access Edge Computing (MEC) Service Provider. Both node-centric and data-centric relationships may exist. While the first (i.e., node-centric) refers to the trust relationship between two nodes, the second (i.e., data-centric) refers to the trust relationship between a node and data (e.g., between a vehicle and a CAM). In addition, as part of the CONNECT trust model, both referral and direct (functional) trust relationships are considered. Towards this direction, *the TAF should accommodate assessing trust based both on direct trust relationships but also using referral relationships that enable leveraging trust assessments (or opinions) that have already been made by other entities*. The first (i.e., the direct trust relationships) can be both node- and data-centric, whereas the latter (i.e., the referral trust relationships) are always node-centric. **Remarks:** (1) Please note that we refer to generalizability in relation to the architecture and mode of operation of the TAF. Namely, the TAF should be able to be adopted even in the context of new (or never before encountered)

	scenarios. However, the appropriate trust models need to be defined for such scenarios, and generalizability should not be confused with the need to have defined trust models for all scenarios to be encountered. (2) While assessing data-centric relationships, there is always an inherent dependency to the trust level of the node producing this data that also needs to be accounted for.	
Connected To Other Requirements	FR.TR.6	
KPIs	**Description**	**Value**
	Number of CCAM use cases	**=3 heterogeneous CCAM use cases,** capturing in-vehicle, vehicle to vehicle and vehicle to MEC (and vice versa) scenarios. Thus, all possible trust relationships in these use cases will be instantiated and evaluated, meaning referral (node-centric) and direct (node-centric and data-centric) trust relationships.
ID	**FR.TR.2 (Mandatory)**	
Title	**Run-time Performance**	
Actors/Components Involved	Trust Assessment Framework (TAF)	
Description	**Background:** Due to the dynamic aspects of the systems, as already discussed in FR.TR.1, a trust assessment framework that dynamically assesses and calculates opinions, shall operate in real-time under strict time requirements. This is further emphasised and has been considered as a critical factor in CCAM environments, where numerous (safety-critical) applications need to operate within strict time constraints[20]. Namely, here we focus on real-time applications, like the Intersection Movement Assist (IMA), with the highest requirements on safety and security, especially because any failures to deliver relevant data under very strict time constraints can have a direct impact on the safety profile of the CCAM actors. **Description:** The CONNECT's Trust Assessment Framework (TAF) is intended to be built for complex systems from the CCAM domain that are **time-** and **safety-critical**. *As a result, the TAF should be able to assess the trustworthiness of an entity within strict time requirements*. This is done by each component of the TAF being accountable for doing efficient, optimised, and real-time calculations, with adding minimal to no overhead when all the components are integrated together as part of the TAF. Please note that these requirements related to time- and safety-criticality will vary between applications. For example, in different CCAM applications, such as Intersection Movement Assist (IMA), decisions have to be taken fast. When IMA needs to make decisions that are made based on the output of the TAF, the TAF has to be fast enough not to cause any delays in the decision-making process, since such delays could lead to safety issues. **Remarks:** In relation to the time- and safety-critical requirements there might be some applications (e.g., the IMA) where decisions need to be made fast, while there might be other applications where the highest level of certainty in the trust assessment is required. The latter might require a longer time for the assessment to be completed. As a result, the TAF should be able to produce a	

[20] ETSI TR 103 299, "Intelligent Transport Systems: Cooperative Adaptive Cruise Control (CAAC); Pre-standardization Study", 2019, [Available Online:] https://www.etsi.org/deliver/etsi_tr/103200_103299/103299/02.01.01_60/tr_103299v020101p.pdf

	real-time trust opinion within the time frame allowed which, in turn, will dictate the trust sources that the TAF can use. If, for instance, there is not enough time to get "fresh" evidence, then the TAF can use the previous ones to create the trust opinion but with a higher level of uncertainty.	
Connected To Other Requirements	FR.TR.3	
KPIs	**Description**	**Value**
	Latency of standalone trustworthiness level assessment execution by the TAF	**<=100 ms** delay when the TAF is instantiated and executed as part of the application software stack in the target system (outside the CONNECT TEE) **<=200 ms** delay when the TAF is instantiated and executed within the CONNECT TEE. These calculations include the exchange of information between the application triggering the TAF (either running outside or inside the CONNECT TEE), in essence the application that sends the trust assessment request, for which the TAF replies back with the trust level. Note that this does not include the time needed for the collection, processing and communication from the entities (in-vehicle or other vehicle), while we are further excluding any network latency caused from the reception of trust sources from other neighbouring vehicles or the MEC. The focus is only on the timing requirements of the TAF operation and calculation of the trust level
ID	**FR.TR.3 (Mandatory)**	
Title	**Scalability**	
Actors/Components Involved	Trust Assessment Framework	
Description	**Background:** As discussed in FR.TR.1, the V2X landscape is composed of multiple entities. For example, the CCAM domain may include Intersection Movement Assist comprising vehicles, Multi-access Edge Computer (MEC), and other entities. Hence, these systems are dynamic in the sense that the actors which comprise them change over time. Moreover, the number of actors in any CCAM system also changes and can sometimes reach a large scale. There might be scenarios with a very small number of vehicles at an intersection managed by, for example, a MEC, and scenarios with an extremely high number, such as in rush hour. **Description:** In CONNECT, the TAF should be *scalable in order to assess the trustworthiness levels of all involved vehicular nodes and the data exchanged between the nodes*, within strict time requirements, even if the number of nodes is very high. The TAF should also be able to assess trustworthiness of every new node which enters the system, even if there are many other nodes in the system already. This would imply that the TAF (either instantiated in the vehicle or the backend infrastructure) needs to be able to quickly update, analyse, and break down large trust models representing all vehicular nodes and data needed for a certain CCAM application. Irrespective of the number of nodes in the model, the TAF needs to be able to calculate the necessary Actual Trustworthiness Levels (ATLs) in a short period of time. CONNECT offers this	

	scalability, leveraging the Federated TAF, as well as the Digital Twin, overcoming the barriers of traditional centralised infrastructures.	
Connected To Other Requirements	FR.TR.2	
KPIs	**Description**	**Value**
	Number of nodes supported	For small scale environment (i.e., IMA) testing to be done: **<=10 nodes** (Single TAF) For large scale environment (i.e., IMA) testing to be done: **<=50 nodes** either vehicle or MEC-running (Federated TAF)
	Timeframe to complete trust assessment	**<=100 ms** delay when the TAF is instantiated and executed as part of the application software stack in the target system (outside the CONNECT TEE) **<=200 ms** delay when the TAF is instantiated and executed within the CONNECT TEE. These calculations include the exchange of information between the application triggering the TAF (either running outside or inside the CONNECT TEE), in essence the application that sends the trust assessment request, for which the TAF replies back with the trust level. Note that this does not include the time needed for the collection, processing and communication from the entities (in-vehicle or other vehicle), while we are further excluding any network latency caused from the reception of trust sources from other neighbouring vehicles or the MEC. The focus is only on the timing requirements of the TAF operation and calculation of the trust level
ID	**FR.TR.4 (Mandatory)**	
Title	**Correctness**	
Actors/Components Involved	Trust Assessment Framework (TAF)	
Description	**Background:** It becomes apparent that qualitative, informal trustworthiness assessments are insufficient for making informed decisions on whether an entity can be trusted. Instead, measurable, and quantifiable metrics need to be defined. To this end, such metrics need to be defined to enable the Trust Assessment Framework (TAF) to determine whether an entity can indeed be trusted. The decisions rendered by the TAF must align with the actual trustworthiness of the entity in question. For example, if an entity is malicious and therefore deliberately provides false position data to another entity, the TAF of the receiving entity should decide that the position is not trustworthy. Hence, to make the decision on the trustworthiness of an entity, the level of trustworthiness of a certain trustor towards this entity (at a point in time) is necessary, where this entity is referred to as a trustee in this context. The entity for which the level of trustworthiness is determined is specified in a proposition. This level of trustworthiness is a numeric value referred to as the Actual Trustworthiness Level (ATL). To derive the ATL, trust sources that provide	

<table>
<tr><td></td><td colspan="2">evidence for a trust object are required. As part of this requirement, it is assumed that these trust sources are not compromised and provide correct evidence.

However, calculating the ATL of an entity is not sufficient to decide whether it can be trusted. Therefore, in addition to the ATL, there needs to be a Required Trustworthiness Level (RTL) that reflects the level of trustworthiness of an entity required in order to be characterised as trustworthy. By comparing the ATL and RTL, the TAF can decide whether to trust the corresponding entity.

Description: CONNECT's TAF must be able to <u>produce a correct ATL for a proposition</u>. For this purpose, the trust model is used, which contains all trust relationships relevant for calculating the ATL. For each of these trust relationships, the TAF takes trust sources into account, based on which it calculates a trust opinion for the trust relationship. Based on the individual trust opinions, the ATL for the proposition is calculated. When the outputs of trust sources change, such as when an Intrusion Detection System detects malicious activities within the trustee, this should be reflected directly in the ATL.

In addition to the ATL, the RTL is also necessary in the TAF. The RTL (further discussed in chapters 2.2.2 and 3.2.3) is determined based on a separate mechanism outside the TAF and is then provided as an input to the TAF so that a decision can be made about the trustworthiness of the proposition.

The ATL and the RTL must be correct because these two parameters directly influence the decision about trustworthiness of the proposition. Thus, if either the ATL or the RTL is incorrect, the TAF could decide that an entity is trustworthy, although it is not, or vice versa.

Remarks: The ATL and the RTL are dynamic in the sense that they can change during runtime. For example, new attacks identified for a particular system may affect the ATL and RTL, requiring them to be adjusted in order for the TAF to provide correct outputs.</td></tr>
<tr><td>Connected To Other Requirements</td><td colspan="2">FR.TR.5</td></tr>
<tr><td rowspan="4">KPIs</td><td>Description</td><td>Value</td></tr>
<tr><td colspan="2">To evaluate whether the TAF provides the correct result that an entity (node and data item) is trustworthy or not, a scenario-based evaluation will be conducted for all envisioned use cases.</td></tr>
<tr><td>TAF result correctness for 1st scenario:

All entities are trustworthy (i.e., not compromised by an attacker)</td><td>Since in this scenario, all entities are trustworthy, the output of the TAF for all propositions should be that the corresponding entities are also trustworthy.

For example, if a trustworthy vehicle sends its non-compromised position to a MEC, the TAF of the MEC should decide that the position of the vehicle is trustworthy.</td></tr>
<tr><td>TAF result correctness for 2nd scenario:

One or several entities are not trustworthy because they have been compromised by an attacker.</td><td>In this case, the results of the TAF for the propositions containing these entities should be that the corresponding entities are untrustworthy.

For example, if a vehicle with a compromised GNSS sensor sends an incorrect position of its location to a MEC, the TAF of the MEC should decide that the position of the vehicle is not trustworthy. In this way, it is possible to assess whether the output of the TAF is correct in both scenarios, and thus whether the TAF is working as intended.</td></tr>
</table>

		We expect that in **>= 70 % of successful attacks**, which represent the existence of a compromised entity in the CCAM system, propositions regarding the trust level of this compromised entity **should fail the ATL>RTL** test. This would result in the TAF deciding that the corresponding proposition is **not trustworthy**, and thus allowing a reaction by the CCAM system.
ID	**FR.TR.5 (Mandatory)**	
Title	**Robustness and Resilience**	
Actors/Components Involved	Trust Assessment Framework (TAF)	
Description	**Background:** All things considered, the trust assessment is gradually becoming a very, if not one of the most, critical elements for ensuring security and safety within the automotive sector. Its goal is to help distinguish between trustworthy and untrustworthy entities. Nonetheless, as we delve deeper into the realm of trust assessment, we must also diligently evaluate the resilience and robustness of this very process. First, as part of our TAF there are different components that enable this complex trust assessment process (i.e., the Trust Model Manager -TMM, the trust model, the Trust Sources Manager - TSM, the Trustworthiness Level Expression Engine - TLEE, etc). Second, the TAF runs in the same host environments as the other CCAM applications - either on the MEC or the vehicle itself. Thus, they can also be the target of an attack for disrupting the normal operation. Third, although TAF is considered as part of the Trusted Computing Base (TCB) of an entity, and it is safeguarded by adequate mechanisms enabled by the underlying Root of Trust (i.e., it is protected by the security mechanisms of the TEE), it is still susceptible to attacks that could potentially affect its normal operation. Finally, different instances of the TAF are running in different nodes that do a decentralised and distributed trust assessment (also referred to as TAF federation in D3.1). All the above-mentioned aspects form various attack vectors, and by leveraging various techniques, attackers can perform several attacks on the TAF. For example, in an “on-off attack”, when the trust values of malicious nodes performing the attack are significantly reduced, attackers can perform good behaviours to increase their trust values over a period of time. And when their trust values reach a certain level, they again begin to execute malicious behaviours. **Description:** The TAF in CONNECT should include mechanisms to increase resilience against possible attacks on its operations. Undoubtedly, there are many ways to attack the TAF itself, resulting in the TAF providing incorrect or no output. Such attacks could affect each component of the TAF, and they can change the trust relationships, including the trust model, directly affecting the robustness of the trust model, the trust sources considered, or the trust opinions between two entities. Other attacks could aim at the federation of TAFs for different entities and all the aspects that the federation includes, e.g., communication, sharing the trust model or parts thereof, or sharing trust opinions among different TAF instantiations in vehicles or MEC. Therefore, the TAF should include mechanisms to increase resilience against possible attacks on its operations.	
Connected To Other Requirements	FR.OC.2, FR.OC.3, FR.OC.4	

KPIs	**Description**	
	TAF resilience against simulated attacks trying to alter the operation of the internal building blocks	
ID	**FR.TR.6 (Mandatory)**	
Title	**Flexibility of Trust Sources**	
Actors/Components Involved	Trust Assessment Framework (TAF)	
Description	**Background:** As mentioned in FR.TR.1**,** the Zero-Trust principle is one of the fundamental concepts for the TAF. Chapter 2 of the present deliverable sheds light to the Zero Trust paradigm, explaining that in the Zero-Trust principle all entities (a node or a data item) are considered to be possibly untrustworthy at the beginning. Thus, no initial trust between the entities shall be assumed, but the trust between the entities shall be continuously evaluated based on evidence. This evidence is provided by heterogeneous trust sources. **Description:** CONNECT envisages a landscape that incorporates a broad range of trust sources into the TAF; hence, depending on the use case, different entities are involved for which the trustworthiness has to be assessed. Accordingly, depending on the entity, different types of evidence can be provided. For example, one vehicle might contain an Intrusion Detection System (IDS) that can be used as a trust source, while another vehicle might contain an automotive network firewall instead. Therefore, the TAF should be flexible about which trust sources are used for the trust assessment. In order to align with the zero-trust notion, *the evidence collected must be verifiable* so that it can be verified that the information provided in the trust sources are indeed correct (produced and signed by a valid Root-of-Trust), and not just claimed by a malicious node, or modified by an attacker during transmission. All security and safety mechanisms can be used as trust sources. For example, trust sources can be hardware components, such as a TEE or an HSM. In addition, software components can also serve as trust sources, such as misbehaviour detection or IDSs. Trust sources provide either positive or negative evidence for an entity w.r.t. a specific property. For example, trust sources could provide evidence that the integrity-property of data provided by a node has not been compromised. The trust sources are provided as input to the TAF, based on which the TAF can determine the trustworthiness of an entity. **Remarks:** The consideration of multiple trust sources should not violate vehicle privacy in the sense that a vehicle could be identified based on the trust sources used as a quasi-identifier.	
Connected To Other Requirements	FR.TR.1, FR.PR.1, FR.PR.3	
KPIs	**Description**	**Value**
	Trust Sources	**>=3 different trust sources are included in the use cases**. These trust sources are collected so that they are verifiable through trustworthiness claims.
ID	**FR.TR.7 (old TR.12) (Optional)**	
Title	**Service Management over multiple domains**	

Actors/Components Involved	Trust Assessment Framework (TAF)
Description	**Background:** In the context of edge computing, each MEC-enabled service provider may operate within its own trust domain, which represents a specific level of security and trust, based on the type of security controls and capabilities (i.e., secure boot), that are deployed into the infrastructure that hosts the MEC, operated by a specific network operator. These security capabilities ensure that one MEC domain achieves the desired level of security and trustworthiness and that its resources are protected under a specific set of characteristics, in terms of security. Examples of such levels are defined by ETSI [157] and depict the level of trust for the internal building blocks, starting from secure boot moving all the way to the attestation of the network plane. This service management over multiple domains with various trust levels is a critical enabler for the secure interconnection of services. Imagine for example a platooning service, where multiple vehicles (i.e., trucks) travel in convoys to enhance fuel efficiency and reduce air resistance. The MEC collects and processes real-time data about speed, distance, and braking intentions, creating a 'Platoon Coordination Model' (PCM). This model ensures safe and synchronised manoeuvres, allowing each vehicle to optimise its performance. This communication between platoon members enhances road safety, traffic flow, and fuel efficiency. The MEC enables such communication, data exchange, and decision-making among the vehicles involved in the platoon. **Description:** CONNECTs vision for day 3 of CCAM automated operations includes a single service which can be deployed over multiple domains (i.e., a traffic control management system). This service should be able to expand over multiple domains in different regions of a city. These domains may be operated by different network operators. In this particular case, exchange of information should be possible between services that are instantiated into different domains, hence exhibit a different trust level. CONNECT should provide a framework that allows the assessment of the trustworthiness (i.e., secure and authenticated) communication channel between domains exhibiting different trust levels. This improves the overall system's portability and robustness, while maintaining the necessary security measures along with support for legacy systems. It enables dynamic and transparent communication between MECs that operate in different trust domains, with different security controls. This functionality further ensures the system’s interoperability among MEC domains that may have access to different types of security controls. Hence CONNECT should be able to support the provision of evidence, similarly to the trustworthy evidence as provided by the vehicles, to evaluate the level of trust per domain so that: Towards this direction CONNECT shall provide the evidence in a verifiable means so that: 1. a STAKEHOLDER may establish a secure connection to a SERVICE hosted in a MEC provided by a network operator. The stakeholder is able to connect to an edge service and is able to verify the end-to-end authenticity (=identity of service), confidentiality, and integrity of the connection, and the STAKEHOLDER can ensure that its privacy requirements remain satisfied. In essence, these claims provide evidence of a MEC A's operational integrity, including factors such as reliability, availability, and adherence to operational standards, thus they serve as proof of a MEC A's security posture, providing useful information to a MEC B. 2. a STAKEHOLDER may verify that the security and privacy guarantee of an identified SERVICE satisfy its own security and privacy requirements

	under acceptable trust assumptions. This covers both extremes: i) a service is deployed as part of one MEC domain, which means that it is managed by one MEC operator, having access to a specific set of security controls that can guarantee a specific trust level or ii) a service which expands between multiple MEC domains that exhibit different trust levels depending on the type of security controls they have deployed. This second category opens further questions regarding the privacy of an imminent handover from one domain to another.
Connected To Other Requirements	FR.PR.04 (Privacy): Privacy can be achieved via (a) anonymous services or (b) guaranteed protection and removal of PII - e.g., protecting and deleting PII inside a STAKEHOLDER-verified TEE.

8.1.2 Security & Operational Assurance Requirements

Apart from the Trust Assessment Requirements, there is the need to define more traditional security requirements to ensure protection against various threats and vulnerabilities, thus cybersecurity and system reliability guarantees for the emerging CCAM landscape. Traditional security requirements encompass fundamental aspects such as confidentiality, integrity, availability, authentication, and non-repudiation. These address core security concerns, including safeguarding sensitive data, ensuring data integrity, maintaining system availability, verifying user identities, and preventing repudiation of actions.

By splitting security requirements into these categories, organisations and system designers can effectively address foundational security needs and then broader them considering also the trust-related issues, creating resilient and trustworthy systems for the ever-evolving AVs landscape.

In addition to the security requirements, we explore operational assurance requirements offering a solid categorization between the two, based on the following distinction: i) security specifications, which cover the more traditional security requirements, ii) runtime operational correctness, which focuses on the correctness of the collected evidence as attributes to offer trust guarantees, as well as iii) function isolation and migration to ensure that critical functions of the trust framework as well as the key-related functions are protected.

8.1.2.1 Security Specifications

ID	FR.SR.1 (Mandatory)
Title	**Dynamic Credential Management**
Actors/Components Involved	ECUs, Zonal Controllers, Attestation and Integrity Verification, Trustworthiness Claims Generator, Key Management System, TAF, MEC.
Description	**Background:** The current V2X landscape is based on multiple certificates. Starting from certificates of authentication of validity of the in-vehicle components (i.e., trustworthiness evidence), going all the way up to certificates provided by the current PKI system with the parallel use of pseudonyms to provide an identifier while protecting the privacy of the vehicle. Evidently, apart from the certificates used to verify for the identity of the vehicles, due to the data-centric approach on trust, certificates are also employed to verify the data provenance; a fact that adds further complexity to the system. *Therefore, there is a need for a dynamic credential management system, which supports both the identity related certificates and the certificates of provenance needed*. The latter are consumed by the trust assessment framework as evidence, enabling the execution of certain services. These trustworthiness assessments will vary over the time regarding the data, since new information arrives from the trust sources.

	Currently, the PKI system that is used to ensure (as much as possible) the privacy of vehicles as they send CAM/CPM messages relies on pseudonyms. These are issued to the vehicle in advance and can then be used as required. While this does give some assurance that any message received by a vehicle, or the MEC, came from a valid vehicle, this says nothing about the trustworthiness of the data received. **Description:** The dynamic credential management (DCM) system will support both certificates of authentication for the validity of the in-vehicle components, as well as and the issuance of trustworthiness claims (TC) used to certify the provenance of the data. The certificates for the validity of the in-vehicle components will make use of pseudonyms issued by the PKI, to conceal the identity of the vehicle, while certificates of provenance from the data collected by the in-vehicle components are used for the assessments of the trustworthiness. These trustworthiness assessments will vary over time as the system changes, for example, the misbehaviour detection system may identify discrepancies in the GNSS data, and this will be reflected in the trustworthiness level reported with that data. This use of TCs will extend to the MEC, which will also issue them alongside any data that it provides. The use of TCs allows parties to issue (publicly) *verifiable statements* that can be backed by *trustworthy evidence*, as enabled by the underlying Chain of Trust. These statements will report on the current state of the system and, where appropriate, will be supported by cryptographic primitives used to achieve authentication, confidentiality, integrity and authorship of statements. To do this, the DCM system will closely interact with the relevant key management systems both in the vehicle and the MEC (ref: FR.SR.3). CONNECT will investigate the design and issuance of the following types of certificates: • **certificates for the validity of the in-vehicle components**, making use of pseudonyms issued by the PKI, to conceal the identity of the vehicle. • **certificates of provenance from the data collected by the in-vehicle components** and are used for the assessments of the trustworthiness. **Remarks:** Key revocation, particularly for the privacy preserving signatures we envisage using for the vehicle's VCs, is important, but not part of the CONNECT research programme. Thus, we will rely on current well-established solutions [158].	
Connected To Other Requirements	FR.SR.2, FR.SR.3, FR.SR.4	
KPIs	**Description**	**Value**
	Size of the Trustworthiness Claims TCs (they need to be included in the CAM/CPM messages)	It should be ideally **< 30 bytes,** so that it can fit into the existing security header of the standardised CAM definition. In the case where the size of the TC >30 bytes, then extensive testing will be performed to evaluate the impact of increasing the size of the header and thus, the size of the message, as it pertains to the bandwidth needed for the communication between the

		CCAM actors. In this context, a detailed analysis will also be conducted on the integration of different omission strategies, similar to what has already been investigated in the context of omitting PKI based identity certificates to test the frequency based on which such claims should be sent.
	Signing and verification times for the TCs (for a typical vehicle OBU).	**100 signing/verifications per second** as per ETSI specifications.
ID	**FR.SR.2 (Mandatory)**	
Title	**Secure and Efficient Cryptography**	
Actors/Components Involved	ECU, Vehicle, MEC, TAF, Zonal controllers	
Description	**Background:** Starting from the work assumption that *security without efficiency is meaningless*, priority will be given to the latter. While safety has traditionally been the primary focus in vehicular systems, specifically when it comes to V2I and V2X communications, there is now a growing consensus that security is also crucial. Consequently, *it is necessary to prioritise the design of schemes that enhance security without compromising safety, to maintain vehicular system integrity and reliability, while ensuring individual well-being and system efficiency.* V2X and V2I technologies aim at increasing road safety, improving traffic flow, and overall providing transportation efficiency. To accommodate the diverse landscape of CONNECT, which involves a large amount of data and various types of computing devices, it is essential to employ efficient cryptographic algorithms. These algorithms should be capable of handling the heterogeneity present in terms of the different types of messages (such as CCAM, CPM, and security-related messages) and the varying capabilities of the devices involved. In addition to addressing security concerns, privacy aspects are also taken into consideration, particularly in the context of V2V (vehicle-to-vehicle) communication. The objective is to safeguard the identity of the vehicle, ensuring that it remains concealed and protected from unauthorised access or tracking. Traditionally, to achieve anonymization in this scenario, a public key infrastructure (PKI) has been employed with the support of pseudonyms. The PKI framework generates, distributes, and verifies pseudonyms to protect vehicles' identities. However, the existing schemes, such as PKI, mentioned above, must be reassessed in light of the integration of additional entities into the overall architecture. This reassessment is necessary to ensure that the confidentiality, integrity, and privacy requirements are adequately addressed. For instance, the addition of the Multi-Access Edge Computing (MEC) introduces a new entity into the system that plays a crucial role in providing more precise services with reduced latency. While this advancement enables the realisation of advanced Cooperative, Connected, and Automated Mobility (CCAM) services on Day 2 and Day 3, it also introduces a potential impact on the security of the V2X ecosystem. Therefore, *it is imperative to evaluate and adapt the existing schemes to accommodate the presence of the MEC, as well as other entities that may be added, while maintaining the security of the V2X ecosystem.*	

	These crypto requirements are not limited to the time, but further consider the size of the signature as well as the size of the certificate. **Description:** In this complex and heterogeneous V2X ecosystem, the crypto mechanisms must protect the whole lifecycle of both the data, as well as the vehicle and its building blocks that produce the data. Consequently, confidentiality and integrity requirements apply both on the communication and the devices, which constitute, in essence, the sources of trust. The information stemming from the communication of the different entities and the devices, is used by the Trust Assessment Framework (TAF) to take trust-related decisions. These sources of trust may be linked to the integrity of the devices that provide a CCAM service or the data used in a CCAM service, for example. Hence, the trust sources, which further participate in the attestation, must be protected with crypto mechanisms, to conceal: i) the *types of signatures that are used to verify the integrity* within and outside the vehicle, where privacy protection is also crucial; and ii) *the type of encryption*, to achieve confidentiality. CONNECT aims at studying, thus proposing new *lightweight crypto schemes*, to cover all V2X security requirements both on the communication, as well as on a device level. Observing that different parts of the system impose different constraints, crypto algorithms will be integrated to be efficient in situ, while also guaranteed to be interoperable between different parts of the framework. Algorithms as such, are not only evaluated in terms of *time*, but should also consider *the size of the signature and the certificate*, to propose the idea schemes per case. CONNECT will investigate the design of lightweight crypto for the following operations: ● **enable the integrity of communications both within and outside the vehicle,** further considering the different levels of privacy, especially outside the vehicle (i.e., leveraging the existing standardised pseudonym-based signatures of PKI schemes) ● **enable the configuration integrity of every in-vehicle computer and ECU** (i.e., leveraging the in-vehicle manager), **further considering the privacy protection**, through anonymization, of relevant information (i.e., identities) included in CCAM, CPM and security related messages, which are shared with other entities. CONNECT envisages the design of more advanced crypto schemes (i.e., Direct Anonymous Attestation), adopting the need of zero trust paradigm. **Remarks:** To achieve the vision of zero trust all crypto operations need to leverage the *root of trust capabilities* (i.e., CONNECT's TEE Guard) to provide CCAM actors with certificate self-issuance capabilities as well as the *dynamic credential management system for continuous authentication and authorization* between CCAM actors, as a means for the verifiable exchange of either identifying attributes (i.e., for authorization) or other trust related attributes for the dynamic trust assessment. Since the vision of CONNECT is based on a hierarchical in-vehicle topology (i.e., ECUs connected to Zone Controller which is controlled by the in-vehicle manager), the design schemes to achieve the necessary requirements need to be able to take under consideration the available resources (i.e., considering both symmetric and asymmetric capable ECU with limited resources, as well as in-vehicle computers with more resources). Therefore, CONNECT's priority is to provide crypto agility in the sense that different design schemes will be designed, capturing the capabilities of all ECUs in an interoperable way.

Connected To Other Requirements	FR.SR.1, FR.SR.3, FR.SR.4	
KPIs	**Description**	**Value**
	Number of crypto operations (i.e., signature encryption per second)	To support **100 signing/verifications per second**
	Crypto agility (i.e., in terms of supporting different types of crypto primitives for different ECU types) This can include simple breakdown between symmetric/asymmetric but also to more advanced crypto including threshold signature or DAA for the communication integrity of the produced TCs.	**>=3** crypto primitives that can be supported
	Size of signatures and certificates (i.e., in terms of overhead introduced due to the volume of data)	**<= 30%** overhead introduced by the size of crypto structures (i.e., signatures and certificates), associated to trust related information, compared to the ETSI standardised certificates regarding identity management
	Computational resources (i.e., in terms of overhead introduced due to the CPU cycles for the crypto operation execution).	**<= 30%** overhead introduced by the trusted computing mechanisms provided by the CONNECT TEE Guard. This will include the detailed benchmarking of all crypto operations, needed to support the CONNECT trust assessment when instantiated and executed both inside and outside the employed RoT.
ID	**FR.SR.3 (Mandatory)**	
Title	**Flexible and Reliable Key Management**	
Actors/Components Involved	ECU, Vehicle, MEC, TAF, Zonal controllers	
Description	**Background:** The process of key management encompasses the secure generation, distribution, operation, and deletion of cryptographic keys. Some keys can be generated and distributed *during vehicle production*, while others need to be produced or exchanged *during vehicle operation* in the field. Many cybersecurity mechanisms used to protect CCAM systems against threats are based on cryptographic algorithms in combination with cryptographic keys. In the context of in-vehicle systems, key management might seem straightforward since the vehicle manufacturer is primarily responsible. However, manufacturers need to *coordinate and integrate key management with their suppliers*, including hardware suppliers and Tier-1/ECU suppliers, to create a comprehensive key management system. In addition, CCAM systems employ *pseudonymous identities* to protect the real identities of road users, which demands the management of *numerous keys for different digital identities by each vehicle*. These aspects result in complex key management systems which are able to cover the huge number of keys, the different involved stakeholders, and the different key lifecycles.	

<table>
<tr><td></td><td colspan="2">This complexity can be showcased via the VRU use case in the “As-Is" stage. In this use case, the in-vehicle side includes 9x long-term keys to authenticate each single ECU with cryptographic capabilities (A- and S-ECUs), 9x short-term session keys, 9x long-term keys for secure boot processes, 520x pseudonym public and private keys for V2X communication (assuming a set of 20 pseudonyms with a validity period of 2 weeks over the course of 1 year. With the introduction of CONNECT concepts, even more keys will be needed and managed.

Description: Effective key management is a crucial aspect in the management of keys for all CCAM actors, with a particular emphasis on in-vehicle key management systems. The primary objective is to guarantee the confidentiality and authenticity of the data that is exchanged between various components, thereby enhancing the overall security of the CCAM system. Consequently, the establishment of a secure and reliable communication framework becomes imperative. In order to guarantee the integrity and confidentiality of communication between various components, it is crucial to establish suitable identity keys during the manufacturing process of vehicles. These identity keys serve as a foundation for deriving authentication and encryption keys, thereby ensuring secure and authenticated communication.

In the context of CONNECT key management, specific rules must be established to safeguard the confidentiality, integrity, availability, and authentication of the key sources. Research is required to determine how to efficiently manage the entire lifecycle of the necessary keys. This includes exploring the possibility of establishing links between some keys, such as deriving keys from each other using the trusted component's key hierarchy to reduce the key overhead.

Several types of keys that might be of interest are considered for key management, including the Endorsement Key, which serves as a component identity key for secure onboarding in the system. Additionally, attestation keys and keys per CAM/CPM ID are relevant for maintaining the security and trustworthiness of the CCAM ecosystem.

CONNECT will investigate the design of lightweight crypto for the following operations:
• to establish suitable identity keys during the manufacturing process of vehicles to provide confidentiality and authenticity of the data that is exchanged between various components.
• establishing links between some keys using the trusted component's key hierarchy such as deriving keys from each other to reduce the key overhead.

Remarks: Approaches also in the context of binding or linking keys together as part of the key hierarchy will also need to be investigated to have efficient key management.</td></tr>
<tr><td>Connected To Other Requirements</td><td colspan="2">FR.SR.1, FR.SR.2, FR.SR.4</td></tr>
<tr><td rowspan="2">KPIs</td><td>Description</td><td>Value</td></tr>
<tr><td>Overhead against using keys with and without key restriction usage policies.</td><td>around 25% overhead should be introduced by the key restriction usage policy manager introduced by the employed TEE Guard. This in turn will allow the detailed benchmarking of both local and</td></tr>
</table>

		remote attestation processes that CONNECT will provide for extracting information evidence as trust sources.
	Efficiency of key hierarchy construction of different types of keys.	**<= 60 ms**; this considers the construction of the appropriate key hierarchies comprising all of the necessary crypto primitives and keys, needed to support the entire lifecycle of a system (i.e., from its authentication and onboarding to its application participation and trust related evidence secure communication)
	Types of keys to be supported and maintained	**> 4 keys** (i.e., identity key, integration key, authentication key, attestation key, etc.)
ID	**FR.SR.4 (Mandatory)**	
Title	**Secure Data Handling and Provenance**	
Actors/Components Involved	ECU, Vehicle, MEC, TAF, Zonal controllers	
Description	**Background:** In a fully decentralised environment, such as the one envisioned in the CONNECT framework multiple data sources and data processing units are distributed across various layers of the application stack, spanning from the vehicle to the Multi-Access Edge Computing (MEC) and the backend digital twin. Achieving both security and operational assurance in such a dynamic system poses unique challenges. As previously described in chapter 2, this heterogenous decentralised environment introduces a complexity when it comes to the trust relationships and the trust domain. To form such trust connections, there is a need for stronger evidence for the *data provenance* and the *secure data handling*. This need stems from the data exchanged between the trust sources, which need to be also trusted; hence need to be associated with strong verifiable evidence. In this context the primary concern is accurately determining the location and origin of sensitive data that can be used as sources of trust, due to the system's dynamic nature. The *variety of the trust properties*, spanning from integrity to resilience to robustness, as mentioned in chapter 3 and further elaborated in D3.1, shall be further considered. This variety and heterogeneity are representative of the trust sources that a system needs to collect. Hence, it is imperative to have mechanisms for data provenance and secure data handling, both in the context of the application generating or processing the data, and the location. For instance, inside the vehicle, components should be able to check the integrity and provenance of the data that they receive, so, for example, a steering controller should not accept instructions unless it can verify their source. Outside the vehicle, message authentication plays a crucial role in ensuring security in V2X communications, as highlighted in [159]. While modern solutions that prioritise anonymity are valuable, they may overlook the importance of considering linkability in certain cases. With MEC applications and network functions able to operate anywhere within the infrastructure, it is imperative to improve data provenance. This involves attaching assertions that audit the data lifecycle's integrity and correctness, especially for safety-critical applications. To accomplish this, the environment	

model must be constructed by integrating and validating data from multiple sources, including vehicles, Roadside Units (RSUs), and MEC. By establishing a trustworthy and auditable data provenance mechanism, the system can ensure the veracity and traceability of vital data, thereby augmenting the security and operational assurance of the V2X ecosystem.

Description: Based on the aforementioned considerations, a crucial security requirement in the context of V2X systems, such as CONNECT, involves establishing mechanisms that provide *runtime evidence* for trust assessment and ensure appropriate data associations. These mechanisms should enable authenticated entities and components to trace and link data back to their sources (when necessary), facilitating the assessment of the trustworthiness of participating entities (i.e., data origin), thus enabling node and data-centric trust relationships. Ensuring this linkage is essential when collecting trust properties/evidence, as it ensures the integrity and authenticity of data are tied to specific entities before trust relationships can be established.

Nevertheless, linkability is not always a desired property in modern systems due to privacy considerations. *Towards this direction and in order to provide the desired trust relationships and accountability within the system, while aligning with the privacy profile of the entities, strictly controlled linkability is performed*. This means that only authenticated and authorised entities and components should be able to link the evidence back to the data source. The integration of appropriate cryptographic primitives allows for the deployment of controlled linkability, safeguarding the privacy of entities while maintaining trust mechanisms.

In addition to verifying the integrity and authenticity of trust properties, evidence is also employed to verify that data have been processed by certified applications. This further enhances the trustworthiness of the data and ensures that processing occurs only through authorised channels, instilling confidence in the entire V2X system.

As an example, consider a scenario where a car undergoes a trust assessment within the Multi-Access Edge Computing (MEC) environment, and its trustworthiness claims are unsuccessful during attestation. In such cases, it becomes crucial for either the Trust Assessment Framework (TAF) or the Original Equipment Manufacturer (OEM) to be able to trace back the data to the specific vehicle involved. This linkage allows for the identification of the Electronic Control Unit (ECU) that failed the attestation process, enabling appropriate actions to be taken to address the security or operational concerns associated with that specific component. By establishing appropriate associations between runtime evidence and the entities involved, the CONNECT framework ensures that the necessary information can be traced back to its origin, facilitating effective troubleshooting and remediation processes.

The following cases are identified for both uplink and downlink:

- **Verifiability of the data used for the trust assessment (i.e., message application payload in the uplink):** *Both the data itself as well as the source of the data should be trusted*. Ensuring the authenticity and integrity of *data collected* and transmitted from the vehicle to the services, (i.e., Misbehaviour Detection), as well as establishing a *clear linkage* to the specific data source are hence two crucial aspects. The verifiability of data varies based on the trust relationships being assessed. For instance, when assessing the trust relationship between the vehicle and the MEC, the verification is directed towards the MEC side, to validate both the integrity and authenticity of the message application payload as well as verify the source of the data. This validation ensures that the received message indeed originated from a

	legitimate source and remains unaltered. Nevertheless, since the linkability with the source of the data is not always desired due to privacy considerations, strictly controlled linkability shall be performed. Hence, there is a tradeoff between privacy and data provenance. For example, the DAA-A scheme enables attribute verification with zero knowledge of the source whatsoever. • **Integrity and attribution of the message in the downlink**: This requirement focuses on maintaining the integrity and attribution of the application payload during the transmission from the MEC service to the vehicle. It ensures that the message received by the vehicle has not been altered or modified in transit, maintaining its integrity and authenticity. Additionally, it establishes the attribution of the message to the specific MEC service, ensuring that the recipient vehicle can trust the source of the information received.
Connected To Other Requirements	FR.SR.1, FR.SR.2, FR.SR.3

KPIs	**Description**	**Value**
	Controlled linkability (i.e., in terms of the time needed to integrate the necessary crypto primitives such as link token, as part of a TC, so that only authorised CCAM actors, like such as OEMs, can link back to an in-vehicle system)	**< 2sec** This is of particular importance in the case of an ECU with failed attestation evidence for which the OEM should be able to link back to the ECU id from the received TC, containing the harmonised attributes.
	Vehicle privacy exposure due to the communication of trust related information	**FALSE** (i.e., there should not be any gain for an adversary monitoring the system to deanonymize vehicles or link actions back to vehicles by overhearing the transmission of the trust related information)

8.1.2.2 Runtime Operational Correctness Evidence as Attributes for Trust Provisioning

ID	**FR.OC.1 (Mandatory)**
Title	**Common Trusted Computing Protocols**
Actors/Components Involved	Trusted Computing Base
Description	**Background**: The increased usage of digital software and radio links to the outside world, particularly the Internet, makes modern vehicles more susceptible to malicious actors [160]. A single compromised sender can, by using standardised protocols (e.g., for CCAM messages), reach and thus attack numerous receivers. Since a successful attack on a car could endanger human lives, the development of vehicular IT systems should call for the adoption of strict security measures. Related works in the field of V2X employ a Trusted Computing Base (using for example, a Trusted Execution Environment) to provide protection of in-vehicle ECUs and their communications and hence ensure the integrity of data, such as that coming from sensors from its collection till its transmission from the vehicle [161][162]. Note that, this integrity protection mechanism must consider

	the possibility of both software and hardware compromise. In chapter 8.1.4 we give an overview of the TCB and how it is used to protect a device from attack, or malfunction. In outline TCB of a device is the software stack and hardware components that are required for it to function correctly and guarantee the security/privacy of the given function, service or requirement that it supports. While a TCB implemented in software running in a TEE or secure enclave may be protected as a whole, it may also be built on top of a smaller "Root of Trust (RoT)". This is usually a vendor-provided hardware and firmware feature that allows the assessment and validation of application software that is part of the TCB. Examples of roots of trust that allow protected execution of application software (and also building of TCBs) are ARM TrustZone, Intel Software Guard Extensions, or AMD Secure Encrypted Virtualization. Other roots of trust, like trusted platform modules (TPM) or hardware security modules (HSM) do not provide protected execution of applications but do allow a TCB (containing these applications) to be built upon them. To enable hardware-backed trust assessment, CCAM actors that provide safety-critical functionalities should be equipped with a hardware root-of-trust so as to be able to provide the necessary guarantees on the operational correctness as well as enhanced crypto primitives for the confidentiality and integrity of sensitive data. A hardware RoT will provide evidence on the trust properties required for assessing the level of trust for the target actor and is likely to achieve a higher level of assurance as compared to one that is only implemented in software. However, in a highly heterogeneous landscape, consisting of different Original Equipment Manufacturer (OEM), RoT choices may vary both in terms of software vs hardware, as well as design choices across different vendors (i.e., TEE, ARM TrustZone, IntelSGX, automotive HSMs, etc.). In CONNECT it is essential to establish a set of protocols that is flexible and adaptable enough to work on top of different types of trust anchors. It should not be dependent on a specific model or brand of trusted software/hardware; instead, it should be able to support any trusted components that satisfy the defined properties and requirements. The CONNECT protocols should support interoperability, secure communication between different systems, remote attestation and secure update of the software protected by the TCB whatever the underlying hardware may be. **Description**: In the context of the CONNECT project, *it is essential to establish a set of protocols that is flexible and adaptable enough to accommodate different kinds of trusted components*. While the TEE may be vendor-specific, the software protected by it (i.e., the TCB) should not be overly constrained by the specific model or brand hardware TEE. Instead, it should be developed to support any trusted component that satisfies the defined properties and requirements. To ensure a consistent and standardised approach to security, all CCAM actors participating in the CONNECT project should be able to provide or be described by trust evidence that can be processed by the Trust Assessment Framework. In the context of CONNECT, a trusted component is deemed to be part of the TCB if it supports the properties outlined in section 8.1.4 of the present deliverable. These properties include: ✓ secure storage and ✓ secure boot mechanisms which are essential for safeguarding sensitive data and ensuring the integrity of the system during the boot process. By adopting trusted components that meet these specifications, the CONNECT TCB is able to integrate a wide variety of technologies and implementations while maintaining the desired level of security.

	By being agnostic regarding the type of trusted component integrated, CONNECT seeks to foster collaboration and innovation within the CCAM ecosystem. It enables the use of a variety of trusted hardware and software solutions, allowing for flexibility and adaptability as technology evolves. This strategy ensures that the CONNECT system is future-proof and can accommodate advancements in security technologies without being bound to a particular vendor or technology stack.	
Connected To Other Requirements	TR.4, FR.OC.4	
KPIs	**Description**	**Value**
	Granularity of Levels of Assurance (LoA) that can be achieved by various RoTs and essentially the common trusted computing base	**>= 5 LoA**; CONNECT will adopt and build on top of the classification of LoA specified by ETSI in the context of virtualized infrastructures. The same classification of LoA for the MEC will also be employed by CONNECT and an equivalent classification will also be provided for capturing the LoA for the vehicles.
	Number of operations supported by such a TCB (i.e., secure storage, secure boot, key management, etc.).	**>= 3 operations**
ID	**FR.OC.2 (Mandatory)**	
Title	**Operational Assurance & Configuration Integrity**	
Actors/Components Involved	Trust Assessment Framework	
Description	**Background:** To ensure proper configuration and provide verifiable evidence of secure operation at runtime, the V2X system must incorporate robust security enablers to support attestation processes. Ensuring the integrity and authenticity of all messages, particularly in the context of CCAM communications, is crucial in order to protect road users from potential safety hazards that may arise from compromised information. In accordance with the discussion in **FR.SR.1**, it is necessary for every message to be signed with a private key. This signature serves the purpose of generating a certificate as proof-of-authenticity; hence ensuring the message's integrity for the intended recipient. The current V2X framework is dependent on multiple certificates for the purpose of verifying the authenticity of the numerous in-vehicle components and confirming the origin of collected data. In a V2X system that provides high-criticality services, the establishment of a secure communication channel is of utmost importance in order to safeguard the safety and integrity of road users utilising the network. Nevertheless, in addition to ensuring secure communication, *it is imperative for the system to also address the continuous verification of software components operating on edge nodes in order to guarantee the integrity and trustworthiness of the entire system*. The successful execution of software and services on edge nodes requires the establishment of robust mechanisms for continuous integrity monitoring. *These mechanisms are responsible for verifying the reliability and accuracy of the software and services throughout their entire runtime*. The requirement for runtime verifiable evidence grows as a result of the continually shifting nature of the V2X ecosystem. Software components are exposed to a range of threats and attacks while they are in operation. It is crucial to promptly identify and address any potential tampering or unauthorised modifications to ensure the	

	security and integrity of the components. Through continuous monitoring of the integrity of software components, the V2X system possesses the capability to detect modifications that have the potential to undermine the system's functionality or compromise its security. To accomplish continuous integrity monitoring, the V2X system may employ *attestation processes that verify the authenticity and integrity of software components during runtime*. Attestation enablers and cryptographic techniques have the capability to produce evidence that verifies the reliability of these components. The utilisation of this dynamic assessment enables the system to make informed determinations regarding the trustworthiness of individual software components. Moreover, runtime verifiable evidence contributes to establishing security assurance within the V2X ecosystem. By ensuring that only trusted and unmodified components are permitted to execute and interact with one another, the system can maintain a robust security posture. This is particularly crucial for high-criticality services that have a direct impact on the safety of road users. **Description:** Security enablers play a critical role in *attestation processes, ensuring that components in the V2X ecosystem maintain the correct state both in terms of their configuration and their operational behaviour at runtime*. By providing verifiable evidence, these enablers establish and maintain trust among the different actors within the V2X system. Verifiable evidence ensures that each actor can carry out its tasks with transparency and accountability, promoting the overall security and integrity of operations. Attesting to the integrity of components involves measurement and verification processes, which ultimately lead to establishing information security assurance. To assess the trustworthiness of specific components, a scale of specific Levels of Assurance (LoA) should be defined. These LoAs will serve as a basis for determining the level of confidence in the integrity of each component. CONNECT, as part of the V2X ecosystem, focuses on *providing runtime integrity mechanisms*. These will involve collecting and assessing runtime data from system components, depending on the devices involved the assessments can be performed either locally on the edge device or remotely. Regardless of whether the root of trust is integrated into hardware or software, it must offer essential functionalities to ensure the V2X system's security: ✓ **Dynamic Assessment through Verifiable Evidence**: The platform dynamically evaluates the reliability of software building blocks, allowing only trusted components to execute and interact within the V2X ecosystem. This process establishes security assurance and assesses the integrity of each component, providing measurable and verifiable information. ✓ **Levels of Assurance**: To facilitate trustworthiness assessment, a common scale will be used, defining different "Levels of Assurance" (LoA) for various components within the CCAM ecosystem. This classification enables clear distinctions in trust levels, promoting effective decision-making and appropriate security measures for different components. **Remarks:** The attestation evidence being provided should be available in a timely way. To meet any timing constraints the attestation evidence may be cached and, in this case, when using the evidence in the TAF it may be given a lower confidence level depending on the time elapsed since the attestation evidence was obtained.	
Connected To Other Requirements	FR.SR.1, FR.SR.2, FR.SR.3, FR.OC.1, FR.OC.4	
KPIs	**Description**	**Value**
	Time needed for the execution of the local attestation assuming the provision of authenticated runtime measurements	**< 200ms** when the attestation process is instantiated and executed outside the CONNECT TEEs and

		< 10% overhead, when the attestation process is instantiated and executed inside the TEEs.
	Time needed for the construction and signing of the TCs	**< 900 ms**
ID	**FR.OC.3 (Mandatory)**	
Title	**Integrity Verification of CCAM Components**	
Actors/Components Involved	Trust Assessment Framework	
Description	**Background:** The CCAM functions are of utmost importance in safeguarding the safety of road users. This is because the decisions made by CCAM directly impact the driving behaviour of vehicles, such as in the case of Cooperative Adaptive Cruise Control (C-ACC), or indirectly guide drivers to take necessary actions, as seen in collision warning systems. The reliability of these functions is paramount and relies on the intended and certified behaviour of their implementing components. To safeguard this reliability, it is essential to protect the integrity and availability of these components against present and future cybersecurity threats. When the integrity of a component is compromised as a result of manipulation or cyber-attacks, its reliability cannot be assured, thus necessitating a diminished level of trustworthiness. In light of these circumstances, it is imperative for CCAM systems to promptly execute remedial actions, including the augmentation of safety buffers and the implementation of plausibility checks. In an alternative scenario, individuals have the option to communicate the compromised situation to the driver, thereby providing them with necessary data to make informed decisions. In more extreme cases where safe operation becomes unattainable, individuals may also have the ability to disengage the function. **Description:** Safety-critical CCAM services require robust protection from other functions operating within the same software host environment, in order to maintain the integrity of their operation and interaction with other components. To achieve this, the underlying root-of-trust (RoT) must be capable of detecting runtime attacks that do not modify the static state of the targeted component. *Additionally, the CCAM components must be able to respond effectively to changes in the trust level of the execution platform to ensure operational assurance of CCAM service execution.* The aforementioned requirement calls for the *creation of effective mechanisms that can swiftly adapt to alterations in trust levels within CCAM components*. These components encompass both software and hardware assets, as well as ECUs and sensors that furnish vital data for CCAM applications. The RoT is a crucial element in facilitating the identification of runtime attacks that have the potential to alter the static state of a component, thereby impacting the integrity attribute of the trustworthiness level of the CCAM. In order to effectively implement policies in response to changes in trust levels, it is imperative to develop a range of diverse mechanisms. These policies may encompass measures such as deactivating the compromised ECU, reverting to a secure state, or, in more intricate situations, transferring the ECU's state to an nearby ECU possessing the required level of trust. It is crucial that these actions are carried out in a manner that does not cause any disruption to the service and does not compromise its safety. CONNECT will investigate the design of mechanisms for the following operations: ✓ implement **policies** effectively in response to changes in trust levels	
Connected To Other Requirements	TR.4, FR.OC.2, FR.OC.3, FR.SR.3	
KPIs	**Description**	**Value**

	Time needed for the execution of configuration integrity verification	**< 100ms** when the attestation process is instantiated and executed outside the CONNECT TEEs and **< 25% overhead**, when the attestation process is instantiated and executed inside the TEEs and leverages the established key restriction usage policies enabling local configuration integrity verification
ID	**FR.OC.4 (Mandatory)**	
Title	**Chain of Trust Creation**	
Actors/Components Involved	Trust Assessment Framework	
Description	**Background:** As described in requirement **FR.OC.1**, components in the system will all have a Root of Trust that will support and attest to the component's Trusted Computing Base (TCB). These attestations allow us to establish a hierarchical chain of trust that ensures the integrity of the systems software and hardware components, thereby enhancing the overall security posture. The concept of the chain of trust is a hierarchical series of trusted components within a system that ensures the integrity and security of its components, both in terms of software and hardware. It is based on a succession of verifications and attestations, beginning with the ECUs, then the zonal controllers and then the secure containers in the vehicle computer itself (see Figure 26). As part of establishing the chain of trust, we need to create a secure environment in which only trusted and validated components may operate and interact with one another. By building this hierarchical trust relationship (i.e., also mentioned in the requirement **FR.SR.2**) and the establishment of a chain of trust within the system we provide a framework where trusted entities can delegate their authority to other trusted entities, enabling them to act on their behalf. In this context, the requirement suggests that an ECU (Electronic Control Unit) may authorise another trusted entity to perform certain actions on its behalf. The ECU can expand its reach and capabilities by delegating duties to another trustworthy entity. This delegation enables a more flexible and distributed system architecture in which various groups can interact and contribute to the system's overall functionality and security. For example, if an ECU needs to perform a specific operation but lacks the necessary resources or capabilities, it can delegate the task to another trusted entity that possesses the required expertise or resources (i.e., the Digital Twin). This delegation ensures that the operation is still carried out by a trusted entity, maintaining the overall security and trustworthiness of the system. **Description:** In the context of building a chain of trust for any given service, all CONNECT entities and actors involved should actively participate. This can be achieved through direct involvement or delegation, where one party acts on behalf of another. For example, when the Digital Twin or the In-Vehicle Manager initiates an operation based on input received from an Electronic Control Unit (ECU), there should be a sufficient level of trust in these entities. During the execution of an operation, all actors involved should be able to provide evidence of their level of assurance. This includes demonstrating their trustworthiness from the moment of their trusted launch and configuration to their ongoing runtime attestation. By establishing and maintaining trust throughout the operation, the system can ensure that each actor and entity is reliable and can be trusted to carry out their designated tasks.	

	In essence, this trust should be verifiable and based on reliable evidence, allowing for transparency and accountability within the system. By implementing a robust chain of trust, the V2X ecosystem can enhance security, ensure the integrity of operations, and build confidence in the overall reliability and trustworthiness of the system. Hence the property that CONNECT aims to collect here is: ✓ **verifiable evidence** for the hierarchical trust model for all operations **Remarks:** Although the chain of trust is most easily viewed as something built within the vehicle, it can be extended outside of the vehicle to include the MEC and its components as well. How useful this is is not clear as the vehicle will be switching from one MEC to another as it moves around. This applies even more strongly when considering including other vehicles in the chain of trust as they are coming and going all the time.	
Connected To Other Requirements	FR.OC.1, FR.SR.2	
KPIs	**Description**	**Value**
	Storage of trust-related information to the Blockchain for auditability and certifiability.	**<= 5sec,** since this is not a real-time operation.
	The AIV should be able to request transmission of fresh raw evidence by the ECUs depicting the current state (so as to be used as a trust source).	**>= 1 Mbit/sec** data transfer rate for the raw evidence
	Simulate low bandwidth channel with high message loss.	**<1 sec** for the BC transaction
ID	**FR.OC.5 (Mandatory)**	
Title	**Secure Measurement/Attribute Extraction**	
Actors/Components Involved	Trust Model	
Description	**Background:** Remote attestation of system integrity is an essential part of trusted computing. Current remote attestation techniques should provide integrity proofs of both static and dynamic system properties in order to offer holistic security, through integrity verification. For the dynamic attestation, dynamic properties are used to verify the runtime integrity of the system, hence providing the integrity evidence. Monitoring runtime data and execution streams requires actively tracking and assessing the data and execution flows within the system while it is in operation. This monitoring enables the collection of useful information about the behaviour and integrity of system components. Potential anomalies, security breaches, or deviations from intended behaviour can be detected by tracing the execution stream and monitoring the runtime data. However, according to [163], creating a dynamic attestation system may present some issues. Firstly, the dynamic and temporal nature of objects and attributes, makes it difficult to identify and deduce their "known" good states. In static attestation the known good state is measured using static objects' cryptographic checksums. Another point raised is that in the case of the dynamic attestation, access and verification to large amounts of information is required, which raises the issue of efficiency and scalability. Hence, it is of utmost importance to design such schemes to collect useful information, while minimising the overhead, thus ensuring that the monitoring does not reduce the system's responsiveness or efficiency. Same notion also applies for the scalability of the system, to ensure that large volumes of data will not affect the system's efficiency.	

	Description: CONNECT should provide support for runtime data and execution stream monitoring and introspection. This capability is crucial for efficiently tracing the system properties required to establish the level of trust in the system. It emphasises the need for dynamic tracing functionalities as an integral part of the underlying root of trust. The need emphasises that this monitoring and introspection capability be efficient, allowing the system to effectively trace the required system properties. These properties serve as evidence to determine the system's level of trust, ensuring that it runs reliably and securely. Furthermore, the requirement emphasises the significance of dynamic tracing functionalities as an underlying root of trust. The capacity to trace and monitor system operations in real-time, allowing for the detection of security vulnerabilities or abnormalities as they occur, is referred to as dynamic tracing. The system can efficiently respond to emerging threats and maintain a high degree of trustworthiness by embedding dynamic tracing into the root of trust. Hence the property that CONNECT aims to collect here is: ✓ **verifiable traces** (i.e., runtime data and execution streams)	
Connected To Other Requirements	TR.4, FR.OC.2, FR.OC3, FR.OC.4	
KPIs	**Description**	**Value**
	Efficiency of tracing and device state monitoring	**<= 500 ms**
ID	**FR.OC.6 (Mandatory)**	
Title	**Secure Remote Asset Management and Reconfiguration Effectiveness**	
Actors/Components Involved	Trust Assessment Framework	
Description	**Background:** The capability to remotely update vehicles is a critical aspect of ensuring the long-term safety and security of both vehicles and road users. Unlike consumer electronic devices that have relatively short product lifetimes, vehicles can remain in operation for multiple decades. A typical vehicle model may be in production for 6 to 7 years, and these vehicles can continue to be used in the field for around 15 to 20 years. Over such extended periods, it is highly likely that the cybersecurity measures implemented during the vehicle's design phase will become inadequate to counter emerging and unforeseen cyber threats. Recognizing the importance of addressing this challenge, the United Nations Economic Commission of Europe (UN ECE) has taken significant steps to address vehicle cybersecurity. They have included cybersecurity as a mandatory aspect of the vehicle model's type approval process, known as homologation. This inclusion is established through two regulations: "UN R155 Uniform provisions concerning the approval of vehicles with regards to cyber security and cyber security management system" and "UN R156 Uniform provisions concerning the approval of vehicles with regards to software update and software updates management system." The UN R155 regulation specifies the requirements for a vehicle's cybersecurity management, emphasising the need for cybersecurity maintenance and updates. This ensures that vehicles can adapt to evolving cybersecurity threats and challenges throughout their long-running product lifetimes. The regulation aims to ensure that vehicle manufacturers implement robust cybersecurity measures and are prepared to address potential vulnerabilities with timely updates. Similarly, the UN R156 regulation focuses on software updates management, addressing the process by which vehicles receive and apply software updates. This	

	regulation ensures that vehicles can receive necessary software updates securely and efficiently, providing a means for manufacturers to address software-related vulnerabilities or introduce new features and improvements over time. By incorporating cybersecurity and software update management requirements into the type of approval process, the UN ECE aims to enhance the resilience of vehicles against cyber threats, ultimately contributing to safer and more secure road transport systems. The ability to remotely update vehicles allows for ongoing protection and improvement of cybersecurity measures, even as the threat landscape evolves, and new challenges arise. **Description:** CONNECT shall *provide a secure remote asset management and reconfiguration capability to enable the seamless and secure update of services and safety-critical functions*. This capability shall ensure that the trusted state of the overall service graph chain is not compromised during the remote upgrade process. The system shall support the reconfiguration of cooperative algorithms, allowing them to adapt to new mobility patterns and novel attacks while maintaining the highest level of security. To achieve this, the V2X system shall *incorporate mechanisms that facilitate secure and authenticated remote upgrades, ensuring that only authorised and verified updates are applied to the system*. The remote asset management shall be designed to prevent unauthorised access and tampering, thereby safeguarding the integrity and confidentiality of the system's components and data. The system's reconfiguration capabilities shall allow for the dynamic adjustment of cooperative algorithms to respond effectively to changing mobility patterns and emerging security threats. This adaptive behaviour shall be performed while preserving the trustworthiness and reliability of the CCAM services. Furthermore, the secure remote asset management and reconfiguration effectiveness shall be *seamlessly integrated with the stateful migration functionality*, as presented in the system architecture. This integration shall allow for smooth transitioning between different system states during reconfiguration processes, ensuring continuous service availability and maintaining the overall security posture of the V2X ecosystem.	
Connected To Other Requirements	TR.4	
KPIs	**Description**	**Value**
	Update and adaptation of the trust models capturing the newly updated asset.	**TRUE**
	Deployment (and revocation of old) of new system configurations including also possible SW upgrades safeguarded by the CONNECT TEE Guard.	**TRUE**

8.1.2.3 Function Isolation and Migration

ID	FR.SF.1 (Mandatory)
Title	**Dynamic awareness on potential vulnerabilities and threats and complete overview of the deployed environment**
Actors/Components Involved	Trust Assessment Framework
Description	**Background:** The complexity of modern V2X (Vehicle-to-Everything) systems is on the rise, as they now encompass highly distributed environments that

	consist of numerous interconnected and diverse devices. Transportation systems play a pivotal role in guaranteeing the safety and efficacy of transportation operations, rendering them highly susceptible to cyberattacks. Consequently, it is imperative to fulfil rigorous security prerequisites in order to ensure the preservation and dependability of V2X communications and services. To fulfil the security requirements, the inclusion of *a dedicated and resilient security monitoring platform is imperative within any V2X ecosystem*. The primary function of this monitoring platform is to consistently identify and address potential risks, evaluate weaknesses, and analyse the activities of network traffic and system events in a live environment. By implementing this measure, it functions as the primary safeguard against potential cyber risks and guarantees the overall security robustness of the V2X ecosystem. The security monitoring platform plays a crucial role in ensuring the safety of road users by identifying potential security breaches that may compromise the functionality of V2X systems, which include Collision Avoidance, Cooperative Adaptive Cruise Control (C-ACC), and Intersection Collision Warning. The preservation of user trust and confidence in V2X technologies necessitates a prompt and resolute reaction to security incidents. The pivotal nature of the System Administrator's role in this particular context cannot be overstated. In the role of the security service operator, the System Administrator requires the ability to perform real-time evaluations of the security status of the complete V2X system. This system encompasses a wide range of interconnected elements, including Roadside Units (RSUs), Multi-Access Edge Computing (MECs) nodes, and In-Vehicle Managers. Each of these devices plays a crucial role in enhancing the overall functionality of the V2X system. *However, due to their diverse characteristics, they also introduce distinct security challenges*. The implementation of prompt corrective actions is of paramount significance when confronted with emerging threats or vulnerabilities. The ability of the System Administrator to promptly and efficiently address potential risks is crucial in maintaining the resilience of the V2X system and ensuring the provision of secure and dependable services to individuals utilising roadways. **Description:** CONNECT utilises virtualization technologies to provide a comprehensive and user-friendly graphical depiction of the monitored V2X ecosystem. The utilisation of a graphical representation facilitates the process of monitoring and comprehending the interrelationships among the various devices present within the ecosystem, encompassing RSUs, MECs, and In-Vehicle Managers. This visual representation is of utmost importance for the System Administrator, as it enables them to efficiently evaluate cyber risks associated with essential services. This is particularly crucial in cases where new vulnerabilities or threats are identified either through CONNECT's advanced monitoring mechanisms or reported by the community in relation to recognized assets. CONNECT incorporates an advanced risk assessment tool that serves a crucial function in the ongoing analysis and surveillance of the V2X environment. The aforementioned tool effectively collects relevant information from the different actors involved in CCAM, integrating essential data that forms the basis for generating comprehensive risk reports. The risk assessment process facilitates the acquisition of comprehensive knowledge by the System Administrator regarding the security condition of the V2X deployment. This knowledge empowers the System Administrator to make prompt and well-informed decisions regarding corrective actions and attestation procedures. In order to maintain a leading position in automotive security, CONNECT's risk assessment solutions strictly adhere to industry best practices and standards. CONNECT ensures that its risk assessment capabilities remain at the forefront

<table>
<tr><td></td><td colspan="2">of emerging threats and challenges by adopting dynamic risk assessment methodologies in automotive environments. The adoption of a proactive approach is crucial within the dynamic V2X landscape, characterised by the constant evolution of cyber risks.

Furthermore, the outcome of the risk assessment procedure plays a crucial role in facilitating the formal verification of Hardware and Software co-design. The integration of hardware and software components within the V2X ecosystem is crucial for ensuring their harmonious functioning, thereby enhancing the overall security of the system.

The agility of CONNECT's awareness is indicative of its capacity to rapidly obtain and analyse events from the monitored topology in a nearly instantaneous manner. The prompt identification and timely resolution of potential threats through this rapid response capability effectively mitigate the risk of any exploitations that may compromise the integrity and safety of the V2X system.

As a result, CONNECT will offer the following capabilities:
✓ Risk assessment whose output will enable HW and SW formal verification co-design.
✓ Dynamicity on the awareness to acquire and process events from a monitored topology in near real-time manner</td></tr>
<tr><td>Connected To Other Requirements</td><td colspan="2">TR.4</td></tr>
<tr><td rowspan="3">KPIs</td><td>Description</td><td>Value</td></tr>
<tr><td>Dynamic Risk Assessment based on the identification of new threats</td><td><= 2sec considering the identification of a new threat (by a security administrator) based on monitored evidence collected as part of a failed attestation process that indicates a possible risk.</td></tr>
<tr><td>Recalculation of RTL considering the identification of new risks</td><td><= 3sec</td></tr>
<tr><td>ID</td><td colspan="2">FR.SF.2 (Mandatory)</td></tr>
<tr><td>Title</td><td colspan="2">Stateful Function Upgrade</td></tr>
<tr><td>Actors/Components Involved</td><td colspan="2">Trusted Execution Environment</td></tr>
<tr><td>Description</td><td colspan="2">Background: Security-critical components, specifically in the CCAM ecosystem, bear the responsibility of executing distinct functions or services that hold significant importance for multiple stakeholders; hence, play a crucial role in ensuring the security and integrity of the overall operation. Consequently, it is imperative for these components to uphold a state of critical security throughout their operational lifespan.

Nevertheless, in certain situations, it may be necessary to update security-critical components in order to improve their functionality or mitigate vulnerabilities. These updates could involve installing a new software release to improve the service or fix potential issues. In such cases, it is essential to carefully manage the update process to ensure that critical aspects of the component's prior state are preserved.

For instance, consider a smart motor control unit in a vehicle. While updating the software of this control unit, it is crucial to retain certain configuration settings</td></tr>
</table>

	and status data that were present before the update. This is necessary to ensure that the motor control unit continues to function correctly and securely, even after the update. An additional example can be found in the form of an odometer, which is required to precisely document the distance covered by the automobile. Preservation of the existing odometer reading, and any associated cryptographic keys is imperative during the firmware update process. By implementing this measure, it guarantees that the odometer is able to consistently fulfil its purpose of accurately measuring distance, while simultaneously upholding the integrity and security of the recorded data. It is evident that the inclusion of security-critical components is imperative for ensuring the comprehensive security of a system, and it is crucial to uphold their optimal performance throughout the update process. The preservation of essential elements of the system's previous state guarantees its continued secure and efficient operation, even following the implementation of required updates. **Description:** In the context of CONNECT and more specifically the software updates related to safety-critical functions, it is imperative to guarantee the secure execution of the update procedure, while also avoiding any disruption to crucial components of the application state on the device. The main goal is to facilitate the ability to "upgrade" while maintaining specific elements of the device's state and guaranteeing a seamless transition to the new software iteration. In order to accomplish this, the software stack on the device is categorised into two primary divisions: version-specific state and version-persistent state. The version-specific state refers to the state data that is intricately linked to the particular software version being executed on the device. Throughout the process of updating, the installation of new software occurs, which consequently necessitates the generation or installation of version-specific state to ensure compatibility with the updated software. The persistent state, on the other hand, refers to a specific type of state that encompasses vital information or configurations that need to be retained during software updates in order to ensure the continued functionality and integrity of the device. The aforementioned data is characterised as persistent and necessitates migration from the previous iteration to the subsequent iteration as part of the update procedure. The secure update mechanism must effectively address consequently, two primary aspects: ✓ **Installation of New Software:** The update process should allow for the installation of new software versions on the device without compromising its security or stability. This ensures that the device can benefit from the latest improvements and fixes. ✓ **Managing Version-Specific State:** Upon installation of the updated software, it is necessary to generate or install any version-specific state that is required to support the new software. This guarantees that the device functions accurately in conjunction with the updated software version. **Remarks:** An example user story that drives this requirement is a controlled upgrade of the TEEguard key storage to remove a newly discovered bug in the TEEguard application running in the TEE. Naturally, the keys managed by the TEEguard application need to be continuously protected and most need to be migrated to the new software version. More formally, the goal is to upgrade from one version of the application to the next signed version such that the keys remain usable inside the TEE without ever being exposed outside the TEE.

<table>
<tr><td></td><td colspan="2">Naturally, in addition to updating the software inside the TEE, one still needs to ensure that a larger update of many components maintains compatibility. I.e. the TEEguard update may be part of a larger update where many components are upgraded while maintaining their compatibility.

To achieve this requirement, the following aspects need to be considered:
✓ Controlled Upgrade: The TEEguard's functionality and key security should not be disrupted during the upgrade process. This means ensuring that the new version is securely signed and that all necessary measures are in place to protect the keys during the upgrade.
✓ Seamless Key Migration: The TEE must seamlessly migrate cryptographic keys from the current TEEguard to the new version. This ensures that the keys remain securely stored and accessible only to authorised processes within the TEE.
✓ Compatibility Maintenance: As part of the larger system update, the TEEguard upgrade should be coordinated with the updates of other components to maintain compatibility. This ensures that the entire system continues to function as expected after the upgrade.</td></tr>
<tr><td>Connected To Other Requirements</td><td colspan="2">TR.4</td></tr>
<tr><td>Impactful Attacks and Mitigation Measures</td><td colspan="2">To ensure that private credentials do not leak (see requirement in Section 3.5), we require that an upgrade provides well-defined security guarantees for the state of the component:
✓ A malicious actor may try to install bogus firmware to leak secret state or corrupt services. We require “authenticity of the upgrade” - that only authorised software updates can be installed in a component,
✓ A malicious actor may try to install an outdated firmware to potentially re-introduce vulnerabilities that have been fixed in the past. We require “downgrade prevention” - where the service prevents an outdated software version from being installed. This may require interaction with a third-party service.</td></tr>
<tr><td rowspan="2">KPIs</td><td>Description</td><td>Value</td></tr>
<tr><td>Upgrading the function in one ECU</td><td>< 750 ms</td></tr>
<tr><td>ID</td><td colspan="2">FR.SF.3 (Mandatory)</td></tr>
<tr><td>Title</td><td colspan="2">Stateful Function Migration</td></tr>
<tr><td>Actors/Components Involved</td><td colspan="2">Trusted Execution Environment</td></tr>
<tr><td>Description</td><td colspan="2">Background: As mentioned in FR.SR.2, security-critical components within the CCAM ecosystem must ensure safe and reliable operations and must maintain a state of critical security throughout their lifespan. Nevertheless, there are circumstances in which it becomes necessary to transfer security-critical operations from one node to another within the V2X ecosystem. The migration process may be motivated by a range of factors, including the desire to enhance performance, optimise resource utilisation, or adapt to evolving environmental circumstances.

Preserving the critical security state of these functions without compromise is of utmost importance during their migration. This implies that the security context, encompassing cryptographic keys, configuration settings, and other crucial data, must smoothly migrate to the new node. The migration process must ensure the uninterrupted operation and integrity of the security-critical function.</td></tr>
</table>

	The collaborative service in charge of managing intersections is a good case study in the importance of maintaining the security context during function migration. In specific circumstances, it may be necessary to transfer the service from a vehicle to a mobile edge cloud (MEC) in order to utilise increased communication bandwidth and enhance collaboration. Throughout the process of migration, it is imperative for the service to maintain its state of security and preserve its cryptographic keys. This is necessary in order to guarantee the uninterrupted flow of secure communication and effective coordination with other nodes within the V2X network. Another example involves a safety-critical function responsible for real-time hazard detection. If this function needs to be migrated to a more capable node to improve its responsiveness and accuracy, it is vital to transfer its critical security state to the new location. This ensures that the safety-critical function can continue to operate effectively while maintaining the necessary security measures. **Description:** CONNECT will implement the capability of migration to allow seamless transfer of services from one node to another within the V2X ecosystem. The term "migrate" encompasses a defined set of criteria that facilitate the smooth transfer of services from one node to another within the V2X ecosystem. These requirements ensure that the migration process is secure, reliable, and generic enough to be integrated into any type of security policy enforcement mechanism. The following are the essential criteria: ✓ **Migration of Authentic States:** The system must ensure that the migrated state is authentic, meaning that it comes from a trusted and verified source. This involves verifying the origin and integrity of the state before and after the migration process. ✓ **Integrity Preservation**: Throughout the migration process, the integrity of the state should be maintained. This means that the state should not be altered, tampered with, or corrupted during its transfer from one node to another. ✓ **Confidentiality Preservation**: If the state contains sensitive information, it is essential to maintain its confidentiality during the migration. Proper encryption and decryption mechanisms should be employed to protect the state from unauthorised access or disclosure. ✓ **Migration of Cryptographic Keys**: Depending on the type of function being migrated, cryptographic keys may need to be transferred along with the state. The system should facilitate the secure migration of cryptographic keys, ensuring that their confidentiality and integrity are preserved. ✓ **Flexibility and Generic Integration**: The migration capability should be generic and flexible enough to integrate with various security policy enforcement mechanisms. It should not be tightly coupled to a specific system or architecture, enabling easy deployment and adaptability across different environments. Part of this migration process involves the vital task of version-persistent state migration. The system must ensure the smooth transition of version-specific data, configurations, and critical information from the previous software version to the updated one. By accomplishing this process effectively, the device can continue its operation without any interruptions, adhering to its intended functionality and maintaining the necessary security measures. ✓ **Version-Persistent State Migration to** guarantee the preservation of crucial data or configurations, allowing the device to operate without any interruptions, in accordance with its intended functionality. **Remarks:** In some scenarios, we require a singleton service. I.e., that only one instance of a service can run at any point in time.

	An example user story that drives this requirement is a collaborative service (such as intersection management) that requires high-bandwidth collaboration. To enable higher bandwidth, the service is migrated from the vehicle to the MEC and then continues collaborating with the same secured state - while profiting from the enhanced communication bandwidth of the MEC.	
Connected To Other Requirements	TR.4, FR.SR.2	
KPIs	**Description**	**Value**
	Establishing a similar function on another box.	**< 1 sec** Note that CONNECT focuses only on the downtime/availability - NOT the preparation phase for the setup and synchronisation between the different ECUs. Hence this metric reflects the time needed to transmit state over CAN Bus - in-vehicle network latency, excluding network delays and latency (especially during migration between Vehicle and MEC).

8.1.3 MEC Operational & Security Requirements

As it pertains to CONNECT, the edge computing technology (based on relevant standards compliant implementation) is leveraged to deploy, support, and test the trust assessment mechanisms and the collection of the verifiable evidence based on which the trust evaluation will be conducted (e.g., trust calculations, attestation etc). It should be clarified that the project does not claim breakthrough innovation on the employed/demonstrated MEC (networking) technology. However, it uses the edge computing primitives (such as maintenance of data locality/proximity, fast response, computation, and virtualization) as a CONNECT concept *enabler* (for our trusted CCAM vision).

Therefore, in contrast to other sections, we herein recognize a set of requirements that are characterised as 'baseline'. This means that they constitute de-facto characteristics (i.e., background features) of the MEC technology, irrespective of what further requirements the CONNECT concept may pose. Those requirements prescribe the needed technology for the typical MEC operation (as defined mainly by standardisation documents mentioned below) and are evaluated differently compared to the rest, somewhat indirectly; by measuring regular network KPIs (such as end-to-end latency) we argue that those 'baseline' requirements are essentially met.

Then, we also identify some mandatory and nice-to-have MEC requirements closely related to the CONNECT framework. Hereafter we refer to the baseline, nice-to-have and mandatory requirements as *operational requirements*.

Both, the operational requirements engineering-work and the corresponding CONNECT prototypes (to be developed in the context of WP5) are well-aligned with the ETSI MEC specifications [164][165][166] which are broadly adopted by the networking and industrial community. In contrast, when it comes to the MEC security requirements, the relevant literature is less mature in terms of standardised concepts. To cover them, we mainly resort to publicly available documents from the same organisation [167].

In the following table (i.e., Table 27) we clarify the involved terms in MEC technology (maintaining an alignment with ETSI documents) where we also show a basic MEC overview [168] in the rightmost column.

Table 27 - Terms in MEC technology (according to ETSI)

Term	Description	(ETSI-introduced) MEC overview
MEC system	Umbrella term to include all the following modules and their	

	management. That is carried-out by the MEC orchestrator as its core component having an overview of the complete system.
MEC host	The MEC host (i.e., a computing infrastructure or device that is part of the MEC network) is an entity that contains a MEC platform and a Virtualisation infrastructure which provides compute, storage, and network resources, for the purpose of running MEC applications.
MEC platform	The MEC platform is the collection of essential functionality required to run MEC applications on a particular Virtualisation infrastructure and enable them to provide and consume MEC services. The MEC platform can also provide services.
MEC applications	MEC applications are instantiated on the Virtualisation infrastructure of the MEC host based on configuration or requests validated by the MEC management.

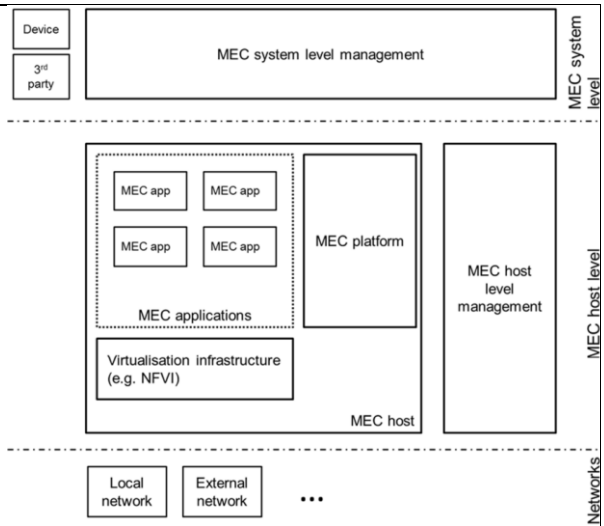

Finally, it is to be noted that CONNECT, especially in its experimental part (i.e., WP5, WP6) may potentially consider different (physical) locations keeping the same MEC platform. Different MEC platforms are outside our scope mainly due to the nature of the CONNECT consortium (lacking for instance an MNO or an actual edge computing facility provider).

ID	FR.MEC.1 (Baseline)
Title	**Operational requirements on (MEC) application lifecycle and application environment**
Actors/Components Involved	CONNECT MEC
Description	**Background:** Virtualization technology [169] (i.e., the one that allows the creation of applications using resources that are traditionally bound to hardware) enables the utilisation of a physical machine's full capacity by virtually distributing its capabilities to many users or environments. As such, it constitutes the de-facto approach to networking of distributed applications in connected systems (not limited to the automotive setting). Along these lines, the adopted edge computing infrastructure should follow the core principles of virtualization (architecture and design wise) to facilitate an agile framework capable of hosting and managing both the necessary trust and security extensions, which enable the dynamic trust assessment of frameworks, such as the one envisioned by CONNECT, but also orchestrating the workloads related to the various CCAM services. Additionally, the network-side requirements from the perspective of the applications (e.g., data rate and end-to-end latency) should also be ensured. **Description:** The CONNECT MEC system should adhere to a number of relevant operational requirements to capture the above needs. Most of them have

<table>
<tr><td></td><td>been introduced and established in the ETSI MEC specifications with which the CONNECT MEC technology is typically aligned.

1. The MEC system should reuse the NFV high-level functional architectural framework and design philosophy of virtualized network functions and services and of the supporting infrastructure, as described in the NFV architecture framework of ETSI GS NFV 002.
2. The MEC system should provide capabilities to interact with the 5G core network on behalf of (CONNECT or CCAM) applications, to influence on the traffic routing and policy control of UPF (re-)selection and allow the corresponding user traffic to be routed to the applications running on MEC host. Based on this information the MEC system should support selection of a MEC host or MEC hosts and the instantiation (and/or relocation) of an application on the selected MEC host or hosts.
3. The MEC system shall be able to collect and expose performance data regarding the virtualisation environment of the MEC host related to MEC services and application (e.g., related to telemetry monitoring such as compute utilisation of the MEC hosts). Such data can aid in intelligent orchestration decisions for the various CONNECT services, taking into account e.g., the workload of MEC hosts and the service requirements for load balancing.
4. The MEC management platform shall support the instantiation, termination, and modification (i.e., life cycle management operations, LCM) of an application on a MEC host when required.
5. The MEC management shall be able to identify which features and MEC services a MEC application requires to run. This allows the MEC system to decide whether and on which MEC host to instantiate the application.
It shall be possible to deploy MEC applications on different MEC hosts in a seamless manner, without a specific adaptation to the application.</td></tr>
<tr><td>Connected To Other Requirements</td><td>Baseline requirements are to be viewed as the de facto MEC features establishing a functional basis over which the rest of the CONNECT (MEC) requirements are expressed.</td></tr>
<tr><td>KPIs</td><td>Introducing KPIs[21] to measure (most of the) baseline requirements may not be straightforward. Those requirements essentially express basic features and capabilities for the operation of the involved technology. As such, we mainly resort to typical network KPIs that are expected to assist us gain indirect evidence of having met the considered requirement. In other cases, we employ indicators capturing performance of lifecycle operations or energy footprint of the CONNECT MEC technology. Particularly:
<table>
<tr><th>Description</th><th>Value</th></tr>
<tr><td>Container LCM Operations [referring to Instantiation, termination, modification time]
Instantiation and termination time of containers.</td><td>Expected value for container start/stop functions is about 30 seconds. In case containers depend on external services (e.g., other containers or services inside or outside the CONNECT system) the start/stop time will be determined accordingly.
One relevant dimension of interest is the comparison of the instantiation/termination time needed with and without secure containers.</td></tr>
</table></td></tr>
</table>

[21] The KPI values for the MEC requirements are derived in line with our prior experimentation experience (mainly through the 5G-LOGINNOV project) and a careful look-up of reference documents (e.g., deliverables of other EU-funded projects). In certain cases where we lack insights (pertaining to the above combination), we only provide best estimates which would be updated with our preliminary experimental results.

	Container LCM Operations [referring to Instantiation, termination, modification time] Modification:	This metric will heavily rely on the type of the modification, e.g., cpu scaling (up/down) operations expected completion time is **about 30 seconds** when the original host avails the needed capacity. A relevant modification instance may involve the task-offloading in the context of the STMD use-case where the relevant task, when offloaded to the MEC, might need container resources of different characteristics (compared to the original in-vehicle containerization).
	Evaluation of **overhead imposed by varying levels of CONNECT TEE** assurances (i.e., different achievable assurance levels are expected to be achieved in line with the available SW/HW TEE combinations) given application specific requirements (e.g., data should reach the vehicle within a given time threshold).	This will be measured via the ratio of CONNECT service resource usage (e.g., compute) for the various controls and features included in the CONNECT TEE configurations. In certain cases, a systematic breakdown of the achievable/induced values associated with the SW/HW features of the CONNECT TEE will constitute a CONNECT 'benchmarking' result (that may become particularly important for future research, especially when safety critical CCAM applications are involved).
	E2E latency: measured round-trip-time (RTT) from the moment the IP ICMP Echo Request packet leaves the source host (e.g., vehicle) until the IP ICMP Echo Reply is received from the destination host (e.g., MEC).	As measured via the 5G-LOGINNOV project (https://5g-loginnov.eu/) which exploits a commercial private (non-standalone) 5G network and edge computing, RTT values averaged **< 20ms**. However, in CONNECT we expect higher values as the RTT will traverse a public network including also Internet delays.
	Energy efficiency: the energy efficiency of the CONNECT MEC system will be measured to showcase the extra power consumed via the varying levels of CONNECT TEE assurances.	These measurements are dependent on the **hardware** used (e.g, rack server hosting the MEC platform) and the application itself, hence we cannot provide an expected value at this stage. Our aim is to carry-out benchmarking measurements on the power consumption of the various security controls and tracing capabilities instantiated as part of the CONNECT TEE configuration.
	CONNECT Application Layer Data-rate [Uplink (UL) and downlink (DL) measurements]	The relevant measured flows would include data from legacy C-ITS messages plus the trustworthiness claims added via CONNECT services. The specific data flows may include either the CONNECT claims piggybacked to the legacy C-ITS messages or sent somewhat asynchronously, triggered by Trustworthiness Assessment Request (issued by an application and referring to the data associated with those C-ITS messages)

		[indicative values[22]: UL: **93.6 Mb/s** tcp, **55.3 Mb/s** udp and DL: **92.9 Mb/s** tcp, **55.1 Mb/s** udp]. Local testbed from ICCS]
ID	**FR.MEC.2 (Optional)**	
Title	**Operational Requirements for mobility support**	
Actors/Components Involved	CONNECT MEC, Vehicle platforms, UEs	
Description	**Background:** As discussed in FR.TR.6, the edge computing technologies enable service providers to operate within a certain own trust domain that supports certain security capabilities; hence adhere to specific trust levels. Moreover, the pivotal role of edge computing becomes apparent when considering the mobility requirements of vehicles, particularly in scenarios involving extensive geographic coverage. Edge computing seamlessly facilitates the transition and handover processes across diverse edge locations. In the context of AVs, it's crucial to note that traditional cloud computing solutions would be inadequate due to the substantial round-trip time (RTT) resulting from considerable distances between servers and user devices. The deployment of edge computing emerges as the optimal solution to ensure the timely and efficient delivery of critical data and services. **Description:** The CONNECT MEC system should be able to maintain connectivity between a UE and an application instance when the UE performs a handover to another cell associated with the same/different MEC host and same/different provider with the levels of trust without breaching the security and privacy requirements. Privacy concerns are already known and, in some cases, accepted when it comes to handover between operators. Apart from the already known implications regarding unlikability and untraceability, the case of information exchange during the service migration in the Digital Twin should be further researched. As described in [170] connectivity may serve a variety of purposes. AVs, for example, gather data about their immediate surroundings, sharing these insights with the MEC. Within the MEC, a specialised service aggregates this data, creating a comprehensive 'Shared World Model' (SWM). Each AV may utilise this SWM to 'see' beyond the limitations of its own onboard sensors, enhancing situational awareness and contributing to the collective intelligence of the CCAM system. Connectivity is also crucial in platooning, a technique where vehicles travel together in convoys, enhancing fuel efficiency and reduced air resistance. The MEC collects and processes real-time data about speed, distance, and braking intentions, creating a 'Platoon Coordination Model' (PCM). This model ensures safe and synchronised manoeuvres, allowing each vehicle to optimise its performance. This communication between platoon members enhances road safety, traffic flow, and fuel efficiency. CONNECT MEC enables this communication, data exchange, and decision-making among the vehicles involved in the platoon.	
Connected To Other Requirements	This mobility support requirement can be expressed provided that (first) the baseline MEC requirements are fulfilled. Then, some relevance can be identified to: FR.TR.6 and FR.OC.6	

[22] measured using a Sierra Wireless LTE Cat. 6 mPCIe module

KPIs	Description	Value
	Mobility interruption time [defined as the time whereby UE cannot exchange user plane packets with any MEC host.	The expected value is **<10 sec** (in accordance to 5G-Mobix project[23] classification of cloud-assisted automated driving use-case). To be measured on a 'theoretical basis', by emulating the interrupt time and comparing it to the relevant use case needs. (e.g., via Linux Traffic Control system that controls the kernel packet scheduler).

ID	FR.MEC.3 (Baseline)	
Title	**Operational Requirements for MEC services**	
Actors/Components Involved	CONNECT MEC	
Description	**Background:** The edge computing services are expected to continuously exchange information with the MEC platform (see IMA and SMTD use case as examples). Relevant mechanisms to ensure the availability and discovery of the services residing at the MEC, the authentication of (their) users as well as the authenticity of the involved messages need to be in-place. **Description:** The CONNECT MEC system should adhere to a number of relevant operational requirements to capture the above needs to cover the access control aspect. These needs are summarised as follows: 1. The MEC platform shall allow authentication and authorization of providers and consumers of MEC services. 2. The MEC platform shall have the capability to provide MEC services that can be consumed by authorised MEC applications. 3. When necessary, the MEC system shall allow operators to dynamically control the access of running MEC applications to certain services. 4. The MEC platform shall allow MEC services to announce their availability. The platform shall allow the discovery of available MEC services. 5. The MEC management shall support the relocation of a MEC application instance from one MEC host to a different host within the system. That would allow (opportunities for) the efficient management, placement, load-balancing of virtualised resources. 6. The MEC system shall be able to move MEC application instances between MEC hosts in order to continue to satisfy the requirements of the MEC application.	
Connected To Other Requirements	Baseline requirements are to be viewed as the de facto MEC features establishing a functional basis over which the rest of the CONNECT (MEC) requirements are expressed. CONNECT requirements that may be relevant are: FR.SR.4, FR.OC.2.	
	Description	**Value**

[23] https://www.5g-mobix.com/assets/files/5G-MOBIX-D2.5-Initial-evaluation-KPIs-and-metrics-V1.4.pdf

KPIs	Authentication application-level latency (Time to create an application secure/authenticated channel)	**tens to a few hundred ms.** This is shaped also by delays induced at the lower levels of the stack such as legacy TLS handshakes under low congestion/normal traffic conditions.
	Discovery Latency (i.e., the time it takes for a client to receive a response to a MEC service discovery request). Note that the MEC service discovery request is a process of finding optimal edge application server endpoints for a client device to connect to.	Expected value is **> 100ms** (under low congestion/normal traffic conditions) which is the average latency of a single HTTP GET request. Note that the response time of this request may vary depending on the network latency and the number of available service endpoints.

ID	**FR.MEC.4 (Mandatory)**
Title	**Operational Requirements for (applications) connectivity**
Actors/Components Involved	CONNECT MEC, 3rd party services
Description	**Background:** Edge computing services, especially in the Secure Mobility Trust Domain (SMTD) domain, enable many applications and services. Notably, these services, that may execute trust, CCAM, and HD data management duties on the edge computing hosts, are not confined to a single location, but are instead distributed across multiple hosts or servers. This dynamic and distributed nature of edge computing highlights the adaptability and versatility of the ecosystem, ensuring the availability and accessibility of essential services to meet the diverse requirements of CCAM. **Description:** The CONNECT MEC system should adhere to a number of relevant operational requirements to capture the above needs of connectivity. These needs are summarised as follows: 1. The MEC system shall support two (or more) instances of a MEC application running on different MEC hosts to communicate with each other. 2. The MEC platform shall be able to allow an authorised MEC application to communicate with third-party servers located in external networks. That is to ensure that external services (e.g., cloud based CCAM services) can offer data/inputs to the MEC applications. For example, information regarding observations from other vehicles or geographically targeted advertising, could be offered by a third party.
Connected To Other Requirements	The MEC baseline requirements need to be fulfilled.

KPIs	**Description**	**Value**
	Inter-containers communication latency: It reflects the	Expected value **<100ms.** (measured as IP ICMP Echo Request packet leaving the source host (e.g., container A, host A) until the IP

	latency requirements between container communications	ICMP Echo Reply is received by the destination host (e.g., container B, host B).

8.1.4.1 MEC Security Requirements

The threats pertaining to the MEC can be common to most of its cases of use and the threat factors can be broadly categorised based on various areas of vulnerabilities related to platform integrity, virtualization and application-programming interfaces (APIs). In brief, there is an extended attack surface (related to the MEC infrastructure, deployed software or even accommodated traffic over the established interfaces) which needs to be protected. As such, it may span from the mobile backhaul inheriting vulnerabilities, the public internet in-coming attacks and clearly any hosted (native or third party) application or RESTful APIs.

This wide security scope goes beyond the CONNECT trust focus as well as the automotive set-up and thus, we eventually limit it to a more 'confined' area yet capable to serve the project purposes. As such, we adopt the assumption that all involved security requirements that are directly relevant to NFV technology (e.g., the involved microservices have been correctly/securely configured) mainly constitute software engineering challenges considered fulfilled. Another important dimension of this assumption relates to the K8s deployment over the MEC host whereby the deployment of an application to K8s clusters would typically include the pulling of images from a Docker (or other public/private) registry as specified in the application's manifest file. Likewise, we consider any relevant registry/repository to be trusted and assume integrity for the downloaded data. The project will not invest effort to comprehensively measure and analyse them.

Furthermore, we do not consider security (or privacy-related) concerns that emerge in case of multiple MEC platforms and inter-MEC communications (mainly due to the CONNECT consortium line-up and the lack of an MNO or edge computing provider to offer such insights). Along these lines, we put forward a set of requirements that cover on the one hand basic MEC functionalities and on the other relate to NFV cyber-threats (for the containers per se) that become relevant in MEC deployments. For the following entries we rely on reference documents and standards [171][167][172][47] and references therein.

ID	**SR.MEC.1 (Mandatory)**
Title	**Security Requirements on MEC service authorization and access (authorised service access)**
Actors/Components Involved	Vehicle, OEM, Mobile Network Operator, Service Provider, Trust Assessment Framework, CONNET TEE Guard, Risk Assessment Engine, Misbehaviour Detection Service
Description	**Background:** The concept of edge computing (see paragraph 2.2.4) relies heavily on the provision of computing functionalities stemming from locations closer to the user. The relevant edge services (to realise that functionality) together with the on-top running applications (which most oftenly provide contextual/use-case related information) constitute the main MEC stack which can be completed by external (third party) services. In this setting, all involved interactions need to meet security requirements that ensure the MEC services authenticity and deployed applications privileges (i.e., authorisation) as well as the confidentiality/integrity of the involved data flows. **Description:** In the context of CONNECT in particular, as reflected especially in the SMTD and IMA use-case) the CONNECT MEC offers a variety of applications and services to receive off-loaded tasks (from the vehicles) or host

	misbehaviour detection computations. All applications deployed in the CONNECT MEC platform need to have been authenticated while the corresponding data exchanges to be subject to standard controls that ensure confidentiality and integrity. More specifically the following requirements apply: 1. The MEC platform should authenticate all MEC application instances and only provide them with the information for which the application is authorised. MEC specifications mandate the use of the OAuth 2.0 for authorization of access to RESTful MEC service APIs defined by ETSI ISG MEC. The implementation of the OAuth 2.0 authorization protocol uses the client credentials grant type according to IETF RFC 6749 and with bearer tokens according to IETF RFC 6750. In orchestration terms, (app) containers (to be deployed) may not be even launched unless it is ensured (e.g. through signature checking) that they realise certified applications. 2. Sensitive data exchanges (i.e., transmissions) between MEC components (across VNFs) should be sufficiently encrypted. This includes secure communications at the transport layer, supporting confidentiality and data integrity of all messages by using e.g., TLS on each interface. On a more evolved note which mainly relates to the features of hosted applications, the integrity requirements of the applications that are to be deployed at the MEC should be aligned with the above encryption primitives adopted for the data exchanges between MEC components. 3. The access to the information regarding MEC service availability and related interfaces shall only be allowed to authenticated and authorised MEC applications. 4. Access to information about each MEC service shall be separately authorised. Separate authorization shall be possible for registering MEC services, and for obtaining information about registered MEC services
Connected To Other Requirements	At the typical end-points of the considered communication (say MEC and vehicle), there need to have mechanisms in-place to ensure the computation of CONNECT guarantees in line with Secure Data Handling and Provenance of SR.4

KPIs	**Description**	**Value**
	Authentication and authorization in a MEC application: In general it can vary widely based on several factors including network latency (typically low especially in a well-optimised 5G network settings and MEC), service processing time and the server load, the complexity of the employed authentications mechanisms (e.g., simple, with basic username and password checks, or complex, with multi-factor authentication, token generation/validation, or interfacing with external authentication providers), and other overheads such as TLS handshakes.	A benchmarking study (to be driven by the project security experts) will constitute an interesting CONNECT contribution to specifying practical (lower) time bounds for the considered latency which will be subsequently assessed (end-to-end) at application level.

	Additionally, the expected values are to be heavily determined by the corresponding operations needed to realise the CONNECT trustworthiness claims. One part of the considered latency is to be attributed to the involved code execution time needed for self-issued verifiable credentials and presentation.	
	Additionally, the expected values are to be heavily determined by the corresponding operations needed to realise the CONNECT trustworthiness claims. One part of the considered latency is to be attributed to the involved code execution time needed for self-issued verifiable credentials and presentation.	A benchmarking study (to be driven by the project security experts) will constitute an interesting CONNECT contribution to specifying practical (lower) time bounds for the considered latency which will be subsequently assessed (end-to-end) at application level. Hence, the proposed KPI will isolate the separate components that contribute to the overall latency and present a holistic measurement study of all involved sub-components. End to end (authentication/authorisation) latency at MEC application level [with expected value in the range of 10 to 200ms plus CONNECT functions overhead].
ID	**SR.MEC.2 (Mandatory)**	
Title	**Security Requirements on virtualization and containerization technology (employed -among others- at the MEC)**	
Actors/Components Involved	Vehicle, OEM, Mobile Network Operator, Service Provider, Trust Assessment Framework, CONNET TEE Guard, Risk Assessment Engine, Misbehaviour Detection Service	
Description	**Background:** As discussed, (see FR.MEC.1) edge computing relies on virtualisation (NFV) technologies to realise the concept of offering computational resources in the vicinity of the end-user. In that sense, the containerised applications (services or even the MEC platform) to be deployed in the MEC platform need to be protected from potential threats in order to ensure their nominal operation. Hence, the appropriate security requirements should be defined in order to protect these services from adversaries. **Description:** Since CONNECT relies heavily on the MEC infrastructure to support the trust assessment as well as other services are explored in the use cases, the security of this infrastructure should be ensured with specific measures. Potential vulnerabilities in virtualization may result in MEC deployed software malfunction affecting every MEC-related CONNECT use case and essentially impacting the whole realisation of the CONNECT concept. Two of the most prominent vulnerabilities with high impact are the following: 1. The MEC system can be susceptible to a number of threats emerging from these virtualization technologies, e.g., possible contamination of shared hardware resources, the noisy neighbour problem, i.e., shared	

<table>
<tr><td></td><td colspan="2">resources might be monopolised by a neighbouring container, abuse of privilege elevation of containers with higher levels of privileges, use of open-source APIs, etc. Vulnerabilities in the MEC virtualization platform can include compromise of the underlying system (FW, Bootloader, Host OS/Hypervisor), inadequate isolation of resources in OS/container layers and vulnerabilities specific to cloud technologies used in MEC implementation.
2. Common Software Environment: if a vulnerability or zero-day exploit was found in software used across multiple virtualized NFs then an attacker might be able to exploit all of these NFs with the same attack. The vulnerability might allow the attacker multiple access points into the MEC network/service, or may allow them to propagate through the network
Therefore, the CONNECT MEC system should provide adequate security against these two attacks.</td></tr>
<tr><td>Connected To Other Requirements</td><td colspan="2">FR.SR.2, FR.SR.3, FR.OC.1, FR.OC.2, FR.OC.3, FR.OC.4, FR.OC.5, FR.OC.6</td></tr>
<tr><td rowspan="2">KPIs</td><td>Description</td><td>Value</td></tr>
<tr><td>Degree of coverage of the (defined) virtualization/containerization threats.</td><td>A number of attack vectors will be defined in D3.2. Those will constitute a basis against which the CONNECT containers security controls will be assessed.
That will enable the evaluation of the proposed KPI (to be carried-out by consortium security experts). Define it in D2.2 on M24.</td></tr>
<tr><td>ID</td><td colspan="2">SR.MEC.3 (Nice to have)</td></tr>
<tr><td>Title</td><td colspan="2">Security Requirements and compatibility with (available/standardised) Application Programming Interfaces (APIs)</td></tr>
<tr><td>Actors/Components Involved</td><td colspan="2">Vehicle, OEM, Mobile Network Operator, Service Provider, Trust Assessment Framework, CONNET TEE Guard, Risk Assessment Engine, Misbehaviour Detection Service</td></tr>
<tr><td>Description</td><td colspan="2">Background: The functionality that can be realised by edge computing applications is heavily based on the available information at the edge platform. This may include information sent by the end-user devices (UEs) or supported IoT devices as well as available network (conditions) and contextual information. The latter piece of information is typically provided (to deployed applications) through MEC standardised interfaces.
APIs facilitate the communication and transmission of data between various system components. These APIs must have strict security mechanisms to secure data, manage access, and mitigate threats. By adding security standards to APIs, a trustworthy ecosystem in which data is protected, is proposed. In addition, such standardisation streamlines development processes, promotes interoperability, and increases scalability of systems, increasing the overall connectivity of the system in a secure and reliable way.
Description: The CONNECT concept (as reflected especially in the SMTD use-case) draws on the task-offloading concept to carefully select and assign non-safety-critical computations to the infrastructure (e.g., MEC) under a zero-trust hypothesis.</td></tr>
</table>

	To realise the relevant offloading decision (see D5.1), information about the network radio conditions, positioning data of the users or (UE) subscription and relevant cellular network (i.e., PC5 interface) configuration data, may be needed. CONNECT applications will in that case resort to standardised interfaces to obtain the information; this implies on the one hand the CONNECT compliance to standardised APIs and on the other hand the need to meet the corresponding security (authentication) requirements of those APIs. Relevant standardised interface may include APIs (Radio Network Information API [ref. MEC-012], Location API [ref, MEC-013], V2X API [ref. MEC-030], etc.) serve as conduits that expose applications for third-party integration; as a consequence of that, also APIs are potentially susceptible to attacks like any other software. From a software development point of view, compliance with the above APIs should be ensured during the interfaces design and implementation phases. For example, OAuth 2.0 based on X.509 client certificates are used for authorization of access to REST MEC service APIs defined by ETSI ISG MEC.
Connected To Other Requirements	The MEC baseline requirements (i.e., FR.MEC.1, FR.MEC.3) need to be fulfilled.

8.1.4 Trusted Computing Base

Since of the main visions of CONNECT is to also be **agnostic** on the type of technologies, and especially Trusted Computing technologies to be employed, in what follows we describe the baseline of characteristics that a Root of Trust (RoT) needs to exhibit in order to able to support the Trusted Computing Base of CONNECT.

NIST[24] defines the Trusted Computing Base (TCB) as "*the totality of protection mechanisms within a computer system, including hardware, firmware, and software, the combination of which is responsible for enforcing a security policy*". For us, it is the part of the system (firmware and software) that can be measured and attested to. In order to achieve this, we need a Root of Trust (RoT), a part of the system that is assumed to be trustworthy and upon which all other results are built. Typically, on bootup of the firmware, the system will initialise the TCB and, using the RoT, measure the TCB's software components and store the results. This provides a **static assessment of the system**. In some devices the measurements continue as the software is executed and this provides extra confidence about the trustworthiness of the system. When required the RoT reports on the measurements, signing them to ensure their authenticity. The final part of this attestation system is a verifier who has access to the expected measurement values for a system in an unaltered (good) state, they can compare these with measured values and report on the results. The attestation evidence from a device [2] [3] is used to provide the trust sources that are used by the TAF to assess the trust level of that device. This evidence can just be the static measurements of the TCB but can also be supplemented with runtime measurements as well. These attestation results are combined with those from other components when generating the harmonised attributes used to report on the integrity of the devices used to generate data being sent outside of the vehicle.

Note: In order to guarantee the security of the TCB we need to keep it as small as possible. There is a trade-off here, as we also wish to ensure the security of the applications that we are running. One compromise is to also attest to these applications, but not to include them in the TCB. If an application is compromised, we can stop it running and either fix the problem or migrate the application to another processor on the system. If the TCB is compromised, then nothing can be guaranteed anymore. Having provided this overview, we now go into more details, particularly regarding the RoT.

8.1.4.1 The Root of Trust

[24] https://csrc.nist.gov/glossary/term/trusted_computing_base

Together with the firmware executed at start-up, the RoT plays an important role in attesting to the state of the system. The RoT can be divided into three subcomponents:

- ✓ **Root of Trust for Measurement (RTM):** This part of the RoT will measure the different software components of the TCB before loading and running them, and where appropriate continue to assess the software as it runs. For the static case these measurements will be hashes of the executables, while for the dynamic case the measurements (execution traces) will vary depending upon the device and method chosen, but will result in attestation evidence that can later be compared to the expected results and verified. Static measurements start with the firmware, then the OS (if there is one) and so on until all of the TCB has been measured.
- ✓ **Root of Trust for Storage (RTS):** For the RoT to be trusted, we need to be sure that the measurements and the keys used to sign the attestation reports cannot be tampered with. We therefore need some non-volatile RAM that can only be accessed by the RoT. This controlled memory can also be used to store other sensitive information, particularly the other keys that will be used to ensure integrity, to ensure confidentiality and for access control (see FR.SR.3).
- ✓ **Root of Trust for Reporting (RTR):** When challenged, this component of the RoT signs and reports on the attestation measurements (either static or runtime, depending on the type of challenge received). To disable the replaying of previous reports a number of mechanisms can be used: the challenger can send a nonce, the RTR can use a monotonic counter, or a timestamp can be used.

We now list the basic functionalities that any RoT in our system should provide to ensure that the protocols that we develop can be implemented and used by any of them.

8.1.4.2 Root of Trust Base Functionalities

The set of basic functionalities that must be provided by the underlying component, acting as a root-of-trust, are the following. These constitute the functional basis that will be leveraged for the design and implementation of the CONNECT attestation variants and calculation of verifiable claims for achieving the previously described functional requirements:

RoT Basic Functionalities

- [**TR.ROT. 1**] It should have a unique identifier (**M**andatory).
- [**TR.ROT. 2**] It should provide a pseudorandom number generator (**M**andatory).
- [**TR.ROT. 3**] It should provide protected storage. This will be used to store essential keys and for a small set of platform configuration registers used to hold the software measurements and traces extracted by the CONNECT Attestation Toolkit. (**M**andatory)

  Note: Ideally, this storage will be in non-volatile Random Access Memory (NVRAM). If no NVRAM is available the data could be stored encrypted in external memory, but this might cause issues if, for example, the power was removed/failed after an update and before writing to external memory. In this case, the data stored locally (in RAM) and externally will be out of sync and this will cause a problem when the system is restarted, and the local copy is restored from external memory. When protected storage is provided by encrypting sensitive data and writing it to external memory great care must be taken to ensure that everything remains in sync, this is a challenge.
- [**TR.ROT. 4**] It should support protocol and algorithm agility for being able to choose among various cryptographic primitives and protocols for providing the necessary CONNECT security services (**D**esirable).
- [**TR.ROT. 5**] It should support enhanced authorization on which process, running on a device, can access the trusted component (**M**andatory).

- [**TR.ROT. 6**] It should support software measurement and secure measurement reporting (Quote) – (**M**andatory);

- [**TR.ROT. 7**] Support remote attestation functionalities and sealing and binding operations (**M**andatory).

TC Performance and Cost Effectiveness

- [**TR.ROT. 8**] It should be feasible to implement the CONNECT attestation trust extensions, using the chosen trusted component, on platform with restricted memory, while providing an acceptable performance (**M**andatory).

- [**TR.ROT. 9**] Selected trusted components should be chosen in such a way so that it is possible to enhance the performance of the CONNECT cryptographic primitives and algorithms (**D**esirable).

- [**TR.ROT.10**] The selected trusted component should be able to support the secure implementation of the CONNECT lightweight crypto algorithms on an identified platform (**M**andatory).

- [**TR.ROT.11**] Integration and testing of the trusted component functionalities, including adequate support for the CONNECT software stack (**D**esirable).

- [**TR.ROT.12**] The use of CONNECT trusted component to attest (during run-time) code snippets, running in an embedded system, should be substantially better than the (hardware-based) state of the art ones (**M**andatory).

- [**TR.ROT.13**] Allow support for some legacy primitives/protocols (**D**esirable).

Implementation

- [**TR.ROT.14**] The implementation of easily extendable software abstractions (e.g., interesting DICE abstraction supporting post-quantum algorithms), may facilitate research in the future (**D**esirable).

- [**TR.ROT.15**] Trying to integrate CONNECT attestation APIs with any type of trusted component (**D**esirable).

Secure Storage

- [**TR.ROT.17**] The (external) memory is protected against rollbacks, i.e. an attacker is unable to replay earlier versions of the memory. (**D**esirable).

8.1.4.3 Examples of Implementations of Roots of Trust

As described above a TCB relies upon a root of trust to measure the software components of the TCB, store those results and report on them when requested. A system''s RoT can be implemented in number of ways; some examples are those using a:

- **Trusted Platform Module.** The trusted platform module (TPM) is a crypto processor that provides tamper resistance and secure storage for data (e.g., measurements) and keys. It is specified by the Trusted Computing Group and is found in a number of guises: a hardware device installed on the systems motherboard, an integrated version, included as part of another chip, a virtual TPM used by virtual machines and a software TPM. In terms of tamper resistance, the hardware device is the strongest, while the software TPM the is the weakest (although a software TPM could be run inside a TEE thus increasing its tamper resistance). As part of the system's initialisation process the software of the TCB is measured and the results stored in the TPM. These can later be reported on as required.

- **Hardware Security Module**. Hardware security modules (HSM) provide storage for keys and measurements. They can be expensive standalone devices, but most relevant here are on chip HSMs which can store keys, store and report on measurements and manage a secure boot process. An example of an on chip HSM is that on the Infineon AURIX microcontroller.
- **Trusted Execution Environment.** A trusted execution environment (TEE) allows data and keys to be stored and binaries executed in a hardware-protected environment. Software-only attacks outside of this protected environment cannot affect the protected keys, data and software. One example is Intel's Software Guard Extensions (Intel SGX) where a so-called enclave is protected against all software of the machine (including operating systems, hypervisor, kernels, …). This can be used to develop the system's TCB.

More details regarding the benefits, limitations and analysis of features for the two types of Roots of Trust (i.e., hardware-based vs software-based) will be discussed in D4.1.

Different types of such secure elements (SEs) and trust anchors can be used at the same time to support the different types of requirements that a component may need to exhibit. For instance, in the context of C-ACC where all the different types of ECUs that comprise a vehicle are considered, it can be the case that a subset of these ECUs have access to a fully-fledged execution environment (i.e., these are usually the zonal controllers and the main V2X OBU), while dedicated sensors, actuators and ECUs (i.e., LIDAR) may have access to a single HSM due to the time restrictions in the decision making process. Hence the second category does not have access to a TEE-rich OS, which adds an overhead. The different implementations of RoT might be present in a vehicle at the same time, hence the importance for a Trust Assessment Framework as CONNECT to enable the interaction with all different types of trusted components. This serves as a guideline to all of our designs on the type of security and attestation controls based on which the trust assessment will operates. As we will see in D4.1 all of the designs need to be agnostic on the type of RoT and will operate based on the type of evidence that different RoT can provide in a verifiable manner.

A key requirement of CONNECT is to allow software-extensibility of the Root of Trust. While this is supported by multiple RoTs, we decided to focus on using Intel Software Guard Extensions (Intel SGX) for the large ECUs and the MEC. In addition, automotive ECUs will support a wide range of RoTs (usually HSMs) that can be assessed using our trust module.

8.1.4.4 Trusted Computing Bases in CONNECT

While the RoT provides the basic capability to protect and attest to running software, one goal is to minimize its functionality to minimize the attack surface. I.e., in most cases, the Root of Trust only provides a few basic hardware/firmware functions that allows running and attesting software. While this is sufficient for basic protection, it is hard to use for developers.

As a consequence, we decided to deploy and extend the following two tools that allow developers to run a wide range of existing software inside the TCB:

Gramine (https://github.com/gramineproject/gramine): To allow developers to quickly migrate linux applications into such an environment, the library-OS Gramine provides a linux-compatible run-time that provides most standard Linux system calls. As a consequence, a developer of a linux application that is to be protected can be recompiled for Gramine and then executed using Gramine and protected within an Intel SGX enclave. The result is that the protected application is then shielded against software attacks from untrusted areas outside the enclave. Gramine has its own mechanisms for attestation independent of the underlying system and is used to provide the TCB for its applications and for runtime measurements of software running outside of the enclave.

Secure Containers https://github.com/confidential-containers/enclave-cc : To simplify deployment and management of applications, an application and all its dependencies can be packaged in a so-called container (e.g. Docker). This ensures a standard interface between the application and the hosting operating system. In CONNECT we plan to deploy secure containers that package a security-critical application inside Gramine that is then packaged inside a container. As a consequence, secure applications inside secure containers can be managed seamlessly along

with ordinary applications. In this scenario, Gramine provides a transparent layer of security while the enclave-cc containers add manageability on top.

8.2 CONNECT Non-Functional Requirements

It is essential to make a distinction between the two primary types of requirements, which are known as functional requirements and non-functional requirements. The functional requirements of a system offer explicit descriptions of how the system should behave and outline the functions that are required for the system to accomplish the goals that it has set for itself. Non-functional requirements, on the other hand, are those that convey the characteristics that influence the behaviour of the system but do not express its specific capabilities. These non-functional criteria could take into account a wide range of factors, including issues related to users' privacy.

8.2.1 Privacy Requirements over the Edge-Cloud CCAM Continuum

Privacy and trust concerns are of utmost importance in the CCAM field, particularly with respect to vehicles and all other road users (specifically vulnerable road users). Standardised protocols, like the vehicular PKI, that function as base for trust in V2X communication have effectively addressed the privacy-respecting identity management of vehicles, incorporating measures such as digital certificates and robust authentication mechanisms. Nevertheless, the issue goes beyond the mere tracing of cars and encompasses the privacy implications arising from a broader set of data and the interactions between all actors (vehicles, roadside users, MEC infrastructure, etc.)

Recently, the Mozilla Foundation researched how 25 major car brands collect and share deeply personal data and published the results as part of their *Privacy Not Included buyer's guide [25]. The research looked on how personal data is being gathered by sensors, microphones, cameras, and the phones and devices drivers connect to their cars, as well as by car apps, company websites, dealerships, and vehicle telematics. They found that brands then share or sell this data to third parties. Car brands can also take much of this data and use it to develop inferences about a driver's intelligence, abilities, characteristics, and preferences, making vehicles 'Privacy Nightmare on Wheels', as the report concludes.

In order to address the complex privacy concerns at hand, it is imperative to establish comprehensive frameworks that encompass many aspects of privacy. Additionally, the development of data anonymization techniques, powerful encryption protocols, and privacy controls to provide unlinkability and untraceability is crucial. In parallel, it is essential to develop the aforementioned schemes without negatively affecting the performance of the system. For example, privacy enhancing technologies like changing pseudonyms may introduce additional challenges to trust management (for example, long-term reputations might not work), hence privacy should be examined from an operational standpoint as well.

ID	FR.PR.1 (Optional)
Title	Unlinkability of Representation Artefacts or Repeated System Interactions
Actors/Components Involved	Operators of communication components and operators of components of trust assessment framework
Description	**Background:** Identifiers and representation artefacts of various kinds are used to identify entities in a system. V2X communications also rely on the usage of identifiers to monitor activities and entities within the network. For instance, the MAC address identifies a vehicle in the physical layer component. Apart from practical issues, these identifiers may be leveraged to unlawfully monitor the movements and activities of vehicles or people; thus breach their privacy. To tackle this issue, the support of non-permanent pseudonyms, used in the place of permanent identifiers, is being adopted. Pseudonyms refer to dynamic

[25] https://foundation.mozilla.org/en/blog/privacy-nightmare-on-wheels-every-car-brand-reviewed-by-mozilla-including-ford-volkswagen-and-toyota-flunks-privacy-test/

ID	FR.PR.1 (Optional)
Title	Unlinkability of Representation Artefacts or Repeated System Interactions
	identities that are periodically modified with the aim of augmenting location privacy. The guidelines for this are delineated in established standards such as ETSI TS 103 941. However, even a changing pseudonym may identify a vehicle in the V2V communication component, if the MAC address is not updated or concealed. Since the objective of a pseudonym change (as in ETSI TS 103 941 clause 6.2.2) is to provide location privacy, it is necessary to timely update the corresponding MAC address when pseudonyms undergo modifications. Data Protection Impact Assessments (DPIAs) are frequently carried out in order to detect and mitigate privacy issues associated with identifiers and representation artefacts inside V2X systems. By conducting DPIAs, enterprises may verify that their privacy measures, such as pseudonymization, adequately protect the privacy of persons and entities in the V2X ecosystem, while still enabling secure and efficient communication. **Description:** The requirement for ITS systems is to prevent linkability of the observable representation artefacts of the distributed trust layer to vehicles. In this context, the TAF should not create additional risks to privacy of persons participating in CCAM systems beyond those that already exist in those systems without TAF. The TAF collects information about trust relationships in the CCAM system and aggregates them in Trust Models. To establish atomic trust opinions, trust sources will also access and analyse CCAM data communicated, e.g., in CAM messages. CONNECT and its Trust Assessment Framework adds and manages new observable representations and identifiers, and their impact needs to be investigated in the context of the overall system and potentially new privacy controls need to be identified and applied. **Remark:** We illustrate this with an example: trust models inherently encode information about certain entities interacting in a CCAM system at certain locations and points in time. In order to mitigate threats of location privacy, ETSI ITS standards foresee the use of changing pseudonyms. This mechanism must not be weakened by maintaining linkability of entities by CONNECT components beyond a pseudonym change. Ensuring this might reduce the performance of the TAF, as long-term node-centric trust values cannot be established. This impact needs to be investigated. ✓ Such representation artefacts can also extend to identify other aspects of a CCAM actor, such as the trust level. All these need to be considered for not breaching the privacy of the vehicle. ✓ This should be targeting all layers of a CCAM service/actor lifecycle and layers.
ID	FR.PR.2 (Optional)
Title	Unlinkability of Data Provenance
Actors/Components Involved	Operators of MEC, and Operators of V2X infrastructure
Description	**Background:** In the context of ensuring the overall trustworthiness of the V2X infrastructure, the collection, and analysis of data serve as crucial evidence for the detection of misbehaviour or normal operation, as discussed in FR.SR.4. However, the intricacies stemming from the heterogeneous and decentralised nature of this environment, along with the variety of trust properties, introduces a significant layer of complexity, particularly concerning trust relationships and the trust domain.

ID	FR.PR.1 (Optional)
Title	Unlinkability of Representation Artefacts or Repeated System Interactions
	Hence, managing provenance information, which essentially involves establishing confidence that data comes from the correct source, becomes a challenge. This challenge is compounded by the need to directly link provenance information to the data used in the trust assessment process. However, this linkage must be executed in a manner that preserves privacy, ensuring that individuals' sensitive information remains safeguarded. In essence, this linkage is employed to instil confidence in the data's authenticity without infringing upon the critical principles of privacy protection. Striking this balance between data integrity and privacy is a nuanced and essential aspect of building trust in the V2X ecosystem. **Description:** CONNECT integrates the best practices established by industry standards such as ITU-T Y.3602 (Big data – Functional requirements for data provenance) and ISO/IEC 5181 (Security and privacy - Data provenance). These standards offer a comprehensive framework outlining the essential elements required for managing data provenance effectively. They provide a clear blueprint for specifying the data flows and prerequisites for provenance information. One of the primary objectives is to achieve data provenance while maintaining unlinkability, ensuring that data sources can be traced without compromising the privacy of individuals or entities. By adhering to these standards, CONNECT not only aligns with industry guidelines but also prioritises privacy preservation through meticulously controlled linkability, particularly in situations where accountability is paramount. For instance, let us consider a scenario within the MEC environment, where a vehicle undergoes a trust assessment, and its claims of trustworthiness fail during the attestation process. In such cases, it becomes imperative for either the TAF or the OEM to establish a traceable link back to the specific vehicle involved. This linkage is instrumental in identifying the ECU whose attestation failed, thus enabling the initiation of targeted measures to address security or operational issues associated with that particular component. By implementing the data provenance principles outlined in industry standards, CONNECT facilitates these critical traceability capabilities while safeguarding the privacy of the individuals and entities involved, thereby promoting both security and privacy within the V2X ecosystem. **Remark:** ISO/IEC 5181 Security and privacy - Data provenance, has just started.
ID	FR.PR.3 (Mandatory)
Title	**Attributes Related to Vehicle Trustworthiness should not Create Privacy Threats with Medium or High Level beyond those Already in Existence in the CCAM Ecosystem**
Actors/Components Involved	Operators of communication components and operators of components of trust assessment framework
Description	**Background:** Apart from the identity and the unlikability of the data to the subject, as discussed in FR.PR.1 and FR.PR.2, it is crucial to make sure that the data that is used as trustworthiness evidence also do not reveal any type of information that could potentially affect the privacy. For example, identifying the vehicle through, for example, vehicle fingerprinting, based on information from trustworthiness evidence, should not be possible. On the other hand, since this information is consumed by the trust assessment, it should be verifiable. Hence, there needs to be an interplay between trust and privacy related to the information extracted by the vehicles. Towards this direction, there is extended research for implications and mitigations of identifiable attributes (i.e., vehicle location, speed, even the attestation result). Short term pseudonyms (i.e., DAA) have been proposed as a means against breaching anonymity, while supporting controlled linkability.

ID	FR.PR.1 (Optional)
Title	Unlinkability of Representation Artefacts or Repeated System Interactions
	Description: Within the CONNECT framework, the process of collecting and analysing data for trust assessments is highly targeted, aiming to extract only the most relevant and meaningful trust-related information. The objective is to obtain data that is essential for establishing trust relationships within the CCAM ecosystem, while avoiding the collection of excessive or unnecessary information. This approach is crucial to prevent potential privacy concerns associated with the identification of the vehicle or its users. It should not be possible for the attacker to profile the vehicle by having access to TAF information (i.e., evidence related to trust sources - output of misbehaviour detection). This information should not help the attacker gain any additional information to help de-anonymise the vehicle. One of the primary considerations is to ensure that the information accessible within the TAF, particularly the evidence derived from trust sources and the output of misbehaviour detection, does not inadvertently assist potential attackers in profiling or de-anonymizing the vehicle. In essence, even if an attacker gains access to TAF-related information, it should not provide them with any additional insights or data that could compromise the anonymity of the vehicle or its users. This level of protection is vital for maintaining the privacy of V2X participants. To fulfil this requirement, while processing high volumes of different information originating from different vehicles, there is a need for *harmonisation in the collection and processing of the multiple trust-related attributes, while maintaining the privacy guarantees.* The harmonisation process, in essence, obfuscates (i.e., either by hiding or by grouping together) the trustworthiness evidence (of a CCAM actor) based on which the trust assessment was performed, enhancing the privacy profile of the vehicle and avoiding additional identification of the vehicle (i.e., fingerprinting) from the exchange of trust related information. Details about the needs and motivation of harmonisation and how CONNECT is approaching this problem can be found in D5.1. Apart from the obfuscation though, the harmonisation includes the secure construction of the necessary data models (i.e., YANG) that can disclose only the necessary trust level of an entity (i.e., vehicle), with the associated proof of ownership of the trust attributes, based on which the trust level was calculated, without revealing any details about the trust sources. Hence, these attributes are carefully managed, to ensure that the information necessary for trust assessments is accessible without compromising the privacy. This balance between acquiring trust-related data and preserving privacy is a critical aspect for the CONNECT framework, aligning with the broader goals of security and privacy protection within the V2X ecosystem. **Remark:** The concept of harmonised attributes is not only crucial within the CONNECT framework but also demands widespread acceptance and implementation by all stakeholders in the CCAM ecosystem, including OEMs. It's imperative that these harmonised attributes, which facilitate the collection of trust-related evidence, are embraced universally to ensure consistent and effective trust assessments while upholding privacy standards. Trustworthiness profiles play a pivotal role, as universally accepted standards, in the calculation of trust levels, encompassing trust models and the types of evidence involved. This harmonisation effort is a fundamental cornerstone in establishing trustworthiness profiles that are recognized and embraced by all relevant stakeholders. It aligns seamlessly with the existing standards and, notably, with the goals of GAIA-X's[26] trust and federated identity management systems. These initiatives aim to define comprehensive, generic trustworthiness profiles that can ultimately contribute to the creation of an expansive international data repository for trustworthiness profiles specifically tailored for the needs of CCAM ecosystems. In

[26] https://gaia-x.eu/

<table>
<tr><td>ID</td><td colspan="2">FR.PR.1 (Optional)</td></tr>
<tr><td>Title</td><td colspan="2">Unlinkability of Representation Artefacts or Repeated System Interactions</td></tr>
<tr><td></td><td colspan="2">essence, harmonisation is a critical step towards developing a globally accepted framework for trust within the realm of connected and autonomous mobility.
This approach not only enhances the security and reliability of V2X systems but also safeguards the privacy rights of all participants, aligning with the overarching goals of trust, security, and privacy in the connected mobility landscape.</td></tr>
<tr><td rowspan="4">KPIs</td><td>Description</td><td>Value</td></tr>
<tr><td>Attributes disclosed as part of the trustworthiness evidence should not enable vehicle fingerprinting.</td><td>TRUE</td></tr>
<tr><td>Controlled linkability should be provided only to authenticated entities (i.e., OEMs) in case of indication of risks of internal vehicle components.</td><td>TRUE</td></tr>
<tr><td>Secure construction of TCs including harmonised attributes</td><td>< 200 ms</td></tr>
<tr><td>ID</td><td colspan="2">FR.PR.4 (Optional)</td></tr>
<tr><td>Title</td><td colspan="2">Privacy Preservation in Multi-MNO Service Domains</td></tr>
<tr><td>Actors/Components Involved</td><td colspan="2">Operators of MEC, and Operators of V2X infrastructure</td></tr>
<tr><td>Description</td><td colspan="2">Background: The introduction of MECs and handovers between MECs should not introduce any privacy threats of medium or high risk. In particular, handovers between MECs should not allow extensive tracking of vehicles throughout the CCAM system.
The implementation of MEC infrastructure and the execution of handovers between MECs consist of important technical progressions within the domain of CCAM, as discussed in FR.MEC.2. Nevertheless, it is crucial to ensure that these technological progressions do not unintentionally reveal privacy-related vulnerabilities, which have the potential to undermine the privacy and security of vehicles and their passengers.
One such problem pertains to the handovers between MEC nodes. When a vehicle transitions from one MEC to another, it is imperative to prevent the occurrence of pervasive vehicle monitoring inside the CCAM system.
Privacy concerns in this context relate to vulnerabilities or situations in which malicious entities or unauthenticated individuals may exploit the MEC infrastructure or handover procedures to acquire sensitive data or track the movement of vehicles. The spectrum of threats incorporates a variety of severity levels, ranging from very insignificant issues to more severe risks that have the potential to lead to privacy violations.
Description: CONNECT examines the topic of migration of tasks as well as task offloading from one MEC origin-to-MEC destination to achieve both the aspect of efficiency as well as the one of privacy-preservation. In addition, CONNECT places a strong emphasis on identity confidentiality and privacy, ensuring that the identities of vehicles and users are safeguarded throughout the CCAM system. This proactive approach to identity management ensures data confidentiality by preventing access to unauthorised parties. With these capabilities, CONNECT not only improves the overall trustworthiness of CCAM operations, but also strengthens privacy protections that are essential to the success of connected mobility.
Remark: MEC and MNOs involve the MNO-level being aware of the identifier of a UE, such as a mobile device or connected vehicle. This knowledge is derived from regular interactions between UEs and the MNO's network infrastructure, where each</td></tr>
</table>

ID	FR.PR.1 (Optional)	
Title	Unlinkability of Representation Artefacts or Repeated System Interactions	
	UE is assigned a unique identifier for network communication and administration. However, the potential for an adversary at the orchestration level to link the MNO's identity with the UE's PKI identity, is a significant concern. In MEC and V2X scenarios, the PKI identity is used to secure communications and establish trust. With the MNO joining the ecosystem, we need to re-evaluate the results of the risk assessment and examine whether additional technical or organisational measures should be taken. 5GAA has already identified this problem and suggests moving away from measures to guarantee that "no single entity" [should be able to track a vehicle] and rather establish rules and policies on the operation level [173]. However, it still remains open to understand whether and how much privacy protection would be reduced by following this approach.	
ID	FR.PR.5 (Mandatory)	
Title	**Trust Information and Assessment Lifecycle Management and User Acceptance**	
Actors/Components Involved	Operators of MEC, and Operators of V2X infrastructure	
Description	**Background:** As aforementioned in Chapter 3, there is a **human dimension of trust**, which is affected also by the **perception of the privacy aspect** (i.e., which is a human not a technical attribute). Obviously, it is crucial to manage, thus, mitigate, threats related to privacy, especially in contexts where sensitive data or personal information is involved. One important aspect of privacy threat mitigation is **ensuring that solutions and measures taken to protect privacy are communicated effectively to the system's users**. The General Data Protection Regulation (GDPR) establishes a shared understanding of privacy standards and offers a structure for fostering trust via the promotion of openness and the adoption of ethical principles in the processing of data. Nevertheless, trust is a complex notion that is impacted by several elements beyond the mere act of disclosing information. Sustaining trust entails more than just adhering to legal obligations; it necessitates the cultivation of a culture that promotes responsible use of data. **Description:** CONNECT considers the interplay between privacy and trust. Information that can identify a person, or potentially identify a person when combine with other information, must be recognised as posing a risk to privacy. Where potential privacy risks have been identified, particular steps need to be taken such that the potential to identify people is removed or reduced. Where that is not possible, then extra protections need to be put in place to limit the people and institutions who can access and use this information - only those with a justified reason to access this private, or potentially private information, and with the technical protections to control access to that information are given access. If private or potentially private information is needed for a particular justified purpose, then steps need to be in place to anonymise and disaggregate that information once the particular purpose has been achieved. CCAM users to have access to communications that can verifiably assure them that privacy is being protected.	
KPIs	**Description**	**Value**
	Provision of empirical studies (based on interviews conducted with OEMs and automotive vendors) on the user acceptance benefit by increasing the perceived level of trust offered by CCAM services.	**TRUE**

ID	FR.PR.1 (Optional)	
Title	Unlinkability of Representation Artefacts or Repeated System Interactions	
	GDPR analysis that the type of (trust-related) information exchanged within CONNECT do not pose additional privacy implications.	**TRUE**

9 Summary and Conclusion

The present deliverable established the foundational operational, security and privacy requirements for the CCAM landscape, serving as the cornerstone for the design of the CONNECT framework. It placed significant emphasis on the two key building blocks of CONNECT: i) *the Trust Assessment Framework (TAF),* responsible for estimating trustworthiness levels for both nodes and data, and ii) the *Trusted Execution Environment (TEE) extensions*, that provide an isolated and secure environment to ensure a trustworthy basis for critical operations such as key generation and trust calculation. These operations are comprehensively addressed on both the vehicle and Multi-Access Edge Computing (MEC) sides, **providing a holistic approach to address CCAM trust assessment challenges, considering observations from diverse and heterogeneous sources**. The provision of the Federated TAF for collecting opinions and observations from multiple sources (e.g., vehicles, RSUs) thus offering a consolidated view of the trust landscape, was further described, and depicted through the architecture. The combination of the aforementioned pillars captures the strictest security requirements, providing chip-to-cloud assurance while covering the verifiability and protection of CCAM and trust-related information from the observation point up to MEC processing, **elevating vehicle-wide trust quantification to CCAM-wide trust evaluation**. In addition to the security and trust-related requirements, CONNECT further considered and introduced the privacy-preserving exchange of trustworthiness claims outside a vehicle, either to another vehicle or to the infrastructure. The outcome of this deliverable is a comprehensive view of the overall CONNECT architecture, its individual components and their interactions, as well as the functional and non-functional requirements of the entire framework.

More details regarding the internal functionalities, interactions and interfaces of the individual components will be analysed in the upcoming deliverables of WP3, WP4 and WP5. More specifically, the Federated TAF will be further explored and reported during M18, while the TAF-DT (Digital Twin) will be developed and thoroughly analysed by M30. Regarding the TEE extensions, a first implementation focusing on the in-vehicle operations will be ready by M18, while the chip-to-cloud notion, will be documented during M30. Lastly, in the final version of this deliverable (i.e., D2.2) an elaborated ethical analysis, considering the user acceptance of such technologies based on the perceived trust will be documented, along with the view of CCAM the threat landscape.

List of Abbreviations

ASD	Active (MEC) Service Directory
AND	Active V2X Node Directory
AV	Autonomous Vehicle
C-ACC	Cooperative Adaptive Cruise Control
CAM	Connected and Automated Mobility
CCAM	Cooperative, Connected and Automated Mobility
CPM	Collective Perception Message
DAA	Direct Anonymous Attestation
DT	Digital Twin
ECU	Electronic Control Unit
GNSS	Global Navigation Satellite System
HSM	Hardware Security Modules
IDS	International dataspace
IMU	Inertial Measurement Unit
ITS	Intelligent Transport System
ITS-Ss	Intelligent Transport System Stations
LoA	Level of Assurance
MD	Misbehaviour Detection
MEC	Multi-access Edge Cloud
MNO	Mobile Network Operator
MR	Misbehaviour Report
NTM	(V2X) Node Trustworthiness Message
NTS	(V2X) Node Trustworthiness assessment Service
OBU	On-Board Unit
OEM	Original Equipment Manufacturer
PKI	Public Key Infrastructure
RoT	Root of Trust
RSU	Roadside Unit
RTK	Real Time Kinematic
SMTD	Slow Moving Traffic Detection
TAF	Trust Assessment Framework
TC	Trustworthiness Claims
TCB	Trusted Computing Base
TCH	Trustworthiness Claims Handler
TCG	Trusted Computing Group
TEE	Trusted Execution Environment
TLEE	Trustworthiness Level Evaluation Engine
TMM	Trust Model Manager
TPM	Trusted Platform Module
TSM	Trust Source Manager
V2I	Vehicle to Infrastructure

V2V	Vehicle to Vehicle
V2X	Vehicle to Everything
VC	Verifiable Claims
VPKI	Vehicular Public Key Infrastructure

References

[1] CAR 2 CAR Communication Consortium. (2019). Guidance for Day 2 and Beyond Roadmap. [Online] https://www.car-2-car.org/fileadmin/documents/General_Documents/C2CCC_WP_2072_RoadmapDay2AndBeyond.pdf

[2] 5GAA. (2022). C-V2X Use Cases and Service Level Requirements — Volume III. [Online] https://5gaa.org/content/uploads/2023/01/5gaa-tr-c-v2x-use-cases-and-service-level-requirements-vol-iii.pdf

[3] ETSI MEC ISG, “Mobile Edge Computing (MEC); Framework and Reference Architecture”, ETSI DGS MEC 003, 2016.

[4] Yoshizawa, T., Singelée, D., Muehlberg, J. T., Delbruel, S., Taherkordi, A., Hughes, D., & Preneel, B. (2023). A survey of security and privacy issues in v2x communication systems. *ACM Computing Surveys*, *55*(9), 1-36.

[5] Brecht, B., Therriault, D., Weimerskirch, A., Whyte, W., Kumar, V., Hehn, T., & Goudy, R. (2018). A security credential management system for V2X communications. *IEEE Transactions on Intelligent Transportation Systems*, *19*(12), 3850-3871.

[6] EU Commission, "Certificate Policy for Deployment and Operation of EU Cooperative Intelligent Transport Systems (C-ITS)", 2018.

[7] China Communications Standards Association (CCSA), “Technical Requirement of Security Certificate Management System for LTE-based Vehicular Communication”, [Online] http://www.ccsa.org.cn

[8] 5GAA Automotive Association,” White Paper: 5GAA Efficient Security Provisioning System", 2020.

[9] Wolf, M., & Gendrullis, T. (2012). Design, implementation, and evaluation of a vehicular hardware security module. In Information Security and Cryptology-ICISC 2011: 14th International Conference, Seoul, Korea, November 30-December 2, 2011. Revised Selected Papers 14 (pp. 302-318). Springer Berlin Heidelberg.

[10] Brickell, E., Camenisch, J., & Chen, L. (2004, October). Direct anonymous attestation. In Proceedings of the 11th ACM conference on Computer and communications security (pp. 132-145).

[11] Whitefield, J., Chen, L., Giannetsos, T., Schneider, S., & Treharne, H. (2017, November). Privacy-enhanced capabilities for vanets using direct anonymous attestation. In 2017 IEEE Vehicular Networking Conference (VNC) (pp. 123-130). IEEE.

[12] Larsen, B., Giannetsos, T., Krontiris, I., & Goldman, K. (2021, June). Direct anonymous attestation on the road: Efficient and privacy-preserving revocation in c-its. In Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks (pp. 48-59).

[13] Angelogianni, A., Krontiris, I., & Giannetsos, T. (2023, April). Comparative Evaluation of PKI and DAA-based Architectures for V2X Communication Security. In 2023 IEEE Vehicular Networking Conference (VNC) (pp. 199-206). IEEE.

[14] Kurdi, H., Alshayban, B., Altoaimy, L., & Alsalamah, S. (2018). TrustyFeer: A subjective logic trust model for smart city peer-to-peer federated clouds. Wireless Communications and Mobile Computing, 2018.

[15] Garlichs, K., Willecke, A., Wegner, M., & Wolf, L. C. (2019, October). TriP: Misbehavior detection for dynamic platoons using trust. In 2019 IEEE Intelligent Transportation Systems Conference (ITSC) (pp. 455-460). IEEE.

[16] Cheng, M., Yin, C., Zhang, J., Nazarian, S., Deshmukh, J., & Bogdan, P. (2021, May). A general trust framework for multi-agent systems. In Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (pp. 332-340).

[17] Shen, C., Zhang, H., Wang, H., Wang, J., Zhao, B., Yan, F., ... & Xu, M. (2010). Research on trusted computing and its development. Science China Information Sciences, 53, 405-433.

[18] Jøsang, A. Subjective Logic 2016 Cham Springer 10.1007.

[19] European Telecommunications Standards Institute (ETSI). (2014). ETSI GS NFV-SEC 003 V1.1.1 (2014-12), Network Functions Virtualisation (NFV); NFV Security; Security and Trust Guidance, Whitepaper.

[20] Gambetta, D. (1988). Trust: Making and breaking cooperative relations.

[21] van der Heijden, R. W., Dietzel, S., Leinmüller, T., & Kargl, F. (2018). Survey on misbehavior detection in cooperative intelligent transportation systems. IEEE Communications Surveys & Tutorials, 21(1), 779-811.

[22] Haidar, F., Kamel, J., Jemaa, I. B., Kaiser, A., Lonc, B., & Urien, P. (2020, May). Dare: a reports dataset for global misbehavior authority evaluation in c-its. In 2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring) (pp. 1-6). IEEE.

[23] Schmidt, R. K., Leinmüller, T., Schoch, E., Held, A., & Schäfer, G. (2008, June). Vehicle behavior analysis to enhance security in vanets. In Proceedings of the 4th IEEE Vehicle-to-Vehicle Communications Workshop (V2VCOM2008).

[24] Pouyan, A. A., & Alimohammadi, M. (2014). Sybil attack detection in vehicular networks. Computer Science and Information Technology, 2(4), 197-202.

[25] Nguyen, V. L., Lin, P. C., & Hwang, R. H. (2020). Enhancing misbehavior detection in 5G vehicle-to-vehicle communications. IEEE Transactions on Vehicular Technology, 69(9), 9417-9430.

[26] Bißmeyer, N., Mauthofer, S., Bayarou, K. M., & Kargl, F. (2012, November). Assessment of node trustworthiness in vanets using data plausibility checks with particle filters. In 2012 IEEE Vehicular Networking Conference (VNC) (pp. 78-85). IEEE.

[27] Allig, C., Leinmüller, T., Mittal, P., & Wanielik, G. (2019, December). Trustworthiness estimation of entities within collective perception. In 2019 IEEE Vehicular Networking Conference (VNC) (pp. 1-8). IEEE.

[28] Raya, M., Papadimitratos, P., Gligor, V. D., & Hubaux, J. P. (2008, April). On data-centric trust establishment in ephemeral ad hoc networks. In IEEE INFOCOM 2008-The 27th Conference on Computer Communications (pp. 1238-1246). IEEE.
[29] Müller, J., Meuser, T., Steinmetz, R., & Buchholz, M. (2019, July). A trust management and misbehaviour detection mechanism for multi-agent systems and its application to intelligent transportation systems. In 2019 IEEE 15th International Conference on Control and Automation (ICCA) (pp. 325-331). IEEE.
[30] Zhang, J., Jemaa, I. B., & Nashashibi, F. (2022, December). Trust Management Framework for Misbehavior Detection in Collective Perception Services. In 2022 17th International Conference on Control, Automation, Robotics and Vision (ICARCV) (pp. 596-603). IEEE.
[31] S. Gyawali and Y. Qian, "Misbehavior detection using machine learning in vehicular communication networks," in ICC 2019 - 2019 IEEE International Conference on Communications (ICC), 2019, pp. 1–6.
[32] Ercan, S., Ayaida, M., & Messai, N. (2021). Misbehavior detection for position falsification attacks in VANETs using machine learning. IEEE Access, 10, 1893-1904.
[33] Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. Information fusion, 14(1), 28-44.
[34] Varghese, B., Wang, N., Barbhuiya, S., Kilpatrick, P., & Nikolopoulos, D. S. (2016, November). Challenges and opportunities in edge computing. In 2016 IEEE international conference on smart cloud (SmartCloud) (pp. 20-26). IEEE.
[35] Cruz, P., Achir, N., & Viana, A. C. (2022). On the edge of the deployment: A survey on multi-access edge computing. ACM Computing Surveys, 55(5), 1-34.
[36] Tun, Y. K., Alsenwi, M., Pandey, S. R., Zaw, C. W., & Hong, C. S. (2019, September). Energy efficient multi-tenant resource slicing in virtualized multi-access edge computing. In 2019 20th Asia-Pacific Network Operations and Management Symposium (APNOMS) (pp. 1-4). IEEE.
[37] Nguyen, T. D., Nguyen, V., Pham, V. N., Huynh, L. N., Hossain, M. D., & Huh, E. N. (2020). Modeling data redundancy and cost-aware task allocation in MEC-enabled Internet-of-Vehicles applications. IEEE Internet of Things Journal, 8(3), 1687-1701.
[38] Jheng, Y. S., Wu, M. L., Yang, T. W., Chou, C. F., & Chang, C. (2021, October). A Systematic Resource Management for VR Streaming on MECs. In 2021 30th Wireless and Optical Communications Conference (WOCC) (pp. 36-37). IEEE.
[39] Xu, X., Liu, Z., Bilal, M., Vimal, S., & Song, H. (2022). Computation offloading and service caching for intelligent transportation systems with digital twin. IEEE Transactions on Intelligent Transportation Systems, 23(11), 20757-20772.
[40] Wu, Y., Zhang, K., & Zhang, Y. (2021). Digital twin networks: A survey. IEEE Internet of Things Journal, 8(18), 13789-13804.
[41] Wang, S., Tuor, T., Salonidis, T., Leung, K. K., Makaya, C., He, T., & Chan, K. (2018, April). When edge meets learning: Adaptive control for resource-constrained distributed machine learning. In IEEE INFOCOM 2018-IEEE conference on computer communications (pp. 63-71). IEEE.
[42] D. Chen, C. S. Hong, Y. Zha, Y. Zhang, X. Liu and Z. Han, "FedSVRG Based Communication Efficient Scheme for Federated Learning in MEC Networks," in IEEE Transactions on Vehicular Technology, vol. 70, no. 7, pp. 7300-7304, July 2021.
[43] Ranaweera, P., Jurcut, A. D., & Liyanage, M. (2021). Survey on multi-access edge computing security and privacy. IEEE Communications Surveys & Tutorials, 23(2), 1078-1124.
[44] Lal, S., Taleb, T., & Dutta, A. (2017). NFV: Security threats and best practices. IEEE Communications Magazine, 55(8), 211-217.
[45] Jia, X., He, D., Kumar, N., & Choo, K. K. R. (2019). A provably secure and efficient identity-based anonymous authentication scheme for mobile edge computing. IEEE Systems Journal, 14(1), 560-571.
[46] Gai, K., Wu, Y., Zhu, L., Xu, L., & Zhang, Y. (2019). Permissioned blockchain and edge computing empowered privacy-preserving smart grid networks. IEEE Internet of Things Journal, 6(5), 7992-8004.
[47] 5GAA Automotive Association. (2021). MEC for Automotive in Multi-Operator Scenarios: Technical Report.
[48] Podium Project [URL] https://podium-project.eu/
[49] Kung, A., Baudoin, C., & Tobich, K. (2022). Report of TWG Digital Twins: Landscape of Digital Twins (1.0). Zenodo. https://doi.org/10.5281/zenodo.6556917
[50] ISO/IEC. (2023). ISO/IEC 30186 ED1: Digital twin – Maturity model and guidance for a maturity assessment. [Online] https://www.iec.ch/dyn/www/f?p=103:38:316658826824133::::FSP_ORG_ID,FSP_APEX_PAGE,FSP_PROJECT_ID:20486,23,111892
[51] Stark, R., & Damerau, T. (2019). Digital Twin. In Springer Berlin Heidelberg (Ed.), Berlin, Heidelberg: 1–8. http://dx.doi.org/10.1007/978-3-642-35950-7_16870-1.
[52] Guo, J., Bilal, M., Qiu, Y., Qian, C., Xu, X., & Choo, K. K. R. (2022). Survey on digital twins for Internet of Vehicles: Fundamentals, challenges, and opportunities. Digital Communications and Networks.
[53] Y. Dai and Y. Zhang, "Adaptive Digital Twin for Vehicular Edge Computing and Networks," in Journal of Communications and Information Networks, vol. 7, no. 1, pp. 48-59, March 2022, doi: 10.23919/JCIN.2022.9745481.
[54] Dobre, C., Shin, H., & Duong, T. Q. (2022). URLLC edge networks with joint optimal user association, task offloading and resource allocation: A digital twin approach. IEEE Transactions on Communications, 70(11), 7669–7682.
[55] Sun, W., Zhang, H., Wang, R., & Zhang, Y. (2020). Reducing offloading latency for digital twin edge networks in 6G. IEEE Transactions on Vehicular Technology, 69(10), 12240–12251.

[56] Dai, Y., Zhang, K., Maharjan, S., & Zhang, Y. (2021). Deep reinforcement learning for stochastic computation offloading in digital twin networks. IEEE Transactions on Industrial Informatics, 17(7), 4968–4977.

[57] Chen, Y., Zhao, F., Chen, X., & Wu, Y. (2022, May). Efficient multi-vehicle task offloading for mobile edge computing in 6G networks. IEEE Transactions on Vehicular Technology, 71(5), 4584–4595.

[58] Cloudflight, "Learnings from the Digital Twins Data Architecture of Tesla". [Online] https://www.cloudflight.io/en/blog/learnings-from-the-digital-twins-data-architecture-of-tesla/
[59] 5GCarmen Project. [Online] https://5gcarmen.eu/

[60] Wang, Z., Han, K., & Tiwari, P. (2022, June). Digital Twin-Assisted Cooperative Driving at Non-Signalized Intersections. IEEE Transactions on Intelligent Vehicles, 7(2), 198-209. doi: 10.1109/TIV.2021.3100465.

[61] He, C., Luan, T. H., Lu, R., Su, Z., & Dong, M. (Year, Month). Security and Privacy in Vehicular Digital Twin Networks: Challenges and Solutions. IEEE Wireless Communications. doi: 10.1109/MWC.002.2200015.

[62] Mulligan, D. P., Petri, G., Spinale, N., Stockwell, G., & Vincent, H. J. M. (Year). Confidential Computing—a brave new world. In 2021 International Symposium on Secure and Private Execution Environment Design (SEED), (pp. 132-138). Washington, DC, USA. doi: 10.1109/SEED51797.2021.00025.

[63] Baier, V. E., March, J. G., & Saetren, H. (1986). Implementation and ambiguity. Scandinavian journal of management studies, 2(3-4), 197-212.
[64] Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. Human factors, 46(1), 50-80.
[65] O'Hara, K. (2012). A general definition of trust. University of Southampton. [Online] https://eprints.soton.ac.uk/341800/
[66] McLeod, C. (2021). Trust. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy (Fall 2021). Metaphysics Research Lab, Stanford University. [Online] https://plato.stanford.edu/archives/fall2021/entriesrust/
[67] Hardin, R. (2002). Trust and trustworthiness. Russell Sage Foundation.
[68] Kohn, S. C., de Visser, E. J., Wiese, E., Lee, Y. C., & Shaw, T. H. (2021). Measurement of trust in automation: A narrative review and reference guide. Frontiers in psychology, 12, 604977.

[69] ISO. (n.d.). ISO/IEC TS 5723:2022—Trustworthiness—Vocabulary. Retrieved May 26, 2023, from https://www.iso.org/standard/81608.html

[70] ISO. (n.d.). ISO/IEC 22624:2020(en), Information technology—Cloud computing—Taxonomy based data handling for cloud services. Retrieved June 7, 2023, from https://www.iso.org/obp/ui#iso:std:iso-iec:22624:ed-1:v1:en:term:3.2

[71] International Telecommunication Union (ITU-T). (n.d.). Y.3057: A trust index model for information and communication technology infrastructures and services. Retrieved May 26, 2023, from https://www.itu.int/rec/T-REC-Y.3057-202112-P

[72] Fernandez, L. D., & Gomez, G. E. Trustworthy Autonomous Vehicles. EUR 30942 EN, Publications Office of the European Union, Luxembourg, 2021, ISBN 978-92-76-46055-8, doi:10.2760/120385, JRC127051.
[73] Adnan, N., Nordin, S. M., bin Bahruddin, M. A., & Ali, M. (2018). How trust can drive forward the user acceptance to the technology? In-vehicle technology for autonomous vehicle. Transportation research part A: policy and practice, 118, 819-836.
[74] Zhang, T., Tao, D., Qu, X., Zhang, X., Lin, R., & Zhang, W. (2019). The roles of initial trust and perceived risk in public's acceptance of automated vehicles. Transportation research part C: emerging technologies, 98, 207-220.
[75] Choi, J. K., & Ji, Y. G. (2015). Investigating the importance of trust on adopting an autonomous vehicle. International Journal of Human-Computer Interaction, 31(10), 692-702.
[76] Kenesei, Z., Ásványi, K., Kökény, L., Jászberényi, M., Miskolczi, M., Gyulavári, T., & Syahrivar, J. (2022). Trust and perceived risk: How different manifestations affect the adoption of autonomous vehicles. Transportation research part A: policy and practice, 164, 379-393.
[77] Benleulmi, A. Z., & Blecker, T. (2017). Investigating the factors influencing the acceptance of fully autonomous cars. In Digitalization in Supply Chain Management and Logistics: Smart and Digital Solutions for an Industry 4.0 Environment. Proceedings of the Hamburg International Conference of Logistics (HICL), Vol. 23 (pp. 99-115). Berlin: epubli GmbH.
[78] Herrenkind, B., Brendel, A. B., Nastjuk, I., Greve, M., & Kolbe, L. M. (2019). Investigating end-user acceptance of autonomous electric buses to accelerate diffusion. Transportation Research Part D: Transport and Environment, 74, 255-276.
[79] Eiser, J. R., Miles, S., & Frewer, L. J. (2002). Trust, perceived risk, and attitudes toward food technologies 1. Journal of applied social psychology, 32(11), 2423-2433.
[80] Liu, P., Xu, Z., & Zhao, X. (2019). Road tests of self-driving vehicles: Affective and cognitive pathways in acceptance formation. Transportation research part A: policy and practice, 124, 354-369.
[81] Kaye, S. A., Somoray, K., Rodwell, D., & Lewis, I. (2021). Users' acceptance of private automated vehicles: A systematic review and meta-analysis. Journal of safety research, 79, 352-367.
[82] Milford, S. R., Elger, B. S., & Shaw, D. M. (2023). Believe me! Why tesla's recent alleged malfunction further highlights the need for transparent dialogue. Frontiers in Future Transportation, 4, 1137469.

[83] ISO/IEC/IEEE. (2022). ISO/IEC/IEEE 42010:2022 Software, systems and enterprise — Architecture description. https://www.iso.org/standard/74393.html

[84] ISO/IEC JTC 1/SC41. (2023). Best practices and guidelines for Reference Architectures (RA). ISO/IEC JTC 1/SC41 N 2306.

[85] ISO/IEC. (n.d.). ISO/IEC 30141 ED2 Internet of Things (IoT) - Reference architecture. https://www.iec.ch/dyn/www/f?p=103:38:412808078329758::::FSP_ORG_ID,FSP_APEX_PAGE,FSP_PROJECT_ID:20486,23,104064#

[86] ISO/IEC. (n.d.). ISO/IEC 30188 Digital Twin - Reference architecture, under development. https://www.iec.ch/dyn/www/f?p=103:38:713337145614036::::FSP_ORG_ID,FSP_APEX_PAGE,FSP_PROJECT_ID:20486,23,104896

[87] 5GAA Automotive Association. (2021). MEC for Automotive in Multi-Operator Scenarios.

[88] ETSI (European Telecommunications Standards Institute). (2017). Network Functions Virtualisation (NFV); Trust; Report on Attestation Technologies and Practices for Secure Deployments.

[89] ETSI (European Telecommunications Standards Institute). (2023). ETSI TS 103 759, "Intelligent Transport Systems (ITS); Security; Misbehaviour Reporting service; Release 2".

[90] ETSI (European Telecommunications Standards Institute). (2010). Intelligent Transport Systems (ITS); Communication Architecture. ETSI EN 302 665 V1.1.1.

[91] Hobert, L., Festag, A., Llatser, I., Altomare, L., Visintainer, F., & Kovacs, A. (2015). Enhancements of V2X communication in support of cooperative autonomous driving. IEEE communications magazine, 53(12), 64-70.
[92] Xia, N., & Yang, C. S. (2017). Vehicular communications: Standards and challenges. In Internet of Vehicles. Technologies and Services for Smart Cities: 4th International Conference, IOV 2017, Kanazawa, Japan, November 22-25, 2017, Proceedings 4 (pp. 1-12). Springer International Publishing.
[93] Abdel Hakeem, S. A., Hady, A. A., & Kim, H. (2020). 5G-V2X: Standardization, architecture, use cases, network-slicing, and edge-computing. Wireless Networks, 26, 6015-6041.
[94] Kamel, J., Jemaa, I. B., Kaiser, A., Cantat, L., & Urien, P. (2019, December). Misbehavior Detection in C-ITS: A comparative approach of local detection mechanisms. In 2019 IEEE Vehicular Networking Conference (VNC) (pp. 1-8). IEEE.
[95] Kamel, J., Ansari, M. R., Petit, J., Kaiser, A., Jemaa, I. B., & Urien, P. (2020). Simulation framework for misbehavior detection in vehicular networks. IEEE transactions on vehicular technology, 69(6), 6631-6643.

[96] ETSI (European Telecommunications Standards Institute). (2021). ETSI TS 102 940 V2.1.1 (2021-07): Intelligent Transport Systems (ITS); Security; ITS communications security architecture and security management; Release 2.

[97] Pei, Y., Biswas, S., Fussell, D. S., & Pingali, K. (2019). An elementary introduction to Kalman filtering. Communications of the ACM, 62(11), 122-133.

[98] ETSI (European Telecommunications Standards Institute). (2019). ETSI TS 102 941 V1.3.1 (2019-02): Intelligent Transport Systems (ITS); Security; Trust and Privacy Management.

[99] ETSI (European Telecommunications Standards Institute). (2021). ETSI TS 103 097 V2.1.1 (2021-10): Intelligent Transport Systems (ITS); Security; Security header and certificate formats; Release 2.

[100] ETSI (European Telecommunications Standards Institute). ETSI TS 103 324 Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Collective Perception Service; Release 2, 2023.

[101] European Telecommunication Standards Institute. (2019). ETSI EN 302 637-3 V1.3.1: Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specifications of Decentralised Environmental Notification Basic Service [European Standard].

[102] Institute of Electrical and Electronics Engineers (IEEE), International Organization for Standardization (ISO), International Electrotechnical Commission (IEC). (2011). IEEE/ISO/IEC 29148-2011, "Systems and software engineering - Life cycle processes - Requirements engineering".
[103] Beck, K., Beedle, M., Van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., ... & Thomas, D. (2001). The agile manifesto.
[104] Pandey, D., Suman, U., & Ramani, A. K. (2010, October). An effective requirement engineering process model for software development and requirements management. In 2010 International Conference on Advances in Recent Technologies in Communication and Computing (pp. 287-291). IEEE.
[105] Cohn, M. (2004). User stories applied: For agile software development. Addison-Wesley Professional.
[106] The CONNECT Consortium, “D3.1 – Architectural Specification of CONNECT Trust Assessment Framework, Operation and Interaction”, July 2023.
[107] The CONNECT Consortium, “D5.1 – Distributed Processing and CCAM Trust Functions Offloading & Data Space Modelling”, November 2023.
[108] H. B. Debes and T. Giannetsos, "Segregating Keys from noncense: Timely Exfil of Ephemeral Keys from Embedded Systems," *2021 17th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, Pafos, Cyprus, 2021, pp. 92-101, doi: 10.1109/DCOSS52077.2021.00029.
[109] The CONNECT Consortium, “D2.2 – Operational Landscape, Requirements and Reference Architecture – Final Version”, August 2024.
[110] The CONNECT Consortium, “D4.2 – Virtualization- and Edge-based Security and Trust Extensions (First Release)”, February 2024.

[111] The CONNECT Consortium, “D4.1 – Conceptual Architecture of Customizable TEE and Attestation Models Specifications”, December 2023.
[112] The CONNECT Consortium, “D3.2 – CONNECT Trust & Risk Assessment and CAD Twinning Framework – Initial Version”, February 2024.

[113] Confidential Containers. (n.d.). Enclave-CC: Design Documentation. GitHub. https://github.com/confidential-containers/enclave-cc/blob/main/docs/design.md

[114] Confidential Containers. (n.d.). Enclave-CC. GitHub. https://github.com/confidential-containers/enclave-cc
[115] Feiri, M. (2016). Scalable Broadcast Authentication for V2V Communication. [PhD Thesis - Research UT, graduation UT, University of Twente]. Twente University Press (TUP). https://doi.org/10.3990/1.9789036542005.

[116] Prisco, R., & Yung, M. (Eds.). (2006). Security and Cryptography for Networks: 5th International Conference, SCN 2006, Maiori, Italy, September 6-8, 2006. Proceedings. Springer-Verlag Berlin Heidelberg.
[117] Chen, L., & Urian, R. (2015). DAA-A: Direct anonymous attestation with attributes. In Trust and Trustworthy Computing: 8th International Conference, TRUST 2015, August 24-26, 2015, Proceedings 8 (pp. 228-245).
[118] Heini Bergsson Debes and Thanassis Giannetsos. 2022. ZEKRO: Zero-Knowledge Proof of Integrity Conformance. In Proceedings of the 17th International Conference on Availability, Reliability and Security (ARES '22). Association for Computing Machinery, New York, NY, USA, Article 35, 1–10
[119] Bernstein, D. (2014). Containers and cloud: From lxc to docker to kubernetes. IEEE cloud computing, 1(3), 81-84.

[120] Confidential Containers. enclave-cc. GitHub. https://github.com/confidential-containers/enclave-cc

[121] Public Key Infrastructure Consortium. PKI Consortium. https://pkic.org/

[122] International Organization for Standardization (ISO). (2021). ISO 23247-1:2021 Automation systems and integration — Digital twin framework for manufacturing — Part 1: Overview and general principles. https://www.iso.org/standard/75066.html

[123] International Organization for Standardization (ISO). (2021). ISO 23247-2:2021 Automation systems and integration — Digital twin framework for manufacturing — Part 2: Reference architecture. https://www.iso.org/standard/78743.html

[124] International Organization for Standardization (ISO). (2021). ISO 23247-3:2021 Automation systems and integration — Digital twin framework for manufacturing — Part 3: Digital representation of manufacturing elements. https://www.iso.org/standard/78744.html

[125] International Organization for Standardization (ISO). (2021). ISO 23247-4:2021 Automation systems and integration — Digital twin framework for manufacturing — Part 4: Information exchange. https://www.iso.org/standard/78745.html

[126] International Organization for Standardization (ISO), International Electrotechnical Commission (IEC), Institute of Electrical and Electronics Engineers (IEEE). (2022). ISO/IEC/IEEE 42010:2022 Software, systems and enterprise — Architecture description. https://www.iso.org/standard/74393.html

[127] Roth, G. A., Abate, D., Abate, K. H., Abay, S. M., Abbafati, C., Abbasi, N., ... & Borschmann, R. (2018). Global, regional, and national age-sex-specific mortality for 282 causes of death in 195 countries and territories, 1980–2017: a systematic analysis for the Global Burden of Disease Study 2017. The Lancet, 392(10159), 1736-1788.
[128] Watzenig, D., & Horn, M. (Eds.). (2016). Automated driving: safer and more efficient future driving. Springer.
[129] CAR 2 CAR Communication Consortium. (2019). Guidance for Day 2 and beyond Roadmap. https://www.car-2-car.org/fileadmin/documents/General_Documents/C2CCC_WP_2072_RoadmapDay2AndBeyond.pdf
[130] Soto, I., Calderon, M., Amador, O., & Urueña, M. (2022). A survey on road safety and traffic efficiency vehicular applications based on C-V2X technologies. Vehicular Communications, 33, 100428.

[131] ETSI MEC specifications. (n.d.). https://www.etsi.org/committee/1425-mec

[132] Milford, S. R., Elger, B. S., & Shaw, D. M. (2023). Believe me! Why tesla's recent alleged malfunction further highlights the need for transparent dialogue. Frontiers in Future Transportation, 4, 1137469.
[133] Amanullah, M. A., Loke, S. W., Baruwal Chhetri, M., & Doss, R. (2023). A Taxonomy and Analysis of Misbehaviour Detection in Cooperative Intelligent Transport Systems: A Systematic Review. ACM Computing Surveys, 56(1), 1-38.
[134] 5GAA Automotive Association, C-V2X Use Cases and Service Level Requirements Volume I (2020-10)[https://5gaa.org/content/uploads/2020/12/5GAA_T-200111_TR_C-V2X_Use_Cases_and_Service_Level_Requirements_Vol_I-V3.pdf]
[135] TSI (European Telecommunications Standards Institute). (2018). ETSI TS 101 539-2 V1.1.1 (2018-06) Intelligent Transport Systems (ITS); V2X Applications; Part 2: Intersection Collision Risk Warning (ICRW) application requirements specification [https://www.etsi.org/deliver/etsi_ts/101500_101599/10153902/01.01.01_60/ts_10153902v010101p.pdf]
[136] CAR 2 CAR Communication Consortium. (2021). For Day 2 and Beyond Roadmap. [https://www.car-2-car.org/fileadmin/documents/General_Documents/C2CCC_WP_2072_RoadmapDay2AndBeyond_V1.2.pdf]

[137] ETSI (European Telecommunications Standards Institute). (2022). ETSI TS 103 900 V2.0.0 (2022-07): Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Specification of Cooperative Awareness Basic Service; Release 2.

[138] 5GAA Automotive Association. (2022). Misbehaviour Detection.

[139] van der Heijden, R. W., Dietzel, S., Leinmüller, T., & Kargl, F. (2018). Survey on misbehavior detection in cooperative intelligent transportation systems. IEEE Communications Surveys & Tutorials, 21(1), 779-811.
[140] Asselin-Miller, N., Biedka, M., Gibson, G., Kirsch, F., Hill, N., White, B., & Uddin, K. (2016). Study on the deployment of C-ITS in Europe: Final Report. Report for DG MOVE MOVE/C, 3, 2014-794.
[141] Schrank, D., Eisele, B., & Lomax, T. (2019) Urban Mobility Report, Texas A&M Transp. Inst., Texas A&M Univ., College Station, USA, Rep., Aug. 2019.

[142] Milanés, V., Shladover, S. E., Spring, J., Nowakowski, C., Kawazoe, H., & Nakamura, M. (2013). Cooperative adaptive cruise control in real traffic situations. IEEE Transactions on intelligent transportation systems, 15(1), 296-305.
[143] Wang, Z., Wu, G., & Barth, M. J. (2018, November). A review on cooperative adaptive cruise control (CACC) systems: Architectures, controls, and applications. In 2018 21st International Conference on Intelligent Transportation Systems (ITSC) (pp. 2884-2891). IEEE.

[144] ETSI (European Telecommunications Standards Institute). (2019). ETSI TR 103 299 V2.1.1 (2019-06): Intelligent Transport Systems (ITS); Cooperative Adaptive Cruise Control (CACC); Pre-standardization study.

[145] Maul, M., Becker, G., & Bernhard, U. (2018). Service-oriented EE zone architecture key elements for new market segments. ATZelektronik worldwide, 13(1), 36-41.

[146] International Organization for Standardization (ISO). (17987). ISO 17987 Road vehicles — Local Interconnect Network (LIN).

[147] International Organization for Standardization (ISO). (11898). ISO 11898 Road vehicles — Controller area network (CAN).

[148] International Organization for Standardization (ISO). (17458). ISO 17458 Road vehicles — FlexRay communications system.

[149] Institute of Electrical and Electronics Engineers (IEEE). (2018). IEEE Standard 802.1AE-2018 - IEEE Standard for Local and metropolitan area networks-Media Access Control (MAC) Security.

[150] Trkulja, N., Hermann, A., Petrovska, A., Kiening, A., Ferraz de Lucena, A. R., & Kargl, F. (2023). In-vehicle trust assessment framework.
[151] European Telecommunication Standards Institute (ETSI). (2019). ETSI TR 103 562: Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Analysis of the Collective Perception Service (CPS); Release 2.
[152] European Telecommunication Standards Institute (ETSI). (2019). ETSI EN 302 637-2 V1.4.1: Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service.
[153] International Telecommunication Union (ITU-T). (2002). ITU-T X.680: Series X: data networks and open system communications; OSI networking and system aspects – Abstract Syntax Notation One (ASN.1); Information technology – Abstract Syntax Notation One (ASN.1): Specification of basic notation.
[154] European Telecommunication Standards Institute (ETSI). (2014). ETSI EN 302 895: Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Local Dynamic Map (LDM).
[155] Institute of Electrical and Electronics Engineers Standards Association (IEEE). (2022). IEEE Standard for Wireless Access in Vehicular Environments – Security Services for Application and Management Messages.
[156] European Telecommunication Standards Institute (ETSI). (2020). ETSI TR 103 460 V2.1.1: Intelligent Transport Systems (ITS); Security; Pre-standardization study on Misbehaviour Detection; Release 2.
[157] European Telecommunications Standards Institute,"Network Functions Virtualisation (NFV); Trust; Report on Attestation Technologies and Practices for Secure Deployments ", ETSI GR NFV-SEC 007, 2017. [Available] https://www.etsi.org/deliver/etsi_gr/NFV-SEC/001_099/007/01.01.01_60/gr_nfv-sec007v010101p.pdf
[158] Benjamin Larsen, Thanassis Giannetsos, Ioannis Krontiris, and Kenneth Goldman. 2021. Direct anonymous attestation on the road: efficient and privacy-preserving revocation in C-ITS. In Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '21).
[159] Ghosal, A., & Conti, M. (2020). Security issues and challenges in V2X: A survey. Computer Networks, 169, 107093.
[160] Elkhail, A. A., Refat, R. U. D., Habre, R., Hafeez, A., Bacha, A., & Malik, H. (2021). Vehicle security: A survey of security issues and vulnerabilities, malware attacks and defenses. IEEE Access, 9, 162401-162437.
[161] Jangid, M., & Lin, Z. (2022, January). Towards a tee-based v2v protocol for connected and autonomous vehicles. In Workshop on Automotive and Autonomous Vehicle Security (AutoSec).
[162] Yoshizawa, T., Singelée, D., Muehlberg, J. T., Delbruel, S., Taherkordi, A., Hughes, D., & Preneel, B. (2023). A survey of security and privacy issues in v2x communication systems. ACM Computing Surveys, 55(9), 1-36.
[163] Kil, C., Sezer, E. C., Azab, A. M., Ning, P., & Zhang, X. (2009, June). Remote attestation to dynamic system properties: Towards providing complete system integrity evidence. In 2009 IEEE/IFIP International Conference on Dependable Systems & Networks (pp. 115-124). IEEE.

[164] European Telecommunications Standards Institute (ETSI). (2022). ETSI GS MEC 009 V3.2.1 (2022-07) Multi-access Edge Computing (MEC); General principles, patterns and common aspects of MEC Service APIs.

[165] European Telecommunications Standards Institute (ETSI). (2022). ETSI GS MEC 003 V3.1.1 (2022-03) Multi-access Edge Computing (MEC); Framework and Reference Architecture.

[166] European Telecommunications Standards Institute (ETSI). (2022). ETSI GS MEC 030 V2.2.1 (2022-05) Multi-access Edge Computing (MEC); V2X Information Service API.

[167] Sabella, D., Reznik, A., Nayak, K. R., Lopez, D., Li, F., Kleber, U., Leadbeater, A., Maloor, K., Baskaran, S. B. M., Cominardi, L., & Cristina C. (2022). MEC security; Status of standards support and future evolutions, 2nd edition – September 2022. ETSI White Paper No. #46.

[168] European Telecommunications Standards Institute (ETSI). (2022). ETSI GS MEC 003 V3.1.1 (2022-03) Multi-access Edge Computing (MEC); Framework and Reference Architecture.

[169] Han, B., Gopalakrishnan, V., Ji, L., & Lee, S. (2015). Network function virtualization: Challenges and opportunities for innovations. IEEE communications magazine, 53(2), 90-97.
[170] Eijnden, E.K. van den (2020) Optimal Handover in MEC for an Automotive Application.

[171] European Union Agency for Cybersecurity (ENISA). (2020). ENISA Threat Landscape for 5G Networks Report.

[172] 3rd Generation Partnership Project (3GPP). (2018). 3GPP 33.848 Study on security impacts of virtualisation.

[173] 5GAA Automotive Association. (2019). 5GAA Efficient Security Provisioning System. White Paper.