CONNECT

CCAM TRUST & RESILIENCE

D2.2:Operational Landscape, Requirements and Reference Architecture - Final Version

Project number:	101069688
Project acronym:	**CONNECT**
Project title:	Continuous and Efficient Cooperative Trust Management for Resilient CCAM
Project Start Date:	1st September, 2022
Duration:	36 months
Programme:	HORIZON-CL5-2021-D6-01-04
Deliverable Type:	R – Report
Reference Number:	D6-01-04 / D2.2 / 1.00 December 20, 2024
Workpackage:	WP 2
Due Date:	August 31, 2024
Actual Submission Date:	December 20, 2024
Responsible Organisation:	INTEL
Editor:	Matthias Schunter
Dissemination Level:	PU - Public
Revision:	1.00 December 20, 2024
Abstract:	The *CONNECT* project addresses security, privacy, and trust in Cooperative, Connected, and Automated Mobility (CCAM) systems. D2.2 extends the overall architecture outlined in D2.1 and focuses on ethical analysis, the final architecture, detailed threat modelling and requirements, and the definition of the *CONNECT* Minimal Viable Platform. By prototyping this architecture, *CONNECT* will ensure secure collaboration across vehicles, MEC, and cloud, providing a foundation for privacy-preserving, trustworthy mobility solutions.
Keywords:	Architecture Specification, Functional Components, Interfaces & APIs, Requirements Analysis, Trust Assessment, Use Cases, User Stories, MVP TEE, MEC, CCAM

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
Matthias Schunter(INTEL)

Contributors (ordered according to beneficiary numbers)

Anna Angelogianni, Nikolaos Fotos, Thanassis Giannetsos, Stefanos Vasileiadis (UBITECH)
Ana Petrovska, Ioannis Krontiris, Koffi Ismael Ouattara, Theo Dimitrakos (HUAWEI)
Vasilis Milionis, Vangelis Kosmatos. Pavlos Basaras, Panagiotis Pantazopoulos (ICCS)
Benjamin Erb, Artur Hermann, Frank Kargl, Nataša Trkulja (UULM)
Antonio Kung, Guillaume Mockly (TRIALOG)
Anderson Ferraz de Lucena, Alexander Kiening (DENSO)
Dmitrii Kuvaiskii, Sergej Schumilo, Matthias Schunter (INTEL)
Konstantinos Latanis (SUITE5)
Ilias Aliferis (UNISYSTEMS)
Adam Henschke (UTWENTE)
Marco Rapelli, Claudio Casetti, Guido Marchetto, Carla Fabiana Chiasserini (POLITO)
Francesca Bassi, Ines Ben Jemma (IRTSX)
Christopher Newton (SURREY)

Executive Summary

The growing adoption of Cooperative, Connected, and Automated Mobility (CCAM) systems is transforming the transportation landscape, promising enhanced safety, efficiency, and environmental sustainability. These advancements, driven by technologies such as autonomous driving, perception sharing, and coordinated driving, aim to revolutionize mobility. However, with the increasing complexity and interconnectivity of CCAM ecosystems, addressing security, privacy, and trustworthiness becomes imperative to ensure their reliable and ethical operation.

The *CONNECT* project tackles these challenges by introducing innovative security and trustworthiness assessment mechanisms. Its overarching goal is to establish a framework that facilitates secure data exchanges, dynamic trust evaluations, and reliable decision-making within CCAM ecosystems. By leveraging trusted computing anchors and focusing on the Multi-access Edge Computing (MEC) and Edge domains, CONNECT prioritizes dynamic trust assessment to ensure security and privacy across the Edge-Cloud continuum.

This deliverable builds upon the work performed during the second year of the project and outlines the final view of the CONNECT architecture. It delves into the functional and non-functional requirements essential for security and trust within CCAM systems, emphasizing dynamic trust modeling tailored to the unique needs of each CCAM service. Additionally, the deliverable incorporates:

- **Ethical Analysis**: Evaluating trust management activities to address potential biases and align with ethical standards, particularly in relation to Trustworthy AI.
- **Threat Modeling**: Using the STRIDE model to identify key threats and enhance secure and privacy-preserving operations within CCAM ecosystems.
- **Technical and Use Case Requirements**: Consolidating requirements and key performance indicators (KPIs) and mapping them to specific use cases to prioritize their importance.

The CONNECT framework establishes a foundation for verifiable, evidence-based trust assessment, enabling secure collaboration among vehicles, MEC, and cloud components. Moreover, it addresses broader ethical considerations, ensuring that trust assessment mechanisms align with societal and regulatory expectations. This deliverable provides a comprehensive roadmap for implementing a robust security and trust framework within CCAM ecosystems. By integrating threat modeling, ethical considerations, and trust evaluation, CONNECT is poised to set new benchmarks for secure, privacy-preserving, and trustworthy mobility solutions.

Contents

List of Figures

List of Tables

Chapter 1

Introduction and Overview

The future of transportation is rapidly changing with the emergence of *Connected, Co-operative and Automated Mobility (CCAM)* systems. These technologies offer innovative solutions to challenges and improve commuter experiences through autonomous driving, perception sharing, path planning, real-time local updates, and coordinated driving. The core of this shift is improving road safety and traffic control, while reducing transportation times and fuel expenses. CCAM technologies have the potential to revolutionize passenger mobility and their interaction with the environment. However, as CCAM systems become more integrated into everyday life, it is crucial to prioritize reliability, security, and privacy. Given the diverse range of stakeholders in the ecosystem, such as sensors, vehicles, and traffic conditions, the attack surface is broad and multifaceted. Security and trustworthiness are key properties of CCAM systems, and it is essential to develop effective protection and trust evaluation methods before their widespread usage.

The CONNECT project aims to facilitate the development and adoption of CCAM ecosystems by introducing novel security and trustworthiness assessment mechanisms. Through CONNECT, participating entities can define a trust model and assess dynamically the trust relationships, establishing trust for cooperatively executing safety-critical decisions. It will provide secure data exchange between data sources in the CCAM ecosystem that lack or have inadequate pre-existing trust relationships, and enable reliable outsourcing of tasks to the Multi-access Edge Computing (MEC) and the cloud. CONNECT extends its scope beyond traditional safety requirements to cover the overall security and trustworthiness of CCAM ecosystems by focusing on trust assessment. By enhancing the overall level of trust in the system, participants in the ecosystem can be trusted, and the data they exchange is secure and trustworthy.

1.1 Scope and Purpose

The deliverable builds on top of D2.1 [4] and the developments after the first release of the framework to outline the final view of the *CONNECT* architecture, elaborating on the functional components, and their interconnections, providing in-depth insights into the CONNECT framework's design. It further elaborates on the functional (i.e., technical) and non-functional requirements, prioritizing security and safety in CCAM ecosystems further considering the privacy aspects across the Edge-Cloud CCAM Continuum.

Trust is a key focus, as it is a crucial element in any CCAM ecosystem. Trust models play a critical role in capturing the security and privacy needs of each CCAM service. The present deliverable

places emphasis on the Multi-access Edge Computing (MEC) and Edge domains, elaborating on the the security extensions offered by *CONNECT* leveraging the trusted computing anchors. The insights and characteristics presented serve as a reference for the CONNECT project, leading to the dynamic assessment of trust within complex and ever-evolving CCAM ecosystems. With this deliverable, CONNECT establishes the ground for the development of a complete framework that incorporates security, privacy, and trustworthiness.

Moreover, the deliverable integrates an ethical analysis of trust management activities, specifically regarding the processes of trust decision-making and enforcement. This analysis evaluates potential biases and the broader ethical implications of these mechanisms, ensuring alignment with ethical standards, particularly in light of ongoing standardization efforts in Trustworthy AI. The ethical dimension ensures that the framework not only secures CCAM systems but also upholds ethical principles in its operation.

The deliverable further incorporates detailed threat modeling for CCAM ecosystems, identifying the key threat categories leveraging the STRIDE model. By addressing these risks, the framework aims to enhance not only security but also the ethical handling of sensitive data and decision-making processes in CCAM. The combination of threat modeling and trust management ensures a comprehensive approach to maintaining secure and privacy-preserving operations in next-generation mobility systems.

1.2 Deliverable Structure

The remainder of this deliverable is structured as follows: **Chapter 2** provides the ethical analysis of a trust assessment framework such as the one proposed by *CONNECT* for the *CCAM* landscape. **Chapter 3** illustrates the final view of the *CONNECT* architecture, providing the necessary details on the key components that enable the verifiable, evidence-based, trust assessment across the Vehicle-*Multi-access Edge Computing (MEC)*-Cloud continuum. **Chapter 4** delves into the *CCAM* threat landscape considered by *CONNECT* in terms of the trust assessment framework designs. **Chapter 5** presents the consolidated technical requirements and KPIs of *CONNECT* while **Chapter 6** maps the technical requirements per use case (UC) providing a categorisation of High (H), Medium (M) or Low (L) in terms of importance for the given UC. Finally, **Chapter 7** draws the conclusions.

Chapter 2

Ethical Analysis in Trust Management Activities

2.1 Towards Ethical and Legal analysis of Trust Management in CCAM

Trust is an essential element of both conventional driving and driving within the context of CCAM (Cooperative, Connected, and Automated Mobility). It is also an essential factor in the intricate decision-making processes of road systems. Undoubtedly, **trust in driving systems is a multifaceted concept that includes several factors extending beyond technical reliability**. For a vehicle to be considered trustworthy, it must **operate as expected manner; thus ensure safety**. In the context of connected, autonomous vehicles, **trustworthiness also requires the protection of personal data, secure communication with other vehicles, and resilience against cyber-attacks**. This expands the challenge for CCAM systems, which must integrate additional **ethical considerations—such as privacy, information security, and cybersecurity**—alongside traditional values like safety. Addressing these challenges demands a clear definition of trust and trustworthiness, especially in the CCAM domain. This involves both conceptual and ethical analyses to identify the properties/values that make a CCAM system worthy of trust. Analytic philosophy provides tools to clarify and articulate these concepts, ensuring they are well-defined and practically applicable. For instance, defining trust in ways that can guide the design of technologies and be communicated effectively to stakeholders is essential for creating systems that inspire confidence. Ethical analysis complements this conceptual work by addressing why trust and related properties/value matter. Ethics moves beyond merely describing the world, to envisioning how it should be. For example, while acknowledging the frequency of vehicle crashes is a factual observation, ethical reasoning provides the rationale for reducing crashes—emphasizing the moral imperative to avoid harm, save lives, and minimize costs. Similarly, in CCAM, identifying trust-enhancing features requires not just understanding their function but explaining why they are desirable and how they align with societal values.

In *CONNECT* **the goal is to combine ethical principles, expert insights, and technical considerations to create a framework that makes *CCAM* systems trustworthy**. This requires moving beyond vague claims about improving trust, to developing a comprehensive methodology for assessing trustworthiness. *CONNECT* aims to define trust explicitly, link it to ethical values, and provide practical steps for ensuring that CCAM systems meet these criteria. This chapter tackles three key challenges: first, defining trust in a clear and actionable manner; second, link-

ing trust to ethical values; and third, providing a practical procedure for evaluating and ensuring trustworthy CCAM systems. Towards this direction the present chapter is structured as follows: Section 2.2 sets out the general methodology that *CONNECT* has developed to give a definition for trust and trustworthiness. Section 2.3 details what *CONNECT* has done so far, and gives an assessment list for trustworthy CCAM, to allow an evaluation of trust and trustworthiness that engages with and includes ethical values. Section 2.4 sets out steps of stakeholder engagement, to validate the approach being taken. Section 2.5 then provides suggestions for future steps towards a recommendation of a trust assessment.

2.2 *CONNECT* Methodology Towards Assessment of Trustworthiness and Trust

As detailed in Section 2.3.2, *CONNECT* is developing the Assessment List for Trustworthy CCAM Components and Case Examples (ALTCCE), a structured methodology for evaluating trust and trustworthiness in CCAM systems. The ALTCCE comprises four main steps, each broken down into sub-steps with specific questions and processes designed to guide the assessment systematically.

Step 1: Specify trust and trustworthiness

This involves the specification of trust and trustworthiness. In order to do this, *CONNECT* has developed a general definition of trust and trustworthiness.

General Definition of Trust: TA,B – A relies on B to exhibit R(X) in C to a degree D with a warrant W

General Definition of Trustworthiness: TwB,A – A has W that they can rely on B to exhibit behaviour R(X) in Context C to Degree D

Step 2: Mapping Trust Relationships

These general definitions of trust and trustworthiness can be used to specify the trust relations and what the trustee is expected to do. To map the relations, one needs to identify the following:

1. A: Who is the trustor?
2. B: Who is the trustee?
3. R(X): What behaviour is A expecting B to do?
4. C: What is the context of the trust relationship?
5. D: To what degree is B being trusted?
6. W: What is the warrant, such as reasons or evidence, that A has for trusting B?

Following these steps, a specific trust relationship, and a description of the trustworthiness can be established. The specific components (A, B, R(X), C, D, and W) are explained in Section 2.3.1.

Step 3: Setting Trustworthiness Standards Through Assessment List for Trustworthy CCAM

This details the particular trust relations in relation to seven key ethical values. In order to asses the trustworthiness of a CCAM component, process, or system, a user must assess if particular values are relevant to the particular component, process, or system being developed. Moreover, trustworthiness must be considered across a range of different ethical values. The seven key values are drawn from the EU High Level Ethics Group (HLEG) on AI [28] (High Level Expert Group on Artificial Intelligence 2019), and are as follows:

1. Human Agency and Oversight;
2. Technical Robustness and Safety;
3. Privacy and Data Governance;
4. Transparency;
5. Diversity, Non-discrimination and Fairness;
6. Societal and Environmental Well-being;
7. Accountability.

Section 2.3.2 explains the origins of these values, places them in the context of CCAM trust assessments, and provides specific CCAM relevant questions for each one. We note here that not all values will be relevant for every single component, process, or CCAM system. As such it is possible that particular values will not be part of the assessment list. For instance there are particular CCAM components and processes that do not involve the generation, collection, communication, storage, or use of personally identifiable information. As such, it is unlikely that privacy would be included in the final assessment. However, as part of this method, each of these values (and the questions given in Section 2.3.2) need to be checked to see if they are relevant.

Step 4: Validating and Ensuring the achievement of set trustworthiness standards

Having clarified the trust and trustworthiness relations, mapped the relations, and set the standards required, the final parts of the *CONNECT* framework involve validation and ensuring the trustworthiness standards. That is, Step 3 identifies what conditions CCAM need to meet in order to be worthy of trust, and Step 4 is ensures that those standards are actually met. The two sub-steps of validation and ensurance should occur in a dynamic mode, each guiding the other.

Validation engages key stakeholders to review the ALTCCAM, to see how particular components, processes, and systems relate to the seven key areas, and if specific questions from the ALTCCAM need to be adapted or updated to adequately guide the assessment. At the same time as validation, specific measures must be developed and applied to ensure that a component, process, or system actually meets the standards set by Step 3 and validated by Step 4.

Validation, detailed in Section 2.4, requires stakeholder groups to be identified, and stakeholders to be engaged to work through the ALTCCAM. They are to consider each of the seven key requirements, and to update or extend the questions as needed. Following the HLEG, the ALTCCAM is intended to be a dynamic and evolving process. Moreover, as the technologies and applications of CCAM themselves develop, the seven key requirements will need to be revisited and updated.

Ensuring that the standards have been met requires a range of different skills. For instance, *CONNECT*ś Trust Assessment Framework (TAF) is designed to assess the Actual

Trust Level (ATL) of specific CCAM features; particularly related to security. This features are evaluated leveraging the three use case scenarios: i) Intersection Movement Assist (IMA), ii) Cooperative Adaptive Cruise Control (C-ACC), and iii) Slow Moving vehicle Traffic Detection (SMTD). But to ensure that a component, process, or system is worthy of trust, that ATL must meet or surpass the Required Trust Level, (RTL). The RTL is derived from the seven key requirements of trustworthy CCAM. The stakeholders involved in the validation will be essential to developing how the seven key requirements are converted or translated into RTL, such that a meaningful comparison between ATL and RTL can be made.

2.3 Towards Consolidating Trustworthy CCAM Ecosystems

CONNECT sought to clarify, sharpen, and communicate a general definition of trust that could be used, applied, and adapted to specific instances where trust is required. We sought to produce an account of trust and trustworthiness that can then be applied autonomous and information handling systems like CCAM in ways that are practically useful. The context of the system would demand that the account of trust and trustworthiness followed here is applicable to a diverse set of relationships – for example, relationships involving two users, one user and a part of the autonomous system, and two technological parts of the autonomous or information handling system (or two components).

2.3.1 Trust and Trustworthiness

Clarifying Trust
At a broad general level, trust may be thought of as a species of reliance or reliability, such that to say A trusts B is to say that A relies on B (which in turn implies that A expects something of B).

A trusts B (TA,B) – A expects something of B (1)

A more specific account of trust may view this as three-part relation, such that trust involves reliance on an entity for a goal/outcome X.

A trusts B (TA,B) – A relies on B for a goal X [1]. (2)

Another way to frame this “reliance” is to say that A expects B to perform the goal X. Here, the goal X is, at the minimum known prior to trust formation/establishment (for example, B has already stated intentions to perform/achieve X) and is ideally valued/desired by A. The latter feature, that is A’s valuing X, may be important to distinguish trust from mere reliance (and show why trust is a special species of reliance) – if B states an intention to perform a goal X, which A does not value, A may still come to rely on B (or have expectations that B would do X) without trusting B. In some cases, A’s trust (or expectations from) B may depend not only on what the outcome (X) of B’s actions (performance) is, but also on how this outcome was achieved (process), and why this goal (or outcome) was chosen (purpose). These 3 Ps (performance, purpose, process) are argued to be the basis of trust in automation by [25]. To better capture these three dimensions as the basis of trust, we can write,

TA,B – A relies on B to exhibit behaviour R(X) (3)

where R needs to conform to A's expectations regarding the outcome (performance), procedure (or process), and purpose. The 3 Ps are captured under the umbrella term 'behaviour'. The expectations that A has can be predictive and/or normative (where A has the attitude that B should exhibit R(X) along such lines). Further, since trust can be context-dependent, the above formulation can be modified to represent trust as a quadruple relation [27]:

TA,B – A relies on B to exhibit R(X) in C (4)

where C is some context C in which B exhibits the behaviour R(X) Trust may not always be binary, and rather, be a matter of degree. To reflect this, (4) can be modified to:

TA,B – A relies on B to exhibit R(X) in C to a degree D (5)

Further, A may trust B based on some warrant (or evidence) that A may have had prior to (and as a reason for) trust formation.

TA,B – A relies on B to exhibit R(X) in C to a degree D with a warrant W (6)

We 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 by an agent Z about B's ability (and willingness) to exhibit R(X) in context C made available to A (referral or transitive trust)
- Trust Anchors (for example, legal regulations that incentivize B to exhibit R(X) or disincentivize/prohibit B to deviate from exhibiting R(X))

One may also add a temporal dimension here, such that *(6)* is applicable only in or at time T. There are at least two ways to conceive this temporal dimension:

- The warrant W is only applicable at time T
- The context C, for which there is the warrant W, is only stable at time T

So, following these steps, we have produced in *(6)* a General Definition of Trust (GDoT), which is

"TA,B – A relies on B to exhibit R(X) in C to a degree D with a warrant W"

The GDoT has been designed to be generally applicable for a range of uses and purposes. When developing this for specific uses, it will need to be adapted. For instance, in *CONNECT* we are using the more specific definition of trust as

"A Trusts B implies that A has expectations that B will have the property of being Trustworthy"

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. 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 (see: [4]). Following this, *CONNECT* considered the conceptual and practical relations between trust and trustworthiness.

Clarifying Trustworthiness

In general, while trust depends on the expectations (or the attitude) that A has towards B, trustworthiness is a property of B. More specifically, trustworthiness is a measure of a trustee's ability (and willingness) to exhibit some behaviour R(X) that may be relied upon by the trustor. The analysis of trust offered in the previous section, where one can define trust from a general to a specific level, can also be followed for trustworthiness. For example, on a broad general level, an entity B is trustworthy for an entity A if B is worthy of trust for A, that is if A can rely on B (or in other words, if A's expectations from B are warranted). Following this, *(1)* can be modified for trustworthiness in the following way:

> *B is trustworthy for A (TwB,A)) – if A can rely on B (or B can meet A's expectations) (7)*

Similarly, at a specific level, where trust and trustworthiness are conceived as quadruple relations,

> *TwB,A – A can rely on B to exhibit behaviour R(X) in Context C (8)*

In an ideal scenario, the Warrant W available to A to trust B completely guarantees TwB,A in *(8)*. That is, the evidence A has regarding B's ability (and willingness) to R(X) in C leaves no room for uncertainty and perfectly aligns with B's actual ability (and willingness) to R(X) in C. One can also think of it as a situation where there is no gap between actual trustworthiness and perceived trustworthiness (which is based on the warrant W). This ideal scenario would also imply that TwB,A is measurable/observable either to A or to some other agent Z who can make its assessment available to A. In the case where an independent agent Z is involved in this assessment on behalf of, or for, A, it is important that this assessment takes into account A's expectations/requirements since trustworthiness is subjective (i.e. here one is interested in whether B is trustworthy for A, and not if B is trustworthy in general or for an unspecified agent.)

The relationship between TA,B, TwB,A, and W can be captured in at least three ways:

1. *TA,B aligns with TwB,A because of W – that is trust between A and B is well calibrated as B is actually worthy of trust for A and A has sufficient evidence W to evaluate/know this.*
2. *TA,B and TB,A are misaligned such that while TB,A is true (and possibly backed up by W), A fails to trust B (β error)*
3. *TA,B and TB,A are misaligned such that while TB,A is false (and possibly backed up by W), A still trusts B (α error)*

Drawing from the GDoT, and bearing in mind the issues of whether A is warranted in trusting B, we offer a general definition of trustworthiness, given here as GDoTW.

TwB,A – A has W that they can rely on B to exhibit behaviour R(X) in Context C to Degree D (9)

As with the GDoT, the GDoTW is concerned with a trust relation between A and B, in which A has warrant to rely on B to exhibit particular behaviour in a specific context to a particular degree. The GDoTW repeats the same information as the GDoT, in that they both consider and use relations between A and B, A's expectations about B's capacity to reliably exhibit behavour R(X) in a given context to a given degree. However, the GDoTW is now focused on whether B is worthy of that trust. For instance, a driver might trust their vehicle to be safe, but it may not be able to navigate an intersection without significant risk of crashing. A may trust B, but B is not worthy of that trust. By reorganising the GDoT into the GDoTW, the emphasis is shifted onto whether B is worthy of A's trust; what matters in the GDoTW is whether B is worthy of A's trust, rather than whether A trusts B or not.

Similar to the above account, the ISO/IEC 5723 [18] defines trustworthiness as the "ability to meet stakeholders' expectations in a verifiable way". Further, [18] provides some "characteristics" of trustworthiness. From the perspective of the account laid out here, such as in *(8)*, these characteristics can be conceived as dimensions/properties of R(X), and along which the trustworthiness of B can be assessed. These characteristics are: Accountability, accuracy, authenticity, availability, controllability, integrity, privacy, quality, reliability, resilience, robustness, safety, security, transparency, and usability.

In summary, we have produced the GDoT, providing clarity on what is involved in a trust relation at a general level. We then took the GDoT and restructured it as the GDoTW, such that the focus was on whether the trustee is worthy of trust. This work on the basic conceptualisation of trust and trustworthiness is intended to produce general definitions for the concepts. In other *CONNECT* deliverables, we have then adapted the general definitions for specific applications. Having completed the conceptualisation – answering the question of what we mean when talking about trust, giving a clear, sharply defined, and explicable account of trust. The next step is to show how trust in CCAM relates to ethical values, discussed in Section 2.3.2.

2.3.2 Assessment List for Trustworthy CCAM

On the account discussed above, trustworthiness is a function of some further value. It is not enough to say that B is worthy of trust, or that a CCAM component, process, or action is worthy of trust. We must be clear that B or the CCAM component, process, or action is worthy of trust with regard to some further value. To say that a CCAM vehicle is trustworthy is shorthand to say that it is safe, or informationally secure, because it can navigate an intersection or engage in adaptive cruise control without collision, or can handle system critical information securely and recognise if such information has been compromised. The Trust Assessment Framework (TAF) being developed in *CONNECT* is a tool by which evidence is produced to give warrant for a judgment on the trustworthiness of the component, process, or action.

The challenge here is that, with such complicated systems as CCAM, by looking only at the technical features that give the foundation for whether A is warranted in trusting B, we have an incredibly complicated and interdependent way of discussing trustworthiness. These technical features are essential, but part of the purpose of *CONNECT* is to provide general and system level ways of understanding and assessing trustworthiness. To this end, in parallel with the development of the TAF, we have also developed methods for assessing trustworthiness that start

with the values that make CCAM worthy of trust. In short, to answer the question of the ethical values that make a CCAM vehicle worthy of trust we have drawn from a range of values to give an overall judgment of whether a connected vehicle is worthy of trust.

Our approach here has been to draw from work produced by the EU High Level Ethics Group (HLEG) on AI [28] and The Assessment List for Trustworthy Artificial Intelligence (ALTAI), following the work of David Fernández Llorca and Emilia Gómez [26]. This approach was taken as the HLEG on AI has conducted an analysis of 350+ stakeholders involved in AI and trust. Our approach has been to draw from the HLEG on AI's work, particularly as it informed the ALTAI, and adapt this to *CONNECT* and cyber-physical systems like CCAM. Like the GDoT and the GDoTW, we have sought to draw from the general and apply that to the specific conditions of *CONNECT* and to CCAM.

The HLEG for AI identified seven key requirements that must be met for a system to be considered trustworthy. As with our GDoTW, the emphasis here is on (a) what makes a system worthy of trust, and (b) the generalisable values – considered by the HLEG for AI as key requirements – that explain why a system should be trusted. In line with the description of ethics given earlier, this is an ethics-based approach, in that the seven key requirements provide reasons why we should consider if a system should be trusted. The seven key requirements given by the HELG for AI are:

1. **Human agency and oversight**: AI systems should empower human beings, allowing them to make informed decisions and fostering their fundamental rights. At the same time, proper oversight mechanisms need to be ensured, which can be achieved through human-in-the-loop, human-on-the-loop, and human-in-command approaches

2. **Technical Robustness and safety**: AI systems need to be resilient and secure. They need to be safe, ensuring a fall back plan in case something goes wrong, as well as being accurate, reliable and reproducible. That is the only way to ensure that also unintentional harm can be minimized and prevented.

3. **Privacy and data governance**: besides ensuring full respect for privacy and data protection, adequate data governance mechanisms must also be ensured, taking into account the quality and integrity of the data, and ensuring legitimised access to data.

4. **Transparency**: the data, system and AI business models should be transparent. Traceability mechanisms can help achieving this. Moreover, AI systems and their decisions should be explained in a manner adapted to the stakeholder concerned. Humans need to be aware that they are interacting with an AI system, and must be informed of the system's capabilities and limitations.

5. **Diversity, non-discrimination and fairness**: Unfair bias must be avoided, as it could have multiple negative implications, from the marginalization of vulnerable groups, to the exacerbation of prejudice and discrimination. Fostering diversity, AI systems should be accessible to all, regardless of any disability, and involve relevant stakeholders throughout their entire life circle.

6. **Societal and environmental well-being**: AI systems should benefit all human beings, including future generations. It must hence be ensured that they are sustainable and environmentally friendly. Moreover, they should take into account the environment, including other living beings, and their social and societal impact should be carefully considered.

7. **Accountability**: Mechanisms should be put in place to ensure responsibility and accountability for AI systems and their outcomes. Auditability, which enables the assessment of algorithms, data and design processes plays a key role therein, especially in critical applications. Moreover, adequate an accessible redress should be ensured [28] (High Level Expert Group on Artificial Intelligence 2019).

We note here that these seven key requirements are developed in relation to AI. While AI is a background feature to certain *CONNECT* processes, *CONNECT*ś focus on connected vehicles rather than AI means that the seven key requirements need to be adapted to *CONNECT* and to CCAM. However, the main values involved are directly relevant to both connected and autonomous vehicles, and provide an invaluable foundation from which to build a set of methods to assess whether a CCAM component, process, action, or indeed CCAM vehicle or system are worthy of trust.

In this process, outlined below, we follow the seven key requirements in the order that the HLEG for AI presented them. This means that the following analysis and adaptation of the requirements starts with considerations that are more applicable to the specific technical issues of *CONNECT* but then gradually expand out to be applicable to CCAM, and wider cyber-physical systems more generally.

The work of the HLEG for AI was subsequently developed into the ALTAI, a tool to "help organisations identify how proposed AI systems might generate risks, and to identify whether and what kind of active measures may need to be taken to avoid and minimise those risks" (High Level Expert Group on Artificial Intelligence 2020). Following the ALTAI, each of the seven key requirements is tested through a series of questions and prompts, that encourage developers and stakeholders to reflect on the ways that particular features of their systems impact upon and should consider key ethical values.

This Assessment List for Trustworthy CCAM Components and Case Examples (ALTCCE) is best completed involving specialists familiar with either a specific component/aspect of *CONNECT* and/or with one of the three case examples. It is concerned with clarifying seven key requirements for trustworthy CCAM, to make an assessment of whether the Automated Information Handling system (AIH system) is worthy of trust in relation to one/all of the seven key requirements. These are listed below, and given CCAM relevant descriptions prior to each set of assessment questions. We note that this assessment list draws directly from the EC HLEG ALTAI, often directly using the language and terminology of the ALTAI. The purpose of the ALTCCE is to follow the guide set by the ALTAI, and to adapt it for CCAM. We do not give every reference where sections are direct quotations of the ALTAI.

The ALTCCE is underpinned by Fundamental Rights. Fundamental rights are concerned with the recognition humans as beings that are worthy not just of considerations in CCAM, but that these considerations should play a significant role in how we design, implement, and oversee CCAM. Like traditional vehicles, CCAM pose significant risks to the physical safety of individuals, whilst also offering these individuals significant freedom of movement. The ability for a person to move around their environment as they desire is perhaps so deep and fundamental that it can easily be overlooked. Whether it is a traditional vehicle or CCAM, being able to decide upon, and act upon, one's physical location is significant expression of individual freedom, and enables many other fundamental rights. At the same time, in contrast to traditional vehicles, CCAM also create risks to fundamental rights through the collection of information, the capacity for remote individuals or systems to control people's movements, and to organise the ways that individual vehicles operate in integrated systems. The basic idea of a fundamental right is that it is something that all

humans have, that it is something that carries significant weight in how we consider interactions and ordering relations between individuals, and that it can only be overridden in situations with very compelling reasons. For instance, we would likely consider things like basic physical safety, free movement, and privacy to be fundamental rights. Each of these should be considered in the design and implementation of CCAM. However, these fundamental rights may at times sit in tension with each other - a right to physical safety might be in tension with a right to free movement. Likewise, a right to privacy might be in tension with rights to physical safety or free movement.

For instance, if we are considering dynamic cruise control, one of the most important considerations would need to be the fundamental right of basic physical safety. Any CCAM system involving dynamic cruise control worthy of trust would need to run such that the vehicle operates safely, and would not put passengers or other road users at a level of physical risk beyond that normally accepted in standard road usage. Or, if we are considering intersection management, one of the most important considerations would need to the fundamental right of free movement. If a CCAM vehicle was wanting to (safely) navigate an intersection whilst turning right through moving traffic, then any CCAM system worthy of trust would need to run such that all vehicles moving through the intersection are actually able to travel towards their destination. If, for instance, traffic was such that the right-turning vehicle would either have to wait two minutes or turn left, the CCAM system here would not be trustworthy if the vehicle was redirected to turn left. Finally, if we are considering basic cybersecurity concerns, a further consideration of CCAM is that the information produced, used, communicated, and stored is not only reliable in terms of data accuracy and integrity, but also that any unnecessary personally identifying information is not produced, used, collected, or stored, and that any any such personally identifiable information that is needed by the CCAM system is in fact necessary, and where necessary, is protected through encryption and storage technologies.

In the following ALTCCE we offer a set of questions by which users can understand and recognise when key values are at stake, and a method of ensuring and assuring that the key values are being effectively considered in the design of *CONNECT* and CCAM. It is important to note that some of requirements may expose or even generate tensions between values. It is beyond the scope of *CONNECT* and the ALTCCE to resolve these tensions. This is in part because resolving such tensions is project beyond the scope of *CONNECT* but also because these tensions are going to arise in, and need to be understood in, particular and specific contexts. To simply say that physical safety is more important than privacy, or that privacy is more important than free movement will not effectively guide practice.

At a general level, fundamental rights encompass rights such as human dignity and non-discrimination, as well as rights in relation to data protection and privacy, to name just some examples. Following the ALTAI, the ALTCCE strongly recommends that particular fundamental rights are considered at the outset. Prior to self-assessing an Automated Information Handling (AIH) system with this Assessment List, a fundamental rights impact assessment (FRIA) should be performed. A FRIA could include questions such as the following – drawing on specific articles in the Charter and the European Convention on Human Rights (ECHR) its protocols and the European Social Charter.

Discriminations

Does the Automated Information Handling (AIH) system potentially negatively discriminate against people on the basis of any of the following grounds (non-exhaustively): sex, race, colour, ethnic or social origin, genetic features, language, religion or belief, political or any

other opinion, membership of a national minority, property, birth, disability, age or sexual orientation?

Have you put in place processes to test and monitor for potential negative discrimination (bias) during the development, deployment and use phases of the AIH system?

Have you put in place processes to address and rectify for potential negative discrimination (bias) in the AIH system?

Children

Does the AIH system respect the rights of the child, for example with respect to child protection and taking the child's best interests into account?

Have you put in place processes to address and rectify for potential harm to children by the AIH system?

Have you put in place processes to test and monitor for potential harm to children during the development, deployment and use phases of the AIH system?

GDPR

Does the AIH system protect personal data relating to individuals in line with GDPR?

Have you put in place processes to assess in detail the need for a data protection impact assessment, including an assessment of the necessity and proportionality of the processing operations in relation to their purpose, with respect to the development, deployment and use phases of the AIH system?

Have you put in place measures envisaged to address the risks, including safeguards, security measures and mechanisms to ensure the protection of personal data with respect to the development, deployment and use phases of the AIH system?

See the section on Privacy and Data Governance in this Assessment List, and available guidance from the European Data Protection Supervisor.

Freedoms

Does the AIH system respect the freedom of expression and information and/or freedom of assembly and association?

Have you put in place processes to test and monitor for potential infringement on freedom of expression and information, and/or freedom of assembly and association, during the development, deployment and use phases of the AIH system?

Have you put in place processes to address and rectify for potential infringement on freedom of expression and information, and/or freedom of assembly and association, in the AIH system?

2.3.2.1 REQUIREMENT #1 Human Agency and Oversight

When considering human agency and oversight for AIH in CCAM, there are two sorts of agency that may be involved. First is concerned with the agency of the vehicle's driver, passenger, or other road user (we will refer to them as the 'driver' but mean to include any person using the vehicle and/or road). Second is the agency of the person developing the CCAM component or system, such as a computer engineer, software designer, cybersecurity expert and so on (we will refer to this group as the 'developer', but mean the full range of those people engaged with the

CCAM components, system, and so on) . While agency is relevant to both sets (driver and developer), ultimately, what is of most importance is the user's agency. This is because the developer's agency - while important - is largely a function of other considerations (such as accountability, see key requirement 7). In contrast, the driver's agency is directly impacted by CCAM vehicle. The right to free movement, for instance, is a right relevant to the driver of a vehicle, that is either supported or diminished by the operation of that CCAM vehicle. That said, while the ultimate concern and focus of CCAM is on drivers, passengers, and other road users, the decisions made in the development, design, and implementation of CCAM that impact human agency are largely the province of the developers, designers, and other expert stakeholders working for Original Equipment Manufacturers (OEMs) government oversight bodies, even legislators. As such, while the subjects of human agency are drivers, the developers, are the ones who need to ensure that the requirement to human agency and oversight is met. It is also necessary to highlight that simply having a right to free movement, expressed by a requirement to human agency is not absolute.

First, and perhaps most obviously, a CCAM vehicle and wider CCAM system would need to take into account the **physical safety of all road users**. Simply because a particular driver wants to arrive at their destination as quickly as possible does not mean that they can ignore road safety, traffic lights, and other established rules and norms of safe driving. Perhaps more controversially, simply because a particular driver wants to arrive at their destination as quickly as possible does not mean that they can exceed the speed limit, even if there are no other vehicles or road users nearby. While the fundamental rights are important, and human agency is absolutely vital, this is not an absolute right. The requirement to human agency for CCAM require considerations beyond basic rights and responsibilities while using the road. Consider that many functions of CCAM will require considerations of maintenance - especially security updates, the right to repair, responsibilities and protections of personal data following vehicle sales and decommissioning. To look at human agency in relation to security updates, there is a tension between a driver's right to accept/reject security updates on their vehicle, and the basic requirements of vehicle road worthiness and/or the security risks that an unprotected vehicle poses to the CCAM system.

Our suggestion here is that, like existing road worthiness requirements, some system critical features should be either updated or kept to a minimum level, given the risks that they pose to other road users, whilst other features that are less critical to safety and system security may be decided by the driver. However, such decisions and their implications would need to be clearly communicated to the drivers.

At a general level, AIH systems in CCAM present ways to support human agency, as well as potential undermining of or violation of human agency. **Trustworthy CCAM will create an environment where humans are freed from the need to attend to all factors in driving, but at the same time, they are not prevented from making ethically relevant decisions impacting either decisions while driving, or the general features of the AIH.** For instance, AIH systems should be able to safely guide a vehicle through an intersection, engage in adaptive cruise control while the vehicle is in motion, and be safe from unknown and unwanted access to and changes to system critical information. Moreover, AIH systems in CCAM should be designed in such a way as to allow human oversight of specific and general aspects of AIH.

Human Agency and Autonomy

This subsection deals with the effect Automated Information Handling (AIH) systems can have on human behaviour in CCAM. It deals with the effect of AIH systems that are support intersection management, adaptive cruise control, and data misbehaviour and management. It also deals with the effect on human perception and expectation when confronted with AIH systems that present

information on the trustworthiness of CCAM. Finally, it deals with the effect of AIH systems on trust and (in)dependence. All of these concerns must be taken into account by CCAM developers. Table 2.1 summarises these concerns.

Human Oversight

This subsection helps to self-assess necessary oversight measures through governance mechanisms such as human-in-the-loop (HITL), human-on-the-loop (HOTL), or human-in-command (HIC) approaches. Human-in-the-loop refers to the capability for human intervention in every decision cycle of the system. Human-on-the-loop refers to the capability for human intervention during the design cycle of the system and monitoring the system's operation. Human-in-command refers to the capability to oversee the overall activity of the AIH system (including its broader economic, societal, legal and ethical impact) and the ability to decide when and how to use the AIH system in any particular situation. The latter can include the decision not to use an AIH system in a particular situation to establish levels of human discretion during the use of the system, or to ensure the ability to override a decision made by an AIH system. Making decisions about whether the CCAM system is HITL, HOTL, or HIC have a deep impact on the resulting human agency. But, as with human agency and autonomy, there are perhaps design decisions which should afford or maximise the driver as a HITL, or ought to be limited to the developers, at HIC. There are two important considerations here. First, for developers to recognise how their design decisions track to HITL, HOTL, and HIC mechanisms. Second, for developers to be aware of, and able to communicate, the reasons why particular governance decisions were made. That is, if a particular CCAM design feature enables HITL decision making where a driver has significant autonomy over the vehicle's operation, or the design feature enables HIC decision making where a driver's autonomy over a vehicle's operation is significantly limited but a designer's autonomy is significantly enhanced, then what reasons can the developers give to support and justify deciding one way or the other? Table 2.1 provides a summary of relevant questions.

Table 2.1: Human Agency and Autonomy - Human Oversight

Human Agency and Autonomy
Is the Automated Information Handling (AIH) system designed to interact, guide or take decisions by human end-users that affect humans or society?
Could the Automated Information Handling (AIH) system generate confusion for some or all end-users or subjects on whether a decision, content, advice or outcome is the result of an automated information handling decision? • Are end-users or other subjects adequately made aware that a decision, content, advice or outcome is the result of an AIH decision?
Could the AIH system generate confusion for some or all end-users or subjects on whether they are interacting with a human or AIH system? • Are end-users or subjects informed that they are interacting with a system in which certain critical decisions are the result of automated information handling?
Could the AIH system affect human autonomy by generating over-reliance by end-users?

(Continued on next page)

(Continued from previous page)

Human Agency and Autonomy
• Did you put in place procedures to avoid that end-users over-rely on the AIH system?
Could the AIH system affect human autonomy by interfering with the end-user's decision-making process in any other unintended and undesirable way? • Did you put in place any procedure to avoid that the AIH system inadvertently affects human autonomy?
Does the AIH system simulate social interaction with or between end-users or subjects?
Human Oversight
Please determine whether the AIH system (choose as many as appropriate): • Is a component, operation or feature of the CCAM system; • Is overseen by a Human-in-the-Loop; • Is overseen by a Human-on-the-Loop; • Is overseen by a Human-in-Command.
Have the humans (human-in-the-loop, human-on-the-loop, human-in-command) been given specific training on how to exercise oversight?
Did you establish any detection and response mechanisms for undesirable adverse effects of the AIH system for the end-user or subject?
Did you ensure a 'stop button' or procedure to safely abort an operation when needed?
Did you take any specific oversight and control measures to reflect the self-learning or autonomous nature of the AIH system?

Possible human actions (Human-on-the-loop)

The human on the loop might be involved in, or be required to intervene in the following actions:

- Software updates
- System or hardware maintenance
- System modification (Tuning)
- Accept the sharing of personal data
- Turning off the function to take over control manually.

2.3.2.2 REQUIREMENT #2 Technical Robustness and Safety

CCAM rely on technical robustness to both ensure, and assure, users of their reliability and safety. A trustworthy CCAM vehicle and system is one in which the components work as expected, can

be identified if they cease working as expected, all to produce an outcome where the physical safety of all road users is both ensured and assured. That is, that all road users are in fact safe, and that there are mechanisms in place to not just bring about that physical safety, but there are also mechanisms to record and communicate system processes. This is to assure users that if any safety breaches occur, then the relevant users have a way of knowing that such breaches have occurred, that physical safety and information security are protected, and that the relevant practices and policies are being followed to mitigate the harms and reduce the chances of them happening again.

A crucial requirement for achieving Trustworthy AIH systems is their dependability (the ability of the overall CCAM system services in which trust is justified), robustness - produces stable results when facing changes such that it doesn't break or fail easily, and resilience means recovering quickly from challenges. Technical robustness requires that AIH systems are developed with a preventative approach to risks and that they behave reliably and as intended while minimising unintentional and unexpected harm as well as preventing it where possible. This should also apply in the event of potential changes in their operating environment or the presence of other agents (human or artificial) that may interact with the AI system in an adversarial manner. The questions in this section address four main issues: 1) security; 2) safety; 3) accuracy; and 4) reliability, fall-back plans and reproducibility. What systems are in place to recognise if there is malicious activity, and if malicious activity occurs, how the system recover from the attack in a timely manner. Table 2.2 summarises the relevant questions.

Accuracy

The centrepiece of *CONNECT* is the trust assessment framework, the TAF. The TAF relies upon accurate data to produce an assessment of whether a particular operation or decision is trustworthy. Table 2.2 summarises the relevant questions.

Reliability, Fall-back plans and Reproducibility Additionally to the accuracy, ensuring reliable system behavior, robust fall-back plans for failure scenarios, and reproducible outcomes for audit and validation is necessary in susch a Trust Assessment Framework.

Table 2.2: Resilience to Attack - Security and General Safety - Accuracy - Reliability, Fall-back plans and Reproducibility

Resilience to Attack and Security
Could the AIH system have adversarial, critical, or damaging effects (e.g., to passengers, other road users, or effects on the wider CCAM system) in case of risks or threats such as design or technical faults, defects, outages, attacks, misuse, inappropriate or malicious use?
Is the AIH system certified for cybersecurity (e.g., the certification scheme created by the Cybersecurity Act in Europe) or is it compliant with specific security standards?
How exposed is the AI system to cyber-attacks? • Did you assess potential forms of attacks to which the AIH system could be vulnerable?

Did you consider different types of vulnerabilities and potential entry points for attacks such as: • Data poisoning (i.e., manipulation of training data); • Model evasion (i.e., classifying the data according to the attacker's will); • Model inversion (i.e., inferring the model parameters).
Did you put measures in place to ensure the integrity, robustness, and overall security of the AIH system against potential attacks over its lifecycle?
Did you red-team or pentest the system?
Did you inform end-users of the duration of security coverage and updates? • What is the expected time frame within which you provide security updates for the AIH system?
General Safety
Did you define risks, risk metrics, and risk levels of the AIH system in each specific use case? • Did you put in place a process to continuously measure and assess risks? • Did you inform end-users and subjects of existing or potential risks?
Did you identify the possible threats to the AIH system (design faults, technical faults, environmental threats) and the possible consequences? • Did you assess the risk of possible malicious use, misuse, or inappropriate use of the AIH system? • Did you define safety criticality levels (e.g., related to human integrity) of the possible consequences of faults or misuse of the AIH system?
Did you assess the dependency of a critical AIH system's decisions on its stable and reliable behavior? • Did you align the reliability/testing requirements to the appropriate levels of stability and reliability?
Did you plan fault tolerance via, e.g., a duplicated system or another parallel system (AIH-based or 'conventional')?
Did you develop a mechanism to evaluate when the AIH system has been changed to merit a new review of its technical robustness and safety?
Accuracy
Could a low level of accuracy of the AIH system result in critical, adversarial, or damaging consequences?

Did you put in place measures to ensure that the data (including training data) used to develop the AIH system is: • Up-to-date, • Of high quality, • Complete, and • Representative of the environment the system will be deployed in?
Did you establish a series of steps to monitor and document the AIH system's accuracy?
Did you consider whether the AIH system's operation could invalidate the data or assumptions it was trained on, and how this might lead to adversarial effects?
Did you implement processes to ensure that the expected level of accuracy of the AIH system is properly communicated to end-users and/or subjects?
Reliability, Fall-back plans and Reproducibility
Could the AIH system cause critical, adversarial, or damaging consequences (e.g. per- taining to human safety or overall trust in CCAM systems) in case of low reliability and/or reproducibility? • Did you put in place a well-defined process (such as the TAF) to monitor if the AIH system is meeting the intended goals? • Did you test whether specific contexts or conditions need to be taken into account to ensure reproducibility?
Did you put in place verification and validation methods and documentation (e.g. logging) to evaluate and ensure different aspects of the AIH system's reliability and reproducibility? • Did you clearly document and operationalize processes for the testing and verification of the reliability and reproducibility of the AIH system?
Did you define tested failsafe fallback plans to address AIH system errors of whatever origin and put governance procedures in place to trigger them?
Did you put in place a proper procedure for handling the cases where the AIH system yields results with a low confidence score?
Is your AIH system using (online) continual learning? • Did you consider potential negative consequences from the AI system learning novel or unusual methods to score well on its objective function?

2.3.2.3 REQUIREMENT #3 Privacy and Data Governance

Closely linked to the principle of prevention of harm is privacy, a fundamental right particularly affected by AIH systems. Prevention of harm to privacy also necessitates adequate data governance that covers the quality and integrity of the data used, its relevance in light of the domain in which the AIH systems will be deployed, its access protocols and the capability to process data in a manner that protects privacy.

Privacy is typically considered to be an issue where information is produced, used, or applied to people. For instance, does the intersection management system or adaptive cruise control use, store, communicate, or produce any information that can be easily related to a person as a source or target of that information? Does the AIH use, store, communicate, or produce information about a person that can be accessed by unauthorised actors due to cyber security vulnerabilities?

Table 2.3 helps to self-assess the impact of the AIH system's on privacy and data protection, which are fundamental rights that are closely related to each other and to the fundamental right to the integrity of the person, which covers the respect for a person's mental and physical integrity as well as the adherence of the AIH system('s use) to various elements concerning data protection.

Table 2.3: Privacy and Data Governance

Privacy
Did you consider the impact of the CCAM system on the right to privacy and the right to data protection?
Depending on the use case, did you establish mechanisms that allow flagging issues related to privacy concerning the CCAM system?
Have you considered if your product/system utilizes personal information, and do you take the requisite steps to ensure that such data is processed in a lawful and ethical manner? • This includes personal information that can be derived from the normal operation of the CCAM system.
Have you considered legal and ethical implications of any non-personal information your system might be using?
Data Governance
Is your AIH system being trained, or was it developed, by using or processing personal data (including special categories of personal data)?
Did you put in place any of the following measures, some of which are mandatory under the General Data Protection Regulation (GDPR), or a non-European equivalent? • Data Protection Impact Assessment (DPIA) • Designate a Data Protection Officer (DPO) and include them at an early stage in the development, procurement, or use phase of the AIH system • Oversight mechanisms for data processing (including limiting access to qualified personnel, mechanisms for logging data access and making modifications) • Measures to achieve privacy-by-design and default (e.g., encryption, pseudonymisation, aggregation, anonymisation) • Data minimisation, in particular personal data (including special categories of data)
Did you implement the right to withdraw consent, the right to object, and the right to be forgotten into the development of the AIH system?
Did you consider the privacy and data protection implications of data collected, generated, or processed over the course of the AIH system's life cycle?
Did you consider the privacy and data protection implications of the AI system's non-personal training data or other processed non-personal data?
Did you align the AIH system with relevant standards (e.g., ISO25, IEEE26) or widely adopted protocols for (daily) data management and governance?

2.3.2.4 REQUIREMENT #4 Transparency

AIH like *CONNECT*ś CCAM rely explicitly on the collection and use of data to ensure and assure the component, operation, vehicle's trusworthiness. But such processes require transparency around how that data is generated, communicated, used, and stored. A crucial component of achieving Trustworthy CCAM is transparency which encompasses three elements: 1) traceability, 2) explainability and 3) open communication about the limitations of the AIH system.

Traceability

A crucial component of achieving Trustworthy CCAM is transparency which encompasses three elements: 1) traceability, 2) explainability and 3) open communication about the limitations of the AIH system.

The traceablity part in Table 2.4 helps to self-assess whether the processes of the development of the AIH system, i.e. the data and processes that yield the AIH system's decisions, is properly documented to allow for traceability, increase transparency and, ultimately, build trust in CCAM in society.

Explainability

The explainability part in Table 2.4 helps to self-assess the explainability of the AIH system. The questions refer to the ability to explain both the technical processes of the AI system and the reasoning behind the decisions or predictions that the AIH system makes. Explainability is crucial for building and maintaining users' trust in AI systems. AIH driven decisions – to the extent possible – must be explained and understood to those directly and indirectly affected, in order to allow for contesting of such decisions. An explanation as to why a model has generated a particular output or de- cision (and what combination of input factors contributed to that) is not always possible. These cases are referred to as 'blackboxes' and require special attention. In those circumstances, other explainability measures (e.g. traceability, auditability and transparent communication on the AIH system's capabilities) may be required, provided that the AIH system as a whole respects fun- damental rights. The degree to which explainability is needed depends on the context and the severity of the consequences of erroneous or otherwise inaccurate output to human life.

Communication

The communication part in Table 2.4 helps to self-assess whether the AIH system's capabilities and limitations have been communicated to the users in a manner appropriate to the use case at hand. This could encompass communication of the AIH system's level of accuracy as well as its limitations.

Table 2.4: Traceability - Explainability - Communication

Traceability
Did you put in place measures that address the traceability of the AIH system during its entire lifecycle?

(Continued on next page)

(Continued from previous page)

Traceability
• Did you put in place measures to continuously assess the quality of the input data to the AIH system? • Can you trace back which data was used by the AIH system to make a certain decision(s) or recommendation(s)? • Can you trace back which AIH model or rules led to the decision(s) or recommendation(s) of the AIH system? • Did you put in place measures to continuously assess the quality of the output(s) of the AIH system? • Did you put adequate logging practices in place to record the decision(s) or recommendation(s) of the AIH system?
Explainability
Is there the capacity to explain the decision(s) of the AIH system to the users?
Will downstream producers and service providers have the capacity to survey the key stakeholders (such as OEMS) to assess if they understand the decision(s) of the AIH system?
Communication
Did you have mechanisms in place that will allow users to be informed about the purpose, criteria, and limitations of the decision(s) generated by the AIH system? • Can you communicate the benefits of the AIH system to users? • Can you communicate the technical limitations and potential risks of the AIH system to users, such as its level of accuracy and/or error rates? • Did you provide appropriate training material and disclaimers to users on how to adequately use the AIH system?

2.3.2.5 REQUIREMENT #5 Diversity, Non-discrimination and Fairness

As we move from more technical features of AIH to wider values like fairness, the applicability of these requirements to CCAM is less obvious. However, there are two points that need to be highlighted here. First, while a requirement to diversity, non-discrimination, and fairness, may not be directly relevant to *CONNECT* this requirement still needs to be considered in relation to wider CCAM. Second, a full ethical accounting of CCAM would still need to have processes to ensure and assure users that the system is fair - that is, it needs to be checked for the recognition of diversity, the potential for bias and discrimination to occur in CCAM, and that a principle of fairness is properly considered by developers. To show how this requirement might be relevant for CCAM, and must be properly considered, think of a scenario where three vehicles are approaching an intersection. Vehicle 1 is a connected vehicle, and the driver has willingly opted to maximum information sharing with the information management system of the intersection. Vehicle 2 is a

connected vehicle, but the driver is concerned about privacy and information security, and so has set their vehicle to only share the complete minimum information needed by the intersection information management system. Vehicle 3 is a traditional vehicle and has no technical capacity to share any information with the intersection information management system. On this description, it would be likely that Vehicle 1 might move through the intersection more quickly than Vehicle 2, and Vehicle 3 might be slowed or in fact be prevented from navigating through the intersection at all. This then raises issues of respect for diversity, non-discrimination, and fairness. On the one hand, it seems a rejection of the notion of diversity, becoming discriminatory and unfair to Vehicles 2 and 3 that they are slowed or prevented from using the intersection, while Vehicle 1 is privileged. This in part because we ought to take the requirement to privacy seriously, and in part because it is overdemanding and impractical to assume that all drivers are willing or even able to buy connected vehicles. On the other hand, the driver of Vehicle 2, and perhaps Vehicle 3, are 'free riding' off Vehicle 1. They do not want to share their vehicle's information, but they want the efficiency and safety that connected vehicles bring. So it is perhaps unfair of them to benefit from the informational openness of other road users like Vehicle 1. The point here is not to offer solutions to this problem, but instead to demonstrate how the requirement here for diversity, non-discrimination and fairness is applicable to CCAM and AIH more generally.

Avoidance of Unfair Bias

In order to achieve Trustworthy CCAM, we must enable inclusion and diversity throughout the entire AIH system's life cycle. AIH systems (both for development and operation) may suffer from the inclusion of inadvertent historic bias, incompleteness, and bad governance models. The continuation of such biases could lead to unintended (in)direct prejudice and discrimination against certain groups or people, (including people exercising their rights not to share information, accepting software updates - should) potentially exacerbating prejudice and marginalisation. Harm can also result from the intentional exploitation of (consumer) biases or by engaging in unfair competition, such as privileging one set of road users over others. Identifiable and discriminatory bias should be removed in the collection phase where possible. AIH systems should be user-centric and designed in a way that allows all people to use CCAM products or services, regardless of their age, gender, abilities or characteristics. Accessibility to this technology for persons with disabilities, which are present in all societal groups, is of particular importance. The question is whether the differential treatment is justified or not?

Accessibility and Universal Design

Particularly in business-to-consumer domains, AIH systems should be user-centric and designed in a way that allows all people to use CCAM products or services, regardless of their age, gender, abilities or characteristics. Accessibility to this technology for persons with disabilities, which are present in all societal groups, is of particular importance. AIH systems should not have a one-size-fits-all approach and should consider Universal Design principles addressing the widest possible range of users, following relevant accessibility standards. This will enable equitable access and active participation of all people in existing and emerging computer-mediated human activities and with regard to assistive technologies.

Stakeholder Participation

In order to develop Trustworthy CCAM, it is advisable to consult stakeholders who may directly or indirectly be affected by the CCAM system throughout its life cycle. It is beneficial to solicit regular feedback even after deployment and set up longer term mechanisms for stakeholder participation, for example by ensuring workers information, consultation and participation throughout the whole process of implementing CCAM systems at organisations. Table 2.5 provides a summary of all relevant questions.

Table 2.5: Avoidance of Unfair Bias - Accessibility and Universal Design - Stakeholder Participation

Avoidance of Unfair Bias
Did you establish a strategy or a set of procedures to avoid creating or reinforcing unfair bias in the AIH/CCAM system, both regarding the use of input data as well as for any algorithm design?
Did you consider diversity and representativeness of end-users and/or subjects in the data? • Did you test for specific target groups or problematic use cases? • Did you research and use publicly available technical tools, that are state-of-the-art, to improve your understanding of the data, model, and performance? • Did you assess and put in place processes to test and monitor for potential biases during the entire lifecycle of the AIH system (e.g., biases due to possible limitations stemming from the composition of the used data sets, such as lack of diversity or non-representativeness)? • Where relevant, did you consider diversity and representativeness of end-users and/or subjects in the data?
Did you put in place educational and awareness initiatives to help AIH designers and AIH developers be more aware of the possible bias they can inject in designing and developing the AIH/CCAM system?
Did you ensure a mechanism that allows for the flagging of issues related to bias, discrimination, or poor performance of the AIH system? • Did you establish clear steps and ways of communicating on how and to whom such issues can be raised? • Did you identify the subjects that could potentially be (in)directly affected by the AI system, in addition to the (end-)users and/or subjects?
Is your definition of fairness commonly used and implemented in any phase of the process of setting up the AIH/CCAM system? • Did you consider other definitions of fairness before choosing this one? • Did you consult with the impacted communities about the correct definition of fairness, i.e., representatives of elderly persons or persons with disabilities? • Did you ensure a quantitative analysis or metrics to measure and test the applied definition of fairness? • Did you establish mechanisms to ensure fairness in your AIH/CCAM system?
Accessibility and Universal Design
Did you ensure that the CCAM system corresponds to the variety of preferences and abilities in society?
Did you assess whether the CCAM system's user interface is usable by those with special needs or disabilities or those at risk of exclusion?

(Continued on next page)

(Continued from previous page)

Accessibility and Universal Design
Did you ensure that information about, and the CCAM system's user interface of, the CCAM system is accessible and usable also to users of assistive technologies (such as screen readers)?
Did you involve or consult with end-users or subjects in need of assistive technology during the planning and development phase of the CCAM system?
Did you ensure that Universal Design principles are taken into account during every step of the planning and development process, if applicable?
Did you take the impact of the CCAM system on the potential end-users and/or subjects into account? • Did you assess whether the team involved in building the AIH/CCAM system engaged with the possible target end-users and/or subjects? • Did you assess whether there could be groups who might be disproportionately affected by the outcomes of the AIH/CCAM system? • Did you assess the risk of the possible unfairness of the system on the end-user's or subject's communities?
Stakeholder Participation
Did you consider a mechanism to include the participation of the widest range of possible stakeholders in the CCAM system's design and development?

2.3.2.6 REQUIREMENT#6 Societal and Environmental Well-being

In line with the principles of fairness and prevention of harm, the broader society, other sentient beings and the environment should be considered as stakeholders throughout the CCAM system's life cycle. Ubiquitous exposure to CCAM systems in all areas of our lives may alter our conception of social agency, or negatively impact our social relationships and attachment. While CCAM systems can be used to enhance social mobility, they can equally contribute to inequality if particular people are excluded from using CCAM. This could equally affect peoples' physical and mental well-being. The effects of CCAM systems must therefore be carefully monitored and considered. Sustainability and ecological responsibility of CCAM systems should be encouraged, and research should be fostered into AI solutions addressing areas of global concern, for instance the Sustainable Development Goals.

Environmental Well-being

This subsection helps to self-assess the (potential) positive and negative impacts of the CCAM system on the environment. CCAM systems, even if they promise to help tackle some of the most pressing societal concerns, e.g. climate change, must work in the most environmentally friendly way possible. The CCAM system's development, deployment and use process, as well as its entire supply chain, should be assessed in this regard (e.g. via a critical examination of the resource usage and energy consumption during training, opting for less net negative choices). Measures to secure the environmental friendliness of an CCAM's system's entire supply chain should be encouraged.

Impact on Work and Skills

CCAM systems may fundamentally alter the work sphere. They should support humans in the

working environment, and aim for the creation of meaningful work. This subsection helps self-assess the impact of the CCAM system and its use in a working environment on workers, the relationship between workers and employers, and on skills.

Impact on Society at large or Democracy

This subsection helps to self-assess the impact of a CCAM system from a societal perspective, taking into account its effect on institutions, democracy and society at large. The use of CCAM systems should be given careful consideration, particularly in situations relating to the democratic processes, including not only political decision-making but also electoral contexts (e.g. when AI systems amplify fake news, segregate the electorate, facilitate totalitarian behaviour, etc.). As with the requirement to diversity, non-discrimination, and fairness, the relations between CCAM and social impacts are at first hard to see. However, consider the relations between population surveillance and CCAM. If CCAM does not not properly respect individual and collective privacy, then it is possible for vehicle manufacturers to collect large amounts of data on individuals and groups. Likewise, it is possible for governments to either collect or access this personal and group data. Such surveillance can have widespread negative social and democratic impacts through 'chilling'. The basic concern with chilling is that if people know or suspect that they are under surveillance, then their behaviours will change. CCAM have the potential to add significant physical aspects to this chilling, through the potential to remotely limit where vehicles and drivers may go. For instance, if there is the capacity for remote control of CCAM vehicles, then the basic right of free movement is potentially limited. A further element of chilling and remote control of CCAM is that it is not just enough for manufacturers and governments to be prevented from, or limited in what they can do here. People must also be assured that no such surveillance is occurring, and that no such remote control is available. The negative impacts on society and democracy are not just whether people are under surveillance or can have their physical movements controlled, but whether people know and believe that no such surveillance or remote control is occurring. Table 2.6 provides a summary of all relevant questions.

Table 2.6: Environmental Well-being - Impact on Work and Skills - Impact on Society at Large or Democracy

Environmental Well-being
Are there potential negative impacts of the CCAM system on the environment? • Which potential impact(s) do you identify?
Where possible, did you establish mechanisms to evaluate the environmental impact of the CCAM system's development, deployment and/or use (for example, the amount of energy used and carbon emissions)? • Did you define measures to reduce the environmental impact of the CCAM system throughout its lifecycle?
Impact on Work and Skills
Does the CCAM system impact human work and work arrangements?
Did you pave the way for the introduction of the CCAM system in your organisation by informing and consulting with impacted workers and their representatives (trade unions, (European) work councils) in advance?

(Continued on next page)

(Continued from previous page)

Impact on Work and Skills
Did you adopt measures to ensure that the impacts of the CCAM system on human work are well understood? • Did you ensure that workers understand how the CCAM system operates, which capabilities it has and which it does not have?
Could the CCAM system create the risk of de-skilling of the workforce? • Did you take measures to counteract de-skilling risks?
Does the system promote or require new (digital) skills? • Did you provide training opportunities and materials for re- and up-skilling?
Impact on Society at Large or Democracy
Could the CCAM system have a negative impact on society at large or democracy? • Did you assess the societal impact of the AI system's use beyond the (end-)user and subject, such as potentially indirectly affected stakeholders or society at large? • Did you take action to minimize potential societal harm of the CCAM system? • Did you take measures that ensure that the AI system does not negatively impact democracy?

2.3.2.7 REQUIREMENT#7 Accountability

Requirement 7 is final requirement for trustworthy AIH, and is placed last by design. This is not because accountability is the least important, but because it is the most general, and draws from and relates to the other requirements. This requirement is a system level process, which is concerned with ensuring and assuring users about who did what, who is responsible for what, and to give guidance on what happens to the relevant actors when something unwanted has happened. While accountability takes in considerations about components and particular processes, it is about the overall AIH system's functioning. Moreover, it is hard to have a system that is worthy of trust without a higher level consideration of accountability within that system. For instance, while safety in intersection management is vital to consider, if there are no systems in place to ensure that that intersection is safe, and no systems to hold relevant people to account if they have failed at their role in ensuring and assuring the safety of users, then the CCAM system will neither be trusted, nor is trustworthy. The principle of accountability necessitates that mechanisms be put in place to ensure responsibility for the development, deployment and/or use of AIH systems for CCAM. This topic is closely related to risk management, identifying and mitigating risks in a transparent way that can be explained to and audited by third parties. When unjust or adverse impacts occur, accessible mechanisms for accountability should be in place that ensure an adequate possibility of redress.

Auditability

This subsection helps to self-assess the existing or necessary level that would be required for

an evaluation of the AIH system by internal and external auditors. The possibility to conduct evaluations as well as to access records on said evaluations can contribute to Trustworthy AIH. In applications affecting fundamental rights, including safety-critical applications, AIH systems should be able to be independently audited. This does not necessarily imply that information about business models and intellectual property related to the AIH system must always be openly available.

Risk Management

Both the ability to report on actions or decisions that contribute to the AIH system's outcome, and to respond to the consequences of such an outcome, must be ensured. Identifying, assessing, documenting and minimising the potential negative impacts of AIH systems is especially crucial for those (in)directly affected. Due protection must be available for whistle-blowers, NGOs, trade unions or other entities when reporting legitimate concerns about an AIH system.

When implementing the above requirements, tensions may arise between them, which may lead to inevitable trade-offs. Such trade-offs should be addressed in a rational and methodological manner within the state of the art. This entails that relevant interests and values implicated by the AIH system should be identified and that, if conflict arises, trade-offs should be explicitly acknowledged and evaluated in terms of their risk to safety and ethical principles, including fundamental rights. Any decision about which trade-off to make should be well reasoned and properly documented. When adverse impact occurs, accessible mechanisms should be foreseen that ensure adequate redress. Table 2.7 provides a summary of all relevant questions.

Table 2.7: Auditability - Risk Management

Auditability
Did you establish mechanisms that facilitate the AIH system's auditability (e.g., traceability of the development process, the sourcing of training data, and the logging of the AIH system's processes, outcomes, positive and negative impact)?
Did you ensure that the AIH system can be audited by independent third parties?
Risk Management
Did you foresee any kind of external guidance or third-party auditing processes to oversee ethical concerns and accountability measures? • Does the involvement of these third parties go beyond the development phase?
Did you organize risk training, and if so, does this also inform about the potential legal framework applicable to the AIH system?
Did you consider establishing an AIH ethics review board or a similar mechanism to discuss the overall accountability and ethics practices, including potential unclear grey areas?

(Continued on next page)

(Continued from previous page)

Risk Management
Did you establish a process to discuss and continuously monitor and assess the AIH system's adherence to this Assessment List for Trustworthy AIH CCE (ALTCCE)? • Does this process include identification and documentation of conflicts between the six aforementioned requirements or between different ethical principles and explanation of the 'trade-off' decisions made? • Did you provide appropriate training to those involved in such a process and does this also cover the legal framework applicable to the AIH system?
Did you establish a process for third parties (e.g., suppliers, end-users, subjects, distributors/vendors, or workers) to report potential vulnerabilities, risks, or biases in the AIH system? • Does this process foster revision of the risk management process?
For applications that can adversely affect individuals, have redress by design mechanisms been put in place?

2.4 Validating and Ensuring the achievement of set trustworthiness standards

The final part of the *CONNECT* trust method involves the enrolment of, and reflections from, key stakeholders. To engage in this process, there are three sub-steps: identification of key stakeholders, validation of ALTCCAM by stakeholders, and ensuring that the standards have been met (and if they not met, what to do). The identification of key stakeholders requires that representatives of particular groups specially related to CCAM are identified. There are four general groups that need to be considered as potential contributing stakeholders: experts in CCAM design, experts in CCAM development and production, experts in CCAM oversight and accountability, and wider non-expert users and community members affected by CCAM.

The first set – experts in CCAM design - is those individuals who are involved in, or have the necessary technical expertise to understand, the core theoretical and technical features of CCAM components, processes, and systems. For instance, software engineers involved in designing the trust assessment framework, the production of outputs for actual trust levels, cybersecurity experts tasked with assessing risks of malicious intrusion etc. More specifically for *CONNECT* this is the members of the *CONNECT* project.

The second set – experts in CCAM development and production – is those individuals working for organisations who use the CCAM components, processes, and systems in their production of CCAM. For instance, this is OEMs (original equipment manufacturers), automotive vendors, security service providers and so on. The second set can loosely be considered users but not consumers as end users. For *CONNECT* this is Tier 1 CCAM researchers such as DENSO, Tier 2 researchers such as CRF, automotive vendors such as Toyota, and security service providers such as Intel.

The third set – experts in CCAM oversight and accountability – is those institutions tasked with setting and assessing CCAM relevant standards and policy making. Here, this set of experts can

assess the trustworthiness of CCAM by reference to existing or new standards or policies. For instance, standards bodies involved in CCAM safety or IoT cybersecurity would be able to review whether particular CCAM features allow for trustworthy assessments given car safety, cybersecurity protocols and so on. They would also include experts in privacy and data management, and other state, national, and international legislation so individuals familiar with GDPR, the European AI act, and so on would need to be involved in the trustworthiness assessments to check whether the components, processes, and systems meet the specific conditions of GDPR, the AI act, or have the potential to meet those standards and policies when put into production. The assessments here, and inclusion of these stakeholders are beyond the scope of *CONNECT* but would be part of a wider CCAM trust assessment.

The final set - wider non-expert users and community members affected by CCAM – would involve wider community consultation. This set of people can loosely be considered those people who are non-CCAM experts who would either use the end CCAM product, or potentially be affected by the end CCAM product. In short, it would be consumers and CCAM customers, and other road users. This set is very broad and can potentially be very large, consisting in whole populations. The assessments here and inclusion of these stakeholders are beyond the scope of *CONNECT* but would be part of a wider CCAM trust assessment in the lead up to, and following roll out of end user CCAM vehicles.

In terms of validation of the ALTCCAM, the first two stakeholder sets will be enrolled in short testing and surveys of the ALTCCAM. They will go through the general questions given in Step 1 and Step 2, and then the specific questions of Step 3, given in Section 2.3. The purpose of these validation steps is to test that the questions are applicable, and understandable. The *CONNECT* trustworthiness approach is intended to be dynamic, and responsive to stakeholder feedback and advice.

At the same time as the validation process, the trustworthiness standards are clarified and assessed for particular CCAM components, processes, and systems by the stakeholders. This goes to the aim of recognising and then translating the seven ethical values identified in the HLEG seven key areas into a *RTL (Required Trustworthiness Level)*. Given *CONNECT*ś focus on the TAF as a way of representing the *ATL (Actual Trustworthiness Level)*, the RTL for particular key areas would be presented in a way that accords with the ATL. The two stakeholder sets would also consider and offer suggestions how the remaining key areas could be represented as RTLs.

For Stakeholder Set 1, the validation and evaluation steps would track to the following questions:

1. Which of the 7 key requirements are relevant to your work on *CONNECT* and which are relevant to *CONNECT* generally?
2. Rank them in importance to your work/to *CONNECT*.Note that this question is then repeated for all seven key requirements (i.e. for Requirement #1 Human Agency and Oversight)
    - Is human agency and oversight relevant to your work on *CONNECT* ? If yes, in what way?
    - Is human agency and oversight relevant to *CONNECT* generally? If yes, in what way?
3. Are there any particular values/areas that are missing?

For Stakeholder Set 2, the validation and evaluation steps would track to the following questions:

1. Which of the 7 key requirements are relevant to *CONNECT* and which are relevant to CCAM generally?

2. Rank them in importance to *CONNECT* and rank them in importance for CCAM. Note that this question is then repeated for all seven key requirements (i.e. for Requirement #1 Human Agency and Oversight)
    - Is human agency and oversight relevant to *CONNECT* ? If yes, in what way?
    - Is human agency and oversight relevant to CCAM generally? If yes, in what way?
3. Are there any particular values/areas that are missing?

Having started with these preliminary questions the validation and evaluation step has the stakeholders go through the questionnaire for the trust mapping and seven key requirements for ALTC-CAM, given in Section 2.3. This will produce a set of assessments of the *CONNECT* work on CCAM, as well as a validation and feedback on how to update the questions and suggestions for ways to translate the seven key requirements into RTLs that can be utilised and adopted as standards to check the ATLs against.

2.5 Future Steps towards recommendation of a Trust Assessment Methodology

CONNECT is focused on developing a TAF, through consideration of three particular use cases. As such the scope of *CONNECT* is limited by design. This trust method is designed with the intention that it is applicable to CCAM and even IoT more generally, but that application requires significant adaptation. The intention here is that *CONNECT* provides a method to support the ethical analysis of trust management in CCAM. The four steps described in Section 2.3 for this ethical analysis provide a road map of how that is being used in *CONNECT* and how it can be achieved in CCAM and IoT systems more generally.

The next step for *CONNECT* is to engage stakeholder sets one and two for the validation and assessment of *CONNECT* and of the method for CCAM. The scope here is limited – we are concerned with validation of the method and the ethical assessment of trustworthiness of *CONNECT* through stakeholder groups one and two. Future work would then involve incorporation of the stakeholder feedback into the questionnaires to ensure their validity, and to give preliminary evaluations of the ethical trustworthiness of *CONNECT*Ḟollowing that, future work would involve identification and engagement with stakeholder groups three and four, with and more systematised and scientifically grounded surveys of stakeholder group 2.

Chapter 3

CONNECT Final Architectural Overview

3.1 CCAM Functionality of the Edge/Cloud Continuum

In Deliverable D2.1 [4], the *CONNECT* consortium presented a comprehensive analysis of the state-of-the-art, primarily focusing on the architectural design and components necessary to provide a dynamic and flexible Trust Assessment Framework (TAF). This framework was aimed at supporting the advanced operations needed for collective perception (*CCAM* Day 2 application), co-ordinated manoeuvres (*CCAM* Day 3 application) and beyond. Thus, such a framework requires effectively handling diverse and contradictory data (i.e., cases where one trust source reports normal operation while another flags potential misbehavior for the same entity). Additionally, the framework must address data with varying levels of uncertainty, ensuring reliable trust decisions in complex environments. To meet these challenges, ***CONNECT* envisions a technical solution that integrates advanced hardware-based trusted computing primitives, allowing *CCAM* ecosystems to operate within trustworthy, heterogeneous communication environments with a powerful trust assessment framework (*Trust Assessment Framework (TAF)*)**. This enables the continuous trust assessment of all resources and workloads deployed across this complex landscape, while facilitating the realization of operations beyond Day 2 applications that rely on the seamless exchange of rich information between vehicles and back end systems.

Admittedly, such systems go beyond traditional cloud platforms, aligning with the vision of disaggregating services across the entire compute continuum — from cloud to edge —enabling the optimal resource utilization and low-latency processing closer to the data source. In this context, edge computing is instrumental, ensuring low-latency, high-efficiency data processing by bringing computational and storage resources nearer to the source of data generation. In this deliverable, we shift our focus towards the edge computing infrastructure (called *MEC* for ease of reference) aspect of the architecture. To address the demands of such dynamic and distributed environments, **the *CONNECT* framework extends its capabilities to the *MEC* infrastructure. This extension enables more comprehensive trust assessment calculations, utilizing trustworthiness evidence received from multiple sources** (i.e., vehicles in the vicinity). Thus, the MEC infrastructure not only benefits from access to a wider array of data but also leverages its vast computational resources to support tasks. This support is critical in two scenarios: when vehicles lack the necessary resources to perform these tasks independently, or when more advanced computations are required, benefiting from the complementary information available at the MEC. The distributed nature of *MEC* allows for localized trust decisions, whenever necessary, reducing reliance on distant, centralized systems, and therefore minimizing latency. This is especially

valuable in safety-critical applications, where timely and secure trust evaluations are paramount for the seamless functioning of autonomous and connected vehicles. By integrating trust assessment capabilities into the MEC infrastructure, *CONNECT* ensures support across all layers of the framework, while adopting a unified methodology for both the calculation and extraction of trustworthiness evidence, leveraging advanced Trusted Computing capabilities.

The primary objective of *CONNECT* is to provide tools and security controls that enable the bootstrapping and dynamic assessment of trust levels across all *CCAM* hardware (HW) and software (SW) resources — from applications and execution environments to vehicle hardware, and *MEC* infrastructures. As the *CCAM* service ecosystem integrates diverse subsystems, the complexity of trust relationships grows significantly. This necessitates a comprehensive trust assessment framework capable of functioning seamlessly at both the in-vehicle level and within the MEC infrastructure. This framework can create trust opinions, based on either locally collected trustworthiness evidence or evidence received by other vehicles in the vicinity. **Towards this direction, *CONNECT* facilitates the extraction of trustworthiness evidence from the edge computing infrastructure, complementing the proposed *CONNECT* trust assessment framework, by incorporating data from all available trust sources. This ensures that trust evaluations reflect a holistic understanding of data trustworthiness across the ecosystem.** The in-vehicle TAF and the MEC infrastructure can exchange their trust evaluations, enabling collaborative trust assessments.

The present deliverable extends the architectural insights from [4] by focusing on the integration of *MEC*-specific components, which are essential for realizing a robust trust assessment mechanism within *MEC* environments. The *CONNECT* framework consists of two phases: Setup and Runtime. The Setup phase involves establishing *CONNECT*-related components, defining trust model templates, and deploying trust-related information and components as part of the *CONNECT* Trusted Computing Base. The Runtime phase involves the operation of *CONNECT* attestation schemes and security controls, such as the Misbehavior Detection service, which securely monitor trustworthiness evidence and provide to the *TAF*. In the following sections we will describe the updates that took place after the release of Deliverable [4], elaborating on the details of the updated framework and providing information about the MEC-related functionalities.

3.2 Updates to the first release of *CONNECT* Conceptual Architecture

The following table summarises the architectural updates that took place after the submission of [4] and the execution of the first integration activities.

Aspect	First Release of *CONNECT* Framework	Updated *CONNECT* Framework

Privacy Requirements for Verifiable Evidence	In the first release of the *CONNECT* framework there was a single mode of operation as it pertains to the privacy requirements of the verifiable evidence	The updated framework introduces two different levels of privacy protection, providing flexibility based on the application's needs. These profiles include: (i) No privacy protection, where attributes in the *Trustworthiness Claim (TC)* can be verified directly using signatures, and (ii) Enhanced privacy protection, which involves using harmonized attributes in a *TC* that is then signed using a Threshold Direct Anonymous Attestation (DAA) scheme.
Attestation Process (TCB)	The Attestation Agent and Tracer, part of the *CONNECT Trusted Computing Base (TCB)*, executed their versions without undergoing prior validation.	A new component, the Verifiable Policy Enforcer (VPE), is introduced into the *CONNECT TCB*. This component ensures that the versions of both the Attestation Agent and the Tracer are valid and verified before execution. This adds an additional layer of security by ensuring only verified versions are being executed.
Security Extensions for Cloud-native applications	The first version of the framework did not consider security extensions for cloud-native applications.	In the updated framework, security extensions are supported, specifically providing the necessary mechanisms for Kubernetes Key Management to securely launch the enclave to the MEC-based infrastructure
Trusted vs Confidential Computing	The original framework supported trust by utilizing enclaves, which are isolated execution environments that provide integrity and confidentiality guarantees.	The updated version extends this by also supporting trust and confidentiality through confidential containers, which provide similar security guarantees as enclaves but offer more flexibility.
Migration Capabilities	The initial framework mentioned static migration, but there were no details on the exact process where the migration of virtual machines or containers was more rigid and less dynamic.	The updated framework enables both static and live migration capabilities of an application migratable state (protected as part of an enclave) between either in-vehicle components or vehicles that demonstrate the required trust levels (i.e., one *Zonal controller (ZC)* to another).
Task Offloading Capabilities	The initial framework mentioned task offloading, but there were no details on the exact process.	The updated framework enables task offloading. Without loss of generality, it has been specified for the offloading an automotive video analytics task.
MEC Trust Assessment	Initially, the MEC infrastructure only assessed the trustworthiness of the vehicle(s) based on their transmitted information.	In the updated framework, trust assessment capabilities are extended. The MEC can now also assess its own trustworthiness (self-assessment), in addition to assessing vehicles, adding a more comprehensive approach to trust management.
ECUs Trustworthiness	In the earlier release, *electronic control unit (ECU)*s such as the *A-ECU* and *S-ECU* were considered equally trustworthy, without distinction between their capabilities, while the proposed schemes were focused on the *A-ECU* which possess advanced capabilities.	The updated framework introduces different trust weights for *ECU*s based on their capabilities. For example, the *A-ECU* (with advanced capabilities) might be considered more trustworthy than the *S-ECU* (with fewer capabilities).

Risk Assessment Framework	The initial release provided a high-level description of the Risk Assessment Framework, without clear alignment with standardized methods.	The updated version aligns the Risk Assessment Framework with standardized methodologies, such as Threat Analysis and Risk Assessment (TARA) defined in ISO/SAE 21434, ensuring compliance with industry standards and improving the accuracy of risk assessments.
Trust Assessment Framework	The first release was based on a preliminary architecture of the *CONNECT* trust assessment framework which were enhanced and became more mature as the WP3 work progressed.	The updated version incorporates a refined architecture of the standalone TAF including also aspects of the federated and digital twin TAF modalities.

Table 3.1: CONNECT Framework in different Releases

3.3 The *CONNECT* Conceptual Architecture (FINAL)

3.3.1 List of *CONNECT* Functional Components - Building Blocks

The following section provides a list of the *CONNECT* functional components, coupled with a high-level description of their functionalities, while Figure 3.1 illustrates the final view of the *CONNECT* Architecture.

1. **Master Compute Node**: It acts as the orchestrator of this distributed CCAM service architecture. *CONNECT* as a framework remains agnostic to the underlying virtualisation and orchestration technologies. However for the demonstration and evaluation activities the orchestration and deployment of these services is facilitated through the Master Compute Node following the Kubernetes virtualization and orchestration technology. More detailed description on the *CONNECT* Kubernetes infrastructure are available in D5.2 [5]. The main role of this component is to perform cluster management and configuration of resources within the cluster, deployed across all points of the CONNECT infrastructure. Relevant requirements defined by a Service provider in the form of Service Level Agreements (SLAs), can be incorporated.

2. **Management and Orchestration (MANO) stack**: It is a framework designed to dynamically control and manage 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 deployed cluster.

   This functionality supports CONNECT's trust-related operations deployed on the infrastructure, especially at the edge computing servers. Both the Master Compute Node and the

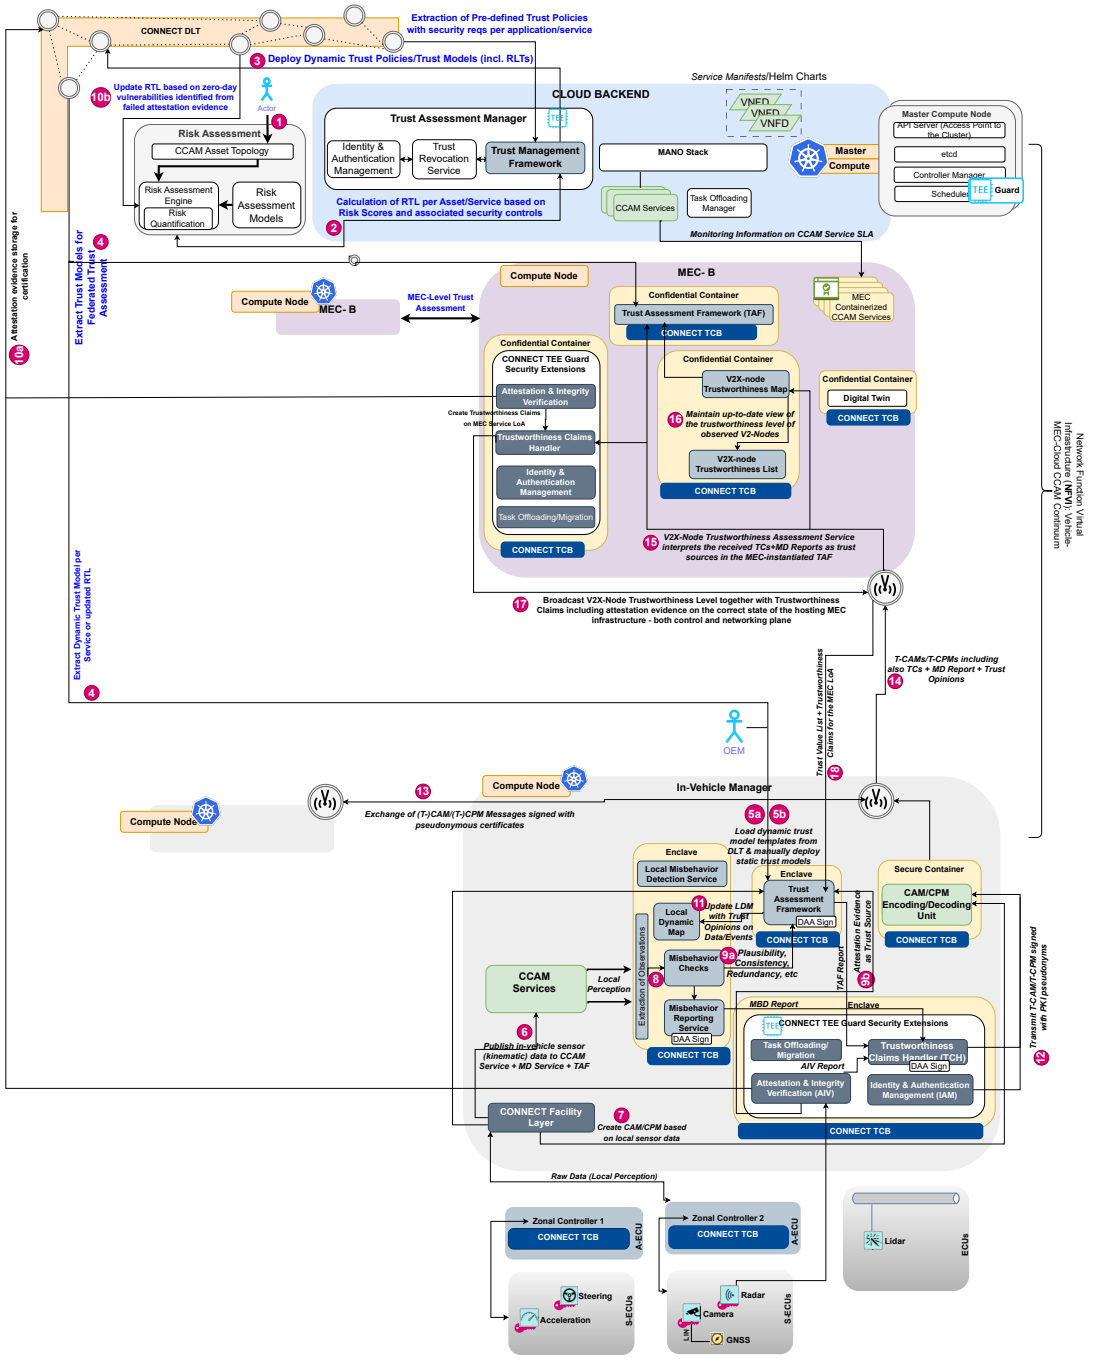

Figure 3.1: CONNECT Final Architecture

MANO Software Stack are responsible for setting up and managing the entire lifecycle of MEC-deployed services, as well as the CONNECT-related services deployed on the vehicle's main on-board unit.

3. **Risk Assessment**: Within the CONNECT project, the Risk Assessment component plays a pivotal role in evaluating, quantifying and managing potential security risks (in a secure and privacy preserving manner), 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. Such information is essential for the Risk Assessment component to perform its calculations effectively.

   The Asset Modelling and Visualization component (which is part of the Risk Assessment Component) processes the input 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 (details on the risk calculation methodology are provided in Chapter 7.2 of D3.2[7]).

   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 (TAM)**: 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 Level (RTL). The RTL is providing the benchmark to which Actual Trustworthiness Levels (ATLs) will be compared and is defined based on the minimum trust required for operation to resume securely / safely. This process is essentially a methodology that will be followed by the Security Administrator (for instance) 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 dissemination to authorized entities, such as other vehicles, edge infrastructure components, that rely on this information for decision-making or trust verification. In addition to the RTL, Trust Policies and Trust Models are also 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 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 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 enforced 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 available 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 central 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), initiatives the evaluation of the specific proposition by the TLEE (see the bullet below), and provides the result either to the requesting application or to the TDE for decision taking.
    - **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 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 a TAF and its state. More specifically the Digital Twin replicates the TAF, including its TMs and TSs in the MEC. This allows a vehicle to outsource

trust assessment process to the MEC where the TAF-DT is expected to run inside a TEE so that its data and state can be kept confidential from the host. The CONNECT Digital Twin facilitates in-depth analysis, scenario testing, and informed decision-making, contributing to the optimization and advancement of CCAM capabilities.

7. **Task Offloading Service**: This service enables the distribution and delegation of computational tasks from in-vehicle systems to the infrastructure e.g., the MEC servers, 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 edge 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 important trust-related state information. 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. Additionally, failed attestation evidence derived from the Vehicles are stored on the DLT so that OEMs may conduct further research.

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 computer. 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-compliant 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 Claims 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 Assessment reports 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) -see below, entry 11 of the list- based on Node Trust Level (NTL). Additionally, the TCH plays a crucial role in facilitating privacy-preserving exchanges by ensuring that vehicle fingerprinting is not possible, while also guaranteeing the trustworthiness of the exchanged messages.
- **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. ***CONNECT* Trusted Computing Base (TCB)**: The components that provide the *CONNECT* TCB are the following:

    - **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.
    - **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 are documented in the D4.1 [2] and D4.2 [8].

CONNECT adopts and builds upon the classification, as defined by ETSI [12], to map the extracted traces to specific assurance levels for the virtualised MECinfrastructure, specifically for the needs of the automotive sector. The ETSI-defined LoA uses numbering from 0-5, to represent a scale of relative trust, where a greater number denotes a higher level of trust. Based on this classification, the following scaling for *CONNECT* is defined:

  - **LoA 0**: denoting the complete absence of any form of integrity verification.
  - **LoA 1**: covering the local integrity verification (i.e., based on signatures) of the hardware and virtualization platform's (hypervisor) during boot and application loading. No proof of integrity is offered. Integrity status is derived from platform state after the end of boot and application load processes.
  - **LoA 2**: Adding to LoA 1 the remote attestation of the hardware and virtualization platform integrity. Measurements of boot time and application load time are considered.
  - **LoA 3**: Adding to LoA 2, LoA 3 includes the local verification of the Kubernetes orchestrator and API server, comprising the network plane, (i.e., based on signatures) as they are loaded on startup.
  - **LoA 4**: Adding to LoA 3 the remote attestation of Kubernetes orchestrator and API server, comprising the network plane. Boot time measurements and application load time measurements should be used.
  - **LoA 5**: Adding to LoA 4 the remote verification of the Kubernetes orchestrator and API server (comprising the network place) integrity state during run-time (i.e. post load time).

- **Tracer**: Another core component of *CONNECT*'s TCB is the Tracer. It has two parts; the first operating in the *untrusted/normal* world, inspecting safety-critical software components, while the other part runs in the *trusted* world, where the Tracer's secret key is stored.

  Regarding the part that is executed in the untrusted world, the Tracer is responsible for continuously fetching new traces by monitoring processes and routines that are executed within the *untrusted* world of each container/device. Its primary scope is the collection of essential information for attestation methods employed in *CONNECT*, for ensuring integrity. The tracer, in essence, is capturing hashes of configuration properties from safety-critical untrusted processes and routines. *This monitoring in the untrusted world is not considered part of the TCB*.

  Nevertheless, there is a part of the Tracer's execution that is executed in the trusted world, within the TCB. This part entails the cryptographic-related operations, and more specifically: i) the decoding of the raw security measurements, ii) the calculation of the real-time configuration hash and iii) the generation of the digital signature over the configuration hash based on the secret key. The collected traces are signed by the Tracer in the trusted world and are sent to the Key Manager to perform the required operations.

It has to be noted here, that the Tracer comes with a pre-shared key pair that acts as a Root-ID key of the Tracer. The public part of this Root-ID key is known by the Identity Authentication Management (IAM) component. During the Secure Configuration of all Tracer-enabled devices/components, the IAM sends the public key of the respective Tracer to the Key Manager of each device/component. This process establishes a shared key, bound to the underlying hardware RoT, enhancing the security of communication and ensuring the integrity of the Tracer's attestation capabilities.

- **Verifiable Policy Enforcer (VPE)**: Acts as the entity responsible for ratifying the validity of the Key Restriction Usage Policy being enforced. Its primary function is to prevent potential attackers from having multiple obsolete policies active. For instance, if a past version of an application is associated with a specific policy with known vulnerabilities, the VPE ensures that this policy is identified and characterised as obsolete, preventing unauthorised or outdated policies from influencing the system's security posture and allowing an attacker to exploit them. Authorised by the *IAM*, the VPE monitors the correctness (i.e., integrity) of the software versions executed by the Tracer and the Attestation Agent (i.e., parts of *TCB*) within the enclave, and if the versions are successfully verified, then the VPE authorises the enforcement of a specific key restriction usage policy. This verification process involves the use of a key obtained (and sealed) from the IAM. If this key is correctly unwrapped and signed by the VPE, then this signifies that the policies are satisfied. More details on the flow of this newly designed protocol can be found in [8].

11. **Misbehaviour Detection**: The components included in the Misbehaviour Detection within CONNECT are as follows:

    - **Local Misbehaviour Detection Service (Local MDS)**: It examines CAMs, CPMs, and other in-vehicle observations to identify any conflicts (in the involved data) 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 the construction of the *Verifiable Presentation (VP)* exhibiting the necessary level of abstraction for the zero-knowledge sharing of trust-related evidence. 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 (see item 12). 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.

12. **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. A non-exhaustive list of such information that is managed within the context of the envisioned use cases (i.e., IMA, C-ACC and SMTD) is as follows:

    - **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.

13. **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.

14. **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.

3.4 *CONNECT* Information Flows

This section offers a comprehensive, step-by-step analysis of the information flows explored by *CONNECT* to capture both *CCAM* service-related and Trust Assessment-related processes. As previously highlighted, the primary goal of *CONNECT* is to establish a Trust Assessment Framework (TAF) based on verifiable trustworthiness evidence, supporting secure operations beyond Day 2 scenarios. To that end, it is important to clarify that the trust assessments detailed in Table 3.2 are integral to the framework's execution. The subsequent sections will provide further insight into the various trust assessment approaches discussed, thereby deepening the understanding of their distinct attributes and application within the framework.

No	Assessment Approach	Details
1	Vehicle trust assessment on itself	The in-vehicle *TAF* proceeds with its internal trust assessment based on its Trust Model (TM) and the locally available trustworthiness evidence on its own state.
2	Vehicle trust assessment on other vehicles	The in-vehicle *TAF* conducts a trust assessment over another vehicle in the vicinity, leveraging its own Trust Model (TM), local trustworthiness evidence along with the *TC*s as received from the T-CAM and T-CPM message about a remote vehicles.
3	Vehicle trust assessment on data	The in-vehicle *TAF* proceeds with a trust assessment over a specific data item, leveraging its own Trust Model (TM), local trustworthiness evidence along with the *TC*s as received from the T-CAM and T-CPM message about data perceived locally and/or received from remote vehicles.
4	MEC trust assessment on itself	The *MEC*-based *TAF* proceeds with its internal trust assessment based on its Trust Model (TM) and the extracted trustworthiness evidence.
5	MEC trust assessment on vehicle(s) [has its own (TM)]	The *MEC*-based *TAF* proceeds with a trust assessment over a vehicle, leveraging its own Trust Model (TM) along with the *TC*s as received from the T-CAM and T-CPM message.

6	DT-TAF on MEC [has the vehicle's TM]	A vehicle offloads the trust assessment task to the *MEC*. Hence the assessment that takes place on the *MEC* leverages the Trust Models of the vehicle.

Table 3.2: CONNECT Framework's different flavours of trust assessment

3.4.1 Vehicle-related flow for enabling Trust Assessments

In the in-vehicle context, two parallel processes occur: i) the CCAM application/service flow and ii) the *CONNECT*enabled continuous trust assessment flow.

3.4.1.1 CCAM application/service flow and in-vehicle relevance

As it pertains to the first flow concerning the CCAM application/service, this is mainly related to the Application Layer. More particularly, this layer includes components such as i) the *CCAM* Services portfolio (i.e., VRU Protection, Trajectory/manoeuvre sharing, coordination/negotiation, Intersection Movement Assistance (IMA), Advanced C-ACC), ii) the CAM/CPM Encoding/Decoding Component and iii) Misbehaviour Detection Services (MBD), executed within the in-vehicle computer. The CCAM application flow begins with a message from the Application Layer to the *Facility Layer (FL)*, requesting kinematic data (step 1). The *FL*'s role involves associating raw sensor data with corresponding device identifiers, thereby enabling the Application Layer to generate the vehicle's Local Perception (LP). To achieve this, the *FL* can either directly obtain kinematic data from available devices such as *A-ECU*s and *S-ECU*s (steps 2 and 5) in simpler topologies **(option #1)** or, in more complex topologies—where *Zonal controller (ZC)*s or *A-ECU*s manage multiple *S-ECU*s—the *FL* forwards the kinematic data request to the *S-ECU* via the *A-ECU* that oversees it (step 6) **(option #2)**. The *S-ECU*, upon receiving the request, gathers the relevant kinematic data related to CCAM vehicle observations and sends it to the requesting *A-ECU* (step 9). The *A-ECU* then combines its own kinematic data with the received data and transmits it back to the *FL* (step 10).

This kinematic data includes raw information on the vehicle's position, speed, direction, and other parameters. The *FL* aggregates this data and delivers it to the Application Layer, where it is used to construct the vehicle's LP, representing its understanding of the surrounding environment. This LP supports decision-making processes, such as obstacle avoidance, and can also be integrated into the CAM/CPM payload for sharing with nearby vehicles. Additionally, the Local *Mis-behaviour Detector (MBD)* Service analyzes the LP to identify discrepancies or conflicts in the data (step 13), ensuring the reliability and trustworthiness of the CCAM ecosystem. The MBD report is sent back to the Application Layer (step 12), further enhancing its decision-making capabilities. The aforementioned flows are illustrated in Figure 3.2.

3.4.1.2 *CONNECT* Enabled In-Vehicle Continuous Trust Assessment Flow *(for the vehicle's self-assessment)*

The second parallel flow centres on executing the trust evaluation procedure within the in-vehicle topology. This process begins with a Trust Assessment Request (TAR) sent by the Application Layer to the Trust Assessment Framework (TAF) (step 1). Upon receiving the request, the TAF utilizes its Trust Model (TM) to determine the required evidence and proceeds to collect it from

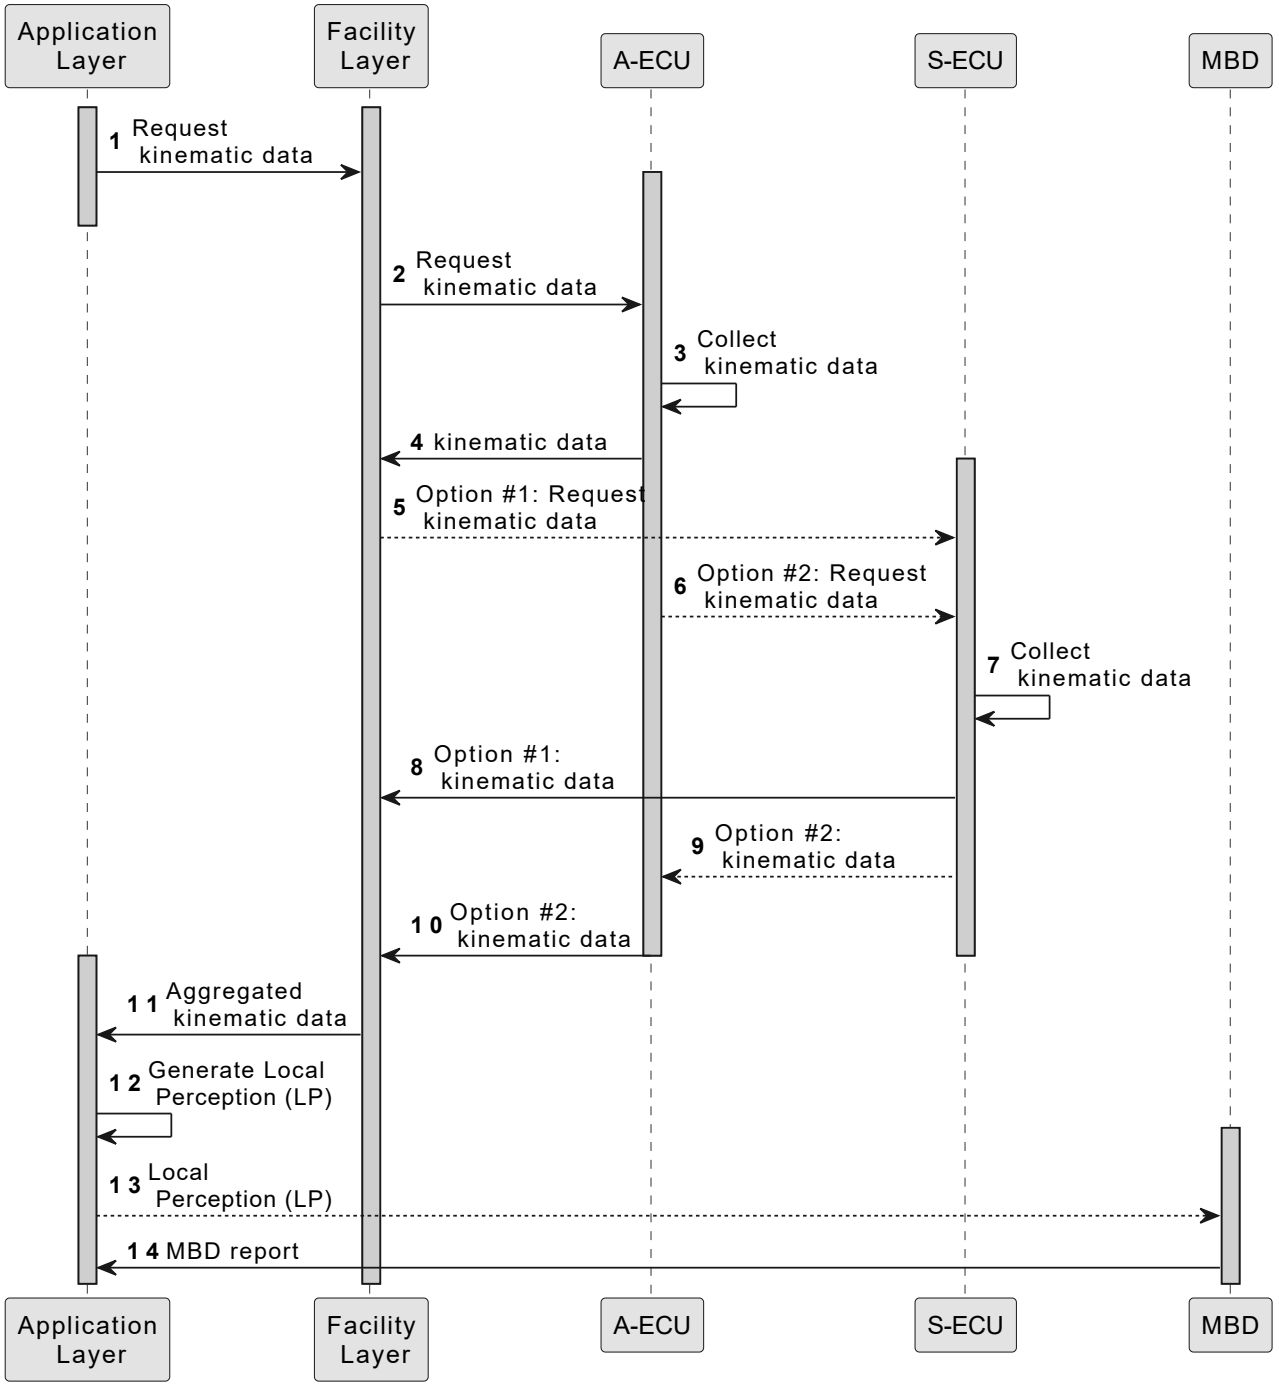

Figure 3.2: CONNECT In Vehicle CCAM application flow

Trust Sources (TS). The TM defines the involved entities (e.g., *Zonal controller (ZC)*, *ECU*s) and the tasks for which trustworthiness evidence are needed (e.g., Configuration Integrity Verification (CIV), Secure Boot, etc.). Evidence collection can occur in two modes: (i) **synchronously** (pull-based evidence collection), where the TAF directly queries the trust source for the necessary evidence, or (ii) **asynchronously** (subscription-based evidence collection), where the TAF subscribes to receive ongoing updates from the trust source, often facilitated through mechanisms such as a Kafka channel. More details on the exact definitions of the messages exchanged per TAF mode is available in D6.1 [6].

In a scenario where the TM dictates the need for attestation evidence collection, the *TAF* will

request them from the *AIV* (Step 2). For protection against replay and other types of attacks, this request also includes a nonce and a signature. Upon receiving the request, the *AIV* initiates the evidence collection process by instructing the *Zonal controller (ZC)* or the *A-ECU* to perform their attestation tasks (step 3). In cases where the *A-ECU* oversees other components, it may further trigger the attestation tasks for the *S-ECU*s under its control (step 5). For the *A-ECU*, the Attestation Agent collects and verifies traces associated with various properties, such as Configuration Integrity Verification (CIV), leveraging the advanced capabilities provided by the underlying Root of Trust (RoT) (Step 4). Meanwhile, within the *S-ECU*, the Attestation Agent focuses on verifying the state of the device, such as ensuring the integrity of the Secure Boot process (step 6). The signed reports from the Attestation Agents of the *S-ECU*s are transmitted to the *A-ECU* (step 7), which aggregates them and sends the aggregated signature shares to the *AIV* for verification (step 8). The *AIV* first verifies the threshold shares received from multiple *A-ECU*s (step 9) and proceeds with verifying the aggregated signature shares originating from the *S-ECU*s (step 10). These steps lead to the construction of the attestation report (step 11), which the *AIV* then shares with the TAF (step 12). Additionally, this report is signed with a Direct Anonymous Attestation (DAA) local key and sent to the *TCH* (step 13).

With all the necessary evidence in hand (from the AIV and possibly the MBD) the TAF is able to conduct its Trust Assessment (TA) and provide a DAA signed report to the TCH (step 14). Similarly, the *MBD* may send its own report, signed with DAA local key, to the *TCH* (step 15). The outputs handled by the *TCH* are governed by privacy requirements stipulated in the Trust Model (TM), which defines two privacy levels: (i) sharing attributes without explicit privacy considerations and (ii) sharing attributes with enhanced privacy protection, leveraging threshold DAA signatures. The latter approach provides a higher level of abstraction to safeguard vehicle privacy and mitigate risks such as vehicle fingerprinting attacks. The *TCH* plays a key role in harmonizing the received Trustworthiness Claims (TCs) from the *AIV*, *MBD*, and TAF (step 16). It then generates a DAA credential incorporating these attributes (step 17). This credential is transmitted to the *IAM*, where it is transformed into a Verifiable Credential (*Verifiable Credential (VC)*) and signed with a PKI-based pseudonym (step 18). The aforementioned information is then forwarded to the CAM/CPM Encoder/Serialiser (step 19). The final step prior to transmitting the TCs to the MEC or other vehicles involves serializing the T-CAM/T-CPM (i.e., *TC* enabled, CAM or CPM) message (step 20). The aforementioned flows are illustrated in Figure 3.3.

3.4.1.3 Vehicle-to-Vehicle Trust Assessment Flow *(for trust assessments on other vehicles based on received information OR for assessing specific data received from other vehicles)*

In addition to the trust assessment process that relies on evidence gathered from the vehicle's own sensors, the vehicle's *TAF* can also evaluate trust based on *Trustworthiness (claim-enabled) Cooperative Awareness Message (T-CAM)* and *Trustworthiness (claim-enabled) Collective Perception Messages (T-CPM)* messages received from nearby vehicles as part of the standalone TAF, while it can further leverage a Trust Opinion (TO) or ATL as exported by another TAF. These assessments enable the vehicle to form trust opinions about the behavior and trustworthiness of other vehicles in its vicinity, thereby enhancing both security and collaborative operations within the vehicular network. More specifically:

Option #1: Trustworthiness Claims exchanged among Standalone TAFs: This mode is designed to independently analyse the trustworthiness of individual entities within a localised context. The standalone *TAF* operates autonomously in a specific environment, such as a vehicle

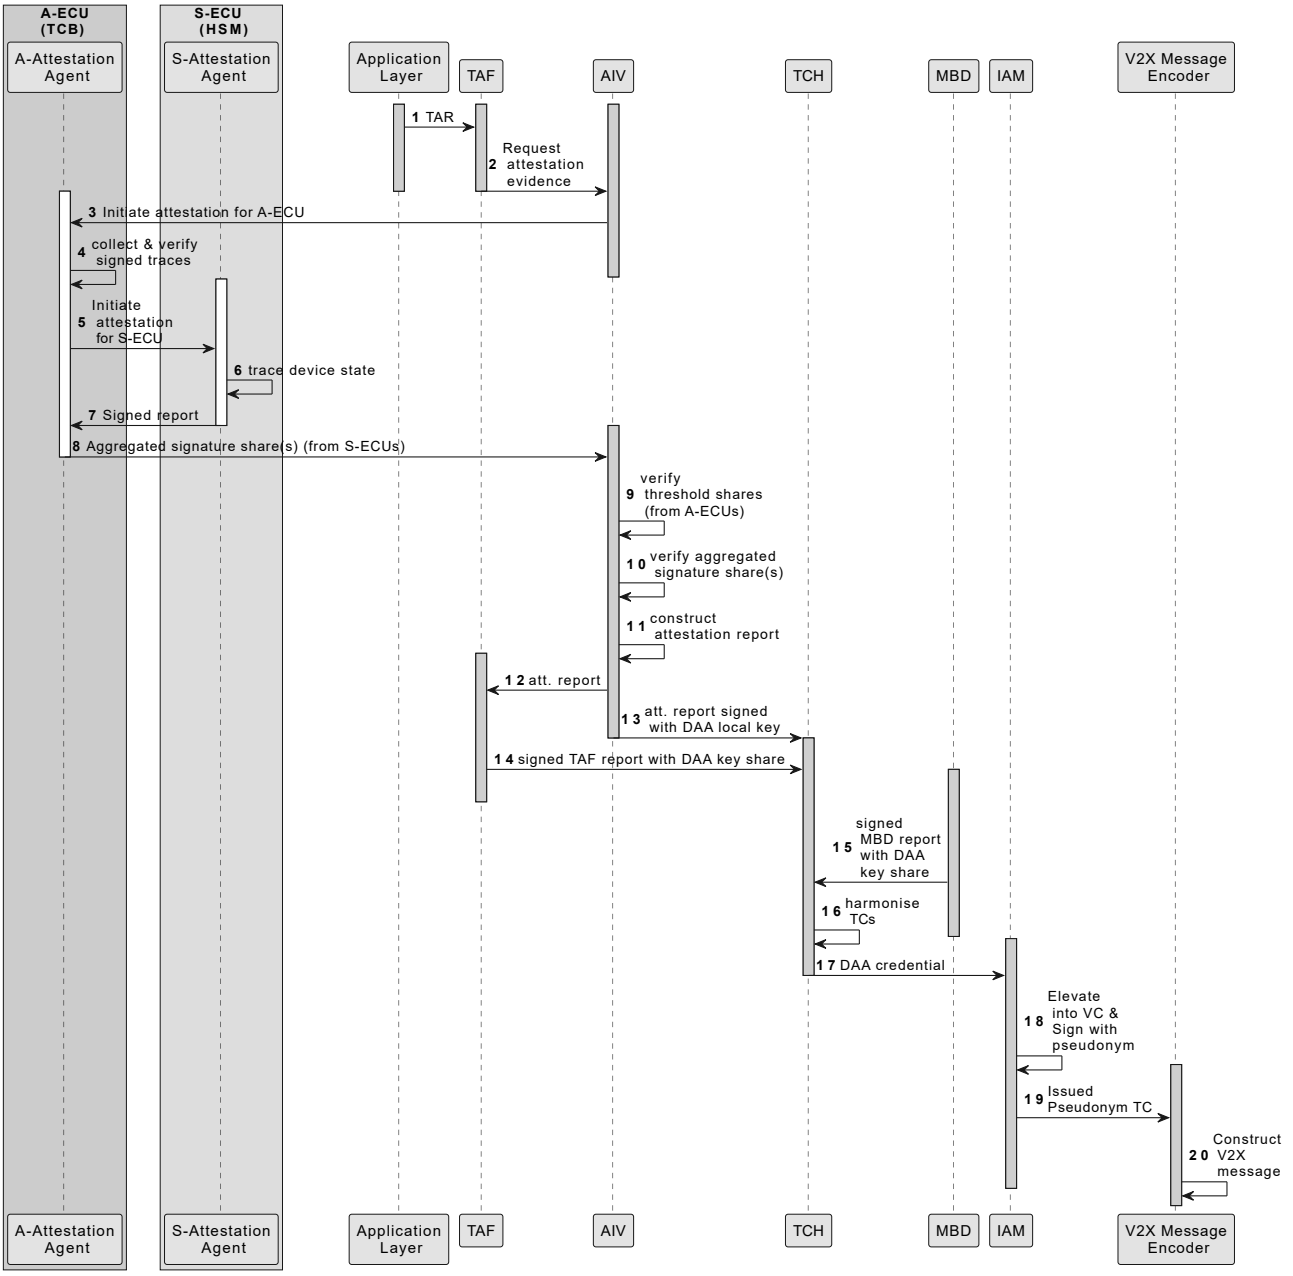

Figure 3.3: CONNECT In Vehicle Trust Assessment flow

or MEC, evaluating trust for both local and remote entities solely from its own perspective using internal assessments.

In this case, the T-CAM or T-CPM message are received via the V2X Communication Interface in the vehicle or the MEC. These messages are then transmitted the *IAM* for pseudonym verification. If the pseudonym is successfully verified, the message is subsequently forwarded to the *TCH* and if the DAA signature is successfully verified then the TCH afterwards the TAF-related information of the TC to the *TAF* to leverage it for its assessments.

Option #2: Trust Opinions or Actual Trustworthiness Levels (ATLs) exchanged as part of the Federated TAF: This mode expands the scope of trust assessment by enabling different TAFs to exchange trustworthiness information. The federated *TAF* works collaboratively with other *TAF*s deployed across the entire CCAM ecosystem, including other vehicles and MECs. Through

this cooperation, it exchanges (upon request) trust-related information, such as trust models and opinions, to form a comprehensive understanding of trust across a distributed CCAM "System-of-Systems." This federated approach enables a more holistic and unified trust assessment by incorporating diverse perspectives.

In this context, the vehicle's *TAF* interacts with the *TAF* of another vehicle by exchanging critical trust-related data. This includes an Atomic Trust Opinion (ATO) for a specific trust source, a Trust Opinion (TO) representing a trust relationship, or a calculated Actual Trustworthiness Level (ATL), which may be provided upon request by the initiating *TAF*. Further details on this process are provided in Deliverable D3.2 [7], with additional refinements outlined in D3.3.

3.4.2 MEC-level Trust Assessment

Vehicles may share kinematic data along with the *TC*s with other vehicles as well as the *MEC* infrastructure. The MEC, in turn, uses the information gathered from multiple vehicles to generate a more comprehensive representation of the environment. Notifications regarding road events, based on this enhanced perception, can be shared with the vehicles in the vicinity through *Decentralised Environmental Notification Messages (DENM)* messages. In addition to kinematic information, DENM messages, if extended, may also include trust-related information, reflecting the MEC's assessment of nearby vehicles. These messages are referred to as *Trustworthiness (claim-enabled) Decentralised Environmental Notification Messages*. Hence, apart from the trust assessment of the vehicles, the MEC may also provide its own infrastructure's trustworthiness.

It should be clarified that MEC-level based assessment scope is threefold:

1. **Trust Assessment of the Vehicle(s) based on the Received T-CAM/T-CPM Messages or the Trust Opinion/Actual Trustworthiness Level (ATL)**: The MEC-TAF can construct a more refined representation of the environment by utilizing T-CAM and T-CPM messages received from multiple vehicles, which include both trust-related information and kinematic data. As a result, the MEC-based TAF gains a more accurate and comprehensive understanding of the CCAM environment. Its purpose is to provide a trust opinion on a specific entity (e.g., Vehicle A), which can then be used by another vehicle (e.g., Vehicle B) in combination with its own assessment to support informed decision-making. Another option beyond the exchange of Trustworthiness Claims involves the exchange of Trust Opinions (TOs) and Actual Trustworthiness Level (ATL) as calculated by other TAFs leveraging the Federated TAF option.Towards this direction one TAF may directly request another TAF for an opinion.

2. **Trust Assessment of the MEC Infrastructure**: The MEC-based TAF also evaluates the trustworthiness of the MEC infrastructure, particularly where critical services (e.g., traffic control) are deployed. For this assessment, the MEC-TAF requests a report from the MEC-based AIV, which is responsible for collecting and verifying attestation evidence of the container(s) hosting the service(s). Based on this report, the MEC-TAF conducts its infrastructure trust evaluation.

3. **DT-TAF Trust Assessments**: In this approach, the vehicle delegates its trust assessment tasks to the *MEC* infrastructure, where the DT-TAF resides. The DT-TAF, acting as an extension of the vehicle's own *TAF* leveraging its Trust Models and Trust Sources. By offloading trust assessment functions to the MEC, the vehicle may leverage the computational resources of the edge for its trust evaluations.

The subsequent sections will examine the steps related to each trust assessment option, offering a thorough analysis of their scope, flow and internal mechanisms used per case.

3.4.2.1 MEC-to-Vehicle Trust Assessment Flow *(for MEC's trust assessment of vehicles using its own trust model)*

In this scenario, the goal is for the MEC to assess one or more vehicles based on the received T-CAM/T-CPM messages. The process begins when the MEC receives these messages from the vehicle(s) via the V2X Communication Interface (steps 1 and 2). These messages not only include kinematic data but also contain the *TC*, which encompasses: i) the vehicle's *TAF* report, ii) the vehicle's *MBD* report, and iii) the vehicle's *AIV* report, all encapsulated within a *VP*.

Similar to the In-Vehicle Trust Assessment described in Section 3.4.1, the MEC's trust assessment follows distinct flows for: i) CCAM application/service flows and ii) *CONNECT*-enabled continuous trust assessment flows, that are illustrated in Figure 3.4. Both flows begin with verifying the PKI pseudonym used to sign the T-CAM/T-CPM messages. To initiate this verification, the received messages are forwarded to the MEC's *IAM*, a *CONNECT* component responsible for verifying PKI-provided signatures (step 3 and 4). Once the pseudonym is successfully verified, the process proceeds with the following steps:

CCAM application/service flow: After the *IAM* successfully verifies the T-CAM and T-CPM messages, the messages are sent to the Application Layer where the CCAM service is being executed, for its decision-making. Since the MEC-based *MBD* is integrated within CCAM services, upon receiving the report from the *TCH* via the *IAM*, it validates it and performs its own checks for potential misbehavior in the received CAM/CPM messages. Following these checks, the results are forwarded to the MEC-based *TAF* for incorporation into its evaluations, as further described in the next paragraph.

CONNECT* enabled MEC Continuous Trust Assessment Flow *(for MEC's self-assessment): The MEC's assessment of the vehicle can either rely on i) the standalone TAF option, where the vehicle shares its *TC*, or ii) the Federated TAF, where the vehicle shares a Trust Opinion based on a Trust Source or a Trust Relationship.

Option #1 Trustworthiness Claims Exchanged Among Standalone TAFs: In this flow, the IAM first verifies the PKI signature and forwards the DAA credential to the TCH for verification (steps 5 and 6). If the DAA signature is successfully verified then the vehicle's TAF report (included in the TC) is transmitted to the MEC-based TAF for consumption in its own trust assessments (step 7).

Option #2 Trust Opinions or Actual Trustworthiness Levels (ATLs) exchanged as part of the Federated TAF: The Federated TAF option provides an alternative approach for trust assessment, aligned with the rationale used by the vehicle's TAF as discussed in Section 3.4.1.3. In this mode, upon request (step 8), the vehicle's TAF directly shares its Atomic Trust Opinion (ATO) for a specific Trust Source or a Trust Opinion (TO) for a trust relationship with the TAF on the MEC (step 9), which then incorporates this information into its assessment process.

Dynamic Map Update: In addition to the TAF report, the Trustworthiness Claim (TC) includes the Misbehavior Detection (MBD) report. This information is forwarded from the TCH to the MEC-MBD module (step 10), enabling an update of the local dynamic map (step 11). The MBD report is also shared with both the MEC-TAF (step 12) and the MEC Application

Layer (step 13). It is important to note that *within the* CONNECT *framework, the MEC Application Layer does not perform decision-making based solely on the MBD report at this stage. Instead, it awaits the trust assessment outcome from the MEC-TAF (step 15).* Consequently, the trust assessment results, combined with the MBD report, are used to inform decisions (step 17), ensuring a comprehensive evaluation of trust and misbehavior data.

Trust Assessment: The MEC-TAF initiates its trust assessment by analyzing all relevant information received from the TCH as received from the vehicles, the MEC-MBD as well as the MEC-AIV (the latter is elaborated in Section 3.4.2.2) (step 14). The result of this process is a TAF report, which can serve two purposes: it may be directly used by the MEC Application Layer, particularly for safety-critical applications (step 15), or it may be transmitted to the vehicle's *IAM*, which holds the authority to enforce related policies. In the latter case, where the trust information needs to be shared with nearby vehicles, the TAF forwards the report to the TCH (step 16). The MEC's *TCH* constructs a Trustworthiness Claim (TC) based on this information and then forwards it to the MEC's *IAM* for signing (step 19). In addition to the TC, the MEC Application Layer further sends the Notification derived from its decision-making (step 18). Unlike in-vehicle processes, the MEC-TCH does not include DAA credentials, as anonymity is not required at this stage for the MEC. Once signed by the *IAM*, the data is transmitted via the V2X Communication Interface for dissemination to nearby vehicles (steps 20 and 21). Upon reception, this information is processed by the receiving vehicle's V2X encoder and passed to its *IAM* and Application Layer for integration into its decision-making processes. The aforementioned flow is illustrated in Figure 3.4.

3.4.2.2 *CONNECT* Enabled MEC Continuous Trust Assessment Flow *(for MEC's self-assessment)*

Similarly to the vehicle case, the MEC is capable of performing its own internal trust assessment. This functionality can be realized in multiple scenarios, such as when a new T-CAM/T-CPM message is received, during the creation of a new DENM message—where the MEC's TAF report is included—or whenever an application or the Digital Twin (DT) initiates an offloading task. In the context of offloading tasks, the trust assessment ensures that the vehicle is informed in advance about the trustworthiness of the MEC node designated for task execution. This assessment guarantees secure task offloading and execution, further detailed in Section 3.4.3.1.

In this case, the MEC Application Layer initiates the MEC Trust Assessment (TA) to the TAF with the deployment of each service (step 1). The TAF, guided by its internal Trust Model (TM), sends a Request For Evidence (RFE) to the MEC-*AIV* (step 2). The *AIV* then initiates the collection of evidence from both the underlying infrastructure and the (virtualised) CCAM services that are executed. The attestation report is transmitted from the MEC-AIV to the MEC-TAF (step 3), enabling the latter to proceed with the MEC's trust assessment (step 5). Simultaneously, the report is also sent to the MEC Application Layer to support decision-making. A notification, based on the Application Layer's decision, is integrated into the V2X message that will be broadcasted to vehicles (steps 7 and 8).

In addition to the attestation report, the Actual Trustworthiness Level (ATL) is sent from the MEC-TAF to the MEC-TCH (step 9), where it is incorporated into the Trustworthiness Claim (TC) transmitted to the vehicle. This ensures that the vehicle receives detailed and reliable trust information about the MEC, enabling secure and trustworthy interactions. The MEC-TCH forwards the TC to

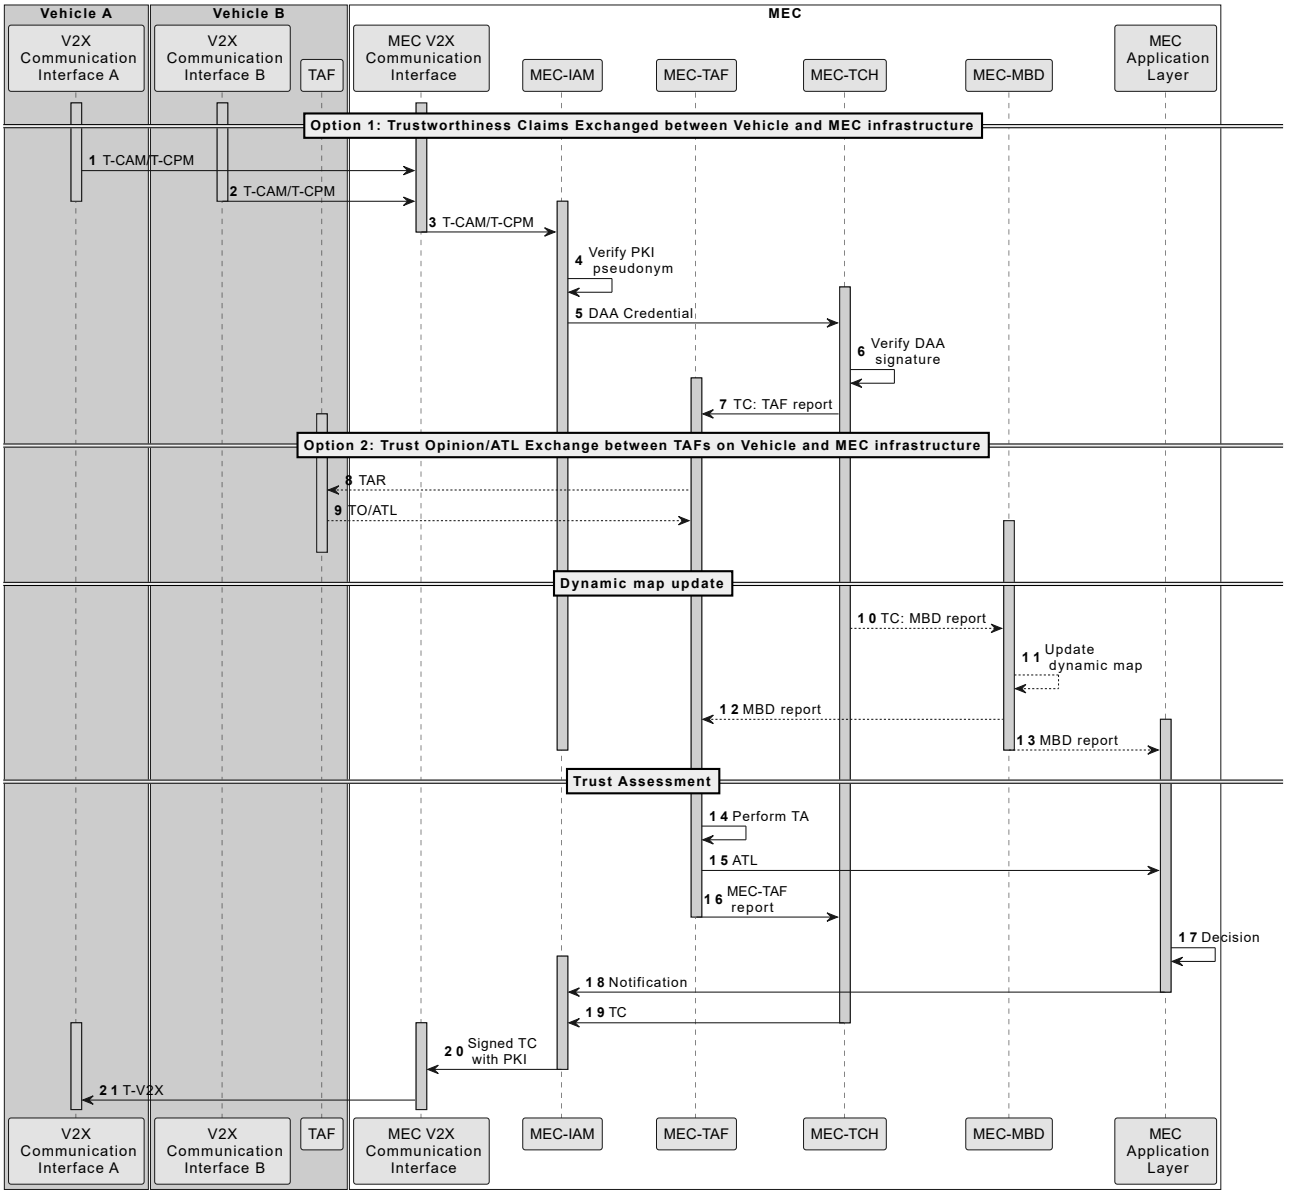

Figure 3.4: CONNECT MEC Trust Assessment flow based on vehicle's TCs or TOs

the MEC-IAM for signing (step 11). Finally, the signed message is transmitted to the vehicle via the V2X Communication interface (steps 12 and 13). This process is illustrated in Figure 3.5.

3.4.2.3 DT-TAF Enabled MEC Trust Assessment Flow *(MEC's trust assessment using the vehicle's trust model)*

In addition to the standalone and federated TAF options, *CONNECT* introduces the capability to offload the trust assessment function to its Digital Twin (DT) counterpart at the MEC level. This enhancement enables the entire TAF operation for the vehicle to be executed at the MEC, leveraging the vehicle's Trust Models and Trust Sources for trust evaluation. Such an approach ensures seamless trust assessment for the vehicle, providing it with a TAF report even under constrained resource conditions. Further details on the DT-internal architecture and operations is described in Section 3.4.4.

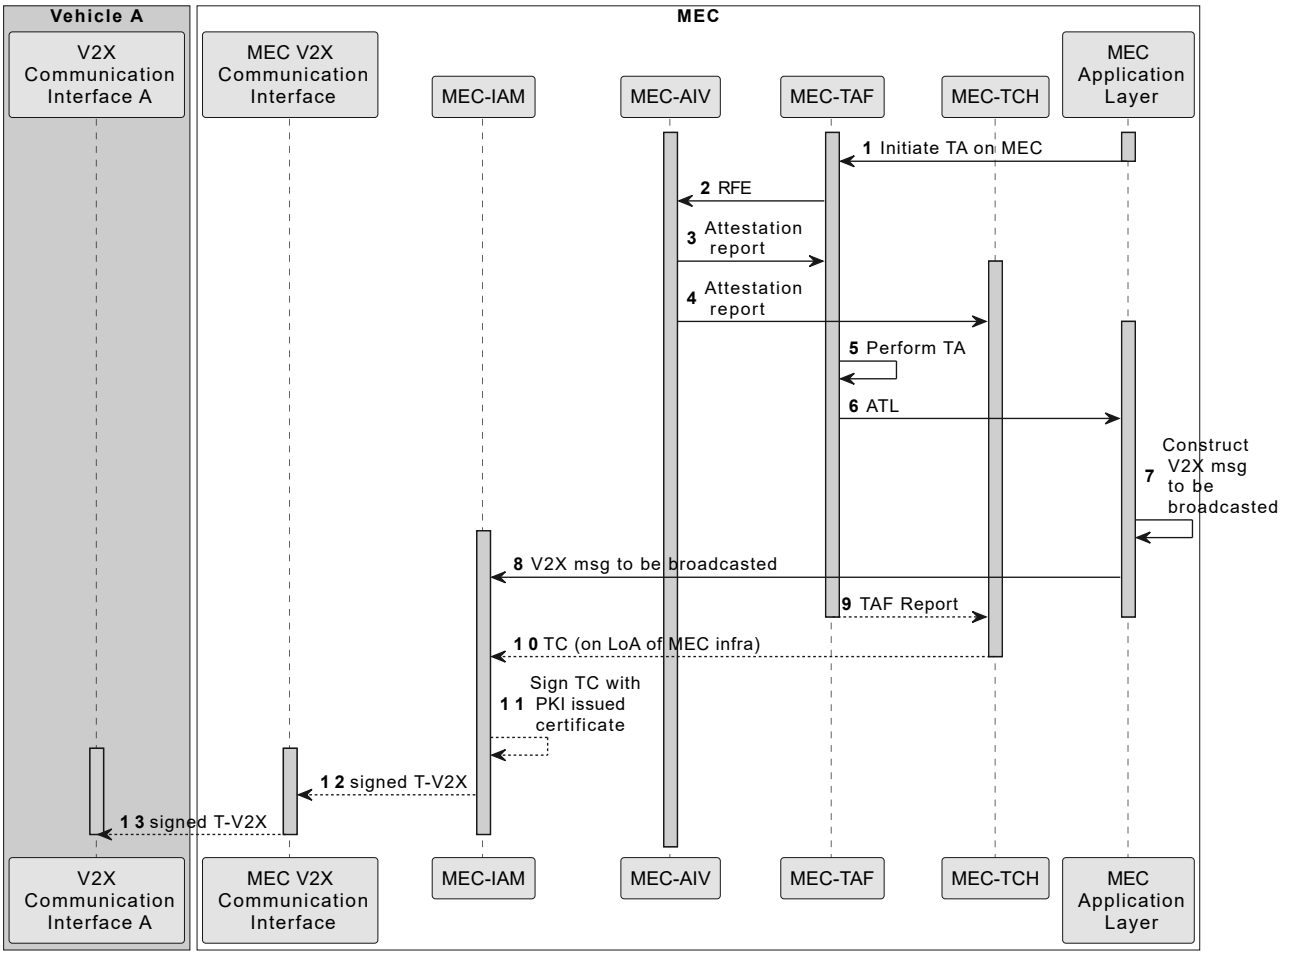

Figure 3.5: CONNECT MEC internal Trust Assessment flow

In this scenario, the vehicle's Application Layer first checks its internal policies and initiates the offloading process accordingly. However, before offloading the task, the MEC's Actual Trustworthiness Level (ATL) must be verified. To accomplish this, the vehicle's TAF sends a Trust Assessment Request (TAR) to the MEC TAF through the TCH. The MEC's internal Trust Assessment process adheres to the rationale described in Section 3.4.2.2. Upon receiving the ATL of the MEC, the vehicle's Task Offloading Module evaluates the trustworthiness of the MEC node. Based on this evaluation, it determines whether to proceed with migrating the task, ensuring both secure and trustworthy task execution.

3.4.3 Specifying the CONNECT Task-Offloading pipeline and its MEC relevance

In addition to the previously discussed flows that facilitate the trust assessment process at both the vehicle and MEC levels, another critical feature of *CONNECT* is its **secure task offloading capabilities**. These capabilities are specifically designed to address the challenges posed by resource constraints on the vehicle side in CCAM scenarios. Task offloading in *CONNECT* serves two primary objectives: i) to **enhance service provision**, leveraging the computational and collaborative resources of the *MEC* and ii) to **optimise resource utilisation**, leveraging additional computational and communication resources available to the *MEC*, thereby maximizing overall

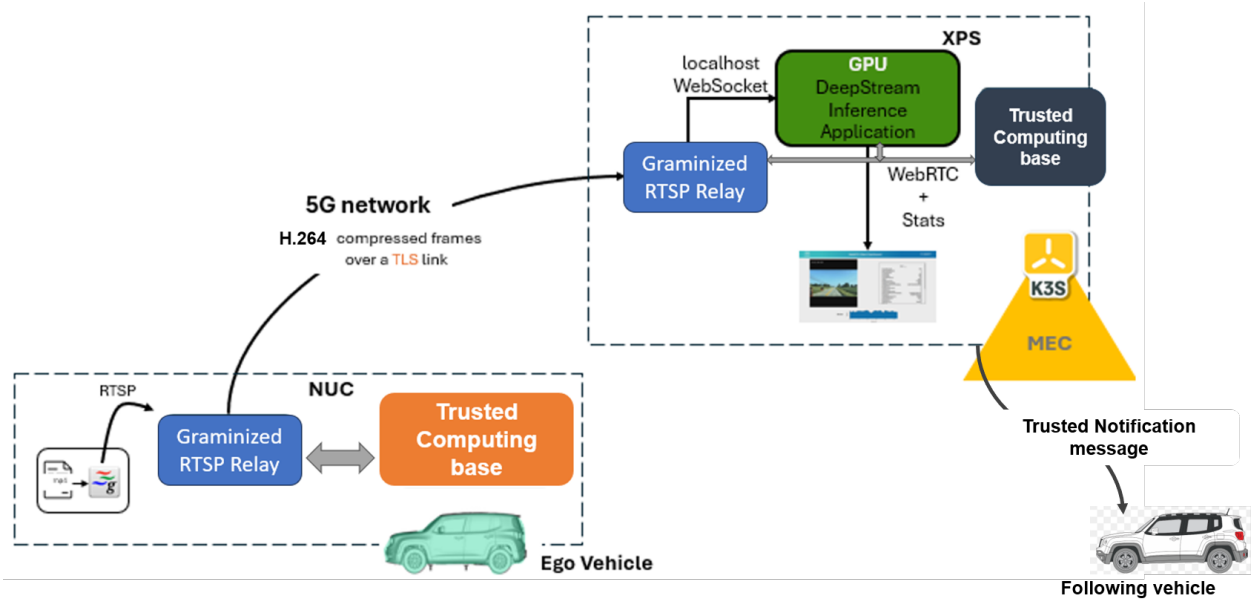

Figure 3.6: The CONNECT task-offloading scenario for video analytics service (fully in line with the SMTD use-case)

system efficiency and balancing workload distribution across the infrastructure.

A key-point in the *CONNECT* concept involves task-offloading to the *MEC* infrastructure carried-out following our guiding zero trust principle This principle assumes that the infrastructure, including the MEC, is inherently untrustworthy and requires trust to be dynamically established during the process. While task-offloading has been extensively explored in simulation environments [10], its real-world implementation—particularly with integrated trust mechanisms—remains largely unexplored. *CONNECT* investigates task offloading in two key areas: i) **resource-intensive operations**, where offloading of operations that require significant computational resources allows the vehicle to perform other safety critical tasks and ii) **trust assessment tasks**, where delegating trust-related computations to the MEC, enhances system's efficiency and security. This section focuses on the first case, addressing resource-intensive operations, while the second case, involving trust assessment tasks, is analysed in Section 3.4.4.

For demonstration and evaluation, the Slow-Moving Traffic Detection (SMTD) use case was selected as a representative scenario of the first category. However, the proposed methodology is generalisable to any high-demand application where offloading to the MEC could improve performance. To provide further details, the selected offloading scenario, depicted in Fig. 3.6, features an ego vehicle equipped with a dedicated camera that streams video frames securely to a machine learning (ML) application hosted at the *MEC*. Video streaming is inherently resource-intensive, making it an ideal case for testing network capabilities and showcasing the *CONNECT* framework's ability to handle high data demands (frames per second, FPS) while maintaining trustworthiness. The CONNECT MEC is capable of providing GPU accelerated decoding and inference, both shown to result in higher FPS [23] and lower latency [21]. The intelligent application (see green module) after being appropriately (pre)trained with relevant data, infers the presence of vehicles in-front of the ego (in the same or other lanes). This inference complemented with appropriate estimation of the distance allows the edge application to identify a potential slow-moving vehicle/road congestion (ahead) and accordingly, inform/recommend, in a trustworthy way, the following ego vehicles to change lane.

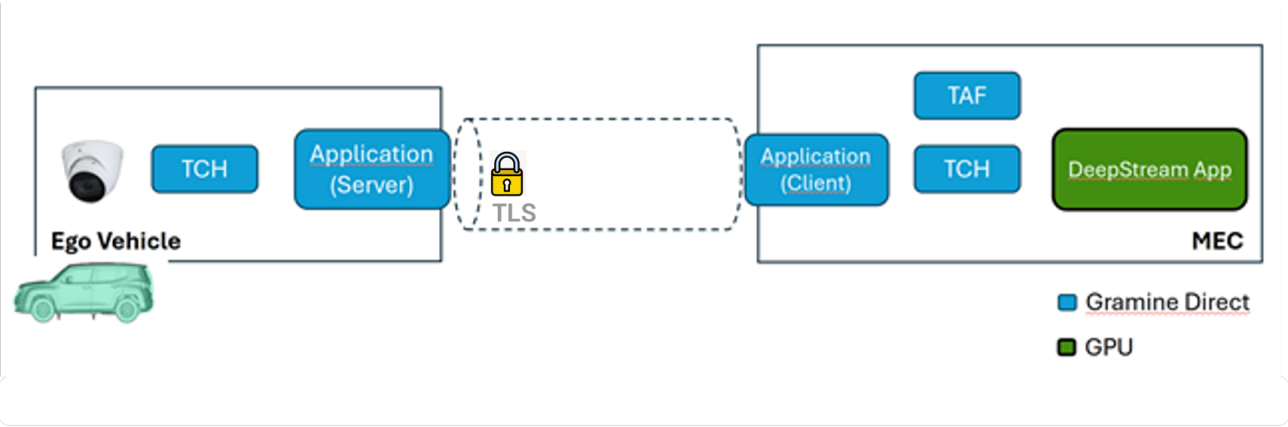

Figure 3.7: The CONNECT task-offloading basic (vehicle-MEC) architecture

This task-offloading scenario serves as an adaptation of the 5GAA "Lane Change Warning" use case [1]. Unlike the original 5GAA case, the *CONNECT* implementation emphasizes the pivotal role of edge computing in facilitating the decision-making process. The *CONNECT* framework further ensures that both the video-analytics data streamed to the MEC and the notifications sent to trailing vehicles, maintain a high degree of trustworthiness. This integration of trust mechanisms highlights the innovation and robustness of the *CONNECT* task-offloading pipeline in real-world scenarios.

It shall be noted that all trustworthiness evidence and relevant computations from the trust sources are provided with the support of *CONNECT TCB* (see Fig. 3.6). The *CONNECT* TCB is present in both the Vehicle Computer and the ECUs, as well as the MEC infrastructure. The acquired evidence may comprise of attestation, misbehavior and trust assessment reports, which are leveraged both by the vehicle to assess the trustworthiness of the MEC infrastructure where the task is to be offloaded and by the MEC to evaluate the trustworthiness of the vehicle and its data. This dual assessment ensures a secure and reliable offloading process, adhering to the zero-trust principle inherent in the *CONNECT* framework.

3.4.3.1 The CONNECT task-offloading pipeline architecture

Architecture-wise, Fig. 3.7 provides an overview of the basic offloading scenario architecture, including the "typical" *CONNECT* modules (i.e., the TCH and the TAF) together with the communication channel between the Vehicle and the MEC. In the current setup, a TAF instance is primarily deployed on the MEC, necessitating the transmission of evidence from the Vehicle's TCH to the MEC-based TAF via the communication channel. At a second stage, a TAF instance will also be deployed on the ego vehicle's On-Board Unit (OBU), and the involved trust model will be updated to reflect this configuration. Moreover, the MEC-based TAF instance MEC will be instantiated as part of the vehicle's Digital Twin (DT). For further details on the DT integration, refer to Section 3.4.4.

In *CONNECT*, a common assumption made to simplify the evaluated Trust Models is that the vehicle inherently trusts its onboard sensors, such as the RGB Camera. However, *CONNECT* aims to relax this assumption by incorporating mechanisms to assess the trustworthiness of the data

[1] Refer to the 5GAA technical report entitled C-V2X Use Cases and Service Level Requirements Volume I, Oct. 2020

originating from these onboard sensors. In the described scenario, the vehicle's computer serves a dual purpose: it hosts the *CONNECT* trust-related components while also executing critical CCAM services. The trust-related components are securely deployed within the trusted boundaries of isolated enclaves using Gramine, ensuring robust protection for sensitive operations. In contrast, the CCAM services typically operate in the "untrusted world" of the vehicle's computer and interact with their corresponding counterparts on the MEC. On the MEC side, trust-related services are equally safeguarded, operating within confidential containers to maintain a secure environment for critical interactions and computations. More details on the secure launching and deployment of (confidential) containers are described in D4.2 [8].

Although CCAM services are generally not executed within confidential containers, for the offloading task in the SMTD use case, the CCAM service is deployed within a confidential container. This setup ensures secure streaming of video data from the vehicle to its counterpart on the *CONNECT* MEC. While the CCAM services themselves do not operate within Gramine, they are still protected through this secure containerization framework. As depicted in Figure 3.7, the vehicle sensor lies outside the trusted environment provided by Gramine (indicated by the blue-colored modules). However, the software responsible for processing the sensor data and streaming video via the Real-Time Streaming Protocol (RTSP) remains securely protected, maintaining the integrity and confidentiality of the offloaded task.

Trustworthiness Claims (TCs) are exchanged securely, ensuring the protection of the vehicle's identity against fingerprinting attacks. This is achieved through the harmonization functionality provided by the TCH, as previously described. For transmitting video data, a WebSocket connection —demonstrated to perform effectively within trusted execution environments— is utilized to securely stream the data to the MEC. Concurrently, the vehicle's TCH module uploads the relevant evidence to the *CONNECT* MEC, facilitating the trust assessment operation. This process enables the MEC-based TAF to determine the Actual Trustworthiness Level (ATL) of the vehicle using the received evidence. Similar to the vehicle-side, the MEC trusts its GPU accelerators (i.e., the DeepStream App at the GPU) and hosts a confidential container that will receive the data from the Vehicle (i.e., Application Client) in an enclave. Finally, since data are streamed from a Gramine Enclave and received in another, encrypting the communication channel creates a secure end-to-end communication between the Vehicle and the MEC. To achieve this Transport Layer Security (TLS) encryption is employed.

3.4.3.2 Basic CONNECT offloading data-flows

In this paragraph we highlight the flow starting from the RGB Camera sensor on the Vehicle to the GPU Accelerated Inference operation on the CONNECT MEC discussing the interactions between the involved applications and the CONNECT Platform modules.

At the MEC side (see Fig. 3.8), the SMTD application is responsible for listening to incoming WebSocket connections. Once a connection is received from the corresponding application in the vehicle, a TLS handshake is initiated. During this process, the vehicle and MEC applications exchange public keys, and each side uses its private key along with the received public key to derive a shared encryption key. This ensures that all subsequent communication between the vehicle and the MEC remains secure, utilizing the encrypted WebSocket connection. As a next step, the Application on the MEC creates a session with the TAF at the MEC, requesting to initialise a specific trust model. Then the Vehicle Application, provided that trustworthiness is ensured (with guarantees for the MEC and/or the vehicle), launches an ffmpeg process to capture the RTSP video produced by the trusted RGB Camera. The RTSP video (under the H.264

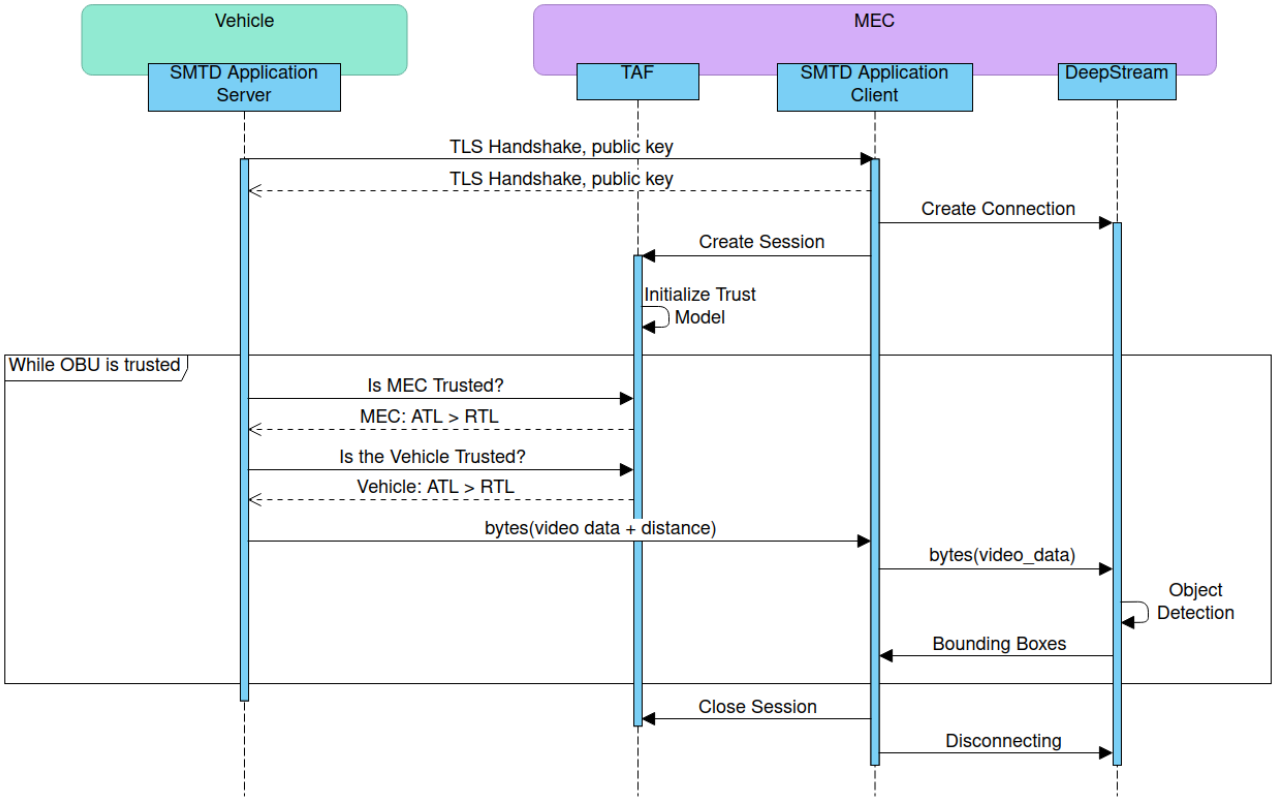

Figure 3.8: Vehicle application and TAF (at the MEC) interaction

format) is encrypted before transmission through the WebSocket connection. The transmitted data also includes a distance measurement from a sensor on-board of the Vehicle. The receiving SMTD Application client forwards the data to the DeepStream module for it to carry out the object detection.

3.4.3.3 MEC Secure Service Deployment through Enclave-cc

For the *CONNECT* MEC, we focus on containerized workloads that are managed using Kubernetes. By default, this technology does not provide Trusted Execution Environment (TEE)-backed security functions such as attestation or sealing. Driven by this limitation, in *CONNECT* we leverage `enclave-cc` to add support for hardware-backed trusted execution to the existing container framework. To achieve this, the ”encapsulated” TEE must effectively manage the lifecycle of the container on a local machine, including functionalities such as hibernation and restoration. In addition to supporting these local operations, we have extended the Intel Gramine library OS to facilitate remote migration and local upgrades. This enhancement enables us to leverage these extensions within the `enclave-cc` framework.

Figure 3.9 illustrates the general flow of migrating a `enclave-cc`-enhanced container. The process begins when the controller initiates a migration from a source vehicle to a target MEC, which involves the necessary key exchanges to authorize the migration. Once authorized, `enclave-cc` utilizes the Gramine migration services to create a clone of the service. If the service is a singleton that must not exist in multiple instances, it employs the *CONNECT* secure monotonic counter service (see D4.2 [8]) to atomically transition the service to the target machine. The Gramine migration service, as outlined in D4.2, ensures that (a) migration is permitted only to an authorized

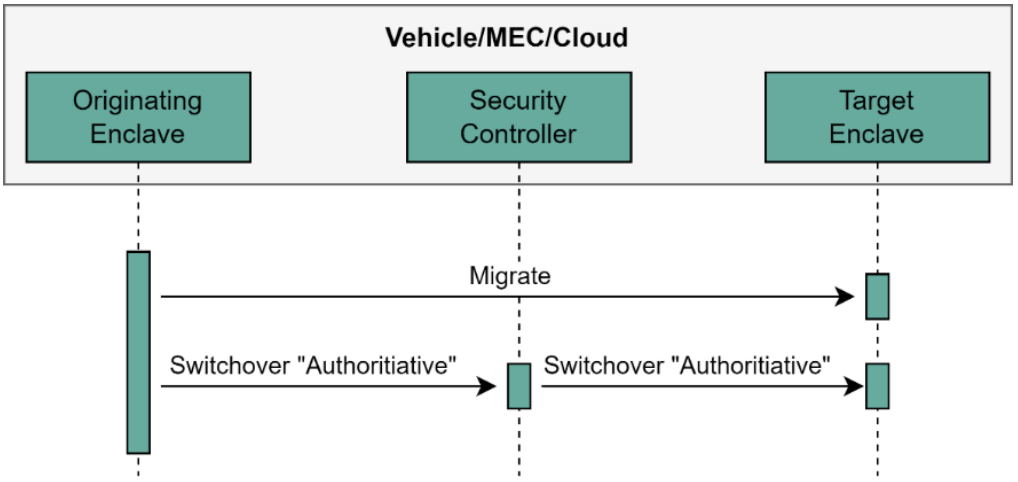

Figure 3.9: *CONNECT* `enclave-cc` Migration.

target machine, (b) the integrity and confidentiality of the state are maintained during transit, and (c) only the designated number of clones is operational at any given time.

3.4.4 *CONNECT* Digital Twin for trust-task offloading

In *CONNECT* the Digital Twin (DT) is envisioned as a mechanism to offload resource-intensive tasks from the vehicle to the MEC. Such offloading serves dual purposes: i) **optimizing resource utilization** and ii) **enhancing trust calculations**. The second is achieved by introducing multiple DT instances, each managed by different operators (i.e., discounting operator). This multi-instance approach ensures a more comprehensive and reliable trust assessment, aligning with the collaborative and distributed architecture of the CONNECT framework. The previous sections have outlined the process of establishing a trust relationship between the vehicle and the MEC and creating a secure communication channel to enable collaborative operations between these two entities. However, the Digital Twin (DT) paradigm introduces an additional layer of complexity through the requirement for **state synchronization** between the twinned object and its replica. Thus, this section focuses on the trust-task offloading process, emphasizing how the DT supports a component like the TAF, which necessitates synchronization.

Specifically, to facilitate trust assessments in the MEC on behalf of the vehicle, **the vehicle's Trust Models (TMs) must be accessible in its MEC-based counterpart**. This requires the secure transfer of the TMs to the MEC, with mechanisms to ensure they are **continuously updated** while the offloading capability remains active. Given the sensitive nature of the Trust Models, a strong trust relationship between the vehicle and the MEC is essential for ensuring that the vehicle can trust the results of these requests, making this scenario an ideal case for demonstrating the *CONNECT* framework. The DT eventually exposes an interface which allows remote instances of the TAF to be started as containerized applications hosted in the MEC infrastructure. In what follows these interfaces and management components will be referred to as "TAF-DT" and the remote containerized TAF instances will be called "Remote TAF". Figure 3.10 shows the main data flows the use of a Remote TAF introduces into the CONNECT architecture. The Digital Twin relies on the V2X message encoding/decoding layer to integrate the various components seam-

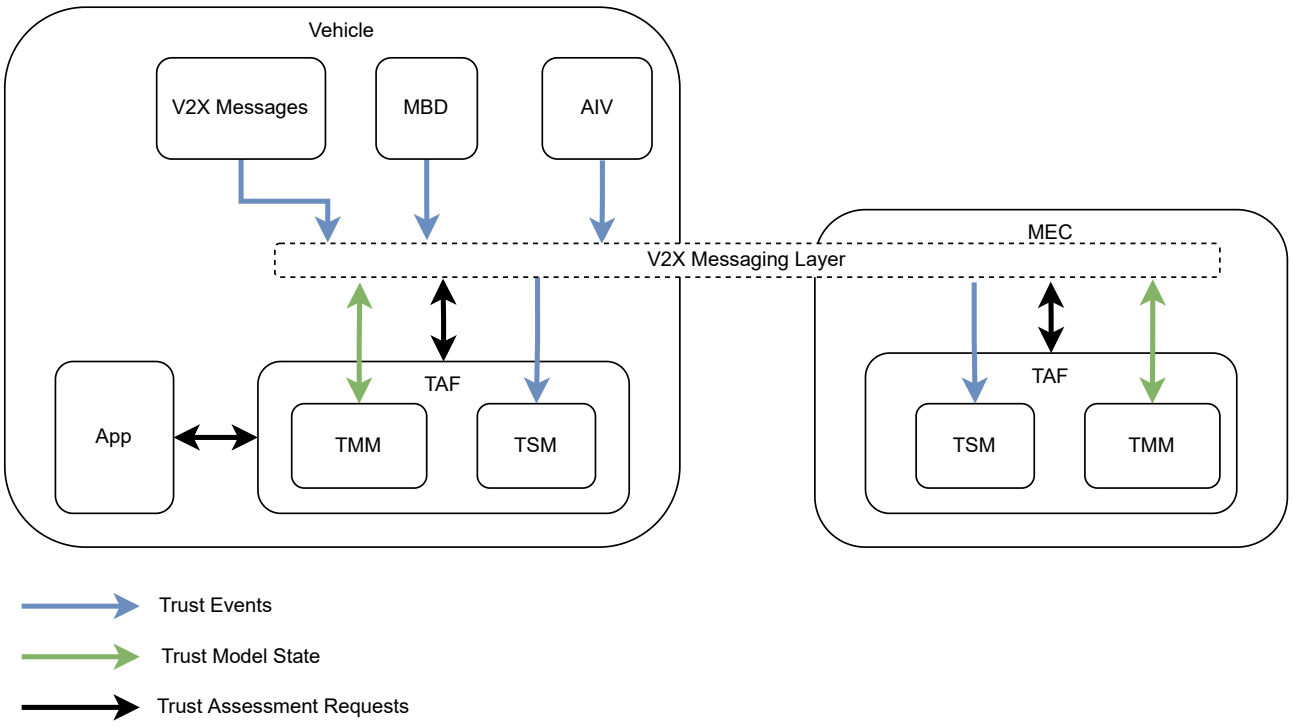

Figure 3.10: Information flows supporting a Remote TAF

lessly. This setup establishes a secondary TAF instance that synchronizes the state of its Trust Model Manager (TMM) with the vehicle's TMM via this communication layer. Additionally, the Remote TAF processes events from the Trust Sources to which the vehicle's TAF is subscribed. Trust Assessment Requests and Responses can also be transferred through the V2X messages layer, enabling the Remote TAF to respond to application requests originating from the vehicle. Further details on the message definitions and associated mechanisms will be presented in the forthcoming deliverables D3.3 and D6.2.

3.4.4.1 DT establishment for Trust Management Offloading

The following section assumes that a trusted communication channel has been established between the TAF-DT in the MEC and local TAF in the vehicle through the V2X messages layer, as described in section 3.4.3.1. This communication channel could be a direct TLS connection or any other way to transfer messages from the local TAF to the remote TAF and will be used to initiate the establishment process; thus, transfer relevant state information and transmit the request and response for offloading. This communication channel may be subdivided into different streams (ex: one stream for establishment and state transfer, another for offloading request/responses) to respect the separation of concerns.

The DT Establishment process, initiated by the local TAF, allows the TAF-DT to collect the information needed to set up and start a remote TAF. This information should cover at least: i) **the Trust Models used and current trust state of the local TAF** and ii) **the Trust Sources used by the local TAF** and implementation-wise the channels used to request or subscribe to trust events. Once this information is provided, the TAF-DT will start an instance of the TAF. This TAF will then initialise its TMM with the trust model information provided and its TSM with the trust sources configuration of the local TAF. Once this process is complete, the TAF-DT will have a

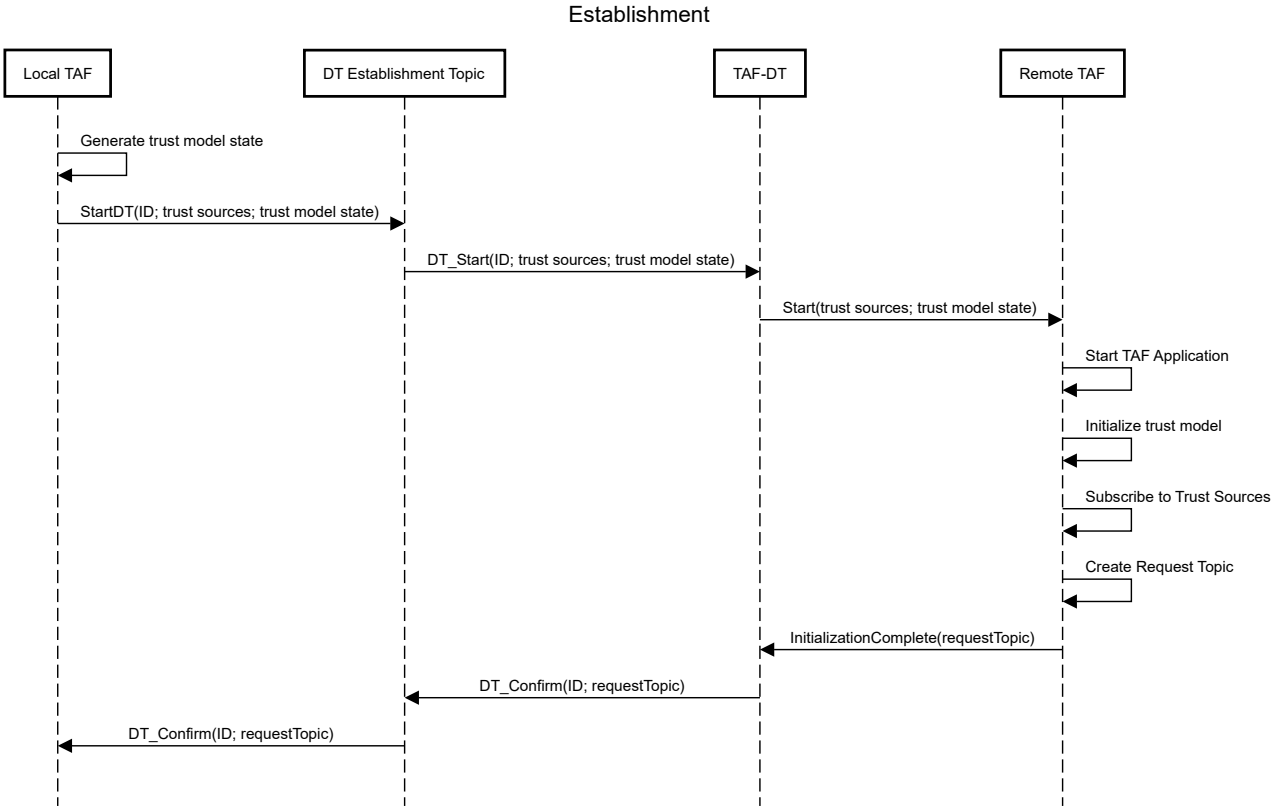

Figure 3.11: Digital Twin Establishment

channel through which the Trust Assessment Service of the Remote TAF can be queried and will inform the local TAF of the success or failure of the process and give it the information needed to access the remote TAF Trust Assessment Service channel. This process is summarized in figure 3.11.

Once the remote TAF is started, it can receive Trust Assessment Requests through the same interfaces as the local TAF. As a result, offloading can happen transparently from the point of view of the application sending Trust Assessment Requests by having the local TAF act as a proxy, forwarding requests and responses to and from the remote TAF.

3.4.4.2 Synchronisation and Trust Sources Management

Different strategies for synchronisation can be implemented. In what follows we outline the synchronisation capabilities required for a full-fledged DT, while providing a summary of the planned implementation approach within the context of *CONNECT*. The strategy for a fully-fledged DT should address the following key aspects:

- The offloading functionality for the DT will only be available during ”offloading sections” which define a state of the local TAF where it can synchronise its state with the remote TAF. Every offloading section starts and ends with a synchronisation point.
- At a synchronisation point, a TAF serialises its trust models state and sends it to another TAF

- When entering an offloading section, the local TAF stops updating its trust models, synchronises with the DT and will forward trust assessment requests to the remote TAF.
- During an offloading section, the remote TAF receives trust events from the trust sources and updates the trust models it received.
- When leaving an offloading section, the local TAF signals to the remote TAF which then stops updating its trust models and synchronises with the local TAF
- The local TAF starts updating the trust models and resumes normal operation.

The main advantage of such a strategy is that it allows to focus on the core support function of offloading computations without having to worry about the divergences, caused by having local and remote trust models being updated in parallel. Downsides include the necessity to pause the local TAF instance updates, in order to prevent potential divergences of the trust models, making it so it cannot answer requests without offloading them and some potentially expensive synchronisation phases when entering and leaving an offloading section. In addition to these, there are some potential challenges that need to be considered, as presented in Table 3.3.

An alternative approach involves streaming event updates for the trust models from the local TAF to the remote TAF, enabling parallel operation. In this method, the local TAF continues processing evidence collected from Trust Sources and updates its trust models, while transmitting atomic event updates to the remote TAF using an event-sourcing pattern. This approach ensures that the trust models eventually converge, mitigating issues associated with explicit synchronization of entire trust models. However, it introduces potential local inconsistencies between the two states, which would need to be resolved.

As an intermediate milestone toward achieving a fully capable DT within the *CONNECT* framework, we focus on a partial implementation that emphasizes evidence synchronization between the two TAFs. This approach facilitates the offloading of the Trust Level Expression Engine (TLEE) to the MEC-based DT, demonstrating the core capabilities of the DT concept.

3.4.4.3 Additional modes of operation for the Digital Twin

The previous section describes the case of a Digital Twin used by the vehicle to expand its capabilities, however a DTs can also be used by the Road Infrastructure. Another use case for a Digital Twin in the context of managing trust relationships would be to combine the ability use a vehicle's TAF twin with the federation capability described in the next section. This federation of twins could be used to enhance trust assessments by having a more complete view of trust relationships across participants in the CCAM ecosystem without hitting the penalties in term of latency and additional load placed on the vehicles that would occur when federating the actual vehicles.

As this exchange of information between twins implies a divergence in the trust models between the local TAF and the TAF-DT, because the local TAF may not systematically use the federated mode, it would need to operate in parallel to the mode described previously. This option could build on the solutions proposed in *CONNECT* but would be the subject of further research.

Table 3.3: Challenges and Solutions for Trust Synchronisation in CONNECT Framework

Challenge	Description and Mitigation
Atomicity of synchronisation with respect to trust events and model updates	The synchronisation process is decomposed into the following steps: 1. Collect trust models state. 2. Serialize state into messages. 3. Transmit messages. 4. Deserialize messages into state. 5. Apply state to trust models. During this process, the trust sources may still be producing trust events which will need to be applied by the receiving party. The synchronisation mechanism should ensure that these events are correctly applied. A way to resolve this would be to include a timestamp when freezing the trust models and apply every trust event produced after this timestamp
Loss of connectivity between the DT and the local TAF	This may cause the remote TAF to miss events and the local TAF would be de-synchronised and unable to rebuild state. Storing the events produced by trust sources during an offloading section and applying them a posteriori if the digital twin cannot be reached would mitigate this issue.
Addition of trust models during offloading	The TAF is able to add new trust models during its execution. If this happens during an offloading section, the additional trust models would have to be transferred and applied to the local TAF at the end of the offloading section, making synchronisation more complex.

3.4.5 Trust Assessment Framework (TAF)

As previously explained, the *TAF* operates in several modalities: i) standalone, ii) federated, and iii) DT-TAF. While the DT-TAF was discussed in detail in the preceding section, this section will focus on the first two modalities. In what follows we describe the information flows between the TAF and the external components that the TAF is interacting with. These external components are the application, the TCH, the AIV, the MBD, the CONNECT DTL and possibly another TAF (federated TAF).

3.4.5.1 Standalone TAF

Session Initiation: In the session instantiation, the application provides several parameters to the TAF, e.g. the trust model template that should be used. The TAF then internally prepares session handling for this session and eventually responds with the ID of the session.

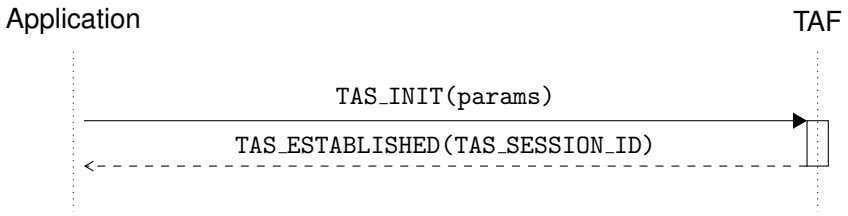

Session Tear-Down: After a session has been established, the application can tear-down the created session. This allows the TAF to free all resources allocated for this session.

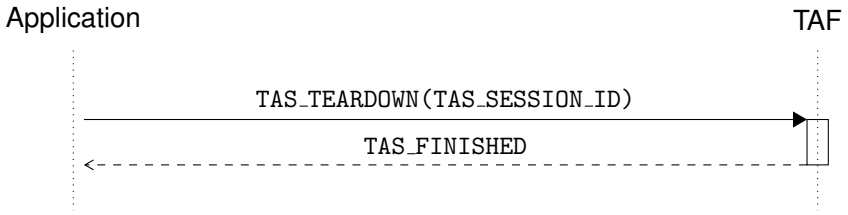

Regular Trust Assessment Request: Once a session has been established, the application can send Trust Assessment Requests (TARs) to the TAF. The message format of a TAR is described in D6.1 [6]. The TAF will respond to a TAR in an best-effort way and will (if not specified otherwise in the options of the TAR) potentially use cached results that have been calculated recently.

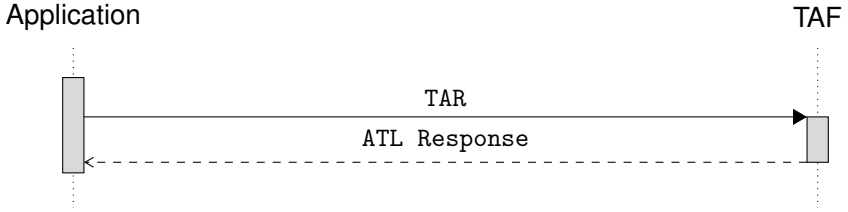

Subscription for Trust Level Changes: In addition to the regular trust assessment request, applications can also subscribe to changes of the ATLs. In the subscribe request several options can be provided. With these options, it can be specified whether the application should be notified if the ATL has changed or only if the ATL has changed and therefore exceeds or falls below the specified threshold.

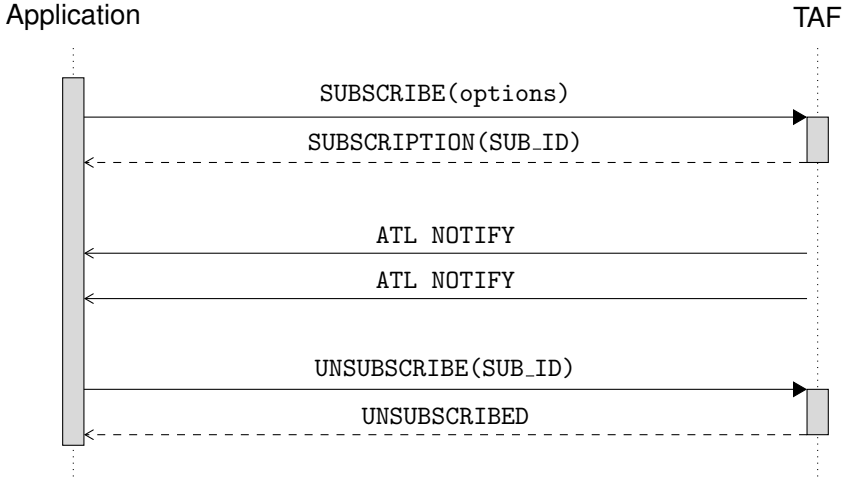

Pull-based Evidence Collection: The TAF communicates with trust sources to receive evidence. Depending on the trust source, the TAF can communicate with this trust source in different ways. An example of pull-based evidence collection is the AIV. The TAF can communicate with the AIV in a pull-based manner. In this case, the TAF sends a request to the AIV including the parameters for which entities which type of evidence is requested.

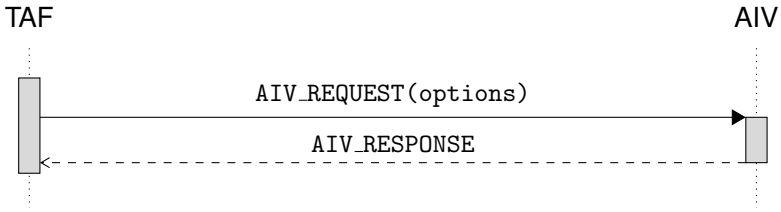

Subscription-based Evidence Collection: In addition to the pull-based evidence collection, it is also possible to subscribe to a trust source. An example for subscription-based evidence collection is the AIV or MBD. The TAF can subscribe to the AIV and MBD. In this way, the TAF is notified as soon as new evidence updates exist. If the TAF is not interested in the trust source anymore, it can unsubscribe from it. The following sequence diagram shows the subscription-based evidence collection for the MBD. The interactions are exactly the same for the AIV.

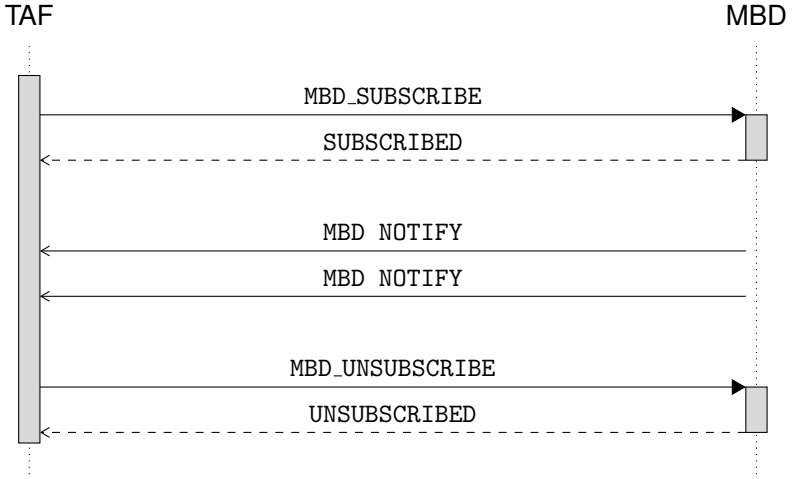

One-Way Evidence Notifications A third option for evidence collection is to receive evidence without a request and without a subscription. An example for One-Way evidence notifications is the TCH. This is the case, Verifiable Presentations are sent from the TCH to the TAF.

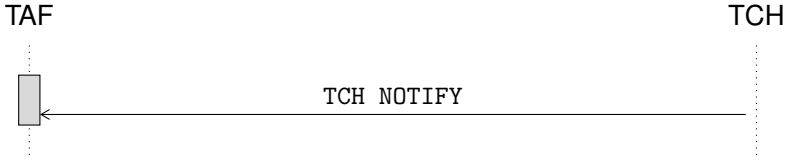

CONNECT DLT: The CONNECT DLT represents a remote repository for the retrieval of useful information. Such information could include Trust Model Templates (TMT) or Required

Trustworthiness Levels (*RTL (Required Trustworthiness Level)s*). From the perspective of the *TAF*, the CONNECT DLT functions as a distributed, read-only key-value store where *Trust Model Templates (TMTs)* and *RTL (Required Trustworthiness Level)s* are indexed by unique identifiers and version numbers. We assume following use cases in which the TAF could interact with the DLT:

1. retrieving a specific Trust Model Template based on its ID and
2. retrieving a Trust Model Template based on its ID *and* referring to a specific version
3. retrieving dynamic updates of the RTL

In the first two cases, the TAF initially receives a request to establish a trust assessment session that references a Trust Model Template not yet know by that TAF instance. Hence, the CONNECT DLT is used to retrieve the template. In case the session request only specifies the Trust Model Template identifier, we assume that the latest version of that template available in the DLT is to be used. If a specific version number for that template is used, the TAF will query for that specific version of the template. If the specified template (or version number) is unknown to the TAF, the remote resolution will fail and the TAF cannot instantiate the session and corresponding trust model instances. The third scenario involves setting up interfaces for continuous interaction with the DLT to retrieve dynamic updates of the RTL. **This interface and functionality, related to the RTL will be validated in subsequent phases as part of the *CONNECT* framework's implementation**.

3.4.5.2 Federated TAF

Querying other TAF instances: The TAF contains a Trust Assessment Query (TAQ) Interface that enables the TAF to query trust assessments from other TAF instances. These TAF instances can run on two different entities, e.g., two different vehicles, two different MECs, or one on a vehicle and one on a MEC. This interface facilitates interactions between different *TAF* instances, forming a foundational component for federated behavior. By querying another *TAF*, a local *TAF* can enrich its own perspective by incorporating assessments generated by the queried instance. This can be useful if the TAF does not have the computing capacity to calculate the corresponding trust opinion or does not have access to the evidence required to calculate the trust opinion.

A TAQ contains a selector to identify trust values the querying TAF is interested in. A selector specifies a Trust Model Template type and potential paths in corresponding trust model instances with optional attributes (i.e., to identify certain trust models and contained nodes). By using wildcards, a selector allows for broader queries, e.g., queries that can ask for any trust model instance in which a node with a certain ID is part of. The queried TAF then returns a potentially empty list of trust values satisfying the selector.

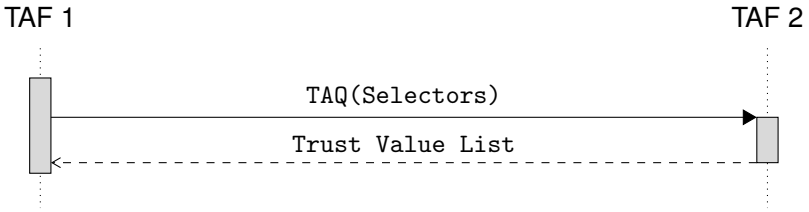

An example of two different trust models is shown in Figure 3.12. If TAF 1 is interested in the trust opinion for the trust relationship between node C and node D, this is specified in the query part of the TAQ. The corresponding trust opinion that TAF 2 has stored in its trust model of this trust relationship is then transmitted to TAF 1 as a response. TAF 1 can then include the trust opinion in this trust model and calculate the corresponding ATL.

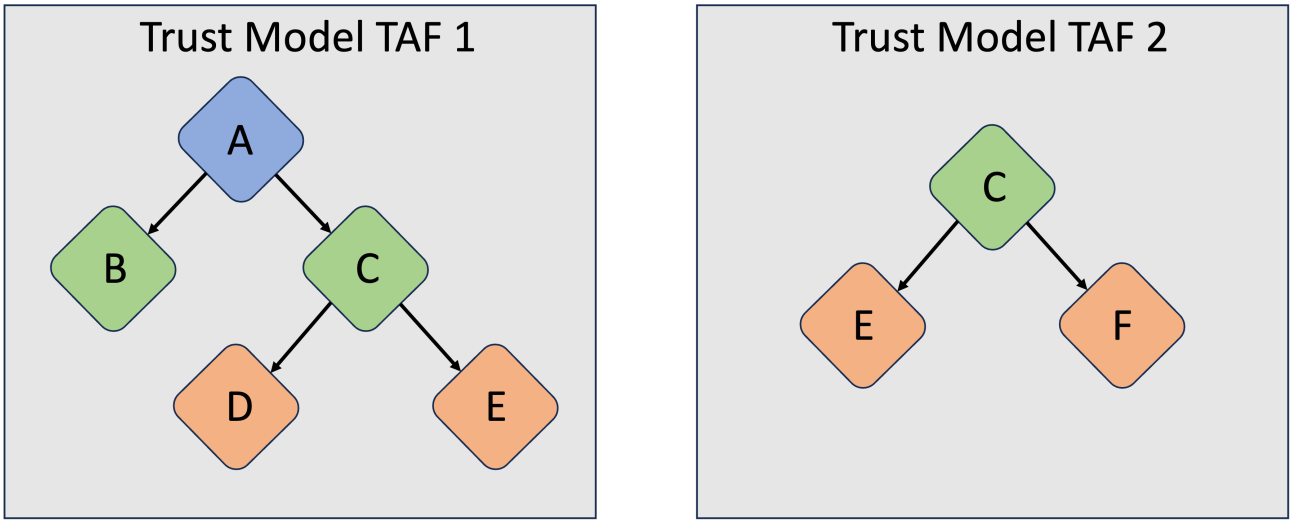

Figure 3.12: Examples of two trust models used in two different TAF instances

3.4.6 *CONNECT* Attestation Technologies as a Trust Assessment Enabler

3.4.6.1 In-vehicle Trust Evidence Monitoring

The attestation mechanisms supported by the *CONNECT* play a critical role in enabling trust assessments within the system. The operational flows for these mechanisms are detailed in D4.1 [2] and D4.2 [8]. It shall be noted that even though the crypto operations in the aforementioned deliverables are focusing on the Configuration Integrity Verification (CIV), the same process is followed for the rest of the claims that can be reported by the Trust Sources, including Access Control, Secure Boot, Secure Over the Air (OTA) updates, Application Isolation and Control Flow Integrity.

As illustrated in Figure 3.3, upon triggering by the *AIV* the process begins with traces (acting as evidence) being extracted by the untrusted part of the Tracer that resides in the *A-ECU*. These traces, which may reflect both static and runtime properties, are signed by the trusted part of the Tracer, before being sent to the Attestation Agent.

The Attestation Agent performs a verification of these traces and generates a zero-knowledge proof, meaning that no sensitive trace details are revealed externally, outside the *TCB*-only a proof of correctness is shared. This ensures that while the integrity of the evidence is verified, the privacy of the underlying data remains intact. Once the attestation evidence, comprising both the Attestation Agent's signature and the zero-knowledge proof of correctness, is prepared, it is transmitted to the *AIV* for verification. Various *A-ECU* components individually report their evidence to the *AIV*, using a threshold signature scheme. Upon successful verification, the attestation report is forwarded to the *TCH* for further processing.

Similarly, the *S-ECU* that possess capabilities to extract solely Secure Boot and Access Control traces and signs them symmetrically. This report is sent from the *S-ECU* to the *A-ECU*. The latter may receive similar reports from multiple *S-ECU*s, which aggregates and sends to the *AIV* for verification.

Once the *AIV* report is received by the *TCH*, it undergoes harmonization, further leveraging the threshold Direct Anonymous Attestation (DAA) scheme. This harmonization enhances the overall security of the attestation process while preserving the privacy of the entities involved, particularly when data is shared beyond the confines of the vehicle. By employing DAA, *CONNECT* ensures both the integrity of the attestation evidence and the protection of the identities of the participants. The outcome of this process is forwarded to the *IAM*.

3.4.6.2 MEC Infrastructure Trust Level Monitoring

Similar to the vehicle, the MEC-based infrastructure utilizes a *TCB*, including components like the Tracer and Attestation Agents, to extract evidence from the underlying infrastructure. Additionally, a specialized μProbe is deployed within the containerized service to support the extraction of traces at the service level. This dual-layer extraction approach allows the MEC to gather evidence both from its core infrastructure and from the services it hosts, enhancing the overall trust assessment process for secure and reliable operations, further adhering to the classification proposed by ETSI [15].

In the *CONNECT* framework, the *ETSI Levels of Assurance (LoA)* categorization is applied to classify the trustworthiness of various components, such as virtualised service(s) or even platforms. These LoAs range from LoA 0, which requires no integrity checking, to LoA 5, which demands remote verification of the infrastructure network, virtualization layer, and the integrity of VNF software packages during runtime. Currently, *CONNECT* is achieving LoA 2, which involves remote verification of hardware and virtualization platform integrity, while LoA 3, requiring local verification of VNF software packages, is under investigation. More information on these scheme will be analysed in D4.3 [9]. The flow of the MEC-based attestation is illustrated in Figure 3.13.

The process of extracting attestation evidence is triggered by the TAF (Step 1). Both the trace extraction and the attestation process is taking place both for the containarised service(s), leveraging the μProbe, as well as the infrastructure, leveraging the Security Probe. The flow goes as follows, the *AIV* receives a request from the *TAF*; along with this information it requests information about the topology from the *IAM* (Step 2). The Security Probe receives a request from the *AIV* (Step 3); thus initiates its attestation process. The traces are collected by the untrusted part of the Tracer that is part of the Security Probe and signed by its trusted part before being sent to the Attestation Agent for verification (Step 4). Both the Tracer and the Attestation Agent are parts of the Security Probe. The Attestation Agent then validates these traces and generates a zero-knowledge appraisal, meaning that the exact trace details are not disclosed to any external entity outside the *TCB*. Only a proof of correctness is provided, ensuring the integrity of the process while maintaining confidentiality. If this request further specifies the containerised service, then the Security Probe will trigger the attestation for the μProbe as well (Step 5). The same process described in Step 4 is taking place also for the containerised service(s), leveraging the μProbe (Step 6). In this case, the Security Probe is tasked with the verification of the attestation evidence as received from the μProbe (Step 7-8). The signed report is sent to the *AIV* and then to the *TCH*.

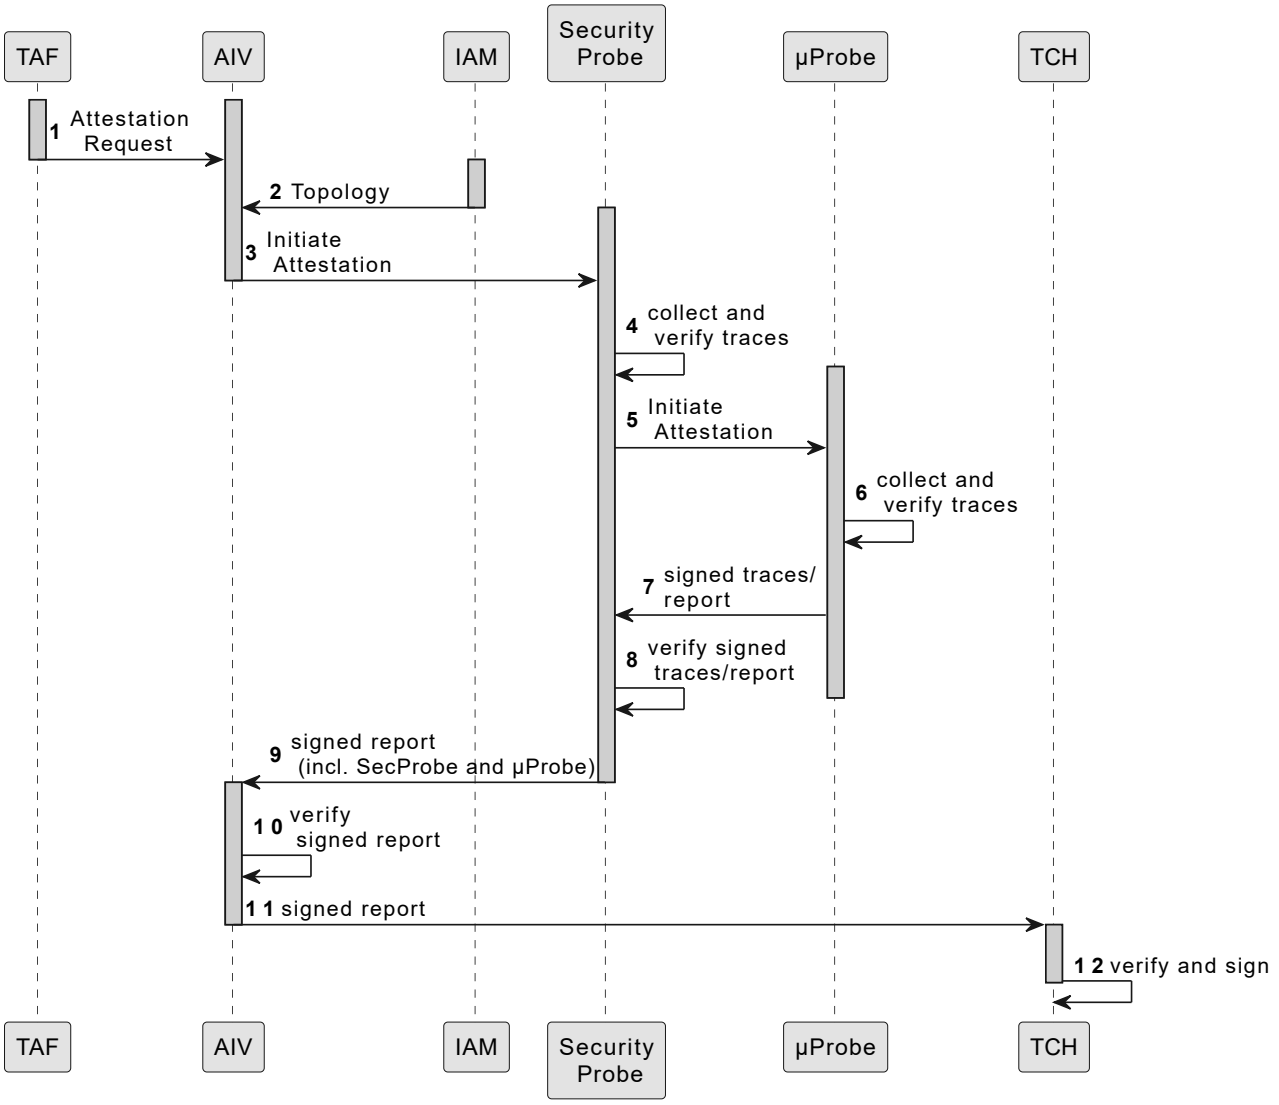

Figure 3.13: CONNECT MEC Monitoring

3.4.7 Distributed Ledger Technology (DLT) for Trust Policy Enforcement and Evidence Management

The *CONNECT* DLT serves as a pivotal mechanism for storing and retrieving information within the *CONNECT* framework, enabling the establishment and maintenance of trustworthiness among vehicles. It facilitates different data workflows, while ensuring key properties such as auditability, immutability and certifiability of data sharing transactions among the *CONNECT* entities. These attributes collectively ensure the trustworthiness of vehicles under the *CONNECT* framework. The architecture and components of the *CONNECT* DLT are described in detail in deliverable D5.2 [5]. Figure 3.14 illustrates the five distinct data workflows of the *CONNECT* DLT, with each workflow color-coded to represent different phases of the trustworthiness establishment process. The figure highlights the interactions between the *CONNECT* DLT components and other entities within the *CONNECT* framework. Below, we describe the steps involved in each data workflow.

3.4.7.1 Setup phase – Trust Model Template (including RTL) deployment from TAM to TAF Agents through DLT

The setup phase ensures that the necessary trust-related information is accessible to relevant entities through the *Distributed Ledger Technology (DLT)*. This process is initiated by the *TA (Trust Assessment Manager)*, which operates at the cloud level and is responsible for storing *TMT* – including the *RTL (Required Trustworthiness Level)* – in the Private Ledger. For this, the TAM must possess a valid certificate issued by the Privacy Certification Authority (CA) during the enrolment phase (Step 1). In addition to the certificate, the TAM must be authorized by the *Attribute-based Access Control (ABAC)* service, a component of the *Security Context Broker (SCB)*. This authorization is based on verification using a *VP* provided by the TAM. Once the TAM has been authorized, it forwards the *TMTs* that include the *RTL (Required Trustworthiness Level)s* to the CONNECT *DLT* via the *SCB* (Step 2).

The *CONNECT* framework leverages a Private Ledger to store trust-related data (i.e., Trust Model Templates that include the Required Trust Level). This approach ensures tamper-proof storage and controlled access to sensitive information. The Blockchain Peer writes the information into the ledger using the relevant chaincode (Step 3), maintaining the integrity and confidentiality of the trust-related operations.

Once the *TMTs* (including the *RTL (Required Trustworthiness Level)s*) are successfully stored in the ledger, a notification is sent from the Blockchain Peer to the TAF Agents located both at the *MEC* infrastructure and the vehicle (Step 4). These TAF Agents, using their *VP*s, are then authorized through ABAC to access the CONNECT DLT and retrieve the *RTL (Required Trustworthiness Level)s* included in the *TMTs* from the Private Ledger (Steps 5, 6, 7). Depending on whether the service pertains to MEC or *CCAM*, the next steps—(B and C for MEC, D and E for CCAM)—are followed.

3.4.7.2 Failed attestation evidence storage on the DLT from the MEC

After the deployment of the smart contracts that include trust-related information (i.e., RTL) as derived from the Trust Assessment Manager (TAM), the MEC's Trust Assessment Framework (TAF) Agent sends a *Trust Assessment Request (TAR)* to the *AIV* (Step 8). The latter then initiates the collection and verification of attestation evidence for the MEC infrastructure and the deployed services (Step 9). In the case of a failed attestation, the AIV stores the respective attestation evidence to the *CONNECT* Distributed Ledger Technology (DLT).

To do this, the AIV needs to be authorised first through the Attribute-Based Access Control (ABAC) (Step 10). This is achieved leveraging its *VP*. The Security Context Broker (SCB), acting as a mediator to the DLT, forwards the failed attestation evidence to a Blockchain Peer in the CONNECT DLT, where it is stored in the Private Ledger using a smart contract (chaincode) (Step 11, 12, 13). It needs to be noted here that the raw attestation evidence is encrypted using Attribute-Based Encryption (ABE) and stored in Off-Chain Storage (Step 14). A database pointer to the location of the encrypted evidence is sent back to the Blockchain Peer and stored in the Private Ledger, updating the block for the corresponding failed attestation report (Step 15).

3.4.7.3 Access to MEC failed attestation evidence from authorised third parties

After the failed attestation evidence have been stored on the DLT, a Security Administrator may access them via the Security Context Broker (SCB) to investigate potential new threats or vulnerabilities (Step 16). To perform this task, the Security Administrator, needs to be authorized first through the Attribute-Based Access Control (ABAC) mechanisms leveraging its *VP* (Step 17). After the successful authorisation, the Blockchain Peer retrieves the encrypted failed attestation evidence from Off-Chain Storage, leveraging the (database location) pointer that is available on-chain (Step 18). Hence, the evidence is retrieved from the Blockchain Peer and forwarded through the SCB to the Security Administrator (Step 19). With the necessary attribute-based decryption keys, which are created through the CONNECT Trusted Computing Base (TCB), the Security Administrator, as an authenticated and authorized entity, decrypts the evidence and processes the raw attestation traces for potential risks (Steps 20, 21).

If new vulnerabilities are identified, they are uploaded through the SCB to the Private Ledger of the CONNECT DLT (Step 22). The Trust Assessment Manager (TAM) retrieves these vulnerabilities (Step 23) and collaborates with the Risk Assessment (RA) component to recalculate the Required Trust Level (RTL) for the MEC services (Step 24). The updated RLT and trust templates are stored in the Private Ledger through the SCB (Step 25) and consequently forwarded to the TAF Agent at the MEC (Step 26).

3.4.7.4 Failed attestation evidence storage on the DLT from vehicle

As described for the MEC in subsection 3.4.7.2, attestation evidence are also collected in the In-vehicle topology. After the storage of the *TMTs* from the Trust Assessment Manager (TAM), the TAF Agent of the In-Vehicle Manager may leverage this information for a new trust assessment. Hence, the *TAF* triggers the *AIV* to initiate the attestation process (Step 27). The latter collects and verifies the attestation evidence from the *ECU*s (Step 28). In the case of a failed attestation, the *AIV* stores to the DLT the respective evidence (Step 29). To do so, the *AIV* requires to be authorized first, through Attribute-Based Access Control (ABAC) using *VP*.

Once authorized, the AIV submits the failed attestation evidence to the Security Context Broker (SCB) (Step 30), which forwards it to a Blockchain Peer in the CONNECT DLT. The Blockchain Peer stores the failed attestation report in the Private Ledger using a smart contract (chaincode) (Step 31). The raw attestation evidence is encrypted using the Attribute-Based Encryption (ABE) service (Step 32) and stored in Off-Chain Storage (Step 33). A database pointer for the failed attestation evidence is returned to the Blockchain Peer and recorded in the Private Ledger, amending the relevant block in the data transaction (Step 34).

3.4.7.5 Access to CCAM failed attestation evidence from authorised third parties

Subsequently, an *Original Equipment manufacturer (OEM)* may request access to the failed CCAM attestation evidence from the *CONNECT* DLT via the SCB, to assess potential threats or vulnerabilities (Step 35). The OEM, authorized by ABAC with its *VP* (Step 36), retrieves the encrypted evidence from the Blockchain Peer. The latter uses the corresponding (database location) pointer from the Private Ledger to retrieve the encrypted failed CCAM attestation evidence from the Off-Chain Storage (Step 37), and later transmits them to the *OEM*, through the SCB (Step 38).

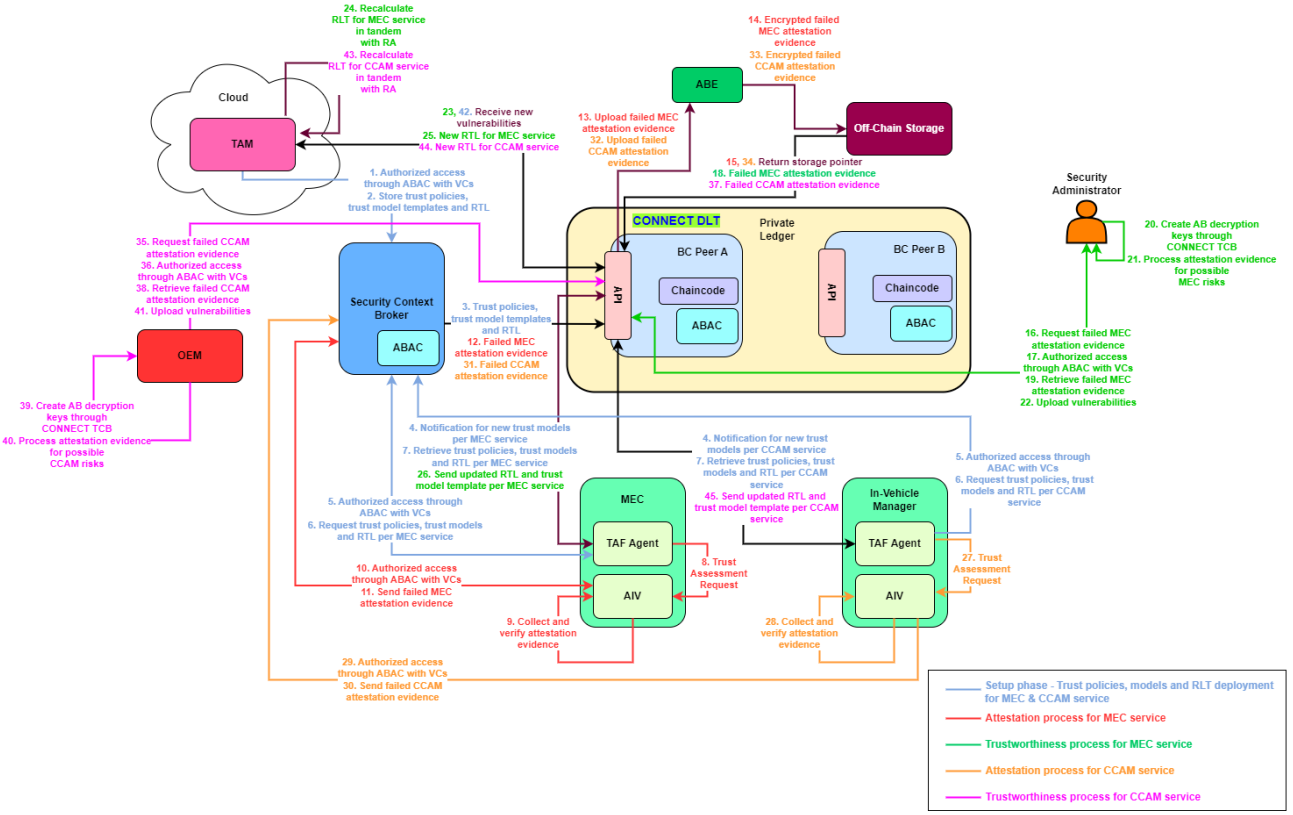

Figure 3.14: CONNECT DLT Architecture

The *OEM* decrypts this information using attribute-based decryption keys generated through the *CONNECT* Trusted Computing Base (TCB) (Step 39), and processes the evidence for CCAM risks(Step 40). If new vulnerabilities are identified, the *OEM* leverages the DLT to store them (Step 41).

The TAM retrieves these vulnerabilities as identified by the *OEM*, through the SCB and collaborates with the Risk Assessment (RA) component to recalculate the Required Trust Level (RLT) for the CCAM service (Step 43). The new RLT and *TMTs* are stored in the Private Ledger (Step 44) and forwarded to the TAF Agent at the vehicle (Step 45).

3.4.8 CCAM Continuum Key & Identity Management

In CONNECT the security of each device or container in the system is provided by a consistent set of properties and algorithms which, for whatever hardware is used, form the *TEE Security Guard Extensions (TEE-SGE)* and the *Trusted Computing Base (TCB)*.

TEE-Security Guard Extensions As their name suggests the TEE-SGE are built upon a Trusted Execution Environment and provide the core functionalities required by the other components in the system. These are: i) the Identity Authentication Management (IAM), ii) the Attestation Integrity Verification (AIV) and iii) the Trustworthiness Claims Handler (TCH). Note that as we do not have the same privacy requirements for messages generated by the MEC the TCH is not required.

In the *CONNECT* implementation the TEE functionalities are provided by Gramine a software shell that sits above Intel-SGX and makes the deployment of applications inside SGX more

straightforward. In the vehicle the TEE-SGE components run directly in a Gramine based SGX enclave, while in the MEC the TEE-SGE components run in a confidential container which uses Gramine and SGX to provide the underlying TEE.

Trusted Computing Base The TCB runs on each device and in each application container and in all cases it provides: i) an Attestation Agent, ii) a Verifiable Policy Enforcer (VPE), iii) an Attestation Tracer, iv) a Migration Agent and v) a Key Management System. These different security components are described in detail for the vehicle and the MEC in deliverables D4.2 [8], D5.1 [3] and D5.2 [5].

When a device is on-boarded, or a secure application container deployed on a device, their information (IDs and (public) key(s)) are provided to the IAM which can use them or make them available to other devices to enable information from that device to be validated. Other keys, for example those for migrating an application from one enclave to another, are derived as necessary. The IAM also maintains a list of other keys and certificates, in particular those from the PKI that it uses to sign verifiable presentations that are sent to other components in the system.

Other keys used by the TAF, MBD and TCH to generate their verifiable credentials (VCs) and by any authorities that issue VCs for the attributes used for encryption and access control are stored in a data registry that is readable by all devices that need it. Typically this data registry is a blockchain as in addition to providing a distributed database, it provides a tamper-proof record.

In the first release of the CONNECT framework not all of these facilities are in place. Some, like the provision of a blockchain data registry for storing decentralised identifiers for the CONNECT VCs, will be added in the next release. Others, like the how to maintain a vehicle's privacy when sending information outside of the vehicle need more discussion.

Keys	Description & Need
Vehicle Master Key	The master key, K_M, is stored in the IAM and is used when generating other keys in the system, for example, the keys used when integrating ECUs into the vehicle. It can be considered as the vehicle's identity key.
S-ECU keys	These are symmetric keys used for integration, software update and protecting communications with other components in the vehicle. They are described in Table 5.2 in D4.2. S-ECUs have no root of trust and therefore no TCB.
A-ECU keys	These are keys used for integration, software update and protecting communications with other components in the vehicle. They are described in Table 5.3 in D4.2. A-ECUs will have a TCB and will therefore have further keys used for configuration and integrity verification (see below).
TCB keys - Configuration and Integrity Verification	CIV extends the 'normal' attestation to include confirmation that the correct software configuration is also being used. It uses key restriction policies to report its results in a zero knowledge manner, a challenge (nonce) can only be signed and returned if the system configuration is correct. To do this, the scheme extends the on-boarding process to include the generation of the necessary attestation keys. The scheme is described in D4.2 and the set of keys used are given in Table 5.4.
TCB keys - Software Migration	For software migration a checkpoint and restore mechanism is used. In order to do this securely the state of the application is encrypted before transferring to the new host. The encryption key used to do this is obtained from the (trusted) key broker service.
Kubernetes Keys	These keys are used when installing a Kubernetes POD. The key is also used: (1) for communication of the POD with the Orchestrator, (2) to protect telemetry from the POD while it is running and (3) when there is an update to the POD deployed by the Orchestrator (this update is encrypted based on this key).

Credential Keys	Verifiable Credentials and Verifiable Presentations are used extensively in CONNECT (see Chapter 5 in D5.1). Underlying their use is a data registry which contains the decentralised identifiers (DIDs) for the components using the system. The different components will hold their private keys and the DIDs will hold the public keys needed for verification.
Attribute Keys	Attribute keys are used in CONNECT for Attribute-based Encryption (ABE) and Attribute-based Access Control (ABAC). Attribute keys are issued by the blockchain's Security Context Broker. While there is some description of these mechanisms in Chapter 5 of D5.2, they will be described in detail and examples provided in D5.3.
Integrity Keys (i.e., MREnclave, MRSigner)	Managed by the RoT to protect and safeguard the integrity of the enclaves and enable the exchange of data between enclavised applications through an authenticated and encrypted communication channel.

Table 3.4: Cryptographic Keys

Privacy Requirements In D2.1 [4], we described how a set of harmonised attributes could be used to ensure the privacy of vehicles when they send their attestation results outside of the vehicle. This was to avoid fingerprinting, i.e. matching the information sent against the known configurations of devices inside of vehicles. This is not part of the first release of the *CONNECT* framework as more work is needed to understand how the harmonised attributes can be incorporated into the trust models. So, currently the attestation data is provided directly, there is no attempt at concealment. In the future there are several possibilities: (a) we use the harmonised attributes already described, or (b) we investigate other more flexible options. In (a) many of the attributes are pass/fail, for example, if a single ECU fails on one of the tests the whole system fails for that test. This feels too restrictive, so we will also explore other possibilities. In either case, work will be needed to adapt our trust models for these new 'measurements'.

Chapter 4

Threat Modelling in CCAM

The *CONNECT* project aims to enhance the security and trustworthiness of *CCAM* ecosystems by developing a Trust Assessment Framework (TAF). This framework is designed not only to extract verifiable evidence during runtime for the evaluation of the Actual Trustworthiness Level (ATL), but also to proactively identify and assess potential vulnerabilities and risks. This action allows to dynamically update the Required Trustworthiness Level (RTL), ensuring that defences remain robust and adaptive to emerging attack vectors. This continuous assessment and adjustment of the RTL enhances the resilience of the system, safeguarding it against evolving threats in real-time.

The threat model of *CONNECT* considers both external and internal threats, adhering to the zero-trust principle. This dual focus is crucial, especially in the context of emerging technologies like Multi-Access Edge Computing (MEC) and their implications for autonomous vehicles. The TAF's requirements in terms of threat coverage are based on real-world security challenges, with privacy considerations also playing a crucial role. In this chapter we will provide a comprehensive threat analysis based on the STRIDE and leveraging TARA methodology while we will further analyse the TAF considerations for the identified threats. It shall be clarified though that *CONNECT* TAF focuses on security rather than safety; hence, risk related to safety are not explored.

4.1 Threat Categories and Classes

In the automotive industry, the development of safe and secure systems is guided by established standards like ISO 26262 and SAE J3061, which focus on functional safety and cybersecurity, respectively. These standards provide critical frameworks for creating dependable automotive systems. To complement these, a specialized methodology for threat and risk analysis tailored to the unique challenges of automotive systems is essential. Various approaches, such as EVITA, HEAVENS, SHIELD, and TVRA, have been proposed in the literature to address these needs. However, Threat Analysis and Risk Assessment (TARA) has emerged as the most widely adopted and effective methodology in the automotive sector.

TARA, standardized in ISO/SAE 21434 [19],, integrates best practices for conducting threat analysis and risk assessments in automotive cybersecurity. It offers a structured approach to identifying, analyzing, and mitigating potential threats across complex vehicular systems, ensuring the highest level of protection for both data and vehicle operations. As elaborated in D3.2 [7] TARA is performed on an Item, which includes information about its function, attack surface,

item boundaries, and operational environment. The TARA activities include asset identification, threat scenario identification, impact rating, attack path analysis, attack feasibility rating, risk level determination, and risk treatment decision.

Microsoft's STRIDE model has also significantly influenced threat modeling practices, laying the foundation for TARA's threat identification and categorization. STRIDE (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege) serves as a widely recognized framework that aids in categorizing security threats and simplifying threat modeling efforts. Its integration into TARA ensures that the automotive industry can apply a proven and methodical approach to identifying and addressing security vulnerabilities early in the development process. Table 4.1 outlines the six STRIDE categories, each representing a critical threat type that needs to be assessed and mitigated within automotive systems:

1. **Spoofing (S)**: Unauthorized access through falsified identities or credentials.
2. **Tampering (T)**: Unauthorized modification of data or system components.
3. **Repudiation (R)**: Denial of actions or operations to avoid responsibility.
4. **Information Disclosure (I)**: Unintended leakage or exposure of sensitive information.
5. **Denial of Service (D)**: Disruption of system availability, rendering it inaccessible.
6. **Elevation of Privilege (E)**: Gaining higher-level access than permitted, compromising the system's security structure.

These categories serve as the foundation for understanding and addressing the various threats that automotive systems face, helping to ensure that the security, and trustworthiness of connected vehicles remain intact throughout their operational lifecycle. Furthermore this table provides a mapping of the STRIDE categories to the 5 pillars of information assurance in terms of security property violation: Confidentiality Integrity Availability Authenticity and Non-repudiation. ***It shall be emphasized though that attacks against availability are not considered***.

Category	Description	Security property violation
Spoofing	It involves gaining unauthorised access to another user's authentication data, such as a login and password, and exploiting it for illicit purposes.	Authenticity
Tampering	It involves malicious modification of data, such as including unauthorised changes in the stored database or modification of a data item being transiting over a network	Integrity
Repudiation	It is related to users who deny having acted and there is no way to prove otherwise by other parties. For example, a user acts to enable the traceability of prohibited operations. Non-repudiation is the ability of a system to counter repudiation threats. For instance, a system has evidence to prove which user has created a purchase of an item, such as an authentication log.	Non-repudiation

Information Disclosure	It refers to the access of information by unauthorised users, whether real users or intruders. For instance, users may be able to read a file to which they were not given access.	Confidentiality
Denial of Service (DoS)	It refers to attacks to deny service to valid users. It interferes with the system's availability and reliability.	Availability
Elevation of Privilege	It refers to an unprivileged user gaining privileged access to the system, providing sufficient access to compromise or destroy the entire system, for example.	Authorisation

Table 4.1: STRIDE threat model categories

4.2 Threat cartography for CCAM ecosystems

Having identified the threat categories leveraging the STRIDE framework, the next step is to map these threats to the specific challenges in the *CCAM* landscape, as envisioned by *CONNECT*. This involves breaking down the potential threats into distinct classes that align with the unique components of CCAM systems. For this purpose, we have categorized the threats into the following three primary classes:

Threats on Sensors: Targeting the sensors that provide critical input to autonomous and connected vehicles, including manipulation or disruption of sensor data, leading to compromised decision-making or information leakage.

Threats on In-Vehicle Communication Systems: Focused on the internal vehicle networks, such as CAN bus, where unauthorized access, data tampering, or denial-of-service attacks could impact the vehicle's internal communications and operations.

Threats on V2X Communication: Involving vehicle-to-everything (V2X) communication, where attacks such as spoofing or data interception can disrupt the interaction between vehicles and external entities, including other vehicles, infrastructure, and cloud services.

The following tables outline the specific threats identified within each of these categories, providing a detailed overview of potential attack vectors and their implications within the CCAM ecosystem. These tables form the basis for a comprehensive threat analysis, facilitating the development of security measures and trust assessment mechanisms in the *CONNECT* framework.

Table 4.2: Threats on sensors

Threats on sensors	
Title	**Unauthorised access**
Description	Malicious actors can exploit vulnerabilities in the vehicle's internal network to gain control over sensor data.
Property violation	Authorisation.

Security Measures	Security controls like secure sensor interface, secure boot, firewalls, intrusion detection systems are put in place to reduce risk of compromise.
Title	**Tampering on sensor's data**
Description	Modification of sensor's data by an internal or external attacker to cause unpredictable vehicle behaviour.
Property violation	Integrity.
Security Measures	Security controls like firmware integrity, blockchain for data integrity, digital signatures, real-time sensor monitoring are put in place to reduce risk of integrity violations.
Title	**Information disclosure on sensor's data**
Description	Sensitive data collected by sensors can be exposed to unauthorised (internal or remote) parties.
Property violation	Confidentiality
Security Measures	Encryption of sensor data, access control mechanisms.
Title	**Spoofing**
Description	Attacks impersonate sensor's data or legit vehicle sensor to gain unauthorised access or control.
Property violation	Authenticity.
Security Measures	Sensor authentication, signal encryption.
Title	**Interference**
Description	Attacker interfere physically or digitally the sensor perception to tamper data or cause error.
Property violation	Denial of service, availability, integrity.
Security Measures	Sensor authentication, redundancy, signal encryption.

Table 4.3: Threats on in-vehicle communication system

Threats on in-vehicle communication system	
Title	**Man-in-the-middle attack**
Description	An attacker intercepts and potentially modifies communication between vehicle components.
Property violation	Authenticity.
Security Measures	Digital signatures, Authentication protocols
Title	**Message injection**
Description	An attacker injects false messages into the network to impersonate a legit component.

Property violation	Authenticity.
Security Measures	Digital signatures, Authentication protocols
Title	**Message modification**
Description	An attacker changes the context of the messages exchanged between the vehicle's components.
Property violation	Integrity.
Security Measures	Secure communication protocols, message integrity check
Title	**In-vehicle computer firmware modification**
Description	Unauthorised changes to the firmware of ECU to modify vehicle behaviour.
Property violation	Integrity.
Security Measures	Secure communication protocols, Secure over-the-air update
Title	**CAN bus flooding**
Description	Attacker floods the CAN bus with messages, causing traffic to be delayed or lost.
Property violation	Availability.
Security Measures	Network traffic filtering, prioritisation of critical messages, Intrusion detection system

Table 4.4: Threats on V2X communication

Threats on V2X communication	
Title	**Vehicle impersonation**
Description	An attacker pretends to be another vehicle or infrastructure to send false information.
Property violation	Authorisation
Security Measures	Mutual authentication, digital certificates for V2X devices, digital signature.
Title	**False data injection**
Description	An attacker injects incorrect data into the V2X network and denies responsibility
Property violation	Non-repudiation.
Security Measures	Digital signature on V2X messages
Title	**Perception data modification used in V2X message encoding**
Description	An attacker injects a malicious code which modifies data provided by the perception module to the V2X encoding module

Property violation	Integrity
Security Measures	Secure perception interface at the sender, Misbehaviour detection at the receiver
Title	**Compromised V2X message encoding module**
Description	An attacker injects a malicious code in the V2X encoding module to change the data fields in a CAM or a CPM message
Property violation	Integrity
Security Measures	Misbehaviour detection at the receiver
Title	**Sensor fault generates false sensor data that are used in encoding V2X message**
Description	A sensor reports false readings (due to a poor calibration for instance), has a high false positive rate or a high false negative rate (due to incorrect parametrization of the algorithms) and there is no sensor redundancy of the perceived area
Property violation	Integrity
Security Measures	Misbehaviour detection at the receiver
Title	**Attacker generates false sensor data that are used in encoding V2X message**
Description	A sensor reports false readings due to physical manipulation by an attacker (camera blurring, change sensor pose, etc). Thus, the sensor has a high false positive rate or a high false negative rate and there is no sensor redundancy of the perceived area
Property violation	Integrity
Security Measures	Misbehaviour detection at the receiver

4.3 TAF Consideration against the identified threats

This section examines how the Trust Assessment approach is expected to support mitigation of the identified threats, in accordance with the categorization outlined in the preceding tables. The Trust Assessment Framework is designed to detect and mitigate recognized threats in real time, leveraging verifiable trustworthiness evidence to provide ongoing monitoring. Through the identification of irregularities, assessment of trustworthiness (utilizing the *ATL (Actual Trustworthiness Level)*), and prompt response to potential threats, the TAF can dynamically modify the *RTL (Required Trustworthiness Level)* as necessary, hence providing resilient defences. The assessment of the TAF in WP3 will be essential, as it will undergo testing across multiple use cases. The TAF's flexibility in addressing various threat scenarios, including sensors, in-vehicle communication, and V2X communication, will guarantee reliability throughout the CCAM domain. Aligning threat mitigation with the TAF enhances the ecosystem's security posture against both current and emerging threats. The following tables analyse the TAF considerations per identified threat.

These considerations will guide evaluation of the TAF in WP3 and inside the use-cases.

Please note that certain threats may be excluded from the evaluation activities. This decision stems from our initial focus on the most prominent attack vectors to assess the applicability of the complex trust models proposed by *CONNECT*. As the first instance of such a comprehensive framework, the *CONNECT* TAF prioritizes demonstrating its effectiveness against a focused set of major threat vectors, as defined by the selected use cases. The following tables elaborate on the TAF considerations per threat category, while specifying the Use Case where each of them will be evaluated.

Table 4.5: Threats on sensors

Threats on sensors	
Title	**Unauthorised access**
Security Measures	Mutual authentication, digital certificates for V2X devices, digital signature.
Mitigating effect of TAF	If the proposed security and integrity measures from Table 4.2 are implemented, the added integrity protection can be expressed and quantified by a trust source and this contributes to overall trustworthiness within the trust model.
Implemented in which UC	Cooperative Adaptive Cruise Control (C-ACC), Slow Moving Traffic Detection (SMTD)
Title	**Tampering on sensor's data**
Security Measures	Firmware integrity, blockchain for data integrity, digital signatures, real-time sensor monitoring.
Mitigating effect of TAF	If the proposed security and integrity measures from Table 4.2 are implemented, the added integrity protection can be expressed and quantified by a trust source and this contributes to overall trustworthiness within the trust model.
Implemented in which UC	C-ACC
Title	**Information disclosure on sensor's data**
Security Measures	Encryption of sensor data, access control mechanisms.
Mitigating effect of TAF	As the scope of the TAF that CONNECT investigates will focus on integrity, TAF will only have an indirect mitigation effect that it can verify if other mitigations are in place.
Implemented in which UC	None
Title	**Spoofing**
Description	Attacks impersonate sensor's data or legit vehicle sensor to gain unauthorised access or control.
Security Measures	Sensor authentication, signal encryption.
Mitigating effect of TAF	If the proposed security and integrity measures from 4.2 are implemented, the added integrity protection can be expressed and quantified by a trust source and this contributes to overall trustworthiness within the trust model.

Implemented in which UC	C-ACC, SMTD
Title	**Interference**
Security Measures	Sensor authentication, redundancy, signal encryption.
Mitigating effect of TAF	If the proposed security and integrity measures from Table 4.2 are implemented, the added integrity protection can be expressed and quantified by a trust source and this contributes to overall trustworthiness within the trust model.
Implemented in which UC	None

In summary, the TAF will enable assessment of the degree of integrity protection of sensor data and to express this as a trust opinion in a trust model.

Table 4.6: Threats on in-vehicle communication system

Threats on in-vehicle communication system	
Title	**Man-in-the-middle attack**
Security Measures	Digital signatures, Authentication protocols
Mitigating effect of TAF	As the scope of the TAF that CONNECT investigates will focus on integrity, TAF will only have an indirect mitigation effect that it can verify if other mitigations are in place. A missing signature should lead to dropping of messages and therefore not be reflected in the TAF.
Implemented in which UC	C-ACC
Title	**Message injection**
Security Measures	Digital signatures, Authentication protocols
Mitigating effect of TAF	As the scope of the TAF that CONNECT investigates will focus on integrity, TAF will only have an indirect mitigation effect that it can verify if other mitigations are in place. A missing signature should lead to dropping of messages and therefore not be reflected in the TAF.
Implemented in which UC	C-ACC
Title	**Message modification**
Security Measures	Secure communication protocols, message integrity check
Mitigating effect of TAF	If the proposed security and integrity measures from Table 4.3 are implemented, the added integrity protection can be expressed and quantified by a trust source and this contributes to overall trustworthiness within the trust model.
Implemented in which UC	Intersection Movement Assistance (IMA), C-ACC
Title	**In-vehicle computer firmware modification**

Mitigating effect of TAF	If the proposed security and integrity measures from Table 4.3 are implemented, the added integrity protection can be expressed and quantified by a trust source and this contributes to overall trustworthiness within the trust model.
Implemented in which UC	C-ACC, SMTD
Title	**CAN bus flooding**
Security Measures	Network traffic filtering, prioritisation of critical messages, Intrusion detection system
Mitigating effect of TAF	As the scope of the TAF that CONNECT investigates will focus on integrity, TAF will only have an indirect mitigation effect that it can verify if other mitigations are in place.
Implemented in which UC	None

Table 4.7: Threats on V2X communication

Threats on V2X communication	
Title	**Vehicle impersonation**
Security Measures	Mutual authentication, digital certificates for V2X devices, digital signature.
Mitigating effect of TAF	As the scope of the TAF that CONNECT investigates will focus on integrity, TAF will only have an indirect mitigation effect that it can verify if other mitigations are in place. A missing signature should lead to dropping of messages and therefore not be reflected in the TAF. Beyond, the TAF could also be used as part of trust-based access control decisions to allow a more fine-grained authorization compared to classical DAC and RBAC.
Implemented in which UC	IMA, SMTD
Title	**False data injection**
Security Measures	Digital signature on V2X messages
Mitigating effect of TAF	As the scope of the TAF that CONNECT investigates will focus on integrity, TAF will only have an indirect mitigation effect that it can verify if other mitigations are in place. A missing signature should lead to dropping of messages and therefore not be reflected in the TAF.
Implemented in which UC	IMA
Title	**Perception data modification used in V2X message encoding**
Security Measures	Secure perception interface at the sender, Misbehaviour detection at the receiver
Mitigating effect of TAF	The TAF allows to integrate MBD output with other information of trustworthiness. It is therefore expected that more accurate decisions on trustworthiness of V2X message can be achieved.

Implemented in which UC	IMA, SMTD
Title	**Compromised V2X message encoding module**
Security Measures	Misbehaviour detection at the receiver
Mitigating effect of TAF	The TAF allows to integrate MBD output with other information of trustworthiness. It is therefore expected that more accurate decisions on trustworthiness of V2X message can be achieved.
Implemented in which UC	SMTD
Title	**Sensor fault generates false sensor data that are used in encoding V2X message**
Security Measures	Misbehaviour detection at the receiver
Mitigating effect of TAF	A TAF could already be deployed on the sender side to detect sensor faults as pointed out in Table 4.5. In addition, a TAF on the receiver side allows to integrate MBD output with other information of trustworthiness. It is therefore expected that more accurate decisions on trustworthiness of V2X message can be achieved.
Implemented in which UC	SMTD
Title	**Attacker generates false sensor data that are used in encoding V2X message**
Security Measures	Misbehaviour detection at the receiver
Mitigating effect of TAF	The TAF allows to integrate MBD output with other information of trustworthiness. It is therefore expected that more accurate decisions on trustworthiness of V2X message can be achieved.
Implemented in which UC	IMA

In summary, the TAF will enable assessment of the degree of integrity protection of communicated data (in or inter-vehicle/V2X) and to express this as a trust opinion in a trust model. In combination with MBD, it will also allow the integration of MBD with different other sources of trustworthiness evidence (like VCs) to allow for more accurate trust decisions than relying on MBD alone. As far as confidentiality or availability are concerned, this is out-of-scope for the TAF within the context of *CONNECT*. However, this does not imply that a similar approach could not support trustworthiness reasoning for confidentiality or availability.

4.4 Towards Trustworthiness Profiles for CCAM ecosystems

As introduced in D2.1 [4], current standards efforts focus on developing trustworthiness profiles, where trustworthiness is defined as the probability that the trustee will meet the trustor's expectations in a given context. These expectations are shaped by the specific scenario, methodology, and objectives of the trust relationship. To establish trust, it is essential to achieve a shared understanding of these expectations and implement an assurance strategy that provides confidence

in their fulfilment. Standards have started addressing trustworthiness profiles by introducing the overall architecture and corresponding architecture profiles.

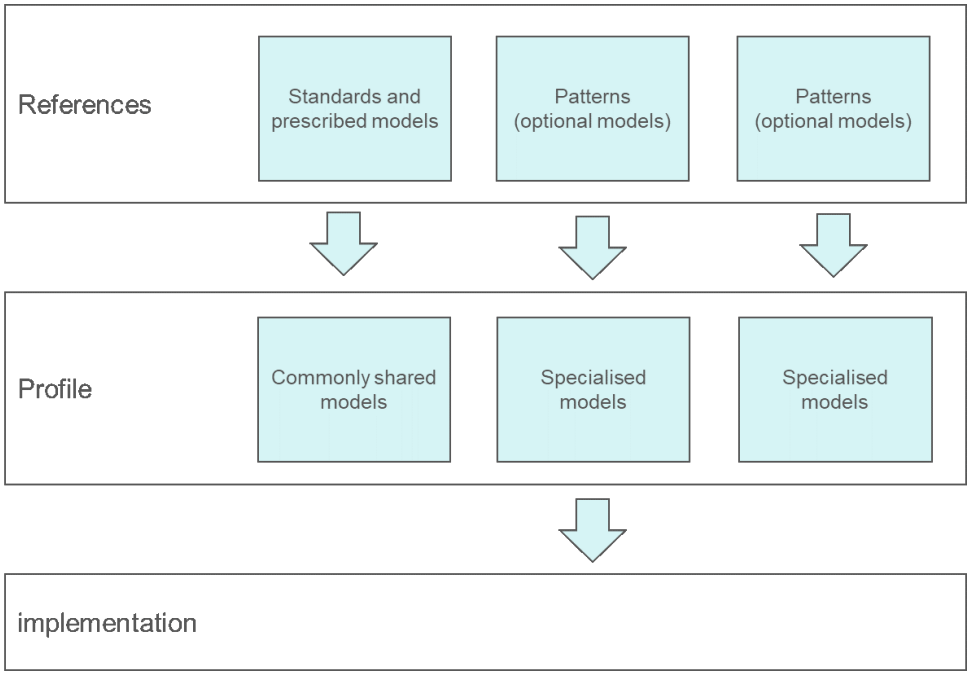

Figure 4.1: Creation of an architecture profile

The process of creating an architecture profile builds on the foundational work of ISO/IEC JTC 1/SC 41 (IoT and digital twin) and ISO/IEC JTC 1/SC 7/WG 42 (software and systems engineering, architecture), as outlined in standards such as ISO/IEC 40141 (guidance on reference architecture), ISO/IEC/IEEE 42024 (architecture fundamentals), and ISO/IEC/IEEE 42042 (reference architecture). The process involves selecting a set of references—comprising standards, prescribed models, and optional patterns—to define a profile best-tailored for the target application domain (i.e., safety-critical CCAM systems). This profile is made up of commonly shared and specialized models, ensuring that any implementation aligns with the defined profile. For example, a profile related to the CONNECT activities may include prescribed models like the ISO/IEC 30188 digital twin reference architecture associated with data patterns and mechanisms implemented by CONNECT digital twin-based instantiated at the MEC. Figure 4.1 illustrates this process.

In this context, Figure 4.2 further refines the process for defining trustworthiness profiles. We have to highlight than although current standards predominantly concentrate on safety and security criteria necessary for a trustworthiness profile, as detailed in D2.1 [4], *CONNECT* is among the pioneers in addressing an enhanced definition of trust that encompasses the full spectrum of requirements outlined in various standards such as ETSI and ITU-T. Operationalizing this conceptual trustworthiness model is a critical step, involving several activities depicted in Figure 4.3: listing trustworthiness characteristics and their associated requirements, identifying goals and related capabilities, assessing risks and their treatments, and defining assurance requirements in terms of claims, arguments, and evidence.

The initial step in defining trustworthiness characteristics is identifying the relationships among the properties of trust pertinent to all components within the CCAM continuum. This process enables a semantic-level understanding of the factors contributing to trust and their interdependencies. By identifying these relationships, it becomes possible to prioritize aspects most critical

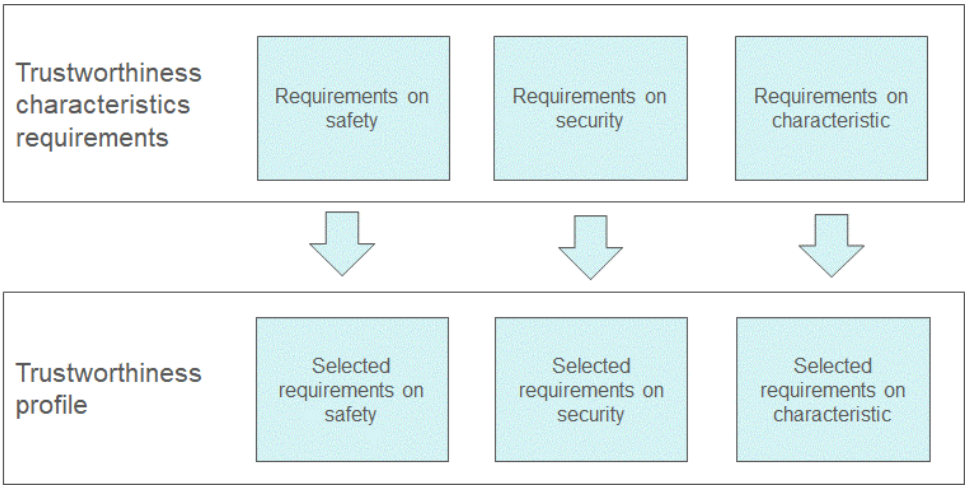

Figure 4.2: Creation of a trustworthiness profile

to industrial stakeholders—such as over-the-air software updates for safety-critical CCAM applications like Cooperative Adaptive Cruise Control (C-ACC) in *CONNECT* —as well as regulatory compliance and the needs of other target groups. An example of additional factors that can guide the identification of such relationships has already been documented as part of ISO 26262 series. For instance, ISO/SAE 21434 evaluates security in relation to key criteria, including safety, financial impact, operational considerations, and privacy, ensuring a holistic approach to trustworthiness.

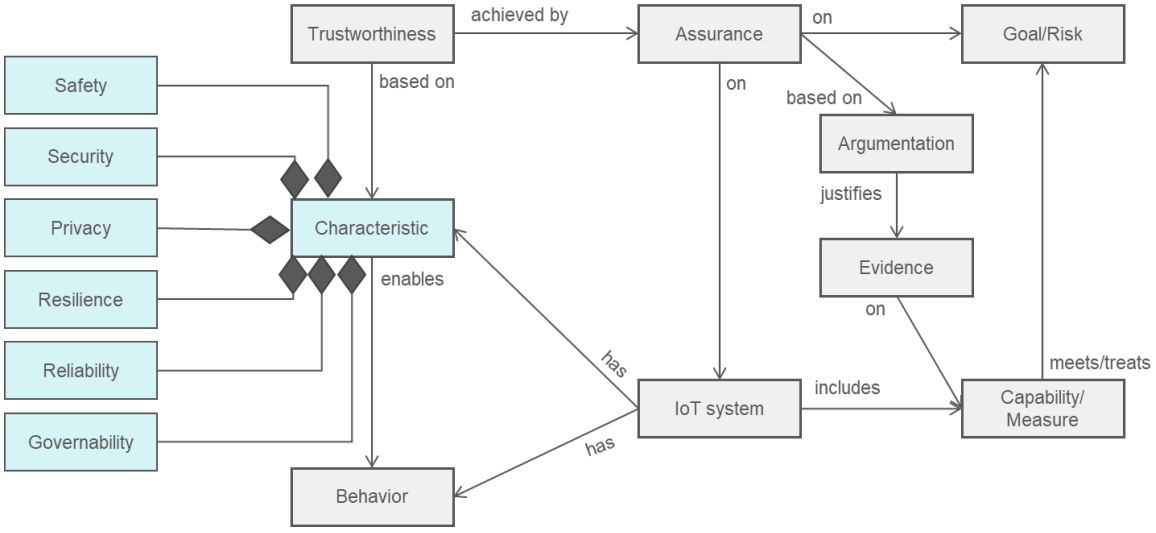

Figure 4.3: Conceptual model for trustworthiness

Following these criteria and aiming to contribute to standards, ***CONNECT* envisions to establish such trustworthiness profiles for the CCAM sector. These profiles capture and represent semantic-level relationships across all CCAM layers**, encompassing additional characteristics such as robustness, resilience and availability of a CCAM. This approach ensures the definition of application-level trustworthiness profiles, abstracting the translation of such requirements in a holistic manner so as to represent all elements in the automotive supply chain. Moreover, *CONNECT* defines the methodologies and mechanisms for realising this vision into real-world CCAM systems, materialising the goal to support a secure transition to Level 3 autonomous driving capabilities. The results of this process will be documented in ISO/IEC 27568, where

CONNECT will be showcased as a practical use case for instantiating trustworthiness profiles.

Chapter 5

CONNECT Technical Requirements and KPIs

This section outlines the technical requirements and Key Performance Indicators (KPIs) for the Trust Assessment, Security and Operational Correctness, MEC Operational and Security, and Privacy components within the *CONNECT* framework. The Trust Assessment Framework (TAF) is required to provide accurate, real-time trust evaluations across vehicles and MEC environments, ensuring flexibility and scalability. Security and runtime operational correctness focus on robustness, resilience through cryptographic mechanisms, credential management, and evidence correctness, leveraging Trusted Computing Base functionalities to maintain trustworthiness under varying conditions. For MEC, the emphasis is on low-latency performance, operational effectiveness, integrity, and secure data handling to support safety-critical CCAM applications. Privacy requirements are centred on data confidentiality, minimal disclosure, and compliance with regulations such as GDPR. Corresponding KPIs evaluate performance overhead, resource utilization, and reliability to ensure that *CONNECT*ś enablers meet the demands of real-time, trustworthy, and secure CCAM operations. The following sections will delve into the specific technical requirements of *CONNECT*.

5.1 Trust Assessment Requirements

This subsection presents the requirements that the Trust Assessment Framework (TAF) must meet, as previously defined in D2.1 [4]. As a core component of the *CONNECT* framework, the TAF evaluates the trustworthiness of nodes and data exchanged within CCAM applications across vehicles, Multi-access Edge Computing (MEC), and the cloud. These requirements emphasize the essential characteristics necessary to address CONNECT use cases and safety-critical CCAM applications, ensuring measurable and reliable outcomes. Beyond technical functionality, the TAF also considers trust relationship establishment, user confidence in system reliability, transparency, and compliance with privacy regulations. The six key requirements for the TAF are: i) Generalizability, ii) Runtime Performance, iii) Scalability, iv) Correctness, v) Robustness and Resilience, and vi) Flexibility of Trust Sources.

Table 5.1: FR.TR.1 Generalizability

FR.TR.1 (Mandatory)		
Title	Generalizability	
Actors 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.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.2: FR.TR.2 Performance

FR.TR.2 (Mandatory)		
Title	Run-time Performance	
Actors Involved	Trust Assessment Framework	
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 [13]. 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 **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	$\leq$ **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). $\leq$ **200 ms** delay when the TAF is instantiated and executed within the CONNECT TEE. These calculations represent the time it takes for the TAF to respond with an Actual Trustworthiness Level (ATL) or a Trust Decision (TD) after it receives a Trust Assessment Request (TAR) from an application. Note that this does not include the time needed to collect and process data communicated by other entities (in-vehicle components or another vehicle). We also exclude any network latency caused by the reception of trust sources from neighbouring vehicles or the MEC. The requirement focuses only on the TAF internal operation and the trust level calculation after all the necessary input is there.

Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.3: FR.TR.3 Scalability

FR.TR.3 (Mandatory)		
Title	Scalability	
Actors 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: $\leq$ **10 nodes** (Single TAF) For large scale environment (i.e., IMA) testing to be done: $\leq$ **50 nodes** either vehicle or MEC-running (Federated TAF)
	Timeframe to complete trust assessment	$\leq$ **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) $\leq$ **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.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.4: FR.TR.4 Correctness

FR.TR.4 (Mandatory)		
Title	Correctness	
Actors Involved	Trust Assessment Framework	
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 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 produce a correct ATL for a proposition. 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 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.	
Connected to other requirements		
KPIs	**Description**	**Value**
	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.	
	TAF result correctness for **1st scenario**: All entities are trustworthy (i.e., not compromised by an attacker)	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.

	TAF result correctness for **2nd scenario**: One or several entities are not trustworthy because they have been compromised by an attacker.	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. We expect that in $\geq$ **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.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.5: FR.TR.5 Robustness and Resilience

FR.TR.5 (Mandatory)	
Title	Robustness and Resilience
Actors Involved	Trust Assessment Framework
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	All TAF internal components to be able to be instantiated inside a TEE.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.6: FR.TR.6 Flexibility of Trust Sources

FR.TR.6 (Mandatory)		
Title	Flexibility of Trust Sources	
Actors Involved	Trust Assessment Framework	
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, 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**

KPIs	Trust Sources	$\geq$ **3 different trust sources are included in the use cases**. These trust sources are collected so that they are verifiable through trustworthiness claims.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

5.2 Security and Operational Assurance Requirements

In addition to the Trust Assessment Requirements, it is crucial to define traditional security requirements to protect against a variety of threats and vulnerabilities, ensuring cybersecurity and system reliability within the emerging CCAM landscape. These traditional security requirements focus on core principles such as confidentiality, integrity, availability, authentication, and non-repudiation. They address essential concerns, such as safeguarding sensitive data, ensuring data integrity, maintaining system availability, verifying user identities, and preventing the repudiation of actions. By categorizing security requirements in this way, organizations and system designers can first address foundational security needs, then extend these considerations to trust-related issues, fostering resilient and trustworthy systems in the rapidly evolving autonomous vehicle (AV) ecosystem.

Alongside these security requirements, operational assurance requirements are explored, providing a clear distinction between security specifications and operational guarantees. This includes: i) security specifications, which encompass traditional security concerns; ii) runtime operational correctness, which ensures the accuracy of evidence used for trust guarantees; and iii) function isolation and migration, which protect critical trust framework functions and key-related operations. This comprehensive approach enables the creation of robust, secure, and trustworthy systems in the context of connected and automated mobility.

In what follows all requirements are based on the assumption of a Root of Trust (RoT) capable of supporting the capabilities defined in D2.1 (e.g., RoT for measurement and related functionalities). However, to maintain focus on the KPIs, we have chosen not to replicate the detailed property descriptions here.

5.2.1 Security Specifications

Table 5.7: FR.SR.1 Dynamic Credential Management

FR.SR.1 (Mandatory)	
Title	Dynamic Credential Management
Actors 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 [24].
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.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2.	

Table 5.8: FR.SR.2 Secure and Efficient Cryptography

FR.SR.2 (Mandatory)	
Title	Secure and Efficient Cryptography
Actors 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)	$\leq$ **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).	$\leq$ **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.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.9: FR.SR.3 Flexible and Reliable Key Management

FR.SR.3 (Mandatory)	
Title	Flexible and Reliable Key Management
Actors 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. 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.	
Connected to other requirements	FR.SR.1, FR.SR.2, FR.SR.4	
KPIs	**Description**	**Value**
	Overhead against using keys with and without key restriction usage policies.	**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 remote attestation processes that CONNECT will provide for extracting information evidence as trust sources.
	Efficiency of key hierarchy construction of different types of keys.	$\leq$ **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.)
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.10: FR.SR.4 Secure Data Handling and Provenance

FR.SR.4 (Mandatory)	
Title	Secure Data Handling and Provenance
Actors 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 [16]. 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)
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

5.2.2 Runtime Operational Correctness Evidence as Attributes for Trust Provisioning

Table 5.11: FR.OC.1 Common Trusted Computing Protocols

FR.OC.1 (Mandatory)	
Title	Common Trusted Computing Protocols
Actors 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 [ ]. 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 [20][29]. Note that, this integrity protection mechanism must consider the possibility of both software and hardware compromise. In chapter 8.1.4 in D2.1 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 in D2.1 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**

KPIs	Granularity of Levels of Assurance (LoA) that can be achieved by various RoTs and essentially the common trusted computing base	$\geq$ **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.).	$\geq$ **3 operations**
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.12: FR.OC.2 Operational Assurance & Configuration Integrity

FR.OC.2 (Mandatory)	
Title	Operational Assurance & Configuration Integrity
Actors Involved	Trusted Computing Base
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**
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.13: FR.OC.3 Integrity Verification of CCAM Components

FR.OC.3 (Mandatory)	
Title	Integrity Verification of CCAM Components
Actors 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**: Time needed for the execution of configuration integrity verification **Value**: $<$ **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
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.14: FR.OC.4 Chain of Trust Creation

FR.OC.4 (Mandatory)	
Title	Chain of Trust Creation
Actors 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).

<table>
<tr><td></td><td>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. 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.</td></tr>
<tr><td>Connected to other requirements</td><td>FR.OC.1, FR.SR.2</td></tr>
<tr><td>KPIs</td><td>
<table>
<tr><th>Description</th><th>Value</th></tr>
<tr><td>Storage of trust-related information to the Blockchain for auditability and certifiability.</td><td>$\leq$ 5sec, since this is not a real-time operation.</td></tr>
<tr><td>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).</td><td>$\geq$ 1 Mbit/sec data transfer rate for the raw evidence</td></tr>
</table>
</td></tr>
</table>

	Simulate low bandwidth channel with high message loss.	$<$ **1 sec** for the BC transaction
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.15: FR.OC.5 Secure Measurement/Attribute Extraction

FR.OC.5 (Mandatory)		
Title	Secure Measurement/Attribute Extraction	
Actors 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 [22], 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	$\leq$ **500 ms**
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.16: FR.OC.6 Secure Remote Asset Management and Reconfiguration Effectiveness

FR.OC.6 (Mandatory)	
Title	Secure Remote Asset Management and Reconfiguration Effectiveness
Actors 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**
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

5.2.3 Function Isolation and Migration

Table 5.17: FR.SF.1 Dynamic awareness on potential vulnerabilities and threats and complete overview of the deployed environment

FR.SF.1 (Mandatory)	
Title	Dynamic awareness on potential vulnerabilities and threats and complete overview of the deployed environment
Actors 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 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
Connected to other requirements	TR.4
KPIs	**Description** / **Value** Dynamic Risk Assessment based on the identification of new threats: $\leq$ **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. Recalculation of RTL considering the identification of new risks: $\leq$ **3 sec**
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.18: FR.SF.2 Stateful Function Upgrade

FR.SF.2 (Mandatory)	
Title	Stateful Function Upgrade
Actors Involved	Trusted Execution Environment

Description	**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 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.

	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.
Connected to other requirements	TR.4
KPIs	**Description**: Upgrading the function in one ECU — **Value**: $<$ **750 ms**
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.19: FR.SF.3 StatefulStateful Function Migration

FR.SF.3 (Mandatory)	
Title	Stateful Function Migration
Actors Involved	Trusted Execution Environment
Description	**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. 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” enco mpasses 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).

Current Status: Evaluation is ongoing and final assessment will be documented in D6.2

5.3 MEC Operational and Security Requirements

CONNECT leverages edge computing technology to support trust assessment mechanisms and verifiable evidence collection for safety-critical CCAM applications. While not focused on MEC innovation, the project uses key MEC features like data locality, low latency, real-time computation, and virtualization to enable trusted CCAM operations. MEC requirements are categorized into baseline, mandatory, and nice-to-have operational requirements, with the baseline features aligned with ETSI specifications and evaluated through network KPIs (e.g., latency). Mandatory MEC security requirements are based on evolving standards, ensuring real-time trust evaluations and secure data handling. CONNECT's approach integrates MEC technology within a secure, scalable framework for CCAM systems.

5.3.1 MEC Operational Requirements

Table 5.20: FR.MEC.1

FR.MEC.1 (Baseline)	
Title	Operational requirements on (MEC) application lifecycle and application environment
Actors Involved	CONNECT MEC
Description	**Background**: Virtualization technology [17] (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 been introduced and established in the ETSI MEC specifications with which the CONNECT MEC technology is typically aligned. • 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. • 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.

<table>
<tr><td></td><td>

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.
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.
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 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:</td></tr>
<tr><td>KPIs</td><td>
<table>
<tr><th>Description</th><th>Value</th></tr>
<tr><td>Container LCM Operations [referring to Instantiation, termination, modification time]</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.</td></tr>
<tr><td>Container LCM Operations [referring to Instantiation, termination, modification time]</td><td>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).</td></tr>
<tr><td>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).</td><td>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 messages will also include the Trustworthiness Claims (target value for overhead posed $<$15%)</td></tr>
<tr><td>CONNECT Application Layer Data-rate [Uplink (UL) and downlink (DL) measurements]</td><td>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 : 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]</td></tr>
</table>
</td></tr>
<tr><td>Current Status</td><td>Evaluation is ongoing and final assessment will be documented in D6.2</td></tr>
</table>

Table 5.21: FR.MEC.2

FR.MEC.2 (Optional)		
Title	Operational Requirements for mobility support	
Actors 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 [11] 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). his 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	
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 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).
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.22: FR.MEC.3

FR.MEC.3 (Baseline)	
Title	Operational Requirements for MEC services
Actors 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.
KPIs	**Description** / **Value** 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
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.23: FR.MEC.4

FR.MEC.4 (Baseline)	
Title	Operational Requirements for (applications) connectivity
Actors 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 latency requirements between container communications — Expected value <**100 ms**. (measured as IP ICMP Echo Request packet leaving the source host (e.g., container A, host A) until the IP ICMP Echo Reply is received by the destination host (e.g., container B, host B).
	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
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

5.3.2 MEC Security Requirements

Table 5.24: SR.MEC.1

SR.MEC.1 (Mandatory)	
Title	Security Requirements on MEC service authorization and access (authorised service access)
Actors 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.

<table>
<tr><td></td><td colspan="2">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</td></tr>
<tr><td>Connected to other requirements</td><td colspan="2">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</td></tr>
<tr><td>KPIs</td><td>Description</td><td>Value</td></tr>
<tr><td></td><td>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.
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.</td><td>Authentication leveraging both CONNECT adopted mechanisms (i.e., PKI) as well as CONNECT developed authentication controls (i.e., leveraging the newly-developed trust extension, should take $<$ than 2 sec (excluding network latency) .
We have to highlight that this operation focus on the authentication and onboarding of the vehicle into one administrative domain managed by one single MNO (not considering accessing a service expanding between different administrative domains).</td></tr>
</table>

	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.	The proposed KPI will also consider the overhead posed by *CONNECT* trustworthiness claims that need to be exchanged during authentication. In the context of size during the authentication handshake (i.e., $<$ **20%**) so that it does not incur packet fragmentation. In the context of latency at MEC application-level [with expected value in the range of **10 to 200ms**].
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.25: SR.MEC.2

SR.MEC.1 (Mandatory)		
Title	Security Requirements on virtualization and containerization technology (employed -among others- at the MEC)	
Actors 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 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.	
Connected to other requirements	FR.SR.2, FR.SR.3, FR.OC.1, FR.OC.2, FR.OC.3, FR.OC.4, FR.OC.5, FR.OC.6	
KPIs	**Description**	**Value**
	Degree of coverage of the (defined) virtualization/containerization threats.	The focus of this KPI will be on capturing attacks against integrity violations of the enclavised containers ($>$**95%** of attacks targeting both the secure launch and operational integrity of the confidential containers).

Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

5.4 Privacy Requirements

Privacy and trust are critical concerns in the CCAM domain, especially regarding vehicles and vulnerable road users. While standardized protocols like vehicular PKI have addressed privacy-respecting identity management through digital certificates and authentication, privacy challenges extend beyond vehicle tracking. A recent study by the Mozilla Foundation revealed how car brands collect and share deeply personal data through sensors, cameras, apps, and devices, sometimes selling this data to third parties. To mitigate these privacy risks, comprehensive frameworks are needed, incorporating data anonymization, encryption, and unlinkability techniques. However, it's essential to balance privacy with system performance, as privacy-enhancing technologies (e.g., pseudonym changes) may complicate trust management and long-term reputation building.

Table 5.26: FR.PR.1

FR.PR.1 (Optional)	
Title	Unlinkability of Representation Artefacts or Repeated System Interactions
Actors 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 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
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.27: FR.PR.2

FR.PR.2 (Optional)	
Title	Unlinkability of Data Provenance
Actors 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. 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.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.28: FR.PR.3

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 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.

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 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 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.

KPIs

Description	Value
Attributes disclosed as part of the trustworthiness evidence should not enable vehicle fingerprinting.	**TRUE**

	Controlled linkability should be provided only to authenticated entities (i.e., OEMs) in case of indication of risks of internal vehicle components.	**TRUE**
	Secure construction of TCs including harmonised attributes	$<$ **200 ms**
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Table 5.29: FR.PR.4

FR.PR.4 (Optional)	
Title	Privacy Preservation in Multi-MNO Service Domains
Actors Involved	Operators of MEC, and Operators of V2X infrastructure
Description	**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 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 [ ]. However, it still remains open to understand whether and how much privacy protection would be reduced by following this approach.
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2

Table 5.30: FR.PR.5

FR.PR.5 (Mandatory)		
Title	Trust Information and Assessment Lifecycle Management and User Acceptance	
Actors 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**
	GDPR analysis that the type of (trust-related) information exchanged within CONNECT do not pose additional privacy implications.	**TRUE**
Current Status	Evaluation is ongoing and final assessment will be documented in D6.2	

Chapter 6

Definition of the *CONNECT* Minimal Viable Platform (MVP)

In the present chapter, the tables that map the use case requirements to the *CONNECT* framework specification are provided, as defined in D2.1 and evaluated in the context of D6.1 and D6.2. The outcome highlight the *CONNECT* framework's pivotal role in securing computations and connections across diverse scenarios. Guided by the principle of "Never Trust, Always Verify," the envisioned *CONNECT* architecture aims to transform emerging CCAM ecosystems into a distributed, collaborative infrastructures with embedded trust across the entire continuum, stemming from far edge (i.e., vehicles to cloud-edge). The framework further leverages trusted computing and blockchain technologies to address the requirements of various vertical industry sectors, enabling efficient, reliable, and secure extraction and sharing of threat intelligence. It also supports perceived zero-trust environments characterized by heterogeneous device hardware densities and high mobility demands.

6.1 *CONNECT* Use Cases Overview

6.1.1 Cooperative Adaptive Cruise Control (C-ACC)

Co-operative Adaptive Cruise Control (C-ACC) uses *Vehicle-to-Everything (V2X)* to extend the traditional *Adaptive Cruise Control (ACC)*. It uses sensor data from surrounding vehicles and infrastructure, in addition to its sensors, to derive the safer optimal distance to the preceding vehicle, considering the current speed, acceleration and position of the involved vehicles. As a demonstrator, this use case intends to evaluate the in-vehicle trustworthiness, from a component and sensor data perspective, for upcoming and imminent driving situations.

6.1.2 Intersection Movement Assist (IMA)

The Intersection Movement Assist application leverages the V2X shared data to estimate potential collision risks in intersection zones. More precisely, it uses Cooperative Perception Messages (CPM) received from V2X sources, which contain the kinematic description of the perceived data in the environment. When a potential collision risk is detected with the perceived objects, the IMA

raises a collision warning. The IMA use case shows the necessity to evaluate reliably the V2X data and sources trustworthiness for a reliable collision risk assessment in the intersection areas.

6.1.3 Slow-Moving Traffic Detection (SMTD)

Slow-Moving Traffic Detection use case leverages V2X messages to detect and inform regarding the presence and position of a slow moving vehicle in front. The presence and position of this slow vehicle is detected through sensors of the vehicle behind, which is equipped both with ADAS sensors (e.g., cameras, radars, lidars, etc.) and with V2X capabilities. The V2X-equipped vehicle will create CAMs for signalling its own presence and position and CPMs for signalling the presence and position of the slow moving vehicle in front. Each of those messages are sent to a Traffic Control Center, monitoring the situation on the road. It is important that all the messages sent are associated with a trustworthiness level to assess the goodness of the message flows.

6.2 Mapping of UC requirements to *CONNECT* technical requirements

ID	Title	IMA	C-ACC	SMTD
Trust Assessment				
FR.TR.1	Generalizability	L	L	M
FR.TR.2	**Performance**	**H**	**H**	**H**
FR.TR.3	**Scalability**	**H**	**M**	**H**
FR.TR.4	**Correctness**	**H**	**H**	**M**
FR.TR.5	Robustness and Resilience	H	H	L
FR.TR.6	Flexibility of Trust Sources	H	L	M
Security				
FR.SR.1	**Dynamic Credential Management**	**M**	**H**	**H**
FR.SR.2	Secure and Efficient Cryptography	L	H	L
FR.SR.3	Flexible and Reliable Key Management	M	H	M
FR.SR.4	**Secure Data Handling and Provenance**	**H**	**H**	**M**
Runtime Operational Correctness				
FR.OC.1	Common Trusted Computing Protocols	M	H	L
FR.OC.2	Operational Assurance & Configuration Integrity	H	H	L
FR.OC.3	**Integrity Verification of CCAM Components**	**H**	**H**	**H**
FR.OC.4	Chain of Trust Creation	H	H	L
FR.OC.5	Secure Measurement/Attribute Extraction	L	H	L
FR.OC.6	Secure Remote Asset Management and Reconfiguration Effectiveness	L	H	L
Function Isolation and Migration				

FR.SF.1	Dynamic awareness on potential vulnerabilities and threats and complete overview of the deployed environment	L	H	M
FR.SF.2	Stateful Function Upgrade	L	H	M
FR.SF.3	Stateful Function Migration	L	H	H
MEC-Functional				
FR.MEC.1	Operational requirements on (MEC) application lifecycle and application environment	H	N/A	H
FR.MEC.2	**Operational Requirements for mobility support**	**H**	N/A	**H**
FR.MEC.3	**Operational Requirements for MEC services**	**H**	N/A	**H**
FR.MEC.4	**Operational Requirements for (applications) connectivity**	**H**	N/A	**H**
MEC-Security				
SR.MEC.1	Security Requirements on MEC service authorization and access (authorised service access)	M	N/A	M
SR.MEC.2	Security Requirements on virtualization and containerization technology (employed -among others- at the MEC)	N/A	N/A	L
SR.MEC.3	Security Requirements and compatibility with (available/standardised) Application Programming Interfaces (APIs)	L	N/A	L
(non-functional) Privacy				
FR.PR.1	Unlinkability of Representation Artefacts or Repeated System Interactions	L	L	L
FR.PR.2	Unlinkability of Data Provenance	L	N/A	L
FR.PR.3	Attributes Related to Vehicle Trustworthiness should not Create Privacy Threats with Medium or High Level beyond those Already in Existence in the CCAM Ecosystem	M	M	M
FR.PR.4	Privacy Preservation in Multi-MNO Service Domains	N/A	N/A	M
FR.PR.5	Trust Information and Assessment Lifecycle Management and User Acceptance	M	N/A	L

Table 6.1: Requirement priorities by use case

6.3 *CONNECT* Most Valuable Product (MVP)

The *CONNECT* Most Valuable Product (MVP), as derived from Table 6.1 consists of a total of 9 “must-have” requirements, along with 17 additional “must-have” requirements for the underlying

trusted component TR.ROT, as outlined in D2.1. Additionally, there are 22 “should-have” requirements that further enhance the platform’s capabilities. The nine “must-have” requirements are proposed as the foundational core of the platform, designed to meet the most critical operational needs and represent the essential *CONNECT* services utilized in the use cases. These core requirements will be validated comprehensively within the context of the use cases to ensure their effectiveness and reliability.

More specifically, the MVP derives from the common requirements across the three use cases (IMA, C-ACC, and SMTD) as identified in Table 6.1. The overall goal of *CONNECT* is to develop a robust and high-performance Trust Assessment Framework (TAF) that ensures operational correctness, security, and privacy, while also maintaining the flexibility to accommodate future enhancements. A key focus of this product is its ability to deliver real-time trust evaluations, supporting a high level of i) performance (FR.TR.2), ii) scalability (FR.TR.3), and iii) correctness (FR.TR.4) in trust assessments across diverse CCAM environments.

In addition to the core operation of the TAF framework, *CONNECT* ensures the authenticity (FR.SR.1) and configuration integrity (FR.OC.3) of in-vehicle components, as well as the provenance of the collected evidence (FR.SR.4). This is achieved through verifiable credentials and dynamic credential management mechanisms (FR.SR.1), which provide reliable and trustworthy evaluations of nodes and data within the system. Configuration integrity is further maintained by monitoring CCAM components through defined policies, ensuring that components operate in a trusted and consistent state throughout their lifecycle. Additionally, the framework emphasizes secure data handling, ensuring that sensitive information is properly protected and traceable across its entire lifecycle. This is achieved through controlled linkability (FR.SR.4), which ensures that all data is securely associated with its origin, while maintaining privacy and accountability within the system.

The operational effectiveness of the MEC infrastructure is optimized to provide mobility support (FR.MEC.2), ensuring seamless integration with dynamic and mobile environments. This optimization enables the framework to support low-latency and secure services (FR.MEC.3), which are essential for real-time data processing and communication in CCAM systems. Furthermore, the MEC infrastructure meets the operational requirements for CCAM applications (FR.MEC.4), ensuring that the system is responsive and reliable under diverse conditions.

Chapter 7

Conclusion and Outlook

The present deliverable marks the completion of a milestone of the CONNECT project, presenting the finalized architecture of the framework alongside its key components and operational workflows. It addresses the key challenges of dynamic trust evaluation and secure data exchange, applying CONNECT's vision to enhance security, privacy, and trustworthiness in Cooperative, Connected, and Automated Mobility (CCAM) systems. These efforts provide a strong foundation for advancing CCAM ecosystems and their deployment in real-world scenarios. Trusted computing anchors and novel trust modeling enable dynamic, evidence-based trust assessments tailored to CCAM services. In addition, the deliverable integrates an ethical analysis, addressing potential biases in trust decision-making and aligning the framework with Trustworthy AI principles and regulatory standards.

Another key aspect of this deliverable is its comprehensive threat modeling, based on the STRIDE model, which identifies and addresses risks specific to CCAM systems. This analysis bolsters the framework's resilience, ensuring it remains robust against emerging security challenges. The refined technical and use case requirements, along with updated KPIs, further demonstrate CONNECT's focus on meeting practical needs and ensuring its framework remains adaptable and effective. This deliverable also builds upon the first framework release by providing updates and refinements to components, information flows, and operational pipelines. The integration of synchronous and asynchronous processes further enhances the framework's ability to support dynamic operations across the CCAM landscape. These improvements not only validate the initial design but also lay the foundation for upcoming advancements, including the incorporation of Distributed Ledger Technology (DLT) for dynamic updates and the use of indirect evidence in trust assessments.

CONNECT takes a comprehensive approach by embedding security and trust mechanisms throughout the CCAM ecosystem. By addressing both technical robustness and ethical considerations, the framework is designed to adapt to the complexity of CCAM environments while maintaining its focus on privacy and security. This approach ensures the CONNECT framework is not only innovative but also aligned with broader societal and regulatory expectations. Looking ahead, the next phases of the project will focus on extending the component evaluation and use case validation activities. In conclusion, this deliverable underscores the significant progress made in CONNECT, setting a clear path forward for achieving its vision of a secure, trustworthy, and dynamic CCAM framework. By combining technical excellence with ethical oversight, CONNECT is well-positioned to address the challenges of next-generation mobility and deliver a framework that is both innovative and reliable.

Appendix A

Glossary and User Roles

ABAC Attribute-based Access Control.

ACC Adaptive Cruise Control.

A-ECU An *A-ECU* is an *ECU* with a *Trusted Execution Environment (TEE)* providing secure storage for keys and other data. it is able to do asymmetric and symmetric cryptography.

AIV Attestation and Integrity Verification.

ATL (Actual Trustworthiness Level) The ATL reflects the result of an evaluation of a specific (atomic or complex) proposition for a specific scope provided by the TLEE. It quantifies the extent to which a certain node or data can be considered trustworthy based on the available behavioural evidence.

C-ACC Co-operative Adaptive Cruise Control.

CCAM The European Commission has on 30th of November 2016 adopted a European Strategy on Cooperative Intelligent Transport Systems (C-ITS), a milestone initiative towards cooperative, connected and automated mobility. The objective of the C-ITS Strategy is to facilitate the convergence of investments and regulatory frameworks across the EU, in order to see deployment of mature C-ITS services in 2019 and beyond [13].

DENM The DENM is standardised in ETSI EN 302 637-3 V1.3.1 [14]. 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..

DLT Distributed Ledger Technology.

ECU An electronic control unit (ECU), also known as an electronic control module (ECM). In automotive electronics it is an embedded system that controls one or more of the electrical systems or subsystems in a car or other motor vehicle.

FL The Facility Layer (FL) manages the (kinematic) data stemming from the in-vehicle sensors by relaying them to all components of the Vehicle Manager that have subscribed to receive them. These are essentially the: (i) CAM/CPM Encoder/Decoder component that will start the construction of the respective V2X messages (e.g., CAM, CPM) to be broadcasted either following the currently ETSI specified standards or including also the Verifiable Presentations (VPs) comprising the trust-related information outputed by the TCH (i.e., T-CAM/T-CPM messages), (ii) CCAM Application module for constructing the local view of the vehicle's vicinity towards supporting the decisions making process of the service (e.g., breaking, changing lanes, etc.), (iii) the Misbehavior Detection service for checking the vercity of the measures kinematic data through a series of plausibility checks, and (iv) Trust Assessment Framework (TAF) for associating a Trust Opinion to each data object. It is important to highlight that the FL doesn't perform any processing or checks to the received data. Any verification controls, especially for asserting to the integrity of the data and its safety in the context of been signed and processed by only "*certified*" applications is performed by the IAM.

IAM Identity and Authentication Management.

LoA ETSI Levels of Assurance.

MBD The Mis-behaviour Detector (MBD) component monitors the data from the vehicle and from elsewhere (from CPM/CAM messages) and looks for anomalies. If these are detected is sends mis-behaviour reports to the TAF and outside of the vehicle. Reports for the TAF will be 'normally' signed, while those being sent outside will be anonymously signed.

MEC The MEC serves a number of functions. It makes more powerful computing resources available to vehicles. These resources are provided close to the edge of the network so that calculations can be 'outsourced' by the vehicles and still meet the necessary low latency requirements. It can also combine information from vehicles in its vicinity to produce a more detailed map of their positions and trajectories and feed this back to them together with its assessment of their trustworthiness.

OEM An *Original Equipment Manufacturer*. In the context of CONNECT the OEM is the vehicle manufacturer who, often in association with a *Tier 1 Supplier* supplier, designs, assembles, markets and sells the vehicle.

RTL (Required Trustworthiness Level) Is a technical measure representing the required trust level in a system/entity, calculated at design-time. More specifically it identified the minimum level of belief and the maximum level of acceptable uncertainty and/or disbelief for a vehicle function or data either created or received. The RTL is influenced by factors such as risk analysis, technical demands, and requirements or demands forced by requirements to obtain technical approval to produce the car, for example. Although RTL cannot rely on actual evidence, it can be derived from risk assessment tools, assumptions, and observing expected cybersecurity scenarios. It is used as a mean for comparison (i.e, threshold) with the ATL to derive to a trust decision..

SCB Security Context Broker.

S-ECU An *S-ECU* is an *ECU* with secure storage for keys and other data, possibly a *System on Chip (SoC)* with an HSM. It can only do symmetric crypto.

SGX *Intel SGX* is a hardware feature of Intel CPUs that provides a *TEE* for user-space applications on Intel CPUs. The goal is to protect an application from unauthorized access or modification by any component outside the TEE. I.e. neither the operating system nor other untrusted applications should be able to breach the confidentiality or integrity of the protected application.

SoC A *system on a chip or system-on-chip (SoC)* is an integrated circuit that integrates most or all components of a computer or other electronic system. These components almost always include on-chip central processing unit (CPU), memory interfaces, input/output devices and interfaces, and secondary storage interfaces, often alongside other components such as radio modems and a graphics processing unit (GPU) – all on a single substrate or microchip.

TA (Trust Assessment Manager) The TAM is a component inside the TAF which orchestrates the overall process of trust assessment.

TAR The Trust Assessment Request (TAR) is the triggering point, initiated by a CCAM application, for requesting from the CONNECT Trust Assessment Framework a trust opinion on the data exchanged within this specific service.

T-DENM In CONNECT T-DENM are regular DENM messages, extended such that they can accommodate *TC*. *TC* are included in a periodic fashion in T-DENM produced by a given station, so that not all T-DENM by that station are expected to include *TC*..

T-CPM In CONNECT T-CPMs are regular CPM messages, extended such that they can accommodate *TC*. *TC* 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 *TC*..

T-CAM In CONNECT T-CAMs are regular CAM messages, extended such that they can accommodate *TC* for the purposes of CONNECT. *TC* 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 *TC* ..

TAF The TAF component does the trust assessments and forms trust opinions on the vehicle and data. The trust opinion on the data is sent outside the vehicle and needs to be anonymously signed.

TC Is a structure that contains the harmonised trustworthiness evidence and is included in the form of a VP in the CAM/CPM messages that are disseminated from a vehicle to another or from a vehicle to the infrastructure.

TCB The *Trusted Computing Base (TCB)* of a computer system is the set of all hardware, firmware, and/or software components that are critical to its security, in the sense that bugs or vulnerabilities occurring inside the TCB might jeopardize the security properties of the entire system. By contrast, parts of a computer system that lie outside the TCB must not be able to misbehave in a way that would leak any more privileges than are granted to them in accordance to the system's security policy.

TCH The Trustworthiness Claims Handler (TCH) is the component responsible for sharing all trust-related information outside the Vehicle in a privacy-preserving manner. This data bundle (encoded in the context of a VP) comprises Trustworthiness Claims (TCs), the Trust Opinion (produced by the TAF) and the Misbehavior Report (produced by the MBD). The

TC is usually produced (by the Attester) so as to provide trustworthiness evidence ("Trust Source") that can be used for appraising the trustworthiness level of the Attester in a **measurable** and **verifiable** manner. Measurable reflects the ability of the TAF to assess an attribute of the Attester against a pre-defined metric (e.g., RTL) while verifiability highlights the need for all claims to have integrity, freshness and to be provably & non-reputably bound to the identity of the original Attester. Examples sets of TCs might include (among other attributes) evidence on system properties including: (i) integrity in the context that all transited devices (e.g., ECUs) have booted with known hardware and firmware; (ii) safety meaning that all transited devices are from a set of vendors and are running certified software applications containing the latest patches and (iii) communication integrity.

TEE A *Trusted Execution Environment* allows to execute applications while enforcing well-defined security policies for a given application. An example is *Intel Software Guard Extensions (Intel SGX)*.

Tier 2 A *Tier 2 supplier* provides components to the *Tier 1* suppliers and is the next level in the supply chain. Tier 2 suppliers may not just provide components for the automotive industry, but other industries as well. For CONNECT we focus on the suppliers of ECUs (micro-controllers) used in the vehicle and their role in providing identity keys for them.

Tier 1 A *Tier 1 supplier* directly supplies *OEMs* with components that are ready for installation into the vehicle. They work closely with the *OEM* at all stages of a vehicle's development. The Tier 1 supplier may well work with several manufacturers on the development of their vehicles. The Tier 1 supplier will obtain the components that they need from *Tier 2 Supplier* suppliers.

TMT Trust Model Template.

V2X Vehicle-to-Everything.

VC Vehicle Communication (V2X) This provides communication facilities for the vehicle. Connectivity – automotive ethernet, 5G, V2X.

VP The Verifiable Presentation (VP) is the data structure used for disclosing only a subset of the trust-related information needed for the receiving entity to evaluate the trust level of the originator. This allows the *TCH* to construct data bundles that hold the Trust Opinion, Misbehavior Report and "abstracted" attestation assertions, as described in D5.1 [3].

Zonal controller (ZC) The *Zonal controller (ZC)* is an *A-ECU* that acts as a gateway between the ECUs and the vehicle computer. As an *A-ECU* they will have a *TEE* providing secure storage for keys and other data and will be able to do asymmetric and symmetric cryptography.

References

[1] A. Baier. Trust and antitrust. *Ethics*, 96(2):231–260, 1986.

[2] The CONNECT Consortium. Conceptual architecture of customisable tee & attestations. Deliverable D4.1, Project 101069688 within HORIZON-CL5-2021-D6-01, Dec. 2023.

[3] The CONNECT Consortium. Distributed processing and CCAM trust functions offloading & data space modelling. Deliverable D5.1, Project 101069688 within HORIZON-CL5-2021-D6-01, Nov. 2023.

[4] The CONNECT Consortium. Operational landscape, requirements and reference architecture - initial version. Deliverable D2.1, Project 101069688 within HORIZON-CL5-2021-D6-01, Nov. 2023.

[5] The CONNECT Consortium. Distributed processing, fast offloading and MEC-enabled orchestrator. Deliverable D5.2, Project 101069688 within HORIZON-CL5-2021-D6-01, Mar. 2024.

[6] The CONNECT Consortium. Integrated framework (first release) and use case analysis. Deliverable D6.1, Project 101069688 within HORIZON-CL5-2021-D6-01, May 2024.

[7] The CONNECT Consortium. Trust & risk assessment and CAD twinning framework (initial version). Deliverable D3.2, Project 101069688 within HORIZON-CL5-2021-D6-01, February 2024.

[8] The CONNECT Consortium. Virtualization- and edge-based security and trust extensions (first release). Deliverable D4.2, Project 101069688 within HORIZON-CL5-2021-D6-01, Jan. 2024.

[9] The CONNECT Consortium. Virtualization- and edge-based security and trust extensions (final release). Deliverable D4.3, Project 101069688 within HORIZON-CL5-2021-D6-01, Mar. 2025.

[10] Shi Dong, Junxiao Tang, Khushnood Abbas, Ruizhe Hou, Joarder Kamruzzaman, Leszek Rutkowski, and Rajkumar Buyya. Task offloading strategies for mobile edge computing: A survey. *Computer Networks*, 254:110791, 2024.

[11] EK Eijnden. Optimal handover in mec for an automotive application. Master's thesis, University of Twente, 2020.

[12] ETSI. Network Functions Virtualisation (NFV); Trust; Report on Attestation Technologies and Practices for Secure Deployments. Technical Report GR NFV-SEC 007 V1.1.1, ETSI, October 2017.

[13] ETSI. Intelligent Transport Systems (ITS); Cooperative Adaptive Cruise Control (CACC); Pre-standardization study. Technical Report ETSI TR 103 299 V2.1.1, ETSI Technical Report, 2019. Accessed: 2024-10-21.

[14] European Telecommunications Standards Institute. 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, ETSI, 2019.

[15] European Telecommunications Standards Institute (ETSI). Network Functions Virtualisation (NFV) Trust; Report on Attestation Technologies and Practices for Secure Deployments. ETSI Group Report GR NFV-SEC 007 V1.1.1, ETSI, October 2017. https://www.etsi.org/deliver/etsi_gr/NFV-SEC/001_099/007/01.01.01_60/gr_nfv-sec007v010101p.pdf.

[16] Amrita Ghosal and Mauro Conti. Security issues and challenges in v2x: A survey. *Computer Networks*, 169:107093, 2020.

[17] Bo Han, Vijay Gopalakrishnan, Lusheng Ji, and Seungjoon Lee. Network function virtualization: Challenges and opportunities for innovations. *IEEE communications magazine*, 53(2):90–97, 2015.

[18] ISO/IEC TS 5723:2022. Trustworthiness — vocabulary. ISO, 2022. Retrieved February 2, 2023, from https://www.iso.org/standard/81608.html.

[19] ISO/SAE 21434:2021(E). Road vehicles– cybersecurity engineeting. Standard, ISO & SAE International, Geneva, CH, ago 2021.

[20] Mohit Jangid et al. Towards a tee-based v2v protocol for connected and autonomous vehicles. In *Workshop on Automotive and Autonomous Vehicle Security (AutoSec) 2022*, 2022.

[21] Sami Kekki, Walter Featherstone, Yonggang Fang, Pekka Kuure, Alice Li, Anurag Ranjan, Debashish Purkayastha, Feng Jiangping, Danny Frydman, Gianluca Verin, et al. Etsi white paper: Mec in 5g networks. *The European Telecommunications Standards Institute (ETSI), Tech. Rep. ETSI White Paper*, 28, 2018.

[22] Chongkyung Kil, Emre C Sezer, Ahmed M Azab, Peng Ning, and Xiaolan Zhang. Remote attestation to dynamic system properties: Towards providing complete system integrity evidence. In *2009 IEEE/IFIP International Conference on Dependable Systems & Networks*, pages 115–124. IEEE, 2009.

[23] Geun-Yong Kim, Ryangsoo Kim, Sungchang Kim, Ki-Dong Nam, Sung-Uk Rha, and Jung-Hyun Yoon. Dnn inference offloading for object detection in 5g multi-access edge computing. In *2021 International Conference on Information and Communication Technology Convergence (ICTC)*, pages 389–392, 2021.

[24] Benjamin Larsen, Thanassis Giannetsos, Ioannis Krontiris, and Kenneth Goldman. 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, page 48–59, New York, NY, USA, 2021. Association for Computing Machinery.

[25] J. D. Lee and K. A. See. Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 2004.

[26] F. D. Llorca and G. E. Gutierrez. Trustworthy autonomous vehicles. Technical Report EUR 30942 EN, Publications Office of the European Union, Luxembourg, 2021. JRC127051.

[27] K. O'Hara. A general definition of trust. University of Southampton, August 2012.

[28] Independent High-Level Expert Group on Artificial Intelligence. Ethics guidelines for trustworthy ai. European Commission, 2019. Retrieved from https://digital-strategy.ec.europa.eu/en/policies/expert-group-ai.

[29] Takahito Yoshizawa, Dave Singelée, Jan Tobias Muehlberg, Stéphane Delbruel, Amir Taherkordi, Danny Hughes, and Bart Preneel. A survey of security and privacy issues in v2x communication systems. *ACM Computing Surveys*, 55(9):1–36, 2023.