Misbehaviour Detection for V2X: Operational Aspects

5GAA Automotive Association
White Paper

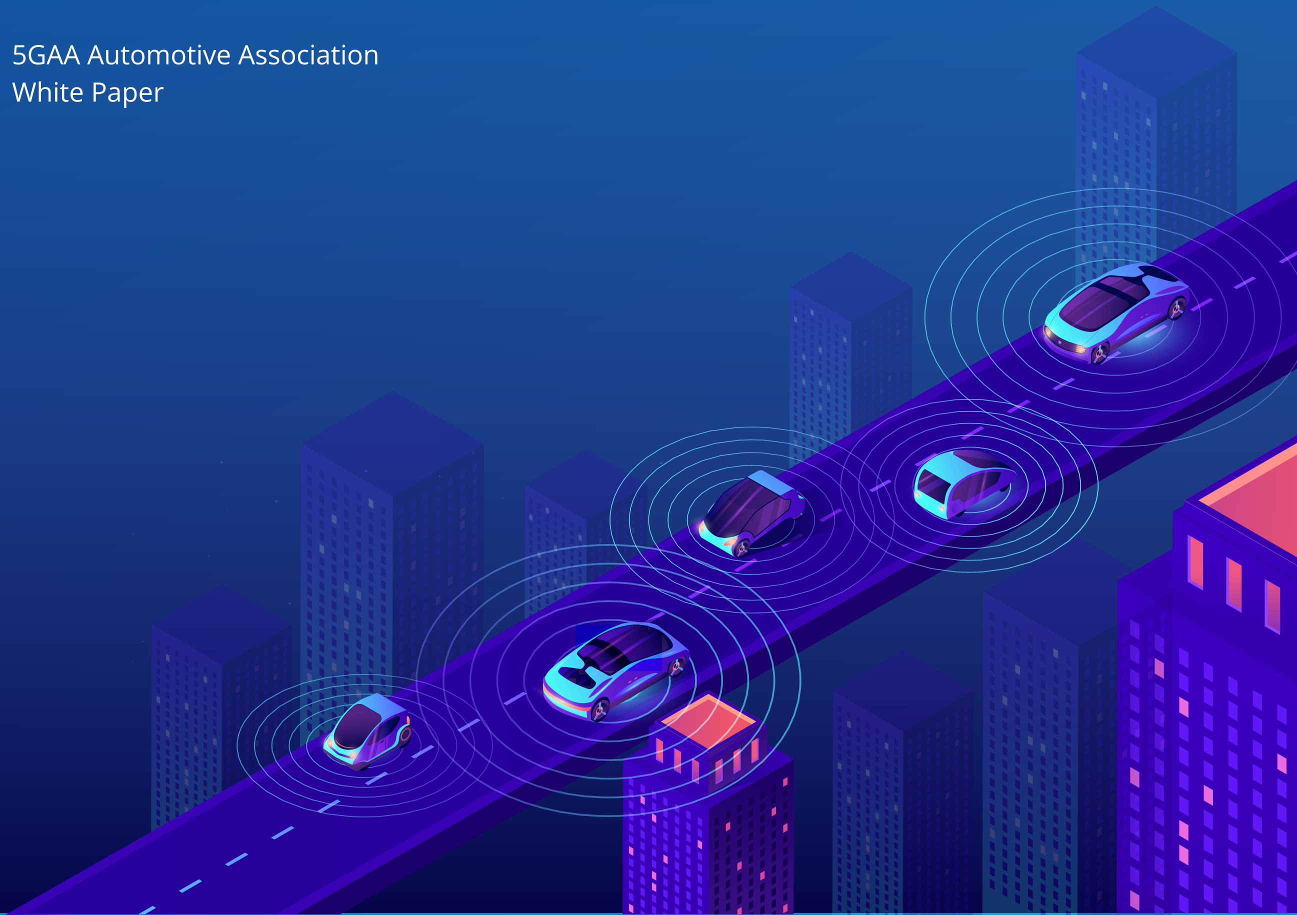

CONTACT INFORMATION:
Executive Manager – Thomas Linget
Email: liaison@5gaa.org

MAILING ADDRESS:
5GAA c/o MCI Munich
Neumarkter Str. 21
81673 München, Germany
www.5gaa.org

VERSION:	1.0
DATE OF PUBLICATION:	29 May 2024
DOCUMENT TYPE:	White Paper
EXTERNAL PUBLICATION:	Yes
DATE OF APPROVAL BY 5GAA BOARD:	12 February 2024

Contents

1 Scope

The scope of this work item is to organise the current state of the art for vehicle-to-everything (V2X) misbehaviour management (local and global detection, reporting, remediation) in a form accessible to non-experts of the field. This white paper aims to serve as a starting point and as an input document for any future technical or policy specifications on V2X misbehaviour management. Here and in the rest of this white paper, V2X refers to direct broadcast communication. Misbehaviour management for V2X network (using cellular and backend) based communication is out of scope.

One important aspect of a misbehaviour management system is how misbehaviour by a particular sender is remediated when it is discovered. There are a number of remediation techniques that could be used, of which permanent revocation is perhaps the best known. V2X systems being deployed have not yet specified the conditions to determine when a particular remediation technique is to be used. A lack of consistent remediation conditions can lead to outcomes that are perceived as unfair. This work does not attempt to specify those but attempts to describe what conditions might be used by policy organisations to define clear remediation conditions.

2 References

The following documents contain provisions which, through reference in this text, constitute provisions of the present document.

- References are either specific (identified by date of publication, edition number, version number, etc.) or nonspecific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies.

[1] [5GAA-MBD] 5GAA White Paper on Misbehaviour Detection (2022-07). https://5gaa.org/content/uploads/2022/07/5GAA-Misbehaviour-detection-Final.pdf

[2] [Brecht+2018] Brecht, B., Therriault, D., Weimerskirch, A., Whyte, W., Kumar, V., Hehn, T., Goudy, R. A Security Credential Management System for V2X Communications. IEEE Trans. Intell. Transp. Syst. 19(12): 3850-3871 (2018).

[3] [ETSI-302665] ETSI EN 302 665 V1.1.1 (2010-09). Intelligent Transport Systems (ITS); Communications Architecture

[4] [ETSI-102940] ETSI TS 102 940 V2.1.1 (2021-07). Intelligent Transport Systems (ITS); Security; ITS communications security architecture and security management; Release 2. https://www.etsi.org/deliver/etsi_ts/102900_102999/102940/02.01.01_60/ts_102940v020101p.pdf

[5] [ETSI-102941] ETSI TS 102 941 V2.2.1 (2022-11). Intelligent Transport Systems (ITS); Security; Trust and Privacy Management; Release 2. https://www.etsi.org/deliver/etsi_ts/102900_102999/102941/02.02.01_60/ts_102941v020201p.pdf

[6] [ETSI-103097] ETSI TS 103 097 V2.1.1 (2021-10). Intelligent Transport Systems (ITS); Security; Security header and certificate formats; Release 2. https://www.etsi.org/deliver/etsi_ts/103000_103099/103097/02.01.01_60/ts_103097v020101p.pdf

[7] [ETSI-103759] ETSI TS 103 759 V2.1.1 (2023-01). Intelligent Transport Systems (ITS); Security; Misbehaviour Reporting service; Release 2. https://www.etsi.org/deliver/etsi_ts/103700_103799/103759/02.01.01_60/ts_103759v020101p.pdf

[8] [IEEE-1609.2] IEEE 1609.2-2022. IEEE Approved Draft Standard for Wireless Access in Vehicular Environments--Security Services for Applications and Management Messages. https://standards.ieee.org/ieee/1609.2/10258/

[9] [IEEE-1609.2.1] IEEE 1609.2.1-2022. IEEE Standard for Wireless Access in Vehicular Environments (WAVE) - Certificate Management Interfaces for End Entities. https://standards.ieee.org/ieee/1609.2.1/10728/

3 Preliminaries

3.1 Terms and definitions

Many terms used in this document are explained in the WG2 document 5GAA_A-170188_V2XDEF_TR, '5GAA V2X Terms and Definitions'. The following definitions also apply:

ITS object: An ITS object (ITSO) (e.g., a roadside unit (RSU), an onboard unit (OBU), etc.) is a computing/communication system that creates/sends/receives V2X messages (conformant or not to the relevant message specification).

Misbehaviour: Misbehaviour within the V2X system refers to behavior that impedes an ITS object's ability to obtain an accurate understanding of the ground truth in its vicinity.

Misbehaviour authority: A component of the V2X ecosystem that receives reports of malicious or potentially malicious application activities, analyses them, and determines whether to take mitigating actions.

Misbehaviour management: Misbehaviour management refers to the entire lifecycle of misbehaviour, which includes local misbehaviour management on an ITS object, misbehaviour report transmission from an ITS object to the backend, and then misbehaviour management at the backend.

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

ASN.1	Abstract Syntax Notation One
BSM	Basic Safety Message
CAM	Cooperative Awareness Message
CRL	Certificate Revocation List
DENM	Decentralised Environmental Notification Message
DoS	Denial of Service
EE	End Entity
ETSI	European Telecommunications Standards Institute
GMBD	Global Misbehaviour Detection
HW	Hardware
IEEE	Institute of Electrical and Electronics Engineers
IR	Immediate Response
ISO	International Organisation for Standardisation
ITS	Intelligent Transportation System
ITSO	ITS Object
ITS-S	ITS Station
LMBD	Local Misbehaviour Detection
MA	Misbehaviour Authority
MBD	Misbehaviour Detection

MBDR	Misbehaviour Detection and Remediation
MBMS	Misbehaviour Management System
MPR	Minimum Performance Requirements
OBU	Onboard Unit
OEM	Original Equipment Manufacturer
OTA	Over-The-Air
PKI	Public Key Infrastructure
RA	Registration Authority
RC	Reinstatement Requirement Category
RSU	Roadside Unit
SCMS	Security Credential Management System
ST	Suspension Technique
SW	Software
V2X	Vehicle-to-Everything

4 Introduction

An intelligent transportation systems (ITS) object (ITSO) (e.g., a roadside unit (RSU), an onboard unit (OBU), etc.) is a computing/communication system that creates/sends/receives V2X messages (conformant or not to the relevant message specification). Note that an ITSO is a broader class of devices or systems than an intelligent transportation system station (ITS-S) (cf. ETSI EN 302 665). ITSO is defined to capture vehicle-to-everything (V2X) misbehaviour involving messages that may or may not conform to any V2X message specification.

For purposes of this document, misbehaviour refers to behaviour by one ITS object that impacts another ITS object's ability to obtain an accurate understanding of the ground truth in its vicinity. This includes willful or inadvertent transmission of bad data, meaning data that can result in bad driving or information outcomes if it is believed to be true by a receiver. It also covers behaviour such as channel jamming and other Denial of Service (DoS) attacks, and attempts by senders to execute commands or have requests responded to where the sender is not entitled to that response by the receiver. Examples include:

- A vehicle sending incorrect data (position, speed, acceleration, etc.) that results in a receiving vehicle miscalculating the kinematics of the sender and hence, either raising a false alert to its driver, or worse causing a safey incident.
- A regular passenger vehicle pretending to be an emergency response vehicle and sending a signal preemption request to a traffic signal.
- A transmission device creating ghost vehicles by sending properly generated V2X messages.

There can be many types of misbehaviour targeting different parts/aspects of the V2X ecosystem. The primary focus of this document are misbehaviours that involve ITS objects as the sender and/or receiver of misbehaving messages. Here and throughout this document, V2X misbehaviour refers to only those misbehaviour in the V2X ecosystem that involve ITS objects. It is important to detect and manage misbehaviours in a timely manner, because a persistent and/or widespread misbehaviour can negatively impact the potential benefits of V2X communications thereby discouraging honest users from participating in the system because it gives them false warnings or bad outcomes.

Misbehaviour management refers to the entire lifecycle of misbehaviour, which includes local misbehaviour management on an ITS object, misbehaviour report transmission from an ITS object to the backend, and then misbehaviour management at the backend. The misbehaviour management system (MBMS) proposed in [5GAA-MBD, ETSI-103759] and shown here in Figure 1 has three main components of the backend misbehaviour management, namely, misbehaviour preprocessing, misbehaviour authority (which in turn includes misbehaviour investigation and misbehaviour analysis), and misbehaviour remediation. Similarly, the ITSO side of misbehaviour management has five components: local misbehaviour detection, context storage, misbehaviour reporting, local misbehaviour reaction, and local misbehaviour remediation. Only the misbehaviour remediation components (highlighted in orange) of the backend and

ITSO misbehaviour management are within the scope of the current document.

When an ITSO detects and reports misbehaviour locally, the backend first analyses the report and determines if the misbehaviour indeed took place and who was responsible for it. Once the backend has made that determination, the next step is for the backend to take action. There are a number of remediation techniques available, including a software/hardware (HW/SW) update, pausing certificate issuance, and certificate revocation. It is important that remediation actions are applied fairly (across the V2X population) and proportionately (to the damage caused by the misbehaviour). However, currently there are no specifications of conditions applied to determine when a particular remediation technique is to be used. As such, a lack of consistent remediation conditions can lead to outcomes that may be unfair/inadequate. The main goal of this document is to provide technical guidance for developing policies and procedures around backend misbehaviour remediation.

The rest of this document is organised as follows: Sections 5 through 7 cover the core contributions, where Section 5 details the different remediation options available for V2X misbehaviour management, Section 6 explores different ways to classify misbehaviour, and Section 7 discusses different approaches to mapping remediations that are appropriate for any given misbehaviour. Finally, Section 8 concludes the white paper and briefly discusses open problems and the next steps for misbehaviour management.

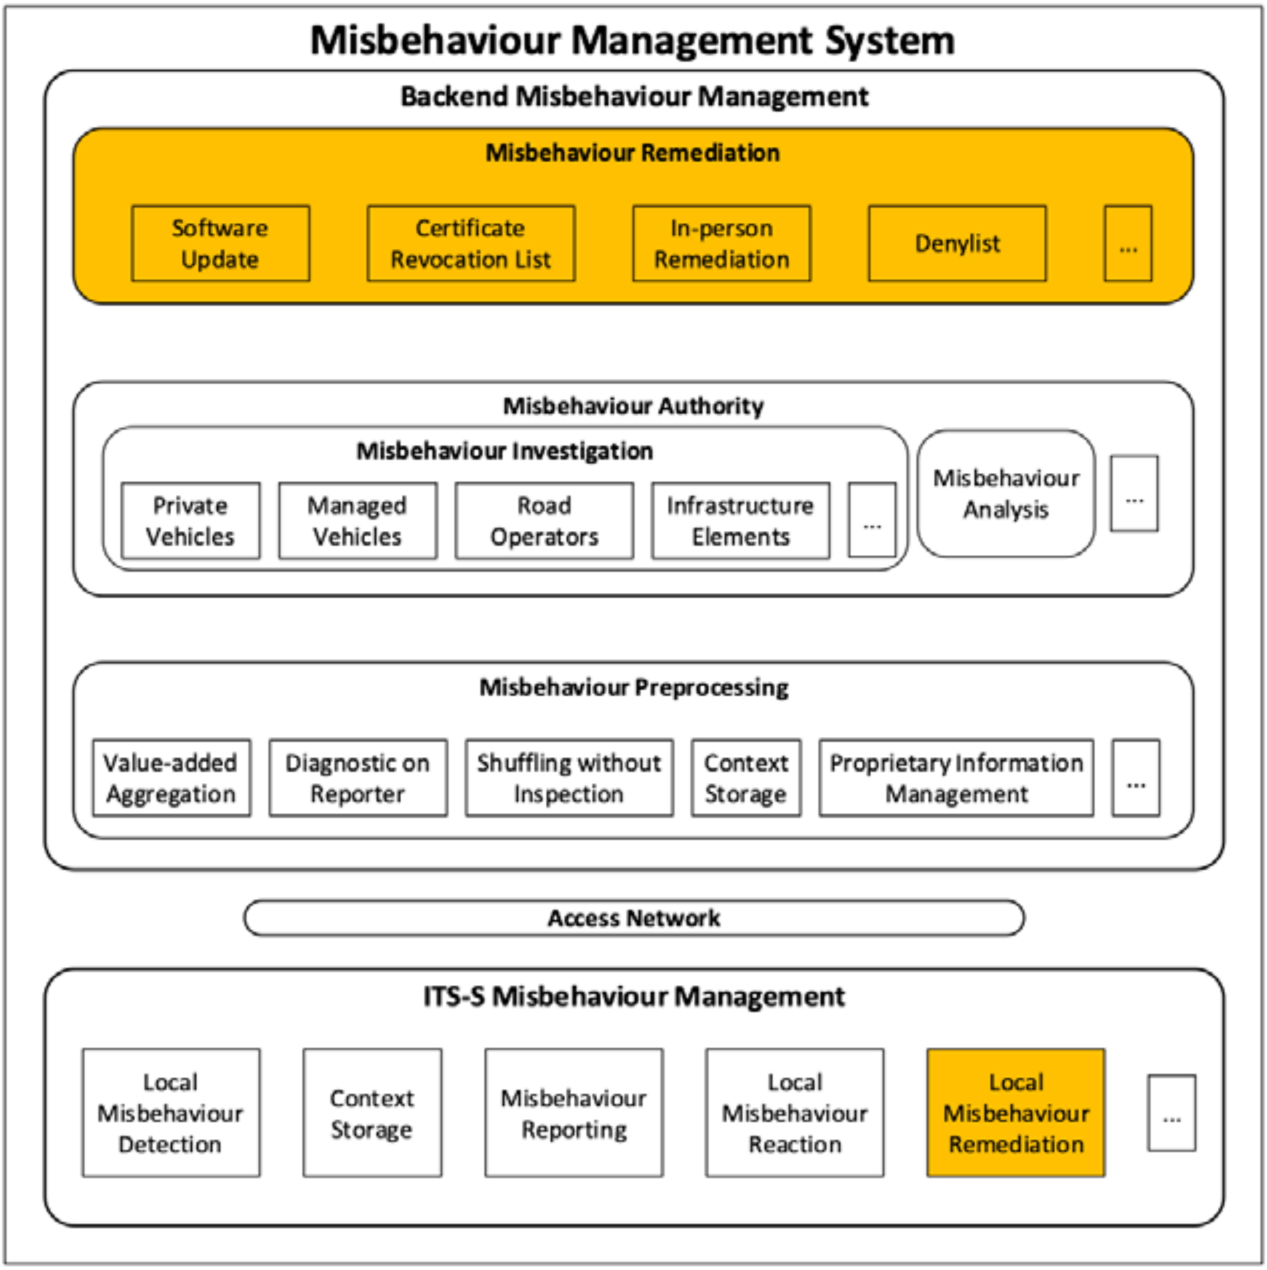
Misbehaviour Management System
Backend Misbehaviour Management
Misbehaviour Remediation
Software Update
Certificate Revocation List
In-person Remediation
Denylist
...
Misbehaviour Authority
Misbehaviour Investigation
Private Vehicles
Managed Vehicles
Road Operators
Infrastructure Elements
...
Misbehaviour Analysis
...
Misbehaviour Preprocessing
Value-added Aggregation
Diagnostic on Reporter
Shuffling without Inspection
Context Storage
Proprietary Information Management
...
Access Network
ITS-S Misbehaviour Management
Local Misbehaviour Detection
Context Storage
Misbehaviour Reporting
Local Misbehaviour Reaction
Local Misbehaviour Remediation
...

Figure 1: Misbehaviour management system

5 Remediation classification

5.1 General

As shown in Figure 2, at a high level misbehaviour remediation techniques can be divided into the following four categories:

1. Sender local: Remediation is **fully contained** locally at the end entity (EE) and affects the **send-side** behaviour, e.g., self-diagnostics.
2. Receiver local: Remediation is **fully contained** locally at the EE and affects the **receive-side** behaviour, e.g., ignore lists.
3. Sender non-local: Remediation is **not contained** locally at the EE and affects the **send-side** behaviour, e.g., software updates.
4. Receiver non-local: Remediation is **not contained** locally at the EE and affects the **receive-side** behaviour, e.g., certificate revocation lists (CRLs).

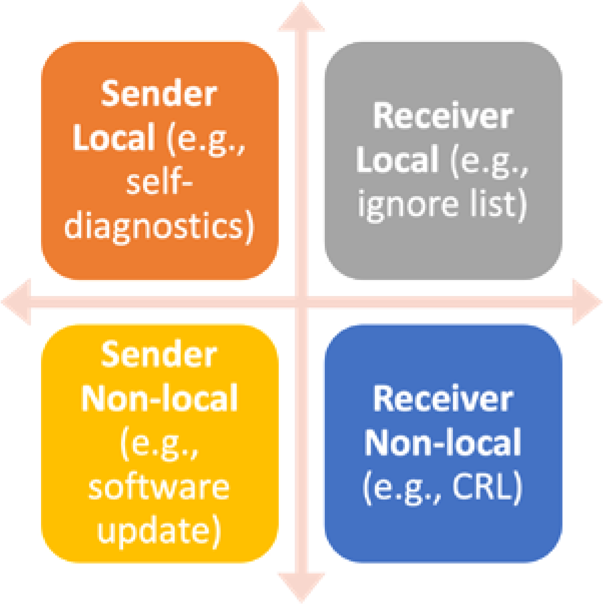
Sender Local (e.g., self-diagnostics)
Receiver Local (e.g., ignore list)
Sender Non-local (e.g., software update)
Receiver Non-local (e.g., CRL)

Figure 2: Misbehaviour remediation – bird's eye view

The focus of this work is non-local remediation as that requires interactions among the different components of the misbehaviour management system (MBMS) and EEs, which may further necessitate standardisation efforts, stakeholder best practices, and policy interventions.

Once the MBMS has determined the existence and source (i.e., the misbehaving EE) of a misbehaviour, remediation can begin. Misbehaviour remediation is an inherently iterative and non-deterministic process. Assume there are a certain number of steps for misbehaviour remediation. For any specific misbehaviour: the order of the steps may differ, some of them may be processed multiple times, and some of the steps may be omitted.

NOTE: How the MBMS determines the existence and source of a misbehaviour is out of scope.

What follows is a list of remediation steps. One step is implicit (continuous evaluation of the system to see what step needs to be taken next), and therefore it is called ‘step 0’. As explained above, the steps listed below may not be taken in order and certain steps may be taken more than once.

- **R0: evaluation**: Continuously evaluate the V2X ecosystem to see if any remediation step is needed (see 5.2 for further discussion).
- **R1: update**: Update the software/firmware/hardware of the misbehaving EE (see 5.3).
- **R2: suspension**: Suspend the misbehaving EE from participating in V2X activities. The suspension can be implemented using different techniques depending on the type of misbehaviour, availability of the technique, etc. (see 5.4).

5.2 R0: Evaluation

5.2.1 General

The MBMS needs to continuously monitor misbehaving EEs and evaluate their impact on the V2X ecosystem. Based on the result, the MBMS determines if any remediation step is needed.

There are misbehaviour scenarios where the cost of doing something might outweigh the damage caused by the misbehaviour itself. So, doing nothing in such scenarios may be the suitable response. Some example scenarios where this may be appropriate are:

- **Easily detectable**: The misbehaviour can be easily detected and filtered out by the EE locally. An EE’s ability to detect misbehaviour can vary widely depending on how well/ill-equipped the EE is, so the minimum performance requirements (MPR) with respect to misbehaviour detection may need to be considered.
- **Mild**: The impact of misbehaviour is very low, such that it does not reduce the benefits of the V2X system in any meaningful/noticeable way. The intent is not to do a full cost-benefit analysis, which may be resource-intensive and unfeasible, but to use expert judgment and prior experience.
- **Temporary**: The misbehaviour occurs for a brief period, much smaller than the time required for a full misbehaviour detection and remediation (MBDR) cycle (i.e., it can take a week, and the misbehaviour lasts less than a day).
- **Localised**: The misbehaviour is limited to a relatively small geographic region or EE population, where the size is in comparison to the geographic region/EE population affected by the remediation, (e.g., the misbehaviour affects a particular vehicle manufacturer, year or model but remediation efforts span

all models and years).

5.2.2 Other considerations

The MBMS needs to monitor misbehaving EEs even after remediation steps have been taken against them. This is especially important if:

- The cause and/or effect of misbehaviour was miscalculated/revised by the MBMS.
- The size of the EE population affected by the misbehaviour was miscalculated/ revised by the MBMS.
- Similar misbehaviour is observed in a different class of EEs in future.
- ...

5.3 R1: Update

5.3.1 General

This is a non-local, send-side remediation technique. It involves updates of some kind and depends on the reason of the misbehaviour. Here, it is important to distinguish between several categories from the EE's point of view:

- Accidentally malfunctioning EE: An EE has an error which leads to wrong data in V2X messages e.g., GPS sensor.
- Design error in HW or SW: The EE has a design fault which leads to the sending of wrong data under specific circumstances.
- Cybersecurity incident: An attacker hijacked the EE.
- Misuse: Somebody used the V2X function in an unintended way.

The update process for vehicles and roadside units (RSU) will be similar except that most RSUs will not be able to visit a repair shop. In this case maintenance must be carried out remotely or on-site by the operator.

1) Update categories for vehicles:
   a) Update HW and/or SW for a single car in the repair shop either due to yearly inspection or by request.
   b) Remote SW update of a vehicle to the most current SW version (OTA).
   c) Implement a fix for a SW bug and distribute it to the relevant vehicles (OTA).
   d) Implement a fix in HW or SW and deploy the new version in newly produced vehicles.
   e) Implement a fix in HW and exchange it in all relevant vehicles in the repair shop because authorities requested it.
2) Update categories for infrastructure / RSU:

a) Update HW and/or SW for a single RSU through maintenance by the operator.

b) Remote SW update of a RSU to the most current SW version (over-the-air, OTA).

c) Implement a fix for a SW bug and distribute it to the relevant RSU (OTA).

d) Apply a fix to the HW or SW and implement the new version in newly produced RSUs.

e) Implement a fix in the HW and exchange it in all relevant RSUs because authorities requested it.

5.3.2 Process and state of the art

Figure 3 shows a typical update process for vehicles. Below are some relevant points about hardware and software updates.

- For the backend misbehaviour management, it is important to know if the misbehaviour is associated with a single event affecting, say, one compromised device, or if the error relates to all devices of a series.
  - A single erroneous device could be easily repaired in a workshop.
  - A defect in the series means the manufacturer must find the root cause and develop an update.
- The costs of a hardware update are significantly higher than for a software update.

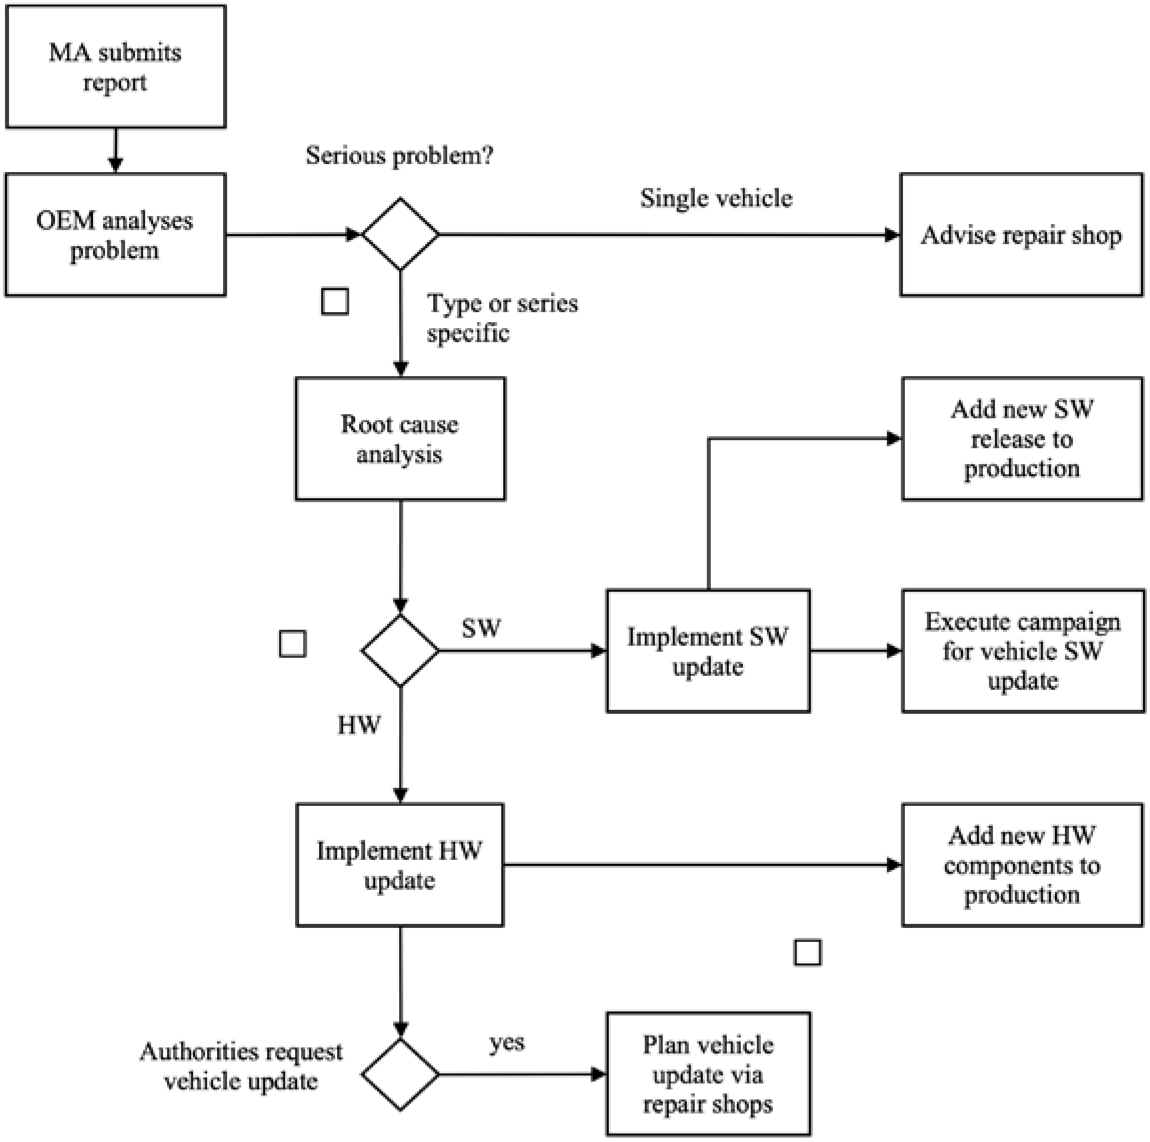
MA submits report
OEM analyses problem
Serious problem?
Single vehicle
Advise repair shop
Type or series specific
Root cause analysis
Add new SW release to production
SW
Implement SW update
Execute campaign for vehicle SW update
HW
Implement HW update
Add new HW components to production
Authorities request vehicle update
yes
Plan vehicle update via repair shops

Figure 3: Update process example for vehicles

5.3.3 Other considerations

Updates are costly for the manufacturers or operators. Therefore, misbehaviour needs to be classified to justify the costs of updates. Depending on the severity of the misbehaviour, the update strategies can be ranked accordingly:

1. No update necessary
2. Repair defective EE
3. Update SW in the production process
4. Update SW in the production process and for operated EEs
5. Update HW in the production process
6. Update HW in the production process and for operated EEs

5.4 R2: Suspension

5.4.1 General

In any misbehaviour scenario other than those covered above, the appropriate remediation would be to suspend the EE's participation in at least those V2X activities where the misbehaviour was determined to be happening. Some example scenarios where R2 may be appropriate are:

- Recurring: The same EE (or its security credentials used by someone) misbehaves repeatedly at multiple locations and/or different times.
- Safety: The misbehaviour compromises the safety of road users.
- Usability: The misbehaviour renders the V2X system unusable or significantly reduces the capacity for V2X participants.
- Malicious: The misbehaviour is determined to be due to malicious actions as opposed to an EE malfunction.

Below, four suspension techniques (STs) are elaborated. This is not an exhaustive list, and is intended to illustrate the range of options available to V2X deployers.

- ST0: Certificate issuance pause
- ST1: Partial certificate revocation
- ST2: Full certificate revocation
- ST3: Alternative mechanisms

The requirements for a suspended EE to be reinstated can vary depending on several factors, such as the type of misbehaviour and suspension, security and certificate policies for the region, and the design of the V2X public key infrastructure (PKI). A few categories of reinstatement requirement categories (RCs) are listed below.

- RC1: Suspension with automatic reinstatement after some time.
- RC2: Suspension with remote verification of correct operation before reinstatement, e.g., perform a software update that runs diagnostics and verifies that the system is good.
- RC3: Suspension with local non-invasive verification of correct operation before reinstatement, e.g., a technician checks in person that the system is good.
- RC4: Suspension with local invasive verification of correct operation, e.g., a technician replaces certain hardware components in the EE and verifies that the system is good.

NOTE: Further discussion on the topic is out of scope.

5.4.2 ST0: Certificate issuance pause

5.4.2.1 General

In this suspension technique, the EE is denied access to new certificates by temporarily pausing the certificate generation and/or issuance (blacklisting). At this stage ST0 looks identical to R0 from an EE's point of view. Only when the EE attempts to request/download a new set of certificates will the suspension become evident (i.e., certificate issuance is paused). An EE may want to request new certificates for various reasons:

- Previously issued certificates have expired or soon will.
- Certificate parameters have changed, and the EE needs the new ones.
- …

5.4.2.2 Usage

Misbehaviour scenarios where this technique will be useful are similar to those for R0, with the main distinguishing factor being the misbehaviour time window:

- If it is shorter or the same as the total certificate validity duration an EE is allowed per download, then ST0 is not any more effective than R0, hence there is no point in using ST0.
- Otherwise, ST0 should be used.

5.4.2.3 Reinstatement

When the EE is deemed fit to resume its V2X activities, it can be reinstated by resuming the certificate generation and issuance.

5.4.3 ST1: Partial certificate revocation

5.4.3.1 General

One of the ways to implement a partial certificate revocation is do it individually. Individual certificates can be revoked by listing their hashes on the certificate revocation list (CRL). See Clauses 7 and 7.3.5 in IEEE Std 1609.2-2022 for more details on CRLs and specifically hash-based revocation, respectively.

Other ways of implementing partial certificate revocation may be introduced/standardised in future, e.g., see 'Privacy-Preserving Method for Temporarily Linking/Revoking Pseudonym Certificates in VANETs', by Marcos Antonio Simplicio Junior et al. (https://eprint.iacr.org/2018/185).

5.4.3.2 Usage

Misbehaviour scenarios where this technique will be useful are those where the damage caused by the misbehaviour might outweigh the overall cost incurred by the V2X system in implementing it, including:

- Incremental cost of generating and distributing the added CRL entries by the MBMS to all the EEs.

- Incremental cost of verifying the added CRL entries for all V2X recipients who will receive that CRL.
- Cost of reinstating the misbehaving EE.

5.4.3.3 Reinstatement

As certificates are revoked partially (e.g., individually) in this technique, certificates that are not on a CRL and are otherwise valid can be used by the EE. So, to reinstate a suspended EE that is deemed fit to resume its V2X activities, the EE needs to be issued new certificates.

5.4.4 ST2: Full certificate revocation

5.4.4.1 General

In this suspension technique, all the certificates in the possession of an EE from a fixed time onwards are revoked. One of the ways of implementing a full certificate revocation is by listing the appropriate 'linkage seeds' on the CRL. See Clause 5.1.3 in IEEE Std 1609.2-2022 for more details on linkage seeds and chains, and consult Clauses 7 and 7.3.7 in the same standard for greater detail on CRLs and specifically linkage-based revocation, respectively.

5.4.4.2 Usage

Scenarios where CRLs will be useful are where the damage caused by the misbehaviour might outweigh the overall cost incurred by the V2X system in implementing the technique. Costs for ST2 are very similar to that of ST1, except for reinstatement (see 5.4.4.3 for more details).

5.4.4.3 Reinstatement

As certificates are revoked via the linkage chain in this technique, to reinstate a suspended EE that is deemed fit to resume its V2X activities, a new linkage chain needs to be used. This is significantly more complicated and resource-intensive than reinstating under ST1.

5.4.5 ST3: Alternative mechanisms

5.4.5.1 General

A recognised alternative to certificate revocation is to use so-called activation codes. In this suspension technique, the misbehaving EE can be prevented or blocked from unlocking certificates during a defined time period (see Clause 9.4 of IEEE Std 1609.2.1-2022 for more details on 'activation codes' and 'unlocking values'). This is different from both ST1 and ST2 in that ST3 only affects the misbehaving EE, whereas in both ST1 and ST2 every receiving EE is impacted by having to download and process larger CRLs. However ST3 comes with its own overhead in the form of periodic broadcasts/downloads of activation codes. For more details on the performance and benefits analysis of ST3, see IEEE Std 1609.2.1-2022 and relevant references therein.

5.4.5.2 Usage

Again, this technique is deemed useful where the damage caused by the misbehaviour might outweigh the overall cost incurred by the V2X system in implementing it. Costs for ST3 vary significantly from that of ST1 or ST2, and include:

- Base costs for all EEs, i.e., there is a non-zero cost even when there are no suspended EEs in the system (NOTE: There is also a base cost for ST1 and ST2 that depends on the CRL update frequency).
- Incremental costs of suspending an EE depend on the communication model used for providing activation codes to EEs:
  - Broadcast: If activation codes are provided via a broadcast medium, then the incremental costs depend indirectly on the number of suspended EEs.
  - Two-way communication: If activation codes are provided via a two-way communication medium, then there is no incremental cost of suspension.
- Cost of reinstating the misbehaving EE.

5.4.5.3 Reinstatement

As EEs have all their certificates on the device – albeit in a locked form – in order to reinstate a suspended EE that is deemed fit to resume its V2X activities, the EE needs to be given its ‘unlocking value’. This is unlikely to incur any significant cost.

5.4.5.4 Other considerations

One of the main benefits of ST3 is that it scales well with the number of suspensions. While in the case of ST1/ST2 EEs may be overwhelmed by large-scale suspensions (a CRL with millions of entries could slow down even the most capable EEs), because EEs have two-way communication for downloading activation codes, ST3 is less likely to affect an EE’s performance, even when every other EE is suspended.

6 Misbehaviour classification

6.1 General

We first explore different ways to classify misbehaviour and then down select from those that are relevant for misbehaviour remediation. Here is a (non-exhaustive) list of bases for misbehaviour classification:

- Evidence
- Impact
- Local detection
- Footprint and duration

6.2 Evidence

This classification is useful in the context of misbehaviour reporting, as evidence needs to be included in the misbehaviour report for the Misbehaviour Authority to be able to determine the cause/source of misbehaviour. ETSI TS 103 759 divides misbehaviour into five classes:

- Class 1: Implausible values within the incoming message.
- Class 2: Inconsistencies between the incoming message and previous messages of the same type emitted from the same ITS object.
- Class 3: Inconsistencies between the incoming message and information about the local environment from the 'ego vehicle'.
- Class 4: Inconsistencies between the incoming message and the onboard sensors' perception.
- Class 5: Inconsistencies between the incoming message and previous messages of other types from the ITS object or messages (of the same type or not) emitted by other ITS objects.

6.3 Impact

Individual instances of misbehaviour can be classified according to their impact on the V2X ecosystem. Four levels in increasing order of severity are:

Very low	Misbehaviour can easily be filtered out by the V2X application on the end entity resulting in a very low to no impact on the system.
Low	Misbehaviour causes (or, has the potential to cause) non-safety related events, e.g., vehicle slows down, misses green light phase, shows false warning to the driver, etc.
Medium	Misbehaviour causes (or, has the potential to cause) safety related events but does not result in any physical damages and/or injuries to living beings, e.g., unnecessary emergency brake, collision avoidance measure, etc.
High	Misbehaviour causes (or, has the potential to cause) safety related events and results in physical damages and/or injuries to living beings.

NOTE: This document does not define actual impacts, it just provides guidance on how they should be defined. This document also does not stipulate who oversees impact definition: it could be policymakers, MAs or some other entity. The impacts can also change over time, so the policy around impact definition needs to take that evolving nature of impacts into account.

6.4 Local detection

Misbehaviour can be classified based on how easy or difficult is it to locally detect by the end entity. Four levels in the increasing order of severity are:

Very easy	Misbehaviour can be detected locally by the end entity without the need for advanced/dedicated detection software or hardware, i.e., it can be detected even by the least equipped end entities.
Easy	Misbehaviour can be detected locally by the end entity with the use of advanced detection software but otherwise does not require sophisticated hardware like cameras and sensors.
Moderate	Misbehaviour can be detected locally by the end entity only with the use of advanced/dedicated detection software and hardware, i.e., it can be detected only by highly equipped end entities.
Difficult	Misbehaviour cannot be detected locally by a single end entity, i.e., either multiple end entities need to collaborate among themselves, or an end entity needs to collaborate with the backend to detect such a misbehaviour.

6.5 Footprint and duration

Misbehaviour can be classified according to how widespread and persistent it is. Four levels in increasing order of severity are:

Local and temporary	Misbehaviour is confined to a small region (e.g., an intersection or short stretch of road) or a small V2X population (e.g., a few hundreds or thousands as opposed to millions), and lasts a short period of time (e.g., up to a few days).
Local and persistent	Misbehaviour is confined to a small region/population and continues for a long period of time (e.g., a few weeks or more).
Global and temporary	Misbehaviour is spread over a large region (e.g., cities that are far apart, different states, or even different countries) or a large V2X population (e.g., a few million), and lasts a short period of time.
Global and persistent	Misbehaviour is spread over a large region/population and continues for a long period of time.

6.6 Overall severity ratings

Since misbehaviour classification based on evidence is mainly useful for misbehaviour reporting, and seems orthogonal to the severity of a misbehaviour, for the overall severity ratings only the other three bases are considered: impact, local detection, footprint, and duration.

NOTE: The misbehaviour classifications presented in this document, including the overall ratings below, are intended to be used as an example/suggestion as opposed to a technical specification.

For rows 1 through 16 since the impact is very low, the base overall rating is very low with the following exceptions:

- Rows 4 and 8 are rated low (i.e., a level higher than the base rating) because the misbehaviours are global and persistent and such misbehaviours have the potential to create unwanted noise in the system.
- Rows 12 and 16 are rated medium (i.e., two levels higher than the base rating) due to the higher difficulty (moderate or difficult) in detection as well as the misbehaviours being global and persistent.

Row	Impact	Detection	Footprint and duration	Overall
1	Very low	Very easy	Local and temporary	Very low
2	Very low	Very easy	Local and persistent	Very low
3	Very low	Very easy	Global and temporary	Very low
4	Very low	Very easy	Global and persistent	Low
5	Very low	Easy	Local and temporary	Very low
6	Very low	Easy	Local and persistent	Very low
7	Very low	Easy	Global and temporary	Very low

8	Very low	Easy	Global and persistent	Low
9	Very low	Moderate	Local and temporary	Very low
10	Very low	Moderate	Local and persistent	Very low
11	Very low	Moderate	Global and temporary	Very low
12	Very low	Moderate	Global and persistent	Medium
13	Very low	Difficult	Local and temporary	Very low
14	Very low	Difficult	Local and persistent	Very low
15	Very low	Difficult	Global and temporary	Very low
16	Very low	Difficult	Global and persistent	Medium

For rows 17 through 32 since the impact is low, the base overall rating is low with the following exceptions:

- Rows 17-19 are rated very low (i.e., a level lower than the base rating) because of very easy detection paired with misbehaviours being either local or temporary (or, both local and temporary).
- Similarly, row 21 is rated very low (i.e., a level lower than the base rating) because the misbehaviour is local and temporary.
- Rows 28 and 32 are rated medium (i.e., a level higher than the base rating) because the misbehaviours are global and persistent.

Row	Impact	Detection	Footprint and duration	Overall
17	Low	Very easy	Local and temporary	Very low
18	Low	Very easy	Local and persistent	Very low
19	Low	Very easy	Global and temporary	Very low
20	Low	Very easy	Global and persistent	Low
21	Low	Easy	Local and temporary	Very low
22	Low	Easy	Local and persistent	Low
23	Low	Easy	Global and temporary	Low
24	Low	Easy	Global and persistent	Low
25	Low	Moderate	Local and temporary	Low
26	Low	Moderate	Local and persistent	Low
27	Low	Moderate	Global and temporary	Low
28	Low	Moderate	Global and persistent	Medium
29	Low	Difficult	Local and temporary	Low
30	Low	Difficult	Local and persistent	Low
31	Low	Difficult	Global and temporary	Low
32	Low	Difficult	Global and persistent	Medium

For rows 33 through 48 since the impact is medium, the base overall rating is medium with the following exceptions:

- Rows 33 – 35 and 37 – 39 are rated low (i.e., a level lower than the base rating) because of the ease (very easy or easy) of detection and the misbehaviours being either local or temporary (or, both local and temporary).
- Row 41 is also rated low (i.e., a level lower than the base rating) because the misbehaviour is local and temporary.
- Row 48 is rated high (i.e., a level higher than the base rating) because the misbehaviour has the highest ratings in the other two columns: detection, footprint and duration.

Row	Impact	Detection	Footprint and duration	Overall
33	Medium	Very easy	Local and temporary	Low
34	Medium	Very easy	Local and persistent	Low
35	Medium	Very easy	Global and temporary	Low
36	Medium	Very easy	Global and persistent	Medium
37	Medium	Easy	Local and temporary	Low
38	Medium	Easy	Local and persistent	Low
39	Medium	Easy	Global and temporary	Low
40	Medium	Easy	Global and persistent	Medium
41	Medium	Moderate	Local and temporary	Low
42	Medium	Moderate	Local and persistent	Medium
43	Medium	Moderate	Global and temporary	Medium
44	Medium	Moderate	Global and persistent	Medium
45	Medium	Difficult	Local and temporary	Medium
46	Medium	Difficult	Local and persistent	Medium
47	Medium	Difficult	Global and temporary	Medium
48	Medium	Difficult	Global and persistent	High

For rows 49 through 64 since the impact is high, the base overall rating is high with the following exceptions:

- Rows 49-56 are rated medium (i.e., a level lower than the base rating) because of the ease (very easy or easy) of detection.
- Rows 57, 58, 61, 62 are also rated medium (i.e., a level lower than the base rating) because all these misbehaviours have a local footprint.

Row	Impact	Detection	Footprint and duration	Overall
49	High	Very easy	Local and temporary	Medium
50	High	Very easy	Local and persistent	Medium
51	High	Very easy	Global and temporary	Medium
52	High	Very easy	Global and persistent	Medium
53	High	Easy	Local and temporary	Medium
54	High	Easy	Local and persistent	Medium
55	High	Easy	Global and temporary	Medium
56	High	Easy	Global and persistent	Medium
57	High	Moderate	Local and temporary	Medium
58	High	Moderate	Local and persistent	Medium
59	High	Moderate	Global and temporary	High
60	High	Moderate	Global and persistent	High
61	High	Difficult	Local and temporary	Medium
62	High	Difficult	Local and persistent	Medium
63	High	Difficult	Global and temporary	High

64	High	Difficult	Global and persistent	High

7 Mapping remediation to misbehaviour

Section 5 listed the different remediation techniques currently available to manage misbehaviour in a V2X system, and Section 6 suggested a way to assign an overall severity rating (very low, low, medium, high) to a given misbehaviour. This document suggests the following approaches for the next task of mapping remediation to misbehaviour:

- A1: Natural progression – triggers for progression could be time based, behaviour based, or something else

  1. Do nothing at first (R0).

  2. If the misbehaviour by the same device continues, update software/firmware (R1) if available, otherwise go to the next step.

  3. If step 2 does not help, suspend the device (R2) – starting with ST0, then ST1, ST2, and so on.

- A2: Severity based

  1. For each misbehaviour severity level determine the appropriate remediation techniques.

  2. Given a misbehaviour, determine its severity level.

  3. Apply the appropriate remediation corresponding to the misbehaviour severity level.

- Combination of A1 and A2

  1. For each misbehaviour severity level determine a base remediation technique, e.g., if the misbehaviour severity is very low the base remediation can be R0, but if the misbehaviour severity is high the base remediation can be ST1 or ST2 within R2.

  2. Given a misbehaviour, determine its severity level.

  3. Apply the base remediation technique corresponding to the misbehaviour severity level.

  4. If the misbehaviour by the same device continues, go to the next level of remediation, e.g., if the base remediation was R1, apply R2 starting with ST0, then ST1, ST2, and so on.

8 Conclusion and next steps

This white paper provides technical guidance on the different options for misbehaviour remediation, and the conditions (misbehaviour classes) under which each of the remediation options may be applied. With that said, this is just a first step in the direction of specifying policies and procedures for V2X misbehaviour remediation. Some immediate next steps and directions for future work are as follows.

- Draft specification: The stakeholders of the V2X ecosystem should take this white paper as a starting point, and explore/finetune the processes of misbehaviour classification (Section 6) and mapping remediation to misbehaviour (Section 7). This should naturally lead to a first draft of policies and procedures for V2X misbehaviour remediation.

- Reinstatement: This white paper only briefly touches on the topic of reinstatement (Section 5.4). What is needed is a specification of policies and procedures for the reinstatement of a suspended V2X device. A reasonable starting point is: to be reinstated, a suspended V2X device would have to demonstrate compliance to requirements similar to those for device bootstrap/initialisation.

- Suspension without misbehaviour: This white paper considers remediation only when misbehaviour of some sort has occurred or been reported to the backend misbehaviour management system. There are scenarios where a device suspension may be warranted even if there was no misbehaviour reporting, e.g., a malicious actor that has extracted security credentials of a V2X device and posted it on a website. Policies and procedures for V2X misbehaviour remediation should also consider such scenarios.

- Technical specifications: As the above policies and procedures are developed, it is possible that newer techniques (or, refinements to older techniques) will be needed. This white paper identifies two such topics:

  o Additional information for misbehaviour report: Certain information like impact level, difficulty of local detection, footprint and duration of a misbehaviour may need to be included in a misbehaviour report to help the backend misbehaviour management better estimate the severity of a misbehaviour.

  o New revocation mechanisms: As pointed out in Sections 5.4.3 and 5.4.5, for certain types of misbehaviour it may be necessary to standardise new ways of revoking certificates.

5GAA is a multi-industry association to develop, test and promote communications solutions, initiate their standardisation and accelerate their commercial availability and global market penetration to address societal need. For more information such as a complete mission statement and a list of members please see https:/5gaa.org

5GAA
Automotive Association