Safety Meets Security: Using IEC 62443 for a Highly Automated Road Vehicle

Dominik Püllen(✉), Nikolaos Anagnostopoulos, Tolga Arul, and Stefan Katzenbeisser

University of Passau, Passau, Germany
dominik.puellen@uni-passau.de

Abstract. In this work, we conduct and discuss a consensus-based risk analysis for a novel architecture of a driverless and electric prototype vehicle. While well-established safety standards like ISO 26262 provide frameworks to systematically assess risks of hazardous operational situations, the automotive security field has emerged only in the last years. Today, SAE J3061 provides recommendations and high-level guiding principles of how to incorporate security into vehicle systems. ISO/SAE 21434 is a novel automotive security standard, which, however, is still under development. Therefore, we treat the aforementioned architecture as a single Industrial Automation and Control System (IACS) and provide an implementation of the IEC 62443 series. We collaboratively identify threats in a three-round process and define a scoring scheme for automotive risks. As a result, we obtain a tailored bundle of compensating security mechanisms. Based on our work, we suggest improvements for future automotive security standards when it comes to the co-engineering of safety and security.

Keywords: IEC 62443 · Security · Safety · Risk mitigation

1 Introduction

The increasing connectivity and the growing computational power of road vehicles come along with great potential, but likewise lead to security concerns as demonstrated by prior works [6,15,25]. Beside new security challenges, the field of safety is also affected by vehicle automation, because a human being cannot be assumed anymore as a fallback layer. As modern road vehicles are typically complex cyber-physical systems and need to meet legal requirements, a standardized process for risk identification and mitigation is typically applied. Currently, the ISO/SAE 21434 [5] is the most promising candidate for an automotive security standard. It provides risk assessment methods for Intelligent Transportation Systems (ITS). After identifying and decomposing threat scenarios into attack paths, Cybersecurity Assurance Levels (CAL) indicate the estimated security requirements for given items. Moreover, SAE J3061 [3], published in 2016, provides general guidelines for the development of secure automotive components.

A. Casimiro et al. (Eds.): SAFECOMP 2020, LNCS 12234, pp. 325–340, 2020.
https://doi.org/10.1007/978-3-030-54549-9_22

It is inspired by the ISO 26262 [1] safety standard and reuses techniques from existing security models such as EVITA [10] and HEAVENS [2]. Schmittner et al. [20] demonstrate the security analysis of an automotive communication gateway by applying the concept phase of SAE J3061. They derive high-level security requirements, using the Confidentiality, Integrity, and Availability (CIA) triad. While their work focuses on a single component, our objective is to analyze an automated vehicle as a whole. For that purpose, we apply the IEC 62443 standards [4] to a recently announced novel vehicular architecture [24]. We argue, that IEC 62443 overlaps with the main idea of the unpublished ISO/SAE 21434. That is, it provides a risk-based security analysis process, takes into account interfaces to external components (e.g., V2X), identifies and assesses threats, and eventually uses Security Levels (SL) to describe security requirements. When it comes to threats, Petit and Shladover [17] identified 12 sources of potential attacks on automated vehicles and evaluated each one regarding its feasibility, occurrence probability, consequences and mitigation techniques. Beside the goal of a systematic security requirement analysis, our research question concerns the possibility of co-engineering security and safety demands in a vehicular system. In the following, we aim at sharing our lessons learned and suggest improvements for future automotive security standards.

2 Overview of a Novel Vehicular Architecture

In 2018, seven German universities and industry partners announced the development of four fully automated and driverless vehicles [24]. These vehicles are supposed to serve as an evaluation platform for new concepts in various fields, such as automation, modularization, verification, validation, safety, and security. Unlike contemporary vehicles, that typically consist of dozens of Electronic Control Units (ECUs), the novel E/E architecture follows a centralized approach, which is inspired by the human nervous system. That is, four *sensor modules* collect and preprocess radar, Lidar, and camera data. They hand them over to the *cerebrum*, which is responsible for the trajectory and for behavioral planning based on the sensor data. The *brainstem*, in turn, implements and tracks the trajectory and instructs the *spinal cord* to eventually move the vehicles. The latter provides all necessary steering angles and both braking and acceleration torques to four *dynamic modules*, which drive the wheels. In case of failures, the brainstem reflexively enforces an emergency trajectory, by which a safe halt is usually triggered. All aforementioned modules are connected over BroadR-Reach in a ring topology, allowing communication even if a switch breakes down. The dynamic modules are additionally wired over FlexRay, which serves as a supplementary fallback layer. In total, 26 ultrasonic and 2 radar sensors, denoted as *platform sensors*, allow for near-field sensing and are directly connected to the brainstem over CAN. For in-vehicle communication, the Automotive Service-Oriented software Architecture (ASOA) [11] is deployed, a new modular framework, that enables flexible communication, fast and secure updates of ECUs, and easy replacement of hardware components.

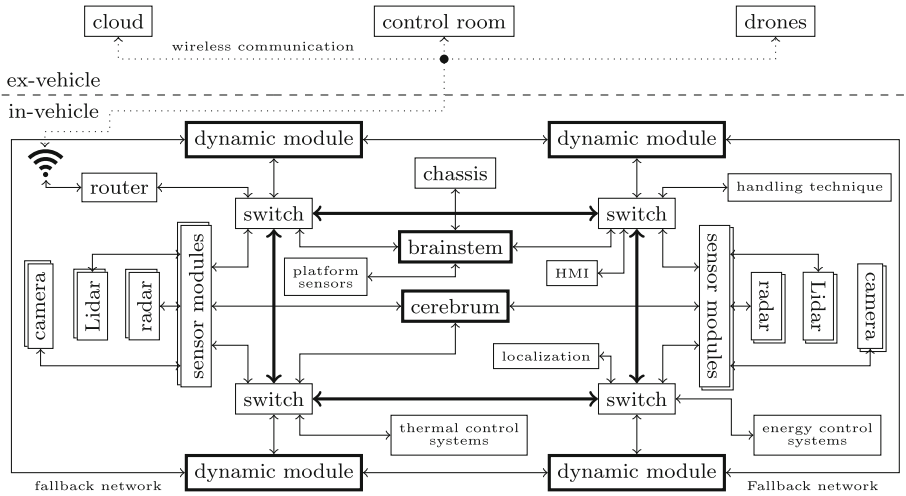

Fig. 1. Overview of the proposed vehicular architecture [12]

Beside the prototype vehicles, a new infrastructural concept provides environmental information such as traffic updates via V2I communication. A *cloud* serves as a collective memory, such that vehicles can incorporate predictive driving behavior by learning from each other. *Drones* collect and share additional traffic information which helps the vehicles to create a realistic environmental model. This information is fed into trajectory planning algorithms on the cerebrum. A *control room* enables the remote control by a human in case automatic maneuvering is not possible anymore. It is only called into action in exceptional situations, in order to meet European legal requirements.

In the following, the term *vehicular architecture* refers to the novel E/E architecture in combination with its external components as illustrated in Fig. 1.

3 Introduction to IEC 62443

The IEC 62443[4] is a series of standards and technical reports that provide a structured risk assessment and mitigation process for Industrial and Automation Control Systems (IACS), alongside with management guidances, policies, and terminology. An IACS typically describes a complex system consisting of various computing units, sensors, actuators, temporarily connected devices, and a human interface, that all collaboratively work on the outcome of a specific product. The overall objective is to identify threats, assess resulting risks, and come up with protection techniques.

As shown in Fig. 2, the actual risk assessment process is described in IEC 62443-3-2 in consecutive steps, denoted as Zone and Conduit Requirements (ZCR). In the first step, all relevant assets of the System Under Consideration

(SUC) are identified (ZCR 1). A high-level security analysis (ZCR 2) gives indication about the worst-case unmitigated risk on each asset and whether further investigation is necessary. Based on the results of this high-level analysis, the SUC is partitioned into zones and conduits (ZCR 3), whereas a zone contains assets with the same or similar security requirements. A conduit is a special zone type, that connects two other zones and therefore, usually describes a network. In ZCR 4, the tolerable risk $r^{\mathrm{tol,max}}$ of each zone is compared with the unmitigated risk r^{u}. If $r^{\mathrm{tol,max}}$ is below r^{u}, no further action is required. Otherwise, a detailed security risk assessment follows in ZCR 5.

The main objective of ZCR 5 is to iteratively reduce the unmitigated security risk of identified *threats* ($\mathcal{T}$) by applying compensating countermeasures. Threats are associated with seven Foundational Requirements ($\mathcal{FR}$). That is, *Identification and Authentication Control* (IAC), *Use Control* (UC), *System Integrity* (SI), *Data Confidentiality* (DC), *Restricted Data Flow* (RDF), *Timely Response to Events* (TRE), and *Resource Availability* (RA). Since the security of a system refers to the mitigation of threats, an exhaustive list of threats and exploitable vulnerabilities is crucial (ZCR 5.1–5.2). Both the impact and the likelihood of each threat (ZCR 5.3–5.4) is determined, in order to compute the unmitigated security risk r^{u} of each threat (ZCR 5.5). Based on these results, a target security level SL-T for each zone is computed. IEC 62443 differentiates between four levels, SL-1, SL-2, SL-3, and SL-4. While SL-0 is implicitly defined as no requirements, SL-1 demands for protection against coincidental violations. SL-2 - SL-4 cover intentional violation with an increasing level of skills, resources, and motivation. Both impact and likelihood are reevaluated (ZCR 5.9) after applying changes to the SUC, e.g., after introducing a new countermeasure. Ideally, this leads to a reduction of the residual risk (ZCR 5.10). A reassessment of the

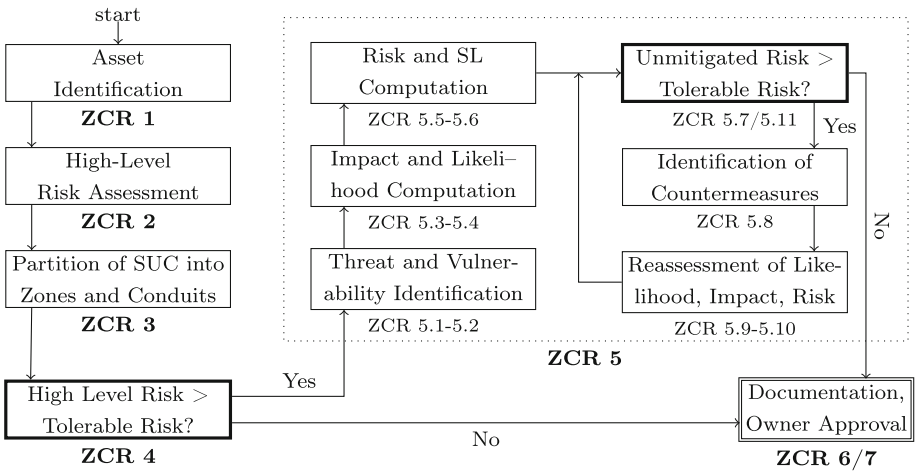

Fig. 2. Simplified workflow of IEC 62443-3-2

high level risk, however, is not caused. Once the unmitigated risk of all threats is below the tolerable risk, the SUC is considered secure.

4 Application of IEC 62443

We consider the presented vehicular architecture as our System Under Consideration (SUC). Since it shares most IACS properties like sensors, actuators, and computing units, we use the IEC 62443 standards for a full security risk analysis. We demonstrate how to implement the generic guidelines ZCR 1–5 of IEC 62443-3-2 with the ultimate goal to create a tailored bundle of security means for a secure and safe operation of the automated vehicles. Our findings may inspire future automotive security standards. As discussed in Sect. 5, all assessments are the results of an expert committee.

4.1 High-Level Risk Analysis (ZCR 1 - ZCR 4)

In **ZCR 1**, our expert team identified a total number of 19 assets in the SUC, i.e., functional components with a potential impact on security and safety. These include both in-vehicle assets (e.g., brainstem, radar) and external ones (e.g., drones, control room).

In **ZCR 2**, a high-level security risk analysis was performed for each asset a_i. For this, IEC 62443-3-2 requires to assess the high-level likelihood $L_{a_i}^{\mathrm{HL}}$ and the high-level impact $I_{a_i}^{\mathrm{HL}}$ of a potential attack on a_i. Since it does not state how this is supposed to happen, we apply a multi-criteria decision making process. More precisely, we implement a Simple Additive Weighting (SAW) approach [19], where predefined evaluation criteria are scored and then ranked according to their importance. As demonstrated in the subsequent paragraphs, evaluation criteria for both likelihood and impact are represented as vectors $\boldsymbol{L}_{a_i}^{\mathrm{HL}} = (L_1\ L_2\ ...\ L_n)^\intercal$ and $\boldsymbol{I}_{a_i}^{\mathrm{HL}} = (I_1\ I_2\ ...\ I_m)^\intercal$, respectively. The ranking of each criterion is done with the normalized weight matrices $\boldsymbol{W}_{\mathrm{L}}$ and $\boldsymbol{W}_{\mathrm{I}}$, respectively. We compute the scores $L_{a_i}^{\mathrm{HL}}$ and $I_{a_i}^{\mathrm{HL}}$ by summing up the products of each score and its weight, i.e., $L_{a_i}^{\mathrm{HL}} = \boldsymbol{L}_{a_i}^{\mathrm{HL}} \cdot \boldsymbol{W}_{\mathrm{L}}$, respectively $I_{a_i}^{\mathrm{HL}} = \boldsymbol{I}_{a_i}^{\mathrm{HL}} \cdot \boldsymbol{W}_{\mathrm{I}}$, where $\cdot$ is the dot product. Due to normalization, a score of $L_{a_i}^{\mathrm{HL}}$ =1 indicates the highest possible likelihood, while $I_{a_i}^{\mathrm{HL}}$ =1 stands for the worst-case impact. The high-level risk $r_{a_i}^{\mathrm{HL}} = (I_{a_i}^{\mathrm{HL}}, L_{a_i}^{\mathrm{HL}},)$ is mapped to a risk class, using the weighted normalized decision matrix in Table 3. We argue that such a scoring scheme is compliant with current automotive guidelines, as SAE J3061 recommends additive scoring for the assessment of impact factors.

Impact: We describe $\boldsymbol{L}_{a_i}^{\mathrm{HL}} = (\mathrm{PS}\,\mathrm{FL}\,\mathrm{OL})^\intercal$ as a vector of three impact criteria, where PS represents *Passenger Safety*, FL *Financial Loss*, and OL *Operational Limitations* [10]. Each criterion is independently scored by the experts, who use a set of exclusive parameters ($\mathcal{P}$) for this task. Based on its severity, each parameter is mapped on a distinct integer value according to the rules of SAW.

That is, the least severe parameter is associated with 1, which is then incremented by 1 for subsequent parameters. For instance, we differentiate between $\mathcal{P}_{\text{PS}} = \{\textit{fatal}, \textit{seriously injured}, \textit{slightly injured}, \textit{no injuries}\}$ for passenger safety. The parameter $p_0 = \textit{no injuries}$ is associated with 1, $p_1 = \textit{slightly injured}$ with 2, $p_2 = \textit{seriously injured}$ with 3, and $p_3 = \textit{fatal}$ with 4. Following SAW, we normalize these integer scores with $\widetilde{\text{p}_i} = \frac{\min \text{p}_j \forall p_j \in \mathcal{P}}{\text{p}_i}$ and then rank them with predefined weights. In our case, we use probabilistic weights, yielding 0.3 for both operational limitations and financial loss. As we consider passenger safety the most valuable criterion for an automotive system, we prioritize it with a weight of 0.4. We define $\mathcal{P}_{\text{OL}} = \{\textit{massive}, \textit{high}, \textit{medium}, \textit{low}, \textit{none}\}$ to assess operational limitations of a potential attack. A *massive* limitation occurs if all traffic comes to an halt. This, for instance, may happen if the control room is taken over by an adversary. *High* limitations lead to traffic jams in a designated area, e.g., when sending fake traffic information. *Medium* constraints occur in case a vehicle can only operate at reduced speed, e.g., when hijacking or deceiving sensors. Finally, *low* limitations are the result of hijacking non-critical assets such as the chassis. Regarding the financial loss, we distinguish between four monetary classes as shown in Table 1. The so-called *Value of a Statistical Life* (VSL) [22] served as a reference value to determine these classes. The VSL indicates the mortality risk reduction benefit for the U.S. government. In 2016, the U.S. Department of Transportation indicated a VSL of $9.6 million. Since the VSL is not a universal constant, we assume VSL = $10M for simplicity. Table 1 shows the normalized and weighted scores for each impact parameter.

Table 1. Impact criteria with their normalized and weighted scores

Passenger Safety (PS)	Fatal **0.4**	Seriously injured **0.2**	Slightly injured **0.134**	None **0.1**	
Operational Limitation (OL)	massive **0.3**	high **0.15**	medium **0.1**	low **0.085**	none **0.075**
Financial Loss (FL)	]$10M, ∞] **0.3**	]$10K, $10M[**0.15**	]$0, $10K] **0.1**	$0 **0.075**	

During our impact assessments, we encountered the problem of *transitive attack relations*. Theoretically, every asset a_i may be accountable for a worst-case attack if an adversary manipulates a safety-critical asset a_j through a_i, $i \neq j$. As a consequence, all assets would receive the highest impact score, which eventually could result in over-engineering. Therefore, for the assessment of a_i, we focus solely on its functional description, without considering propagating side effects. This, however, does not mean that transitive attacks are left out from the risk analysis, since they are covered by conduits in later steps.

Likelihood: We describe the high-level likelihood $\boldsymbol{L}_{a_i}^{\text{HL}} = (\text{IC WC P EI B TP})^\intercal$ as a six dimensional Boolean vector. That is, we decide for each asset a_i, whether an *Internet Connection* (IC) can be established, a *Wireless Communication*

Table 2. High-level assessments (ZCR 2) and SUC partitioning (ZCR 3)

Asset	Weighted Impact			Weighted Likelihood						HL Risk Class	Zone
	PC	OL	FL	IC	WC	P	EI	B	TP		
Control room	.4	.3	.3	.286	.238	.19	.143	0	0	ex.high	Z_F
Brainstem	.4	.15	.4	.286	0	.19	0	.095	0	ex.high	Z_A
Dynamic module	.4	.15	.15	.286	0	.19	0	.095	0	ex.high	Z_E
Radar	.2	.1	.15	0	0	0	0	.095	.048	Medium	Z_B
Sensor module	.134	.1	.15	.286	0	.19	.143	.095	0	High	Z_A
Chassis	.1	.075	.15	0	0	0	0	.095	0	Low	Z_C
⋮		⋮				⋮				⋮	⋮
HMI	.1	.1	.1	.286	.238	.19	0	.095	.048	High	Z_D

(WC) is possible, a_i is *re-Programmable* (P), a_i has *External Interfaces* (EI) such as ODB2, USB, a_i is directly connected to the in-vehicle *Bus* (B), and whether a_i is produced by a *Third Party* (TP). For instance, $\boldsymbol{L}_{a_i}^{\mathrm{HL}} = (1\ 0\ 1\ 0\ 1\ 0)^\intercal$ describes a re-programmable asset, which is connected to the in-vehicle bus and has the ability to establish an Internet connection. The order of the above criteria implicitly shows their ranking, i.e., to what extent they facilitate an attack. Similar to the high-level impact, we assign each criterion a distinct integer. As an Internet connection enables a potential attack the most, it receives the largest value of 6. For each subsequent criterion, we subtract 1 from the value, such that *Third Party* is eventually associated with 1. After normalization, we obtain $\boldsymbol{W}_{\mathrm{L}} = (0.286\ 0.238\ 0.19\ 0.143\ 0.095\ 0.048)^\intercal$. Table 2 gives an overview of the weighted evaluation criteria for both impact and likelihood of a selected number of assets.

In **ZCR 3**, we partition the SUC into zones and conduits, using the results of the high-level security analysis. We obtain nine zones and conduits $\mathcal{ZC} = \{Z_A, Z_B, ..., Z_F, C_A, C_B, C_C\}$. Instead of putting all assets with the same high-level risk into one zone, we additionally differentiate between safety-critical and remote assets. For instance, Z_A consists of highly safety-critical in-vehicle assets (brainstem, cerebrum, sensors, router), while the (remote and safety-critical) control room resides in Z_F. In that way, we are able to better address specific safety and security demands. Although the dynamic modules are highly safety-critical,

Table 3. Weighted normalized risk matrix with acceptance ranges

Grey cells: tolerable risk for ZCR 4			**Impact** ($I_{a_i}^{\mathbf{HL}}$)				
			negligible [.235,.26]	**minor** [.27,.384[	**major** [.384,.584[	**critical** [.584,.7[	**catastrophic** [.7,1]
Likelihood ($L_{a_i}^{\mathrm{HL}}$)	**trivial**	[0,.048]	ex. low	low	medium	high	high
	unlikely	[.048,.191]	ex. low	low	medium	high	high
	possible	[.191,.428]	low	medium	high	high	ex. high
	likely	].428,.714]	medium	medium	high	ex. high	ex. high
	certain	].714,1]	medium	high	ex. high	ex. high	ex. high

they are put into the dedicated zone Z_E, because they are additionally connected between each other over a fallback bus (cf., Fig. 1) and thus, have a significantly different attack surface. The fallback bus is treated in a dedicated conduit (C_C) as well. C_A describes the in-vehicle Ethernet-based communication and C_B is the wireless network that connects external components such as the cloud and the control room. Table 2 shows the zones to which an asset belongs.

In **ZCR 4**, a detailed security analysis follows for a zone $Z_i \in \mathcal{ZC}$, if there is an asset a_i with a high-level risk $r_{a_i}^{\mathrm{HL}} > r_{Z_i}^{\mathrm{tol,max}}$, where $r^{\mathrm{tol,max}}$ denotes the maximum tolerable risk. Since IEC 62443-3-2 does not prescribe how to determine $r^{\mathrm{tol,max}}$, we define both a maximum tolerable impact $I^{\mathrm{tol,max}}$ and likelihood $L^{\mathrm{tol,max}}$. We exclude passenger harm and financial loss, but find low operational limitations tolerable, resulting in $I^{\mathrm{tol,max}} = \boldsymbol{I}_{a_i}^{\mathrm{HL}} \cdot \boldsymbol{W}_{\mathrm{I}} = 0.1 + 0.085 + 0.075 = 0.26$. Similarly, we define the tolerable likelihood. That is, we only consider it non-critical if an asset is manufactured by a third party and/or is connected to the vehicle bus, while *all* other criteria are excluded. These considerations lead to $L_{a_i}^{\mathrm{tol,max}} = \boldsymbol{L}_{a_i}^{\mathrm{HL}} \cdot \boldsymbol{W}_{\mathrm{L}} = (0\ 0\ 0\ 0\ 1\ 1)^{\intercal} \cdot (0\ 0\ 0\ 0\ 0.095\ 0.048)^{\intercal} = \mathbf{0.143}$. The grayed fields of the risk matrix in Table 3 correspond to the tolerable risk. Since no asset has a tolerable high-level risk, a detailed security analysis is required for all zones and conduits.

4.2 Detailed Risk Analysis (ZCR 5.1 - ZCR 5.10)

The objective of **ZCR 5** is to move the unmitigated risk $r_{Z_i}^{\mathrm{u}}$ of potential threats in a zone $Z_i \in \mathcal{ZC}$ below the maximum tolerable zone risk $r_{Z_i}^{\mathrm{tol,max}}$. This is achieved by applying compensating security countermeasures, which lower $r_{Z_i}^{\mathrm{u}}$ and thus, move the achieved security level SL-A$_{Z_i}$ closer to the target security level SL-T$_{Z_i}$. A security level measures security demands arising from risks, whereas a risk results from a threat on a given asset in combination with at least one vulnerability. Thus, a crucial step for a reasonable risk analysis is the thorough determination of a threat and adversary model, taking into account known vulnerabilities and both the impact and likelihood of the identified threats.

Threat Modeling: A core prerequisite of a risk analysis is an exhaustive list of threats $\mathcal{T}$, since compensating security techniques may not protect the SUC from unidentified threats. During threat identification, we face two core problems: First, it remains impossible to prove completeness for $\mathcal{T}$, even though numerous identification techniques, such as *CIA*, *STRIDE*, and *Threat Trees* have been proposed [21]. Since threats are identified by the expert committee, we claim to have diverse views on the SUC and to obtain a reasonable number of threats. Additionally, we acknowledge the work by Petit and Shladover [17], who identified potential attack surfaces on road vehicles, that inspired our threat identification. Second, a collaborative threat identification process requires a common notion of a threat, when it comes to the *granularity* level. For instance, t_1 = *"The attacker triggers the vehicle brakes."* and t_2 = *"A network man-in-the-middle attacker injects forged braking commands."* are both candidates for threats with

the same outcome. However, t_1 is phrased on a purely functional level, while t_2 already addresses *one* potential attack scenario. While the author of t_1 may view the SUC at a coarser granularity level, he potentially misses attack vectors, as more than one vector can lead to the same outcome. In order to obtain threats with a comparable granularity level, we apply a three-round iterative threat identification process, as illustrated in Fig. 3. In a first round, we identify *top-level* threats on a purely functional level, having in mind *what* consequences are possible. After that, each top-level threat is split up into intermediate threats, taking into account *how* they can be realized, i.e., a precise attack vector. Since an attack vector can be used to realize more than one attack, an intermediate threat may appear multiple times. For instance, threats $t_{0\text{-}2}$ and $t_{2\text{-}2}$ in Table 4 are identical and are thus treated equally in succeeding steps.

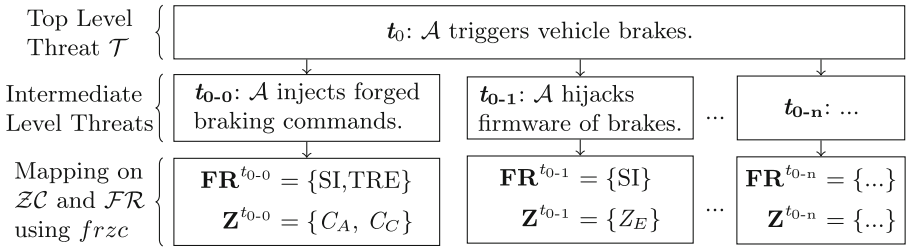

Fig. 3. Three-round iterative threat identification process

In the last round, each intermediate threat is associated with at least one zone or conduit $Z_i \in \mathcal{ZC}$ and at least one foundational requirement $fr \in \mathcal{FR}$, resulting in $\mathcal{T}_{Z_i}^{fr} \subseteq \mathcal{T}$. We formally model $\mathcal{T}_{Z_i}^{fr} = \{t \in \mathcal{T} : \pi_1(frzc(t)) = fr \wedge \pi_2(frzc(t)) = z\}$ with $frzc : \mathcal{T} \rightarrow (\mathcal{FR}^{\leq 7} \times \mathcal{ZC}^{\leq 9})$ and π the projection operation. In that way, we identified 63 intermediate threats.

Computation of Zone-Based Security Levels: Based on the identified threats, we derive a target security level SL-T$_{Z_i}$ for each zone $Z_i \in \mathcal{ZC}$ (**ZCR 5.6**). The security model HEAVENS [2] combines a threat level with the impact level to derive a security level. In contrast, IEC 62443-3-2 has no prescribed method to compute a security level. It only recommends to either represent SL-T$_{Z_i}$ as a scalar or as a vector. A scalar value minimizes the effort during verification, because the total number of possible states is kept low. In turn, a scalar may lead to over-engineering, since it does not allow a fine-grained requirement analysis. For instance, a zone requiring a confidentiality level of SL-4 would obtain SL-T$_{Z_i}$ as an overall security level, although precautions regarding other security goals may not be necessary. We express the security level of a zone $Z_i \in \mathcal{ZC}$ as a seven dimensional vector, where each dimension takes into account the unmitigated risks $r_{t_i}^{\text{u}} \in \mathcal{T}_{Z_i}^{fr}$ for a given $fr \in \mathcal{FR}$. In other words, we assign a security level to each foundational requirement and thereby, express to what

extent it is affected in Z_i. Precisely, we determine $\mathrm{SL}_{Z_i} = (\mathrm{SL}^{\mathrm{IAC}}_{Z_i}\, \mathrm{SL}^{\mathrm{UC}}_{Z_i} \ldots \mathrm{SL}^{\mathrm{RA}}_{Z_i})^{\intercal}$ where $\mathrm{SL}^{\mathrm{fr}}_{Z_i} = \max(< r^{\mathrm{u}}_{t_0}, r^{\mathrm{u}}_{t_1}, ..., r^{\mathrm{u}}_{t_n} >)$, $r^{\mathrm{u}}_{t_i} \in \mathcal{T}^{fr}_{Z_i}$, $fr \in \mathcal{FR}$. As an example, for the achieved security level of zone Z_A, we obtain $\text{SL-A}_{ZA} = (\text{SL-2}\,\text{SL-3}\,\text{SL-4}\,\text{SL-3}\,\text{SL-2}\,\text{SL-4}\,\text{SL-3})^{\intercal}$.

Similar to the high-level analysis, we express the unmitigated risk $r^{\mathrm{u}}_{t_i}$ of a threat t_i as a combination of likelihood (ZCR 5.4) and impact (ZCR 5.3). We suggest a *cascading parameter approach*, that evaluates likelihood and impact depending on further sizes like *vulnerabilities* and *attacker's capabilities*. This means, any change in one of those parameters immediately propagates to $r^{\mathrm{u}}_{t_i}$. Precisely, the likelihood for a successful threat t_i is determined by the required capabilities for its implementation and by exploitable vulnerabilities (ZCR 5.2). We acknowledge, that the idea of incorporating the attacker's capabilities into the threat likelihood has been already proposed in [16]. Similar to the HEAVENS project[2], we model the attacker's capabilities $\mathcal{AC}$ as three factors, *experience*

Table 4. Mapping of intermediate threats on zones and foundational requirements

Top Layer Threat	No.	Intermediate threat	Zones& conduits				Functional requirements				Risk class
			Z_A	Z_B	...	C_C	IAC	SI	...	RA	
Adversary prevents braking	$t_{0\text{-}0}$	Injection of fake braking commands			x	x		x	x		ex.high
	$t_{0\text{-}1}$	Illegal firmware on dynamic modules			x			x			Low
	$t_{0\text{-}2}$	Manipulation of brainstem firmware	x				x	x			high
	$t_{0\text{-}3}$	Malicious device to vehicle bus			x	x		x		x	ex.high
	$t_{0\text{-}4}$	Impersonation of control room			x		x				High
Manip-ulation of position	$t_{1\text{-}0}$	Injection of forged position data	x		x		x				ex.high
	$t_{1\text{-}1}$	GPS jamming			x					x	High
	$t_{1\text{-}2}$	Illegal firmware on localization ECU		x	x			x			Medium
Passenger Imprison-ment	$t_{2\text{-}0}$	Malicious firmware to door control			x		x				Low
	$t_{2\text{-}1}$	DoS on door control				x		x	x	x	ex.high
	$t_{2\text{-}2}$	Manipulation of brainstem firmware	x				x	x			High
Manipula-tion of traffic data	$t_{3\text{-}0}$	Sending fake traffic data to cloud			x		x	x			ex.high
	$t_{3\text{-}1}$	Map poisoning on cerebrum	x					x			Medium
	$t_{3\text{-}2}$	Forging traffic data from drone			x			x	x		High
...	...	...	...	...	...	...	...	...	...	...	...

$\mathcal{E}$, *knowledge* $\mathcal{K}$, and *resources* $\mathcal{R}$, i.e., $\mathcal{AC} = \mathcal{E} \times \mathcal{K} \times \mathcal{R}$. We score each of them independently for all threats t_i. Afterwards, we assign each identified vulnerability a value from $\mathcal{V} = \{severe, medium, negligible\}$. A severe vulnerability, for instance, may be a broken cryptographic protocol or a zero-day exploit, while a medium one requires user privileges. A negligible vulnerability is mostly theoretical, such as quantum attacks. Regarding the impact of t_i, we use the same criteria as we did for the high-level analysis, i.e., by ranking a threat according *Personal Safety*, *Operational Limitations*, and *Financial Loss*. Eventually, we receive a tuple for unmitigated risk $r^{\mathrm{u}}_{t_i}$ of every threat $t_i \in \mathcal{T}$ as illustrated in Table 4, allowing us to compute a target security level for all zones.

4.3 Threat Mitigation and Results

After having computed a security level SL-A$_{Zi}$ for each zone $Z_i \in \mathcal{ZC}$, we iteratively identify and apply compensating countermeasures, such that the unmitigated risk $r^{\mathrm{u}}_{t_i}$ of threat $t_i \in \mathcal{T}^{fr}_{Z_i}, \forall fr \in \mathcal{FR}$ shrinks below the tolerable risk $r^{\mathrm{tol,max}}_{Z_i}$. For this purpose, we constantly compare $r^{\mathrm{u}}_{t_i} \leq r^{\mathrm{tol,max}}_{Z_i}$ by reassessing all parameters, that impact likelihood and impact (ZCR 5.7-5.10). For instance, a compensating countermeasure for the foundational requirement *System Integrity* in zone Z_A is data authentication. Assuming the verifiable authenticity of in-vehicle traffic, the required capabilities to inject fake braking commands without being recognized (c.f. threat $t_{0\text{-}0}$ in Table 4) rise significantly, since an attacker would need to circumvent cryptographic protection. This, in turn, makes $t_{0\text{-}0}$ less likely and consequently, $r^{\mathrm{u}}_{t_{0\text{-}0}}$ decreases. Beforehand, the expert committee defines a tolerable risk $r^{\mathrm{tol,max}}_{Z_i,fr}$ for each zone $Z_i \in \mathcal{ZC}$ and foundational requirement $fr \in \mathcal{FR}$. Precisely, they determine the tuple $r^{\mathrm{tol,max}}_{Z_i,fr} = (I^{\mathrm{tol,max}}_{Z_i,fr}, L^{\mathrm{tol,max}}_{Z_i,fr})$ and compare it with all $r^{\mathrm{u}}_{t_i}$, $t_i \in \mathcal{T}^{fr}_{Z_i}$. For instance, the maximum tolerable likelihood of the foundational requirement *Use Control* for zone Z_F (control room) is set to *extremely low*, because only individuals with assigned privileges are allowed to remotely control a vehicle. According to Table 3, this leads to $L^{\mathrm{tol,max}}_{C_F,\mathrm{UC}} = 0.191$. As the high-level analysis in Sect. 4.1 has revealed, the malicious operation of the control room can lead to life-threatening situations. Therefore, we set the maximum tolerable impact $I^{\mathrm{tol,max}}_{Z_F,\mathrm{UC}}$ to *extremely low*, i.e., $I^{\mathrm{tol,max}}_{Z_F,\mathrm{UC}} = 0.26$. As a result, we obtain $r^{\mathrm{tol,max}}_{Z_F,\mathrm{UC}} = (I^{\mathrm{tol,max}}_{Z_F,\mathrm{UC}}, L^{\mathrm{tol,max}}_{Z_F,\mathrm{UC}}) = (0.26, 0.191)$.

Depending on the security level, IEC 62443-3-3 provides countermeasures for each foundational requirement. However, we argue, that most of them are not directly applicable to our SUC, since they have not been designed for automotive challenges. That is, *real-time behavior*, *resource-constrained* control units, and a high *reliability*. For example, a security level SL-2 of the foundational requirement *Identification and Authentification Control* demands for public key infrastructure certificates. This, however, is hardly applicable to in-vehicle communication, because public key certificates lead to unacceptable overhead, as they come along with long certification chains and demanding cryptographic operations. As a consequence, we looked for alternative, more lightweight, and resource-saving solutions. In particular, the work of El-Rewini et al. [9] inspired

us, as they provide an extensive survey on automotive security frameworks. As a result, we obtain a detailed list of security mitigation techniques for each zone, that protect against the 63 identified threats. They can be summarized as follows:

Authenticated In-Vehicle Communication: A core prerequisite of safety-critical in-vehicle conduits and zones (Z_A, Z_E, C_A, C_B) is the ability to verify the authenticity of data streams. In this way, the injection of fake commands becomes difficult. A promising alternative technique are implicit certificates in combination with distinct physical memory characteristics [18], avoiding potentially long and resource-consuming certification chains. Related work [9] illustrates additional solutions for the wide variety of in-vehicle communication protocols.

Integrity of In-Vehicle Control Units: Since both the SUC and contemporary road vehicles possess an increasingly large number of external interfaces, and the ability to remotely update control units, adversaries, residing inside the ECUs, must be prevented. In our case, we propose to treat the brainstem as a trust anchor, that verifies the integrity of all control units before the vehicle start. For this purpose, we suggest *Remote Attestation* (RA), a technique, allowing to prove the integrity of a device to a third party. Kohnhäuser et al. [13] show how to use RA in the automotive domain.

Malicious Behavior Detection: During the security analysis, high-risk threats on safety-critical assets were associated with the foundational requirement *Timely Response to Events* (TRE). Specifically, adversaries connecting to the in-vehicle bus may be able to flood the in-vehicle network (DoS attack) and thus to cause failures. We find an anomaly-based intrusion detection system [8] for safety-critical in-vehicle traffic (i.e., conduits C_A, C_B) a suitable compensating countermeasure. Also, inter-vehicle communication (i.e., zone C_C) is prone to DoS attacks, for which, however, many mitigation frameworks have been presented [23].

Data Separation: The initial design of the SUC provides a single in-vehicle bus for all data flows. Consequently, user input and potentially safety-critical data streams are mixed, which may lead to the delayed transmission of safety-critical demands. As our risk analysis revealed threats affecting the foundational requirement *Restricted Data Flow* for in-vehicle traffic, the presented vehicular architecture requires means to separate data flows. Both physical and virtual data separation achieve this goal. For our SUC, we configure VLAN priority levels for the Ethernet-based in-vehicle network and use the arbitration logic of the CAN bus. The FlexRay network inherently realizes a TDMA-based schedule, allowing to reserve dedicated time slots for critical data.

Access Control: An integral part of the SUC is the control room, that enables human remote control in case of emergency situations. In order to distinguish between legitimate and illegitimate remote control, an access control system is

necessary at the vehicle's edge. This is highlighted by the risk analysis, that exposes the importance of user authentication, in particular when it comes to the communication between the vehicle and its exterior. Since the router is the only gateway to the external world, access control mechanisms have to be implemented in the corresponding zone (Z_A). This includes a strict deny-by-default policy and mutual identity checks.

5 Discussion

Our analysis particularly highlights the demand for system integrity and timely responses, since a significant number of threats are mapped on the corresponding foundational requirements. Both software and communication integrity are key factors for a safe driving state. This evidence coincides with related work [17], that considers the injection of fake messages as one of the severest attacks on modern vehicles.

Although our analysis yields effective means to protect against the identified threats, we lack techniques to handle actual security incidents during vehicle operation. We need means to assess them in real-time and to adopt appropriate (safety) measures. We plan to address this problem in future work. Regarding our security requirement analysis, we want to stress the following points:

Quality of Assessments: For our consensus-based risk analysis, we presented a scoring scheme, using Simple Additive Weighting (SAW) as a decision support system. In order to obtain reasonable and consistent assessments, and to ensure a broad insight into the SUC, we engaged an expert committee, consisting of eight computer scientists and mathematicians from the *Securing Engineering Lab*[1] of the *Technische Univeristät Darmstadt*. As a first step, the experts have been thoroughly introduced to the novel vehicular architecture in a Q&A session. Afterwards, we established the presented set of evaluation criteria based on related work and empirical values. As both the threat identification process and all assessments have been jointly accomplished by the expert committee, we argue to properly address subjectivity and vagueness. However, we admit, that a higher committee heterogeneity in terms of educational background may yield even better results. Regarding the proposed scoring scheme, we currently assign a fixed probabilistic weight to each evaluation criterion. We consider the Fuzzy Analytic Hierarchy Process (FAHP) [7] an effective alternative to determine preference weight, but leave this for future work.

Safety and Security Co-Engineering: We pursued the question of what changes are necessary for a safety-aware security risk analysis in the automotive domain. As safety and security demands may contradict, the possibility of prioritization is crucial. We find the mapping of risks onto zones and foundational requirements a promising technique, because it allows fine-grained solutions for large-scale systems. The partition of the SUC into zones and conduits should take safety criteria into consideration. In addition, we suggest the following adaptions:

[1] http://www.seceng.de.

- So far, the seven foundational requirements are purely security-related. We suggest extending them with safety requirements such as *reliability*, *redundancy*, and *real-time behavior*. By doing so, the unmitigated risk of a threat or of a hazardous situation would take both safety and security dimensions into account. The presented vector-driven approach allows prioritization, by pointing out which foundational requirement is affected most by a set of threats. Appropriate countermeasures can be deduced in that way.
- The countermeasures listed in IEC 62443-3-3 need to be adjusted for the automotive domain. Instead of user-oriented, potentially computationally heavy systems (e.g., PKI, multi-factor authentication, ...), lightweight and resource-saving techniques (e.g., implicit certification, hardware-based security, ...) are worthwhile. There has been extensive work on automotive security with numerous frameworks, covering many automotive challenges [8,9,14,23], that should be included in a future standard.
- A common set of evaluation criteria and a consistent scoring scheme for automotive systems is desirable, in order to make analysis results comparable. We presented a scheme, that incorporates both security and safety criteria for risk assessment.

5.1 Comparison to ISO/SAE DIS 21434

The high-level risk analysis and the subsequent partition into zones and conduits allow for efficient identification of relevant assets. Besides, the analysis process becomes more scalable, since uncritical assets are excluded from further steps. The use of foundational requirements as a reference point enables the clear establishment of mitigation techniques.

While the detailed risk analysis of IEC 62443-3-2 begins with the identification of threats, the novel ISO/SAE DIS 21434 starts from potential damage scenarios and traces them back to attack paths. More precisely, the risk assessment methods are comprised of seven steps (I-VII). Initially, damage scenarios are identified, which may occur through compromised assets (I). A damage scenario is triggered by a set of adverse actions, a so-called threat scenario, which are enumerated in (II). The impact of each damage scenario is assessed according to four core categories of consequences, *Safety*, *Financial*, *Operational*, and *Privacy* (III). Subsequently, each threat scenario is decomposed into attack paths in a top-down or bottom-up approach (IV). The feasibility of each path is assessed according to a pre-defined scale (V), resulting in a risk value (VI) for each threat scenario, which also incorporates the impact of the damage scenario. Finally, risk reduction methods shall be realized (VII). In case the risk for a threat scenario has to be reduced, a Cybersecurity Assurance Level (CAL) reveals requirements for the affected item.

At first glance, the risk analysis process of ISO/SAE DIS 21434 and IEC 62443 have little in common. On closer inspection, however, both standards do share similar concepts. The CAL is similar to the SL-T value, which is only determined if the risk is too large. Instead of our iterative threat identification process and conduits, attack paths cover propagating effects of adverse

actions. While the idea of decomposing threat scenarios into attack paths is the most promising feature of ISO/SAE 21434, our work reveals requirements that are not yet met by ISO/SAE 21434. Unlike IEC 62443, the novel automotive standard prescribes assessment criteria and parameters. However, it insists on neither underlying cybersecurity requirements nor mitigation techniques, contrary to IEC 62443. For the sake of a common minimum security perception, suggestions of countermeasures for specific CALs would be helpful, in particular, because road vehicles are generally subjected to the same safety and legal requirements. Also, consistent scoring schemes and a dedicated process to identify relevant critical assets of a potentially complex architecture would be desirable.

6 Conclusion

In this work, we presented a consensus-based implementation of the generic IEC 62443 cybersecurity standard for a novel vehicular architecture. In particular, we identified risk evaluation criteria and developed an additive scoring scheme to assess automotive risks. Furthermore, we introduced a hierarchical threat model for a collaborative threat identification process. We used a cascading parameter approach to express risks as zone-based vectors, yielding fine-grained security levels, that express security requirements. We conclude that especially data and software integrity, the separation of safety-critical commands, and the ability to detect anomalies are crucial for automated vehicles. Based on our lessons learned, we find as essential for a future standard the systematic partition of a potentially complex vehicular architecture into relevant assets, the computation of security levels with regard to cybersecurity reference goals, the treatment of transitive adverse actions, and the suggestion of mitigation techniques. We also make suggestions on how to incorporate safety requirements into a future standard. In particular, a safety-aware automotive security standard should use a redefined set of foundational requirements, including safety objectives such as reliability, redundancy, and real-time behavior. Although IEC 62443 has been originally designed for IACS, we promote its applicability to the automotive domain in combination with the adaptions suggested in this work.

Acknowledgement. This work has been accomplished within the project "UNICARagil" (FKZ 16EMO0392). We acknowledge the financial support for the project by the Federal Ministry of Education and Research of Germany (BMBF). We also thank the *Security Engineering Group* for their support in assessing, scoring, and ranking security parameters.

References

1. ISO 26262 Road vehicles - Functional Safety. Standard, International Organization for Standardization (2011)
2. Healing vulnerabilities to enhance software security and safety (HEAVENS) project (2016). https://research.chalmers.se/en/project/5809. Accessed 25 Feb 2020

3. SAE J3061 - Cybersecurity Guidebook for Cyber-Physical Vehicle Systems. Standard, Society of Automotive Engineers (2016)
4. ISA-62443 Security for Industrial Automation and Control Systems. Standard, International Society of Automaton (2017)
5. ISO/SAE DIS 21434:2020(E): Road vehicles - cybersecurity engineering (2020)
6. Checkoway, S., et al.: Comprehensive experimental analyses of automotive attack surfaces. In: USENIX Security Symposium, vol. 4, pp. 447–462 (2011)
7. Chen, V.Y., et al.: Fuzzy mcdm approach for selecting the best environment-watershed plan. Appl. Soft Comput. **11**(1), 265–275 (2011)
8. Cho, K.T., Shin, K.G.: Fingerprinting electronic control units for vehicle intrusion detection. In: 25th USENIX Security Symposium, pp. 911–927 (2016)
9. El-Rewini, Z., et al.: Cybersecurity challenges in vehicular communications. Veh. Commun., 100214 (2019)
10. Henniger, O., et al.: Securing vehicular on-board IT systems: the Evita project. In: VDI/VW Automotive Security Conference, p. 41 (2009)
11. Kampmann, A., et al.: A dynamic service-oriented software architecture for highly automated vehicles. In: 2019 ITSC, pp. 2101–2108. IEEE (2019)
12. Keilhoff, D., et al.: UNICARagil – new architectures for disruptive vehicle concepts. 19. Internationales Stuttgarter Symposium. P, pp. 830–842. Springer, Wiesbaden (2019). https://doi.org/10.1007/978-3-658-25939-6_65
13. Kohnhäuser, F., et al.: Ensuring the safe and secure operation of electronic control units in road vehicles. In: 2019 IEEE Security and Privacy Workshops (SPW)
14. Mejri, M., et al.: Survey on VANET security challenges and possible cryptographic solutions. Veh. Commun. **1**(2), 53–66 (2014)
15. Nie, S., Liu, L., Du, Y.: Free-fall: Hacking tesla from wireless to CAN bus, pp. 1–16. Briefing, Black Hat USA (2017)
16. Ben Othmane, L., et al.: Incorporating attacker capabilities in risk estimation and mitigation. Comput. Secur. **51**, 41–61 (2015)
17. Petit, J., Shladover, S.E.: Potential cyberattacks on automated vehicles. IEEE Trans. Intell. Transp. Syst. **16**(2), 546–556 (2014)
18. Püllen, D., et al.: Using implicit certification to efficiently establish authenticated group keys for in-vehicle networks. In: 2019 IEEE VNC, pp. 1–8 (2019)
19. Putra, D.W.T., Punggara, A.A.: Comparison analysis of Simple Additive Weighting (SAW) and weigthed product (WP) in decision support systems, p. 01003 (2018)
20. Schmittner, C., Ma, Z., Reyes, C., Dillinger, O., Puschner, P.: Using SAE J3061 for automotive security requirement engineering. In: Skavhaug, A., Guiochet, J., Schoitsch, E., Bitsch, F. (eds.) SAFECOMP 2016. LNCS, vol. 9923, pp. 157–170. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45480-1_13
21. Shevchenko, N., Chick, T.A., O'Riordan, P., Scanlon, T.P., Woody, C.: Threat modeling: a summary of available methods, July 2018
22. U.S. DoT: Revised departmental guidance 2016: Treatment of the value of preventing fatalities and injuries in preparing economic analyses (2016)
23. Verma, K., et al.: Prevention of DoS attacks in VANET. Wireless Pers. Commun. **73**(1), 95–126 (2013)
24. Woopen, T., et al.: UNICARagil-disruptive modular architectures for agile, automated vehicle concepts. Aachener Kolloquium GbR (2018)
25. Yan, C., Xu, W., Liu, J.: Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle. DEF CON **24** (2016)