



The Automotive BlackBox: Towards a Standardization of Automotive Digital Forensics

Downloaded from: <https://research.chalmers.se>, 2024-04-27 06:12 UTC

Citation for the original published paper (version of record):

Strandberg, K., Arnljung, U., Olovsson, T. (2023). The Automotive BlackBox: Towards a Standardization of Automotive Digital Forensics. WIFS 2023 - IEEE Workshop on Information Forensics and Security. <http://dx.doi.org/10.1109/WIFS58808.2023.10375003>

N.B. When citing this work, cite the original published paper.

© 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, or reuse of any copyrighted component of this work in other works.

The Automotive BlackBox: Towards a Standardization of Automotive Digital Forensics

Kim Strandberg^{*†}, Ulf Arnljung^{*}, Tomas Olovsson[†]

^{*}Volvo Car Corporation, Sweden {firstname.lastname}@volvocars.com

[†]Chalmers University of Technology, Sweden, {firstname.lastname}@chalmers.se

Abstract—There is a trend toward increased cyberattacks on vehicles. Aligned, forensics requirements and standards are emerging. Digital forensics refers to identifying, preserving, verifying, analyzing, documenting, and finally presenting digital evidence with high confidence in its admissibility, thus ensuring forensics soundness. However, current automotive regulations and standards, such as the United Nations Regulation No. 155 and the International Organization for Standardization standard 21434, provide no details or guidelines. Vehicular data is often extracted using tools unsuitable for digital forensics, thus lacking forensics soundness. The data storage is generally not resistant to tampering and often lacks adequate cybersecurity mechanisms.

Digital forensics is a relatively new field within the automotive domain, where most of the existing self-monitoring and diagnostic systems only monitor safety-related events. To support a forensic investigation, automotive systems must be extended to securely log and store additional information, especially those related to security events. There is no standardization for automotive digital forensics that defines requirements, needed components, and techniques for the automotive domain. In this paper, we identify and propose requirements for automotive digital forensics and present the Automotive BlackBox, an architecture guiding the design of an automotive digital forensic-enabled vehicle.

Index Terms—vehicle forensics, automotive forensics, forensics architecture, forensics guidelines, automotive security

I. INTRODUCTION

The complexity of vehicles is increasing at a high pace. A vehicle today can contain around 150 Electronic Control Units (ECUs) and has various connection interfaces, which inherently implies a large amount of data exchange between many entities such as sensors, actuators, ECUs, the Internet, and infrastructure. If this data is assessed satisfactorily through automated processes, a wealth of significant information can be provided to stakeholders such as law enforcement, insurance companies, and manufacturers.

Increased complexity increases the risk of system vulnerabilities and, consequently, the number of potential attack vectors. At the same time, the increased connectivity gives a higher potential to find exploits related to vulnerabilities due to a larger attack surface. For instance, a buffer overflow vulnerability in software (i.e., an attack vector) can be exploited due to an increased attack surface (e.g., a connection interface), enabling the capability to execute arbitrary code and thus potentially disrupt vital in-vehicle functions. Moreover, attacks can be associated with life-threatening hazards due to their potential to affect safety-critical systems such as brakes, steering, and engine control. Thus, such attacks are highly relevant to identifying and tracing in a post-incident digital forensic investigation. It has been shown several times that vehicle cyberattacks have to be taken seriously, e.g., practical attacks in [1], [2] that demonstrates the fragility

of automotive systems and the susceptibility to malicious actions to disrupt and modify these systems. For instance, in [1], the firmware was extracted and reverse-engineered to understand hardware features which enabled them to add new functionalities related to their attacks, such as remote access persistence via the cellular connection. Moreover, they managed to add malicious code to a vehicle telematics unit which automatically erased any evidence of its existence after a crash; thus, there was no post-incident available data related to a potentially life-threatening code. Due to the continuous increase in complexity and connectivity, it is only logical to assume that cyber-attacks against vehicles will continue to rise and be even more prevalent. Thus establishing guidelines for forensic automotive design to enable the detection and post-analysis of cyberattacks is imperative.

However, Automotive Digital Forensics (ADF) is a relatively new field within the automotive domain. Most existing self-monitoring and diagnostic systems only monitor safety-related events, such as the status of brakes, seat belts, and airbag deployments, via an Event Data Recorder (EDR). Current vehicle EDRs are used mainly to record limited events under a few seconds before and during a crash, while, e.g., flight data recorders can record hundreds of parameters for many hours. Numerous vehicle manufacturers already transmit EDR-related data to a central location, such as the GM's OnStar, in the occurrence of a crash [1]. To support a forensic investigation, these systems must be extended to log additional information, especially cybersecurity-related, e.g., cyberattacks. A satisfactory ADF solution must consider in-vehicle data and its surroundings from an individual, vehicle fleet, and infrastructure perspective.

In [3], four main stakeholders are identified for ADF, namely: Law Enforcement (LE), Vehicle Manufacture (VM), Vehicle Drivers (VD), and Insurance Companies (IC). LE refers to, e.g., the police and related legal systems. VM requires ADF data for fault-tracing, e.g., to distinguish hardware and software failures with non-malicious origin from cybersecurity incidents, e.g., attacks from threat actors. VD might try to remove or manipulate forensic evidence with the intent to hide traces of crime, whereas IC are interested in insurance cases and cost and risk profiling/statistics. Stakeholders, e.g., LE, IC, and VM, must establish a trustworthy and admissible chain of events to derive the cause of accidents concerning malicious actors, e.g., hackers and terrorists, and non-malicious actors/origins, e.g., weather and animals on the road.

Contributions. We present the Automotive BlackBox, an architecture guiding the design of an automotive digital forensic-

enabled vehicle. We highlight challenges, identify forensic components, and propose a standard data format, techniques, goals, and requirements in an architectural automotive context, considering current and upcoming regulations and standards. Based on our previous work, a systematic literature review of the area [3], our contributions are novel and relevant for automotive digital forensics investigations.

II. CHALLENGES

Quite a few challenges need to be considered when establishing requirements for ADF. A modern vehicle consists of many devices running various operative systems. Furthermore, they communicate over many different protocols internally and with the outside world via Vehicle-2-everything (V2X) communication. An immense amount of data is continuously transmitted, e.g., with safety-critical systems, including brakes, steering, and acceleration. In previous work, we identified 16 categories of forensically relevant data and stated the required security properties for the data [3]. However, modern vehicles only log a fraction of forensically relevant data, and manual approaches are often used to manage the data.

A vehicle has various devices where the data is spread out in multiple places in a distributed fashion, e.g., different ECUs, networks, and the cloud. Locating all devices containing relevant data is challenging since the vehicle's proprietary architecture. Data can be stored in Virtual Machines (VMs), where data in registry entries and temporary files can be erased when turning off or rebooting the machine. Thus, there is a need to extend and automate data collection covering all relevant data. However, currently, there are no standardized interfaces for information extraction and no standardized format for storage. For instance, sometimes desoldering memory chips are required to extract data. Moreover, forensic investigations require following an established process, a scientifically proven methodology, and using validated tools and techniques to maintain the chain of custody, but that is currently only sometimes the case since the lack of standardization within ADF forces OEMs to develop and use their tools and strategies for fault tracing and data collection.

Manual approaches for managing the steady increase in forensically relevant data, considering data collection, extraction, and analysis, are time-consuming. Sufficient pertinent data needs to be improved, and the security mechanism needs to be more robust in ensuring trustable data. For instance, many legacy systems and protocols currently lack satisfactory security features. In many cases, there is a need for more performance, better storage capacity, and increased data security to enable a reasonable level for ADF. The related cost of fixing these issues is challenging. There is also multi-jurisdictional litigation to consider, sometimes contradictory, e.g., privacy regulations [4] versus requirements for data collection for forensics investigations [5]. There are requirements to secure data [5], which makes data availability for forensics challenging due to the inaccessibility of secret keys for decryption. Moreover, security techniques required for forensics

might negatively affect requirements for safety-related time-critical systems [6]. Thus, requirements for privacy, forensics, cybersecurity, and safety regulations sometimes conflict.

In summary, we conclude the following main challenges: (i) Only a fraction of logging and analysis is currently performed on available data. Moreover, data is spread out in various places making identification and retrieval time-consuming. Thus, an increase in automated data collection for all relevant data is needed. (ii) Due to, e.g., cost and performance restraints in current vehicles, trustable data is often not ensured, nor are common security properties fulfilled for digital evidence. Thus, there are requirements to secure potential evidence better. (iii) Regulations, standards, and common guidelines within ADF concerning, e.g., forensics processes, data collection, management, formats, and tools must be evolved and revised. Thus, there is a necessity to standardize ADF to ensure forensic soundness. (iv) Regulations in different fields and countries must be revised to align, i.e., privacy, forensics, security, and safety. Variations of forensic solutions must be considered in different countries. The following sections consider challenges i-iii and leave iv as further work.

III. DIGITAL FORENSICS PRINCIPLES

Digital forensics is a field with strong dependencies on information security. It is imperative to ensure available trustable data. Although digital forensics mainly emphasizes post-incident, cybersecurity aims to mitigate threats to forensic data, such as removing and manipulating digital evidence.

Digital forensics includes the collection and investigation of data, generally about crime. Security techniques must be used for the data to be admissible in a court of law. Generally accepted principles of a digital forensic investigation apply to ADF. However, the naming and number of steps might differ between methodologies, although the core concepts are usually the same [3], namely: (i) *Identification*. Has a crime occurred? What data is relevant, and where is the data stored? What resources (e.g., tools and experts) are needed? (ii) *Preservation*. How can we preserve data integrity (e.g., running devices, remote access, extraction, and anti-forensics)? (iii) *Acquisition and verification*. How can we extract (e.g., imaging, log files, live acquisition) and validate the data's authenticity (e.g., signatures and hashes)? (iv) *Analysis*. What type of information is relevant to assess? (v) *Reporting*. How can we document all parts of the forensic investigation and its related result to be admissible in a court of law? Moreover, law enforcement guidelines need to be considered, such as the four Association of Chief Police Officers (ACPO) principles [7].

A forensic investigation requires establishing trust in the chain of events, where the life cycle of the data must be considered. Thus, processes for handling forensic data are needed as technical solutions to collect and secure forensics data. Any inadequacies in these two can potentially devastate the forensic case, e.g., making the data invalid for the investigation.

IV. THE AUTOMOTIVE BLACKBOX

As mentioned in Section I and II, the forensics mechanisms of today's vehicles, e.g., EDRs, are insufficient for ADF. A more comprehensive approach is necessary to align with current and upcoming standards and regulations. In the remainder of this section, we state an attacker model, Automotive Digital Forensics Goals (ADFG), requirements, technical details, and a reference architecture for ADF.

Attacker Model. We consider the six threat actors as stated in [8], namely, the Financial Actor (FA), the Foreign Country (FC), the Cyber Terrorist (CT), the Insider (IN), the Hacktivist (HA), and the Script Kiddie (SK). We assume a common agenda to perform various cybercrime targeting vehicles with the potential to affect the driver, passengers, and objects in the vicinity using the vehicle as leverage. However, the main objective is to hide, delete or manipulate digital evidence, such as digital traces of crimes, to obstruct or prevent forensic investigation.

Automotive Digital Forensics Goals and Requirements. Based on the attacker model and previously mentioned principles and challenges, we establish six ADFG. ADFG-1 is a general rule based on the availability and trust of digital evidence, and ADFG-2 to ADFG-6 are more specific based on accepted forensic principles. Challenges are summarized at the end of Section II, where ADFG-1 assesses challenges (i) and (ii), and ADFG-2 to ADFG-6 assess challenge (iii). With the attacker model in mind, these ADFG are further mapped into specific requirements inspired by a standardized approach for argument notation to demonstrate coverage [9] (cf. Table I and Figure 1). For instance, ADFG-1 is general and about *Availability and Trust* and needs R1-R7. ADFG-2 is more specific and concerns *Data Identification* and needs R2, R4-R6, and R9-R10. As emphasized in previous work by us [3], we adopt the well-known CIA security triad extended with two other properties, NP, where the first four are prerequisites for securing vehicle forensic data and the fifth for personal data. An explanation of these properties follows. *Confidentiality(C)* guarantees that only authorized entities can access and disclose data. *Privacy(P)* concerns personal data, such as traffic violations, location data, and synced data from external devices, e.g., text messages and phone records. Therefore, such data must be protected according to local laws and regulations [4]. *Authenticity* is a form of *integrity(I)* ensuring data origin and is imperative for forensic investigations. *Availability(A)*, e.g., in the event of a crash, must be ensured, and secure and tamper-proof storage guaranteed. *Non-repudiation(N)* ensures that occurrences of events and their origin can not be denied. Therefore, *authenticity* and *integrity* are required for *non-repudiation*. An explanation of the six ADFG follows.

ADFG-1: Availability and Trust of Digital Evidence. A prerequisite for ADFG-2 to ADFG-5 is available and authentic data. Thus, we identify requirements for technical solutions to detect and securely store forensically relevant events, including fulfilling R1 (cf. Table I) and the CIANP properties for digital evidence where applicable. **ADFG-2: Identification.**

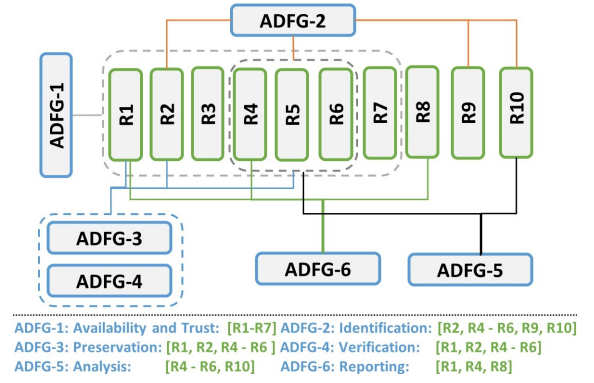


Fig. 1. Mapping of Automotive Digital Forensics Goals to Requirements

The first step is to identify what has happened. Has there been a crash? Can anyone describe the incident? What is the most relevant data to assess? To identify evidence, a prerequisite is satisfactory data collection and filtering. An Intrusion Detection System (IDS) shall, therefore, detect and securely store events related to anomalies and predefined patterns. Moreover, other forensically relevant events shall be considered. For instance, time for braking, acceleration, seat-belt traction, airbag deployment, weather conditions, location data, and detected warnings (e.g., tired driver, lane assist, V2X data) can all be relevant to contribute to establishing the cause of an incident. **ADFG-3: Preservation.** How can we guarantee integrity and privacy during data collection? How can we ensure that relevant data is recovered? Can the engine be turned off? Is there a risk that data can be erased by a perpetrator remotely? Potential evidence shall be stored securely, considering the CIANP properties. **ADFG-4: Verification.** How can we validate the authenticity of the data? Evidence shall be stored in a standard format that includes the potential to validate time, integrity, and origin. **ADFG-5: Analysis.** What data is relevant to assess concerning the crime investigated? Forensic data shall be identifiable concerning the type of data and the order of occurrence. For instance, detecting anomalies in the network aligned with normal events such as opening and closing doors, speed, braking, and location data. Data collection and analysis shall be automated, and manual work reduced to a minimum. For instance, incorporate Artificial Intelligence (AI) and Machine Learning (ML) approaches for automated data management. **ADFG-6: Reporting.** How can we document the evidence and ensure admissibility in legal proceedings? It shall be possible to identify relevant data for a predefined period concerning the type and order of events in relation to a potential crime. Data shall be verifiable concerning authenticity with a detailed timeline for the events.

A. Technical details

IDS shall detect *Indicators of Attacks (IoA)* and store *Indicators of Compromise (IoC)*. The main distinction between IoAs and IoCs is that the former are ongoing events of potential attacks, while the latter are events indicating a previous com-

TABLE I
ADF REQUIREMENTS

Requirement R1: fulfilment of CIANP. R1a. Confidentiality. R1b. Integrity. R1c. Availability. R1d. Non-Repudiation. R1e. Privacy
Requirement R2: secure logging, storage and extraction. R2a. There shall be mechanisms that guarantee the authenticity of logged and stored data. R1 and mechanisms for preventing data modification, tampering, and deletion shall be considered. R2b. Storage shall be constructed with physical integrity in mind, thus, to survive crashes and physical violence. R2c. Forensically relevant events shall be securely stored for fault tracing and post-incident investigations. For a list of relevant data to consider, we refer to [3]. R2d. A secure physical extraction interface shall exist, requiring mutual authentication to extract forensic images.
Requirement R3: infrastructure and communication. The infrastructure, cryptographic algorithms, and key material shall follow best security practices.
Requirement R4: common format and tools. Forensic data shall have a common format. The format shall be verifiable and contain information about the logical order of occurrence. The tools used shall adhere to standardized, accepted, and regulated digital forensics processes.
Requirement R5: time. It shall be possible to trace the logical order for events according to a time value, e.g., the logical and clock time. Thus, the forensic system requires trust in a time server and an agreement on the logical order of events.
Requirement R6: redundancy. Relevant redundancy shall be used to ensure that data is authentic and available. For instance, the same data stored in different sources, such as in-vehicle and cloud data, shall be possible to verify its identical and detect deviations.
Requirement R7: secure boot. State-of-the-art secure boot protection mechanisms shall be used where applicable, e.g., manipulations in relevant entities, such as the IDS, SIEM, and the Automotive Blackbox, shall be detected.
Requirement R8: least privilege. Data shall only be available to authorized entities.
Requirement R9: Intrusion Detection/Prevention Systems. IDSs/IPs shall detect and react to anomalies from normal communication patterns and known attacks, e.g., maintain secure logging of relevant events (R2).
Requirement R10: threat intelligence. Learning about attacks to keep pace with attackers shall be possible, for instance, using honeypots and analyzing, correlating, and mapping data from multiple sources.

promise. Thus, one or many IoAs can give rise to IoCs, where the latter is most relevant from an ADF perspective. IDSs can record anomalies from a predefined pattern (anomaly-based IDS) and detect specific signatures (specification-based IDS). The former is more suitable for detecting unknown attacks, and the latter is better at detecting known attacks. A higher rate of false positives is usually the case for the former and false negatives for the latter. Thus, a hybrid approach is beneficial to increase coverage. IoCs from IDS can be forensic evidence of potential network and ECU breaches, e.g., unusual traffic and other deviations. An example can be that the speed should always be zero when the vehicle is in parking mode. Any mismatch in specific signals or vehicle status can indicate IoAs or IoCs. Other examples are failed authorization attempts (e.g., attempted access of privilege mode via debug ports), invalid software signatures during updates, or the secure boot process.

IoCs and IoAs from IDSs are managed by a *Security Information and Event Management (SIEM)* along with other detected relevant events, such as safety-related, e.g., braking, acceleration, steering, engine control, airbag release, and seat belt traction. Examples of non-safety-related events are software update events, location, opening/closing of doors, and executed diagnostics. V2X communication with infrastructure, other vehicles, and external devices can be forensically relevant. We do not aim to provide a complete list of forensically relevant data. Still, we refer to our previous work [3], which identified relevant ADF data. SIEM offers real-time monitoring, analysis, data collection, and storage of events and logs from various sources. It creates an in-depth overview

of previous and ongoing events for threat management and auditing purposes, e.g., threat mitigation, fault tracing, and ADF. AI and ML approaches automate management, e.g., rating alerts. SIEM data and automated analysis are further transferred to a Cyber Incident Response Team (CIRT) for further analysis and decision-making.

B. Architecture

We propose a core architecture with domain separation according to Figure 2, where hybrid IDS components detect and securely log events. As shown, sensors can detect specified ECU events and anomalies in communication, further sent to SIEM for automatic analysis and secure storage according to a predefined format. We propose to use a hybrid SIEM, divided into a local (L) and cloud (C) part, adaptable where analysis occurs dependent on performance and cost restraints. For instance, L-SIEM can be implemented with measures during cyberattacks, e.g., log and analyze, while others are offloaded to C-SIEM, which takes further decisions and generate fleet responses. Another option, if performance/cost is an issue, is to run the local part in log-only mode, i.e., only log events and create images for a defined period. C-SIEM and the CIRT entirely perform the analysis for the latter case.

The L-SIEM stores the events in the correct order in the Automotive BlackBox, including time, and a counter that keeps track of the number of occurrences for each event. A pre-shared certificate between C-SIEM and L-SIEM can be used for the key-wrapping of symmetric keys to ensure secure storage/transfer. For such a case, L-SIEM can create a list of symmetric keys further used to encrypt images of forensic data. In turn, keys are encrypted with an asymmetric public key from the public part of the shared certificate. Only C-SIEM and the CIRT team can access the corresponding private key. Thus, the required symmetric keys are kept in escrow to protect user privacy. For instance, if a malicious entity manages to extract or manipulate digital evidence stored in the Automotive BlackBox, it is still encrypted, ensuring confidentiality/privacy, and signed, ensuring integrity/non-repudiation.

In summary, we propose that L-SIEM use the public part of a pre-stored encryption certificate and the private part of a pre-stored signing certificate, where the C-SIEM has access to the corresponding part for decryption and validation. Thus, any authorized entity, such as a forensic investigator, must request the symmetric decryption keys from the CIRT to access and disclose potential digital evidence. The L-SIEM securely, i.e., encrypted and mutually authenticated, uploads forensics data, i.e., the image, to the cloud for a specified time interval and a.s.a.p. if a vehicle is out of range from connectivity. The C-SIEM verifies the image signature before storage. The C-SIEM and the *Automotive BlackBox* have identical redundant data for a defined time. Thus, concerning that period, it can be compared for potential deviations.

Due to cost restraints within the automotive, we propose using a circular memory buffer, i.e., a first-in-first-out (FIFO) approach where old periods are overwritten by new periods. The downside is that data might be lacking beyond that

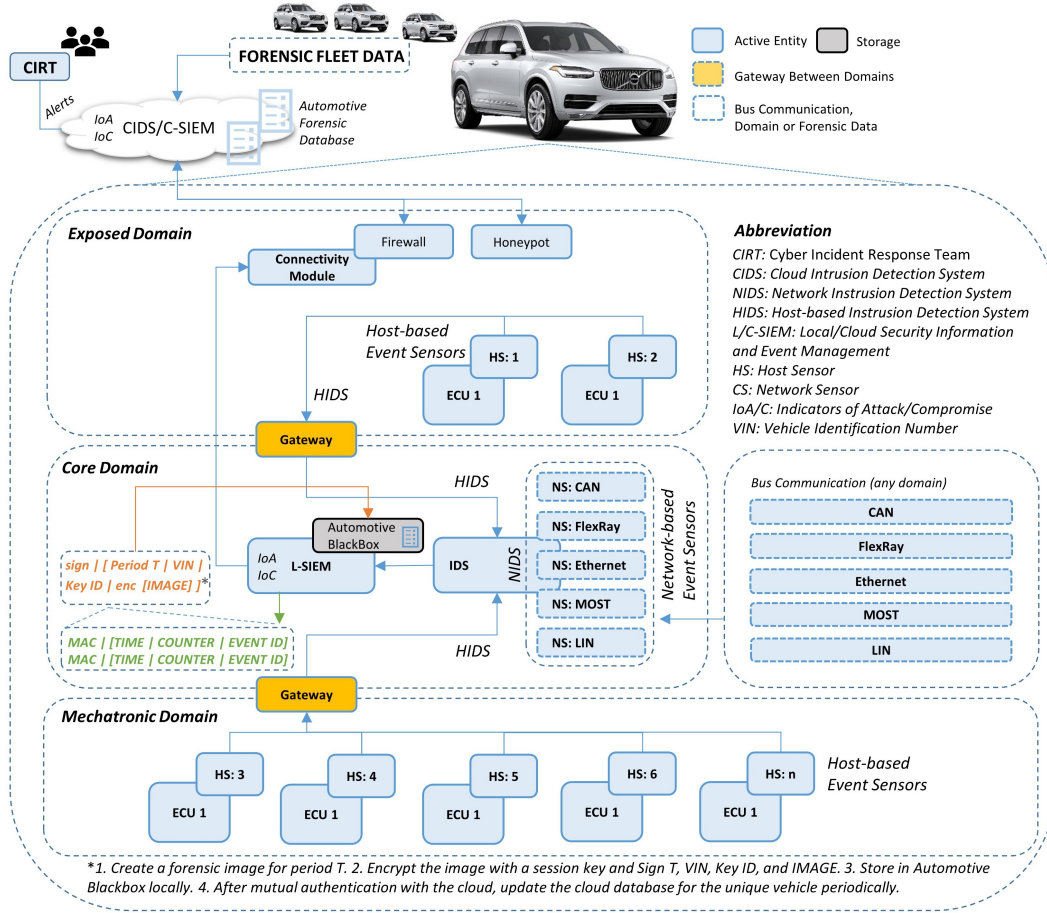


Fig. 2. The Automotive BlackBox within a centralized architecture context

period, for instance, due to a long time without connectivity. Additionally, as shown in Figure 2 we propose using a forensic honeypot, which attracts attackers to learn about their intentions and attack types to analyze, investigate and mitigate future attacks. Honeypots must be adapted regularly, e.g., via secure software updates [10], [11], to lower the risk that threat actors learn it's not a real system. We propose that relevant events are analyzed to acquire a status of the vehicle fleet's health, such as awareness of ongoing large-scale cyberattacks. We acknowledge the cost constraints within the automotive industry, and although beneficial, honeypots and similar solutions might not always be feasible. Also, note that our solutions cover mainly forensically relevant events from the ECUs and communication buses in the vehicle. However, specific synced data from external devices (cf. Table III in [3]) can contain relevant but privacy-sensitive data, such as messages and call logs. Our approach does not cover automated retrieval and analysis of such data, but our proposal can be extended by transferring it to the Automotive BlackBox and further to C-SIEM for processing. However, aligning with local laws and regulations concerning privacy-sensitive data is important. Moreover, we acknowledge that some ECUs might not be able to have a host-based sensor and still contain

relevant data. Still, our approach aims to automate the data collection and analysis process as much as possible to limit manual work.

Standard Data Format and Key Management. We propose the format as visualized in Figure 2, which contains the following attributes for each event. MAC is a key-based cryptographic hash over the rest of the event values. TIME, real or logical time, depends on the available source to synchronize time between different devices in the vehicle. COUNTER the number of occurrences of the same events under a predefined period. EVENT ID, the identification number of the actual event taking place.

L-SIEM creates an image for a predefined period of events, generates a symmetric key, and encrypts the image with this key. Identification data for a period T, VIN, and encryption key ID (not the actual key) is added to the image metadata, whereafter, a hash is calculated over the data and signed with an in-vehicle pre-stored certificate. The symmetric key used to encrypt the image is further encrypted with the public part of a pre-stored certificate in the vehicle and added to a key manifest along with the key ID. Key manifest is stored with encrypted images in the Automotive Blackbox, further periodically synchronized with and stored in the cloud. The private part of the certificate is securely stored and accessed

at C-SIEM to decrypt symmetric keys for further decryption of image files enabling automatic direct analysis by CIRT.

V. DISCUSSION AND FUTURE WORK

Infrastructure development differs in countries and locations, where cost and transfer speed can be challenging. There might be storage limitations in the vehicle, where a satisfactory storage size might be too costly. Low storage means that only a limited time can be saved in-vehicle. Constraints in connectivity, transfer speed, and cost might lead to that important data can be lost. As previously mentioned, there can be many distributed ECUs and sensors, and ensuring the authentic order/timing of events is challenging. Entities might suddenly stop generating alerts and must be detected. For instance, units can be disabled by hardware failures originating from malfunction or cyberattacks. However, having, e.g., a heartbeat signal from devices might not be possible due to performance restraints. From a fleet perspective, it is valid to be able to correlate time between events, for instance, speed, acceleration, and braking between involved vehicles, something that C-SIEM can automate. Enabling the collection of potential digital evidence and still adhering to privacy regulations is difficult. Data might reveal sensitive information about other individuals than intended via external communication or when correlating data. Anonymizing data and, at the same time, being able to connect it to individuals potentially involved in a crime is both contradictory and challenging.

Using AI, ML, and blockchain technology in automated data collection and analysis is promising for future research. Challenge iv (cf. Section II) aligning and revising different regulations and standards, i.e., privacy, forensics, security, and safety, emphasizing ADF is important, as studying the cost impact of new architectures. The chain of custody needs to be fulfilled to ensure forensic soundness. Trust in keys for encryption, signing, and MAC values is imperative to guarantee the CIANP properties. More work is needed to analyze potential attack vectors, e.g., vulnerabilities in key management, process isolation, and virtualization technologies such as trusted execution environments and containers.

VI. RELATED WORK

In 2006, NIST released SP 800-86, a document for practical guidance on performing computer and network forensics. SP 800-86 defines digital forensics as the science of identification, collection, examination, and analysis of data while preserving the integrity of the information and maintaining a strict chain of custody. The ISO 27037, yet another standard for digital forensics, was established in 2012 and further reviewed and confirmed in 2018. ISO 27037 provides digital evidence identification, collection, acquisition, and preservation guidelines. In 2004, NIST published SP 800-72 (PDA Forensics), and later in 2007, SP 800-101r1 (Mobile Device Forensics) provided guidelines for tool usage and procedures about PDAs and mobile devices. However, these documents are not automotive-specific, thus, do not provide satisfactory guidance within this area. In [3], we introduce the area of ADF. We perform

an extensive systematic literature review where we consider over 300 publications. We further group relevant papers into surveys, technical solutions, and focus categories. We also assess the cybersecurity aspect of the technical solutions by discussing and mapping them to cybersecurity attributes where applicable. Furthermore, we detail the type of forensically relevant data mentioned that was considered and how it needs to be secured. However, to the best of our knowledge, there is no previous work that extensively details goals and general requirements in an architectural context with the aim to guide ADF design with current and upcoming regulations in mind. Thus, our contributions are both novel and important.

VII. CONCLUSION

We have introduced the Automotive BlackBox, an architecture for automotive digital forensics, including components, standard data format, techniques, goals, and requirements. We have identified and highlighted challenges, such as the lack of existing regulations, standards, and common guidelines, and considered them when establishing our architecture. The identified goals are inspired by accepted digital forensics principles and have been further mapped to specific automotive requirements via a standardized approach for argument notation to ensure broad coverage. Furthermore, we have presented detailed guidelines, including a conceptual architectural description, key management, and data formats. Considering current and upcoming regulations, our contributions are useful in guiding the design of automotive and similar systems within a digital forensics context.

Acknowledgment. This research was supported by the CyReV project (2019-03071) funded by VINNOVA, the Swedish Governmental Agency for Innovation Systems.

REFERENCES

- [1] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno *et al.*, "Experimental security analysis of a modern automobile," in *2010 IEEE Symposium on Security and Privacy*, 2010, pp. 447–462.
- [2] C. Miller and C. Valasek, "Remote exploitation of an unaltered passenger vehicle," *Black Hat USA*, 2015.
- [3] K. Strandberg, N. Nowdehi, and T. Olovsson, "A systematic literature review on automotive digital forensics: Challenges, technical solutions and data collection," *IEEE Transactions on Intelligent Vehicles*, pp. 1–19, 2022.
- [4] Proton Technologies AG, "Complete guide to GDPR compliance," <https://gdpr.eu/>, 2020, accessed: 2023-03-27.
- [5] United Nations, "UN Regulation No. 155," 2022.
- [6] "ISO 26262:2011 Road Vehicles – Functional Safety," International Organization for Standardization (ISO), Standard, 2011.
- [7] Association of Chief Police Officers, "Acpo good practice guide for digital evidence," 2012.
- [8] K. Strandberg, T. Rosenstatter, R. Jolak, N. Nowdehi, and T. Olovsson, "Resilient shield: Reinforcing the resilience of vehicles against security threats," in *2021 IEEE 93th Vehicular Technology Conference*, 04 2021, pp. 1–7.
- [9] T. Kelly and R. Weaver, "The goal structuring notation—a safety argument notation," *Proc Dependable Syst Networks Workshop Assurance Cases*, 01 2004.
- [10] K. Strandberg, D. K. Oka, and T. Olovsson, "Unisuf: a unified software update framework for vehicles utilizing isolation techniques and trusted execution environments," in *19th escar Europe*, 2021.
- [11] K. Strandberg, U. Arnljung, T. Olovsson, and D. K. Oka, "Secure vehicle software updates: Requirements for a reference architecture," in *2023 IEEE 97th Vehicular Technology Conference*, 2023.