

Threat Analysis and Risk Assessment in Automotive Cyber Security

David Ward, Ireri Ibarra and Alastair Ruddle
MIRA Ltd

ABSTRACT

The process of hazard analysis and risk assessment (H&R or HARA) is well-established in standards and methods for functional safety, such as the automotive functional safety standard ISO 26262. Considering the parallel discipline of cyber security, it is necessary to establish an analogous process of threat analysis and risk assessment (T&R) in order to identify potential security attacks and the risk associated with these attacks if they were successful.

While functional safety H&R processes could be used for threat analysis, these methods need extension and adaptation to the cyber security domain. This paper describes how such a method has been developed based on the approach described in ISO 26262 and the related MISRA Safety Analysis Guidelines. In particular key differences are described in the understanding of the severity of a security attack, and the factors that contribute to the probability of a successful attack. However it also acknowledges that some threats may contribute to a safety-relevant hazard.

The paper will also explore some potential future directions, such as how the T&R can be used to support an assurance case for cyber security.

CITATION: Ward, D., Ibarra, I. and Ruddle, A., "Threat Analysis and Risk Assessment in Automotive Cyber Security," *SAE Int. J. Passeng. Cars – Electron. Electr. Syst.* 6(2):2013, doi:10.4271/2013-01-1415.

INTRODUCTION

Functional safety is defined in very general terms as the part of the overall safety of a system that depends on it operating correctly in response to its inputs. More specifically, the automotive standard for functional safety, ISO 26262 [1], is concerned with preventing hazards that may result from malfunctions of electronic systems. In ISO 26262, a “hazard” is defined as a potential source of harm, and the definition of “harm” is restricted to physical injury or damage to the health of a person.

In contrast in IEC 61508 [2] (the generic standard for functional safety which ISO 26262 claims to be derived from), harm is more widely defined and can include environmental or economic effects. This reflects the origins of IEC 61508 in the industrial process control sector, but is a reminder that in more general terms the definition of “harm” can be wider than just personal injury.

An area of increasing importance for the automotive domain is “cyber security”. Electronic control systems in vehicles now rely heavily on networked communications between individual electronic controllers. While historically an individual vehicle has operated in isolation, the possibility of communications with the in-vehicle network from

“outside” has become a reality. The legitimate motivations for such communications include

- “End of line” programming, where the software for a specific vehicle configuration is loaded (or “flashed”) into its electronic controllers at the end of the vehicle's construction;
- Service actions, such as reading diagnostic data through the on-board diagnostic (OBD) port and “reflashing” of software to apply updates in the field;
- The use of consumer devices (also known as “nomadic” devices) which typically interface with the vehicle's information and entertainment (“infotainment”) systems through a standardized wired (USB) or wireless (Bluetooth) connection;
- Vehicle-to-X (V2X) communications, where vehicles communicate with each other and with the transport infrastructure through a mobile *ad hoc* network.

Inevitably the existence of these external communications interfaces also creates the possibility for non-legitimate use, whether accidentally or by persons with nefarious motivations. This means it is necessary to secure the networks against unintended intrusion and use. Note that “security” requirements for road vehicles already exist but are

traditionally concerned with anti-theft provisions, e.g. [3]. Consequently, and also reflecting the wider association with IT security, this paper uses the term “cyber security” to refer to the security of vehicle networks and communications.

From a functional safety perspective, standards such as ISO 26262 describe how to design an electrical or electronic (E/E) system to provide adequate assurance that hazards will not be caused by failures in the system. A central tenet of systems safety engineering in this context is that it is not possible to completely eliminate all possible causes of hazards; instead system safety engineering seeks to determine that sufficient measures have been taken to reduce the risks associated with these hazards to an acceptable level (called “freedom from unacceptable risk” in ISO 26262). Risk is a function of both the probability of harm occurring and the severity associated with that harm. It therefore follows that hazards with higher associated potential harm will require more stringent controls for reducing the probability that the hazard occurs. In terms of E/E systems this is commonly addressed through safety integrity, and standards often have discrete levels of safety integrity such as SIL (IEC 61508) or ASIL (ISO 26262) in order to set targets on the systems. Both SIL and ASIL are associated with graded levels of increasing rigor in the specification, development and verification of safety-related E/E systems, with increased rigor applied to systems where there is a greater need for risk reduction through preventing the failures of the system that can lead to the hazard.

Functional safety standards have a process for evaluating the risk associated with a system and specifying the requirements to reduce that risk to an acceptable level. In outline the process is typically as shown in [Figure 1](#) below.




Figure 1. Typical functional safety process

The process of performing a hazard analysis and risk assessment (sometimes referred to as an H&R, or a HARA) is concerned with the first three steps of this process. It is well-established in standards and guidelines for functional safety such as IEC 61508 [2], ISO 26262 [1] and the MISRA Safety Analysis guidelines [4].

The standard IEC 15408 [5] is concerned with security evaluation for IT products, but does not explicitly address the possible safety implications of security breaches for safety-related control systems. A further limitation is that it does not provide a framework for risk analysis. Methods for evaluating the probability of a successful attack (described as “attack potential”) are described in IEC 18045 [6], but the severity of the impact is not evaluated to allow risk to be assessed. Risk analysis in an IT security context is outlined in [7] and described in more detail elsewhere (e.g. [8], [9]).

In IEC 15408 the concept of “evaluation assurance levels” (EAL) has a similar role for security considerations to the SIL and ASIL categories used in the functional safety context. The EALs are similarly associated with graded levels of increasing development rigor.

EXPANDING H&R TO T&R

In the second edition of IEC 61508 [2] published in 2010, there is reference to consideration of cyber security issues with regards to their potential impact on functional safety. However, cyber security threats that are not safety-related, such as those affecting privacy or financial transactions, are beyond the scope of IEC 61508.

In the European collaborative research project EVITA [10], a unified approach to risk analysis specifically for cyber security and security-related functional safety threats to in-vehicle networks was developed. The overall aim of the project was to prototype a toolkit of techniques and components to ensure the security of in-vehicle systems, including security hardware, software, and analysis methods, as well as an evaluation of the associated legal aspects. The EVITA project was focused on in-vehicle networks, as other projects have addressed issues relating to the security of V2X communications. However the scope of cyber security threats in EVITA considered those wider issues noted above that are outside the scope of IEC 61508 Edition 2.

The MISRA “Development Guidelines for Vehicle Based Software” [11] identified the need to protect vehicle software from unauthorized access that could compromise the performance of safety-related systems, as well as to provide detection of such tampering. As an historical observation, the concept of (safety) integrity levels introduced in this document was in part inspired by earlier work on security evaluation assurance levels.

The similarities between the EAL and SIL/ASIL concepts suggested the potential for developing a unified approach for automotive functional safety and cyber security [12]. Similar observations have also been made with regard to security and safety analysis in other fields, such as mobile *ad-hoc* networks [13] and defense applications [14].

Unifying the engineering processes for functional safety and cyber security may offer potential benefits in terms of reduced development costs through:

- Re-use of risk analysis for those applications where cyber security may also have possible functional safety implications;
- Deployment of common solutions to achieve both safety integrity and security integrity;
- Sharing of evidence in an assurance case.

While alignment is needed over all aspects of the lifecycle, the work in EVITA specifically concentrated on a unified approach to threat analysis and risk assessment [15]. The term “threat” is used generically to cover both functional safety hazards (which may result in physical harm), and cyber security breaches that result in other types of quantifiable losses.

In developing such a unified process it is necessary to consider what “risk” means in the context of cyber security, and specifically how to interpret the severity and probability factors associated with a threat.

Severity

The starting point for the severity classification is the categories given in ISO 26262. This defines three classes of severity (S1, S2 and S3), which are defined in terms of the estimated personal injury that could result from the hazard. A fourth class (S0) is reserved for hazards that could result in some physical damage but no significant personal injury. In practical use, the severity classes are often calibrated with respect to a defined scale such as the Abbreviated Injury Scale (AIS).

An important restriction of scope of ISO 26262 is that any technical failure resulting in a hazard can only affect a single vehicle. In contrast, a cyber security threat that has a potential functional safety consequence may affect multiple vehicles. It was therefore necessary to extend the definitions of the severity classes to include the potential consequences for multiple vehicles, and to include a fourth class, S4, for threats that have the potential to cause the most severe consequences in multiple vehicles. It should be noted that such a higher classification of severity already exists in the MISRA Safety Analysis guidelines and is reserved for cases involving multiple vehicles.

As well as considering functional safety related consequences of cyber security threats, it was also necessary to extend the definition of “severity” to cover non-safety consequences.

Breaches in the security of vehicle information or functions could lead to possible issues for stakeholders in four main areas:

- Privacy - unwanted or unauthorized acquisition of data relating to vehicle/driver activity, vehicle/driver identity data, or vehicle/sub-system design and implementation.
- Financial - unwanted or unauthorized commercial transactions.
- Operational - unwanted or unauthorized interference with on-board vehicle systems or V2X communications that may

impact on the operational performance of vehicles and/or intelligent transportation systems (ITS) (without affecting physical safety).

- Safety - unwanted or unauthorized interference with onboard vehicle systems or V2X communications that may impact on the safe operation of vehicles and/or ITS.

An important consequence of considering these four aspects collectively is that the severity of an attack, which is a measure of the impact of the harm for the stakeholders, becomes a “severity vector” with four components relating to safety, privacy, financial and operational threats. Each of these components may have different ratings, rather than a single severity parameter for all aspects. For example, it is possible that an attack has little or no impact on safety, but presents significant risks in terms of compromised driver privacy or loss of reputation for vehicle manufacturers. Thus, the individual components may translate to quite different relative risk levels, depending on the probability measures that are applied to assess the associated risk level.

Table 1 shows how the safety-related component of severity is classified. Entries in *italics* are extensions to the definitions of ISO 26262.

Table 1. Safety-related component of severity

Class	Safety-related severity
S0	No injuries
S1	Light or moderate injuries
S2	Severe and life-threatening injuries (survival probable) <i>Light or moderate injuries for multiple vehicles</i>
S3	Life threatening (survival uncertain) or fatal injuries <i>Severe injuries for multiple vehicles</i>
S4	<i>Life threatening or fatal injuries for multiple vehicles</i>

Tables 2, 3, 4 show how the non-safety-related components of severity are classified.

Table 2. Privacy-related component of severity

Class	Privacy-related severity
S0	No unauthorized access to data.
S1	Anonymous data only (neither specific driver nor vehicle data)
S2	Identification of vehicle or driver Anonymous data for multiple vehicles
S3	Driver or vehicle tracking Identification of driver or vehicle, for multiple vehicles
S4	Driver or vehicle tracking for multiple vehicles