

CPNI

Centre for the Protection
of National Infrastructure

CAE:

COMBINED APPROACH
TO DEVELOPING
SECURITY-INFORMED
SAFETY ASSURANCE



CONTENTS

1 Introduction	3
2 Signposting	4
3 Cyber security risk assessment process	5
4 Assurance cases	7
4.1 Requirements and policies layer (L0)	7
4.2 Architectural layer (L1)	7
4.2.1 Subclaims.....	9
4.2.2 Evidence	9
4.2.3 Implementation-derived requirements	10
4.3 Implementation layer (L2).....	10
5 Mapping cases to the risk assessment process	11
6 Acknowledgements	11

FIGURES

Figure 1: Location of this guide in the set of resources	4
Figure 2: General structure of architecture level case.....	8

TABLES

Table 1: Cyber security assessment summary.....	6
Table 2: Layers of assurance mapping.....	11

INTRODUCTION

This document provides generic guidance on developing security-informed safety assurance using a combined approach. The approach has two main components: the development of an engineering cyber security risk assessment process and a layered assurance case approach. Both components are mapped to each other and used together to achieve security-informed safety assurance.

This work supported the development of the CPNI 'Rail Code of Practice for Security-Informed Safety: A Good Practice Guide' and the BSI PAS 11281 [1] on connected automotive ecosystems, impact of security on safety. Additional detailed guidance can be found in these documents.

The detailed description of each of the components is provided in Sections 3 and 4 respectively. Section 5 shows the mapping between them to complete the picture of how these components are used together.

Further to this guidance, practical examples are provided to help focus, illustrate and communicate the approach in:

- 'Worked Example: Requirements and Policies Assurance Case'; and
- 'Worked Example: Architecture and Implementation Assurance Case'.

The approach includes a generic systems-driven risk assessment approach informed by component level analyses. As with all fields, risk assessment is evolving, some recent cyber-related perspectives can be found in NCSC guidance [2].

[1] PAS 11281:2018, Connected automotive ecosystems – Impact of security on safety – Code of practice

[2] NCSC: Introducing component-driven and system-driven risk assessments, Version 1.0, December 2017

02. SIGNPOSTING

This is the first detailed generic guide in CPNI's resources for security-informed safety assurance. Figure 1 below shows its location in the set of guides (highlighted in red).




Figure 1: Location of this guide in the set of resources

03.

CYBER SECURITY RISK ASSESSMENT PROCESS

The first component of the approach is the development of a generic cyber security risk assessment process. This document provides a high-level overview of the process, more detailed information is available in 'Risk Assessment Process' guidance.

The methodology has a defined series of steps developed from those used in the standard information assurance approach. These steps are set out in Table 1 below.

Step	Brief description
Step 1 – Establish system context and scope of assessment	Describe the system to be assessed and its relationship with other systems and the environment. Identify the services provided by the system and the system assets. Agree on the scope of and motivation for the assessment and identify the stakeholders and their communication needs. Identify the type of decisions being supported by the assessment.
Step 2 – Configure risk assessment	<p>Identify any existing analyses, e.g. safety cases or business continuity assessments that provide details of the system, the impact of failure and the mitigations that are in place. Characterise the maturity of the systems or project and the key uncertainties.</p> <p>Ensure that the risk assessment is focused on the kinds of threats that are of concern. Define possible threat sources and identify potential threat scenarios. Refine generic capability and impact levels for the systems being assessed. Identify risk criteria.</p> <p>Refine and focus system models in the light of the threat scenarios and existing analyses to ensure that they are at the right level of detail for an effective security-informed risk analysis.</p>
Step 3 – Analyse policy interactions	Undertake an analysis of policy issues considering interactions between safety requirements and security policies. Resolve any conflicts, show that the trade-offs are satisfactory and document the decisions made.
Step 4 – Preliminary risk analysis	Undertake architecture-based risk analysis, identifying potential hazards and consequences and relevant vulnerabilities and causes together with any intrinsic mitigations and controls. Consider doubts and uncertainties, data and evidence needs. Identify intrinsic and engineered defence in depth and resilience.
Step 5 – Identify specific attack scenarios	Refine preliminary risk analysis to identify specific attack scenarios. Focus on large consequence events and differences with respect to the existing system.

Step 6 – Focused risk analysis	Prioritise attack scenarios according to the capabilities required and the potential consequences of the attack. As with the previous step, the focus is on large consequence events and differences with respect to the existing system.
Step 7 – Finalise risk assessment	Finalise risk assessment by reviewing implications and options arising from focused risk analysis. Review defence in depth and undertake sensitivity and uncertainty analysis. Consider whether the design threat assumptions are appropriate. Identify additional mitigations and controls.
Step 8 – Report results	Report the results of the risk assessment to stakeholders at the appropriate level of detail.

Table 1: Cyber security assessment summary

The main differences between this risk assessment process and the standard approach are summarised below:

- The approach to threat assessment (Step 2) is slightly different. Without access to intelligence data, it is not possible to assess the actual threat, but it is still useful to identify potential threat scenarios in order to ensure that the risk assessment is focused on the kinds of threats that are of concern.
- Similarly, when it comes to prioritising risk (Step 6), it is not possible to judge the likelihood of an attack from a particular threat source without access to intelligence data, but we can assess the capabilities and level of access to the system that a threat agent would need in order to launch a successful attack. Thus, the attack scenarios are ranked according to required capabilities and potential impact rather than likelihood and impact.
- Step 3 has been introduced to explicitly look at the safety and security interactions at a policy and requirements level.
- In Step 4 an architecture-based approach is used (similar to a conventional hazard analysis [3] or failure modes and effects analysis [4]) where consideration is given to cyber attacks on individual subsystems and the impact of loss of integrity and availability of the subsystem on the overall service.

- In Step 6, the resilience of the system to such service failures has to be taken into account when assessing the consequential impact.

A security-informed assessment of the safety risks has to recognise the problem that the frequency of attacks is unknown and changing. It also needs to be structured so that the intelligence assessment that estimates such likelihoods can be supplied by those with the necessary intelligence or authority to make such judgements. The process described therefore parameterises the risk assessment on the capabilities that an attacker would need in order to achieve a safety impact failure. Even without access to intelligence data it is still useful to identify potential threat scenarios in order to ensure that the risk assessment is focused on the kinds of threats that are of concern. We considered the capabilities of potential threat sources and identified a range of capability levels of potential threat sources adapted from UK guidance on technical risk assessment [5].

[3] IEC61882:2002 Hazard and operability studies (HAZOP studies) - Application Guide 2002

[4] ESA, Failure Modes, Effects and Criticality Analysis (FMECA). D. European Space Agency. ECSS-Q-30-02A, 1991

[5] CESG, HMG Information Assurance Standard No 1 and 2 Supplement, Technical Risk Assessment and Risk Treatment, Issue 1.0, April 2012, https://en.wikipedia.org/wiki/HMG_Infosec_Standard_No.1

4.1 REQUIREMENTS AND POLICIES LAYER (L0)

The second component of the approach is the use of structured assurance cases for communicating and building confidence in the safety and security properties of the system. Structured assurance cases are used in a wide range of industrial domains, but the practice set out in this guidance is based on a concept of Claims, Arguments and Evidence (CAE), which can be related to the approach developed by Toulmin[6]. CAE supports the description of how sophisticated engineering arguments are actually made: the key elements of CAE are described in 'CAE One Page Mini-Guide' with information on CAE Blocks available in 'CAE Blocks and Connection Rules'.

In addition to CAE and CAE Blocks we need a structuring mechanism for dealing with the complexity of real cases. The idea of compositionality and layered assurance was raised by Rushby and DeLong [7] and has been adopted as an approach by this guidance. The goals of the approach are to build assured systems from compositions of previously assured components, while being able to derive the system level properties (e.g. safety and security) systematically from the properties of the components. Abstraction is one of the key structuring mechanisms for layered assurance, with three levels of abstraction used when creating security-informed safety cases. The layers, or layers of assurance, described below can be applied recursively:

- Requirements and policies layer (L0) – the highest level of abstraction where the system represents its requirements, and defines safety and security policies and their interaction;
- Architectural layer (L1) – the intermediate level where the abstract system components and architecture are analysed; and
- Implementation layer (L2) – the detailed level where the implementation of specific components and their integration within the specific system architecture are scrutinised.

A brief description of each abstraction level is provided below.

4.2 ARCHITECTURAL LAYER (L1)

Systems are built from components within an architecture. The architecture specifies the way in which components are connected as well as the interfaces available and access methods and protocols used throughout the system.

The components and the architecture play equally important roles in achieving the objectives and enforcing critical system properties. To evaluate the whole system assurance, the contribution of both the architecture and system components needs to be considered.

At the L1 level the general system architecture should be analysed to identify the various parts that contribute to achieving the critical properties of the system. Different components can have a greater or lesser contribution. At this level the focus should be on the major parts that have the greatest effect on the critical system properties.

The general engineering process analyses the importance of components by using both a top-down approach, mapping the critical properties to specific components, and a bottom-up approach, analysing the failure modes of the components. Such analysis is covered in various domain-specific standards and is independent of the security-informed analysis.

In addition to the criticalities defined by the engineering approach and derived from the conventional analysis at L0, there is also a need to analyse the security-informed aspects of the system. This requires investigating the failure behaviour of the system, examining the trust relationship between its components, and considering whether any new controls need to be introduced or any of the existing criticalities of components adjusted.

Therefore, at the L1 level the critical properties need to be revisited and an analysis conducted to identify any derived properties, assess which components play important roles in enforcing them, and identify any areas that might need additional security controls.

This is done on the basis of a security-informed hazard analysis in a hazard and operability study (HAZOP). It is a widely used approach to hazard analysis in a variety of industries, including chemical, nuclear and railway. The central activity of a HAZOP is to identify the hazards posed to a system by examining some abstract, typically architecture-based, representation of the system. In considering security aspects of safety there are a number of ways of enhancing the safety HAZOP by a combination of desktop review and focused workshop. This security-informed HAZOP provides an opportunity for a structured and informed discussion about the security risks associated

with the system. 'Security-informed HAZOP' provides details on how to plan and conduct such a study. There is a wide range of other security and safety analysis techniques but a security-informed HAZOP has been selected as it has been used in a number of actual projects to leverage the safety analysis from a security perspective.


At the end of the study, when the full range of attack capabilities is considered, the hazards and system vulnerabilities are identified. The identified potential attack scenarios are summarised and linked to the hazards where the attacks are graded according to the capability level required to implement it. These are evaluated for their potential impact and linked to critical recommendations and assurance activities in order to mitigate the risk. This can be done to a certain capability level that is decided prior to the study. Combining the security-informed HAZOP and the penetration testing gives the possibility of an informed assessment of the security of the system with a final set of recommendations.

In order to construct a case, at the L1 level, the following things need to be taken into account:

- the output from the L0 level of abstraction;
- the main identified and revised critical safety and security properties of the system;
- components that play essential roles in enforcing the critical properties;
- a high-level architecture of the system representing components and their interaction; and
- the dynamic aspect to consider possible changes to the system in the future.

The general structure of the case constructed at this level of abstraction is shown in Figure 2. Note that the term 'components' is used in a general sense at this level of abstraction and can mean service, procedure, etc.

Figure 2: General structure of architecture level case



The purple blocks with a double rectangle outline shown in Figure 2 indicate that the claims about the critical properties of the system are enforced by various components within the system architecture. At the next level of abstraction it is necessary to look into the design and implementation of those components to ensure they are effective in their role as determined by the architecture (see 4.2.1) and provide the required evidence that the claims are satisfied (see 4.2.2).

The top-level claim is that the L1 system architecture should be able to meet the requirements and constraints established in L0, the 'critical properties' in Figure 2. From a security perspective, the architecture should include features to minimise service malfunctions due to accidental and deliberate threats (see 4.2.3).

4.2.1 SUBCLAIMS

In developing a specific case, the top-level claim would typically be split into subclaims related to different service attributes, for example:

- the functions assigned to the system architecture will provide the specified service function;
- the system design is capable of meeting the service-level agreement (SLA) targets, i.e. the non-functional properties;
- the system design can withstand attacks via service interfaces, the external environment and on internal components up to the specified capability level; and
- the other functions or properties the system should have to enable other systems to be secure.

In applying these concepts, the decomposition of claims into subclaims should be justified. The role of side-claims in achieving this is discussed in 'CAE Blocks and Connection Rules Guide'.

4.2.2 EVIDENCE

As described above, it is necessary to ensure that there is evidence that the claims are satisfied.

4.2.2.1 DEFINITION OF THE SYSTEM ARCHITECTURE

This defines the main system components:

- system components;
- system interfaces;
- defences against external threats, via service interfaces, resources, external dependencies; and
- defences against internal threats, via personnel, compromised components.

4.2.2.2 FUNCTIONAL ASSIGNMENT

Evidence of the functions assigned to the system includes:

- functions assigned to components;
- traceability of functions to service requirements; and
- verification that there are no additional service functions.

4.2.2.3 SECURITY-INFORMED HAZARD ANALYSIS

A security-informed hazard analysis is a systematic review of the system architecture in order to identify:

- potential hazardous failures at the system output interfaces;
- potential causes of hazardous failures (for a security-informed hazard analysis, this will include deliberate attacks);
- mitigations within the system to prevent hazardous failures; and
- recommendations for additional mitigations against the causes of hazardous failures (including deliberate attacks).

A possible means of performing a security-informed hazard analysis can be found in 'Security-informed HAZOP'.

4.2.2.4 SERVICE AVAILABILITY ANALYSIS

The SLA for the service needs to be translated into requirements for the system. The analysis should typically be based on an assumption of random component failure. The analysis should take into account:

- component reliability targets;
- internal redundancy;
- failure detection functions within the system;

- assumptions of failure independence between components; and
- availability assumptions for external resources.

4.2.2.5 PERFORMANCE ANALYSIS

Service capacity and throughput requirements have to be correctly translated into performance requirements for system components. The analysis should show that the component performance capabilities can satisfy the service performance requirements.

4.2.3 IMPLEMENTATION-DERIVED REQUIREMENTS

The consideration of security-informed safety at the architectural level will place requirements on system implementation including:

- component reliability targets;
- component performance targets;
- component security targets;
- component segregation and independence requirements; and
- additional design and procedural requirements to reduce the likelihood of hazardous failures.

4.3 IMPLEMENTATION LAYER (L2)

This level moves down from the abstract components and architecture to their specific implementations. It is important to make sure that the implemented components within the specific system architecture really enforce the critical properties of the system. Even though a system may have a strong architectural design which looks very convincing at the L1 level, the specific implementation may break that abstract architecture, introduce new unwanted properties or not provide the implementation of some critical functions.

A detailed CAE structure should be developed, capturing the implementation of the specific components. This completes the security-informed safety case in the sense that arguments and evidence are provided to support all the claims made about the components enforcing critical system properties.

Developing the implementation-level case involves:

- using the output from the L1 level of abstraction;
- analysing the implementation details of every critical component;
- creating an argument structure and elaborating the evidence to show that all the critical properties of the system are enforced; and
- documenting the results and providing traceability to the appropriate L0 and L1 security-informed safety case elements.

The case created at the L2 level of abstraction is based on two types of technical information:

- General technical information produced and supplied with the components as part of the normal development process. This information should provide evidence that components have been implemented to specification and implement their required attributes.
- Context-specific technical details derived from the analysis of the specific system implementation, and the implementation and integration of the components in a particular context.

The case completed at this level should provide evidence that the implementations of the components integrated into a system according to the specific architecture really enforce the critical properties of the system, and do not introduce any additional properties that are considered unwanted for the system.

05.

MAPPING CASES TO THE RISK ASSESSMENT PROCESS

One advantage of the assurance case approach, and indeed a requirement of it, is that it can be mapped to a variety of project-specific processes as well as to a variety of different stages of the project: some might be at a procurement stage, others might be adapting legacy systems and yet others might be novel future systems. For example, the development of a security-informed safety case could be linked to specific engineering and assurance processes such as:

- system and product development lifecycles;
- specific lifecycles for safety and security, reliability, availability, maintainability, and safety; and
- risk management lifecycles.

The mapping between the layers of assurance approach to creating security-informed safety cases and the cyber security risk assessment process is shown in Table 2.

Step of the cyber risk assessment process	Role of layers L0, L1, L2
Step 1 – Establish system context and scope of assessment	Addressed at L0
Step 2 – Configure risk assessment	Addressed at L0 Reviewed at L1 and L2 for additional relevant detail
Step 3 – Analyse policy interactions	Addressed at L0
Step 4 – Preliminary risk analysis	Architecture-based assessment at L1 and refined at L2
Step 5 – Identify specific attack scenarios	Architecture-based assessment at L1 and refined at L2 (initial threat scenarios part of Step 2 and L0)
Step 6 – Focused risk analysis	L1 and refined for L2
Step 7 – Finalise risk assessment	L1 and refined for L2
Step 8 – Report results	This is undertaken progressively with CAE providing core of documentation.

Table 2: Layers of assurance mapping

A checklist of policy and requirements issues that should be considered has been produced and is available in the 'Risk Assessment Process Guide'. Some of these will be resolved at the requirements and policies level but others will set policies and constraints that shape the development of the case at the architectural and implementation layers.

06.

ACKNOWLEDGEMENTS

This document is based on material developed in previous CPNI projects and published research by Adelard.

Disclaimer

This guide has been prepared by CPNI and is intended to support the implementation of security-informed safety assurance and the Claims, Arguments and Evidence (CAE) methodology. This document is provided on an information basis only, and whilst CPNI has used all reasonable care in producing it, CPNI provides no warranty as to its accuracy or completeness.

To the fullest extent permitted by law, CPNI accepts no liability whatsoever for any expense, liability, loss, damage, claim or proceedings incurred or arising as a result of any error or omission in the report or arising from any person acting, refraining from acting, relying upon or otherwise using the [report]. You should make your own judgment with regard to the use of this document and seek independent professional advice on your particular circumstances.

No Endorsement

Reference to any specific commercial product, process or service by trade name, trademark, manufacturer or otherwise, does not constitute or imply its endorsement, recommendation or favour by CPNI. The views and opinions of authors expressed within this document shall not be used for advertising or product endorsement purposes.

© Crown Copyright 2022.

You may use or reuse this content without prior permission but must adhere to and accept the terms of the Open Government Licence for public sector information. You must acknowledge CPNI the source of the content and include a link to the Open Government Licence wherever possible. Authorisation to reproduce a third party's copyright material must be obtained from the copyright holders concerned.