

CPNI

Centre for the Protection
of National Infrastructure



WORKED EXAMPLE:
ARCHITECTURE AND
IMPLEMENTATION
ASSURANCE CASE

CONTENTS

- 1 Introduction3
- 2 Signposting.....4
- 3 Case study analysis 5
 - 3.1 Step 4 – Preliminary risk analysis 5
 - 3.2 Step 5 – Identify specific attack scenarios 5
 - 3.3 Step 6 – Focused risk analysis 5
 - 3.4 Step 7 – Finalise risk assessment 6
 - 3.5 Assurance case at architecture and implementation layers 6
 - 3.6 Summary of the analysis 17
- 4 Discussion..... 18
- 5 Acknowledgements 18

FIGURES

- Figure 1: Location of this guide in the set of resources4
- Figure 2: Decomposition by CIA attributes7
- Figure 3: Discussion of system availability 8
- Figure 4: Discussion of data integrity at different stages..... 9
- Figure 5: Data in use integrity considerations 10
- Figure 6: Data in motion integrity considerations11
- Figure 7: Data at rest considerations 12
- Figure 8: Discussion of system confidentiality 13
- Figure 9: Future-related branch of the case with various types of changes considered 14

TABLES

- Table 1: Analysis of possible attack interfaces – example table 6
- Table 2: Analysis and evidence requirements for the architectural layer – examples 16
- Table 3: Summary of steps 4 – 7 of the analysis 18

INTRODUCTION

This document provides a worked example of an assurance case analysing the architecture and implementation details. It can be used as a practical guide to illustrate the second part of the process of developing security-informed safety cases using a combined approach as set out in ‘Combined Approach to Developing Security-Informed Safety Assurance’.

As prerequisites, this guidance relies on good knowledge of the Claims, Arguments and Evidence (CAE) concepts and their application to the development of assurance cases. It also assumes the reader is familiar with the other guides available on the CAE approach available on the CPNI website.

The focus of this guidance is on the architecture and implementation layers only (L1 and L2) described in ‘Combined Approach to Developing Security-Informed Safety Assurance’. The preceding requirements and policies layer (L0) is the focus of ‘Worked Example: Requirements and Policies Assurance Case’.

The development of the architecture and implementation parts of assurance cases are done by working through the following four steps of the cyber security risk assessment process:

Step 4 – Preliminary risk analysis

Step 5 – Identify specific attack scenarios

Step 6 – Focused risk analysis

Step 7 – Finalise risk assessment

The combination of the layered assurance and cyber security risk assessment process facilitates a thorough analysis of the system, helping to develop a good understanding of the technical detail and identify issues that may need to be addressed.

There is one last step of the risk assessment process:

Step 8 – Report results

This task is undertaken progressively when developing CAE rather than at the end of the process as this helps to provide the core of documentation.

The guidance document illustrates the practical application of the approach to a case study of a transport advisory system. The example has been anonymised to ensure confidentiality of the actual system and organisation.

The guidance is aimed at experienced practitioners who have understanding of safety and security aspects but would like to see a real-world example on how the security-informed safety cases are developed in practice.

02. SIGNPOSTING

This is the second example-based guide in our stack of resources for security-informed safety assurance. Figure 1 below shows its location in the set of guides (highlighted in red).

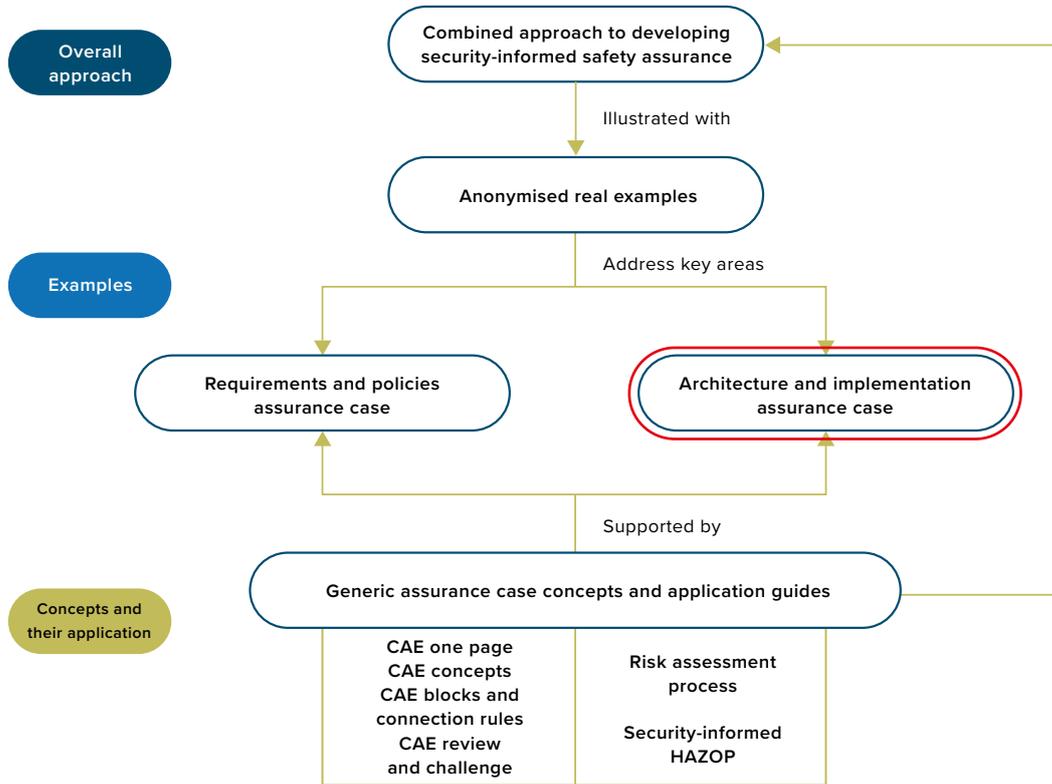


Figure 1: Location of this guide in the set of resources

CASE STUDY ANALYSIS

In this section the application of the last four steps of the risk assessment methodology and the construction of the architecture and implementation cases are discussed in detail.

3.1 STEP 4 – PRELIMINARY RISK ANALYSIS

In order to conduct a risk assessment on the transport advisory system (TRAS), it was necessary to involve a multi-disciplinary team of stakeholders. For this, an appropriate model of the system also had to be generated based on the description of the architecture, requirements and uses (see ‘Worked Example: Requirements and Policies Assurance Case’). The model was used at a workshop that was held to analyse potential attacks via external and internal network connections to the TRAS database. The workshop took the form of a security-informed Hazop attended by experts from the transport industry, system developers and hosting service providers. A briefing note was distributed prior to the workshop based on the guidance contained in the ‘Security-informed Hazop’ guide.

During the workshop, each element of the system was reviewed systematically, using a set of guidewords to prompt the experts to identify potential hazards. The experts were asked to identify:

- causes of a potential malfunction;
- potential consequences of the malfunction;
- any system features that can detect or mitigate the malfunction; and
- any follow-up activities.

The goal of the Hazop was to identify potential attacks on the advisory system that could lead to potential disruptions or hazardous situations, i.e. violation of the Material Safety Uses (MSU), and to suggest some additional controls and assurance activities that would provide confidence that the system was protected against such attacks.

The security hazard analysis indicated the following hazard classes on the TRAS boundary:

- No connection;
- Wrong data; and
- Corrupted data.

If those hazards were realised, a potential dangerous situation could occur, e.g. a vehicle not safe for use could be released to traffic, which could lead to a collision. Another severe consequence (albeit not safety-relevant) could be a severe disruption of operation.

Considering the criticality of the system and the potential consequences on operation and safety, an attack Capability C is initially considered for the recommendation of controls.

3.2 STEP 5 – IDENTIFY SPECIFIC ATTACK SCENARIOS

Proceeding from the preliminary risk analysis, specific attack scenarios were brainstormed and captured during the security-informed Hazop workshop. The focus of the discussion was on TRAS events with critical consequences, e.g. attacks that may cause various safety issues or major disruption to the TRAS service. Any new functionality within TRAS which may enable additional attacks different from the attacks on the systems that are currently in place were also given a special consideration.

Details of the attacks are omitted from this document, but addressed service users themselves, social engineering, web interfaces, specially crafted attacks using features of the systems as well as brute force attacks on the service-hosting organisation.

These scenarios were considered in more detail at the next step of the analysis. Other specific attack scenarios will be identified during the penetration testing that is to be conducted for the TRAS system. The new attack scenarios from the penetration testing along with the results and potential weaknesses of the tests will be captured and fed back into the hazard analysis and the assessment of the credibility of attacks for the attacker capability of concern.

3.3 STEP 6 – FOCUSED RISK ANALYSIS

The attack scenarios identified in the previous step of the analysis were prioritised 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.

The summary of the scope of impact of each potential attack, the related hazards and classification of the required attack capability to achieve the hazard were summarised in a table along with the recommendations, to protect against such attacks. In addition the wider safety impact of the attacks on selected MSUs was defined. The details are omitted from this document for confidentiality reasons.

3.4 STEP 7 – FINALISE RISK ASSESSMENT

The risk assessment process was finalised based on the results of the security-informed Hazop workshop. The analysis of each of the potential attack interfaces examined during the workshop is captured in a table such as Table 1, along with the additional mitigations and controls identified and recommended for future implementation.

Interface	Analysis	Recommendations
Reference to interface	Analysis of attacks to interfaces (omitted for confidentiality reasons)	Examples might be: R: TRAS disaster recovery needs to cope with progressive pollution attack. R: Recovery plans / procedures need to be regularly exercised.

Table 1: Analysis of possible attack interfaces – example table

Additionally, the 20 critical controls provided in a document produced by the Center for Internet Security [1] were analysed with respect to the TRAS system and any gaps noted (the findings are excluded from this document).

3.5 ASSURANCE CASE AT ARCHITECTURE AND IMPLEMENTATION LAYERS

Steps 4 – 7 cover both the architecture and implementation levels of the security-informed safety case.

At the architecture level the focus is on the risk analysis, building on the existing safety and business continuity analyses that may be available.

At the architectural layer, L1, the components and the architecture of the system, which play important roles in achieving system objectives and enforcing the critical properties of the system, are analysed. To address security considerations, various methods of security can be applied at this stage, for example, a guideword-based approach derived from the safety Hazop analysis was used in this

case study. The primary focus was on confidentiality, integrity and availability security attributes (CIA). The overall approach and the guidewords that helped to identify security-related issues are described in the ‘Security-informed Hazop’ guidance.

The threats and controls captured during the security-informed Hazop analysis were fed into the security-informed safety case developed at this stage of analysis.

The L1 case is started with a top level claim: ‘System components and architecture guarantee satisfaction of safety and security requirements’. This claim is factored into two subclaims, one about whether the safety and security requirements are satisfied by the architecture and its components now, and one about whether the requirements will continue to be satisfied in the future.

As in this case the main interest is in CIA, the claim that the requirements are satisfied is now decomposed into sub-claims about availability, integrity, confidentiality.

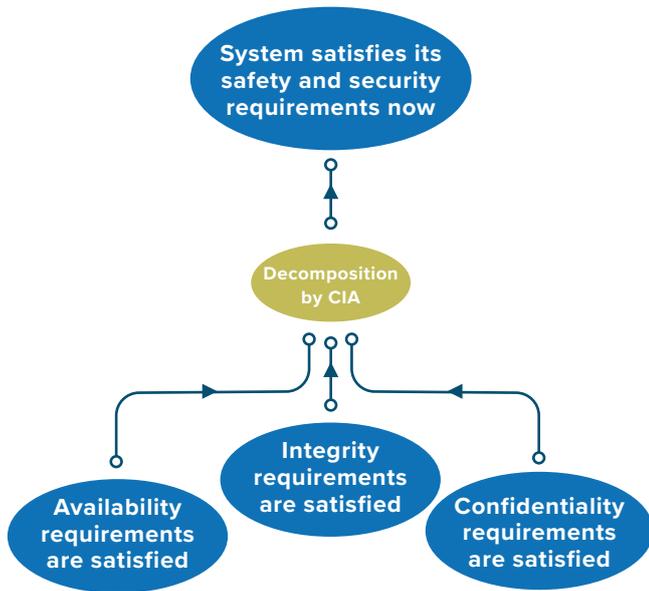


Figure 2: Decomposition by CIA attributes

Note that the blue ovals represent claims and the rectangular box the argument step that links lower level claims to the top claim. For a summary of the notation and concepts see the ‘CAE One Page Mini-Guide’.

The availability requirement should first be concreted into a claim with a specific availability value. For this, the requirement has to be analysed from both the safety and operational perspective. Discussion with engineers confirmed that the safety requirement is less strict than the operational one in this particular system, therefore the availability value from the service level agreement (SLA) was used in the concreted claim.

The different types of components contributing to the overall system availability then need to be considered. The concreted claim is therefore expanded into sub-claims concerning data availability, availability of software, hardware and channels, and their interaction (correct configuration). It is possible to progress down the data availability by looking into the types of data protection employed in the system: data storage redundancy and backups, as well as the solution of how the data can be reached/retrieved in the event of a failure.

These considerations, supported by some of the evidence, are illustrated in Figure 3.

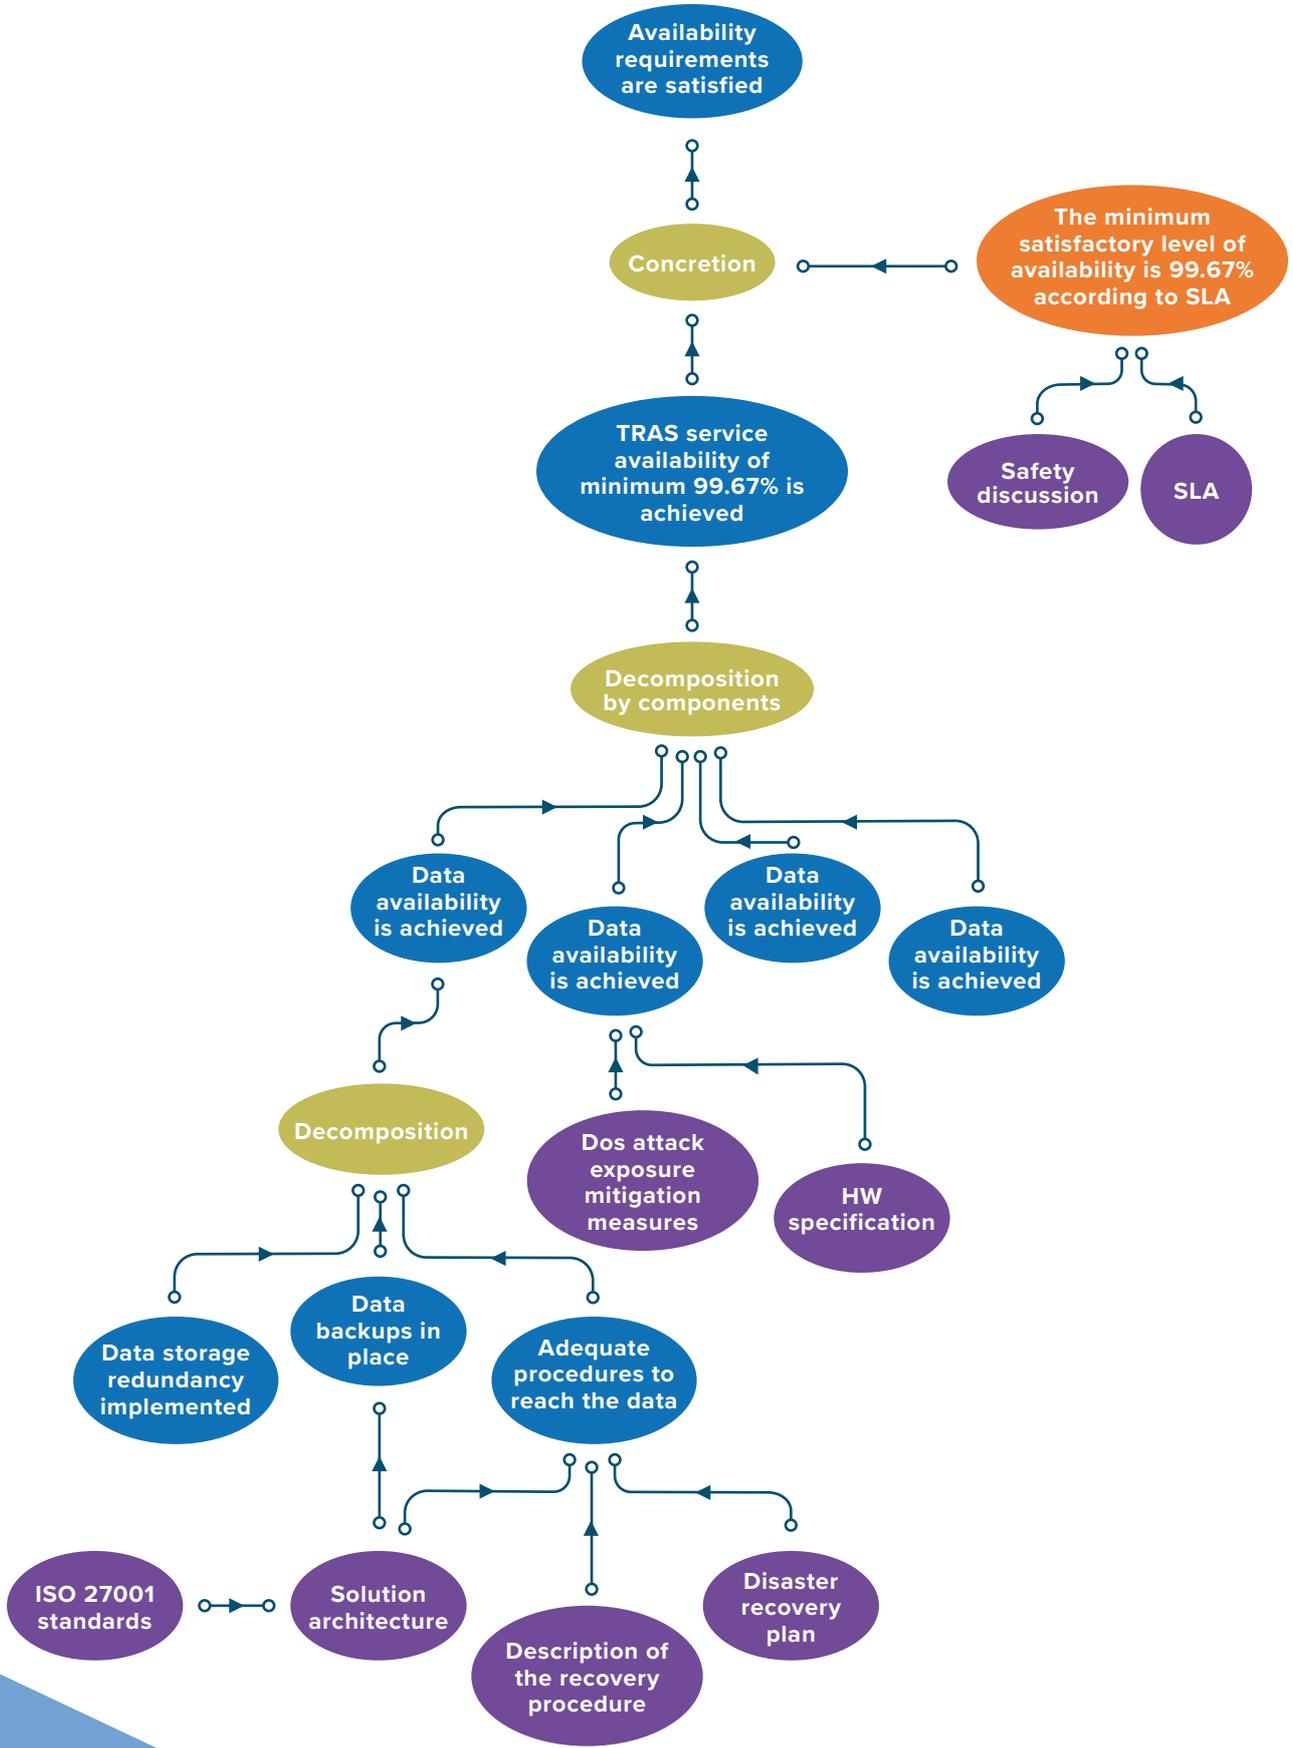


Figure 3: Discussion of system availability

03.

The integrity-related branch starts with a more precise definition of integrity: accuracy and consistency of the TRAS data. It is important to ensure that the data initially transferred to TRAS from the previous systems are accurate, not corrupted and not modified. The transition schedule, migration report and all the related evidence showing the transition is performed successfully are to be provided by the system developers. The claim that the data are accurate and consistent during the lifecycle is essential and requires a detailed consideration. During the lifecycle, data can be stored in several persistent or non-persistent locations, as well as be in transit between various elements of the system. Classically, three states of digital data are distinguished: data at rest, data in motion, and data in use. In terms of TRAS, these states will be defined as follows:

- 'data in use' refers to active data in a non-persistent state on a user's machine or in a thirdparty tool while the data are being created or modified
- 'data in motion' is a state in which data are being transferred to and from the TRAS database
- 'data at rest' refers to the data stored in the TRAS database in the inactive state (when they are not being used by anyone)

It is necessary to analyse each of these states, therefore, the claim about the data being accurate and consistent during the lifecycle of the system is factored into three subclaims, each for the specific data state. The top part of the integrity-related branch is shown in Figure 4.

Data in use includes data being created/updated by TRAS users or applications accessing TRAS via various endpoints, as well as data being processed by the TRAS itself. Data in this state is susceptible to various types of threats. The most vulnerable points are at the endpoints where users and applications can access and interact with the data. It is important to ensure that all the users manipulating the TRAS data are knowledgeable and trusted. An additional control-like validation of any user input before the submission is also good to have. Any automatic data updates made by other applications need to be correct and trusted as well. And of course, there should be no interference to the input by any third party or malicious software. For this, user devices should be protected and strong user authentication, identity management and permissions control implemented. Some of these considerations and controls are illustrated in Figure 5.

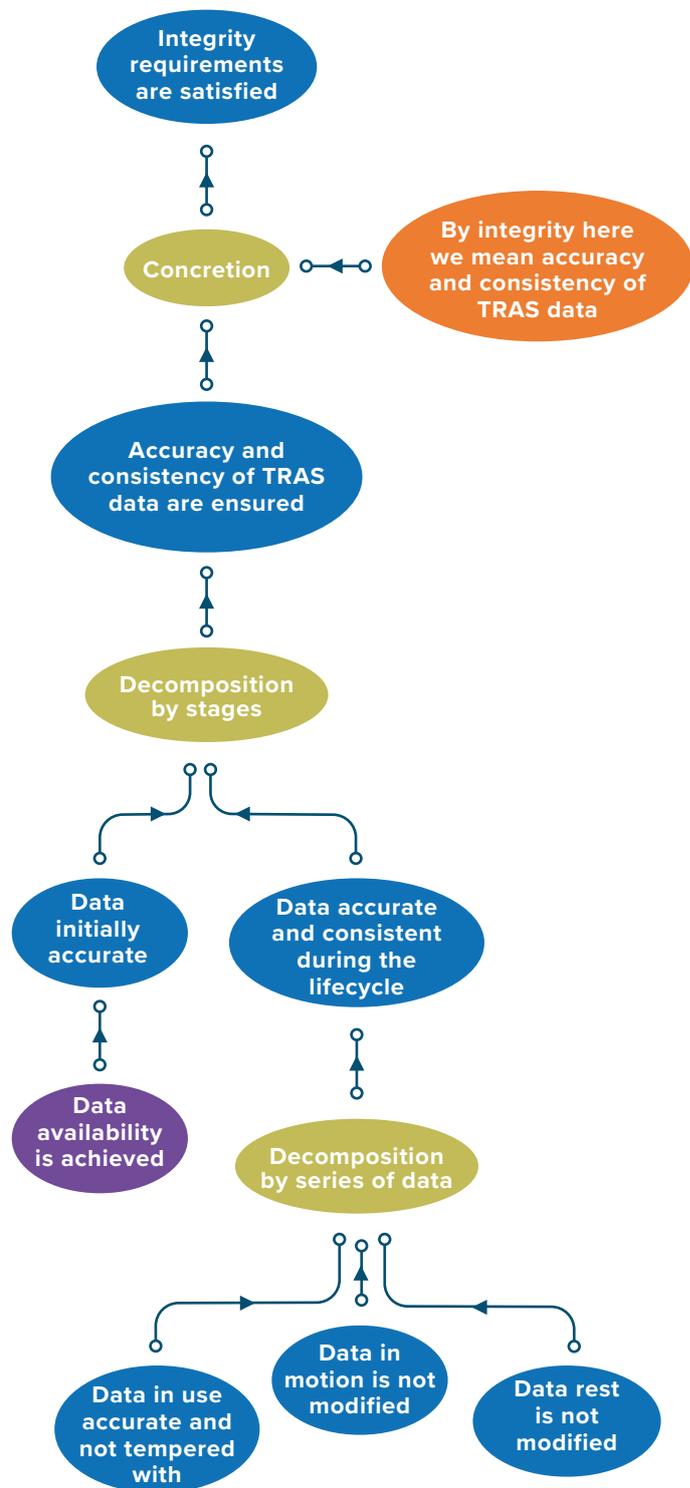


Figure 4: Discussion of data integrity at different stages

Data in motion is the next state to consider. To ensure the data is not modified while travelling to or from the database it is necessary to check that links and interfaces are secure and are correctly implemented by the TRAS software code. It is also necessary to make sure configurations for network devices such as firewalls, routers and switches are secure and no malicious code is in place on the way to the database. This discussion is illustrated by Figure 6.

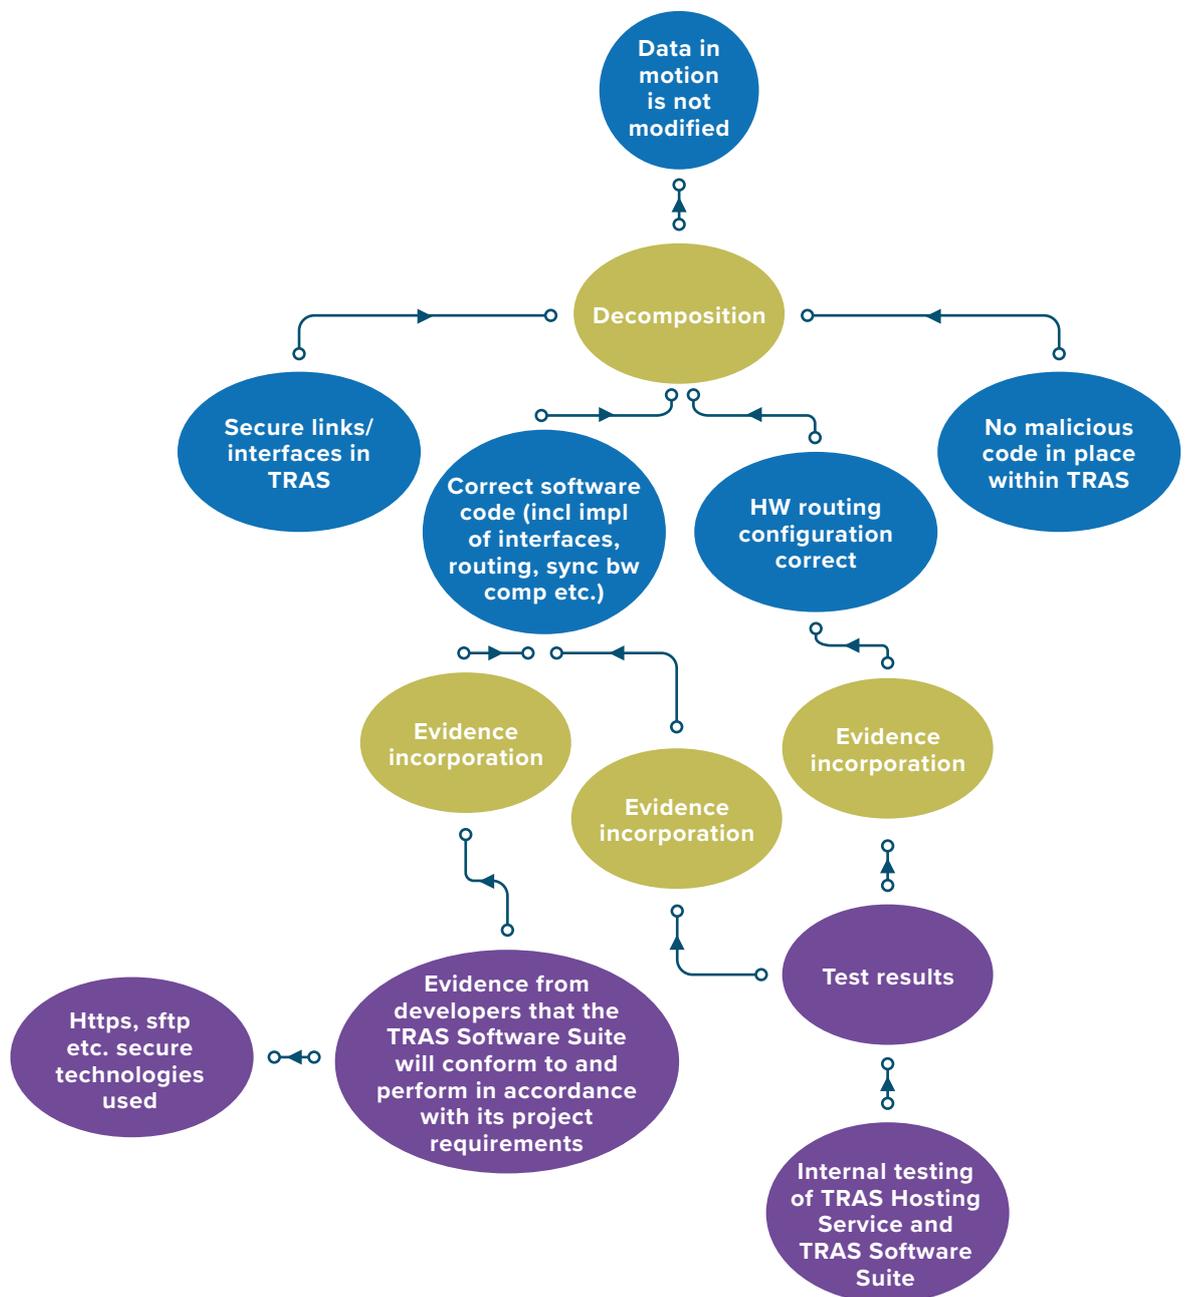


Figure 6: Data in motion integrity considerations

Finally, data at rest, which has reached the destination and is passively stored in the TRAS database, needs to be protected from modifications by employing digital and physical access controls. Of course, no malicious code should be within TRAS. Ideally, the data should also be encrypted. The possible decomposition is illustrated in Figure 7.

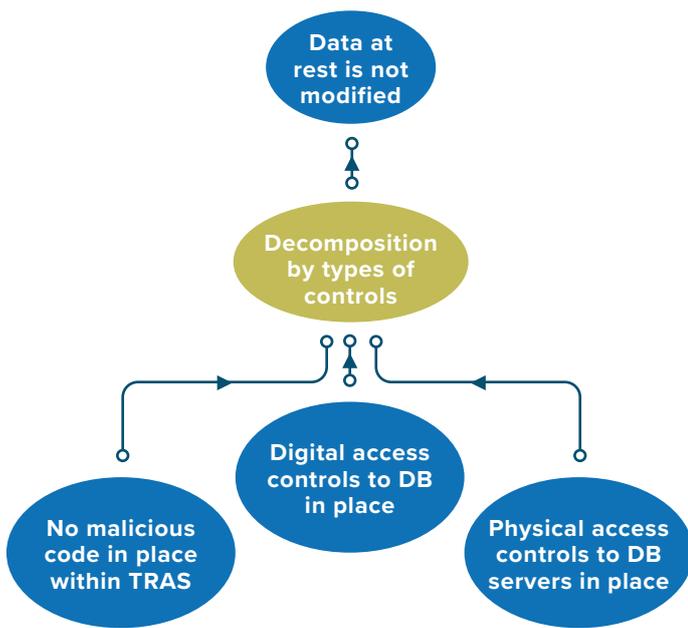


Figure 7: Data at rest considerations

In terms of confidentiality, most of the TRAS data have a business critical value to the organisations rather than a safety critical value. However, leakage of some data can potentially have safety implications. For example, disclosure of information about vehicle locations and schedules can enable physical attacks on the vehicles. Therefore, it is worth taking into account confidentiality as part of the security-informed safety case. In order to prevent leakage of the TRAS data, we need to make sure there is proper data separation between organisations using TRAS and also that there is no leakage of information to the outside world. The former relies on the correct implementation of specific TRAS features, such as authorisation of web users, additional access control set up for the TRAS report writers and users accessing the warehouse, correct TRAS software. The leakage of data to the outside world is prevented by many controls that have already been considered when analysing the data integrity, e.g. strong user authentication, identity management and permissions control, digital and physical access controls to the TRAS database. Additionally, organisations may request TRAS users to sign legal agreements that they will not share any TRAS information with any third party. The structure of the confidentiality related fragment is shown in Figure 8.

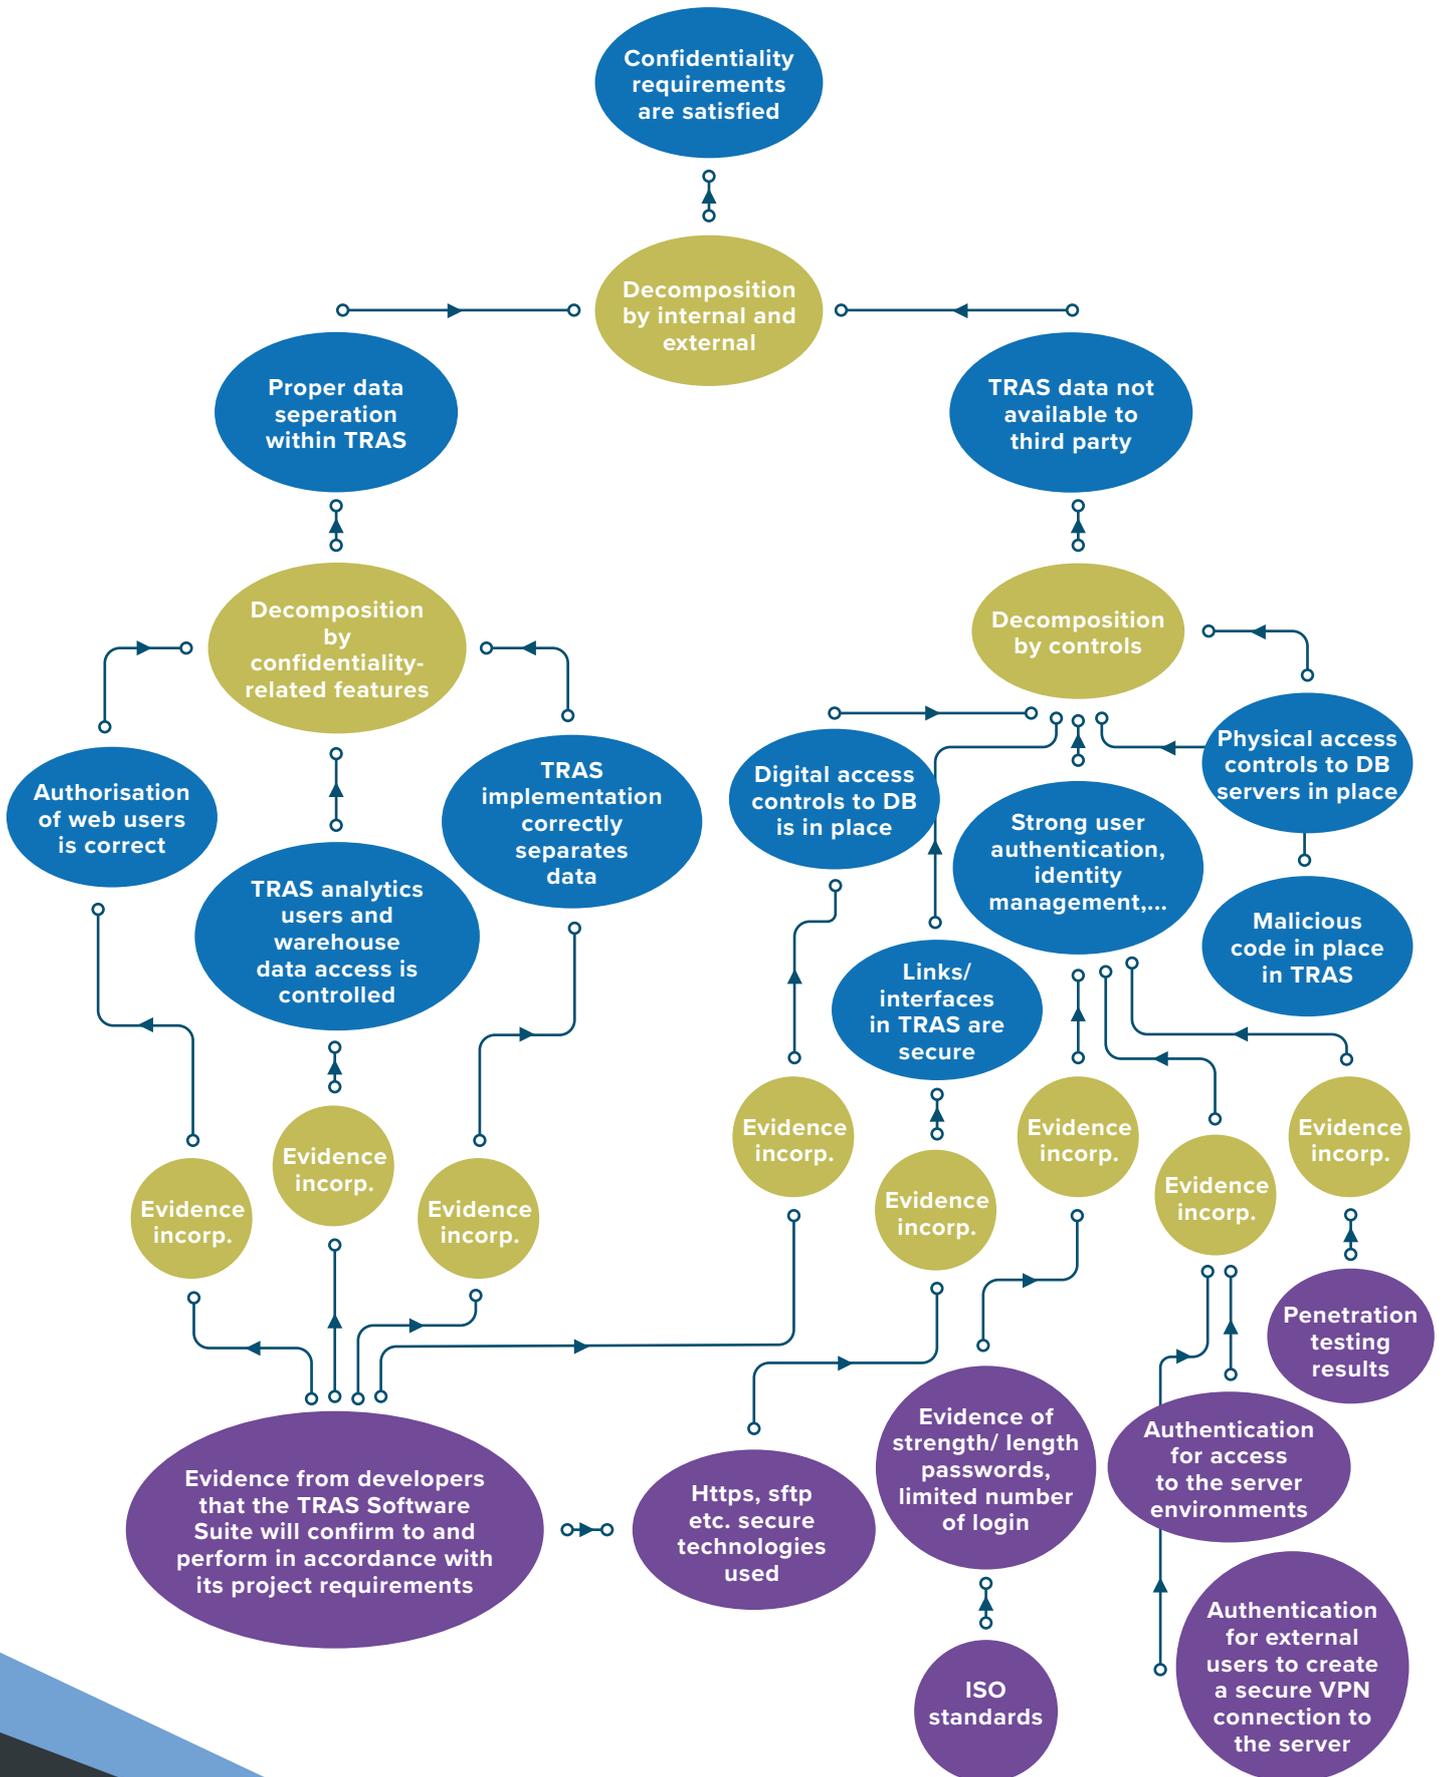


Figure 8: Discussion of system confidentiality

With reference to the claim about system requirements being satisfied in the future, concretion is used to clarify that 'in the future', means 'after any change'. To support such a claim it is necessary to consider all possible future changes or events that the device should deal with (e.g. component failures, changes to environment, etc.). Some of these events will be handled by component-level fault tolerance and recovery mechanisms and some will be handled by escalating the device's fault handling to another system or device. A part of the future-related branch is shown in Figure 9.

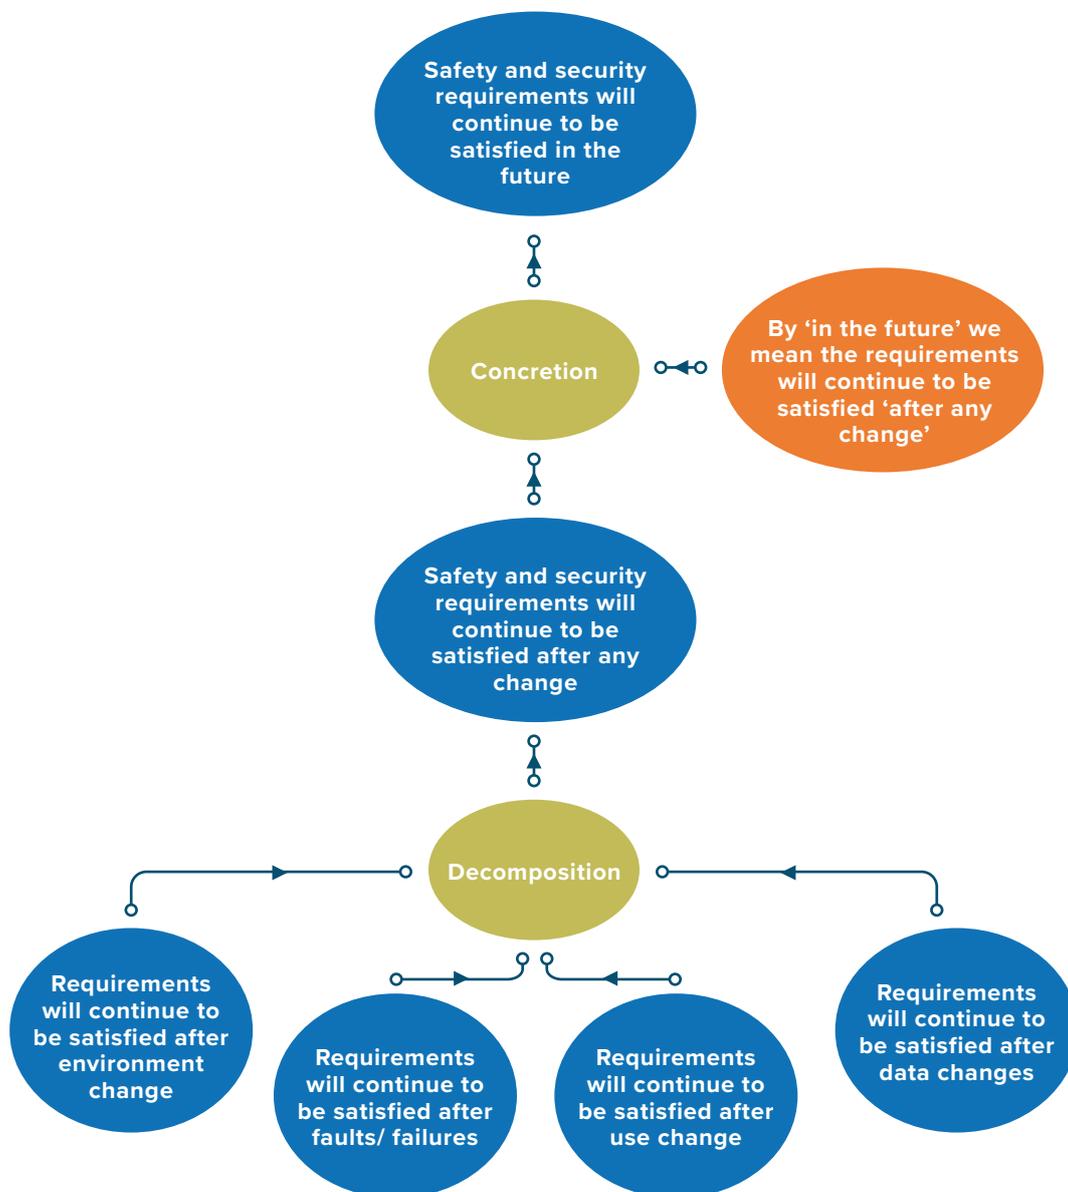


Figure 9: Future-related branch of the case with various types of changes considered

The change management procedure is not very well covered in the current TRAS documentation. This part of the case will require further development and the collection of supporting evidence.

This and other observations related to the supporting evidence at L1 are summarised in the table below.

Observations from L1 case	Analysis and recommendations
<p>Claim about a specific level of TRAS service availability (minimum 99.67%) needs to be supported by evidence with calculations performed.</p>	<p>There is evidence of redundancy and backup solutions in place, description of the recovery procedure and how it is tested, information on DoS attack exposure mitigations. HW and channels specification, configuration details were not provided. No calculations on how the chosen solution guarantees the required 99.67% of availability were presented.</p>
<p>Claim about data being initially accurate needs to be supported by evidence about the transition/migration procedure identified and followed properly.</p>	<p>A TRAS Data Migration Report and the transition schedule were provided. Need another evidence showing that the migration of data and users was successful. The process should be fully documented with respect to each transition and an audit document should be supplied. R: Provide evidence of successful data migration.</p>
<p>Claims about data in use related to the correct ongoing manipulation by users and apps at the endpoints need to be supported by evidence showing that the users are knowledgeable and trusted, input from all the applications trusted and there are no modifications of the input.</p>	<p>Evidence of the user authentication procedures in place (for both web and virtual private network connections) is provided. The strength/length of passwords, limited number of login attempts etc. policies are in place. Penetration testing results were at the time in progress. Not provided: Evidence of trainings for the TRAS users, any legal agreements signed by users that they will not share private data, evidence that any input from other tools is trusted. Evidence provided for the user authentication shows that the protection is not strong enough for this type of system. Additional means like two-factor authentication, IP white lists etc. should be considered for implementation.</p>

<p>Claims about data in motion require evidence showing that the links and interfaces in TRAS are secure and are correctly implemented by the TRAS software code, hardware routing configuration is correct and no malicious code is in place.</p>	<p>Evidence secure technologies used for the links (https, sftp etc.) is provided.</p> <p>The developer shall procure that the TRAS Software Suite will conform to and perform in accordance with the project requirements.</p> <p>Testing reports shall be provided by the developer for the TRAS Hosting Service and TRAS Software Suite.</p> <p>TRAS doesn't have any malware detection capabilities within it so absence of malicious code needs to be ensured by other means.</p>
<p>Claims about TRAS data confidentiality should be supported by evidence of proper data separation within TRAS and evidence that TRAS data is not available to third party.</p>	<p>This is covered by the evidence about the user authentication provided, as well as testing and conforming to the requirements report expected from the developer. The authentication should be improved by implementing additional controls.</p>
<p>Claims about safety and security requirements being satisfied in the future require evidence of the ongoing monitoring, intrusion detection, patching, and other activities dealing with both benign and malicious changes to system, data, use or environment.</p>	<p>Evidence of restore from backup procedures is provided.</p> <p>Evidence of the intrusion detection system available in the data centre is provided. Discussion on how administrators detect changes, manual procedures to deal with disruptions, user support services in place held. Planning for some of the upcoming changes (closer integration with TRAS) is provided.</p> <p>No detection capabilities available within TRAS.</p> <p>No evidence of the logging feature provided.</p> <p>Change management document describing the approach to dealing with various changes would be useful.</p>

Table 2: Analysis and evidence requirements for the architectural layer – examples

In this case study the system is still under development and the implementation details are not available for the analysis. Therefore, the L2 implementation layer is dealt with by developing recommendations for assurance activities – a verification and test strategy – for the different design and procedural controls that have been identified. Specific issues have been identified at L2.

At the L2 level, which is the detailed implementation level, all the technical information available about the actual system implementation should be introduced. This would help to analyse whether the critical safety and

security properties are really achieved when the system is implemented.

The implementation case should consider:

- how the controls identified at the architecture level have been addressed within the implementation (or alternative mitigations identified); and
- whether the connectivity of the components used in the architectural analysis is respected by the implementation or whether there are additional connections and interdependencies introduced by the implementation.

3.6 SUMMARY OF THE ANALYSIS

To sum up, the four steps of the analysis applied to the case study concerned different levels of detail of the risk analysis and specific attack scenarios.

In terms of the assurance case, these steps were initially followed at L1, providing the architecture-based assessment, and then refined at L2, taking into account the implementation detail. At both L1 and L2 it might also be necessary to review the results and assumptions of the previous steps with respect to the new detail that is now available.

In this phase, two analysis activities were combined:

- a risk analysis based on the safety documentation and a security-informed Hazop workshop; and
- the development of an L1 – L2 assurance case to support the Hazop recommendations and provide further analyses of the evidence required.

The steps with their application to the case study are summarised in Table 3.

Step	Description	Case Study
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.	The industry partner hosted a safety workshop to identify system hazards and consequences of failure. The analysis of this is included in Section 3.1 and forms the basis for security analysis: a mixture of desktop and meeting based security-informed hazard analysis as defined in relevant guidance [10].
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.	Specific attack scenarios were identified during the security-informed Hazop workshop and captured in Section 3.2. Additionally, penetration testing was identified as a future activity.

03.

<p>Step 6 – Focused risk analysis</p>	<p>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.</p>	<p>Based on the security-informed Hazop workshop and associated analysis, the attack scenarios and potential consequences of the attack were prioritised and corresponding recommendations summarised.</p>
<p>Step 7 – Finalise risk assessment</p>	<p>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.</p>	<p>The risk assessment process was finalised after the security-informed Hazop workshop with additional mitigations and controls identified and recommended for future implementation.</p>

Table 3: Summary of steps 4 – 7 of the analysis

Additionally, the step 8 – report the results of the risk assessment process was carried out in parallel to update the stakeholders at the appropriate level of detail during the process.

04. DISCUSSION

This has provided an anonymised example based on a real project that was done with realistic budget and time constraints.

In applying this approach to another example, this example could be augmented with a more detailed use of the CAE Blocks. For example, the argument blocks in Figure 9 could be examined using side-claims (see ‘CAE Blocks and Connection Rules Guide’) and the sources of change

considered justified. Other enhancements might be to use the guidance on review and challenge and complete a sentencing statement (see ‘CAE Review and Challenge’) for the example.

05. ACKNOWLEDGEMENTS

This document is based on material developed in earlier 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.