

SafSec: Commonalities Between Safety and Security Assurance

Samantha Lautieri, David Cooper, and David Jackson
Praxis Critical Systems
Bath, England
www.praxis-cs.co.uk
www.safsec.com

Abstract

Many systems, particularly in the military domain, must be certified or accredited by both safety and security authorities. Current practice argues safety and security accreditations separately. A research project called SafSec has been investigating a combined approach to safety and security argumentation, and has shown that there can be practical benefits in performing a combined analysis and documenting a combined argument for both safety and security.

1 Introduction

Where a computer-based system is required to meet rigorous standards of dependability, certification and approval costs can form a substantial proportion of the overall development costs. When such a system is maintained in service for an extended period, the cost of maintaining these approvals through in-service modifications and changes in operating environment escalates this element of cost still further. In an effort to manage and reduce certification and approval costs, the Defence Procurement Agency sponsored the SafSec (Safety and Security) project, which aimed to support safety and security accreditation of complex computer-based systems, particularly those now being deployed as Integrated Modular Avionics (IMA) systems.

SafSec focussed on two major issues: identifying and exploiting commonalities between the various disparate certification processes that an IMA system may be subject to, and providing a framework for certification of *modular* systems – those composed of standard components which are re-used in different configurations by a variety of applications.

This paper illustrates how commonalities exist in safety and security certification. When the commonalities are exploited, the effort and cost involved will be reduced and, if undertaken through a modular approach, issues of obsolescence will also be minimised, and possibly removed.

The resulting approach is called the SafSec Methodology and is the result of two years of research and case studies, involving a large number of stakeholders from the development, procurement and approval communities.

2 Background and Motivation

The acceptance into service of an Integrated Modular Avionic (IMA) system (ARINC 1997) presents a number of challenges which are not unique but which are perhaps more stringent in the avionics domain than in many others. The primary challenge is the need to satisfy a number of different accreditation bodies that a system is fit-for-purpose before operational clearance will be granted – this will generally include both a *safety* certification and a *security* accreditation. The second major challenge is the need to support modular certification; where components are shared between applications, we wish to be able to re-use elements of the evidence offered to support their acceptance.

Defence Standard 00-56 (MoD 1996) is the key UK MoD requirement for safety management; security accreditation will typically require meeting an approved standard such as the Common Criteria (ISO 1999). Neither of these standards, as issued at the start of the SafSec project two years ago, was entirely suitable for dealing with modular certification¹. Methods for certification need to support the certification of modules in isolation and support certification of combinations of such modules, rather than expecting certifiers to handle large complex systems as monolithic items for certification. This becomes a key issue in the on-going maintenance of certification – changing a single element in a modular system should be straightforward, and we do not wish to have to revisit the acceptance case for the whole system whenever a single substitution is made.

Although the detailed requirements of safety and security acceptance are often different, sufficient commonality is visible in the acceptance processes to encourage us to seek cost savings by eliminating duplicate effort. Certifiers need to be presented with convincing, objective arguments that the system has the safety and security properties that are required. The methodology therefore must be based on the presentation of direct arguments, and supporting evidence, that systems have the necessary properties and behaviour, and don't have any undesirable properties, to be safe and secure. Underlying acceptance by either community is a demand for good engineering practice in matters such as requirements traceability, verification and validation, configuration management and change control.

Note that although the SafSec project was initiated in the context of IMA systems, the challenges described here are generally applicable to a much wider domain, and we have already received substantial interest from other domains.

¹ The recent drafts of the new issue of DefStan 00-56 adopt a more flexible approach than the current issue, and are thus supportive of the goals of SafSec. The MoD team working on the new issue were included in the stakeholders consulted by the SafSec project.

3 The Goal of Common Certification

Conventional wisdom says that the safety and security domains have significant differences, and attempts to harmonise their work fail. Accreditation authorities demand different arguments, different evidence, presented differently, and focussed on different issues. Harmonisation is unattainable.

We think not.

Work we have carried out, supported by a range of stakeholders in the MoD, in military contractors and in the approval authorities, indicates that despite a number of differences between safety and security, there are still considerable benefits to be gained through a combined approach. Indeed, a case study currently underway is gaining these benefits and showing how safety and security can work together in a practical setting on a real project.

Our harmonisation is based on identifying common concepts between safety and security, and showing how the different analysis approaches can be seen as facets of a common, general analysis. This allows the different analyses to be documented in a way that highlights the common areas, and encourages the diverse teams to share information and insights. This common representation allows us to see overlap at four levels:

- system loss (e.g. death, security leak)
- cause
- mitigation
- evidence

The more the overlap, the greater the chance for re-use and savings between safety and security.

4 The SafSec Project

The results of the SafSec Project are captured in the Standard and Guidance Documents (Praxis 2004a, Praxis 2004b), which present an integrated methodology satisfying both safety and security certifiers.

The methodology illustrates a means for certification of both safety and security properties and is based on the identification of risks, and the justification that these risks have been adequately mitigated in the design and use of the system. Both safety and security fields require arguments and evidence to be provided that adequate measures have been taken to mitigate the risks, with the extent of the evidence required defined by the level of criticality of the mitigation measures. This commonality of approach provides an opportunity to apply a common method, and hence realise savings in effort and time.

One of the advantages of new modular approaches to system architecture, where properties as well as functions of modules are defined, is the feasibility of evolving systems, as technologies or requirements change, with limited impact on the design or implementation of systems or modules. If these benefits are to be realised for critical systems, then a way of structuring the process of certification, to minimise the need for re-certification when evolution occurs, is required.

One of the SafSec Methodology principles came from the realisation that both fields work from a base concept of *Risk* assessment and management. Identifying and mitigating risk is an essential driver for development processes. However, this concept is rarely given a central role in current certification approaches based on procedural frameworks.

Another emerging principle was the need to consider *Properties* of systems alongside the functions they perform, if safety and security are of concern. This leads to the idea that properties (expressed as objectives and assurance requirements (Hawes and Steinacker 1997)) should be central to the design process as well as function. Although this approach of focusing on properties can be applied to other non-functional aspects of design, SafSec has restricted its attention to just safety and security aspects.

The three technical areas which are central to the Methodology are those concerned with the processes of risk management, argument and evidence production during development, and modular certification. The three components of the Methodology combine as illustrated in Figure 1, and fit into a wider framework.

The Unified Risk Management Process takes account of safety hazards and security threats, together with the operational requirements of the target system, to produce a risk model alongside the architectural model of the system.

The Risk Directed Design Process uses the risk model, together with the architectural model of the system, to define the dependability properties of all the system modules in parallel with their functional properties, and produce the arguments supporting traceability of this process.

The Modular Certification Process takes the module's functional and dependability properties, and uses supporting arguments and evidence to justify certification of a module. As modules can be composed from collections of other modules, multiple applications of this process contribute to system certification. The result of the process is a set of safety and security certificates, and information that can support operational acceptance.

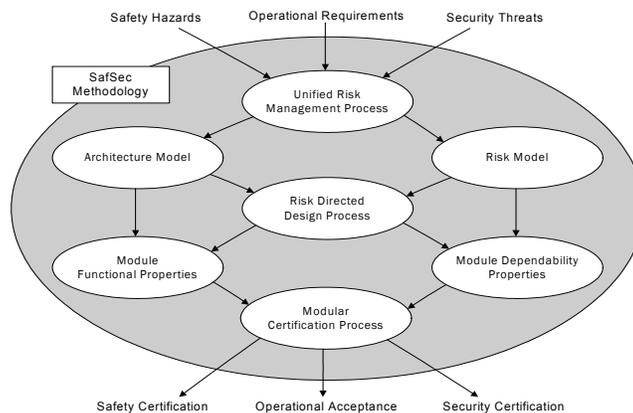


Figure 1: The SafSec Methodology

These three components of the methodology fit together, as indicated in Figure 1, and provide a core set of methods that support the key processes of requirements identification, modular design, and modular certification, based on a common understanding of the need to use dependability properties as one of the prime structuring principles.

5 Illustrating the Commonalities in Safety and Security Assurance

In this section we consider a simple example showing how safety and security analysis can be viewed in a common way, and hence overlaps and areas for re-use can be identified.

We will look at a (simplified) safety analysis, presented reasonably traditionally, together with an equally simplified security analysis. From these analyses we shall consider the likely mitigations that may be put forward, one set derived from the safety analysis and one set from the security analysis. We shall then view these analyses in a common SafSec framework, and suggest how common elements may be identified, leading to improved mitigations and re-use of arguments and evidence.

5.1 Introducing the Example

The example we will consider is derived from the Allied Standard Avionics Architecture Council (ASAAC) designs for modular avionics systems. The system consists of a variety of (safety and security related) software tasks running on a networked collection of identical processing elements. The distribution of tasks among processors, and the configuration of input-output and communications links, is controlled by an *Application Manager* process. The Application Manager configures the system according to one of several pre-defined *blueprints*, which are stored in a database. The Application Manager can reconfigure the system dynamically to take account of changes in hardware availability (e.g. if a processor fails).

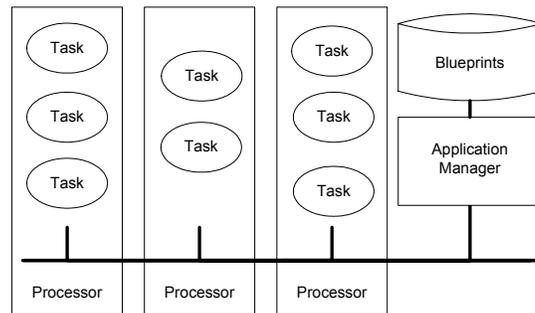


Figure 2: Example Architecture

Our illustration will focus on the *re-configurability* of the system. In order to cope with system failure or damage, the aircraft avionics are designed to allow tasks to run on different processors with different communication paths, and to allow the distribution to change dynamically during flight. If, for example, one processor were to be damaged in flight, then the tasks running on it would be migrated to other processors and the communications re-routed. We shall consider only the case of a pre-determined set of allowed configurations, called *blueprints*. If one blueprint becomes inappropriate, a different blueprint may be called up, leading to a different configuration of tasks, processors and resources.

We shall consider the safety and security analyses centred on the difficulties encountered in invoking a change of blueprint in flight. The boundary of the system being analysed is that of the *avionics system*, within the overall aircraft.

5.2 Conventional Analysis

Conventionally, a safety analysis would involve Hazard Analysis, workshops, Fault Tree analysis, FMEAs, etc. A fragment of the safety analysis for an IMA-supported aircraft might identify the following:

Safety Analysis

- Accident: Death of aircrew.
- Hazard: avionics and flight control systems catastrophically mis-configured.
- Cause/Consequence: Hazard leads to inoperable avionics, including flight control, which leads to an inoperable aircraft, which leads to total loss of aircraft and hence death (if ejection and recovery systems fail).
- Causes of Hazard (derived from FTA):
 - blueprint data as read corrupted or otherwise incorrect (and hence does not work safely).
 - Or: inability to read blueprint data successfully during system reconfiguration.
 - Or: installation of blueprint does not complete within the intended time.

Security analysis has its own techniques, focussing on the information assets at risk, the threat agents, and the means of attack. A small selection of the security analysis of such a system might identify the following attributes:

Security Analysis

- Information assets: secure information held in avionics system, e.g. crypto keys.
- Threat agents: enemy agents with direct access to the flight systems.
- Attack: obtain direct access to flight systems by bringing aircraft down in enemy territory, by corrupting blueprint data to make aircraft un-flyable on re-configuration.

Clearly these two analyses have produced different outputs and have brought different expertise to bear on the problem. However, in amongst the genuine differences there are hidden commonalities that may be hard to detect. If left to work in isolation, the safety and security teams may move on to propose the following mitigations:

Safety Mitigations

- 1) Include some form of checksum on the data to ensure that the read was successful.
- 2) Ensure that the blueprint data is read successfully before the old configuration is removed.
- 3) Carry out timing analysis to demonstrate that the worst-case execution time for each blueprint is within the allowed time window. Rely on the blueprint manager to install only blueprints that have been successfully analysed for timing behaviour.

Security Mitigations

- 1) Cryptographically sign the blueprint data at source.
- 2) Ensure that the blueprint data is read successfully before the old configuration is removed.
- 3) On read, check that the signature is still valid, and has been signed by an authorised source. This ensures that the read was successful, and that no accidental or malicious corruption has occurred.

We can see the overlap in these two analyses, but they each come from a slightly different angle and propose slightly different mitigations. The analyses are hard to compare, and we are in danger of implementing two overlapping mitigations, such as implementing both a checksum and a cryptographic signature, when one would do. Of course, in this simple example one would expect this overlap to be spotted. But in real systems, with hundreds of safety and security risks and multiple independent teams such confusions are more likely to arise and pass unnoticed.

5.3 Combined Analysis

Consider instead a combined analysis derived from inputs from both safety and security.

Combined Analysis

- Loss: Death of aircrew
Or: Loss of secure information to enemy troops.
- Caused by: aircraft crashes in enemy territory.
- Caused by: aircraft inoperable.
- Caused by: avionics and flight control systems catastrophically mis-configured.

- Caused by: blueprint data as read has been corrupted (by accident or maliciously) or otherwise incorrect (and hence does not work correctly).
 - Or: inability to read blueprint data successfully during system reconfiguration.
 - Or: installation of blueprint does not complete within the intended time.

This demonstrates clearly that the two potential losses: of aircrew and of sensitive information, are both due to the same underlying risk: loss of aircraft. This loss of aircraft in turn is caused by the same issues and risks with dynamic reconfiguration. When we get down to the level of technical problems like data corruption we can see that safety and security look at the same basic issues with slightly different emphasis. For example, safety considers first and foremost the risk of accidental corruption of data, whereas security focuses on malicious attacks. Both result in corrupted data, and a single mitigation, if selected correctly, can be used to address both aspects. If safety and security are considered jointly, we can derive the following combined mitigations.

Combined Mitigations

- 1) Cryptographically sign the blueprint data at source.
- 2) Ensure that the blueprint data is read successfully before the old configuration is removed.
- 3) Carry out timing analysis to demonstrate that the worst-case execution time for each blueprint is within the allowed time window. Rely on the source of the blueprints to sign only blueprints that have been successfully analysed for timing behaviour.
- 4) On read, check that the signature is still valid, and has been signed by an authorised source. This ensures that the read was successful, that no accidental or malicious corruption has occurred, and that the blueprints are certified as having successfully passed a timing analysis.

5.4 Re-use

During development, both safety and security authorities need evidence to support the arguments relating to the losses and mitigations, and evidence that the system development has followed required development standards. By combining the safety and security analysis it will be possible to re-use the same evidence for both authorities – within limits. For example, there is a reasonably complex argument that needs to be made to show that timing constraints are met from the combination of timing analysis at source, cryptographic signing at source, and signature checking at configuration time. Such an argument can be constructed and documented in, say, a GSN form (Goal Structuring Notation, Kelly 1999). The argument will be supported by evidence that demonstrates its validity. This will include, for example, contractual conditions to ensure that blueprint developers will perform timing analysis; technical evidence demonstrating the effectiveness of

such an analysis, and design information relating to the management of the cryptographic keys used for signature and verification. The whole argument and its supporting evidence will be of interest to both security and safety authorities.

There will, of course, be cases where one authority is not interested in some of the arguments, or possibly needs the arguments presented in a certain, specific way. In these cases there will need to be specific safety or specific security arguments presented. Does this conflict with our claim that finding the commonalities is cost effective? Does the occasional need to pull safety and security apart negate the benefits of merging them together?

We believe not.

Experience on a case study currently running with an MoD contractor suggests that the benefits of merging safety and security in the analysis phase greatly compensates for later having to present the same information in different styles. In fact, by concentrating first on the task of analysis, divorced from the idiosyncrasies of individual authorities' needs for presentation of evidence, we have achieved more effective analysis and clearer internal project documentation.

6 Revisiting the Goal of Common Certification

This example indicates how searching for the commonalities between safety and security can lead to a common presentation of information, despite safety and security focussing on different issues and using different analysis technique to arrive at the information. A single “cause-effect chain” presentation can capture both security- and safety-relevant information.

In this example there is no overlap in terms of the system loss: the accident identified as death of crew is different from the asset compromise identified as information leakage. But following down the cause-effect chain there is a lot of overlap. Indeed, the underlying causes of accidental and malicious alteration to data apply to both safety and security, although they tend to be given different weight by the two disciplines. Note that the notion of a combined analysis is independent of the manner in which the analysis is carried out – it applies equally to ‘top-down’ analysis (eg using fault- or attack-trees) and to bottom-up analysis (eg using event trees).

There is overlap of mitigation, too, and the system design may be simpler because the overlap has been identified.

For certification, the safety and security authorities may demand different types of evidence. But within the common framework it is easier to see how evidence can be re-used, and may even lead to certification authorities converging on their demands.

6.1 Implementing Common Certification

To facilitate realising the possible benefits of common certification, the SafSec standard and guidance documents provide a framework for planning development and assurance activities so as to meet the requirements of both assessment communities. The approach taken is objective-based: the SafSec documentation

identifies targets that a developer should satisfy rather than specific processes to be followed. The targets address four main areas:

- identification of losses and the causal relationships between them;
- definition of protection and mitigation measures and the associated assurance requirements;
- implementation and maintenance of systems which meet the identified requirements; and
- verification that the requirements are satisfied to the necessary assurance levels.

In some areas, the SafSec standard defines objectives at several levels – in deciding to implement the standard, we may choose to address a high-level objective directly, or to accept the breakdown proposed by the SafSec standard and meet a number of more detailed lower-level objectives.

The SafSec guidance supports assessment by both safety and security communities by providing a mapping between the SafSec objectives and their respective domain standards.

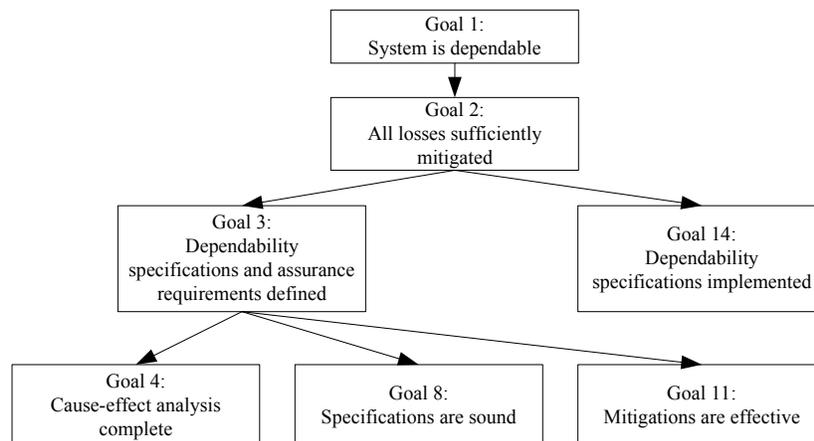


Figure 3: Simplified top-level goal structure for the SafSec Standard

The common framework provided by this argument structure also provides an important support for the certification of modular systems.

The process of identifying unified and explicit dependability properties can be carried out at component level. Extending the analysis of cause and effect relationships across modules allows safety and security properties to be expressed as *contracts* between components. Decomposition of safety arguments in this way has already been studied (Kelly 2001) – the SafSec framework allows this approach to be extended to exploit commonality between safety and security properties.

7 Conclusions

This paper has provided an overview of the work of the SafSec project, which has derived a new approach to the certification of highly modular safe or secure systems, such as proposed for advanced avionic architectures, based on the construction of safety and security arguments and the collection of evidence supporting those arguments.

Investigations into the opportunities and problems presented by modular architectures, and into the potential for the exploitation of commonality between the safety and security certification processes, have resulted in the definition of a framework and methodology which should provide scope for reducing the effort, cost and timescales associated with certification of a wide range of modular systems, including those that are safety- or security-critical.

The project has defined the SafSec Methodology, which combines:

- a unified approach to risk assessment for safety and security,
- a risk directed design process, which includes risk mitigation decisions in the design process and produces substantiated arguments to support them,
- a process supporting certification of modules within a modular architecture.

This paper presents the primary arguments for the usefulness and utility of this methodology, as a means of exploiting the inherent commonality between safety and security certification processes, and through modular certification realises the potential presented by modular architectural approaches.

References

- ARINC (1997). ARINC 651-1 Design Guidance for Integrated Modular Avionics, ARINC Incorporated, Annapolis, November 1997.
- Hawes and Steinacker (1997). Combining Assessment Techniques from Security and Safety to Assure IT System Dependability—The SQUALE Approach, VIS97 security conference, Freiburg, Germany.
- ISO (1999). ISO 15408, Common Criteria for Information Technology Security Evaluation. International Standardisation Organisation August 1999 (Version 2.1).
- Kelly, T P (1999). Arguing Safety – A Systematic Approach to Safety Case Management, DPhil Thesis, York University, Department of Computer Science Report YCST 99/05.
- Kelly, T P (2001). Concepts and Principles of Compositional Safety Case Construction, University of York, COSMA/2001/1/1, May 2001.
- MoD (1996). UK Ministry of Defence, Defence Standard 00-56, Safety Management Requirements for Defence Systems, Parts 1 and 2, Issue 2.
- Praxis (2004a). SafSec Standard Material S.P1199.50.2, Issue 2.6, May 2004.
- Praxis (2004b). SafSec Guidance Material S.P1199.50.3, Issue 2.6, May 2004.

Acknowledgement

We are grateful to Mark Suitters FBG Strike (4) DPA, for his support and funding for this work, and all the SafSec stakeholders for their invaluable contributions.

