

From Safety Cases to Security Cases

R D Alexander, R D Hawkins, T P Kelly

Department of Computer Science, University of York

Deramore Lane, York, YO10 5GH, UK

Abstract *Assurance cases are widely used in the safety domain, where they provide a way to justify the safety of a system and render that justification open to review. Assurance cases have not been widely used in security, but there is guidance available and there have been some promising experiments. There are a number of differences between safety and security which have implications for how we create security cases, but they do not appear to be insurmountable. It appears that the process of creating a security case is compatible with typical evaluation processes, and will have additional benefits in terms of training and corporate memory. In this paper we discuss some of the implications and challenges of applying the practice of assurance case construction from the safety domain to the security domain.*

1 Assurance cases are a powerful way to capture arguments about system properties

An assurance case is a structured argument that a system has some properties we desire; that it is safe, or reliable, or secure against attack. Defining *safety* cases in particular, Kelly says “A *safety case should communicate a clear, comprehensive and defensible argument that a system is acceptably safe to operate in a particular context.*” (Kelly 1998). We can generalise this easily to properties other than safety.

The argument in an assurance case shows how a high-level claim (e.g. “the system is adequately secure”) is ultimately supported by detailed evidence of particular low-level properties (e.g. some statistical testing data targeting a key security requirement on one small software component). An assurance case will refer to a range of evidence items, and will show how different types of evidence (e.g. evidence of good process, results of manual design review, and results of component testing) combine to give us confidence in higher-level properties (Hawkins and Kelly 2010). In other words, the assurance case captures the *rationale* of why the evidence we have produced (the results of the low-level analysis activities we

have carried out) gives us reason to believe the high-level claim (which is, ultimately, what we are interested in). Evidence on its own is never enough – even the strongest, most relevant evidence provides only partial support for a given high-level claim, and making that connection between evidence and claims requires an act of subjective judgment. Assurance cases make that subjective reasoning explicit. An assurance case does not replace *any* specific technique for analysis or for generating evidence. What it does do is show the connection between those techniques and the high level claims we want to make. An assurance case also provides a way of capturing assumptions about a system – these may describe a particular context, including a specific version of the system, a description of its environment (including particular threats present in that environment) and a description of how the system will be used.

Security and safety are open problems – the domain is open-ended and ultimately somewhat subjective. Security is worse in this respect, because of increased uncertainty about the capabilities and knowledge of attackers. At any time, new information may arise (about our system or about its environment) and thus it is necessary to review the reasoning about the safety or security of the system. Assurance cases provide a record of what that reasoning was.

In a sense, whenever we honestly claim that a system is acceptably safe or secure, we have an implicit assurance case; we have some mental model behind that claim that we could probably describe if asked. By making an *explicit* assurance case, however, we open up that mental model to review and criticism by others, and we record our reasoning so that others can learn from it. For example, if someone wants to make a change to the system, they can use the assurance case to help assess the impact of that change.

Acceptance of an assurance case is not a mechanical process; it requires subjective assessment by a customer, regulator or evaluator. This does not mean that it is wholly arbitrary and idiosyncratic; we can have rigour in assurance cases by drawing on the work that philosophers – such as Toulmin (Toulmin 1958) – have done on informal logic. See (Kelly and Weaver 2004) and (Bishop et al. 2004) for an explanation of this. Key techniques include assurance case patterns (to capture best practice in argument structure), systematic review processes, and the use of appropriate notations.

It is typical, in safety, for the developer of a system to create a safety case for that system and submit the safety case to the customer or regulator for their assessment. In many safety regimes this is required by standards. If the developer is not required to produce a safety case, it is possible for a third party to produce one as part of a safety evaluation or assessment process. In particular, a high-level assurance case can provide a strong tool for understanding how the safety efforts taken by a developer fit together to create a safe system.

2 Assurance cases have a strong track record in safety

Safety cases have been widely used for over 30 years and are now a mature approach; for example, they are required by the MOD for all equipment acquisitions, under Def Stan 00-056 (MOD 2015).

One benefit of producing a safety case, especially early in development, is that it can help engineers to think about where they need particular arguments, justification and evidence. It can also lead to early recognition of deficiencies in the system design or the process around it. For example, in one case study carried out by researchers at the University of York, creating a safety case revealed a lack of traceability between high-level requirements and evidence about the implementation. Similar results occurred in a security case example produced by Ankrum (see Ankrum 2008 and section 3.1.2 of this paper).

For safety cases to be effective however, they have to be approached with intent to genuinely capture the safety (or not) of the system; failing that, they need to be assessed by evaluators who are competent and willing to find flaws in them. A safety case produced and accepted as part of a box-ticking exercise will not be effective; one of the findings of the report into the Nimrod accident (Haddon-Cave 2009), was that the Nimrod safety case had been such an activity. Kelly, in (Kelly 2008) identifies, through experience of reviewing safety case, several different ways in which safety cases may be ineffective.

3 There are good reasons to apply assurance cases in the security domain

We explained in sections 1 and 2 that assurance cases are widespread in safety. In this section, we will review the potential for applying them in the security domain. In section 3.1 we look at whether (and how) security assurance cases can be created. In section 3.2 we discuss the differences between safety and security, and in section 3.3 we look at the possible benefits from security cases, given typical security practices and current challenges. In section 3.4 we review some of the practicalities of moving to a security case approach.

3.1 It is practical to create security assurance cases

3.1.1 Methods and guidance are available

There is already some published guidance on creating security cases. Basic advice is provided by Goodenough, Lipson and others – see (Goodenough et al. 2007) and (Lipson et al. 2008) – as part of the “Build Security In” initiative in the USA. The advice there is not particularly detailed, but it is clear and practical. Their motive for proposing security cases is the increasing complexity of systems. They note that many people have responded to increasing system complexity by proposing an empirical, post-hoc approach (treat the developed system as a natural phenomenon, and assess its security after deployment by observing the number of security flaws uncovered). They reject this: they don’t see how it can provide the level of confidence needed for high-security systems. Their approach is aimed at vendors; they want vendors to instrument their development processes with evidence-generating activities, and use a security case to capture the result.

A more detailed process is provided by the SAFSEC standard (Dobbing and Lautieri 2006), which was developed by Altran Praxis as a way to unify safety and security cases. In the SAFSEC approach, safety and security risks are treated equivalently – a combined set of mitigations is proposed and a single assurance case is produced that argues that the system will be *safe and secure*. The safety side is based on Def Stan 00-56, and the security side on release 2 of the Common Criteria. It is broadly compatible with similar work on safety-security unification by the Industry Avionics Working Group (IAWG) (IAWG 2007).

The SAFSEC standard provides a strong analogy between safety and security in an assurance case framework – it explains where safety-security commonality exists and provides a specific process to exploit that. As such, it could be followed directly to create a security case. It does not, however, address the cultural, epistemic and economic challenges that we discuss in Section 3.2.

3.1.2 There is some published experience with security cases

There have been a number of security cases created in small-scale case studies. In (Ankrum 2005), Todd Ankrum of MITRE briefly outlines a case study he carried out for the US National Security Agency. The system was a secure enclave project by the NSA’s research division, which used several authentication systems and provided an access log. The software had been developed to Common Criteria EAL 5, used formal methods, and extensive documentation was available includ-

ing a vendor-supplied Security Target document and a NSA-supplied Protection Profile.

When Ankrum and his colleagues produced a security case for this system, they found a number of problems. Despite the obvious rigour with which the project had been developed and documented, it was not possible to fully argue that the system met its security goals. First, although most security threats could be traced to requirements that mitigated them, at least one threat had no such requirements. Second, once the threat-requirements connection was explicitly made, it was not clear that the requirements for each threat sufficiently mitigated the threat; presumably the vendor's process had not made this relationship explicit, so it had not been apparent before. Finally, Ankrum's security case attempted to argue that all security enforcing functions had their dependencies met, and it became apparent that this was not true in all cases.

Lautieri and her colleagues at Altran Praxis describe a case study in (Lautieri et al. 2005) of a Command and Control system, for which they produced a combined safety and security argument. Their paper explains how they create a modular case for this system, and notes that the combination of safety and security was quite acceptable to certifying authorities that were only concerned with one of them; combining the two domains did not cause a communication problem.

Ankrum has also illustrated how an assurance case can capture the implicit argument in a standard that doesn't demand an assurance case. Specifically, in (Ankrum 2005) he briefly outlines an argument that a product has achieved Common Criteria EAL 4, by starting with the claim "Product meets EAL 4" and arguing that it does everything the Common Criteria requires to support that claim. Ankrum's approach has some weaknesses; for example, it is primarily a process rather than product argument, and it is perhaps better described as a compliance case rather than a true security case. The process he gives for creating the case is perhaps over-mechanized; realistic security case creation will involve more subjective judgement than that. It has the value, however, of showing how non-assurance-case standards can be adapted.

The concept of security cases has received increasing attention in the literature in recent years. Knight (Knight 2015) suggests that existing security verification techniques such as proof need to be supported and contextualised within the framework of rigorous (but informal) reasoning that assurance cases offer. He (He and Johnson 2012) present some positive experiences of using security cases (structured using the Goal Structuring Notation) in the Healthcare IT domain. Other papers have continued to explore how best to structure the security case. For example, Yamamoto et al. (Yamamoto et al. 2013) suggests a structure based upon the Common Criteria. There are also some positive experiences in the sub-domain of 'security informed safety cases' (Netkachova et al. 2015) which specifically address the extension of existing safety case practice to include appropriate consideration of security-related safety failings.

3.2 There are significant differences between safety and security

There is no question that safety and security are separated by different goals. The open question is whether the same means (in this paper, assurance cases) can be used to demonstrate achievement of those goals. Cockram and Lautieri say in (Cockram and Lautieri 2007) that the two domains *can* be served by the same assurance case mechanism: both can use cause-effect models (e.g. fault trees or attack trees), both can derive requirements to mitigate the problems thus identified, and both can argue in an assurance case that those requirements are implemented and that they will achieve the mitigation that is needed. Lautieri et al. (Lautieri et al. 2007) also give an example where separate requirements for safety and security were merged to create a single requirement that served the needs of both domains.

Beyond the basics discussed by Cockram and Lautieri, there are a range of theoretical and practical differences that we need to consider.

3.2.1 Theoretical Differences

The obvious difference between safety and security is the presence of an intelligent adversary; as Anderson (Anderson 2008) puts it, safety deals with Murphy's Law while security deals with Satan's Law. Safety is mostly concerned with the predictable (or random) behaviour of the non-human, non-goal-seeking world, and with adaptive behaviour by humans (e.g. seeking to make their job easier) that is not aimed at reducing safety per se (although it often does as a side effect). Security has to deal with agents whose goal is to compromise systems. These attackers may be systematically probing a system for vulnerabilities (rather than acting randomly) and may realise when they have breached one defence and move to exploit that (whereas non-human phenomena cannot respond to such feedback, and non-malicious humans may try to undo their actions if they realise they have bypassed a defence).

Another contrast is that a lot of effort in traditional system safety has gone into assigning probabilities to basic events (e.g. random failure of a valve) and computing the probabilities of accidents stemming from combinations of those events. Assigning probabilities to the action of unknown intelligent adversaries is dubious – our uncertainty there is *epistemic* (due to lack of knowledge) rather than *aleatory* (due to chance).

Safety does, however, already have to worry about epistemic uncertainty. This has become very apparent over the last 30 years as systemic faults (faults in the design that will cause failures under certain circumstances) have overtaken random failures as the cause of accidents. Partly, this is due to the great strides that system safety has made in handling random failures through component redundancy and architectures that support it. It has also come, however, from increasing

use of software. There are methods for assigning failure probabilities to software – e.g. see Bishop (Bishop 2005) – but they do not have great credibility.

Increasingly, software safety is approached in terms of the evidence we can generate and the confidence it gives us that the software’s behaviour will be safe in context. Where the link between some evidence and the claims it supports is inadequate, we have an *assurance deficit*. The significance of assurance deficits has been recognised in safety – e.g. see (Hawkins and Kelly 2009) and (Hawkins et al. 2009) – and they may be even more significant for security because of increased uncertainty about attacker’s goals, capabilities and actual actions.

Because of this high uncertainty about attacker behaviour, it is common in security-critical development to utilise a variety of security measures that are not responses to specific threats. These measures instead provide a degree of protection against a whole class of attacks. For example, if one process has to be given temporarily elevated privileges, it is common to give it those privileges for only the minimum amount of time. In safety, there is concern about such generic “hardening” – a safety feature added without good rationale may merely be adding to the complexity of the system, and thus increasing our chance of introducing an error and not discovering it.

Well-justified defences against common errors are, however, common in safety. For example, safety-critical programming language subsets (such as MISRA C) often forbid dynamic memory allocation, thereby showing that errors from running out of heap space cannot occur. Similarly, they often forbid recursive calls to functions, allowing the size of the stack to be bounded and thus ruling out all stack overflow errors.

Creating an assurance case can help to capture the rationale for such hardening, and thus distinguish between features that give a clear benefit, and features for which the benefit is uncertain. If it is not possible to justify the inclusion of a security feature (a valid place for it in the argument cannot be found), then you may be spending effort on measures with no actual security impact. Just like safety, all security is a trade-off, and every software feature has a cost (if nothing else, in the complexity of the resulting architecture and hence in the chance for an evaluator to miss a vulnerability).

Security evaluators often use static analysis for vulnerability signature detection – not to detect specific violations of known requirements, but rather to find “holes” in the system that an attacker might be able to exploit (Gutgarts and Temin 2010). This is quite compatible with assurance cases, and can be incorporated into the argument as a complement to claims about dealing with identified threats. In (Hawkins and Kelly 2009), the authors provide a set of argument patterns for software. Although these patterns provide arguments regarding safety, they could be adapted for software security. In such cases the use of static analysis for vulnerability signature detection fits neatly into the part of the argument where it is shown that additional hazardous contributions at the code level have been identi-

fied. In other words, it is argued that you have looked well enough for common low-level vulnerabilities that there is adequate confidence in their absence. “Adequate” can be scaled to a level that suits the security criticality of the system.

A final theoretical difference is that security-critical software often has to adapt quickly as attack patterns change (Carter 2010). This has implications for the use of assurance cases, because a case may have to change if the software changes. Even if it turns out that there is no impact on the case, the case maintainer will have to spend time reviewing the case to confirm this. Assurance cases can therefore increase the cost of making changes, and introduce delays. On the other hand, we need to understand the security impact of any changes we make; if we don’t, we may fix one vulnerability but create a worse one in the process. If we have an assurance case, then we can use it to trace a low-level change up to its implications in terms of our high-level security claims. The challenge of making assurance cases more dynamic in response to changes to the system or its environment is an area of increasing interest for safety cases, and work in this area will also benefit security cases.

3.2.2 Practical Differences

The previous section talked about the fundamental distinctions between safety and security – those that are likely to endure over time. In this section, we will look at the practical differences, which may be accidents of history; they cannot be ignored, but it may be possible to change them. A number of these distinctions were identified in a meeting between representatives of the safety and security communities in the UK – Carter reports the results of this in (Carter 2010).

Perhaps the biggest practical concern is the process maturity of security-critical practice. In safety-critical software, even at lower criticality levels, mature development processes are the norm: developers use good processes for requirements management, test planning and configuration control. They monitor their processes for weaknesses, and correct them when they find them. In security-critical projects of an ordinary (commercial-grade) standard, this is not always the case. Requirements may be implicit (requiring evaluators to define their own security targets), testing may be unstructured (making it difficult to relate test schedules to specific requirements) and configuration control may be poor (making it difficult to relate review or test results to specific software versions).

The process problems in security-critical software are not necessarily unique to the domain; more likely, they are present because vendors have not had the economic incentives to do better. Poor process can make it difficult to produce an assurance case. It may be that the system is adequately secure, but because of poor process it is not possible to get adequate evidence of this. Similarly, it may be that we cannot build our case adequately because we cannot understand the security

requirements. These are problems, certainly, but they are not problems with assurance cases: they are problems with the system being evaluated. If we cannot understand the system well enough to build an assurance case, then we are not in a position to say that it is secure. Problems with the assurance case may be the “mine canaries” that alert us to problems with the system.

It has been suggested that safety has more problems with requirements, security more with low-level defects in implementation. For example, this view was reported by Carter (Carter 2010) and seems to be an assumption of Lipson (Lipson and Weinstock 2008). This could be a fundamental distinction, but it may just be a side-effect of immature processes in most security-critical software. Once your development process can reliably implement the requirements that you set out to implement, then getting the requirements right is the one place where problems can still manifest (as requirements engineering is ultimately unbounded, there will always be requirements errors). As process maturity improves in security-sensitive software vendors then this difference may vanish.

It has been noted that there are some specific differences between software security standards and software safety standards, in terms of the development practices they demand. For example, King noted in (King 2009) that an avionics system certified to the highest level of DO-178 (RTCA 1992) may not meet Common Criteria EAL 5 or above, because EAL 5 requires that the “developer mathematically prove the security properties of the [software]”, whereas DO-178B does not. These are likely incidental, rather than fundamental differences, and agreement could be reached between standards (the inevitable politics aside).

Overall, we suggest that security evaluators remain aware of these practical differences, but experiment with assurance cases based on the safety case model, and see what obstacles they encounter in practice. There is a wealth of experience with assurance cases in the safety domain – where obstacles appear in security, there may already be known solutions.

A final note on practical challenges – safety-critical systems are becoming increasingly software-controlled and increasingly connected. As a consequence, there is an increasing threat of malicious outsiders causing safety-critical failures. If the security-critical software domain does not adopt assurance cases, then the safety-critical software industry will have to extend safety cases to include security. We will, therefore, need to resolve the problems above, and this will be easier if assurance cases are adopted by the security community as well.

3.3 Security cases can help security practice

3.3.1 There are benefits for handling complexity

As noted earlier, the major role of assurance cases is to combine evidence from diverse analyses and show how they complement each other. For example, code review alone provides limited assurance – high assurance will require a range of complementary techniques targeted at fairly specific types of attack. This might include code review, static analysis of code-level requirements (e.g. pre- and post-conditions), static analysis for vulnerability signatures, and statistical or systematic tests.

Assurance cases provide a way to connect system-level properties (e.g. security of certain data) to low-level requirements and analyses. For example, it can be difficult for an evaluator to implicitly maintain the connection between the code in a given source file and the system-level goals to protect certain data or prevent denial of service.

When effective, assurance cases can focus evaluator attention on critical parts of the system at the expense of others. In a world of finite effort and growing software complexity, this is necessary. It is likely that a given evaluator already has some way of doing this – for example, they may concentrate primarily on security enforcing functions, rather than trying to review all of the code. Creating an assurance case allows you to put this in context and see how much assurance such a strategy really gives you.

If a vendor provides a security case, this focussing aspect allows you to treat the case as a summary of the vendor’s security thinking – what they have concentrated on, and what they have not. It may be safe to assume that some vulnerabilities will appear in the aspects that they have not addressed, which means that an evaluator can concentrate on attacking those areas. This prioritisation would be difficult to do with only an implicit security case.

Where problems are encountered in assurance because of limited expertise, it may be possible to use an assurance case to “bracket” those problems ready to hand them over to a third-party expert. For example, we may have a portion of a software system that deals with radio communications according to a complex industry standard. As a specialist in software security evaluation, we may not be able to understand that part of the software well enough to assess its security. What we can do is make a number of claims about that component, claims that if true would allow us to support the overall security claims of the system. We can then ask a third-party domain expert to check those claims against the specialist software.

One tool for managing complexity in assurance is modular safety cases (Kelly 2005). For example, if there is an operating system that is often used in multiple systems, a security case module may be created for it. This module would make certain claims about the operating system (for example, that there was definitely no way for a process to elevate its privileges without appropriate authorisation) subject to certain dependencies (for example, that all loaded device drivers could be trusted). This argument module could then be used as part of a security case for any system using that operating system. Similarly, modules might be created for common hardware or for common support applications.

The modular case approach has been explored in other domains (Cockram and Lautieri 2007) with the aim of reducing certification update costs. Both the IAWG (IAWG 2007) and SAFSEC (Dobbing and Lautieri 2006) processes support some degree of modularity. If an evaluator is creating assurance cases, they may be able to create modules for common OSs, hardware and software components, thus simplifying the process and reducing costs.

The security of a system obviously depends on the context in which it operates. The assurance case for a system may need to change if the system's context changes, for example if the system needs to communicate with a new peer system. Thomas (Thomas 2010) raises this as a potential obstacle to the adoption of assurance cases. It is a genuine problem, although it is one faced in the safety domain as well, particularly as the complexity of safety-critical software increases. Modular assurance cases may help to deal with this challenge.

3.3.2 There are benefits for justifying decisions

When an evaluator believes that a system is insecure, they can demand that the vendor generate additional evidence (e.g. run more tests), or they can demand changes to the design, or they can impose restrictions on how the system is used. If the evaluator is all-powerful then this is not a problem. Realistically, evaluators have to justify their position against claims by vendors that the change will be very expensive. If the problem is an objective vulnerability (something that can be seen e.g. in the code, and is easy to exploit) then this may be straightforward. If the concern is an assurance deficit – perhaps a lack of confidence that a certain requirement is satisfied – the justification needs to be rather more subtle. Assurance cases can help with providing such justification by linking analysis activities and claims about the system design to high-level security properties.

3.3.3 They can enhance typical evaluation processes

As noted in section 1, assurance cases can be created retrospectively. Indeed, if the vendor does not produce an assurance case then the evaluator can. This is similar to when an evaluator creates a security target document because the vendor did not supply one (or supplied an inadequate one). There are potential traps in this activity, particularly if the evaluator is unwilling to reach the conclusion that the system is not secure – see the Nimrod report again (Haddon-Cave 2009) for an analogous example in safety. The cost of creating an assurance case retrospectively may be high, but doing so should also lead to some specific benefits: in creating a retrospective security case an evaluator creates a paper trail for their reasoning – they will capture a justification for why they think the system is adequately secure (or why it is not).

The first benefit of this that if we later need to re-evaluate the product (for example, after a major change or security incident), we know what the original evaluator did and can repeat only that which has been invalidated by the change.

A second benefit of such records is that they can provide a training tool and a mechanism of corporate memory. Junior evaluators can look at existing security cases to understand what kinds of security processes and evidence are acceptable to the evaluating organisation – they can see what was included and what was emphasized. Similarly, if an experienced evaluator encounters an unfamiliar type of system or an unfamiliar set of security requirements, they can look over past assurance cases for similar systems, and see what their peers did. There are existing techniques, such as assurance case patterns, for distilling this kind of information into a generic form.

The provision of a security case (whether by vendor or evaluator) may help coordinate teams of evaluators. The case could provide a central structure around which activity is organised – one evaluator could be assigned to each abstraction tier, or to a major component, or to a high-level requirement. They can also record their results in terms of the assurance case structure (so, for example, an evaluator working at the source code level may note problems there, and trace their significance back up to security claims made at the system level).

The costs and benefits of assurance cases are likely to vary with the level of security claimed. At lower levels, the assurance being sought may be relatively modest, which will make it easier to create a compelling security case. On the other hand, the product may be off-the-shelf, and when this is combined with low process maturity it may make it hard to create the case at all. As noted above, of course, if the evaluator requires good justification of security then inability to create a case is grounds for rejecting the product.

As we move to higher security levels, the product is more likely to be bespoke and the evaluator is likely to be involved from an early stage. This may make it easier to justify a vendor-created security case, or one produced through vendor-

evaluator collaboration. The risk at higher levels is that the case will need to claim a very high level of assurance, and this may be difficult to justify. As ever, if a security case cannot be made compelling, then it may well be that the system is not secure.

3.4 There are practical challenges in moving to security cases

As noted in section 3.1.2, there are no public descriptions of major applications of security cases. It is therefore difficult to say what the costs and timescales will be. However, we can draw on experiences of moving from procedural safety standards to safety case approaches. The major need is for expertise: at least the evaluators, and maybe the vendors as well, will need to learn to produce and assess assurance cases. This expertise does not need to be spread uniformly throughout the population, but a sufficient number of staff will need to be skilled at judging assurance cases and fluent in GSN. There are published accounts of moving from a prescriptive standard to a safety case approach, and from moving from a textual case to a GSN case. For example, see Chinneck et al (Chinneck et al. 2004).

The move to assurance cases does not always have a major disruptive effect on the products being assured, or on the techniques used to evaluate them. In the first instance, assurance cases merely provide a way to relate what is already done to the goals that are already held (although those goals may previously have been implicit). Their impact will not, for example, be comparable to moving a vendor's software development from procedural to object-oriented. What they *may* do, in terms of disruptive influence, is reveal that there are weaknesses in the products or in the evaluation of them. That is, after all, the point of an assurance case: it should allow a third party to assess whether the system has the properties that are claimed for it. When an assurance case does reveal problems, then evaluators and vendors may, of course, decide that they need to change their processes and tools. *That* may be disruptive.

If an evaluator already produces a security target document, then an assurance case is partly an extension of that - it maps threats onto requirements, and requirements onto justification of their adequacy and evidence that they have been implemented. If the vendor already follows the requirements of the Common Criteria, particularly the refinement structures required by the higher EALs, then there may be a clear mapping onto the existing software safety case patterns, (Hawkins and Kelly 2009), which use a similar structure of refinement through 'tiers' of development.

4 Conclusion – there is potential, cost, and risk

Adopting assurance cases can offer benefits to security assurance practice, potentially increasing the rigour and transparency of the security evaluations. However, there will be costs to adopting assurance cases (e.g. associated with training), and there is an element of financial risk here. Assurance cases are widely used in safety, but there has been limited use in security thus far. To gain the benefits from assurance cases, both developers and regulators will need to develop significant in-house expertise, and address some of the challenges specific to security that we have highlighted in this paper. Finally, it is worth recognising that some of us have expressed concerns that the preparation and presentation of a ‘full and frank’ security case (that highlights, for example, contextual assumptions) could itself present a security risk, i.e. it could aid an attacker greatly in the formulation of their attack vectors. These concerns bring us back to long-standing debates concerning the effectiveness of *security by obscurity*.

Acknowledgments We acknowledge the financial support provided by the UK CESG for the original study underlying this paper.

References

- Anderson R. Security Engineering: A Guide to Building Dependable Distributed Systems. 2nd ed: John Wiley & Sons; 2008.
- Ankrum TS, Kromholz AH. Structured Assurance Cases: Three Common Standards. 2005. www.asq509.org/ht/action/GetDocumentAction/id/2132
- Bishop P, Bloomfield R, Guerra S. The future of goal-based assurance cases. Proceedings of the Workshop on Assurance Cases - supplemental Volume of the 2004 International Conference on Dependable Systems and Networks, 2004.
- Bishop P. SILs and Software. Safety Critical Systems Club Newsletter. 2005;14:2.
- Carter A-L. Safety-Critical Versus Security-Critical Software. British Computer Society; 2010.
- Chinneck P, Pumfrey D, Kelly T. Turning Up the HEAT on Safety Case Construction. Proceedings of the Twelfth Safety-critical Systems Symposium, 2004.
- Cockram TJ, Lautieri SR. Combining Security and Safety Principles in Practice. 2nd IET International Conference on System Safety, London, 2007.
- Dobbing B, Lautieri S. SafSec Methodology: Standard. 3.1 ed., Altran Praxis; 2006.
- Goodenough J, Lipson H, Weinstock C. Arguing Security - Creating Security Assurance Cases. 2007. <https://buildsecurityin.us-cert.gov/bsi/articles/knowledge/assurance/643-BSI.html>
- Gutgarts PB, Temin A. Security-critical versus safety-critical software. IEEE International Conference on Technologies for Homeland Security (HST), Waltham, MA 2010.
- Haddon-Cave C. The Nimrod Review. 2009.
- Hawkins R, Kelly T, Knight J, Graydon P. A New Approach to Creating Clear Safety Arguments. Proceedings of the Safety-critical Systems Symposium, 2011.
- Hawkins R, Kelly T. A Structured Approach to Selecting and Justifying Software Safety Evidence. Proceedings of the 5th IET System Safety Conference, Manchester, 2010.
- Hawkins R, Kelly T. A Systematic Approach for Developing Software Safety Arguments. Proceedings of the 27th International Systems Safety Conference, 2009.

- Hawkins R, Kelly T. Software Safety Assurance - What Is Sufficient? Proceedings of the 4th IET System Safety Conference, London, UK, 2009.
- He, Ying, and C. W. Johnson. "Generic security cases for information system security in healthcare systems." *System Safety, incorporating the Cyber Security Conference 2012, 7th IET International Conference on*. IET, 2012.
- IAWG. Short Study: SafSec Coherence. Industrial Avionics Working Group; 2007.
- Kelly T, Weaver R. The Goal Structuring Notation - A Safety Argument Notation. Proceedings of the Dependable Systems and Networks 2004 Workshop on Assurance Cases, 2004.
- Kelly T. Are 'Safety Cases' Working? Safety Critical Systems Club Newsletter. 2008;17:2
- Kelly T. Arguing Safety - A Systematic Approach to Managing Safety Cases [PhD Thesis]: University of York; 1998.
- Kelly T. Using Software Architecture Techniques to Support the Modular Certification of Safety-Critical Systems. Proceedings of the Eleventh Australian Workshop on Safety-Related Programmable Systems, Melbourne, Australia, 2005.
- King T. Two different realms: RTOS support for safety-critical vs. security-critical systems 2009. <http://www.vmcritical.com/articles/id/?4031>
- Knight, J., 2015. The importance of security cases: Proof is good, but not enough. *IEEE Security & Privacy*, 13(4), pp.73-75.
- Lautieri S, Cooper D, Jackson D. SafSec: Commonalities Between Safety and Security Assurance. Thirteenth Safety Critical Systems Symposium, Southampton, 2005.
- Lipson H, Weinstock C. Evidence of Assurance: Laying the Foundation for a Credible Security Case. 2008. <https://buildsecurityin.us-cert.gov/bsi/articles/knowledge/assurance/973-BSI.html>
- MoD Defence Standard 00-056 Issue 6 - Safety Management Requirements for Defence Systems. Ministry of Defence; 2015.
- Netkachova, K., Müller, K., Paulitsch, M. and Bloomfield, R., 2015, September. A layered approach to architecting security-informed safety cases (applied to an avionics case study). In 2015 IEEE/AIAA 34th Digital Avionics Systems Conference (DASC) (pp. 1-36). IEEE.
- RTCA. DO-178B: Software Considerations in Airborne Systems and Equipment Certification. 1992
- Thomas R. Software Assurance Using Structured Assurance Case Models. Journal of Research of the National Institute of Standards and Technology. 2010.
- Toulmin S. The Uses of Argument: Cambridge University Press; 1958.
- Yamamoto, S., Kaneko, T. and Tanaka, H., 2013, March. A proposal on security case based on common criteria. In Information and Communication Technology-EurAsia Conference (pp. 331-336). Springer Berlin Heidelberg.