City Research Online

City, University of London Institutional Repository

Citation: Netkachova, K. & Bloomfield, R. E. (2016). Security-Informed Safety. Computer, 49(6), pp. 98-102. doi: 10.1109/MC.2016.158

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/15106/>

Link to published version: <https://doi.org/10.1109/MC.2016.158>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Security informed safety: if it's not secure, it's not safe.

Kate Netkachova and Robin E Bloomfield, Adelard LLP and City University London

Society relies on the safe functioning of computer based networks and systems whether it is in transportation, in energy production, banking or in medical devices. In some sectors, notably high hazard ones, achieving and assuring safety is a relatively mature undertaking - although of course we must not be complacent [20,21]. The advent of cyber issues brings enormous challenges and changes to the traditional engineering tempo and approach. This is exacerbated by the increasing sophistication of attackers, the commoditisation of low-end attacks, the increasing vulnerabilities of digital systems as well as their connectivity - both designed and inadvertent. In our research and practice we have been considering the impact of cyber issues on safety critical and safety related computer systems¹. This article shares some of the issues and lessons learned.

Safety engineering

Safety systems are engineered depending on their criticality which ranges from less critical safety-related systems, whose malfunction may potentially compromise safety and might lead to accidents with marginal or negligible severity, to highly critical systems, whose failure or malfunction can result in death and serious injuries to people, damage to property or the environment². Achieving and assuring safety is a specialist activity and rigorous safety analysis is performed to identify and mitigate hazards. A wide spectrum of analysis techniques is deployed to achieve and assure the system. Techniques such as static source code analysis, statistical testing, reliability modelling, event trees, FMECA, FTA, formal methods and proofs are used at different stages of the system lifecycle to analyse system failures and minimise safety risks.

If it is not secure it's not safe

When dealing with safety and security terminology is important as the different communities can use the same terms to mean different concepts, and have different terms for different concepts. In this short article we clarify the difference by noting that

- safety is concerned with protecting the environment from the system whereas
- security is concerned with protecting the system from the environment.

Traditionally, security and safety have been treated as separate disciplines, with their own regulation, standards, culture and engineering. This approach is increasingly becoming infeasible as there is a growing realization that security and safety are closely interconnected: it is no longer acceptable to assume that a safety system is immune from malware because it is built using bespoke hardware and software, or that it cannot be attacked because it is separated from the outside world by an "air gap". In reality, the existence of the air gap is often a myth [1]. A safety justification, or safety case, is incomplete and unconvincing without a consideration of the impact of security.

Impact on safety of cyber

We have used the Claims, Argument, Evidence (CAE) framework (See summary in [Table 1](#) and [3][4]) to analyse the impact of security on a safety justification or safety cases and we found that a

---

¹ We use the term "safety systems" to cover all of these categories

²

significant portion of a security-informed safety case will need to change to address security explicitly [2]. In some instances this will lead to substantial changes to the design, the implementation process and the justification. For example, the following areas are particularly significant from a security perspective and need more scrutiny in a security-informed justification of a safety system:

- Supply chain integrity.
- Malicious events post deployment, that will also change in nature and scope as the threat environment changes.
- Weakening of security controls as the capability of the attacker and technology changes. This may have major impact on proposed lifetime of installed equipment and design for refurbishment and change.
- Security considerations are likely to challenge the effectiveness and independence of safety barriers.
- Design changes to address user interactions, training, configuration, and vulnerabilities. This might lead to additional functional requirements that implement security controls.
- Possible exploitation of the device/service to attack itself or others.

Don't pay twice for assurance

There are technical drivers to integrate security into safety analyses – because of the interactions and trade-offs that are necessary to consider. For example at the requirements stage we might need to consider how the security aspects of the information flow policy under attack or degraded plant conditions impact the safety or we might need to consider at the architecture level whether that highly critical third party component does have sufficient security provenance given its supply chain.

There are a variety of initiatives to integrate security into hazard analyses. We have been using security (or cyber) informed Hazops to assess architectures of industrial systems and in this we adapt this well known approaches with additional security guidewords and an enhanced multi disciplinary team. Another area where there is common ground between security and safety is in static analysis of code. Both security and safety perspectives are needed to assess the likelihood of vulnerabilities being exploited and the effectiveness and consequences of their mitigations.

There are also business drivers to integrating safety and security, as stakeholders do not want to pay twice for assurance, or worse, find they have conflicts between safety and security that significantly impact project timescales and require considerable rework or re-architecting of the system.

Need resilience as well as safety

Another finding from our analysis of the security of industrial safety critical systems is that availability becomes an issue as many systems (e.g. in rail transport, power plants) are designed to fail-stop. This safety bias makes denial of service attacks relatively easy. When we combine this with the difficulty of understanding the design basis threats and the attractiveness of many embedded system targets we must plan for incident recovery and adaptation: in other words systems need to be resilient. This is a particularly true for critical infrastructures.

Most critical infrastructures are reliant on the software-based information systems, which control their operation, monitor activities, and provide real-time response to incidents and events. The recent attack on the Ukrainian power grid demonstrated how vulnerable critical infrastructures can be to malicious actions.

On the 23rd December 2015, in Ukraine a destructive malware wrecked computers of several regional distribution power companies and wiped out sensitive control systems for parts of the power grid, causing power outages and blackouts. According to the public record, an unauthorized intrusion disconnected 7 substations (110 kV) and 23 (35 kV) substations leading to an outage for 80,000 customers. The attack was very well-coordinated and comprised of multiple different elements, including a denial-of-service to the phone systems, a direct interaction from the adversary, and the malware itself installed on workstations and servers to enable the attack [5].

This crafted attack can be considered the third public example of targeted intrusions leading to outages or physical damage available to date, along with the well-known Stuxnet and the German Steelworks facility attacks [19]. Fortunately, significant cyber-attacks causing outages and physical damage are still relatively rare. More often we see cyber attacks on corporate IT, data and operational technology. These attacks lead to financial losses, violations of privacy, reputation issues and also affect people and technology that support industrial processes opening up opportunities for designing future attacks with more severe consequences.

One recent example of such potential attack-enabling intrusion was an incident with the Israeli power grid. In January 2016, one of the employees at Israel's Electricity Authority, a government department in the country's Ministry of Energy, opened a phishing email and was infected with ransomware which subsequently spread to other computers in the network. If the problem hadn't been quickly identified and resolved, it could have easily resulted in outages and other serious consequences using control gained over power grid components [7].

Model-based approach needed to deal with scale and tempo

The drivers for a more model-based approach are a need to

- Address the scale and connectivity of systems
- Deal with uncertainties in structure and to understand and evaluate systemic risks
- Interpret and analyses incidents to guide mitigation and recovery strategies
- Provide rebuttal and commentary on events as appropriate.
- For many complex systems, especially the critical ones, it is often impossible to perform live analysis. Instead, a model of a system operating in a simulated environment is constructed.

There are many approaches to modelling infrastructure and interdependencies [12,13]. In our approach [8,9] – Probabilistic Interdependency Analysis (PIA) – we look at both qualitative and quantitative aspects of assessment. The models are in part probabilistic and part deterministic and include appropriate service models, documenting assumptions about resources, environmental impact, threats and any other factors. The modelled systems are studied with an operational environment where cyber-attacks are introduced with an explicit adversary model.

A key concept of the PIA methodology is representing the system components as continuous-time state machines. The simulation of the state machines by the PIA tool produces series of events that are then aggregated to calculate the metric of interest. Typically, the metrics are various “loss functions”, e.g. the number of failed components, the duration of non-working state of a particular component or a combined characteristic of many components' states. Statistical analysis of the metric data is enabled by repeating the simulation multiple times. Although our models are abstract and less detailed than some design models they are themselves complex software that produces non-intuitive results.

Need to trust models

This type of model-based approach and probabilistic design are fundamental to the evaluation of critical infrastructures and we need to be sure we can trust these models. The models are complicated and rely on complex software for their calculations. We have experimented with using the CAE assurance framework to support an analysis of their trustworthiness.

In this approach we first focus on defining the precise claims we are making about the systems and making the assumptions about the adversary environment explicit. This highlights the need to consider various types of attacks, defining them in terms of capability, frequency and to justify that they adequately represent the possible attacks on the system.

Having established the claim we are making of the real system under its design-basis attacks we substitute it with a claim about a model under the simulated attacks. When such a substitution is made, it is essential to justify the argument that the model is adequate for the specific purpose it is being used for. Not only should the model of the system adequately represent the actual system, the model of the usage should be realistic and the model of the environment should be adequate. The evidence we use can be varied, e.g. scientific papers, insider knowledge, external expert analysis, validation testing of the tools, benchmarks and so forth.

The CAE framework supports the justification by helping elicit the claims we are making, identifying the arguments we are using to support or refute the claims and indicating what evidence we have, if any to justify what we are claiming. It provides the basis for challenge and peer review.

The security-informed safety case would require us to provide convincing evidence that the models of the system, its usage and the environment are realistic and represent those in real life.

Need to respond to claim and counter-claim

In gathering such evidence, a detailed analysis is performed and it is not uncommon to untangle various claims and lines of argument that are not necessarily true. One well-known example of such a claim is the presentation by Hugo Teso at the Hack in the Box security summit in April 2013 [11] claiming that it is possible to hijack airplanes with an Android phone. Using a flight simulator, Teso showed off the ability to control an airplane remotely, by sending radio signals to its flight-management system to change the direction, speed and altitude of the plane. The detailed analysis of the models in a proper security-informed safety case would show that the simulated use and the modelled environment are not realistic and would help with the rebuttal that the hack demonstrated using the flight simulator would not be possible on the actual certified flight systems.

Conclusion

The engineering of computer based safety systems needs to explicitly address cyber issues and from the work we have been involved with cyber has a very significant impact on the claims we make about systems, the arguments and evidence we use to justify and challenge them. We have found that the CAE framework provides an approach to evaluate the impact and trade-offs and support innovation as we move to a more model based approach to support system complexity and the assurance tempo that is now needed.

The key elements of the Claims, Argument, Evidence (CAE) approach are:

Claims, which are assertions put forward for general acceptance. They are typically statements about a property of the system or some subsystem. Claims that are asserted as true without justification become assumptions and claims supporting an argument are called subclaims.

Arguments link the evidence to the claim. They are the “statements indicating the general ways of arguing being applied in a particular case and implicitly relied on and whose trustworthiness is well established”[Toulmin], together with the validation for the scientific and engineering laws used. In an engineering context arguments should be explicit.

Evidence that is used as the basis of the justification of the claim.

Sources of evidence may include the design, the development process, prior field experience, testing (including statistical testing), source code analysis or formal analysis.

In order to support the use of CAE, a graphical notation is used to describe the interrelationship of the claims, argument and evidence.

In practice the desired top claims we wish to make such as “the system is adequately secure” are too vague or are not directly supported or refuted by evidence. It is therefore necessary to develop them into subclaims until the final nodes of the assessment can be directly supported (or refuted) with evidence. The basic concepts of CAE are supported by an international standard [3] and industry guidance [4]. A recent comprehensive review [15,16] and analysis of assurance cases is provided by John Rushby in [14].

In the light of an empirical analysis of actual safety cases, we identified a number of basic building blocks (CAE blocks) that can form the basis for describing the assessment. The blocks are:

- Concretion blocks;
- Substitution blocks;
- Decomposition blocks;
- Calculation blocks;
- Evidence Incorporation blocks.

The resulting CAE structure supports the assessment being made, but in addition, there will be important narrative and analyses explaining and detailing the claims and arguments being made. Narrative is an essential part of the assessment.

Table 1: The CAE Framework

Bibliography

1. See for example, the recent report https://www.chathamhouse.org/sites/files/chathamhouse/field/field_document/20151005_CyberSecurityNuclearBaylorBruntLivingstoneUpdate.pdf or DHS evidence “Hearing Before

- The Subcommittee On National Security, Homeland Defense And Foreign Operations Of The Committee On Oversight And Government Reform House Of Representatives One Hundred Twelfth Congress First Session, May 25, 2011, Serial No. 112–55”
2. Bloomfield, R. E., Netkachova, K. & Stroud, R. (2013). Security-Informed Safety: If it's not secure, it's not safe. Paper presented at the 5th International Workshop on Software Engineering for Resilient Systems (SERENE 2013), 03rd - 04th October 2013, Kiev, Ukraine
  3. Bishop, P.G., Bloomfield, R.E.: A Methodology for Safety Case Development. In: Redmill, F., Anderson, T. (eds.) *Industrial Perspectives of Safety-critical Systems: Proceedings of the Sixth Safety-Critical Systems Symposium*, Birmingham 1998, pp. 194–203. Springer, London (1998)
  4. ISO/IEC 15026-2:2011. Systems and software engineering — Systems and software assurance, Part 2: Assurance case (2011)
  5. <https://ics.sans.org/blog/2016/01/09/confirmation-of-a-coordinated-attack-on-the-ukrainian-power-grid>
  6. <http://edition.cnn.com/2016/02/11/politics/ukraine-power-grid-attack-russia-us/>
  7. <https://ics.sans.org/blog/2016/01/27/context-for-the-claim-of-a-cyber-attack-on-the-israeli-electric-grid>
  8. Bloomfield, R.E., et al.: Preliminary Interdependency Analysis (PIA): Method and tool support. p. 56. Adelard LLP (2010)
  9. Bloomfield, R.E., Chozos, N., Nobles, P: Infrastructure interdependency analysis: Requirements, capabilities and strategy. Adelard document reference: d418/12101/3, issue 1 (2009).
  10. Netkachova, K., Bloomfield, R., Popov, P., Netkachov, O. “Using Structured Assurance Case Approach to Analyse Security and Reliability of Critical Infrastructures”, *Lecture Notes in Computer Science, Computer Safety Reliability and Security*, pp 345-354, doi:10.1007/978-3-319-24249-1_30
  11. http://edition.cnn.com/2013/04/11/tech/mobile/phone-hijack-plane/index.html?hpt=hp_t2
  12. Bloomfield, R., N. Chozos, and P. Nobles, *Infrastructure interdependency analysis: Introductory research review*. 2009, Adelard LLP.
  13. Utne, I.B., P. Hokstad, and J. Vatn, *A method for risk modeling of interdependencies in critical infrastructures*. *Reliability Engineering and System Safety*, 2011. **96**(6): p. 671-678
  14. [Rushby, J. The Interpretation and Evaluation of Assurance Cases, Technical Report SRI-CSL-15-01, July 2105
  15. Bloomfield, R. and Netkachova, K. “Building blocks for assurance cases”, S 2014 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), 3-6 Nov. 2014 doi: 10.1109/ISSREW.2014.72
  16. Netkachova, K., Netkachov, O., Bloomfield, R. “Tool Support for Assurance Case Building Blocks, Providing a Helping Hand with CAE”, *Lecture Notes in Computer Science, Computer Safety Reliability and Security*, pp 62-71, doi: 10.1007/978-3-319-24249-1_6
  17. See for example https://ics.sans.org/media/ICS-CPPE-case-Study-2-German-Steelworks_Facility.pdf, <http://www.wired.com/2015/01/german-steel-mill-hack-destruction/> and the original BSI report
  18. Automotive software issues. Toyota
<http://www.edn.com/design/automotive/4423428/Toyota-s-killer-firmware--Bad-design-and-its-consequencesand and Koopmans analysis>
https://users.ece.cmu.edu/~koopman/pubs/koopman14_toyota_ua_slides.pdf

19. Problems with medical infusion pumps led to this FDA initiative
<http://www.fda.gov/medicaldevices/productsandmedicalprocedures/generalhospitaldeviceandsupplies/infusionpumps/ucm202501.htm>