

Model-based Security Risk Analysis for Networked Embedded Systems

Maria Vasilevskaya and Simin Nadjm-Tehrani

Dept. of Computer and Information Science
Linköping University, Sweden

Email: [maria.vasilevskaya, simin.nadjm-tehrani]@liu.se

Abstract. Finding a balance between functional and non-functional requirements and resources in embedded systems has always been a challenge. What brings this challenge into a sharper focus is that embedded devices are increasingly deployed in many networked applications, some of which will form the backbone of the critical information infrastructures on which we all depend. The Security-Enhanced Embedded system Development (SEED) process has proposed a set of tools that bridge the two islands of expertise, the engineers specialised in embedded systems development and the security experts. This paper identifies a gap in the tool chain that links the identification of assets to be protected to the associated security risks seen from different stakeholder perspectives. The needed tool support for systematic prioritisation of identified assets, and the selection of security building blocks at design stage based on a risk picture of different stakeholders, are characterised. The ideas are illustrated in a smart metering infrastructure scenario.

1 Introduction

Meeting the security needs of the society and the privacy needs of the individual users of networked information systems is a subject for current active discussion. While the generics of this challenging problem are being discussed by a spectrum of scholars in an interdisciplinary manner, the technical development of new types of systems and infrastructures is ongoing in parallel, with more applications realised as networked embedded systems. The forthcoming vehicular networks and smart grid infrastructures are examples of such a technological development with economic sectors driving the development, waiting for the societal and regulatory dimensions to catch up.

Embedded systems add new challenges to the existing map of security landscape since embedded systems were until very recently isolated from the rest of information infrastructures, and their potential threat to societal and personal security was both limited and local. With the advent of Internet of Things (IoT) and higher rate of absorption of embedded devices in current applications, this premise no longer holds. The earlier adopted approach of “adding on security” which was already shown to be not an effective technique for enterprise systems is definitely not an option in future IoT security. Hence, it is essential to address

systematic approaches to development of networked applications that include embedded devices.

A recent European project, SecFutur, combines reusable building blocks and a systematic process for constructing security-enhanced embedded systems [1]. In particular, we recognise that the security experts are largely outnumbered by the embedded systems engineers, and that the combination of the two expertise in every variant of networked embedded systems – a sector highly driven by economic returns, cost, size, or other form factors – is difficult to achieve in an efficient manner. We have therefore proposed a new process – called Security-Enhanced Embedded system Development (SEED) – that exploits security experts’ knowledge in ontological repositories, to help a developer of an embedded networked system with no/little access to security expertise [2].

The proposed process starts with a (UML) functional model of a networked system on the one hand, and the knowledge captured about the security requirements in a *domain* on the other hand. To bridge the embedded systems and security worlds, we employ domain-specific modelling and ontology technologies. This process is supported by tools that 1) systematically search for involved assets in the functional models, and 2) systematically find countermeasures through the ontology-based repositories.

Our study of the gaps in SEED points towards the need for a link between existing risk analysis techniques and model-based system development process. The current paper asks new questions about their applicability in a critical infrastructure context, namely:

- How can the asset-driven assessment of required security properties be complemented by tools and methods that prioritise and select relevant assets?
- How can the stakeholder perspectives be utilised in deciding a higher or lower level of security within the design exploration space?
- Once the relevant assets are identified, how is the selection of security building blocks to protect them affected by the same stakeholder perspectives?

Section 2 describes the proposed method for prioritisation of assets and selection of countermeasures by linking to stakeholder profiles that addresses the questions stated above. We briefly illustrate the motivation for planned tool extensions by describing its application to a part of the smart grid infrastructure in Section 3. Exposure to the mathematical base of this method requires a lengthy account that is outside the scope of this paper. However, this paper shows that the step from that quantification to actual selection of countermeasures is highly dependent on a new component, namely the stakeholder perspective. Section 4 sums up the paper and provides directions for future work.

2 Linking Stakeholders and Risks

Critical information infrastructures have a lot of stakeholders whose preferences should be accounted for when deciding the appropriate level of security to demand during design stages. In the telecommunications sector for example, there

is a growing number of end user devices all with their own characteristics (incorporating software and hardware from many different vendors), a number of network operators (wired and wireless), a number of communication system vendors (supporting various access technologies and incarnations of the same standards), as well as regulatory authorities that have their national interests. Similar characteristics are emerging in the vehicular telematics networks with both entertainment and automotive value-added functions emerging side-by-side with functions that enhance societal interests (e.g. the e-Call standard proposed in the European Union). Our earlier application of SEED has been in the smart grid domain, undergoing similar multi-perspective development.

This section refines the step that associates a measure for security. This measure rests on two components: assets automatically extracted within a system model and a set of stakeholders for the considered application. As an example of assets, our security ontology includes two types of assets: data that is in storage and data that is in transit between subsystems. Focusing on confidentiality and integrity as security goals, our current challenge is how to associate a “number” that characterises the absence of protection, e.g. integrity loss associated with an asset.

Here we envisage that the classic notion of risk [3] can be exploited. More specifically, we will consider confidentiality loss and integrity loss in the vein of a risk that needs to be averted. Hence, we will associate with the metric the two elements *likelihood* and *consequence*. While the simpler part is association of consequence with an asset, the association of likelihood of a breach of security in our model-based vision is computationally intensive.

The consequence assessment part of risk evaluation is typically carried out in consultation with stakeholders. For example, tools like CORAS [4] are formed around eliciting the costs of breaching security in connection with each asset. The notion of cost varies from one application domain to another and from one asset to another, but also from one stakeholder to another. There are also different costs depending on which security goal is violated. For example, for a utility provider as stakeholder the breach of integrity for end user electricity measurements are usually associated with high costs, while customer privacy (confidentiality) may have a lower relative priority. The right hand side of Fig. 1 visualises this idea. This stakeholder-parameterised version of consequence assessment can then be used in the decision process arriving at which asset(s) to prioritise for protection, or even in business decisions like who should bear the initial costs of an investment in a given security solution.

To compute the likelihood element of the risk is a more elaborate and demanding activity. First, a given system design and selected platform should be coupled with relevant attack models to obtain the likelihood to violate a certain security property. The left hand side of Fig. 1 depicts this process.

In order to support the computation of the likelihood in an efficient manner one has to choose a suitable formalism. Our current work [5] involves modelling attacks as directed acyclic graphs (i.e. attack trees) so that the combination of attacker behaviour and operations of a system leading to manipulation of assets

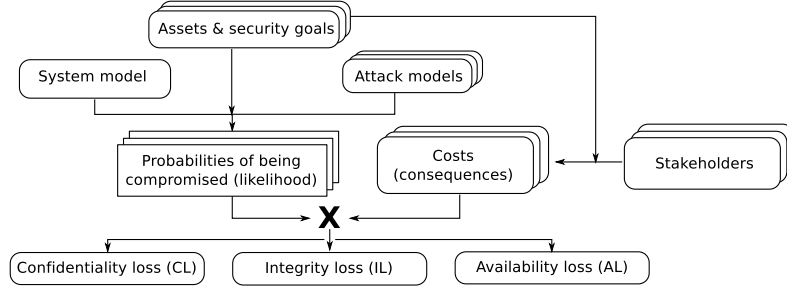


Fig. 1. Supporting the focus on relevant assets and security goals

can be probabilistically evaluated. An attack tree is an established representation of attack scenarios [6] formally defined [7] and extended with stochastic and time semantics [8, 9]. Similarly, a system is formalised as a stochastic semi-Markov model that is an intuitive and powerful tool to represent dynamic aspects of a system behaviour. Finally, the combination of the likelihood and consequence are combined for each asset, and the outcome can be used as a means of ranking/filtering the important assets and the less relevant ones.

3 Smart Grid Illustration

We illustrate the novelty of the proposed asset selection and prioritisation on the smart metering infrastructure called Trusted Sensor Network (TSN) [1]. The TSN is built of a set of metering devices referred to as Trusted Sensor Module (TSM), database servers, client applications, and a communication infrastructure. The main goal of this system is to measure energy consumption at households and to associate measurements with the clients' data for billing purposes.

The overall specification of this case study consists of seven main scenarios that have a range of diverse security considerations. Consequently, there are many assets identified in these scenarios, e.g. measurements (meter readings), a set of user account data (customer, administrator, operator), a set of certificates (calibration, installation), communication configurations, functional settings, commands, control messages, etc. Additionally, as any large system the metering infrastructure has many stakeholders.

Let us assume that a realisation of tools and techniques mentioned in the previous section enables a per-asset characterisation of integrity/confidentiality loss seen from the perspective of different stakeholders. In this section, we focus on three assets, namely measurements (denoted by A_1), certificates (A_2), and commands (A_3). We also consider three distinct stakeholders, i.e. end users, the utility provider, and the national regulatory agency.

Violation of confidentiality and integrity of these assets has different consequences for different stakeholders. For example, for a utility provider, breach of

the integrity of measurements is usually associated with high costs. A systematic misuse of the metering device can lead to manipulations at large scale and result in economic losses. However, the breach of confidentiality for the same measurement data is of a lower priority. Obviously, the picture is different for the user as a stakeholder. One can consider the national regulatory agency to be mainly interested in the availability dimension of the electricity supply and thereby, the breach of confidentiality of the measurement data has a lower consequence. On the other hand, a large scale manipulation of the commands issued to the sensor nodes, can be used in a scenario where national security is threatened.

Application of SEED allows systematically identifying the presence of above assets within a system model. Here, we propose that the calculated metrics introduced in Section 2 for all assets can be organised in a *stakeholder security profile* that shows losses for a stakeholder with respect to each asset. These profiles can be visualised as plots depicted in Fig. 2. Here, the selected assets are listed along the x-axis, and the y-axis shows the calculated confidentiality loss.




Fig. 2. Stakeholder security profile view

Next, guided by SEED, a system engineer selects a set of security building blocks (SBBs) to reduce the potential loss of security for stakeholders. Obviously, integration of any new functionality into a system will imply extra costs. In order to incorporate these costs and to distribute them among stakeholders, we need to evaluate how each stakeholder benefits when a certain SBB is integrated. We propose that the added benefit is expressed as a reduction effect that an SBB brings in terms of confidentiality (integrity, availability) loss for each asset.

As illustration, we consider three SBBs selected within the SecFutur project to be integrated into the TSM device: secure storage, anomaly detection, and secure communication. Secure storage and security communication reduce the likelihood of breaching integrity and confidentiality of stored data and transmitted data respectively. The anomaly detection, already shown to be viable in a prototype of the TSM [10], aims to reduce the likelihood of integrity loss for measurements stored in the device. Reduction effect of implemented SBBs is visualised in Fig. 2 as dashed arrows that shift the initial confidentiality loss (black dots) to lower values (grey dots). The placement of the dots and the scale of the reduction (the size of arrows) is a relative placement to visualise the intended use of the suggested techniques. This way, a system designer can analyse

which stakeholders benefit most from integration of which SBBs and consider the cost-benefit trade-off for the implementation appropriately.

4 Summary and Future Work

Society depends on critical infrastructures for its vital functions, and these increasingly rely on embedded devices for their continued operation. The shift from the proprietary, isolated development of such networked applications towards large scale integration of off the shelf units necessitates a new mindset.

Our earlier work on SEED lays out a workflow for systematic identification of security needs of a system and selection of a suitable set of security mechanisms. In this paper we have characterised a missing part of the puzzle – the justification for prioritising assets as input to selection of security mechanisms. We suggested a bridge towards the traditional concepts from risk analysis, made specific in terms of integrity, confidentiality, or availability loss. This paper outlines the path to support the missing technology. Our ongoing work creates the mathematical underpinnings for the calculation of integrity/confidentiality loss using semi-Markov models [5] and we will provide tools to support the mentioned activities in future works.

References

1. The SecFutur project: Design of Secure and Energy-efficient Embedded Systems for Future Internet Application. <http://www.secfutur.eu>.
2. Maria Vasilevskaya, Linda Ariani Gunawan, Simin Nadjm-Tehrani, and Peter Herrmann. Integrating Security Mechanisms into Embedded Systems by Domain-specific Modelling. *Journal of Security and Communication Networks*, Wiley, 2013.
3. Christopher Alberts and Audrey Dorofee. *Managing Information Security Risks: The Octave Approach*. SEI Series in Software Engineering, Addison-Wesley, 2003.
4. Folker den Braber, Ida Hogganvik, Soldal Lund, Ketik Stølen, and Fredrik Vraalsen. Model-Based Security Analysis in Seven Steps – a Guided Tour to the CORAS Method. *BT Technology Journal*, 2007.
5. Maria Vasilevskaya and Simin Nadjm-Tehrani. Quantifying Risks to Data Assets Using Formal Metrics in Embedded System Designs. Submitted, 2014.
6. Barbara Kordy, Ludovic Piètre-Cambacédès, and Patrick Schweitzer. DAG-Based Attack and Defense Modeling: Don't Miss the Forest for the Attack Trees. *Computer Science Review*, Elsevier, 2014.
7. Sjouke Mauw and Martijn Oostdijk. Foundations of Attack Trees. In *Information Security and Cryptology (ICISC)*, Springer Berlin Heidelberg, 2006.
8. Florian Arnold, Holger Hermanns, Reza Pulungan, and Mariëlle Stoelinga. Time-dependent analysis of attacks. In *Principles of Security and Trust (POST)*, Springer Berlin Heidelberg, 2014.
9. Jaafar Almasizadeh and Mohammad Abdollahi Azgomi. A Stochastic Model of Attack Process for the Evaluation of Security Metrics. *Journal Computer Networks*, Elsevier, 2013.
10. Massimiliano Raciti and Simin Nadjm-Tehrani. Embedded Cyber-Physical Anomaly Detection in Smart Meters. In *Critical Information Infrastructures Security (CRITIS)*, Springer Berlin Heidelberg. Springer Berlin Heidelberg, 2012.