CHAPTER 16

An Example for Safety and Security

Safety standards require one to look at an IACS system, identify the hazards, perform a risk analysis on them, and if the risk of a specific hazard is unacceptable, to introduce supplementary fsystem functions called safety functions to mitigate the risk to the point at which it becomes acceptable. Cybersecurity standards require one to divide the physical system architecture into areas called zones and communications channels called conduits between them. A zone is a collection of assets (physical or logical objects) which share common cybersecurity requirements; a conduit is a cybersecurity-protected group of communication channels between two zones.

The "Siemens Proposal" of August 2017 suggests that cybersecurity and safety can be treated largely separately for IACS, that an environment shall be secured within which safety engineering can be performed without regard to additional cybersecurity measures [7]. An alternative view says that safety and cybersecurity issues are inevitably intertwined in IACS. An example of a view consistent with this is that of the guidance document VDE-AR-E 2802-10-1, which says to distinguish between cybersecurity for safety functions, and cybersecurity for teleological functionality[1]. Both views promote cybersecurity by design

I have participated in recent discussion of the "Siemens View" (i.e., the various

1 A plant is usually build to serve a specific purpose. An electricity-generation plant is built to generate electricity. A pumping plant is build to pump fluid, usually water, from one place to another. Functionality which is directly intended to serve this purpose I term teleological functionality. In German, for example in VDE-AR-E 2802-10-1 it is called "business- economic functionality" [8].

views expressed in the Siemens Proposal) and the contrasting view presented in the guidelines VDE-AR-E 2802-10-1 [8]. I am not sure that a discussion on the level of principles will lead to resolution of differences. I suggest that a discussion on how cybersecurity by design interacts (or does not) with safety requirements on concrete examples is likely to be more fruitful. In that spirit, here is an example.

16.1 An Example

System. A cogeneration plant with four industrial turbines. Each turbine is housed in a separate containment building. The reasons for separation are damage-limitation in case of accident, respectively availability. The operator and its insurer have experience with the damage caused when a running turbine disintegrates, and have agreed not to expose their turbines to common-cause failure by putting, say, two in each containment. Concerning availability, with two turbines in a containment building, required maintenance on one turbine would mean taking out half your capacity, because operational safety would require all turbines in one building to be shut down when personnel are present. With one turbine per containment, only a quarter of capacity is taken out for maintenance on one turbine rather than half.

Control. The control system for the turbines, the IACS, is a moderately complex piece of SW running on diverse digital-HW elements. It is distributed, in that various parts of the functionality are necessarily located in different parts of the plant, as follows. There are sensors and actuators on turbines, which are housed in separate buildings. Requirements for load balancing necessitate that the control systems for separate turbines are coordinated, which requires some communication paths between a load-balancing controller and the kit installed on/around each individual turbine. Since the turbines are in physically-separated containment buildings, this involves a computer network with distributed components.

Centralised Command Processing. The central IACS processing, for example, is housed in a facility protected from the consequences of turbine disintegration (again, for reasons of availability). It is centralised for well-understood reasons:

System updates (SW “maintenance”) need only be accomplished once, on one SW system, and not, say, four times as would be the case if each turbine had its

own control, risking diverse and possibly incompatible control states.

Load balancing cannot be achieved without considering all four turbines as a single system, so there needs to be at least one system which controls some aspects of all four turbines and their interactions as a unity.

A central IACS is necessarily considered as one zone, whereas distributed control would require not only four zones, but conduits between them and coordinated cybersecurity maintenance. There are well-known engineering reasons stemming from high-reliability and high-availability considerations for using one carefully-designed, built and maintained control system rather than so-called "n-version" duplicated systems (see, for example, Y.C. (Bob) Yeh, Design Considerations in Boeing 777 Fly-By-Wire Computers, in Proceedings of the Third International IEEE High-Assurance Systems Engineering Symposium, 1998, available from https://citemaster.net/get/02baa57e-f4b3-11e3-b859-00163e009cc7/yeh98_777-fbw.pdf ; or Gregg F. Bartley, Boeing B-777: Fly-By-Wire Flight Controls, in The Avionics Handbook, CRC Press 2001. Available at http://www.davi.ws/avionics/TheAvionicsHandbook_Cap_11.pdf).

Distributed, networked system. The control system is dependent on feedback on the state of the turbines and generation equipment. This information comes primarily from sensors placed on or near the hardware, which is not where the main control processor sits, as explained above. Thus is the IACS distributed. A distributed system requires a network to connect its physically distributed parts. This network requires not only cabling for the pure electrical transmission of information, but repeaters to maintain signal strength, in particular in the electromagnetically-intense environment which typically accompanies large-scale electricity generation equipment. These repeaters would typically also be involved in auxiliary processing, for example,

- multiplexing/concentration of signals from diverse system elements;
- maybe also some information simplification and preliminary processing;
- ECC – bit-level error checking and correction.

Such repeaters are typically all the same (model and version) from one supplier, originally configured by the system integrator and maintained under responsibility of the plant operator. They are typically called "intelligent switches". Two important

reasons for uniformity are

- reduction of complexity, and
- relative ease of maintenance (maintenance here largely consists of SW and firmware updates).

Control commands issued by the central control processor pass through the network, and the "intelligent switches" on their way to the actuators. I take it as an assumption that

(*) there are control commands which can theoretically be issued by the central processor which would cause hazard conditions/accidents with the turbines in many turbine states.

16.2 Analytical Requirements

The international standard governing safety concerning E/E/PE subsystems of a process plant is IEC 61511 [2]. Concerning software-based systems, IEC 61511 defers largely to IEC 61508. Cybersecurity requirements in IEC 61508 and IEC 61511 have been listed in Chapter 13. To remind the reader,

- n IEC 61508:2010, there are two subclauses in Part 1 (the general system part) which refer to security [1], namely
- IEC 61508-1:2010 subclause 7.4.2.3 says that if malevolent or unauthorised action is reasonably foreseeable, a "*security threats analysis*" shall be carried out.
- IEC 61508-1:2010 subclause 7.5.2.2 says that, if threats are identified, a "*vulnerability analysis*" shall be carried out.

The terms "security threats analysis" and "vulnerability analysis" are left undefined; neither is there any reference to sources which explain how such analyses are to be carried out and what their purposes are. Neither, surprisingly, is there any requirement to do anything about threats or vulnerabilities which are thereby discovered (presuming that the point of an analysis, at minimum, is to discover something).

Our analysis of the networking shows up three threats:

- MITM attacks
- DoS attacks on the network

- compromising the integrity of the software/firmware running the repeaters. This is often called “unauthorised modification”, but this raises the question of what “unauthorised” means; further, if the authorisation regime is faulty, such integrity-compromising modifications may well be “authorised”. For example, a trojan inadvertently or maliciously introduced into the supply-chain can get into an “authorised” software update from the software supplier (call this the “supply-chain problem”).

Vulnerability Scenario.

- Say a vulnerability becomes known in the intelligent switches. It enables typical man-in-the-middle behaviour as an exploit. In particular, in this installation, commands from the central processor are intercepted by switch software, and the exploit enables arbitrary commands to be forwarded in their place. Alternatively, sensor data can similarly be intercepted, modified and forwarded, causing the central processor to issue inappropriate control commands.
- A vulnerability in one of the units which puts messages on the network (switches, but also sensors and maybe other devices) could allow DoS attacks on the network.
- Compromised integrity of the software/firmware on the switches could enable them to behave differently than specified, altering message contents or intended receivers, for example.

Risk. The risk is the highest possible, given assumption (*): common-cause destructive failure of all four turbines.

16.3 Some Exercises

We may consider some exercises to be performed on the example. I encourage the reader to perform the exercises herself, before reading my commentary below.

1. Analyse, formulate requirements for, and (high-level-)design the switching system using the Siemens View.
2. Analyse, formulate requirements for, and (high-level-)design the switching system using VDE-AR-E 2802-10-1 guidance.

3. Analyse, formulate requirements for, and (high-level-)design the switching system using safety-and-cybersecurity engineering best practice. (Hint. You are welcome to think of using technology from other engineering areas: see, for example, [4], or, more recently, [6].)

16.4 Some Solutions

16.4.1 The Siemens View

Roughly speaking, the view entails a two-step cybersecurity&safety process, that

- security considerations (from, say, IEC 62443 [3]) are applied first, to enable a "cybersecure environment" (or "Security Environment" as I understand it is being called).
- In this cybersecure environment, the safety functions may be designed and implemented without further attention paid to cybersecurity.

I understand from talking to Siemens engineers that there is an important caveat, namely that this separation shall be pursued as far as possible, but it is acknowledged that there are some situations in which the separation will not be as perfect as envisaged above. This reminds me of Einstein's dictum "*Everything should be made as simple as possible, but no simpler*". Adapted, it could read "*Everything should be separated as far as possible, but no further*". I take the caveat "*but no further*" to be non-trivial.

Let us assume that the first step in the ideal separation has been performed, that the cybersecurity of the network and repeaters is assured, according to principles, say from IEC 62443. Let us look at the second step. For safety, what is required is:

- lossless transmission of messages sensor $\longrightarrow$ control processor and control processor $\longrightarrow$ sensor
- content-changeless transmission of such messages
- multiple pathways sender $\longrightarrow$ receiver in case a repeater or transmission cable is non-functioning

This allows the following technology:

- message content transmitted in clear text

- visible addressing (and other forms of sensitive metadata)
- opportunistic message transmission (e.g., Ethernet) instead of round-robin/time-triggered protocols (opportunistic is generally "lighter-weight" when messages are not time-dense)
- ECC to address inadvertent change of content and rectification

In addition, requirements on the software/firmware driving the repeater are:

- Functional correctness of any installed code for message disassembly/reassembly

A non-requirement on the software/firmware is:

- No repeating validation of the continued integrity of the code

This solution fulfils the cybersecurity-then-safety Siemens View. However, I find it highly unsatisfactory. I cannot think of any design in a cybersecurity-sensitive environment in which I would consider it a good idea to have

- message content transmitted in clear text
- visible addressing (and other forms of sensitive metadata)
- No repeating validation of the continued integrity of the code

Sending messages and metadata in clear text is an invitation to MITM. A way in which this could be inhibited at minimal damage to the model is via end-to-end message encryption, including metadata. But this means metadata, including sender and receiver, are part of the encrypted transmission and it is hard to see how receivers could identify messages for themselves, or repeaters could engage in disassembly/reassembly of messages for multiplexing and validation purposes, unless there were shared keys amongst all operative nodes on the network. This in itself requires symmetric encryption and dependable distributed key management.

How the continued integrity of the installed operational code is to be assured, without regular validation, is not clear to me.

This approach seems to obviate any need to consider DoS attacks at the safety-design stage. Such attacks may be assumed to be stopped at the boundary of the cybersecure environment to which the network belongs. It is defined to be someone else's problem as far as safety engineering is concerned. I am not at all sure this would reflect best practice in the design of resilient networks.

16.4.2 The VDE-AR-E- 2802-10-1 Approach

The “2802 Approach” (as I will call it) proceeds as follows. We are assuming that the hazard analysis identifies non-veridical message transmission, respectively non-dependable/veridical disassembly/reassembly of messages as a high-severity hazard. Given the capacity of MITM attacks to achieve this state of affairs, there is derived a cybersecurity-for-safety requirement that MITM attacks and other attacks on the required properties of lossless transmission and content-changeless transmission in the message-passing functionality should be inhibited. The requirement on veridical transmission and disassembly/reassembly of messages requires a safety function according to the precepts of IEC 61508, which should be appropriately attack-resistant. This can be achieved straightforwardly using techniques hinted at in Exercise 3, namely

- no addresses in metadata,
- no addressing inferable from timing-slot analysis,
- encrypted content,
- individual sender-receiver keys, including keys at disassembly/reassembly points,
- “heartbeat” signals to indicate network continuity.

Regular validation of the integrity of installed operational code is indicated. This seems a clear example of a necessary security requirement.

Given the possibility of DoS attacks on the repeaters, the simplest solution might be to install a parallel network based on the same cybersecurity-engineering principles but different cybersecurity infrastructure. A second network would need only minimal functionality if there were a possibility to take the primary network down and reinstall/restart it with renewed cybersecurity parameters (e.g., change of keys). For a DoS attack on the primary network, if it is designed and built according to the principles above, entails that some of its cybersecurity has been compromised. Hence reinstalling with presumed-uncompromised cybersecurity parameters should inhibit continued DoS exploitation.

Designing such a reduced-functionality backup network is clearly a matter for safety engineering. Only safety analysis can establish what reduced functionality is appropriate for the length of time it takes to detect, react and restore the primary network. The need for a reduced-functionality network comes from safety considera-

tions allied with recognition of the possibility of a DoS attack. This seems to me to be a clear example of an inevitably intertwined requirement.

16.5 Conclusions

1. It is not clear at time of writing that or how the “Siemens View” guidance would guide implementers to best practice. Such practice as follows from the Siemens View could lead a safety engineer to implement contraindicated network technology without any hint of the contraindication. The view identifies no need for continued-integrity checks on repeater software/firmware. It also rules out DoS attacks by fiat. They are assumed solved at the boundary of the security environment, within which the entire network lies. A reasonable engineering solution to DoS attacks does not arise.
2. VDE-AR-E 2802-10-1 guidance on security-for-safety leads to identification of effective threat mitigation. It indicates the need for best-practice networking solutions which would mitigate MITM attacks through the switches. Continuing operational validation of the repeater software/firmware is indicated as a matter of cybersecurity best practice. DoS attacks on the switches may be mitigated through a reduced-functionality backup network and reinstallation using changed cybersecurity parameters of the primary network.
3. Personnel knowledgeable in dependable-network technology are required to identify threats and threat mitigations. Non-specialist engineers do not necessarily know about SAFEbus principles or principles of TTP in general, let alone which technologies are suitable for threat mitigation.

Bibliography

[1] International Electrotechnical Commission, *IEC 61508, Functional safety of electrical/electronic/programmable electronic safety-related systems*, 2nd Edition, 7 parts, 2010.

[2] International Electrotechnical Commission, *IEC 61511, Functional safety - Safety instrumented systems for the process industry sector*, 2nd Edition, 3 parts, plus AMD1:2017, 2016/2017.

[3] International Electrotechnical Commission, *IEC 62443 Industrial communication networks - Network and system security*, many parts, various dates.

[4] Kenneth Hoyme and Kevin Driscoll, *SAFEbus (for avionics)*, IEEE Aerospace and Electronic Systems Magazine, 8:34 – 39, 1993. Doi 10.1109/62.199819.

[5] Roman Obermaisser (ed.), *Time-Triggered Communication*, CRC Press, 2011.

[6] Michael Paulitsch and Kevin Driscoll, *SAFEbus*, Chapter 7 of [5]. Available from http://www.crcnetbase.com/doi/abs/10.1201/b11155-8

[7] Siemens AG, *Regulation and Conformity Assessment of Security Relevant Products and Solutions– a Siemens Perspective*, August 2017.

[8] VDE Verband der Elektrotechnik Elektronik Informationstechnik e.V., *VDE-AR-E 2802-10-1: Zusammenhang zwischen funktionaler Sicherheit und Informationssicherheit am Beispiel der Industrieautomation – Teil 1: Grundlagen; Relation between functional safety and IT security on the example of industrial automation – Part 1: Basic principles*, 2017.