Security for safety-critical control rooms

How to integrate two disconnected schools of thought

Christian Flachberger (*Author*)
Chief Information Security Officer
Frequentis AG
Vienna, Austria
christian.flachberger@frequentis.com

Andreas Gerstinger (*Author*)
Safety Manager
Frequentis AG
Vienna, Austria
andreas.gerstinger@frequentis.com

IT security threats have become a root cause for safety hazards. At the same time, some established IT security best practices contradict safety requirements. Operators of safety-critical control rooms may find themselves in a dilemma: should they disregard internal regulations by deploying non-certified software or configurations, or should they ignore critical security patches and run the risk of a serious incident? This paper proposes an approach to tackle the challenge of integrating safety and security of control rooms in the transportation domain. Areas of conflict and resulting challenges are described and a proposal for a harmonized approach is elaborated. This harmonized approach is based on three pillars: (a) on the security side moving away from a pure compliance based approach towards a risk based approach, (b) on the safety side moving away from a static safety understanding towards an understanding taking changing security threats into account, and (c) on practical engineering level on an architecture, which allows to apply specific safety- and security regimes to different parts of the system. This paper describes, how this approach can be implemented during different phases of the system life-cycle (design, development, integration, release, operation)

Air traffic management; railway operation, vessel traffic, control room, control centre, safety, security

I. Introduction

Operators of critical infrastructures from the transportation domain such as air traffic management, railways and vessel traffic services have a strong focus on safety and productivity. They are responsible for keeping airplanes safe (in the air, during landing and on the ground), for ensuring safety of vessels and for safe operation of high-speed trains. Successful cyber attacks can disrupt such critical procedures. Security threats have become a root cause for safety hazards, hence there is no safety without security (see Fig. 1). Existing safety assurance procedures should therefore be considered incomplete if they do not call for appropriate measures to mitigate security risks. At the same time, several established IT security best practices contradict safety requirements.

New legal frameworks, such as the NIS Directive in Europe (EU 2016/1148) put new liabilities on infrastructure operators and they now find themselves in a dilemma. For example, in certain scenarios, software-assurance regulations may conflict with security best practices around deploying critical system updates as soon as they are available. Should system operators disregard internal regulations by deploying non-certified software, or should they ignore critical security patches and run the risk of a serious incident?

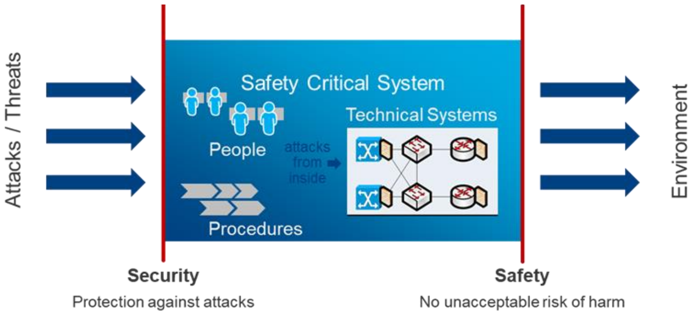

Fig. 1. The relationship between security and safety

A. *The impact of platform standardisation and system networking*

Voice communication and control systems as well as many other applications found in control rooms used for safety-critical or mission-critical applications were typically based on proprietary hardware and software running in complete isolation. Having effectively no connection to the outside world ("air gap"), such systems were less exposed to IT security risks—and even where connections existed, the high degree of customization gave external parties little chance of finding exploitable vulnerabilities.

Today, an increasing number of organizations in safety-critical industries are migrating communication systems to IP-based solutions running on commercial off-the-shelf (COTS) hardware to take advantage of significantly lower costs for acquisition and operation. These open solutions typically offer greater flexibility and usability, but not without potential downsides. Systems running on standardized platforms face threats of wide-spread, highly sophisticated malware or targeted attacks. At the same time, as systems become connected with other systems both internal and external, the potential attack surface is growing rapidly. This explains, why today safety cannot be achieved without security.

II. Challenges

Although security is essential for safety, several established IT security best practices contradict safety requirements. Achieving both goals at the same time is not necessarily easy. Let's consider some areas where these two topics can clash.

A. *"Once safe always safe" versus security adjustments on a daily basis*

Before a safety-critical system is put into operation, a formal safety assessment has to be performed concluding into a safety case. Once, system safety has been approved, the system is "sealed". If a change is needed, the safety case must be evaluated again including formal tests and analysis. In contrast to this static concept of safety, security is very dynamic: new attack vectors are identified every day and keeping a system secure means to change the system on an ongoing basis to harden it against known attack vectors.

B. *Usability*

It is a standard security practice to protect accounts by implementing two-factor authentication or a strong password policy combined with a lockout mechanism, which adds increasing delay times after entering a wrong password multiple times. From a safety point of view, such mechanisms can delay reaction of operators especially in emergency situations and might even lead to a denial-of-service situation if the system locks down because of too many wrong inputs.

C. *"Fail safe" versus "fail secure"*

This conflict is about two opposite philosophies: in case of an alert - all doors in a building are opened to allow people to leave the building (fail safe) - or all doors are locked to prohibit attackers from entering the building (fail secure). This consideration can be transferred to the IT world: does a firewall in the case of a failure deny all or allow all? Shall antivirus software be allowed to stop suspicious processes without human intervention?

D. *Redundancy and diversity versus attach surface minimzation*

Redundancy and diversity both increase safety e.g. by providing alternative communication links via diverse technology. At the same time this decreases security due to the growing attack surface through different technologies with different sorts of vulnerabilities.

However, it shall be noted, that there are also significant commonalities: Safety and security are both focused on the identification and treatment of risks, and the avoidance of faults, failures, vulnerabilities and incidents. The corresponding activities have to start early in the lifecycle, both safety and security cannot be added to a system as an afterthought – they have to be built-in. Safety and security require an appropriate culture, continuous training and strong management attention. These similarities between safety and security make a common and harmonized approach, as presented in this paper, plausible.

A comprehensive survey of approaches to combine safety and security is contained in [1] and a specific example of a combined safety and security process in a safety-critical company is presented in [2].

III. State of standardisation

Tackling such challenges is (still) hampered by today's state of standardization: As there was minimal potential for an IT security risk to have a tangible impact on safety, safety management was historically treated as a completely separate topic from IT security, with virtually no coordination at the architectural or organizational levels. Safety and security were two disconnected schools of thought which is still reflected by two disconnected worlds of standardization for safety (represented e.g. by IEC 61508 or by EUROCAE ED-153 and ED-109A specifically for air traffic management) and security (represented e.g. by NIST SP 800, BSI Grundschutz or ISO 27001). Only recently, new standards started to emerge which integrate safety and security practices to a common concept. Examples are IEC TR 63069, IEC TR 63074, partly IEC 62443, or the recently released standard EUROCAE ED-205 for air traffic management ground systems.

IV. Approach

Integrating safety and security requires a new way of thinking: moving away from an undifferentiated and pure compliance-based approach towards a differentiated and risk-based approach. The problem with a pure compliance-based approach is its limited practical feasibility and limited effectiveness. The attempt to assure security by check-marking a list of security best practices often leads to conflicts with safety requirements and at the same time may omit relevant security threats. A different approach is needed to integrate safety and security. This can be achieved by segmenting a system into different zones, where specific safety- and security regimes are applied for each zone to mitigate the overall safety and security risk.

Such an approach supports the development of an overall safety and security architecture while minimizing risks. This approach is inline with current developments of standardization, for example the standard EUROCAE ED-205 for air traffic management ground systems and other standards, which were mentioned above. These standards are driven by the insight, that protecting critical data and safety-critical processes requires, to a certain extent, different security practices and solutions. In reality, IT security best practices and solutions are now complemented by OT (operational technology) security best practices and solutions focusing on safety-critical processes. In addition to the many established IT security consultancy companies more and more security consultants specialized in OT security can be seen in the market.

We propose a harmonized approach, which is based on three pillars: (a) at the security side on moving away from an undifferentiated and pure compliance-based security approach towards a differentiated and risk-based approach; (b) on the safety side on moving away from a static safety understanding ("once safe, always safe") towards an understanding which considers security threats as possible root causes for safety hazards within the assurance model and (c) on practical engineering level on a segmentation of the system into different protection zones where specific safety and security regimes can then be applied for each zone to mitigate the overall safety and security risk in an adequate manner.

This approach paves the way towards an integrated safety and security approach, as it is for example shown in the standard EUROCAE ED-205 for air traffic management ground systems. For many organizations, it will remain advisable to maintain separated responsibilities for safety and security on organizational level to ensure the required focus on both topics. However, new procedures and a new culture of collaboration between these dedicated departments shall be established. This will finally allow - together with new, upcoming standards - to overcome the conflicts and to integrate these two disconnected schools of thought.

V. System architecture

In safety-critical environments we suggest the introduction of protection zones to apply the appropriate security concepts to the correct places within the users' network. Protection zones are defined as a collection of hardware, software and personnel with a common trust level. In many cases, three different protection zones with sufficient isolation between them is adequate: The internal zone with no direct connections to other systems, the shared zone with connections to other "trusted" networks and the public zone with connections to a non-trusted environment (e.g. public network). In a simplified form, safety-critical functionality is located in the internal zone and security best practices focusing on safety are applied here, while functionality requiring high connectivity is located in the public zone and IT security best practices focusing on data protection are applied in this zone.

Systems operated by an end-user are usually composed of a number of subsystems from different vendors. When Frequentis delivers a system, it usually comprises different protection zones with adequate isolation between them. When such a subsystem is integrated into the overall system on site, it is important to respect the defined protection zones and to connect networks and interfaces only as foreseen to trusted or non-trusted environments. Security needs to be ensured on a system level and this is a responsibility of the system operator.

To enhance security of a subsystem inside the perimeters which are separating the different protection zones, the principle of "complete mediation" may be applied. This principle says, that "every access to every object must be checked for authority." When this principle is applied, not only users are authenticated and authorized when they log in, but authorization also takes place inside the system itself between individual software services every time they exchange information with each other. This concept is important if a system is distributed and it is hard to define reliable perimeters: the paradigm in this case would be: "trust no network". To make this principle compatible with safety requirements - in case of a failure the reaction could be alerting only, or alerting plus blocking after a defined grace period, or alerting and immediate blocking. The applied mechanism depends on the protection zone.

VI. Life cycle - security as a process

State-of-the-art technical systems must be designed for safety and security from the beginning. A secure development lifecycle covers the phases design, development, integration, verification, and release. A security architecture is defined based on assumptions on the later operational environment of the system; security requirements are defined, implemented during development and integration, and tested. During the release phase, the responsibility for keeping the system secure is moving from the vendor to the operator of the system.

In the maintenance phase, the system operator needs to establish a security governance and security processes for keeping the system secure during its lifetime and system vendors provide the required support.

The activities to be done during operation can be broken down into four categories :

1. Risk management and governance
2. Protection
3. Defense
4. Resilience

VII. Security collaboration during the operation phase

Every required task for maintaining the security of a system needs to be done by somebody. Therefore, system operators should implement an Information Security Management System (ISMS). Support for agreements should ensure the required security support services and can be shaped in different ways for security tasks to be split between system operator, vendors and integrators according to individual preferences (see Fig. 2).

Some of the system operators run their own full-blown operational security team and only would ask for minimal generic security support services. Others may operate customized systems instead of standard products or may want to maintain their product releases for an extended lifetime. In these cases, customized security support services will be required. Other system operators may not want to operate their own operational security team and may want to purchase extended security support services. Whatever the specific collaboration scenario would be, all stakeholders share a common goal - to enable the system operator to prove due care to the regulator or authority at any given time.

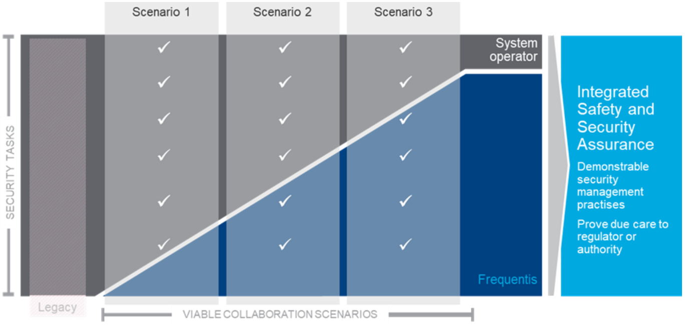
Scenario 1
Scenario 2
Scenario 3
System operator
SECURITY TASKS
✓
Frequentis
Legacy
VIABLE COLLABORATION SCENARIOS
Integrated Safety and Security Assurance
Demonstrable security management practises
Prove due care to regulator or authority

Fig. 2. Security collaboration - different scenarios for sharing duties

VIII. Maintaining security

Every modern enterprise should consider their company a fortress (Fig. 3). There's something of importance that needs to be kept safe – the crown jewels. In order to keep the crown jewels safe, you must build a perimeter and in many cases this perimeter needs to have more than one layer. Always keep in mind that attacks don’t always come from outside of the fortress, there could be traitors within. Unlike a medieval fortress, which is built once to last forever, today’s modern fortress must constantly be evolving to address evolving attack threats. And lastly there should always be a secure escape route in case all else fails and the crown jewels need to be evacuated.

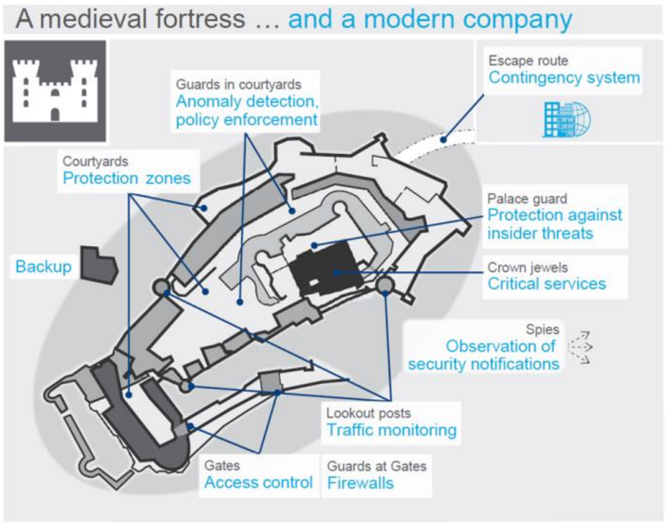

Fig. 3. Operational security tasks (a modern company compared to a medieval fortress)

A secure system which is operated in a secure way can be compared to a fortress. Walls and different courtyards are surrounding the buildings inside the fortress. Gates are connecting them. It is obvious that his architecture only provides security, if it is operated in a secure way.

Authentication and authorization needs to be done at the outer gates, but also between the courtyards (this is also referred to as defense in depth) and at the entrance to each individual building (this reflects the principle of complete mediation). Applied authorization lists need to be kept up to date. In distributed systems, where it is not easy to define and protect perimeters, these principles gain additional importance.

Guards need to be placed at the gates. They can be compared to firewalls. Firewalls need to be managed and new indicators of compromise need to be implemented into the filter rules when they get known.

Guards should also be placed in the courtyards, for example to detect enemy soldiers if they are smuggled into the fortress by a farmer delivering hay to the stables. These guards can be compared to intrusion detection systems or an anomaly detection which is performed by regular checking of log-files.

A fortress also comprises look-out posts, which can be compared to network traffic monitoring to detect threats before an intrusion has happened.

But even more pro-activity is needed: The lord of the fortress usually sends out spies to gather intelligence about new attack methods before he is hit by the attack unprepared. He might gain intelligence that another fortress was successfully attacked through the sewer system which provided a hidden access to the inside. In this case, he would immediately check his sewer system if it is wide enough to allow passing and if so, he would install iron bars. Transferred to the IT world, the system operator would regularly analyze security notifications distributed by CERTs (computer emergency response teams) to detect unknown vulnerabilities of technologies (e.g. an unprotected sewer system) and to mitigate the risk by applying security patches (e.g. installation of iron bars).

Finally, it may happen, that the look out posts detect an approaching foreign army with overwhelming power. In order to cope with this situation, every fortress is equipped with an underground escape route to evacuate people and valuable assets to another secure place. In the IT world, this means that every system can be put out of order by an overwhelming attack and it is necessary to prepare for this situation by having a contingency system at hand. Backup and recovery mechanisms allow to recover the main system again after an attack.

IX. Legacy Systems

A security health check is recommended for legacy systems. Often, for systems delivered years ago, security governance and operational security management were only partly implemented by the system operators. Such systems may still provide state of the art functionality and productivity. However, security is at risk.

Most legacy systems can be brought to an acceptable secure state. Frequentis recommends performing a security health check to determine the current status of legacy systems in use. It is important to choose a security consultant with domain expertise who understands safety and security as well as IT and OT security concepts to get a feasible and affordable solution. The international standard IEC 62443 for “Industrial Security” is a good basis for a security health check in such an environment.

X. Conclusion

Safety requires security, this is not a topic that can be overlooked. The magnitude for impact to an organization if some form of intrusion occurs can result in negative financial or physical outcomes or for the brand reputation. In safety-critical industries this could even result in a loss of life. The security of systems must be managed to comply with basic requirements for due care and requires a change in the way organizations and suppliers work together toward an increased level of collaboration.

References

[1] Siwar Kriaa, Ludovic Pietre-Cambacedes, Marc Bouissou, Yoran Halgand. A survey of approaches combining safety and security for industrial control systems. Reliability Engineering and SystemSafety 139 (2015) 156–178. Elsevier Ltd, 2015.

[2] Claudia Braun, Andreas Gerstinger, Gabriele Schedl. System Safety & Security: Establishing a Holistic Assurance Process for Safety-critical Systems. 37th International System Safety Conference, Norfolk, VA, (to be published), 2019.