

White paper: System security in a safety critical context

A common interest and collaborative effort of system operators and suppliers

Most control room operators have a strong focus on safety: air traffic management during all phases of operation, emergency services answering calls and coordinating first responders, vessel traffic and railway operation. Successful cyber-attacks can disrupt such critical procedures. IT security threats are one possible root cause for safety hazards.

Legislators globally have reacted to this problem by enacting new laws targeting Cybersecurity of vital infrastructures. Examples are the NIS Directive in Europe, the National Cybersecurity and Critical Infrastructure Protection Act in the U.S, the Security of Critical Infrastructure Act in Australia or the Cybersecurity Act in Singapore. These laws place new obligations and liabilities on the infrastructure operators to manage cyber-risks adequately. At the same time, integrating IT security practices with safety requirements is not necessarily easy as many common security practises conflict with safety. At times, organisations may feel that they are in a no-win scenario.

This white paper provides an orientation as to how security can be achieved in a safety-critical context. Additionally, it describes, what needs to be done to keep systems secure during their lifetime, what the industry can contribute and how the security can be shaped.



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, such systems were less exposed to IT security risks—and even where connections existed, the high degree of customisation gave external parties little chance of finding exploitable vulnerabilities.

Today, an increasing number of organisations 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 standardised platforms face threats of widespread, 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.

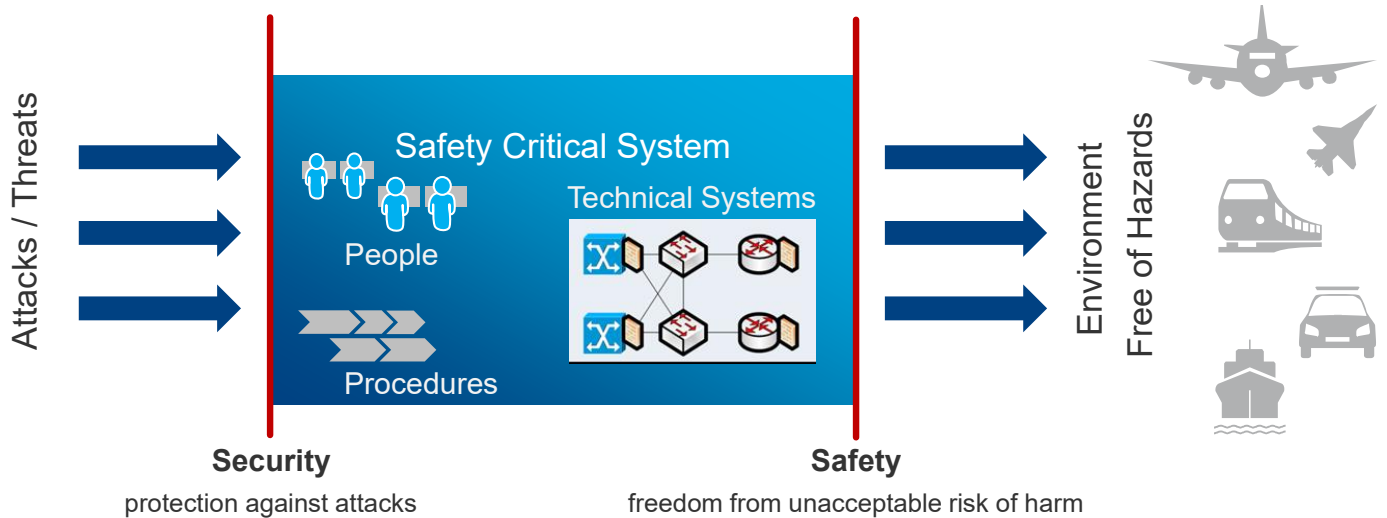
Challenges of integrating safety and security


Although security is essential for safety, several established IT security best practises 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.

Challenge 1: Once safe - always safe versus security adjustments daily

Before a safety critical system is put into operation, a formal safety assessment must be performed concluding in a safety case. Once, system safety has been approved, the system is "sealed". If a change is needed, the safety

Figure 1: The relationship between security and 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.

Challenge 2: Usability

It is a standard security practise 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.

Challenge 3: Redundancy and diversity versus attack surface minimisation

Redundancy and diversity increases 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.

Challenge 4: Antivirus

If antivirus is allowed to stop suspicious processes without human intervention, it could potentially stop safety critical functions after an update of detection patterns.

Tackling such challenges is still additionally hampered by todays' state of standardisation: 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 organisational levels. Safety and security were two disconnected schools of thought, which is still reflected by two disconnected worlds of standardisation 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 practises 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.

Figure 2: Four challenges for integrating safety and security



Tackling the challenges

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 practises 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 minimising risks. This approach is in line with current developments of standardisation, for example the new 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 practises and solutions. In reality, IT security best practises and solutions are now complemented by OT (operational technology) security best practises and solutions focusing on safety critical processes. In addition to the many established IT security consultancy companies more and more security consultants specialised in OT security can be seen in the market.

System architecture supporting safety and security

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 practises focusing on safety are applied here, while functionality requiring high connectivity is located in the public zone and IT security best practises 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 in order to ensure security also on the subsystem 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 authorised when they log in, but authorisation 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.

Lifecycle - security as a process

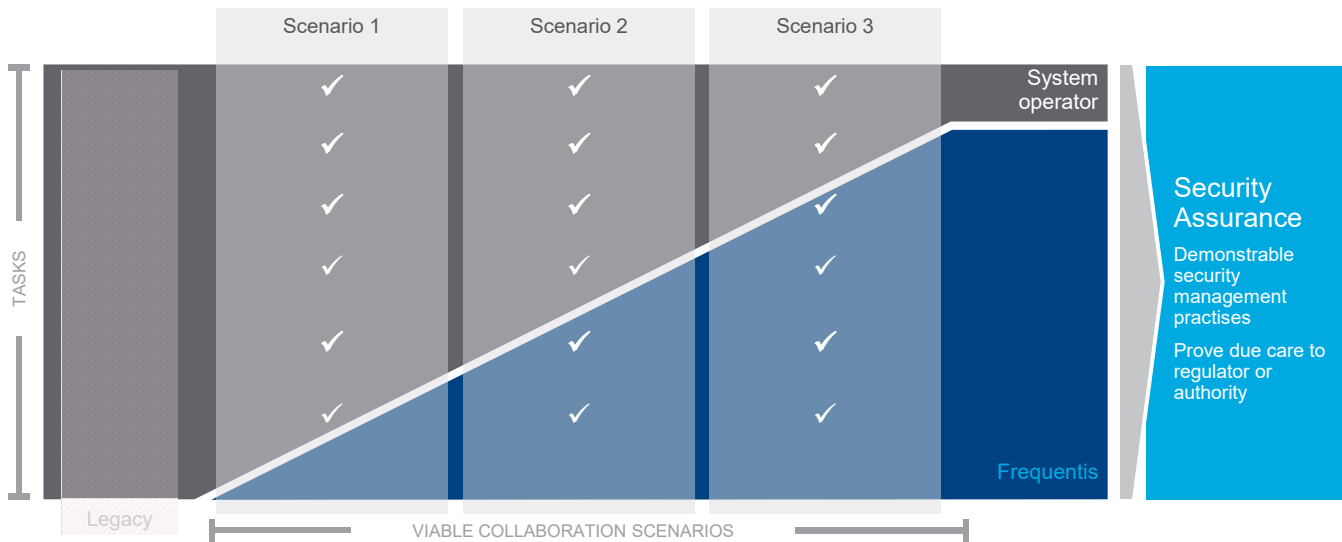
State-of-the-art technical systems must be designed for safety and security from the beginning. A secure development lifecycle covers the phases design, development, integration, verification, and release. A security architecture is defined based on assumptions on the later operational environment of 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 categories¹:

- Risk management and governance
- Protection
- Defence
- Resilience.

Figure 3: Security collaboration - different scenarios for sharing of duties



Security collaboration during the maintenance 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 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.

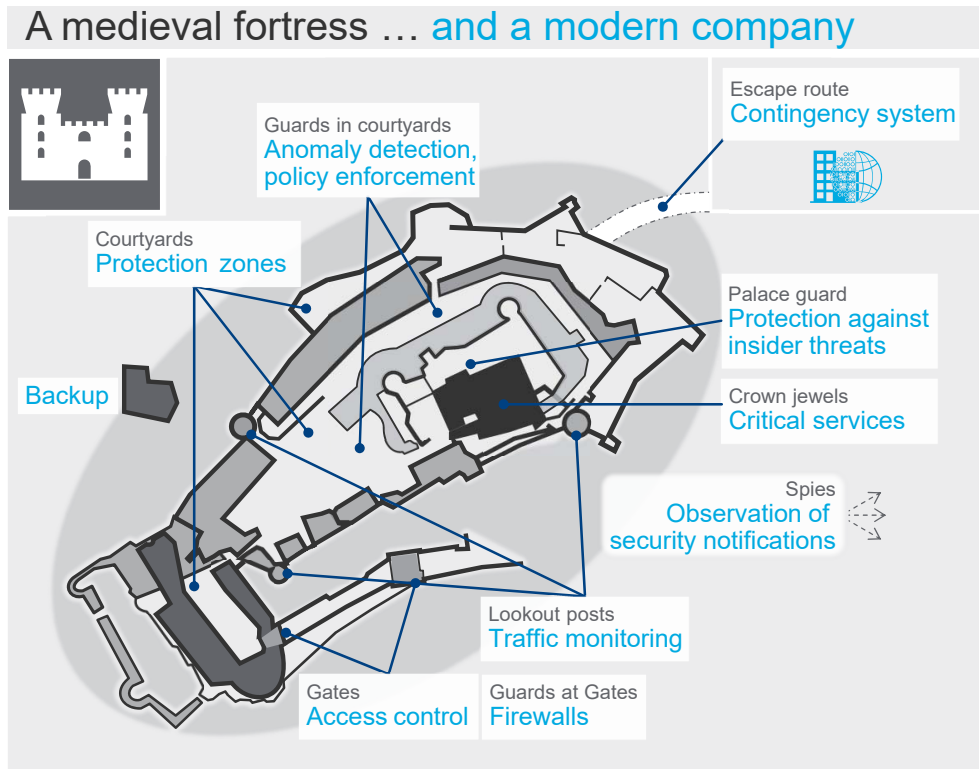
Some of the system operators run their own full-blown operational security team and only would ask for minimal generic security support services. Others may operate customised systems instead of standard products or may want to maintain their product releases for an extended lifetime. In these cases, customised 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.

Building your fortress

Every modern enterprise should consider their company a fortress. 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.

A secure system which is operated in a secure way can be compared to a fortress. Walls and different courtyards surround the buildings inside the fortress and gates connect them. It is obvious that this architecture only provides security if it is operated in a secure way. Authentication and authorisation need to be done at the outer gates, but also between the courtyards and ideally at the entrance to each individual building (this reflects the principle of complete mediation). Applied authorisation 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.

Figure 4: Building your fortress



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. These guards can be compared to intrusion detection systems or an anomaly detection which is performed by regular checking of log-files.

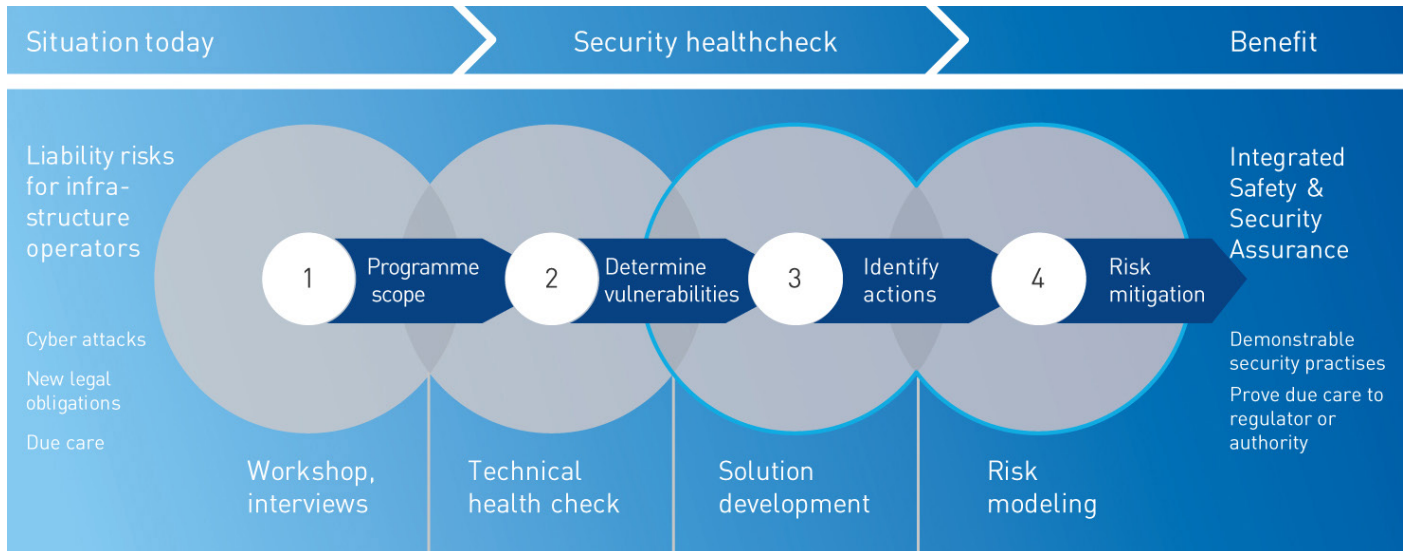
A fortress also comprises lookout posts, which can be compared to network traffic monitoring to detect threats before an intrusion has happened.

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 analyse 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 lookout 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 for the recovery of the main system after an attack.

Figure 5: Recommended procedure for a security health check



Where to begin?

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 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.

1 The activities shown here represent a common baseline of the most important laws and standards. Depending on the legislation in a specific country, additional requirements may apply. The intention of this whitepaper is not to provide a complete list of all applicable requirements worldwide, but to provide a comprehensive baseline.

Conclusion

Safety requires security; this is not a topic that can be overlooked. The magnitude for impact to an organisation 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 organisations and suppliers work together toward an increased level of collaboration. Security collaboration can be shaped in very different ways.

Contact Frequentis to learn more about the future management of system security and how we can help to ensure your organisation is safe and secure.

FREQUENTIS AG

Innovationsstraße 1
1100 Vienna, Austria
Tel: +43-1-811 50-0
www.frequentis.com

The information contained in this publication is for general information purposes only. The technical specifications and requirements are correct at the time of publication. Frequentis accepts no liability for any error or omission. Typing and printing errors reserved. The information in this publication may not be used without the express written permission of the copyright holder.