Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)
Edited by Mário P. Brito, Terje Aven, Piero Baraldi, Marko Čepin and Enrico Zio
©2023 ESREL2023 Organizers. *Published by* Research Publishing, Singapore.
doi: 10.3850/978-981-18-8071-1_P190-cd

A pragmatic mission-centric approach to ICT risk and security – Autonomous vehicles as a case

Federico Mancini

Norwegian Defence Research Establishment (FFI), Kjeller, Norway. E-mail: federico.mancini@ffi.no

Cyber security in the military domain has long been characterized by the focus on data confidentiality protection through strong system isolation and cryptography. Current warfare is more and more dependent on quick and distributed access to information, both from open and closed sources. In this setting, too much focus on reducing the risk of data leakage may lead to security solutions that, in practice, hinder access to critical information and thus reduce the actual operative effect of the Armed Forces. Additionally, as most military platforms are becoming more digitalized and interconnected, availability and integrity of steering systems should also be taking into account when securing cyber-physical systems, irrespective of their classification. This article discusses the challenge of how to integrate these concerns in a comprehensive risk management approach where trade-offs between competing security needs can be analyzed in a more systematic and traceable way.

Keywords: Risk, Metrics, Cyber security, Autonomous systems, Mission-centric security.

1. Introduction

In the military domain, cyber security has long been characterized by its focus on data confidentiality protection and strict prescriptive requirements aimed at reducing the risk of unauthorized data access through strong logical and physical isolation of the systems handling the data. Modern warfare has been challenging this approach to security.

Most military platforms and processes are becoming highly digitalized and cyberspace has been recognized as a domain of operations. As access to the right information at the right time is a fundamental requirement to achieve information superiority and conduct successful operations, the need of protecting information confidentiality must now be weighed against other operational concerns. Additionally, as most physical platforms like planes, ships, and underwater and ground vehicles are becoming unmanned or even autonomous, cyber security must also encompass the protection of physical assets.

In this context, there may emerge conflicting requirements concerning how different types of data and systems need to be protected because of the different high-level assets they support. For instance, one would want to protect the integrity of sensor data on a self-driving vehicle to prevent malicious spoofing, but the added overhead due to the authentication mechanisms might cause a delay in the vehicle's control systems and lead to a safety hazard. In lack of an optimal solution that can eliminate both types of risk, some trade-off needs to be accepted. What can be defined as an acceptable risk depends on the risk appetite of the risk owner, and it varies based on the operational context in which the vehicle is used.

We propose an approach that can help formulate risk and security controls at system level as a function of high-level organizational assets expressed in a language familiar to military decision makers: operative effect, classified information, and physical assets. This allows the formulation of structured arguments for how different security concerns compare to each other within the context of a given mission, and can constitute a better starting point to identify and evaluate alternative security trade-offs at the system level. Firstly we review some related work about the basic concepts of ICT risk and security that constitute the building blocks of our approach in Section 2. In Section 3 we describe the model underlying our approach and how to use it assets and their value. How to use it to conduct more comprehensive risk and security assessments is covered in Section 4. The

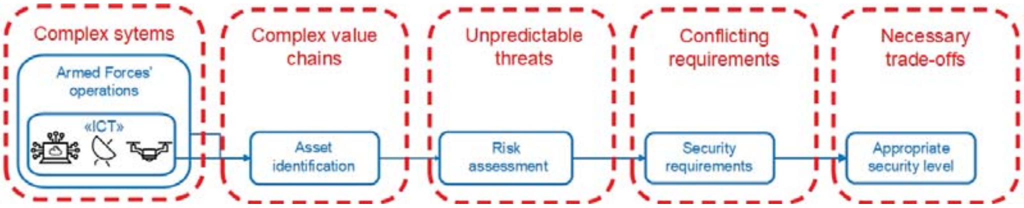

Fig. 1. The general approach to derive a risk-based security posture.

assessment of an actual military capability using our approach is discussed in Section 5. Finally, in Section 6, we summarize the main contributions of our approach and its current shortcomings.

2. Related Work

An *appropriate level of security* is achieved when the level of risk one's assets are exposed to is acceptable. What is acceptable is usually determined by a combination of external regulatory requirements and the *risk appetite* of the stakeholder(s) owning the assets. Most standards and guidelines agree that the best way to achieve this is by identifying what are the assets, assessing the risk they are exposed to, and, based on this, derive and validate appropriate security controls. For ICT security, the assets are the confidentiality, integrity and availability (C,I,A) of data and, by extension, the information systems handling the data. This approach is summarized in a very simplified form in Figure 1, as it should actually be a continuous iterative process and not a linear one. However, despite the fact that the main steps may be clear, their actual implementation is far from straightforward. The underlying reason is an inevitable limitation of risk and security assessments: we cannot precisely predict all future adverse events no matter how much information we may have at our disposal; and limited resources prevent us from implementing security controls against all possible threats. Each of the steps in the process has its own challenges that contribute to increase the uncertainty around risk and security in different ways, as summarized by the red text in the figure.

Even before any assessment, one needs to understand the ICT system in question by creating some kind of model and defining the context in which it is to be used. Already here, because the complexity of most of today's ICT systems is far too great to allow for an exhausting modeling, some uncertainty will arise. Architectural frameworks like the NATO Architectural Framework NATO C3 Board (2018) try to handle this problem through a structured decomposition of the system and the organization using it into more manageable components and relations. There are some shortcomings in this kind of approach, especially when used in the context of risk and security Grov et al. (2019), but this work is also a step in filling those gaps.

The second step requires to identify the critical assets that need to be protected. The challenge here is that it is very hard to quantify the value of an asset, as it is created by complex chains of dependencies that are difficult to understand completely. For the Armed Forces an asset has been synonymous with classified information, and value a synonymous with classification level. Now, The Norwegian Security Act Ministry of Justice and Public Security (2018) has defined that all ICT systems, data and physical objects and infrastructures that support some Fundamental National Function (FNF) are assets Ministry of Justice and Public Security (2018). In order to identify their value, a more structured model of these assets will be needed. A coherent hierarchy of assets to describe these dependencies from an organizational level down to ICT systems is defined in Endregard and Nystuen (2023), and we use it as the base of our model.

The third step, which is the classical risk assessment, is probably the more challenging and where more uncertainty arises. On the one hand, we have the uncertainty due to the complexity of the system and asset models that are to be assessed. On the other hand, predicting an opponent's motivation and course of action is much harder than modeling accidental threats. Consequence and likelihood are often indicated as the two main aspects of risk, but since likelihood is so hard to calculate for deliberate threats, an alternative model with threat, value and vulnerability has been proposed Maal et al. (2017). Both are actually inadequate to handle risk in modern ICT systems since they do not systematically consider uncertainty, but we tend to favor the latter as our approach focuses on value, and likelihood is very hard to determine for deliberate threats. We argue, namely, that being able to quantify the value of assets in a coherent way throughout the asset hierarchy is the key to harmonize the way risk appetite, risk of asset degradation and the effect of security are assessed.

The fourth step is often very prescriptive and the connection to risk appetite and risk assessment is only superficial. There might exist alternative combinations of security controls that can give equivalent levels of security, but where some assets are prioritized over others because of limited resources. Without being able to measure the effect of security on the risk associated with different assets, it is not possible to generate and compare these alternatives. In physical systems this process can be supported through simulations and physical measurements, but we are not aware of methods for ICT security that are as effective.

This is why, in the last step, it is common to use compliance with pre-defined lists of controls NIST (2020) to prove that an appropriate level of security has been achieved, together with some certification process for selected critical components Common Criteria (2022). This approach does not scale very well as systems and organization become more complex, and does not give any particular guarantee that high-level assets are adequately protected as a whole.

Other frameworks suggest ways to connect system and mission risk and security Carter et al. (2018); Rheaume (2019), but they still do not provide a way to estimate asset value in a way that can be used to guide risk and security assessments throughout the entire management cycle.

3. Underlying Model and Value Assessment

As already discussed in Section 2, in order to assess the value of system assets as a function of the value of the high-level assets they support, we need to build a more structured model of the dependencies between ICT systems and the rest of the organization, which in our case is the Norwegian Armed Forces. The model is shown in Figure 2.

At the system level, we know that the important assets are the C,I,A of data and the components handling it, while at a higher level, it is defined in the Norwegian Security Act that the critical assets are the operative effect of the Armed Forces, classified information, and physical assets like strategic installations or critical infrastructure, that support the FNFs. In order to connect the two types of asset, we use *ICT-based functions* to break down missions into tasks that can be directly associated with some ICT system component(s) that implements them. For instance, the *Information exchange* function can be directly associated with a radio or network infrastructure. Existing military taxonomies can be used to define these functions NATO C3 Board (2021).

The other important function of the model is to define a way to measure the value of the assets so that the effect of threats and security controls can be quantified as a function of it. Risk appetite can also be defined through the same metrics by setting a threshold for what reduction in value is acceptable.

Classified data is measured by using the classification levels TOP SECRET, SECRET, CONFIDENTIAL, RESTRICTED, and UNCLASSIFIED. These are indirect measures of value, as they represent the potential damage for national security if the confidentiality of the data is compromised, rather than the actual value for the Armed Forces. For now, we leave it as it is as these classifications already come with clear security requirements that need to be fulfilled.

Unlike commercial organizations, economic loss is not a good way to measure the degradation of operative effect in the military. One possibility

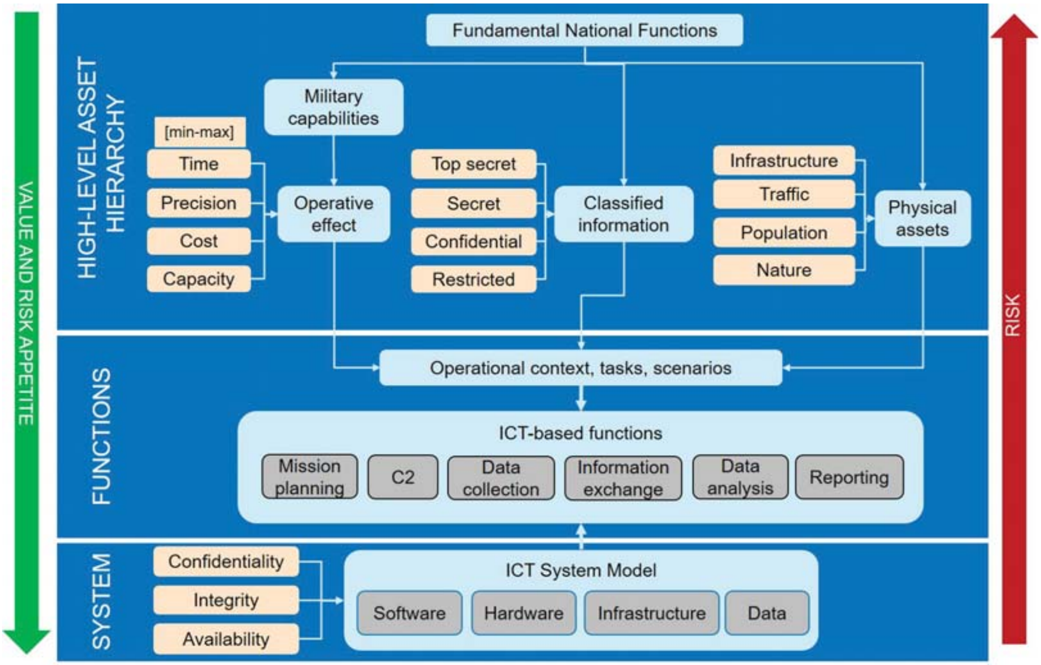

Fig. 2. A simplified overview of the model on which the approach described in this paper is based. It is in practice an architectural model focused on assets, which are modeled and linked on all layers through a pre-defined hierarchy. Their value can be defined through quantifiable parameters (in yellow) and propagated in top-down fashion through the model from high-level assets to system level assets.

we are investigating is to define some classes of measurable parameters that can be used to characterize military capabilities. Some examples of such parameter classes can be *Time*, *Precision*, *Cost*, and *Capacity*, which can be further refined for specific capabilities. For instance, Time can mean how long it takes to complete or be ready for a mission, or time to acquire some information. Cost can refer to the cost of the disposable weapons or ammunition needed to complete the mission, or the cost of training and paying for the personnel needed to operate the mission systems. Precision can refer to the the rate of target hits at a given distance for a weapon or the resolution of a sensor. Capacity could be the range a drone can fly with a given payload before battery gives out or how much data a radio link can transfer under certain conditions.

Physical assets' value is usually already estimated because of their dependency to other FNFs or safety regulation. For instance, the value of a certain area in which a mission will be conducted could have high value because of the high density of *Traffic* or *Population*, or the presence of critical *Infrastructure* or some *Natural* resources. The ICT system itself could be a physical asset to protect in some cases.

The advantage of describing high-level and system assets in a comprehensive and coherent model, is that different stakeholders in the organization can recognize those they are familiar with, and the model can be used to fill the gaps and identify the dependencies between them. So far, technical assets where given priority, without understanding their actual value for the organization, and vice-versa. Safety and security assessments would also be conducted independently from each other although they would consider the same systems, producing sub-optimal solutions. By using this model we can assess the value of high-level assets for each type of military capability in a top-down manner, and use it to inform the assessment of the data and components at system level that actually have a value for the organization as a whole.

4. Risk and Security

While value is assessed in a top-down manner, risk is aggregated bottom-up. To assess the overall risk associated with an asset, we use the threat, value and vulnerability model, rather than consequence and likelihood. Thus, the value assigned to the assets as discussed in the previous section, can be directly used to inform the risk process at system level. As for the threat and vulnerability components of risk, they are not in the scope of the paper, but we use a threat assessment model developed specifically for unmanned and autonomous military vehicles Mancini et al. (2021); Mancini, F. et al. (2023) to support the case study in Section 5. For other types of ICT systems there exist other models and threat catalogs that can be used.

What we want to focus on here, is how to manage risk associated with multiple and potentially competing high-level assets. A comprehensive threat analysis of a system will likely expose many potential ways to compromise it, but each of them needs to be associated with the right system-level asset and aggregated upwards the asset hierarchy to define the risk for each of the high-level assets. For instance, the fact that a malware on a drone manages to read some secret data, poses no direct risk to physical assets. Similarly, spoofing GPS on the drone to make it report the wrong position will not pose a direct risk to classified data stored on it, but it could have a huge impact on the mission. The model described in the previous section lays the foundation to perform this kind of analysis.

It is also important to understand how far the assessed risk is from what is acceptable (risk appetite) in order to design adequate security. As mentioned at the end of the previous section, we use the value assessment to estimate risk appetite, so it is possible to perform a direct comparison as shown in Figure 3. The arrows within the triangles represent the risk axes of each asset, which is zero at the center (green colored) and increases towards the angles (red colored). The risk associated with the assets are points on the respective axis, which form an inner triangle when joined together. The green inner triangle on the middle left figure represents the risk the stakeholder is willing to accept (risk appetite). The red inner triangle in the middle right figure is the risk estimated for the assets based on the risk assessment of the system. The two can then be used to measure how different combinations of security controls (the black arrows) can affect risk and possibly move it within the green area. An example is given in the next section.

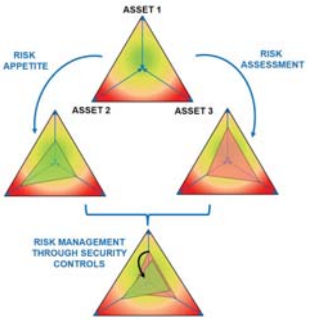

Fig. 3. Example of how the risk associated with three different high-level assets could be visualized.

One new challenge that emerges as multiple high-level assets are dependent on the same system, but require different types of protection, is that established security controls typically used to protect confidentiality, may be not very effective, or even counterproductive, in protecting safety or operative effect. This is especially clear in cyber-physical systems, as noted in the example given in the Introduction. Since the approach we presented in Section 3 to measure value can also be used to measure the effect of security on the asset, it can be used to identify the most cost-effective combinations of security controls. Whether the effect of a security control can be measured with reasonable precision, however, is not certain.

5. Autonomous Mine Hunting as a Case

The case presented here, deals with the modernization of the Mine Countermeasure (MCM) capability of the Norwegian Armed Forces. In particular, it assumes a new concept where the vessels conducting the mission are unmanned and at least partially autonomous as those shown in Figure 4.

Fig. 4. Illustration of two unmanned and autonomous vessels, one surface and one underwater, searching for mines.©FFI

This capability can be used for various types of missions, but here we focus on the continuous monitoring of sailing routes along the Norwegian coast to make sure they are mine-free and safe.

The operation consists of four main phases. Initially the sea area to search for mines is defined and the vessel should autonomously calculate the best route for an extensive search given, for instance, weather and sea-bottom conditions. Once the mission starts, the vessel will use its sensors, like sonar and camera, to collect data, and possibly adjust its route during navigation based on findings and need for data re-acquisition. The raw sensor data is then processed against databases with mine signatures or through recognition algorithms trained on such signatures. For each mine that is identified with enough confidence, its position and type will be reported. Both mine signatures and high-resolution sea-maps are typically classified. The assessments for this case are those shown in Figure 5.

This capability has been first evaluated in the context of the FNFs and what is acceptable in terms of its operative effect (OP. EFF. in the figure) has been estimated. This constitutes the risk appetite for this high-level asset as shown in the top-left triangle in Figure 5. Roughly, we can say that this is a mission that is conducted as a routine task in a low-threat scenario and that is not particularly time critical, therefore the operational effect can be relatively low and more risk is accepted. Specifically, we measured the operative effect by putting more weight on the Precision parameter, which in this case indicates the number of correctly classified mines, but less on Time used to clear one area, and medium weight on Cost in terms of how many fewer trained operators autonomy would require.

Once an operational concept and its context has been defined, it becomes clear also which classified information would be necessary to use on autonomous vessels and the potential physical assets that could be damaged during an operation, so we could estimate the risk appetite also for these other two high-level assets. The reasoning is that the confidentiality of classified data on the vessels would be critical for future mine-clearing missions in more demanding scenarios, and the consequence of losing it would not be acceptable. An adversary knowing which mines we can recognize could, namely, purposefully change their design and make them invisible to our vessels in the future. So the risk appetite for classified information (CL. INF.) is very low. Similarly, since the autonomous vessels would operate close to busy sailing routes and possibly critical underwater infrastructure, we would also want to take as little risk as possible when it comes to physical assets (PHY. ASS.).

The value assessment of the high-level assets is then used to identify the risk appetite and the system level assets in the autonomous vessels and their relative value as described in Section 3. This information is used in the risk assessment of the vessels together with a threat analysis conducted with the help of an especially designed threat model and catalog developed by FFI Mancini et al. (2021).

The risk assessment at system level is aggregated upwards for each high-level asset and summarized in the top-right triangle in red in Figure 5. It is clear that the overall risk of using fully autonomous vehicles in this kind of mission is not acceptable.

When it comes to classified information, even in a low-threat scenario, we estimated that the intrinsic vulnerabilities of an autonomous vessel used as described in the operational concept pose an unacceptable risk because of the high value of the information. Classified data stored on the

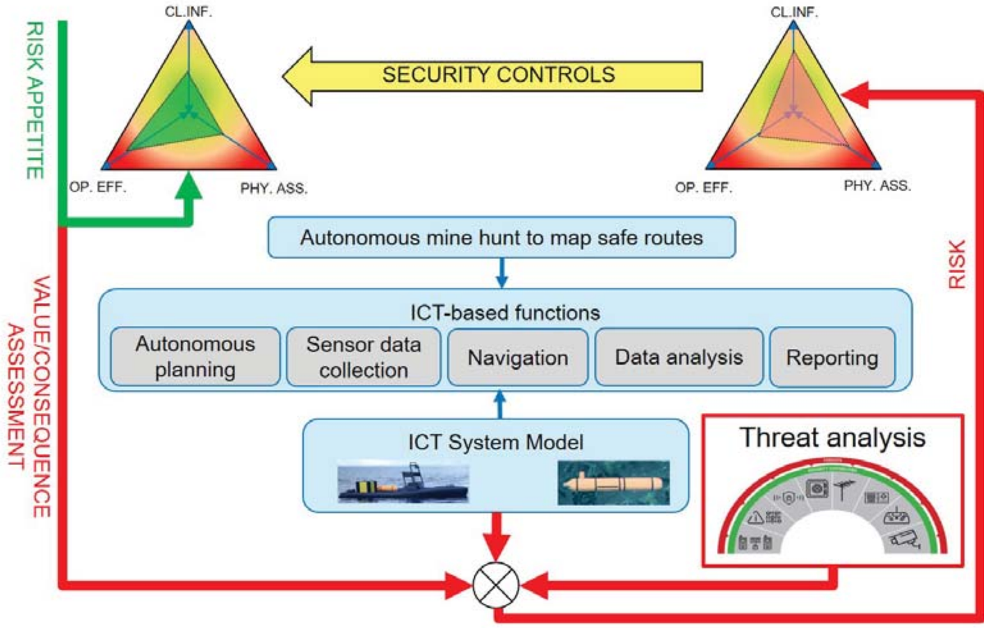

Fig. 5. Example of application of our approach in to assess the risk appetite and actual risk associated with a particular MCM mission type. This will constitute the starting point to derive appropriate security controls to reduce the risk within what has been defined as acceptable.

system could be stolen by an opponent capable and motivated enough, since they would have both sufficient time and opportunity to act on an uncontrolled vessel out on a long mission. Safety-wise, the risk is again characterized by the high value of the physical assets in the area where the vessel will operate, and the potential vulnerability of the autonomous algorithms steering the vessels, which still lack adequate robustness and reliability. However, as these algorithms improve, risk could decrease. Finally, the risk for a degraded operative effect is actually very low when everything is automated as Time and Cost improve while Precision is assumed to be almost as good as human analysts, if not better in some situations. This risk associated with this high-level asset is therefore the only one within what is acceptable.

The next step, which we do not describe here, is to apply security controls to try and reduce the risk within what is acceptable. As we discussed in Section 4, applying controls to reduce the risk on one asset type, could be detrimental to other assets. In this case, however, it is possible to sacrifice some operative effect in order to reduce the risk for the other two high-level assets. For instance, one could accept a compromise where only navigation and data collection is automated, but data analysis is done on land in a more secure location. This would increase the operation time and cost, but still within what is acceptable, while reducing drastically the risk that confidential data is lost. As for safety, intense supervised testing of the navigation algorithms could decrease the uncertainty around their reliability and lead to a new risk assessment which is within an acceptable range.

6. Conclusions and Future Work

In this paper we presented a pragmatic approach to integrate technical risk assessments of ICT systems with the overall risk management process of an organization, with focus on the Defense

Sector. Our claim, based on empirical evidences matured through years of experience in assessing risk and security for various defense projects, is that in order to achieve an appropriate level of ICT security in an organization, the value of system-level assets should be defined in a top-down manner as a function of the value of the high-level organizational assets. The reason being that this would allow to assess the level of security of ICT systems within the risk appetite of the organization. A necessary requirement to achieve this, is to define appropriate metrics to measure the value of different kinds of assets, as we illustrated in our case.

Thus, the novelty is not in the methods used to assess and model risk and security, but rather in the idea to use a coherent definition of value throughout a layered and structured model of the organization that allows to operationalize these methods in a more holistic way and to produce traceable and more understandable assessments for the stakeholders. Furthermore, considering multiple high-level assets simultaneously and using them to derive the value of system-level assets, solves in part also the problem of dealing with cyber security, safety and operation security in isolation from each other as it is often the case with more bottom-up and system-focused approaches.

The approach presented here is, at the moment, only a proof-of-concept and some gaps need to be filled before it can be used operationally. The main aspects that need to be developed next are: the metrics to measure the value of the high-level assets; a way to describe uncertainty and tie it to risk and security controls; and some tools to automate parts of the process. Especially without adequate tools, even if the theoretical foundation is correct and in place, continuous risk assessments and adjustments to the security posture will be almost impossible to perform in an effective manner because of the large amount of information and relationships that need to be handled.

References

Carter, B. T., G. Bakirtzis, C. R. Elks, and C. H. Fleming (2018). A systems approach for eliciting mission-centric security requirements. In *2018 Annual IEEE International Systems Conference (SysCon)*, pp. 1–8.

Common Criteria (2022). Common criteria for information technology security evaluation - part 1: Introduction and general model. Technical report, CCMB-2022-11-001.

Endregard, M. and K. O. Nystuen (2023). A pragmatic capability-based framework for national security risk governance. In *Proceedings of ESREL'23*. Research Publishing Services.

Grov, G., F. Mancini, and E. M. S. Mestl (2019). Challenges for risk and security modelling in enterprise architecture. In *The Practice of Enterprise Modeling: 12th IFIP Working Conference, PoEM 2019, Luxembourg, Luxembourg, November 27–29, 2019, Proceedings 12*, pp. 215–225. Springer.

Maal, M., O. Busmundrud, and M. Endregard (2017). Methodology for security risk assessments: Is there a best practice. *Risk, reliability and safety: Innovation theory and practice. London: Routledge*, 860–866.

Mancini, F., B. Greve, S. Bruvoll, and J. H. Wiik (2021). A threat model and security capabilities for autonomous military vehicles – exploring the challenges of designing and integrating security. (UO) 21/00428, FFI report.

Mancini, F. et al. (2023). Securing unmanned and autonomous vehicles for mission assurance. Technical Report Technical Report RDP STO-TR-IST-164 (NATO UNCLASSIFIED), NATO STO.

Ministry of Justice and Public Security (2018). The security act - lov-2018-06-01-24.

NATO C3 Board (2018). Nato architecture framework version 4. Technical report, AC/322-D(2018)0002-REV1.

NATO C3 Board (2021). C3 taxonomy baseline 5.0. Technical report, AC/322-D(2021)0017.

NIST (2020). Security and privacy controlsfor information systems and organizations. Technical report, NIST Special Publication 800-53 Revision 5.

Rheaume, F. (2019). Risk-based cyber mission assurance model, process and metrics. In *24th International Command and Control Research and Technology Symposium (ICCRTS)*.