24th ICCRTS - Managing Cyber Risk to Mission

Topic 1

Cyber Risk to Mission

Risk-based cyber mission assurance model, process and metrics

Francois Rheaume

Defence Research and Development Canada – Valcartier

Building 24, 2459 de la Bravoure Road, Quebec, QC

G3J 1X5

Francois.Rheaume@drdc-rddc.gc.ca

Abstract

The objective behind cyber mission assurance is to ensure that missions can be performed successfully despite operating in a cyber contested environment. This requires the ability to not only assess potential cybersecurity events, but also to assess their impacts in the first place, and to develop resilience to both the events and their impacts. Resilience is the ability to avoid, withstand or recover from potential adverse events and their impacts.

Building from existing guidelines and frameworks, this paper presents a cohesive set of tools that project managers can use to develop a cyber mission assurance program, define requirements or build a cyber mission assurance capacity. The goal is not to reinvent the CMA concepts but rather to provide a structured way to decompose the necessary CMA activities, to execute them and to measure their results. Three complementary elements are described: a layered model that structures types of risks and their relations, a process that assesses the risks and that develops the resilience, and a set of metrics to measure the effectiveness and performance of cyber mission assurance in projects. Attempts at measuring the state of cyber resilience alone are not enough; stakeholders must first measure their state of awareness about the risks of operating in the cyber space. Only on the basis of this awareness can the state of resilience be measured. The presented process and metrics, along with the underlying model, explicitly manage this correlation, therefore supporting informed decision-making during all phases of the life cycle of systems.

1. Introduction

Cyber is a complex multi-dimensional space. From hardware to software, many layers and technologies exist and they evolve constantly. If individual elements and systems of the cyber space are becoming more and more difficult to understand, so are their interactions. The challenge for organizations is to get enough understanding of their technologies, the threat behind them and the impacts on their missions, to react appropriately. Nowadays, this is increasingly important in a world where technology acts not only in the information world, like in traditional enterprise networks, but also in the physical world, like in military platforms such as aircrafts, vehicles and ships, or smaller devices that have all become heavily dependent on cyber technologies.

To deal with the potential failure of their technologies and increase mission success, organizations must first become aware of their technology dependence. This is important in that it allows them to predict the impacts of potential technology failures on their mission. They must then take actions to prevent the attacks capable of causing the predicted failures and, in assurance, they must ensure mission continuity if attacks do take place. This is what Cyber Mission Assurance (CMA) is meant to be: ensure that the mission can be accomplished successfully despite the risks of cyber-attacks.

Existing definitions of CMA are mostly derived from the Mission Assurance domain and, in many cases, in a military context [1][2][3][4]. Two principal elements come out of those definitions: 1) CMA is a risk management process, and 2) the goal is to achieve resilience. Like CMA, cyber resilience has seen a growing interest in the recent years, which has led to many definitions of the term [5][6][7][8][9]. Simply put, these can be summed up to: ‘Cyber resiliency is the ability to anticipate, withstand, recover from, and adapt to adverse conditions’[8].

In this paper, we will define CMA in terms of the 5 core security functions defined in the cybersecurity framework of the National Institute of Standards and Technology (NIST): Identify, Protect, Detect, Respond and Recover [10]. Resilience in itself will be decomposed into the Protect, Detect, Respond and

Recover functions, while the Identify function will provide the awareness necessary to develop the resilience requirements.

At present, CMA and resilience concepts are relatively mature and well-stated. Military organizations already know what CMA requires: a mission-focused continuous risk management process that supports decision-making aimed at improving resilience and increasing the probability of mission success [11][12]. The goal of this paper is not to reinvent the CMA concepts but rather to provide a structured way to decompose the necessary CMA activities, to execute them and to measure their results and the state of accomplishment. This is achieved by spelling out a CMA model, a CMA process and CMA metrics, respectively.

To assist organizations in developing and harmonizing the necessary CMA activities, a CMA model is presented in Section 2. The model is adapted from MITRE's Crown Jewel Analysis [13][14] and the NIST Risk Management Framework [15][16]. It facilitates orderly communications and results integration in the different activities and during the life cycle of projects and systems.

Based on the presented CMA model, a CMA process is described in Section 3. The process combines existing approaches, such as the ones in [17][18][10], and decomposes the underlying activities in a cohesive manner. It is suggested that cyber resiliency engineering may be viewed as a specialty discipline of systems security engineering [8]. This paper pushes the idea further, where it extends the traditional analysis of losses of Confidentiality, Integrity and Availability by including analysis of impacts on the mission aspects, and where solutions are not restricted only to technical solutions but also include operational solutions. The result is a risk-based cyber mission assurance process that analyses the mission criticality, analyses risks and mitigates them. Before launching such a process, a project manager should first seek to identify and abide by any higher-level authority laws and policies that may supersede some cited herein.

Finally, metrics are proposed in Section 4 to measure CMA. A review of existing CMA metrics has been conducted in [5]. This paper innovates in that it suggests that measuring the state of resilience is not enough, stakeholders must first measure their state of awareness about the problems and then measure the state of resilience with respect to what they know about the problems. To this task, CMA effectiveness and CMA performance metrics are introduced.

In an effort to unify the terminology and harmonize communications during the life cycle of systems, the terms 'controls', 'measures', 'mitigations' or 'solutions' will be referred to as 'requirements' and in relation to engineering concepts.

2. CMA Model

To be mission-focused, a risk management process must rely on a model that describes the relationships between the mission elements and the technology elements. To goal behind the elements of the model and their relationships is to help with the decomposition of the risk management activities and the communications of their results.

A simple CMA model is shown in Figure 1. It is inspired from the dependency model defined in MITRE's Crown Jewel Analysis [13][14] and the description of the three-tiered views in the NIST Risk Management Framework [15][16]. It relates technology assets, system functions, operational capabilities and mission objectives and allows description of the impacts of cyber threats at each layer of the model. Technology assets are the physical and logical components of systems, while system functions are the physical and logical products or effects that systems realize. A capability represents "the ability to achieve a desired effect under specified standards and conditions through combinations of means and ways to

perform a set of tasks" [19] or, equivalently, ''the ability to execute a specified course of action'' [20]. A capability could be strategic, operational or tactical and can be categorized in hierarchical ways. Military organizations usually define their capabilities in their doctrines [21][22][20] or as part of capability-based planning activities and frameworks [23][24][25][26]. For instance, the Royal Canadian Air Force (RCAF) has defined the capability of intelligence, surveillance and reconnaissance (ISR) [27], which corresponds to the operational level. Capabilities are necessary to achieve mission objectives, which can also be defined at the strategic, operational and tactical level. For instance, the Canadian Armed Forces are engaged in a number of operations, such as Operation DRIFTNET, whose mission objective is 'to stop drift netting and other forms of illegal fishing' [28].

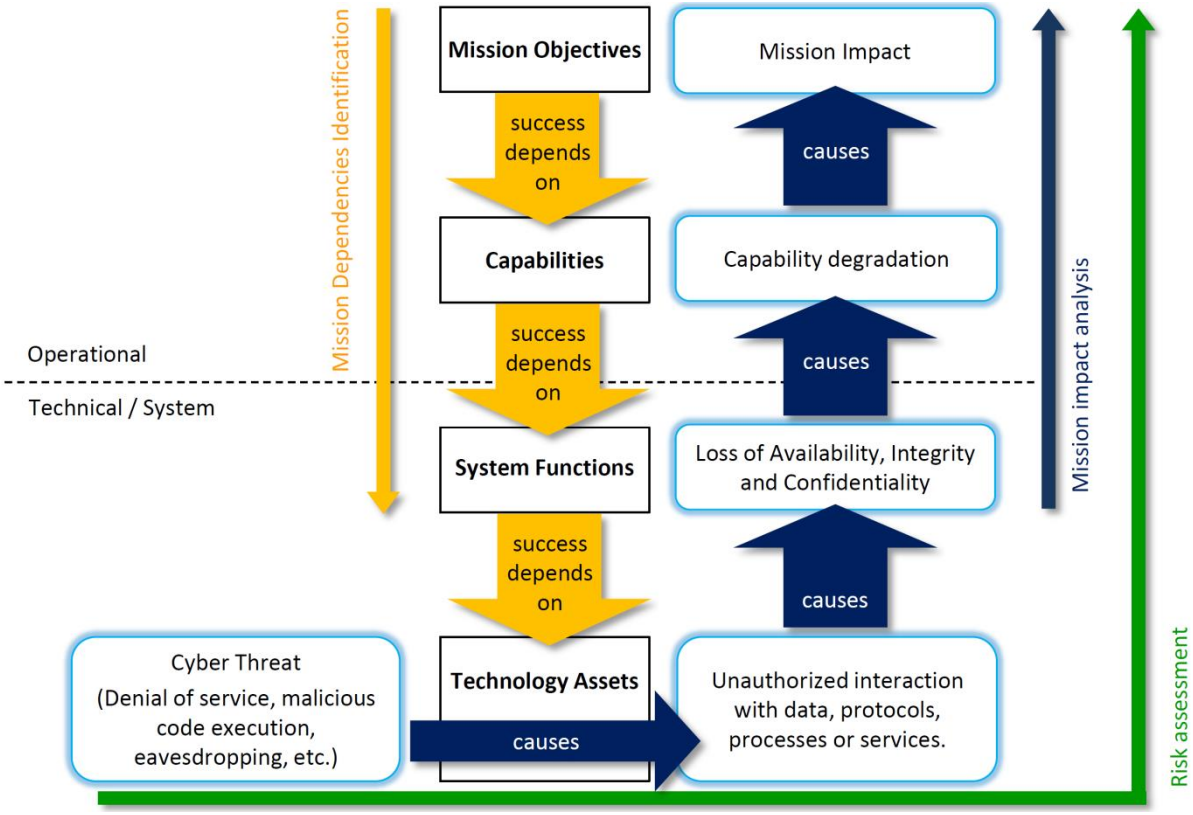

Figure 1: Cyber Mission Assurance model.

Through its four layers, the CMA model shown in Figure 1 also helps communicate CMA needs and results during the different life cycle phases of systems, therefore contributing in ensuring CMA continuity. The typical system life cycle stages include concept, development, production, utilization, support, and retirement [29]. From the perspective of military organizations, these stages can be summed up to acquisition (concept, development and production), utilization and support. Acquisition can be the result of contracts or in-house engineering. Utilization puts the system into execution as part of one or many missions, which may involve operational planning activities [30]. Support has to do with day-to-day operation of the systems and includes maintenance and support tasks. Routine training on secure system operation, for instance, is in the support phase. Like the CMA model, the defined CMA processes and metrics must allow for harmonized risk communications between those life cycle phases.

3. CMA Process

A CMA process represents the activities necessary to achieve CMA and their interaction. With a layered CMA model defined, the activities can be decomposed in a cohesive manner, where the outputs of one

activity are understandable for other activities. To benefit from knowledge and lessons learned from the past, and to allow for efficient communications of CMA activities among different players, CMA processes should leverage from recognized frameworks, guidelines and best practices that are either currently used or have been used, both at the operational level and at the technical level. For instance, better organizational coherence is achieved if CMA adopts the same impact levels used by the capability planning and operational community. If the latter use a four-level scale, for instance, then it would only complicate things if CMA was to use a different scale (e.g. three or five levels).

In the cybersecurity area, one of the most widely adopted frameworks is NIST's cybersecurity framework (CSF) [10]. One of the characteristic of the CSF is that its five core security functions, identify, protect, detect, respond and recover, are in-line with the CMA objectives. In fact, the first function, identify, develops the awareness onto which protect, detect, respond and recovery functions can be built. It turns out that the four latter functions correspond to the definition of resilience, while 'identify' has to do with the assessment of risks.

In terms of engineering, the System Security Engineering (SSE) processes described in [31][8] use a structured language and follow recognized engineering standards [29]. Most particularly, the security-related technical processes are oriented toward the definition, implementation, verification and validation of requirements, before the final products are operated and maintained. A mission analysis process is also included and that could be implemented following MITRE's Crown Jewel Analysis [13][14].

Although NIST SSE and the CSF provide sound orientations on 'what to do', they do not dive into 'how to do it'. Put the notions, concepts and directions aside, engineering and cybersecurity are complex multi-facets fields that are best learned through practice and experience. From a practical perspective, the aviation industry has practiced safety engineering for decades, where processes recognized for certification have been used and tailored over time [32][33]. This involves a risk management process where failure conditions are identified, mitigating measures are defined, implemented, verified and validated against the defined safety needs. Over the years, the aviation industry has become heavily reliant on electronic systems. As a consequence and in response to the threat of intentional unauthorized electronic interaction to aircraft safety, and building from years of experience in safety engineering, guidance on airworthiness security has been developed and put into practice to comply with accreditation and certification criteria. This includes the Airworthiness Security Process Specification [17][34] and their accompanying documents [35][36][37][38]. One of the key aspects of the specification is the use of a two-stage risk assessment approach, which consists in a preliminary risk assessment and a full risk assessment. This approach helps with decomposing the problem complexity and communicating risks, starting from an initial high-level standpoint and finishing with a refined low-level testing phase. The two risk assessment stages can also be aligned with the system engineering process, where the preliminary stage aligns with the definition of the architecture and design, and where the full stage aligns with the implementation and configuration of the developed product. The two stages also serve during operation and maintenance, where preliminary risks are assessed based on the actual architecture and design, and where full risks are assessed based on hands-on testing activities.

In accordance with the presented CMA model and based on an integration of the highlighted properties of the above-mentioned references, three main interacting CMA activities are introduced in Figure 2: Mission Criticality Analysis and Asset Valuation (MCAAV), Risk Assessment (RA) and Resiliency Development (RD). The process as a whole achieves two complementary and interrelated objectives: awareness of risks and development of resilience. To help organizations with integration of the process into their normal activities, such as system acquisition, operations and support, and operational planning, the activities are aligned with system engineering processes [29]. The risks are assessed with consideration of the mission criticalities of the assets, where threat scenarios assessed at the technical level can have their impact evaluated at the mission level. The assessed risks not only help to define

requirements, during the preliminary risk assessment, but they also participate in the verification and validation of the requirements, during the full risk assessment where the verification and validation tests are used to evaluate the residual risks. The three activities are described in the next paragraphs.

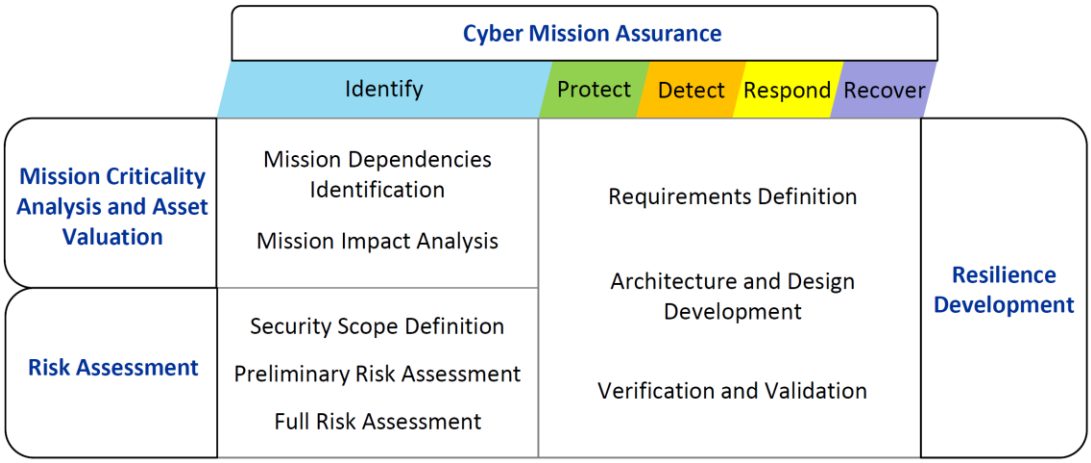

Figure 2: Risk-based Cyber Mission Assurance Process (RCMAP).

3.1. Mission criticality analysis and asset valuation

Mission criticality analysis and asset valuation (MCAAV) aims at determining the degree to which cyber mission assurance is needed. It predicts mission impacts caused by potential losses of the assets of an organization, therefore identifying assets that are most critical to the mission accomplishment. The provided mission impacts prepare for risk assessment and are a first step in getting aware of the risks of operating in a cyber contested environment. This activity is essential to the development of resilience of the mission assets [39].

As shown in Figure 3, MCAAV includes the identification of mission dependencies and the analysis of mission impacts in the events of losses of system functions.

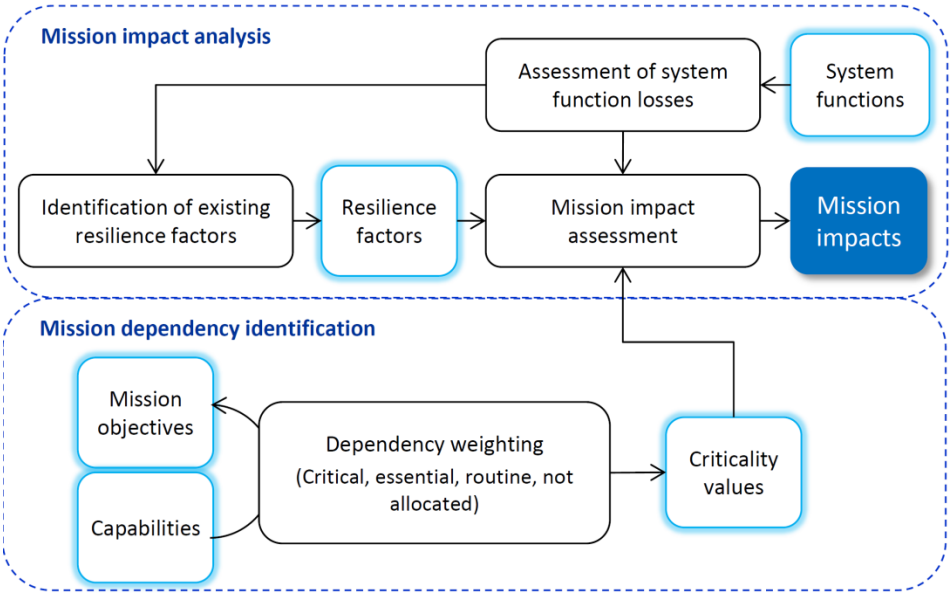

Figure 3: Mission criticality analysis and asset valuation.

For mission dependency identification, a simplification of the method presented in [40][3] is suggested, where only the dependencies between mission objectives and capabilities are identified, and according to criticality levels known to the operational community [41] [42] and described in Table 1: Critical, Essential, Routine or Not Allocated.

Table 1: Criticality levels of a capability to the mission objectives.

Criticality	Description
Critical	A capability that is absolutely necessary to achieve mission success. The mission objectives cannot be met without precise and timely support of the capability.
Essential	A capability that is necessary to achieve mission success. The execution of the mission will be severely impacted if the capability is not employed or cannot achieve its tasks or functions.
Routine	A capability that is required to execute to achieve mission success but has either a routine supporting task or function or a very low likelihood of employment.
Not Allocated	A capability that is not required for the mission objectives being considered.

For mission impact analysis, traditionally the security categorization of information systems [15][16][43] has characterized adverse impacts by security objective (e.g., confidentiality, integrity, or availability). The same is suggested for CMA with the distinction that the definition of losses of security objectives must take into account the existence of cyber-physical systems, which, in contrast to information systems, produce physical effects instead of information.

Based on the definitions of losses of Confidentiality, Integrity and Availability, in Table 2, the method consists in analysing the system function losses and then assessing their impacts on the capabilities used in the mission. This requires identifying the system functions in the first place. At this stage, the system functions correspond to the high-level functions[1] that systems execute. In many cases, the system functions are found after the name of the systems that execute them. A list of those systems could be found, for instance, in statements of operating requirements, which provide a high-level view of the required systems. For instance, a military aircraft has navigation and display, flight management, avionics networking, diagnostic and Identification Friend or Foe (IFF) functions, among others.

Table 2: Definition of losses of confidentiality, integrity and availability.

Security Objective	Loss definition
Confidentiality	A loss of confidentiality is the unauthorized disclosure of information or discovery physical effect.
Integrity	A loss of integrity is the unauthorized modification or destruction of information, or the unauthorized modification of a physical effect.

[1] 'function' may sometimes be referred to as 'task'.

Availability	A loss of availability is the disruption of access to or use of information or an information system, or the disruption or stoppage of a physical effect.

Mission impact analysis first determines the impact of the losses of Confidentiality, Integrity or Availability of the system functions on the capabilities. For instance, a system function loss could be the 'loss of integrity of the IFF function'. For each system function loss, the impact on each of the capabilities identified for the project must be rated. This exercise must consider the existing resilience factors that contribute in attenuating the impact of the function loss. For instance, can the failing function be halted to prevent or attenuate the impacts? Can operators detect the function loss or its effects, and how? Can they detect it in time? If they are able to detect the function loss or its effects, can they respond to it time, before the effect is produced and the capability is impacted? What are the available responses and do they engender indirect impacts (e.g. increased pressure on resources)? Are there equivalent, compensating or redundant functions to take over the execution of the impacted capabilities? Can they be actuated in time? The answers will help determine the real impacts of the function losses.

Mission impact analysis is concerned with resiliency at the operational level (strategies, operations or tactics). Although resiliency at the system level can be considered roughly at a high level (e.g. consideration of redundant systems), the resiliency taking place in the technological assets is considered during the risk assessment activity of the process.

Considering the mission dependencies and the assessed impacts of system function losses on capabilities, a mission impact can be determined. For instance, if a critical capability is highly impacted by the function loss, the resulting mission impact may be Catastrophic (see Table 3), but if a routine capability is highly impacted, the mission impact may only be Marginal. The overall mission impact of a system function loss, considering all impacted capabilities, may be determined using the maximum rule over all capabilities, for instance.

Table 3: Example of mission impact levels, adapted from [18][44].

Severity Level	Impact of a system function loss
Catastrophic	The function loss causes a total failure of the mission. Objectives of the mission are not accomplished.
Critical	The function loss causes a significant degradation of the mission. Only few mission objectives are met and with a significantly reduced effectiveness.
Marginal	The function loss causes a limited degradation of the mission. Mission objectives are met, but with reduced effectiveness.
Negligible	Little or no adverse impact on mission objectives.

3.2. Risk assessment

Risk Assessment (RA) is about the definition of cyber threat scenarios at the system-level and correlation of them with system function losses and their mission impacts. This activity includes the security scope definition, which identifies the technology asset(s) and their security aspects, and the risk assessment itself, which assesses cyber threat scenarios to the assets and evaluates their impacts. Risk assessment

involves vulnerability assessment and is performed by cyber security experts along with the collaboration of systems engineers, analysts, developers, technicians and operators.

The scope definition activity consists in the identification of the assets that require risk assessment and their description, including their attack surface and their security environment. The attack surface of a system represents its exposures by which a threat actor can perform unauthorized or potentially harmful activities [45]. It can be decomposed into a number of categories, including the physical, sensing/electromagnetic, logical, personnel and indirect (support and supply chain) attack surfaces. The logical surface, which is composed of hardware and software, is the most complex of them.

The preliminary RA prepares for a full RA and consists in a higher-level analysis of the vulnerabilities in the architecture and design of the assets under assessment, as well as in the procedures, the rules and the policies related to their operation and maintenance [17][35]. Sometimes only partial information is available depending on the project progress. Preliminary RA identifies preliminary threat scenarios that the described attack surface reveals. A threat scenario is made of a series of threat events, which are found based on an assessment of the vulnerabilities of the assets. At the preliminary stage, threat events can be seen as tactics and techniques. Examples can be consulted in NIST's Guide for Conducting Risk Assessments [18], the Common Attack Pattern Enumeration and Classification (CAPEC) dictionary [46] and the European Union Agency For Network And Information Security (ENISA) threat taxonomy [47], among others. Threat trends and intelligence, when available, may also contribute in identifying threats relevant to the CMA project. For instance, 'buffer manipulation' of a particular asset, might be a threat event [46].

The full RA wraps up the preliminary assessment with hands-on security verification activities, such as vulnerability scanning, penetration testing, application fuzzing, red teaming and reverse engineering. Using the vulnerabilities discovered during these activities, the adversarial tactics and techniques defined during the preliminary RA can be completed with the lower-level procedures described according to series of actions and vulnerability exploitations of the implemented, installed or in-service assets. Figure 4 shows the two stages of threat scenarios and their relationships to tactics, techniques and procedures.

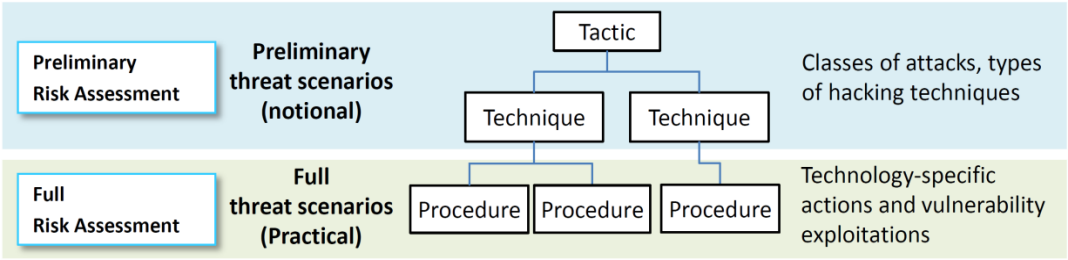

Figure 4: Two levels of threat scenarios.

Risk are assessed by associating the threat scenarios to the system function losses that they can cause, in terms of loss of Confidentiality, Integrity and/or Availability, and the related mission impact obtained during mission impact analysis. The risks that threat scenarios cause to the mission can be defined in terms of the estimated likelihood of threat scenarios and the impact of their related system function loss(es):

$$\text{Mission risk} = \text{likelihood(threat scenarios)} \times \text{impact(system function loss)}.$$

where the impact can be obtained using Table 3, for instance, and where various methods an concepts exist for the estimating likelihoods of threat scenarios, such as in [35][37][18]. When threat intelligence is available, threat scenarios can also consider a threat actor, described in terms of its intent and capabilities. Methods for calculating risks from likelihood and impact are provided in [18]. A risk scoring matrix

example is shown in Table 4. The impact levels fits with the mission impact levels described in Table 3. The matrix is used by the Canadian Forces for risk management in general [44]. Using it for CMA would allow them harmonizing the results with other risk management areas (e.g., natural hazards, safety).

Table 4: Risk scoring matrix.

		Impact			
		Negligible	Marginal	Critical	Catastrophic
Threat Likelihood	Very High				
	High				
	Medium				
	Low				
	Very Low				

Risk degree	Color code
Extremely High	
High	
Moderate	
Low	

To help with enumerating and structuring the different test activities in risk assessments, the classification found in NIST's Technical Guide to Information Security Testing and Assessment, which defines testing methodologies and address the common processes associated to them [48], was extended to take into consideration all the resilience functions (protect, detect, respond, recover). According to this, four classes of CMA tests can be defined at a high-level:

- Reviews: Set of techniques used to evaluate systems, applications, networks, policies, and procedures to discover vulnerabilities, and that are generally conducted manually. This includes documentation, log, ruleset, and system configuration reviews; network sniffing; and file integrity checking.
- Target Identification and Analysis: Set of techniques that can identify systems, ports, services, and potential vulnerabilities, and that may be performed manually but that are generally performed using automated tools. This includes network discovery, network port and service identification, vulnerability scanning, wireless scanning, and application security examination.
- Target Vulnerability Validation: Set of techniques that corroborate the existence of vulnerabilities, and that may be performed manually or by using automatic tools, depending on the specific technique used and the skills of the test team. Target vulnerability validation techniques include password cracking, penetration testing, social engineering, and application security testing.
- Resilience testing: Set of techniques that test the detect, respond and recover capabilities assuming failures to prevent cyber-attacks at the technical level and the related system function losses and their mission impacts assessed during the MCAAV activity. Ideally, the related activities should involve developers, operators and testers altogether.

Test activities that cover the four classes are enumerated in Table 5, including references for further description. When conducting these activities, security experts are tasked to identify vulnerabilities, document threat scenarios and evaluate their risks in terms of likelihood and impact. Risk decisions also need to be documented, including mitigating measures.

Table 5: Test activities for risk assessment.

Reviews activities
Documentation review [48] Reviews for technical accuracy and completeness, including security policies, requirements, architectures, designs, standard operating procedures, system security plans, authorization agreements, memoranda of understanding and agreement for system interconnections, and incident response plans.
Security audits [49] Audit implies that we are measuring things against a fixed, pre-determined, rigorous set of standards. These audits are almost always done with detailed checklists.
Ruleset review [48] Reveals holes in ruleset-based security controls.
System configuration review [48] Evaluates the strength of system configuration. Validates that systems are configured in accordance with hardening policy.
Network sniffing [48] Monitors network traffic on the local segment to capture information such as active systems, operating systems, communication protocols, services, and applications. Verifies encryption of communications.
File integrity checking [48] Identifies changes to important files; can also identify certain forms of unwanted files, such as well-known attacker tools.
Log review [48] Provides historical information on system use, configuration, and modification. Could reveal potential problems and policy deviations.
Target Identification and Analysis activities
Network scanning [48] Network discovery: Discovers active devices. Identifies communication paths and facilitates determination of network architectures; Network port and service identification: Discovers active devices. Discovers open ports and associated services/ applications.
Vulnerability scanning [48] Identifies hosts and open ports. Identifies known vulnerabilities. Often provides advice on mitigating discovered vulnerabilities.
Static Analysis [50] Analysis of software without actually executing programs. Static analysis is performed by an automated software tool and should not be confused with human analysis of software security architectural reviews, which involve human code reviews.
Fuzzing [50] Attack simulation in which unexpected data is fed to the system through an open interface, and the behavior of the system is then monitored.
Reverse engineering [51][52] Discover hardware, firmware and software flaws by deconstructing and deducing architecture and design features.
Wireless scanning [48] Identifies unauthorized wireless devices within range of the scanners. Discovers wireless signals outside of an

organization's perimeter. Detects potential backdoors and other security violations.
Forensics [53] Identifies, collects, examines, and analyses data to investigate crimes and internal policy violations, reconstruct computer security incidents, troubleshoot operational problems, or recover from accidental system damage. Although forensic is more a response activity than a test activity, its results can reveal vulnerabilities or threat that can be considered for risk assessment.
Target Vulnerability Validation activities
Penetration testing [48] Tests security using the same methodologies and tools that attackers employ. Verifies vulnerabilities. Demonstrates how vulnerabilities can be exploited iteratively to gain greater access. This category includes network penetration testing [49], red teaming [54] and other activities listed in this table that can be put together to discover, verify and exploit vulnerabilities.
Cryptanalysis attack [49] Focuses on bypassing or breaking the encryption of data stored on a local system or across the network.
Physical security test [49] Looks for flaws in the physical security practices of an organization. Testers may attempt to gain access to buildings and rooms, or to take laptops or other assets out of target facilities. Dumpster diving tests are a variation of a physical security analysis.
Hardware hacking [55] Includes information gathering and reconnaissance, external and internal analysis of the device, identification of communication interfaces, data acquisition using hardware communication techniques and software exploitation using hardware exploitation methods, including backdooring.
Radio hacking [55] Wireless communications analysis and exploitation.
Social engineering [48] Allows testing of both procedures and the human element (user awareness).
Password cracking [48] Identifies weak passwords and password policies.
Resilience testing activities
Chaos engineering experiments [56][57][58][59] Chaos engineering is the discipline of experimenting on a software system in production in order to build confidence in the system's capability to withstand turbulent and unexpected conditions.
Fault injection [60][61] Accomplishment of controlled experiments where the observation of the system's behavior in presence of faults is induced explicitly by the injection of faults in the system. [60]
Purple teaming [59] [62] The goal of the Purple Teaming is the collaboration of offensive and defensive tactics: the offensive team should use all TTPs available by the attacker while the defensive team, by trying to detect and respond to the attacks, should assess the detection and response capabilities based on their obtained results, and identify areas that need improvement.
Incident management capabilities reviews [63] Document analysis, interviews and on-site observations to assess the actual incident management capabilities against the incident management plan.

Incident management tests and exercises [64][65][66]

Tests and exercises ranging from table top exercises (paper-driven analysis of scripted incident scenarios) to live simulations (real incident scenarios) to verify the incident responses in terms of preparation, detection, containment, eradication and recovery.

3.3. Resilience development

In the context of CMA, resilience development is dependent on the management of risks, which implies risk decisions as a first step. Note that risk decisions could also be referred to as risk requirements, in that a decision will necessarily lead to something that must be done (or not done). Traditionally risk requirements have been separated into the following categories [16]:

1. Risk acceptance: the risk is accepted as it is without action. This is normally the preferred choice when the costs of the other options, e.g., mitigating the risks by technical or operational measures, is deemed too high compared to the benefit.

2. Risk transfer or sharing: Risk are shared or transferred to second or third parties. This mostly applies to the financial aspect of risk management, where it resorts to taking insurances or buying warranties. Although it is a common choice in the civil world, this is less true for the military. The government is actually its own insurer and although warranties can be asked in terms of material replacement or financial compensations, they cannot cover mission success.

3. Risk avoidance: the risk is avoided by disabling the functions or aborting the actions that cause it, or by circumventing the conditions for which it can exist. This decision can be seen as way to develop resilience in cases where it does not affect to continuity of the mission.

4. Risk mitigation: Actions are taken to lower the risk to an acceptable level.

Since their goal is to ensure mission success, resilience requirements form a subset of the above risk requirements categories composed of risk mitigation and risk avoidance types of requirements. Resilience can take place at any layer in the CMA model (Figure 1). It can be a property of a particular technology, system, operational capability or mission. Although this is not the topic of this paper, resilience could also extend to an organization, region or nation. Whatever the scope, the objective is the same: manage the assessed risks to ensure mission continuity. Resilience requirements can be decomposed into the protect, detect, respond and recover functions defined in [10]. These functions can play out at any layer in the CMA model:

- Technical measures consist in technology or system solutions to protect, detect, respond or recover from cyber-attacks and their mission impacts, and

- Operational measures consist in capabilities or mission objectives to protect, detect, respond or recover from cyber-attacks and their mission impacts.

Resilience development focuses on the definition, validation and verification of resilience requirements. These stages are part of the security engineering processes in [31]. If the latter addresses the system/technical layer of the CMA model, the definition, validation and verification processes can be extended to operational resilience measures, with the distinction that experts in operational activities are required instead of technical experts.

CMA acts on the mitigation of mission risks caused by operating in the cyberspace. Resilience requirements are defined to avoid or mitigate the mission risks and ensure that mission objectives can be

successfully met. Given that a maximum tolerable risk level is determined, the definition of resilience requirements must consider each assessed risk rated as unacceptable to ensure mission success, starting with the higher ones. To perform this, the definition procedure should first look for the system function losses that cause the highest mission impacts. Assessed threat scenarios associated to each system function loss (i.e., loss of Confidentiality, Integrity or Availability) should then be searched for. Given the system function loss under consideration, the related threat scenarios and their mission impacts, the task is to find out, through analysis (e.g. cost-benefit), a set of resilience measures that will mitigate the risk associated to the system function loss, i.e. reduce the likelihoods of the threat scenarios or reduce their impacts. Table 6 shows the categories of resiliency measures to help achieve this task. Under equal degrees of effectiveness, technical solutions may be prioritized first over operational solutions to minimize costs. In particular, this applies when operational solutions ask for supplemental systems and technology, which can cause additional risks on top of the additional costs. If no technical solution comes out as effective enough to mitigate the risks with reasonable cost, then operational solutions could be searched for, spreading from administrative/management procedures (e.g., training, maintenance, etc.), orders and directives, and military capabilities (tactical, operational or strategic).

At the technical level, references exist to help with the definition of resilience requirements, such as the cyber resiliency techniques and approaches in [8] (e.g., diversity, deception and unpredictability), as well as NIST CSF's framework core and its informative references [10]. At the operational level, the resilience management model presented in [67] provides practices that guides on the development of resilience by experts on military strategies, operations and tactics.

Table 6: Categories of resilience measures to mitigate a residual risk.

CMA layer	Protect	Detect	Respond	Recover
Technical	Technology solutions or system-level measures to prevent or diminish the likelihood of occurrence of the threat events or system function loss. (e.g., access control, encryption)	Technology solutions or system-level measures to detect the occurrence of the threat events or system function loss. (e.g., intrusion detection systems, log analysis)	Technology solutions or system-level measures to respond to the occurrence of the threat events or system function loss. (e.g., forensics tools)	Technology solutions or system-level measures to recover from the occurrence of the threat events and restore the secure execution of the system function. (e.g., backup systems, reboot techniques)
Operational	New procedures, directives, orders, capabilities or mission objectives to prevent or diminish the likelihood of occurrence of the threat events or system function loss. (e.g., cybersecurity training, preventive attacks against enemy targets)	New procedures, directives, orders, capabilities or mission objectives to detect the loss of system function and allow responding and/or recovering before the executed capabilities are impacted. (e.g., facility/platform/system inspections)	New procedures, directives, orders, capabilities or mission objectives to respond to the loss of system function before the executed capabilities are impacted. (e.g., system shut-down directives, evacuation procedures, return to base measure)	New procedures, directives, orders, capabilities or mission objectives to recover from the loss of system function and return to the execution of the capabilities in a timely fashion. (e.g., backup platforms)

Resilience requirements can be categorized into the protect, detect, respond and recover categories, each with estimate of the costs and resources necessary, time to implementation. Each requirement mitigates a particular set of assessed risks. Some requirements mitigate only a specific risk, while other requirements

mitigate a collection of risks. Overall, the determination of the best set of requirements asks for a cost-benefit analysis, which should not only be constrained to CMA but also consider other classes of requirements (e.g. safety, maintainability, etc.), in conformity with the organization's acquisition, operations and support, or mission planning processes.

The resilience requirements are normally gathered into a requirement traceability matrix [68][69], and are accompanied with the test cases or test methods for verification and validation. The requirement traceability matrix will inform on the development status of the defined requirements, from design, implementation and up to verification and validation. In terms of risk management, the requirements should also be documented in relation to the assessed risks that they mitigate, including the residual risks.

Ideally, an organization would reach a stage where all risks are identified and all the proper resilience measures are put into operation. This would represent full CMA achievement. However, the reality is that the definition, implementation and testing of resilience measures is a long process that is subject to various constraints. Moreover, the cyberspace is in constant change. New threats arise every day and new vulnerabilities are exposed as systems are reconfigured, updated or operated in new ways. This is why it is almost impossible to fully complete CMA at any specific point time. CMA is never finished. It is a continuous catch-up process. Risks must be continuously identified and organization must make their best to develop resilience under limited time and resources.

On top of the CMA process, metrics are needed to measure to state of achievement of CMA and support decision-making during the different phase of the life cycle of systems. CMA metrics are presented in the next section.

4. CMA Metrics

Whether project stakeholders are concerned with the acquisition, utilization or support of military assets, they need to know if these assets allow them continuing their mission even in the presence of cyber-attacks, i.e. be cyber resilient [4], and what remains to be done if it is not the case. After they have done what needs to be done, they need to check that it works as intended. Finally, they need to have an idea of the degree of resilience that they achieve, i.e., are they still facing some risks and to what amount? While in previous works is was observed that 'no single cyber resiliency metric or set of metrics will work for all environments' [70], the proposed CMA metrics aims at fulfilling the needs of any environments and for any kind of projects, including acquisition, operation and support, and operational planning projects, for instance.

To answer the aforementioned goals, two types of CMA metrics are introduced: CMA effectiveness and CMA performance. CMA effectiveness metrics represent the progress of CMA accomplishment. It tells if enough has been done under predefined conditions (e.g., time and cost). This supposes that objectives are defined, which can be determined in terms of the necessary CMA activities to be achieved. These activities are defined inside the CMA process and according to the CMA model that organizations or project stakeholders have adopted. Rather than a measure of the progress of accomplishment, CMA performance measures the results of the activity accomplished, in terms of the achieved state of resilience and residual risks. In short, CMA effectiveness measures the amount of work accomplished while CMA performance measures the results of the work accomplished.

An overview of the proposed CMA metrics is shown in Figure 5, where the metrics are aligned with the five CMA functions: identify, protect, detect, respond and recover. The metrics are further decomposed in Figure 6, which presents a hierarchical view.

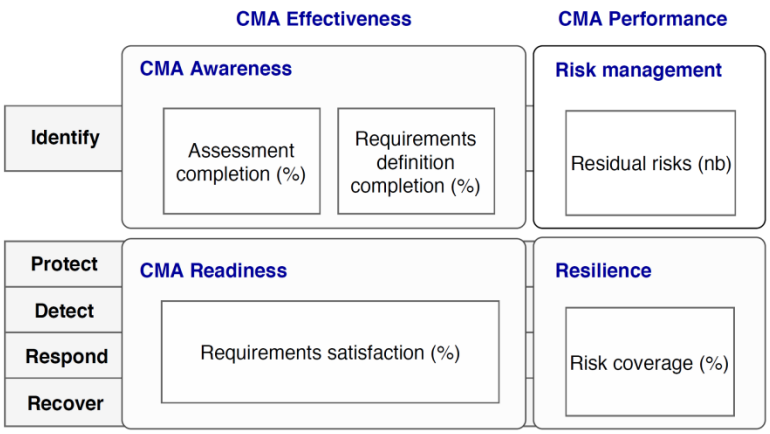

Figure 5: Overview of risk-based CMA metrics.

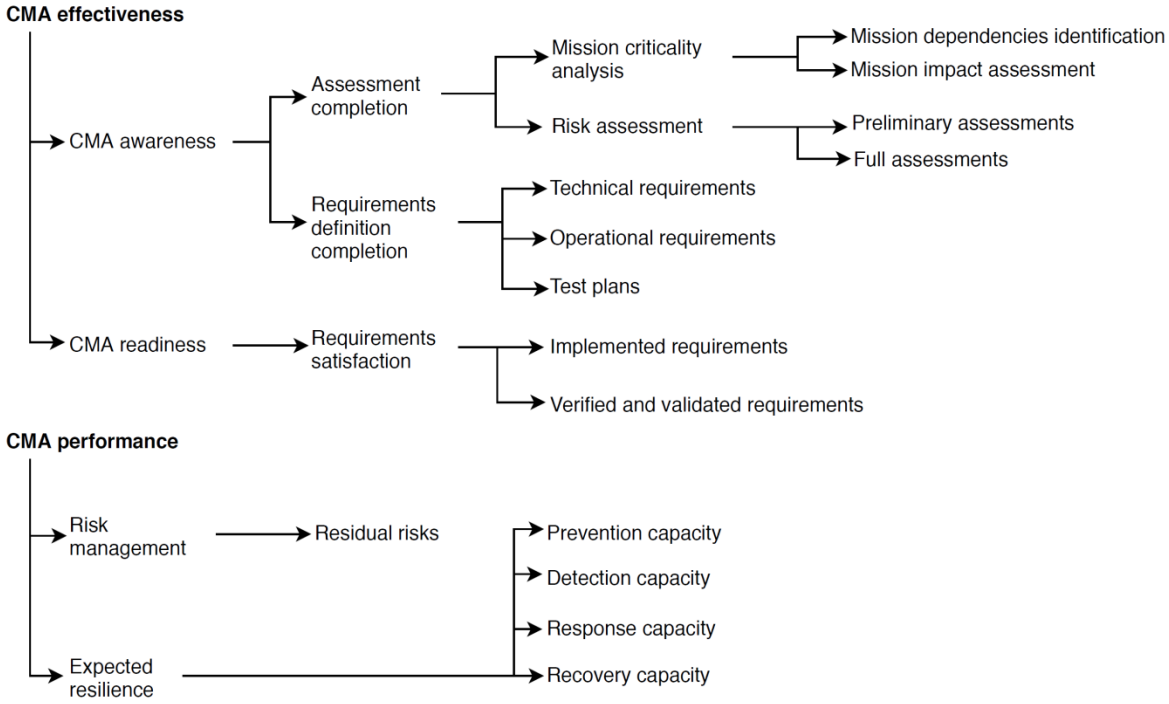

Figure 6: Hierarchical view of risk-based CMA metrics.

4.1. CMA effectiveness

CMA effectiveness measures the state of accomplishment of CMA. It can be expressed by the ratio of CMA activities achieved to the total number of CMA activities in a project. The ratios can be weighted by cost and time, for instance. CMA activities are what composes a CMA process and are applied to assets related to the project. An asset can be a person, system, platform, weapon system, unit or department, for instance. The state of accomplishment of CMA can be expressed in terms of the CMA functions (identify, protect, detect, respond and recover) or more simply in terms of the degrees of CMA awareness (identify) and CMA readiness (protect, detect, respond, recover) that a project has.

CMA awareness measures the degree to which a project is aware of the problems that need to be addressed and the solutions to solve them. In CMA, the problems are expressed in terms of risks and the solutions are expressed in terms of the resilience requirements to solve the risks. A low awareness ratio is measured when many assets have not been assessed, which means that the risks that they pose to the mission are unknown, or when the requirements to counter the assessed risks are not elicited, which means that the required measures to achieve CMA are unknown. A high awareness ratio is measured when all assets have been risk assessed and when all the requirements to counter the assessed risks are defined.

CMA readiness measures the degree to which a project is prepared to manage the assessed risks. This can be represented by the ratio of the number of resilience requirements satisfied to the total number of resilience requirements defined. A low readiness, e.g., below fifty percent, means that few of the defined requirements have been satisfied and therefore the related project is not ready to counter the risks that have been assessed. A high readiness, e.g., near a hundred percent, means that all of the defined requirements have been implemented, verified and validated, and therefore the projects and the related assets are ready to counter the risks that have been assessed.

CMA effectiveness metrics are defined in Table 7, in accordance to the CMA activities presented in the previous section. The metrics are expressed in terms of ratios where the total number of assets in a project is the denominator. The measured CMA effectiveness can influence decisions, such as determining whether an asset can be employed in missions or in environments known to have sophisticated threat actors with specific cyber capabilities and intents, and under what conditions and constraints it can be used.

Table 7: CMA effectiveness metrics.

CMA Awareness
Assessment completion = [*Number of assessed assets / Total number of assets*]. An assessed asset is an asset of the project for which mission criticality analysis and risk assessment are completed.
Mission criticality analysis = [*Number of assets with completed mission criticality analysis / Total number of assets*]. The mission criticality analysis of an element is completed if the mission dependencies identification and mission impact assessment are completed.
Mission dependencies identification = [*Number of assets with completed mission dependencies identification / Total number of assets*]
Mission impact assessment = [*Number of assets with completed mission impact assessment / Total number of assets*]
Risk assessment = [*Number of assets with completed risk assessment / Total number of assets*]. Risk assessment includes the preliminary and full risk assessment.
Preliminary risk assessment = [*Number of assets with completed preliminary risk assessment / Total number of assets*]
Full risk assessment = [*Number of assets with completed full risk assessment / Total number of assets*]
Requirements definition completion = [*Number of assessed assets for which requirements are defined / Total number of assets*]
Technical requirements **=** [*Number of assessed assets for which technical requirements are defined / Total number of assets*]
Operational requirements **=** [*Number of assessed assets for which operational requirements are defined / Total number of assets*]

Test plans = [*Number of assessed assets for which test plans are defined / Total number of assets*]
CMA Readiness
Requirements satisfaction = [*Number of assets for which requirements are implemented, verified and validated.*]
Implemented requirements = [*Number of assets for which requirements are implemented.*]
Verified and validated requirements = [*Number of assets for which requirements are verified and validated.*]

4.2. CMA performance

CMA performance measures how risks are managed and the expected resilience to those risks. It requires risks to be assessed in the first place. A project with a high CMA performance has low residual risks, which are managed by the necessary resilience measures. A project with a low CMA performance is one that counts high risks and where resilience measures are missing or not effective enough in mitigating the risks. Knowing what the risks left are and their severity, project stakeholders can take decisions on whether to accept, mitigate, transfer or share the responsibility of the risks.

CMA performance metrics are presented in Table 8. The performance of risk management, which is related to the 'Identify' CMA function, is measured by the residual risks metric, which includes risks that have either been decided to be accepted, transferred or shared, and also those that were required to be mitigated but that are not actually fully mitigated. The metric gives the number of residual risks per risk score. This informs on how well risks are managed. For instance, a project could have 3 residual risks rated as 'Low' and 2 residual risks rated as 'High'. The goal is to get a minimum of high residual risks and a maximum of low residual risks. The best performance is achieved when only residual risks of a minimum severity remain, e.g., Low.

The resilience performance is measured by four metrics: prevention capacity, detection capacity, response capacity and recovery capacity. Those metrics inform on the performance of the individual resilience functions: protect, detect, respond and recover, respectively. For instance, the prevention capacity informs on the amount of risks that are effectively countered by prevention measures. This provides stakeholders and analysts with data on the repartition of the resilience functions. The concept behind CMA is to have a balance of measures to cope with different possibilities, including not only the possibility to prevent cyber-attacks but also to react to the eventuality that cyber-attacks succeed in affecting the assets under protection. For example, if the prevention capacity is a hundred percent, meaning that all risks are effectively managed by prevention measures, but that the response capacity is twenty percent, meaning that twenty percent of risks have response measures, than the overall resilience is not so good even if excellent protection is achieved. Part of the effort associated with such performance metrics will be to classify the implemented measures into the prevent, detect, respond and recover categories, for which the distinction might not always be clear. The NIST CSF provides a good lexicon to start with and to support this task. Moreover, the determination that a risk is managed by a particular measure or group of measures assumes that the measure(s) is(are) effective in mitigating the risk. The latter is assessed by subject matter experts (e.g. systems and cybersecurity experts) during the risk assessment activity.

Table 8: CMA performance metrics.

Risk management
Residual risks = *Number of residual risks per risk score (e.g., High, Very High)*
Expected resilience
Prevention capacity = *Number of risks managed with prevention measures / Total number of risks* Detection capacity = *Number of risks managed with detection measures / Total number of risks* Response capacity = *Number of risks managed with response measures / Total number of risks* Recovery capacity = *Number of risks managed with recovery measures / Total number or risks*

5. Discussion

The presented model, process and accompanying metrics provide foundations on which to institute policies, programs and procedures regarding CMA within organizations and their projects. They help in structuring, describing and measuring the necessary actions to achieve CMA. Their principal features are:

1. a layered CMA model that spans from cyber threats to mission impacts,

2. a simple method for identifying dependencies between possessed capabilities and mission objectives,

3. a mission criticality analysis that extends the loss of Confidentiality, Integrity and Availability to cyber-physical systems and in connection with the identified dependencies,

4. a two-stage risk assessment integrated with the verification process of system engineering,

5. a resilience development focused on the mitigation of assessed mission risks and aligned with system engineering practices,

6. CMA effectiveness metrics that allow measurement of the state of accomplishment of CMA within organizations or in projects, and it terms of the degrees of CMA awareness and CMA readiness, and

7. CMA performance metrics that informs on the quality of the developed CMA capacity in terms of risk management and expected resilience.

These are achieved by integration of a multitude of already existing and well-proven approaches from different areas in a cohesive and practical fashion. Based on those foundations, it is suggested to initiate efforts by keeping a small scope until sufficient experience is established, especially when measurements of CMA is involved [70]. As experience is gained, tools will be developed to support the process and their measurements, knowledge will be acquired, lessons will be learned and efficiency will be increased. Based on those gains, the scope can be extended progressively to cover more assets, more types of risks, and eventually CMA can become an enterprise-wide endeavor.

As organizations, projects and systems will age, they will be subjected to a variety of events that will change the condition of CMA. New systems may be acquired or existing ones may be modified. Mission objectives and capabilities executed by the systems may change. New vulnerabilities may arise during utilization and support, as well as new threats and TTPs. In terms of resilience, new methods or measures that did not exist before may appear, while some of the old ones may become obsolete. Each of these events has the potential to decrease the effectiveness and/or performance of CMA in comparison to what

was achieved before it happened. To cope with that, the presented CMA process and the measurement of the related CMA metrics need to be practiced continuously, where procedures are in place to both detect and respond to the events. This would include continuous rounds of mission criticality analysis and asset valuation, risk assessment and resilience development. The guidelines for implementing a continuous CMA process, including the planning of the conduct of the activities and the determination of their triggering events, are topics for future work.

6. References

[1] Harriet Goldman, Rosalie McQuaid and Jeffrey Picciotto, Cyber Resilience for Mission Assurance, IEEE International Conference on Technologies for Homeland Security (HST), November 2011.

[2] Office of the Under Secretary of Defense for Policy, MISSION ASSURANCE (MA) CONSTRUCT, DOD INSTRUCTION 3020.45, August 14, 2018.

[3] MITRE Systems Engineering Guide (SEG), Cyber Mission Assurance, 2014. Online: https://www.mitre.org/publications/systems-engineering-guide/

[4] Brad Bigelow, Mission Assurance: Shifting the Focus of Cyber Defence, 9th International Conference on Cyber Conflict, NATO CCD COE Publications, 2017.

[5] Alexander Kott, Igor Linkov, Cyber Resilience of Systems and Networks, Springer, May 30, 2018.

[6] Therese P. McAllister, Community Resilience Planning Guide for Buildings and Infrastructure Systems, Volume I, NIST Special Publication 1190, May 2016.

[7] Scott Musman and Seli Agbolosu-Amison, A Measurable Definition of Resiliency Using "Mission Risk" as a Metric, Mitre Technical Report, Document number 140047, February 2014.

[8] National Institute of Standards and Technology (NIST), Special Publication 800-160, Volume 2, Systems Security Engineering, Cyber Resiliency Considerations for the Engineering of Trustworthy Secure Systems, 2016.

[9] C2 Agility, Task Group SAS-085 Final Report, STO Technical Report, STO-TR-SAS-085, 2014.

[10] National Institute of Standards and Technology (NIST), Framework for Improving Critical Infrastructure Cybersecurity, 2014.

[11] Department of Defense, DoD Directive 3020.40 Change 1 "DoD Policy and Responsibilities for Critical Infrastructure", July 1, 2010.

[12] Cyber Mission Assurance Working Group (CMA WG), Program Scope, Information Brief to CDS / DM, Director Cyber Force Development, 6 June 2017.

[13] William Heinbockel, Steven Noel, and James Curbo, Mission Dependency Modeling for Cyber Situational Awareness, NATO STO Meeting Proceedings, STO-MP-IST-148, September 2016.

[14] Gary Hastings, Lou Montella, and Jim Watters, MITRE Crown Jewels Analysis Process, MTR090088, April 8 2009.

[15] National Institute of Standards and Technology (NIST), Special Publication 800-37 Revision 1, Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach, 2010.

[16] National Institute of Standards and Technology (NIST), Special Publication 800-39 Revision 1, Managing Information Security Risks, 2011.

[17] RTCA Inc, Airworthiness Security Process Specification, RTCA DO-326A, 2014.

[18] National Institute of Standards and Technology (NIST), Special Publication 800-30 Revision 1, Guide for Conducting Risk Assessments, 2012.

[19] Military Operations Research Society CBP Workshop, Washington, October 2004.

[20] Joint Publication 3-0, Joint Operations, Joint Chiefs of Staff, Incorporating Change 1, United States, October 2018.

[21] Canadian Forces Joint Publication, CFJP 01, Canadian Military Doctrine, April 2004.

[22] Joint Doctrine Publication 0-01, UK Defence Doctrine, Fifth Edition, United Kingdom, November 2014.

[23] Joint Systems and Analysis Group, Technical Panel 3, Guide to Capability-Based Planning, 2004.

[24] National Defence, Capability-based Planning Handbook, Canada, June 2014.

[25] Department of National Defence, Joint Capability Framework, Canada, June 2014.

[26] Chairman of the Joint Chiefs of Staff Instruction, Charter of the Joint Requirements Oversight Council (JROC) and Implementation of the Joint Capabilities Integration and Development System (JCIDS), CJCSI 5123.01H, August 2018.

[27] Royal Canadian Air Force Doctrine: Intelligence, Surveillance and Reconnaissance, 2nd Edition, B-GA-401-002/FP-001, November 2017.

[28] Operation DRIFTNET, Online: https://www.canada.ca/en/department-national-defence/services/operations/military-operations/current-operations/operation-driftnet.html, Access date: 11 April 2019.

[29] Systems and software engineering - Life cycle management -- Part 1: Guidelines for life cycle management, ISO/IEC/IEEE 24748-1:2018, 2018.

[30] Joint Publication 5-0, Joint Planning, Joint Chiefs of Staff, United States, June 2017.

[31] National Institute of Standards and Technology (NIST), Special Publication 800-160, Volume 1, Systems Security Engineering, Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems, 2016.

[32] Guidelines for Development of Civil Aircraft and Systems, ARP4754A, SAE International, 2010.

[33] Guidelines for Development of Civil Aircraft and Systems, ED 79, EUROCAE, 2010.

[34] EUROCAE, Airworthiness Security Process Specification, ED-202A, 2014.

[35] RTCA Inc, Airworthiness Security Methods and Considerations, RTCA DO-356A, 2018.

[36] RTCA Inc, Information Security Guidance for Continuing Airworthiness, RTCA DO-355, 2014.

[37] EUROCAE, Airworthiness Security Methods and Considerations, ED-203, Revision A, June 2018.

[38] EUROCAE, Information Security Guidance for Continuing Airworthiness, ED-204, 2014.

[39] The MITRE Corporation, Cyber Resiliency Framework & Terminology, 2nd Annual Secure and Resilient Cyber Architectures Workshop, Final Report, 2012.

[40] William Heinbockel, Steven Noel, and James Curbo, Mission Dependency Modeling for Cyber Situational Awareness, NATO STO Meeting Proceedings, STO-MP-IST-148, September 2016.

[41] Chief Force Development. Capability Based Planning Handbook, v6.3, Art 943, United States, 15 January 2011, Art 943.

[42] Debra S. Herrmann, Complete Guide to Security and Privacy Metrics: Measuring Regulatory Compliance, Operational Resilience, and ROI, CRC Press, January 22, 2007.

[43] Standards for Security Categorization of Federal Information and Information Systems, Federal Information Processing Standards Publications, FIPS PUB 199, February 2004.

[44] Department of National Defence, Risk Management for CF Operations, Joint Doctrine Manual, B-GJ-005-502/FP-000, Canada, November 2007.

[45] Communications Security Establishment Canada, Enterprise Security Architecture (ESA) Enterprise Threat Assessment, January 2017.

[46] The MITRE Corporation, Common Attack Pattern Enumeration and Classification, Online: https://capec.mitre.org/index.html, Access date: January 8, 2019.

[47] European Union Agency For Network And Information Security, ENISA Threat Taxonomy - A tool for structuring threat information, January 2016.

[48] National Institute of Standards and Technology (NIST), Technical Guide to Information Security Testing and Assessment, Recommendations of the National Institute of Standards and Technology, NIST Special Publication 800-115, September 2008.

[49] SEC560: Network Penetration Testing and Ethical Hacking, The SANS Institute, 2012.

[50] James Ransome and Anmol Misra Core Software Security: Security at the Source, CRC Press, 416 pages, 2013.

[51] Jonas Zaddach and Andrei Costin, Embedded Devices Security Firmware Reverse Engineering, Blackhat USA, 2013.

[52] Ian McLoughlin, Secure embedded systems: the threat of reverse engineering, 14th IEEE International Conference on Parallel and Distributed Systems, 2008.

[53] National Institute of Standards and Technology (NIST), Guide to Integrating Forensic Techniques into Incident Response, Special Publication 800-86, August 2006.

[54] Pascal Brangetto, Emin Çalişkan, Henry Roigas, Cyber Red Teaming, Organisational, technical and legal implications in a military context, NATO Cooperative Cyber Defence Centre of Excellence, 2015.

[55] Aaron Guzman and Aditya Gupta, IoT Penetration Testing Cookbook, Packt Publishing Ltd, 415 pages, 2017.

[56] Principles of Chaos Engineering. Online: principlesofchaos.org. Access date: March 1st, 2019.

[57] Ricardo Martins, Resilience testing: breaking software for added reliability, Talkdesk Engineering, September 5, 2017. Online: https://engineering.talkdesk.com/resilience-testing-breaking-software-for-added-reliability-7f1e60207d06, Access date: March 1st, 2019.

[58] Casey Rosenthal, Lorin Hochstein, Aaron Blohowiak, Nora Jones and Ali Basiri, Chaos engineering, Building confidence in system behavior through experiments, September 26, 2017. Online ebook: https://www.oreilly.com/ideas/chaos-engineering, Access date: March 1st, 2019.

[59] Mattia Reggiani, Purple Teaming: A Security-Testing Collaborative, Posted in Application Security, Penetration Testing, on July 26, 2016.

[60] Haissam Ziade, Rafic Ayoubi and Raoul Velazco, A Survey on Fault Injection Techniques, The International Arab Journal of Information Technology, Volume 1, No. 2, July 2004.

[61] Nuno Silva, Ricardo Barbosa, Joao Carlos Cunha and Marco Vieira, A view on the past and future of fault injection, 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Budapest, Hungary, June 24-27, 2013.

[62] Aaron Rinehart, Purple testing and chaos engineering in security experimentation, 14 June, 2018. Online: https://opensource.com/article/18/6/security-experimentation, Access date: March 1st, 2019.

[63] Audrey Dorofee, Robin Ruefle, Mark Zajicek, David McIntire, Christopher Alberts, Samuel Perl, Carly Lauren Huth and Pennie Walters, Incident Management Capability Assessment, Technical Report, CMU/SEI-2018-TR-007, CERT Division, December 2018.

[64] Jason Kick, Cyber Exercise Playbook, The MITRE Corporation, November 2014.

[65] Steve Markey, Testing Your Computer Security Incident Response Plan, Journal Online, ISACA Journal, Volume 2, 2012.

[66] Patrick Kral, Incident Handler's Handbook, SANS Institute, Information Security, GIAC (GCIH) Gold Certification, 2011.

[67] Richard A. Caralli, Julia H. Allen and David W. White, CERT Resilience Management Model (CERT-RMM): A Maturity Model for Managing Operational Resilience (SEI Series in Software Engineering), 1st Edition, Kindle Edition, Addison-Wesley Professional, 2010.

[68] Requirements Traceability Matrix, Online: https://www2a.cdc.gov/cdcup/library/templates/default.htm, Access date: 9 August 2018.

[69] Harold F. Tipton, M. Krause, Information Security Management Handbook, Fifth Edition, Volume 3, CRC Press, January 13, 2006.

[70] Deborah Bodeau and Richard Graubart, Cyber Resilience Metrics: Key Observations, The MITRE Corporation, 2016.