



HAL
open science

CyberSec4Europe D3.21 - Framework to design and implement adaptive security systems

Alzubair Hassan, Dimitri Van Landuyt, Liliana Pasquale, Manuel Cheminod,
Marko Kompara, Panayiotis Kotzanikolaou, Romain Laborde, Susana
Gonzalez

► To cite this version:

Alzubair Hassan, Dimitri Van Landuyt, Liliana Pasquale, Manuel Cheminod, Marko Kompara, et al.. CyberSec4Europe D3.21 - Framework to design and implement adaptive security systems. [Research Report] D3.21, University college Dublin; KU Leuven; Consiglio Nazionale delle Ricerche; University of Maribor; University of Piraeus Research Centre; IRIT - Institut de Recherche en Informatique de Toulouse; ATOS. 2022. hal-03762517

HAL Id: hal-03762517

<https://hal.science/hal-03762517>

Submitted on 28 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Cyber Security for Europe

D3.21

Framework to design and implement adaptive security systems

Document Identification	
Due date	31 March 2022
Submission date	31 March 2022
Revision	3.0

Related WP	WP3	Dissemination Level	CO
Lead Participant	UCD	Lead Author	Liliana Pasquale (UCD) and Alzubair Hassan (UCD)
Contributing Beneficiaries	ATOS, CNR, KUL, UM, UPRC, UPS-IRIT	Related Deliverables	D3.4

Abstract: The research results presented in this document represent the deliverable D3.21: “Framework to design and implement adaptive security systems” and are the continuation of the deliverable D3.4: “Analysis of key research challenges for adaptive security”. In this deliverable we aim to validate the challenges elicited in D3.4 with practitioners from industry. We also aim to understand how the assets developed by each partner in task T3.5 could be used to engineer an adaptive security system.

In this deliverable, we present the results of a survey conducted with security practitioners from industry. The objective of the survey is to acquire the practitioners’ perspective on the application of adaptive security technologies and identify the research challenges that practitioners perceive to be the most critical in adaptive security. We use a reference architecture of an adaptive security system to identify the contributions of the assets provided by each partner. We showcase how the joint contributions of each asset can support some of the maritime transport use case scenarios elicited in task T5.5 Maritime Transport.

This document is issued within the CyberSec4Europe project. This project has received funding from the European Union's Horizon 2020 Programme under grant agreement no. 830929. This document and its content are the property of the CyberSec4Europe Consortium. All rights relevant to this document are determined by the applicable laws. Access to this document does not grant any right or license on the document or its contents. This document or its contents are not to be used or treated in any manner inconsistent with the rights or interests of the CyberSec4Europe Consortium and are not to be disclosed externally without prior written consent from the CyberSec4Europe Partners. Each CyberSec4Europe Partner may use this document in conformity with the CyberSec4Europe Consortium Grant Agreement provisions and the Consortium Agreement.



The information in this document is provided as is, and no warranty is given or implied that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

Executive Summary

The research results presented in this document extend the deliverable D3.4: “Analysis of key research challenges for adaptive security”. Deliverable D3.4 presents the results of a systematic literature review that was conducted to survey existing research on adaptive security.

In this deliverable we aim to validate the challenges elicited in deliverable D3.4 with practitioners from industry. We also aim to understand how the assets developed by each partner in task T3.5 could be used to engineer an adaptive security system.

In the first part of the deliverable, we present the results of a survey conducted with security practitioners from industry. The survey has the following objectives: a) acquire the practitioners’ perspective on the application of adaptive security technologies and their level of automation; b) understand the involvement of stakeholders in the execution of adaptive security tasks and c) identify the research challenges that practitioners perceive to be the most critical in adaptive security. From the responses collected from the participants we observed that practitioners have limited familiarity with adaptive security technologies, especially with adaptive access control and adaptive authentication technologies. The activities for which participants considered that the input from engineers was most beneficial are security monitoring, risk assessment, security logging and malware detection. The reason is that security monitoring/logging and risk assessment are activities that can hardly be automated and require to be configured or performed entirely by a security engineer.

The top 3 research challenges that practitioners considered to be the most important are related to 1) the provisioning of assurances in adaptive security solutions; 2) the ranking of protection strategies that can be used to prevent or contain threats; 3) the use of machine learning techniques to evaluate effectiveness of protection strategies in the long-term.

In the second part of the deliverable, we use a reference architecture for adaptive security systems to identify the contributions of the assets provided by each partner. We showcase how the joint contributions of each asset can support some of the maritime transport use case scenarios elicited in task T5.5 Maritime Transport. Our assets provide the following contributions:

- *Data-flow-centric threat assessment*: an approach to automate threat elicitation and identify evolving threats when the architecture of the system changes.
- *Adaptive Authentication*: a decision-making technique to automatically select an authentication method that mitigates the security risks and maximizes the satisfaction of security and other requirements, such as performance and usability.
- *Situation-driven risk assessment and security enforcement*: an approach to identify situations and pre-compute security risks based on the utilized assets, threats, vulnerabilities and impacts. Depending on the security risks this approach can enforce situation-specific security controls that effectively mitigate the risks.
- *Adaptive risk assessment*: a novel technique to verify effectiveness of security controls when changing scenarios affect interdependencies between system components.
- *Adaptive incident reporting*: a novel approach to incident reporting that can adaptively change the reporting process and the report template depending on the type of incident and the location of the components affected by the security incidents.

Finally, we identify General Data Protection Regulation (GDPR) compliance issues that can arise in adaptive security systems.

Document information

Contributors

Name	Partner
Alzubair Hassan	UCD
Dimitri Van Landuyt	KUL
Liliana Pasquale	UCD
Manuel Cheminod	CNR
Marko Kompara	UM
Panayiotis Kotzanikolaou	UPRC
Romain Laborde	UPS-IRIT
Susana Gonzalez Zarzosa	ATOS

Reviewers

Name	Partner
Pablo Fernandez Saura	UMU
Alberto Lluch Lafuente	DTU

History

Version	Date	Authors	Comment
0.01	2021-05-26	Liliana Pasquale	Table of contents
0.1	2021-07-21	Liliana Pasquale, Alzubair Hassan, Susana Gonzalez, Zarzosa, Manuel Cheminod, Dimitri Van Landuyt, Marko Kompara, Panayiotis Kotzanikolaou, Romain Laborde	1 st Draft including motivating example, application scenario for each asset and analysis of the results of the industrial survey
0.2	2021-11-16	Panayiotis Kotzanikolaou	Added consolidated maritime motivating example
0.3	2021-12-17	Liliana Pasquale, Alzubair Hassan, Susana Gonzalez, Zarzosa, Dimitri Van Landuyt, Marko Kompara, Panayiotis Kotzanikolaou, Romain Laborde	Each partner added information about the application scenario of each asset, a description of the asset and exemplification of the asset functionalities using the application scenario
0.4	2022-01-14	Manuel Cheminod	Added asset contributions for CNR
1.0	2022-02-14	Liliana Pasquale, Alzubair Hassan	2 nd Consolidated draft of the deliverable
1.1	2022-03-04	Liliana Pasquale, Alzubair Hassan, Susana Gonzalez, Zarzosa, Dimitri Van Landuyt, Marko Kompara, Panayiotis Kotzanikolaou, Romain Laborde	Included novelty of each asset w.r.t. state of the art. The adaptive security architecture was put in the framework of D3.1 and D3.12
2.0	2022-03-18	Liliana Pasquale	3 rd draft of the deliverable including abstract, introduction, executive summary, and conclusion.
3.0	2022-03-25	Liliana Pasquale	Final version of the deliverable
3.0	2022-04-01	Ahad Niknia	Final check, preparation and submission process

Table of Contents

1	Introduction.....	1
1.1	Motivation.....	1
1.2	General Objectives.....	1
1.3	Participants.....	2
1.4	Organization of the Deliverable.....	2
2	Industry Survey on Adaptive Security.....	3
2.1	Objectives.....	3
2.2	Methodology.....	3
2.3	Results.....	4
2.3.1	Practitioners’ familiarity with adaptive security technologies and their level of automation.....	5
2.3.2	Involvement of users and engineers in adaptive security.....	8
2.3.3	Practical research challenges for adaptive security.....	9
3	Motivating Example and Reference Architecture.....	11
3.1	Overview of the Maritime Transport Example.....	11
3.2	A Reference Architecture for Adaptive Security.....	13
4	Security Modelling.....	16
4.1	Data-flow-centric threat assessment.....	16
4.1.1	Asset description.....	16
4.1.2	Application Scenario.....	18
4.1.3	Output.....	21
4.1.4	Asset limitations and future work.....	21
5	Analysis and Planning.....	23
5.1	Adaptive Authentication.....	23
5.1.1	Asset description.....	23
5.1.2	Application scenario.....	24
5.1.3	Adaptive Authentication Framework.....	26
5.1.4	Outputs.....	33
5.1.5	Asset limitations and future work.....	33
5.2	Situation-driven risk assessment and security enforcement framework.....	35
5.2.1	Application scenario.....	35
5.2.2	Asset description.....	36
5.2.3	Outputs.....	49
5.2.4	Asset limitations and future work.....	49
5.3	Adaptive risk assessment.....	50
5.3.1	Application scenario.....	51
5.3.2	Asset description.....	53

5.3.3	Asset limitations and future work	54
6	Execution.....	56
6.1	Adaptive Incident Reporting	56
6.1.1	Application scenario.....	56
6.1.2	Asset description	58
6.1.3	Outputs	66
6.1.4	Asset limitations and future work	66
7	GDPR compliance issues in adaptive security.....	68
7.1	Application scenario	68
7.2	Asset description	69
7.3	DPIA considerations for adaptive security in a maritime environment.....	71
7.4	Outputs.....	74
7.5	Asset limitations and future work	74
8	Conclusion.....	76
	References.....	77

List of Figures

Figure 1. Sectors of participants organizations.	4
Figure 2 Participants' geographical region (on the left), education degree (on the centre), and years of experience (on the right).	4
Figure 3 Familiarity with adaptive security technologies.....	6
Figure 4 Level of Automation of adaptive security technologies.....	7
Figure 5 Activities performed by an adaptive security system.....	8
Figure 6 Adaptive security activities requiring user input.....	8
Figure 7 Adaptive security activities requiring input from an engineer.	9
Figure 8 Maritime transport example.....	11
Figure 9 Adaptive Security Reference Architecture	14
Figure 10 CyberSec4Europe global architecture.	15
Figure 11 Screenshot of the SPARTA graphical editor.....	17
Figure 12 Outcome of threat generation and risk assessment.....	18
Figure 13 DFD of the maritime case.....	19
Figure 14 Illustration of an Elevation of Privilege Threat.....	19
Figure 15 Illustration of a Denial of Service threat.....	19
Figure 16 Threat count and inherent risk calculated.....	19
Figure 17 Illustration of the DFD as a superset of the vessel in different situations.	20
Figure 18 Adaptive Authentication Asset.....	24
Figure 19 Requirements and Contextual Factors.....	29
Figure 20 Extended Feature model of authentication methods.....	30
Figure 21 Security Risks.....	31
Figure 22 Abstract model of the Fuzzy Causal Network.....	32
Figure 23 The security situation of a vessel during the cargo transfer service.....	36
Figure 24 An overview of the MITIGATE methodology.....	38
Figure 25 An overview of the DynSMAUG methodology.....	40
Figure 26 An adaptive situation-driven risk assessment and risk enforcement framework.	41
Figure 27 Decision-tree based situation elicitation.....	42
Figure 28 Relationships among risk related entities and relevant datasets.....	43
Figure 29 Vulnerability level calculation matrix.....	45
Figure 30 Impact level calculation matrix.....	46
Figure 31 Situational impact calculation matrix.....	46
Figure 32 High level view of Scenario S1.....	51
Figure 33 Attack scenario in S2.....	52

Figure 34 Complex S2 scenario including port and other vessels.	53
Figure 35 SYSVER architecture.	54
Figure 36 The AIRE architecture.	59
Figure 37 Incident reporting BPMN.	61
Figure 38 Incident reporting escalation procedure BPMN.	62
Figure 39 AIRE-workflow-enforcement APIs interaction diagram.	63
Figure 40 AIRE-reports-generator template mapping file example.	64
Figure 41 AIRE-thehive-plugin APIs interaction diagram.	65
Figure 42 The main steps of the DPIA template.	70

List of Tables

Table 1 Examples of attacks on maritime transport systems.	13
Table 2 Nodes of the Fuzzy Causal Network.....	31
Table 3 Functions used to aggregate the values of the nodes of the FCN.	33
Table 4 AIRE-workflow-enforcement APIs.	63
Table 5 AIRE-reports-generator service.	64

List of Acronyms

<i>A</i>	AES	Advanced Encryption Standard
	AIRE	Atos Incident Reporting Engine
	AIS	Automatic Identification System
	ATT&CK	Adversary Tactics Techniques and Common Knowledge
<i>B</i>	BPMN	Business Process Modelling Notation
<i>C</i>	CAPEC	Common Attack Pattern Enumeration and Classification
	CSIRT	Computer Security Incident Response Team
	CVE	Common Vulnerability Enumeration
	CVSS	Common Vulnerability Scoring System
<i>D</i>	DFD	Data Flow Diagram
	DPIA	Data Protection Impact Assessment
	DPO	Data Protection Officer
<i>E</i>	ECDIS	Electronic Chart Display Information System
	EDPB	European Data Protection Board
	EMF	Eclipse Modelling Framework
	ENISA	European Union Agency for Cybersecurity
<i>F</i>	FCN	Fuzzy Causal Network
<i>G</i>	GDPR	General Data Protection Regulation
	GNSS	Global Navigation Satellite System
<i>I</i>	ISO	International Organization for Standardization
	ISPS	International Ship and Port Facility Security
<i>L</i>	LINDDUN	Linkability, Identifiability, Non-Repudiation, Detectability, Disclosure of information, Unawareness, Non-compliance
<i>M</i>	MiTM	Man in The Middle
<i>N</i>	NISD	Network and Information Security Directive

	NIST	National Institute of Standards and Technology
	NVD	National Vulnerability Database
<i>O</i>	OCR	Optical Character Recognition
	OES	Operators of Essential Services
	OSINT	Open Source Intelligence
	OTP	One-Time Password
<i>P</i>	PMS	Port Management System
<i>R</i>	RFID	Radio Frequency Identification Devices
<i>S</i>	SCADA	Supervisory Control and Data Acquisition
	SHA	Secure Hash Algorithm
	SIEM	Security Information and Event Manager
	SOAR	Security Orchestration, Automation and Response
	SQL	Standard Query Language
	SSL	Secure Socket Layer
	STRIDE	Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of privilege
	SYSVER	SYStem VERification
<i>T</i>	TLS	Transport Layer Security
	TOS	Terminal Operating System
<i>V</i>	VPN	Virtual Private Network
	VTS	Vessel Traffic Service

1 Introduction

1.1 Motivation

Software systems are traditionally developed by enacting “static” security controls. However, unanticipated changes can occur in the environment where the system operates (e.g., new assets require to be protected, or new threats may emerge over time), in the system itself (previously unknown vulnerabilities are discovered) and/or in the security properties (i.e., confidentiality, integrity, availability and accountability (CIAA) [1]) that a system must satisfy. These changes may render the security controls deployed ineffective, making the system more vulnerable to potential attacks. In addition, the static approach may lead to systems that are difficult to evolve, that incur substantial technical debt over time or become brittle when faced with changes over longer periods of time. To address these issues, different approaches have been proposed in previous research to build *adaptive security systems* [2], which can *self-protect* [3] from the varying risk of harm by adjusting their security controls, in a way that minimally impacts other system requirements.

Surveys [4-6] about adaptive security systems demonstrate that existing research in this domain has only thrived recently. In our previous deliverable D3.4 “Analysis of key research challenges for adaptive security” [7] we surveyed previous work on adaptive security and identified the gaps in engineering adaptive security systems. We investigated application domains where adaptive security systems have been deployed and trends, patterns, and gaps in existing research. We noticed that existing solutions do not consider the application of security controls in both the cyber and physical spaces where the system operates. We also noticed that very limited research was performed to encourage stakeholders’ involvement in the activities of an adaptive security system. We found the need to provide systematic approaches to generate assurances and explanations demonstrating compliance with existing security standards and regulations, even in the presence of adaptation. Finally, we discovered that adaptive security approaches are mainly focused on the prevention of security threats and do not consider other post-incident activities, such as incident containment, investigation, and reporting.

Finally, at the end of D3.4 we provided a set of recommendations for future research directions on adaptive security. These recommendations include:

- The application of adaptive security approaches to cyber-physical systems;
- An explicit consideration of security requirements in the design and development of the activities of an adaptive security system;
- The possibility of involving stakeholders, such as engineers and users, in the execution of the activities of an adaptive security system, when those cannot be automated.

However, in our previous deliverable, we did not gather an industrial view on the challenges and future research directions of adaptive security. We also did not validate with practitioners the recommendations that we provided to guide future research on adaptive security systems.

1.2 General Objectives

The objective of this deliverable is two-fold. First, we aim to acquire an industrial perspective on the application of adaptive security technologies in practice, the involvement of stakeholders, such as users and engineers, in the execution of the activities of an adaptive security system, and the research challenges and directions that practitioners perceive to be most critical in adaptive security. To achieve this aim, in the CyberSec4Europe project we conducted a survey with security practitioners from industry. Second, in this deliverable we aim to assess how the security assets developed by the partners involved in task 3.5 “Adaptive Security” could be applied to engineer the activities of an adaptive security system. To achieve this aim, we use the Maritime Transport use cases developed in task T5.5 as a case study for this deliverable. More precisely, for each security asset we specifically identify an application scenario of the maritime

transport system that it can support. We also identify the activities of the MAPE (Monitoring, Analysis, Planning and Execution) adaptive security loop that each asset can cover. Finally, we identify limitations of each asset that should be addressed in future research.

1.3 Participants

The deliverable was led by Liliana Pasquale (UCD) and was collaboratively edited by Alzubair Hassan (UCD), Susana Gonzalez Zarzosa (ATOS), Manuel Cheminod (CNR), Dimitri Van Landuyt (KUL), Romain Laborde (UPS-IRIT), Panayotis Kotzanikolaou, (UPRC), and Marko Kompara (UM).

1.4 Organization of the Deliverable

The rest of the deliverable is organized as follows. Section 2 discusses the results of the survey conducted with practitioners from industry. Section 3 describes the maritime transport example and the adaptive security reference architecture that we considered to discuss the joint contributions of the assets proposed in this task T3.5. Sections 4 illustrates how some of the assets can support adaptive threat assessment for the elicitation of security requirements of an adaptive security system. Section 5 showcases how some of the assets can support the analysis and planning activities of an adaptive security system. Section 6 describes how some of the assets can support adaptive incident reporting. Section 7 discusses GDPR compliance issues in adaptive security. Finally, Section 8 summarizes the main findings of the deliverable.

2 Industry Survey on Adaptive Security

2.1 Objectives

In this Section we report on the results of the industry survey that we conducted with security practitioners in the CyberSec4Europe project. The aim of the survey is to acquire a perspective from the industry on the application of adaptive security technologies in practice, to understand the involvement of engineers and users in the configuration and management of adaptive security technologies and identify the research challenges and directions that practitioners perceive to be most critical in adaptive security. More precisely, the industry survey aims to:

- Assess practitioner's familiarity with existing adaptive security technologies and their level of automation;
- Understand the involvement of end users and engineers in the configuration, use and management of different tasks supported by adaptive security technologies;
- Identify the research challenges that practitioners deem to be most relevant to support the adoption of adaptive security technologies.

2.2 Methodology

Our survey represents a cross-sectional study [8] targeting individuals who hold, or have held in the past, a security role in a company. We recruited security practitioners by sending personal invitations to the industry partners of the CyberSec4Europe security competence network and to our industrial collaborators. To sample the target population, we used a non-probabilistic sampling method [9], particularly convenience sampling, because the target population is very specific and has limited availability. In other words, we obtained responses only from the participants that were available and willing to take part in the study. To increase participation, we rewarded 10 randomly selected participants with a 50 Euros Amazon voucher.

To validate the survey, we conducted two focus groups, during which the content of the survey was reviewed systematically to assess the completeness and ensure that no duplication was present. Since we developed a new survey instrument in a topic area that has not been researched previously, this was the only form of preliminary validation available [8]. Participants of the focus group include all the authors of this paper who are themselves cybersecurity researchers and practitioners. This ensured that members of the target population were included in the focus group.

The survey is divided into four parts. In the first part of the survey, we collect demographic information about participants, such as geographical region, sector and size of each participant's organization, the security role held by the participants, the years of experience and the level of education held by the participants. In the second part of the survey, we gather participants' familiarity with the notion of adaptive security and adaptive security technologies and the level of automation of adaptive security technologies. In the third part of the survey, we understand the involvement of users and engineers in different tasks performed by an adaptive security system, such as security monitoring, threat detection and the execution of security controls. Finally, in the last part of the survey, we collect the participants' view about what in their opinion are the most critical adaptive security research challenges. The questions of the survey can be viewed online [10].

We collected responses to the survey between October 2020 and January 2021. We obtained a total of 58 responses, from which we ruled out 26 invalid responses. We considered invalid (and therefore excluded) the responses provided in a very short time (< 5 minutes) and those that were incomplete (e.g., only including demographic information). 81.82% of the participants hold or have held a security role in the past. We only kept responses from participants who have never held a security role, under the assumption that they had a higher level of education (PhD or PostDoc).

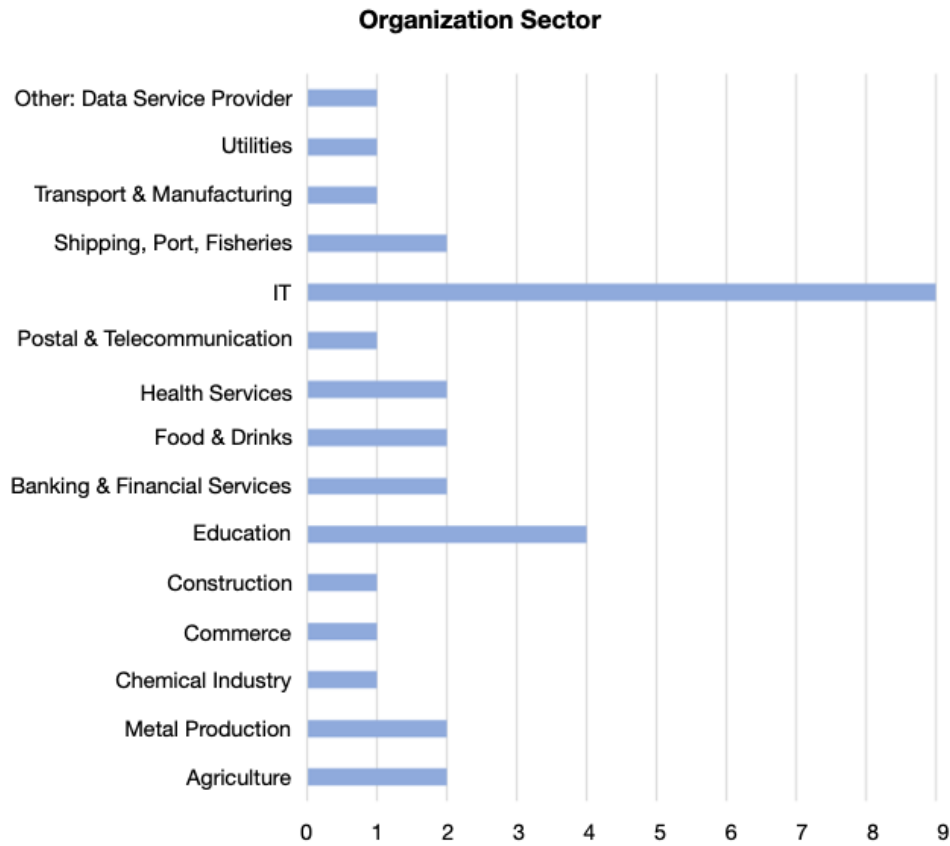


Figure 1. Sectors of participants organizations.

As shown in Figure 1, the participants in our survey were working in different sectors, from IT to education, transport, manufacturing, and health services. As shown on the left of Figure 2, although 75% of the participants worked for an organization based in Europe, we were able to include participants from other geographical regions, such as the USA and Middle East. Finally, as shown in the center and on the right of Figure 2, our participants had varying education degrees and years of experience.

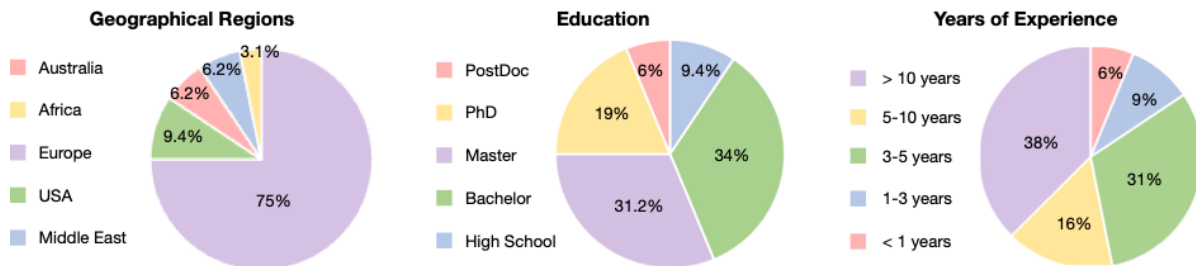


Figure 2 Participants' geographical region (on the left), education degree (on the centre), and years of experience (on the right).

2.3 Results

In this section we report and discuss the results obtained from the industrial survey on adaptive security.

2.3.1 Practitioners' familiarity with adaptive security technologies and their level of automation

In the second part of the survey, we provided the following definition of adaptive security formulated by Gartner in 2017 [10] and asked participants to rate their familiarity with such definition.

Adaptive security is a security model where the monitoring of threats remains continuous and improves as cybersecurity risks change and evolve over time. Adaptive security allows for early detection of security compromises and an automatic, autonomous response when a malicious event occurs.

We observed that 7 (21.9%) participants were very familiar with the notion of adaptive security, 20 (62.5%) participants were somewhat familiar, while 5 (15.6%) participants were not familiar with the notion of adaptive security. We could not identify a statistically significant correlation between any of the demographic information collected about participants and their familiarity with the notion of adaptive security.

When preparing the survey, we conducted 2 focus groups where we brainstormed and identified the following adaptive security technologies:

- **Adaptive Risk Assessment:** provides identification, assessment, and mitigation of real-time risks by combining and correlating security-related information extracted from multiple relevant sources. In addition, such a system enables the operators to simulate risks and take relevant risk mitigation actions.
- **Adaptive Access Control:** can change access rights depending on, for example, user roles, departments, days, times, locations, and security risks. Such a system balances the need to manage risk with the desire to improve the user experience.
- **Adaptive Authentication:** attempts to match the required authentication credentials to the perceived risk of the connection or authorization requested. The objective is to try to reduce the authentication burden on users on the one hand, while enforcing strong authentication where it is most needed, on the other. For example, a user connecting via VPN from his/her home network using a company-managed PC will not be required to present any additional authentication credentials beyond those provided by his/her PC because the connection requested is not perceived to be high-risk. A connection from an unknown WiFi during “odd” hours of the day would require the user to produce additional authentication in the form of a password, OTP, or both, because the connection exhibits risk indicators that elevate the perceived risk.
- **Adaptive Cryptography** allows changing the cryptographic schemes used, for example, to encrypt information (e.g., AES vs SPECK) and support message authentication (e.g., HMAC vs SHA3 vs BLAKE2). The selected cryptographic scheme can depend on different factors, such as the device battery consumption, or the risk of an attack (e.g., man in the middle, session hijacking).
- **Adaptive Intrusion Detection** can change the criteria (e.g., signatures, representation of normal behavior) adopted to identify suspicious activities dynamically. For example, it can use feedback information provided by a human operator or an automated classifier to update the set of signatures adopted to identify an anomaly.
- **Adaptive malware detection** can change the criteria (e.g., anomalous features in the application implementation) adopted to recognize a malware. For example, these criteria can be extracted from the analysis of a set of benign or malicious apps that can be updated continuously.
- **Adaptive penetration testing** can change the types and order in which authorized simulated cyberattacks are conducted against a target computer system, for example, depending on the business risks, the implementation of a target application, information from bug tracking systems and the configuration of the target platform where a target application is deployed.
- **Adaptive SIEM (Security Information and Event Management)** can evolve the correlation rules adopted to catch abnormal behavior, for example, using the input from a human expert or automatically, using machine learning.

- **Adaptive SOAR (Security Orchestration, Automation and Response)** is a platform for orchestrating and automating incident response that can adapt its response to cyber-attacks depending on the security infrastructure, the incident conditions, and the applicable regulations.

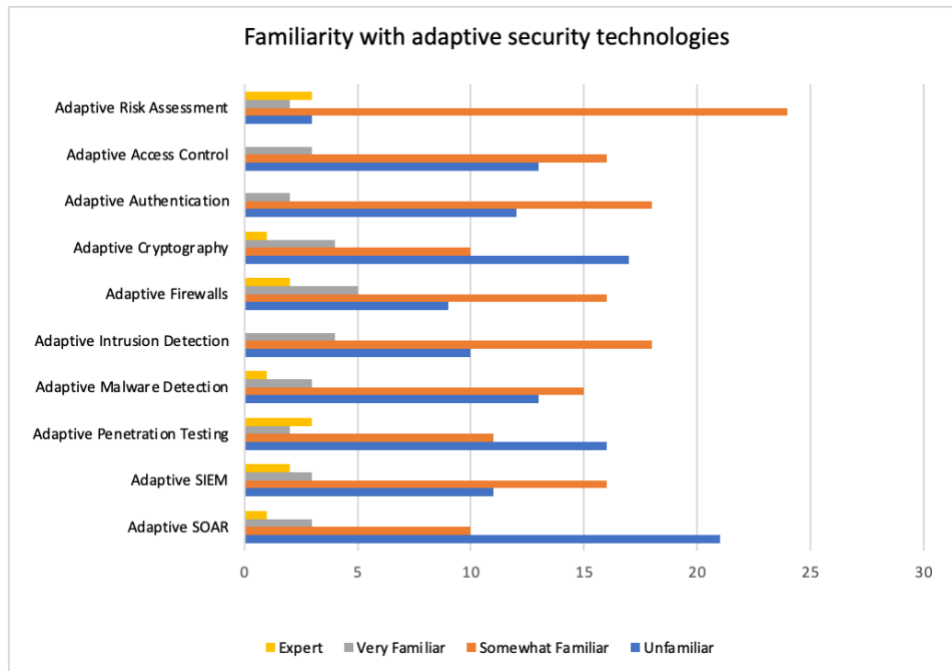


Figure 3 Familiarity with adaptive security technologies

As shown in Figure 3, most participants only had limited or no familiarity with adaptive security technologies. Among the respondents, 3 claimed to be experts in adaptive risk assessment, 2 of them mentioned having acquired expertise by working on the DiESEM EU project¹. Within this project the participants developed a model to show the evolution of risk depending on a specific asset and/or application. 1 participant mentioned having developed research on processing OSINT data and features of internal assets to measure risks dynamically.

No participants claimed to be an expert on adaptive access control and adaptive authentication. 2 participants claimed to be experts on adaptive cryptographic systems. One of them implemented adaptive cryptographic and hashing functions within existing software products, while the other used an adaptive cryptographic system.

2 participants claimed to be experts on adaptive firewalls. One of them built an adaptive firewall that could update its rules daily based on OSINT information. While the other expert participant developed a next generation web application firewall utilizing machine learning techniques to adapt to the latest threats. 2 of the participants who claimed to be very familiar with adaptive firewalls installed and/or used commercial adaptive firewalls, such as CISCO Meraki² and Palo Alto³ in their organization.

No participant claimed to be an expert of adaptive intrusion detection systems. Only 1 participant claimed to be an expert of adaptive malware detection since they implemented such technology in a commercial

¹ <https://cordis.europa.eu/project/id/958339>

² <https://meraki.cisco.com/>

³ <https://www.paloaltonetworks.com/network-security/next-generation-firewall>

product. 3 participants claimed to be experts of adaptive penetration testing although none of them has had opportunities to implement adaptive penetration testing in an existing tool.

2 participants claimed to be experts of adaptive SIEM since they contributed to the implementation of such technology in an existing tool/product. The participants who claimed to be very familiar with adaptive SIEM mentioned to have used software products with adaptive SIEM technology, such Splunk⁴, AlienVault⁵, ArcSight⁶, Elasticsearch⁷ and Demisto⁸. Only 1 participant claimed to be an expert of adaptive SOAR technologies since they were involved in the integration of such technologies in Ueba⁹ and Sentinel¹⁰ software products.

The participants who were at least somewhat familiar with adaptive security technology were also asked to indicate the level of automation of the technology they were exposed to. For each adaptive security technology Figure 4 shows a percentage of respondents who indicated a given automation level (fully automated, semi-automated, not automated). The adaptive security technologies that have the highest support for full automation are adaptive firewalls (32% of respondents), intrusion detection (30% of respondents), malware detection (22% of respondents) and SIEM (20% of respondents). Adaptive access control and SOAR are the only two technologies for which no respondent indicated full automation.

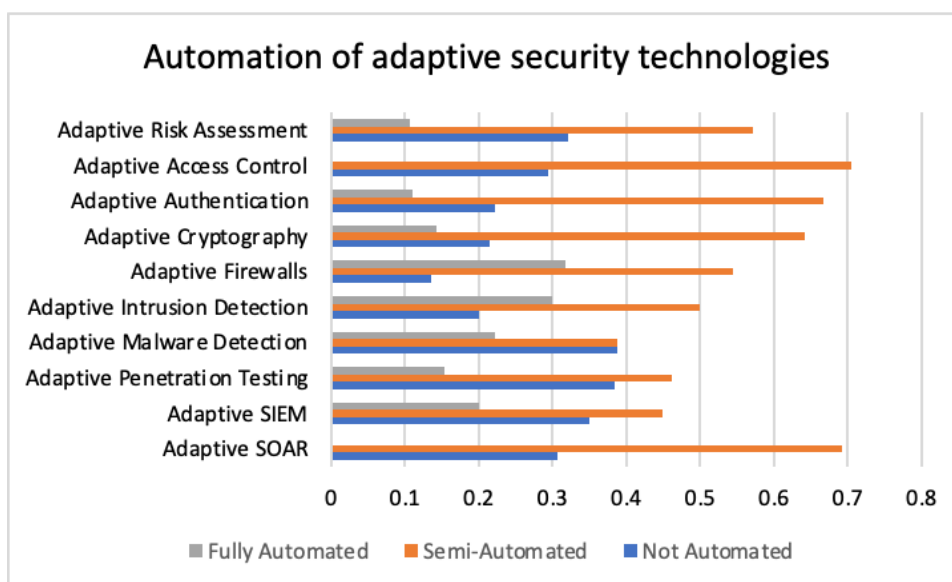


Figure 4 Level of Automation of adaptive security technologies.

Finally, we asked participants to indicate what were the activities performed by the adaptive security systems or products that they have used. We chose to distinguish between the following activities: security monitoring, logging, risk assessment, malware detection and prevention of security breaches, for example, through system hardening and/or isolation. As shown in Figure 5, tool support is mainly provided to perform security monitoring (13 respondents), security logging (11 respondents), risk assessment (11 respondents),

⁴ <https://www.splunk.com/>

⁵ <https://otx.alienvault.com/>

⁶ <https://www.microfocus.com/en-us/cyberres/secops>

⁷ <https://www.elastic.co/>

⁸ <https://apps.paloaltonetworks.com/marketplace/demisto>

⁹ <https://www.exabeam.com/siem-guide/ueba/>

¹⁰ <https://www.exabeam.com/siem-guide/ueba/>

malware detection (9 respondents). Very limited tool support is provided for the design of security policy changes (3 respondents), implementation of baseline systems and posture (3 respondents) and, more generally, post-incident activities, such as incident management and response (2 respondents), incident containment (3 respondents) and investigation (1 respondent).

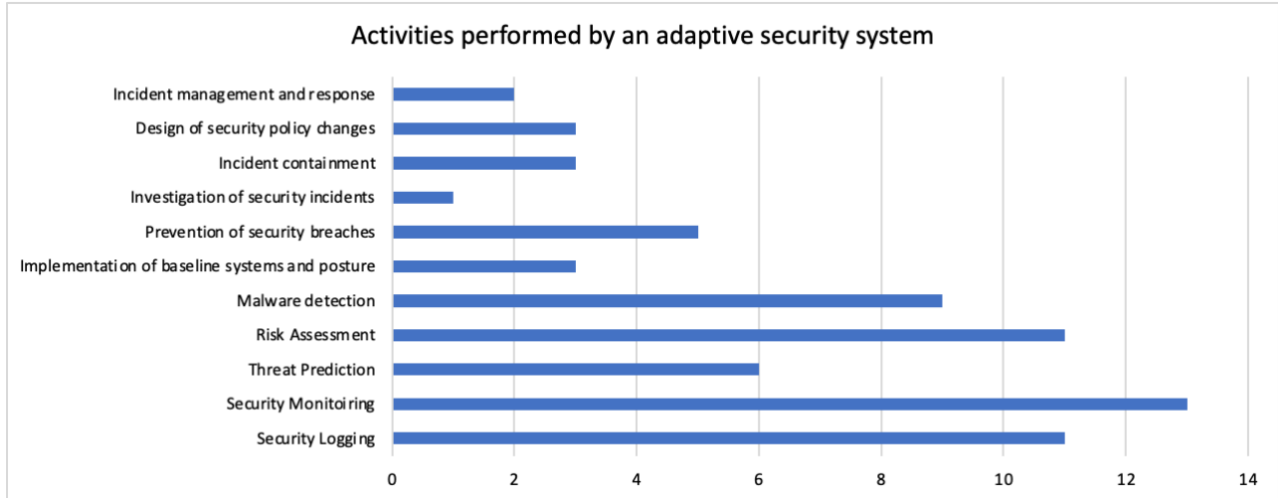


Figure 5 Activities performed by an adaptive security system.

2.3.2 Involvement of users and engineers in adaptive security

In the the third part of our survey, we asked participants to indicate the involvement of stakeholders in adaptive security operations in industry. We focused on users who use an adaptive security system, and engineers, who design, build, and configure an adaptive security system.

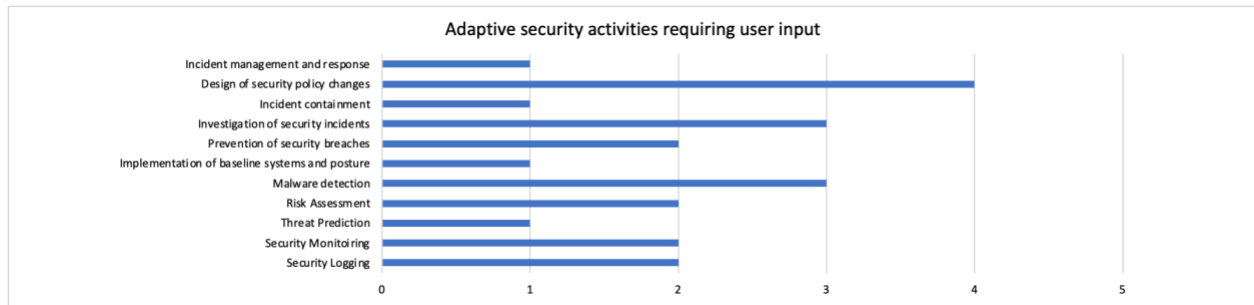


Figure 6 Adaptive security activities requiring user input.

Figure 6 shows the number of participants who indicated that a specific adaptive security activity they were familiar with requires user input. Only 9 respondents indicated that input from the user was required for at least a specific adaptive security activity. Among the activities for which user input was most required, the design of security policy changes was indicated by 4 participants, because information about security policies may require to be provided manually by a user. Investigation of security incidents received 3 responses, mainly because investigations may require personal information about users, such as user ID and geolocation. Malware detection received 3 responses justified by the fact that a user may need to provide information about attack scenarios and detection models. Some respondents indicated that *although users do not provide information, they may need to receive information about risk assessment reports*. Others indicated that *no user input is required because users do not have sufficient security expertise and, instead,*

input to adaptive security activities should be provided by vendors. Other participants also discouraged the provisioning of user input due to their limited security expertise.

Figure 7 shows the number of participants who indicated that a specific adaptive security activity they were familiar with required input from an engineer. All respondents indicated that at least 1 security activity required input from an engineer. The activities for which such input was most required are security monitoring (7 participants), and risk assessment (6 participants), followed by security logging and malware detection (3 participants each). The reason is that security monitoring/logging and risk assessment are activities that can hardly be automated and require to be configured or performed entirely by a security engineer. Regarding malware detection, an engineer can provide information about malware detection mechanisms and have substantial knowledge about the configuration of the network.

Some participants also indicated that the implementation of baseline systems and posture would benefit from the input of an engineer, *who can indicate detection baselines and analytics mode to use.* None of the participants indicated that post-incident activities such as incident containment and prevention required input from an engineer. Two participants indicated that activities related to incident detection and mitigations *would highly benefit by input, to reduce false positives.* Another participant indicated that an engineer could contribute to the implementation of adaptive security policies.

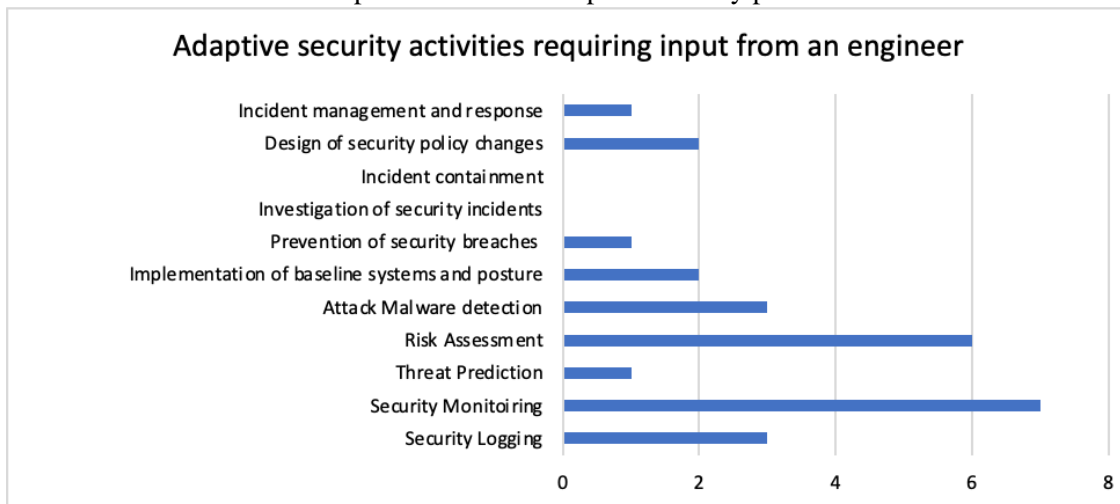


Figure 7 Adaptive security activities requiring input from an engineer.

2.3.3 Practical research challenges for adaptive security

Finally, in the last part of our survey, we asked participants to rank what in their opinion were the most relevant challenges on adaptive security. 6 of the challenges identified in deliverable D3.4 were deemed extremely relevant by at least 8 participants. We ranked them as follows:

1. Existing security solutions should provide assurances to demonstrate that a sufficient level of protection has been achieved (Rated as extremely relevant by 12 participants).
2. Existing security solutions should suggest and rank protection strategies (hardening/isolation) that can be used to prevent/contain threats (Rated as extremely relevant by 10 participants).
3. Existing adaptive security solutions should be equipped with learning capabilities to evaluate whether the enacted security activities (e.g., logging, monitoring, threat prediction, risk assessment, hardening, incident containment) are still effective or need to be changed (Rated as extremely relevant by 10 participants).
4. Decentralized approaches should be investigated to make adaptive security solutions more scalable to large and distributed systems (Rated as extremely relevant by 9 participants).

5. Security engineers should be involved in providing feedback information to support some of the activities of an adaptive security system (Rated as extremely relevant by 9 participants).
6. Adaptive security solutions should be employed and designed to prevent or mitigate interrelated cyber and physical threats in cyber-physical systems (Rated as extremely relevant by 8 participants).

The survey respondents also identified the following challenges for adaptive security:

- Development of techniques to help analysts perform triaging and investigation of security incident by only showing relevant information and security alerts;
- Development of adaptive security solutions that are simple to manage and intuitive to use.

3 Motivating Example and Reference Architecture

This section provides a motivating example for adaptive security and describes a reference architecture for adaptive security systems that are used to introduce the research questions that are investigated in this deliverable.

3.1 Overview of the Maritime Transport Example

Our motivating example is inspired by the maritime transport use cases and demonstrator that are being developed in Task 5.5 [11]. For reasons of simplicity, here we focus on a specific subset of the use cases that are relevant for engineering adaptive security systems.

As shown in Figure 8, most processes in vessels are executed with autonomous or semi-autonomous systems under the control of sophisticated software systems (e.g., Industrial Cyber-Physical Systems, SCADA). Maritime control and navigation systems include the Automatic Identification System (AIS), the Vessel Traffic Service (VTS), and the Electronic Chart Display Information System (ECDIS). The interconnection of all these control systems creates a port-specific SCADA system. AIS is an automatic tracking system mainly used for collision avoidance. It transmits safety related information like course, speed, type of vessel, type of cargo, at-anchor, or underway status. VTS is a marine traffic monitoring system, similar those used in airports, established by port authorities. ECDIS is a navigational chart display that receives data by other control systems, (AIS, GPS, and radars), to allow an officer on deck to navigate the ship. At the port side, the Port Management System (PMS) has a central role; it receives information from the Terminal Operating System (TOS), essential for supply chain management. TOS monitors the location of containers and handling equipment (cranes) through Optical Character Recognition (OCR), Radio Frequency Identification Devices (RFIDs) and GPS systems.

These systems are connected to other maritime operators (e.g., ship owners, trading, transport, and logistics companies) to ensure a seamless and swift data exchange. Vessels at sea are connected via a plethora of communication and data links via satellite communications or conventional radio communications and their navigation is today widely reliant on electronic solutions (e.g., satellite navigation with GPS, Galileo, or electronic chart display information systems, ECDIS).

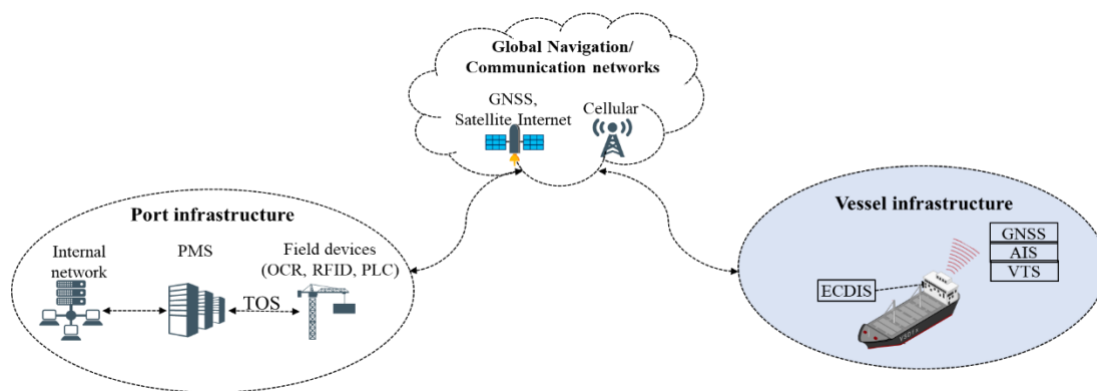


Figure 8 Maritime transport example.

These communication systems play a vital role in maritime operations and provide services falling into the following categories: ship-to-shore, shore-to-ship, and ship-to-ship. However, most of these systems appear to have serious vulnerabilities and threats allowing remote, unauthenticated attackers to compromise their security. Most of these threats and risks are associated with undocumented and/or insecure protocols, weak encryption algorithms and hardcoded credentials.

Representative examples described in [12] are shown in Table 1 and involve:

a) Attacks on maritime electronic navigation systems and Internet services: In [13] attacks against the AIS of existing vessels were presented. By using MiTM attacks, an adversary could hijack and take over AIS communications, tamper with the major online tracking providers and eventually spoof the position of the vessel. Global Navigation Satellite System (GNSS) signals are used even for vessels actively piloted by human operators. But as surface crafts become more autonomous, autopilot systems and dynamic positioning systems are designed under the assumption that GNSS signals are usually available and trustworthy.

In [14] an attack scenario, enabled by a vulnerable on-board mail client (named AmosConnect by Immarsat Solutions) is presented. The vulnerable client could allow unauthenticated attackers to perform blind SQL injection and recover usernames and passwords. Then, with the use of the retrieved credentials, an adversary can remotely execute arbitrary commands with system privileges on the remote system by abusing the Task Manager of the mail client.

In a proof-of-concept attack presented in [15] researchers from University of Texas managed to deviate a maritime surface vessel from its original course, by broadcasting counterfeit civil GPS signals. To remain covert, the spoofed signals were slightly altered. By using search engines like Shodan, a security company named PenTestPartners [16] discovered vulnerable Web interfaces of ship's mission critical systems (e.g., electronic navigation systems). Most of them used weak default passwords, allowed unencrypted HTTP connection without enforcing standard SSL/TLS security and/or were vulnerable to known Web attacks like SQL injection. Various attack scenarios include remotely exploitation of several IT systems of the ship to reveal sensitive information about the ship or the crew and even take control over the ship.

b) Attacks on IoT-enabled port management systems (PMS) and field devices: In a recent study [17] security researchers present an exhaustive analysis of threats and attacks scenarios that include the entire supply chain management such as attacks on Internet-connected port's systems, field devices (OTS, OCR, RFIDs), PLCs and motors that are found mainly installed in yard cranes (ICSA-16-348-05B).

In a real attack incident, an international drug dealer group used hacking techniques that involved the exploitation of the IT systems and services that controlled the movement and location of containers, to illicitly transfer drugs through the port of Antwerp over a two-year period.

Reference	Detailed attack steps	Vulnerabilities exploited	Criticality of attacked system	Potential impact
[13] Demo attacks on AIS vessel tracking system (2013, PoC)	<ol style="list-style-type: none"> 1) Use off-the-shelf H/W to transmit AIVDM messages. 2) Inject spoofed messages in the AIS network (Man-In-The-Water, CPA alerting, signal jamming etc). 3) Force ship to follow a path. 	<ul style="list-style-type: none"> • Weak authentication • AIVDM messages (received data from other vessels) susceptible to tampering • Lack of integrity controls on data context 	The AIS system is directly connected to mission critical systems of a maritime vessel	<p>Pirates introduce fake AIVDM messages to lead a cargo tanker to shallow waters and/or render it invisible.</p> <p>In this scenario the attackers can cause human injuries/fatalities as well as major economic and environmental loss.</p>
[14] Remote attacks on ship's critical systems	<ol style="list-style-type: none"> 1) Search Shodan for exposed web interfaces of 	<ul style="list-style-type: none"> • No/weak authentication mechanisms. 	AmosConnect system is indirectly connected to	Cyber criminals may attempt to take over ship's navigation

<p>using vulnerabilities found on AmosConnect server.</p>	<p>vulnerable systems.</p> <p>2) Recover privileged backdoor account and execute commands with system privileges on the remote system</p> <p>3) Pivot to other segments of the ship's network and locate and takeover ship's mission-critical systems.</p>	<ul style="list-style-type: none"> ● Vulnerable web interfaces. ● Exposure of sensitive data. ● Lack of network segmentation. 	<p>the ship's mission critical systems</p>	<p>systems remotely for ransomware.</p> <p>In this scenario, the attackers can cause human injuries/fatalities as well as major economic and environmental loss.</p>
<p>[17] Attacks on a container port's Internet-connected systems and devices (TOS - OCR - RFID)</p>	<p>1) Use spear phishing techniques gain access to port's internal network.</p> <p>2) Locate vulnerable systems and devices.</p> <p>3) Exploit network and software vulnerabilities to infect the devices.</p>	<ul style="list-style-type: none"> ● No network segmentation/isolation ● Vulnerable network protocols ● Vulnerable/Outdated OS installed ● Lack of authentication 	<p>The infected systems and devices are indirectly connected to the Internet through the company's corporate network</p>	<p>Terrorists infiltrate port's internal networks and infect/remotely control port's OCR - GPS and RFIDs systems to smuggle weapons.</p> <p>In this scenario, the adversaries can harm human lives and cause substantial economic, public trust and confidence loss.</p>

Table 1 Examples of attacks on maritime transport systems.

A vessel is subjected to frequent changes in its position and in the technologies adopted to exchange navigation information. Thus, threats can change dynamically, for example, depending on the proximity of the vessel to other ships and depending on the communication mechanisms used by the vessel to exchange navigation information. Thus, it is necessary to provide approaches that could identify new/changing threats dynamically when the operating conditions of the vessel change. Also, these changes can increase or reduce the security risks and require changing the security controls that are applied in the system. As shown in the examples above, all these attacks were facilitated by lack or insufficient authentication. Finally, if an incident occurs it will be necessary to change the way in which an incident is reported and the information that is shared about the incident, depending on the specific jurisdiction in which the vessel is located.

3.2 A Reference Architecture for Adaptive Security

As shown with the maritime transportation example, it is often not possible to anticipate how security threats can materialise and thus appropriate security countermeasures to prevent them cannot be selected at design time. In contrast, software systems such as these that face unanticipated security threats have to be designed and architected in such a way that they can fundamentally adapt their security countermeasures dynamically, to continue to satisfy some security goals at runtime, i.e. during execution.

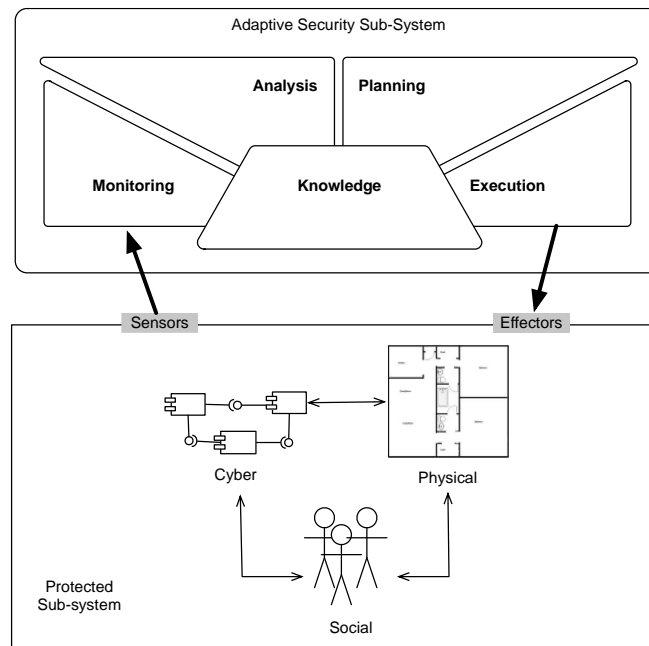


Figure 9 Adaptive Security Reference Architecture

Adaptive security systems [18] are a class of self-adaptive systems able to detect and counteract security threats at runtime. Figure 9 shows a reference architecture for adaptive security systems. Similarly, to other adaptive systems – this architecture generally comprises two sub-systems: an *adaptive security sub-system* is concerned with satisfaction of security goals and manages a *protected sub-system* concerned with the domain functionality. The protected sub-system does not only include software components (cyber), but also physical components and people.

The adaptive security sub-system is structured in accordance with the MAPE-K (Monitor-Analyze-Plan-Execute) loop architecture [19]. This is traditionally considered to be the reference architecture to engineer adaptive systems [20]. It *monitors* (M) the protected sub-system and its operating environment and maintains an updated representation of the protected sub-system and its operating environment at runtime (*Knowledge* - K). The Knowledge also represents information about threats and security requirements. The protecting sub-system uses this representation to *analyze* (A) security threats and assess security risks, and *plan* (P) and *execute* (E) countermeasures aimed to prevent or thwart the threats discovered during analysis. The protected subsystem (e.g., maritime transport system) is the system to be protected, which interacts with the cyber, physical, and social spaces characterizing its operating environment.

A variety of stakeholders [21] are generally involved in the design and use of an (adaptive) security system. For example, software engineers and system operators are those who may need to assess the security posture of the system and decide which security controls are effective to satisfy some security goals. They can elicit security requirements, perform risk assessment, and design, implement, and release security controls that should be enforced. Software users/end users (e.g., passengers and crew operators) directly use and interact with the system to be protected (the protected sub-system). Business units may contribute to the elicitation of security goals in the form of security policies, whereas legal/regulatory units (e.g., port, and custom authorities) may impose security and related regulations (e.g., the GDPR privacy regulation) and standards (e.g., ISO27001) that the system must comply with.

As part of the CyberSec4Europe project, we integrated the activities of an adaptive security system within the CyberSec4Europe global architecture shown in Figure 10 [22]. The activities in green represent the activities of an adaptive security system. Representation of security-relevant knowledge, including security

requirements and threats is performed as part of the Risk & Incident Management and the Security Modelling activities. Monitoring is performed by the Cybersecurity Awareness component, which can be implemented using SIEMs. Analysis is supported by the Risk Analysis and Assessment components which are tasked to identify and assess security threats. Planning is performed by the Incident and Impact Assessment that allows reasoning about the potential impact of incidents and identify security controls to mitigate those incidents. Finally, the Execution activity is supporting by the Reaction component which is tasked to enforce security controls and incident reporting activities after a security incident occurs.

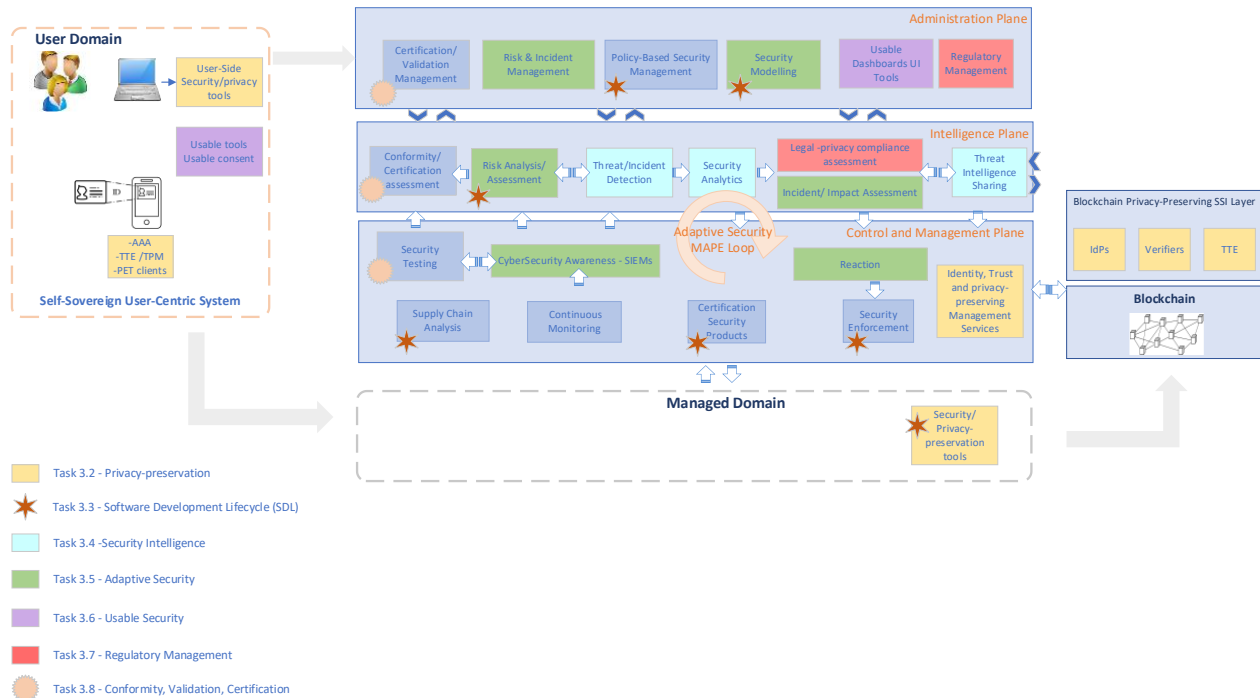


Figure 10 CyberSec4Europe global architecture.

In the rest of the deliverable, we describe the adaptive security assets developed in this task and map them to the activities of the MAPE-K loop.

4 Security Modelling

This Section describes the activities that we conducted to model security-relevant knowledge, particularly security threats.

4.1 Data-flow-centric threat assessment

Threat modelling is a requirement analysis and engineering activity that involves enumerating abuse or misuse cases, i.e., scenarios in which an attacker may succeed in harming or negating the security objectives (confidentiality, integrity, availability) or privacy objectives (unlinkability, transparency, intervenability).

Enumerative threat modeling approaches such as STRIDE for security threats and LINDDUN for privacy threats are algorithmic in the sense that they involve systematically reviewing the possibility of threats over specific parts of the system, typically parts of a Data Flow Diagram (DFD). Traditionally, this enumeration is performed per-element [23, 24] (considering different types of threats over data stores, processes, data flows and entities), but more sophisticated approaches act upon interactions (sender-flow-recipient), while more advanced tools such as SPARTA employ complicated model patterns to identify areas of interest. The quality, accuracy, and completeness of the input model (DFD) is essential to ensure a relevant and reproducible threat analysis.

We first introduce and discuss the SPARTA asset, and then we discuss and illustrate the application of exhaustive threat enumeration to the maritime case below. Finally, we state the key innovation goals for SPARTA to make it more suited to systems such as the maritime case, increasing its adaptive capabilities on the one hand, and its capabilities to deal with adaptive systems on the other hand.

4.1.1 Asset description

SPARTA is an eclipse-based framework in support of threat modeling and risk-based threat analysis. SPARTA automates the threat elicitation step as it generates viable threat scenarios at the basis of a Data Flow Diagram. SPARTA works with enriched system models which include additional information related to security/privacy countermeasures in the system, asset values, possible adversaries, and attacker profiles. Building such enriched models, SPARTA automatically prioritizes the generated threats in terms of the calculated risk [25].

SPARTA is built in the Eclipse EMF (Eclipse Modelling Framework) ecosystem, and includes a graphical designer of DFDs, which can be used to augment the DFD model with metadata such as asset value estimates. A screen shot is shown in Figure 11. Although the maritime transport model is used to showcase the tool here, it will be described in detail in the following sections.

Under the hood, the threat elicitation engine loads a threat catalog, that for each threat type matches model-based patterns (implemented using the advanced model transformation technology of the eclipse ecosystem, VIATRA¹¹). By leveraging the pattern recognition technology, each pattern match is then converted into a threat instance. Subsequently, for each threat instance, a risk is calculated at the basis of multi-component risk quantification schemes based on FAIR¹². The outcomes of these two steps are presented in Figure 12.

Customizability. Threat catalogs can be customized and loaded in specific projects. This allows tailoring the analysis and the identification of threats to the project at hand. Dedicated threat catalogs exist that cover the STRIDE threats (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service,

¹¹ <https://www.eclipse.org/viatra/>

¹² <https://www.fairinstitute.org>

Elevation of Privilege) and LINDDUN threats (Linkability, Identifiability, Non-Repudiation, Detectability, Disclosure of information, Unawareness, Non-compliance).

Different risk assessment schemes are implemented and supported, for example for privacy threats, a characterization of the nature of affected data subjects (e.g., age) is considered, while this makes less sense to calculate security risk.

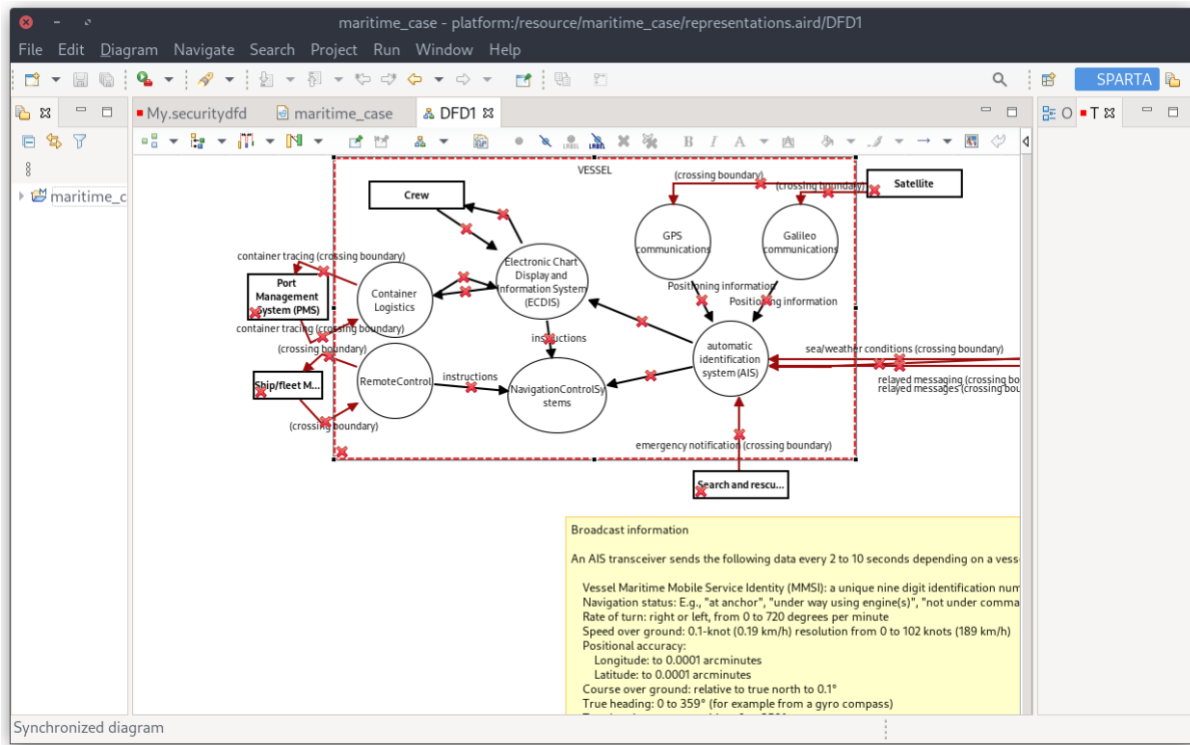


Figure 11 Screenshot of the SPARTA graphical editor.

SPARTA is under active development and is the basis for many related building blocks.

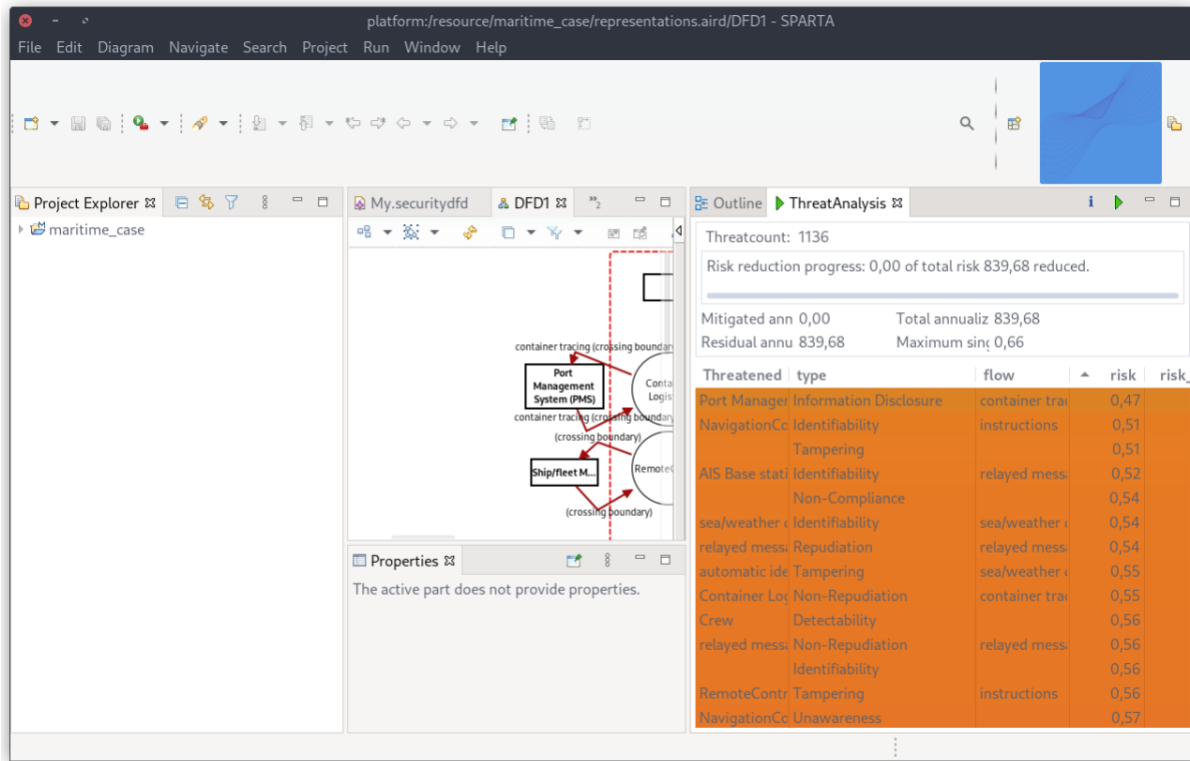


Figure 12 Outcome of threat generation and risk assessment.

4.1.2 Application Scenario

This section discusses the effort to model the maritime case in SPARTA and the ensuing threat analysis.

Step 1a: Model the DFD: a DFD model of the maritime case was created, based upon the characterization of the Maritime Transport example, and augmented with public information about the systems and concepts referred to (e.g., Search and Rescue Transceiver (SART), Automated Identification System (AIS)).

As most of the attacks presented in the descriptions of the Maritime Transport example (cf. Section 3.1) are aimed directly at individual vessels, the DFD is modeled from the perspective of an individual vessel, and the other systems are modeled as external entities. The outcome is presented in Figure 13.

Step 1b: Model enrichment: definition of asset values and instantiation of countermeasures in the model. This creates an overlay of (non-visible) elements that are considered during threat elicitation. In terms of countermeasures, none are selected in a first instance, and thus, a wide array of potential threat scenarios will automatically be considered.

Step 1c: selection of threat catalog: SPARTA allows for customized threat catalogs, and these can be loaded dynamically via the UI. For the Maritime Transport example, we loaded the generic threat catalog that encodes the STRIDE threat categories (Spoofing, Tampering, Repudiation, Information disclosure, Denial of Service and Elevation of Privilege). These threat categories focus on design-level issues/weaknesses and emphasize less on specific attack types (e.g., malware, ransomware).

Step 2: Threat elicitation: This is an automated generation step that yields several threats instantiated in the specific system. For example, at the basis of the above DFD, the following threats are generated by SPARTA:

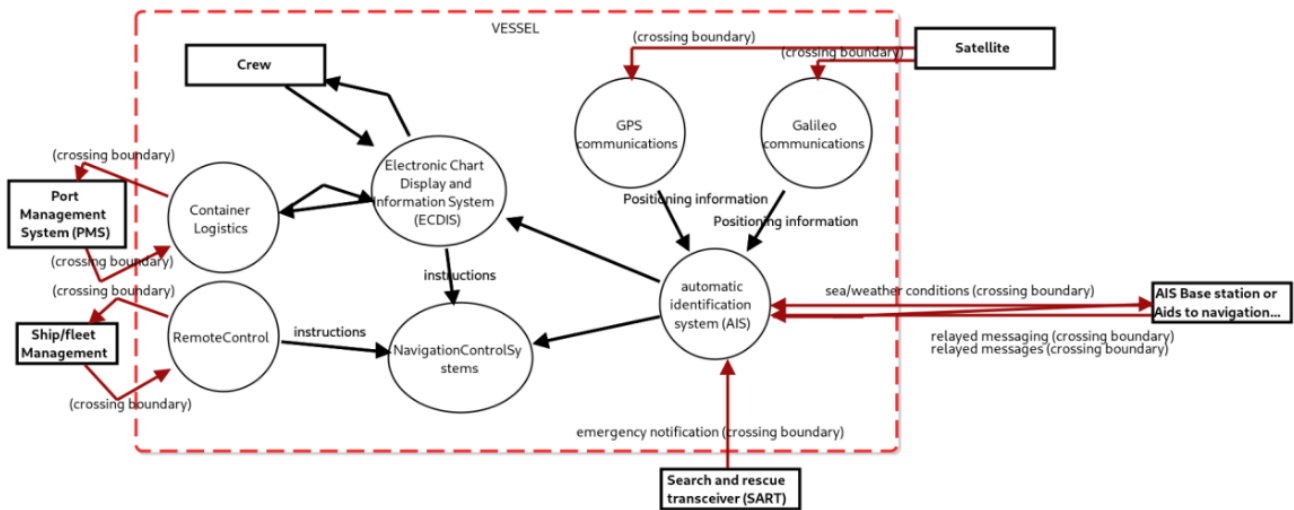


Figure 13 DFD of the maritime case.



Figure 14 Illustration of an Elevation of Privilege Threat.

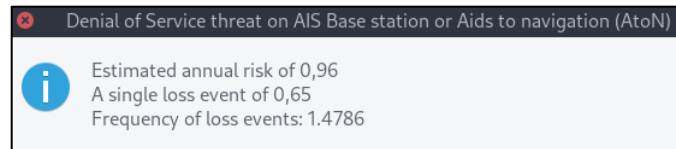


Figure 15 Illustration of a Denial of Service threat.

In total, the automated threat elicitation step yields 1136 distinct threat scenarios at the basis of the above DFD. This is a consequence of a combinatorial explosion, a phenomenon called **threat explosion**.

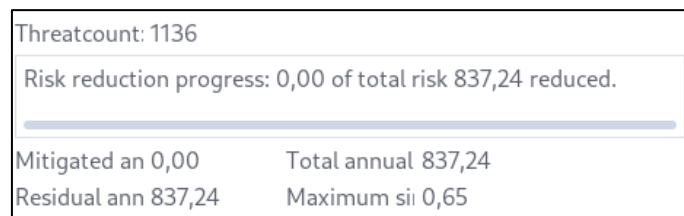


Figure 16 Threat count and inherent risk calculated.

Step 3: Risk-based threat prioritization: To keep the threat management effort practically workable, SPARTA automatically attributes a risk score to each threat, color-codes them in accordance with their priority, and hides those that are ranked beyond a specific threshold (see Figure Figure 14 and Figure 15). This allows the designer to focus on the most stringent issues first.

Step 4: Mitigate threats: in the fourth step, mitigations, countermeasures, and controls can be instantiated to reduce the overall risk. As depicted in Figure 16, in the current model, no such mitigations have been selected yet (risk reduction progress is still zero).

Challenges and roadmap towards adaptive threat management

The DFD depicted in Figure 17 presented a superset view of all systems and communication means of the vessel. However, not all these subsystems will be always equally relevant. The version shown in Figure 17 illustrates this, making explicit and essential distinction between the vessel, at shore, in open sea and when it is or comes in the vicinity of other vessels.

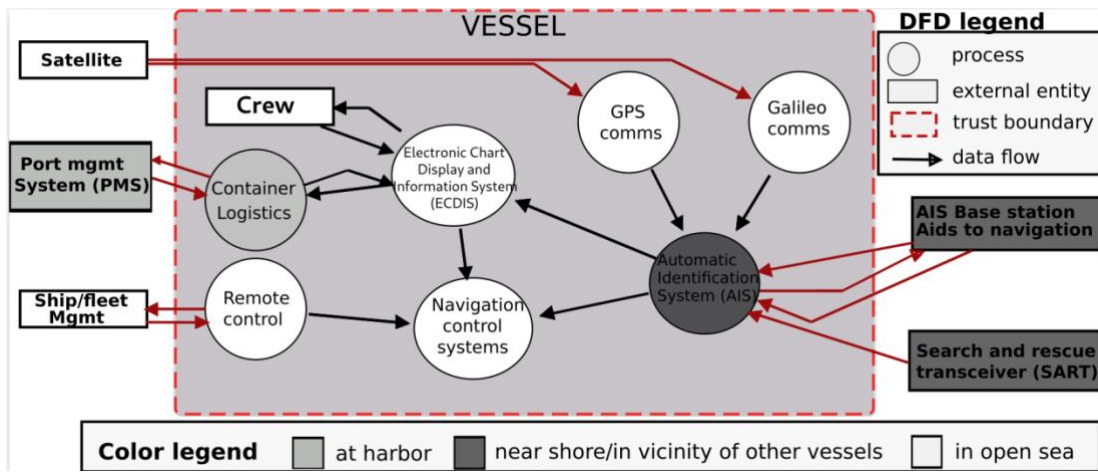


Figure 17 Illustration of the DFD as a superset of the vessel in different situations.

When performing threat analysis in an a-priori design context, the analyst is required to enumerate and anticipate the different operational contexts of the system, and as these contexts may change significantly over time, so will the corresponding threat landscapes.

This is especially problematic in systems for which exhaustive enumeration of these operational context is cumbersome or even infeasible, and therefore, relevant threats may not be identified, prioritized nor mitigated correctly.

In addition, when the system itself changes or evolves over time, novel threats will not be identified nor mitigated unless the entire threat analysis exercise is repeated. Relevant change scenarios we envision in the maritime system case are:

- new subsystems are activated that were not initially present in the analysis (e.g., the installation of on-deck wireless networks to enable personal communication of crew members),
- when new external entities emerge (e.g., new means to communicate to base stations, or new types of base stations),
- as the capabilities of attackers evolve over time (e.g., new vulnerabilities have been published or new types of attacks have been detected) so that threat types that were previously considered impossible have become actual issues.

4.1.3 Output

- A binary release of SPARTA¹³ (UI/editor and threat elicitation engine) is scheduled for release. The release will be accompanied by a repository of different application cases, threat catalogs, risk models and a tutorial for first-time users. The modeled maritime system will be included in the release. Generated threat documentation will be included in the case description, together with instructions on how to perform the analysis again.
- A demonstration video¹⁴
- A publication about the application of SPARTA to the maritime case: [26].

4.1.4 Asset limitations and future work

Two key problem areas have been identified [26]:

A- Modeling adaptive systems in DFDs. As discussed in Section 4.1.1, a DFD convolutes all data flows possible in the system in one model, yet this act incurs a loss of information in terms of the control flow and causality of the depicted data flows. This is essentially hindering in adaptive or context-aware systems that vary their behavior drastically by design, depending on the specific context of the system. The DFD of a vessel from the maritime case study is an example of such a system.

To accomplish this goal, the following enablers will be required:

- to explicitly model and characterize the different situations or contexts or states in which the system can operate, the DFD elements of relevance in these contexts (e.g., which perform the most communication or make key decisions),
- determine mechanisms to express transitions between these states (how the shift occurs from one state to the other), gradual or discrete and how this affects the different risk component inputs and values,
- to allow the threat elicitation to reason across states. A vulnerability or successful attack in one state may lead to bigger problems when a different state is attained. For example, in the maritime case, an attacker that takes over the container management systems while at sea might not cause the full extent of the harm while still at sea, but that harm will only materialize itself when the vessel is near a harbor and more intensively relies on its communications to the shore.

B - Automated DFD derivation and actualization. Maintaining a centralized, comprehensive, and architecture-level abstraction model in terms of data flows (DFDs) of an operational system, through methods of reflection, monitoring, and inspection. The purpose of this model is to keep it as the basis of continuous threat modelling and threat-based risk assessment (SPARTA) and to be able to identify reactively when changes occur that may cause new threats to emerge on the agenda or known threats to become more prominent.

To this end, we envision in future work an adapter-based design that is extensible and capable of integrating different sources of architectural information that can be reified in data flows:

- a) Structure and specification. For example, data model or schema specifications used in data management or data access frameworks, workflow descriptions or business process specifications used in business-process technology. These can be complemented with techniques of static code analysis (e.g., dependency analysis) or reification of abstraction models at the basis of technology-specific annotations (e.g., Spring).
- b) Deployment and configuration. Distributed systems increasingly rely on middleware platforms that require developers to define deployment abstractions (e.g., applications, service interfaces or

¹³ SPARTA release website <https://distrinet.cs.kuleuven.be/software/sparta/>

¹⁴ <https://youtu.be/cdAiaNutfW4>

endpoints, (micro-)services), and this is the level of granularity at which distributed applications are deployed and managed (e.g., unit of scaling, unit of state persistence). This information is available through reflection and via platform-specific deployment descriptors, e.g., in Kubernetes. In addition, virtualization and orchestration systems employ techniques such as Software-defined networking (SDN) which in turn rely on expressive deployment descriptors from which valuable deployment information can be obtained.

- c) Run-time interactions, for example between active objects, threads, and processes. These views contribute information about frequency of data flows/interactions, provide concrete instances of data elements being processed, run-time user sessions, etc. The construction of these views can leverage upon information provided in audit trails, system logs and session management, but also may rely on dynamic analysis of application execution (e.g., based on call graphs to identify run-time interactions, or taint analysis). In addition, intrusion or fraud detection systems are capable of detecting security incidents in an operational system, and these represent run-time occurrences or instances of threat scenarios that may trigger system-wide threat re-evaluation.

5 Analysis and Planning

5.1 Adaptive Authentication

The adaptive authentication asset aims to address the lack or insufficient authentication issue highlighted in the maritime scenario. An adaptive authentication system monitors contextual factors and behavioral features of its users to identify changing security risks. The system can decide to enforce an authentication method to mitigate the security risks and maximize user convenience [27-29].

For example, Hayashi et al.[30] associate a risk level with the location from where a user requests access (home, work, other). They change the authentication method adopted depending on the user's current location. If the user tries to access a service/resource from a previously unknown location, s/he is required to provide additional credentials (e.g., pin, password). Security risks can also be brought by changes in user habits. For example, Gebrie and Abie [31] consider the change in users' daily routines (e.g. walking, eating, sleeping) monitored using wearable devices, to calculate the risk score of an access request. They link the risk score to an abnormal activity and adapt the authentication method accordingly. Similarly, Bakar and Haron [32] analyze the historical records of the users' behavior profile (e.g., login time, location, browser type) and associate a trust score to behavior changes. If the trust score is higher than a given threshold, the user is asked to provide additional credentials to access the required service/resource.

Continuous authentication [33], instead, refers to the activities performed after a user has authenticated successfully, to ensure that the session continues to be held by the legitimate user. It also aims to ensure that the user experience is maximized, for example, by reducing the frequency with which a user is required to re-authenticate. A continuous authentication system usually monitors the user behavior (e.g., applications usage, pressure on touch screens) to identify security risks arising after a user authenticates successfully. For example, Karanikiotis et al. [34] monitor the users' gestures (e.g., swipes) on a mobile device. If the user exhibits abnormal gestures, s/he is classified as an illegitimate user and the mobile device is locked automatically. However, this approach is not suitable when a legitimate user is simply performing a new behavior. In such a situation, continuous authentication should be combined with adaptive authentication. For example, Jorquera et al. [35] uses machine learning to identify whether the owner of a mobile device is legitimate depending on his/her application usage statistics. The system considers the usage statistics falling in the possibly normal category to learn new behaviors, and triggers re-authentication if these statistics fall in one of the anomalous categories.

Previous work on adaptive authentication [27, 28] provides limited guidance on how adaptive authentication systems can be built systematically. Thus, a few open issues remain: i) which requirements are relevant to an adaptive authentication system, ii) how contextual factors can affect the feasibility of authentication methods, and iii) how different authentication methods can affect satisfaction of the requirements. Although previous work on adaptive systems has considered context-driven adaptation (e.g., [36-38]), it has not considered how context can affect the priority of the requirements and the feasibility of authentication methods. These issues fall into the knowledge, analysis and planning in the MAPE-K loop.

5.1.1 Asset description

This section shows how we can use the Adaptive Authentication System asset to support the activities of the MAPE-K loop [19]. As shown in Figure 18, we use the main pillars of adaptive authentication (Requirements, Authentication Methods, and Contextual Factors) to represent and maintain at runtime the Knowledge that is used to configure the activities of the MAPE-K loop. The contextual factors can bring security risks and affect priority of the requirements. They can also make certain authentication methods infeasible. Authentication methods, instead, can mitigate security risks and contribute to the satisfaction of the requirements. During monitoring and analysis, the adaptive authentication system should, respectively, monitor contextual factors and analyze the security risks. During planning, it should identify a feasible authentication method that a) minimizes security risks and b) maximizes the satisfaction of the requirements

considering their trade-offs. During execution, the adaptive authentication system should enforce the selected authentication method.

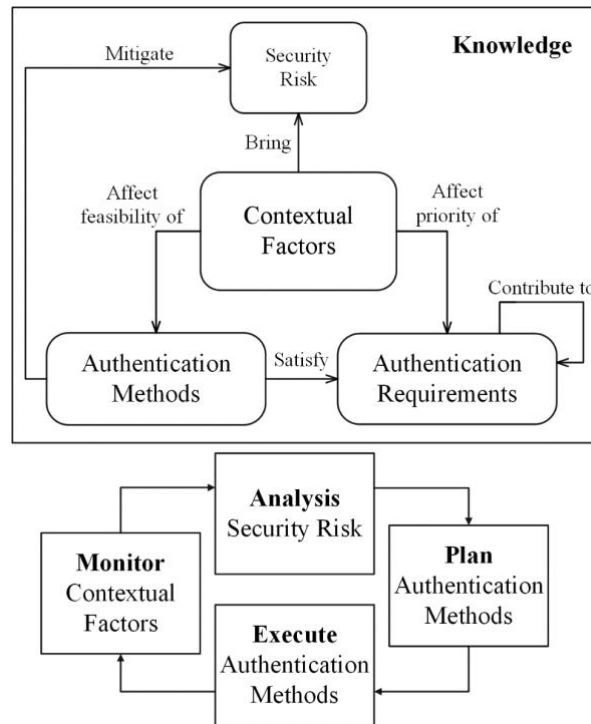


Figure 18 Adaptive Authentication Asset.

This asset provides a holistic framework informed by previous research to characterize the adaptive authentication problem and support the development of an adaptive authentication system. We show how our framework can inform the activities of the MAPE-K loop necessary to build an adaptive authentication system to support maritime transportation case study scenarios. To achieve this aim, we use a contextual goal model [39] to represent the requirements and the impact that contextual factors have on the requirements priorities. Also, we use an extended feature model [1] to represent different features that can be selected to identify an authentication method during planning. We have used a Fuzzy Causal Network (FCN) [40] to analyze the impact of contextual factors on the security risk and identify a suitable authentication method. To support decision-making, we used causal reasoning on FCN to analyze the impact of contextual factors on authentication goals, partial risk. Also, it selects the most suitable authentication method to mitigate the total risk, and satisfy the security goals, and other authentication goals (e.g., usability and performance). Finally, we implemented the fuzzy causal network using Z3 [41], a theorem prover, to choose an effective and optimal authentication method that satisfies the adaptive authentication requirements.

5.1.2 Application scenario

To motivate the adaptive authentication problem, we discuss a set of scenarios in the maritime transportation case study. We use these scenarios to show cases about how our framework can inform the activities of the MAPE-K loop, upon which adaptive authentication systems are built.

Scenario-1 A passenger waiting at the port is using the mobile phone network to access information about his/her ticket. We assume that access to this information is granted using password-based authentication. Afterward, the passenger boards the vessel, and, because of a change of his/her location, s/he switches network, using the vessel Wi-Fi. The use of unsafe network topology can facilitate AmosConnect attacks

[14]. The vulnerable client could allow unauthenticated attackers to perform blind SQL injection and recover usernames and passwords. Then, with the use of the retrieved credentials, an adversary can remotely execute arbitrary commands with system privileges on the remote system by abusing the Task Manager of the mail client. To mitigate the security risk posed by this potential attack, the passenger should be forced to use a stronger authentication method, such as two-factor authentication, when s/he boards the vessel.

Scenario-2 The Liquid Cargo Ship needs to acquire weather and routing information to reach its original course as soon as possible. To achieve this aim, it communicates with the nearest shore or vessel. The nearby vessels can deviate Liquid Cargo Ship from its original course, by broadcasting fraud civil GPS signals, weather, and routing information. To remain covert, the spoofed signals were slightly altered illegitimately. In this scenario, the requirements related to the confidentiality of the weather and routing information and authenticity of the parties sharing information (Liquid Cargo Ship and nearest shore/vessel) have higher priority w.r.t. other requirements that can be relevant to the authentication problem, such as usability and performance. Thus, the Liquid Cargo Ship and the nearest shore/vessel should use, for example, a certificate-based authentication or signcryption-based authentication. In this case, both vessel and shore/vessel should prove their identity to each other before starting to share information which decreases the risk of a proof-of-concept attack [15] and terrorist attack.

Scenario-3 To avoid the collision between the vessels, proximity to other vessels and/or the port is important to be managed. To this aim, the vessels need to exchange with the nearby vessels or port information about their respective distance. The exchange of distance information should happen quickly to allow the vessels to avoid a collision in a timely manner. Thus, performance requirements (e.g., minimize the time necessary to perform authentication) should have a higher priority compared to other requirements, such as security and usability. Using a certificate-based authentication would not be appropriate in this situation since it can require excessive time to verify the identity of the vessels on a remote server. Alternatively, the vessels can use two types of information (e.g., the vessel plate and crew's license) to authenticate with one another. These credentials can be transmitted and verified in a shorter time compared to certificate-based authentication. This authentication method can also avoid typical attacks such as a proof-of-concept attack and terrorist attack because any vessel before communicates with another vessel uses the vessel plate and the crew's license to be sure it communicated with an authorized vehicle.

Scenario-4 This scenario is about the management of access to different parts of the vessel by legitimate users. The vessels crew should be allowed access to critical parts of the vessels and information about the passengers. Besides security requirements, usability requirements are also important in this scenario since the authentication method should minimally distract the crew from their tasks. For example, a biometrics-based authentication (e.g., face or iris recognition) can be ideal in this scenario because it can be automated and requires a little attention from the crew. However, other contextual factors, such as low light may render this authentication method ineffective and require the use of a different one (e.g., token, fingerprint). Moreover, the privacy preferences of the crew can affect the authentication method selection.

As shown in these scenarios, a multitude of contextual factors (e.g., location, network topology, sensitivity of accessed information, proximity with other vessels) can affect the security risk and the priority of the requirements that can be relevant during adaptive authentication (e.g., security, usability, and performance). These requirements can also be conflicting with one another. For example, adopting a strong authentication technique can harm performance (e.g., a certificate-based authentication) and usability requirements (e.g., a password that is very hard to remember). Certain contextual factors (e.g., lighting) can render some authentication methods (e.g., face recognition) ineffective. Moreover, estimating the impact that an authentication method has on the requirements cannot be quantified precisely. Finally, because users can actively engage in authentication, their preferences and privacy also should be considered when an authentication method is selected.

5.1.3 Adaptive Authentication Framework

We reviewed previous work on adaptive authentication and leveraged our authentication scenarios to elicit the main aspects to be considered when building an adaptive authentication system: requirements, authentication methods, contextual factors, and decision-making techniques.

5.1.3.1 Requirements

The requirements of an adaptive authentication system are mainly related to security, privacy, usability, and performance. Most of the adaptive authentication systems (e.g., [29, 31, 32, 35, 42-47]) that we examined adapt the authentication method because of a changing security risk. For example, De Silva et al. [29] link specific changes in the user profile (e.g., location, browser type, mouse behavior, keystroke patterns) to changes in the security risk. When a high-security risk is detected, a stronger authentication method (e.g., two-factor authentication) is enforced. Daud et al. [44] link the user's login attempts to the security risk based on contextual factors, such as the IP address, location, type of browser, and the operating system. In case of an increased risk, this approach applies penalties, for example, it can adopt 2- or 3-factor authentication, it can block authentication for a given period or blacklist a user. Although it has not been considered in previous work on adaptive authentication, an important requirement is authenticity. This requirement is relevant in the scenario shown in **Scenario-2** (see Section 5.1.2) where the selection of a certificate-based authentication is dictated by the need to ensure authenticity of the communicating parties.

Some approaches surveyed, especially those based on user behavior and using physiological credentials, aim to satisfy *privacy* requirements, particularly anonymity and untraceability [48-52]. For example, Xi et al. [49] propose an adaptive anonymous authentication protocol in a V2R topology based on a cryptographic technique called verifiable common secret encoding. This technique uses the cryptographic keys of the communicating users to hide their individual identities. The authentication protocol can also adapt at runtime depending on the level of anonymity required by the users.

Because authentication can be performed by humans, it is also crucial to consider *usability* requirements. These mainly aim to maximize the quality of the user experience during authentication. Usability has been mainly considered in terms of ease of use, for users having different behaviors [35], abilities [53], and ages [54]. Other work [55] has considered usability in terms of transparency, i.e., the system should provide users with explanations justifying why it changed the required authentication method. Usability is also commonly expressed in terms of efficiency and effectiveness of the authentication methods [56]. More precisely, *efficiency* is related to the speed of the authentication method. For example, Jorquera et al. [35] minimize the number of authentication credentials to improve efficiency. Effectiveness is related to the error rate that an authentication method can be prone to. This can be related to the memorability of the credentials (e.g., using a password that is difficult to remember can be ineffective) and also to environmental factors (e.g., noise type and level, lighting level, or temperature) [57]. Other work [47], instead, aims to maximize satisfaction of the *user's preferences*, by allowing a user to select an authentication method for specific applications. This can be relevant when users prefer stronger authentication techniques in specific contexts: work, personal account, and financial [58].

Although performance requirements have been briefly mentioned in previous work [47, 59-61], their distinction with usability requirements has not been defined clearly. From our analysis, *performance* requirements can be about minimize the authentication time. Finally, only a few approaches [35, 47] address the trade-off between the aforementioned requirements, mainly focusing on security and usability requirements.

5.1.3.2 Authentication Methods

The authentication methods that have been used in previous work have optional and mandatory authentication features. It is mandatory to choose a credential type [51], such as something you know (e.g., password, OTP), something you have (e.g., smartcard, token), something you are (e.g., face, iris, fingerprint), or two-factor authentication (e.g., select two credentials). The credential type affects the level

of automation. For example, iris and face recognition have the highest level of automation, since they require the minimum input from the user. Fingerprint-based authentication has a medium level of automation since it requires the user to actively scan his/her finger. Password-based authentication has a low level of automation since it requires the user to remember and input a password. Some authentication features, such as credentials renewal [55, 60] and cryptography type [49] [51], are optional. Others require specific devices to be performed [47, 58] (e.g., smartcard-based authentication requires a reader). Representing the features of an authentication method can help express its impact on the satisfaction of the requirements.

5.1.3.3 Contextual Factors

We group contextual factors depending on whether they affect 1) the security risk and the adaptive authentication requirements or 2) the feasibility of authentication methods.

- **Security risks and requirements**

- **Assets Sensitivity** refers to the criticality of data or applications to which access is requested. Asset sensitivity can increase the priority of security requirements and affect security risks. Thus, some approaches (e.g., [47, 62]) adapt the authentication method depending on the sensitivity of the data to be accessed.
- **Location** refers to the place where a user is authenticating and can have an impact on the security risks. Several approaches have proposed to ask the user for additional credentials, if s/he attempts to access services/resources from an unusual location [32, 43-45].
- **Network Topology** can affect the security risk. Previous work [50] suggests changing authentication method depending on the attacks that can exploit the topology of the network a node is currently connected to.
- **Time** refers to the moment when authentication is performed and can also affect security risks [32, 43, 44]. For example, if a user tries to access an asset at odd times (e.g., outside working hours) s/he can be asked to provide additional credentials during authentication [32, 43] or can be subjected to penalties (e.g., being blocked for some hours or permanently) [44].
- **User Role** (e.g., manager VS regular employee [63]) can affect the security risk. Arfaoui et al. [46] require the nodes of an Internet of Things (IoT) network to adopt an authentication method depending on their role (e.g., IoT gateway, context manager, data consumer) and depending on additional contextual information (e.g., location, time, emergency, normal situation). In the scenario shown in **Scenario-4** (see Section 5.1.2), the role of an actor (e.g., vessels crew) can also increase the priority of the authenticity requirement.
- **Movement of the Nodes** refers to the movement of the nodes within a network. For example, in an IoV network nodes can change their position, requiring authentication to be performed rapidly. As shown in the scenario in **Scenario-3** (see Section 5.1.2), the presence of moving authenticating ship increases the priority of performance requirements. Fayad et al. [59] proposed an adaptive authentication approach where nodes of an IoT network can store their authentication information on the blockchain. This allows authentication to be performed even when the authenticating nodes do not belong to the same network.
- **User Preferences** refer to users favoring specific authentication methods to others [32, 47, 57, 64]. Considering user preferences during adaptive authentication can increase satisfaction of usability requirements.

- **Feasibility of authentication methods**

- **Authentication Devices** refer to the devices (e.g., phone, camera, reader) available to perform authentication. For example, some authentication methods (e.g., RFID) require additional devices (e.g., reader) [65]. In other situations, limited-resources devices may not be able to support

authentication methods that are computationally intensive (e.g., cryptography-based authentication) [35].

- **Proximity** refers to the user's distance from a device and can indicate possession of the device [66]. For example, two-factor authentication can be enabled by sending a PIN to the device a user is close to.
- **Device Position** refers to the relative position of a device w.r.t. its owner (e.g., held in hand or in the pocket). For example, face recognition is not feasible if the device is held in the pocket. Frequent changes of the device position can make gait-based authentication infeasible [67].
- **Network Quality** can affect feasibility of authentication methods (e.g., cryptography-based authentication) that can have overheads in the communication network. For example, the use of a network with limited bandwidth can cause delays and even lead to fatal accidents[50].
- **Environmental Conditions** refer to conditions, such as lighting and noise level. For example, Wojtowicz and Joachimiak [57] propose a system that avoids selecting authentication methods that may not be effective in certain environmental conditions. For example, face recognition and voice recognition are avoided when the lighting level is low and the noise level is high, respectively.

Although we have identified relationships between contextual factors and requirements, existing adaptive authentication approaches have only focused on specific contextual factors relevant to the considered application domain.

5.1.3.4 Building an adaptive authentication system

In this section we show how we can use our framework to build an adaptive authentication system for the maritime transportation case example described in Section 2.1.

To represent the requirements and the impact that contextual factors have on the requirements priorities we use a contextual goal model [68]. Goal models [39] allow decomposing goals into more concrete realizable sub-goals, leading to the identification of requirements (leaf goals). Figure 19 depicts the goal model associated with the scenarios described in Section 5.1.2. Confidentiality, authenticity, and integrity decompose security requirements. Effectiveness and efficiency decompose usability requirements. Minimize authentication time and minimize authentication delay decompose performance requirements. The AND-refinement means that all sub-goals (if necessary) require to be satisfied to satisfy the upper-level goal. We associate specific contextual factors to an increase (+) or a high increase in the priority of the requirements (++) . Because such impacts cannot be quantified precisely, we have used qualitative labels to describe their intensity, by relying on our experience and on previous work. For example, when the vessels crew access critical parts of the vessels the priority of the usability requirements highly increases because the vessels crew needs to perform authentication without distraction, while s/he needs to focus on their task [66]. The movement of the nodes highly increases the priority of the performance requirements because authentication needs to happen in a short time [66]. Asset sensitivity (e.g., nearby vessels or port information) increases the priority of the confidentiality requirements and highly increases the priority of authenticity, since such information should be accessed only by trusted users (e.g., the vessels). When the vessel exchanges distance information with the nearby vessel, the priority of integrity highly increases because tampering such information may cause a crash between vessels.

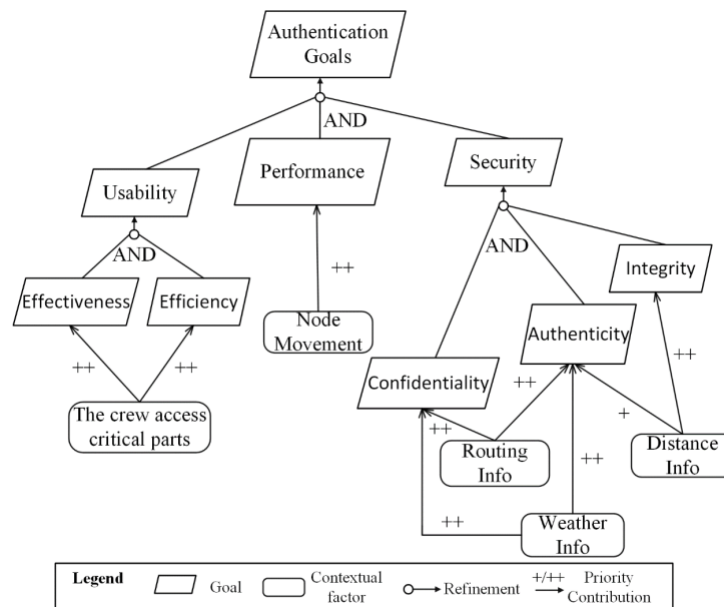


Figure 19 Requirements and Contextual Factors.

We use an extended feature model [69] to represent different features that can be selected to identify an authentication method during planning. Extended feature models have been widely employed in software engineering and in industry to identify and capture variability in systems and products. Figure 20 depicts the extended feature model representing possible authentication methods. The upper part of the figure represents the optional and mandatory authentication features discussed in Section 5.1.3.2. Constraints are provided for some features. For example, credentials should be renewed weekly or monthly and the device type should be chosen based on the credential type. In the lower part of the figure, we link contextual factors to the features that these factors make infeasible (-f). For example, low lighting makes face and iris recognition ineffective [57]. Also, we associate each feature with its contribution to the satisfaction of the different requirements. We use qualitative labels to represent the impact that each feature has on the satisfaction of the different requirements, i.e., positive (+), highly positive (++), negative (-), and highly negative (--). To keep Figure 20 clear, we do not provide satisfaction impacts for all the requirements. For example, the vessel plate and crew's license contribute to minimizing authentication time and delay very positively. However, they are not as effective as using certificate-based authentication to satisfy confidentiality and integrity. Certificate-based authentication, instead, has a negative impact on performance requirements (minimize authentication time and delay). The face and iris features contribute to effectiveness and efficiency highly positively and are also effective in satisfying confidentiality and integrity requirements. Finally, we represent in the diagram shown in Figure 21 the impact that contextual factors have on the security risk. Each partial security risk depends on the likelihood of success of an attack and the harm that such attack can cause. We use qualitative labels to represent the impact of contextual factors on the likelihood of attacks and we use the value of the asset targeted by an attack to estimate the harm.

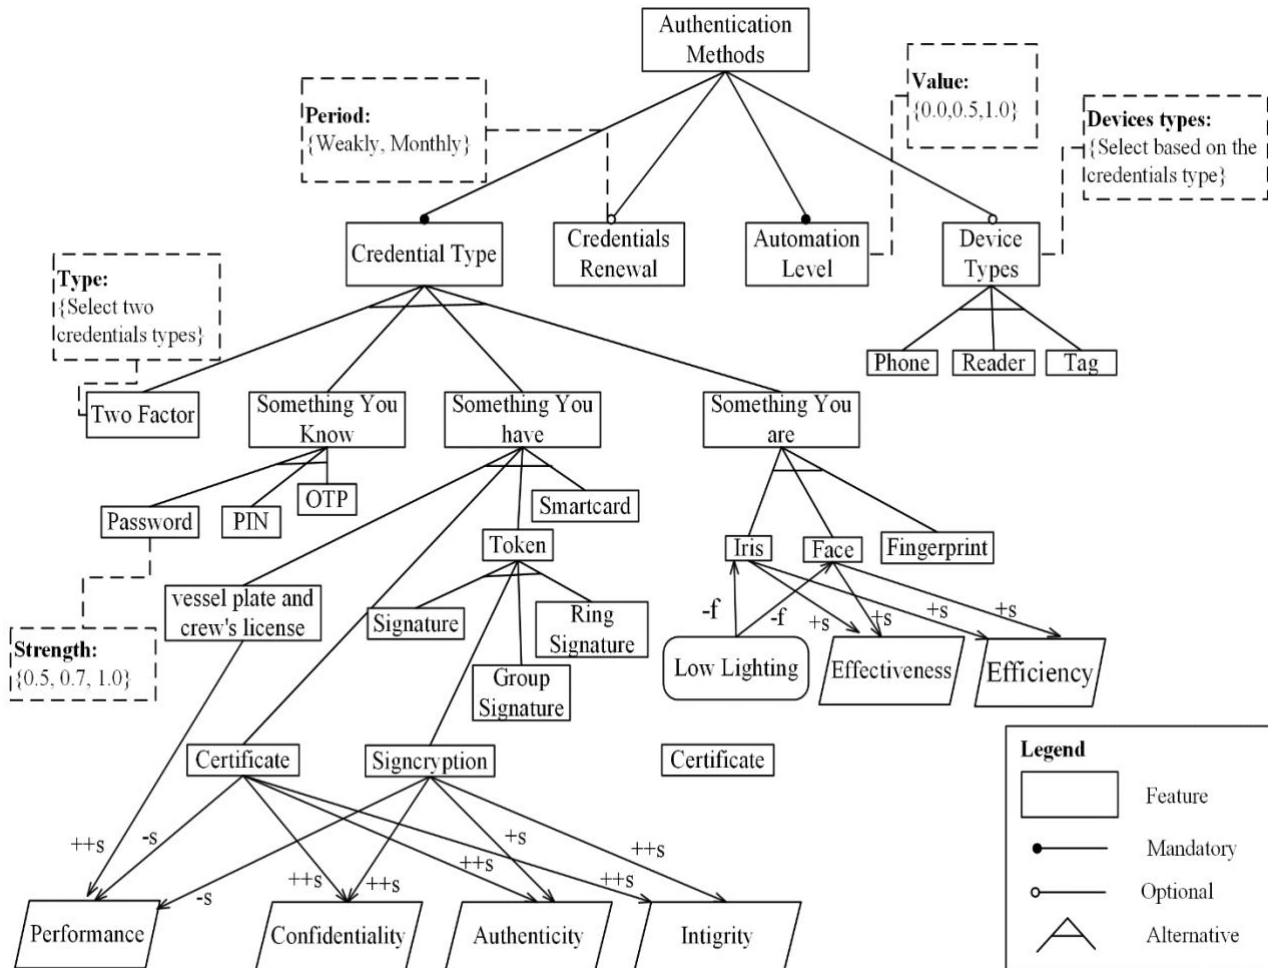


Figure 20 Extended Feature model of authentication methods.

To analyze the impact of contextual factors on the security risk and identify a suitable authentication method we have used Fuzzy Causal Networks (FCN) [40]. FCN are suitable to analyze consequences of context changes on security risks and perform impact analysis of the selection of an authentication method. They have already been used in previous work to support adaptive security [70]. We build a fuzzy causal network from the three models discussed before. Because they are based on fuzzy values, FCN allow to represent uncertain judgement of domain experts, usually expressed in qualitative terms (e.g., 'high', 'medium' and 'low'), for example, regarding the impact of contextual factors on the priority of the requirements or the effectiveness of authentication methods on the satisfaction of the requirements.

Our FCN is built upon the elements and links represented in the diagrams in Figure 19, Figure 20, and Figure 21. Like an influence diagram, our causal network has three types of nodes: chance nodes representing uncertain domain entities significant for causal reasoning (denoted by ovals), decision nodes indicating decisions to be made (denoted by rectangles), and utility nodes corresponding to the fitness value of the network configuration (denoted by a hexagon). Except for the utility node, all the others are represented by fuzzy variables in the range $[0,1]$. Table 2 lists the meaning of the nodes of our FCN. For example, contextual factors can indicate presence of certain factors (e.g., low lighting level) or can be associated with a value indicating the sensitivity of an asset. Figure 22 illustrates the abstract structure of our FCN. In the presented FCN graph, each causal link is labeled with a weight that refers to the causal relationship's strength between two nodes. The actual link could be positive or negative. Where $A^+ \rightarrow B$ refers

to an increase in **A** will cause an increase in **B** and $A^- \rightarrow B$ refers to an increase in **A** will cause decrease in **B** value. The assigned weights can be fuzzy labels, we used quantitative labels in the $[0,1]$ range.

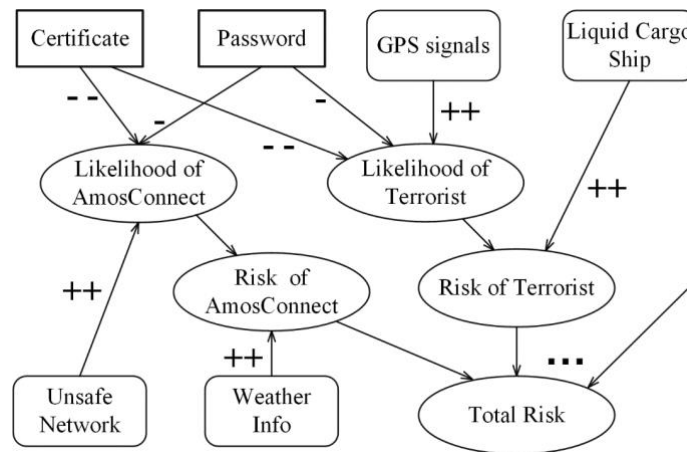


Figure 21 Security Risks.

Node	Meaning	Type
Contextual Factors	Value, Presence	Chance
Attack	Likelihood of success	Chance
Partial Risk	Partial risk of Attack	Chance
Authentication Goals	Satisfaction Level	Chance
Authentication Methods	Strength, Enable/Disable	Decision
Utility	Value	Utility

Table 2 Nodes of the Fuzzy Causal Network.

Chance nodes - Except for the authentication methods, all entities taken from the models in Figures 19-21 are chance nodes in the FCN. The contextual factors can have a negative impact on the authentication methods. For example, low lighting makes some authentication methods (e.g., face and iris) ineffective. Contextual factors can also increase or decrease the priority of authentication goals and security risks. For example, the sensitivity of the traffic information can increase the priority of security requirements and affect security risks by increasing the likelihood of successful attacks (see the link between the contextual factors and the attack node in Figure 21).

The attack nodes have a positive impact on the partial security risk (see the link between attacks and the partial risk). Also, it is affected negatively by the strength of the chosen authentication method strength. For example, the use of unsafe network topology (see Section 5.1.2 **Scenario-2**) highly increases the likelihood of an AmosConnect attack. At the same time this attack could be mitigated by choosing a strong authentication method.

The value of the authentication goal represents its satisfaction level. This is calculated by aggregating the satisfaction levels of the sub-goals. The satisfaction level of the leaf goals is calculated by aggregating the contribution of the selected authentication methods, and, in turn, their satisfaction level propagates upwards

in the goal model following the semantics of AND/OR fuzzy operators (i.e., min/max). The contextual factors can also increase/decrease the priority of the authentication goals (+/-d in Figure 22).

Partial risk and total risk are chance nodes taken from the risk model in Figure 21. A partial risk node is used to estimate the equal risk to each attack, while the total risk aggregates all the partial risks. This estimation is calculated by multiplying the value of the contextual factor representing the targeted asset, and the likelihood of harm (see positive links from contextual factors and attacks to partial risk in Figure 22). For example, the partial risk of a proof-of-concept attack relies on the likelihood of the proof-of-concept attack as well as the value of the ambulance (see Section 5.1.2 **Scenario-2**).

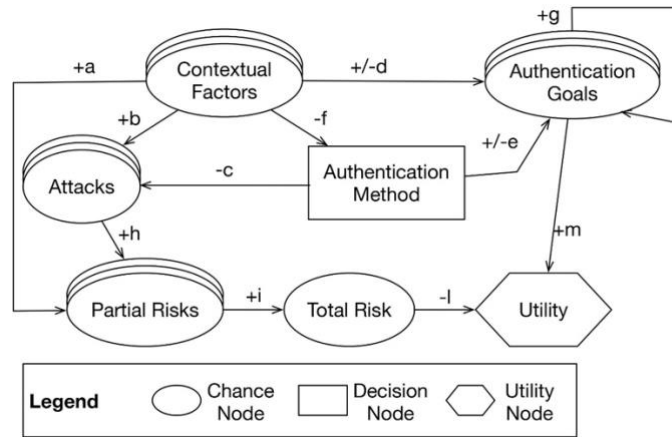


Figure 22 Abstract model of the Fuzzy Causal Network.

Decision nodes - Any authentication method is a decision node in the causal network. Attacks can be mitigated by choosing a suitable authentication method (see the link between authentication method and attacks). For example, certificate-based authentication mitigates impersonation attacks. Moreover, the authentication method may deteriorate or improve the satisfaction of authentication goals. For example, certificate-based authentication has a negative impact on performance. The authentication method value is associated with its strength level in mitigating attacks. Some authentication methods, such as face recognition, are crisp since they can be enabled or disabled, while others, such as password-based authentication have a different level of strength.

The utility node - This node expresses the effectiveness of the chosen authentication method, according to the total risk, the satisfaction of security goals, and other authentication goals (e.g., usability and performance). The utility node should aggregate all these costs and benefits. Benefits depend on how much authentication methods can mitigate the risk, and costs indicate how much they hurt authentication goals. The risk node has a negative impact on utility, while authentication goals have a positive impact. For example, the utility of applying certificate-based authentication benefits authenticity, while a big certificate size impacts performance very negatively. By updating the value of this node, a different utility value can be specified.

To support decision-making, we used causal reasoning on FCN to analyze the impact of contextual factors on authentication goals, partial risk, and select the most suitable authentication method to mitigate the total risk. Table 3 depicts the aggregation functions we employed in each node. The notation $A \rightarrow B$ refers to the set of nodes of type A that are causally impacting B . Minimum, Maximum, Average, and Sum are used as aggregation functions. We used the Sum function to represent the utility node since it is not a chance node (e.g attack, authentication goals). We select one of the other functions (Minimum, Maximum, and Average) to adjust the authentication method hardened level.

Causal Link	Aggregation
$\{CF\} \rightarrow At$	Minimum
$\{CF\} \rightarrow AG$	Maximum
$\{AM\} \rightarrow AG$	Average
$\{CF\}, \{AM\} \rightarrow AG$	Minimum
$\{AM\} \rightarrow At$	Average
$\{CF\} \rightarrow PR$	Maximum
$At \rightarrow PR$	No aggregation (only one attack)
$\{CF\}, At \rightarrow PR$	Minimum
$PR \rightarrow TR$	Maximum or Average
$\{TR\}, \{AG\} \rightarrow U$	Sum

Table 3 Functions used to aggregate the values of the nodes of the FCN.

In some cases, we can use more sets on the left side of the causal link; e.g., $\{CF\}T, \{AM\} \rightarrow AG$. In these cases, the reasoning mechanism initially aggregates the similar type of links and then combines different types. For instance, to evaluate AG , first, the effects of $\{CF\}$ on AG are aggregated by the Maximum function since the priority of the associated goal (e.g., security) is higher than the other authentication goals (e.g., usability and performance). Then the impacts from $\{AM\}$ are combined with the Maximum function because links between authentication method connected to an AG is OR (as shown in Figure 19). Finally, these two impacts are aggregated by the Minimum function, which is more conservative for AG .

Minimum and Maximum could refer to multiplication and addition. To evaluate the partial risk PR , the targeted asset's value should be multiplied by the likelihood of harm, which is translated into the Minimum function. We use the Maximum function for aggregating partial risk to total risk because it refers to adding up partial risks.

Using FCN, we were able to select appropriate authentication methods for the 4 scenarios presented in Section 5.1.2. However, it was cumbersome to provide the full models of authentication goals and assess the impacts between contextual factors, authentication methods and security risks. Also, the time to compute utility of different authentication methods can become intractable when the number of contextual factors and the potential authentication methods increases.

5.1.4 Outputs

The paper titled "Engineering Adaptive Authentication" co-authored with Alzubair Hassan and Bashar Nuseibeh was accepted on 02/07/2021 as a vision paper at the 2nd IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS 2021) [71].

5.1.5 Asset limitations and future work

Considering supporting the MAPE-K loop, still our asset has some limitations. For example, the asset could not monitor the change in all the contextual factors. Due to the multitude of contextual factors, it is not possible to monitor their changes continuously. Thus, there is a need for monitoring approaches able to collect data to detect specific situations (e.g., a vehicle trying to overtake another one, crossing a junction) and identify suitable monitor activities to be performed in those situations. In the execution phase, from our initial study we noticed that there is uncertainty concerning the impact of authentication methods and user's

preferences on the satisfaction of the requirements (e.g., security, usability, performance). Although the impact of an authentication method on the performance requirements can be assessed precisely, the same does not apply to other requirements, such as security and usability. Thus, it will be necessary to identify appropriate techniques to update the impact that authentication methods can have on the satisfaction of security requirements. In addition, in some cases, the decision-making in an adaptive system can depend on the time available. The decision to choose an effective authentication method should be made quickly. A possible way to support this activity is to identify strategies to selectively remove less relevant elements from the models used to support decision making.

5.2 Situation-driven risk assessment and security enforcement framework

Maritime transport is a vast environment, containing internal and external processes, interconnected equipment, and human actors. Our aim is to propose a global security methodology that covers both risk assessment and adaptive security policies deployment. We will extend a maritime specific risk assessment methodology (MITIGATE [72]) to suggest adaptive security controls and integrate it with a situation-driven security management framework (DynSMAUG [73], [74], [75]) to dynamically enforce adaptive security policies implementing the security controls. We will follow a situation-driven approach. Situations allow capturing complex and dynamic constraints (e.g., time, location, workflows, etc.) Moreover, situational awareness is about understanding the context and being able to project in the future to improve decision making. This concept will provide a guide for maritime security assessment and security policy enforcement. The resulting framework will be able to suggest adaptive security controls for various critical functions of the cargo transport service required by each security level of each situation. The situations identified at the risk assessment stage will be formally specified using complex event techniques while situation-driven security control will be translated into situation-driven security policies. The underlying security management infrastructure will then be able to dynamically deploy and enforce security policies making security adaptable to each predefined situation and security level.

5.2.1 Application scenario

Besides the sharp increase in cybersecurity threats, maritime transport is an inherently agile environment, in terms of environmental changes, applicable threat agents and attack vectors. We will utilize a typical maritime transport service, mainly cargo transfer, to demonstrate how environment changes affect the underlying risks of systems and consequently the need for an adaptive security framework.

Cargo transfer is usually initiated by a third party, a merchant who sends a purchase order to the producer. After the contract terms (e.g., pricing, documentation, freight charges) have been agreed, the producer contracts a ship agent to deliver the cargo to the destination port. The ship agent makes the arrangements with the ship owner to assure usage of ships; with customs authorities to arrange for the manifest registration number; with the departure port authority to arrange the ship formalities related to the authorization process from the entry of the ship into the port until its exit and then proceed to load the cargo into the vessel for shipment to the destination port. The ship agent contracts a cargo transport agent and assigns the transfer of cargo from the industry to the departure port. Finally, the ship agent sends the relevant documentation to the importer's local agent who has the responsibility for the ship arrival and the regional procedure of delivering the vessel to the importer (see Figure 23).

Obviously, during the cargo transfer a ship may experience various environmental changes, e.g., route through safe and non-safe sea areas, be in proximity with known, unknown, or even hostile vessels, experience variations in network connectivity etc. Therefore, as the situational changes affect the applicable threats, threat agents and effective attack vectors, they consequently affect the resulting security risks in a dynamic manner.

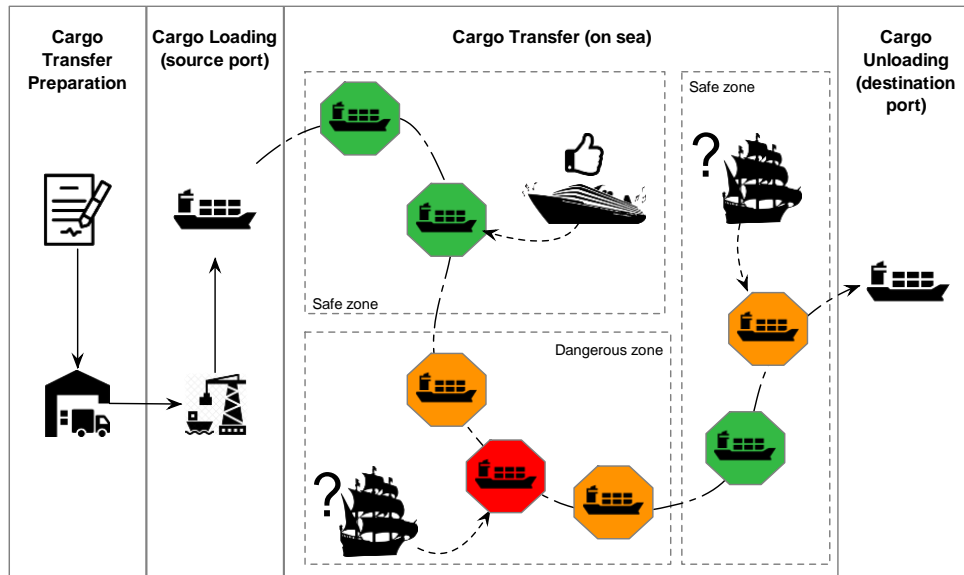


Figure 23 The security situation of a vessel during the cargo transfer service.

In this scenario, we will map assets involved in the cargo transport service, to use it as a test scenario. In our test scenario, situations will be defined, among others, based on the location of the ship (e.g., ship is on port or at the sea) or on information related to the threat status of a ship on route (e.g., high risk areas). This map can later be utilized to unveil interconnectivities of the systems, since it works as a searchable graph, upon which the user can specify entry points and target systems to view potential attack paths. The purpose of this exercise is to unveil potential cyber, physical, and cyber-physical attack paths, throughout the infrastructure used in the context of the cargo transport service. This way, in different situations, different security policies can be generated and dynamically applied under different events defining the situations and various security levels that will be in line with the security levels defined in the International Ship and Port Facility Security (ISPS) code for physical ship safety.

To have a secure, yet functional plan for each security level in each given situation, we procure a central asset-based set of security controls for the entire infrastructure, then proceed to identify which assets are active in each situation and suggest a subset of controls applying specifically to them. Following this process, we avoid overbearing the vessel and port systems with security controls targeted towards inactive attack paths, and we do so following the general instruction of the ISPS code, along with multiple standards (e.g., ISO27001) and taxonomies (e.g., NIST CVE/Intel TAL) utilized in our combined methodologies.

5.2.2 Asset description

A simple solution, followed by existing methodologies is to default to the worst-case scenario by applying the ‘strongest’ security controls, to assure the highest level of authenticity, integrity, availability, confidentiality, non-repudiation and resilience at all situations. Although this policy seems as “being on the safe side”, it affects the operation cost and consequently the actual enforcement of security controls. Applying the strongest security controls is not always possible in maritime systems due to environmental constraints. For example, limitations in network connectivity may prevent the continuous application of security controls that require online verification. In addition, as maritime is a sector with low profit margins, the administrators usually phase limitations in resources for cybersecurity investments. Thus, the security controls should continuously adapt to environmental changes. Risk assessment and mitigation should take into consideration the situational changes and the environmental constraints, in order to dynamically apply

those security controls that can continuously maintain the situational risk below the risk threshold, but in a cost-efficient manner.

To deal with such dynamic security needs, we propose an adaptive security framework that mainly covers the analysis and execute phases of the MAPE loop, by combining dynamic risk assessment and situational driven security policy deployment. We extend a maritime-specific risk assessment methodology (MITIGATE) to suggest adaptive security controls, and integrate it with a situation-driven security management framework (DynSMAUG) to dynamically enforce adaptive security policies implementing security controls. We will first describe the existing assets which are the main building blocks for the proposed adaptive, situation-driven security framework. Then we will describe how these assets are combined and properly extended.

5.2.2.1 Description of MITIGATE Maritime Risk Assessment Methodology

MITIGATE [76] is a collaborative, evidence-driven Maritime Supply Chain Risk Assessment approach. MITIGATE captures various threats arising from the Supply Chain (SC) environment, including threats associated with sectorial infrastructure interdependencies, and the associated cascading effects. MITIGATE extends Medusa methodology [77] and is compliant with the ISO/IEC 27000 information security and the ISO/28000 SC security management international standards. It decomposes in six steps that are gradually undertaken to evaluate risks, as shown in Figure 24.

Step 1. Supply Chain Service (SCS) analysis: aims to identify the boundaries of the risk assessment process, identify the under examination SCS of the maritime transport sector, analyse and model its generic components, i.e., SCS processes, involved business partners and assets operating within the SCS processes. It produces the SCS process models, the SCS asset inventory and the SCS asset interdependencies graph.

Step 2. Cyber threat analysis: aims to identify and assess all cyber threats that are related to the identified assets of the under examination maritime transport SCS. The outcome is a list of cyber threats assessed in a qualitative scale “Very Low” (“VL”), “Low” (“L”), “Medium” (“M”), “High” (“H”), “Very High” (“VH”).

Step 3. Vulnerability analysis: identifies the system vulnerabilities. It supports the estimation of the possibility of a vulnerability exploitation for an asset (individual vulnerability assessment) as well as the cascading effects and propagation of the vulnerability to the interconnected SCS assets concerning the accessibility to a SCS asset (cumulative vulnerability assessment), and the attacker capacity to infiltrate the SCS asset network (propagated vulnerability assessment). It also produces possible attack paths revealed from attack graphs [78].

Step 4. Impact analysis: It estimates the effect that can be expected as a result of the successful exploitation of a vulnerability that resides in a given SCS asset. The step delivers the impact levels for each identified vulnerability to the SCS assets.

Step 5 and 6. Risk Assessment and Mitigation: predicts the risk level for the identified assets, taking into account the produced threat level, vulnerability level and impact level from the previous steps. Using prediction analytics [79, 80], estimates the expected damage due to an attack that exploits multiple vulnerabilities to support the elicitation of the optimal security controls that must be undertaken to minimize

the effect. It also produces attacker and defender strategies along with payoffs estimating the level of damage for each combination of attack and defense strategy.

SCS Analysis	Cyber Threat Analysis	Vulnerability Analysis	Impact Analysis	Risk Assessment	Risk Mitigation
S1.1: Goals & Objectives	S2.1: SCS Cyber Threat Identification	S3.1: Identification of Confirmed Vulnerabilities	S4.1: Individual Asset Impact Assessment	S5.1: Individual Asset Risk Assessment	S6: Risk Mitigation
S1.2: Business Partners	S2.2: SCS Cyber Threat Assessment	S3.2: Identification of Unknown Vulnerabilities	S4.2: Cumulative Impact Assessment	S5.2: Cumulative Risk Assessment	
S1.3: Modelling		S3.3: Individual Vulnerability Assessment	S4.3: Propagated Impact Assessment	S5.3: Propagated Risk Assessment	
		S3.4: Cumulative Vulnerability Assessment			
		S3.5: Propagated Vulnerability Assessment			

Figure 24 An overview of the MITIGATE methodology.

5.2.2.2 Description of DynSMAUG

DynSMAUG is a dynamic security management framework[74]. The specificity of its approach consists in placing the concept of situation at the core of security management. On the one hand, situations capture the dynamic constraints (time, location, risk, etc.) and organize them into a stable and logical concept. Situation-based security policies are simpler and more readable. Also, managing high level policies, close to business, reduces the gap between security requirements and the effective security policy enforced by security devices, and then limits the security policy translation errors. On the other hand, making security policy more independent from technical constraints minimizes the impact of changing security mechanisms and simplifies policy life cycle management.

A situation is a particular time frame of interest with a beginning, a life span and an end [81]. The beginning and the end of a situation can be determined by combining multiple events coming from multiple sensors and occurring at different moments [75]. Indeed, the beginning and the end of a situation involving multiple entities and multiple conditions cannot be limited to simple events captured by one single sensor. Moreover, events being instantaneous, combining multiple events requires complex temporal operators (event ordering, event existence/absence, time windows, etc.) to specify the beginning and end of situations. Complex Event Processing (CEP) provides such features. CEP is “*a defined set of tools and techniques for analyzing and controlling the complex series of interrelated events that drive modern distributed information systems*” [82]. CEP solutions allow specifying complex events through complex event patterns that match incoming event notifications based on their content as well as some ordering relationships on them. Thereby, the beginning and end of the situations elicited during the situation elicitation phase are expressed in a CEP language. The specification of the situations depends on the sensors available in the vessel and their characteristics. Different patterns for describing situations using CEP have been proposed in [74] and [75].

In parallel, situational security controls are refined into situation-based security policies, to be enforced by a situation-based security decision making entity. In our approach, situations are specified and calculated at the situation manager's side. Therefore, the security policy refers to them only. Hence, we represent

situation-based security policies in a generic way as: when situation and some condition then authorization decision and/or obligation(s) where the condition statement is any constraint on any characteristic of the entities involved in the situation as well as the situation itself [74]. This generic approach is flexible enough to express changes of security controls when the situation is shifting to another one using the reactive rules pattern: when situation and situation starts [and some condition] then obligation(s) where the obligations reflect the security controls modifications the security management system will enforce. Situation-based authorization rule is another pattern for specifying adaptive authorization controls: when situation and some condition then authorization decision.

Both the situation specification and the situation-based security policy are injected into the security management system [74].

The actors of the DynSMAUG deployment architecture (Figure 25) are the following:

The *sensors* produce context events. A sensor can be any system available in the target vessel that can trigger context events such a physical button activated by a human, an intrusion detection system, an alarm, a GPS, a proximity sensor, etc.

The *situation manager* continuously calculates situations according to a low-level situation specification. It consumes context events triggered by the sensors and produces situation events. A situation event contains the beginning of the new situation and the end of the last active situation.

The *control center* is the brain of our security deployment framework as it performs the security decision making process. It takes as input situation events and output security decisions, i.e., security controls to be enforced based on the specific situation. Multiple control centers can be deployed for scalability and/or performance reasons. Different strategies can be considered to coordinate decisions [83],[84].

The *actuators* only consume decision events and enforce security controls. Actuators can be any system that can be controlled by a software (e.g., a door that can be locked/unlocked, configurable IT systems, etc.)

The *event broker* is the distribution middleware that transmits all the events between the actors following the publish-subscribe pattern. The broker divides events into three topics: context events, situation events and decision events. The broker also ensures that only authorized actors (sensors, actuators, situation manager and the command center) can assess it.

5.2.2.3 Proposed adaptive, situation-driven methodology

We extend and combine the MITIGATE maritime risk assessment methodology, with the DynSMAUG situation-driven security management framework, to dynamically enforce adaptive security policies. The proposed methodology, illustrated in Figure 26 is comprised of three phases, described in detail below: (1) situations elicitation, (2) situation-based risk management and (3) situation-based policy deployment. Note that although the methodology is specifically crafted for the maritime sector, it is possible to extend it for other critical sectors, by properly adjusting the situation elicitation and other tasks, such as the threat agent mapping, to other sectors.

Phase 1: Situation elicitation

The first phase focuses on eliciting the set of situations in which a vessel can be found. The purpose is to unveil potential cyber, physical, and cyber-physical attack paths, throughout the infrastructure used in the context of the cargo transport service. Going further to specify a situation, we study which threat agent profiles have the required capabilities to exploit each identified attack path, and which are the security policy shortcomings that may allow such an event in a specific time and place:

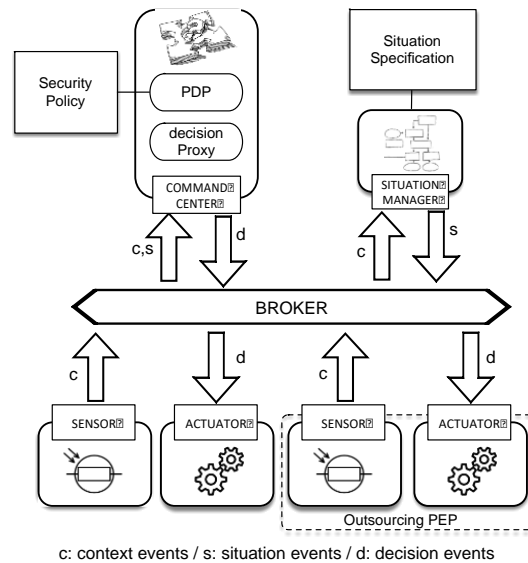


Figure 25 An overview of the DynSMAUG methodology.

Where is the vessel located at this point in time?

Which vessel systems should be active in this specific location? (Hardware/Software)

Which vessel communication channels should be active in this specific location? (related to threat actors, interference from natural phenomena and equipment restrictions)

What does the security policy dictate for human-to-equipment interaction and equipment-to-equipment interaction in each specific location?

Which external threat actors are most active and which internal threat actors percentage (e.g., disgruntled employee) are likely to have sufficient access to initiate an attack towards critical assets in this location?

The purpose and significance of considering situations for maritime cybersecurity, are highly related to the complex nature of the underlying environment. To design a situational maritime security approach, we will take into consideration an abundance of parameters that can be boiled down to the five dimensions of the situation model proposed by Zwaan and Radvansky [85]:

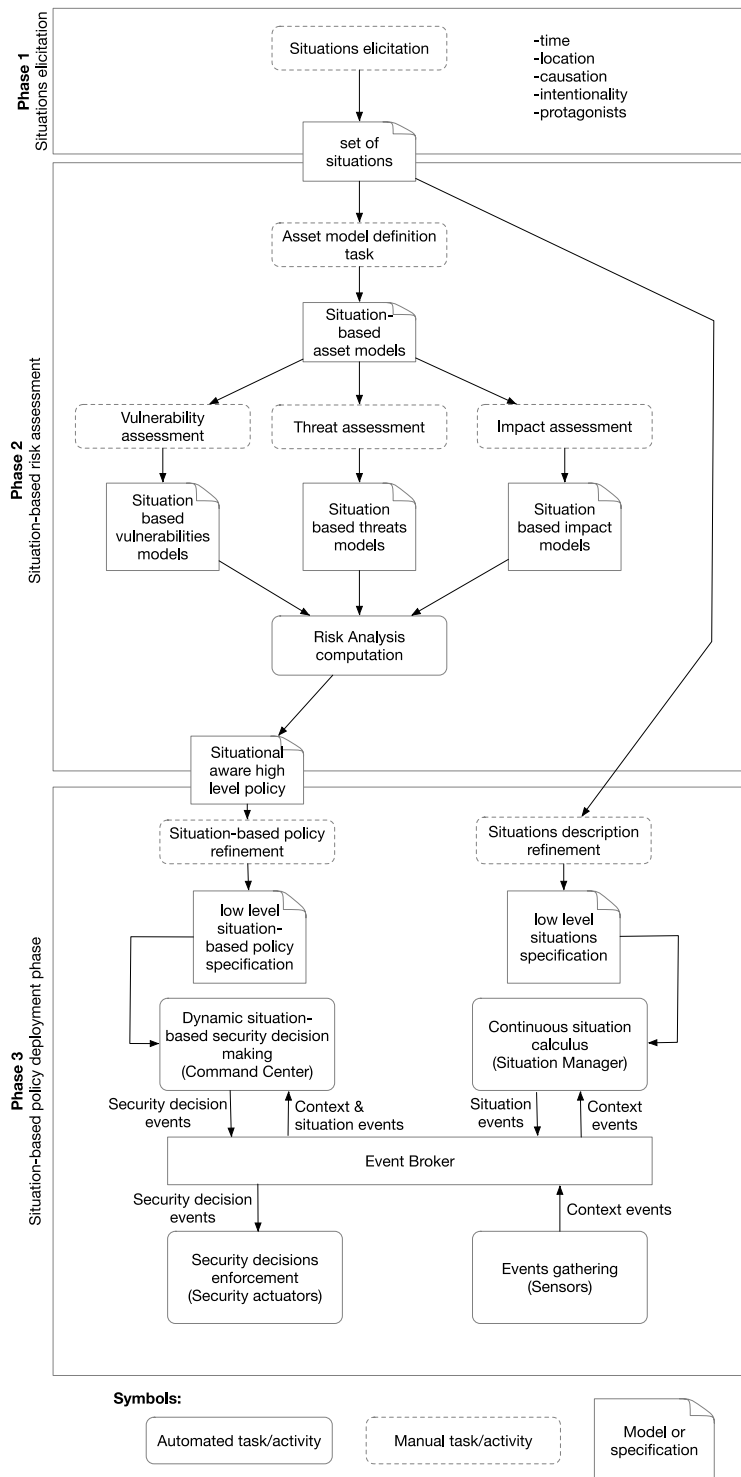


Figure 26 An adaptive situation-driven risk assessment and risk enforcement framework.

The *Protagonists/objects dimension* in the context of maritime transport security requires the study of human and system agents that can potentially interact with the vessel. This includes internal and external actors, such as human or system agents acting on ports, vessels or elsewhere. At the same time, actors may be trusted (e.g., a port official adhering to the protocol) or malicious (e.g., a disgruntled employee or pirates

at open seas)/ Actors may include not only humans be systems also. Active systems/assets and catalogued information also pertains to this dimension.

The *Space dimension* relates to the evolution of the physical locations of the vessel and other protagonists.

The *Time dimension* includes topics related to time periods, maritime transport workflow steps regarding the mission of the vessel, etc.

The *Causation dimension* deals with deducing evidence that can be inferred by other contextual data. For instance, analyzing the speed and trajectory of another vessel may reveal that both vessels will be in physical proximity in the near future.

The *Intentionality dimension* focuses on the goals of the protagonists. Attackers have threat goals while honest parties perform tasks that adhere to their role in the system. Therefore, security policies and procedures that dictate the behavior of humans are studied in this dimension too.

We propose to organize the elicitation of situations by combining the five dimensions of situations and a situation tree structure. A situation tree is a mind map where each level of the tree corresponds to a specific question dedicated to one dimension (e.g., the question 'Where is the vessel?' refers to the space dimension). Sibling nodes in a situation tree are literals representing the possible answers to the question (e.g., the vessel can be *on port* or *at sea*). The leaf nodes are the situation names. The definition of each situation is the path from the root node to the leaf. The resulting situation clause is the conjunction of node literals in the path. By applying our situation elicitation methodology on the cargo transfer service, a concrete set of situations is produced, as shown in Figure 27. For instance, situation S1 means the vessel is on port and loading cargo, while situation S8 corresponds to the vessel is at sea and in a dangerous area and near an unknown vessel.

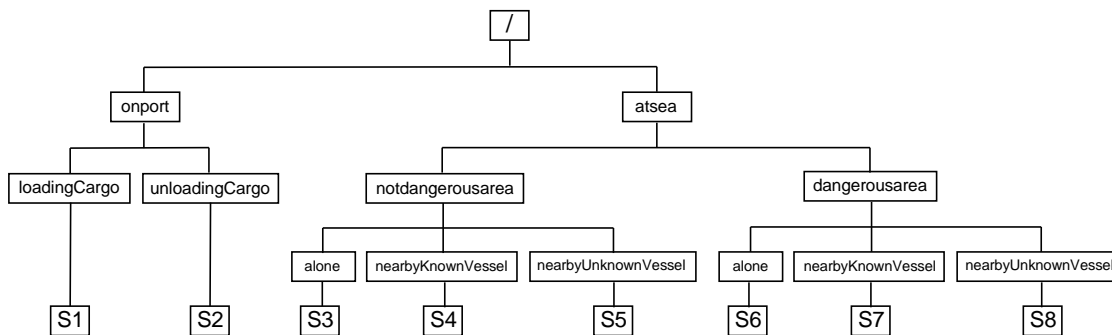


Figure 27 Decision-tree based situation elicitation.

Phase 2: Situation-based risk assessment

Based on the situations defined in the previous phase, all the risk assessment tasks defined in the MITIGATE methodology such as asset modeling, threat, vulnerability, and impact assessment are properly adjusted to each situation, to output a fine-grained, situational risk assessment. In MITIGATE assets, threats, vulnerabilities, and threat agents are instantiated with the use of datasets pulled from open sources provided by widely known organizations like MITRE and the National Institute of Standards and Technology (NIST).

For each asset we identify relevant vulnerabilities based on the Common Vulnerabilities and Exposures (CVE) database. As illustrated in Figure 28, threats are instantiated based on the Common Attack Pattern Enumeration and Classification (CAPEC) catalogue and the Adversary Tactics Techniques and Common Knowledge (ATT&CK) framework, while security controls are derived from MITRE's D3FEND matrix. In addition, new relationships were created by utilizing common characteristics that are utilized in the above datasets.

Situational Asset Model Definition

Situation elicitation is used as an input to define the situation-based risk assessment phase. Based on the defined situations a manual asset modeling process is used to define alternative asset models that represent the different situations. In contrast to risk assessment approaches that do not define situations, e.g. [77], [76], [72], in our methodology, for each situation different asset models are defined to capture the interconnections and the dependencies among assets in different situations.

Step 1 - Service Identification

The first step involves the identification of the available internal maritime services for an organization. A comprehensive list of all maritime services along with their corresponding processes must be generated, such as cargo loading, cargo transfer, cargo unloading, etc. A service is a collection of processes that are part of a specific maritime ecosystem and may depend on external actors. The dependencies between services and business partners, as well as services and processes must be identified, so that the risk assessment can proceed.

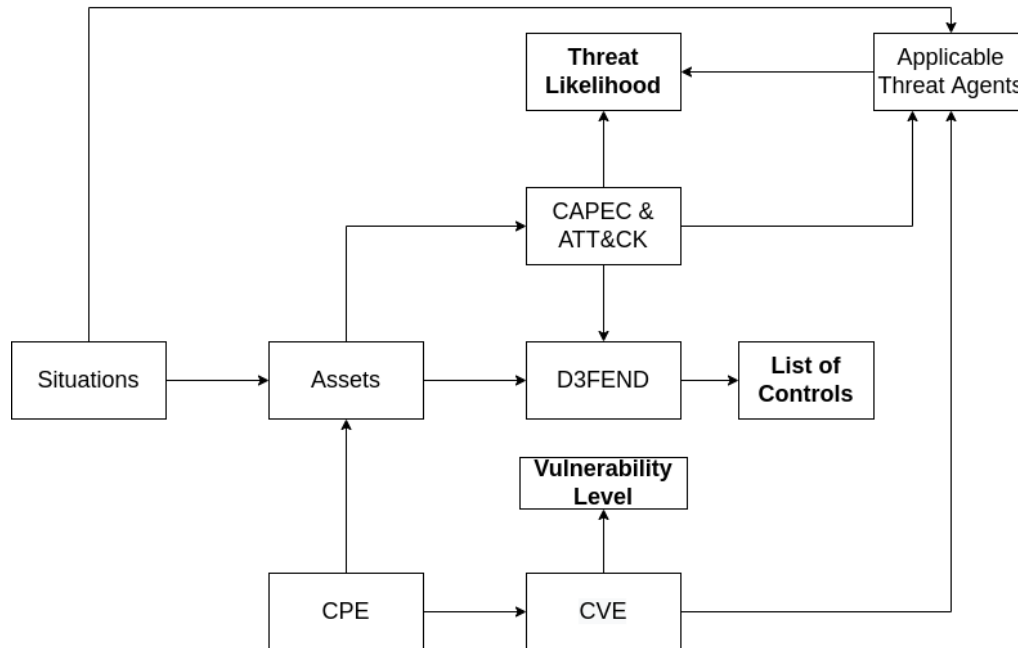


Figure 28 Relationships among risk related entities and relevant datasets.

Step 2 - Asset Identification and Cataloguing

Having identified the available services and processes, the next step is to decompose each process, identify the assets on which it depends and define the asset criticality. These assets would mainly be internal system components that are controlled by the examined organization(s). The available asset types in the context of our methodology are hardware and software assets, where the latter are divided into operating systems and application software. The specific entries are derived from the Common Platform Enumeration (CPE) catalogue. While the underlying MITIGATE methodology defines assets without considering different situations, in the proposed methodology asset identification is specific to each defined situation. A different asset map is performed and catalogued for each situation, in the context of the identified maritime processes.

For example, vessels utilize various communication and navigation systems such as MSW, Ship reporting systems (SRS) and Automatic Identification Systems AIS only when the ship is on route, while the same systems are inactive while the ship is on the port. In the same way, collision avoidance systems may be activated on route or may be in a different state before the departure or upon arrival. Hence, a different asset

map applies to different situations, and assets have a different criticality level for different services and situations, defined in a simple *High, Medium, Low* scale.

Situational Threat Assessment

Having completed the situation-based asset models defined in the previous phase the threat assessment procedure can be implemented utilizing the recorded information. While in MITIGATE [76] threats are defined for each asset and are situation-agnostic, in the proposed methodology we define applicable threat agents that may activate specific threats per situation. This involves the following steps:

Step 1 - Threat Mapping for each situation the involved assets and the applicable threats are mapped. Due to functionality, security and costs, all assets might not be active in all situations. Furthermore, some threats may be easy to activate in some situations only. For example, GPS spoofing is an active threat when a vessel is on route. On the other hand, if a vessel is at port the activation of the same threat will not have any observable impact. Utilizing the threat profiling approach while parsing and searching through a series of known sources, ranging from social media to threat and vulnerability catalogues for references of incidents related to specific CAPEC categories or ATT&CK tactics, further information can be extracted to support the threat level calculation and to build threat characteristics that can be compared to threat agent characteristics.

Step 2 - Threat Agent Mapping

In the second step, threat agents are mapped to each situation. For example, while some threat agents may apply to all situations (e.g., the threat agent ‘cyber-criminal’ may be applied in all the situations defined in Figure 27, while other threat agents may only apply to specific situations (e.g., ‘pirates’ are only applicable to situations S6-S8). To fully identify threat agents, we utilize Intel’s Threat Agent Library (TAL) and inherit the profile characteristics for maritime-specific profile instances. This implementation essentially supports an extra layer of filtering, since it allows the risk assessor to not only map the active threats per situation, but also to match them with the threat agents that are expected to have sufficient capabilities to activate threats. We adopt the approach of [86], where threat agent capabilities are expressed as Common Vulnerability Scoring System (CVSS) v3 vectors. The purpose of this exercise is to map threat agents along with the exact attributes that express their capability of exploiting catalogued vulnerabilities. For example, a threat agent that has been attributed a capability of ‘Adjacent Network’ attack vector, is able to exploit vulnerabilities that require either ‘Network’ or ‘Adjacent Network’ attack vector but cannot exploit vulnerabilities with ‘Local’ or ‘Physical’ attack vector. For threat assessment, different threats agents may apply in each situation and different threat agents can take advantage of different vulnerabilities based on their capabilities, which subsequently affects the underlying vulnerability map enabled by each threat.

Step 3 - Situational Threat Likelihood and Profile Filtering

Combining the output of the previous steps an information map is derived, where a set of attributes that resemble ones used to characterize threat agent profiles in known taxonomies is observed, more specifically resources required, and skills required from CAPEC can be connected to an attacker’s capability characteristics found in TAL. At the same time consequences from the CAPEC catalogue and tactics from ATT&CK can be connected to an attacker’s motivation. Utilizing the catalogued characteristics of both threats and threat agents, we can produce a refined threat likelihood level.

Situational Vulnerability Assessment

Again, as in the threat assessment, the MITIGATE vulnerability assessment is situation-agnostic, in the proposed methodology the vulnerabilities are automatically assessed according to each situation, as described below.

Step 1 - Vulnerability Identification

This step focuses on the identification and assessment of confirmed vulnerabilities of assets, which can be exploited and lead to successful attacks. Our methodology utilizes the National Vulnerability Database

(NVD), which contains over 160.000 detailed entries in a structured format, along with other reliable online sources to identify the characteristics of vulnerabilities. Vulnerabilities are product based, which means that they are targeted towards a specific asset, be it hardware, operating systems, or applications; such components are listed in the Common Platform Enumeration (CPE) catalogue, while their connections to Common Vulnerabilities and Exposures (CVE) entries reside in NVD. Therefore, the list of individual vulnerabilities of the assets from the asset modelling step can be created from existing connections between the two catalogues.

Step 2 - Vulnerability Scoring

Once the vulnerabilities are identified, then it is necessary to determine its exploitability metrics. Our methodology utilizes the CVE's provided by the National Vulnerability Database (NVD), which are recorded along with part of their CVSS v3.1 and v2.0 attributes. The MITIGATE tool utilizes the CVSS v2 exploitability metrics to define a vulnerability level as illustrated in Figure 29. By combining the available exploitability characteristics for each recorded vulnerability (i.e., CVE) a single value is produced to express the vulnerability level based on the CVSS score, as defined in [72].

Step 3 - Vulnerability Score Assessment

AV \ AC	Local			Adjacent			Network		
	Low	Medium	High	Low	Medium	High	Low	Medium	High
Multiple	VL	VL	L	L	L	M	M	M	H
Single	VL	L	M	L	M	H	M	H	VH
None	L	M	M	M	H	H	H	VH	VH

Figure 29 Vulnerability level calculation matrix.

Changes in the asset model along with its interconnections from one situation to another, may affect the vulnerability level of the assets. As the connectivity of networked systems, as well as their physical and logical accessibility depends on the situation, the vulnerability list produced for the same asset in a different situation may differ. For example, when an asset with a vulnerability that has a 'Network' attack vector is not connected to a network with internet connectivity, the vulnerability does not apply in this case and may only be activated by an 'Adjacent network' vector.

Step 4 - Threat Agent Scoring

Since maritime services take place in a vast, dynamic environment, the human actors and more specifically the threat agents active around each service are multiple and may differ. Throughout this step we utilize real life incidents to place threat agent profiles in the context of specific maritime services. Furthermore, to characterize the threat agents with specific attributes that can directly be compared to specific vulnerabilities, we use the approach from our previous work [86]. In [86] we expressed threat agent profiles from the healthcare sector as CVSSv3 capability vectors, which we compared to CVSSv3 vulnerability vectors. We inherit the recorded characteristics for similar profiles in the maritime environment, which presents a final underlying challenge for this step. The vectors need to be translated to CVSSv2 to be compatible with the mitigate solution. The transformation of CVSSv3 vectors to CVSSv2 vectors is based on the approach of [87].

Situational Impact Assessment

Step 1 - Impact Identification

To derive the impact that existing vulnerabilities may cause, the list of vulnerabilities procured by the Situational Vulnerability Assessment is parsed and the impact section of the CVSS vector is extracted to another list. This section contains three values referring to confidentiality integrity and availability impact, which are combined to produce a single value in the next step.

Step 2 - Impact Level Calculation

This step focuses on the Impact level calculation, which measures the effect that can be as the result of the successful exploitation of a vulnerability that resides in a critical asset. In CVSS the three security criteria Confidentiality (C), Integrity (I) and Availability (A) are rated in a three tier scale of *None*, *Low*, *High*. We can define a mapping from the three-tier scale onto a five-tier scale ranging from *Very Low (VL)* to *Very High (VH)* to combine these three characteristics (see Figure 30). This will provide a single estimation for the overall impact of a specific asset/vulnerability combination. As defined in the underlying MITIGATE methodology, the confidentiality, integrity, and availability sub-scores of the recorded vulnerabilities are used as input, to output a single impact level.

Step 3 - Situational Impact

The impact of security attacks may also vary according to the situation. In a typical risk assessment

C \ I \ A	None			Low			High		
	None	Low	High	None	Low	High	None	Low	High
None	VL	VL	L	L	L	M	M	M	H
Low	VL	L	M	L	M	H	M	H	VH
High	L	M	M	M	H	H	H	VH	VH

Figure 30 Impact level calculation matrix.

method impact is based on the 'worst-case' scenario. For example, if the unavailability of a navigation system has very high impact when the vessel is on route, the corresponding impact value will be used while assessing all possible threats that may result in the unavailability of the system in all situations. By applying different impact values according to the situation, more fine-grained risk values will be produced according to the situation. The situational impact values are derived by the combination of the initial impact values and the asset criticality set for each asset in the context of a situation, as shown in Figure 31. Essentially the impact level calculated in the previous step is refined based on the asset criticality per situation to output a situational impact value. Since the same asset may be more important in different situations, a criticality level is defined for each asset per situation. For example, the criticality level of a GPS system may be high while the vessel is on route but may be low when the vessel is on the port. This criticality level is used to weigh the initial impact level defined in the previous step.

Initial Impact	Asset Criticality		
	Low	Medium	High
Very Low	VL	L	L
Low	L	L	M
Medium	L	M	H
High	L	H	H
Very High	M	H	VH

Figure 31 Situational impact calculation matrix.

Situational Risk Assessment

Finally, by combining the situation-based asset models, threat, vulnerability and impact assessment results, the risk analysis engine will output the relevant risks, along with their assessed values, for each different situation.

Let A denote an asset under examination, T a threat identified throughout the Situational Threat Assessment, V a vulnerability identified throughout the Situational Vulnerability Assessment and I its situational Impact. Then the situational risk level caused on asset A by threat T and the Impact I of Vulnerability V in situation S is defined as shown below:

$$R_S(A, T) = T_S(A, T) \otimes V_S(A) \otimes I_S(A)$$

In the above equation, $T_S(A, T)$ represents the threat level calculated throughout the Situational Threat Assessment step, $V_S(A)$ represents the vulnerability level calculated throughout the Situational Vulnerability Assessment step and finally $I_S(A)$ represents the impact calculated throughout the Situational Impact Assessment step, as defined in the above sections. The resulting situational risk level is computed based on the risk table defined in MITIGATE [76], ranging in scale from *Very Low* to *Very High*.

Situational aware high level security policy

As the risk assessment results are dynamically computed for each situation, granular security policies can be defined for different situations. Instead of producing a static list of security controls, expressed as high-level policies, the suggested security controls will be fine-grained based on the different risks that correspond to each situation. The high-level security policy will then be further refined and instantiated to particular security controls in the next phase of the methodology.

Step 1 - Existing Control Identification and Assessment

This step reviews the identified vulnerabilities and threats from the previous phases and identifies the level of mitigation based on the existing controls. First the implemented controls per asset per situation are identified and listed. The existing security controls may provide partial or full mitigation of the effect of existing threats or vulnerabilities. A decision-making process based on the existing risks is implemented to illustrate the functionality of existing controls.

Step 2 - Situational Control Identification and Application

Utilizing the information procured by the previous step, the residual levels of threats, vulnerabilities and risks calculated while incorporating the effect of the initial security controls are mapped. Having identified these values, further security controls that will fully mitigate risks can be suggested. To achieve this result two approaches are considered:

Applicable controls for techniques catalogued in the (ATT&CK) framework are listed in MITRE's D3FEnd Matrix.

Applicable controls for existing vulnerabilities can be found in the NVD's references for each individual vulnerability.

Phase 3: Situation-based policy deployment

The last phase of our methodology consists in enforcing these situational controls. This involves refining the situations elicited in Phase 1 and the high-level situational security controls produced in Phase 2 into low level rules that can be deployed by a security management system.

A situation is a particular time frame of interest with a beginning, a lifespan, and an end [81]. The beginning and the end of a situation can be determined by combining multiple events coming from multiple sensors and occurring at different moments [75]. Indeed, the beginning and the end of a situation involving multiple entities and multiple conditions cannot be limited to simple events captured by one single sensor. Moreover, events being instantaneous, combining multiple events requires complex temporal operators (event ordering, event existence/absence, time windows, etc.) to specify the beginning and end of situations. Complex Event Processing (CEP) provides such features. CEP is a defined set of tools and techniques for analyzing and controlling the complex series of interrelated events that drive modern distributed information systems[82]. CEP solutions allow specifying complex events through complex event patterns that match incoming event notifications based on their content as well as some ordering relationships on them. Thereby, the beginning and end of the situations elicited during the situation elicitation phase are expressed in a CEP

language. The specification of the situations depends on the sensors available in the vessel and their characteristics. Different patterns for describing situations using CEP have been proposed in [74] and [75]. The resulting low level situations specification is then provided to a situation manager that continuously calculates the current situation.

In parallel, situational security controls are refined into situation-based security policies to be later enforced by a situation-based security decision making entity. In our approach, situations are specified and calculated at the situation manager's side. Therefore, the security policy refers to them only. Hence, we represent situation-based security policies in a generic way as: *when situation and some condition then authorization, decision and/or obligation(s)*, where the condition statement is any constraint on any characteristic of the entities involved in the situation as well as the situation itself [74]. This generic approach is flexible enough to express changes of security controls when the situation is shifting to another one using the reactive rules pattern: *when situation and situation starts and some condition then obligation(s)*, where the obligations reflect the security controls modifications the security management system will enforce. Situation-based authorization rule is another pattern for specifying adaptive authorization controls: *when situation and some condition then authorization decision*.

Both the low-level situation specification and the situation-based security policy are injected into the security management system [74]. The actors of our deployment architecture are the following:

The *sensors* produce context events. A sensor can be any system available in the target vessel that can trigger context events such a physical button activated by a human, an intrusion detection system, an alarm, a GPS, a proximity sensor, etc.

The *situation manager* continuously calculates situations according to a low level situation specification. It consumes context events triggered by the sensors and produces situation events. A situation event contains the beginning of the new situation and the end of the last active situation.

The *control center* is the brain of our security deployment framework as it performs the security decision making process. It consumes both context and situation events, takes security decisions based on a situation-based security policy and produces decision events. Multiple control centers can be deployed for scalability and/or performance reasons. Different strategies can be considered to coordinate decisions[83],[84].

The *actuators* only consume decision events and enforce security controls. Actuators can be any system that can be controlled by a software (e.g., a door that can be locked/unlocked, configurable IT systems, etc.)

The *event broker* is the distribution middleware that transmits all the events between the actors following the publish-subscribe pattern. The broker divides events into three topics: context events, situation events and decision events. The broker also ensures that only authorized actors (sensors, actuators, situation manager and the command center) can access it.

To summarize, the main novelty of the proposed methodology is the definition of a dynamic risk assessment and policy enforcement framework, targeted to the maritime transport environment. In contrast to existing maritime specific risk assessment frameworks which are static in terms of security policy enforcement, the proposed framework is dynamic by design. As the security risk of the examined systems is affected by events and situations, the resulting risk level also varies, leading to situation-specific enforcement of security controls. By defining situational asset and threat models, we continuously map active assets to active threat agents for each situation, thus filtering out risk that are not active or very low, in various situations. Following this process, we avoid overbearing the vessel and port systems with security controls targeted towards inactive attack paths. This allows to avoid a policy of 'always defaulting in the highest risk' which in practice may lead to reduced security controls due to lack of resources, lack of efficiency in procedures or other environmental constraints. In addition, the proposed framework supports automation of the security policy enforcement, by allowing the implementation of automated security policies per situation. By sensing events that indicate changes in the situation such as current location or proximity with known or unknown

vessels, it is possible to dynamically adapt the applied security controls to the corresponding situational risk level.

5.2.3 Outputs

- The paper titled "*An Adaptive, Situation-Based Risk Assessment and Security Enforcement Framework for the Maritime Sector*", co-authored by Christos Grigoriadis, Romain Laborde, Antonin Verdier, and Panayiotis Kotzanikolaou, was published in MDPI Sensors journal.
- Christos Grigoriadis, Romain Laborde, Antonin Verdier, and Panayiotis Kotzanikolaou. "An Adaptive, Situation-Based Risk Assessment and Security Enforcement Framework for the Maritime Sector." *Sensors* 22, no. 1 (2022): 238.

5.2.4 Asset limitations and future work

The risk assessment phase is semi-automated, as various steps require manual intervention. For example, to introduce further automation to the risk assessment procedure, further research is required in the context of identifying applicable threat agents and groups in specific geographical coordinates, based on past data of recorded attacks.

Another point that needs further work is assessing the level of confidence/assurance of the situation calculus. Indeed, a calculated current situation might deviate from the actual current situation, due to low quality sensors or due to attacks against the sensors themselves. We need to improve our methodology to cover these risks. Interesting approaches related to the concept *Quality of Context* and *Quality of Situation* [88], [89] may provide useful insight towards this direction. Secondly, specifying complex situations using a rule-based language may be error prone for very complex situations with many entities or context information. Complementary approaches that may handle this problem may include (i) building up a simulation environment for testing/validating situation specifications, and (ii) applying machine learning techniques.

5.3 Adaptive risk assessment

The aim of the SYSVER (SYStem VERification) asset, is to provide a high-level verification of the correctness of security policies in complex and dynamic networked systems. These kinds of systems include large numbers of heterogeneous entities that are interconnected in both physical and logical ways. We focus on the relationships between high-level services provided in one or more systems and their deployment through a complex network of service components. Moreover, we consider the existence of several different types of users of these services, that can be human users of cyber agents. In this type of complex scenario, there are several resources that can be accessed in different ways, and we need to verify that the configuration of all the elements in the system is correct with respect to the security policies defined, and in particular with respect to the access control policies, that specify who can do what on which object.

This problem is particularly relevant when we deal with large systems that are in fact a composition of independent systems that need to interact in complex ways. In this scenario, it is not possible to assume that all the elements in the overall system are owned and controlled by a single entity and, as such, it is not always possible to enforce specific configurations on all the elements. We consider the problem at the composition level where we have several configured systems that are connected in several ways to collaborate and to provide high-level complex services. So, we need to consider the physical connections (communication networks), the interactions between service components, and the access requirements of users (both human, and cyber agents) of the overall system.

It is important to note the necessity of this kind of level reasoning that expand the simpler network reachability verification. In fact, if we consider only the configuration of a network (with its firewalls, routers and so on) we may miss high-level communication paths. For instance, even when the network has been correctly developed with network-segmentation in mind, still some channels are opened if services belonging to different sub-networks are needed to interact. These kinds of channels can also be dynamic in their nature so that it is not enough to consider the underlying network implementation, but it is necessary to reason at the higher-level of logical channels between service components.

The dynamic nature of many of the involved components is critical to this approach, and in fact, the aim of this asset is to support an **adaptive security risk analysis** that can follow the system as it changes.

Our approach is to combine all the available information (both at physical and cyber levels) to build a complete state diagram representing what each user/agent can do in the system, leveraging all his/its capabilities, in a dynamic scenario. For instance, a human user can physically walk between different environments and try to access different resources. But he/she can also leverage some of his/her capabilities (e.g., passwords, authentication tokens, and so on) to interact with service components that provided direct or indirect access to other service components and to other resources. All these kinds of interactions can lead the user/agent to acquire more capabilities and/or to reach remote service components in a dynamic way. The resulting state diagram can be then analysed to understand what the user/agent is able to access and how. The combination of all the state diagrams leads to the overall evaluation of the effective implementation of high-level security policies. Of course, given the context of the adaptive security, this is an evolving process, where any kind of change in some configuration of some elements (for instance, the addition of a service interaction, or a new connection between sub-networks) requires the re-evaluation of the possibly modified state diagrams, to find out if the new scenario has opened new access paths to possibly critical resources, introducing risk elements in the overall system. Another source of changes in the system is also the possible discovery of new potential vulnerabilities affecting one or more system components. In this case, the overall risk analysis is clearly affected by what channels a compromised element could open to a malicious user.

5.3.1 Application scenario

A transport vessel in the maritime scenario is exactly the type of a complex system with heterogeneous elements and processes that we consider. Different kinds of networks, sub-systems, devices, services, coexist on the same ship and different operators require access to different parts of the vessel for different purposes. Both cyber and physical aspects are to be considered in this scenario. Moreover, a transport vessel is inherently a moving object, and this changes the context in which the security of the overall system must be assessed.

This type of complexity requires that the process of the analysis of the security of the ship must consider all the possible interactions between the operators, the different services, and in the different situations the ship could be in.

To provide a high-level description of a vessel system, we consider 3 different communication networks: an “IT network”, an “OT network”, and a “guest network”. The first one represents the main network that connects IT systems, such as communications, navigation, and monitoring. The second one, instead, represents the network where OT devices (e.g., PLCs) are connected and manage the execution of the control process of the ship (a SCADA system from this point of view). The last one, the “guest network” represents the non-critical network that allows the crew to connect through their devices to the services provided by the ship and to the internet. In fact, the three networks provide different services, with different criticalities and different relationships among the components.

In terms of the adaptive scenario, we consider two different situations that can be analyzed and evaluated:

S1) ship is on the sea alone: this is the base case, low risk scenario. A potential vulnerability affecting a node in the IT network could let crew personnel access some resource in the OT network.

S2) ship at the port. High risk scenario. In this case, another set of services and personnel, coming from the Port Management System (PMS), are connected to the system. These services and users could be required to interact with the components in the IT and OT networks of the ship. For instance, to perform cargo management operations from the port, or to allow technicians to perform maintenance procedures onboard, or remotely. In this situation there are two possible attack scenarios to consider: if the ship is compromised, it can affect the subsystems at the port; if, on the contrary, a malicious user has gained control of some service component in the port facilities, it can try to gain access to the resources on the ship. Moreover, at the port it is possible that the ship is in proximity with other ships, which could imply several other potential threats.

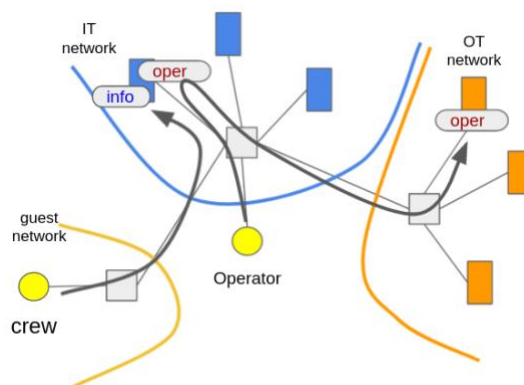


Figure 32 High level view of Scenario S1.

In both the scenarios considered, first, we have to build the models for the networks and the involved resources, including the services and their relationships.

In S1), we consider the single system represented by the vessel, that has several sub-components. In this context, there are connections between the three IT/OT/guest networks that can be leveraged by authorized users but that represent potential hidden communication channels. This scenario is represented, at a high-level point of view, in Figure 32.

In this scenario, an IT network node provides access to a general “info” resource that can be accessed by the crew, from the guest network. The same IT node is also used by an Operator to interact with nodes in the OT network for a critical “oper” operation (eg. navigation, maintenance). The risk in this scenario can be evaluated by checking the correct configuration of the IT node and its deployed services (that should ensure properly configured access control mechanisms). However, a possible source of risk can derive from potential vulnerabilities affecting the node in the IT network. In fact, these types of nodes are mostly general-purpose PC where specific software is installed on one of the most common operating systems (differently from the nodes in the OT networks that could have more specific, and sometimes vendor specific, operating systems).

In such a situation, we consider a vulnerability affecting the IT node in the scenario. This vulnerability allows an attacker able to gain remote access to escalate its privilege and to gain all the capabilities that the node has. The analysis of this situation highlights a possible attack that starts from the guest network and can reach the OT network and one of its critical resources, as shown in Figure 33.

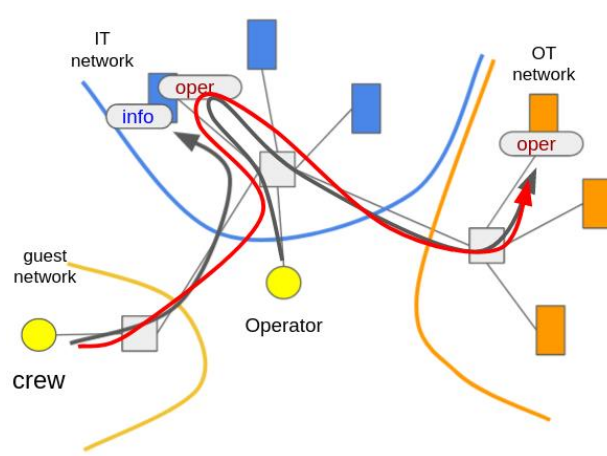


Figure 33 Attack scenario in S2.

In the second scenario (S2), we consider the ship at a port. In this case, the overall system is not more the single ship, but is a composition of all the autonomous systems, as depicted in Figure 34.

maintains its internal sub-systems, is connected in several ways (both physical and cyber) with the Port system, and other ships at the same port. It is worth noting that, for generality, we have depicted all these new systems with the same internal structure, where we have a separation of networks (IT/OT/guest).

The analysis approach in this case must combine the separated system models into a larger overall model and must consider and describe all the new interactions that are spawned for this situation. For instance, Ship 1 must connect to the IT network of the Port to access the process of managing containers at the port (green line in Fig.3). In this case, both IT and OT networks are involved. Moreover, also Ship 2 as to connect in the same way with the port (yellow line), and some paths can share nodes with the connections of Ship 1. Connections between the two ships are also possible, as represented by the red line. It is worth noting that

this could represent an indirect connection, where both ships systems are connected to the same node at the port to access shared information.

In this complex scenario, it is clearly true that the different systems involved are owned by different entities that may not be willing to share the low-level details of their controlled resources. But a high-level description of the services and their foreseen interaction must be available and the SYSVER asset is used proficiently in this case. In fact, the security policies in this case could be kept separated between the systems and the asset can run multiple times for each organization involved, sharing also the very high-level information about the system components.

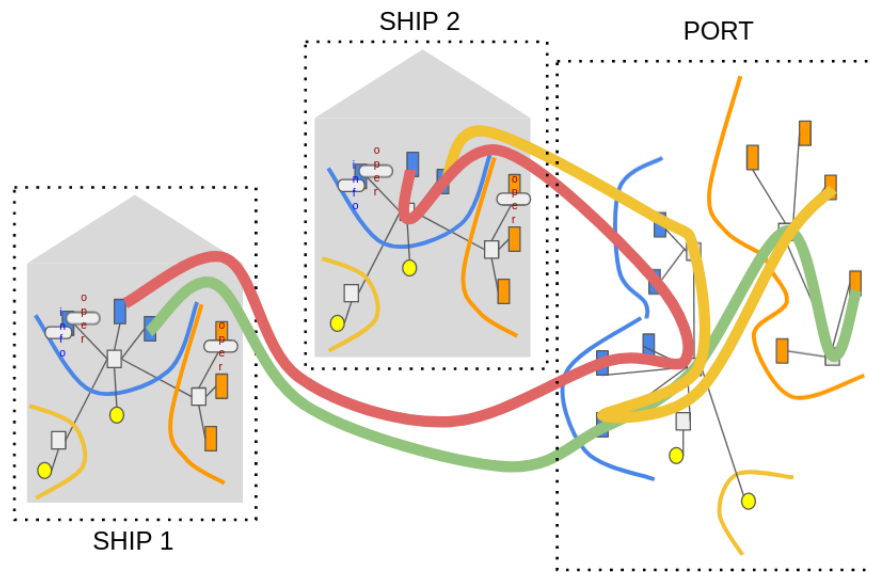


Figure 34 Complex S2 scenario including port and other vessels.

5.3.2 Asset description

The SYSVER asset architecture is shown in Figure 35. The main elements are the “inputs”, the “engine”, the “outputs”, and, at the top of Figure 35, the possible interaction with the asset.

On the “input” side, we see the “System model” that is the combination of several elements:

- the physical/logical resources with their location
- the network, with a broad meaning including both communication network (with firewalls, routers, wifi and wired links, ...) and physical network in the sense of physical environments and their connections (rooms, warehouses, doors with access credentials required, and so on)
- the services defined both at their services components level and at their interaction level. All these elements are described also by their physical/logical location and their access requirements.

The “Agents model”, instead, describes which are the possible users that interact with the system elements and describes their location and their capabilities (credentials, tokens, and so on).

The “Policy model” defines what the agents of the system can or cannot do, by defining the roles and privileges for all the agents involved.

The “Vulnerability model” describes the potential vulnerabilities that can affect elements of the system, and in particular describes the effects of a successfully exploited vulnerability.

At the “engine level”, instead, we have two main components, the “Prolog reasoning engine”, that combines all the information coming from the inputs leveraging a set of “rules” (the “r” database in the figure) that define how the system can evolve depending on the agents’ actions.

The “comparator” evaluates the results of the reasoning engine, and compare them with the Policy model so as to identify *anomalies* between these two inputs, which can correspond to violations of the policy.

The “outputs” of the tool are mainly of two types:

- an Automaton for each of the agents considered
- a high-level report highlighting the overall satisfaction of the security policy

All the single Automaton are then stored in the asset database and can be also further elaborated to highlight specific paths of interest in the state diagram (for instance “all the paths that lead to a specific resource access”).

Finally, the SYSVER asset includes an API that can be accessed remotely, directly (CLI) or through a web interface (that is currently under development with the latest features).

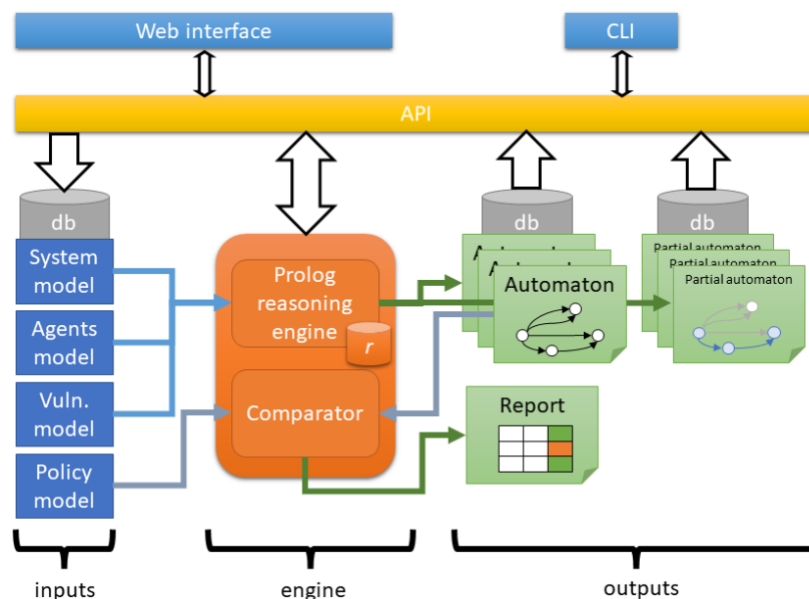


Figure 35 SYSVER architecture.

5.3.3 Asset limitations and future work

Current limitations of the SYSVER assets are two, related to the inputs, and to the reasoning engine with its the outputs:

- regarding the inputs, at the moment, the process of creating the system model from the real system is complex and mostly based on human manual processes. This problem derives also from the heterogeneity of the systems that we consider so that no unique description model exists and is shared. Future work in this direction will further investigate how to collect model elements from data that can be automatically gathered, such as captured network communication or high-level system diagrams.
- regarding the reasoning engine, there are two main limitations: 1) the system is efficient if we assume the monotonicity of agents’ capabilities, but if we discard this assumption this could lead to intractable problems. However, this is usually not a critical issue as it is sound to assume that a user, in particular a malicious one, would not intentionally discard some of his capabilities at some point in time. The second limitation is the difficulty in finding possible solutions to policy violations found by the analysis as any change in the system could fix one problem but could also break some requirements. Future work in this

direction includes using automatic resolution tools (such as Z3 [41]) to ease the construction of possible valid solutions, that can help in selecting fixes to the system.

6 Execution

6.1 Adaptive Incident Reporting

One of the steps of any incident management and response process is to report the incident and not only internally but also with external entities. Security incidents detected, for example in a maritime transport context or in a financial institution, can require the need of mandatory incident reporting to different Supervisory Authorities at different levels (industry, European or national to be compliant with a disparity of applicable regulations. And the lack of harmonization in the procedures leads to the need of producing incident reports adapted to each of the different reporting procedures, templates and formats defined in the different regulatory frameworks.

After analyzing existing open-source incident management tools, we have determined there is a gap related to this mandatory incident reporting capability. In this section we describe how the asset AIRE (Atos Incident Reporting Engine) can help in the incident reporting process, enforcing a predefined and configurable incident reporting workflow and supporting the generation of the required mandatory reports adapted to different templates. This asset has been integrated with one of the existing open-source incident management tools analyzed to complement its provided incident management functionalities with this adaptive incident reporting support.

6.1.1 Application scenario

We have considered the three scenarios already defined for adaptive authentication in section 5.1.2 to analyze in each of them if it would be required to notify the security incident to the competent authorities according to some of the regulatory frameworks that can apply to the maritime scenario.

NIS Directive¹⁵ (adopted by each European Member State at the latest by May 2018 to improve cybersecurity throughout the Union) introduces mandatory incident notification obligations for Operators of Essential Service (OES). Art. 14 (3) of the Directive indicates that “*operators of essential services notify, without undue delay, the competent authority or the CSIRT of incidents having a significant impact on the continuity of the essential services they provide*”.

According to Article 5(2) of the Directive, the criteria for the identification of the operators of essential services are the following ones: “(i) *The entity provides a service which is essential for the maintenance of critical societal and/or economic activities. (ii) The provision of that service depends on network and information systems. (iii) An incident would have significant disruptive effects on the provision of that service.*”. Article 4(4) of the Directive states that an OES is a “*public or private entity of a type referred to in Annex II*” that meets above criteria. In the sector “Transport” and sub-sector “Water transport”, it is included the following type of entities: “(i) *Inland, sea and coastal passenger and freight water transport companies; (ii) Managing bodies of ports including their port facilities; (iii) Operators of vessel traffic services.*”

In summary, according to that directive, we can consider the main stakeholders involved in the maritime scenario (cargo owners, shipping lines, marine services providers, logistics providers or infrastructure providers) as operators of essential services and consequently, any attack “*affecting the availability, authenticity, integrity or confidentiality of networks and information systems*” (as it is defined a NISD incident in the ENISA reference document [90]) with significant impact on the continuity of the service

¹⁵ <https://eur-lex.europa.eu/legal-content/en/TXT/HTML/?uri=CELEX:32016L1148>

should be reported to the national CSIRT (Computer Security Incident Response Team) or to the national NIS Authorities.

Additionally, under Article 33 of the EU General Data Protection Regulation (GDPR)¹⁶, in the case of a personal data breach (defined in the article 4(12) as “*a breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorized disclosure of, or access to, personal data transmitted, stored or otherwise processed*”), the data controller shall “*without undue delay and, where feasible, not later than 72 hours after having become aware of it, notify the personal data breach to the supervisory authority competent*”, unless the personal data breach is unlikely to result in a risk to the rights and freedoms of natural persons. Consequently, if the attack suffered by a European company in the maritime scenario can lead to a personal data breach, it should also be reported according to GDPR procedures.

In both cases, the directives only define the general criteria to be considered and the minimum set of information to be provided, but specific mandatory incident reporting procedures, deadlines and templates are defined at national level.

- **Scenario 1:** Although GDPR does not include specific considerations about passwords, since this is a potential access to personal data, if the use of unsafe network topology results in the execution of an AmosConnect attack and it is considered it can result in a risk to the rights and freedoms of the persons, this attack should be reported by the data controller to the competent authorities. However, this attack probably would not affect the continuity of an essential service so it would not be necessary to generate the report for NIS authorities.
- **Scenario 2:** Any attempt of deviation of a cargo ship from its original course, mainly in a situation where the ship is on sea with other ships at proximity, could be considered an incident with a significant impact for the maritime transport essential service operator and it would require to be notified to the NIS competent authorities.
- **Scenario 3:** A security incident in the authentication process performed by the vessels crew could lead to a personal data breach (since unauthorized persons could have access to sensitive information about the passengers) but also could affect the continuity of the maritime transport essential service itself e.g., in case of taking control over a ship. Those incidents should be reported according to both directives, NIS and GDPR, to different authorities.

In any of those scenarios, once a security incident has been detected, the incident reporting procedure will start, and the different reports required will be generated. Each report will include information about the incident but following the format (e.g., Excel, PDF, Word) and template defined for the specific directive applicable and according to the implementation of that directive at national level. The role of the asset AIRE in those scenarios would be:

- To define and enforce, through the creation and assignment of tasks to roles, the incident reporting process (from the incident data collection to the actual release of the mandatory reports to the Supervisory Authorities) that must be followed when a cyber incident is detected by the port facility, who is the main responsible of the security (according to the ISPS (International Ship and Port Facility Security) code). This will be done through the automatic creation and assignment of tasks to the different roles involved in the incident reporting process using the Incident Management & Response open-source tool TheHive. We can assume the different actors involved in the scenarios are analogous to the ones defined in D5.1 (section 6.3.1 and 6.3.2) [11] for the financial sector: (i) the Incident Management Team (the responsible person or group in the port facility affected by the security incident in charge of analyzing it, collect the required information and open the incident reporting procedure), (ii) the Incident Classification Team (responsible for classifying the incidents opened and determine their severity); (iii) the Incident Reporting Team (in charge of monitoring the evolution of the incidents, preparing the

¹⁶ <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679&qid=1620804229288>

reports and carrying out the actual reporting to the competent authorities); (iv) the Controller (responsible for performing the managerial judgement and giving the authorization for the reporting); and (v) the EU/National Supervisory/Competent Authorities.

If the deadlines established by the regulations to proceed with the reporting arrive and the reports have not been generated and released yet, an escalation procedure will be triggered, and the delay will be notified to the incident contact user defined.

- To generate the reports required to notify the security incident to the competent authorities adapted to the different templates defined by NIS directive and GDPR.

Although the asset AIRE was initially designed to be used by financial institutions in the context of the “Incident Reporting in the financial sector” demonstrator in Task 5.4, it can be used by any company or organization that needs to report security incidents to the competent authorities. In the case of the maritime scenario, this would apply to any operator of essential services in the water transport sector which need to report incidents having a significant impact to be compliant with the NIS Directive, and to any European company suffering a data breach according to the GDPR. These two regulations are already considered in the demonstrator for the financial sector. On the other hand, we assume the companies in the maritime scenario are using the open-source Incident Response Tool TheHive and they have the same roles/teams defined in the financial incident reporting demonstrator. Consequently, the mandatory incident reporting workflow that would be used in the maritime scenario would be the same used in the demonstrator of the financial sector, so it is not necessary to modify the BPMN file used by the asset. We just need to configure the timers associated to each regulation (NIS/GDPR) and the templates to be used for the maritime scenario depending on the nationality of the companies suffering the security incidents, since they are defined at national level.

The asset limitations after adaptive security is adopted in the scenario are the same indicated in the asset description (see section 6.1.4). It is important to highlight that the AIRE asset is not an adaptive SOAR platform for the maritime scenario but a tool to support the companies in the mandatory incident reporting process they need to follow to be compliant with the different regulatory frameworks. In the generation of mandatory incident reports adapted to the different templates and formats and in the enforcement of an incident reporting workflow with notifications adapted to the different deadlines established.

6.1.2 Asset description

The asset AIRE (Atos Incident Reporting Engine) is a tool to support companies or entities in the mandatory security incident reporting process. In particular, the aim of this asset is to address the need to report security incidents adapted to different procedures/methods depending on the applicable regulatory bodies.

The asset AIRE has a modular design composed of different services so it can be more easily adapted or extended to deal with additional regulations or potential changes in the existing ones. This asset was designed to provide some high priority mandatory functional requirements included in CyberSec4Europe in the “Incident Reporting in the Financial Sector” demonstrator. In particular, it covers the following requirements defined in D5.4 [11]: IR-F08 (“*It must request the authorization of the FI operator (Controller) to proceed with the reporting.*”), IR-F09 (“*It must produce the appropriate template in the appropriate format to be sent to the Competent Authority.*”), IR-F13 (“*It must enforce a workflow to be used for reporting purposes*”). Two Springboot microservices¹⁷, the **aire-reports-generators** and the **aire-workflow-enforcement**, compose the AIRE engine to offer those functionalities. Another microservice, the **aire-thehive-plugin**, is the intermediary between the engine and the open-source Security Incident

¹⁷ <https://spring.io/microservices>

Response Platform TheHive¹⁸. This has been the incident management and response tool selected after analyzing and comparing different solutions available in the market and considering the requirements of the demonstrator where the asset was going to be used. Finally, the architecture is completed with a database (that we have called Incident Register Database, where all the information about security incidents is stored) and a web application **aire-dashboard**, that provides the GUI to configure and interact with the asset. Identity and user access management capabilities are provided through the integration with the open-source tool Keycloak¹⁹. Figure 36 summarizes the AIRE asset architecture.

The **aire-workflow-enforcement service** is in charge of enforcing an incident reporting workflow for financial institutions defined through a Business Process Model and Notation (BPMN) file, which tries to cover in a common workflow all the potential phases that can appear in the process of incident reporting and support that each regulatory framework considered can have different deadlines.

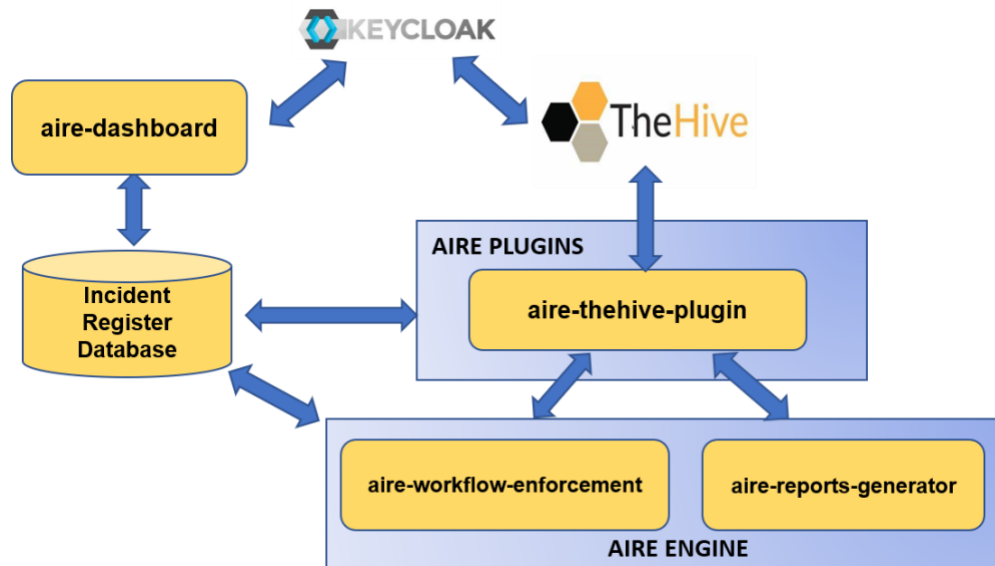


Figure 36 The AIRE architecture.

Consequently, once a new security incident arrives (i.e. it is registered in TheHive and a notification is sent to the service), the responsibility of this service is to create and assign the different tasks that have to be followed by the organization to finally generate the necessary reports that must be sent to the Competent Authorities to be compliant with their mandatory reporting procedures. Depending on the regulation, it is mandatory to send just a report (e.g., under GDPR) or several with specific deadlines. For example, in case of significant institutions the European Central Bank regulatory framework requires the entities to send a first report before 2 hours of classify the security incident as significant, a second report before 10 days from the first report and a final report before 20 days from the second report. Through the AIRE-dashboard, the user can configure the number of reports required for a specific regulation and define different timers associated with them. These timers represent the deadlines associated to the different reporting phases defined for each regulatory framework. The timers need to be related to a specific stage in the incident reporting workflow as defined in the BPMN file. For example, in the case of significant institutions where a first report will need to be sent before 2 hours of an incident classified as significant, it will be defined a

¹⁸ <http://thehive-project.org/>

¹⁹ <https://www.keycloak.org/>

timer with a duration of 2 hours, the report phase where it is triggered (first report), and the workflow stage associated (once classification has been done and it has been confirmed to start data conversion).

To adapt the incident reporting workflow to these specific requirements of each regulation enabled in the asset, the aire-workflow-enforcement integrates a **BPMN parse listener**. Current version of AIRE asset includes an implementation of a BPMN parse listener associated to the BPMN start events, user tasks, service tasks and receive tasks. When the workflow arrives at these stages, it is triggered an escalation procedure in case some timer has been assigned to that stage for some of the active regulations. In particular, the current version of the asset sends a notification to the incident contact user in case a mandatory report has not been sent to the corresponding Supervisory Authorities in the deadline defined by a specific regulation. This escalation procedure is also defined through a BPMN file so it could be adapted in the future to support specific procedures defined by a company.

Figure 37 and Figure 38 show the BPMN files for the Financial Institutions Incident Reporting Workflow and for the Financial Institutions Escalation Procedure followed when some of the regulation milestones arrive.

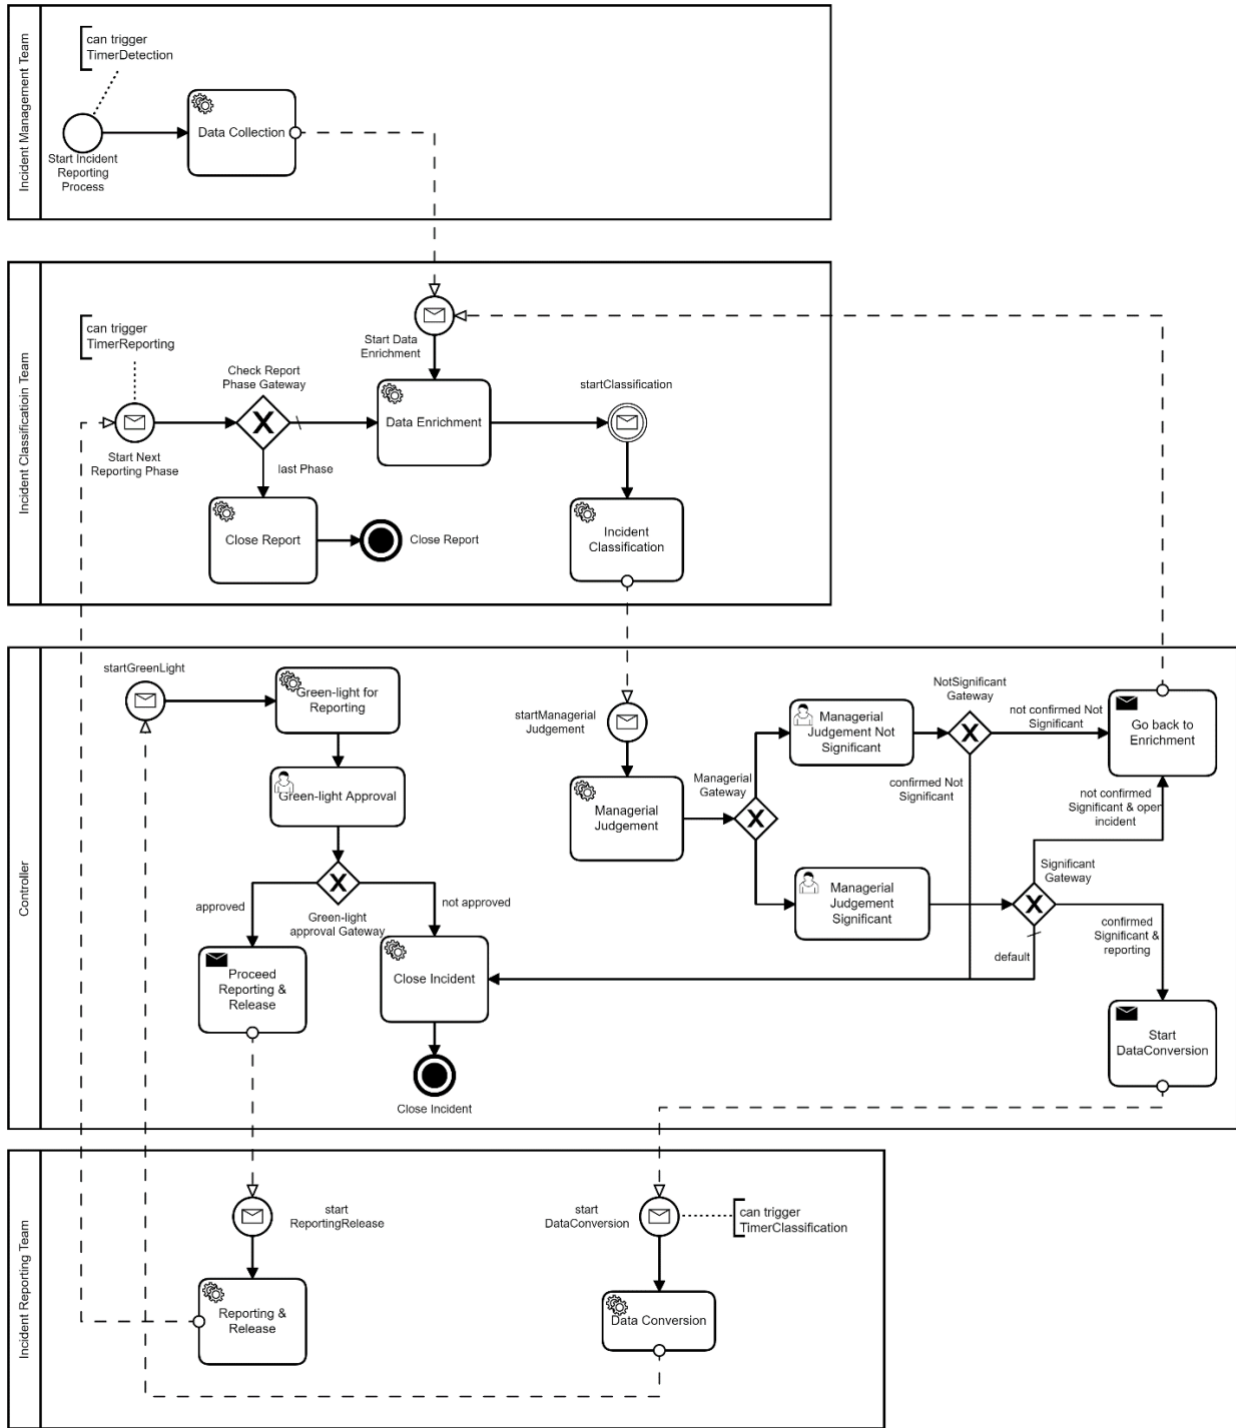


Figure 37 Incident reporting BPMN.

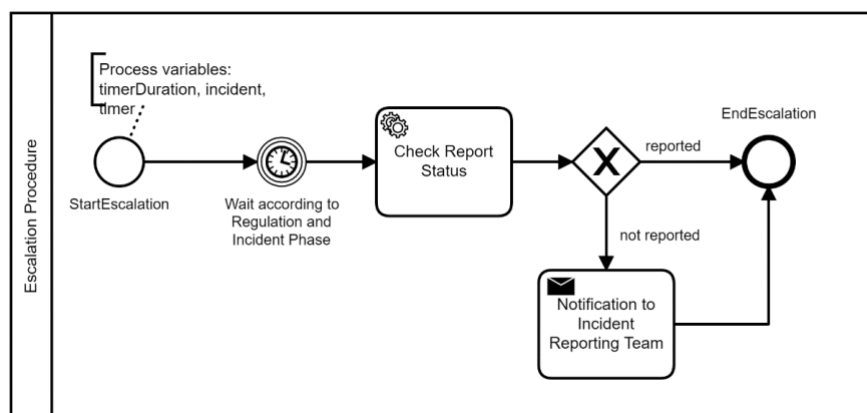


Figure 38 Incident reporting escalation procedure BPMN.

For each of the stages in the BPMN process, the tasks that need to be created in TheHive, the role permissions and the tags and actions associated can be adapted to each situation and defined in a JSON configuration file. They are registered in the database and are configurable through the asset dashboard.

The interaction with the user with Controller role through questionnaires to confirm the classification of the security event and if the reporting process must continue or not (this is what it is known as **managerial judgement**) is also integrated in the incident reporting working using User Tasks (in BPMN terminology). In particular, the current version of the AIRE-workflow-enforcement service includes two different managerial judgement forms as it is described in the document of validation of the incident reporting demonstrator for the financial sector [91]:

- Confirmation of the classification suggested by the platform about the impact of the incident (as Significant or not Significant) and the supervisory authorities that should be notified of the incident based on the criteria and regulations configured in the platform.
- Confirmation for proceeding with the submission of the reports once they have been generated by the platform and reviewed by the Incident Reporting Team.

Through REST APIs, the service receives in real-time requests of new security incidents to be reported, notifications of tasks completed or requests to complete a workflow process. Depending on the user role and the workflow stage, the user will be able of providing information about a security incident, proceed to classify it, confirm the classification (perform the managerial judgement) or carry out the reporting process to the competent authorities selected. The functionality provided by the AIRE-workflow-enforcement service is offered through the creation and assignment of the tasks to the different roles in the open-source tool TheHive related to a new incident registered. The service **AIRE-thehive-plugin** acts as intermediate layer with the AIRE-workflow-enforcement service invoking the REST APIs through JSON over HTTP exposed by TheHive. In this way, AIRE-workflow-enforcement can be more generic and in the future the

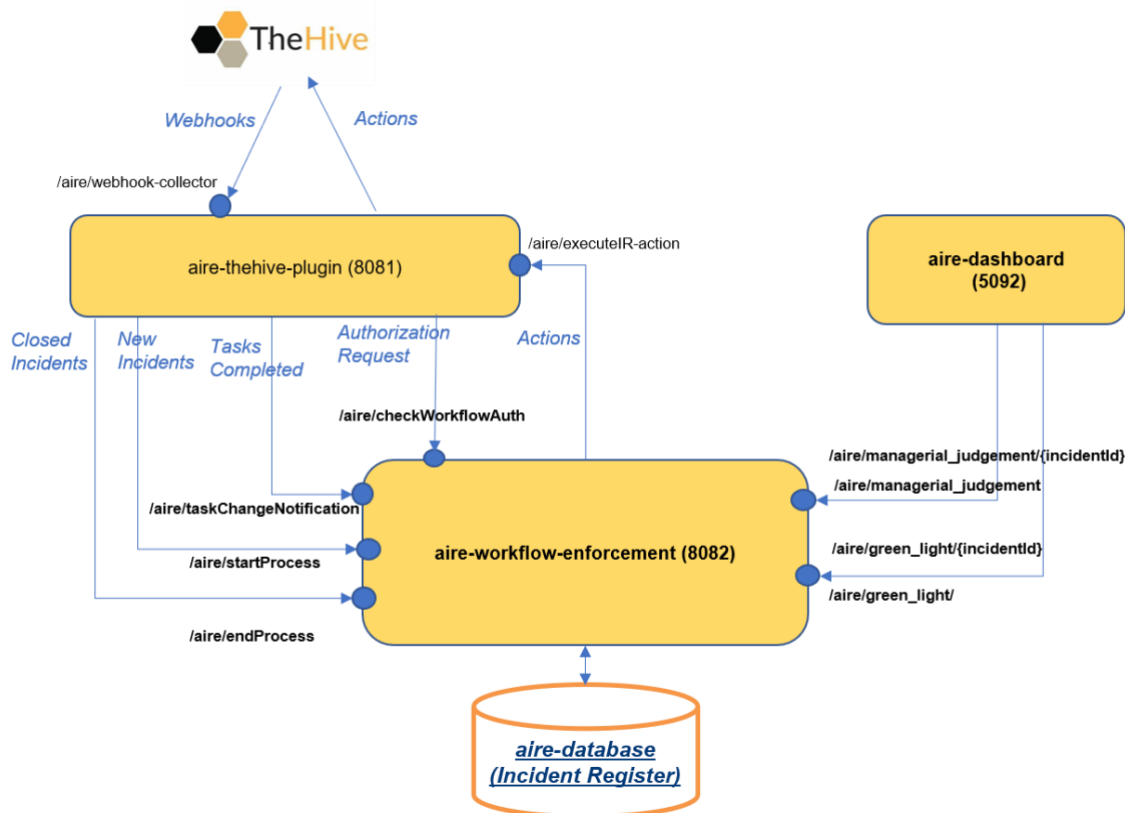


Figure 39 AIRE-workflow-enforcement APIs interaction diagram.

asset could be integrated through other plugins with other incident management tools. Figure 39 and Table 4 summarize the API offered by the AIRE-workflow-enforcement service, which are also described:

HTTP Method	URI	Description
POST	/aire/startProcess	Start AIRE workflow enforcement process when a new incident is registered.
POST	/aire/endProcess	End AIRE workflow enforcement process when a registered incident is closed.
POST	/aire/taskChangeNotification	Notify to trigger the next step in the Incident Reporting Workflow.
POST	/aire/checkWorkflowAuth	Receive a notification to check if a user has permissions to execute an action on an incident in the current workflow stage.
GET	/aire/managerial_judgement/{incidentId}	Get managerial judgement form for an incident.
POST	/aire/managerial_judgement	Submit managerial judgement.
GET	/aire/green_light/{incidentId}	Get green-light form for reporting an incident.
POST	/aire/green_light	Submit green-light managerial judgement

Table 4 AIRE-workflow-enforcement APIs.

The AIRE-reports-generator service is responsible for adapting the information registered about a security incident to the different report templates required by the different regulatory frameworks that apply to a specific company or entity. The AIRE-reports-generator service accesses the information stored about the incident to fill in the different templates and generate the output report files that need to be sent to the Competent Authorities. The Apache POI library for Microsoft Documents²⁰ has been used to work with the Excel and Word files and the Apache PDFBox library²¹ for PDF documents. Both are published under the Apache License v2.0.

InformationEntry	Notification Type	Source	Endpoint	DatabaseFields	Condition	TemplateField	TemplateValue	TemplateFieldType
Type of Report - Individual	M1	database	Report	Report_Type_PSD2_id	ReportTypePSD2.id=2	Type_of_report1_checkbox		
Type of Report - Consolidated	M1	database	Report	Report_Type_PSD2_id	ReportTypePSD2.id=1	Type_of_report2_checkbox		
PSP Name	M1	database	Incident	I_Entity_id(FinancialEntity.E_Name)		PSP_name		
PSP Unique Identification Number, if relevant	M1	database	Incident	I_Entity_id(FinancialEntity.E_BIC)		PSP_UID		
PSP Authorization Number	M1	user						
Head of Group, if applicable	M1	user						
Home Country	M1	database	Incident	I_Entity_id(FinancialEntity.E_Country_id)		Home_country		
Country/Countries affected by the incident	M1	user						
Primary Contact Person	M1	database	FinancialEntity	E_Contact1_id(Contact.C_Name, Contact.C_Surname)		Contact1		
Primary Contact Person E-Mail	M1	database	FinancialEntity	E_Contact1_id(Contact.C_Email)		EmailContact1		
Primary Contact Person Telephone	M1	database	FinancialEntity	E_Contact1_id(Contact.C_Phone)		PhoneContact1		
Secondary Contact Person	M1	database	FinancialEntity	E_Contact2_id(Contact.C_Name, Contact.C_Surname)		Contact2		
Secondary Contact Person E-Mail	M1	database	FinancialEntity	E_Contact2_id(Contact.C_Email)		EmailContact2		
Secondary Contact Person Telephone	M1	database	FinancialEntity	E_Contact2_id(Contact.C_Phone)		PhoneContact2		

Figure 40 AIRE-reports-generator template mapping file example.

The mapping between the location in the database of the information about the incidents and organizations registered in the asset and the fields in the final templates where they need to be included is configured for each template in an Excel file. This mapping Excel file will be loaded together with the associated template through the asset dashboard. In this way, the generation of the reports can be adapted easily to changes in the templates established by the regulations without modifying the service. Figure 40 shows an extract of this excel mapping file.

Table 5 summarizes the API offered by the AIRE-reports-generator service:

HTTP Method	URI	Description
GET	/aire/generateReports/{incidentId}	Generate report templates for a specific incident.

Table 5 AIRE-reports-generator service.

To make the AIRE engine functionalities independent from the incident management and response tool used by the organizations (in our case, the open-source tool TheHive), it has been included an intermediary layer through the service **AIRE-thehive-plugin**. The main functionalities provided by this plugin are the following ones:

- A webhook listener receives notifications when some action is performed by the users in TheHive²² (for example, when a case is created, a task is closed, or a responder is executed).

²⁰ <https://poi.apache.org/>

²¹ <https://pdfbox.apache.org/>

²² More information about TheHive webhooks can be found at: [TheHive Webhooks](#)

- An endpoint that will be used by the services included in the AIRE engine to invoke actions in TheHive using the REST API²³ provided by this tool, such as the creation and assignment of tasks or assign a Tag to a case.
- An endpoint to be invoked from the responders included in TheHive²⁴ to check if the user is authorized to execute that action (e.g., generate a report or perform the event classification) according to the incident reporting workflow in course for a specific incident.
- An endpoint that will be invoked from a Responder in TheHive to generate a report associated with an incident (using the id of the TheHive case).

Figure 41 summarizes the API offered by the AIRE-thehive-plugin service.

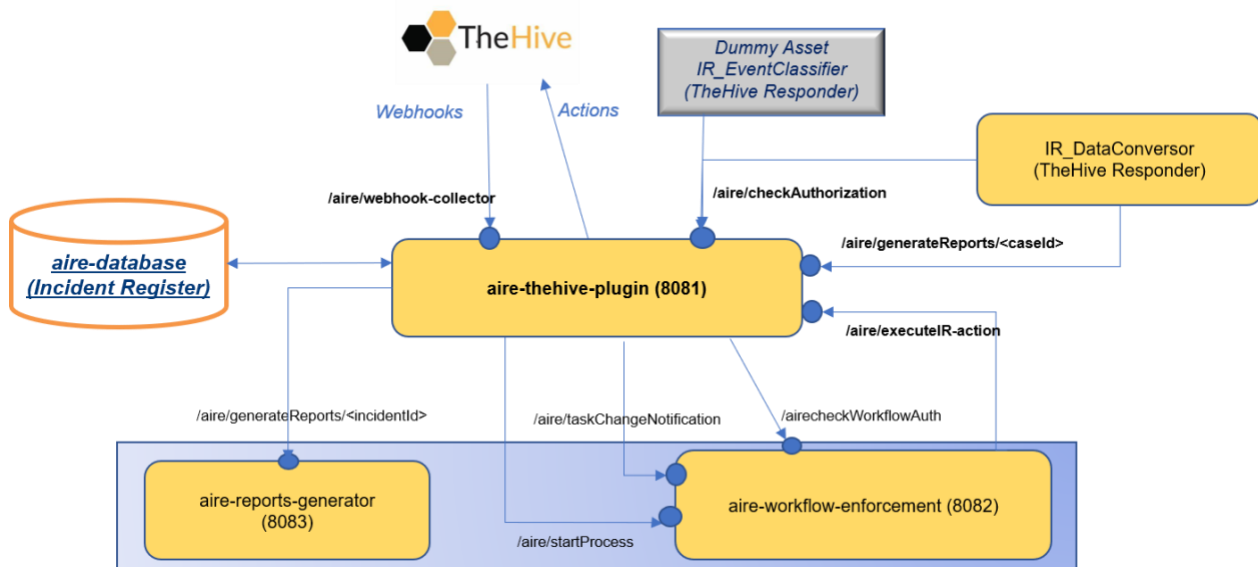


Figure 41 AIRE-thehive-plugin APIs interaction diagram.

In summary, the services included in the Atos Incident Reporting Engine asset provide two functionalities, incident reporting workflow enforcement and generation of adapted incident reports, related to the mandatory procedure of reporting incidents to the Supervisory Authorities we have not found in the incident management and response tools available in the market. And although current version has been integrated with a specific open-source tool (TheHive) and implements a specific BPMN file defined to harmonize a subset the regulations affecting the financial sector, it has been designed in such a way that can be easily extended or adapted to different regulations, changes in the existing ones or applied to different sectors (as explained in section 6.1.1 for the maritime scenario).

²³ More information about REST API exposed by TheHive to invoke these actions can be found at [TheHive API](#)

²⁴ More information about what the Responders are and how they work can be found at: [TheHive Responders](#)

6.1.3 Outputs

A video demonstration of AIRE assets integrated with the Incident Reporting Platform for the financial sector demonstrator developed in the context CyberSec4Europe project is available at the project's YouTube channel²⁵

6.1.4 Asset limitations and future work

The following points describe the main asset limitations regarding supporting adaptive security:

Generation of adaptive mandatory incident reports: Through the definition of Excel mapping files, we have made possible to adapt the information registered about a security incident to automatically generate reports following different templates (as established by the regulatory frameworks or defined by the own companies). However, this is limited to the information about the incidents registered in the database with the current Incident Register data model we have defined in the context of the CyberSec4Europe demonstrator "Incident reporting in the financial sector". This means that if a new template requires that a field has some information not currently included in the data model, it will not be filled in the template. Current data model is based on the mandatory information required by the following regulations: ECB-SSM, PSD2, NIS directive, eIDAS, GDPR and TARGET2, and according to the templates provided for the validation phase by the financial partners involved in the demonstrator of task 5.4 (BBVA in Spain and Intesa SanPaolo in Italy). Consequently, only extensions 'pdf', 'doc', 'docx', 'xls' and 'xlsx' are supported and the asset has been tested only with their templates.

We have also found some limitations in the open-source Apache POI library used in the implementation of the asset, mainly related to the processing of checkboxes and radio buttons. We have solved some of them adapting the templates (e.g., in Excel templates each checkbox must be linked to a hidden cell where the value will be established) and through additional information included in the mapping file, but for example current version does not support to select a checkbox if it has been included in the Excel as ActiveX Control instead of FormControl, if it is a radio button with an option different from Yes/No, or select a checkbox in a Word document template.

Additionally, date and time formats supported in the current version of the asset are the ones defined in the Oracle SimpleDateFormat and they need to be configured through the aire-dashboard GUI associated to each template. For example, "dd/MM/YYYY" for ECB reports and "dd/MM/YYYY, HH:mm" for PSD2

Adaptive workflow enforcement: The adaptive security performed by the asset AIRE in the enforcement of the incident reporting process that need to be followed by an organization to be compliant with the different directives and regulatory frameworks is mainly focused on the triggering of an escalation procedure (in particular, we have implemented the notification by e-mail of the delay in the mandatory reporting) at different and configurable deadlines from different stages in the incident reporting workflow. However, we have found that some regulations are not too clear or concrete on the specifications, to do the enforcement programmatically, and this is a limitation for the asset. In particular, some directives indicate that a report must be sent with "undue delay" but we need to translate it to a specific duration of a timer (e.g., 4 hours as deadline to notify a delay in that case), or that a report needs to be sent once the situation "back to normality" when the asset needs to associate it to some specific stage in the workflow defined.

Additionally, the asset AIRE, as it has been designed to be integrated with an external Incident Response tool such as TheHive through the creation and assignment of tasks in the incident reporting process, is limited to the execution of a same and unique BPMN file for each security incident. This means that we assume a common incident reporting workflow for all the regulations. Thanks to the integration of BPMN parse listeners, the workflow defined is suitable for the mandatory incident reporting procedures required

²⁵ [AIRE demo video 01022022](#)

by all the regulations currently supported by the asset, but we could find in the future some new regulation or organization that requires to modify it.

Finally, we have also found some limitations associated with the tool TheHive. Mainly, in TheHive it is not possible to force that only users with a specific profile (each profile represents a role in the incident reporting process, e.g., controller or incident reporting team) update an incident through their GUI. This means that the asset changes the assignment of the incident and the tasks to a specific user depending on the workflow stage, but we cannot avoid other users without permission performing some action on that incident. To solve in some way this limitation, since the asset includes a webhook listener through the service AIRE-thehive-plugin, all the actions on an incident performed in TheHive are monitored by the asset and if some forbidden action is detected (e.g., close an incident or a task by a user without the required role), it is notified by e-mail to the incident contact user.

7 GDPR compliance issues in adaptive security

The Guidelines for GDPR compliant user experience asset contains two parts, the first is the guidelines themselves, and the second is a template to help perform a Data Protection Impact Assessment (DPIA). The DPIA template was demonstrated in two unique scenarios in deliverables D3.13 and D3.17, which serve as a use case on how this template should be filled out and be a reference for this particular scenario (and/or elements of it). Here we consider special circumstances surrounding adaptive security and the particularity of the maritime environment. To achieve this, we first need a scenario to understand which personal data is processed, as well as information on how exactly the adaptive nature of the security is impacting this specific data (if there is no impact, then this use case will not be very specific to adaptive security). In this section we discuss the challenges associated with adaptive security and the maritime environment.

7.1 Application scenario

This scenario is about the management of access to different parts of the vessel by legitimate users on a cruise ship. The vessel's crew should be allowed access to critical parts of the vessels and information about the passengers. For example, a biometrics-based authentication (e.g., face, fingerprint or iris recognition) can be ideal in this scenario because it can be automated and requires a little attention from the crew. The passengers would use a more traditional method of authentication (e.g., tokens) as using biometric data for what could be only one journey with the carrier could be construed as not proportional to the purpose of authenticating users.

A multitude of contextual factors (e.g., location, network topology, the sensitivity of accessed information, proximity with other vessels) can affect the security risk and the priority of the requirements that can be relevant during adaptive authentication (e.g., security, usability, and performance).

Adaptive security in the given maritime scenario would include processing personal data (e.g., authentication credentials, which, in the case of the crew, is biometric data). Additionally, a DPIA should be considered as many of the services will use personal data of the passengers and personnel on the vessel, and there is a lot of potential for processing of personal data in ways that require a DPIA:

- involves the use of new technologies;
- involves sensitive personal information (e.g., biometric data);
- can affect a large number of data subjects and can involve monitoring of publicly accessible areas (on the vessel);
- There is also automated decision making that affects the crew and passengers (e.g., adaptively changing access privileges).

Other personal data that could potentially be processed in this scenario and would as such need to be included in a DPIA includes, but is not limited to:

- For crew (processed based on the employment contract)
 - Employee number
 - First and last name
 - Cabin
 - Employment title
 - Beginning of service (when and where the employee has joined the cruise)
 - End of service (when and where the employee will leave the cruise)
 - Position/Role on the vessel (one person can have multiple positions)
 - Position/Role on the vessel in case of emergency (one person can have multiple positions)
 - When an emergency is called, employees have additional/different roles – some become firefighters, some rescuers from water, some have to man lifeboats...

- Biometric data (for authentication and access to parts of the vessel and registering coming and leaving the vessel)
- System credentials (for access to vessel system)
- Network credentials (for authentication to the ship network, access to the internet...)
- Time & location of any access (records of access for any of the credentials)
- Access privileges (dependent on the current role and adaptive security)
- For passengers (processed based on a contract, i.e., performance of the service)
 - Passenger number
 - First and last name
 - Cabin
 - Related passengers
 - e.g., children in a separate cabin
 - Date of birth
 - Potentially for access to an onboard Casino
 - Entry port (the port the passenger has joined the cruise)
 - Destination port (the port the passenger will leave the cruise)
 - Bookings (bought tickets or other bookings for events/dinners etc. on the cruise)
 - Type of passenger (e.g., economic, business, 1st class...)
 - Token (for authentication of access to parts of the vessel)
 - Network credentials (for authentication to the ship network, access to the internet...)
 - Time & location of any access (records of access for any of the credentials).

7.2 Asset description

Guidelines for GDPR Compliant User Experience is a deliverable that was produced as D3.6 in the CyberSec4Europe project. As its name implies, it is a collection of guidelines, best practices and recommendations for achieving GDPR compliance. The guidelines help to construct privacy-compliant governance and management practices. However, here we will focus on a specific section of the deliverable, which was designed to serve as a template for the process of performing a Data Protection Impact Assessment (DPIA). A part of this DPIA template was further refined in D3.16. The template is like a to-do list with guidelines on how to perform specific tasks and some pre-prepared structures to support the user. The content of the assessment can vary depending on different circumstances, and the same structure is not always necessarily the best for everybody. That is why the template is fairly soft on the structure and encourages users to change, expand, and upgrade the given template to suit their own organization requirements and circumstances better.

DPIA template is a combination of a guide and pre-prepared content in the form of table templates that personal data controllers can use to perform the DPIA. This solution aims to be primarily of use to the smaller organizations having problems performing or having questions about the assessment's specific steps by giving them a starting point on which they can build.

DPIA is meant to identify and minimize personal data protection risks by systematically analyzing the processing of personal data. Unlike most other risk analyses, DPIA is concentrated on the prevention of harm to data subjects, individuals, and overall society rather than the risk to the organization itself. A DPIA is a legal requirement under the GDPR when the processing is likely to result in a high risk to natural persons' rights and freedoms. This is an excellent example of a condition set by the GDPR for which it is difficult to instinctively know whether it applies or not because there is no definition for "likely to result in a high risk" and the type of issue the enabler is meant to resolve.

The major elements of the DPIA template are presented in Figure 42. The DPIA template aids with the initial decision on the necessity of performing a DPIA. If the circumstances demand the organization to perform the assessment, then the template describes and provides guidelines for the DPIA steps. Before the

processing of personal data can be implemented in the organization, it is important to also make sure all other GDPR requirements are met, which is the purpose of the more broad GDPR compliant user experience enabler. DPIA template contains all the basic information about the assessment as well as many recommendations and good practices on how to perform it. Next, we briefly describe the DPIA template sections where the users can directly utilize the provided content to perform their own assessment.

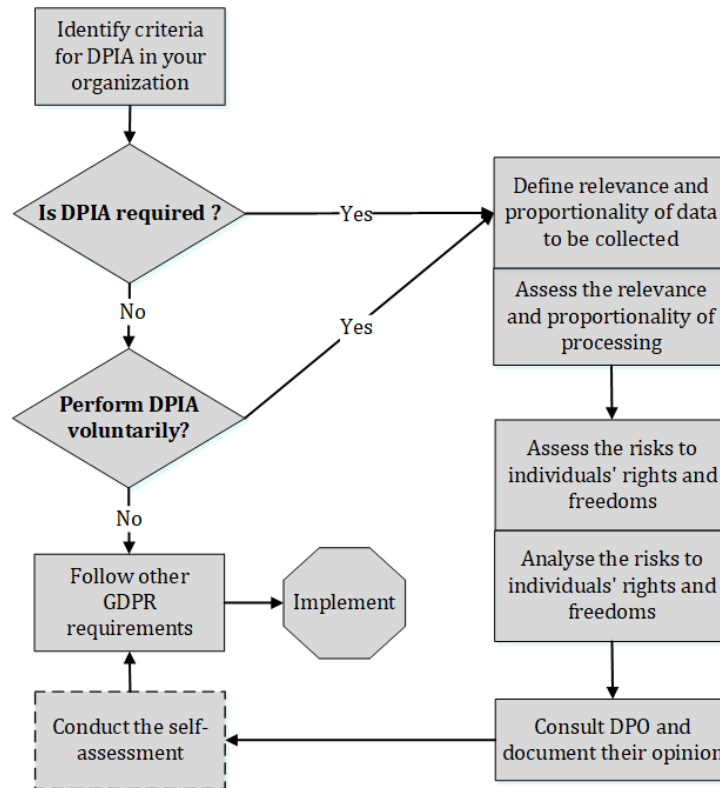


Figure 42 The main steps of the DPIA template.

The first step of performing a DPIA is determining if a DPIA is required. The DPIA template provides a list of criteria based on the GDPR and recommendations given by the Article 29 Working Party and endorsed by the European Data Protection Board. The template requires the users to answer a few questions about the type of processing they intend to perform. Based on the answers, the template enables the user to make an easy judgment about the necessity of the DPIA, given the users' circumstances. This decision process can also be beneficial to show an organization has performed the DPIA voluntarily. Performing DPIA when not necessary can improve the trustworthiness and reputation of an organization, assist in ensuring that the best practices for data security and privacy are being utilized, and helps minimize the organization's liability.

The next major component of the DPIA is focused on risks arising from the processing of personal data. The first step is to establish what type of personal data will be processed, whether data processing is proportional/necessary, for how long it will be stored, and on what legal basis it will be collected. From this information, compliance with the GDPR can be determined. The template helps establish compliance levels based on the collected information. The DPIA template provides instructions on how to measure risk based on the severity and probability of threats. The risk assessment methodology is aligned with the ISO 31000:2018 and ISO 31010:2011 and can be directly used to assign risk levels to all identified threats. The DPIA template provides the users with a template to fill this data into, but more importantly, it already includes a long list of personal data processing threats related to the GDPR requirements. Users can freely

add other threats specific to their organizations, circumstances, used processes, etc. Finally, the risk to individuals' rights and freedoms are evaluated.

GDPR requests that the Data Protection Officer (DPO) provides an opinion on the assessment if one is appointed. The template suggests when a consultation with a DPO might be beneficial. These opinions should be documented.

The final part of the template provides a form for self-assessment. This step of the DPIA is optional and not required by the GDPR. The prepared self-assessment can help organizations track the work they have done and learn from it. Based on their performance, they can improve future work on DPIAs for other processes/services.

There are a few features that set this DPIA template apart from other similar tools. Its simplistic and adaptable design is one of them, but the most novel and for the user critical benefit of the template is the provision of potential risks to include in the risk assessment, with guides on evaluation and in some cases on when a risk is not relevant given specific circumstances. In this deliverable, we look at how a DPIA (regardless of using the DPIA template) would have to change due to the dynamic environment that adaptive security and maritime domain bring into the system. To the best of our knowledge, GDPR and DPIA have previously not been discussed at any significant length in relation to adaptive security.

The DPIA template does not directly support adaptive security or is aimed at any of the specific cybersecurity considerations present in a maritime environment. However, the scenario poses a lot of interesting and unusual questions on how to perform a DPIA in such an environment and with adaptive security, which were not considered in the original GDPR guide and the DPIA template. We would therefore like to expand upon it.

Like we have mentioned, the asset was already demonstrated in deliverables D3.13 and D3.17; therefore, we will focus on the differences one would encounter when performing a DPIA in the adaptive security for the maritime environment. We will discuss some more uncertain and less common situations in performing a DPIA that can be relevant for adaptive security and/or the maritime environment. Here are some of the base considerations:

- As security adapts to circumstances, the risk assessment/impact changes. This brings an issue of changing risk assessments which is an important part of the DPIA.
- How is DPIA different because the vessel sails between countries with possibly different requirements or without GDPR.
- The problem of documenting the purpose of processing personal data, when it changes over time/with circumstances (with changes in adaptive security).
 - An example of such a situation would be to use the data on the age of passengers to build priority evacuation in case of a sinking vessel (this processing is not performed and is unnecessary unless there is a threatening situation to the ship).

7.3 DPIA considerations for adaptive security in a maritime environment

The impact assessment is not mandatory for all controllers and for all processing of personal data, but where it is likely that the type of processing, using new technologies, could, taking into account the nature, scope, circumstances and purposes of the processing, pose a significant risk to the rights and freedoms of individuals. In the composition of the DPIA, it is first necessary to determine whether DPIA is mandatory or only recommended. To determine whether DPIA is mandatory, the controller shall review the relevant legislation applicable to its scope. European Data Protection Board (EDPB) has given recommendations regarding when DPIA is necessary, consisting of multiple criteria. Depending on the criterion, a single or two of them are required for a DPIA to be necessary. In the case of processing biometric data for the purpose of uniquely identifying a natural person, a DPIA has to be carried out if at least one more criterion is also true. At the same time, European Data Protection Board has also stated that these criteria do not have to be

identical across the EU Member States; however, some level of consistency is required. This also brings up another special case in the maritime environment that we will discuss later and is centered around changing legislation as the vessel sails from country to country. In the case of this scenario, we believe a DPIA would be necessary because more than one of the criteria mentioned by the EDPB apply here. A DPIA would be required because biometric data is used, new technologies are used (i.e., adaptive security), automated decisions are being made which use personal data to restrict access to parts of a ship, and systematic monitoring is used on access through the ship. Accountability is an important area and an explicit requirement under the GDPR. As a key accountability tool, a DPIA enables the controller to assess the risks involved as a result of all the different mentioned personal data processing. It is a way of showing that suitable measures have been put in place to address those risks and demonstrate compliance with the GDPR.

First and foremost, the DPIA analysis provides an answer to the question of whether the controller has a legal basis at all for the processing of the personal data of individuals. For the given scenario, it is essential that individuals enter into a contract with the ship's operator (either an employment contract or with payment of the cruise for the guests), which constitutes a legitimate basis for the processing of personal data under article 6. (b) of the GDPR. In this case, processing is lawful if it is " necessary for the performance of a contract to which the data subject is party...". In the given scenario, any other legal basis is not a good fit because it either is not feasible (e.g., legal requirement) or could be problematic (e.g., consent, where users could revoke it mid-cruise).

DPIA is an especially interesting provision of the GDPR for the field of adaptive security, where the level of security generally changes depending on the changes in the current levels of risk present to the system. The need for DPIA is also reliant on the level of risk present in the processing of personal data. It is, therefore, plausible that as adaptive security adapts to the changes in the potential risks and environment, the DPIA should also be revised. With adaptive security, the scope, context, and purposes of processing specific personal data, as well as the underlying risk to that data, could change automatically as security adapts to the changes in the environment. This opens the question of whether a new DPIA should be done every time the system security adapts. Adaptive security is still an obscure field of security and, as such, was, as far as we can tell, not specifically considered, or addressed by any legal body (e.g., EDPB or national Supervisory Authorities) that can give recommendations on how such technologies should comply with the GDPR. DPIA is a process designed to help identify and minimize the data protection risks to the personal data processed in an organization. The DPIA is a type of risk assessment analysis that is very specific to given circumstances, and, as far as we know, there is no automation of this process available (that could automatically adjust the DPIA to the adjustments in the security). To cover such specific circumstances as are present in adaptive security, we recommend including all possible scenarios in the produced DPIA to cover different levels of risks at different levels of security. Defined explanations of when each of these will be in effect should also be documented. Should the DPIA already predict different adjustments, then in our opinion, it is not necessary to conduct a new DPIA every time there is a change in the protection of personal data.

Due to different data privacy legislation in different countries, a ship enters through its voyage, the risk assessment varies, which means it would be necessary to conduct a different DPIA for each of these situations (e.g., legislations). As a ship travels through different countries that also have their own regional legislation, the legislation in individual countries may derogate from the GDPR or be more specific. In accordance with European rules, in the event of abuse of personal data, the legislation of the Member State where the damage occurred, i.e., where the unlawful processing or abuse of personal data has taken place, is applied. Therefore, when making the DPIA, we must also pay attention to the legitimate basis for processing the personal data in accordance with the legislation of each country in which the ship is located at any given time, thus reducing the risk of unlawful processing of personal data. As an example, in accordance with Article 8 of the GDPR, the processing of a child's personal data is lawful when the child is at least 16 years old. Where a child is under the age of 16 years, such processing shall be lawful only if and to the extent that such consent is given or approved by the holder of parental responsibility for the child. Member States may lower this threshold, provided that that age is not less than 13 years. In the process of

travelling between the Member States with different legislations, it would therefore be necessary to conduct a new DPIA for each state with different legislation on the protection of personal data to effectively protect passengers' personal data. To avoid conducting multiple DPIAs, it would be sufficient to assume different options in one DPIA. Specific sections could be made for each country, and anything specific to that country could be included in that section that would only be relevant when in that country. In this way, only the general part (any data protection common to all Member States) and part written under the current location's country would apply.

Both the adaptive security and the travelling of the vessel through different countries cause changes that have to be accounted for in the DPIA. The scenario of adaptive security in the maritime environment is, in this sense very specific, because it doubles the number of causes for a change in a DPIA. Changes in both cases are dynamic, which would require the assessment to be modified or supplemented or upgraded accordingly each time there is a change. To simplify the process, the DPIA can be done in such a way that it already provides for possible adjustments, assesses the protection of personal data in all possible circumstances and shows that the protection of personal data will not be decreased by changes to the security or by changing jurisdictions. It is important for the DPIA to address all possible variations. Processing of personal data under circumstances that were not defined in a DPIA could be considered a breach of the GDPR and the legislation of the country in which the ship is located. If the DPIA can predict all possible combinations and appropriately address the impacts processing would have on the data owners, then one DPIA should suffice to cover all possible variations in the processing of personal data.

Processing of special categories of personal data is generally considered to cause high risk to the user and consequently can require a DPIA. In the presented scenario, the biometric data of employees would be used to gain access to the ship's compartments (i.e., open doors). Biometric device manufacturers often claim that users' privacy is guaranteed because it is not possible to restore data from a biometric template (a format of storing unique traits of an individual's biometric data, for example, a fingerprint). Although this is true, the user's privacy is still not guaranteed since both the fingerprint pattern and its digital pattern are unique identifiers, thus replacing the identity of the individual. The question of reversing the algorithm and the reconstruction of the original data is irrelevant. Key issues from an individual's privacy point of view relate to the use, connectivity, and security of such an identifier. Whenever a latent fingerprint (e.g., from a glass) is obtained, it can be used to create a biometric template with the same method. Using the obtained value, the user can be linked to the original biometric template and identified. Regardless of the format, manner or other changes, that unique bond with the person always remains, even if the amount of detail in the transformation process may decrease. As such, any legal requirements that apply to biometric data also apply to the digital record of those characteristics, which are created based on unique characteristics (no matter how many times and how that record is subsequently transformed).

In this scenario, the crew on board use their fingerprint to record their arrivals and departures from the vessel and to gain access (based on their access privileges) into the ship's closed compartments. Guests on the cruise use smart cards for the same purposes (but have more limited access to the sections of the ship). This allows for tracking of users and is as such considered to be regular and systematic monitoring (WP29 has given their interpretation on when monitoring is regular and systematic [92]), which is one of the reasons a DPIA is required.

Adaptive security can automatically change access privileges for crew and guests based on events affecting the security requirements (e.g., pirates, fire, bad weather, open ocean vs in port, etc.). This brings us to the use of automated decision-making, which is also regulated in the GDPR. Automated decision making is relevant in our scenario because adaptive security could automatically change access to certain areas of the vessels based on the post of a crew member (depending on their role in a given situation – e.g., in a case of a fire, some crew become firefighters while others might be tasked with evacuating the guests etc.) and to the guests (e.g., restrict access to damaged parts of the vessel). The GDPR states that the data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling,

which produces legal effects concerning the data subject or similarly significantly affects them²⁶. The general regulation does not apply to all automated decisions but only to those which have a significant impact on the individual. The European Commission clarified that a decision produces legal effects when the data subject's legal status or their legal rights are impacted or if the processing can significantly affect them by influencing their circumstances, behaviour or choices [93]. Generally, an adaptive security system should not have such significant consequences for the data subject, especially if it is made sure the system cannot prevent any possible contractual obligations the data subject and the data controller had to each other. Additionally, GDPR allows automated decision-making when necessary for entering into, or performance of, a contract between the data subject and a data controller. But GDPR²⁷ also explicitly limits this clause when special categories of personal data are used. Automated decisions involving biometric data (in the case of the crew) could therefore only be performed when "processing is necessary for the purposes of carrying out the obligations and exercising specific rights of the controller or of the data subject in the field of employment"²⁸. However, this section is subject to Union or Member State law. This would therefore need to be considered in the DPIA for each of the Member States on the route of the voyage, or any adaptive changes that affect people on board would have to be confirmed by a person (e.g., officer in charge) before they go into effect. Possible measures to address the data protection risks involved with the automatic decision making could include: informing the data subject about the existence of and the logic involved in the automated, decision-making process, explaining the significance and envisaged consequences of the processing for the data subject, providing the data subject with the means to oppose the decision and allowing the data subject to express their point of view [94].

7.4 Outputs

The guidelines and the DPIA template have been published in D3.6 [95] of the CyberSec4Europe project. A part of the DPIA template (list of potential risks to the rights and freedoms of individuals and how to identify possible non-relevant risks) has been additionally updated and developed in D3.16 [96]. The produced relevant DPIAs are available online^{29,30}. Short videos presenting the two use cases and the resulting DPIAs have also been published on the project Youtube channel³¹.

7.5 Asset limitations and future work

This work is limited to the presented scenario. It can, however, serve as a reference point and a starting point for a discussion on the concerns mentioned. Together with the DPIA template (D3.6), users should have an easier time producing DPIAs for themselves. Additionally, the maritime environment includes specific legislation, constraints, and standards, which are often different between countries. The discussion in this section focuses on the requirements of the GDPR and the performance of a DPIA, and as such does not account for those concerns. In the future, we wish to apply the DPIA template in practice to assist in the preparation of data protection impact assessments for actual businesses.

²⁶ Article 22, Paragraph 1 of the GDPR

²⁷ Article 22, Paragraph 4 of the GDPR

²⁸ Article 9, Paragraph 2(b) of the GDPR

²⁹ [https://cybersec4europe.um.si/Addendum to D3.13 - Student Enrolment Data Protection Impact Assessment.pdf](https://cybersec4europe.um.si/Addendum%20to%20D3.13%20-%20Student%20Enrolment%20Data%20Protection%20Impact%20Assessment.pdf)

³⁰ [https://cybersec4europe.um.si/Addendum to D3.17 – Survey Data Protection Impact Assessment.pdf](https://cybersec4europe.um.si/Addendum%20to%20D3.17%20-%20Survey%20Data%20Protection%20Impact%20Assessment.pdf),

³¹ <https://consent.youtube.com/m?continue=https%3A%2F%2Fwww.youtube.com%2Fchannel%2FUCSAJ78frZjdUTooAC4t6Wuw%3Fcbid%3D1&gl=IE&m=0&pc=yt&uxe=23983171&hl=en&src=1>

8 Conclusion

In this deliverable we presented the results of a survey conducted with security practitioners from industry. The survey was aimed to a) acquire the practitioners' perspective on the application of adaptive security technologies and their level of automation; b) understand the involvement of stakeholders in the execution of adaptive security tasks and c) identify the research challenges that practitioners perceive to be the most critical in adaptive security. From the responses collected from the participants we observed that practitioners have limited familiarity with adaptive security technologies, especially with adaptive access control and adaptive authentication technologies. The activities for which participants considered that the input from engineers was most beneficial were security monitoring, risk assessment, security logging and malware detection. The reason is that security monitoring/logging and risk assessment are activities that can hardly be automated and require to be configured or performed entirely by a security engineer.

The top 3 research challenges that practitioners considered to be the most important are related to 1) the provisioning of assurances in adaptive security solutions; 2) the ranking of protection strategies that can be used to prevent or contain threats; 3) the use of machine learning techniques to evaluate effectiveness of protection strategies in the long-term. Also, the practitioners identified two additional research challenges. Related to the development of adaptive security solutions that are simple to manage and intuitive to use.

In the second part of the deliverable, we use a reference architecture for adaptive security systems to identify the contributions of the assets provided by each partner. We showcase how the joint contributions of each asset can support some of the maritime transport use case scenarios elicited in task T5.5 Maritime Transport. We considered security modelling explicitly, particularly threat modelling. More precisely, we provided an approach to automate threat elicitation and identify evolving threats when the architecture of the system changes. We also supported adaptive authentication to select an authentication method that maximizes satisfaction of the requirements depending on the context and the requirements priorities. We suggested two types of adaptive risk assessment: situation-driven (Section 5.2) and anomaly-driven (Section 5.3). Situation-driven risk assessment identifies situations and pre-compute security risks based on the utilized assets, threats, vulnerabilities, and impacts. While anomaly-driven risk assessment re-evaluates risks depending on anomalies arising from changes in the interdependencies between system components. Moreover, we introduce a novel technique to perform incidents reporting adaptively. Finally, we identified General Data Protection Regulation (GDPR) compliance issues that can arise in adaptive security systems.

The main difficulty that we experienced in this task was to re-think how the assets provided by each partner could be more adaptive depending on varying situations and contexts and, also, how they could also be applied on a cyber-physical system example. In future work, we will improve integrations between assets by exploring synergies between the data-flow-centric threat assessment and the adaptive risk assessment approaches proposed in this task. Also we will explore how to use the situation enforcement framework to enforce adaptive authentication methods. Finally, we will explore how to integrate the approach to execute adaptive incident reporting with the analysis and plan activities of the MAPE-K loop.

References

- [1] C. P. Pfleeger, S. L. Pfleeger, J. Margulies, and Ebscohost, *Security in computing*, Fifth edition. ed. Upper Saddle River, NJ: Prentice Hall (in English), 2015, pp. xxxiii, 910 pages : illustrations (black and white), maps (black and white).
- [2] M. Salehie, L. Pasquale, I. Omoronyia, R. Ali, and B. Nuseibeh, "Requirements-driven adaptive security: Protecting variable assets at runtime.," presented at the *20th IEEE International Requirements Engineering Conference (RE)*, 2012.
- [3] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 1, pp. 41-50, 2003.
- [4] A. Elkhodary and J. Whittle, "A survey of approaches to adaptive application security," in *International Workshop on Software Engineering for Adaptive and Self-Managing Systems (SEAMS'07)*, 2007: IEEE.
- [5] E. Yuan, N. Esfahani, and S. Malek, "A systematic survey of self-protecting software systems," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 8, no. 4, 2014.
- [6] G. Tziakouris, R. Bahsoon, and M. A. Babar, "A Survey on Self-Adaptive Security for Large-scale Open Environments," *ACM Computing Surveys (CSUR)*, vol. 51, no. 5, 2018.
- [7] L. Pasquale, "D3.4 Analysis of Key Research Challenges for Adaptive Security," ed.
- [8] B. Kitchenham and S. L. Pfleeger, "Personal Opinion Surveys," in *Guide to Advanced Empirical Software Engineering*, Springer Ed., 2008 pp. 63-92.
- [9] B. K. a. S. L. Pfleeger, "Principles of survey research: part 5: population and samples," *ACM SIGSOFT Software Engineering Notes*, vol. 27, no. 5, pp. 17-20, 2002.
- [10] "CyberSec4Europe - Industry Survey on Adaptive Security." (accessed).
- [11] A. Sforzin, "CyberSec4Europe D5.1: Requirements Analysis of Demonstration Cases Phase 2," European Commission, 2021.
- [12] I. Stellios, P. Kotzanikolaou, M. Psarakis, C. Alcaraz, and J. Lopez, "A survey of iot-enabled cyberattacks: Assessing attack paths to critical infrastructures and services," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, pp. 3453-3495, 2018.
- [13] M. Balduzzi, K. Wihoit, and A. Pasta, "Hey Captain, where's your ship? Attacking vessel tracking systems for fun and profit," in *11th Annual HITB Security Conference in Asia*, 2013.
- [14] M. Ballano. "AmosConnect: Maritime Communications Security Has Its Flaws." Available: <http://blog.ioactive.com/2017/10/amosconnect-maritime-communications.html> (accessed 2021).
- [15] J. DiRenzo, D. A. Goward, and F. S. Roberts, "The little-known challenge of maritime cyber security," in *2015 6th International Conference on Information, Intelligence, Systems and Applications (IISA)*, 2015: IEEE, pp. 1-5.
- [16] K. Munro. "OSINT from ship satcoms." Available: <https://www.pentestpartners.com/security-blog/osint-from-ship-satcoms/> (accessed 2021).
- [17] P. Beaumont and S. Wolthusen, "Cyber-risks in maritime container ports: An analysis of threats and simulation of impacts," *ISG MSc Information Security thesis series 2017*, 2017.
- [18] E. Yuan, N. Esfahani, and S. Malek, "A systematic survey of self-protecting software systems," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 8, no. 4, pp. 1-41, 2014.
- [19] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41-50, 2003.

- [20] Y. Brun *et al.*, "Engineering self-adaptive systems through feedback loops," in *Software engineering for self-adaptive systems*: Springer, 2009, pp. 48-70.
- [21] S. Maynard, A. Ruighaver, and A. Ahmad, "Stakeholders in security policy development," 2011.
- [22] A. Hita and A. Skarmeta, "D3.12 Common Framework Handbook 2," ed, 2021.
- [23] A. Shostack, *Threat modeling: Designing for security*. John Wiley & Sons, 2014.
- [24] R. Khan, K. McLaughlin, D. Lavery, and S. Sezer, "STRIDE-based threat modeling for cyber-physical systems," in *2017 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, 2017: IEEE, pp. 1-6.
- [25] L. Sion, D. Van Landuyt, K. Yskout, and W. Joosen, "Sparta: Security & privacy architecture through risk-driven threat assessment," in *2018 IEEE International Conference on Software Architecture Companion (ICSA-C)*, 2018: IEEE, pp. 89-92.
- [26] D. Van Landuyt, L. Pasquale, L. Sion, and W. Joosen, "Threat models at run time: the case for reflective and adaptive threat management (NIER track)," in *SEAMS'21: Proceedings of the 16th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2021.
- [27] K. A. A. Bakar and G. R. Haron, "Adaptive authentication: Issues and challenges," in *2013 World Congress on Computer and Information Technology (WCCIT)*, 2013: IEEE, pp. 1-6.
- [28] P. Arias-Cabarcos, C. Krupitzer, and C. Becker, "A survey on Adaptive Authentication," *ACM Computing Surveys (CSUR)*, vol. 52, no. 4, pp. 1-30, 2019.
- [29] H. De Silva, D. C. Wittebron, A. R. Lahiru, K. L. Madumadhavi, L. Rupasinghe, and K. Y. Abeywardena, "AuthDNA: An Adaptive Authentication Service for any Identity Server," in *2019 International Conference on Advancements in Computing (ICAC)*, 2019: IEEE, pp. 369-375.
- [30] E. Hayashi, S. Das, S. Amini, J. Hong, and I. Oakley, "CASA: context-aware scalable authentication," in *Proceedings of the Ninth Symposium on Usable Privacy and Security*, 2013, pp. 1-10.
- [31] M. T. Gebrie and H. Abie, "Risk-based adaptive authentication for Internet of things in smart home eHealth," in *Proceedings of the 11th European Conference on Software Architecture: Companion Proceedings*, 2017, pp. 102-108.
- [32] K. A. A. Bakar and G. R. Haron, "Adaptive authentication based on analysis of user behavior," in *2014 Science and Information Conference*, 2014: IEEE, pp. 601-606.
- [33] S. Eberz, K. B. Rasmussen, V. Lenders, and I. Martinovic, "Evaluating behavioral biometrics for continuous authentication: Challenges and metrics," in *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, 2017, pp. 386-399.
- [34] T. Karanikiotis, M. D. Papamichail, K. C. Chatzidimitriou, N.-C. I. Oikonomou, A. L. Symeonidis, and S. K. Saripalle, "Continuous Implicit Authentication through Touch Traces Modelling," in *2020 IEEE 20th International Conference on Software Quality, Reliability and Security (QRS)*, 2020: IEEE, pp. 111-120.
- [35] J. M. Jorquera Valero *et al.*, "Improving the security and QoE in mobile devices through an intelligent and adaptive continuous authentication system," *Sensors*, vol. 18, no. 11, p. 3769, 2018.
- [36] A. Bucchiarone, R. Kazhamiakin, C. Cappiello, E. Di Nitto, and V. Mazza, "A context-driven adaptation process for service-based applications," in *Proceedings of the 2nd International Workshop on Principles of Engineering Service-Oriented Systems*, 2010, pp. 50-56.
- [37] G. Tamura, N. M. Villegas, H. A. Muller, L. Duchien, and L. Seinturier, "Improving context-awareness in self-adaptation using the DYNAMICO reference model," in *2013 8th International*

- Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*, 2013: IEEE, pp. 153-162.
- [38] L. Kulp, A. Sarcevic, M. Cheng, and R. S. Burd, "Towards Dynamic Checklists: Understanding Contexts of Use and Deriving Requirements for Context-Driven Adaptation," *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 28, no. 2, pp. 1-33, 2021.
- [39] A. Van Lamsweerde, *Requirements engineering: From system goals to UML models to software*. Chichester, UK: John Wiley & Sons, 2009.
- [40] S. Zhou, Z.-Q. Liu, and J. Y. Zhang, "Fuzzy causal networks: general model, inference, and convergence," *IEEE Transactions on Fuzzy Systems*, vol. 14, no. 3, pp. 412-420, 2006.
- [41] L. B. de Moura, Nikolaj, "Z3: An efficient SMT solver," presented at the *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, 2008.
- [42] S. Gupta, A. Buriro, and B. Crispo, "DriverAuth: A risk-based multi-modal biometric-based driver authentication scheme for ride-sharing platforms," *Computers & Security*, vol. 83, pp. 122-139, 2019.
- [43] N. I. Daud, G. R. Haron, and S. S. S. Othman, "Adaptive authentication: Implementing random canvas fingerprinting as user attributes factor," in *2017 IEEE Symposium on Computer Applications & Industrial Electronics (ISCAIE)*, 2017: IEEE, pp. 152-156.
- [44] N. I. Daud, G. R. Haron, and D. Din, "Adaptive Authentication to determine login attempt penalty from multiple input sources," in *2019 IEEE Conference on Application, Information and Network Security (AINS)*, 2019: IEEE, pp. 1-5.
- [45] R. Hulsebosch, M. S. Bargh, G. Lenzini, P. Ebben, and S. M. Iacob, "Context sensitive adaptive authentication," in *European Conference on Smart Sensing and Context*, 2007: Springer, pp. 93-109.
- [46] A. Arfaoui, S. Cherkaoui, A. Kribeche, S. M. Senouci, and M. Hamdi, "Context-Aware Adaptive Authentication and Authorization in Internet of Things," in *ICC 2019-2019 IEEE International Conference on Communications (ICC)*, 2019: IEEE, pp. 1-6.
- [47] O. Riva, C. Qin, K. Strauss, and D. Lymberopoulos, "Progressive authentication: deciding when to authenticate on mobile phones," in *Presented as part of the 21st {USENIX} Security Symposium ({USENIX} Security 12)*, 2012, pp. 301-316.
- [48] K. Sha, Y. Xi, W. Shi, L. Schwiebert, and T. Zhang, "Adaptive privacy-preserving authentication in vehicular networks," in *2006 First International Conference on Communications and Networking in China*, 2006: IEEE, pp. 1-8.
- [49] Y. Xi, K.-W. Sha, W.-S. Shi, L. Schwiebert, and T. Zhang, "Probabilistic adaptive anonymous authentication in vehicular networks," *Journal of Computer Science and Technology*, vol. 23, no. 6, pp. 916-928, 2008.
- [50] S. Sharma and B. Kaushik, "A survey on internet of vehicles: Applications, security issues & solutions," *Vehicular Communications*, vol. 20, p. 100182, 2019.
- [51] A. Hassan, N. Eltayieb, R. Elhabob, and F. Li, "An efficient certificateless user authentication and key exchange protocol for client-server environment," *Journal of Ambient Intelligence and Humanized Computing*, vol. 9, no. 6, pp. 1713-1727, 2018.
- [52] A. Hassan, A. A. Omala, M. Ali, C. Jin, and F. Li, "Identity-based user authenticated key agreement protocol for multi-server environment with anonymity," *Mobile Networks and Applications*, vol. 24, no. 3, pp. 890-902, 2019.

- [53] R. Kainda, I. Flechais, and A. Roscoe, "Security and usability: Analysis and evaluation," in *2010 International Conference on Availability, Reliability and Security*, 2010: IEEE, pp. 275-282.
- [54] J. Nicholson, L. Coventry, and P. Briggs, "Age-related performance issues for PIN and face-based authentication systems," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2013, pp. 323-332.
- [55] S. Ruoti, B. Roberts, and K. Seamons, "Authentication melee: A usability analysis of seven web authentication systems," in *Proceedings of the 24th International Conference on World Wide Web*, 2015, pp. 916-926.
- [56] E. Frøkjær, M. Hertzum, and K. Hornbæk, "Measuring usability: are effectiveness, efficiency, and satisfaction really correlated?," in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, 2000, pp. 345-352.
- [57] A. Wójtowicz and K. Joachimiak, "Model for adaptable context-based biometric authentication for mobile devices," *Personal and Ubiquitous Computing*, vol. 20, no. 2, pp. 195-207, 2016.
- [58] K. Reese, T. Smith, J. Dutson, J. Armknecht, J. Cameron, and K. Seamons, "A usability study of five two-factor authentication methods," in *Fifteenth Symposium on Usable Privacy and Security ({SOUPS} 2019)*, 2019.
- [59] A. Fayad, B. Hammi, and R. Khatoun, "An adaptive authentication and authorization scheme for IoT's gateways: a blockchain based approach," in *2018 Third International Conference on Security of Smart Cities, Industrial Control System and Communications (SSIC)*, 2018: IEEE, pp. 1-7.
- [60] D. Dasgupta, A. Roy, and A. Nag, "Toward the design of adaptive selection strategies for multi-factor authentication," *computers & security*, vol. 63, pp. 85-116, 2016.
- [61] I. You, J. D. Lim, J. N. Kim, H. Ahn, and C. Choi, "Adaptive authentication scheme for mobile devices in proxy MIPv6 networks," *IET Communications*, vol. 10, no. 17, pp. 2319-2327, 2016.
- [62] J. Seifert, A. De Luca, B. Conradi, and H. Hussmann, "Treasurephone: Context-sensitive user data protection on mobile phones," in *International Conference on Pervasive Computing*, 2010: Springer, pp. 130-137.
- [63] D. Goel, E. Kher, S. Joag, V. Mujumdar, M. Griss, and A. K. Dey, "Context-aware authentication framework," in *International Conference on Mobile Computing, Applications, and Services*, 2009: Springer, pp. 26-41.
- [64] A. Forget, S. Chiasson, P. C. Van Oorschot, and R. Biddle, "Improving text passwords through persuasion," in *Proceedings of the 4th symposium on Usable privacy and security*, 2008, pp. 1-12.
- [65] B. Mbarek, M. Ge, and T. Pitner, "Self-adaptive RFID Authentication for Internet of Things," in *International Conference on Advanced Information Networking and Applications*, 2019: Springer, pp. 1094-1105.
- [66] A. Kalamandeen, A. Scannell, E. de Lara, A. Sheth, and A. LaMarca, "Ensemble: cooperative proximity-based authentication," in *Proceedings of the 8th international conference on Mobile systems, applications, and services*, 2010, pp. 331-344.
- [67] A. Primo, V. V. Phoha, R. Kumar, and A. Serwadda, "Context-aware active authentication using smartphone accelerometer measurements," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2014, pp. 98-105.
- [68] R. Ali, F. Dalpiaz, and P. Giorgini, "A goal-based framework for contextual requirements modeling and analysis," *Requirements Engineering*, vol. 15, no. 4, pp. 439-458, 2010.

- [69] K. C. Kang, S. G. Cohen, J. A. Hess, W. E. Novak, and A. S. Peterson, "Feature-oriented domain analysis (FODA) feasibility study," Carnegie-Mellon Univ Pittsburgh Pa Software Engineering Inst, 1990.
- [70] M. Salehie, L. Pasquale, I. Omoronyia, R. Ali, and B. Nuseibeh, "Requirements-driven adaptive security: Protecting variable assets at runtime," in *2012 20th IEEE international requirements engineering conference (RE)*, 2012: IEEE, pp. 111-120.
- [71] A. Hassan, B. Nuseibeh, and L. Pasquale, "Engineering Adaptive Authentication," presented at the *2021 IEEE International Conference on Autonomic Computing and Self-Organizing Systems Companion (ACSOS-C)*, 2021.
- [72] S. Schauer, N. Polemi, and H. Mouratidis, "MITIGATE: a dynamic supply chain cyber risk assessment methodology," *Journal of Transportation Security*, vol. 12, pp. 1–35, 2019.
- [73] R. Laborde, A. Oglaza, F. Barrère, and A. Benzekri, "dynSMAUG: A dynamic security management framework driven by situations," in *2017 1st Cyber Security in Networking Conference (CSNet)*, 2017, pp. 1–8.
- [74] R. Laborde, A. Oglaza, A. S. Wazan, F. Barrère, and A. Benzekri, "A situation-driven framework for dynamic security management," *Annals of Telecommunications*, vol. 74, pp. 185–196, 2019.
- [75] A. Benzekri, R. Laborde, A. Oglaza, D. Rammal, and F. Barrère, "Dynamic security management driven by situations: An Exploratory analysis of logs for the identification of security situations," in *2019 3rd Cyber Security in Networking Conference (CSNet)*, 2019, pp. 66–72.
- [76] S. Papastergiou and N. Polemi, "MITIGATE: a dynamic supply chain cyber risk assessment methodology," in *Smart Trends in Systems, Security and Sustainability*: Springer, 2018, pp. 1–9.
- [77] N. Polemi and P. Kotzanikolaou, "Medusa: a supply chain risk assessment methodology," in *Cyber Security and Privacy Forum*, 2015, pp. 79–90.
- [78] N. Polatidis, M. Pavlidis, and H. Mouratidis, "Cyber-attack path discovery in a dynamic supply chain maritime risk management system," *Computer Standards & Interfaces*, vol. 56, pp. 74–82, 2018.
- [79] E.-M. Kalogeraki, S. Papastergiou, H. Mouratidis, and N. Polemi, "A novel risk assessment methodology for SCADA maritime logistics environments," *Applied Sciences*, vol. 8, p. 1477, 2018.
- [80] S. Rass, S. König, and S. Schauer, "Uncertainty in games: Using probability-distributions as payoffs," in *International Conference on Decision and Game Theory for Security*, 2015, pp. 346–357.
- [81] A. Adi and O. Etzion, "Amit - the situation manager," *The VLDB Journal—The International Journal on Very Large Data Bases*, vol. 13, pp. 177–203, 2004.
- [82] D. Luckham, "The power of events: An introduction to complex event processing in distributed enterprise systems," in *Workshop on Rules and Rule Markup Languages for the Semantic Web*, 2008, p. 3.
- [83] D. W. Chadwick, L. Su, O. Otenko, and R. Laborde, "Coordination between distributed PDPs," in *Seventh IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'06)*, 2006, pp. 10–pp.
- [84] D. W. Chadwick, L. Su, and R. Laborde, "Coordinating access control in grid services," *Concurrency and Computation: Practice and Experience*, vol. 20, pp. 1071–1094, 2008.
- [85] R. A. Zwaan and G. A. Radvansky, "Situation models in language comprehension and memory.," *Psychological bulletin*, vol. 123, p. 162, 1998.

- [86] I. Stelliou, P. Kotzanikolaou, and C. Grigoriadis, "Assessing IoT enabled cyber-physical attack paths against critical systems," *Computers & Security*, vol. 107, p. 102316, 2021.
- [87] C. Grigoriadis, M. Berzovitis, I. Stelliou, and P. Kotzanikolaou, "A Cybersecurity Ontology to Support Risk Information Gathering in Cyber-Physical Systems," in *7th Workshop on the Security of Industrial Control Systems & of Cyber-Physical Systems (CyberICPS 2021)*, 2021.
- [88] S. Chabridon, R. Laborde, T. Desprats, A. Oglaza, P. Marie, and S. M. Marquez, "A survey on addressing privacy together with quality of context for context management in the Internet of Things," *Annals of telecommunications*, vol. 69, pp. 47–62, 2014.
- [89] S. Chabridon, A. Bouzeghoub, A. Ahmed-Nacer, P. Marie, and T. Desprats, "Unified modeling of quality of context and quality of situation for context-aware applications in the internet of things," in *International and Interdisciplinary Conference on Modeling and Using Context*, 2017, pp. 370–374.
- [90] "Reference document on Incident Notification for Operators of Essential Services,"
- [91] A. Sforzin, "D5.3 - Validation Demonstration case Phase 1," European Commission, 2020.
- [92] "Guidelines on Data Protection Officers ('DPOs') (wp243rev.01)." Available: <https://ec.europa.eu/newsroom/article29/items/612048/en> (accessed 2022).
- [93] "Can I be subject to automated individual decision-making, including profiling?" Available: https://ec.europa.eu/info/law/law-topic/data-protection/reform/rights-citizens/my-rights/can-i-be-subject-automated-individual-decision-making-including-profiling_en (accessed 2022).
- [94] "Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679 (wp251rev.01)." Available: <https://ec.europa.eu/newsroom/article29/items/612053/en> (accessed 2022).
- [95] B. Kežmah, "D3.6 Guidelines for GDPR compliant user experience," ed: European Commission.
- [96] L. Outi-Marja, "D3.16 Security Requirements and Risk Conceptualization," ed. European Commission, 2021.