Strategies and Technologies for **Un**ited and **R**esilient Critical **I**nfrastructures and Vital **S**ervices in Pandemic-Stricken **E**urope

D6.5 Cyber-physical resilience tool and training guide V3

Document Identification			
Status	Final	**Due Date**	30/04/2025
Version	1.0	**Submission Date**	30/04/2025

Related WP	WP6	**Document Reference**	D6.5
Related Deliverable(s)	D6.1, D6.2, D6.3, D6.4	**Dissemination Level (*)**	PU
Lead Participant	ATS	**Lead Author**	Susana González Zarzosa
Contributors	ATS, XLB, INS, CAF, PIL, SZ, TS	**Reviewers**	Nikos Avgeros (SQD)
			Fabrice Klein (PMB)

Keywords:

Critical infrastructure, cyber-physical resilience, artificial intelligence, threat intelligence, risk assessment, incident reporting, user manual

This document is issued within the frame and for the purpose of the SUNRISE project. This project has received funding from the European Union's Horizon Europe Programme under Grant Agreement No.101073821. The opinions expressed and arguments employed herein do not necessarily reflect the official views of the European Commission.

The dissemination of this document reflects only the author's view, and the European Commission is not responsible for any use that may be made of the information it contains. This deliverable is subject to final acceptance by the European Commission.

(*) Dissemination level: **(PU)** Public, fully open, e.g., web (Deliverables flagged as public will be automatically published in CORDIS project's page). **(SEN)** Sensitive, limited under the conditions of the Grant Agreement. **(Classified EU-R)** EU RESTRICTED under the Commission Decision No2015/444. **(Classified EU-C)** EU CONFIDENTIAL under the Commission Decision No2015/444. **(Classified EU-S)** EU SECRET under the Commission Decision No2015/444.

Document Information

List of Contributors	
Name	**Partner**
Miguel Martín Pérez	ATS
Susana González Zarzosa	ATS
Aljosa Pasic	ATS
Jorge Martínez Olmo	ATS
Guillermo Yuste	ATS
Justin Činkelj	XLB
Daniel Vladušič	XLB
Lorezo Gorza	INS
Gilda De Marco	INS
Daniele Fabro	CAF
Enrico Di Stefano	CAF
Andreja Markun	TS
Blaž Jemenšek	PIL
Tomaž Ramšak	SZ

Document History			
Version	**Date**	**Change editors**	**Changes**
0.1	07/02/2025	Miguel Martín	First version of the document (ToC).
0.2	07/03/2025	Susana González	Merge of inputs with updates on CPR components
0.3	18/03/2025	Susana González	Updated Executive Summary and chapters 1, 3 and 4. Merged inputs for chapter 5.
0.4	28/03/2025	Susana González	Last inputs for a version ready to be share for peer-review.
0.5	11/04/2025	Susana González	Addressed peer-reviewers' comments.
0.6	14/04/2025	Susana González and Jorge Martínez	Addressed peer-reviewers' comments.
0.7	16/04/2025	Susana González	Addressed changes after proofreading performed by CCL.
0.8	29/04/2025	Susana González	Minor updates
0.9	29/04/2025	Juan Alonso (ATS)	Quality Assessment
1.0	30/04/2025	Aljosa Pasic (ATS)	Final version

Quality Control		
Role	**Who (Partner short name)**	**Approval Date**
Deliverable leader	Susana González Zarzosa (ATS)	29/04/2025
Quality Manager	Juan Alonso (ATS)	29/04/2025
Project Coordinator	Aljosa Pasic (ATS)	30/04/2025

Table of Content

List of Tables

List of Figures

List of Acronyms

Abbreviation / acronym	Description
AIRE	Atos Incident Report Engine
API	Application Programming Interface
ATT&CK	Adversarial Tactics, Techniques and Common Knowledge
BERT	Bidirectional Encoder Representations from Transformers
BGL	BlueGene/L supercomputer
BPMN	Business Process Model and Notation
CERCA	CybEr Risk assessment CAlculator
CERT	Computer Emergency Response Team
CI	Critical Infrastructure
CPR	Cyber Physical Resilience
CPU	Central Processing Unit
CSIRT	Computer Security Incident Response Teams
CTI	Cyber Threat Intelligence
CVE	Common Vulnerabilities and Exposures
D6.1	Deliverable number 1 belonging to WP 6
D6.2	Deliverable number 2 belonging to WP 6
D6.3	Deliverable number 3 belonging to WP 6
D6.4	Deliverable number 4 belonging to WP 6
DB	Database
EC	European Commission
EC2	(Amazon) Elastic Compute Cloud
FTP	File Transfer Protocol
GB	Gigabyte
GDPR	General Data Protection Regulation
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HDD	Hard Disk Drive
HTTP	Hypertext Transfer Protocol
ICT	Information and Communications Technology
IP	Internet Protocol (address)
IRM	Integrated Risk Management
JSON	JavaScript Object Notation
LOMOS	LOg MOnitoring System
MISP	Malware Information Sharing Platform
MS	Milestone – Milestones as defined in DoA
NIS	Network and Information Security Directive
REST	Representational State Transfer
SIEM	Security Information and Event Management
SOC	Security Operations Centre

SQL	Structured Query Language
TIE	Threat Intelligence Engine
TINTED	Threat Intelligence Node inTEgrated with Data enrichment services
TTP	Tactics, Techniques And Procedures
TRL	Technical Readiness Level – TRL1 to TRL9
VM	Virtual Machine
VPN	Virtual Private Network
WP	Work Package

Executive Summary

This deliverable presents the cyber-physical resilience (CPR) tool. The goal of this tool is to improve the cyber-physical security capabilities of critical infrastructure (CI). The importance of improving cyber-security is evident in extreme situations such as a pandemic, where the CIs may face issues such as absenteeism, which reduces the manpower required to defend the increasingly digital and connected critical infrastructure.

The CPR tool (its concept, architecture, initial testing on publicly available data and the piloting scenarios envisaged for validation) was presented in D6.1[1] and previous versions of this Cyber-physical resilience tool and training guide (D6.2[2] and D6.3[3]). Therefore, this final iteration presents the improvements in the tool after its validation and demonstration during the first piloting phase (M21-M22 of the project, June and July 2024). These improvements are:

- Real-world usability improvements and fine-tuning of thresholds in anomaly detection, as observed and requested by CIs;
- Enhancements in anomaly detection for better integration into the various existing systems (such as Wazuh SIEM);
- Identification of threat intelligence data sharing use cases where anonymization and encryption are applicable;
- Improvements in the source confidence scoring module to support CIRCL feeds and in the contextualization of health trends for threat score and better integration with risk assessment;
- An extension of the risk assessment module with a simulation of mitigations mapped to the MITRE DEF3ND framework;
- Adaptation of the incident reporting workflow enforced by the CPR incident reporting module to the NIS2 directive specifications to support early warning, communication of supply chain risk, countermeasures (provided as response to an incident notification), monitoring procedures and vulnerability notification procedures.

The tool was tested with public data and data provided by the CI. Using real data and setting up the facility to receive live data (through FileBeat), we validated the CPR tool (TRL5) via MS8 - First Pilot Trials Complete (where the upgraded and integrated tool was demonstrated and validated, and the first pilot trials were successfully completed). This was reported in D6.4[4]. We will proceed with the final validation and demonstration in M33-M34 of the project (June and July 2025), when the CPR tool is installed locally at the premises of some of the CIs. The final validation results will be reported in D6.6.

The architecture of the CPR tool did not change throughout its evolution – it remains the same as described inD6.2[2]. However, for clarity, we have briefly described the CPR tool's architecture in this deliverable.

Testing the results with the CI provided data logs that strongly suggest we can identify relevant events for cyber-physical resilience in different layers of cyber infrastructure. An example is a problem (bug) in the upper layer (application layer). The anomalies in the infrastructure may be a result of tampering with the application, the virtual resources' infrastructure, or even network issues. Each anomaly may not be directly connected with the cyberattacks. However, their existence or combination with other indicators may show that a threat is imminent. This is an innovative approach. By focusing on a wider threat landscape that includes anomalies or indicators of compromise (IoC) rather than just demonstrably security-related incidents, we may detect more subtle attacks. In this sense, advancements such as threat intelligence scoring or mapping of Indicators of Compromise (IoC) with MITRE ATT&CK matrix techniques improve CI cybersecurity operators' situational awareness. Both detected anomalies and information received from the cyber threat intelligence tool are treated as risk indicators like the already existing indicators, whether they are organisational, technical, physical or workforce related. New input to automated risk (re)assessment also includes temporary conditions related to pandemics, such as changes of priority in impact assessment, availability and awareness

level of workforce, etc. The results of risk assessments are used to help in orientation and decision making, including decisions related to incident reporting. In this way, we follow the OODA loop approach (Observe, Orient, Decide, Act) that contextualizes the available information, while also making sense of newly arrived data and changing circumstances. This is a particularly suitable approach for a volatile, uncertain, complex and often ambiguous inflow of data, such as in the case of cybersecurity decision making systems during temporary conditions like pandemics.

For each module, we describe its deployment in the infrastructure, which serves as guide on how to install the entire CPR tool on the target (CI) infrastructure. We also provide a user manual that describes the configuration and management of the components and the CPR tool itself.

Therefore, the main contributions of this deliverable are the description of the improvements of the CPR tool components based on the results of the first piloting phase, and the updated deployment and user manual, including the integrated CPR dashboard as a single access point to all the functionalities. The piloting activities and their results during the first phase were already reported in D6.4 [4], while the validation of the updated CPR tool already deployed in the CI's premises will be reported in D6.6.

1 Introduction

1.1 Purpose of the document

The purpose of D6.5 is to publicly present the final user guide of the cyber-physical resilience (CPR) tool, initially presented in D6.1 - *Cyber-physical resilience conceptualization* [1], improved in D6.2 - *Cyber-physical resilience tool and training guide V1* [2] and updated in D6.3 - *Cyber-physical resilience tool and training guide V2* [3]. The deliverable presents the tool's architecture and deployment process. The outcomes of the tests conducted on the submodules in a lab environment on public datasets were presented in detail in D6.2 [2]. We summarize D6.2 [2] to be comprehensive and present new results that rely on testing with the sample log files provided by CI operator partners for D6.2. In the annex, we present the final version of the user guide.

This section provides the necessary contextual information, including this document's relation to other project work, its organization, the differences between previous (D6.3 [3]) and this deliverable (D6.5), the glossary adapted in this document, etc.

1.2 Relation to other project work

D6.5 is the final version of the *Cyber-physical resilience tool and training guide* and it follows the second version (D6.3[3]).

This deliverable reports the results of the completed tasks in WP6: *T6.1 Cyber-physical security risk assessment, T6.2 AI-powered log monitoring,* and *T6.3 Incident response and threat intelligence sharing*. Whereas D6.3[3] provided the results of the validation with the real data from the SUNRISE partners, albeit conducted without installing the CPR tool at the CI's premises, this deliverable reports the enhancements and refinements after the activities conducted for the demonstration and validation of the tool during the first piloting phase in M21 and M22 of the project. In addition to the tasks in WP6, D6.5 also reports on T6.4 and T6.5 (integration and demonstration), the updates to the tool's deployment, and the usage guidelines of the CPR tool modules presented in D6.1 [1], D6.2 [2] and D6.3 [3], which include anomaly detection (LOMOS), threat intelligence (TINTED), risk assessment (CERCA), and incident reporting (AIRE).

The CPR tool developed in WP6 contributes to the WP2 strategy for awareness and resilience of CIs through its adaptability to specific situations, such as those experienced during the pandemic. This is achieved through its better detection of anomalous patterns to identify false positives in the context of temporary conditions and events that can happen in the cyber space as well as in a physical context (including workforce absenteeism), when evaluating the risks for a company's infrastructure. The more flexible incident response process supports the challenges of the new NIS2 directive, and fosters collaboration and prevention through secure and more contextualized threat intelligence data sharing. In other words, the CPR tool improves the awareness and resilience of CIs from different perspectives.

Additionally, the CPR tool's development stemmed from the use case scenarios in WP2. This gave us a better understanding of the CI's challenges and gaps during emergency situations such pandemics. These use case scenarios from WP2 enabled us to address the CI needs in WP6. The CPR tool follows requirements elicited in WP3 and reported in D3.2 – *Requirements and designs V2* [5].

The work on the CPR tool reported in this document also relates to other WPs, providing them with a complementary way to physically observe, report and protect CIs. An example is WP7, where the drone's physical security is monitored. The work on the CPR tool provides a way to monitor any possible tampering with potentially dangerous equipment. Information from WP4 related to access control can also enable risk pattern indicators for physical events, as can the health prediction developed from WP5 because it contextualises intelligence data as an input for the CPR risk assessment.

1.3 Changes between D6.3 and D6.5

The differences between D6.3 [3] and its successor D6.5 (this document) are shown in Table 1. In this table, we note new sections and major updates. We also summarize advancements completed.

Table 1: D6.5 - changes in Sections compared to D6.3[3]

Section in D6.5		Section in D6.3 [1]		Difference
Executive Summary		Executive Summary		
1	Introduction	1	Introduction	
1.1	Purpose of the document	1.1	Purpose of the document	Minor update
1.2	Relation to other project work	1.2	Relation to other project work	Minor update
1.4	Structure of the document	1.3	Structure of the document	Minor update
1.3	Changes between D6.3[3] and D6.5	1.4	Differences to D6.3[3]	Major update
1.5	Personas used in the document	1.5	Personas used in the document	Minor update
1.6	Glossary adopted in this document			New section
2	Cyber-Physical Resilience Tool	2	Cyber-Physical Resilience Tool	
2.1	General Context	2.1	General Context	Minor update.
2.2	Architecture	2.2	Architecture	Minor update
		2.3	Operational Prerequisites	Minor update. Moved to section 2.4.
2.3	Deployment	2.4	Deployment	Major update
3	Cyber-Physical Resilience Tool Tested in Pilots	3	Cyber-Physical Resilience Tool Tested in Labs	
3.1	Monitoring System (LOMOS)	3.1	Monitoring System (LOMOS)	Major update
3.2	Threat Intelligence (TINTED)	3.2	Threat Intelligence (TINTED)	Major update
3.3	Risk Assessment (CERCA)	3.3	Risk Assessment (CERCA)	Major update
3.4	Incident Reporting (AIRE)	3.4	Incident Reporting (AIRE)	Major update
3.5	CPR Dashboard			New section
4	Legal Compliance and Testing Methodology	4	Legal Compliance and Testing Methodology	Minor update
5	Pilot trials execution	5	Pilot trials execution	
5.1	Pilot Execution Plans	5.1	Proof of concepts	Minor update
5.2	Description of End-Users' Roles	5.2	Description of End-Users' Roles	No change
5.3	What are the benefits of the Cyber-Physical Resilience tool?			New section

Section in D6.5		Section in D6.3 [1]		Difference
6	Conclusions	6	Conclusions	Minor update
7	References	7	References	Minor update
Annex I	Training guide and user manual	Annex I	Training guide and user manual	
I.I	CPR Dashboard			New section
I.II	Anomaly Detection	I.I	Anomaly Detection	Minor update
I.III	Threat Intelligence	I.II	Threat Intelligence	Minor update
I.IV	Risk Assessment	I.III	Risk Assessment	Minor update
I.V	Incident Reporting	I.IV	Incident Reporting	Minor update

1.4 Structure of the document

This document contains 6 chapters and annex.

Chapter 1 (this chapter) presents the purpose and contextual information of this (D6.5) deliverable.

Chapter 2 presents the overview of the Cyber-Physical Resilience Tool. This chapter focuses on the final architecture and deployment of the four modules. The data, which is required for an easy understanding of the modules and architecture, is summarized for the reader.

Chapter 3 presents the improvements and adaptations of the different modules in the CPR tool after the first demonstration and validation in the pilots. The concrete results of the first piloting phase were reported in D6.4 [4].

Chapter 4 describes possible issues of the pilots due to legal restrictions related to CI and outlines the testing methodology, as these are connected. We expect the possible legal issues to be relevant when we start the actual piloting during the second piloting phase and involve more of the CI technical personnel.

Chapter 5 presents the execution plans for four CI pilot trials related to public administration, water, telecommunication, and transport. It also describes the end-users’ roles.

Chapter 6 outlines the main findings of this deliverable and provides conclusions that will be used to validate our approach and developments.

In summary, the main focus of this deliverable is to provide a final version of the deployment and training guides for the different modules of the CPR tool, and present the latest improvements implemented after the first piloting phase, the latter of which is required for the execution of the final piloting phase.

The **Annex** offers a user manual for the modules presented in previous chapters. The manual is presented in a step-by-step descriptive format with screenshots from demo deployments.

1.5 Glossary adopted in this document

Critical Infrastructure (CI): Power distribution networks, transportation networks, and information and communication systems are all examples of critical infrastructure. The defence of critical assets is indeed essential for ensuring the safety and well-being of the European Union (EU) and its citizens. The electrical grid, transportation systems, and information and communication networks are key examples of what is known as "Critical Infrastructures". These infrastructures are essential to maintain in order to ensure that vital societal functions continue to operate smoothly. Natural disasters, acts of terrorism, and criminal activities all have the potential to cause damage to or destroy essential infrastructure, which may have serious repercussions for both the safety of EU residents and the complete EU.

Critical Assets (CAs): Are the significant resources that support both the social and business parts of an economy. If some of these assets fail, it will bring significant issues for business continuity. This does not mean that the likelihood of failing is high. For planning purposes, each business or organization must identify its critical assets and know the corresponding information about them.

Use Case – A description of the interaction between an actor (e.g. a user or system component) and the system, outlining the sequence of actions or steps taken to achieve a specific goal. It represents a technical or functional task relevant to a particular operational need.

Scenario - Scenarios provide narrative or technical framing for a use case, with two types:

- **Contextual Scenario:** A broader, real-world context such as a pandemic or multi-hazard threat environment that may influence multiple systems.
- **Use Case Scenario:** A focused variant or narrative path within a specific use case, detailing alternative technical or procedural flows.

Pilot: Real-world validation phase of the project, built on defined use cases and scenarios. It involves the deployment and evaluation of solutions under realistic conditions.

WP-Specific Terminology

- **Model**: Within the scope of this document, the term "model" refers to the mathematical structure that uses different algorithms and different libraries/frameworks in order to learn from training data (training process over a dataset) and generate predictions about future events. The term "artifact" is a machine learning term that is used to describe the output created by the training process. Output can be a fully trained model, a model checkpoint, or a file created during the training process.
- **Resilience:** Resilience is a concept that spans different domains (physical, information, cognitive or social) and has different capacities, abilities, or principles regarding different adverse events. In the context of SUNRISE, resilience applies to capacities and abilities of Critical Infrastructures. In WP6, cyber resilience refers to "*The ability to prevent, prepare for, withstand and recover from cyberattacks and incidents*" [17]
- **Cyber-physical:** In this document, the term cyber-physical refers to the CPR tool's ability to process inputs about cyber and physical assets and events, such as from anomaly detection, SIEM or IDS tools within the CIs or any other risk indicator that refers to the physical world (e.g. workforce availability).
- **Temporary conditions:** This term refers to any variable or parameter that changes operational conditions. It is important to consider as a risk indicator (e.g. workforce availability, risks derived from lockdown) or is important for risk assessment (e.g. change of priority, procedure or asset value).
- **Threat intelligence:** This term refers to any information related to the identification, analysis or occurrence of a cyber threat, and is shared with the aim of preventing and mitigating potential incidents, vulnerabilities or cyber-attacks while fostering collaboration among various stakeholders.
- **Risk assessment:** In this document, risk assessment refers to the qualitative (ranging from very low to very high) and quantitative (in terms of potential cost to the organization in the event of an attack) assessment of a set of predefined cyber-physical risk patterns (e.g. Denial of Service, Man-in-the-middle attack, SQL injection attack, etc). This evaluation is based on different type of inputs, including static information (manually configured through a dashboard) and dynamic data from other CPR modules (in particular, anomaly detection and threat intelligence modules), with adaptable probabilities based on the activation of temporary conditions.
- **Anomaly detection:** In the context of this document, anomaly detection refers to identifying patterns, behaviours or events in cyber or cyber-physical systems that deviate from expected norms, potentially indicating a malfunction, misconfiguration, or a cybersecurity incident. These anomalies are detected through continuous monitoring of system or application logs (the logs, typically produced by any system or application, are written by an application developer and are usually contained within a text file). The inspection of such logs is critical for early warning, situational awareness and dynamic risk assessment within Critical Infrastructures.

2 Cyber-Physical Resilience Tool

This chapter describes the Cyber-Physical Resilience Tool developed in SUNRISE, which provides Critical Infrastructure users with an integral security solution to improve and strengthen their cyber-physical resilience. This provides the CI users with a better and more proactive adaptation to new threats and challenge situations such as those in a pandemic. We first provide a general context of the tool with the modules it is composed of and some ethical considerations regarding the use of the data in the tool. Then, each of the modules are described in more detail (section 2.2). This information was already provided in a previous deliverable, D6.3[3], but we included it here to make this deliverable comprehensive. Furthermore, this information has been expanded upon in this deliverable, as it provides more details on the interaction among the different modules via a concrete illustrative use case. Finally, how the CPR tool can be deployed in the Critical Infrastructures (section 2.3) and its requirements are described with detail.

Further details on the technological background for each of the modules and state of the art are provided in D6.1[1]. Details on the results of the validation done during the first piloting phase and the feedback provided by the pilots can be found in D6.4[4]. Concrete implementation examples for the different infrastructures involved in the pilots will be reported after the execution of the second piloting phase in D6.6.

2.1 General Context

The CPR tool contains four modules, which are introduced in the sections below. The modules are:

- The Anomaly detection module (based on LOMOS tool), which utilizes machine learning for log-based anomaly detection.
- The Threat intelligence module (based on TINTED tool), which is used for sharing the threat intelligence information with external stakeholders, including pandemics/health related threat intelligence, the scoring of threats and their sources, and mapping of low-level technical information, such as indicators of compromise (IoC), to higher level techniques, tactics and procedures (TTP).
- The Risk assessment module (based on CERCA tool), which ingests different types of inputs (real time feeds from existing cybersecurity tools such as anomalies or incidents, events related to physical threats, changes in temporary conditions, threat intelligence, organisational and workforce related information etc) and automatically (re)assesses risks, based on pre-defined risk models.
- The Incident reporting module (based on AIRE tool), which is responsible for automating the reporting process to the relevant authorities in the event of a security incident.

The modules are designed to allow the integration of legacy systems that are currently used by CI.

Organizations have a legitimate reason to collect and process personal data with CPR tool (e.g., security monitoring, legal obligations according to NIS2 Directive). The use of CPR tool must be stated in policies, and employees or users should be informed about data collection. In this sense, the CPR tool has purpose limitation, such as threat detection, incident response, and compliance auditing and it cannot be repurposed for other uses without explicit consent or legal justification.

In addition, the CPR tool only collects necessary personal data to fulfill security objectives (e.g. IP addresses or usernames). The process is that the individual logs collected are only used to train the model in the Anomaly Detection module, and later used to infer the behaviour. Essentially, the data, such as in transient with LOMOS (CPR Anomaly Detection module), is only used for model training and inference. After that, it can be discarded, or logs can be anonymized or pseudonymized where possible to reduce privacy risks. Regardless, it is important highlight that the CPR tool will usually be installed on the CI's own infrastructure. This means that the CPR tool has the same level of security as the applications it is trying to protect. Additionally, data retention policies can be applied to automatically delete the data after a predefined period, which is typically based on security needs or legal requirements. The data is protected against unauthorized access, alteration or breaches, but organizations using CPR tool must be able to demonstrate GDPR compliance through documentation and policies.

2.2 Architecture

The CPR tool is designed to enhance the resilience of CI during pandemics by considering both technical and human aspects. It comprises four modules: Detection, Threat Intelligence, Risk Assessment, and Incident Reporting. The tool's architecture and data flow are illustrated in Figure 1.

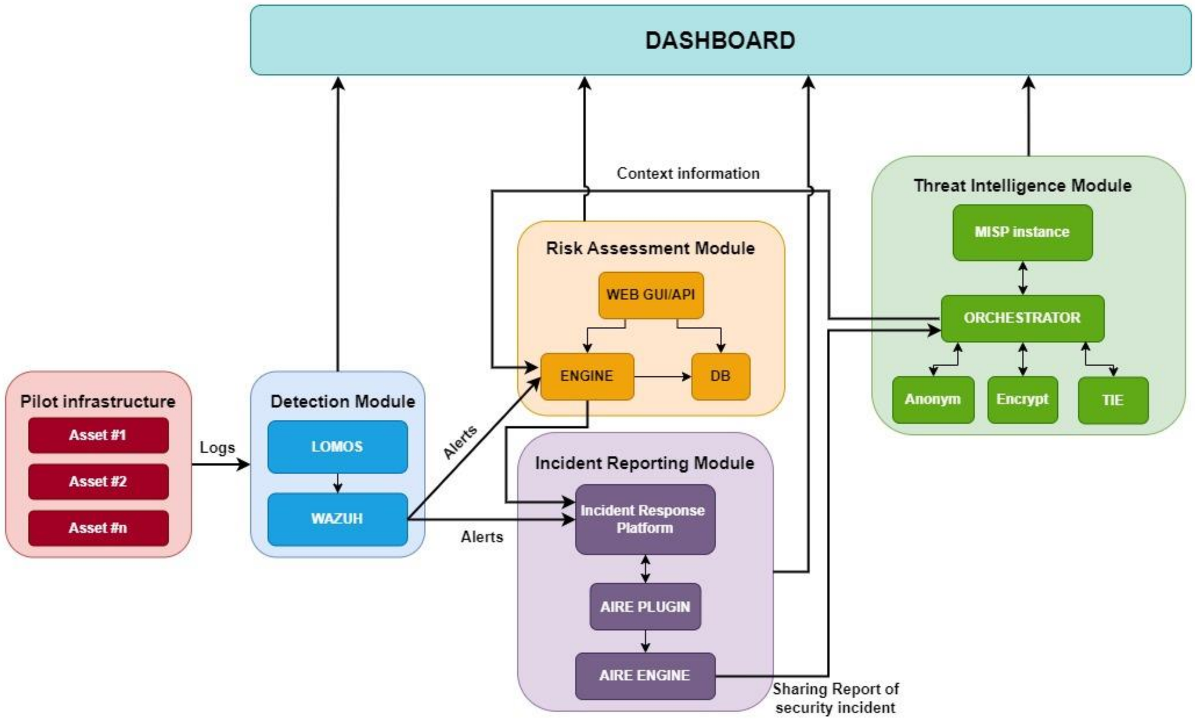

Figure 1. Overall architecture of the CPR tool (extracted from D6.1[1]).

The Detection module oversees CI assets and generates security events and low-level alerts. It comprises the two main components: an anomaly detector named LOMOS, which uses application logs and raw data to train an AI in recognizing normal system behaviour and triggering alerts for deviations, and a SIEM called Wazuh[21], which identifies event sequences related to known threats and raises alarms correlating with LOMOS' output. These alerts and events are sent to the Risk Assessment and Incident Reporting Modules.

The Risk Assessment Module evaluates incoming input to assess the risk associated with the system's assets. The outcome is a risk report that quantifies cyber risk exposure in monetary terms. This report includes the likelihood of security incidents based on many inputs, including real-time cyber environment data, and an estimation of monetary impact related to the incident and based on the currently estimated values of the digital assets, which might change during pandemics in so called "temporary conditions". It is also sent to the Incident Reporting Module, which evaluates events and alarms from the SIEM to determine security incident presence and the severity for potential reporting to authorities. In addition, this module ensures the orderly completion of reporting steps.

The Threat Intelligence Module has two primary roles: secure sharing of Cyber Threat Intelligence (CTI) information and enhancing external threat intelligence data, in order to improve estimations of incident likelihood. It employs a common MISP[22] instance to share relevant attack and threat information. Having notified authorities of a security incident, CIs may choose to share this information with others. Input from the Incident Reporting Module is integrated, ensuring appropriate context in the information flow. Before sharing, data can be encrypted or anonymized. The module employs a threat score generated by the Threat Intelligence Engine (TIE) module, using heuristics on incoming external events to provide context-aware data to the Risk Assessment module.

The logic behind this architecture stems from the need for adaptivity, collaboration and dealing with absenteeism, scenarios that have been also identified and described in WP2 as especially relevant for the temporary operational conditions such as those during pandemics.

During these periods, SUNRISE users observed the fast-changing cybersecurity landscape, systems and behaviours, where the existing static practices were no longer sufficient. The plan-do-check-act (PDCA) approach, for example, follows linear process, which has some drawbacks during these periods with very dynamic changes. Its strength is that it is a simple yet effective way to plan, structure and implement risk controls through corrective and preventive actions. The “Plan” and “Do” phases involved identifying assets, models and impacts, assessing the risks, proposing mitigation etc. The “check” phase involved assessing how well the risks have been controlled (usually by audits). This was separated from the operational cybersecurity management that was focused on real-time monitoring and detection of suspicious events. The main weakness is its static nature, which is inappropriate for steady reduction of the “time to react”, one of the most important metrics for cyber-physical resilience.

More recently, real time risk engines such as the one used in CERCA were introduced, leading to a possibility to trigger risk (re)assessments automatically, such as after the detection of an incident or vulnerability, or when temporary conditions change.

The OODA loop (Observe, Orient, Decide, Act) is another a four-step approach to decision-making, but unlike PDCA, it focuses on the contextualization of the available information, while also making sense of newly arrived data and changing circumstances. It is a particularly suitable approach for a volatile, uncertain, complex and often ambiguous inflow of data, such as in the case of pandemics. Here, rule-based detection, used in cybersecurity operational management tools, is combined with expert-based risk assessment, which prioritizes and optimizes mitigation actions and readjusts risk baselines.

This is why the CPR tool architecture proposes to combine the sensing (“observation”) of external environment (anomalies from LOMOS, incidents from SIEM, threats from TINTED etc.) with a cognitive process of “orientation” to make an optimal decision. Observation works with outside-in sensory information, even if it uses rules for the correlation of events, or signatures for detection of an intrusion. Orientation works with an inside-out created cyber risk landscape, which includes changes in temporary conditions, such as the availability of cybersecurity workforce, or a higher priority and value for health-related digital assets.

Our approach and architecture work better with uncertainties, partial information, diversity or a degree of randomness and disorder. While the availability of incoming data in observation is compulsory, calculated values such as the likelihood of an attack, based on this data, vary depending on the actual context, including confidence in the sources of threat intelligence, timeliness or relevance of incoming information. The same holds for the calculation of impact, which is better aligned with priorities and changes at strategic and tactical level (e.g. due to the interruption in supply chain or unavailability of workforce). As a result of the “Action” phase in OODA loop, there will be a risk of mitigation decision, and the feedback loop goes towards the external observation environment through sharing threat intelligence, which can also help to correct cognitive bias in relation to risks from other CI operators.

2.2.1 Anomaly Detection

The CPR Anomaly Detection module is based on LOMOS. LOMOS is a self-supervised machine-learning system for log-based anomaly detection, comprising a log parser, anomaly detector trainer, and anomaly detector. It extracts log events using Drain, trains models with LogBERT, and assigns anomaly scores per log. The architecture includes components for log storage, parsing, anomaly detection, and visualization. Components like lomos-cli and lomos-api2 facilitate training and inference tasks, see

Figure 2. These components were added to satisfy the needs of the CIs that are part of the SUNRISE project. The system utilizes Celery for job distribution, MLflow for experiment tracking, and Grafana for dashboard development. lomos-api2 offers anomaly detection without direct database access. lomos-cli aids in frequent training with stored parameters for repeatability.

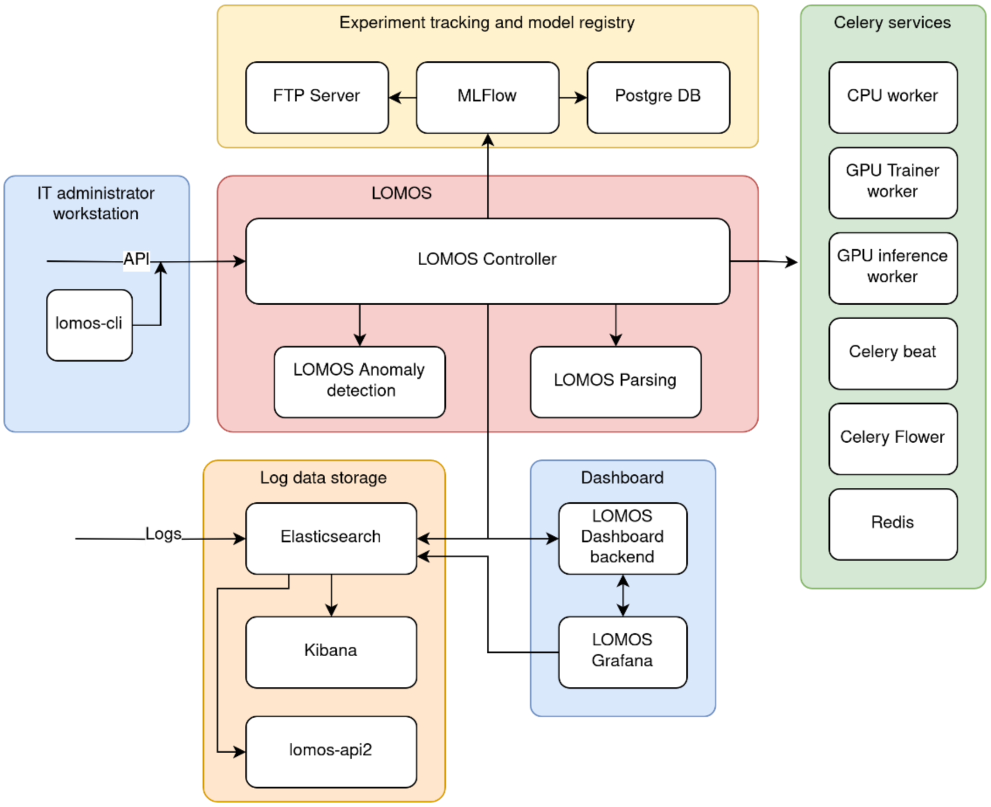

Figure 2. CPR Anomaly Detection module architecture (from D6.1[1]).

The service lomos-api2 was developed to provide a list of detected anomalies in a given time window. It is implemented as a REST API. It isolates the consuming application from the LOMOS internal working. Also, the consuming application does not need direct access to the LOMOS internal database to find anomalies.

During the experimental LOMOS training, we frequently needed to retry experiments, which involves tasks of log parsing, model training and inference. The provided GUI dashboard is primary a tool for all three tasks. However, the CI operator's IT administrator needs to input many parameters, and this makes usage difficult, inefficient and prone to human error.

For frequent LOMOS training, an additional CLI tool named lomos-cli was developed. The required parameters are stored in JSON files and can be committed to the source versioning system (git). This enables experiment repeatability. The credentials required to access LOMOS deployment need to be provided separately. The tool can load credentials from the repository used to deploy LOMOS (env/.env and env/.secrets files).

For the SUNRISE project and its partners, the core of LOMOS is based on the BERT method (see [7][8]). Newer core methods like [15] are in the pipeline to be developed in the future, further improving the method. However, in the landscape of rapidly changing and evolving AI, stability is pursued, especially in the field of cybersecurity.

2.2.2 Threat Intelligence

The CPR Threat Intelligence module, based on TINTED tool, facilitates secure CTI exchange and enhances cyber risk calculation. Its layers include authentication, transport, and privacy/enrichment as shown in Figure 3. It allows end-users to communicate in a decentralized, confidential, and anonymous manner through MISP (Malware Information Sharing Platform). It also includes CTO scoring capabilities to automatically calculate a CTI threat score and a source confidence score that allow users to easily identify a reliable source and to assess how useful it is for them the incoming Indicators of Compromise. Additionally, the IoCs received are automatically enriched with information about related MITRE ATT&CK Techniques.

Figure 3 shows a general overview of the CPR Threat Intelligence module architecture, indicating its main components. Each of the components is described in detail in the following bullet points. An in-depth look at TINTED is provided in D6.2 [2].

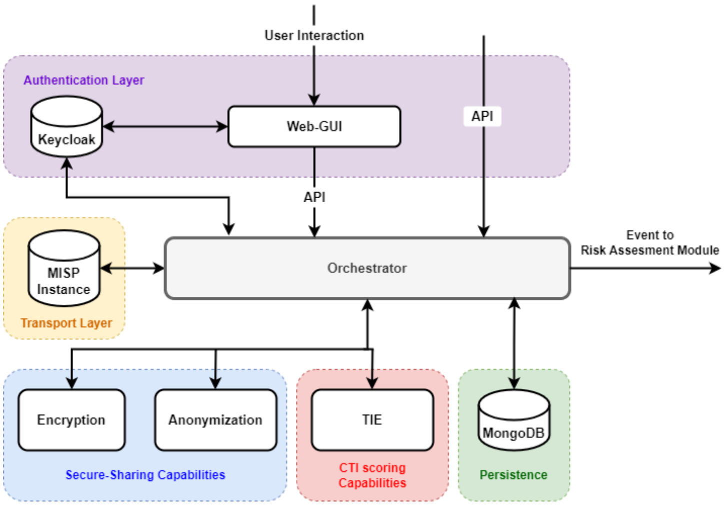

Figure 3. CPR Threat Intelligence Sharing architecture (from D6.1[1]).

- **MISP Instance**: This serves as the transport layer of the platform, ensuring secure CTI exchange between different users and organization with granular sharing controls. The philosophy of MISP is to share information with everyone. However, there are cases in which information sharing can increase the risk, as this CTI may contain sensitive information about the organization, which attackers can exploit. This is one of the issues we address with the TINTED sharing capabilities, by adding greater granularity to the sharing options.
- **Keycloak**: An open-source platform that facilitates Identity and Access Management tasks. TINTED uses these functionalities to implement secure and efficient user authentication and authorization.
- **MongoDB** This handles data storage, storing configuration, and user and infrastructure information.
- **Orchestrator**: The orchestrator is the main component of the TINTED architecture, facilitating communication among modules and providing an API for user interaction. This component is also responsible for sending events with a score over a threshold to the CPR risk assessment module via Kafka.
- **WEB-GUI**: The graphical user interface simplifies tool usage and integrates authentication mechanisms through Keycloak.
- Privacy modules include **Encryption** and **Anonymization**.
- **The Threat Intelligence Engine (TIE)**: This component consists of two modules, the ZMQClient and the HeuristicEngine. The ZMQClient is in charge of subscribing to a ZMQ queue within the MISP instance. It receives a notification every time an event is uploaded, and it sends that event to the HeuristicEngine for its processing. This engine is in charge of calculating a threat score for the IoCs, evaluating its source confidence score and searching for any mapping with MITRE ATT&CK tactics and techniques.

Main innovations beyond the state of the art developed during the project can be summarized as: (i) trustworthy sharing of threat intelligence data using anonymization and encryption; (ii) providing support for monitoring the threat landscape using reliable open-source intelligence feeds through to the automatic processing of the incoming Indicators of Compromise to generate a threat intelligence source confidence scoring and a threat scoring (based on different criteria such as relevance or completeness); (iii) automatic enrichment of the IoCs with information about related attack patterns (using MITRE ATT&CK framework); (iv) contextualization of health trends for threat score with impact on risk assessment.

2.2.3 Risk Assessment

The CPR risk assessment module, based on CERCA, comprises several submodules that generate reports sent to the incident reporting module for evaluation, see Figure 4. The **WEB-GUI/API** facilitates tool configuration, while the **Indicator Value Generator** creates or modifies indicators based on user input and threat intelligence data. The **Triggering Detector** activates the **Risk Model Executors**, which execute algorithms to generate reports on risk per target. These reports are aggregated by the **Aggregator** and stored in **PostgreSQL**. Static information from questionnaires and target configurations, along with dynamic data like alarms, events, vulnerability reports and indicators of compromise, feed into the assessment process. Outputs include a comprehensive Global Report, Reports per Target and Risk, Reports per Model, and Reports per Target, providing insights into the organization's threat landscape at various levels of granularity. An in-depth look into CERCA is provided in D6.2[2].

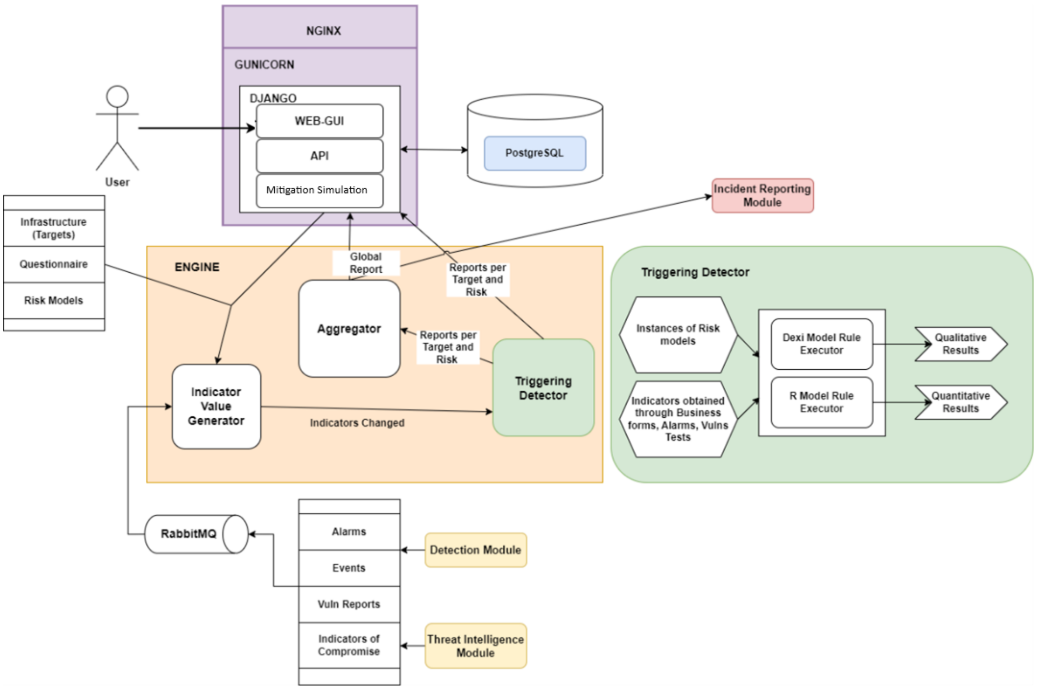

Figure 4. CPR Risk Assessment architecture (from D6.1[1]).

The main innovations beyond the state of the art developed during the project on cyber risk assessment include reduction of uncertainty through qualitative and quantitative risk assessment, integrating temporary conditions considering management of data about absenteeism, physical risks, cyber threat intelligence data, and offering dynamic impact assessment and simulation of mitigations with their impact on the risk assessment.

The mitigation simulation module architecture allows fast simulation and quick iteration. It is deployed inside the dashboard module and benefits from using the same deployment container and access to the Django ORM. New data models were developed to support this functionality, as shown in Figure 5, and with flexibility to adopt new mitigation and simulations strategies in future versions.

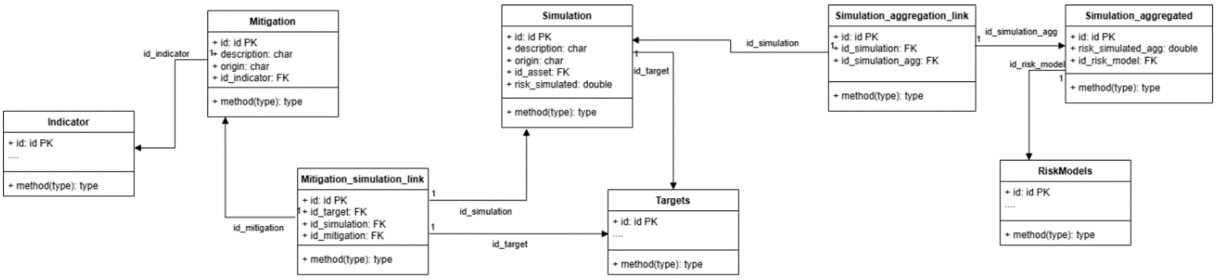

Figure 5. Mitigation simulation data model

2.2.4 Incident Reporting

The incident reporting module included in the SUNRISE Cyber-Physical Resilience (CPR) tool is based on AIRE, the ATOS Incident Reporting Engine tool. It allows companies to manage cybersecurity incidents and to support the reporting to authorities. First, it helps users to manage new incidents through a defined sequence of tasks, saving time by eliminating the confusion of not knowing what to do. Moreover, the module assesses the need to report the incident with respect to various regulations, informing the user about which reports must be submitted. In addition, it ensures that the deadlines for these submissions are met and generates draft reports with the data collected based on the template for each regulation.

The incident reporting tool features a modular design. Its structure, see Figure 6, includes Springboot two main microservices, **aire-reports-generators** and **aire-workflow-enforcement**, along with the **aire-thehive-plugin** service for integration with open-source incident response platform TheHive[23]. The Hive is used to interact with the end-users through tickets assigned to the users with different profiles.

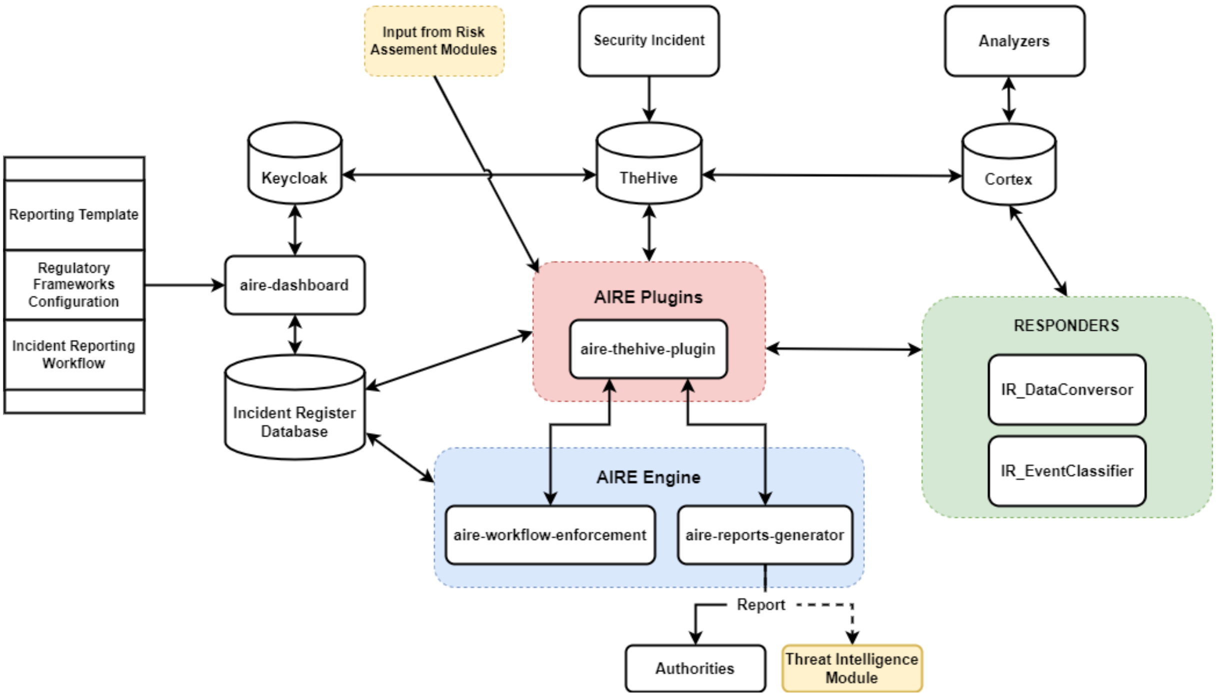

Figure 6. CPR Incident Reporting architecture (from D6.1[1]).

The **aire-workflow-enforcement** orchestrates the mandatory incident reporting workflows that need to be followed by the CIs according to the applicable regulatory frameworks. With this purpose, it follows a BPMN specification that can be adapted to ensure compliance with different regulations and includes managerial oversight. The **aire-reports-generator** service tailors incident data to different report templates mandated by Competent Authorities, utilizing Apache POI for Microsoft formats and Apache PDFBox for PDFs.

The architecture also includes an Incident Register Database and a graphical user interface (**aire-dashboard)** for configuration.

The main innovations beyond the state of the art developed during the project on cyber incident reporting are related to the automation of the incident reporting procedures according to different applicable regulatory frameworks and extending the support of the incident reporting workflow for compliance with the new NIS2 directive. Specifically, it provides improved incident management with risk-based incident classification, early warning with the integration of a countermeasures' evaluation procedure based on the feedback received from CSIRTs, enrichment of incident reports with related threat intelligence data, a vulnerability notification procedure, and support for the communication of supply chain risk.

2.2.5 Dashboard

The CPR tool consolidates modules into a central dashboard for Critical Infrastructure monitoring, providing end-users with an integrated tool that gives them access to all the features to improve their resilience. It includes sections for Log Anomaly Detection (LOMOS), Security Information and Event Management (Wazuh), Threat Intelligence Monitoring (TINTED), Risk Assessment (CERCA), and Incident Reporting (AIRE). Each section offers real-time charts and tables summarizing CI metrics and allows direct access to advanced operation, with options for detailed views and configuration. The CERCA section prioritizes risk scores for monitored assets, while LOMOS showcases anomaly events. SIEM data displays alarms per risk, and TINTED presents event types. AIRE tracks ongoing reports and tasks for regulatory reporting. The dashboard provides a comprehensive overview of CI performance, enhancing decision-making and operational efficiency.

2.2.6 Illustrative Use Case

To clearly describe how the different components of the CPR tool interact with each other and with the different actors, and how the Critical Infrastructures can benefit from the CPR tool, in this section we will present a potential use case where the SESAMO application (included in the Italy public administration pilot, see section 0) is attacked. We will begin with a non-pandemic situation so we can later how the CPR tool enables adaptation to new and challenging circumstances.

The initial Anomaly Detection Phase (represented in the block 1 of Figure 7) shows how the logs generated by the CI's application during the attack are received by LOMOS (CPR Anomaly Detection module). LOMOS detects an abnormal pattern, and an alert is sent to the other CPR modules: CERCA (cyber-physical risk assessment module) and AIRE (incident reporting module). This anomaly detected by LOMOS and the associated incident started in AIRE will appear and can be visualized by the CI Operator through the CPR Dashboard.

When the alert is received CERCA automatically performs a new risk assessment, and the result is that risk is increased from low (the initial state under no attack) to high (represented in block 2 of Figure 7). The new risk will be updated and shown in red to the CI's Operator through the same CPR Dashboard.

Block 3 of Figure 7 represents how the CPR incident reporting module provides support for NIS2 compliance. First, when it receives a new risk assessment, AIRE automatically sends a notification to the supply chain beneficiaries configured for the CI. Furthermore, since the incident is risk-based classified as significant, an early warning is automatically sent to the supervisory authority (CSIRT). They can provide some mitigation countermeasures in response to this early warning.

The CI Operator will interact through the CPR Dashboard (represented in blocks 3 and 4 of Figure 7) to access AIRE and TheHive, and perform all the required tasks for incident management and mandatory incident reporting, including generating the different required templates with the information gathered about the incident. The CI Operator will also interact through the CPR Dashboard to access

CERCA and simulate different mitigations and check how they would impact in the risk assessment. Consequently, through its different modules, the CPR tool provides CIs with decision making support for the implementation of countermeasures (e.g. the application of a vulnerability patch in the SESAMO application). Once the CI Operator determines that countermeasures have been applied and the associated task is closed, AIRE will automatically invoke CERCA for risk assessment re-evaluation. The result is that risk has been decreased. Additionally, via its Threat Intelligence module (see block 5 in Figure 7), the CPR too will enable a secure way to share information through MISP (e.g. with other departments or CSIRTs) about implemented countermeasures.

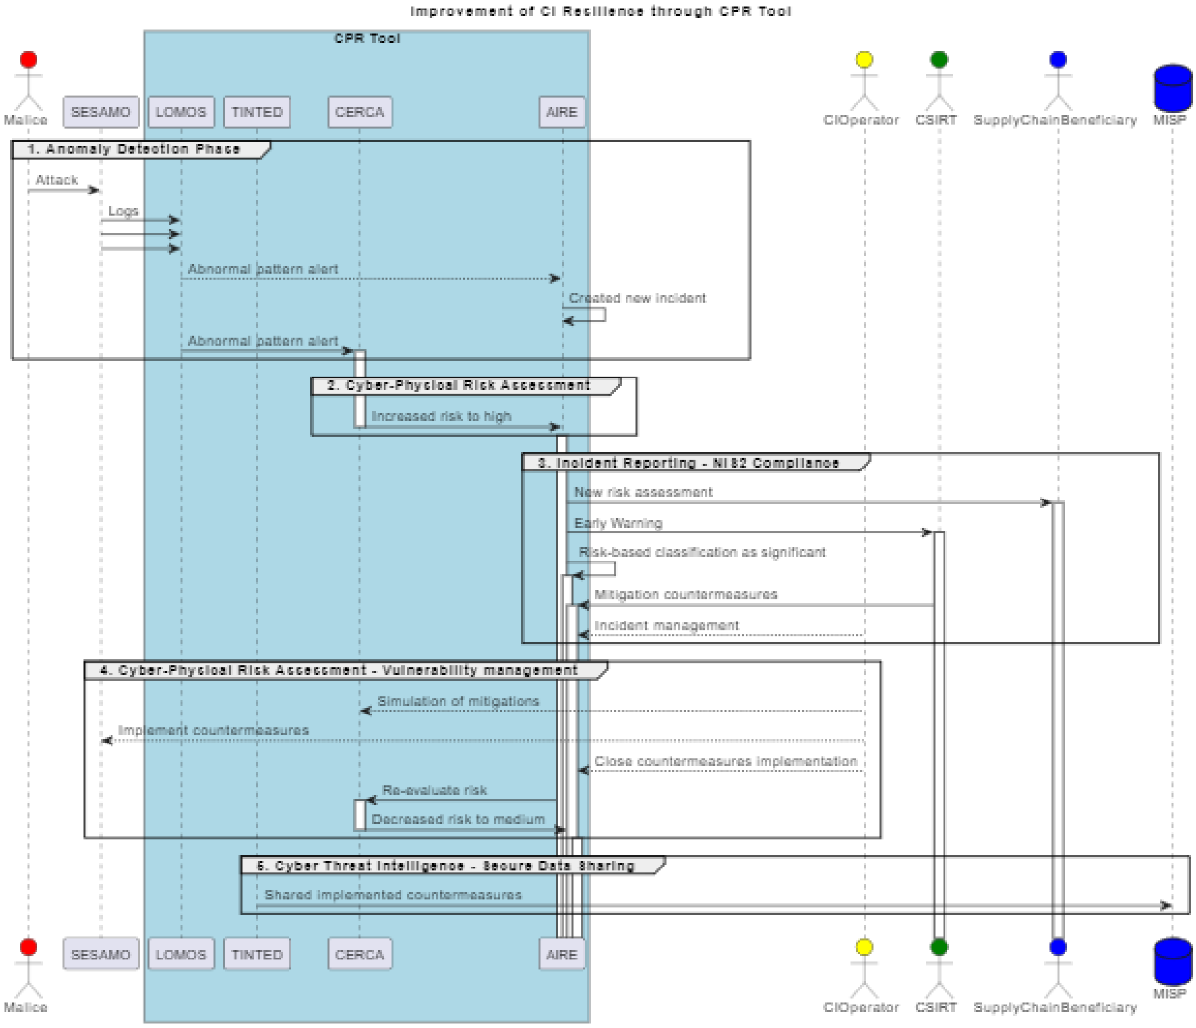
Improvement of CI Resilience through CPR Tool
CPR Tool
Malice
SESAMO
LOMOS
TINTED
CERCA
AIRE
CIOperator
CSIRT
SupplyChainBeneficiary
MISP
1. Anomaly Detection Phase
Attack
Logs
Abnormal pattern alert
Created new incident
Abnormal pattern alert
2. Cyber-Physical Risk Assessment
Increased risk to high
3. Incident Reporting - NIS2 Compliance
New risk assessment
Early Warning
Risk-based classification as significant
Mitigation countermeasures
Incident management
4. Cyber-Physical Risk Assessment - Vulnerability management
Simulation of mitigations
Implement countermeasures
Close countermeasures implementation
Re-evaluate risk
Decreased risk to medium
5. Cyber Threat Intelligence - Secure Data Sharing
Shared implemented countermeasures
Malice
SESAMO
LOMOS
TINTED
CERCA
AIRE
CIOperator
CSIRT
SupplyChainBeneficiary
MISP

Figure 7. Illustrative Use Case with CPR tool (normal use).

Two potential situations during pandemics that impact the risk assessment are lockdowns and workforce absenteeism.

In the case of a lockdown (see block 1 in Figure 8), the increased online activity can be exploited by potential cybercriminals, such as deploying bots that look for invalid session IDs to gain unauthorized access. Extended active sessions increase the risk of hijacking, especially when inactivity or logged out management are inadequately considered in the applications. Additionally, public Wi-Fi or poorly secured home networks are more vulnerable to session sniffing and hijacking attacks. The added value of the CPR tool in this situation is that several temporal conditions can be applied by the CI's Operator to the session hijacking risk model as a response to the increase in risk due to this lockdown situation. Some types of mitigation measures could be related to an increase in session timeout or an idle session, an increase in the number of login attempts before blocking devices or an increase in message timeout. Incident management will notify supply chain beneficiaries of the new risk and CSIRTs, if classified as significant, in a similar way as described previously for any incident.

To be proactive in the detection of potential workforce absenteeism situations, the CPR tool obtains a daily health trend prediction through the Threat Intelligence module, which is sent to CERCA and has an impact on the risk assessment (see block 2 in Figure 8). In this situation, the CI operator can also add temporary conditions to better mitigate this risk.

Finally, block 3 in Figure 8 show how threat intelligence information received through the CPR tool, such as that related to a phishing campaign on a specific sector (such as the health or public sector in relation with the SESAMO application) can also impact in the risk assessment. If the scoring generated by TINTED is over a threshold, the Indicator of Compromise will be sent to CERCA. Risk assessment will be reevaluated and the process shown in Figure 7 will be followed. Additionally, information about phishing can also be securely shared with different trust circles, encrypting or anonymizing sensitive information depending on the recipient.

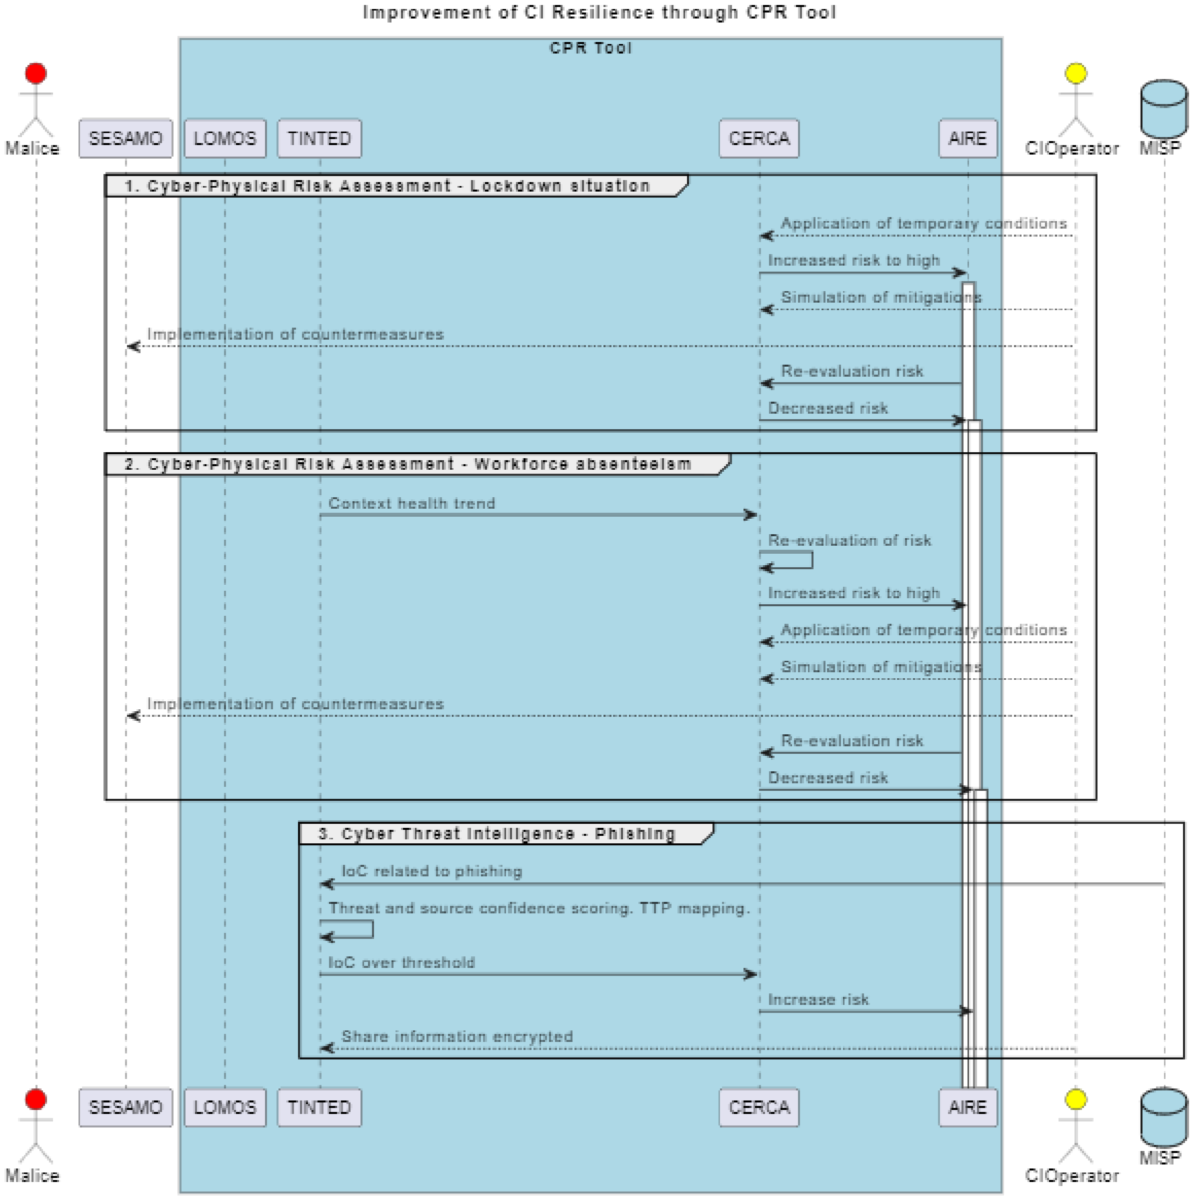

Improvement of CI Resilience through CPR Tool
CPR Tool
Malice
SESAMO
LOMOS
TINTED
CERCA
AIRE
CIOperator
MISP
1. Cyber-Physical Risk Assessment - Lockdown situation
Application of temporary conditions
Increased risk to high
Simulation of mitigations
Implementation of countermeasures
Re-evaluation risk
Decreased risk
2. Cyber-Physical Risk Assessment - Workforce absenteeism
Context health trend
Re-evaluation of risk
Increased risk to high
Application of temporary conditions
Simulation of mitigations
Implementation of countermeasures
Re-evaluation risk
Decreased risk
3. Cyber Threat Intelligence - Phishing
IoC related to phishing
Threat and source confidence scoring. TTP mapping.
IoC over threshold
Increase risk
Share information encrypted
Malice
SESAMO
LOMOS
TINTED
CERCA
AIRE
CIOperator
MISP

Figure 8. Diagrams for CPR tool adaptation to pandemic situations.

2.3 Deployment

2.3.1 Introduction

The CPR tool is composed of five different modules that need to be deployed: the anomaly detection module, the threat intelligence module, the risk assessment module, the incident reporting module and the CPR Dashboard. The details for the deployment of each of these modules is presented in this chapter. Please not that these deployment instructions have been prepared to be read and executed by the IT department, so those sections can be skipped by non-technical end-users.

Considering the hardware needs of the Anomalous Detection Module, it is advisable to use a separate VMs for the deployment of this component. The rest of the modules (Threat Intelligence Module, Risk Assessment Module, Incident Reporting Module and CPR Dashboard) can run in a same VM, but it will need to have enough capacity. The deployment of these modules will also include some dependencies such as the open-source Security Incident Response Platform TheHive and Cortex (and their related databases Elasticsearch and Cassandra), which are used for the Incident Reporting Module, and the deployment of a Kafka server, which is used for communication among the different CPR modules.

However, please note that to operate the Threat Intelligence module, it will be necessary to have a MISP instance deployed and a Keycloak server. The deployment of these tools is not included with the CPR tool deployment scripts. In this document, it is only included the specific configuration required for the CPR tool.

Additionally, information included in this deliverable concerning deployment and training assumes that the Security Information and Event Management (SIEM) tool Wazuh is already deployed in the CI infrastructure.

In general, Wazuh is the most used tool, given its open-source availability. To this end, the connection between LOMOS and Wazuh has been made. However, we can establish connections with other existing SIEMs, such as Splunk (https://www.splunk.com), QRadar (https://www.ibm.com/qradar), FortiSIEM (https://www.fortinet.com/products/siem/fortisiem), etc. It is important to note, that in our experience, Wazuh has been the tool of choice and has most common ground. However, we are open to using other SIEM tools and recognize the need for expanding the default integration efforts if such a need arises.

Information on the installation of the open-source tools used with the CPR tool can be found in the following links:

- MISP Deployment: https://www.circl.lu/doc/misp/get-your-instance/
- Keycloak Deployment: https://www.keycloak.org/operator/basic-deployment
- Wazuh Deployment: https://documentation.wazuh.com/current/installation-guide/index.html
- TheHive: https://github.com/TheHive-Project/TheHive
- Cortex: https://strangebee.com/cortex/
- Kafka: https://kafka.apache.org/

2.3.2 Pre-requisites

All the modules included in the CPR tool are containerized using Docker. Consequently, the deployment environment must have Docker (version 24.06 or newer) to facilitate component deployment.

Testing of these modules has been conducted on Ubuntu 22.04 LTS, Ubuntu 20.04 LTS and Ubuntu 18.04 LTS operating systems.

The current version of the Threat Intelligence Module has been tested with the following considerations:

- MISP v2.4.184.
- MISP API Key used corresponds to a user with the role “admin”.

Although each module in the CPR tool has its own hardware requirements that will be described in more detail in next sections, a checklist is presented below to help users ensure that they have all the necessary elements before beginning the deployment:

□ A VM with at least 8 core CPU, 64 GB RAM and enough hard disk, for the deployment of the Anomaly Detection module.

□ At least 1x GPU installed (11 GB + VRAM) in the VM used for the training mode of the Anomaly Detection Module.

□ A VM with at least 8 core CPU, 64 GB RAM for the deployment of the other CPR tool components.

□ Ubuntu Operating System installed in all the VMs used for the deployment.

□ Docker installed in all the VMs used for the deployment and administrator permissions (e.g. with sudo).

□ MISP instance deployed and accessible from the VM where the Threat Intelligence Module will be deployed.

□ MISP API Key corresponding to a user with the role admin.

□ Keycloak installed and accessible from the VMs where the CPR modules will be deployed.

2.3.3 Operational Prerequisites

2.3.3.1 The content and format of logs for successful anomaly detection

The LOMOS tool is designed to process the human-readable logs. Typically, they are written by the developers and present an intimate knowledge about the application and its state. They are, however, very often discarded and thus not used in assessing the operation of the application.

Another nuance to consider is the origin of the logs. Namely, human readable logs (written to the standard output, which developers standardly use – consider the use of standard tools like Log4J and its derivatives), depict what is happening inside of the application. The more structured, machine-readable logs, typically originate from the outside source, a component or sensor inside the network can sense and sometimes explain the behaviour of the application from the outside. This means that for the anomaly to be detected it is much more useful to follow the logs, emanating from the application itself, from the internal vantage point.

During practical testing (the execution of the pilots), it was identified that the end-users (CI, their CISOs, etc.) understand the meaning of logs quite differently. The following section will explain what word “logs” means in the context of assessing the application’s state with LOMOS.

2.3.3.2 Logs must be human-readable

The LOMOS tool was designed to process human-readable logs. The counterpart to human-readable logs are more structured, machine-readable logs. Machine-readable logs are more suitable for machine processing, and tools for their use and analysis, exist. **LOMOS focuses on human-readable logs only.**

The human-readable logs represent information that is very frequently neglected, but it also represents the actual operation and the actual state, which are messages from the creators of the application. Creators of the application will often write a message (log line) if an unusual situation is detected.

To illustrate what each log type is, two examples are shown below.

Human-readable:

```
2024-03-04T09:58:49.796963Z 0 [Note] Found ca.pem, server-cert.pem and server-key.pem in data directory.
Trying to enable SSL support using them.

2024-03-04T09:58:49.797517Z 0 [Note] Skipping generation of SSL certificates as certificate files are present in
data directory.
```

Machine-readable:

```
{ "timestamp": "2024-03-07T11:07:27+00:00", "remote_addr": "172.18.0.10", "body_bytes_sent": 5,
"request_time": 0.000, "response_status": 200, "request": "POST /broker_auth/api/v2/auth/rabbitmq/topic/
HTTP/1.1", "request_method": "POST", "host": "rproxy","upstream_cache_status": "HIT","upstream_addr": "-
","http_x_forwarded_for": "-","http_referrer": "-", "http_user_agent": "-" }
```

The human-readable logs are generated by an application when the programmer decides there is something interesting to show, as this detail might help later. Typically, the main part of the content is text in natural language.

This is different to structured logs. The example above is from a JSON file; other common formats are CSV or TSV files. The amount of natural language is small. In the example above, there is no natural language in JSON values (only JSON keys are in natural language).

2.3.3.3 Logs should contain a timestamp with high resolution

LOMOS searches for anomalies by looking at N sequential log lines. This requires preservation of the time and the order in which the log lines were generated.

The used database can sort log lines using timestamp, which requires timestamp to be present. If two log lines have an identical timestamp, we cannot tell which line was first. This problem usually happens if a timestamp with small resolution (like 1 second or 1 millisecond) is used - the application can easily generate more than 1 log line in a single millisecond. The application logging should thus be configured with a high-resolution timestamp (microsecond or nanosecond).

A possible workaround is to artificially extend the timestamp if the application cannot easily be reconfigured with a high-resolution timestamp. We do this only if logs were exported to a file, where the order the in the file is the same as the order in time. In this case, three messages with the timestamp “20240102T11:22:33” will be stored as “20240102T11:22:33.000000”, “20240102T11:22:33.000001” and “20240102T11:22:33.000002”. This is difficult to use with log shipping applications (FileBeat sending logs to ElasticSearch, or OpenTelemetryCollector agent). Therefore, it has limited use and will likely not work with automated real-time log shipping.

2.3.4 Anomaly Detection Module

The source code, the built docker images and the docker compose based deployment scripts for LOMOS are stored in the XLB internal GitLab repository. For the piloting, planned at M21 and M22 of the project, a copy of the deployment scripts and the access token for built docker images are required. This means the piloting (i.e. installation of the module at the CI’s premises, at their infrastructure) is very straightforward and does not require any significant changes of the code or significant installation procedures. Optionally XLB can also store a copy of the built docker images to the CI operator controlled docker image registry.

2.3.4.1 Hardware requirements

LOMOS requires a GPU for training the neural network (transformer) model used for anomaly detection. GPUs are not mandatory for inference; however, they significantly improve the throughput. The hardware requirements depend on the mode:

Training mode:

- 8 core CPU.
- 64 GB RAM.
- 1x GPU (11 GB+ VRAM).

Inference mode:

- 8 core CPU.
- 32 GB RAM.

There should be enough disk space to store at least three copies of log data, to enable normal system operation.

2.3.4.2 Installation procedure

The system has two configuration files, one for secrets, and another one for the rest of the parameters. The secrets are located in a separate file to enable sharing and storing non-sensitive configuration parameters in a repository. An example of a non-sensitive parameter configuration file is below:

```
# Configuration

SERVER_IP=10.44.18.214

# Docker

DOCKER_RESTART=unless-stopped

DOCKER_LOGGING_MAX_SIZE=5g

DOCKER_LOGGING_MAX_FILE=4

# Celery

CELERY_BROKER_URL=redis://${SERVER_IP}:6379

CELERY_RESULT_BACKEND=redis://${SERVER_IP}:6379

CELERY_REDIS_SCHEDULER_URL=redis://${SERVER_IP}:6379

REDBEAT_REDIS_URL=redis://${SERVER_IP}:6379

FLOWER_PORT=5555

# Flask

LOMOS_CONTROLLER_URL=http://flask:25000

# Grafana

GRAFANA_URL=https://${GRAFANA_USER}:${GRAFANA_PASS} @${SERVER_IP}/grafana

GRAFANA_URL_NO_CREDENTIALS=https://${SERVER_IP}/grafana

GRAFANA_BACKEND_URL=https://${SERVER_IP}/grafana-backend

GRAFANA_BACKEND_API_PORT=25001

# FTP

FTP_PUBLIC_HOST=${SERVER_IP}

FTP_20=20

FTP_21=21

# MLflow

MLFLOW_PORT=5000

MLFLOW_TRACKING_URI=http://${SERVER_IP}:${MLFLOW_PORT}

# LOMOS

INFERENCE_DEVICE=cpu

LOMOS_PARSING_VERSION=

LOMOS_ANOMALY_DETECTION_VERSION=

LOMOS_ELASTIC_VERSION=

# Nginx

HTTP_PORT=80
```

```
HTTPS_PORT=443
API_PORT=25000
# Elasticsearch
ELASTICSEARCH_PORT=9200
ELASTICSEARCH_URL=http://${SERVER_IP}:${ELASTICSEARCH_PORT}
KIBANA_PORT=5601
KIBANA_URL=http://${SERVER_IP}:${KIBANA_PORT}
```

To deploy the system, execute the *deploy.sh* bash script. The script sets up the environment and starts the system with “docker compose”. The components are grouped into three groups. The first one is the actual LOMOS – the controller, workers, celery services, and dashboard. The second one is the model registry – MLflow with FTP server and Postgres database. The third group consists of Elasticsearch and Kibana. The script can start a new Elasticsearch and model registry deployment, or hook to an existing one.

2.3.5 Threat Intelligence Module

2.3.5.1 Hardware requirements

The suggested hardware requirements are as follows:

- CPUs: 4 or more.
- RAM: 6GB or more.
- Disk space: 64GB or more.

2.3.5.2 Installation procedure

The deployment of TINTED is composed of the following images:

- **tinted-orchestrator**: This is in charge of the communication among the different components of the architecture and interact with the MISP instance.
- **tinted-privacy**: This is in charge of the anonymization of sensitive attributes.
- **tinted-mongodb**: This is the MongoDB database used by TINTED to achieve persistence (e.g. to maintain a register of the mapping between anonymized events and the original ones when the creator organization has been anonymized, for auditing purposes).
- **tinted-app**: This is in charge of the TINTED WEB-GUI.
- **tie-heuristicengine-api**: This is in charge of processing the incoming IoCs and calculation of the different scores.
- **tinted-zmq_client**: The ZeroMQ client to receive incoming events sent through the ZeroMQ server and is configured in the MISP instance.

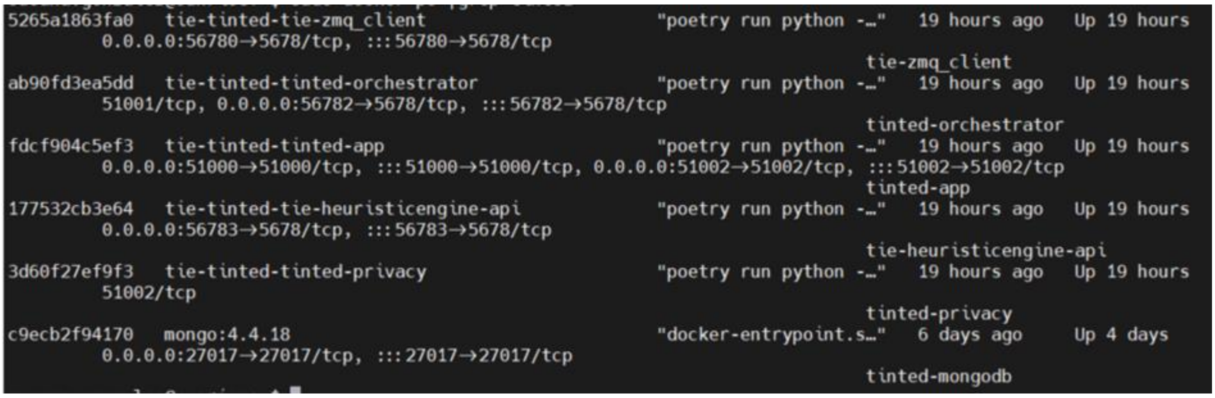

```
5265a1863fa0   tie-tinted-tie-zmq_client                "poetry run python -…"   19 hours ago   Up 19 hours
        0.0.0.0:56780→5678/tcp, :::56780→5678/tcp
                                                                                  tie-zmq_client
ab90fd3ea5dd   tie-tinted-tinted-orchestrator           "poetry run python -…"   19 hours ago   Up 19 hours
        51001/tcp, 0.0.0.0:56782→5678/tcp, :::56782→5678/tcp
                                                                                  tinted-orchestrator
fdcf904c5ef3   tie-tinted-tinted-app                    "poetry run python -…"   19 hours ago   Up 19 hours
        0.0.0.0:51000→51000/tcp, :::51000→51000/tcp, 0.0.0.0:51002→51002/tcp, :::51002→51002/tcp
                                                                                  tinted-app
177532cb3e64   tie-tinted-tie-heuristicengine-api       "poetry run python -…"   19 hours ago   Up 19 hours
        0.0.0.0:56783→5678/tcp, :::56783→5678/tcp
                                                                                  tie-heuristicengine-api
3d60f27ef9f3   tie-tinted-tinted-privacy                "poetry run python -…"   19 hours ago   Up 19 hours
        51002/tcp
                                                                                  tinted-privacy
c9ecb2f94170   mongo:4.4.18                             "docker-entrypoint.s…"   6 days ago     Up 4 days
        0.0.0.0:27017→27017/tcp, :::27017→27017/tcp
                                                                                  tinted-mongodb
```

Figure 9: TINTED Docker images.

The environment is configured in file *.production.env.* In particular, the following variables must be reviewed and configured before proceeding with the deployment:

- MISP_URL and MISP_ADMIN_API_KEY: To configure access to MISP Instance. Please note that the API Key used to interact with MISP needs to belong to a user with administration permissions (role "admin").
- ZMQ_ADDRESS and ZMQ_PORT: To configure access to the ZeroMQ Server
- FEEDING_MODE: To configure if the process will receive feeds from MISP (set up this variable to "MISP"), directly invoking the API provided by the Intelligence Feed sources (set up this variable to "API") or both of them (set up this variable to "ANY").

  NOTE: In case of no external access, please configure this variable to "MISP".
- INTELLIGENCE_FEEDS_FILE: This is the file used to configure the intelligence feeds enabled for scoring. In the deployment this file is mounted as a volume so it can be more easily modified.
- CIRCL_ORG_FEEDS: This environment variable includes the list of creator organizations (separated by commas) of MISP events shared with the CIRCL Feed. For example:

```
CIRCL_ORG_FEEDS="CIRCL,CthulhuSPRL.be,Crimeware,Synovus Financial,CERT-
RLP,CERT-FR_1510,SCTIF,Centre for Cyber security Belgium,MalwareMustDie,The
DFIR Report,wilbursecurity.com,laskowski-tech.com,ESET,THA-CERT,BSK,VK-
Intel,EUROLEA,MiSOC,Hestat"
```

- ENABLE_WHITELIST: Set up to "true" if you want to enable a whitelist overlap score. Please note that this score requires you to have external access to download Ips and domain whitelists. The URLs with these whitelists are configured in WHITELIST_IP_URL and WHITELIST_DOMAIN_URL. By default, these lists will be downloaded when the image is deployed and then at 07:00h. If you want to change the download time, you need to configure WHITELIST_DOWNLOAD_TIME (not enclosed between double quotes). NOTE: In case of no external access, please configure this variable to "False".
- EVENT_SOURCE_ANONYMIZATION: Set up to "true" if you want to anonymize the creator organization of the original event.
- ANONYMIZED_ORG: This variable will have the name of the MISP Organization that will be used when EVENT_SOURCE_ANONYMIZATION=True to replace the original event creator organization.
- REMOVE_ANONYMIZATION_DUPLICATE: To remove the new anonymized event (when the tag data-classification:sensitive-information is present) from the local MISP instance once it has been shared, maintaining only the original one locally.
- ENABLE_KAFKA_OUTPUT: To share MISP events received via Kafka. This needs to be enabled to communicate with the CPR Dashboard.
- COMM_THRESHOLD: This is the score threshold that will be used to send only events with higher value to the Kafka queue.
- KAFKA_ADDRESS: The address where the Kafka server is running

IMPORTANT. File *config/intelligence_feeds.json* must be reviewed before the deployment to enable (set **enabled** to *true*) the intelligence feed sources that want to be considered. It is important to check that the **name** used to identify an intelligence feed in this file is the same configured in the MISP instance for the same feed (if it is included as a MISP feed) as shown in the Event **Info** field, or be part of it in the case of feeds producing multiple events with the same prefix. For example, "name=[AEGIS]" means that an event with the info field "[AEGIS] Weekly Mirai Activity Report – 2024-05-05" will be processed and if the event info field is "[AEGIS] Weekly Trickbot Activity Report – 2024-05-05". The **periodicity** and **interval** fields of the enabled intelligence feeds APIs should be also reviewed. They indicate how often the API will be invoked (e.g. once every 24 hours). More information about this file *intelligence_feeds.json* is provided in Annex section 3.

Follow the steps below for the deployment:

1. Check the environment configuration in file .production.env
2. Check that you have enabled the feeds you are going to consider in file config/intelligence_feeds.json. For example:

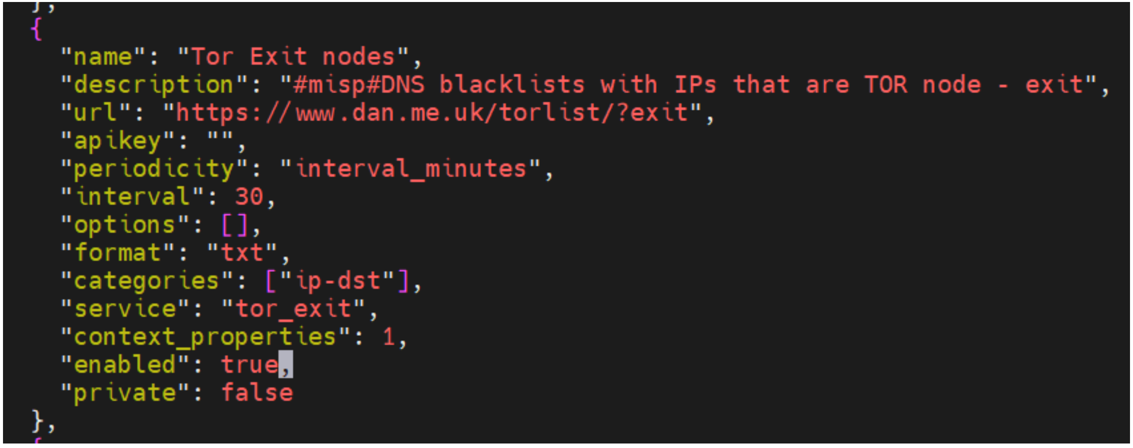

```
{
  "name": "Tor Exit nodes",
  "description": "#misp#DNS blacklists with IPs that are TOR node - exit",
  "url": "https://www.dan.me.uk/torlist/?exit",
  "apikey": "",
  "periodicity": "interval_minutes",
  "interval": 30,
  "options": [],
  "format": "txt",
  "categories": ["ip-dst"],
  "service": "tor_exit",
  "context_properties": 1,
  "enabled": true,
  "private": false
},
```

Figure 10. Example of source feed data.

1. When using encryption, include the configuration for Keycloak in the file docker/app/controller/keycloak_controller.py
2. Configure information about Kafka broker and topics in the file defined in KAFKA_CONFIG env variable (by default, the kafka.conf file under the folder docker/orchestrator)
3. Execute the following command to deploy the images:

```
sudo docker compose -f docker-compose.yml –env-file .production.env up -d
```

NOTE: Please note that in case you change the sources enabled in the intelligence_feeds.json file or the list of organizations in the environment variable CIRCL_ORG_FEEDS, you just need to restart the tie-zmq_client (tinted-scoring in TICX v0.6.5) container to consider the updated list of threat intelligence sources:

```
sudo docker restart tie-zmq_client
```

2.3.5.3 Installation verification

Once deployed, the user can check the logs generated by the containers with the following command:

```
sudo docker ps -a
sudo docker logs -f <container_name>
```

If successful, in the tie-zmq_client logs (see Figure 11), it will be indicated that the service is listening for messages from the ZeroMQ queue and (if enabled) the initialization of Whitelists.

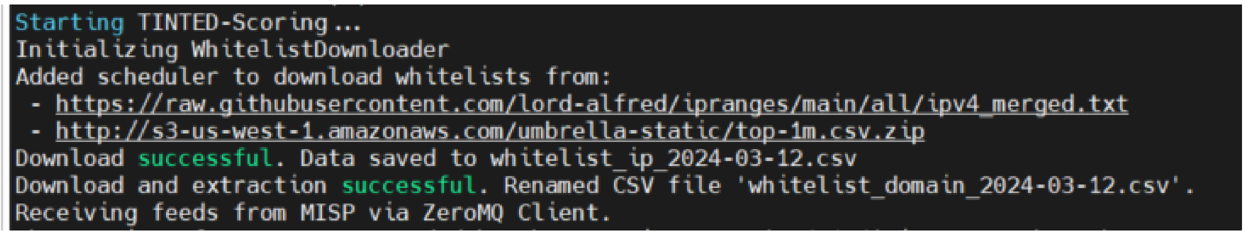

```
Starting TINTED-Scoring...
Initializing WhitelistDownloader
Added scheduler to download whitelists from:
  - https://raw.githubusercontent.com/lord-alfred/ipranges/main/all/ipv4_merged.txt
  - http://s3-us-west-1.amazonaws.com/umbrella-static/top-1m.csv.zip
Download successful. Data saved to whitelist_ip_2024-03-12.csv
Download and extraction successful. Renamed CSV file 'whitelist_domain_2024-03-12.csv'.
Receiving feeds from MISP via ZeroMQ Client.
```

Figure 11: TINTED Scoring Docker log - Initialization.

To test the scoring functionality, you can open one of the existing events corresponding to one of the feeds you have enabled in the file intelligence_feeds.json (e.g. "Tor Exit nodes", checking you have the same name of the event in the field “name” of the intelligence_feeds.json file) and click on “Publish event to ZMQ” as shown in Figure 12:

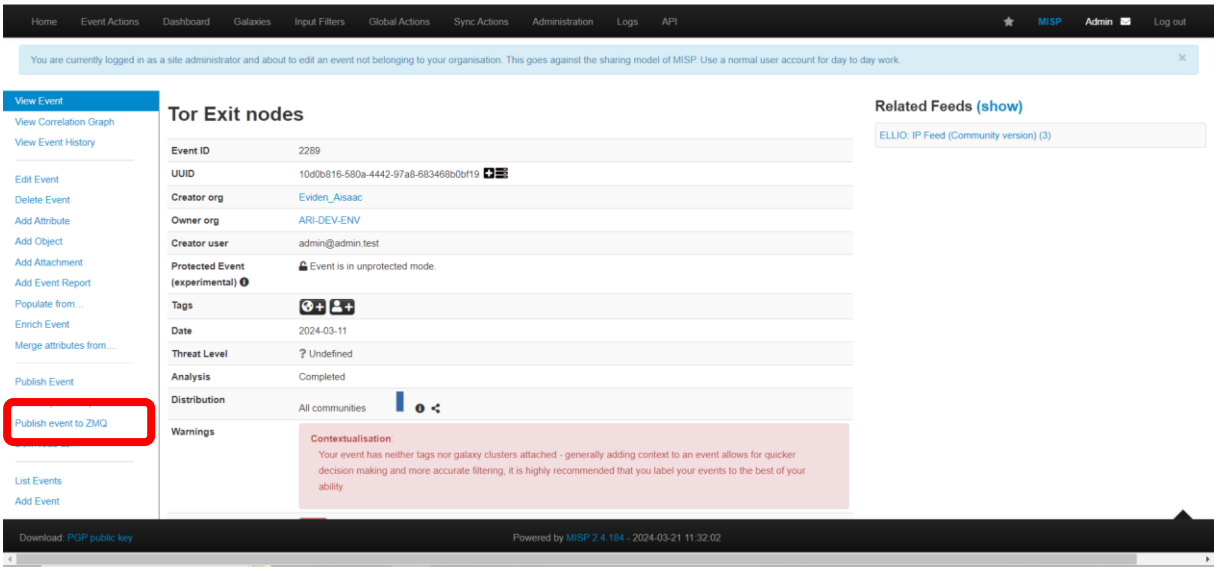

You are currently logged in as a site administrator and about to edit an event not belonging to your organisation. This goes against the sharing model of MISP. Use a normal user account for day to day work.

Tor Exit nodes

Event ID	2289
UUID	10d0b816-580a-4442-97a8-683468b0bf19
Creator org	Eviden_Aisaac
Owner org	ARI-DEV-ENV
Creator user	admin@admin.test
Protected Event (experimental)	Event is in unprotected mode.
Tags	
Date	2024-03-11
Threat Level	? Undefined
Analysis	Completed
Distribution	All communities
Warnings	Contextualisation: Your event has neither tags nor galaxy clusters attached - generally adding context to an event allows for quicker decision making and more accurate filtering, it is highly recommended that you label your events to the best of your ability.

Related Feeds (show)

ELLIO: IP Feed (Community version) (3)

Download: PGP public key

Powered by MISP 2.4.184 - 2024-03-21 11:32:02

Figure 12: Example to publish event to ZMQ

When a new feed is received (remember that only those enabled in the *intelligence_feeds.json* file will be considered), it will appear in the log, as well as the IoCs to be processed for that feed (see Figure 13).

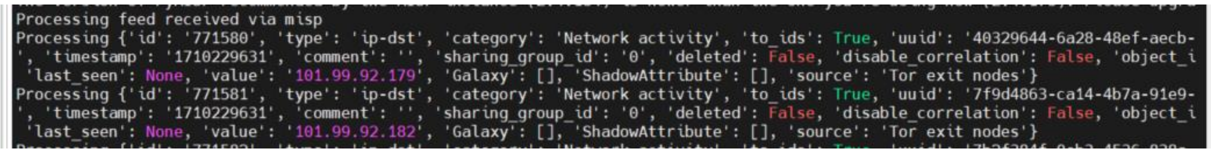

```
Processing feed received via misp
Processing {'id': '771580', 'type': 'ip-dst', 'category': 'Network activity', 'to_ids': True, 'uuid': '40329644-6a28-48ef-aecb-
', 'timestamp': '1710229631', 'comment': '', 'sharing_group_id': '0', 'deleted': False, 'disable_correlation': False, 'object_i
 'last_seen': None, 'value': '101.99.92.179', 'Galaxy': [], 'ShadowAttribute': [], 'source': 'Tor exit nodes'}
Processing {'id': '771581', 'type': 'ip-dst', 'category': 'Network activity', 'to_ids': True, 'uuid': '7f9d4863-ca14-4b7a-91e9-
', 'timestamp': '1710229631', 'comment': '', 'sharing_group_id': '0', 'deleted': False, 'disable_correlation': False, 'object_i
 'last_seen': None, 'value': '101.99.92.182', 'Galaxy': [], 'ShadowAttribute': [], 'source': 'Tor exit nodes'}
```

Figure 13: TINTED Scoring Docker log – IoC Processing.

NOTE: If the event is not processed, please check that it is not included in the “Event Blocklist” (under the “Administration” tab). For example, this may happen if you have previously removed it manually. If the event is blocklisted, just remove it before publishing.

In the tinted-orchestrator logs, it will show the different requests received from the tinted-scoring image and if they are correctly processed (see Figure 14).

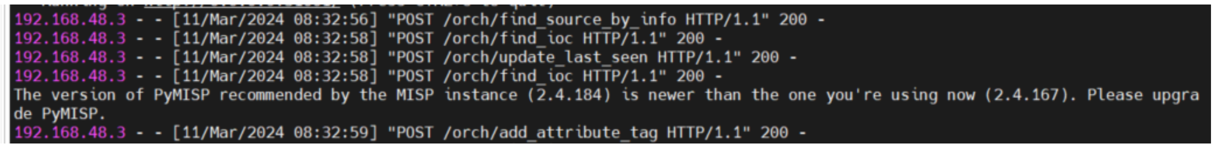

```
192.168.48.3 - - [11/Mar/2024 08:32:56] "POST /orch/find_source_by_info HTTP/1.1" 200 -
192.168.48.3 - - [11/Mar/2024 08:32:58] "POST /orch/find_ioc HTTP/1.1" 200 -
192.168.48.3 - - [11/Mar/2024 08:32:58] "POST /orch/update_last_seen HTTP/1.1" 200 -
192.168.48.3 - - [11/Mar/2024 08:32:58] "POST /orch/find_ioc HTTP/1.1" 200 -
The version of PyMISP recommended by the MISP instance (2.4.184) is newer than the one you're using now (2.4.167). Please upgra
de PyMISP.
192.168.48.3 - - [11/Mar/2024 08:32:59] "POST /orch/add_attribute_tag HTTP/1.1" 200 -
```

Figure 14: TINTED Orchestrator Docker log – POST Requests

If everything works as expected, once you update the event in the MISP Dashboard, it will appear with the Tags “TINTED:Scource-Score” and “admiralty-scale:source-reliability” as shown in Figure 15.

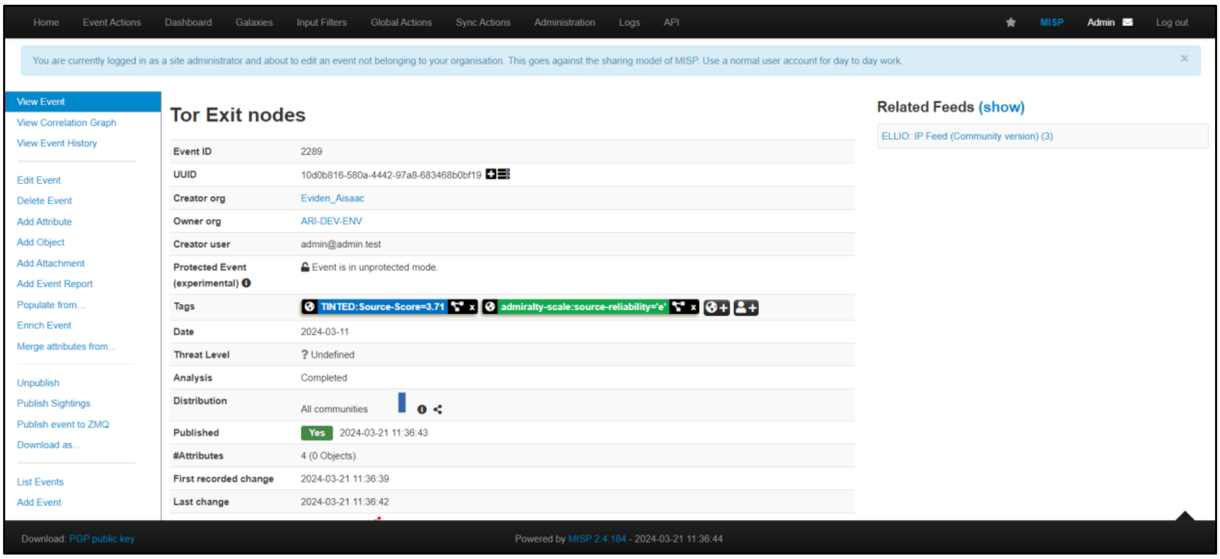

You are currently logged in as a site administrator and about to edit an event not belonging to your organisation. This goes against the sharing model of MISP. Use a normal user account for day to day work.

Tor Exit nodes

Event ID	2289
UUID	10d0b816-580a-4442-97a8-683468b0bf19
Creator org	Eviden_Aisaac
Owner org	ARI-DEV-ENV
Creator user	admin@admin.test
Protected Event (experimental)	Event is in unprotected mode.
Tags	TINTED:Source-Score=3.71 x admiralty-scale:source-reliability='e' x
Date	2024-03-11
Threat Level	? Undefined
Analysis	Completed
Distribution	All communities
Published	Yes 2024-03-21 11:36:43
#Attributes	4 (0 Objects)
First recorded change	2024-03-21 11:36:39
Last change	2024-03-21 11:36:42

Related Feeds (show)

ELLIO: IP Feed (Community version) (3)

Download: PGP public key

Powered by MISP 2.4.184 - 2024-03-21 11:36:44

Figure 15: Example of event with source scoring Tags

To test the anonymization functionality, edit one of the MISP events and add to one of the IoCs with a sensitive ip address and a new tag "**data-classification:sensitive-information**" from the taxonomy library "data-classification" as shown in Figure 16. Then click on "Publish event to ZMQ", as shown in Figure 12.

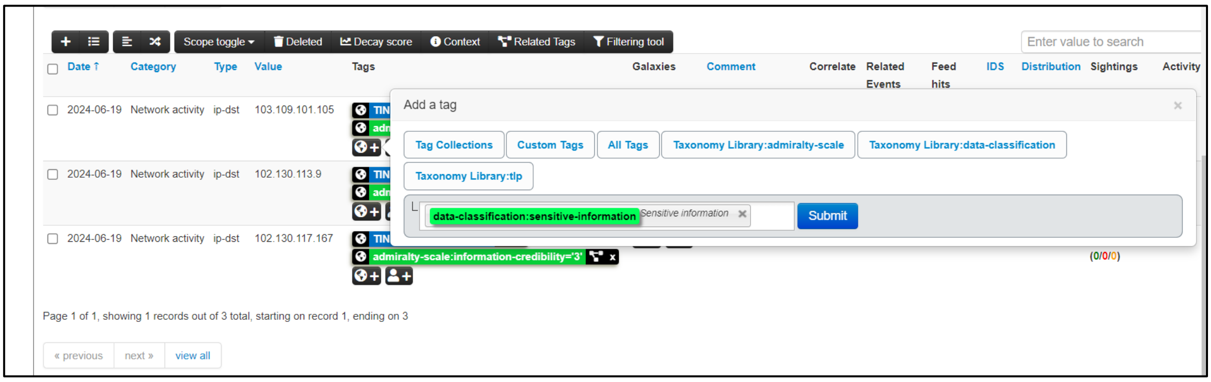

☐	Date ↑	Category	Type	Value	Tags	Galaxies	Comment	Correlate	Related Events	Feed hits	IDS	Distribution	Sightings	Activity
☐	2024-06-19	Network activity	ip-dst	103.109.101.105	TIN adr									
☐	2024-06-19	Network activity	ip-dst	102.130.113.9	TIN adr									
☐	2024-06-19	Network activity	ip-dst	102.130.117.167	TIN admiralty-scale:information-credibility='3' x								(0/0/0)	

Add a tag ×

Tag Collections
Custom Tags
All Tags
Taxonomy Library:admiralty-scale
Taxonomy Library:data-classification
Taxonomy Library:tlp

data-classification:sensitive-information Sensitive information ✕

Page 1 of 1, showing 1 records out of 3 total, starting on record 1, ending on 3

Figure 16: Data-classification tag for anonymization of sensitive data

NOTE: If this taxonomy is not available in the MISP instance, go to the list of taxonomies (under Event Actions menu), filter by "data-classification" and enable it (see Figure 17).

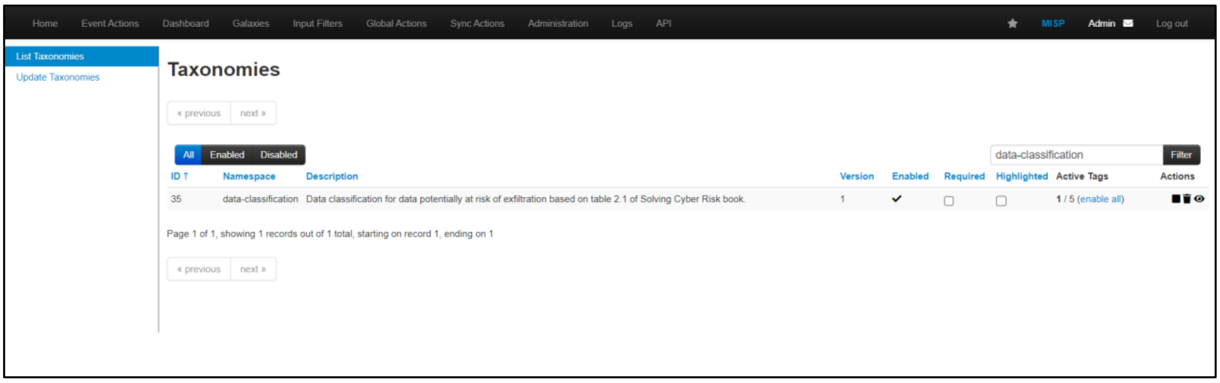

Taxonomies

data-classification

ID ↑	Namespace	Description	Version	Enabled	Required	Highlighted	Active Tags	Actions
35	data-classification	Data classification for data potentially at risk of exfiltration based on table 2.1 of Solving Cyber Risk book.	1	✓	☐	☐	1 / 5 (enable all)	

Page 1 of 1, showing 1 records out of 1 total, starting on record 1, ending on 1

Figure 17: data-classification taxonomy

In the tinted-privacy logs, it will show the different requests received for anonymization and if they are correctly processed (see Figure 18).

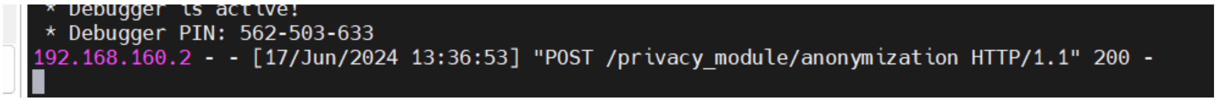

```
 * Debugger is active!
 * Debugger PIN: 562-503-633
192.168.160.2 - - [17/Jun/2024 13:36:53] "POST /privacy_module/anonymization HTTP/1.1" 200 -
```

Figure 18: TINTED Privacy Docker log – POST Requests

If successful, once the new event is published, it will appear with the anonymized sensitive attributes as shown in Figure 19

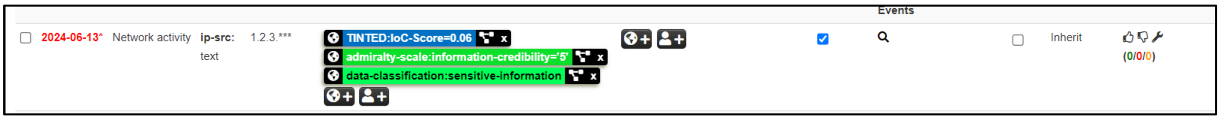

Events

☐ 2024-06-13* Network activity ip-src: 1.2.3.*** text TINTED:IoC-Score=0.06 x admiralty-scale:information-credibility="5" x data-classification:sensitive-information x ☑ ☐ Inherit (0/0/0)

Figure 19: Example of sensitive data anonymized

2.3.5.4 Keycloak configuration

Keycloak is used to provide the authentication layer to TINTED. You will need to create a Realm associated with the CPR Threat Intelligence Sharing Platform (e.g. “cisp”) and a client associated with the CPR TINTED WEB-GUI (Figure 20):

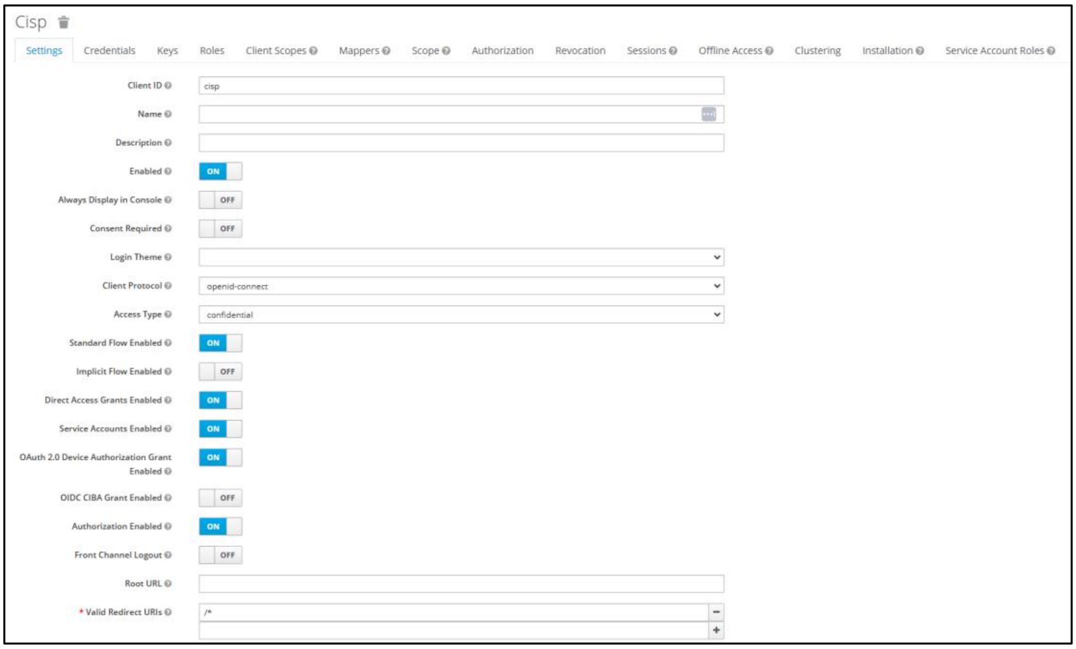

Cisp

Client ID: cisp
Name:
Description:
Enabled: ON
Always Display in Console: OFF
Consent Required: OFF
Login Theme:
Client Protocol: openid-connect
Access Type: confidential
Standard Flow Enabled: ON
Implicit Flow Enabled: OFF
Direct Access Grants Enabled: ON
Service Accounts Enabled: ON
OAuth 2.0 Device Authorization Grant Enabled: ON
OIDC CIBA Grant Enabled: OFF
Authorization Enabled: ON
Front Channel Logout: OFF
Root URL:
* Valid Redirect URIs: /*

Figure 20. Example of Keycloak client.

Three roles must be added to that client (see Figure 21):

- Admin: the user administrator can create/read/update/delete and publish events and users of the system.
- Publisher: The publisher user can create/publish/read and share events within the system.
- Reader: The reader user only can read the events shared with him/her.

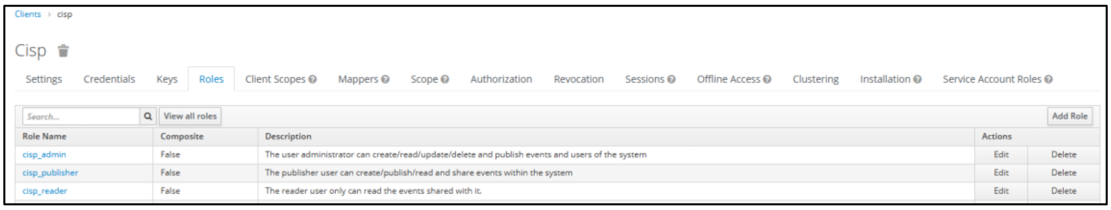

Clients › cisp

Cisp

Role Name	Composite	Description	Actions	
cisp_admin	False	The user administrator can create/read/update/delete and publish events and users of the system	Edit	Delete
cisp_publisher	False	The publisher user can create/publish/read and share events within the system	Edit	Delete
cisp_reader	False	The reader user only can read the events shared with it.	Edit	Delete

Figure 21. Keycloak client roles.

Create Mappers for the client (see Figure 22) including two attributes:

- mispTag (type “User Attribute”)
- passphrase (type “User Attribute”)

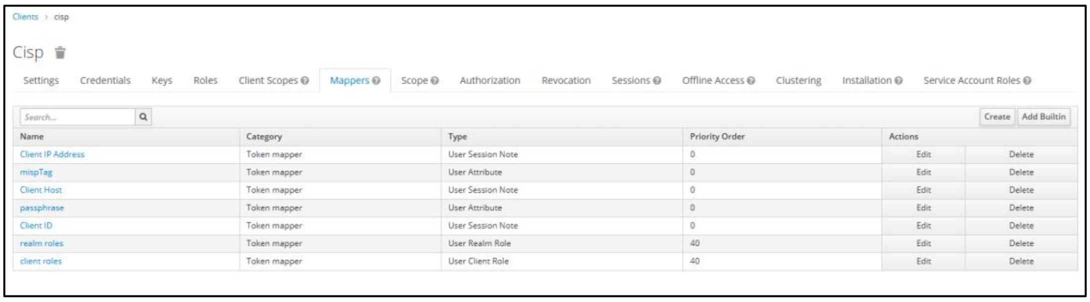

Name	Category	Type	Priority Order	Actions	
Client IP Address	Token mapper	User Session Note	0	Edit	Delete
mispTag	Token mapper	User Attribute	0	Edit	Delete
Client Host	Token mapper	User Session Note	0	Edit	Delete
passphrase	Token mapper	User Attribute	0	Edit	Delete
Client ID	Token mapper	User Session Note	0	Edit	Delete
realm roles	Token mapper	User Realm Role	40	Edit	Delete
client roles	Token mapper	User Client Role	40	Edit	Delete

Figure 22. Keycloak client mappers.

They must be added to ID token, to access token and to userinfo (see Figure 23).

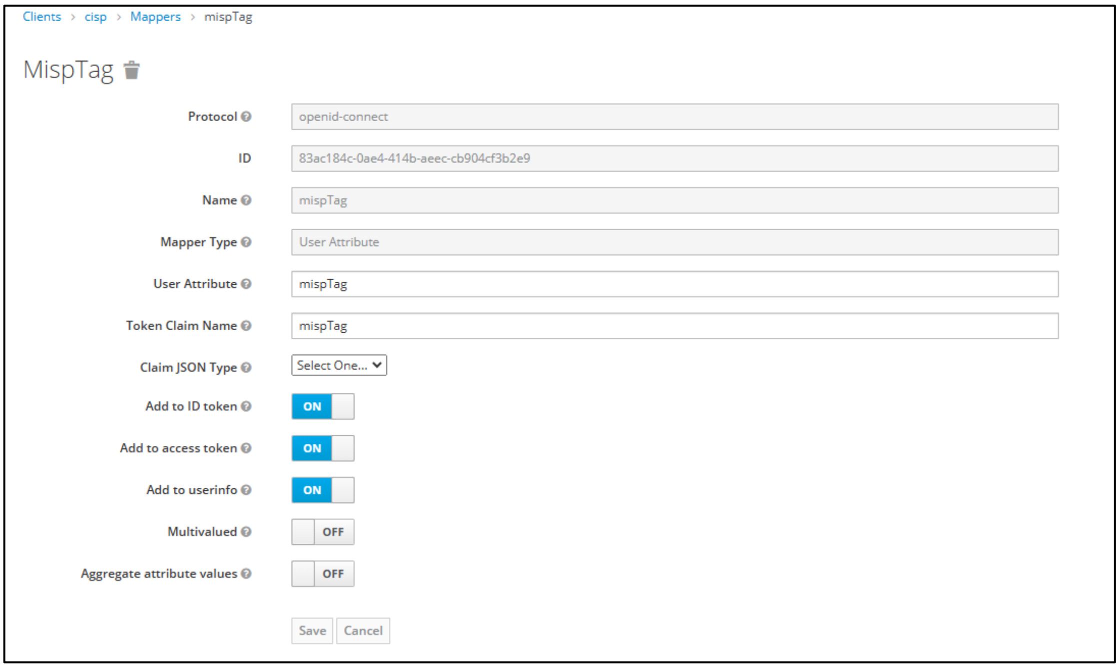

Protocol	openid-connect
ID	83ac184c-0ae4-414b-aeec-cb904cf3b2e9
Name	mispTag
Mapper Type	User Attribute
User Attribute	mispTag
Token Claim Name	mispTag
Claim JSON Type	Select One...
Add to ID token	ON
Add to access token	ON
Add to userinfo	ON
Multivalued	OFF
Aggregate attribute values	OFF

Figure 23. Example of mapper.

Finally, create users with the different roles assigned for data sharing (see Figure 24). For example:

- Administrator (admin)
- Alice (publisher)
- Bob (reader)

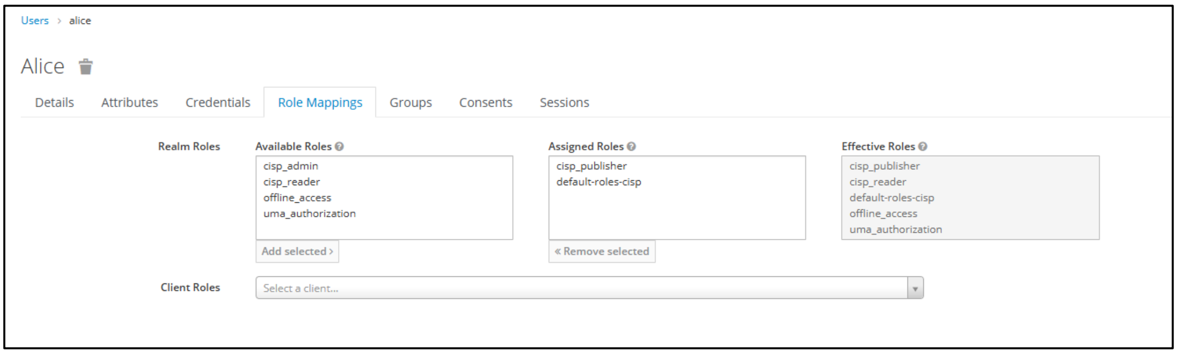

Realm Roles

Available Roles: cisp_admin, cisp_reader, offline_access, uma_authorization

Assigned Roles: cisp_publisher, default-roles-cisp

Effective Roles: cisp_publisher, cisp_reader, default-roles-cisp, offline_access, uma_authorization

Client Roles

Figure 24. Example of Keycloak user with role mapping.

It can be also created groups and assigned roles (see Figure 25). The user can be assigned to these groups.

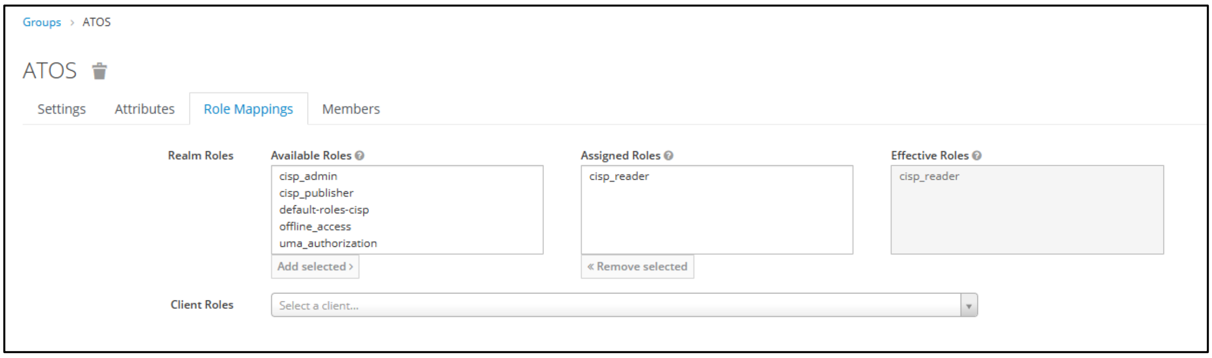

Groups › ATOS

ATOS

Realm Roles

Available Roles: cisp_admin, cisp_publisher, default-roles-cisp, offline_access, uma_authorization

Assigned Roles: cisp_reader

Effective Roles: cisp_reader

Client Roles: Select a client...

Figure 25. Example of Keycloak group with role mapping.

2.3.6 Risk Assessment Module

2.3.6.1 Hardware requirements

The suggested hardware requirements are as follows:

- CPUs: 2 or more.
- RAM: 16 GB or more.
- Disk space: 30 GB or more.

2.3.6.2 Installation procedure

To appropriately set up and deploy the Risk Assessment module, the following applications must be installed [10]:

- The Risk Assessment Engine, which serves as the central element of the tool. It is responsible for executing rules and performing risk assessments. This component is developed as a Python application with multi-threading capabilities.
- The Graphical Interface, which functions as the visualization part of the tool. It is created using the Django framework for Python. This interface offers a control dashboard to exhibit input and output data concerning the security of the target infrastructure.
- The Database, which serves as the storage element of the tool. It is built using SQL and contains data related to risk models, indicators, data processing activities, mitigation measures, and all security-related information about the infrastructure and its components.
- The Message broker, which is responsible for facilitating communication among various components using a RabbitMQ server.
- The Nginx load balancer.

A Docker container is created for each application: Dashboard (Django), Engine (Python), message broker (RabbitMQ), load balancer (Nginx), and the database server (PostgreSQL). These containers are initiated with specific port bindings on the host system, enabling access to internal application components (such as the web application or the broker) from the external network.

Immediately after launching the Dashboard web application, a script is run to populate the use case data. This data is sourced from JSON files located within the directory "*rae_dashboard/rae_dashboard/dashboard/db_test_data/*". This script, in turn, utilizes an internal Python function called "*load_initial_data()*" to parse and load the existing JSON files into the database tables. The PostgreSQL database is stored in a dedicated Docker container.

Essentially, each of the three application directories encompasses source code, configuration files, and a Docker file. This Docker file defines the base image, dependencies, source deployment, and a startup script (commonly referred to as an entry point) for the respective application. To simplify the creation and execution of the Docker containers, a configuration file named "docker-compose.yml" is included in the root directory of the code. This file facilitates building the containers, launching them, or stopping the services:

```
sudo docker-compose up --build --detach --force-recreate
sudo docker-compose down # to stop the running containers
```

Additionally, the accompanying script "*manage_docker_compose.sh*" can achieve the same outcomes:

```
chmod 755 manage_docker_compose.sh
sudo ./manage_docker_compose.sh up build
sudo ./manage_docker_compose.sh down
```

Creating container images from scratch might take some time due to dependency installation. However, these images will be cached at the layer level, leading to quicker recreation of containers in the future. Furthermore, cached images can be exported or uploaded to third-party repositories, like a Docker registry. To obtain a list of currently cached images:

```
sudo docker images
```

Alternatively, a config file named "*docker-compose-images.yml*" has been provided. This script relies solely on existing system images and constructs and launches containers using those images, requiring no source files:

```
sudo docker-compose -f docker-compose-images.yml up --detach
```

2.3.7 Incident Reporting Module

2.3.7.1 Hardware requirements

The Incident Reporting module requires the deployment of TheHive and Cortex. Consequently, hardware requirements for these tools [24] (which also depends on the number of concurrent users) are the same ones required for the CPR tool, that for less than 10 users and deploying all services on the same server are:

- CPUs: 8 cores or more.
- RAM: 16GB or more.

2.3.7.2 Installation Procedure

To deploy AIRE, it is necessary to have a folder **aire-deploy** with the subfolders containing configuration files for the different services and the jar files of the AIRE services (aire-reports-generator-1.0.1-SNAPSHOT.jar, aire-thehive-plugin-1.0.1-SNAPSHOT.jar and aire-workflow-enforcement-1.0.1-SNAPSHOT.jar).

The following steps are required for its installation:

- Copy aire-deploy/aire_dashboard/.envs/.env.docker_run file as aire-deploy/aire_dashboard/.envs/.env and configure it. In particular:
  - configure the DJANGO_ALLOWED_HOSTS to allow access to server from external machines
  - configure information to access to Keycloak (if necessary)
  - configure information about the database
  - configure AIRE_SERVER and THEHIVE_URL endpoints
- Set up execution permissions for AIRE Dashboard **entrypoint.sh** script:

```
# chmod 755 aire-deploy/aire_dashboard/entrypoint.sh
```

- Configuration for TheHive deployment:

  In the AIRE docker-compose.yml file, the deployment of TheHive v4, Cortex and their dependencies is already included, but it is necessary to do the following configuration before deploying them.

  - Generate secret play key for the Hive and copy it into play.http.secret.key in the file aire-deploy/TheHive/application.conf.d/secret.conf
  - Generate keystore file (.p12) using LetsEncrypt Certificate and copy it in folder aire-deploy/TheHive/ssl:

```
# openssl pkcs12 -export -in fullchain.pem -inkey privkey.pem -out sunrise_keystore.p12 -name tomcat -CAfile chain.pem -caname root
```

  - (If necessary) configure the connection with Keycloak in aire-deploy/TheHive/application.conf file.
  - The default max virtual memory areas parameter (vm.max_map_count [65530]) is too low for Elasticsearch, and it is necessary to increase it to at least [262144]. We need to do this in the host machine before deploying the elasticsearch docker image if we want to have it there. Edit /etc/sysctl.conf in the VM to insert "vm.max_map_count = 262144" and execute:

```
#sysctl -p
```

  More details about the installation and configuration of TheHive can be found at: https://docs.strangebee.com/thehive/installation/step-by-step-installation-guide/. Please note that it is necessary to install TheHive v4 since the webhook capabilities are required for the correct working of AIRE.

- Configuration for AIRE services:
  - SSL Certificate generated will be used by TheHive but also by aire-thehive-plugin to access TheHive. Copy it in the folder aire-deploy/aire-thehive-plugin/config/keystore and configure it in the aire-deploy/aire-thehive-plugin/config/application.conf.
  - Check the mail server configuration is correct in aire-deploy/aire-reports-generator/config/application.properties
  - Check the configuration about the database is correct in the files aire-deploy/aire-reports-generator/config/application.properties, aire-deploy/aire-thehive-plugin/config/application.properties and aire-deploy/aire-workflow-enforcement/config/application.properties
- Build and deploy the containers:

```
# sudo docker compose up -d --build
```

- To complete the Cortex installation, follow the steps in https://github.com/TheHive-Project/CortexDocs/blob/master/admin/quick-start.md. This is summarized in the following points:
  - Generate a secret key and include it in the /etc/cortex/application.conf file.
  - Open the Cortex URL to ''Update Database'' and create the administrator account.
  - Login into Cortex with the admin user. Create an Organization (e.g. “atos”) to associate and configure analyzers/responders. Create an Organization Administrator (e.g. “admin_atos”). Create an account "thehive" for TheHive integration. Generate an API Key for this user. It will be used in TheHive configuration.
  - Login to Cortex with the admin_atos user to Refresh analyzers and responders and to Enable them. In the case of AIRE Responders, set up RED as TLP/PAP so it can be applied to any incident.

- To complete TheHive installation:
  - Connect to the docker image and install TheHive4py

```
# sudo docker exec -u 0 -it thehive bash

root@b68ef72e6706:/opt/thehive#  apt-get update

root@b68ef72e6706:/opt/thehive# apt-get install python3-pip

root@b68ef72e6706:/opt/thehive# pip install -U thehive4py
```

  - Change the POST_SITE and POST_AUTH in the script aire-deploy/TheHive/loadCustomFields_TheHive4.py and execute it to create the required custom fields. You need to login as an "admin" user (Member of the "admin" organisation) that has a profile including the manageCustomField permission and consequently, the API key of this user is required to create the custom fields.
  - Change the API_KEY and THEHIVE_URL in the script aire-deploy/TheHive/createUsers.sh, edit your users, and execute it to create users required. The creation of users can be also done through TheHive GUI.

    Important: You first need to create your organization (e.g.ATOS) through the GUI and create the admin user (e.g. admin@atos.net) with the permissions "orgadmin". Generate an API KEY for this admin user and use it in the script. The name of the users must be their email address (e.g. irt@atos.net). By default, if no address is included, it is assumed to be "@thehive.local".

    Default users “imt@thehive.local”, “iclt@thehive.local”, “controller@thehive.local” and “irt@thehive.local” need to be created in TheHive. The tasks created in TheHive through the AIRE asset will be assigned to them. If other users are used, then it is necessary to configure the task with them through the AIRE Dashboard and check that users are created in TheHive.

    It is also necessary to create a user “aire@thehive.local” and “aire@<your_org>” in TheHive with the admin permissions to execute actions in TheHive from AIRE asset. The API Key of the user aire@thehive.local will be included in the parameter “AIRE_api_key” in the application.properties file in the aire-thehive-plugin.

```
            # chmod 755 createUsers.sh
            # ./createUsers.sh
```

  - Activate the webhooks by executing the following script with the admin user of your organisation (with permission manageConfig):

```
# chmod 755 activeWebhooks.sh

# ./activateWebhooks.sh
```

- To connect TheHive with MISP, generate (if necessary) the truststore with the MISP certificate, for example:

```
# keytool -import -alias sunrise_misp -file misp.crt -keypass keypass -keystore sunrise_truststore.jks -storepass storepass
```

  Copy the truststore into the aire-deploy/TheHive/ssl folder and reference it in the MISP block in TheHive/application.properties:

```
play.modules.enabled += org.thp.thehive.connector.misp.MispModule
```

```
misp {
 # Interval between consecutive MISP event  imports  in  hours  (h)  or minutes (m).
 interval: 1 min
 servers: [
   {
     name = "MISP-NAME"
     url = "https://misp.url.instance"
     auth {
       type = key
       key = “your_misp_admin_user_api_key" }
     caseTemplate = "Incident Report"
    tags = [“SUNRISE1”, “SUNRISE2”]
     includedTheHiveOrganisations = ["*"]
     excludedTheHiveOrganisations = []
    #filters
    max-age = 7 days
    max-attributes = 1000
    max-size = 1 MiB }
]
```

- To complete the AIRE installation:
  - Include the API Key of the user ''aire'' in the parameter “AIRE_api_key” in the /opt/aire/config/application.properties file in the aire-thehive-plugin docker container.
  - Restart the docker containers to apply all the changes.

2.3.7.3 Installation verification

To check that the Elasticsearch has been correctly deployed, you can execute (once connected to the docker container):

```
          # curl -XGET 'http://localhost:9200/?pretty=true'
```

If everything is correctly deployed, you should see the following containers up and running:

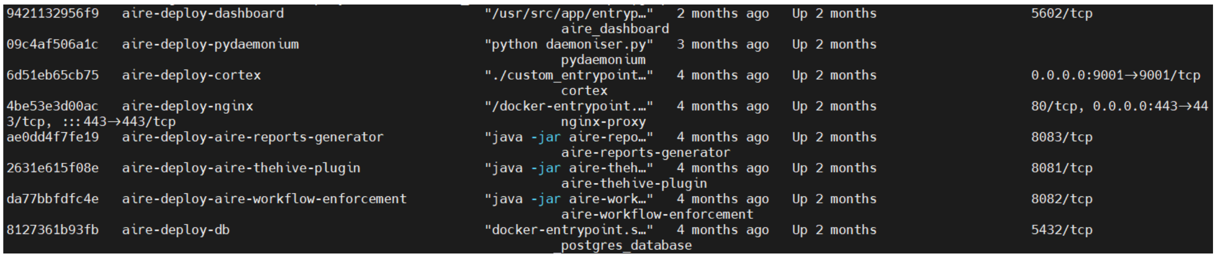

```
9421132956f9   aire-deploy-dashboard                          "/usr/src/app/entryp…"   2 months ago   Up 2 months                     5602/tcp
                                                                        aire_dashboard
09c4af506a1c   aire-deploy-pydaemonium                        "python daemoniser.py"   3 months ago   Up 2 months
                                                                        pydaemonium
6d51eb65cb75   aire-deploy-cortex                             "./custom_entrypoint…"   4 months ago   Up 2 months                     0.0.0.0:9001→9001/tcp
                                                                        cortex
4be53e3d00ac   aire-deploy-nginx                              "/docker-entrypoint.…"   4 months ago   Up 2 months                     80/tcp, 0.0.0.0:443→44
3/tcp, :::443→443/tcp                                                   nginx-proxy
ae0dd4f7fe19   aire-deploy-aire-reports-generator             "java -jar aire-repo…"   4 months ago   Up 2 months                     8083/tcp
                                                                        aire-reports-generator
2631e615f08e   aire-deploy-aire-thehive-plugin                "java -jar aire-theh…"   4 months ago   Up 2 months                     8081/tcp
                                                                        aire-thehive-plugin
da77bbfdfc4e   aire-deploy-aire-workflow-enforcement          "java -jar aire-work…"   4 months ago   Up 2 months                     8082/tcp
                                                                        aire-workflow-enforcement
8127361b93fb   aire-deploy-db                                 "docker-entrypoint.s…"   4 months ago   Up 2 months                     5432/tcp
                                                                        postgres_database
```

Figure 26. AIRE docker containers.

2.3.8 Dashboard

2.3.8.1 Hardware requirements

The CPR tool has been developed as a dockerized application, with few requirements. Serving a single administrator, who supervises the entire system, the container only needs:

- CPUs = 1
- RAM = 256 MiB
- Disk space: 2 GB

2.3.8.2 Installation Procedure

Because the CPR Tool is based on the other components, it is advisable to have the rest of the CPR modules up and running before starting the installation.

The first step is setting the URL of the different CPR components. In the CPR folder, there is one subfolder per component (LOMOS, WAZUH, TINTED, CERCA, and AIRE). All these folders have a file, called api.py, where you need to configure the URLs of each service.

- **LOMOS** (CPR_tool/LOMOS/api.py): This component only has the URL and indexes parameter. Updating them is enough.

```
url = "https://your_lomos_url"
indexes = "test_bgl2k_logs_structured"
```

- **WAZUH** (CPR_tool/WAZUH/api.py): This component only has the URL and port. Updating them is enough. It is possible to use an alias, such as, misp-local or misp_dashboard.

```
url = 'https://your_wazuh_url’
```

- **TINTED** (CPR_tool/WAZUH/api.py): This component only has the URL and port. Updating them is enough. It is possible to use an alias, such as, misp-local or misp_dashboard.

```
url = 'https://your_misp_url’
```

- **CERCA** (CPR_tool/WAZUH/api.py): Update the URL and auth keys from the CERCA service.

```
url = "https://your_cerca_url"
auth = ('your_auth_key', 'your_token')
```

- **AIRE** (CPR_tool/WAZUH/api.py): This component only has the URL and port. Updating them is enough.

```
url = 'https://thehive:9443'
```

When the URLs of the other modules are configured, you just need to build and start the container:

```
sudo docker compose build
sudo docker compose up -d
```

NOTE: The Nginx service needs to be correctly configured to have access to the different services.

By accessing the administration view (initially something like https://localhost/admin with the default Django user root:toor), you can change the root password and then configure users, groups, and pilot keys for the components (see Figure 27).

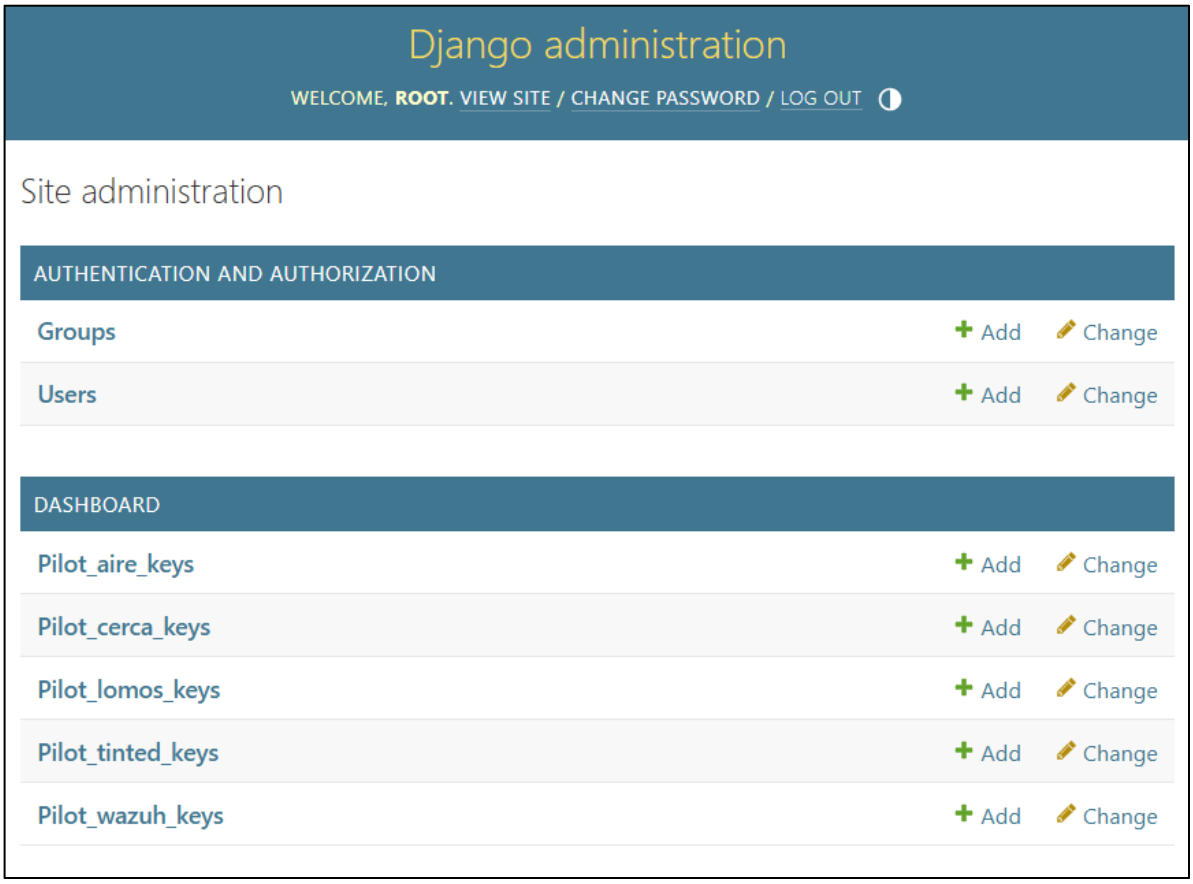

Django administration

WELCOME, **ROOT**. VIEW SITE / CHANGE PASSWORD / LOG OUT

Site administration

AUTHENTICATION AND AUTHORIZATION

Groups	✚ Add	Change
Users	✚ Add	Change

DASHBOARD

Pilot_aire_keys	✚ Add	Change
Pilot_cerca_keys	✚ Add	Change
Pilot_lomos_keys	✚ Add	Change
Pilot_tinted_keys	✚ Add	Change
Pilot_wazuh_keys	✚ Add	Change

Figure 27: CPR Dashboard – Django administration console.

You must create one group for each organization (e.g. sunrise) and as many users as necessary (e.g. admin_sunrise@sunrise.com, imt_sunrise@sunrise.com, irt_sunrise@sunrise.com, iclt_sunrise@sunrise.com, or controller_sunrise@sunrise.com). The usernames must be emails (see Figure 28), but the relationship between users and groups is created automatically. In the worst-case scenario, if a user is created but the corresponding group is not, the system creates the group when the user logs into the system. However, the new group has no associated key.

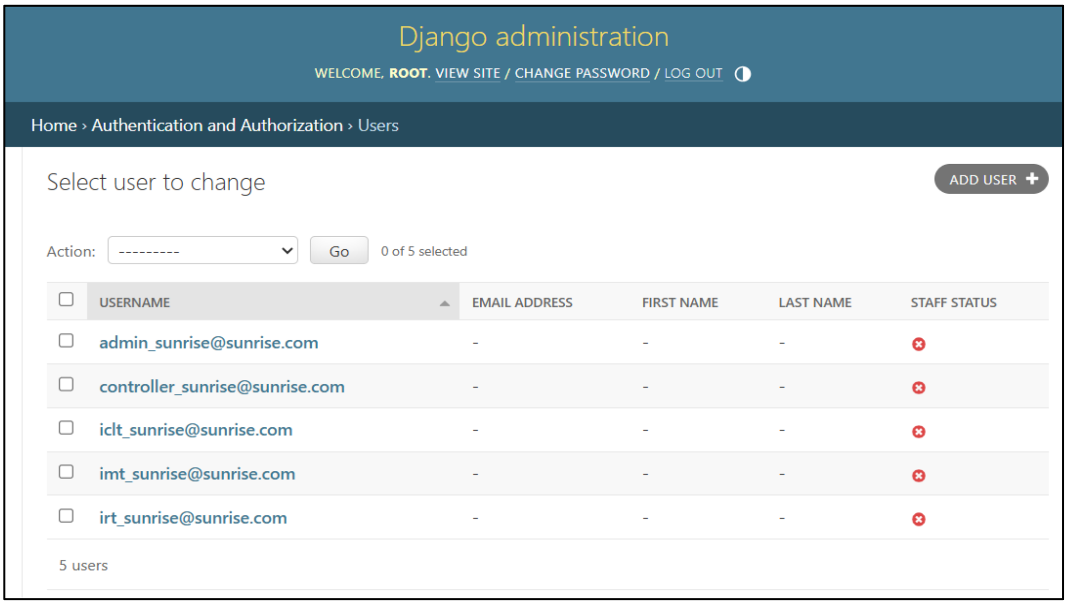

Django administration

WELCOME, ROOT. VIEW SITE / CHANGE PASSWORD / LOG OUT

Home › Authentication and Authorization › Users

Select user to change

ADD USER +

Action: --------- Go 0 of 5 selected

□	USERNAME	EMAIL ADDRESS	FIRST NAME	LAST NAME	STAFF STATUS
□	admin_sunrise@sunrise.com	-	-	-	☒
□	controller_sunrise@sunrise.com	-	-	-	☒
□	iclt_sunrise@sunrise.com	-	-	-	☒
□	imt_sunrise@sunrise.com	-	-	-	☒
□	irt_sunrise@sunrise.com	-	-	-	☒

5 users

Figure 28. CPR Dashboard – Django users administration.

Then, you need to add an API key, obtained from each module, for our organization (see Figure 29).

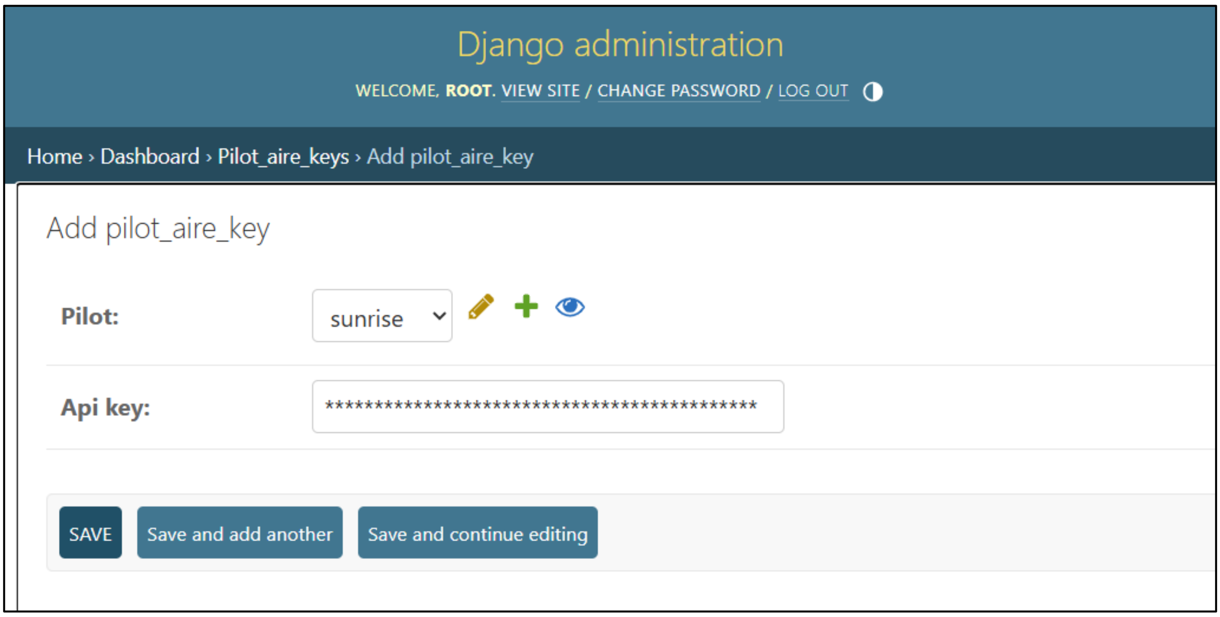

Django administration

WELCOME, ROOT. VIEW SITE / CHANGE PASSWORD / LOG OUT

Home › Dashboard › Pilot_aire_keys › Add pilot_aire_key

Add pilot_aire_key

Pilot: sunrise

Api key: ***

SAVE | Save and add another | Save and continue editing

Figure 29. CPR Dashboard – Api Key Configuration.

2.3.8.3 Advance configuration

The information shown on the control panel can be adjusted to the needs of each organization. To do this, it is only necessary to modify the Model-View-Controller structure of the Django framework[25].

The view (CPR_tool/dashboard/views.py) is what processes the requests from the different pages, and it asks the model layer for the stored information. First, the user and the organization to which he or she belongs are retrieved. Then the organization's API keys are obtained for the different components, with which the different APIs are called.

Part of the model of this application is based on the persistence of the different components. Therefore, the information is obtained by calling the different modules of the system through APIs to retrieve the data. In CPR_tool/{ LOMOS | WAZUH | TINTED | CERCA | AIRE }/api.py you can modify these connections with the different components.

Finally, all the information retrieved is structured in JSON format and passed to the controller, which maps this information to the HTML template using the template language [26]. The template (CPR_tool/dashboard/templates/dashboard/home.html) can be modified to change the main content of the application, such as the labels of the different components or the format of the graphics generated with CanvasJS [27]. While in the base template (CPR_tool/dashboard/templates/base.html), the page header and the links to the different tools can be edited.

3 Cyber-Physical Resilience Tool Tested on Pilots

In the section “Cyber-Physical Resilience Tool Tested in Labs” in D6.3[3], the validation of components, modules and ultimately, the entire CPR tool using the data provided by the CIs and the data from the lab conditions were included. As reported in D6.4[4], the validation during the first piloting phase (M21 and M22 of the project) was done with some demonstrations of the CPR tool. It also used CI-obtained data but has not yet done this in a deployment of the tool in the CI’s operational environment.

The purpose of the first round of pilots’ validation with the end-users was to have their feedback about the usage of the CPR tool, to verify what its different components covered, and for the CI’s end-users to suggest what potential improvements could be implemented to the CPR tool as we reach the last period of the project, when the CPR tool will be deployed in operational environments.

The following subsections provide the enhancements and changes performed in the different components of the CPR tool from the first phase of validation done in the pilots.

We emphasize that an in-depth description of the tools is provided in the previous deliverables, D6.2[2] and D6.3[3].

3.1 Anomaly Detection

The LOMOS (LOg Monitoring System) component is viewed as one of the inputs of the CPR module. This component provides an AI-powered anomaly detection in log data, where we stress that the log data must be humanly readable, raw log data. In-depth description of the requirements for logs are provided in the sections 2.3.3.1 and 0.

The core engine (a BERT-based model) has remained unchanged. During the pilot, we introduced the real-world usability improvements, which were observed and requested by the CIs. These include enhanced utilities, better integration options, and optimizations that simplify deployment in CI environments.

Below we cumulatively describe the relevant enhancements of LOMOS, in order, to provide the best integration into the various existing systems.

- Integration with Wazuh SIEM: LOMOS features the *"lomos-alert"* microservice, which communicates directly with Wazuh. This integration ensures that only high-confidence anomalies are forwarded to Wazuh, reducing unnecessary alerts (reducing alert fatigue) and enhancing incident response efficiency.
- Deployment Enhancements & Installation Re-engineering: New utilities have been added to simplify the setup, configuration, and integration with existing infrastructure. LOMOS is packaged in Docker container. Therefore, Docker Compose can be used in the complex CI environments.
- Scalability & Performance: LOMOS supports multi-instance scaling, allowing it to process high log ingestion rates efficiently. This is particularly relevant for the CI operators with large, distributed infrastructures, such as telecommunications and transport networks.
- Adaptation to GPU Shortages & Alternative Deployment Models: Due to the global GPU shortages, CI operator faced difficulties in acquiring hardware for LOMOS training and inference. LOMOS supports CPU-based inference as a fallback option, ensuring continued operation without requiring a dedicated GPU.
- Pilot demonstrations for Italian and Slovenian CI operators provided valuable feedback. Based on their input, anomaly detection thresholds were fine-tuned.

3.2 Threat Intelligence

During the validation of the CPR Threat Intelligence module for the first piloting phase, most of the main functionalities and enhancements developed in TINTED and already described in D6.3[3] were demonstrated. These included the generation of a threat score for incoming Indicators of Compromise, the generation of a source confidence score for feeds received in the MISP instance and the mapping of the IoCs received with the MITRE ATT&CK framework.

Through discussions with the CI's end-user, and in the context of cyber threat intelligence data sharing and effective collaboration to improve cyber resilience, the following use cases have been identified and categorized according to whether providing anonymization and/or encryption capabilities is required:

Table 2: Threat Intelligence Use Cases

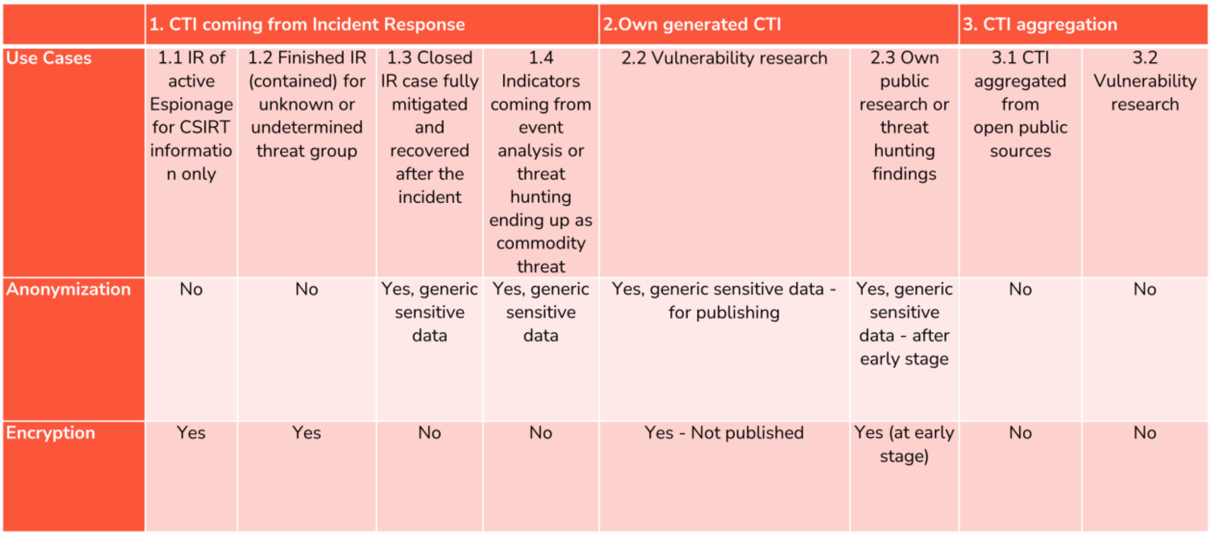

	1. CTI coming from Incident Response				2.Own generated CTI		3. CTI aggregation	
Use Cases	1.1 IR of active Espionage for CSIRT information only	1.2 Finished IR (contained) for unknown or undetermined threat group	1.3 Closed IR case fully mitigated and recovered after the incident	1.4 Indicators coming from event analysis or threat hunting ending up as commodity threat	2.2 Vulnerability research	2.3 Own public research or threat hunting findings	3.1 CTI aggregated from open public sources	3.2 Vulnerability research
Anonymization	No	No	Yes, generic sensitive data	Yes, generic sensitive data	Yes, generic sensitive data - for publishing	Yes, generic sensitive data - after early stage	No	No
Encryption	Yes	Yes	No	No	Yes - Not published	Yes (at early stage)	No	No

Three main groups of use cases have been identified:

(i) cyber threat intelligence (CTI) data that has been generated from an incident response process, where different scenarios can be found based on the status of the incident response (if there is ongoing investigation, it makes sense to use encryption to share information, whereas if the incident is closed, some anonymization is better for sharing sensitive data),

(ii) CTI generated during internal research conducted by the CI about some vulnerability identified, public research or threat hunting finding (encryption could be useful when results have not yet been published and anonymization could be applied to sensitive data for publication), and

(iii) data generated by aggregation of other CTI data (in this last case, no anonymization or encryption is necessary to apply).

This table was used as a starting point during the second piloting phase for the CI end-users' identification of use cases where the establishment of circles of trust with different data sharing policies to share cyber threat intelligence data makes sense and these CPR tool data sharing capabilities can be demonstrated. Currently, two main scenarios are being considered. One is related to sharing information about ongoing phishing campaign attacks (that would be included in the first group in the table) and the other is related to vulnerability management (which is related to the second group in the table).

Regarding the generation of a source confidence score functionality, one improvement resulting from the analysis of the MISP events received by the pilots is related to the processing of events received from CIRCL (Computer Incident Response Centre of Luxembourg) OSINT feed[28]. CIRCL sends threat intelligence data in a slightly different format than most of the default MISP feeds considered initially during the implementation of TINTED (a list of MISP default feeds can be found at [16]). The difference is the CIRCL OSINT feed does not provide a list of indicators of compromise (e.g. IP addresses) that includes a unique MISP event with the name of the feed. It produces a list of MISP events (each of them as a JSON file when the feed is accessed directly through its URL), each of them with different threat intelligence information about the same Indicator of Compromise (that can even be duplicated as provided by different sources). The creator organization can be "CIRCL" but others such as "CthulhuSPRL.be", "Crimeware", "Synovus Financial", "CERT-RLP", "CERT-FR_1510", "SCTIF", etc, follow the same format. The field "Info" included in these events refers to the summary of the information provided, for example "TR-87 – CrowdStrike Agent causing BSOD loop on Windows – Faulty Update on Falcon Sensor". To consider this new type of feed for the generation of the associated source confidence scoring, it is also necessary to provide the list of creator organizations (in the env variable CIRCL_ORG_FEEDS) to the other standard sources (defined in the env variable INTELLIGENCE_FEEDS_FILE, as described in 2.3.5.2). Some of the main changes done for the processing of these events are:

- These events can have MISP attributes with additional and complementary information related to a threat analysis (in particular, attributes with type "text" or "datetime") that are not considered Indicator of Compromises for IoC scoring calculation and are skipped for processing. The main reason is that for the evaluation of the different scoring heuristics, the IoC (e.g. a specific IP address) is searched in the MISP instance to identify if it was already shared by other feeds. In the case of timestamps and generic descriptions, they can be associated to different IoCs, so it makes no sense to search them.
- The completeness heuristic has been modified to consider the number of IoCs provided by the feed, including those belonging to some of the events whose creator organization is included in the environment variable CIRCL_ORG_FEEDS. The reason is because in a CIRCL feed, IoCs are provided in different events with different event Info and different creator Organizations.
- An extensiveness heuristic has also been adapted, since the CIRCL feed usually provides different information about the same IoC in the same event. Consequently, we cannot know the context of the information provided by this feed in advance (as it was done with other types of feeds and established in the intelligence feeds configuration file). It needs to be calculated in a dynamic way, considering the type of information that is present in the event (except attributes present in the MISP event with type "text" or "datetime" that are skipped for scoring are consequently not considered in the total number of IoCs evaluated).
- For evaluation of the whitelist overlap heuristic, it is necessary to know the relationship between whitelisted IoCs in a source with respect to all the IoCs in that source. In the case of events provided by CIRCL, we can have different events from the same creator organization with IoCs. For this reason, we compute in this heuristic the number of IoCs whitelisted that appear in an event with respect to all the IoCs in any of those events coming from the same creator organization. With this approach, we can have different whitelist overlap scores for each of the organizations that provide IoCs with CIRCL and identify, from all the creator organizations participating in CIRCL OSINT Feed, which are more reliable, and which have less whitelisted IoCs.

Finally, improvements have been made in the contextualization of health trends for threat scoring functionality. On the one hand, it has been adapted to individually search for each term included in the list of search keywords and evaluate the average trend. This is because when it looked for a set of keywords, the result returned the number of searches of a keyword with respect to the other keywords. For example, the number of searches about "covid" can be minimal with respect to the number of searches about "flu". However, if the term "covid" is searched individually, its trend can be determined its independently. On the other hand, the function that evaluates the trend slope has been improved, so it can also be valid when the change in the trend is not constant for every time unit, and the result is built as a MISP event with an object type of "report" that is sent as any other IoC shared by TINTED with the CPR risk assessment module through Kafka. Currently, this feature is based on the results of monitoring trends in search queries through the open intelligence source Google Trends as a proof-of-concept for its integration with the CPR risk assessment. However, following the same approach, it could be easily extended in the future to other sources, such as the health prediction generated by the SUNRISE WP5 Demand Prediction tool or some tool that processes Twitter data or asks an AI engine (such as ChatGPT) about the health trend. It is just necessary to have a set of prediction values during a period of time (currently, we are considering 30 days) to calculate the trend or directly send a MISP event to the risk assessment following the same report format with a JSON string in the summary with the trend score (see example in Figure 30).

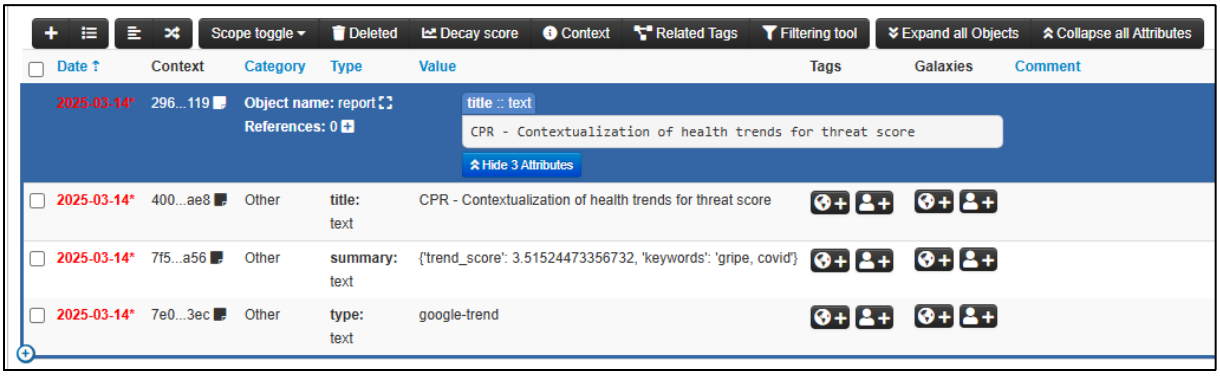

	Date	Context	Category	Type	Value	Tags	Galaxies	Comment
	2025-03-14*	296...119	Object name: report References: 0		title :: text CPR - Contextualization of health trends for threat score			
□	2025-03-14*	400...ae8	Other	title: text	CPR - Contextualization of health trends for threat score			
□	2025-03-14*	7f5...a56	Other	summary: text	{'trend_score': 3.51524473356732, 'keywords': 'gripe, covid'}			
□	2025-03-14*	7e0...3ec	Other	type: text	google-trend			

Figure 30: Contextualization of health trends

3.3 Risk Assessment

Improvements were already reported in D6.3[3], aligned with NIS2, temporary model and threat intelligence contextualisation. They can be summarized in the following bullets:

1. Threat intelligence contextualization: As well as the events reported in D6.3 [3], TINTED now sends information with the contextualization of health trends. This input is utilized as an input to the temporary model as a temporary condition. The value of the trend slope is utilized as a weighted input for the temporary condition modulator.
2. The NIS2 context, mitigations and countermeasures feedback from CSIRT are received by AIRE. The mitigations can be tested in the mitigation simulation module at CERCA, with two approaches:
    a. A mitigation laboratory, where the operator can experiment and simulate different mitigations strategies based on their needs or security policy, in a reactive approach, with various possibilities:
        i. Risk-reduction driven: Reduce the risk to the lowest level.
        ii. Capped cost driven: Reduce the risk within a specific budget to an acceptable level.
        iii. A combination of both, based on the severity of the vulnerabilities and temporary conditions, etc.
    b. An application of mitigations to the risk report and recalculating the risk with the new set of indicators, with some of them returned to a baseline status due to the application of a mitigation or countermeasure.
3. Improved risk reporting: Additional information has been included in risk report, in particular alarms, events and temporary conditions that impacted the risk.
4. Event and alarm persistence: CERCA is now able to store the relevant alarms and events that lead to a risk increase. CERCA also includes the corresponding identifiers to ensure accurate classification and traceability of each alarm. Additionally, the alarms are classified based on severity and this information is now included in the risk report sent to AIRE and is available via GUI and API.
5. Mitigation simulation: CERCA enables the simulation of the effect of applying one or multiple mitigations in a proactive approach, and include practical benefits:
    a. Blue team response capabilities are enhanced, as the mitigation simulation allows for the exploration of defensive strategies based on cost-reduction, on a cost-budget cap limitation or risk reduction.
    b. Ethical red team: Based on the simulated strategies and the selected actor (cost, budget, risk), the ethical red team can focus on assessing the risks related to a particular indicator that needs to be reduced or mitigated.
    c. NIS2: In mitigations sent by CSIRTs in the context of an open incident, operators can see the effect of these mitigations in terms of risk.

d. An improved risk report is sent to AIRE, with the history of what events, alarms and temporary conditions affected the risk, including additional information in the reports.

3.4 Incident Reporting

Improvements already reported in D6.3[3] to support compliance with the NIS2 Directive (communication of supply chain risk, reception of response to an incident notification, attachment of Indicator of Compromises to incident reports and support for vulnerability disclosure) have been refined from the validation done during the first piloting phase and integrated in the default incident reporting workflow enforced by the CPR Incident Reporting module (AIRE). The updated BPMN (Business Process Model and Notation) file with the updated workflow extended in SUNRISE is shown in Figure 31.

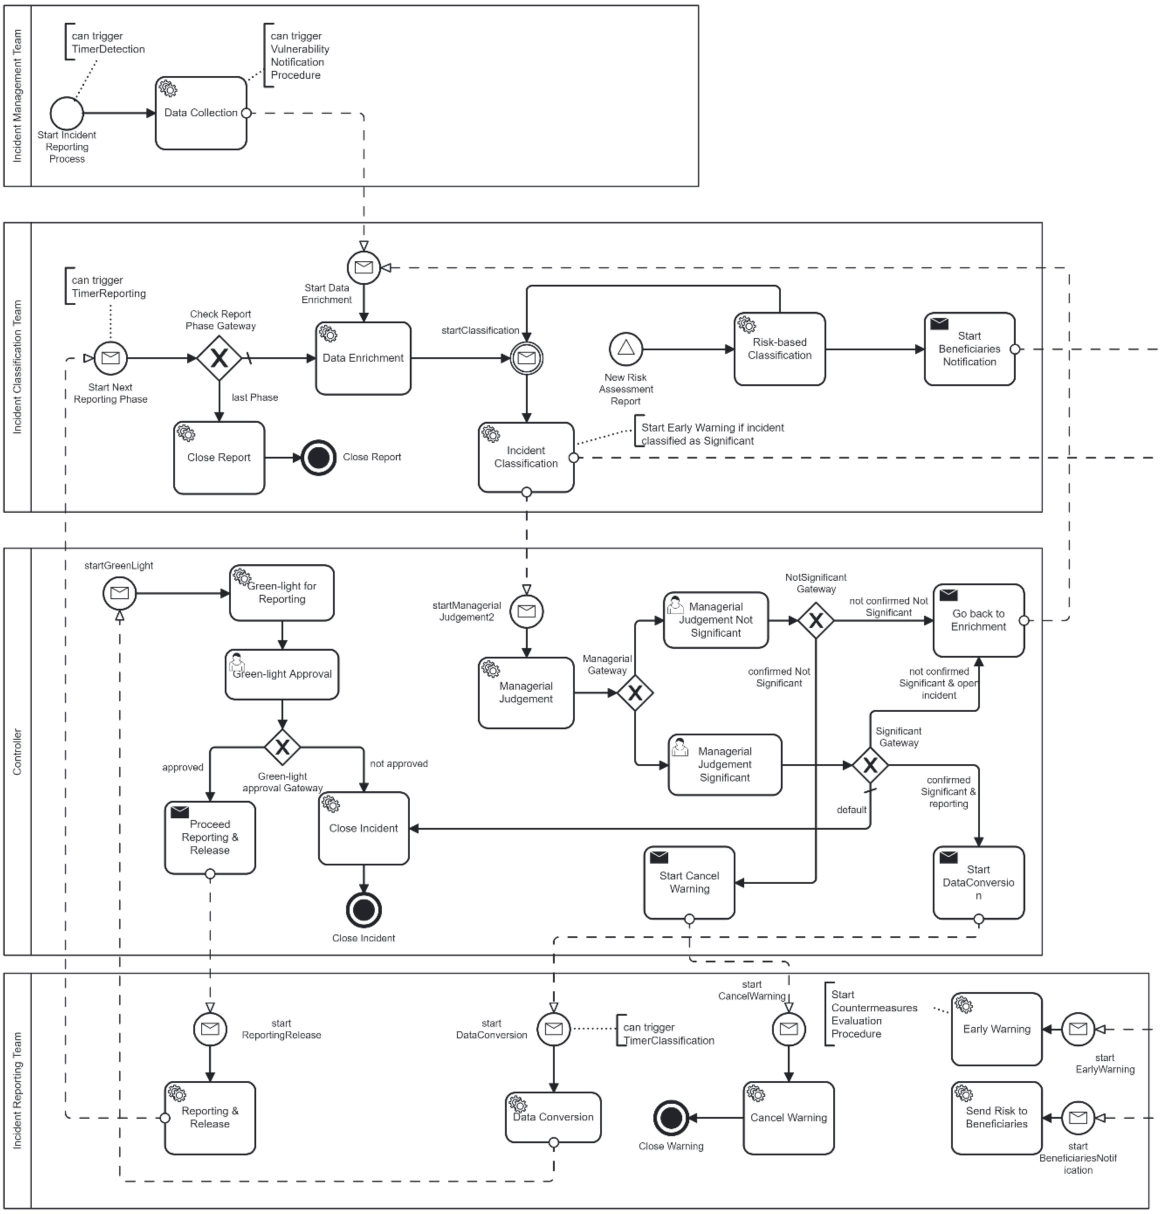
Incident Management Team
can trigger TimerDetection
can trigger Vulnerability Notification Procedure
Data Collection
Start Incident Reporting Process
Incident Classification Team
can trigger TimerReporting
Start Data Enrichment
Check Report Phase Gateway
Data Enrichment
startClassification
Risk-based Classification
Start Beneficiaries Notification
New Risk Assessment Report
Start Next Reporting Phase
last Phase
Close Report
Close Report
Incident Classification
Start Early Warning if incident classified as Significant
Controller
startGreenLight
Green-light for Reporting
Green-light Approval
approved
Green-light approval Gateway
not approved
Proceed Reporting & Release
Close Incident
Close Incident
startManagerialJudgement2
Managerial Judgement
Managerial Gateway
Managerial Judgement Not Significant
NotSignificant Gateway
not confirmed Not Significant
Go back to Enrichment
confirmed Not Significant
not confirmed Significant & oper incident
Managerial Judgement Significant
Significant Gateway
default
confirmed Significant & reporting
Start Cancel Warning
Start DataConversion
Incident Reporting Team
start ReportingRelease
Reporting & Release
start DataConversion
can trigger TimerClassification
Data Conversion
start CancelWarning
Close Warning
Cancel Warning
Start Countermeasures Evaluation Procedure
Early Warning
start EarlyWarning
Send Risk to Beneficiaries
start BeneficiariesNotification

Figure 31: CPR incident reporting workflow supporting NIS2.

Connection with the CPR risk assessment is done through a new service task called “Risk Evaluation Collection” that will gather the risk report generated in order to retrieve the result that will be shown to the end-user. Risk-based classification of a security incident as significant or not is done through the user task “Risk-based Classification” and it automatically starts the process of sending a notification to service beneficiaries through the service task “Send Risk to Beneficiaries”. This initial risk-based classification is integrated with the existing incident classification process. Once this classification is done, it triggers the service task that performs the early warning ("Start Early Warning") and automatically sends a notification email to the Supervisory Authorities, which is configured with this early report of the incident (service task "Early Warning"). This notification can be cancelled if it is determined during the managerial judgement step that the incident is ultimately classified as not significant (service task "Start Cancel Warning" in the Controller pool and service task "Cancel Warning" in the Incident Reporting pool).

After the early warning, and according to NIS2 Directive, CSIRTs must respond within 24 hours with some initial countermeasures to be applied to mitigate the incident or prevent a potential cyber attack (e.g. the application of patches for identified vulnerabilities or the suggestion to close some port in the firewall). This is modelled in the CPR tool through the triggering of the countemeasures monitoring the process shown in Figure 32 after an early warning has started. Basically, in these additional steps of the incident reporting workflow, the information received about countermeasures is registered, it is evaluated for application by the end-user, and finally, the risk assessment is reevaluated.

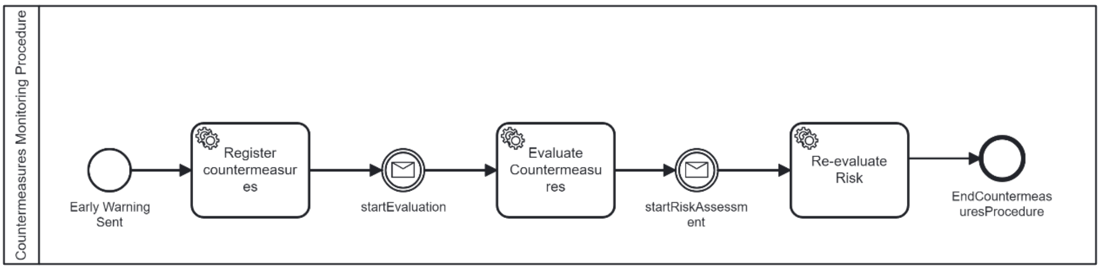

Figure 32: CPR incident reporting – countermeasures monitoring procedure.

Furthermore, since a potential unknown vulnerability can be identified at any time in the data collection task (as reported in D6.3[3]), it is possible to trigger an optional vulnerability notification procedure, as shown in Figure 33. As such, this has been added to the CPR incident reporting workflow.

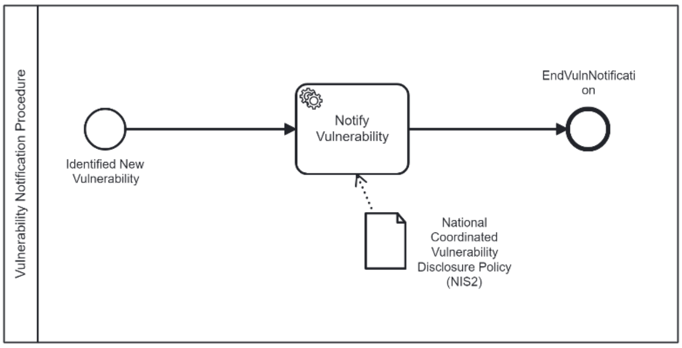
Vulnerability Notification Procedure
Identified New Vulnerability
Notify Vulnerability
EndVulnNotification
National Coordinated Vulnerability Disclosure Policy (NIS2)

Figure 33: CPR incident reporting – vulnerability notification procedure.

Transposition of the NIS2 Directive [7] to national legislation had the deadline 17th October 2024. However, after consultation with the different Critical Infrastructures involved in the project pilots, they indicated that no specific template was defined for reporting and consequently, no changes have been made to the existing templates used during the first piloting phase concerning GDPR and NIS2 regulation frameworks.

3.5 Dashboard

The CPR Dashboard was not yet available during the validation done in the first piloting phase. It has been designed as a unified and integrated dashboard for the CI's end-users that, through a single-sign-on, cannot only provide them with access to the different functionalities of the CPR tool, but also enable them to quickly see what is happening at different levels, including detection inside the infrastructure, risk assessment, external threat context and incident reporting, as is show in Figure 34.

The following is a brief description of the tool: the header of the dashboard has direct access to the different modules, where there are the advanced features and a logout button. Then, there is the graphic that represents the anomalies of the last period, which is generated with the data of the Log Anomaly Detection module. The second row contains the WAZUH alarm graphic and the latest reports from the Risk Assessment module, which displays the list of the risk models selected for the report and the qualitative and quantitative risk of each model. The last row contains the Threat Intelligence Monitoring module, which lists the latest and most relevant IoCs and their associated techniques, and the Incident Reporting module, which lists the open incidents and the task that are currently in progress. All five components can be collapsed, allowing the users to focus on the components most relevant to their roles.

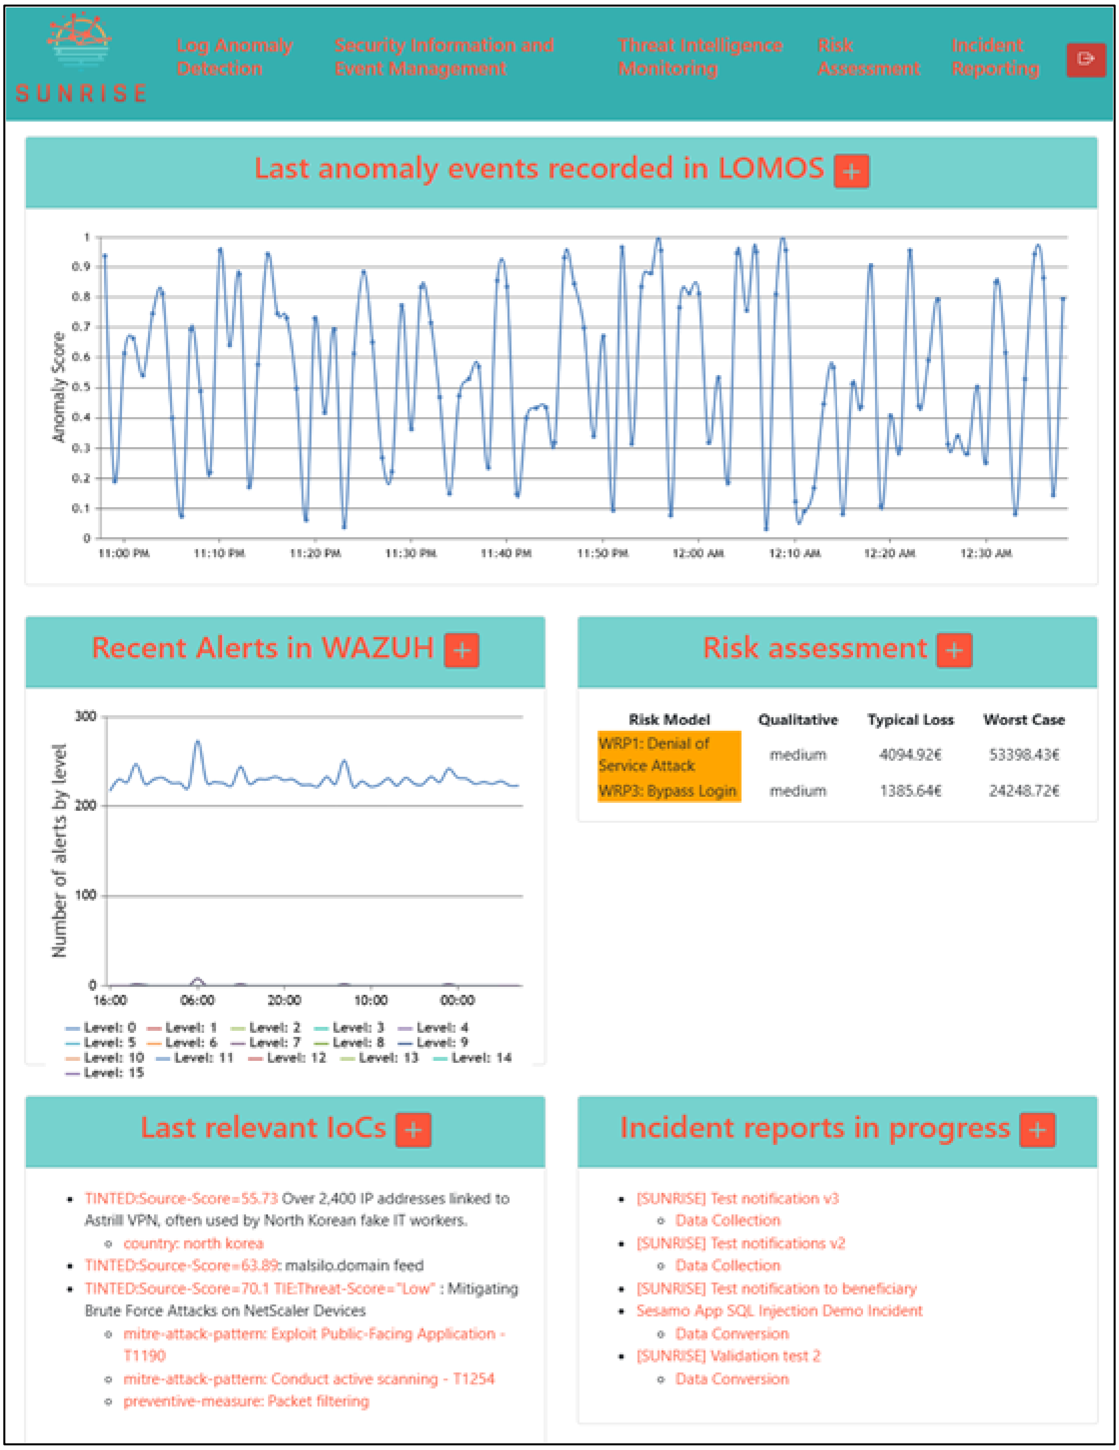

Risk Model	Qualitative	Typical Loss	Worst Case
WRP1: Denial of Service Attack	medium	4094.92€	53398.43€
WRP3: Bypass Login	medium	1385.64€	24248.72€

- TINTED:Source-Score=55.73 Over 2,400 IP addresses linked to Astrill VPN, often used by North Korean fake IT workers.
    - country: north korea
- TINTED:Source-Score=63.89: malsilo.domain feed
- TINTED:Source-Score=70.1 TIE:Threat-Score="Low" : Mitigating Brute Force Attacks on NetScaler Devices
    - mitre-attack-pattern: Exploit Public-Facing Application - T1190
    - mitre-attack-pattern: Conduct active scanning - T1254
    - preventive-measure: Packet filtering

- [SUNRISE] Test notification v3
    - Data Collection
- [SUNRISE] Test notifications v2
    - Data Collection
- [SUNRISE] Test notification to beneficiary
- Sesamo App SQL Injection Demo Incident
    - Data Conversion
- [SUNRISE] Validation test 2
    - Data Conversion

Figure 34: CPR Dashboard.

4 Legal Compliance and Testing Methodology

Section 4 outlines potential legal concerns, such as data sharing restrictions for CI due to legal regulations or limitations on permissible testing activities. In addition, the piloting activities are outlined.

At this stage, no problems have arisen, although they may appear during the second piloting phase. If any concern arises, it will be reported in the Cyber-physical resilience pilot report V2 deliverable (D6.6) and closely monitored in WP9 – Management: Project Coordination and Technical Coordination (WP3 in general, specifically T3.4).

However, considering the WP6 time plan, all CI operators have been notified about the schedule and possible issues, so that the piloting activities are not delayed. The pilot schedule is the following:

- Deployment will be done in the last week of May.
- Testing will be done in the first week of June.
- The collection of feedback is due from the second and to the third week of June.

The testing process will unfold in several meticulous phases. The first two activities will rely on the preparation of a proper environment in terms of hardware and software requirements, which have been detailed in Section 2.3. Initially, deployment will occur within the CI environment, where hardware or virtual machines are prepared, granting access to XLB and ATS to install the CPR tool via VPN. Following installation, XLB and ATS will diligently test the CPR tool with mock data to ensure functionality. Subsequently, the CI operators will connect the log source, initiating the processing phase where the CI operators, XLB and ATS will collaborate to test the functionality. A crucial aspect involves accumulating a sufficient volume of logs, spanning from one hour to a week, to rigorously assess the CPR's efficacy with CI operator data. During this period, both CI operators and XLB will undertake the critical tasks of model training and inference. Once an adequate dataset is amassed for real data analysis, typically within a week, CI operators and XLB will resume model training and inference to ensure accuracy, together with ATS, as the processing of alarms coming from the Anomaly Detection Module is done in CERCA and AIRE. Finally, the CI operators will conduct periodic testing to uphold a continuous performance evaluation and refinement. Each phase reflects a meticulous approach aimed at guaranteeing the reliability and efficacy of the CPR tool within the CI ecosystem. In those CI environments where, by its own nature, it is not possible to deploy the CPR tool, access to deployment in the ATS premises will be provided to the CI’s end-users so they can experiment and test the tool by themselves. An example of such is the Slovenian Railways infrastructure.

5 Pilot trials execution

To be successful, the developed CPR tool needs to provide useful functionality to the CI operators and at the same time reduce the CI operator workload. The pilot trials are planned to show the tool to the CI partners, test the deployment in an operational environment, test the operational procedures, and give the CI operators an opportunity for early, hands-on evaluation.

The main goal of the preliminary test on real data is to see if the CPR components are functional and if there are any problems that could prevent or hinder the installation and subsequent use of the CPR tool in an operational environment. These problems need to be discovered and properly addressed.

The description of the users, their roles within critical infrastructure operators, and their required expertise is provided in Section 5.2.

5.1 Pilot Execution Plans

D6.3[3] presented an initial plan for pilot execution. Those plans are largely unmodified and are included in this document. However, in this deliverable, they have been updated to reflect the pilots where it is not possible to deploy the tool in an operational infrastructure environment.

5.1.1 Italy: Public Administration

INS plans to test the tools developed in WP6 on a specific part of its regional infrastructure. This subset includes the VPN Servers used by INS employees and the SESAMO application, which collects health-related data of citizens in the Region Friuli Venezia Giulia. Access to SESAMO is controlled through federated systems like national SPID and Electronic ID Cards. The testing process will happen in three phases:

Phase 0 (M12): A Proof of Concept will simulate incoming logs from selected applications to identify potential threats to the systems under analysis.

Phase 1 (M23): The tools will be deployed within INS to integrate and aggregate logs with the existing SIEM (Security Information and Event Management) system, which is the Community Edition in the testing environment. Additionally, integrations with the current MISP appliance in the INS infrastructure may be explored to enhance Cyber Threat Intelligence sharing between the Firewall and SIEM. The AIRE functionalities will be integrated to generate incident reports for managing the Incident Response process.

Update on the initial plan for Phase 1: Although initially it was planned to deploy the CPR tool within the INS infrastructure in this phase, the final deployment of the CPR components was done on the ATS/XLB premises and the real logs generated by the INS infrastructure were integrated and aggregated there. The result of these tests is reported in D6.4 [4].

Phase 2 (M34): The integrated CPR tool will be piloted in the operational environment of INS. The tool output will be monitored using real data from the applications under analysis, as well as simulated data for vulnerability tests. The INS Blue Team might oversee these operations.

In summary, INS aims to test the WP6 tool on a specific part of its regional infrastructure, including VPN Servers and the SESAMO application. The testing will progress through phases, beginning with a Proof of Concept, followed by integration with SIEM and MISP, and concluding with operational environment monitoring.

5.1.2 Italy: Water

CAFC intends to test the tools developed by WP6 on its VPN infrastructure to detect and halt suspicious network activities. The VPN system, which has become very important due to the COVID-19 pandemic, was previously used by a limited group of technical users for a decade. However, it has now become the standard means for employees to remotely access various systems.

For training and testing, CAFC will provide logs including network traffic, antispam filter records, antivirus data, and results from ongoing penetration tests.

The testing process will involve the following key steps:

Phase 0 (M12): A Proof of Concept will showcase the tools' ability to identify potential threats to the systems being examined through simulated incoming log data.

Phase 1 (M23): The tools will be implemented within CAFC to demonstrate how logs are acquired and analysed.

Update on the initial plan for Phase 1: Although initially it was planned to deploy the CPR tool within CAFC infrastructure in this phase, the final deployment of the CPR components was done on the ATS/XLB premises and the real logs acquired from the infrastructure were integrated and aggregated there. The result of these tests is reported in D6.4 [4].

Phase 2 (M34): CAFC will carry out a trial of the tools in their actual operational environment. This will involve monitoring the CPR tool output using real data and potentially simulating threats to the VPN infrastructure. The goal is to assess the tools' effectiveness.

To sum up, CAFC plans to test the WP6 tools on its VPN system for detecting suspicious network behaviour. The VPN's significance has increased due to the pandemic, and the testing will proceed through phases involving Proof of Concept, deployment, and operational evaluation.

5.1.3 Slovenia: Telecommunication

TS plans to evaluate the WP6-developed tools on a specific segment of its infrastructure associated with the VALU mobile application's accesses and activities. The testing process will follow an iterative approach, involving the following phases:

Phase 0 (M12): The initial Proof of Concept will showcase the tools' ability to detect potential threats within the analysed systems. This will be done by simulating incoming log data from the selected application.

Phase 1 (M23): The tools will be implemented within TS's environment to demonstrate how logs can be integrated and aggregated between the new tools and the existing SIEM monitoring system (Community Edition) that is already in operation within TS.

Update on the initial plan for Phase 1: Although initially it was planned to deploy the CPR tool within TS’s infrastructure in this phase, final deployment of the CPR components was done on the ATS/XLB premises and the logs acquired from the infrastructure were analyzed. However, since they are not suitable for anomaly detection, they could not be integrated and aggregated. The result of these tests is reported in D6.4 [4].

Phase 2 (M34): In this phase, TS will use the integrated CPR tool to the test environment within its operational setup. The tool' performance will be monitored using both real data from the applications under analysis and simulated data. For instance, simulated vulnerability tests using known patterns might be conducted. These activities could be supervised by TS's advanced payment and cybersecurity teams.

In summary, TS will assess the WP6 tools on a specific infrastructure section linked to the VALU mobile app. The evaluation will occur in iterative stages, comprising Proof of Concept, deployment with log integration, and operational assessment with real and simulated data under the watch of specialized teams.

5.1.4 Slovenia: Transport

As Slovenian transport operators (SZ-SZI), the organization plans to evaluate the WP6-developed tools on a specific component of the railway infrastructure referred to as PRI. This component is the railway traffic monitoring system, known as ISSŽP. The system encompasses numerous services and applications designed to monitor real-time events occurring along the railway infrastructure. These events pertain to both freight and passenger transportation. Certain applications facilitate bidirectional data exchange and communication with external carriers, enabling comprehensive real-time management and monitoring of the railway infrastructure.

The testing process will adhere to an iterative approach via the following stages:

Phase 0 (M12): The initial Proof of Concept will illustrate the tools' capability to identify potential threats within the analysed system. This will involve simulating incoming log data from the chosen system.

Phase 1 (M23): The tools will be implemented within SZ-SZI's environment to showcase the integration and aggregation of logs between the proposed tools and the existing railway monitoring system (ISSŽP).

Update on the initial plan for Phase 1: It was initially planned that tools would be implemented within SZ-SZI's environment but due to SZ-IT security policy, it is not possible to integrate and test the systems directly in the SZ information environment. Consequently, SZ-SZI provided the necessary information regarding SZ's IT infrastructure environment and the different components of the CPR tool were evaluated based on surveys, providing further feedback regarding the tool’s improvement. The result of these tests is reported in D6.4 [4].

Phase 2 (M34): SZ-SZI will undertake a pilot of the tools within its operational setup. The tools' performance will be observed using both actual data sourced from the applications under examination and simulated data. This could encompass attempts to conduct vulnerability tests on the applications using recognized patterns. These activities may be overseen by SZ-SZI's technical department.

Update on the initial plan for Phase 2: SZ-SZI will evaluate the final version of the integrated CPR tool based on the demonstrations done in the training material and testing of the tool as in Phase1, providing the required feedback through surveys. These activities may be overseen by SZ-SZI's technical department to check the suitability of CPR tool for their needs and the application of the tool to SZ-SZI's real environment.

In summary, SZ-SZI, acting as Slovenian transport operators, intends to assess the WP6 tools on a specific railway infrastructure asset known as PRI. This asset encompasses the ISSŽP railway traffic monitoring system with various applications. The assessment process will involve iterative stages of Proof of Concept, deployment with log integration, and operational evaluation under the potential supervision of SZ-SZI's technical department.

5.2 Description of End-Users' Roles

User Organization / ROLE	Profile	Skills	CPR tools interaction Description
Authorities			
Government - institutional level decision maker	Strategic	knowledge of policies, competences	May use reports from CPR dashboard to promote policies or strategies that give authority to CI operators or other institutions to define temporary measures (e.g. increase level of access control, introduce remote working)
CI high level decision maker / operator	Strategic	Knowledge of CI and sector business, knowledge of IT infrastructure and services that support CI	May use reports from CPR dashboard to analyse cyber-physical risks related to regions, sectors, and specific CI. They might also decide on how to generate and share risk indicators, and information about threats, techniques, tactics, and procedures, as well as incident reporting protocols.
CI operator			
CI legal and compliance team			
Legal Manager	Tactical	knowledge of policies, knowledge of laws, strategies	Creates reporting directives and procedures based on the strategies and compliance needs of the CI operators. Adapts reporting templates to the case of the CI, deciding which fields are necessary and relevant or which information should be anonymized before sharing.
Legal User	Operational	knowledge of GDPR/NIS 2 reporting protocols	Analyses whether the incident meets the criteria to be reported. Anonymizes the information shared. Ensures that all necessary information is available.

User Organization / ROLE	Profile	Skills	CPR tools interaction Description
CI risk management team			
Business Risk Manager	Tactical	general risk management skills	Creates cyber risk management procedures based on strategic directives. Defines the cyber infrastructure risk model to assess the effects of a cyber-physical incident and to be able to evaluate the risk/impact. If business or strategic priority is changed during pandemics, adapts data on estimated impact. Provides feedback about tactical level threat information (e.g. observed tactics and techniques of an attacker)
Risk Operator	Operational	Cyber risk analysis	Operates CyberRisk assessment Calculator (CERCA) tool e.g. configuration, manual data inputs where needed, etc. Provides feedback about operational level indicators (e.g. completeness, timeliness etc)
CI Incident Response and security operation team			
IRM/SOC/CERT Manager	Tactical	knowledge of cyber-attacks, knowledge of internal architecture	Defines rules to monitor and detect different threats and decides about response actions, also considering the inputs and priorities received from the Business Risk Manager. It also decides on cyberthreat intelligence (CTI) sharing policies and makes analysis (in collaboration with the other CIs) about adversary tactics, techniques, and procedures (TTP)
IRM/SOC/CERT operator	Operational	knowledge of infrastructure and normal behaviour patterns, knowledge of security countermeasures	Monitors CI cyber assets/services, detecting unusual behaviour. Applies some countermeasures (e.g. fast response such as blocking access) against threats and obtains evidence of incidents for reporting. Shares operational level CTI (indicators of compromise)
IT department			
Administrator	Operational	configuration of CPR tools, access control	Administers CPR tools (LOMOS, CERCA, AIRE…), creating and administrating users

5.3 What are the benefits of the Cyber-Physical Resilience tool?

The report detailing the concrete benefits for end-users of the Critical Infrastructures involved in the project will be done in the next deliverable D6.6 – Cyber-physical resilience pilot report V2, once the CPR tool is deployed within the respective CI's infrastructures and after the execution and validation during the second piloting phase (M34). However, based on the feedback collected from the end-users after the first piloting phase through surveys, we can draw preliminary conclusions indicating a positive general perception from all the participants about the enhanced cybersecurity situational awareness provided by the CPR tool.

When asked about the most relevant situations or challenges they faced during pandemia that directly impacted on cybersecurity operations and incident management where currently they think that the CPR tool could be beneficial for them, most of them highlighted the support to dynamic adaptability provided by the anomaly detection and cyber-physical risk assessment modules integrated in the SUNRISE CPR tool. Specifically, the ability to support potential changes in connectivity patterns, situations of network overload and congestion or the adaptation to changes in priorities and values of digital assets. Additionally, it was also noted as positive the tool's consideration for changes in people behaviour during pandemics and availability of workforce, through the establishment of temporary conditions to include prediction of absenteeism and risk models for phishing. Another significant advantage of the approach followed in the CPR tool seen as positive by end-users is the reduction in the time to fine-tune or make changes in mitigation actions, providing a better reaction time to incidents. Availability of timely and trustworthy cyber threat intelligence data sharing is also other of the top features identified by many end-users. Features included in SUNRISE such as threat intelligence scoring and mapping of Indicators of Compromise with MITRE ATT&CK techniques benefit not only to CI incident response and security operation teams but also to high-level decision-makers. The CPR tool enables them to visualize in a same dashboard the current situation from multiple perspectives, including anomalies identified inside the monitoring infrastructure, the overall risk assessment for the organization, the ongoing incident reporting processes and relevant external threat intelligence data shared e.g. from public feeds, allowing to identify if it comes from trusted sources and if it can be relevant for their organization. Finally, end-users have also considered that the CPR tool can help to solve collaboration issues, e.g. among local and regional authorities, and improve communication with CSIRTs in compliance with current regulatory frameworks, such as the NIS2 directive.

6 Conclusions

The Cyber-physical resilience tool’s design and purpose responds to how critical infrastructures have evolved during recent years, experiencing a wider exposure to attacks, while the pandemic experience revealed new cybersecurity operational challenges and the need for a more dynamic risk assessment.

To address the existing and new challenges and gaps related to the operation of CIs and cyber-physical security in temporary conditions such as pandemics, we have combined several modules in the CPR tool, including anomaly detection, incident reporting, tailored threat intelligence management and semi-automated risk assessment, enhanced by the so-called OODA loop (Observe, Orient, Decide, Act) approach and access to all these functionalities from a unified interface.

This deliverable provides a thorough summary of the SUNRISE Cyber-Physical Resilience Tool (CPR Tool) together with the latest version of its installation and user manual. It extends and updates information given in the previous deliverables: D6.1[1], D6.2[2] and D6.3[3]. While the architecture of the CPR tool with its main components, the piloting plans and the initial validation on real data provided by the end users were already included in former deliverables, this deliverable aims to further explain how the different modules can be integrated into existing operations, highlight the specific benefits for the end-users and outline the latest enhancements based on the feedback received during the initial piloting phase (as reported in detail in D6.4[4]).

The latest developments in the tool have been mainly focused on including usability improvements, better integration among the different modules in the tool and with existing legacy systems (such as SIEM), and adaptation to constraints in the hardware available in the CIs (in particular, in the GPUs). There has also been a big effort to simplify the deployment and configuration of the different modules included in the CPR tool and improve the manuals so the CPR tool can be easily deployed and integrated with the existing infrastructure.

In summary, the Cyber-Physical Resilience tool developed in SUNRISE and presented in this deliverable offers a vital resource to the Critical Infrastructure end-users to enhance their resilience and better prepare them to adapt their cyber security management processes to unforeseen and emergency circumstances. This holistic approach incorporates the initial phase of anomaly and incident detection to the final phase of reporting to the Supervisory Authorities in compliance with regulatory frameworks, with focus on the new NIS2 Directive and including dynamic risk assessment that considers the presence of potential temporary conditions and context information provided from threat intelligence sources.

7 References

[1] **SUNRISE. D6.1** - Cyber-physical resilience conceptualization. Pablo de Juan. 2023.

[2] **SUNRISE. D6.2** - Cyber-physical resilience tool and training guide V1. Tomaž Martinčič. 2023.

[3] **SUNRISE. D6.3** - Cyber-physical resilience tool and training guide V2. Justin Činkelj. 2024.

[4] **SUNRISE D6.4** – Cyber-physical resilience pilot report V1. Susana González Zarzosa. 2024

[5] **SUNRISE. D3.2** - Requirements and designs, V2. George Tsakirakis. 2023.

[6] **P. He, J. Zhu, Z. Zheng and M. R. Lyu**, "Drain: An online log parsing approach with fixed depth tree," 2017 IEEE international conference on web services (ICWS), pp. 33-40, 2017.

[7] **H. Guo, S. Yuan and X. Wu**, "LogBERT: Log Anomaly Detection via BERT," 2021 international joint conference on neural networks (IJCNN), pp. 1-8, 2021.

[8] **J. Devlin, M. W. Chang, K. Lee and K. Toutanova**, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," CoRR, vol. abs/1810.04805, 2018.

[9] **The NIS 2 Directive:** http://data.europa.eu/eli/dir/2022/2555/2022-12-27

[10] **ENSURESEC – D4.5**: Human, cyber & physical business components mapping tool. G. Gonzalez-Granadillo, J. Martinez, J. Torner, R. Diaz, A. Alvarez, J.J. De Vicente. 2021.

[11] **A. Oliner and J. Stearley**, "What Supercomputers Say: A Study of Five System Logs," 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07), pp. 575-584, 2007.

[12] **W. Xu, L. Huang, A. Fox, D. Patterson and M. I. Jordan**, "Detecting large-scale system problems by mining console logs," Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles, pp. 117-132, 2009.

[13] **CyberSec4Europe - D3.21** Framework to design and implement adaptive security systems. L. Pasquale and A. Hassan. 2022.

[14] **CORAS** The CORAS Method (sourceforge.net)

[15] **B. Warner, A. Chaffin, B. Clavié, O. Weller, O. Hallström, S. Taghadouini, A. Gallagher, R. Biswas, F. Ladhak, T. Aarsen, N. Cooper, G. Adams, J. Howard and I. Poli,** "Smarter, Better, Faster, Longer: A Modern Bidirectional Encoder for Fast, Memory Efficient, and Long Context Finetuning and Inference"**:** https://arxiv.org/abs/2412.13663},

[16] M. T. Sharing, **MISP Default Feed**, https://www.misp-project.org/feeds/, retrieved 2024-07-19).

[17] https://websites.fraunhofer.de/CIPedia/index.php/Cyber_Resilience

[18] **Sakkas, Georgios**, "CWA 18024:2023 - Emergency management - Incident situational reporting for critical infrastructures", {https://www.cencenelec.eu/media/CEN-CENELEC/CWAs/RI/cwa18024.pdf}, CEN-CENELEC, 2023.

[19] Corte Analyzers Repository. https://github.com/TheHive-Project/Cortex-Analyzers (last access: 14/04/2025)

[20] **Jelle Ermerins, Niek van Noort, Joao de Novais Marques, Leandro Velasco**. “Scoring model for IoCs by combining open intelligence feeds to reduce false positive”. Security and Network Engineering. February 9, 2020. https://rp.os3.nl/2019-2020/p55/report.pdf

[21] Wazuh https://wazuh.com

[22] MISP https://www.misp-project.org

[23] TheHive https://github.com/TheHive-Project/TheHive

[24] https://docs.strangebee.com/thehive/installation/system-requirements/ and https://docs.strangebee.com/cortex/installation-and-configuration/

[25] https://www.djangoproject.com/

[26] https://docs.djangoproject.com/en/5.1/ref/templates/language/

[27] https://canvasjs.com/

[28] https://www.circl.lu/doc/misp/feed-osint/

Annex I Training guide and user manual

This Annex covers the training guide and user manual for all CPR the tool modules.

Table of Content (Annex I)

List of Tables (Annex I)

List of Figures (Annex I)

1 CPR Dashboard

This section explains how the CPR Dashboard can be used to access the different functionalities provided by the CPR tool. The first step is to login to the dashboard with your user and password (Figure 1).

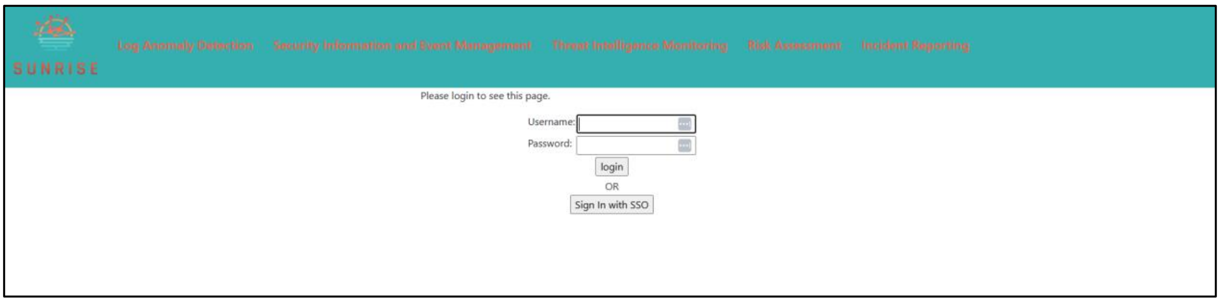

Figure 1. CPR Dashboard Login form.

Once logged in, the user can access to the main dashboard (Figure 2). The graphical interfaces provided by each of the CPR modules can be directly accessed through the menu on the top frame:

- **Log Anomaly Detection**: This gives access to the LOMOS dashboard described in annex section 2.
- **Security Information and Event Management**: This gives access to Wazuh interface.
- **Threat Intelligence Monitoring**: This gives access to the MISP Dashboard that shows all events received. Details about the information provided by TINTED through this dashboard is described in annex section 3.
- **Risk Assessment**: This gives access to the CERCA dashboard described in annex section 4.
- **Incident Reporting**: This gives access to the AIRE dashboard described in annex section 5.

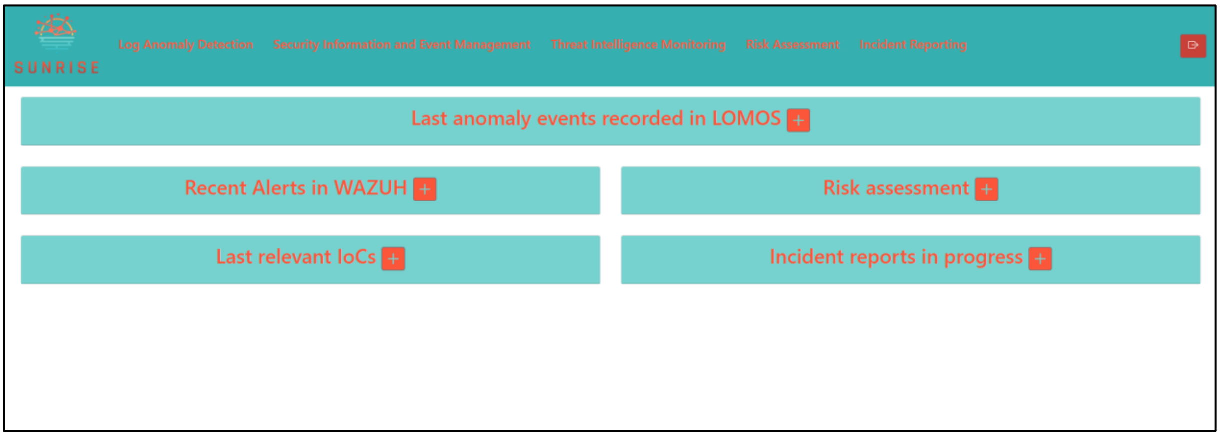
SUNRISE
Log Anomaly Detection
Security Information and Event Management
Threat Intelligence Monitoring
Risk Assessment
Incident Reporting
Last anomaly events recorded in LOMOS +
Recent Alerts in WAZUH +
Risk assessment +
Last relevant IoCs +
Incident reports in progress +

Figure 2. CPR Dashboard Main Panel.

There are also five dropdowns on the main screen with a summary of the information provided by the different CPR components, which provides an overview of the current status of the infrastructure.

- **Last anomaly events recorded in LOMOS** (Figure 3): This chart shows in a timeline, the latest anomaly events detected by LOMOS with its anomaly score (from 0 to 100).

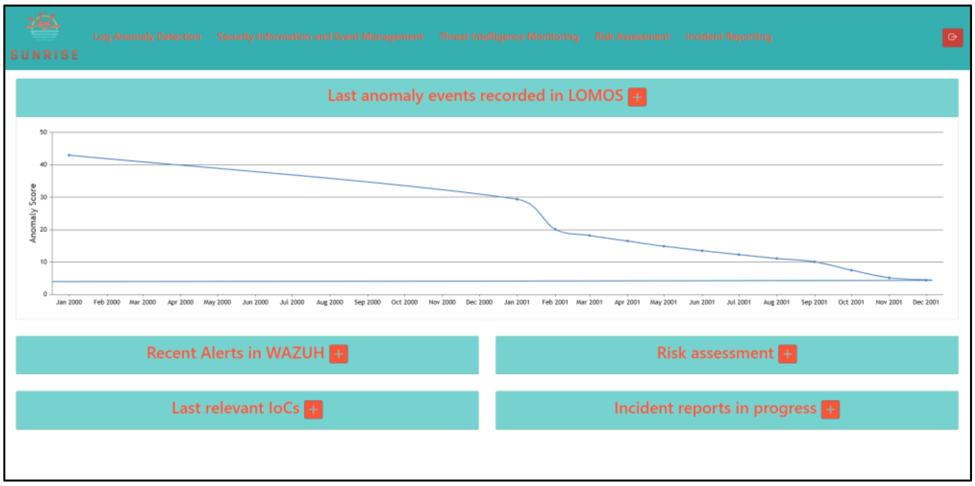

Figure 3. CPR Dashboard – Last anomaly events.

- **Recent alerts in WAZUH** (Figure 4): This chart shows in a timeline the number of the latest alerts detected in Wazuh by level of security.

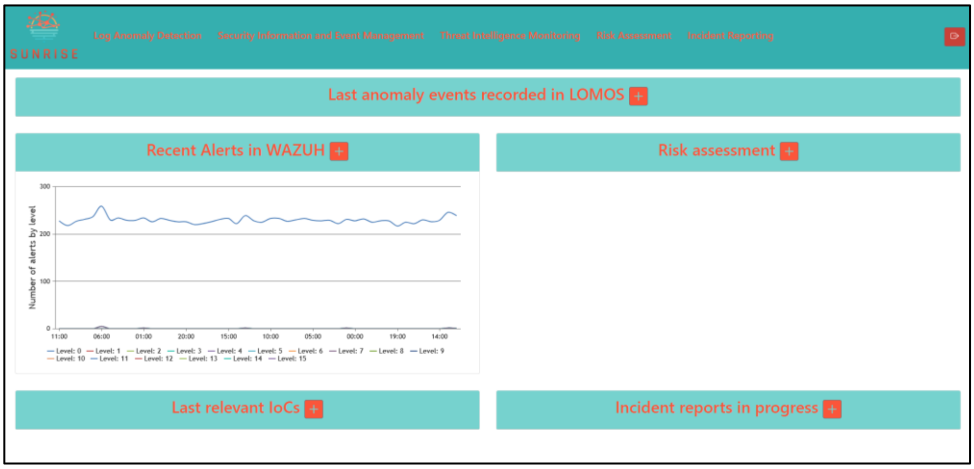

Figure 4. CPR Dashboard – Recent alerts.

- **Risk assessment (**Figure 5**):** This shows, with a colour scheme (green-yellow-red), the list of active risk models enabled in the CPR tool with its qualitative score. It also shows the quantitative risk in economic terms (€) of the typical loss and worst-case scenarios.

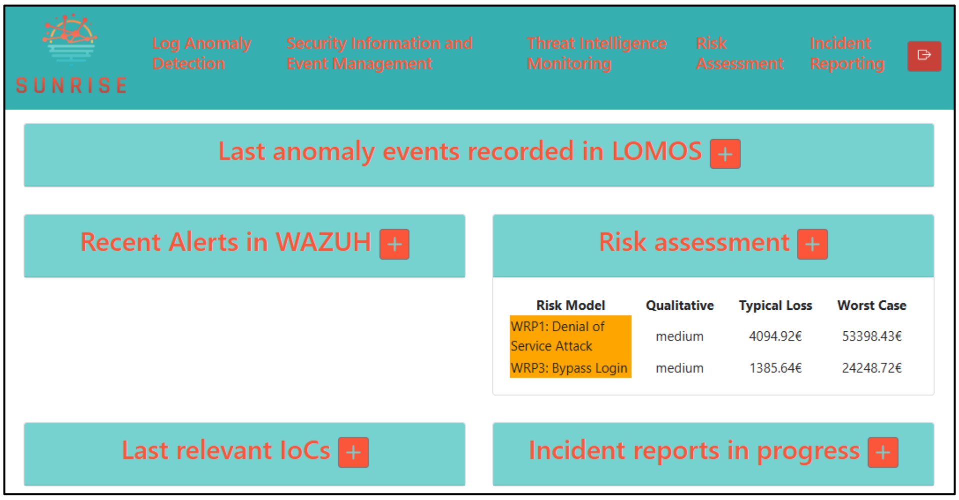
SUNRISE
Log Anomaly Detection
Security Information and Event Management
Threat Intelligence Monitoring
Risk Assessment
Incident Reporting
Last anomaly events recorded in LOMOS +
Recent Alerts in WAZUH +
Risk assessment +
Last relevant IoCs +
Incident reports in progress +

Risk Model	Qualitative	Typical Loss	Worst Case
WRP1: Denial of Service Attack	medium	4094.92€	53398.43€
WRP3: Bypass Login	medium	1385.64€	24248.72€

Figure 5. CPR Dashboard – Risk assessment.

- **Last relevant IoCs** (Figure 6): This shows a list of the latest Indicator of Compromises received in the MISP instance with the threat score and source confidence score calculated by the TINTED component of the CPR tool. If there is some mapping to tactics according to the MITRE ATT&CK framework, it is also shown. Click on the indicators shown for more details to open access to the MISP event in the MISP Dashboard.

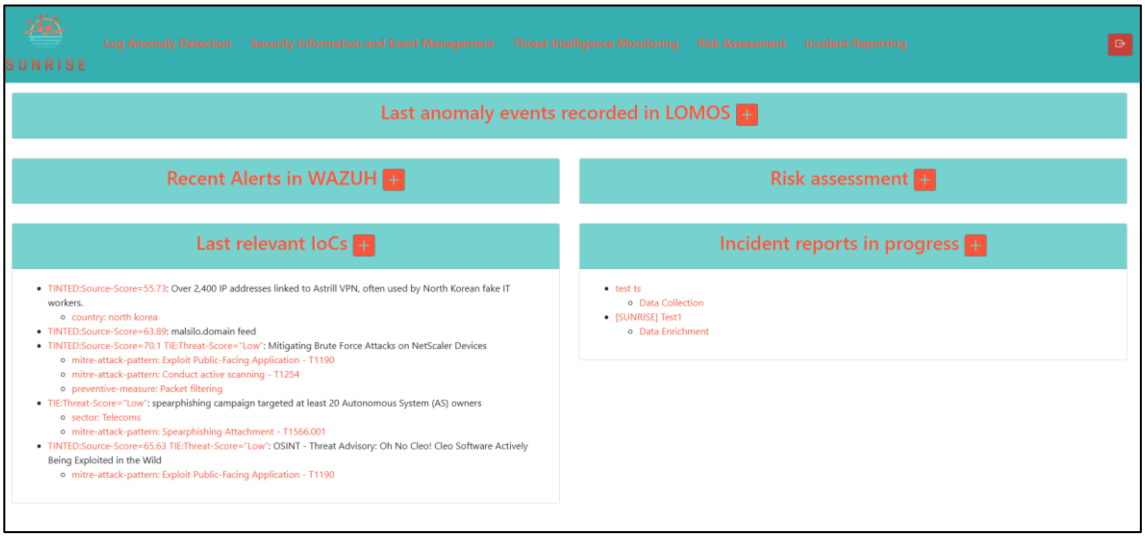
SUNRISE
Log Anomaly Detection
Security Information and Event Management
Threat Intelligence Monitoring
Risk Assessment
Incident Reporting
Last anomaly events recorded in LOMOS +
Recent Alerts in WAZUH +
Risk assessment +
Last relevant IoCs +
Incident reports in progress +

- TINTED:Source-Score=55.73: Over 2,400 IP addresses linked to Astrill VPN, often used by North Korean fake IT workers.
  - country: north korea
- TINTED:Source-Score=63.89: malsilo.domain feed
- TINTED:Source-Score=70.1 TIE:Threat-Score="Low": Mitigating Brute Force Attacks on NetScaler Devices
  - mitre-attack-pattern: Exploit Public-Facing Application - T1190
  - mitre-attack-pattern: Conduct active scanning - T1254
  - preventive-measure: Packet filtering
- TIE:Threat-Score="Low": spearphishing campaign targeted at least 20 Autonomous System (AS) owners
  - sector: Telecoms
  - mitre-attack-pattern: Spearphishing Attachment - T1566.001
- TINTED:Source-Score=65.63 TIE:Threat-Score="Low": OSINT - Threat Advisory: Oh No Cleo! Cleo Software Actively Being Exploited in the Wild
  - mitre-attack-pattern: Exploit Public-Facing Application - T1190

- test ts
  - Data Collection
- [SUNRISE] Test1
  - Data Enrichment

Figure 6. CPR Dashboard – Last relevant IoCs and incident reports in progress.

- **Incident reports in progress** (Figure 6): This shows a list of the open incident reports with the title of the case (as it is in TheHive) and the tasks that are currently opened. Click on these incidents for more information.

2 Anomaly Detection

In this section, we present a step-by-step user guide for working with LOMOS. The first step is model training. The models are then used in the second step, which is log-based anomaly detection inference.

2.1 Setting up data sources

Elasticsearch is often used as a central database for logs from different components of simple or complex systems. It is suitable for handling substantial amounts of data in distributed indices in JSON format, ensuring scalability and availability. Elastic Filebeat is a lightweight agent that can be used to push log data to Elasticsearch. LOMOS is capable of directly interacting with Elasticsearch. It can read data from it, process it, and write the results back.

To set up Filebeat agents, refer to the original and up-to-date documentation by Elastic [5]. Logs from different sources should be stored in separate indices and processed separately by the LOMOS. Elasticsearch has great support for data retention strategies. It is easy to set up the lifecycle policies in Kibana or through the REST endpoint[1].

2.2 Training a log parser

LOMOS can carry out some pre-processing steps on the logs. Logs are sometimes stored in a raw semi-structured format. In such a case, log messages and timestamps must be extracted first.

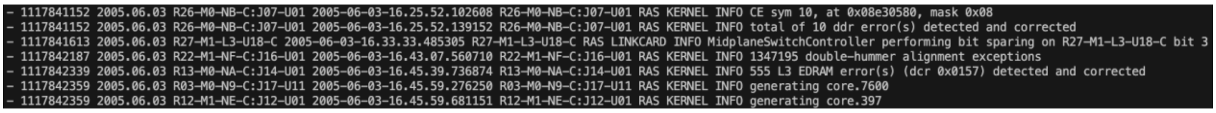

```
- 1117841152 2005.06.03 R26-M0-NB-C:J07-U01 2005-06-03-16.25.52.102608 R26-M0-NB-C:J07-U01 RAS KERNEL INFO CE sym 10, at 0x08e30580, mask 0x08
- 1117841152 2005.06.03 R26-M0-NB-C:J07-U01 2005-06-03-16.25.52.139152 R26-M0-NB-C:J07-U01 RAS KERNEL INFO total of 10 ddr error(s) detected and corrected
- 1117841613 2005.06.03 R27-M1-L3-U18-C 2005-06-03-16.33.33.485305 R27-M1-L3-U18-C RAS LINKCARD INFO MidplaneSwitchController performing bit sparing on R27-M1-L3-U18-C bit 3
- 1117842187 2005.06.03 R22-M1-NF-C:J16-U01 2005-06-03-16.43.07.560710 R22-M1-NF-C:J16-U01 RAS KERNEL INFO 1347195 double-hummer alignment exceptions
- 1117842339 2005.06.03 R13-M0-NA-C:J14-U01 2005-06-03-16.45.39.736874 R13-M0-NA-C:J14-U01 RAS KERNEL INFO 555 L3 EDRAM error(s) (dcr 0x0157) detected and corrected
- 1117842359 2005.06.03 R03-M0-N9-C:J17-U11 2005-06-03-16.45.59.276250 R03-M0-N9-C:J17-U11 RAS KERNEL INFO generating core.7600
- 1117842359 2005.06.03 R12-M1-NE-C:J12-U01 2005-06-03-16.45.59.681151 R12-M1-NE-C:J12-U01 RAS KERNEL INFO generating core.397
```

Figure 7. A sample of raw BGL logs.

The extraction of the relevant data can be performed by setting the regular expressions in the pre-processing step, as shown in Figure 8.

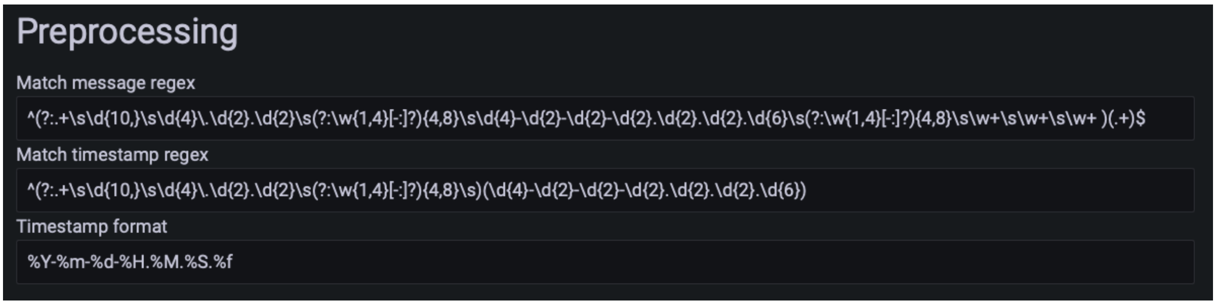

Preprocessing

Match message regex

```
^(?:.+\s\d{10,}\s\d{4}\.\d{2}.\d{2}\s(?:\w{1,4}[-:]?){4,8}\s\d{4}-\d{2}-\d{2}-\d{2}.\d{2}.\d{2}.\d{6}\s(?:\w{1,4}[-:]?){4,8}\s\w+\s\w+\s\w+ )(.+)$
```

Match timestamp regex

```
^(?:.+\s\d{10,}\s\d{4}\.\d{2}.\d{2}\s(?:\w{1,4}[-:]?){4,8}\s)(\d{4}-\d{2}-\d{2}-\d{2}.\d{2}.\d{2}.\d{6})
```

Timestamp format

```
%Y-%m-%d-%H.%M.%S.%f
```

Figure 8. Regular expressions for extracting messages and timestamps from raw logs and the timestamp format.

We generated the regular expressions with the help of the regex101[1] tool (Figure 9 and Figure 10). Based on the sample of the logs, regular expressions could be simpler. However, this would raise the risk of falsely identifying the components of the logs. If the data is already structured in the Elasticsearch index, we can put a regular expression that matches the whole line "^(.+)$" and leave the other two fields empty.

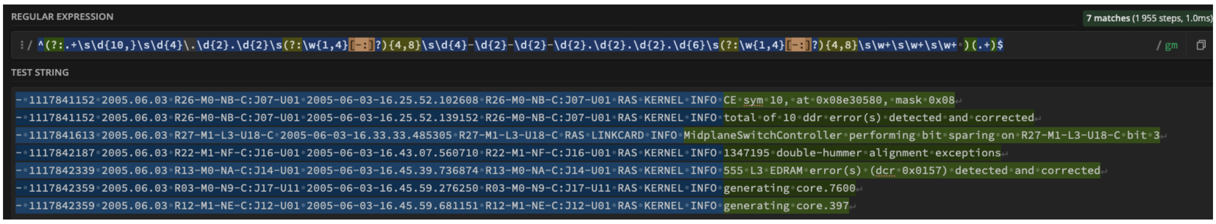

REGULAR EXPRESSION

7 matches (1 955 steps, 1.0ms)

```
/ ^(?:.+\s\d{10,}\s\d{4}\.\d{2}.\d{2}\s(?:\w{1,4}[-:]?){4,8}\s\d{4}-\d{2}-\d{2}-\d{2}.\d{2}.\d{2}.\d{6}\s(?:\w{1,4}[-:]?){4,8}\s\w+\s\w+\s\w+ )(.+)$ / gm
```

TEST STRING

```
- 1117841152 2005.06.03 R26-M0-NB-C:J07-U01 2005-06-03-16.25.52.102608 R26-M0-NB-C:J07-U01 RAS KERNEL INFO CE sym 10, at 0x08e30580, mask 0x08
- 1117841152 2005.06.03 R26-M0-NB-C:J07-U01 2005-06-03-16.25.52.139152 R26-M0-NB-C:J07-U01 RAS KERNEL INFO total of 10 ddr error(s) detected and corrected
- 1117841613 2005.06.03 R27-M1-L3-U18-C 2005-06-03-16.33.33.485305 R27-M1-L3-U18-C RAS LINKCARD INFO MidplaneSwitchController performing bit sparing on R27-M1-L3-U18-C bit 3
- 1117842187 2005.06.03 R22-M1-NF-C:J16-U01 2005-06-03-16.43.07.560710 R22-M1-NF-C:J16-U01 RAS KERNEL INFO 1347195 double-hummer alignment exceptions
- 1117842339 2005.06.03 R13-M0-NA-C:J14-U01 2005-06-03-16.45.39.736874 R13-M0-NA-C:J14-U01 RAS KERNEL INFO 555 L3 EDRAM error(s) (dcr 0x0157) detected and corrected
- 1117842359 2005.06.03 R03-M0-N9-C:J17-U11 2005-06-03-16.45.59.276250 R03-M0-N9-C:J17-U11 RAS KERNEL INFO generating core.7600
- 1117842359 2005.06.03 R12-M1-NE-C:J12-U01 2005-06-03-16.45.59.681151 R12-M1-NE-C:J12-U01 RAS KERNEL INFO generating core.397
```

Figure 9. Regular expression for extracting log message from a raw log.

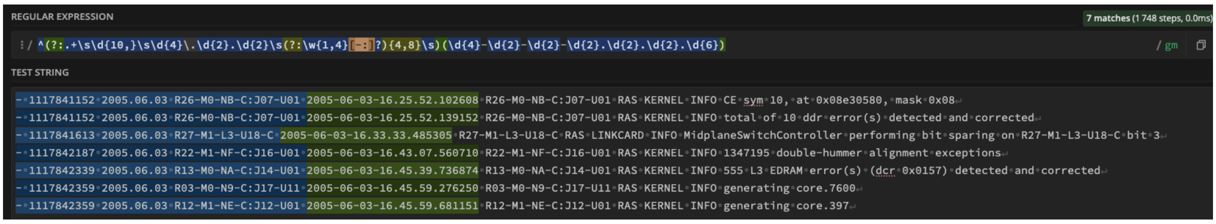

```
^(?:.+\s\d{10,}\s\d{4}\.\d{2}.\d{2}\s(?:\w{1,4}[-:]?){4,8}\s)(\d{4}-\d{2}-\d{2}-\d{2}.\d{2}.\d{2}.\d{6})
```

```
- 1117841152 2005.06.03 R26-M0-NB-C:J07-U01 2005-06-03-16.25.52.102608 R26-M0-NB-C:J07-U01 RAS KERNEL INFO CE sym 10, at 0x08e30580, mask 0x08
- 1117841152 2005.06.03 R26-M0-NB-C:J07-U01 2005-06-03-16.25.52.139152 R26-M0-NB-C:J07-U01 RAS KERNEL INFO total of 10 ddr error(s) detected and corrected
- 1117841613 2005.06.03 R27-M1-L3-U18-C 2005-06-03-16.33.33.485305 R27-M1-L3-U18-C RAS LINKCARD INFO MidplaneSwitchController performing bit sparing on R27-M1-L3-U18-C bit 3
- 1117842187 2005.06.03 R22-M1-NF-C:J16-U01 2005-06-03-16.43.07.560710 R22-M1-NF-C:J16-U01 RAS KERNEL INFO 1347195 double-hummer alignment exceptions
- 1117842339 2005.06.03 R13-M0-NA-C:J14-U01 2005-06-03-16.45.39.736874 R13-M0-NA-C:J14-U01 RAS KERNEL INFO 555 L3 EDRAM error(s) (dcr 0x0157) detected and corrected
- 1117842359 2005.06.03 R03-M0-N9-C:J17-U11 2005-06-03-16.45.59.276250 R03-M0-N9-C:J17-U11 RAS KERNEL INFO generating core.7600
- 1117842359 2005.06.03 R12-M1-NE-C:J12-U01 2005-06-03-16.45.59.681151 R12-M1-NE-C:J12-U01 RAS KERNEL INFO generating core.397
```

Figure 10. Regular expression for extracting timestamp from a raw log.

The next set of parameters is related to the Drain method which extracts log templates from log messages (Figure 11):

The **Similarity threshold** sets the minimal Jaccard index of log message words for them to match into the same log template.

The number of children sets the depth of the Drain search tree. It specifies how many initial words in a log must be an exact match for a log template. Increasing this number generally accelerates the Drain process. Nevertheless, a too-high value could lead to the incorrect identification of potential parameters located at the beginning of a log.

Extra delimiters can add more characters for splitting the message string into words. For example, we can add an underscore, which will be used beside the default space character.

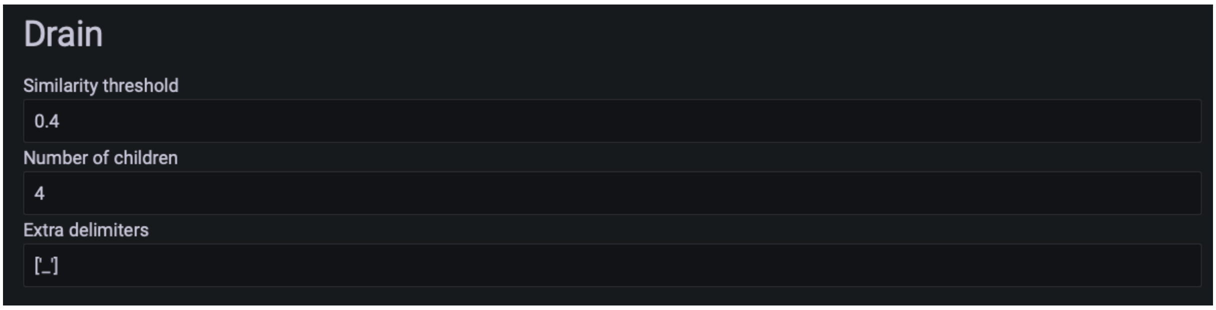

Drain

Similarity threshold
0.4

Number of children
4

Extra delimiters
['_']

Figure 11. Drain parameters.

Next, we add custom regular expressions for masking complex patterns, such as IP addresses or timestamps. Drain is used for automatic parameter extraction; however, it works better if we add some general or tailored regular expressions, as in Figure 12.

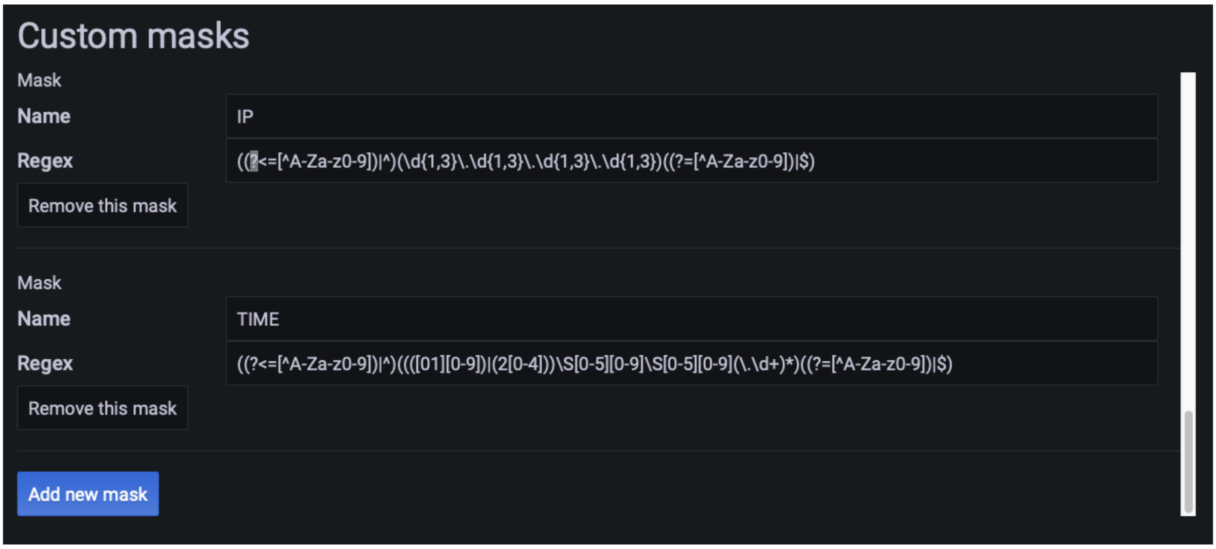

Custom masks

Mask
Name: IP
Regex: ((?<=[^A-Za-z0-9])|^)(\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})((?=[^A-Za-z0-9])|$)

Mask
Name: TIME
Regex: ((?<=[^A-Za-z0-9])|^)((([01][0-9])|(2[0-4]))\S[0-5][0-9]\S[0-5][0-9](\.\d+)*)((?=[^A-Za-z0-9])|$)

Figure 12. Custom masks for complex parameters.

The next section is used for setting the connection to Elasticsearch. We must define the IP address, port, and credentials (Figure 13). Leave the credentials empty if they are not required by the selected Elasticsearch deployment.

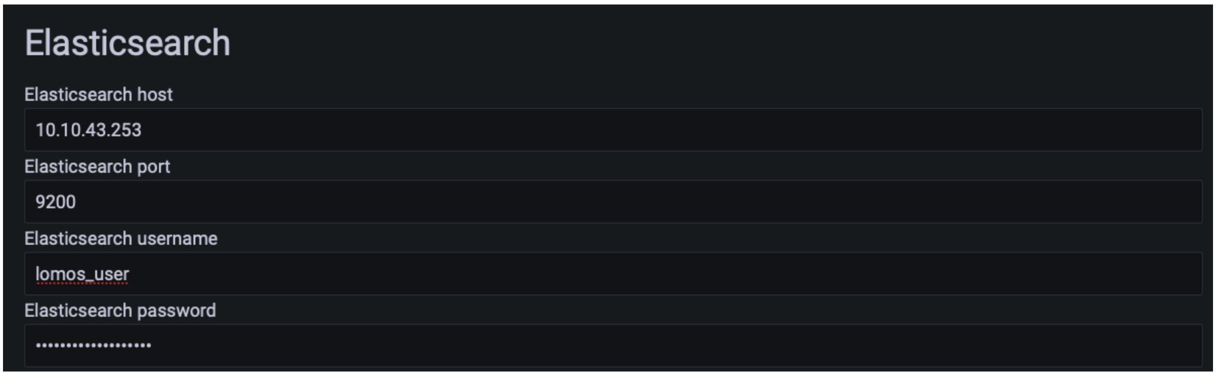

Elasticsearch

Elasticsearch host

10.10.43.253

Elasticsearch port

9200

Elasticsearch username

lomos_user

Elasticsearch password

••••••••••••••••••

Figure 13. Log parser training Elasticsearch connection details and credentials.

Next, we set the source index and message column names (Figure 14).

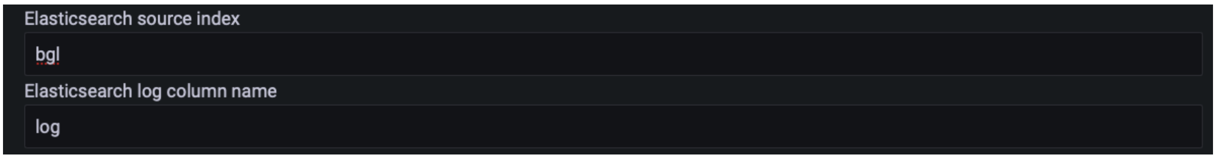

Elasticsearch source index

bgl

Elasticsearch log column name

log

Figure 14. Log parser training source index.

Then, we must select the data with the next set of parameters, as seen in Figure 15. First, select the period we will use for training and the field that will be used for filtering (timestamp or id). Additional Elasticsearch filters can be used to select only the relevant data. The prefix is used to set the name of new indices and is required later in the model training step to reference the parsed data. To start the training, click on the run parser button.

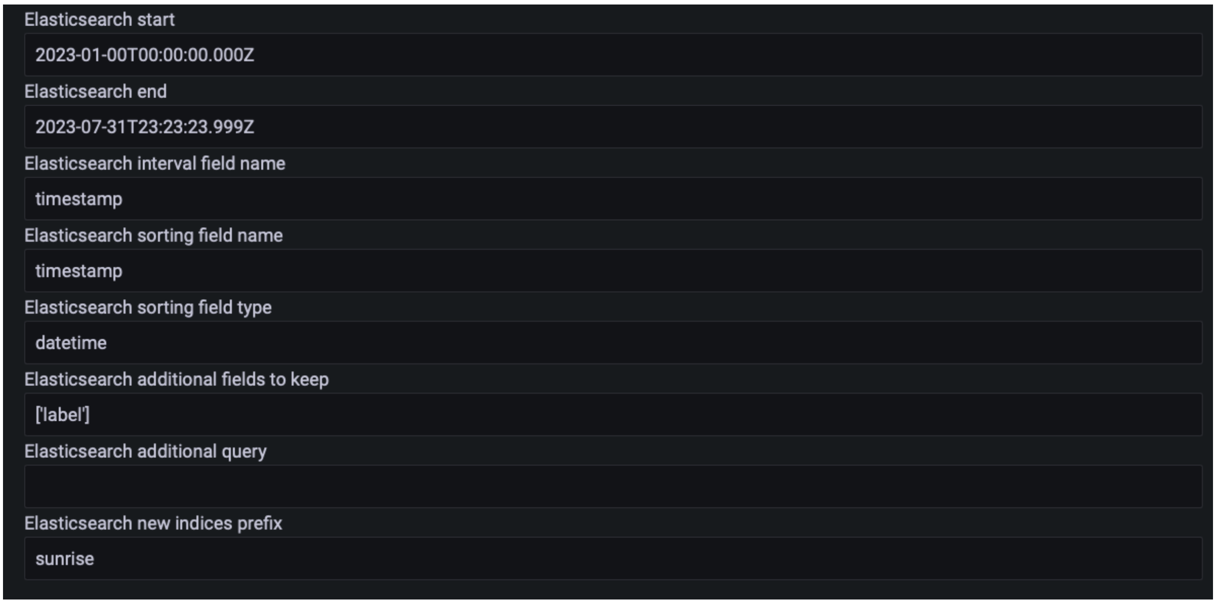

Elasticsearch start

2023-01-00T00:00:00.000Z

Elasticsearch end

2023-07-31T23:23:23.999Z

Elasticsearch interval field name

timestamp

Elasticsearch sorting field name

timestamp

Elasticsearch sorting field type

datetime

Elasticsearch additional fields to keep

['label']

Elasticsearch additional query

Elasticsearch new indices prefix

sunrise

Figure 15. Log parser training period selection and addition filters.

2.3 Inspecting the log parsing results

Once the log parser is trained, it is used on the training data to parse the log templates and push them to a new index in Elasticsearch, where the name of the index is generated based on the prefix set by the user in the previous step, source index name, and "_logs structured" suffix. Another index is created, which stores the unique log templates and relevant statistics describing their frequency and number of detected parameters. The suffix for this index is "_events". The example from the previous step generates "sunrise_bgl_logs_structured" and "sunrise_bgl_events" indices. The data is automatically accessible from the Grafana dashboard, where it can be explored through interactive visualizations. The user can evaluate if the log templates are parsed correctly and proceed with training the model or return to the previous step to adapt the parameters of the log parser and rerun the log parser training.

The first dashboard offers users an overview of the parsed log templates (Figure 16). There are two histograms in the upper row, showing the distributions of the ratio between the automatically extracted parameters to the number of words and masked parameters by the regular expressions to the number of words. The number on the right side shows the number of the unique extracted log templates. The histogram below shows the distribution of log message length (number of words). Finally, there is a table of extracted log templates at the bottom, together with relevant statistics. Users can check the statistics from the charts above for each of the log templates.

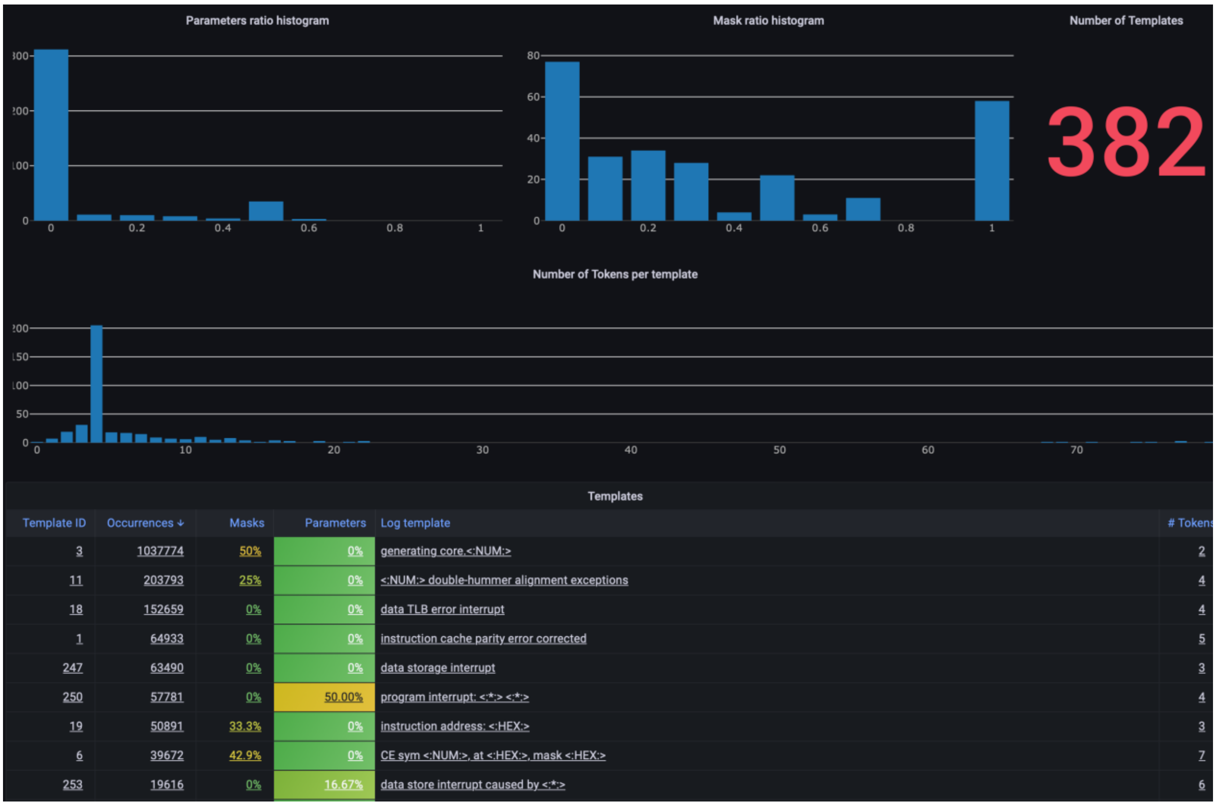

Templates

Template ID	Occurrences ↓	Masks	Parameters	Log template	# Tokens
3	1037774	50%	0%	generating core.<:NUM:>	2
11	203793	25%	0%	<:NUM:> double-hummer alignment exceptions	4
18	152659	0%	0%	data TLB error interrupt	4
1	64933	0%	0%	instruction cache parity error corrected	5
247	63490	0%	0%	data storage interrupt	3
250	57781	0%	50.00%	program interrupt: <:*:> <:*:>	4
19	50891	33.3%	0%	instruction address: <:HEX:>	3
6	39672	42.9%	0%	CE sym <:NUM:>, at <:HEX:>, mask <:HEX:>	7
253	19616	0%	16.67%	data store interrupt caused by <:*:>	6

Figure 16. Extracted log templates overview.

Users can click on the extracted log template to view its details, as seen in Figure 17. Besides the statistics already mentioned above, user can see actual log messages and their occurrences through time.

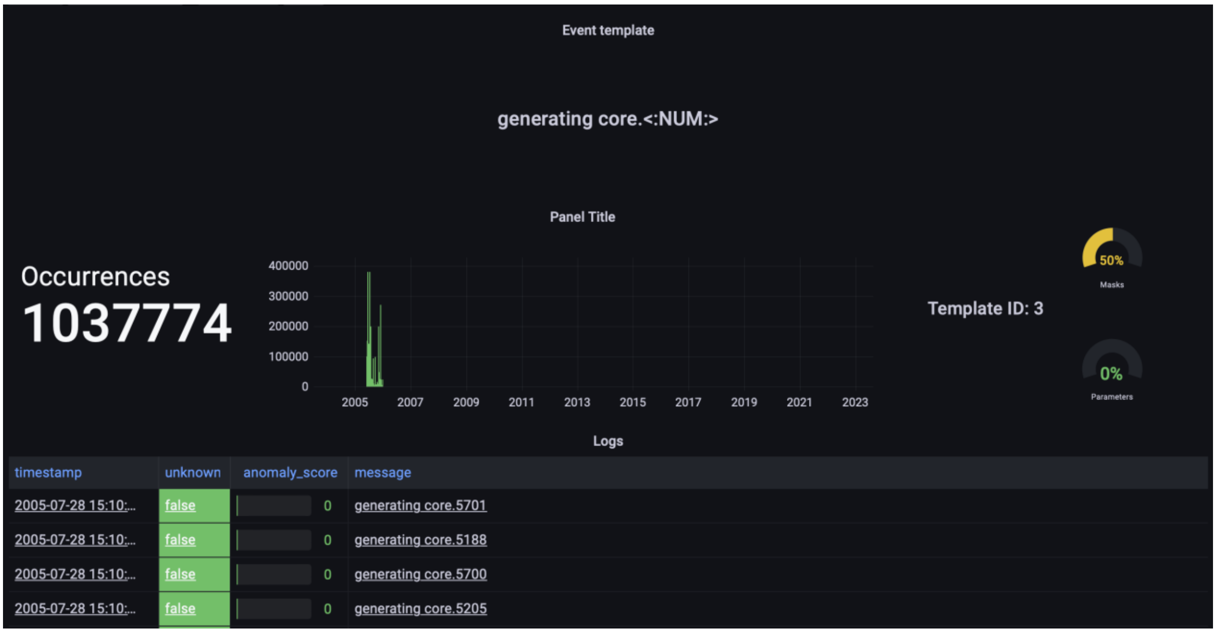

Event template

generating core.<:NUM:>

Panel Title

Occurrences
1037774

Template ID: 3

Logs

timestamp	unknown	anomaly_score	message
2005-07-28 15:10:...	false	0	generating core.5701
2005-07-28 15:10:...	false	0	generating core.5188
2005-07-28 15:10:...	false	0	generating core.5700
2005-07-28 15:10:...	false	0	generating core.5205

Figure 17. Log template details.

2.4 Training an anomaly detection model

When the log templates are properly parsed, we can proceed to train an anomaly detection model. The first part, shown in Figure 18, is dedicated to the Elasticsearch connection settings and the name of the structured logs index, which was explained in the previous step. The next two fields are used for the proper sorting of the logs.

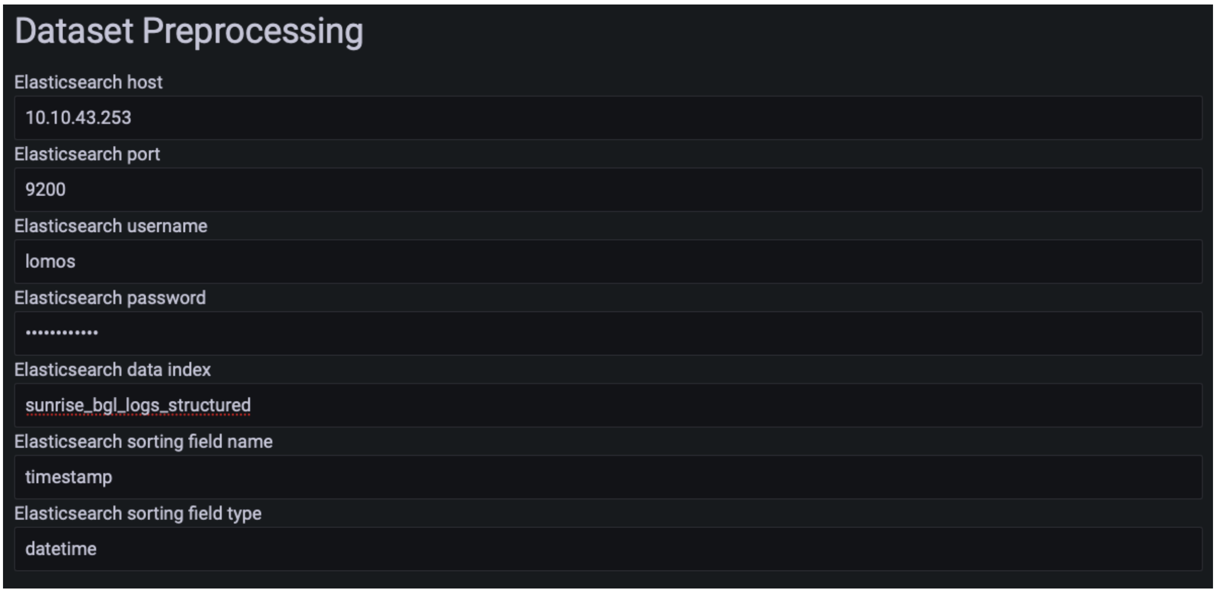

Dataset Preprocessing

Elasticsearch host
10.10.43.253

Elasticsearch port
9200

Elasticsearch username
lomos

Elasticsearch password
••••••••••••

Elasticsearch data index
sunrise_bgl_logs_structured

Elasticsearch sorting field name
timestamp

Elasticsearch sorting field type
datetime

Figure 18. Model training Elasticsearch connection details, credentials, and filters.

To conclude the section related to data, we must select the periods that will be used for training the model, as shown in Figure 19. At least one period is mandatory, but multiple can be set. Such a feature becomes useful when we want to (potentially) skip anomalous data. If we are aware of an anomaly that influenced the logs, we should exclude it from the training set.

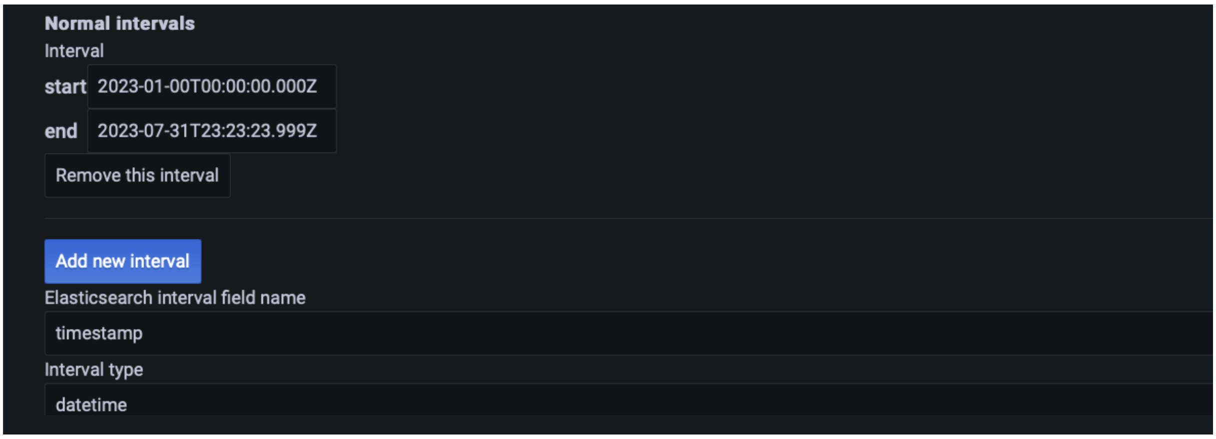

Normal intervals

Interval

start 2023-01-00T00:00:00.000Z

end 2023-07-31T23:23:23.999Z

Elasticsearch interval field name: timestamp

Interval type: datetime

Figure 19. Training data intervals.

The next set of parameters are machine learning hyperparameters (Figure 20). We must set the percentage of data used for training, where the left-out data is used for validation during the training. Next, we set the maximum number of epochs and early stop conditions. Then, we set the number of warm-up epochs, batch size, and window size. The last parameter is the name of the experiment for MLflow tracking. The correct values depend on the amount and complexity of data.

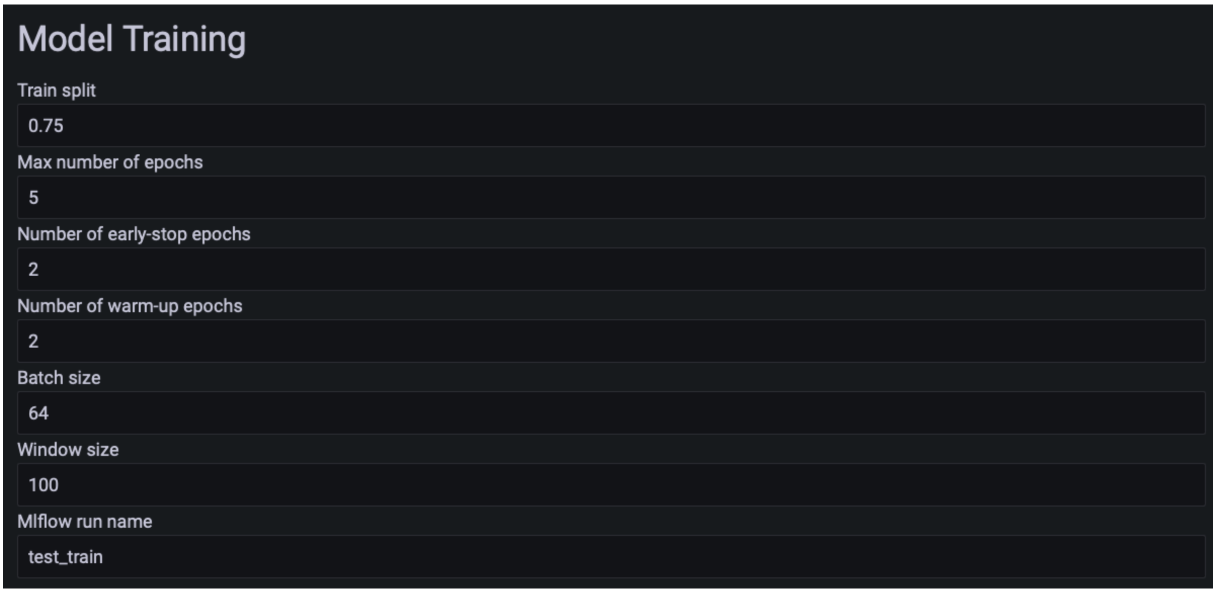

Model Training

Train split: 0.75

Max number of epochs: 5

Number of early-stop epochs: 2

Number of warm-up epochs: 2

Batch size: 64

Window size: 100

Mlflow run name: test_train

Figure 20. Anomaly detection model training hyperparameters.

2.5 Inspecting the training results

The training process can be monitored through the MLflow web dashboard. One of the more important indicators is the training and validation losses. MLflow enables users to explore those metrics in an interactive chart. Both training and validation losses should decay similarly, as seen in Figure 21.

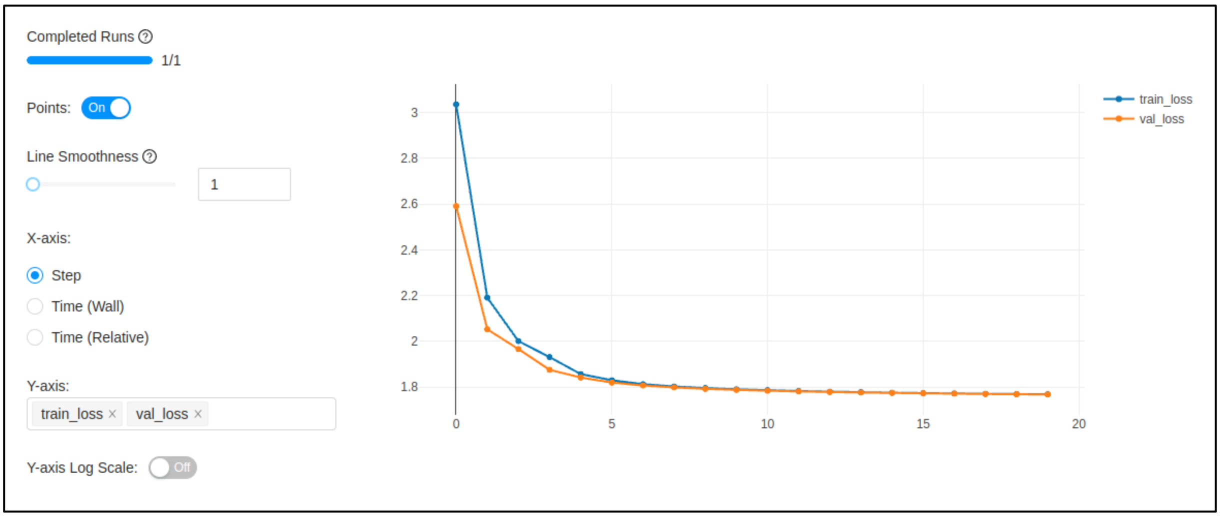

Completed Runs ⓘ
1/1

Points: On

Line Smoothness ⓘ 1

X-axis:
● Step
○ Time (Wall)
○ Time (Relative)

Y-axis:
train_loss × val_loss ×

Y-axis Log Scale: Off

Figure 21. Training anomaly detection loss chart displayed in MLflow.

2.6 Setting up live inference

Once the log template extraction and anomaly detection training phases are concluded, we can use the model for inference of new data. First, we load the parser configuration (Figure 22) by the MLflow experiment ID. The ID can be found in the MLflow web dashboard.

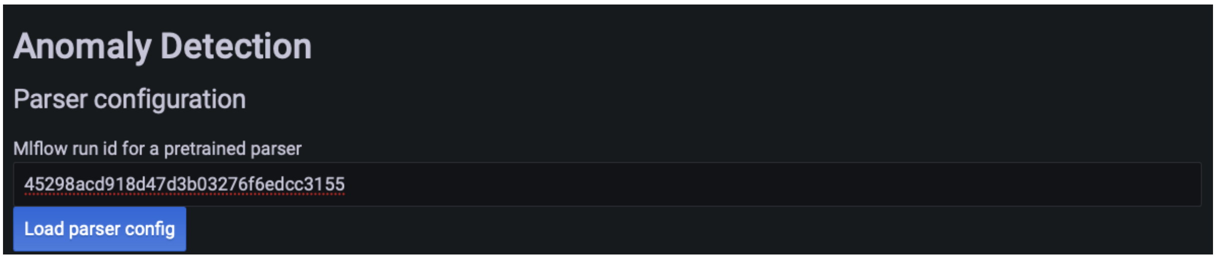

Anomaly Detection

Parser configuration

Mlflow run id for a pretrained parser

45298acd918d47d3b03276f6edcc3155

Figure 22. Pretrained parser MLflow experiment id.

Next, we set the inference task period in seconds. If we want to only execute the job once, we leave the field empty. In Figure 23 we set the job to execute every five minutes.

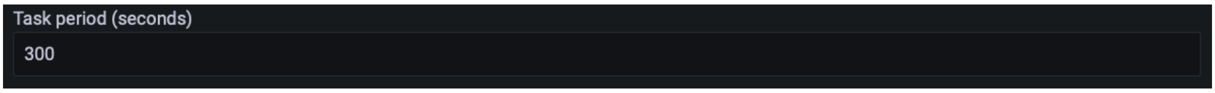

Task period (seconds)

300

Figure 23. Set the inference schedule period.

Then, we again set the Elasticsearch endpoint details and credentials as seen in Figure 24.

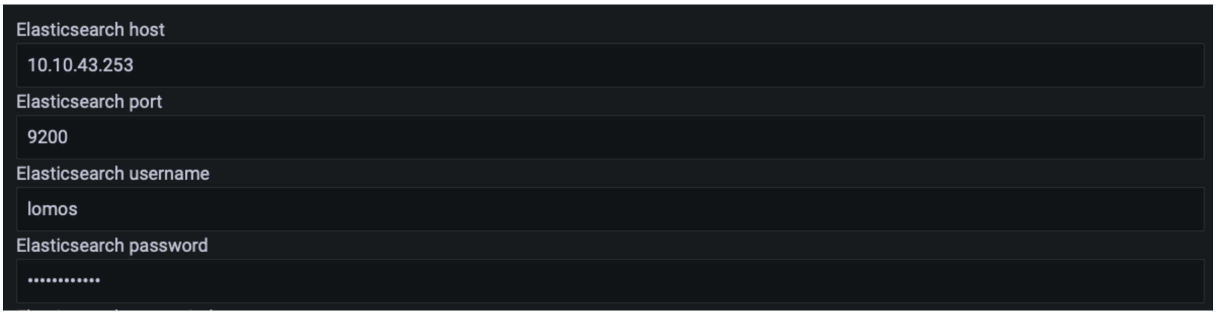

Elasticsearch host

10.10.43.253

Elasticsearch port

9200

Elasticsearch username

lomos

Elasticsearch password

••••••••••••

Figure 24. Anomaly detection inference ElasticSearch endpoint details and credentials.

Next, we set the name of the index with logs and the name of the log message column. Then, we must select the period of data that we will pass through the anomaly detector. We can again use timestamps, a numerical index, or special keywords: “where_left_off” and “now”. Those two keywords are useful for periodical jobs and will ensure the processing of new data at each execution. The configuration example is shown in Figure 25.

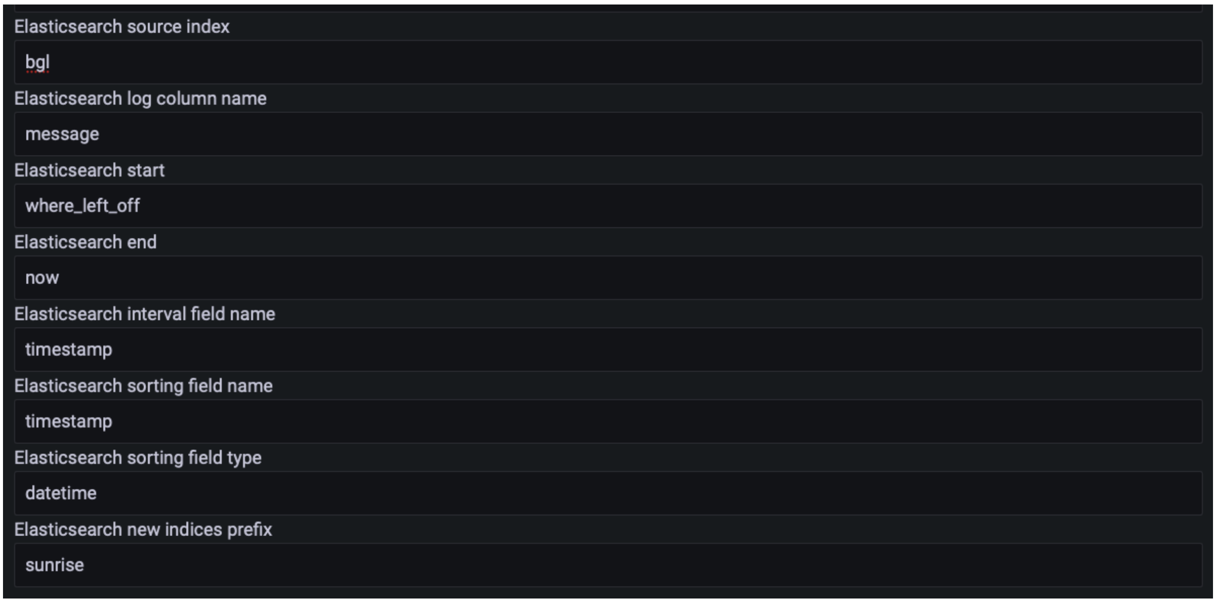

Elasticsearch source index
bgl
Elasticsearch log column name
message
Elasticsearch start
where_left_off
Elasticsearch end
now
Elasticsearch interval field name
timestamp
Elasticsearch sorting field name
timestamp
Elasticsearch sorting field type
datetime
Elasticsearch new indices prefix
sunrise

Figure 25. Elasticsearch index configuration and data filters.

Finally, we set the MLflow run ID of the trained model, batch size, window size, and MLflow run name, as seen in Figure 26. To run the inference, click on the “Run inference” button.

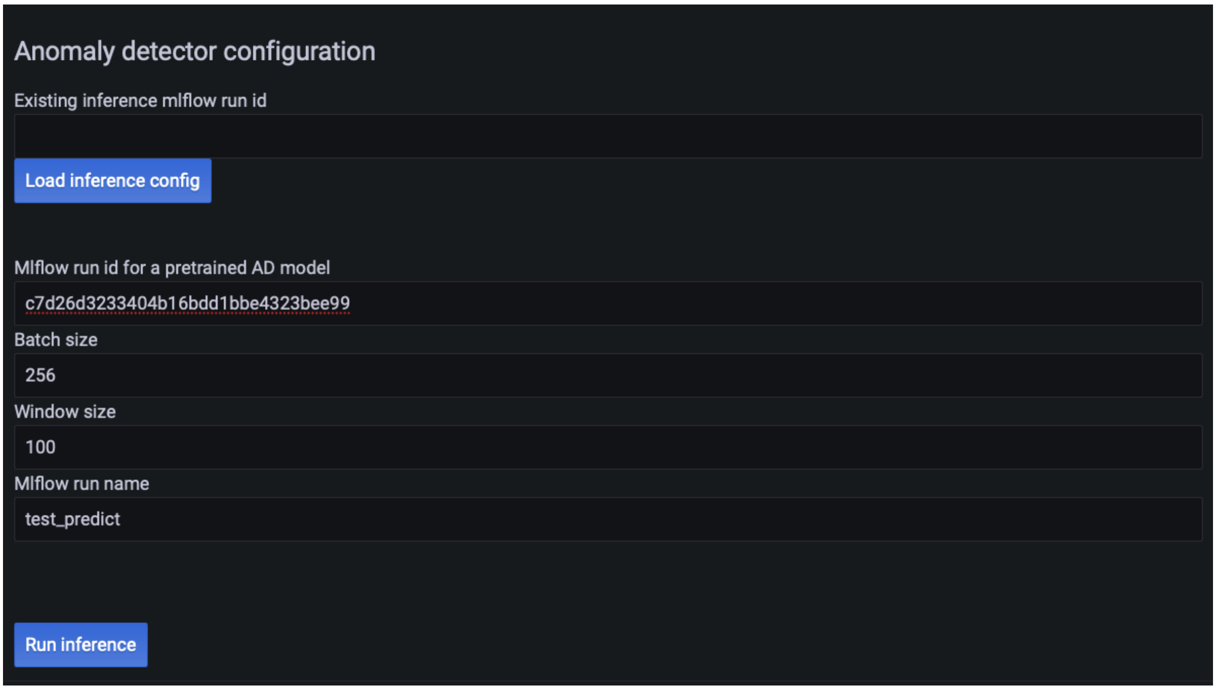

Anomaly detector configuration

Existing inference mlflow run id

Mlflow run id for a pretrained AD model
c7d26d3233404b16bdd1bbe4323bee99
Batch size
256
Window size
100
Mlflow run name
test_predict

Figure 26. Inference model configuration.

2.7 Inspecting live inference results

We finally get to inspect the live results. The default dashboard is presented below, but it is highly customizable since it is based on Grafana. In the first chart (Figure 27), we show the log count through time.

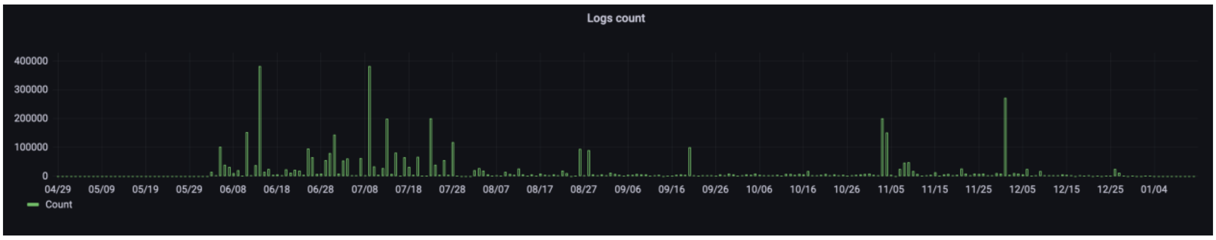

Figure 27. Histogram of logs through time (e.g., per day).

Next, we show the average anomaly score, as seen in Figure 28.

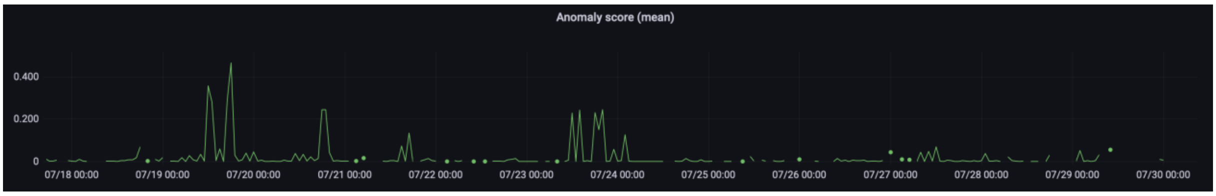

Figure 28. Average anomaly score.

For a better overview of the number of anomalies, charts like those in Figure 29 are useful. This chart shows the count of logs with high anomaly scores. The threshold is customizable and set to 0.7 as a default value.

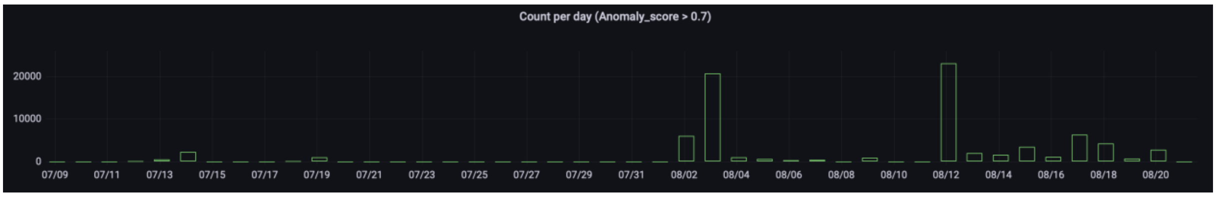

Figure 29. Number of logs with high anomaly score (e.g., above 0.7).

The table at the bottom of the dashboard (Figure 30) shows information about timestamps, log messages, log templates, anomaly scores, and whether the log template was recognized or not.

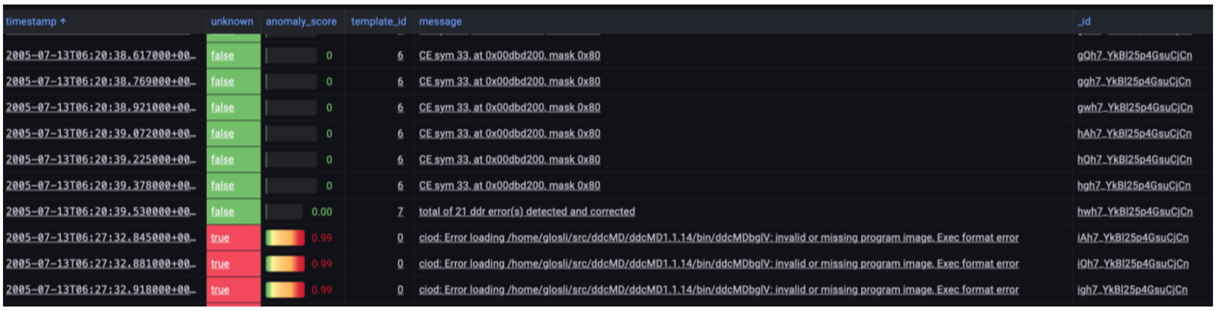

timestamp ↑	unknown	anomaly_score	template_id	message	_id
2005-07-13T06:20:38.617000+00…	false	0	6	CE sym 33, at 0x00dbd200, mask 0x80	gQh7_YkBl25p4GsuCjCn
2005-07-13T06:20:38.769000+00…	false	0	6	CE sym 33, at 0x00dbd200, mask 0x80	ggh7_YkBl25p4GsuCjCn
2005-07-13T06:20:38.921000+00…	false	0	6	CE sym 33, at 0x00dbd200, mask 0x80	gwh7_YkBl25p4GsuCjCn
2005-07-13T06:20:39.072000+00…	false	0	6	CE sym 33, at 0x00dbd200, mask 0x80	hAh7_YkBl25p4GsuCjCn
2005-07-13T06:20:39.225000+00…	false	0	6	CE sym 33, at 0x00dbd200, mask 0x80	hQh7_YkBl25p4GsuCjCn
2005-07-13T06:20:39.378000+00…	false	0	6	CE sym 33, at 0x00dbd200, mask 0x80	hgh7_YkBl25p4GsuCjCn
2005-07-13T06:20:39.530000+00…	false	0.00	7	total of 21 ddr error(s) detected and corrected	hwh7_YkBl25p4GsuCjCn
2005-07-13T06:27:32.845000+00…	true	0.99	0	ciod: Error loading /home/glosli/src/ddcMD/ddcMD1.1.14/bin/ddcMDbglV: invalid or missing program image, Exec format error	iAh7_YkBl25p4GsuCjCn
2005-07-13T06:27:32.881000+00…	true	0.99	0	ciod: Error loading /home/glosli/src/ddcMD/ddcMD1.1.14/bin/ddcMDbglV: invalid or missing program image, Exec format error	iQh7_YkBl25p4GsuCjCn
2005-07-13T06:27:32.918000+00…	true	0.99	0	ciod: Error loading /home/glosli/src/ddcMD/ddcMD1.1.14/bin/ddcMDbglV: invalid or missing program image, Exec format error	igh7_YkBl25p4GsuCjCn

Figure 30. Table of logs with anomaly scores.

Grafana enables users to focus on the periods of interest (e.g. periods with high anomaly scores) by simply selecting the period in any of the charts. This creates a time-based filter. Furthermore, filters based on any other field are supported. For example, users can select to show only logs with anomaly scores above 0.7. Users can then inspect the logs in the table and can react appropriately to address the issues found by the system.

2.8 CLI interaction with LOMOS

The lomos-cli is a CLI tool intended for frequent log parsing, training in inference. It is written in Python and is packaged in a Python package (.whl). It can be installed using standard “pip install”. It is not published on the public pypy.com repository, so the .whl file needs to be downloaded separately.

The needed parameters have 1-to-1 correspondence to GUI dashboard. They are saved in a JSON file. Example JSON files are included with the .whl file. An example of log parsing is shown below:

File sunrise-insiel-sesamo-backend_parse.json

```
{
  "PREPROCESSING": {
   "match_message": "^(?:[\\d\\- \\:\\.\\+]{19,32}) (?:[\\w\\-\\.]+) (.*)$",
   "match_timestamp": "^([\\d\\- \\:\\.\\+]{19,32}) (?:.*)$",
   "timestamp_fmt": "%Y-%m-%d %H:%M:%S.%f",
   "match_log_level": null,
   "match_label": null
  },
  "SNAPSHOT": {
   "snapshot_interval_minutes": 10,
   "compress_state": true
  },
  "MASKING": {
   "mask_prefix": "<:",
   "mask_suffix": ":>",
   "masking": [
     OMMITED-FOR-BREVITY
   ]
  },
  "DRAIN": {
   "max_children": 100,
   "max_clusters": 3000,
   "warning_special_wildcards_ratio_threshold": 0.95,
   "warning_auto_wildcards_ratio_threshold": 0.5,
   "sim_th": 0.4,
   "depth": 4,
   "extra_delimiters": "",
   "parametrize_numeric_tokens": false
  },
  "PROFILING": {
   "enabled": true,
   "report_sec": 30
  },
  "elasticsearch_id_field_name": "_id",
  "export_events_format": "elasticsearch",
  "export_logs_format": "elasticsearch",
  "mode": "training",
  "elasticsearch_host": "PROVIDED-VIA-DEPLOY-ENV",
  "elasticsearch_port": PROVIDED-VIA-DEPLOY-ENV,
  "elasticsearch_source_index": "PROVIDED-VIA-CLI",
  "elasticsearch_log_col_name": "logs_full",
  "elasticsearch_start": "2024-01-17T00:00:00",
  "elasticsearch_end": "2024-02-01T00:00:00",
  "elasticsearch_interval_field_name": "timestamp",
  "elasticsearch_sorting_field_name": "timestamp",
  "elasticsearch_sorting_field_type": "datetime",
  "elasticsearch_additional_query_field": null,
  "elasticsearch_fields_to_keep": ["label", "level"],
  "export_name_prefix": "test",
  "mlflow_run_id_pretrained_drain": null,
  "mlflow_tracking_uri": "",
  "mlflow_run_name": "test_parser",
  "task_name": "train_log_parser"
}
```

In similar way, we need to provide three files:

- sunrise-insiel-sesamo-backend_parse.json (shown above)
- sunrise-insiel-sesamo-backend_train.json
- sunrise-insiel-sesamo-backend_infer.json

With the log data stored in the Elasticsearch index named sesamo_logs, the parsing, training in inference tasks are run as:

```
lomos-cli sunrise-insiel-sesamo-backend parse sesamo_logs
# in mlfow was created mlflow experiment with id=parse_id
lomos-cli sunrise-insiel-sesamo-backend train sesamo_logs
# in mlfow was created mlflow experiment with id=train_id
lomos-cli sunrise-insiel-sesamo-backend infer sesamo_logs parse_id train_id
```

2.9 Access LOMOS results via API

Lomos2-api provides REST API access to LOMOS results. It is implemented as a REST API. It isolates the consuming application from the LOMOS internal working. Also, the consuming application does not need direct access to the LOMOS internal database to find anomalies.

Example REST request:

```
curl http://127.0.0.1:5000/api/top_anomaly?min_anomaly_score=0.7&from_timestamp=2024-01-08T00:00:01.200000Z&to_timestamp=2024-01-18T00:00:01.200000Z"
```

Example REST response:

```
{
  "aggregate": {
    "count": 1,
    "max_anomaly_score": 0.7692307978868484,
    "min_anomaly_score": 0.7692307978868484
  },
  "messages": [
    {
      "_index": "some_app_backend",
      "_type": "_doc",
      "_id": "D01aa40BjZ7m1DvJ_-y0",
      "_score": 1,
      "_source": {
        "timestamp": "2024-01-17T10:06:03.000014",
        "logs_full": "2024-01-17 10:06:03.000014 DEBUG BaseJdbcLogger:159 - <==    Updates: 1",
        "anomaly_score": 0.7692307978868484
      }
    }
  ]
}
```

3 Threat Intelligence

The interaction within the threat intelligence module can be done through two paths: the graphical user interface and the API.

The following shows and describes the sharing capabilities.

3.1 Use of WEB-GUI for secure data sharing

The WEB-GUI shows a login form (Figure 31) that verifies the credentials against the ones stored in Keycloak.

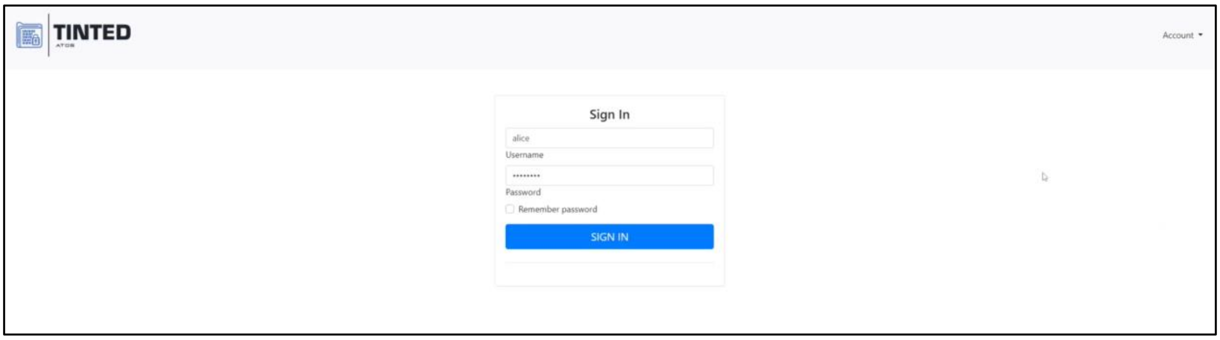

Username: alice

Password: ••••••••

☐ Remember password

Figure 31. TINTED Login.

Supposing that we have logged in as *Alice* for the first time, then we must configure the platform. We have to indicate the different parameters related to the MISP instance such as the URL and the API key (Figure 32), apart from the passphrase for encryption and other information about the events shared within MISP (Figure 33).

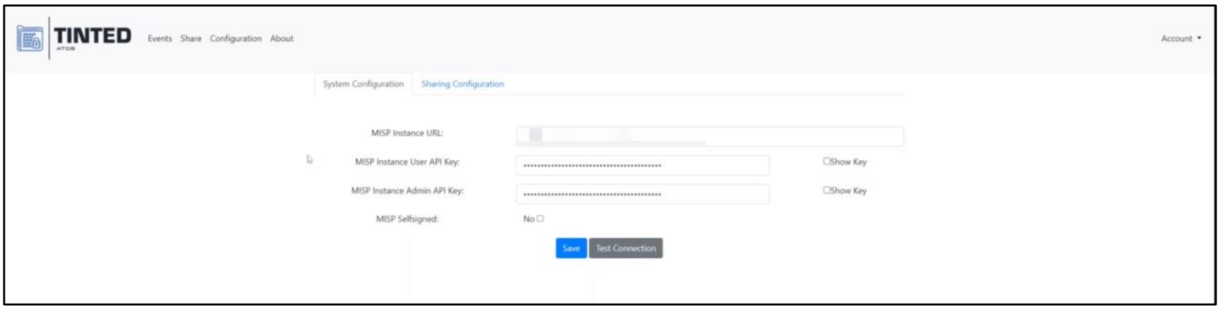

MISP Instance URL:

MISP Instance User API Key: •• ☐Show Key

MISP Instance Admin API Key: •• ☐Show Key

MISP Selfsigned: No ☐

Figure 32. System Configuration.

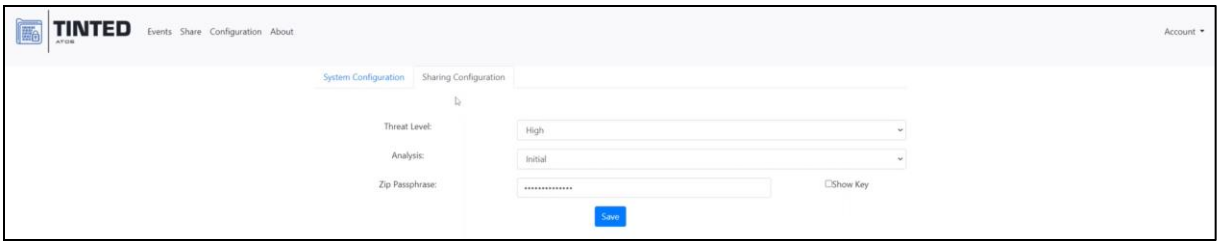

Threat Level: High

Analysis: Initial

Zip Passphrase: •••••••••••••• ☐Show Key

Figure 33. Sharing Configuration.

After configuring the required parameters, we will be redirected to the main page. In this case, as it is the first time that we access the platform we will observe the absence of events (Figure 34).

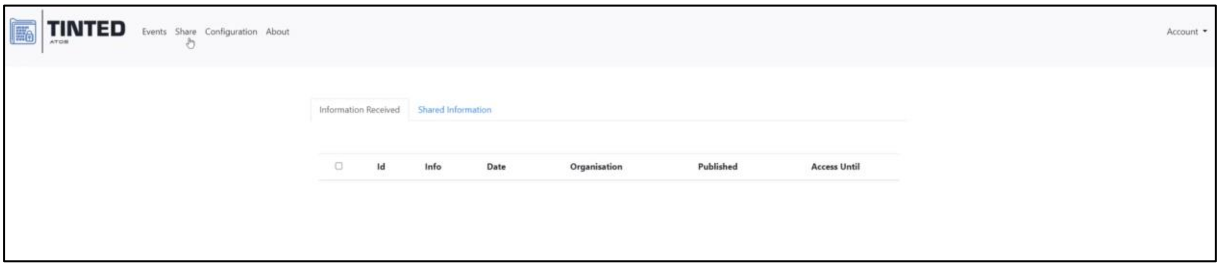

□	Id	Info	Date	Organisation	Published	Access Until

Figure 34. Events on the main page (currently empty).

Moving to the Share webpage, we can see the fields that are available, as illustrated in Figure 35.

For this guideline, we are going to show the process of sending an event from *Alice* to *Bob*.

To initiate the process, a JSON file that follows the MISP event structure is submitted. This file serves to extract the attributes within it, allowing for the selection of desired privacy treatments, which include encryption, anonymization, or maintaining data in cleartext form. Subsequent steps involve populating information fields such as Incident Date (optional, formatted as 'YYYY-MM-DD'), Event Tags (optional keywords describing the event), and Event Info (mandatory event description).

Further actions involve choosing recipients from a Keycloak-loaded list, encompassing various user types like individuals, organizations, sharing groups, or platform roles. Dates are then selected to determine the availability timeframe for the information. Once this period elapses, the event becomes inaccessible on the platform. This information, including the start and end dates and the involved users, is stored in the "sharing_agreement" file, attached to the event as depicted in Figure 37.

Lastly, the MISP objects and attributes are contained within a dynamic table. This table is populated with data sourced from the uploaded JSON file. Figure 35 displays a MISP Event with malicious IPs as attributes. Users have the freedom to append or remove attributes while also choosing the desired transformation type. The default choice is "cleartext," which maintains data as is. Alternatively, data protection options of encryption or anonymization can be selected if desired.

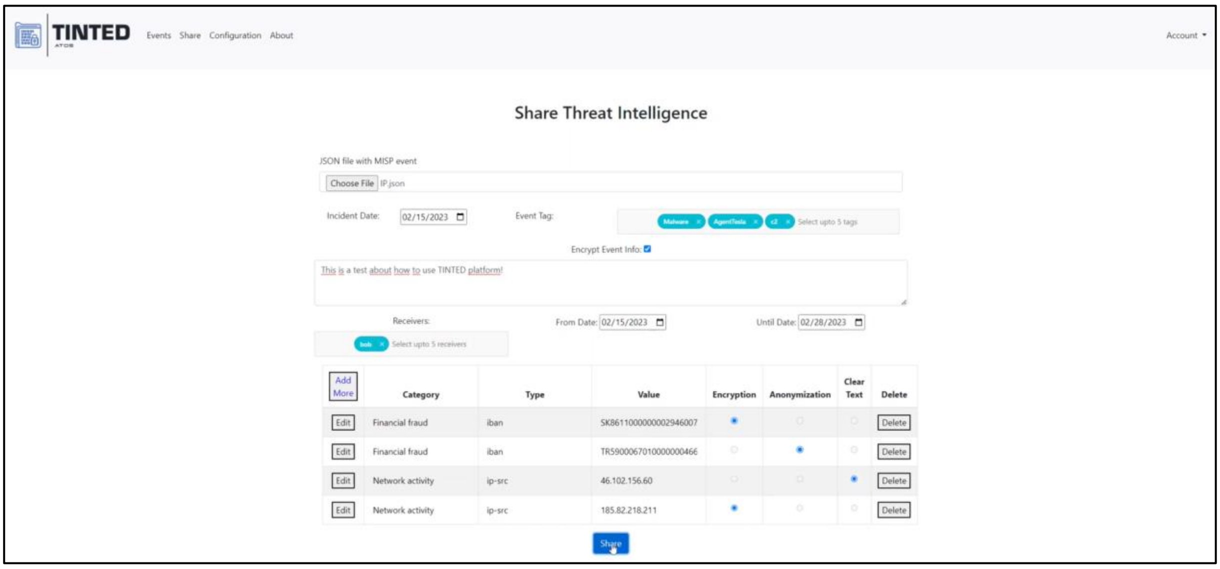

Share Threat Intelligence

JSON file with MISP event: IP.json

Incident Date: 02/15/2023

Event Tag: Malware ×, AgentTesla ×, c2 × Select upto 5 tags

Encrypt Event Info: ☑

This is a test about how to use TINTED platform!

Receivers: bob × Select upto 5 receivers

From Date: 02/15/2023

Until Date: 02/28/2023

Add More	Category	Type	Value	Encryption	Anonymization	Clear Text	Delete
Edit	Financial fraud	iban	SK8611000000002946007	●	○	○	Delete
Edit	Financial fraud	iban	TR590006701000000000466	○	●	○	Delete
Edit	Network activity	ip-src	46.102.156.60	○	○	●	Delete
Edit	Network activity	ip-src	185.82.218.211	●	○	○	Delete

Figure 35. Information Sharing Form.

After sharing the event, we can observe how it has arrived to the MISP instance. Figure 36 shows that the field Event info is encrypted and if we access the event itself (Figure 37), we also see the different privacy transformations that have been carried out.

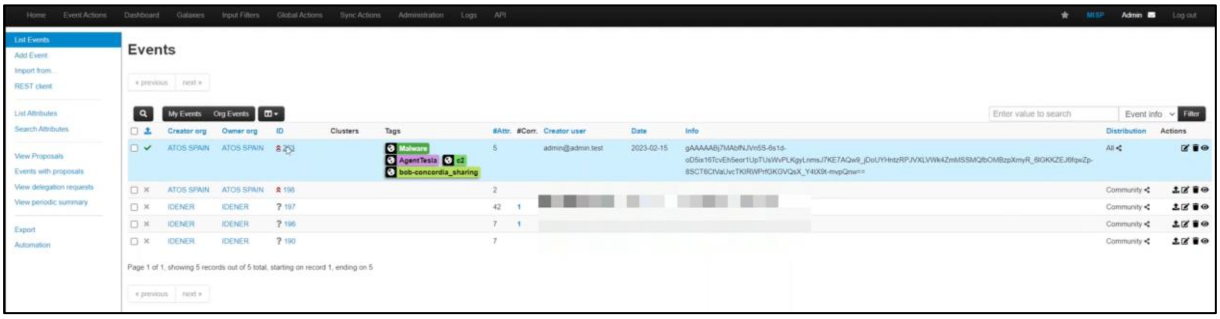

Events

□		Creator org	Owner org	ID	Clusters	Tags	#Attr.	#Corr.	Creator user	Date	Info	Distribution	Actions
□	✓	ATOS SPAIN	ATOS SPAIN	263		Malware AgentTesla c2 bob-concordia_sharing	5		admin@admin.test	2023-02-15	gAAAAABj7MAbfNJVn5S-6s1d-oD5ix16TcvEh5eor1UpTUsWvPLKgyLnmsJ7KE7AQw9_jDoUYHntzRPJVXLVWk4ZmMSSMQfbOMBzpXmyR_6lGKKZEJ6fqwZp-8SCT6CfVaUvcTKlRWPrfGKGVQsX_Y4tX9t-mvpQnw==	All	
□	×	ATOS SPAIN	ATOS SPAIN	198			2					Community	
□	×	IDENER	IDENER	197			42	1	[redacted]			Community	
□	×	IDENER	IDENER	196			7	1	[redacted]			Community	
□	×	IDENER	IDENER	190			7					Community	

Page 1 of 1, showing 5 records out of 5 total, starting on record 1, ending on 5

Figure 36. MISP instance.

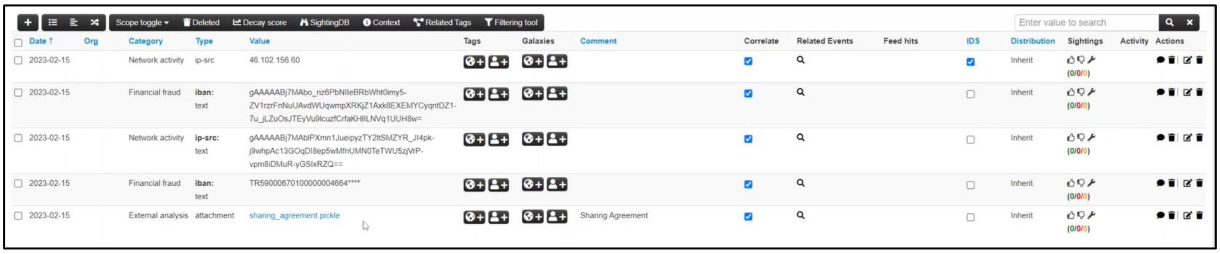

□	Date ↑	Org	Category	Type	Value	Tags	Galaxies	Comment	Correlate	Related Events	Feed hits	IDS	Distribution	Sightings	Activity	Actions
□	2023-02-15		Network activity	ip-src	46.102.156.60				☑			☑	Inherit	(0/0/0)		
□	2023-02-15		Financial fraud	**iban:** text	gAAAAABj7MAbo_riz6PbNlleBRbWht0imy5-ZV1rzrFnNuUAvdWUqwmpXRKjZ1Axk8EXEMYCyqntDZ1-7u_jLZuOsJTEyVu9lcuzfCrfaKHIlLNVq1UUH8w=				☑			□	Inherit	(0/0/0)		
□	2023-02-15		Network activity	**ip-src:** text	gAAAAABj7MAbIPXmn1JueipyzTY2ltSMZYR_Jl4pk-j9whpAc13GOqDI8ep5wMfnUMN0TeTWU5zjVrP-vpm8iDMuR-yGSlxRZQ==				☑			□	Inherit	(0/0/0)		
□	2023-02-15		Financial fraud	**iban:** text	TR590006701000000004664****				☑			□	Inherit	(0/0/0)		
□	2023-02-15		External analysis	attachment	sharing_agreement.pickle			Sharing Agreement	☑			□	Inherit	(0/0/0)		

Figure 37. Individual Event details – MISP.

As the MISP instance does not make a distinction between users, it is crucial to protect the information. In this case, if Bob wants to read the specific information that has been shared with him, he just needs to login to the platform with his credentials (as in Figure 31). Then, he will see that a new event has been received on the dashboard (Figure 38).

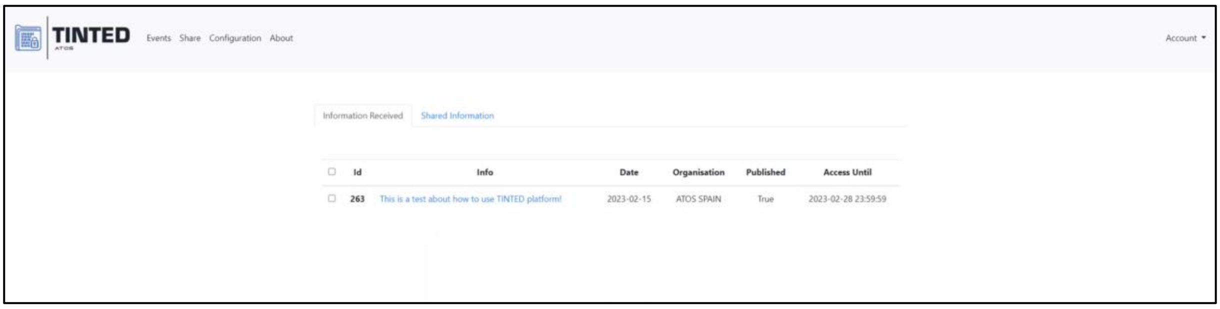

□	Id	Info	Date	Organisation	Published	Access Until
□	263	This is a test about how to use TINTED platform!	2023-02-15	ATOS SPAIN	True	2023-02-28 23:59:59

Figure 38. TINTED Information Received dashboard.

If we click on the event itself, we can get the information after the decryption process (Figure 39).

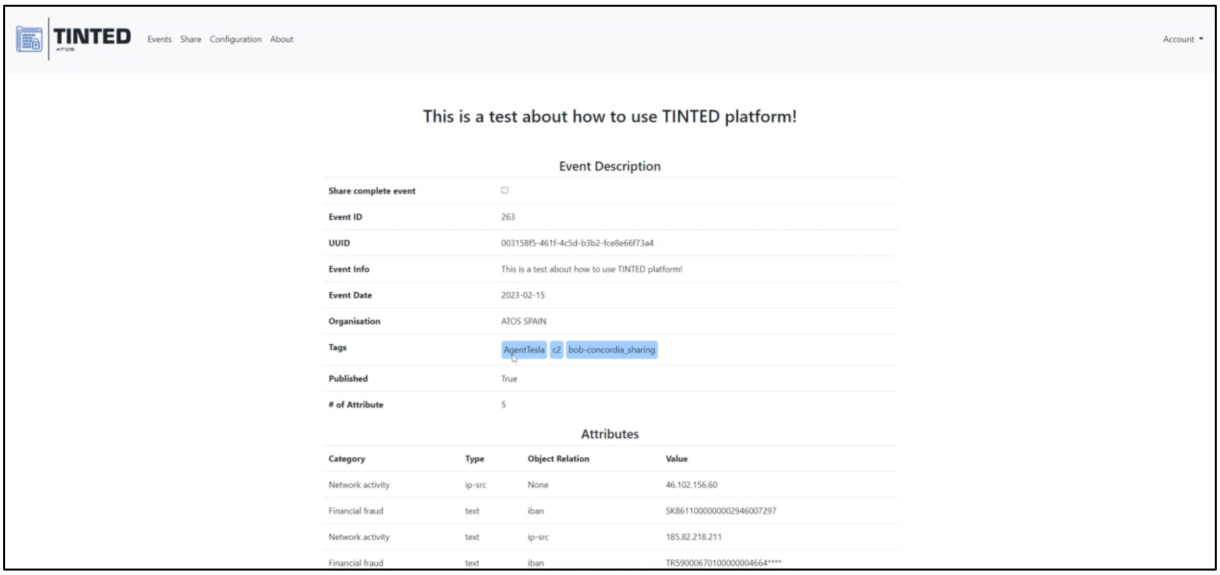

This is a test about how to use TINTED platform!

Event Description

Share complete event	□
Event ID	263
UUID	003158f5-461f-4c5d-b3b2-fce8e66f73a4
Event Info	This is a test about how to use TINTED platform!
Event Date	2023-02-15
Organisation	ATOS SPAIN
Tags	AgentTesla c2 bob-concordia_sharing
Published	True
# of Attribute	5

Attributes

Category	Type	Object Relation	Value
Network activity	ip-src	None	46.102.156.60
Financial fraud	text	iban	SK8611000000002946007297
Network activity	text	ip-src	185.82.218.211
Financial fraud	text	iban	TR590006701000000004664****

Figure 39. Individual Event details – TINTED.

If we want to manage the platform with a privileged role we can set up the administrator role. If the user possesses the administrator role within the platform, they will have the capability to oversee other users. Figure 40 and Figure 41 depict the distinct users registered in TINTED and the functionality to sign up a new user, respectively. This dashboard maintains a dynamic link to the data stored in Keycloak.

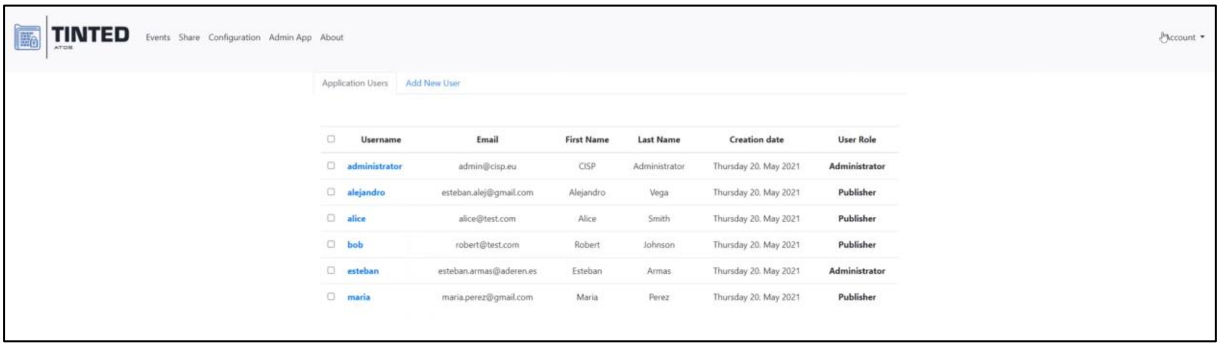

□	Username	Email	First Name	Last Name	Creation date	User Role
□	administrator	admin@cisp.eu	CISP	Administrator	Thursday 20. May 2021	Administrator
□	alejandro	esteban.alej@gmail.com	Alejandro	Vega	Thursday 20. May 2021	Publisher
□	alice	alice@test.com	Alice	Smith	Thursday 20. May 2021	Publisher
□	bob	robert@test.com	Robert	Johnson	Thursday 20. May 2021	Publisher
□	esteban	esteban.armas@aderen.es	Esteban	Armas	Thursday 20. May 2021	Administrator
□	maria	maria.perez@gmail.com	Maria	Perez	Thursday 20. May 2021	Publisher

Figure 40. Administrator Management Dashboard I.

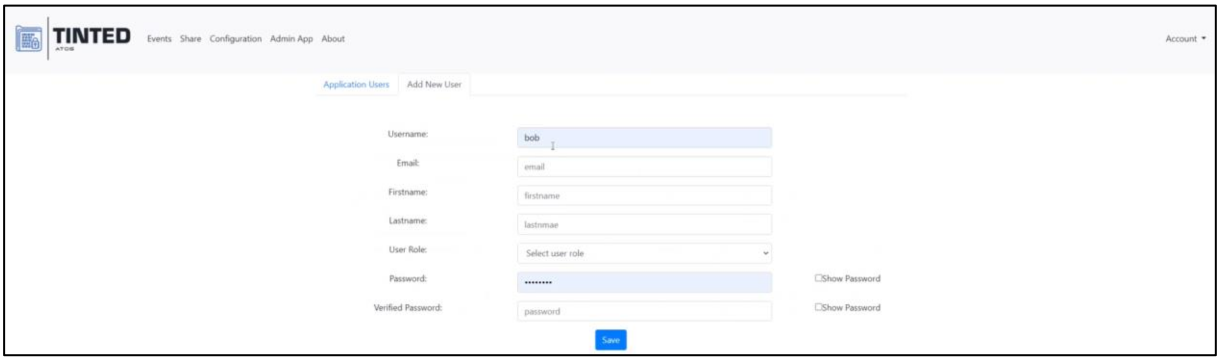

Username: bob

Email: email

Firstname: firstname

Lastname: lastnmae

User Role: Select user role

Password: •••••••• □Show Password

Verified Password: password □Show Password

Figure 41. Administrator Management Dashboard II.

3.2 Use of Application Programming Interface (API) for secure data sharing

As mentioned earlier, the GUI represents one of the two methods through which we can interact with TINTED. Its purpose is to act as an intermediary layer between the user and the orchestrator API. Nevertheless, there exists a direct means of communicating with the orchestrator API. Presently, the orchestrator API operates as a stateless application, requiring configuration parameters inputted through the GUI to be transmitted with each request to the orchestrator API. These parameters encompass:

- Information pertinent to MISP configuration, specifically the URL of the targeted MISP instance for event sharing and its API key.
- Sharing Agreement details, outlining the sender of the event and its intended recipient. The sender must match the user authenticated in the parameters.

This approach allows us to replicate the graphical interface's functionalities using the API.

In cases where we retrieve a list of MISP attributes, we can introduce the "transformation" field to these attributes. This empowers us to determine the transformation to be applied: encryption, anonymization, or "clear-text". Additionally, API requests can incorporate an extra parameter, termed "user_policies". This parameter permits the definition of default behaviour for attributes across one or multiple events. Consequently, we can acquire a list of events from MISP and apply a policy defined by the user. To download a list of events, we can do it through the GUI of the instance, as shown in Figure 42 and Figure 43, or through an API request to MISP, as shown in Figure 44.

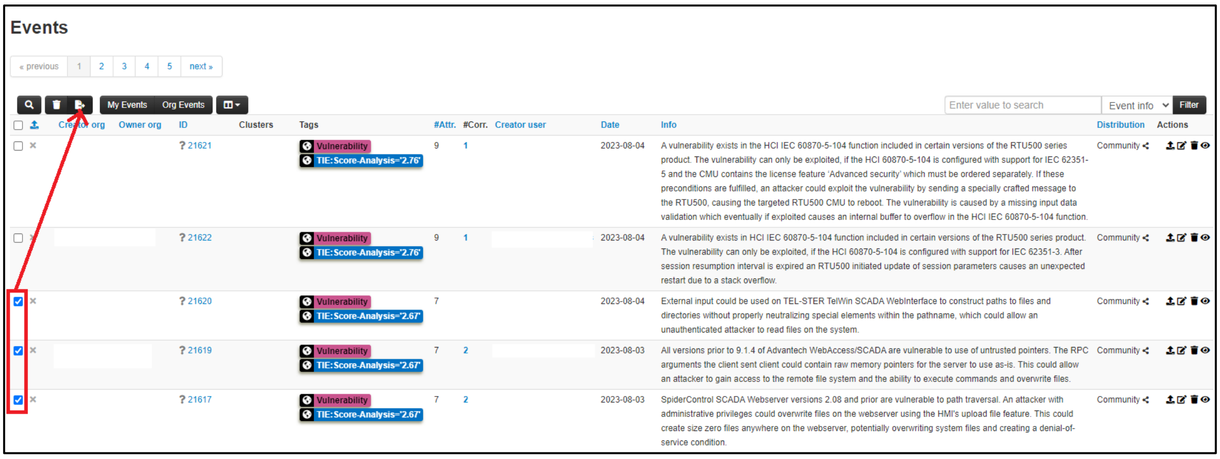

□		Creator org	Owner org	ID	Clusters	Tags	#Attr.	#Corr.	Creator user	Date	Info	Distribution	Actions
□	×			? 21621		Vulnerability TIE:Score-Analysis='2.76'	9	1		2023-08-04	A vulnerability exists in the HCI IEC 60870-5-104 function included in certain versions of the RTU500 series product. The vulnerability can only be exploited, if the HCI 60870-5-104 is configured with support for IEC 62351-5 and the CMU contains the license feature 'Advanced security' which must be ordered separately. If these preconditions are fulfilled, an attacker could exploit the vulnerability by sending a specially crafted message to the RTU500, causing the targeted RTU500 CMU to reboot. The vulnerability is caused by a missing input data validation which eventually if exploited causes an internal buffer to overflow in the HCI IEC 60870-5-104 function.	Community	
□	×	[redacted]		? 21622		Vulnerability TIE:Score-Analysis='2.76'	9	1	[redacted]	2023-08-04	A vulnerability exists in HCI IEC 60870-5-104 function included in certain versions of the RTU500 series product. The vulnerability can only be exploited, if the HCI 60870-5-104 is configured with support for IEC 62351-3. After session resumption interval is expired an RTU500 initiated update of session parameters causes an unexpected restart due to a stack overflow.	Community	
☑	×			? 21620		Vulnerability TIE:Score-Analysis='2.67'	7			2023-08-04	External input could be used on TEL-STER TelWin SCADA WebInterface to construct paths to files and directories without properly neutralizing special elements within the pathname, which could allow an unauthenticated attacker to read files on the system.	Community	
☑	×	[redacted]		? 21619		Vulnerability TIE:Score-Analysis='2.67'	7	2	[redacted]	2023-08-03	All versions prior to 9.1.4 of Advantech WebAccess/SCADA are vulnerable to use of untrusted pointers. The RPC arguments the client sent client could contain raw memory pointers for the server to use as-is. This could allow an attacker to gain access to the remote file system and the ability to execute commands and overwrite files.	Community	
☑	×			? 21617		Vulnerability TIE:Score-Analysis='2.67'	7	2		2023-08-03	SpiderControl SCADA Webserver versions 2.08 and prior are vulnerable to path traversal. An attacker with administrative privileges could overwrite files on the webserver using the HMI's upload file feature. This could create size zero files anywhere on the webserver, potentially overwriting system files and creating a denial-of-service condition.	Community	

Figure 42. Selection of multiple events in MISP instance.

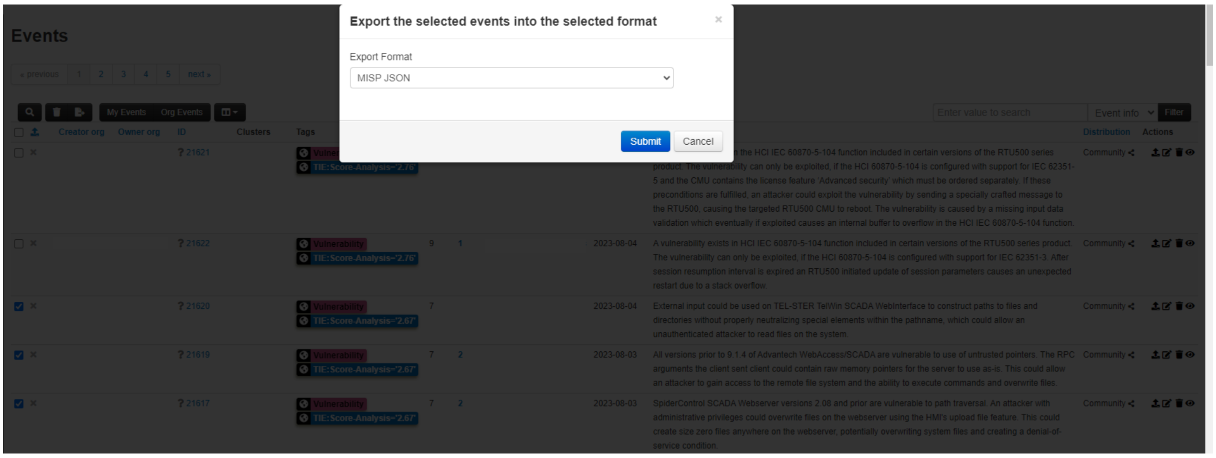

Export the selected events into the selected format ×

Export Format: MISP JSON

Figure 43. Download of events in MISP JSON format.

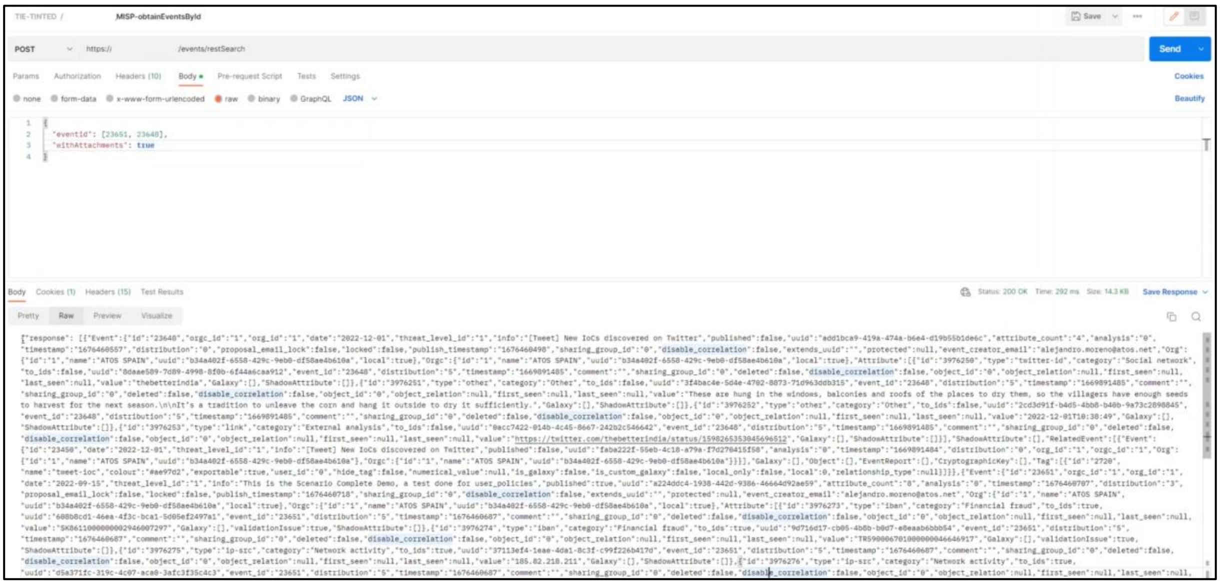

TIE-TINTED / MISP-obtainEventsById
Save
POST
https:// /events/restSearch
Send
Params
Authorization
Headers (10)
Body ●
Pre-request Script
Tests
Settings
Cookies
○ none ○ form-data ○ x-www-form-urlencoded ● raw ○ binary ○ GraphQL JSON
Beautify
Body
Cookies (1)
Headers (15)
Test Results
Status: 200 OK Time: 292 ms Size: 14.3 KB
Save Response
Pretty
Raw
Preview
Visualize

```
{
"eventid": [23651, 23648],
"withAttachments": true
}
```

```
{"response": [{"Event":{"id":"23648","orgc_id":"1","org_id":"1","date":"2022-12-01","threat_level_id":"1","info":"[Tweet] New IoCs discovered on Twitter","published":false,"uuid":"add1bca9-419a-474a-b6e4-d19b55b1de6c","attribute_count":"4","analysis":"0",
"timestamp":"1676460557","distribution":"0","proposal_email_lock":false,"locked":false,"publish_timestamp":"1676460498","sharing_group_id":"0","disable_correlation":false,"extends_uuid":"","protected":null,"event_creator_email":"alejandro.moreno@atos.net","Org":
{"id":"1","name":"ATOS SPAIN","uuid":"b34a402f-6558-429c-9eb0-df58ae4b610a","local":true},"Orgc":{"id":"1","name":"ATOS SPAIN","uuid":"b34a402f-6558-429c-9eb0-df58ae4b610a","local":true},"Attribute":[{"id":"3976250","type":"twitter-id","category":"Social network",
"to_ids":false,"uuid":"8daae589-7d89-4998-8f0b-6f44a6caa912","event_id":"23648","distribution":"5","timestamp":"1669891485","comment":"","sharing_group_id":"0","deleted":false,"disable_correlation":false,"object_id":"0","object_relation":null,"first_seen":null,
"last_seen":null,"value":"thebetterindia","Galaxy":[],"ShadowAttribute":[]},{"id":"3976251","type":"other","category":"Other","to_ids":false,"uuid":"3f4bac4e-5d4e-4702-8873-71d963ddb315","event_id":"23648","distribution":"5","timestamp":"1669891485","comment":"",
"sharing_group_id":"0","deleted":false,"disable_correlation":false,"object_id":"0","object_relation":null,"first_seen":null,"last_seen":null,"value":"These are hung in the windows, balconies and roofs of the places to dry them, so the villagers have enough seeds
to harvest for the next season.\n\nIt's a tradition to unleave the corn and hang it outside to dry it sufficiently.","Galaxy":[],"ShadowAttribute":[]},{"id":"3976252","type":"other","category":"Other","to_ids":false,"uuid":"2cd3d91f-b4d5-4bb8-b40b-9a73c2898845",
"event_id":"23648","distribution":"5","timestamp":"1669891485","comment":"","sharing_group_id":"0","deleted":false,"disable_correlation":false,"object_id":"0","object_relation":null,"first_seen":null,"last_seen":null,"value":"2022-12-01T10:38:49","Galaxy":[],
"ShadowAttribute":[]},{"id":"3976253","type":"link","category":"External analysis","to_ids":false,"uuid":"0acc7422-014b-4c45-8667-242b2c546642","event_id":"23648","distribution":"5","timestamp":"1669891485","comment":"","sharing_group_id":"0","deleted":false,
"disable_correlation":false,"object_id":"0","object_relation":null,"first_seen":null,"last_seen":null,"value":"https://twitter.com/thebetterindia/status/1598265353045696512","Galaxy":[],"ShadowAttribute":[]}],"ShadowAttribute":[],"RelatedEvent":[{"Event":
{"id":"23450","date":"2022-12-01","threat_level_id":"1","info":"[Tweet] New IoCs discovered on Twitter","published":false,"uuid":"faba222f-55eb-4c18-a79a-f7d270415f58","analysis":"0","timestamp":"1669891484","distribution":"0","org_id":"1","orgc_id":"1","Org":
{"id":"1","name":"ATOS SPAIN","uuid":"b34a402f-6558-429c-9eb0-df58ae4b610a"},"Orgc":{"id":"1","name":"ATOS SPAIN","uuid":"b34a402f-6558-429c-9eb0-df58ae4b610a"}}}],"Galaxy":[],"Object":[],"EventReport":[],"CryptographicKey":[],"Tag":[{"id":"2720",
"name":"tweet-ioc","colour":"#ae97d2","exportable":true,"user_id":"0","hide_tag":false,"numerical_value":null,"is_galaxy":false,"is_custom_galaxy":false,"local_only":false,"local":0,"relationship_type":null}]}},{"Event":{"id":"23651","orgc_id":"1","org_id":"1",
"date":"2022-09-15","threat_level_id":"1","info":"This is the Scenario Complete Demo, a test done for user_policies","published":true,"uuid":"a224ddc4-1938-442d-9386-46664d92ae59","attribute_count":"8","analysis":"0","timestamp":"1676460707","distribution":"3",
"proposal_email_lock":false,"locked":false,"publish_timestamp":"1676460718","sharing_group_id":"0","disable_correlation":false,"extends_uuid":"","protected":null,"event_creator_email":"alejandro.moreno@atos.net","Org":{"id":"1","name":"ATOS SPAIN",
"uuid":"b34a402f-6558-429c-9eb0-df58ae4b610a","local":true},"Orgc":{"id":"1","name":"ATOS SPAIN","uuid":"b34a402f-6558-429c-9eb0-df58ae4b610a","local":true},"Attribute":[{"id":"3976273","type":"iban","category":"Financial fraud","to_ids":true,
"uuid":"608b8cd1-46ea-4f3c-bca1-5d05ef2497a1","event_id":"23651","distribution":"5","timestamp":"1676460687","comment":"","sharing_group_id":"0","deleted":false,"disable_correlation":false,"object_id":"0","object_relation":null,"first_seen":null,"last_seen":null,
"value":"SK8611000000002946007297","Galaxy":[],"validationIssue":true,"ShadowAttribute":[]},{"id":"3976274","type":"iban","category":"Financial fraud","to_ids":true,"uuid":"9d716d17-cb05-4b6b-b0d7-e8eaab6bbb54","event_id":"23651","distribution":"5",
"timestamp":"1676460687","comment":"","sharing_group_id":"0","deleted":false,"disable_correlation":false,"object_id":"0","object_relation":null,"first_seen":null,"last_seen":null,"value":"TR590006701000000046646917","Galaxy":[],"validationIssue":true,
"ShadowAttribute":[]},{"id":"3976275","type":"ip-src","category":"Network activity","to_ids":true,"uuid":"37113ef4-1eae-4da1-8c3f-c99f226b417d","event_id":"23651","distribution":"5","timestamp":"1676460687","comment":"","sharing_group_id":"0","deleted":false,
"disable_correlation":false,"object_id":"0","object_relation":null,"first_seen":null,"last_seen":null,"value":"185.82.218.211","Galaxy":[],"ShadowAttribute":[]},{"id":"3976276","type":"ip-src","category":"Network activity","to_ids":true,
"uuid":"d5a371fc-319c-4c07-aca0-3afc3f35c4c3","event_id":"23651","distribution":"5","timestamp":"1676460687","comment":"","sharing_group_id":"0","deleted":false,"disable_correlation":false,"object_id":"0","object_relation":null,"first_seen":null,"last_seen":null,
```

Figure 44. HTTP request to download MISP events.

Once we got the desired events, we must insert them in the HTTP request that is sent to the orchestrator. The format of the request is shown in Figure 45.

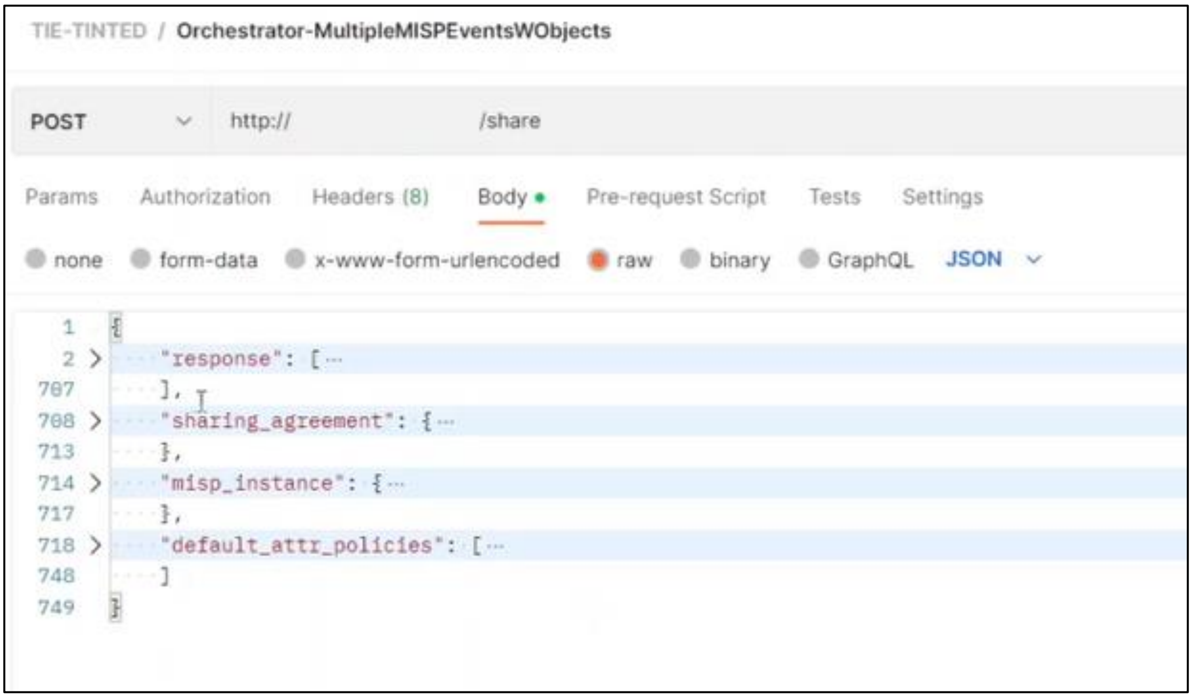

TIE-TINTED / Orchestrator-MultipleMISPEventsWObjects
POST
http:// /share
Params
Authorization
Headers (8)
Body ●
Pre-request Script
Tests
Settings
○ none ○ form-data ○ x-www-form-urlencoded ● raw ○ binary ○ GraphQL JSON

```
{
    "response": [⋯
    ],
    "sharing_agreement": {⋯
    },
    "misp_instance": {⋯
    },
    "default_attr_policies": [⋯
    ]
}
```

Figure 45. Orchestrator HTTP request format.

The response field is populated with the events that we have previously downloaded. Meanwhile the format of the *sharing_agreement* and *misp_instance* fields are shown in Figure 46 and Figure 47, respectively.

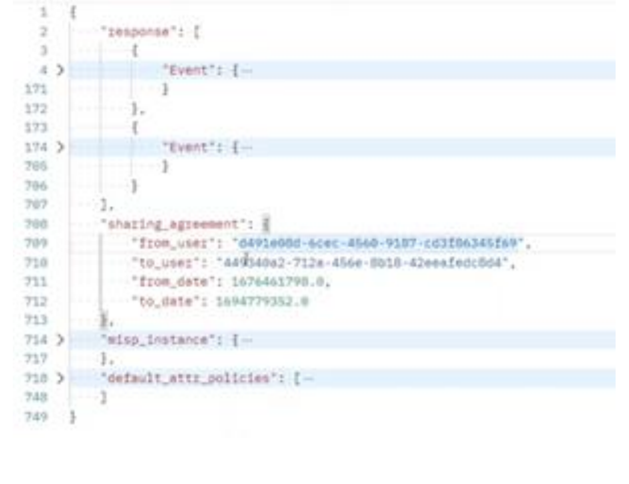

```
{
    "response": [
        {
            "Event": {…
            }
        },
        {
            "Event": {…
            }
        }
    ],
    "sharing_agreement": {
        "from_user": "d491e08d-6cec-4560-9187-cd3f06345f69",
        "to_user": "449340a2-712a-456e-8b18-42eeafedc8d4",
        "from_date": 1676461798.0,
        "to_date": 1694779352.0
    },
    "misp_instance": {…
    },
    "default_attr_policies": […
    ]
}
```

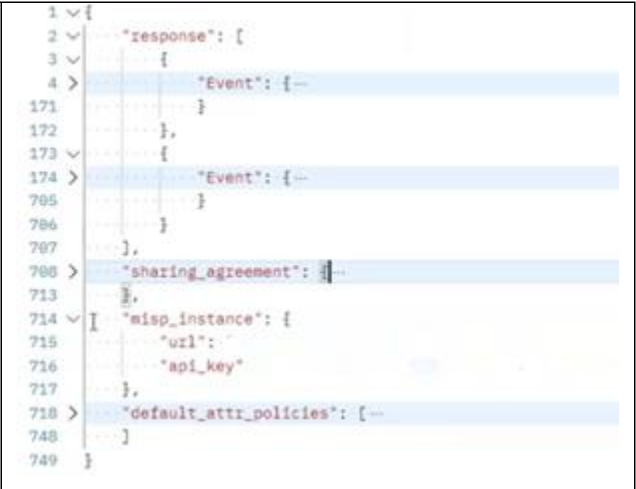

```
{
    "response": [
        {
            "Event": {…
            }
        },
        {
            "Event": {…
            }
        }
    ],
    "sharing_agreement": {…
    },
    "misp_instance": {
        "url": "
        "api_key"
    },
    "default_attr_policies": […
    ]
}
```

Figure 46. Sharing Agreement field.

Figure 47. MISP instance field.

Figure 48 provides an illustration of the user policies.

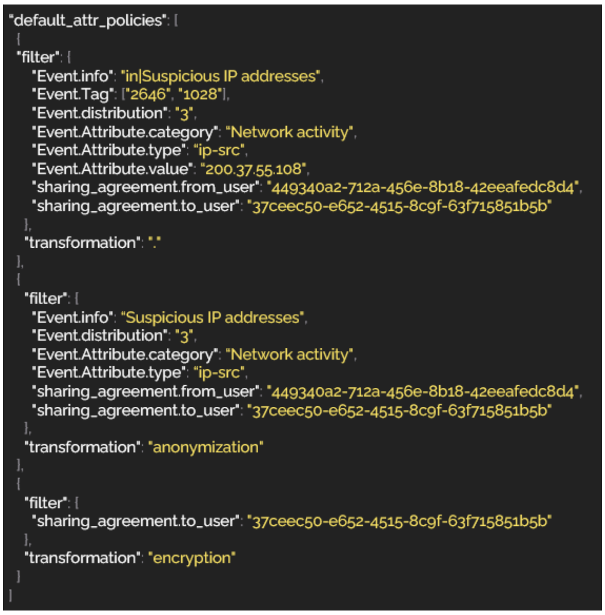

```
"default_attr_policies": [
 {
 "filter": {
   "Event.info": "in|Suspicious IP addresses",
   "Event.Tag": ["2646", "1028"],
   "Event.distribution": "3",
   "Event.Attribute.category": "Network activity",
   "Event.Attribute.type": "ip-src",
   "Event.Attribute.value": "200.37.55.108",
   "sharing_agreement.from_user": "449340a2-712a-456e-8b18-42eeafedc8d4",
   "sharing_agreement.to_user": "37ceec50-e652-4515-8c9f-63f715851b5b"
  },
  "transformation": "."
 },
 {
  "filter": {
   "Event.info": "Suspicious IP addresses",
   "Event.distribution": "3",
   "Event.Attribute.category": "Network activity",
   "Event.Attribute.type": "ip-src",
   "sharing_agreement.from_user": "449340a2-712a-456e-8b18-42eeafedc8d4",
   "sharing_agreement.to_user": "37ceec50-e652-4515-8c9f-63f715851b5b"
  },
  "transformation": "anonymization"
 },
 {
  "filter": {
   "sharing_agreement.to_user": "37ceec50-e652-4515-8c9f-63f715851b5b"
  },
  "transformation": "encryption"
 }
]
```

Figure 48. User policies for privacy sharing.

As observed, the "user_policies" parameter consists of a collection of distinct policies. Each policy comprises a combination of a filter and a transformation. The filter establishes a condition that must be met, while the transformation determines whether the attribute is to be encrypted, anonymized, or kept in its original form ("clear-text"). These policies are executed sequentially. When a policy's filter matches the attribute, the corresponding transformation is applied, and subsequent policies are no longer assessed. To fulfil a filter's requirements, all the specific conditions within it must align. The filter can encompass criteria related to the event, attributes, objects, or sharing agreement fields.

Within the first policy filter's "Event.info" field, the "|" operator is employed to signify an "in" condition type. This indicates that the event's "info" field should contain the substring "Suspicious IP addresses". Alternatively, using "=" followed by the "|" operator would indicate an equality condition, requiring the "info" field of the event to precisely match "Suspicious IP addresses". Additionally, it is significant that transformation values can be set as ".", aside from the options of encryption, anonymization, and "clear-text". This specific transformation denotes that the attribute remains unaltered. This feature accommodates instances when received events have already predetermined transformations, offering the capability to introduce exceptions to established rules. In the given example, it serves as an exception for the value "200.37.55.108" concerning the subsequent policy.

3.3 Scoring capabilities

One of the main components of TINTED is the Threat Intelligence Engine (TIE). It acquires events from the MISP instance and generates a threat score. This score can be divided into its different heuristics, that are on diverse metrics, like timeliness, trending, and completeness. Concurrently, we augment the received event's contextual information, a highly valuable enhancement. This entire procedure is referred to as CTI enrichment, as described in Figure 49.

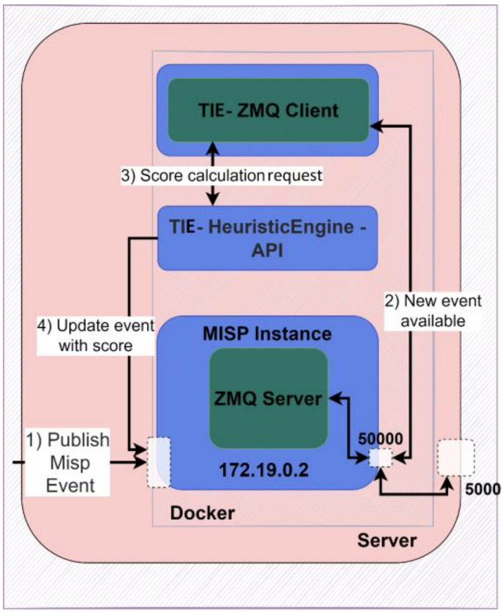

Figure 49. TIE architecture.

The central part of TIE is the HeuristicEngine. It processes API requests containing MISP Events and computes the score by considering the following factors:

- Static data: Information about the infrastructure.
- Dynamic data: Events, alerts, and vulnerability assessments.
- CTI (Cyber Threat Intelligence): The received event itself.

Subsequently, this score is integrated as an Attribute, updating the MISP Event within the MISP Instance.

The second element in TIE's architecture is the ZeroMQ client. A MISP Instance can be configured with a ZeroMQ Server, which the TIE system capitalizes on. The ZMQ Client is established as part of TIE and subscribes to the MISP Instance's ZMQ queue. When a fresh event arrives, the client sends a request to the HeuristicEngine component. This action triggers the execution of heuristic functions that ultimately lead to the score computation.

TIE does not encompass all MISP objects. Instead, it concentrates on four specific objects considered highly valuable in the realm of threat intelligence: **Vulnerability**, **Domain-IP**, **BTC-Address**, and **File**. These four objects comprise various attributes, including required and optional ones, which later contribute to the heuristics' calculations.

Figure 50 shows different MISP events that have arrived at the instance, and they have been processed for threat score calculation. In Figure 51, we can observe one individual event that contains a vulnerability object.

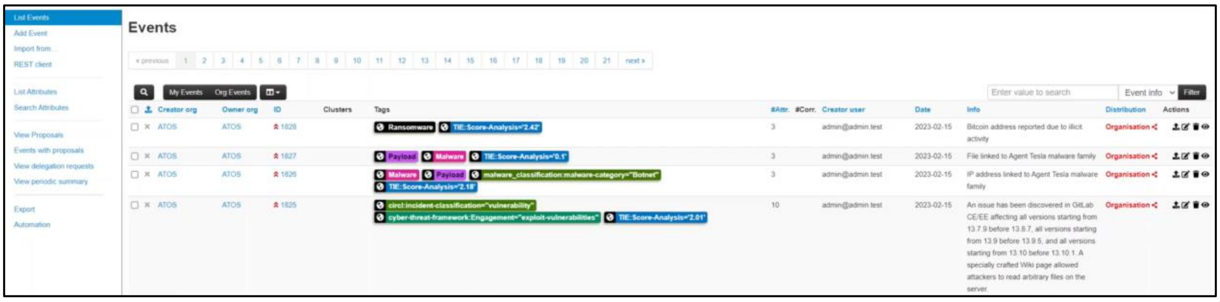

Creator org	Owner org	ID	Clusters	Tags	#Attr.	#Corr.	Creator user	Date	Info	Distribution	Actions
ATOS	ATOS	1828		Ransomware TIE:Score-Analysis="2.42"	3		admin@admin.test	2023-02-15	Bitcoin address reported due to illicit activity	Organisation	
ATOS	ATOS	1827		Payload Malware TIE:Score-Analysis="0.1"	3		admin@admin.test	2023-02-15	File linked to Agent Tesla malware family	Organisation	
ATOS	ATOS	1826		Malware Payload malware_classification:malware-category="Botnet" TIE:Score-Analysis="2.18"	3		admin@admin.test	2023-02-15	IP address linked to Agent Tesla malware family	Organisation	
ATOS	ATOS	1825		circl:incident-classification="vulnerability" cyber-threat-framework:Engagement="exploit-vulnerabilities" TIE:Score-Analysis="2.01"	10		admin@admin.test	2023-02-15	An issue has been discovered in GitLab CE/EE affecting all versions starting from 13.7.9 before 13.8.7, all versions starting from 13.9 before 13.9.5, and all versions starting from 13.10 before 13.10.1. A specially crafted Wiki page allowed attackers to read arbitrary files on the server.	Organisation	

Figure 50. MISP events enriched with TIE's threat score.

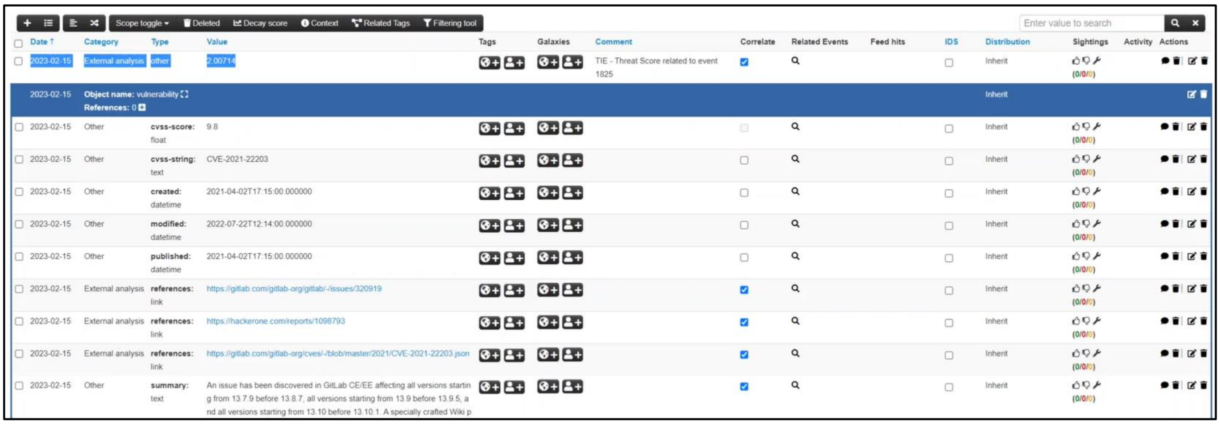

Date	Category	Type	Value	Tags	Galaxies	Comment	Correlate	Related Events	Feed hits	IDS	Distribution	Sightings	Activity	Actions
2023-02-15	External analysis	other	2.00714			TIE - Threat Score related to event 1825	☑			□	Inherit	(0/0/0)		
2023-02-15	Object name: vulnerability References: 0										Inherit			
2023-02-15	Other	cvss-score: float	9.8				□			□	Inherit	(0/0/0)		
2023-02-15	Other	cvss-string: text	CVE-2021-22203				□			□	Inherit	(0/0/0)		
2023-02-15	Other	created: datetime	2021-04-02T17:15:00.000000				□			□	Inherit	(0/0/0)		
2023-02-15	Other	modified: datetime	2022-07-22T12:14:00.000000				□			□	Inherit	(0/0/0)		
2023-02-15	Other	published: datetime	2021-04-02T17:15:00.000000				□			□	Inherit	(0/0/0)		
2023-02-15	External analysis	references: link	https://gitlab.com/gitlab-org/gitlab/-/issues/320919				☑			□	Inherit	(0/0/0)		
2023-02-15	External analysis	references: link	https://hackerone.com/reports/1098793				☑			□	Inherit	(0/0/0)		
2023-02-15	External analysis	references: link	https://gitlab.com/gitlab-org/cves/-/blob/master/2021/CVE-2021-22203.json				☑			□	Inherit	(0/0/0)		
2023-02-15	Other	summary: text	An issue has been discovered in GitLab CE/EE affecting all versions starting from 13.7.9 before 13.8.7, all versions starting from 13.9 before 13.9.5, and all versions starting from 13.10 before 13.10.1. A specially crafted Wiki p				☑			□	Inherit	(0/0/0)		

Figure 51. Vulnerability object.

3.4 Source Confidence Score

The new functionality "Source Confidence Calculation" included in TINTED has been implemented introducing several capabilities to enhance existing functionalities:

- Dynamic Adjustment of Source Confidence: This allows dynamic updates to the confidence level of intelligence sources, ensuring users have access to the most up-to-date data reliability assessment.
- NATO Admiralty Tag Assignment: This assigns a specific "admiralty" tag to each intelligence source, facilitating identification and categorization within the system.
- Scoring Model Integration/Adjustment: This integrates TINTED's existing scoring model with state-of-the-art methodology to provide a comprehensive and accurate evaluation of source reliability.

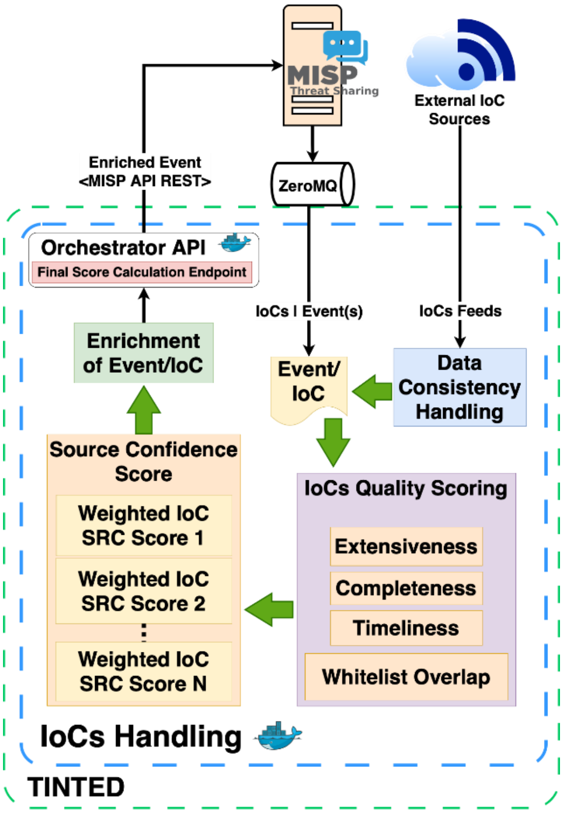

Figure 52: Source Confidence Score Architecture.

This source confidence score functionality uses data received directly through intelligence feeds configured in the MISP instance but also retrieves information with requests to the API provided by these intelligence feeds. The idea of using both is that although in some cases they will provide the same IoCs, sometimes more context information is provided when requested to the API. Consequently, it can operate in three modes, configured via the "FEEDING_MODE" environment variable:

- MISP: Intelligence feeds received from MISP through ZeroMQ Client.
- API: Intelligence feeds received by invoking APIs provided by the sources.
- ANY: Both modes of data collection are enabled.

To leverage this functionality, we must go through the already mentioned .env file described in Deployment section 2.3.5.

3.4.1 Data Collection

The list of intelligence feeds considered for source confidence calculation is the one defined in the file *intelligence_feeds.json*. Each source will be defined by a JSON structure as shown below:

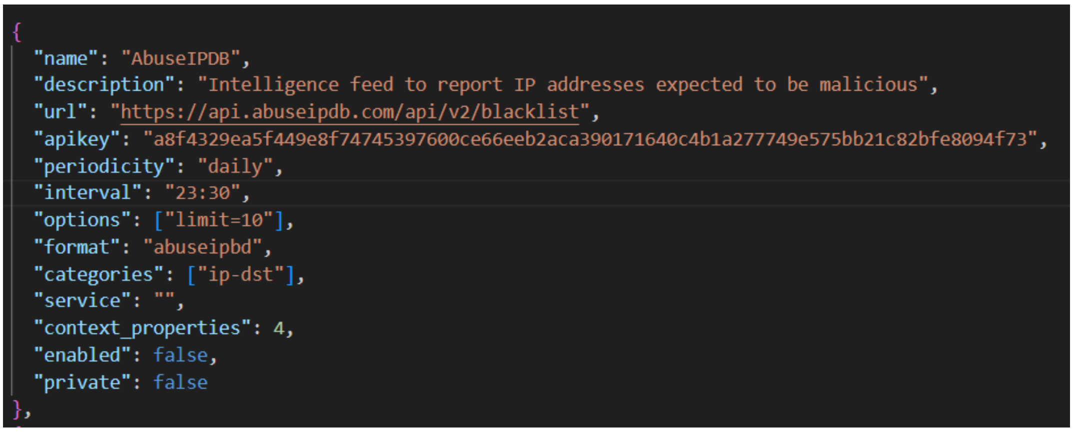

```
{
  "name": "AbuseIPDB",
  "description": "Intelligence feed to report IP addresses expected to be malicious",
  "url": "https://api.abuseipdb.com/api/v2/blacklist",
  "apikey": "a8f4329ea5f449e8f74745397600ce66eeb2aca390171640c4b1a277749e575bb21c82bfe8094f73",
  "periodicity": "daily",
  "interval": "23:30",
  "options": ["limit=10"],
  "format": "abuseipbd",
  "categories": ["ip-dst"],
  "service": "",
  "context_properties": 4,
  "enabled": false,
  "private": false
},
```

Figure 53: Intelligence feed JSON definition

- **Name:** This name will be used to identify an intelligence feed. If the intelligence feed is also included in the MISP instance enabled feeds, this name must match the one configured in the MISP instance and the one shown in the Event "Info" field. This means that it can be the same when configured in the MISP instance or part of it, in the case of multiple events generated by a same threat intelligence data source that have the same prefix.

  There should not be any problem as long as MISP administrators do not change the name given to the feeds. This way we will avoid any inconsistencies between MISP satellites and the central MISPEX instance.

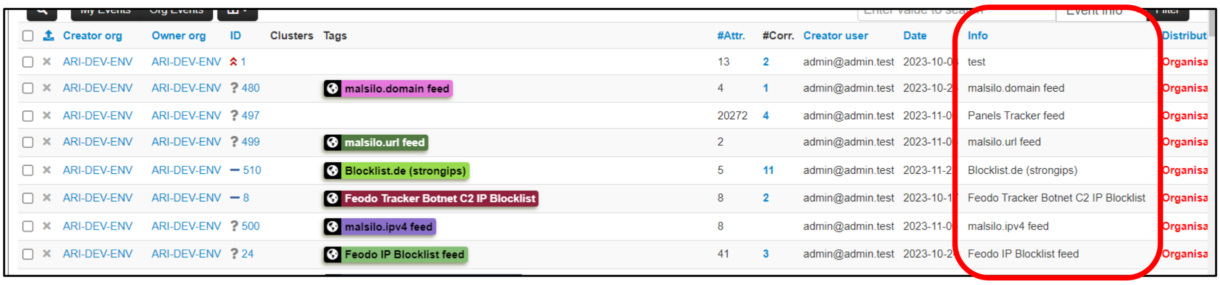

☐		Creator org	Owner org	ID	Clusters	Tags	#Attr.	#Corr.	Creator user	Date	Info	Distribut
☐	×	ARI-DEV-ENV	ARI-DEV-ENV	1			13	2	admin@admin.test	2023-10-0	test	Organisa
☐	×	ARI-DEV-ENV	ARI-DEV-ENV	? 480		malsilo.domain feed	4	1	admin@admin.test	2023-10-2	malsilo.domain feed	Organisa
☐	×	ARI-DEV-ENV	ARI-DEV-ENV	? 497			20272	4	admin@admin.test	2023-11-0	Panels Tracker feed	Organisa
☐	×	ARI-DEV-ENV	ARI-DEV-ENV	? 499		malsilo.url feed	2		admin@admin.test	2023-11-0	malsilo.url feed	Organisa
☐	×	ARI-DEV-ENV	ARI-DEV-ENV	— 510		Blocklist.de (strongips)	5	11	admin@admin.test	2023-11-2	Blocklist.de (strongips)	Organisa
☐	×	ARI-DEV-ENV	ARI-DEV-ENV	— 8		Feodo Tracker Botnet C2 IP Blocklist	8	2	admin@admin.test	2023-10-1	Feodo Tracker Botnet C2 IP Blocklist	Organisa
☐	×	ARI-DEV-ENV	ARI-DEV-ENV	? 500		malsilo.ipv4 feed	8		admin@admin.test	2023-11-0	malsilo.ipv4 feed	Organisa
☐	×	ARI-DEV-ENV	ARI-DEV-ENV	? 24		Feodo IP Blocklist feed	41	3	admin@admin.test	2023-10-2	Feodo IP Blocklist feed	Organisa

Figure 54: Intelligence Feed Name in MISP Dashboard.

- **Description:** A short description of the intelligence feed.
- **URL:** The intelligence feed endpoint that will be invoked to retrieve IoCs.
- **Apikey:** The API Key required to get IoCs from the source (if any). This field is optional.
- **Periodicity:** This indicates how often the intelligence feed URL will be invoked to get IoCs. It can have the following values: "daily", "interval_hours", "interval_minutes" or "interval_days".
- **Interval:** This value is used with the previous one (periodicity) for scheduling periodical retrieval of information.
- **Options:** This includes options to be added in the requests sent to the intelligence feed API. For example: ["limit=10"]. Different options can be added separated by commas.
- **Format:** This specifies the format of the data received from the intelligence feed using its API.
- **Categories:** This contains the type of IoCs provided by the intelligence feed. It is used for the generation of the MISP events or attributes when IoCs are retrieved from threat intelligence sources directly requesting their API. Current supported categories are: "ip-dst", "ip-dst|port", "ip-src", "domain", "domain-ip", "url", "md5", "email" and "filename".
- **Service:** This is an optional field that includes the name of the service related to the IoCs provided by the intelligence feed. For example, it is a tor node or IP address related to ssh brute force attacks. This information is not currently used in the source confidence calculation.
- **Context_properties:** This indicates the number of context properties provided by an intelligence feed. This value is used for extensiveness calculation.
- **Enabled:** This indicates if this intelligence feed will be considered or not in the source confidence calculation.
- **Private:** This indicates if this intelligence feed is a private or public source. Public sources will be published to MISP whereas private ones are not.

3.4.2 IoC Features Calculation

The heuristics involved in the source confidence score are described below:

- **Extensiveness Calculation**

  Extensiveness measures how much context an intelligence feed provides to complement additional information. With extensiveness, we assign a higher source confidence to intelligence feeds that give more context per IoC.

The timestamps provided by us during the processing of IoCs are not considered in the extensiveness calculation. Otherwise, we would be giving advantage to those feeds without timestamp since we are always including this piece of information.

The formula used is based on Equation 4 of the paper [4], where z_s is the total number of IoCs provided by the intelligence feeds, (oi)s is the number of contextual properties provided for IoC_i and n is maximum number of contextual properties.

$$SC_E(s) = \frac{1}{z_s} \sum_{i=0}^{z_s} \frac{(o_i)_s}{n}$$

In the current implementation, we assume that all IoCs provided by the same intelligence feed have the same context properties, as indicated in the configuration of the intelligence feed. This number of context properties is according to the information that can be retrieved when requesting the feed API, since it sometimes provides more context (e.g. country) that is not provided when the IoC is received through a MISP feed.

The context properties currently considered in the extensiveness calculation are:

- Timestamp of IoC publication provided by the feed.
- Timestamp of the last sighting of the IoC.
- Timestamp of the first sighting of the IoC.
- Country.
- Service associated to the IP address.
- Description of threats/malware related to the IoCs.
- Hash of the malware related to a domain or IP address IoC.
- Status of a service with an IP address.
- Number of sightings per IoC.

Timeliness Calculation

Timeliness defines how fast an intelligence feed shares its IoCs compared to other feeds. If a certain intelligence feed shares the same IoCs later than others, the IoCs could be outdated. Considering the slowness in sharing its IoCs, we will assign less source confidence to it.

The formula used is based on Equation 7 of the paper [4][19]:

$$SC_T(s) = \frac{1}{z_s} \sum_{i=0}^{z_s} \frac{min(t_i) - ((t_s)_i - \lambda)}{(t_s)_i - ((t_s)_i - \lambda)} = \frac{1}{z_s} \sum_{i=0}^{z_s} \frac{min(t_i) - (t_s)_i + \lambda}{\lambda}$$

To measure timeliness, we consider the timestamp field and not the first-seen field as sometimes *feed X* sighted an IoC before *feed Y* but the first one that decided to share it was *feed Y*.

Completeness Calculation

Completeness defines how much an intelligence feed contribute to the total collection of IoCs. Please note that this score focuses more on the quantity than on the quality of an intelligence feed, assuming that if it has many IoCs, the feed is more useful.

The formula used is based on Equation 9 of the paper [4]:

$$SC_C(s) = \frac{z_s}{z_{total}}$$

Whitelist Overlap Score Calculation

Whitelist Overlap Score defines how much of the IoCs in an intelligence feed also exist in a trusted whitelist. The formula used is based on Equation 10 of the paper[4], where it has been used ρ=0.1

(assuming that an intelligence feed is trustworthy if 10% of its IoCs are whitelisted) and δ=0.5 (assuming that a small overlap between an intelligence feed and a whitelist should not have a large influence on the whitelist overlap score):

$$SC_W(s) = max(0,\ 1 - (\frac{u_s}{z_s \cdot \rho})^{\frac{1}{\delta}})$$

Currently, this heuristic is calculated using the Cisco Umbrella domain whitelist and considers as whitelist the IP ranges used from cloud providers. These whitelists are downloaded from the URLS configured in the environment variables, WHITELIST_IP_URL and WHITELIST_DOMAIN_URL and stored as csv files to be used in the heuristics. These whitelists are downloaded initially when the service is deployed and then periodically every day at 07:00h. This time can be configured with the environment variable WHITELIST_DOWNLOAD_TIME. The CSV files with the downloaded whitelists are maintained locally, with the number of days configured in the environment variable WHITELIST_DAYS_TO_KEEP (by default, 2 days).

This heuristic requires you to download whitelists, which is not possible in the case of no external access from the server where the service is running. For that reason, the environment variable ENABLE_WHITELIST has been added to enable the calculation of this score (by default, it is disabled). If not enable, this heuristic is not considered in the calculation of the source confidence score.

Source Confidence

The four heuristics presented in the previous section are used in a weighted mean to calculate the source confidence of the gathered intelligence feeds. This is the value that will appear in the tag "***TINTED:Source-Score***" of the MISP event associated with a feed.

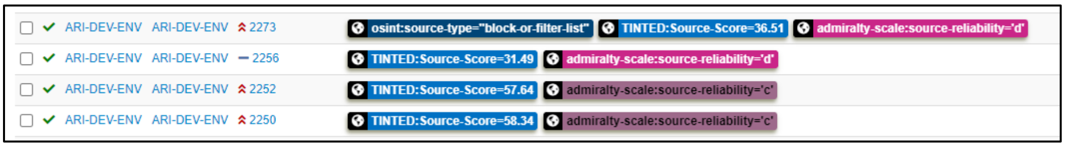

□ ✓	ARI-DEV-ENV	ARI-DEV-ENV	2273	osint:source-type="block-or-filter-list" TINTED:Source-Score=36.51	admiralty-scale:source-reliability='d'
□ ✓	ARI-DEV-ENV	ARI-DEV-ENV	2256	TINTED:Source-Score=31.49	admiralty-scale:source-reliability='d'
□ ✓	ARI-DEV-ENV	ARI-DEV-ENV	2252	TINTED:Source-Score=57.64	admiralty-scale:source-reliability='c'
□ ✓	ARI-DEV-ENV	ARI-DEV-ENV	2250	TINTED:Source-Score=58.34	admiralty-scale:source-reliability='c'

Figure 55: TINTED Source Confidence Tags.

The formula used is based on Equation 11 of the paper [4]:

$$SC(s) = \frac{W_E \cdot SC_E(s) + W_T \cdot SC_T(s) + W_C \cdot SC_C(s) + W_W \cdot SC_W(s)}{W_E + W_T + W_C + W_W}$$

Currently, these weights all have a value of 1 which means that all the heuristics have the same relevance for the source confidence. However, some of the heuristics such as completeness could be less useful than others such as the whitelist overlap score to calculate the source confidence. These weights can be configured through the environment variables WE (weight for extensiveness feature), WT (weight for timeliness feature), WC (weight for completeness feature) and WW (weight for whitelist overlap feature). Some experiments will be performed during Spring 5 to determine the influence of the different features and the weights to be configured for the source confidence calculation.

Each MISP event associated with an intelligence feed is also tagged using the Admiralty taxonomy ("***admiralty-scale:source-reliability***"), a structured classification system designed to enrich event data. This assigned admiralty taxonomy plays a crucial role in calculating the decay of Indicators of Compromise (IoCs). Whenever an IoC is analysed, its decay is influenced by the Admiralty taxonomy value of the associated event. This ensures that the significance of an IoC is dynamically adjusted based on the taxonomy classification. The integration of Admiralty taxonomy into the IoC decay calculation process allows for a more nuanced and context-aware approach to threat intelligence analysis.

The mapping between the TINTED source scoring and the admiralty tag is done according to the following table:

Table 1: TINTED score mapping with Admiralty tag

TINTED Score	Admiralty Tag	Numeric Value
>= 87,5	admiralty-scale:source-reliability="a"	100
>= 62,5	admiralty-scale:source-reliability=”b"	75
>= 37,5	admiralty-scale:source-reliability=”c"	50
< 12,5	admiralty-scale:source-reliability=”d"	25

- **IoC Scoring**

For each IoC in a specific intelligence feed, a separate IoC score is also calculated following the same formula used for the Source Confidence calculation, but only with those heuristics that apply to individual IoCs; these are extensiveness and timeliness. This value will appear in the tag “***TINTED:IoC-Score***" of the MISP attribute with the IoC for the specific MISP event associated with the feed.

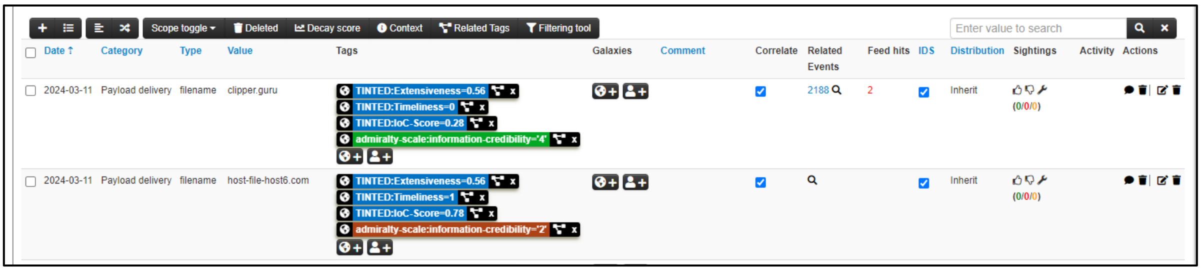

□	Date ↑	Category	Type	Value	Tags	Galaxies	Comment	Correlate	Related Events	Feed hits	IDS	Distribution	Sightings	Activity	Actions
□	2024-03-11	Payload delivery	filename	clipper.guru	TINTED:Extensiveness=0.56 x TINTED:Timeliness=0 x TINTED:IoC-Score=0.28 x admiralty-scale:information-credibility="4" x			☑	2188	2	☑	Inherit	(0/0/0)		
□	2024-03-11	Payload delivery	filename	host-file-host6.com	TINTED:Extensiveness=0.56 x TINTED:Timeliness=1 x TINTED:IoC-Score=0.78 x admiralty-scale:information-credibility="2" x			☑			☑	Inherit	(0/0/0)		

Figure 56: TINTED IoC Scoring Tags.

- **Final Score**

The final score is a combination of the scores for all separate intelligence feeds and the confidence scores of these feeds.

To avoid every feed having the same amount of influence, we want feeds with a higher source confidence to have more influence than feeds with a lower source confidence. To achieve this, the final score needs to be a weighted mean, where the source confidence works as a weight on the final score per feed, as shown in Equation 15 of the paper [4][19]:

$$final\_score = \frac{\sum_{i=0}^{N} source\_confidence_i^2 \cdot score_i}{\sum_{i=0}^{N} source\_confidence_i}$$

The final score has been developed as an endpoint in the TINTED Orchestrator working on demand. This way, the user can retrieve the final score for a specific IoC.

You must send the following request:

Method: POST

URL:
http://IP_ADDRESS_WHERE_ORCH_IS_DEPLOYED:PORT/orch/get_average_ioc_score

Payload:
{
ioc: XXX.XXX.XXX.XXX (IP addresses, domains, hashes and emails)
}

3.5 IoC enrichment with MITRE ATT&CK techniques

Another important feature that has been integrated in TINTED is the ability to enrich IoC with techniques from the MITRE ATT&CK matrix.

The feature uses the same logic as the scoring. For each new event that is received in the MISP instance, the HeuristicEngine, leveraging VirusTotal API endpoints, introduces the techniques associated to those attributes.

This way, whenever CI operators pull limited information from intelligence feeds (Figure 57), they can get more insights about the IoC they received, as shown in Figure 58.

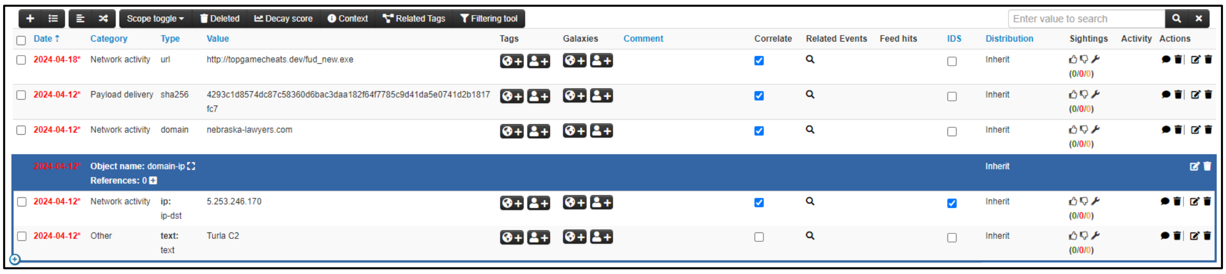

☐	Date ↑	Category	Type	Value	Tags	Galaxies	Comment	Correlate	Related Events	Feed hits	IDS	Distribution	Sightings	Activity	Actions
☐	2024-04-18*	Network activity	url	http://topgamecheats.dev/fud_new.exe				☑			☐	Inherit	(0/0/0)		
☐	2024-04-12*	Payload delivery	sha256	4293c1d8574dc87c58360d6bac3daa182f64f7785c9d41da5e0741d2b1817fc7				☑			☐	Inherit	(0/0/0)		
☐	2024-04-12*	Network activity	domain	nebraska-lawyers.com				☑			☐	Inherit	(0/0/0)		
	2024-04-12*	Object name: domain-ip References: 0										Inherit			
☐	2024-04-12*	Network activity	ip: ip-dst	5.253.246.170				☑			☑	Inherit	(0/0/0)		
☐	2024-04-12*	Other	text: text	Turla C2				☐			☐	Inherit	(0/0/0)		

Figure 57. IoCs before enrichment with MITRE ATT&CK techniques.

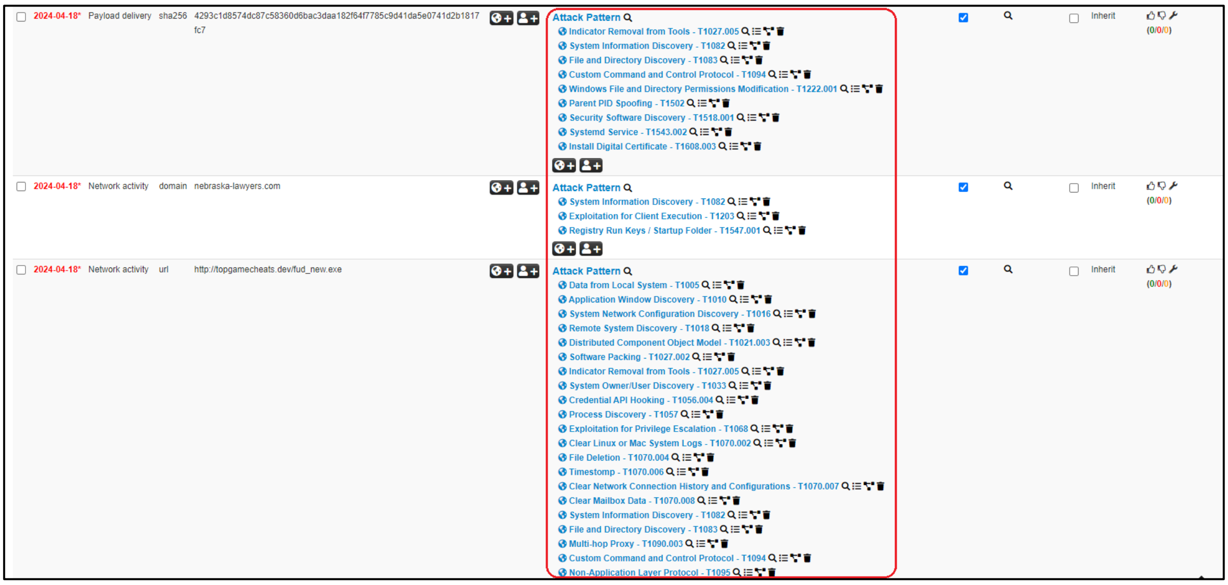

☐	2024-04-18*	Payload delivery	sha256	4293c1d8574dc87c58360d6bac3daa182f64f7785c9d41da5e0741d2b1817fc7	Attack Pattern Indicator Removal from Tools - T1027.005 System Information Discovery - T1082 File and Directory Discovery - T1083 Custom Command and Control Protocol - T1094 Windows File and Directory Permissions Modification - T1222.001 Parent PID Spoofing - T1502 Security Software Discovery - T1518.001 Systemd Service - T1543.002 Install Digital Certificate - T1608.003	☑	☐	Inherit	(0/0/0)
☐	2024-04-18*	Network activity	domain	nebraska-lawyers.com	Attack Pattern System Information Discovery - T1082 Exploitation for Client Execution - T1203 Registry Run Keys / Startup Folder - T1547.001	☑	☐	Inherit	(0/0/0)
☐	2024-04-18*	Network activity	url	http://topgamecheats.dev/fud_new.exe	Attack Pattern Data from Local System - T1005 Application Window Discovery - T1010 System Network Configuration Discovery - T1016 Remote System Discovery - T1018 Distributed Component Object Model - T1021.003 Software Packing - T1027.002 Indicator Removal from Tools - T1027.005 System Owner/User Discovery - T1033 Credential API Hooking - T1056.004 Process Discovery - T1057 Exploitation for Privilege Escalation - T1068 Clear Linux or Mac System Logs - T1070.002 File Deletion - T1070.004 Timestomp - T1070.006 Clear Network Connection History and Configurations - T1070.007 Clear Mailbox Data - T1070.008 System Information Discovery - T1082 File and Directory Discovery - T1083 Multi-hop Proxy - T1090.003 Custom Command and Control Protocol - T1094 Non-Application Layer Protocol - T1095	☑	☐	Inherit	(0/0/0)

Figure 58. IoCs after enrichment with MITRE ATT&CK techniques.

To enable this functionality, the user must set the environment variable "ENABLE_MITRE_ENRICHMENT" in the .env file to True.

3.6 Contextualization of health trends for threat score

One of the SUNRISE's key strategy points, developed under WP2, is workforce absenteeism.

This feature tries to leverage the nature of the threat intelligence module and monitors cyber and *physical* events that may impact CI. Figure 59 shows a graph of searches for some keywords (influenza, fever, and sick leave) that peak around Christmas time (24th-25th December).

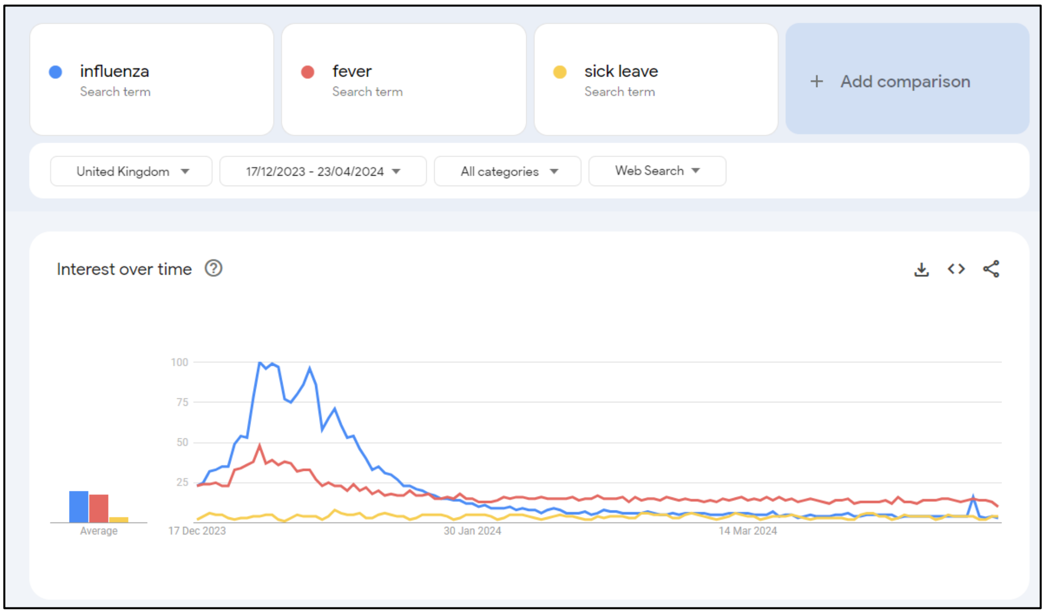

influenza
Search term

fever
Search term

sick leave
Search term

+ Add comparison

United Kingdom ▾ | 17/12/2023 - 23/04/2024 ▾ | All categories ▾ | Web Search ▾

Interest over time

Figure 59. Interest over time for health-related keywords in Google Trends.

This is a good example that shows how special events during the year, especially in a pandemic scenario, can produce an increase in infections. If we can detect anomalies in the number of searches for health topics, we may be able to predict whether our workforce will be unavailable to attend to their workplace due to illness.

The feature can be enabled in the .env file under the variable ENABLE_GOOGLE_TRENDS_MONITORING and requires a minimum set of configurations:

```
[search_query]

physical_events_keywords = influenza, fever, sick leave

[location]

# https://serpapi.com/google-trends-locations

geo = UK
```

First, we need a list of keywords that we want to monitor and then the country where we want to check the number of searches. It is worth noting that the keywords should be written in the same language as the selected country to avoid bias as most of the population search for content in their mother-tongue language. By default, this information is in a file called google_trends_parameters.txt.

By default, the timespan of the search covers the last 30 days, as this is enough to retrieve a sample of the trend. Information will be retrieved daily at the time indicated in the env variable HEALTH_TRENDS_SYNC_TIME. Information retrieved will be stored for its processing in the folder configured in the env variable HEALTH_TRENDS_DOWNLOAD_FOLDER.

If the script detects a high increase in the graph, the Heuristic Engine module will send an alarm to the risk assessment module so that indicators are triggered with their respective risk models.

4 Risk Assessment

4.1 Login and initial configurations

To access the risk assessment module, the user can click on the tab “Risk Assessment” of the CPR Dashboard (see Figure) to directly loge in or, if not integrated with single sign-on through a Keycloak server, the first step is to enter valid credentials to the login form, as shown in Figure 60.

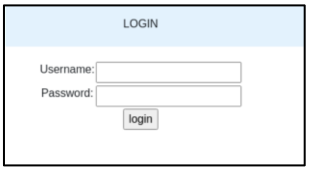

LOGIN

Username:

Password:

login

Figure 60. Login form.

Once logged in, the user will see a dashboard (Figure 61) with their personal information and a button to update their profile.

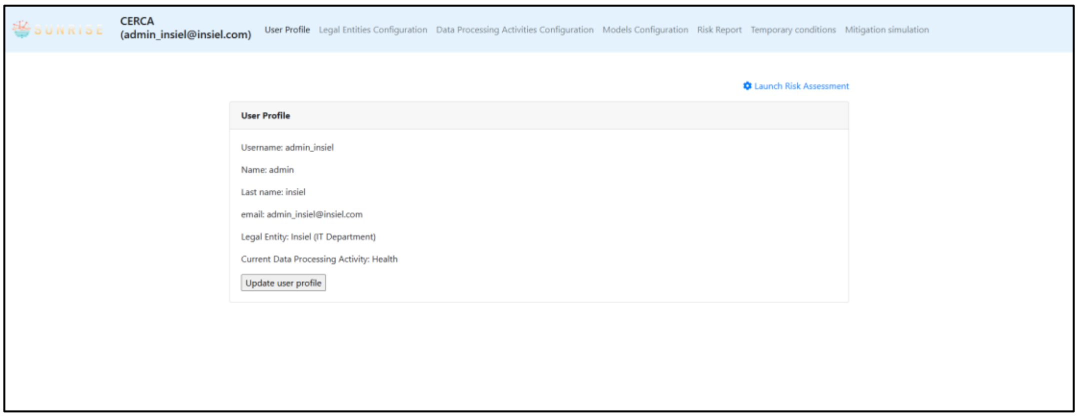

CERCA
(admin_insiel@insiel.com)

User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report Temporary conditions Mitigation simulation

Launch Risk Assessment

User Profile

Username: admin_insiel

Name: admin

Last name: insiel

email: admin_insiel@insiel.com

Legal Entity: Insiel (IT Department)

Current Data Processing Activity: Health

Update user profile

Figure 61. User profile menu.

Following screenshot (see Figure 63) shows the form to select your active processing activity that appears when you click in “Update user profile”.

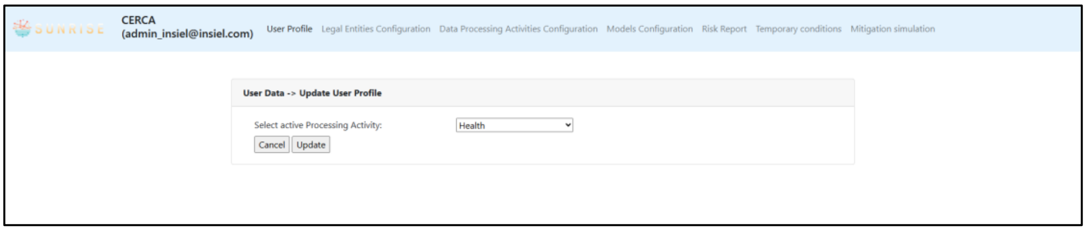

CERCA
(admin_insiel@insiel.com)

User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report Temporary conditions Mitigation simulation

User Data -> Update User Profile

Select active Processing Activity: Health

Cancel Update

Figure 62. Select active data processing activity

The following screenshot (see Figure 63) shows further information about the user legal entity. Regarding Figure 64, we can also analyse the different data processing activities concerning the user legal entity.

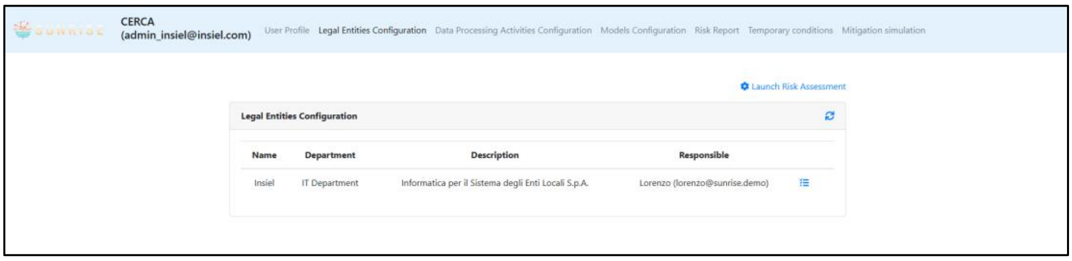

CERCA
(admin_insiel@insiel.com)

User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report Temporary conditions Mitigation simulation

Launch Risk Assessment

Legal Entities Configuration

Name	Department	Description	Responsible	
Insiel	IT Department	Informatica per il Sistema degli Enti Locali S.p.A.	Lorenzo (lorenzo@sunrise.demo)	☰

Figure 63. Legal entities configuration menu.

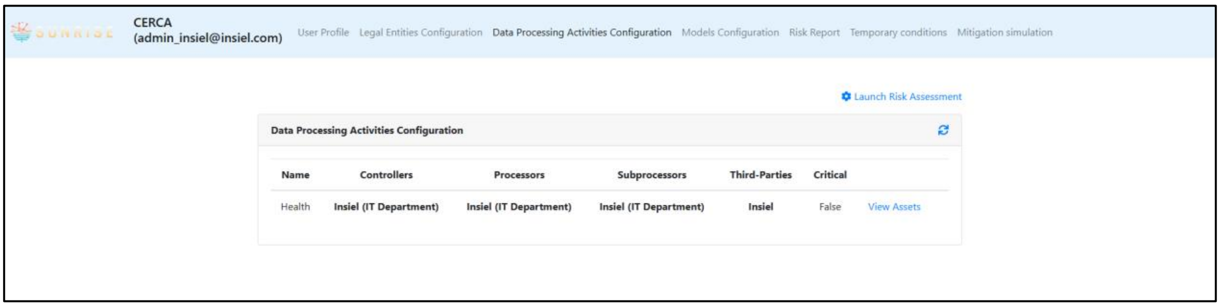

Data Processing Activities Configuration

Name	Controllers	Processors	Subprocessors	Third-Parties	Critical	
Health	Insiel (IT Department)	Insiel (IT Department)	Insiel (IT Department)	Insiel	False	View Assets

Figure 64. Data processing activities configuration

Figure 65 shows the menu where all the assets associated to a specific data processing are displayed following the CIA triad (confidentiality, integrity, and availability).

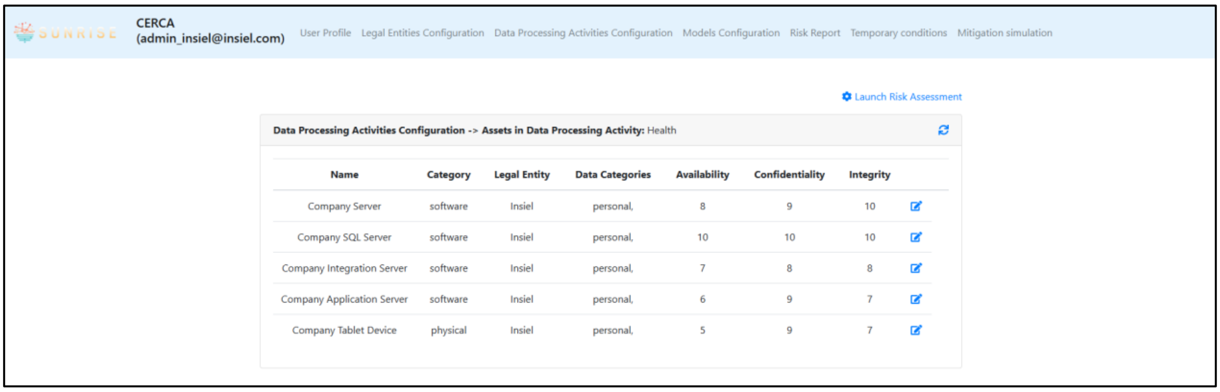

Data Processing Activities Configuration -> Assets in Data Processing Activity: Health

Name	Category	Legal Entity	Data Categories	Availability	Confidentiality	Integrity
Company Server	software	Insiel	personal,	8	9	10
Company SQL Server	software	Insiel	personal,	10	10	10
Company Integration Server	software	Insiel	personal,	7	8	8
Company Application Server	software	Insiel	personal,	6	9	7
Company Tablet Device	physical	Insiel	personal,	5	9	7

Figure 65. Asset view dashboard.

If you click to edit an asset, a new form will appear as shown in Figure 66. The user can characterise the criticality of the asset (in terms of impact on confidentiality, integrity and availability) if the asset is compromised, indicate the personal data processed by that asset and the loss values (in euros) for a typical loss scenario and for the worst scenario. All this information is required for a correct qualitative and quantitative risk assessment.

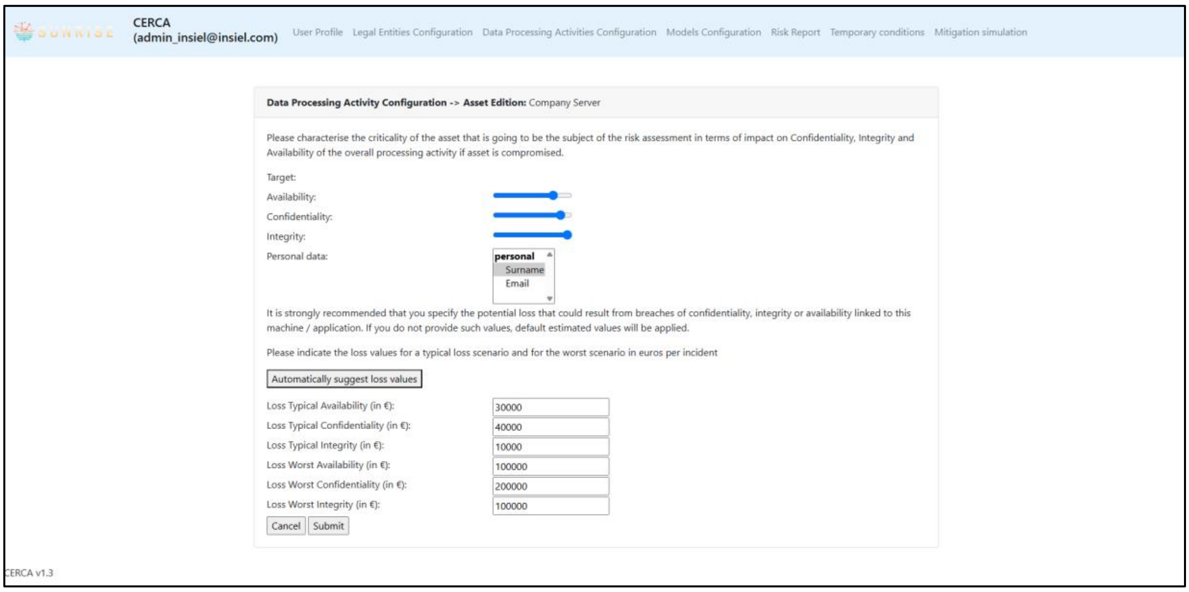
CERCA
(admin_insiel@insiel.com)
User Profile
Legal Entities Configuration
Data Processing Activities Configuration
Models Configuration
Risk Report
Temporary conditions
Mitigation simulation
Automatically suggest loss values
Cancel
Submit
CERCA v1.3

Data Processing Activity Configuration -> Asset Edition: Company Server

Please characterise the criticality of the asset that is going to be the subject of the risk assessment in terms of impact on Confidentiality, Integrity and Availability of the overall processing activity if asset is compromised.

Target:

Availability:

Confidentiality:

Integrity:

Personal data: **personal** Surname Email

It is strongly recommended that you specify the potential loss that could result from breaches of confidentiality, integrity or availability linked to this machine / application. If you do not provide such values, default estimated values will be applied.

Please indicate the loss values for a typical loss scenario and for the worst scenario in euros per incident

Loss Typical Availability (in €):	30000
Loss Typical Confidentiality (in €):	40000
Loss Typical Integrity (in €):	10000
Loss Worst Availability (in €):	100000
Loss Worst Confidentiality (in €):	200000
Loss Worst Integrity (in €):	100000

Figure 66. Asset edition dashboard.

Figure 67 shows a list of threats, risks, and security measures. We can also select a risk model or clear the current selection.

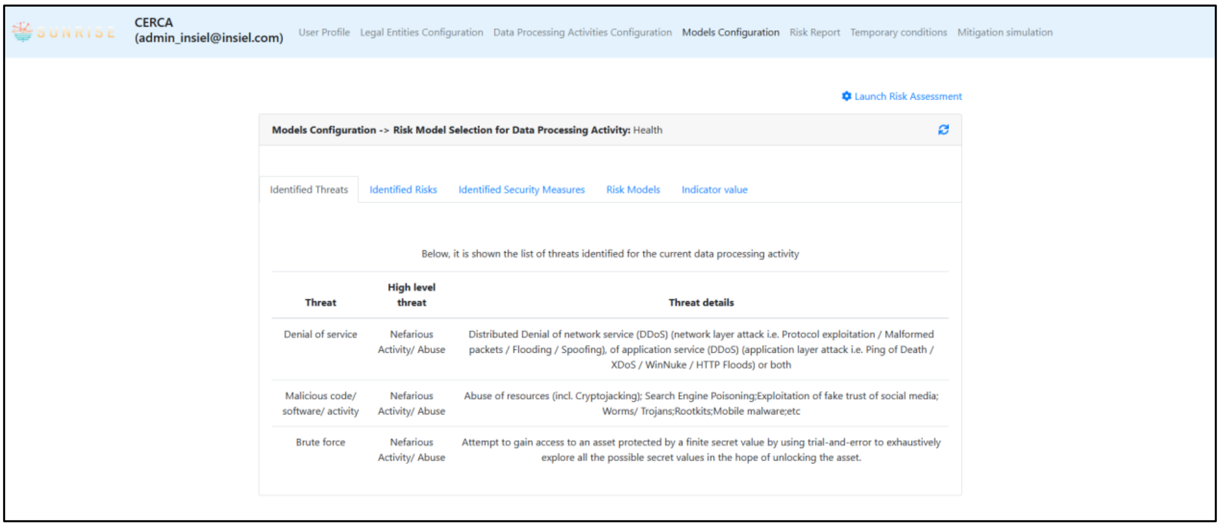

Models Configuration -> Risk Model Selection for Data Processing Activity: Health

Below, it is shown the list of threats identified for the current data processing activity

Threat	High level threat	Threat details
Denial of service	Nefarious Activity/ Abuse	Distributed Denial of network service (DDoS) (network layer attack i.e. Protocol exploitation / Malformed packets / Flooding / Spoofing), of application service (DDoS) (application layer attack i.e. Ping of Death / XDoS / WinNuke / HTTP Floods) or both
Malicious code/ software/ activity	Nefarious Activity/ Abuse	Abuse of resources (incl. Cryptojacking); Search Engine Poisoning;Exploitation of fake trust of social media; Worms/ Trojans;Rootkits;Mobile malware;etc
Brute force	Nefarious Activity/ Abuse	Attempt to gain access to an asset protected by a finite secret value by using trial-and-error to exhaustively explore all the possible secret values in the hope of unlocking the asset.

Figure 67. Model configuration dashboard.

One of the most important inputs that CERCA needs to process the cyber risk calculation is the questionnaire. Figure 68 shows an example of this. This questionnaire is available from the “Legal Entities Configuration” menu.

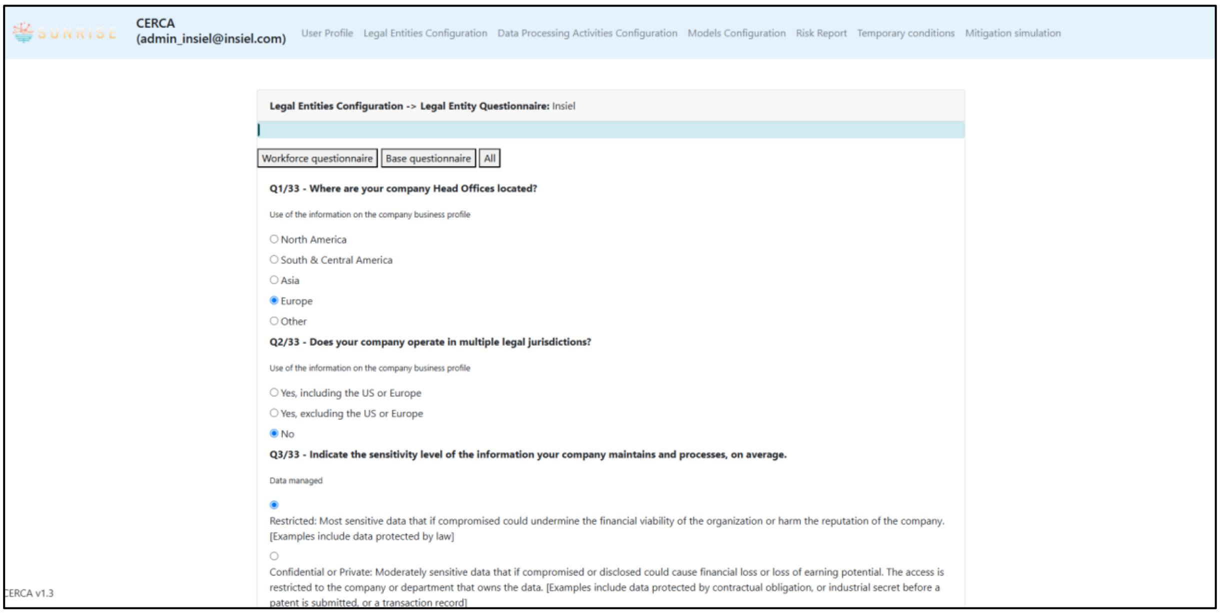

Legal Entities Configuration -> Legal Entity Questionnaire: Insiel

Q1/33 - Where are your company Head Offices located?

Use of the information on the company business profile

○ North America
○ South & Central America
○ Asia
● Europe
○ Other

Q2/33 - Does your company operate in multiple legal jurisdictions?

Use of the information on the company business profile

○ Yes, including the US or Europe
○ Yes, excluding the US or Europe
● No

Q3/33 - Indicate the sensitivity level of the information your company maintains and processes, on average.

Data managed

●
Restricted: Most sensitive data that if compromised could undermine the financial viability of the organization or harm the reputation of the company. [Examples include data protected by law]

○
Confidential or Private: Moderately sensitive data that if compromised or disclosed could cause financial loss or loss of earning potential. The access is restricted to the company or department that owns the data. [Examples include data protected by contractual obligation, or industrial secret before a patent is submitted, or a transaction record]

Figure 68. Questionnaire for risk model.

Legal entities, data processing activities, assets and risk models shown the dashboard to the user must have been previously registered in the system by the administrator user. Once logged to the CERCA dashboard with a user with administration permissions, a dashboard with these configuration options will appear, as shown in Figure 69.

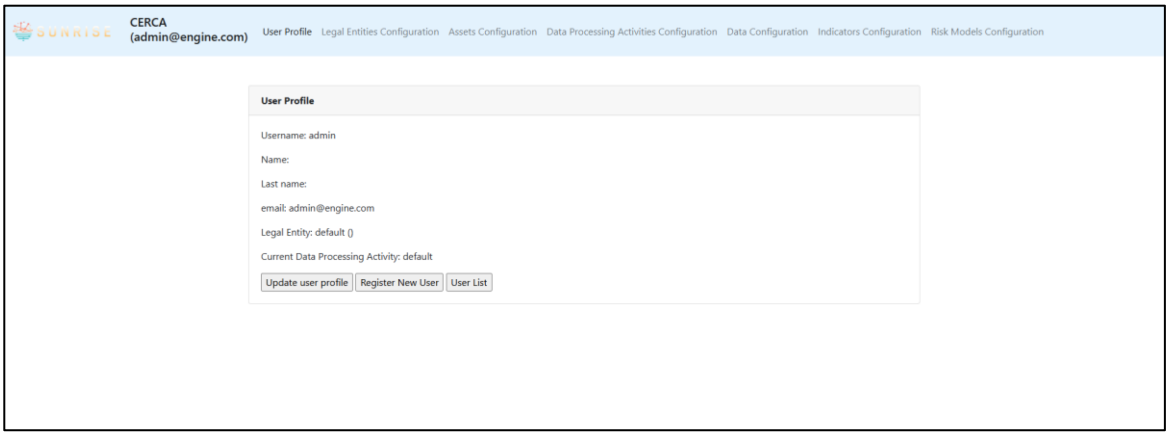

User Profile

Username: admin

Name:

Last name:

email: admin@engine.com

Legal Entity: default ()

Current Data Processing Activity: default

Figure 69. Risk Assessment Administrator dashboard.

If you click on Data Processing Activities Configuration tab, it will be shown all data processing activities registered in the system for all the legal entities configured (see Figure 70).

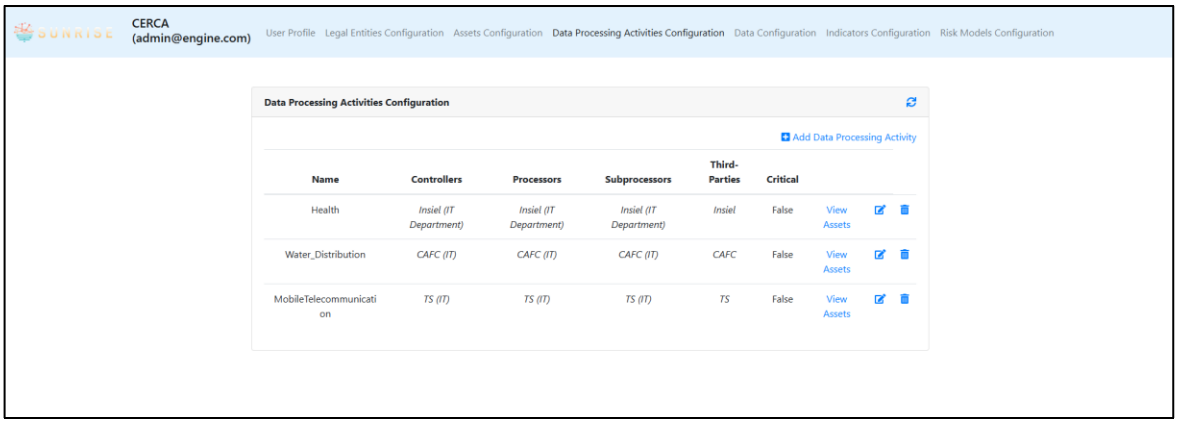
SUNRISE
CERCA (admin@engine.com)
User Profile
Legal Entities Configuration
Assets Configuration
Data Processing Activities Configuration
Data Configuration
Indicators Configuration
Risk Models Configuration
Add Data Processing Activity

Data Processing Activities Configuration

Name	Controllers	Processors	Subprocessors	Third-Parties	Critical			
Health	*Insiel (IT Department)*	*Insiel (IT Department)*	*Insiel (IT Department)*	*Insiel*	False	View Assets		
Water_Distribution	*CAFC (IT)*	*CAFC (IT)*	*CAFC (IT)*	*CAFC*	False	View Assets		
MobileTelecommunication	*TS (IT)*	*TS (IT)*	*TS (IT)*	*TS*	False	View Assets		

Figure 70. Data Processing Activities Configuration dashboard.

Information about assets is configured after clicking “Assets Configuration” tab as shown in Figure 71.

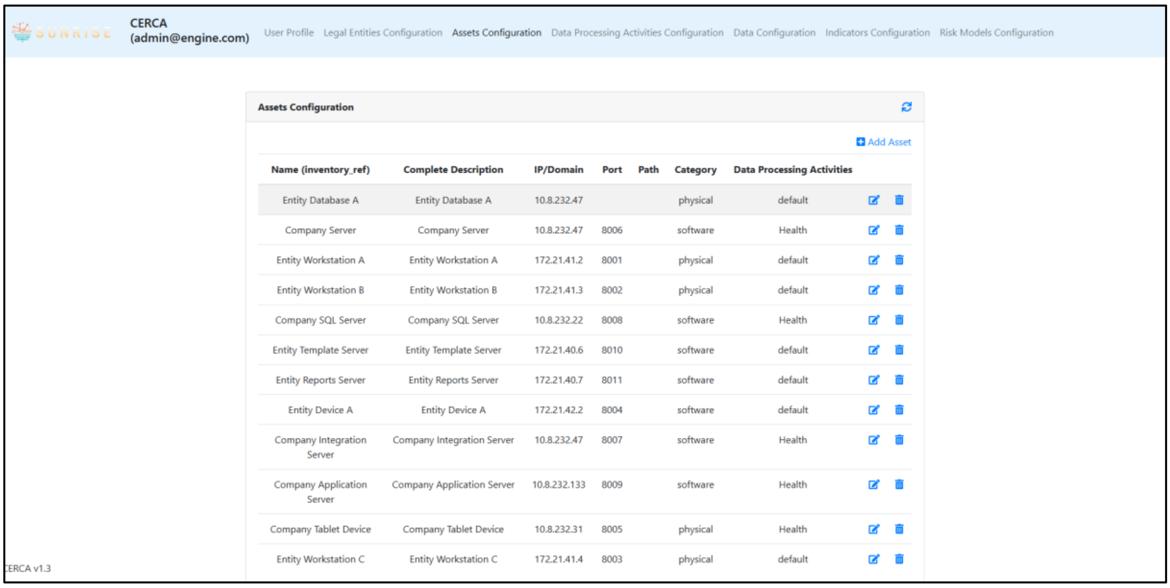

Assets Configuration

Name (inventory_ref)	Complete Description	IP/Domain	Port	Path	Category	Data Processing Activities		
Entity Database A	Entity Database A	10.8.232.47			physical	default		
Company Server	Company Server	10.8.232.47	8006		software	Health		
Entity Workstation A	Entity Workstation A	172.21.41.2	8001		physical	default		
Entity Workstation B	Entity Workstation B	172.21.41.3	8002		physical	default		
Company SQL Server	Company SQL Server	10.8.232.22	8008		software	Health		
Entity Template Server	Entity Template Server	172.21.40.6	8010		software	default		
Entity Reports Server	Entity Reports Server	172.21.40.7	8011		software	default		
Entity Device A	Entity Device A	172.21.42.2	8004		software	default		
Company Integration Server	Company Integration Server	10.8.232.47	8007		software	Health		
Company Application Server	Company Application Server	10.8.232.133	8009		software	Health		
Company Tablet Device	Company Tablet Device	10.8.232.31	8005		physical	Health		
Entity Workstation C	Entity Workstation C	172.21.41.4	8003		physical	default		

Figure 71. Assets Configuration dashboard.

4.2 Risks models selection

After clicking the “Select risk models” hyperlink, we are redirected to another menu, which is shown in Figure 72. The tool can suggest a risk model, but the user is free to choose whatever risk model they think best suits their infrastructure.

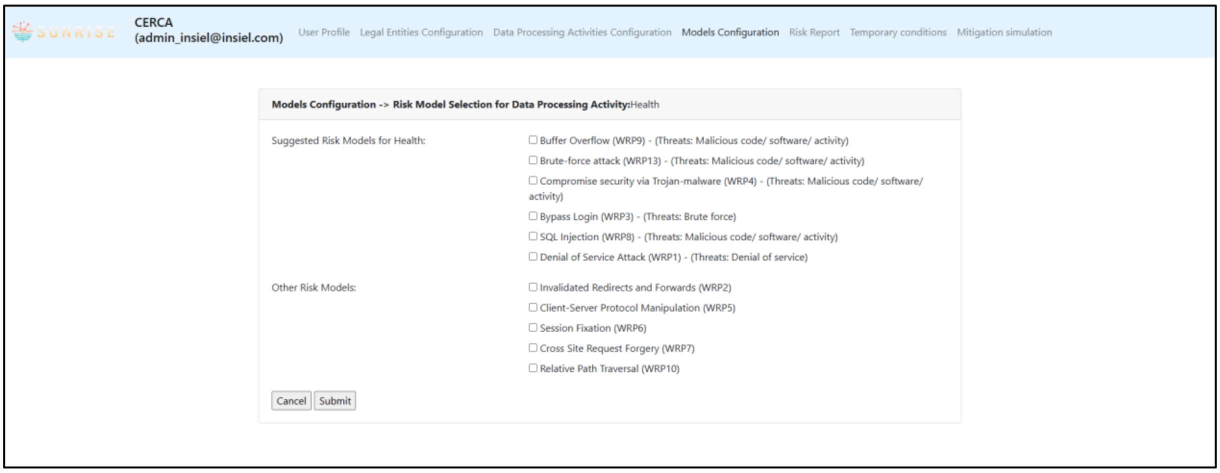

Models Configuration -> Risk Model Selection for Data Processing Activity: Health

Suggested Risk Models for Health:

- □ Buffer Overflow (WRP9) - (Threats: Malicious code/ software/ activity)
- □ Brute-force attack (WRP13) - (Threats: Malicious code/ software/ activity)
- □ Compromise security via Trojan-malware (WRP4) - (Threats: Malicious code/ software/ activity)
- □ Bypass Login (WRP3) - (Threats: Brute force)
- □ SQL Injection (WRP8) - (Threats: Malicious code/ software/ activity)
- □ Denial of Service Attack (WRP1) - (Threats: Denial of service)

Other Risk Models:

- □ Invalidated Redirects and Forwards (WRP2)
- □ Client-Server Protocol Manipulation (WRP5)
- □ Session Fixation (WRP6)
- □ Cross Site Request Forgery (WRP7)
- □ Relative Path Traversal (WRP10)

Figure 72. Risk model selection dashboard.

After selecting one of the risk models, we can observe that the models' configuration menu (Figure 73) is updated with the new information.

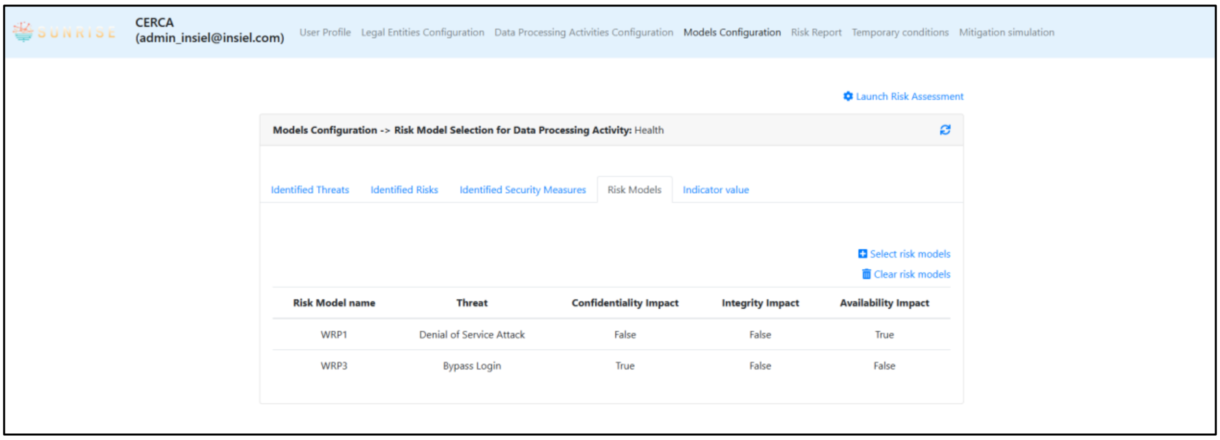

Models Configuration -> Risk Model Selection for Data Processing Activity: Health

Risk Model name	Threat	Confidentiality Impact	Integrity Impact	Availability Impact
WRP1	Denial of Service Attack	False	False	True
WRP3	Bypass Login	True	False	False

Figure 73. Model configuration dashboard updated with risk model.

4.3 Risk report

Finally, after inputting the necessary information and once selected the risk models, we are able to get the analysis performed by the tool. Figure 74 and Figure 75 show that the risk report works in both in a qualitative and quantitative way.

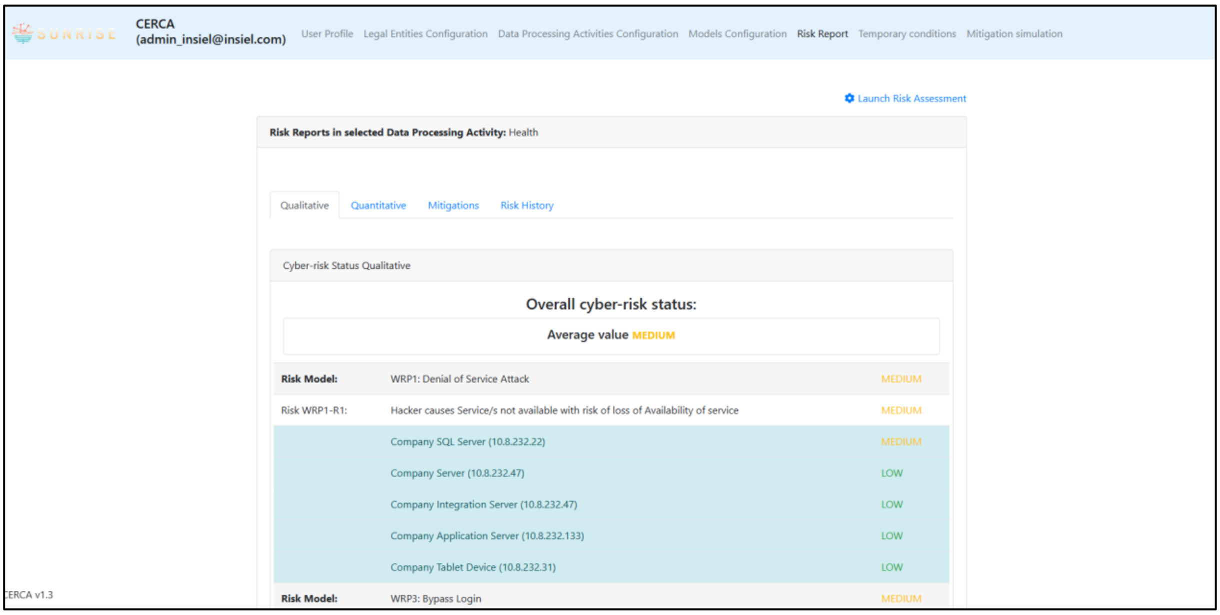

CERCA
(admin_insiel@insiel.com)

Risk Reports in selected Data Processing Activity: Health

Cyber-risk Status Qualitative

Overall cyber-risk status:

Average value MEDIUM

Risk Model:	WRP1: Denial of Service Attack	MEDIUM
Risk WRP1-R1:	Hacker causes Service/s not available with risk of loss of Availability of service	MEDIUM
	Company SQL Server (10.8.232.22)	MEDIUM
	Company Server (10.8.232.47)	LOW
	Company Integration Server (10.8.232.47)	LOW
	Company Application Server (10.8.232.133)	LOW
	Company Tablet Device (10.8.232.31)	LOW
Risk Model:	WRP3: Bypass Login	MEDIUM

Figure 74. Risk Report summary (qualitative assessment).

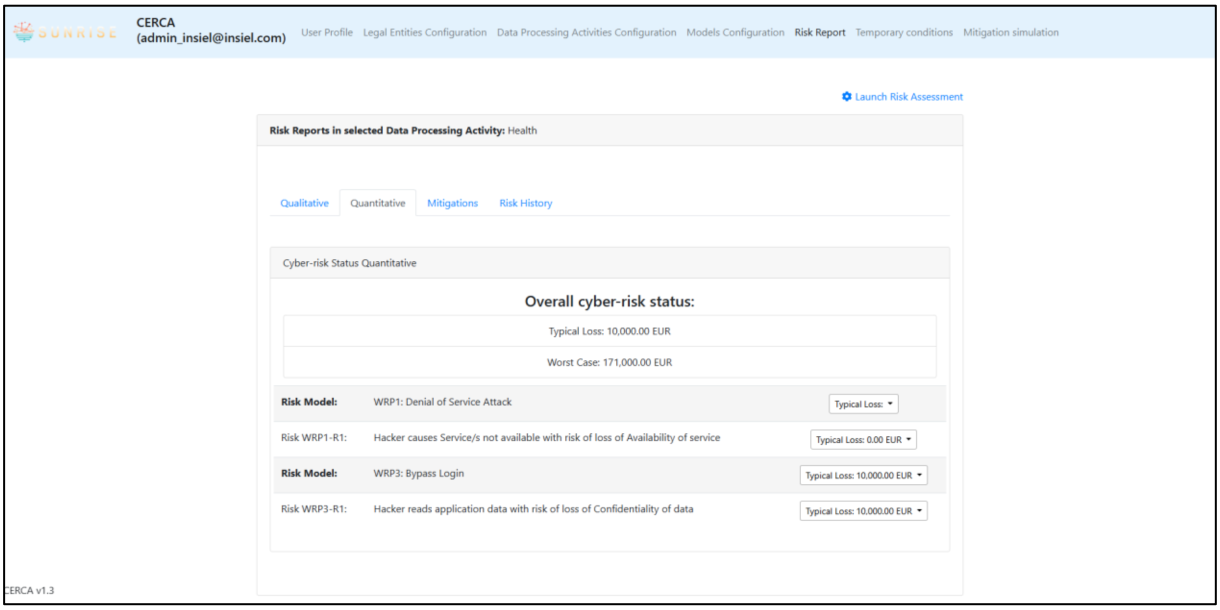

CERCA
(admin_insiel@insiel.com)

Risk Reports in selected Data Processing Activity: Health

Cyber-risk Status Quantitative

Overall cyber-risk status:

Typical Loss: 10,000.00 EUR

Worst Case: 171,000.00 EUR

Risk Model:	WRP1: Denial of Service Attack	Typical Loss: ▾
Risk WRP1-R1:	Hacker causes Service/s not available with risk of loss of Availability of service	Typical Loss: 0.00 EUR ▾
Risk Model:	WRP3: Bypass Login	Typical Loss: 10,000.00 EUR ▾
Risk WRP3-R1:	Hacker reads application data with risk of loss of Confidentiality of data	Typical Loss: 10,000.00 EUR ▾

Figure 75: Risk Report summary (quantitative assessment).

Suggested mitigations for each of the risk models selected are also available in the risk report summary (see Figure 76).

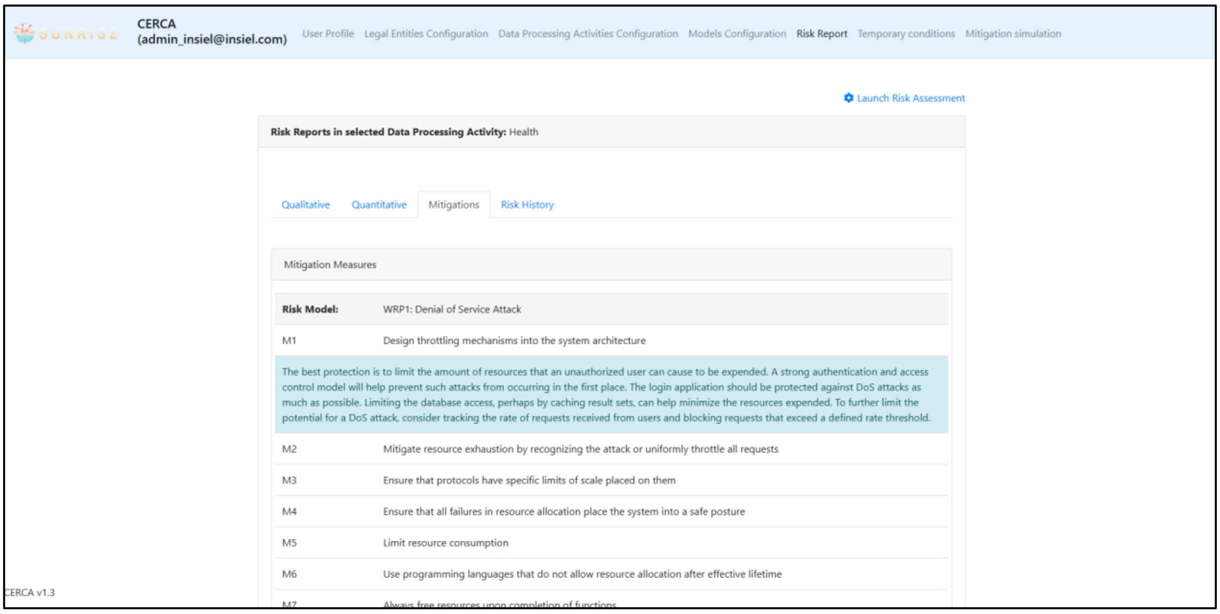

Risk Reports in selected Data Processing Activity: Health

Mitigation Measures

Risk Model:	WRP1: Denial of Service Attack
M1	Design throttling mechanisms into the system architecture
The best protection is to limit the amount of resources that an unauthorized user can cause to be expended. A strong authentication and access control model will help prevent such attacks from occurring in the first place. The login application should be protected against DoS attacks as much as possible. Limiting the database access, perhaps by caching result sets, can help minimize the resources expended. To further limit the potential for a DoS attack, consider tracking the rate of requests received from users and blocking requests that exceed a defined rate threshold.	
M2	Mitigate resource exhaustion by recognizing the attack or uniformly throttle all requests
M3	Ensure that protocols have specific limits of scale placed on them
M4	Ensure that all failures in resource allocation place the system into a safe posture
M5	Limit resource consumption
M6	Use programming languages that do not allow resource allocation after effective lifetime
M7	Always free resources upon completion of functions

Figure 76: Risk Report suggested mitigations.

Risk report is updated any time some of the inputs affecting indicators in the risk models is updates (e.g. some option of the questionnaire is updated, a new alarm is received, or a new threat intelligence data arrives). A new risk assessment can be also triggered clicking "Launch Risk Assessment" button from this risk report screen.

4.4 Indicator values reset

The indicator view and reset to default value functionality shows the current value of the indicators for the selected risk models. It is available in the "Indicator value" tab of the Models Configuration as shown in Figure 77.

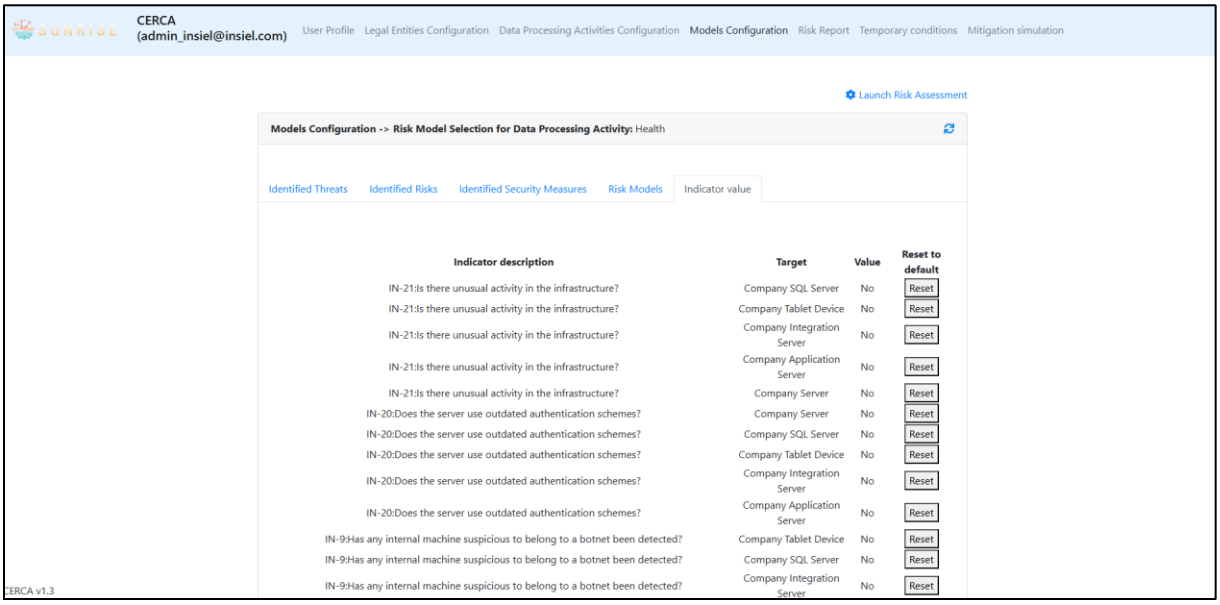

Models Configuration -> Risk Model Selection for Data Processing Activity: Health

Indicator description	Target	Value	Reset to default
IN-21:Is there unusual activity in the infrastructure?	Company SQL Server	No	Reset
IN-21:Is there unusual activity in the infrastructure?	Company Tablet Device	No	Reset
IN-21:Is there unusual activity in the infrastructure?	Company Integration Server	No	Reset
IN-21:Is there unusual activity in the infrastructure?	Company Application Server	No	Reset
IN-21:Is there unusual activity in the infrastructure?	Company Server	No	Reset
IN-20:Does the server use outdated authentication schemes?	Company Server	No	Reset
IN-20:Does the server use outdated authentication schemes?	Company SQL Server	No	Reset
IN-20:Does the server use outdated authentication schemes?	Company Tablet Device	No	Reset
IN-20:Does the server use outdated authentication schemes?	Company Integration Server	No	Reset
IN-20:Does the server use outdated authentication schemes?	Company Application Server	No	Reset
IN-9:Has any internal machine suspicious to belong to a botnet been detected?	Company Tablet Device	No	Reset
IN-9:Has any internal machine suspicious to belong to a botnet been detected?	Company SQL Server	No	Reset
IN-9:Has any internal machine suspicious to belong to a botnet been detected?	Company Integration Server	No	Reset

Figure 77: CERCA Indicators current values and reset.

An indicator may need to be reset to its default value to observe the effect on the risk assessment. In this example “IN-7 Is the account locked or had too many login attempts?” has a value of “Yes” (the default value is “No”) which makes the risk higher than it would be for the default value:

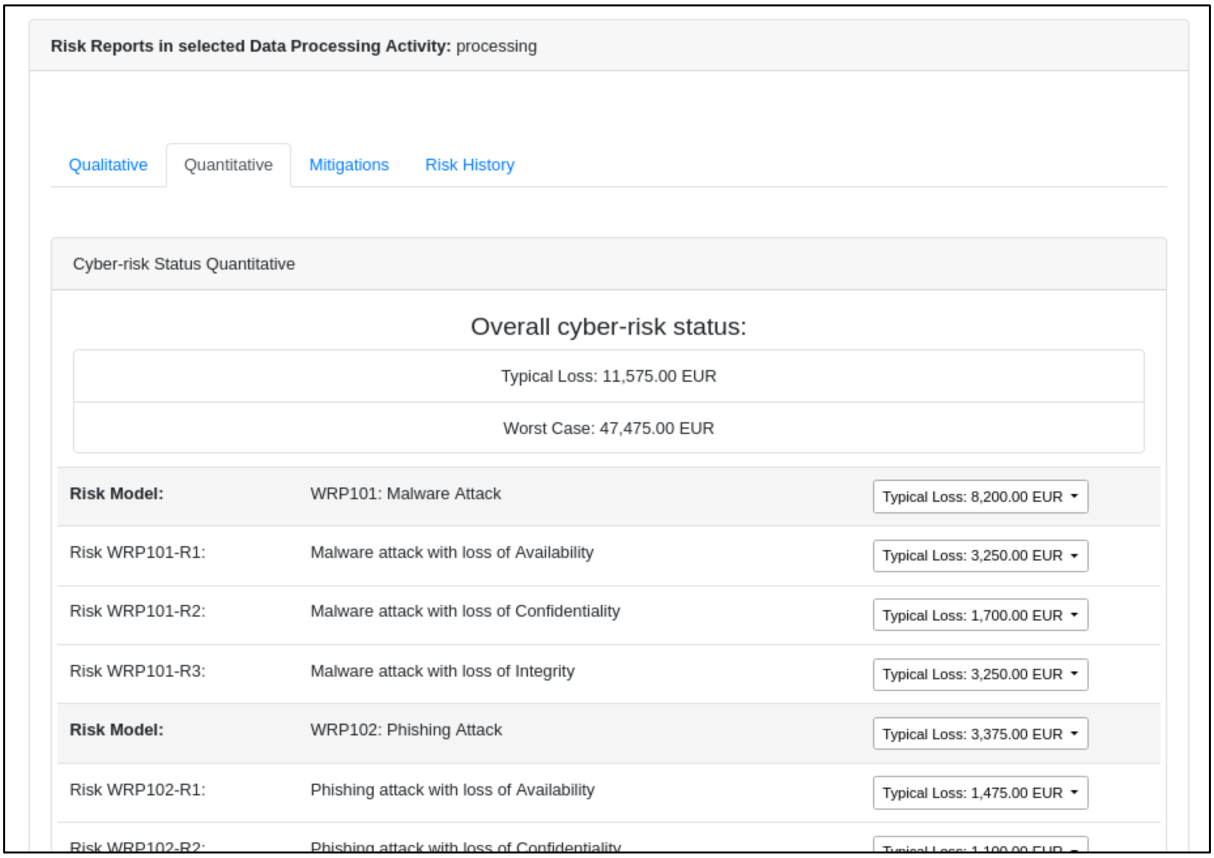

Risk Reports in selected Data Processing Activity: processing

Cyber-risk Status Quantitative

Overall cyber-risk status:

Typical Loss: 11,575.00 EUR

Worst Case: 47,475.00 EUR

Risk Model:	WRP101: Malware Attack	Typical Loss: 8,200.00 EUR
Risk WRP101-R1:	Malware attack with loss of Availability	Typical Loss: 3,250.00 EUR
Risk WRP101-R2:	Malware attack with loss of Confidentiality	Typical Loss: 1,700.00 EUR
Risk WRP101-R3:	Malware attack with loss of Integrity	Typical Loss: 3,250.00 EUR
Risk Model:	WRP102: Phishing Attack	Typical Loss: 3,375.00 EUR
Risk WRP102-R1:	Phishing attack with loss of Availability	Typical Loss: 1,475.00 EUR
Risk WRP102-R2:	Phishing attack with loss of Confidentiality	Typical Loss: 1,100.00 EUR

Figure 78: CERCA risk report before reset.

The value of any indicator can be reset:

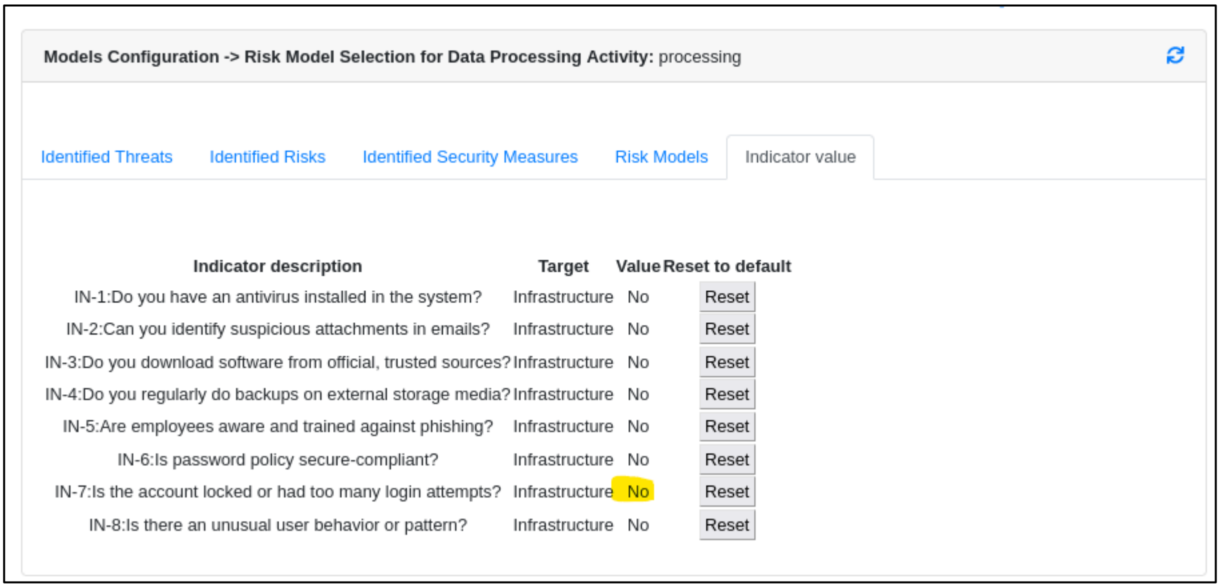

Models Configuration -> Risk Model Selection for Data Processing Activity: processing

Indicator description	Target	Value	Reset to default
IN-1:Do you have an antivirus installed in the system?	Infrastructure	No	Reset
IN-2:Can you identify suspicious attachments in emails?	Infrastructure	No	Reset
IN-3:Do you download software from official, trusted sources?	Infrastructure	No	Reset
IN-4:Do you regularly do backups on external storage media?	Infrastructure	No	Reset
IN-5:Are employees aware and trained against phishing?	Infrastructure	No	Reset
IN-6:Is password policy secure-compliant?	Infrastructure	No	Reset
IN-7:Is the account locked or had too many login attempts?	Infrastructure	No	Reset
IN-8:Is there an unusual user behavior or pattern?	Infrastructure	No	Reset

Figure 79: CERCA indicator reset.

The indicator update is automatically detected, and a new risk report is automatically generated:

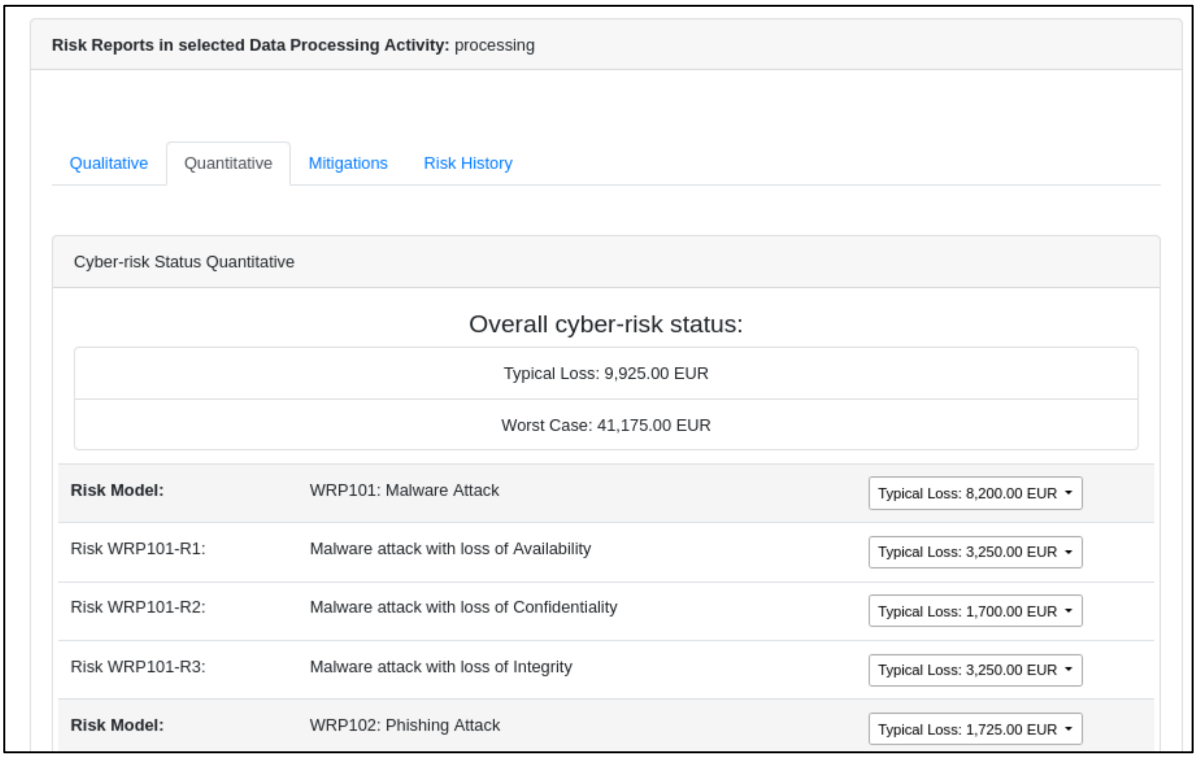

Risk Reports in selected Data Processing Activity: processing

Cyber-risk Status Quantitative

Overall cyber-risk status:

Typical Loss: 9,925.00 EUR

Worst Case: 41,175.00 EUR

Risk Model:	WRP101: Malware Attack	Typical Loss: 8,200.00 EUR
Risk WRP101-R1:	Malware attack with loss of Availability	Typical Loss: 3,250.00 EUR
Risk WRP101-R2:	Malware attack with loss of Confidentiality	Typical Loss: 1,700.00 EUR
Risk WRP101-R3:	Malware attack with loss of Integrity	Typical Loss: 3,250.00 EUR
Risk Model:	WRP102: Phishing Attack	Typical Loss: 1,725.00 EUR

Figure 80: CERCA risk report after reset.

4.5 Workforce questionnaire

The legal entities questionnaire contains the following:

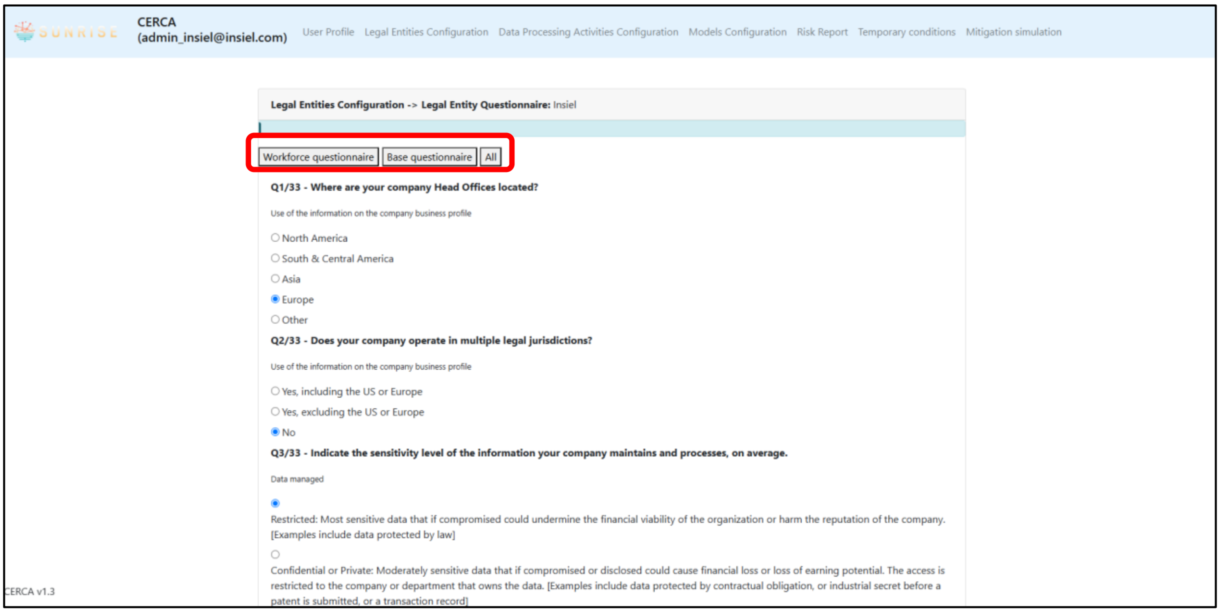

Legal Entities Configuration -> Legal Entity Questionnaire: Insiel

Q1/33 - Where are your company Head Offices located?

Use of the information on the company business profile

○ North America
○ South & Central America
○ Asia
● Europe
○ Other

Q2/33 - Does your company operate in multiple legal jurisdictions?

Use of the information on the company business profile

○ Yes, including the US or Europe
○ Yes, excluding the US or Europe
● No

Q3/33 - Indicate the sensitivity level of the information your company maintains and processes, on average.

Data managed

●
Restricted: Most sensitive data that if compromised could undermine the financial viability of the organization or harm the reputation of the company. [Examples include data protected by law]

○
Confidential or Private: Moderately sensitive data that if compromised or disclosed could cause financial loss or loss of earning potential. The access is restricted to the company or department that owns the data. [Examples include data protected by contractual obligation, or industrial secret before a patent is submitted, or a transaction record]

Figure 81: CERCA questionnaire options.

After clicking the “Workforce questionnaire” button, the questions are filtered to “Workforce” questions:

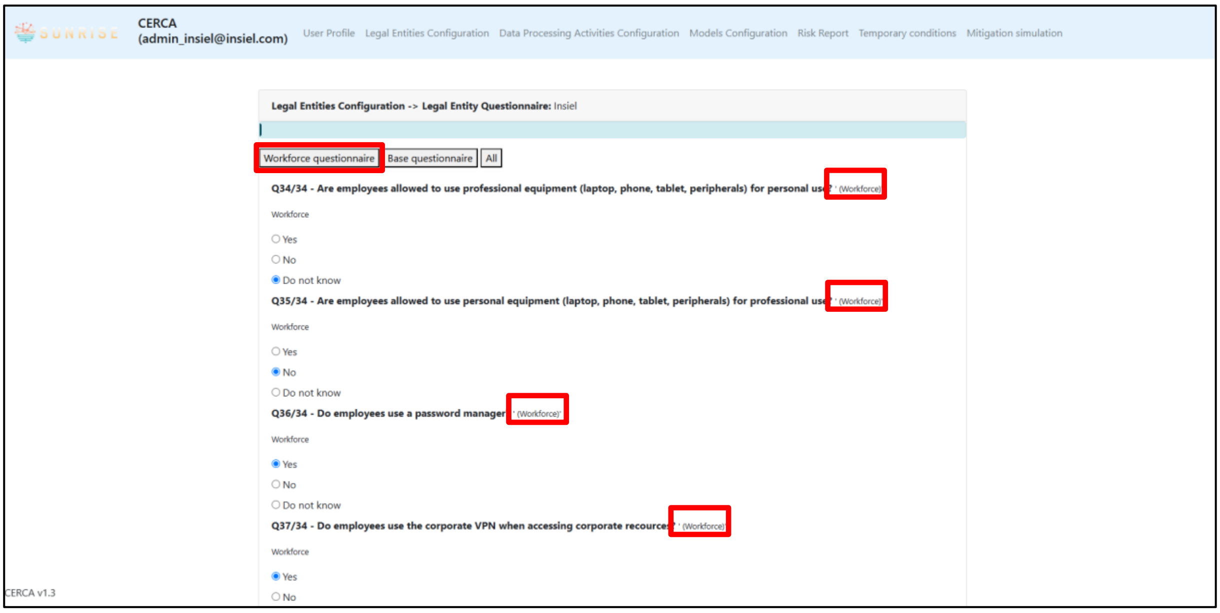

Legal Entities Configuration -> Legal Entity Questionnaire: Insiel

Q34/34 - Are employees allowed to use professional equipment (laptop, phone, tablet, peripherals) for personal use? ' (Workforce)

Workforce

○ Yes
○ No
● Do not know

Q35/34 - Are employees allowed to use personal equipment (laptop, phone, tablet, peripherals) for professional use? ' (Workforce)'

Workforce

○ Yes
● No
○ Do not know

Q36/34 - Do employees use a password manager? ' (Workforce)'

Workforce

● Yes
○ No
○ Do not know

Q37/34 - Do employees use the corporate VPN when accessing corporate recources? ' (Workforce)

Workforce

● Yes
○ No

Figure 82: CERCA Workforce questionnaire.

Questions tagged as “Workforce questionnaire” can be treated and configured separately and affect the risk within the context of workforce conditions.

The temporary conditions view contains configurations and information regarding the workforce.

The “workforce information” tab shows the answers to the workforce questionnaire, its effects on the risk models and the option to edit the answers.

4.6 Temporary conditions

Temporary conditions can be established after clicking the “Temporary conditions” tab as shown in Figure 83.

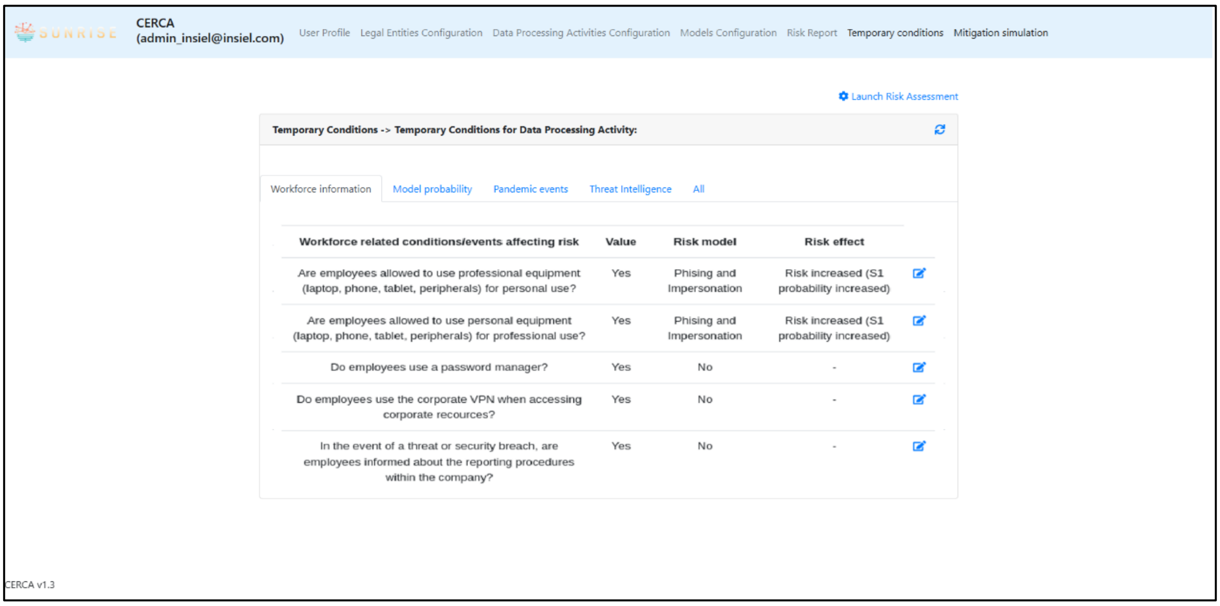

Temporary Conditions -> Temporary Conditions for Data Processing Activity:

Workforce related conditions/events affecting risk	Value	Risk model	Risk effect	
Are employees allowed to use professional equipment (laptop, phone, tablet, peripherals) for personal use?	Yes	Phising and Impersonation	Risk increased (S1 probability increased)	📝
Are employees allowed to use personal equipment (laptop, phone, tablet, peripherals) for professional use?	Yes	Phising and Impersonation	Risk increased (S1 probability increased)	📝
Do employees use a password manager?	Yes	No	-	📝
Do employees use the corporate VPN when accessing corporate recources?	Yes	No	-	📝
In the event of a threat or security breach, are employees informed about the reporting procedures within the company?	Yes	No	-	📝

Figure 83: Workforce information tab.

Probabilities of an attack initiating “Start-1” or S1 and of an attack being a successful “Unwanted incident 1” or U1 (the number after S and U indicates a possible path from “Start” to “Unwanted incident”, N:N relation) are available under the “Model probability” tab and can be edited.

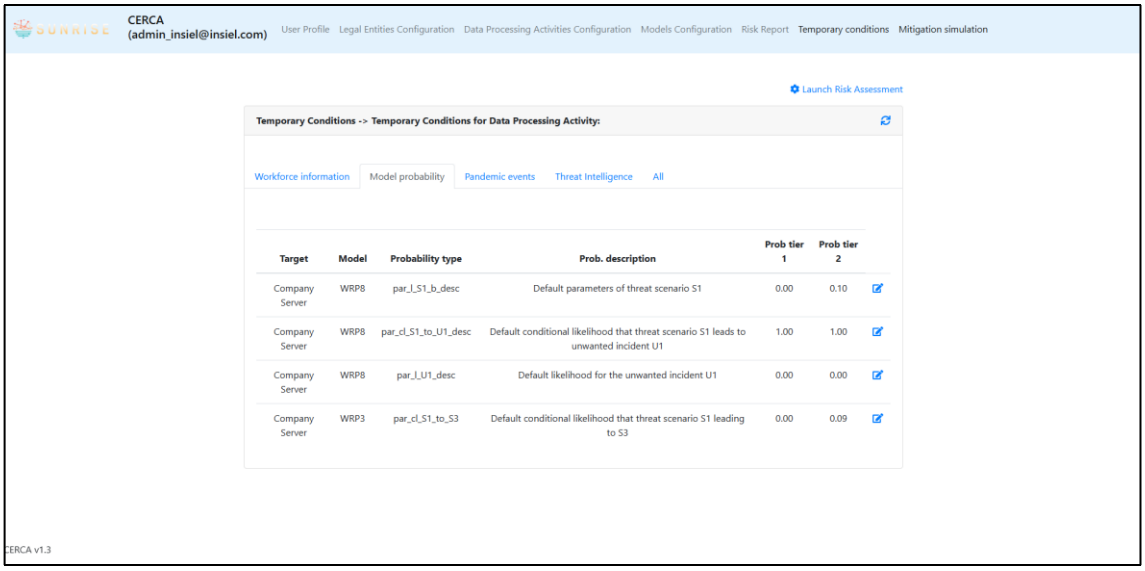

Temporary Conditions -> Temporary Conditions for Data Processing Activity:

Target	Model	Probability type	Prob. description	Prob tier 1	Prob tier 2
Company Server	WRP8	par_l_S1_b_desc	Default parameters of threat scenario S1	0.00	0.10
Company Server	WRP8	par_cl_S1_to_U1_desc	Default conditional likelihood that threat scenario S1 leads to unwanted incident U1	1.00	1.00
Company Server	WRP8	par_l_U1_desc	Default likelihood for the unwanted incident U1	0.00	0.00
Company Server	WRP3	par_cl_S1_to_S3	Default conditional likelihood that threat scenario S1 leading to S3	0.00	0.09

Figure 84: Model probability tab: information per asset.

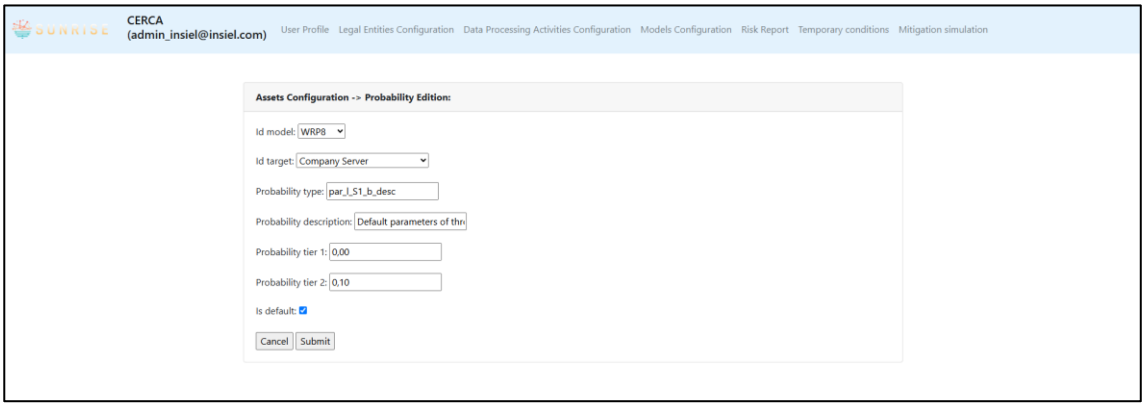

Assets Configuration -> Probability Edition:

Id model: WRP8

Id target: Company Server

Probability type: par_l_S1_b_desc

Probability description: Default parameters of thr

Probability tier 1: 0,00

Probability tier 2: 0,10

Is default: ☑

Figure 85: Probability edition for an asset and model.

The model conditionings Threat Intelligence tab shows conditionings like MISP events affecting a specific target or entity:

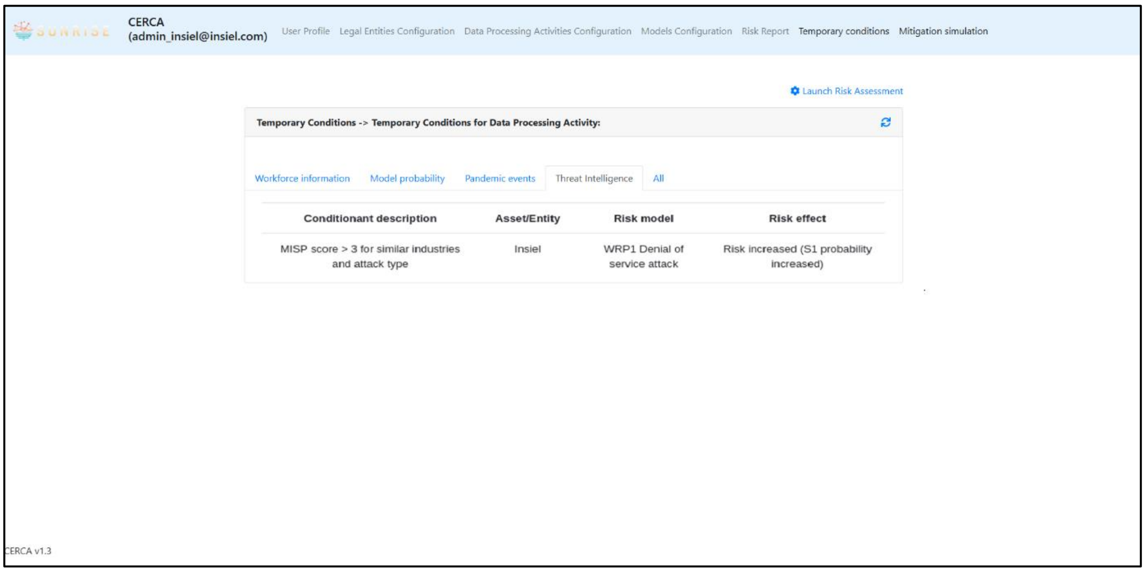

Temporary Conditions -> Temporary Conditions for Data Processing Activity:

Conditionant description	Asset/Entity	Risk model	Risk effect
MISP score > 3 for similar industries and attack type	Insiel	WRP1 Denial of service attack	Risk increased (S1 probability increased)

Figure 86: Model conditionings tab.

4.7 Risk assessment related to physical events

Indicators related to physical events can be added to the risk models to be considered for risk assessment. An example could be some image used by the SUNRISE AI-based inspection tool that could generate an output vector that triggers an indicator update at CERCA and a new risk assessment.

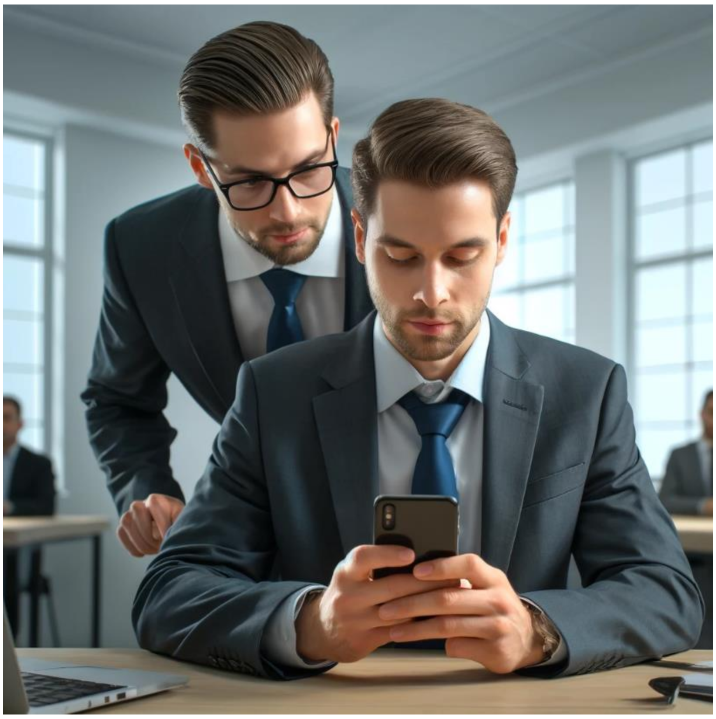

Figure 87: Shoulder surfing example.

The image displayed in Figure 87 as a shoulder surfing example has been generated using artificial intelligence (AI) technology, specifically with Dally3 from OpenAI. This choice was made due to the lack of freely licensable images that adequately represent this concept. During the proof of concept (PoC) phase, real images were used that cannot be included in this document due to copyright restrictions, though these images represent scenarios and results similar to those presented here.

4.8 Mitigation simulation

The result of the mitigation simulation is shown at the dashboard clicking “Mitigation Simulation” tab as shown in Figure 88. A simulation is a run with the following characteristics:

- A mitigation nullifies the risk introduced by an indicator or temporary condition
- A mitigation can affect several risk models by affecting a shared indicator
- For each selected risk model, a simulation is run for each target with a set of possible values for the indicators that are affected by the mitigations
- The results show simulated risks per asset and aggregated, with the combination of indicators and mitigations that lead to each risk

In Figure 88 , Mitigations 1 and 2 (MIT_1 and MIT_2), affect indicators 34 and 35 respectively. The risk is simulated for a specific risk model and target. With this feature, the user can check the potential risk reduction if a mitigation is to be applied.

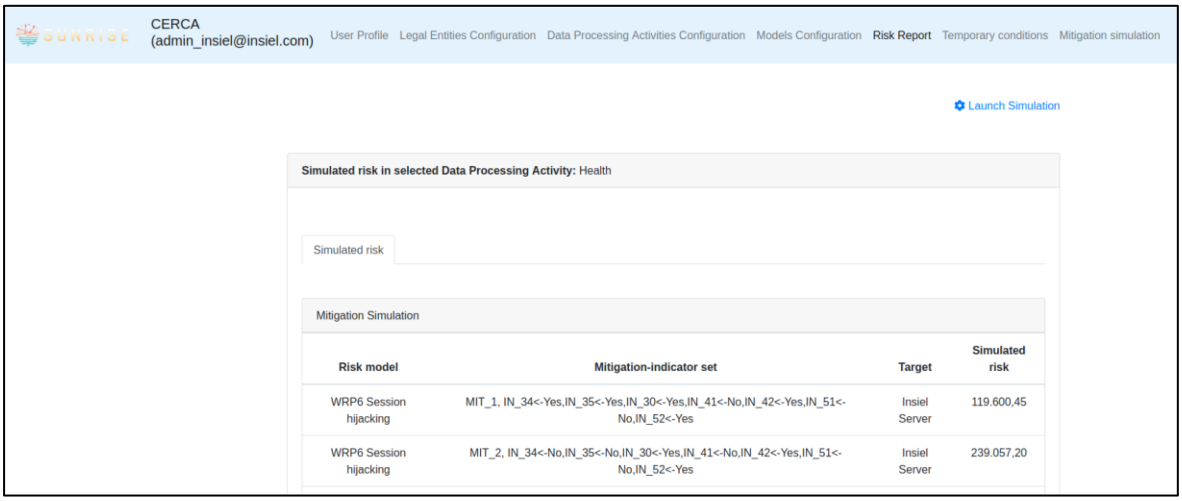

CERCA (admin_insiel@insiel.com)

Simulated risk in selected Data Processing Activity: Health

Mitigation Simulation

Risk model	Mitigation-indicator set	Target	Simulated risk
WRP6 Session hijacking	MIT_1, IN_34<-Yes,IN_35<-Yes,IN_30<-Yes,IN_41<-No,IN_42<-Yes,IN_51<-No,IN_52<-Yes	Insiel Server	119.600,45
WRP6 Session hijacking	MIT_2, IN_34<-No,IN_35<-No,IN_30<-Yes,IN_41<-No,IN_42<-Yes,IN_51<-No,IN_52<-Yes	Insiel Server	239.057,20

Figure 88 Mitigation simulation with simulated risk

5 Incident Reporting

This manual is an updated version of the initial manual of the ATOS Incident Reporting Engine (AIRE) tool produced during the CyberSec4Europe project [13] This manual incorporates with new functionalities that have been added to the tool and updated screenshots extracted from the deployment of the tool in SUNRISE pilots.

The Incident Reporting module (AIRE) has two main interaction modes. The first is a graphical interface where users can create new incidents or transform the alarms into new incidents, follow the evolution of the incidents, and create and manage reports for authorities. The second is an API that allows the unattended ingestion of the output from other assets. In addition, AIRE has an integrated connection with MISP, which allows it to receive threat alarms from other similar infrastructures.

Next subsections present step-by-step instructions to show first how to configure the module for a specific Critical Infrastructure and then how end-users will interact with the tool to perform regular tasks.

5.1 Setting up Incident Reporting module

Before working with the CPR Incident Reporting module, it is necessary that the user with Administrator role (IT Department) registers the organization, users, contact addresses, regulations, etc. To do so, login as a user with administration permissions. The access is done clicking "Incident Reporting" tab in the CPR Dashboard (Figure 1) through single sign-on. Then you will see the list of tabs for the different configurations, see Figure 89:

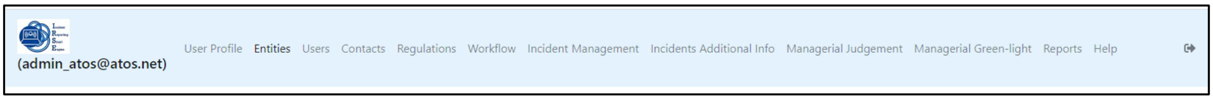

Figure 89. AIRE navigation bar.

- **Entities:** This shows the list of organizations that are registered on AIRE. See Figure 90:

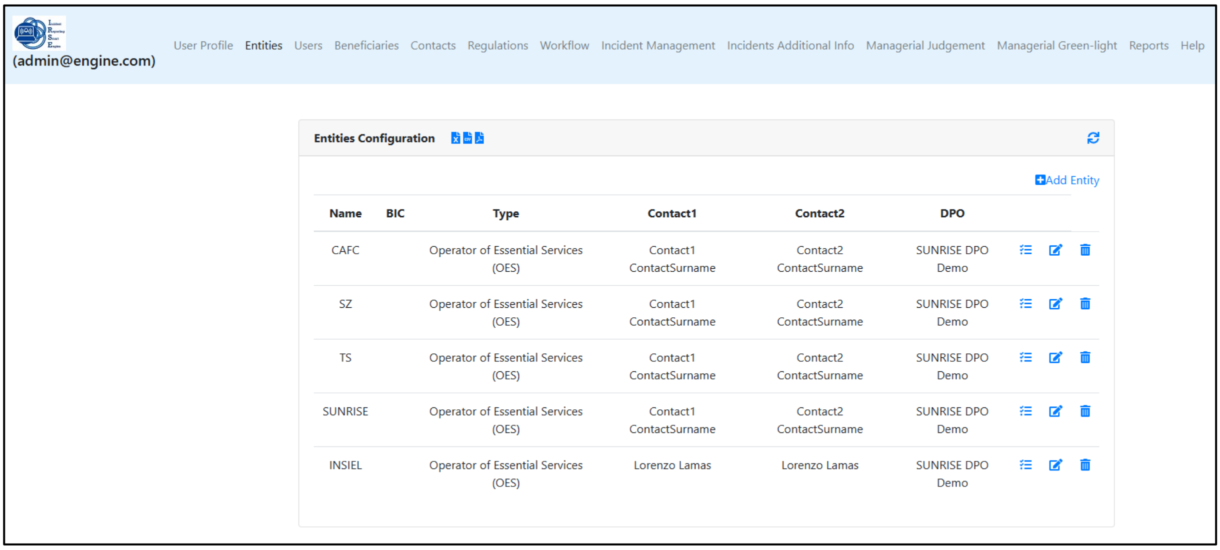

Name	BIC	Type	Contact1	Contact2	DPO
CAFC		Operator of Essential Services (OES)	Contact1 ContactSurname	Contact2 ContactSurname	SUNRISE DPO Demo
SZ		Operator of Essential Services (OES)	Contact1 ContactSurname	Contact2 ContactSurname	SUNRISE DPO Demo
TS		Operator of Essential Services (OES)	Contact1 ContactSurname	Contact2 ContactSurname	SUNRISE DPO Demo
SUNRISE		Operator of Essential Services (OES)	Contact1 ContactSurname	Contact2 ContactSurname	SUNRISE DPO Demo
INSIEL		Operator of Essential Services (OES)	Lorenzo Lamas	Lorenzo Lamas	SUNRISE DPO Demo

Figure 90. Entities Configuration List.

Click the *Add Entity* or the edit button to open the form that allows you to create or modify a new organization (see Figure 91). It is important to configure "TheHive API Key" with the API key of the user "aire@<organization>" (with org-admin" profile), since this will be used by AIRE to interact with TheHive.

The access to TheHive GUI is available from the tab "Incident Management" and it is necessary to login with the admin user for the organization (or the global TheHive admin user) to obtain this information.

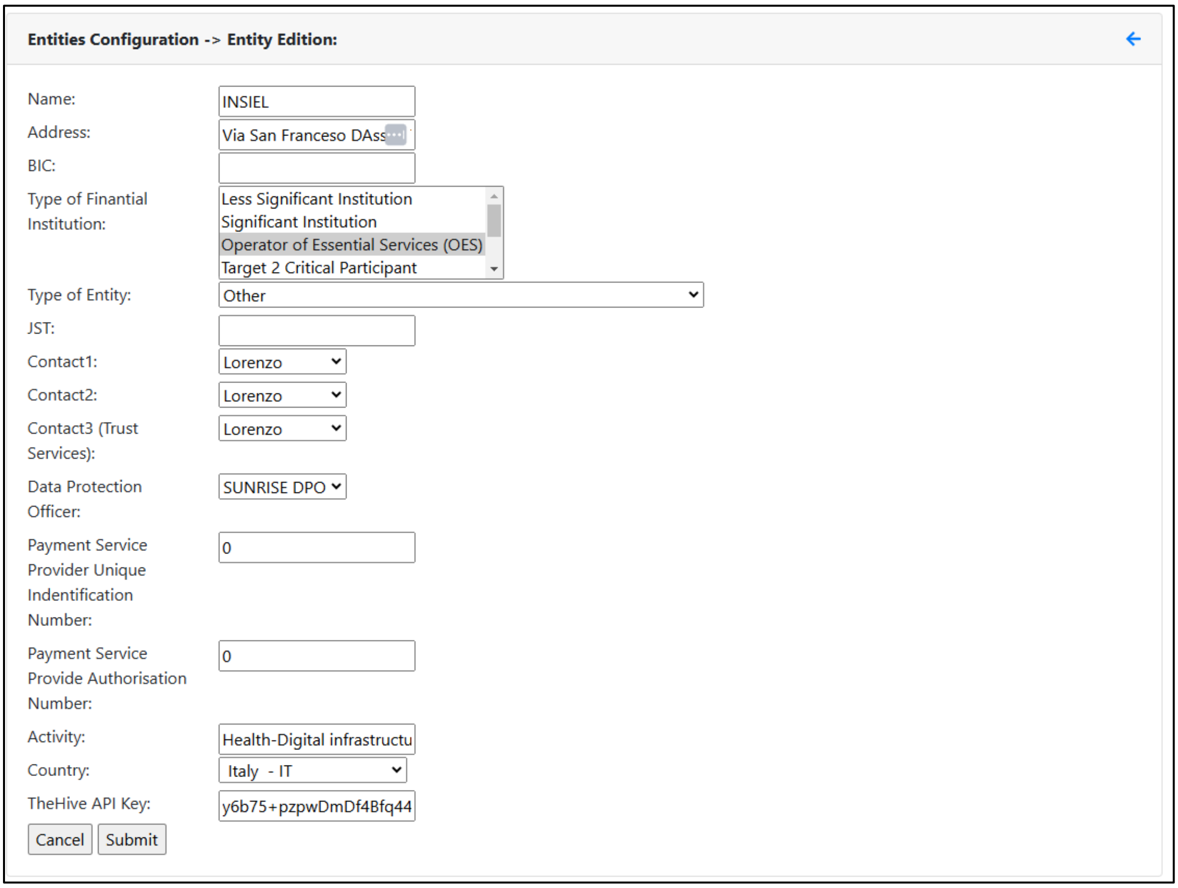

Entities Configuration -> Entity Edition:

Name:	INSIEL
Address:	Via San Franceso DAss
BIC:	
Type of Finantial Institution:	Less Significant Institution Significant Institution Operator of Essential Services (OES) Target 2 Critical Participant
Type of Entity:	Other
JST:	
Contact1:	Lorenzo
Contact2:	Lorenzo
Contact3 (Trust Services):	Lorenzo
Data Protection Officer:	SUNRISE DPO
Payment Service Provider Unique Indentification Number:	0
Payment Service Provide Authorisation Number:	0
Activity:	Health-Digital infrastructu
Country:	Italy - IT
TheHive API Key:	y6b75+pzpwDmDf4Bfq44

Figure 91. New entity form.

Then, choose the regulation that the organization must meet (as shown in Figure 92). Only reports associated to enabled regulations will be generated by the platform. This information should be provided to the administrator by the **CI legal and compliance team.**

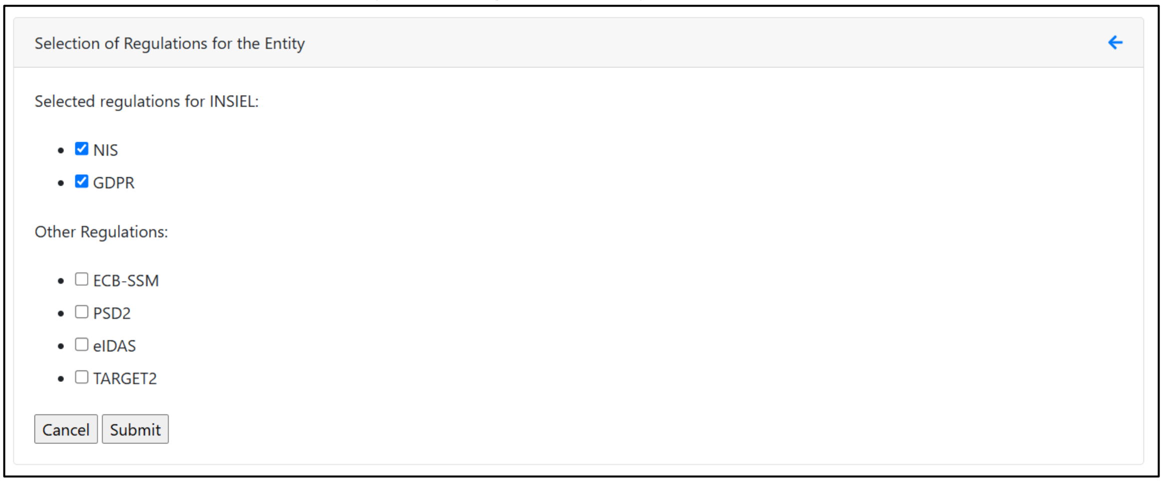

Selection of Regulations for the Entity

Selected regulations for INSIEL:

- ☑ NIS
- ☑ GDPR

Other Regulations:

- ☐ ECB-SSM
- ☐ PSD2
- ☐ eIDAS
- ☐ TARGET2

Figure 92: Regulation for entities.

- **Users:** This shows the list of registered users with their information, including contact information, position and function (role). This page (see Figure 93) allows you to create new users (click on *Add User* button), or edit the information on the created users.

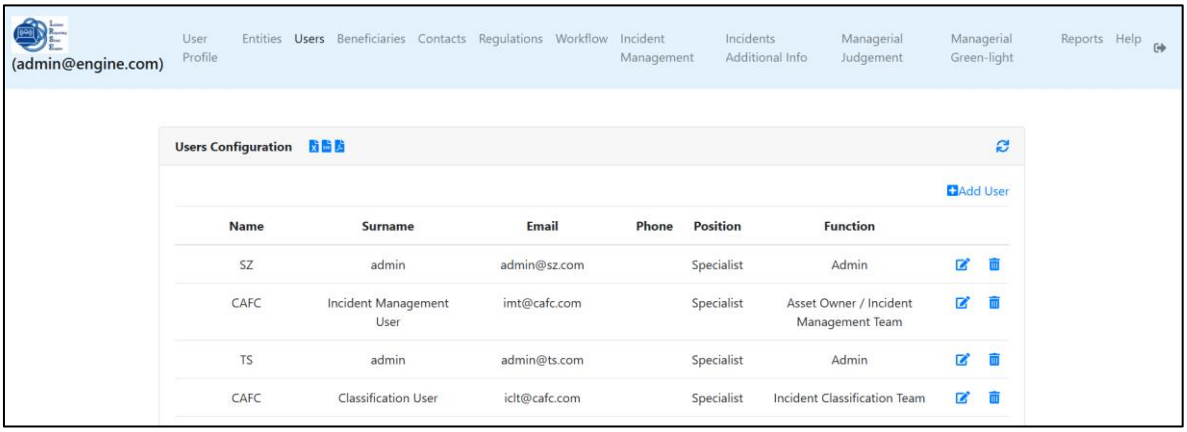

Name	Surname	Email	Phone	Position	Function
SZ	admin	admin@sz.com		Specialist	Admin
CAFC	Incident Management User	imt@cafc.com		Specialist	Asset Owner / Incident Management Team
TS	admin	admin@ts.com		Specialist	Admin
CAFC	Classification User	iclt@cafc.com		Specialist	Incident Classification Team

Figure 93: List of registered users.

Following configurations are related to mandatory incident reporting process and compliance with the different regulatory frameworks. Consequently, to set up correctly the following points, the administrator will require the support of the **CI legal and compliance team.**

- **Beneficiaries:** This shows the list of service beneficiaries that will receive a notification when a new incident is registered or when an ongoing incident is closed. See Figure 94.

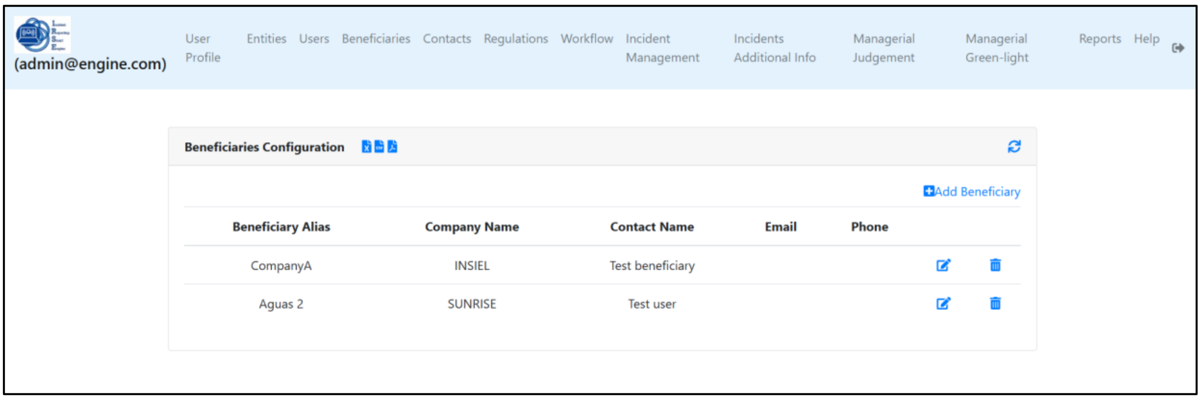

Beneficiary Alias	Company Name	Contact Name	Email	Phone
CompanyA	INSIEL	Test beneficiary		
Aguas 2	SUNRISE	Test user		

Figure 94. List of service beneficiaries.

- **Contacts:** This view (see Figure 95) allows you to create the list of all contacts that will be included in the incident reports generated. There are the following contact types:
  - Contact1: primary contact
  - Contact2: secondary contact
  - Contact3: associated to the trust officer
  - DPO: Data Protection Officer

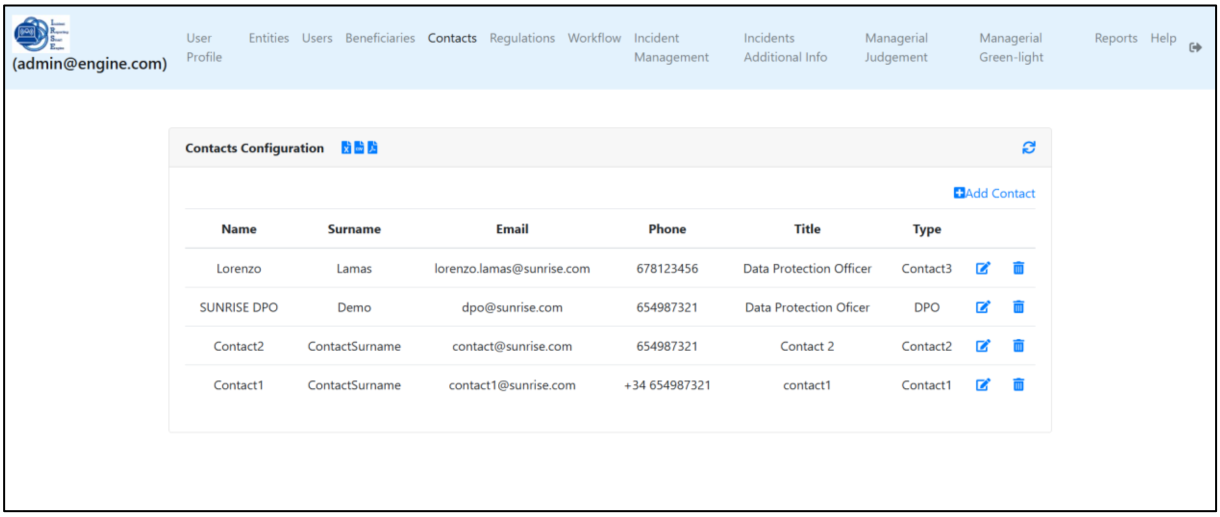

Contacts Configuration

Name	Surname	Email	Phone	Title	Type
Lorenzo	Lamas	lorenzo.lamas@sunrise.com	678123456	Data Protection Officer	Contact3
SUNRISE DPO	Demo	dpo@sunrise.com	654987321	Data Protection Oficer	DPO
Contact2	ContactSurname	contact@sunrise.com	654987321	Contact 2	Contact2
Contact1	ContactSurname	contact1@sunrise.com	+34 654987321	contact1	Contact1

Figure 95: List of configured contacts.

- **Regulations:** This is the view for Regulations for Mandatory Incident Reporting. This section manages the different regulations enabled in the tool and has several subsections:
  - **Regulations:** List of registered regulations, as shown in Figure 96.

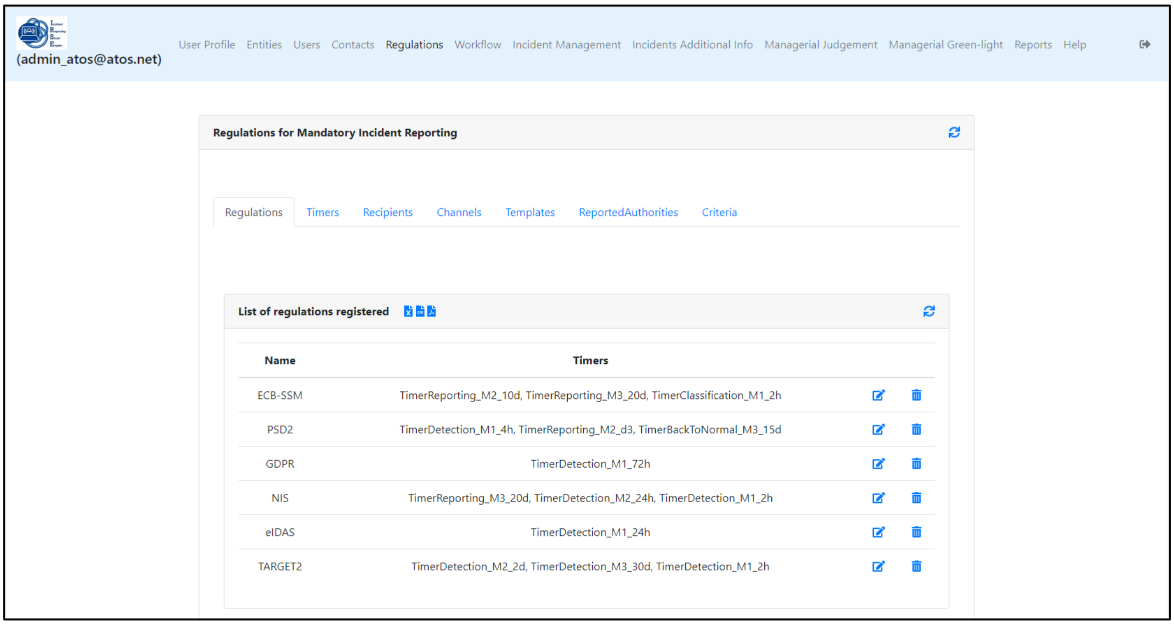

Regulations for Mandatory Incident Reporting

List of regulations registered

Name	Timers
ECB-SSM	TimerReporting_M2_10d, TimerReporting_M3_20d, TimerClassification_M1_2h
PSD2	TimerDetection_M1_4h, TimerReporting_M2_d3, TimerBackToNormal_M3_15d
GDPR	TimerDetection_M1_72h
NIS	TimerReporting_M3_20d, TimerDetection_M2_24h, TimerDetection_M1_2h
eIDAS	TimerDetection_M1_24h
TARGET2	TimerDetection_M2_2d, TimerDetection_M3_30d, TimerDetection_M1_2h

Figure 96: List of regulations registered.

For each regulation registered in the platform, as shown in Figure 97, it is necessary to indicate:

- The last phase of reporting: Depending on the reports that must be disseminated according to a specific directive or regulation, the following will be selected:
  - M1 (Initial) if only one report is required.
  - M2 (interim) in case a first and a second reports are required.
  - M3 (final) in case three mandatory reports (first, interim, and final) are necessary.
- The Timers that will be triggered with the regulations (see the next point about Timers).

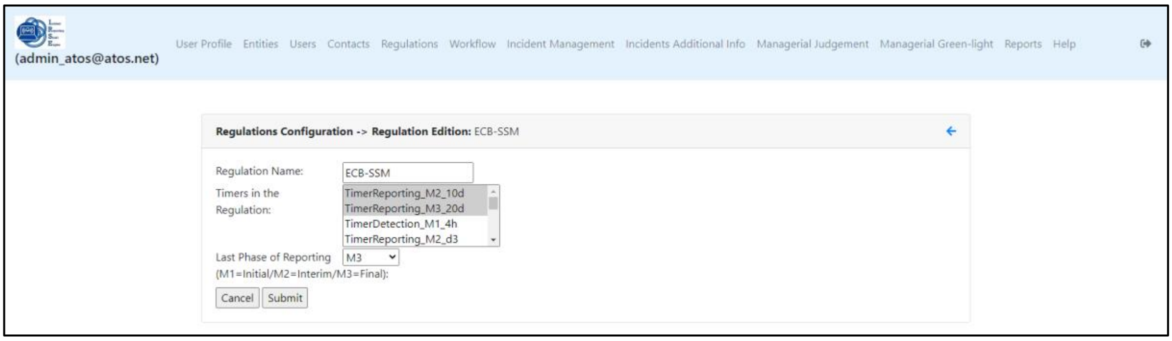

(admin_atos@atos.net)

Regulations Configuration -> Regulation Edition: ECB-SSM

Regulation Name: ECB-SSM

Timers in the Regulation: TimerReporting_M2_10d
TimerReporting_M3_20d
TimerDetection_M1_4h
TimerReporting_M2_d3

Last Phase of Reporting (M1=Initial/M2=Interim/M3=Final): M3

Figure 97: New regulation form.

- **Timers:** This indicates when a notification needs to be sent to the incident contact user (the email configured as Contact User in the TheHive template), as shown in Figure 98. This is relevant in the case where a mandatory report has not been sent to the corresponding Supervisory Authorities in the deadline defined by the specific regulation.

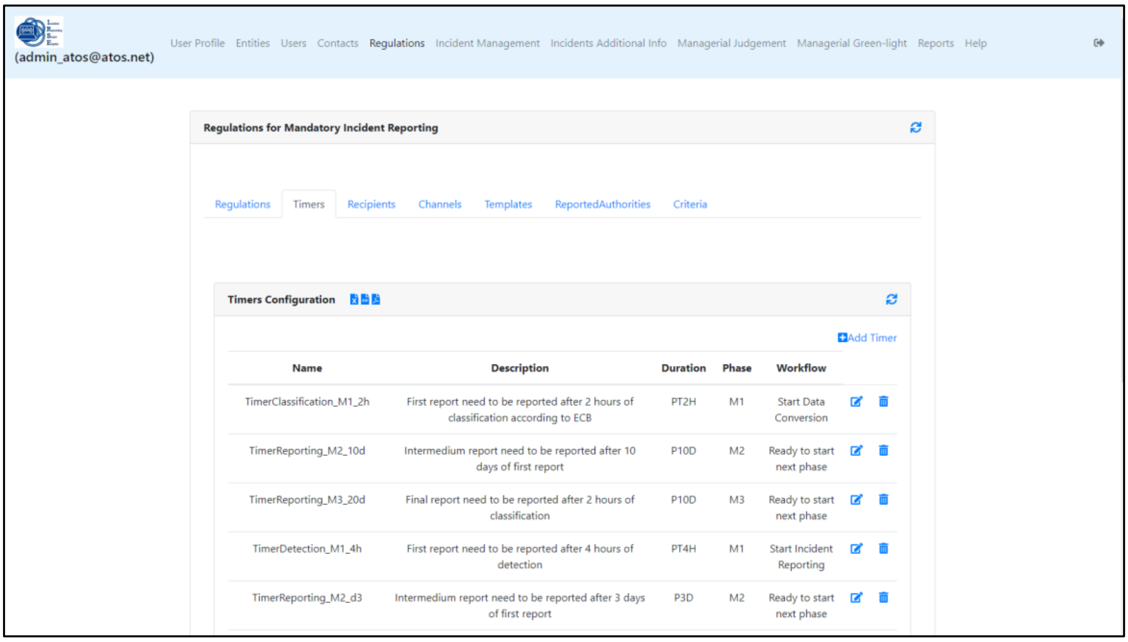

(admin_atos@atos.net)

Regulations for Mandatory Incident Reporting

Timers Configuration

Name	Description	Duration	Phase	Workflow
TimerClassification_M1_2h	First report need to be reported after 2 hours of classification according to ECB	PT2H	M1	Start Data Conversion
TimerReporting_M2_10d	Intermedium report need to be reported after 10 days of first report	P10D	M2	Ready to start next phase
TimerReporting_M3_20d	Final report need to be reported after 2 hours of classification	P10D	M3	Ready to start next phase
TimerDetection_M1_4h	First report need to be reported after 4 hours of detection	PT4H	M1	Start Incident Reporting
TimerReporting_M2_d3	Intermedium report need to be reported after 3 days of first report	P3D	M2	Ready to start next phase

Figure 98. List of timers.

The Timer Edition, as shown in Figure 99, enables the creation of new timers or the modification of existing ones. The *timer duration* field defines the time windows within the report that must be sent and uses the ISO 8601 durations format [3]. The *Report phase* defines the phase in which the report should be sent. The *Workflow stage* defines the event that triggers the current timer.

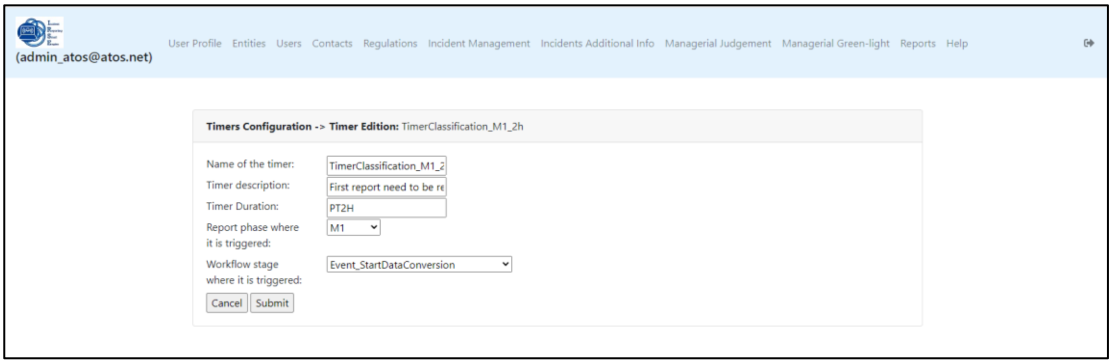

Timers Configuration -> Timer Edition: TimerClassification_M1_2h

Name of the timer:	TimerClassification_M1_2
Timer description:	First report need to be re
Timer Duration:	PT2H
Report phase where it is triggered:	M1
Workflow stage where it is triggered:	Event_StartDataConversion

Figure 99. Timer edition form.

AIRE sends an email to the responsible incident contact user when a report is not sent on time. The contact user email must be configured for each incident in the information collected through TheHive.

- **Recipients:** This is associated to each entity and regulation, as shown in Figure 100 and Figure 101. This will have also an associated Channel, which will be shown once the report has been generated by the platform as a suggestion of a channel (e.g. email address) that need to be followed for the reporting. Only those channels configured as enabled will trigger that automatically an email is sent with the report, but by default they will be disabled.

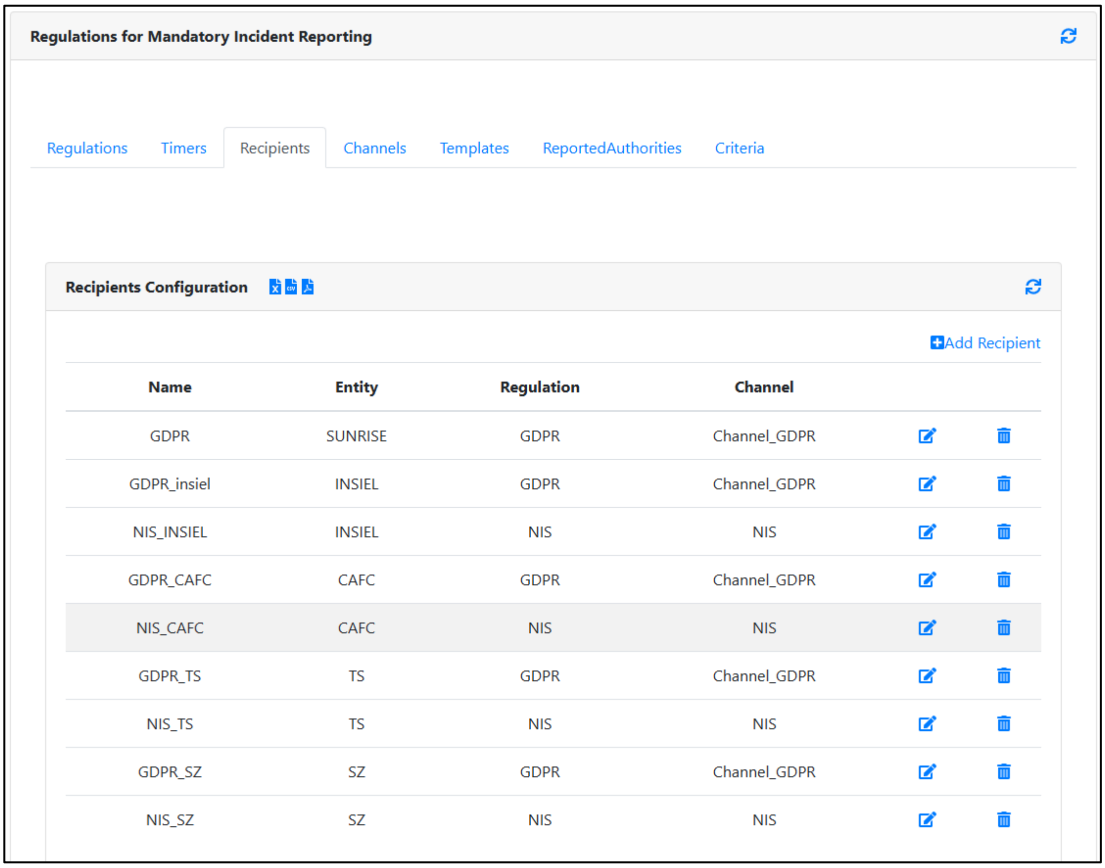

Name	Entity	Regulation	Channel
GDPR	SUNRISE	GDPR	Channel_GDPR
GDPR_insiel	INSIEL	GDPR	Channel_GDPR
NIS_INSIEL	INSIEL	NIS	NIS
GDPR_CAFC	CAFC	GDPR	Channel_GDPR
NIS_CAFC	CAFC	NIS	NIS
GDPR_TS	TS	GDPR	Channel_GDPR
NIS_TS	TS	NIS	NIS
GDPR_SZ	SZ	GDPR	Channel_GDPR
NIS_SZ	SZ	NIS	NIS

Figure 100. List of Recipient Configurations.

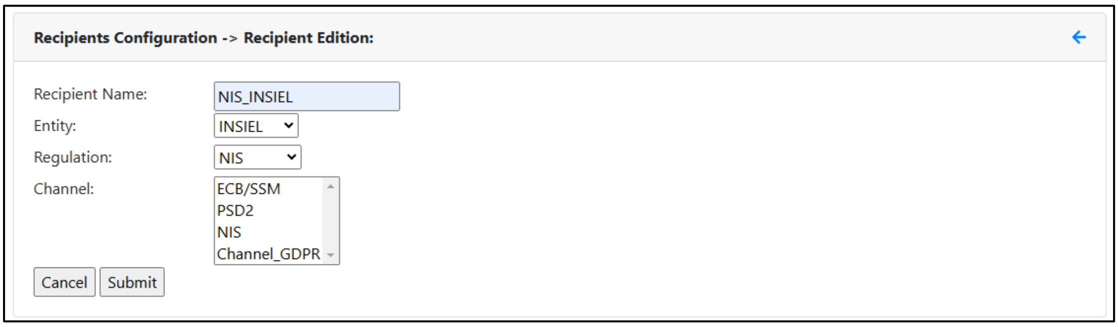

Recipients Configuration -> Recipient Edition:

Recipient Name:	NIS_INSIEL
Entity:	INSIEL
Regulation:	NIS
Channel:	ECB/SSM PSD2 NIS Channel_GDPR

Figure 101. Recipient Edition Form.

- **Channels:** The information registered here, as shown in Figure 102, will be used in the Recipients and shown to the user as a suggestion when the reports are ready for revision and releasing.

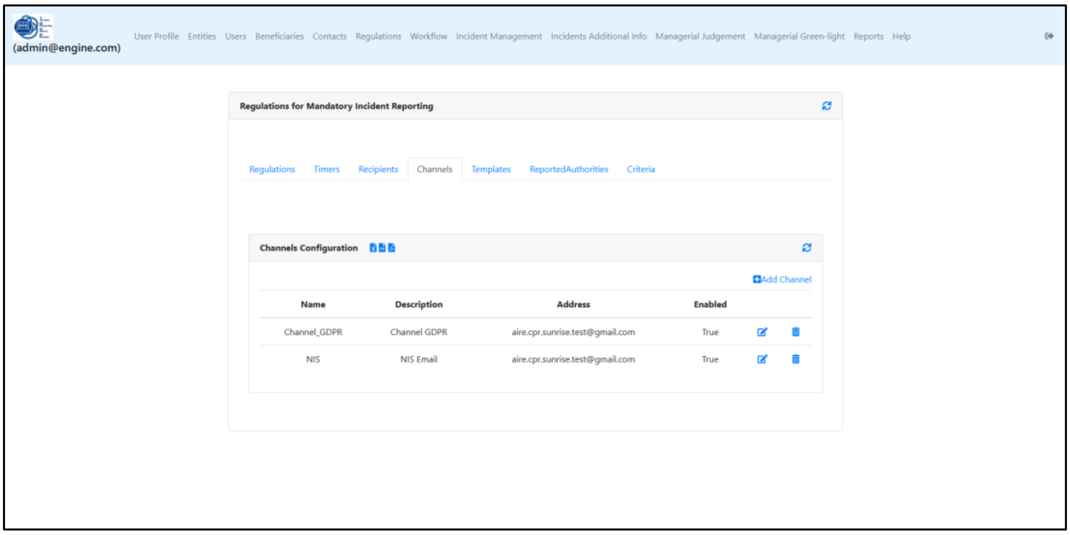

Regulations for Mandatory Incident Reporting

Channels Configuration

Name	Description	Address	Enabled		
Channel_GDPR	Channel GDPR	aire.cpr.sunrise.test@gmail.com	True	📝	🗑
NIS	NIS Email	aire.cpr.sunrise.test@gmail.com	True	📝	🗑

Figure 102. List of communication channels.

The *Add Channel* button enables the addition or modification of channels, as shown in Figure 103. If the channel is set up as “Enabled”, this means that once the user with the controller role approves the reporting, the generated report will automatically send the emails configured in these enabled channels to the Supervisory Authorities.

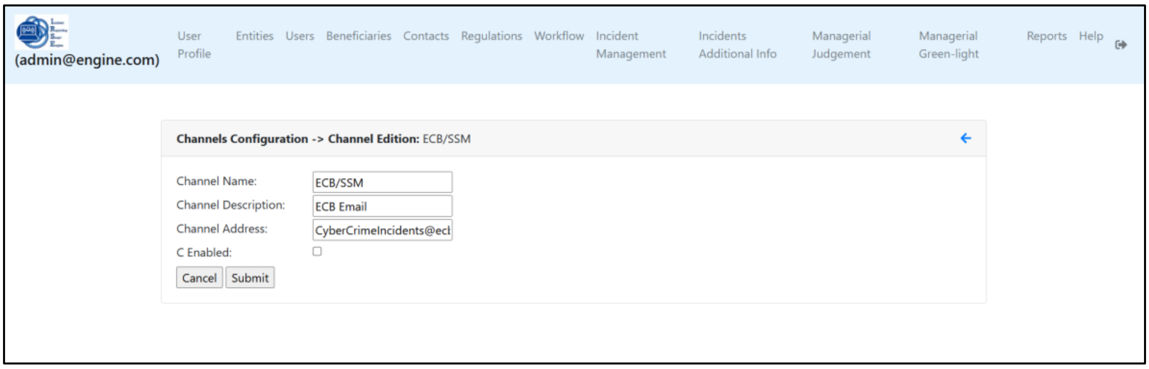

Channels Configuration -> Channel Edition: ECB/SSM

Channel Name:	ECB/SSM
Channel Description:	ECB Email
Channel Address:	CyberCrimeIncidents@ecl
C Enabled:	□

Figure 103: Channel edition.

- **Templates:** The template used by the platform for the generation of the report needs to be associated with a regulation and a recipient. The formats currently supported are EXCEL, PDF, and WORD. The template data format is the format used by the platform to populate information about the times in the reports. The date and time formats used are the ones defined in SimpleDateFormat[6].

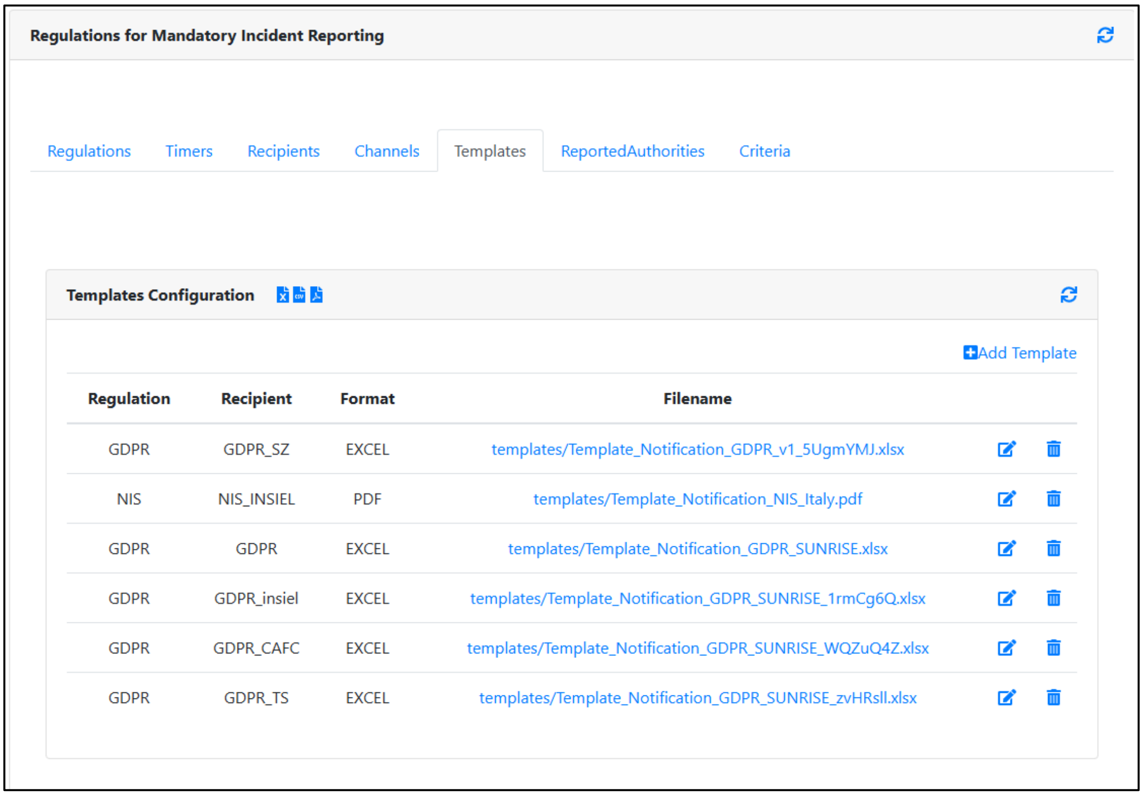

Regulations for Mandatory Incident Reporting

Templates Configuration

Regulation	Recipient	Format	Filename		
GDPR	GDPR_SZ	EXCEL	templates/Template_Notification_GDPR_v1_5UgmYMJ.xlsx	✎	🗑
NIS	NIS_INSIEL	PDF	templates/Template_Notification_NIS_Italy.pdf	✎	🗑
GDPR	GDPR	EXCEL	templates/Template_Notification_GDPR_SUNRISE.xlsx	✎	🗑
GDPR	GDPR_insiel	EXCEL	templates/Template_Notification_GDPR_SUNRISE_1rmCg6Q.xlsx	✎	🗑
GDPR	GDPR_CAFC	EXCEL	templates/Template_Notification_GDPR_SUNRISE_WQZuQ4Z.xlsx	✎	🗑
GDPR	GDPR_TS	EXCEL	templates/Template_Notification_GDPR_SUNRISE_zvHRsll.xlsx	✎	🗑

Figure 104. List of Templates.

Each *Template* is associated with a *Regulation*, a *Recipient*, a *Report template file*, and a *Template mapping file*. The *Report template file* is the base for the generated report and accepts the extensions *pdf*, *doc*, *docx*, *xls*, and *xlsx*; the extension must be specified in the *Template format* field. The *Template mapping file* identifies which information from the Incident Register database need to be used in each field of the report template file, and the *Template date format* specifies what the date format is. Figure 105 shows an example of this.

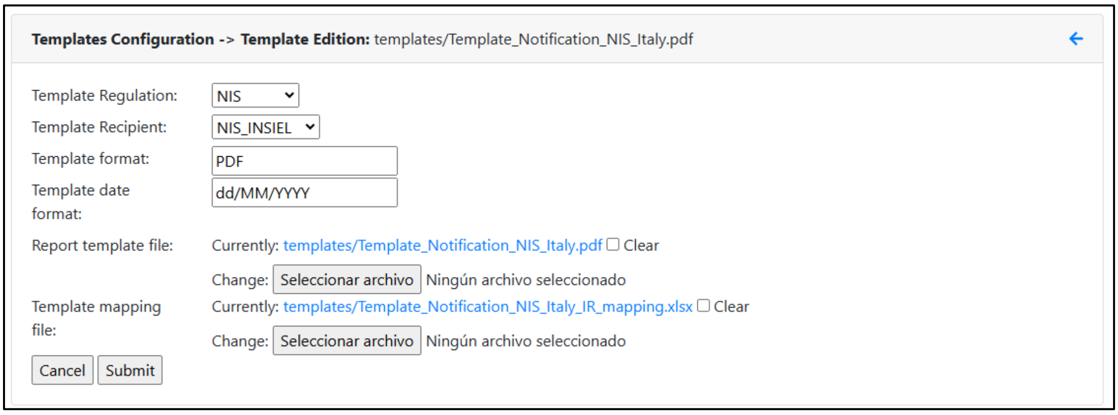

Templates Configuration -> Template Edition: templates/Template_Notification_NIS_Italy.pdf

Template Regulation:	NIS
Template Recipient:	NIS_INSIEL
Template format:	PDF
Template date format:	dd/MM/YYYY
Report template file:	Currently: templates/Template_Notification_NIS_Italy.pdf □ Clear Change: Seleccionar archivo Ningún archivo seleccionado
Template mapping file:	Currently: templates/Template_Notification_NIS_Italy_IR_mapping.xlsx □ Clear Change: Seleccionar archivo Ningún archivo seleccionado

Figure 105. Template Edition Form.

- **Reported Authorities:** In this menu, as shown in Figure 106, it is necessary to associate each reported authority with the regulations or specifications that require a mandatory report to be sent to it. This information will be used in the Managerial Judgement process to suggest the reported authorities that need to be notified when the criteria and thresholds associated with those specifications are matched.

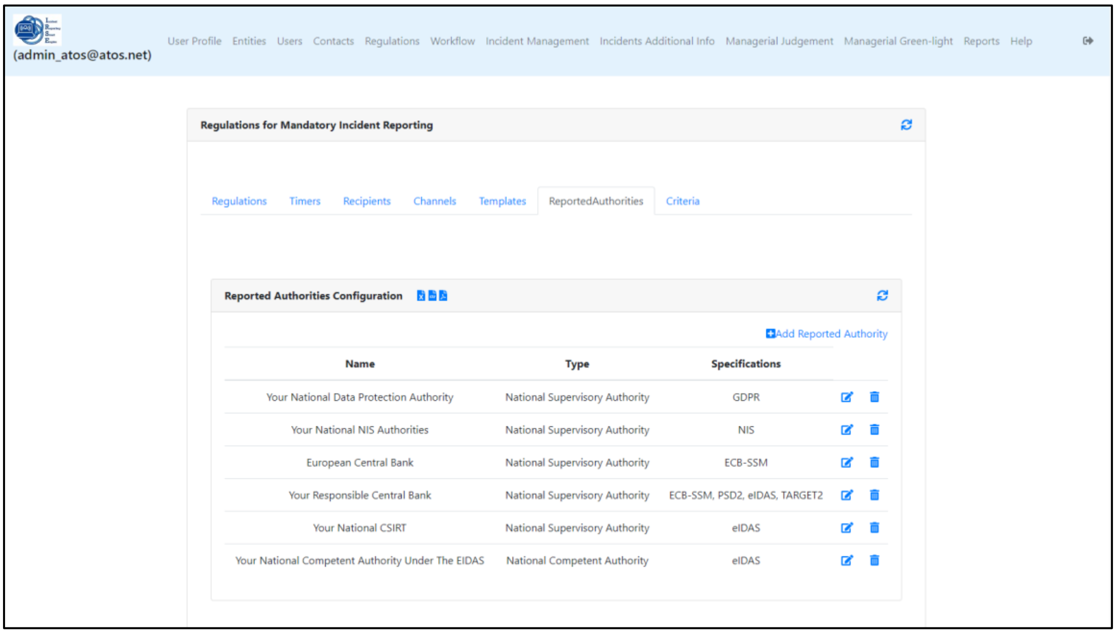

Regulations for Mandatory Incident Reporting

Reported Authorities Configuration

Name	Type	Specifications
Your National Data Protection Authority	National Supervisory Authority	GDPR
Your National NIS Authorities	National Supervisory Authority	NIS
European Central Bank	National Supervisory Authority	ECB-SSM
Your Responsible Central Bank	National Supervisory Authority	ECB-SSM, PSD2, eIDAS, TARGET2
Your National CSIRT	National Supervisory Authority	eIDAS
Your National Competent Authority Under The EIDAS	National Competent Authority	eIDAS

Figure 106: Reported Authorities Configuration.

The Reported Authorities Editor, as shown in Figure 107, enables the creation or modification of the entries of Reported Authorities.

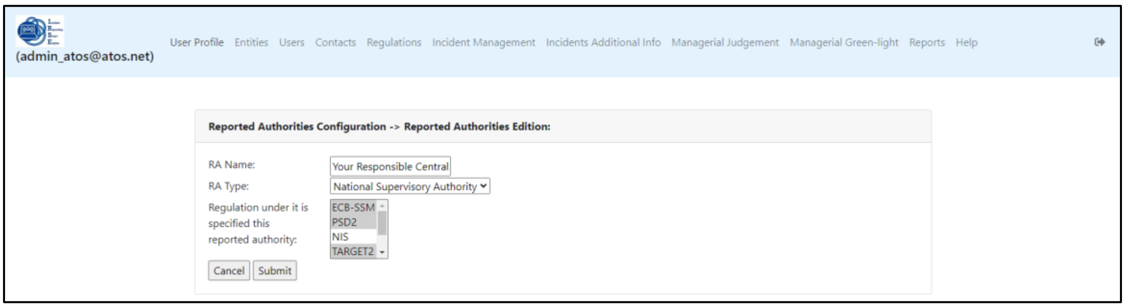

Reported Authorities Configuration -> Reported Authorities Edition:

RA Name:	Your Responsible Central
RA Type:	National Supervisory Authority
Regulation under it is specified this reported authority:	ECB-SSM PSD2 NIS TARGET2

Figure 107. Reported Authorities Editor.

- **Criteria:** Figure 108 shows the criteria supported by the platform for the event classification. The current version of the tool does not support customization of the criteria included in the regulations. The criteria included here are preloaded and just included for information purposes. However, they are not considered in real-time by the responder IR Event Classifier included in the demonstrator. Consequently, if any of them are removed or any more are added, the changes will not be considered by the classifier.

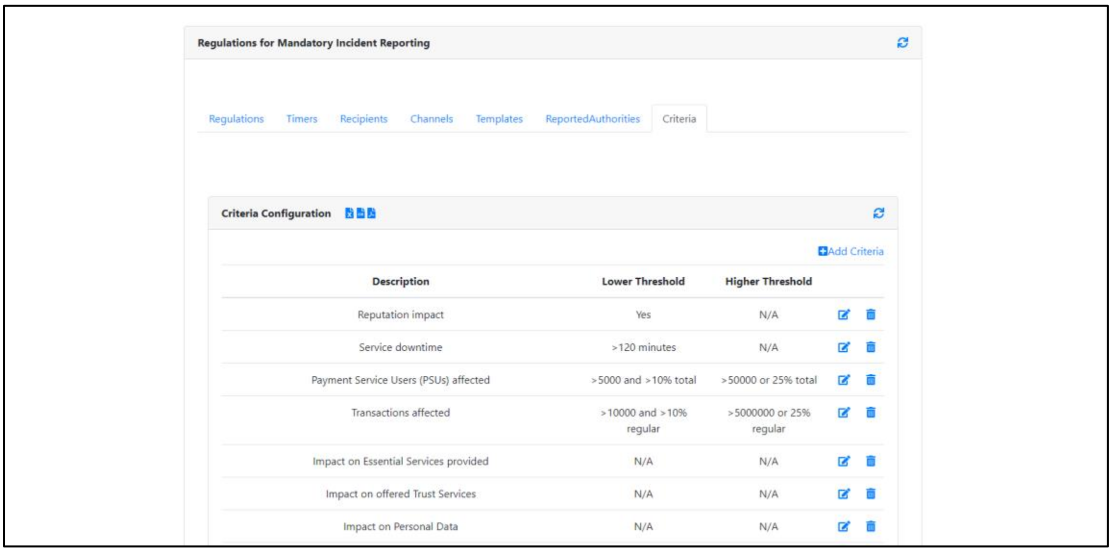

Regulations for Mandatory Incident Reporting

Criteria Configuration

Description	Lower Threshold	Higher Threshold
Reputation impact	Yes	N/A
Service downtime	>120 minutes	N/A
Payment Service Users (PSUs) affected	>5000 and >10% total	>50000 or 25% total
Transactions affected	>10000 and >10% regular	>5000000 or 25% regular
Impact on Essential Services provided	N/A	N/A
Impact on offered Trust Services	N/A	N/A
Impact on Personal Data	N/A	N/A

Figure 108. List of Criteria.

Figure 109 displays the *Criteria Edition* View, where the user can set the lower and higher thresholds.

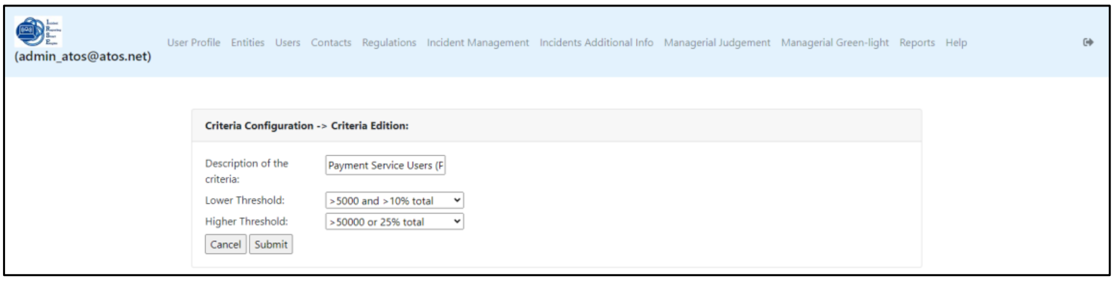

(admin_atos@atos.net)

Criteria Configuration -> Criteria Edition:

Description of the criteria: Payment Service Users (F

Lower Threshold: >5000 and >10% total

Higher Threshold: >50000 or 25% total

Figure 109. Criteria Edition.

More information about regulations can be found under the *Help* menu, as shown in Figure 110:

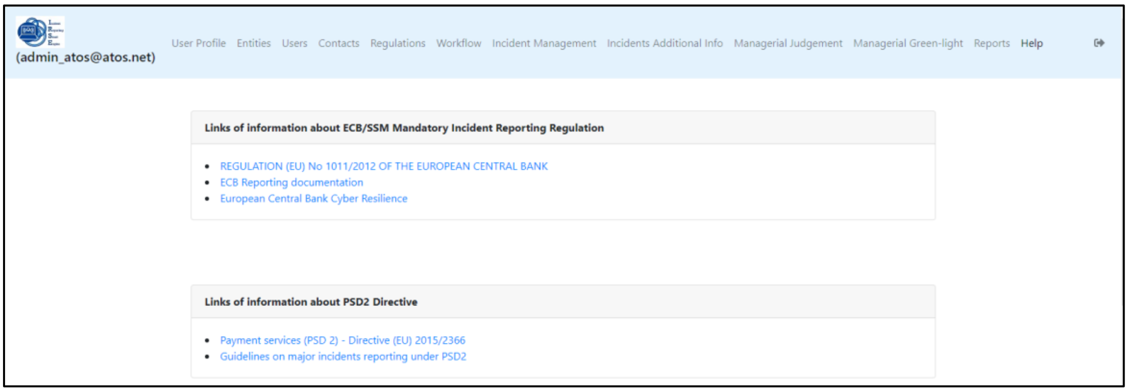

(admin_atos@atos.net)

Links of information about ECB/SSM Mandatory Incident Reporting Regulation

- REGULATION (EU) No 1011/2012 OF THE EUROPEAN CENTRAL BANK
- ECB Reporting documentation
- European Central Bank Cyber Resilience

Links of information about PSD2 Directive

- Payment services (PSD 2) - Directive (EU) 2015/2366
- Guidelines on major incidents reporting under PSD2

Figure 110. Help menu.

5.2 Register a new incident.

Once AIRE has been configured by the administrator, users can manage incidents and send reports. This section describes how a user can register a new incident in the tool.

The incidents are managed by the TheHive[7] tool. When the "Incident Management" is clicked, a pop-up window will be shown with the graphical interface provided by TheHive, as shown in Figure 111.

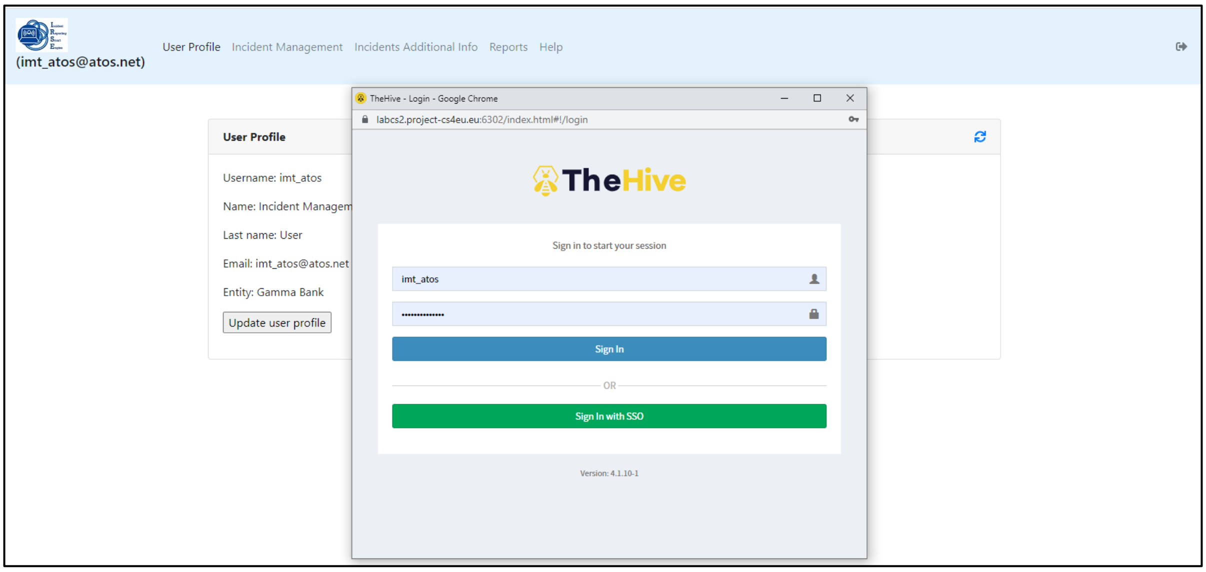

(imt_atos@atos.net)

User Profile | Incident Management | Incidents Additional Info | Reports | Help

User Profile

Username: imt_atos

Name: Incident Managem

Last name: User

Email: imt_atos@atos.net

Entity: Gamma Bank

Update user profile

TheHive - Login - Google Chrome

labcs2.project-cs4eu.eu:6302/index.html#!/login

TheHive

Sign in to start your session

imt_atos

••••••••••••••

Sign In

OR

Sign In with SSO

Version: 4.1.10-1

Figure 111. TheHive login interface.

After signing in, the TheHive graphical interface is embedded within the AIRE interface, as shown in Figure 112, where the user can perform several actions related to incident management.

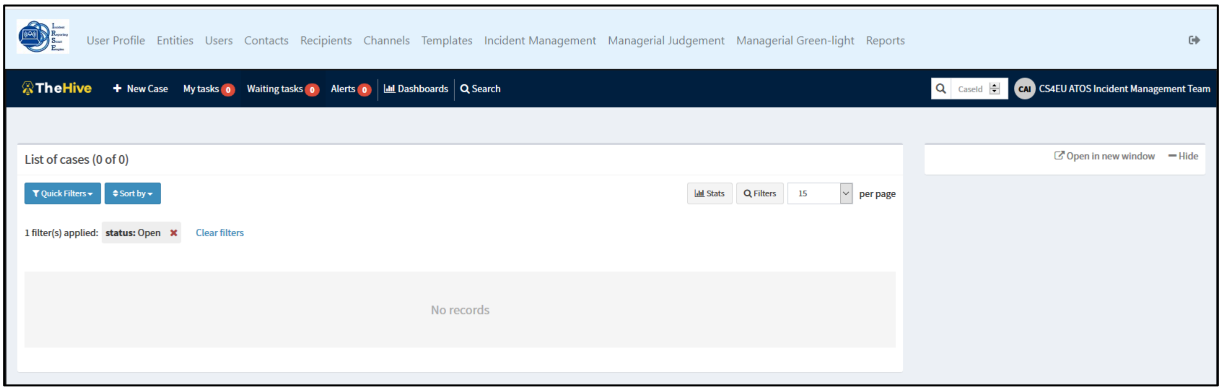

User Profile | Entities | Users | Contacts | Recipients | Channels | Templates | Incident Management | Managerial Judgement | Managerial Green-light | Reports

TheHive | + New Case | My tasks 0 | Waiting tasks 0 | Alerts 0 | Dashboards | Search | CaseId | CS4EU ATOS Incident Management Team

List of cases (0 of 0)

Quick Filters | Sort by | Stats | Filters | 15 | per page

1 filter(s) applied: status: Open ✖ Clear filters

No records

Open in new window | Hide

Figure 112: TheHive embedded in AIRE interface.

Click on **New Case** button to start the process of creating a case for an incident. This may be an empty case or use a predefined template, as shown in Figure 113. For the correct working of the CPR tool, the template “Incident Report” must be selected since it includes the fields required for the mandatory incident reporting, as shown in Figure 114.

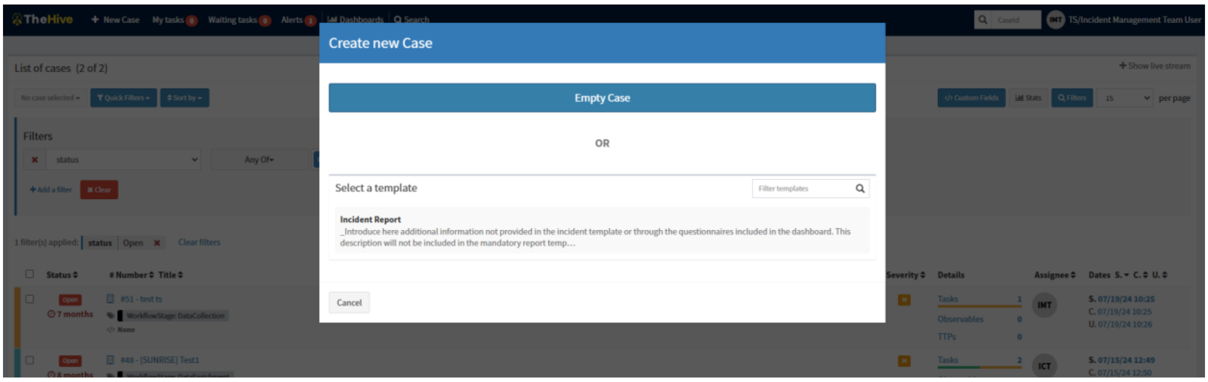

Incident Report
_Introduce here additional information not provided in the incident template or through the questionnaires included in the dashboard. This description will not be included in the mandatory report temp…

Figure 113: Template selection for a new Case.

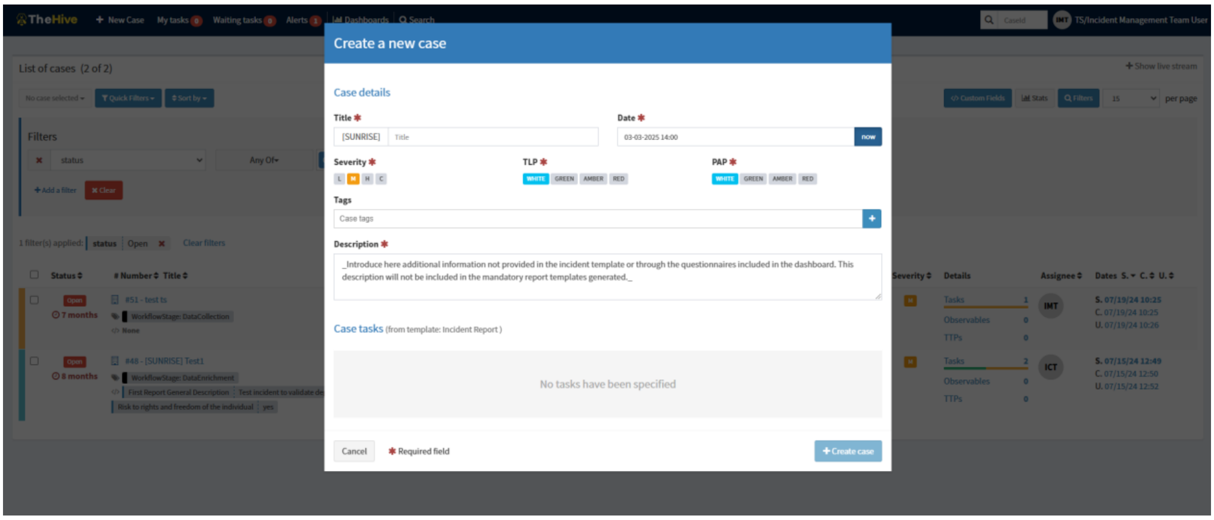

Title * [SUNRISE] Title

Date * 03-03-2025 14:00

Severity * M

TLP * WHITE

PAP * WHITE

Tags: Case tags

Description * _Introduce here additional information not provided in the incident template or through the questionnaires included in the dashboard. This description will not be included in the mandatory report templates generated._

Case tasks (from template: Incident Report)

No tasks have been specified

* Required field

Figure 114: New case form with Incident Report template.

This action creates an empty case in TheHive, as shown in Figure 115:

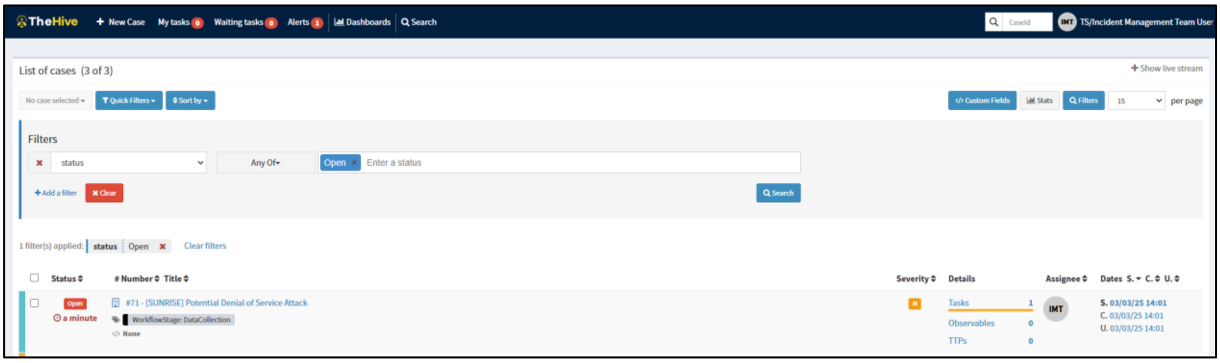

List of cases (3 of 3)

Filters: status | Any Of | Open | Enter a status

1 filter(s) applied: status | Open

Status	# Number	Title	Severity	Details	Assignee	Dates S. C. U.
Open a minute	#71 - [SUNRISE] Potential Denial of Service Attack WorkflowStage: DataCollection None		M	Tasks 1 Observables 0 TTPs 0	IMT	S. 03/03/25 14:01 C. 03/03/25 14:01 U. 03/03/25 14:01

Figure 115. New incident in the list of cases.

Figure 116 displays the details of the case with all the default fields to gather information related to the incident that will be used for incident reporting. Please note that the email included in the field “**Contact User**” will be the one used for notification purposes e.g. in case of delays in the deadline in the reporting process.

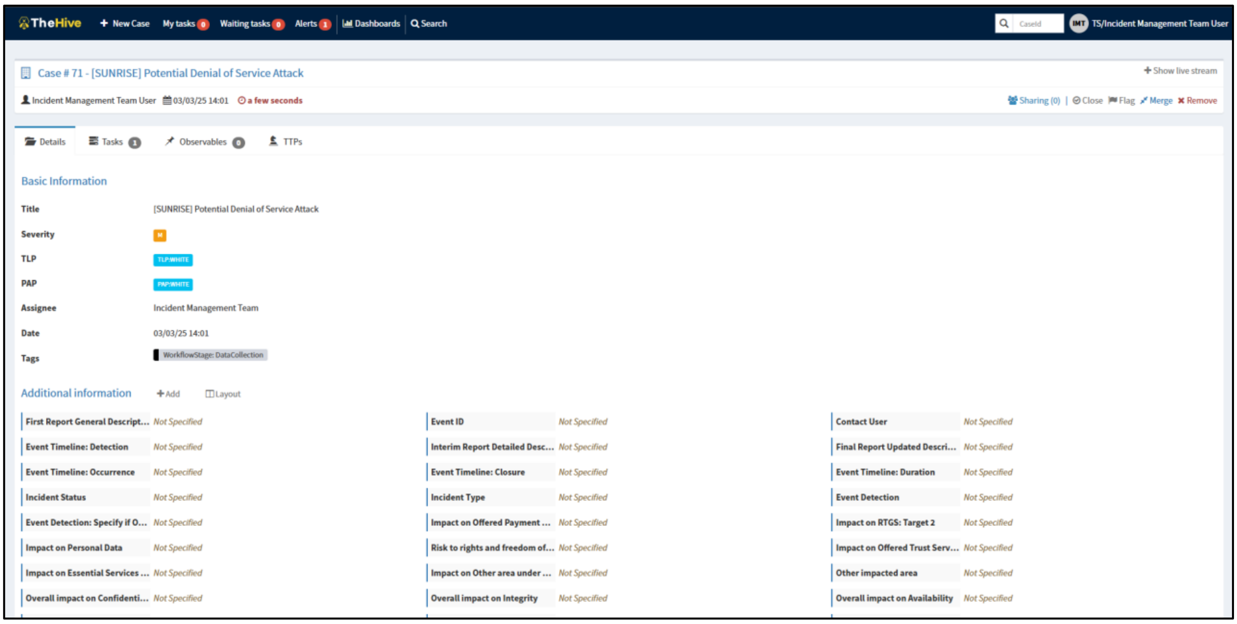

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 a few seconds

Basic Information

Title	[SUNRISE] Potential Denial of Service Attack
Severity	M
TLP	TLP:WHITE
PAP	PAP:WHITE
Assignee	Incident Management Team
Date	03/03/25 14:01
Tags	WorkflowStage: DataCollection

Additional information

First Report General Descript…	Not Specified	Event ID	Not Specified	Contact User	Not Specified
Event Timeline: Detection	Not Specified	Interim Report Detailed Desc…	Not Specified	Final Report Updated Descri…	Not Specified
Event Timeline: Occurrence	Not Specified	Event Timeline: Closure	Not Specified	Event Timeline: Duration	Not Specified
Incident Status	Not Specified	Incident Type	Not Specified	Event Detection	Not Specified
Event Detection: Specify if O…	Not Specified	Impact on Offered Payment …	Not Specified	Impact on RTGS: Target 2	Not Specified
Impact on Personal Data	Not Specified	Risk to rights and freedom of…	Not Specified	Impact on Offered Trust Serv…	Not Specified
Impact on Essential Services …	Not Specified	Impact on Other area under …	Not Specified	Other impacted area	Not Specified
Overall impact on Confidenti…	Not Specified	Overall impact on Integrity	Not Specified	Overall impact on Availability	Not Specified

Figure 116. Details of a case.

In the *Tasks* tab there are the pending tasks for the case, as shown in Figure 117. The *Data Collection* task is the first task automatically created by AIRE when a new incident is registered and assigned to the *IMT (Incident Management Team)* group. In the context of SUNRISE Critical Infrastructures, this group corresponds with the users belonging to the **CI Incident Response and security operation team** (see end-users roles in section 5.2).

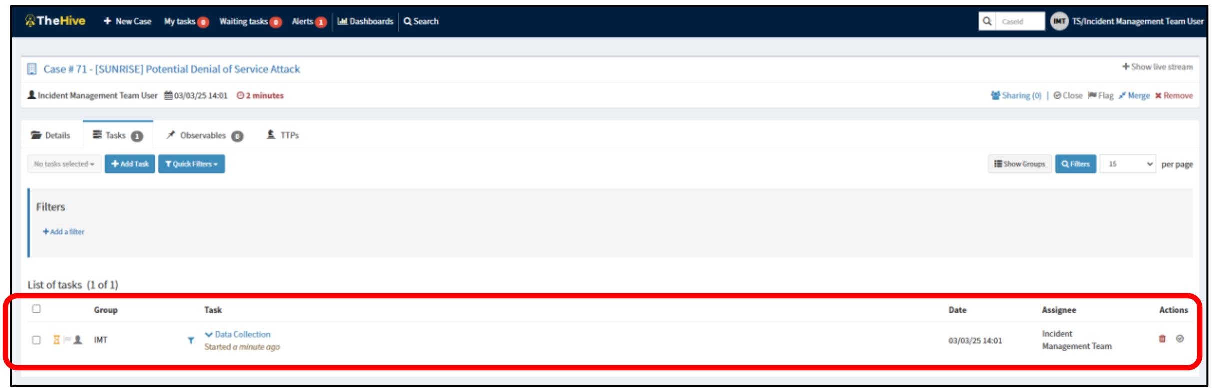

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 2 minutes

List of tasks (1 of 1)

□	Group	Task	Date	Assignee	Actions
□	IMT	Data Collection Started *a minute ago*	03/03/25 14:01	Incident Management Team	

Figure 117: Tasks associated with a case.

In the tab *Reports* of the AIRE graphical interface, as shown in Figure 118, a new report will be open:

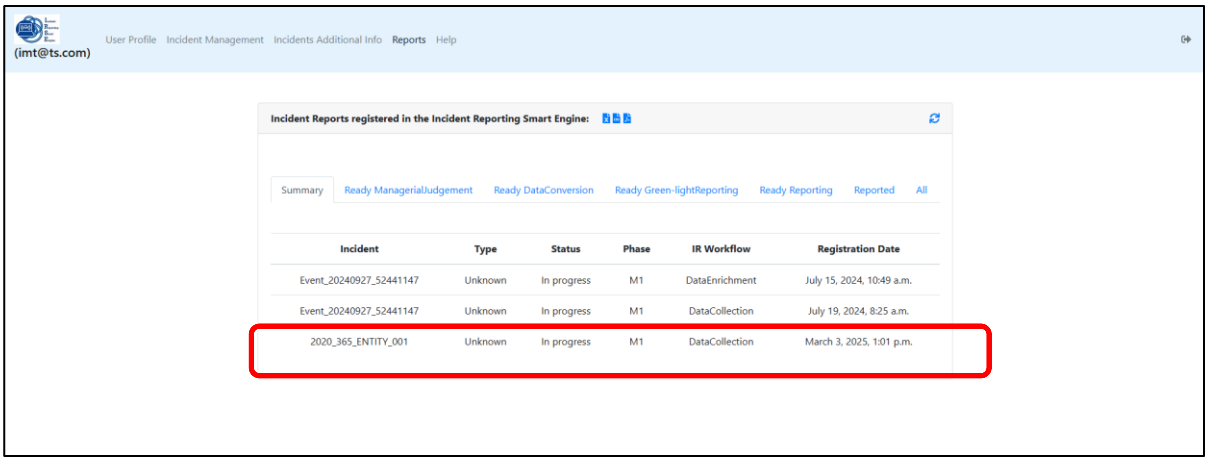

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Type	Status	Phase	IR Workflow	Registration Date
Event_20240927_52441147	Unknown	In progress	M1	DataEnrichment	July 15, 2024, 10:49 a.m.
Event_20240927_52441147	Unknown	In progress	M1	DataCollection	July 19, 2024, 8:25 a.m.
2020_365_ENTITY_001	Unknown	In progress	M1	DataCollection	March 3, 2025, 1:01 p.m.

Figure 118: List of Reports

When a new incident is registered, an email notification is automatically sent to the service beneficiaries (as initially configured as show in Figure 94) in compliance with NIS2 Directive, as shown in Figure 119.

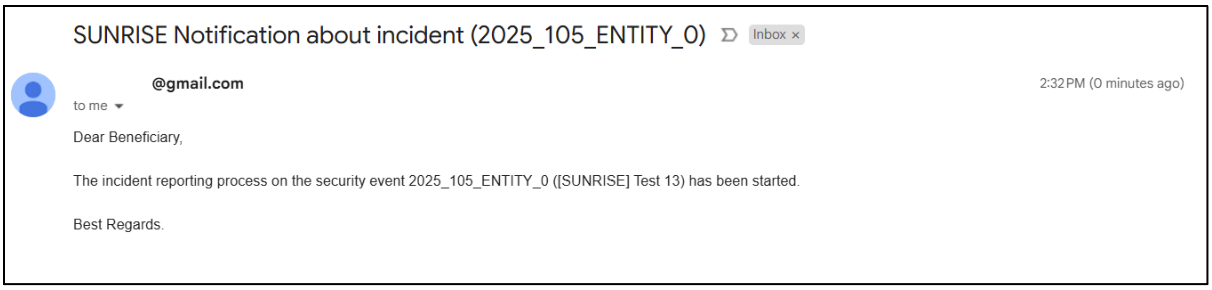

SUNRISE Notification about incident (2025_105_ENTITY_0)

@gmail.com

to me

2:32 PM (0 minutes ago)

Dear Beneficiary,

The incident reporting process on the security event 2025_105_ENTITY_0 ([SUNRISE] Test 13) has been started.

Best Regards.

Figure 119: Notification sent to service beneficiaries

NOTE: When a *New Case* is registered, since no information has been provided yet, the event ID will be assigned by default with the format <yyyy>_<dayoftheyear>_<ENTITY>_<caseId> (e.g. “2025_105_CAFC_12”). Once the information is included through TheHive and the Event ID is changed, it will be also reflected in the dashboard. See Figure 120 below for an example of this:

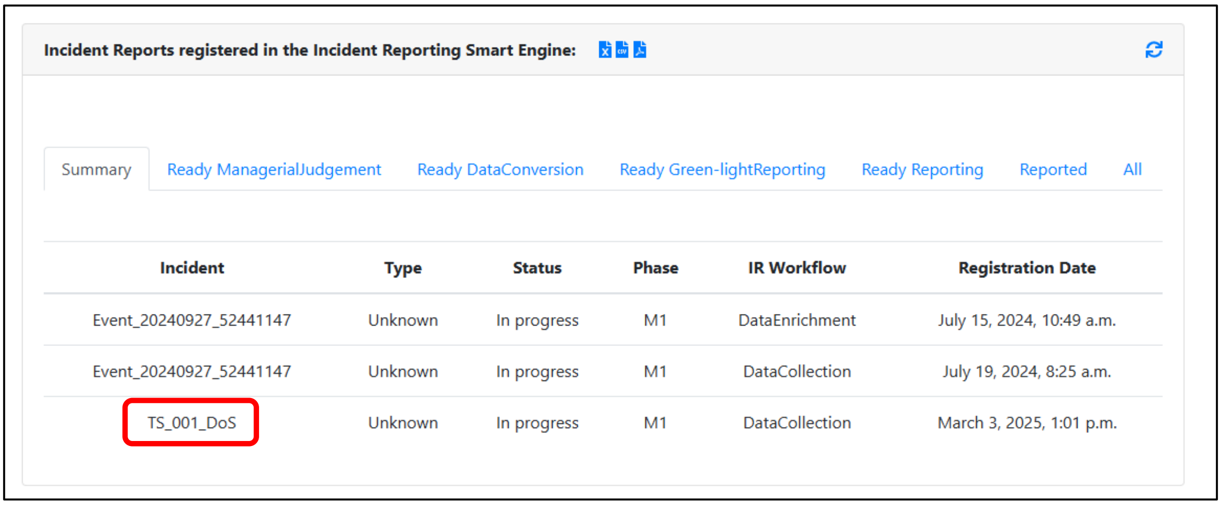

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Type	Status	Phase	IR Workflow	Registration Date
Event_20240927_52441147	Unknown	In progress	M1	DataEnrichment	July 15, 2024, 10:49 a.m.
Event_20240927_52441147	Unknown	In progress	M1	DataCollection	July 19, 2024, 8:25 a.m.
TS_001_DoS	Unknown	In progress	M1	DataCollection	March 3, 2025, 1:01 p.m.

Figure 120: Report with incident EventID assigned

5.3 Work on assigned tasks.

During the incident management and mandatory incident reporting process for compliance with the applicable regulatory frameworks, there are different steps that need to be followed by the Critical Infrastructures. Each of them is associated in the CPR tool to a different task that will be created by AIRE in the open-source tool TheHive and that will be assigned to a different group of users, depending on the end-user role that must work on it.

As it has been shown, the first task automatically created and assigned to the **CI Incident Response and security operation team** is ***Data Collection***. In this task, the operators must include through the template presented in TheHive (as shown in Figure 116) and through the *Incident Additional Info* tab of the AIRE dashboard (as shown in Figure 123) the information available about the incident.

The **Task Actions** button in TheHive GUI allows you to close, resume, or delete the task of a case, as shown in Figure 121.

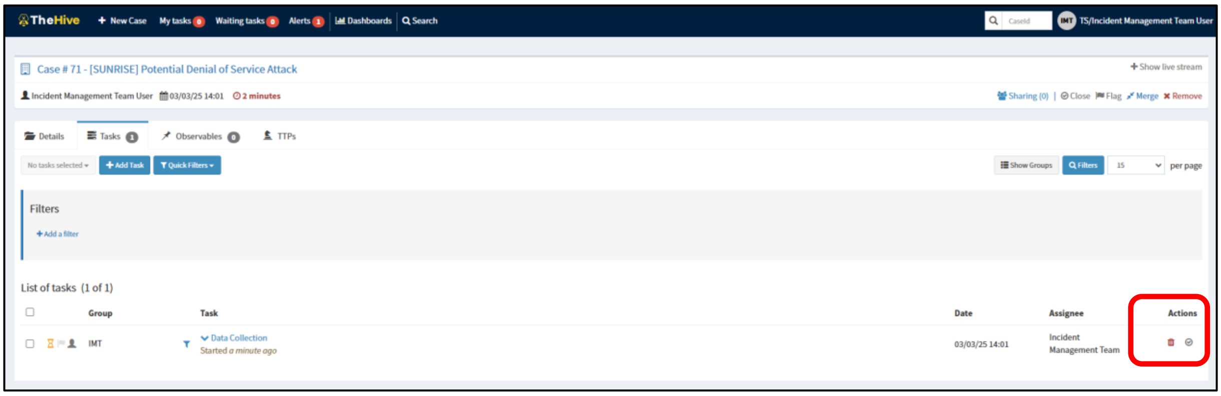

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 2 minutes

List of tasks (1 of 1)

□	Group	Task	Date	Assignee	Actions
□	IMT	Data Collection Started *a minute ago*	03/03/25 14:01	Incident Management Team	

Figure 121: Task actions of a case.

5.4 I.V.IV Close a task and progress in the incident reporting process.

Once a task has been completed by the assigned team, it must be closed in TheHive (e.g. once information about the incident about a new incident has been included by the operators, the Data Collection task can be closed (see Figure 122)).

When a task is closed, it is automatically created the following one in the mandatory incident reporting process and assigned to a new group (end-user role). For example, when the Data Collection task is closed, a new task ***Data Enrichment*** is created, and the workflow of the associated report changes to *Enrichment.*

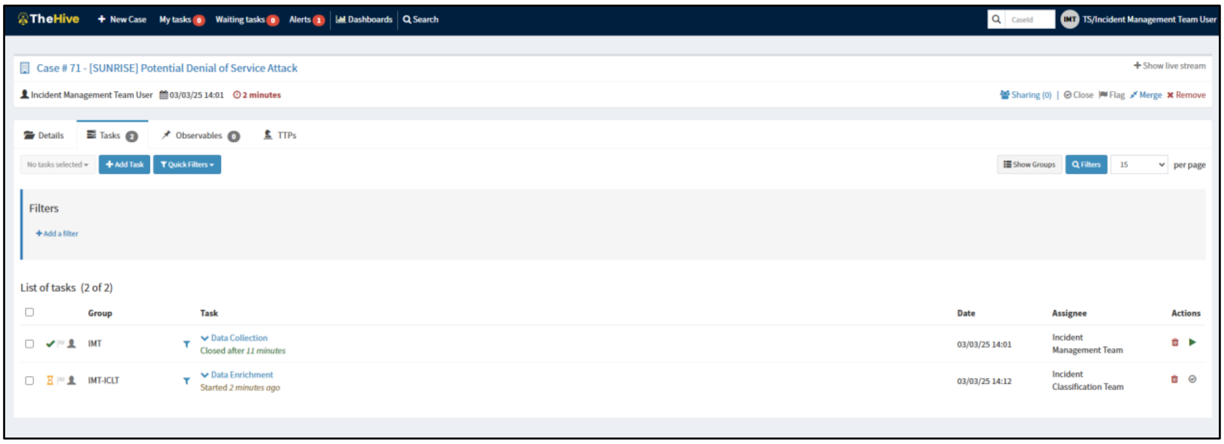

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 2 minutes

List of tasks (2 of 2)

□	Group	Task	Date	Assignee	Actions
□	IMT	Data Collection Closed after *11 minutes*	03/03/25 14:01	Incident Management Team	
□	IMT-ICLT	Data Enrichment Started *2 minutes ago*	03/03/25 14:12	Incident Classification Team	

Figure 122: New task Data Enrichment automatically created.

5.5 Register additional information about an incident.

Through the tab **Incident Additional Info** of AIRE dashboard, operators can include additional information about an incident. There are different tabs corresponding with different additional information:

- The *All Incidents* tab displays the list of all the incidents registered, as shown in Figure 123.

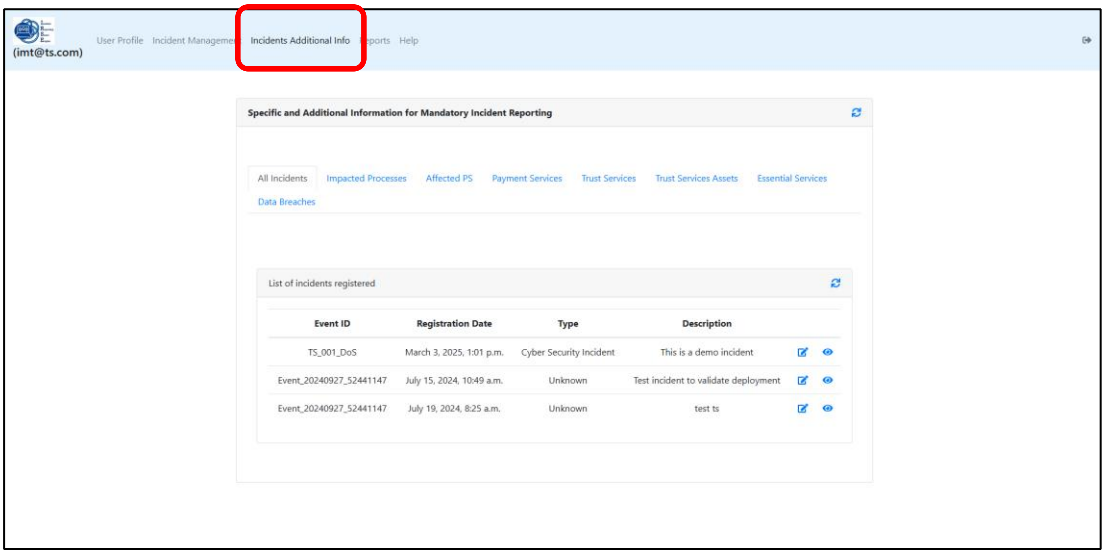

Specific and Additional Information for Mandatory Incident Reporting

List of incidents registered

Event ID	Registration Date	Type	Description
TS_001_DoS	March 3, 2025, 1:01 p.m.	Cyber Security Incident	This is a demo incident
Event_20240927_52441147	July 15, 2024, 10:49 a.m.	Unknown	Test incident to validate deployment
Event_20240927_52441147	July 19, 2024, 8:25 a.m.	Unknown	test ts

Figure 123: List of incidents registered in “Incidents Additional Info”.

Through this dashboard, a user can enrich the data by editing the information about a registered incident, as shown in Figure 124.

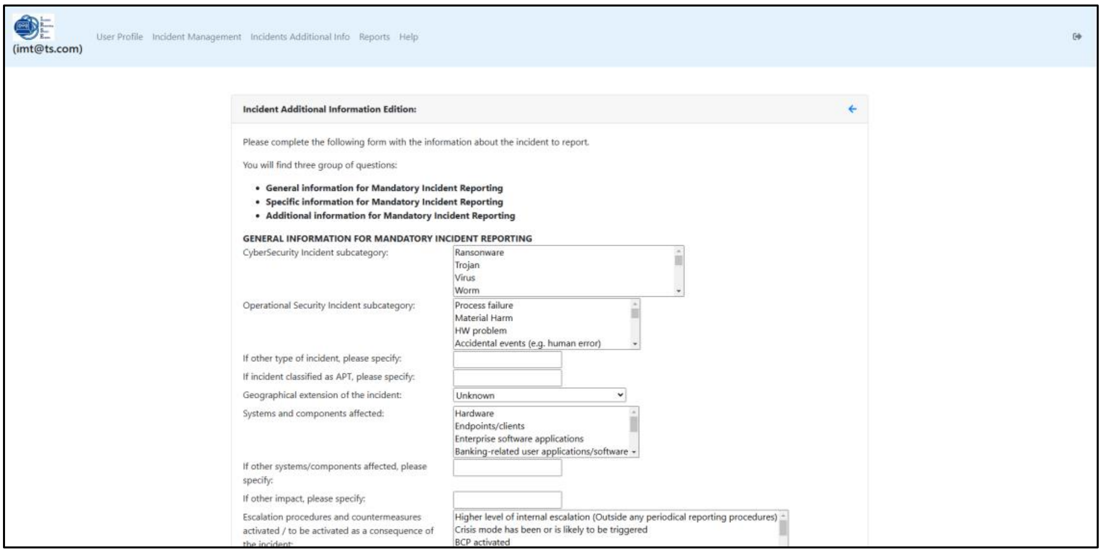

Incident Additional Information Edition:

Please complete the following form with the information about the incident to report.

You will find three group of questions:

- **General information for Mandatory Incident Reporting**
- **Specific information for Mandatory Incident Reporting**
- **Additional information for Mandatory Incident Reporting**

GENERAL INFORMATION FOR MANDATORY INCIDENT REPORTING

CyberSecurity Incident subcategory:	Ransonware Trojan Virus Worm
Operational Security Incident subcategory:	Process failure Material Harm HW problem Accidental events (e.g. human error)
If other type of incident, please specify:	
If incident classified as APT, please specify:	
Geographical extension of the incident:	Unknown
Systems and components affected:	Hardware Endpoints/clients Enterprise software applications Banking-related user applications/software
If other systems/components affected, please specify:	
If other impact, please specify:	
Escalation procedures and countermeasures activated / to be activated as a consequence of the incident:	Higher level of internal escalation (Outside any periodical reporting procedures) Crisis mode has been or is likely to be triggered BCP activated

Figure 124. Incident Additional Information Edition.

The *eye* icon in Figure 123 shows the logs registered that are related to the security event lifecycle, as shown in Figure 125:

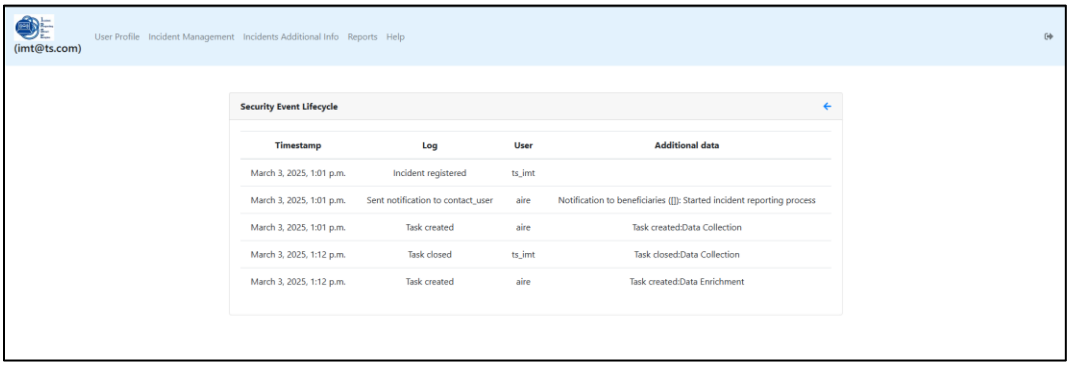

Security Event Lifecycle

Timestamp	Log	User	Additional data
March 3, 2025, 1:01 p.m.	Incident registered	ts_imt	
March 3, 2025, 1:01 p.m.	Sent notification to contact_user	aire	Notification to beneficiaries ([]): Started incident reporting process
March 3, 2025, 1:01 p.m.	Task created	aire	Task created:Data Collection
March 3, 2025, 1:12 p.m.	Task closed	ts_imt	Task closed:Data Collection
March 3, 2025, 1:12 p.m.	Task created	aire	Task created:Data Enrichment

Figure 125. Security Event Lifecycle.

The other tabs in Figure 123 (the Incident Additional Info tab) enable the management of the elements that are affected by an incident, such as services, assets, processes, or data. These elements are:

- ***Essential Services***: When the organization is a provider of essential services, they will be defined in this menu. A name needs to be assigned so it can be assigned to the incident. Figure 127 displays the edition/creation form for this sub-menu.

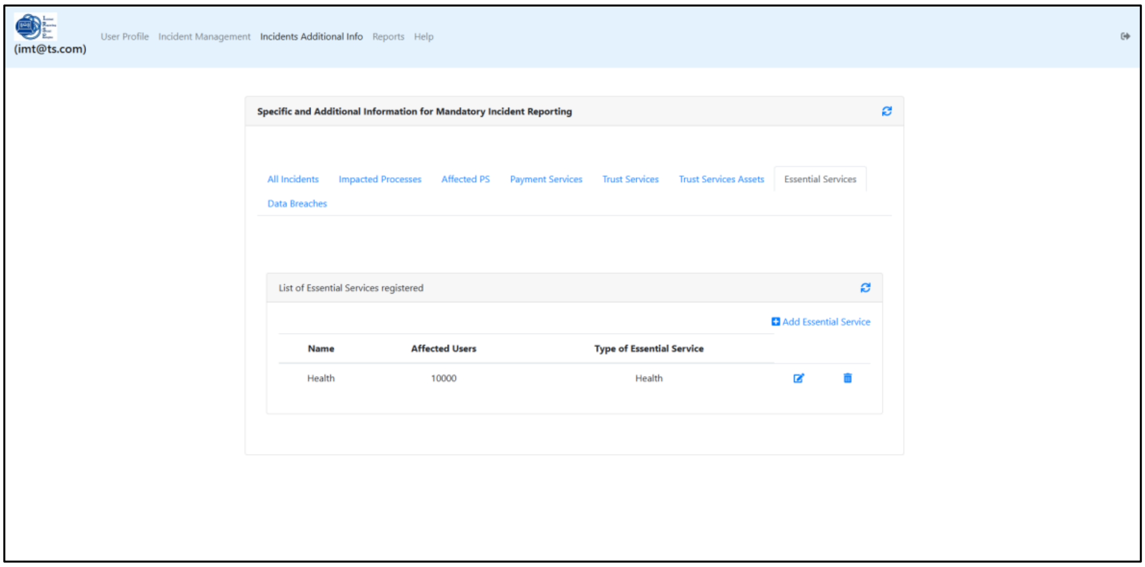

Specific and Additional Information for Mandatory Incident Reporting

List of Essential Services registered

Name	Affected Users	Type of Essential Service
Health	10000	Health

Figure 126. List of Essential Services.

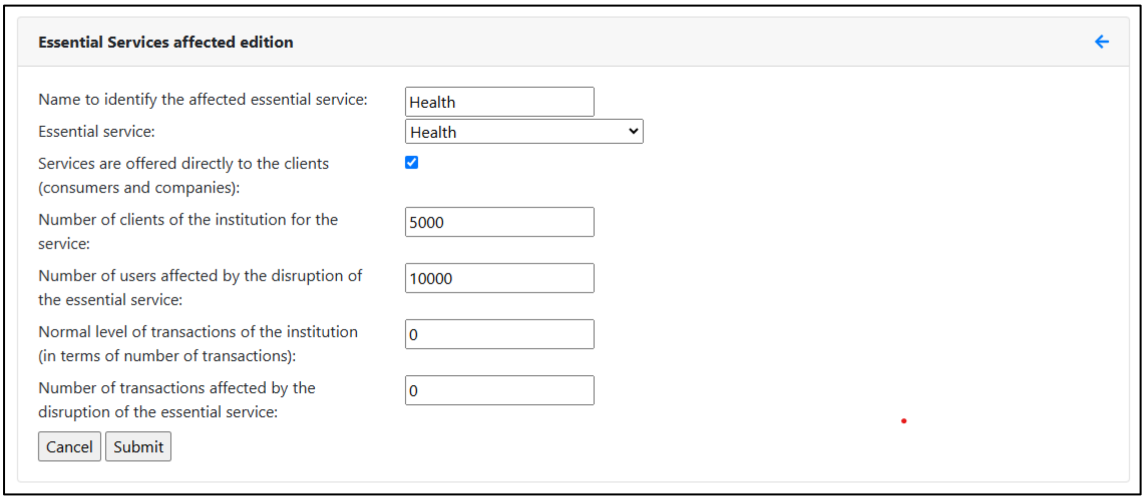

Essential Services affected edition

Name to identify the affected essential service:	Health
Essential service:	Health
Services are offered directly to the clients (consumers and companies):	☑
Number of clients of the institution for the service:	5000
Number of users affected by the disruption of the essential service:	10000
Normal level of transactions of the institution (in terms of number of transactions):	0
Number of transactions affected by the disruption of the essential service:	0

Figure 127. Essential Services edition.

- ***Trust Services Assets***: These are external services that support the trust protocols for Critical Infrastructures. Figure 128 displays the list of Trust Services Assets registered. Figure 129 shows the edition form of one Trust Service Asset.

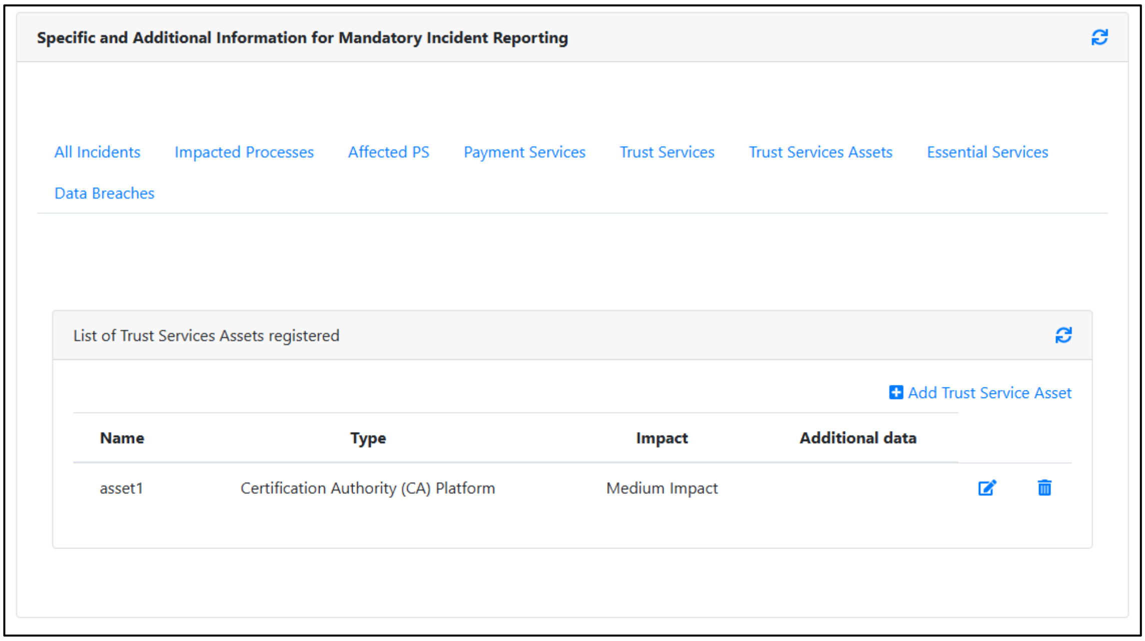

Specific and Additional Information for Mandatory Incident Reporting

List of Trust Services Assets registered

Name	Type	Impact	Additional data
asset1	Certification Authority (CA) Platform	Medium Impact	

Figure 128. List of Trust Services Assets registered.

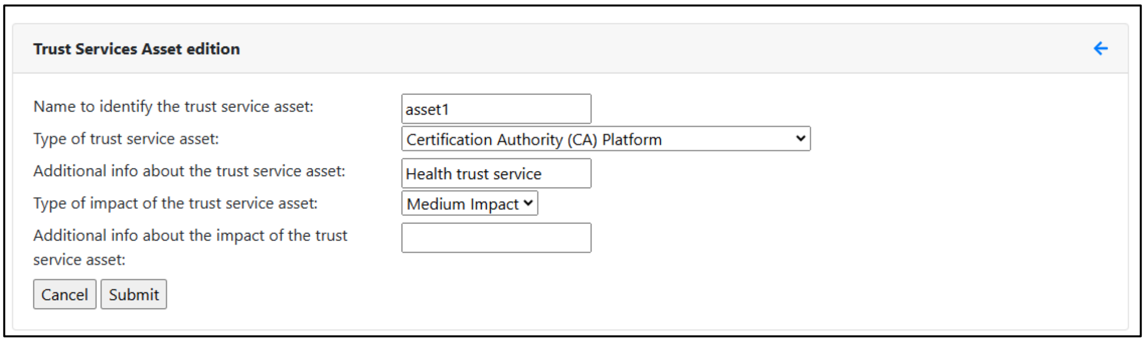

Trust Services Asset edition

Name to identify the trust service asset:	asset1
Type of trust service asset:	Certification Authority (CA) Platform
Additional info about the trust service asset:	Health trust service
Type of impact of the trust service asset:	Medium Impact
Additional info about the impact of the trust service asset:	

Figure 129. Trust Services Asset edition.

- ***Trust Services***: These are internal services that support the trust protocols for Critical Infrastructures. Figure 130 shows the list of Trust Services registered. Figure 131 shows the edition form of one Trust Service Asset.

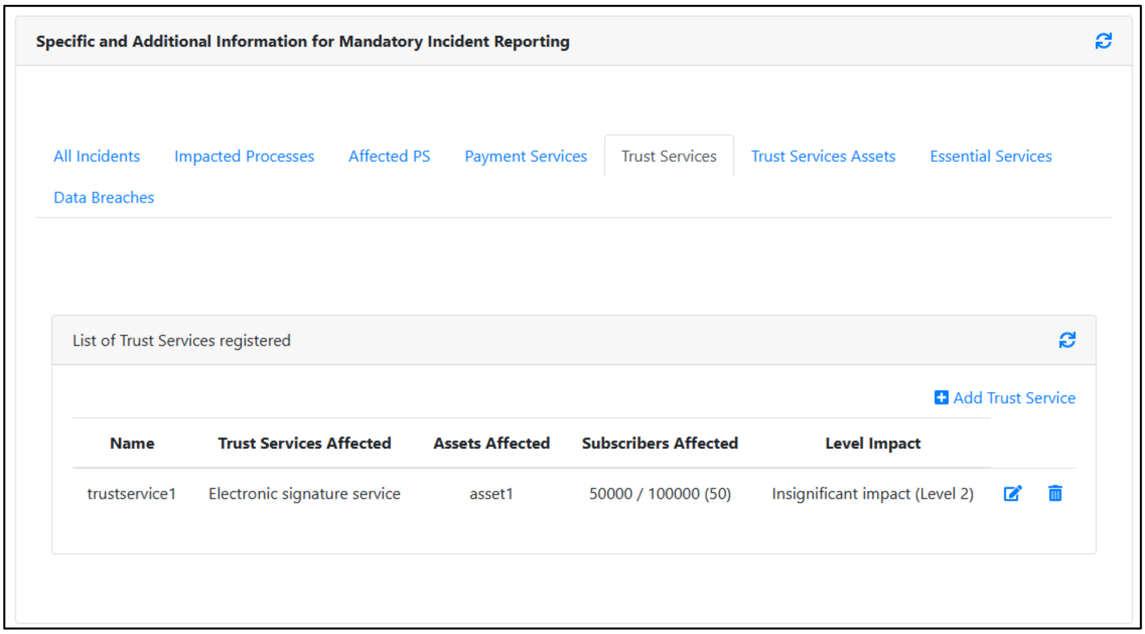

Specific and Additional Information for Mandatory Incident Reporting

List of Trust Services registered

Name	Trust Services Affected	Assets Affected	Subscribers Affected	Level Impact
trustservice1	Electronic signature service	asset1	50000 / 100000 (50)	Insignificant impact (Level 2)

Figure 130. List of Trust Services registered.

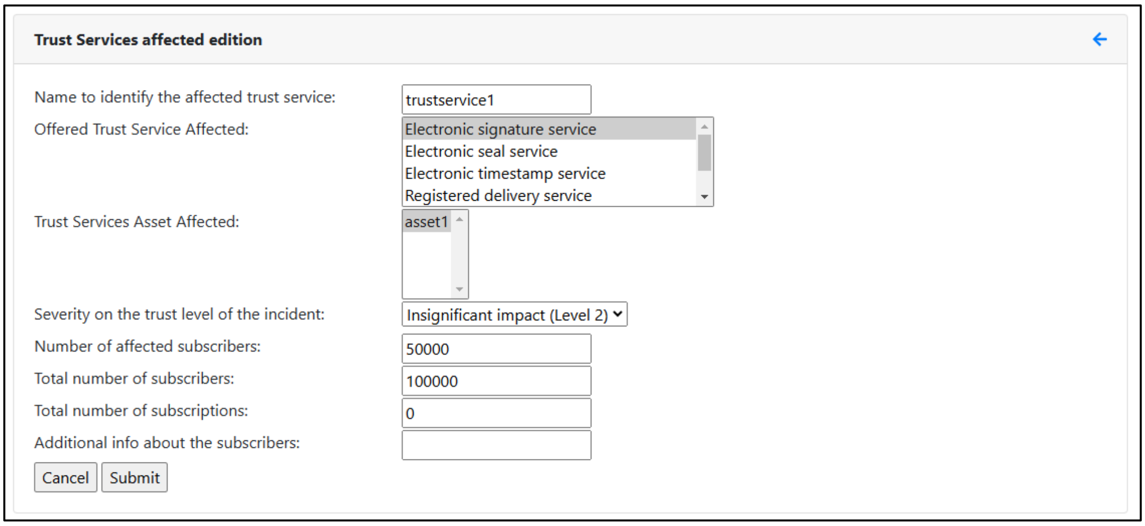

Trust Services affected edition

Name to identify the affected trust service:	trustservice1
Offered Trust Service Affected:	Electronic signature service Electronic seal service Electronic timestamp service Registered delivery service
Trust Services Asset Affected:	asset1
Severity on the trust level of the incident:	Insignificant impact (Level 2)
Number of affected subscribers:	50000
Total number of subscribers:	100000
Total number of subscriptions:	0
Additional info about the subscribers:	

Figure 131. Trust Services affected edition.

- ***Impacted Processes***: Figure 132 shows the list of processes impacted by an attack, the impact, and the recovery time. Figure 133 shows the form to add or create the *Processes Affected*.

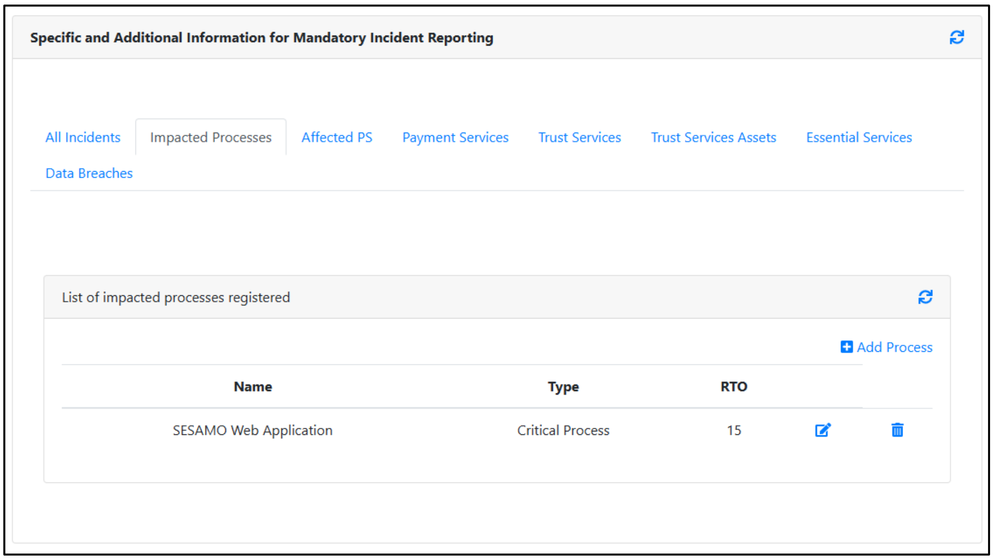

Specific and Additional Information for Mandatory Incident Reporting

List of impacted processes registered

Name	Type	RTO
SESAMO Web Application	Critical Process	15

Figure 132. List of Impacted Processes.

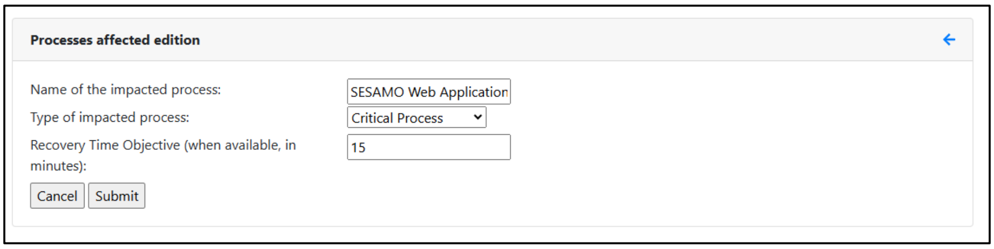

Processes affected edition

Name of the impacted process: SESAMO Web Application

Type of impacted process: Critical Process

Recovery Time Objective (when available, in minutes): 15

Figure 133. Processes Affected edition.

- Data Breaches: This lists the *Personal Data Breaches* with the type, data category and the number of affected subjects, as shown in Figure 134. Figure 135 displays the form to add or edit the information about a *Personal Data Breach*.

  NOTE: The association of a data breach with an incident is done by selecting from the list of active incidents in that menu.

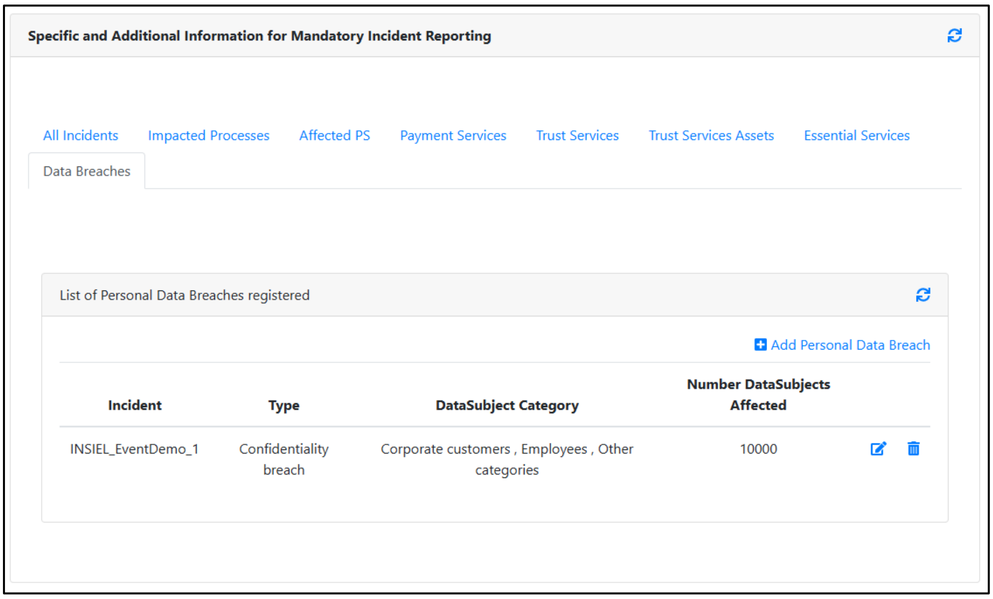

Specific and Additional Information for Mandatory Incident Reporting

List of Personal Data Breaches registered

Incident	Type	DataSubject Category	Number DataSubjects Affected
INSIEL_EventDemo_1	Confidentiality breach	Corporate customers , Employees , Other categories	10000

Figure 134. List of Personal Data Breaches registered.

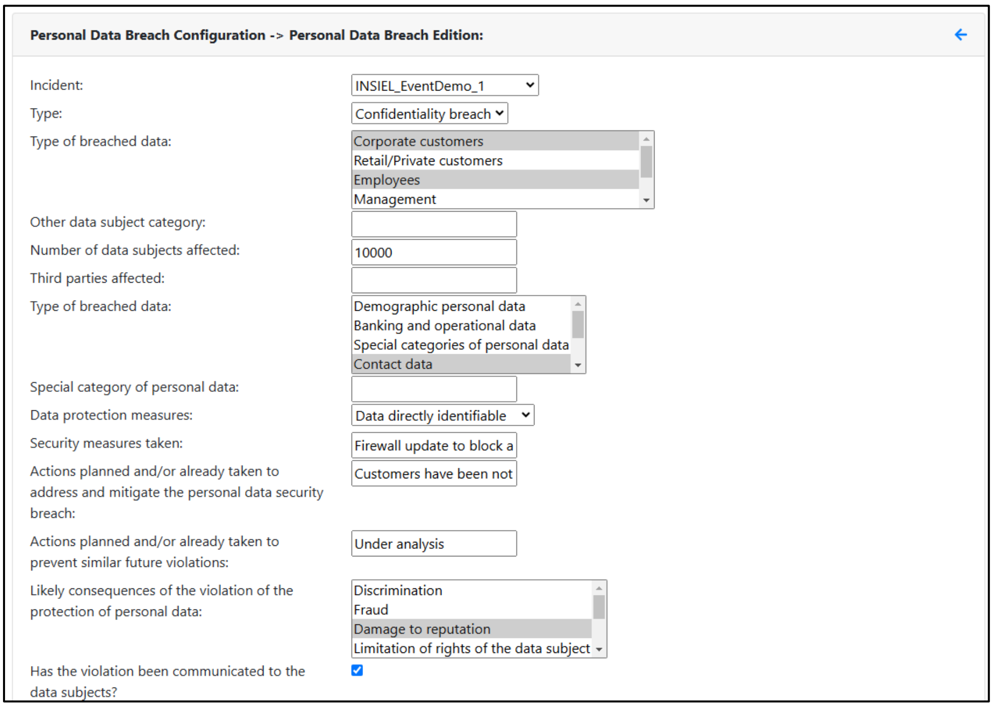

Personal Data Breach Configuration -> Personal Data Breach Edition:

Incident:	INSIEL_EventDemo_1
Type:	Confidentiality breach
Type of breached data:	Corporate customers Retail/Private customers Employees Management
Other data subject category:	
Number of data subjects affected:	10000
Third parties affected:	
Type of breached data:	Demographic personal data Banking and operational data Special categories of personal data Contact data
Special category of personal data:	
Data protection measures:	Data directly identifiable
Security measures taken:	Firewall update to block a
Actions planned and/or already taken to address and mitigate the personal data security breach:	Customers have been not
Actions planned and/or already taken to prevent similar future violations:	Under analysis
Likely consequences of the violation of the protection of personal data:	Discrimination Fraud Damage to reputation Limitation of rights of the data subject
Has the violation been communicated to the data subjects?	☑

Figure 135. Personal Data Breach Edition.

NOTE: The association of a data breach with an incident is done by selecting from the list of active incidents in that menu.

5.6 Register observables

Through the graphical interface provided by TheHive and integrated in AIRE dashboard, it is possible to include information about the incident as *Observables* and run Cortex analysers on them, as shown in Figure 136. This is done going in TheHive dashboard to **Observables** tab once an incident (*Case* in TheHive terminology) is opened:

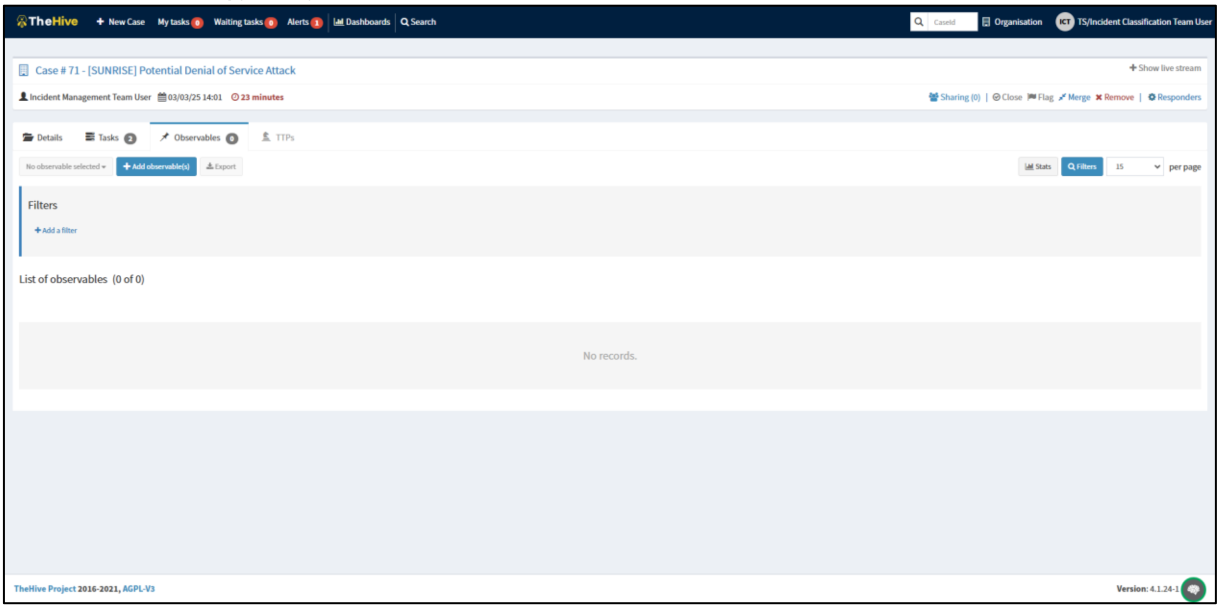

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 23 minutes

Filters

List of observables (0 of 0)

No records.

TheHive Project 2016-2021, AGPL-V3

Version: 4.1.24-1

Figure 136. List of observables.

The button *Add observable* opens a form to create new observables elements, as shown in Figure 137.

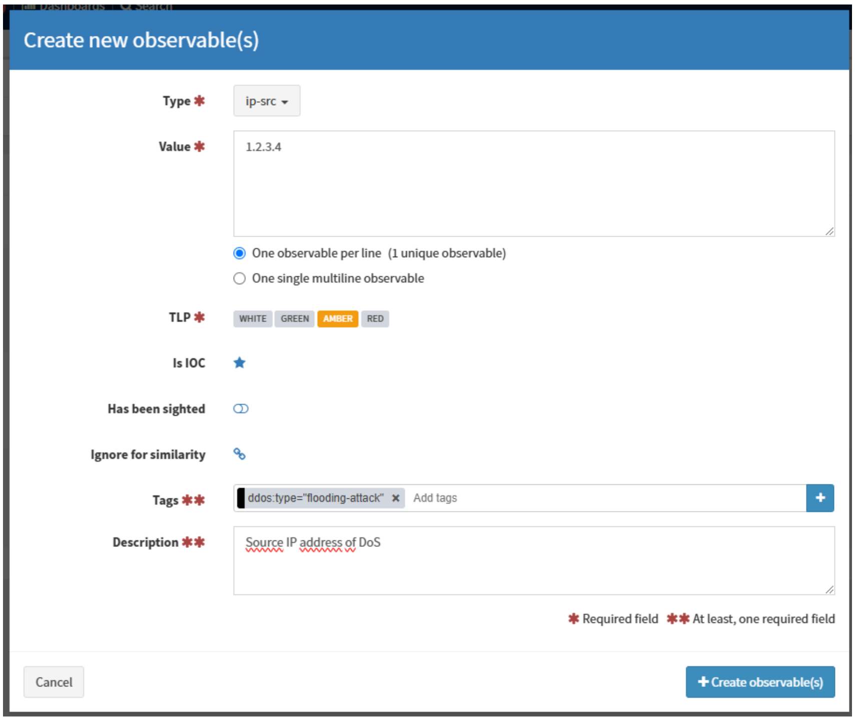

Create new observable(s)

Type ✱ ip-src

Value ✱ 1.2.3.4

● One observable per line (1 unique observable)

○ One single multiline observable

TLP ✱ WHITE GREEN AMBER RED

Is IOC ★

Has been sighted

Ignore for similarity

Tags ✱✱ ddos:type="flooding-attack" × Add tags

Description ✱✱ Source IP address of DoS

✱ Required field ✱✱ At least, one required field

Figure 137. Create new observable form.

After selecting some observables, the *Run analysers* button allows us to run different analysers, as shown in Figure 138. More information about Cortex Analyzers can be found at [4].

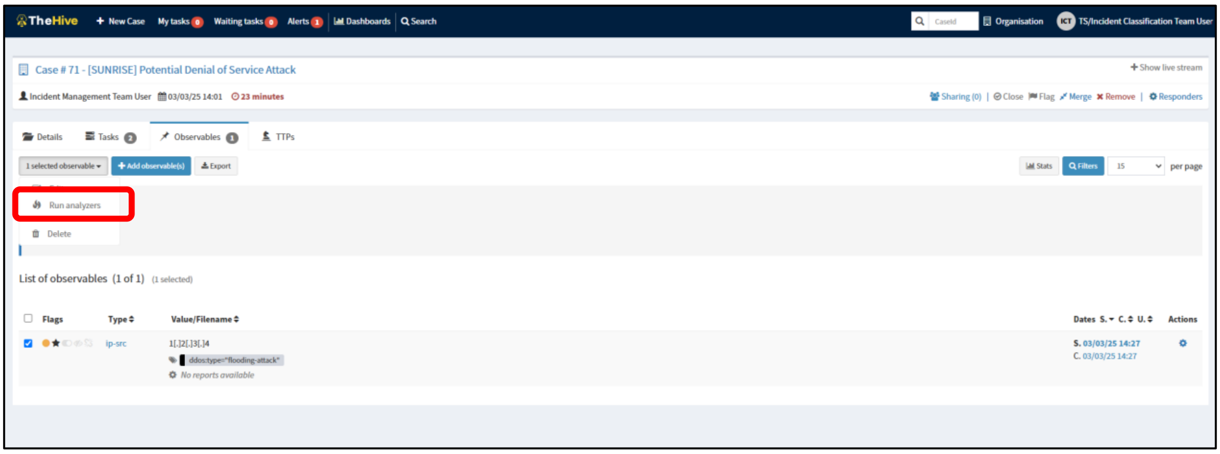

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 23 minutes

List of observables (1 of 1) (1 selected)

☐	Flags	Type	Value/Filename	Dates S. C. U.	Actions
☑	● ★	ip-src	1[.]2[.]3[.]4 ddos:type="flooding-attack" No reports available	S. 03/03/25 14:27 C. 03/03/25 14:27	⚙

Figure 138: Run analyzers option.

Observables can be also added to an incident automatically from the alerts received. These alerts in the CPR tool can come from the Anomaly Detection module or from the Threat Intelligence module. They will appear in TheHive dashboard under Alerts tab, such as the one shown in Figure 139.

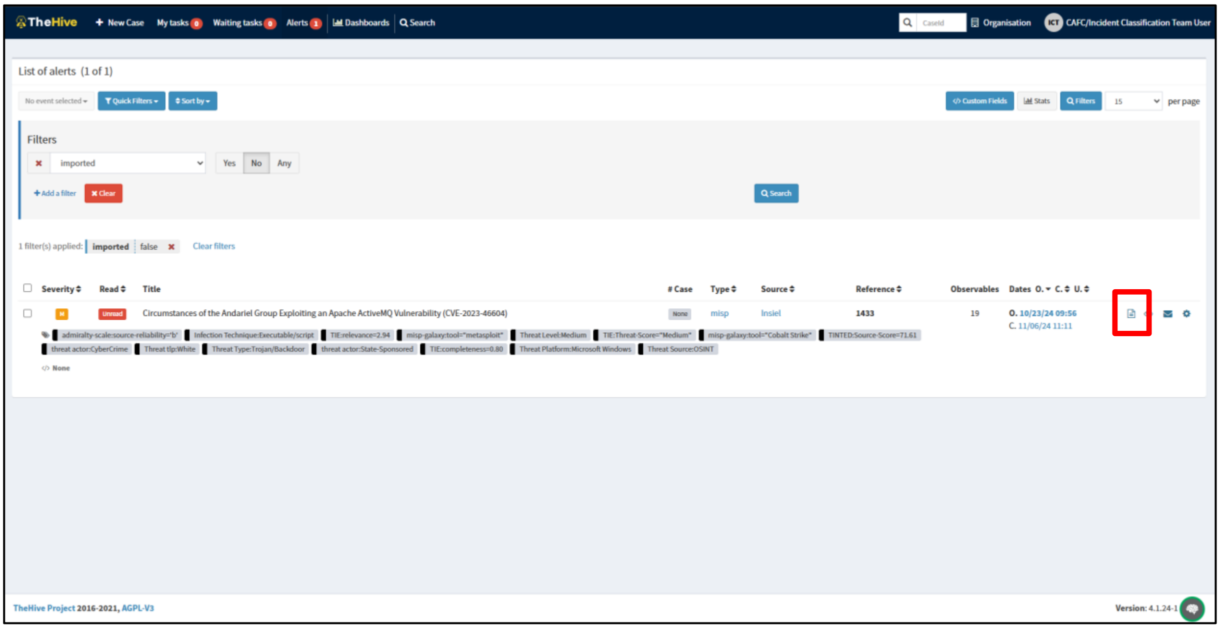

List of alerts (1 of 1)

Filters

imported: Yes [✓ No] Any

1 filter(s) applied: imported | false

	Severity	Read	Title	# Case	Type	Source	Reference	Observables	Dates O. C. U.
□	M	Unread	Circumstances of the Andariel Group Exploiting an Apache ActiveMQ Vulnerability (CVE-2023-46604)	None	misp	Insiel	1433	19	O. 10/23/24 09:56 C. 11/06/24 11:11

admiralty-scale:source-reliability='b' | Infection Technique:Executable/script | TIE:relevance=2.94 | misp-galaxy:tool="metasploit" | Threat Level:Medium | TIE:Threat-Score="Medium" | misp-galaxy:tool="Cobalt Strike" | TINTED:Source-Score=71.61 | threat actor:CyberCrime | Threat tlp:White | Threat Type:Trojan/Backdoor | threat actor:State-Sponsored | TIE:completeness=0.80 | Threat Platform:Microsoft Windows | Threat Source:OSINT

None

Figure 139: Alert received in the Incident Reporting module.

If the user clicks on the “Preview and Import” button (see in red in the previous Figure), it is opened a new window with all the information provided in this alert (in this case, all the attributes of the MISP event as shown in Figure 140).

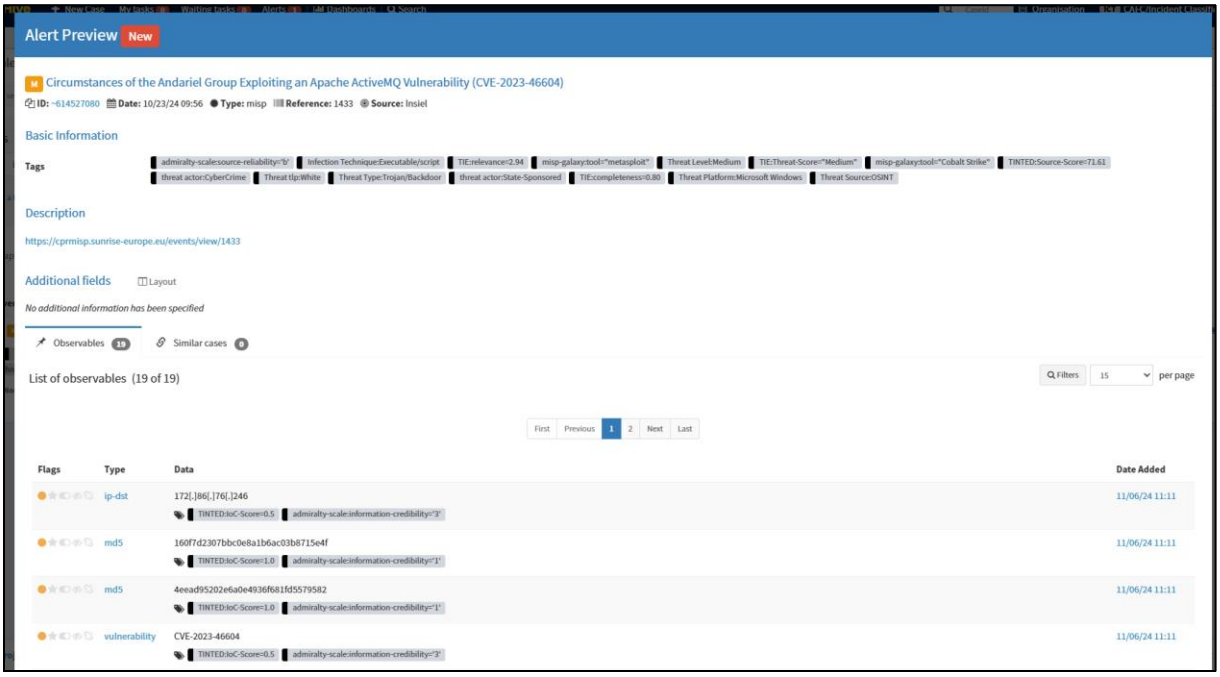

M Circumstances of the Andariel Group Exploiting an Apache ActiveMQ Vulnerability (CVE-2023-46604)

ID: ~614527080 Date: 10/23/24 09:56 Type: misp Reference: 1433 Source: Insiel

Basic Information

Tags: admiralty-scale:source-reliability='b' | Infection Technique:Executable/script | TIE:relevance=2.94 | misp-galaxy:tool="metasploit" | Threat Level:Medium | TIE:Threat-Score="Medium" | misp-galaxy:tool="Cobalt Strike" | TINTED:Source-Score=71.61 | threat actor:CyberCrime | Threat tlp:White | Threat Type:Trojan/Backdoor | threat actor:State-Sponsored | TIE:completeness=0.80 | Threat Platform:Microsoft Windows | Threat Source:OSINT

Description

https://cprmisp.sunrise-europe.eu/events/view/1433

Additional fields

No additional information has been specified

List of observables (19 of 19)

Flags	Type	Data	Date Added
●	ip-dst	172[.]86[.]76[.]246 TINTED:IoC-Score=0.5 | admiralty-scale:information-credibility='3'	11/06/24 11:11
●	md5	160f7d2307bbc0e8a1b6ac03b8715e4f TINTED:IoC-Score=1.0 | admiralty-scale:information-credibility='1'	11/06/24 11:11
●	md5	4eead95202e6a0e4936f681fd5579582 TINTED:IoC-Score=1.0 | admiralty-scale:information-credibility='1'	11/06/24 11:11
●	vulnerability	CVE-2023-46604 TINTED:IoC-Score=0.5 | admiralty-scale:information-credibility='3'	11/06/24 11:11

Figure 140: Import Observables from alert received.

They can be imported as observables in a new incident (in this case, a new incident reporting process would start as described in previous section to register a new incident) or merged into a case (if the user clicks on this button) to be added to an existing incident.

5.7 Vulnerability Management

If at any point during the data collection or data enrichment phases, it is identified that the incident is related to a new vulnerability, it can be indicated in the template with the associated field as shown in Figure 141.

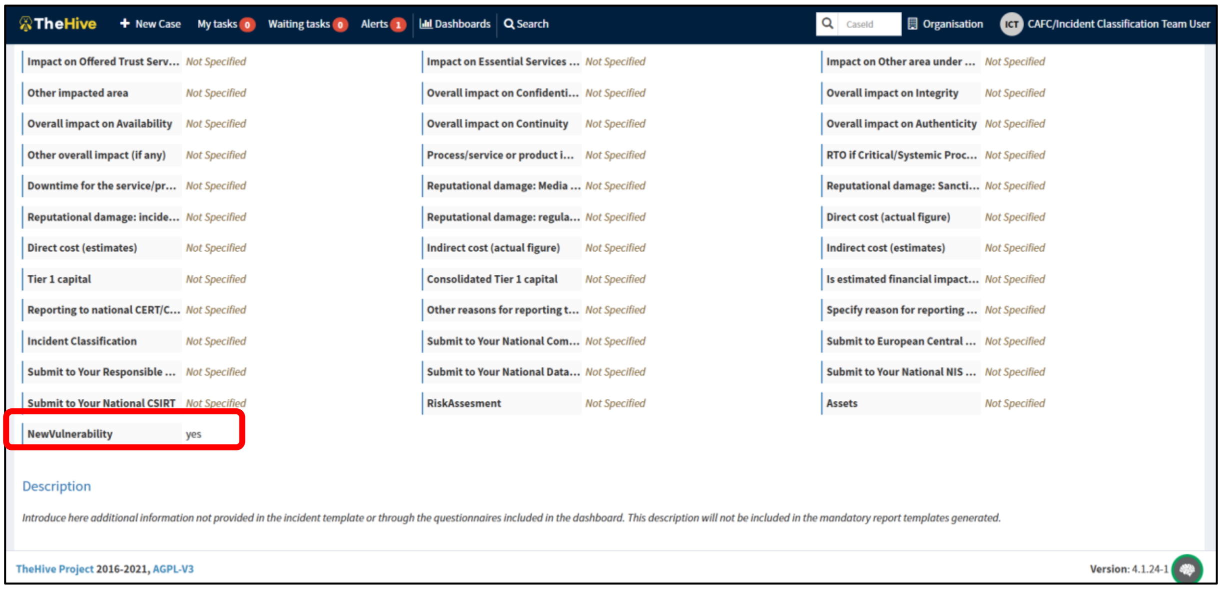

Impact on Offered Trust Serv...	*Not Specified*	Impact on Essential Services ...	*Not Specified*	Impact on Other area under ...	*Not Specified*
Other impacted area	*Not Specified*	Overall impact on Confidenti...	*Not Specified*	Overall impact on Integrity	*Not Specified*
Overall impact on Availability	*Not Specified*	Overall impact on Continuity	*Not Specified*	Overall impact on Authenticity	*Not Specified*
Other overall impact (if any)	*Not Specified*	Process/service or product i...	*Not Specified*	RTO if Critical/Systemic Proc...	*Not Specified*
Downtime for the service/pr...	*Not Specified*	Reputational damage: Media ...	*Not Specified*	Reputational damage: Sancti...	*Not Specified*
Reputational damage: incide...	*Not Specified*	Reputational damage: regula...	*Not Specified*	Direct cost (actual figure)	*Not Specified*
Direct cost (estimates)	*Not Specified*	Indirect cost (actual figure)	*Not Specified*	Indirect cost (estimates)	*Not Specified*
Tier 1 capital	*Not Specified*	Consolidated Tier 1 capital	*Not Specified*	Is estimated financial impact...	*Not Specified*
Reporting to national CERT/C...	*Not Specified*	Other reasons for reporting t...	*Not Specified*	Specify reason for reporting ...	*Not Specified*
Incident Classification	*Not Specified*	Submit to Your National Com...	*Not Specified*	Submit to European Central ...	*Not Specified*
Submit to Your Responsible ...	*Not Specified*	Submit to Your National Data...	*Not Specified*	Submit to Your National NIS ...	*Not Specified*
Submit to Your National CSIRT	*Not Specified*	RiskAssesment	*Not Specified*	Assets	*Not Specified*
NewVulnerability	yes				

Description

Introduce here additional information not provided in the incident template or through the questionnaires included in the dashboard. This description will not be included in the mandatory report templates generated.

TheHive Project 2016-2021, AGPL-V3 Version: 4.1.24-1

Figure 141: New vulnerability identified.

In compliance with NIS2, when this happens, the Incident Reporting module of the CPR tool automatically will trigger a Vulnerability Notification procedure (see workflow in Figure 33) and an additional task Notify Vulnerability will be created in The Hive to deal with it (see Figure 142).

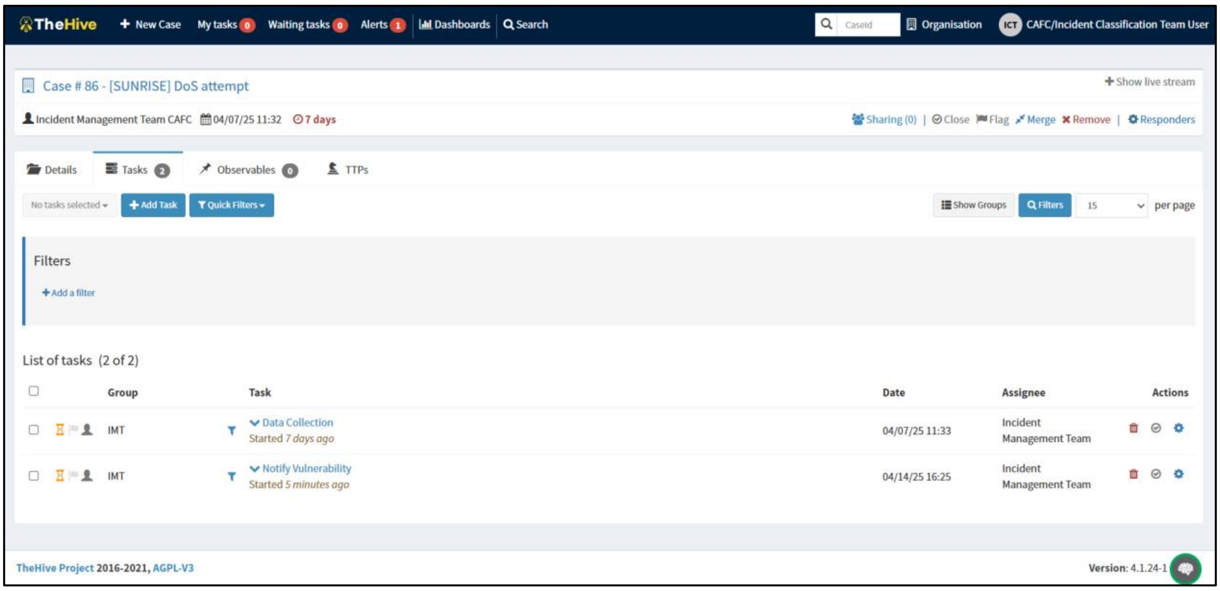

Case # 86 - [SUNRISE] DoS attempt

Incident Management Team CAFC 04/07/25 11:32 7 days

List of tasks (2 of 2)

□	Group	Task	Date	Assignee	Actions
□	IMT	Data Collection Started *7 days ago*	04/07/25 11:33	Incident Management Team	
□	IMT	Notify Vulnerability Started *5 minutes ago*	04/14/25 16:25	Incident Management Team	

TheHive Project 2016-2021, AGPL-V3 Version: 4.1.24-1

Figure 142: Notify Vulnerability task.

5.8 Incident Classification

The ***Incident Classification* Task** is automatically created in TheHive when the *Data Enrichment* Task is completed, as shown in Figure 143. During this task, data available about the incident is analyzed to determine if the incident must be classified as Significant (and consequently notified to the Supervisory Authorities) or not. At this stage, the ongoing report passes to the next IR workflow too, as shown in Figure 144. Consequently, this task is assigned to the ICLT (Incident Classification Team) group, which corresponds in SUNRISE to the end-users belonging to the **CI legal and compliance team**, as described in Section 5.2.

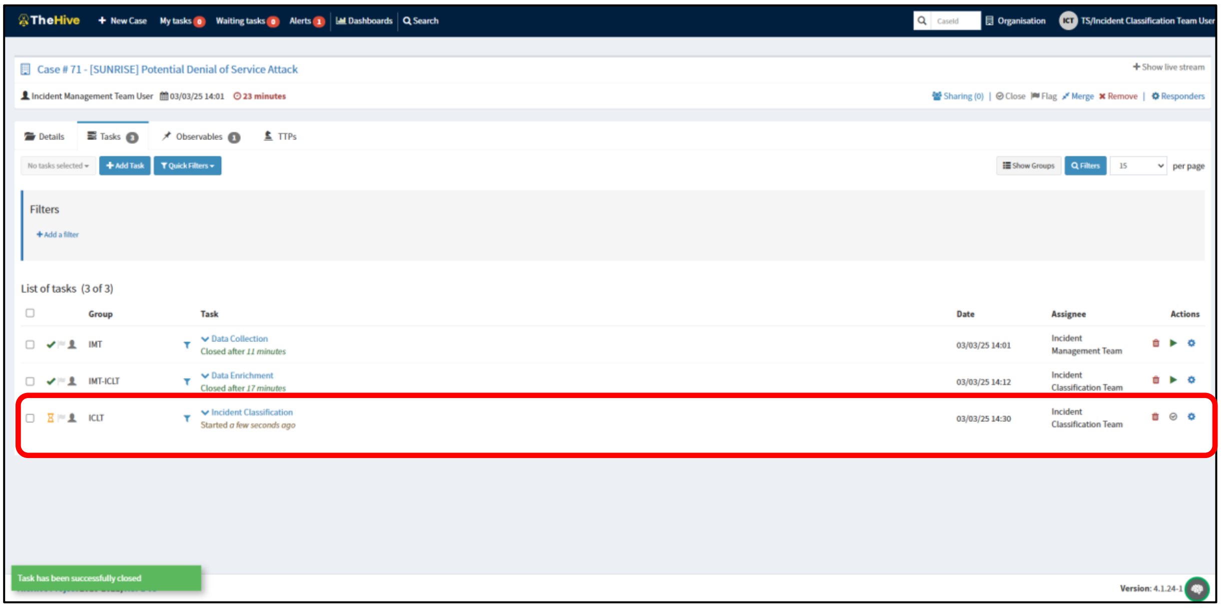

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 23 minutes

List of tasks (3 of 3)

□	Group	Task	Date	Assignee	Actions
□	✓ IMT	Data Collection Closed after *11 minutes*	03/03/25 14:01	Incident Management Team	
□	✓ IMT-ICLT	Data Enrichment Closed after *17 minutes*	03/03/25 14:12	Incident Classification Team	
□	⌛ ICLT	Incident Classification Started *a few seconds ago*	03/03/25 14:30	Incident Classification Team	

Figure 143. Automated creation of Incident Classification task.

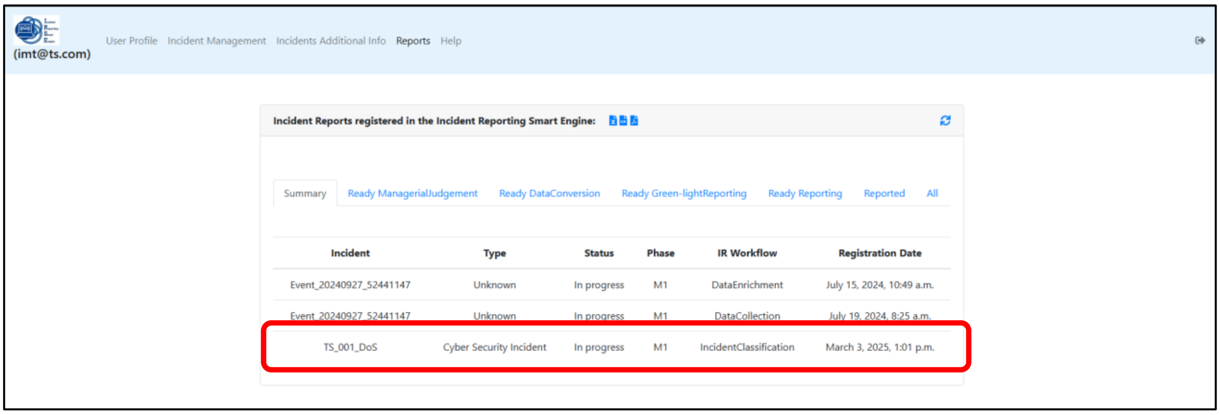

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Type	Status	Phase	IR Workflow	Registration Date
Event_20240927_52441147	Unknown	In progress	M1	DataEnrichment	July 15, 2024, 10:49 a.m.
Event_20240927_52441147	Unknown	In progress	M1	DataCollection	July 19, 2024, 8:25 a.m.
TS_001_DoS	Cyber Security Incident	In progress	M1	IncidentClassification	March 3, 2025, 1:01 p.m.

Figure 144. List of Incident Reports registered.

To support this task of event classification and process the information introduced in the template, the user can invoke the Event Classification *Responder* included with AIRE (which is located in the upper-right corner in TheHive GUI on the page with the Case details), as shown in Figure 145. More information about Cortex Responders and how they work can be found at [4].

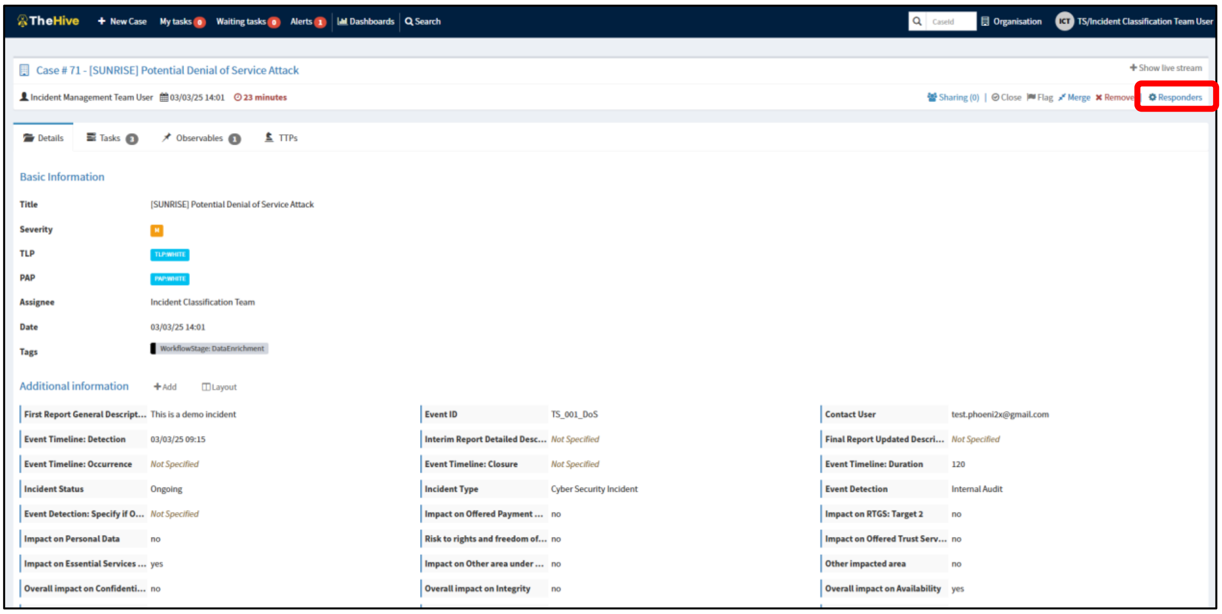

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 23 minutes

Basic Information

Title	[SUNRISE] Potential Denial of Service Attack
Severity	M
TLP	TLP:WHITE
PAP	PAP:WHITE
Assignee	Incident Classification Team
Date	03/03/25 14:01
Tags	WorkflowStage: DataEnrichment

Additional information

First Report General Descript...	This is a demo incident	Event ID	TS_001_DoS	Contact User	test.phoeni2x@gmail.com
Event Timeline: Detection	03/03/25 09:15	Interim Report Detailed Desc...	*Not Specified*	Final Report Updated Descri...	*Not Specified*
Event Timeline: Occurrence	*Not Specified*	Event Timeline: Closure	*Not Specified*	Event Timeline: Duration	120
Incident Status	Ongoing	Incident Type	Cyber Security Incident	Event Detection	Internal Audit
Event Detection: Specify if O...	*Not Specified*	Impact on Offered Payment ...	no	Impact on RTGS: Target 2	no
Impact on Personal Data	no	Risk to rights and freedom of...	no	Impact on Offered Trust Serv...	no
Impact on Essential Services ...	yes	Impact on Other area under ...	no	Other impacted area	no
Overall impact on Confidenti...	no	Overall impact on Integrity	no	Overall impact on Availability	yes

Figure 145. Responders in an Incident Detail View.

The Responder is integrated with the AIRE asset, as shown in Figure 146. As such, the role of the user who executes it will be checked. It will only get a result when it belongs to the CI legal and compliance team.

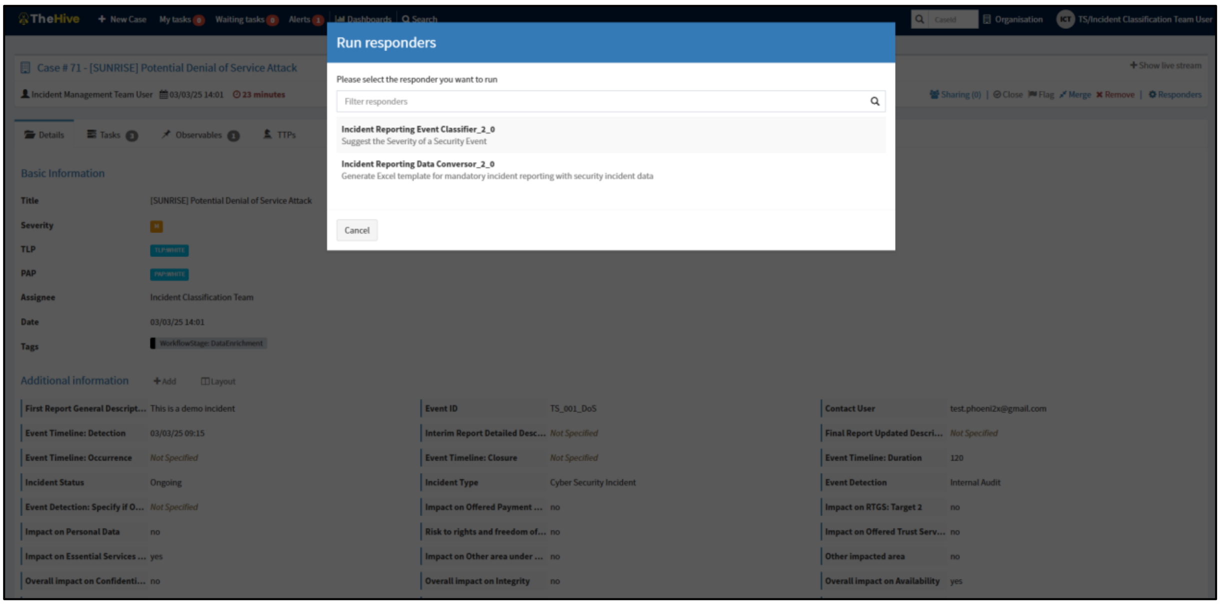

Run responders

Please select the responder you want to run

Incident Reporting Event Classifier_2_0
Suggest the Severity of a Security Event

Incident Reporting Data Conversor_2_0
Generate Excel template for mandatory incident reporting with security incident data

Figure 146: Selection of one responder to run for an incident case.

The result of the execution of the responder with the classification will be available at the bottom of the incident page, as shown in Figure 147, and the action button shows the output, as shown in Figure 148.

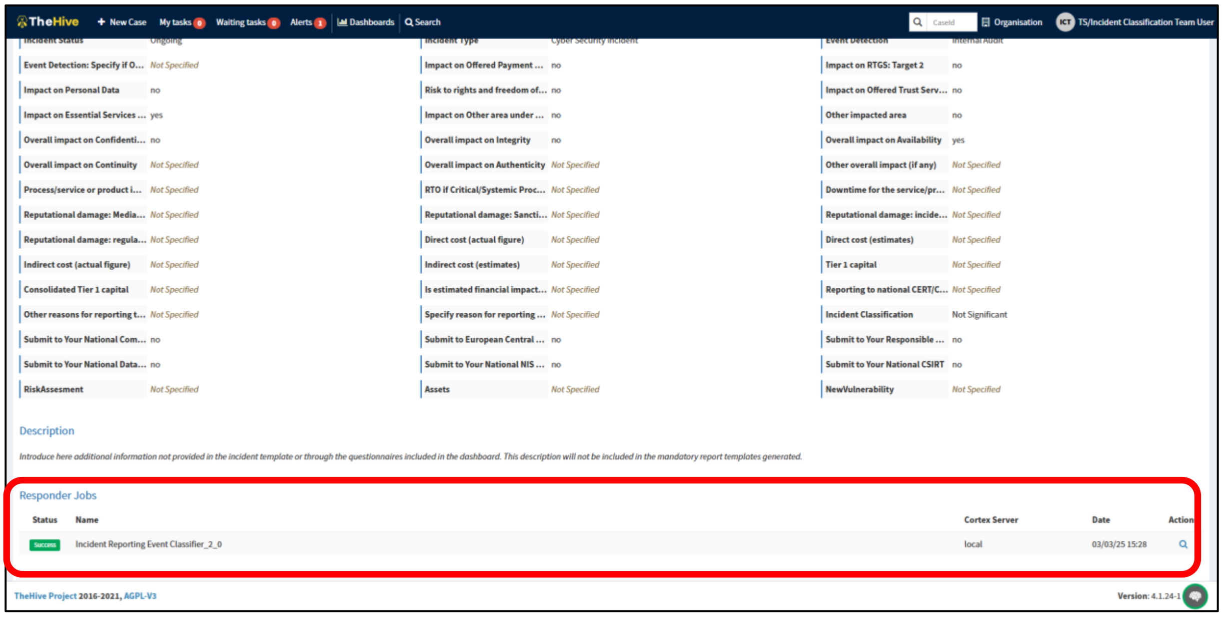

Incident Status	Ongoing	Incident Type	Cyber Security Incident	Event Detection	Internal Audit
Event Detection: Specify if O...	Not Specified	Impact on Offered Payment ...	no	Impact on RTGS: Target 2	no
Impact on Personal Data	no	Risk to rights and freedom of...	no	Impact on Offered Trust Serv...	no
Impact on Essential Services ...	yes	Impact on Other area under ...	no	Other impacted area	no
Overall impact on Confidenti...	no	Overall impact on Integrity	no	Overall impact on Availability	yes
Overall impact on Continuity	Not Specified	Overall impact on Authenticity	Not Specified	Other overall impact (if any)	Not Specified
Process/service or product i...	Not Specified	RTO if Critical/Systemic Proc...	Not Specified	Downtime for the service/pr...	Not Specified
Reputational damage: Media...	Not Specified	Reputational damage: Sancti...	Not Specified	Reputational damage: incide...	Not Specified
Reputational damage: regula...	Not Specified	Direct cost (actual figure)	Not Specified	Direct cost (estimates)	Not Specified
Indirect cost (actual figure)	Not Specified	Indirect cost (estimates)	Not Specified	Tier 1 capital	Not Specified
Consolidated Tier 1 capital	Not Specified	Is estimated financial impact...	Not Specified	Reporting to national CERT/C...	Not Specified
Other reasons for reporting t...	Not Specified	Specify reason for reporting ...	Not Specified	Incident Classification	Not Significant
Submit to Your National Com...	no	Submit to European Central ...	no	Submit to Your Responsible ...	no
Submit to Your National Data...	no	Submit to Your National NIS ...	no	Submit to Your National CSIRT	no
RiskAssesment	Not Specified	Assets	Not Specified	NewVulnerability	Not Specified

Description

Introduce here additional information not provided in the incident template or through the questionnaires included in the dashboard. This description will not be included in the mandatory report templates generated.

Responder Jobs

Status	Name	Cortex Server	Date	Action
Success	Incident Reporting Event Classifier_2_0	local	03/03/25 15:28	🔍

TheHive Project 2016-2021, AGPL-V3

Version: 4.1.24-1

Figure 147: List of Responder Jobs with status.

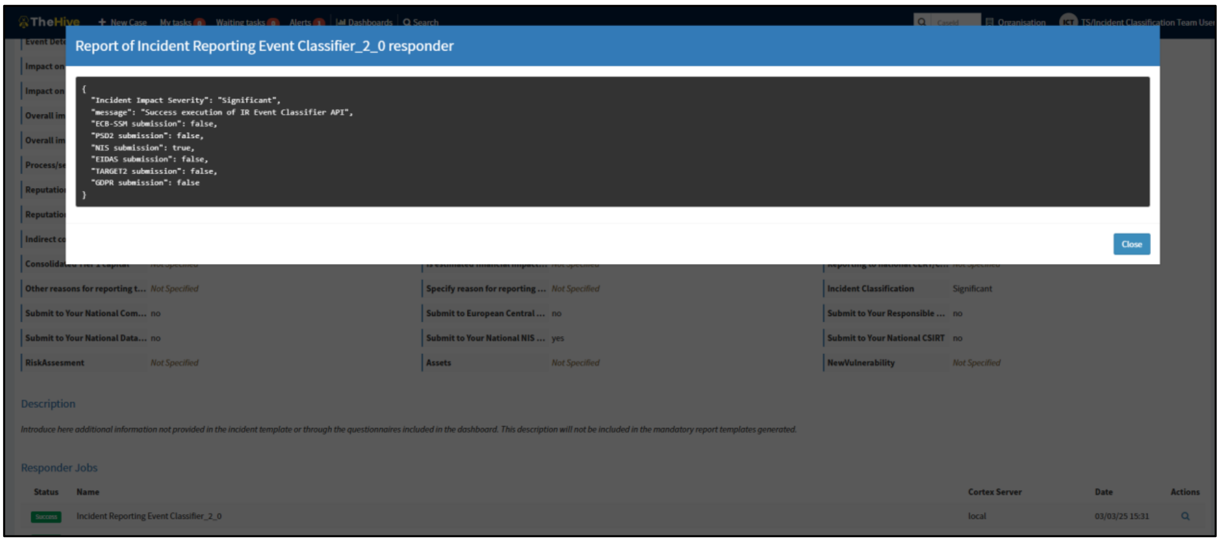

Report of Incident Reporting Event Classifier_2_0 responder

```
{
  "Incident Impact Severity": "Significant",
  "message": "Success execution of IR Event Classifier API",
  "ECB-SSM submission": false,
  "PSD2 submission": false,
  "NIS submission": true,
  "EIDAS submission": false,
  "TARGET2 submission": false,
  "GDPR submission": false
}
```

Other reasons for reporting t...	Not Specified	Specify reason for reporting ...	Not Specified	Incident Classification	Significant
Submit to Your National Com...	no	Submit to European Central ...	no	Submit to Your Responsible ...	no
Submit to Your National Data...	no	Submit to Your National NIS ...	yes	Submit to Your National CSIRT	no
RiskAssesment	Not Specified	Assets	Not Specified	NewVulnerability	Not Specified

Description

Introduce here additional information not provided in the incident template or through the questionnaires included in the dashboard. This description will not be included in the mandatory report templates generated.

Responder Jobs

Status	Name	Cortex Server	Date	Actions
Success	Incident Reporting Event Classifier_2_0	local	03/03/25 15:31	🔍

Figure 148: Output example for incident classification.

This information will also be automatically updated in the tags of the case, as shown in Figure 149, and in the fields of the template (see Figure 150), so the user can check the suggestion and modify it if he/she considers it. The **risk-based classification** performed by the Cyber-physical risk assessment module (CERCA) of the CPR tool is also automatically added by AIRE to these tags to be considered by the CI's end-user.

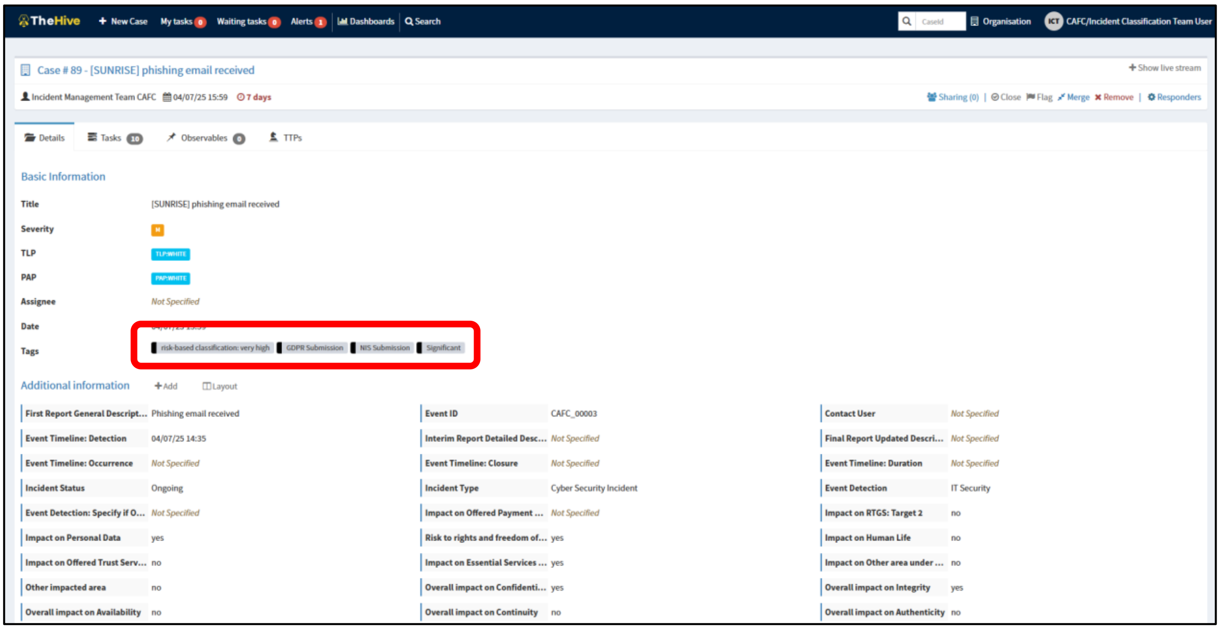

Case # 89 - [SUNRISE] phishing email received

Incident Management Team CAFC 04/07/25 15:59 7 days

Basic Information

Title	[SUNRISE] phishing email received
Severity	M
TLP	TLP:WHITE
PAP	PAP:WHITE
Assignee	Not Specified
Date	04/07/25 15:59
Tags	risk-based classification: very high GDPR Submission NIS Submission Significant

Additional information

First Report General Descript...	Phishing email received	Event ID	CAFC_00003	Contact User	Not Specified
Event Timeline: Detection	04/07/25 14:35	Interim Report Detailed Desc...	Not Specified	Final Report Updated Descri...	Not Specified
Event Timeline: Occurrence	Not Specified	Event Timeline: Closure	Not Specified	Event Timeline: Duration	Not Specified
Incident Status	Ongoing	Incident Type	Cyber Security Incident	Event Detection	IT Security
Event Detection: Specify if O...	Not Specified	Impact on Offered Payment ...	Not Specified	Impact on RTGS: Target 2	no
Impact on Personal Data	yes	Risk to rights and freedom of...	yes	Impact on Human Life	no
Impact on Offered Trust Serv...	no	Impact on Essential Services ...	yes	Impact on Other area under ...	no
Other impacted area	no	Overall impact on Confidenti...	yes	Overall impact on Integrity	yes
Overall impact on Availability	no	Overall impact on Continuity	no	Overall impact on Authenticity	no

Figure 149. Tags of an Incident Case.

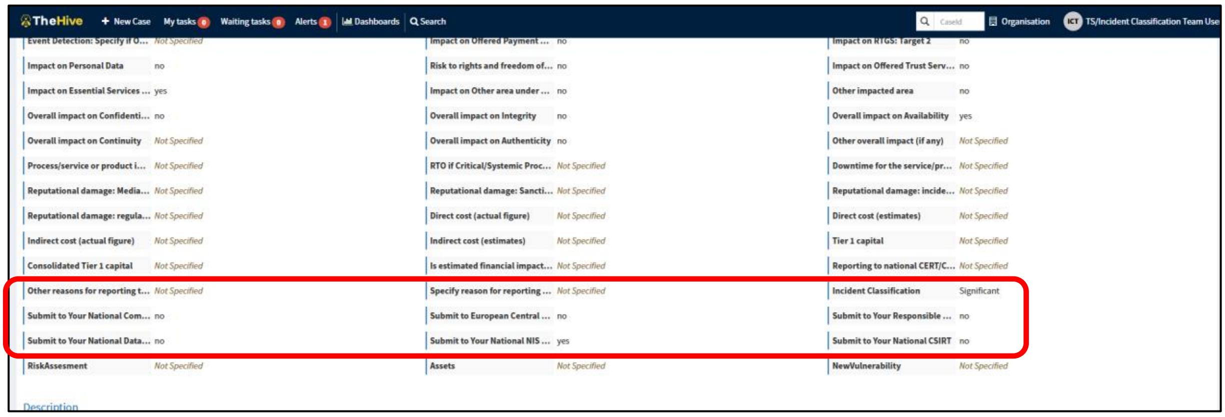

Event Detection: Specify if O...	Not Specified	Impact on Offered Payment ...	no	Impact on RTGS: Target 2	no
Impact on Personal Data	no	Risk to rights and freedom of...	no	Impact on Offered Trust Serv...	no
Impact on Essential Services ...	yes	Impact on Other area under ...	no	Other impacted area	no
Overall impact on Confidenti...	no	Overall impact on Integrity	no	Overall impact on Availability	yes
Overall impact on Continuity	Not Specified	Overall impact on Authenticity	no	Other overall impact (if any)	Not Specified
Process/service or product i...	Not Specified	RTO if Critical/Systemic Proc...	Not Specified	Downtime for the service/pr...	Not Specified
Reputational damage: Media...	Not Specified	Reputational damage: Sancti...	Not Specified	Reputational damage: incide...	Not Specified
Reputational damage: regula...	Not Specified	Direct cost (actual figure)	Not Specified	Direct cost (estimates)	Not Specified
Indirect cost (actual figure)	Not Specified	Indirect cost (estimates)	Not Specified	Tier 1 capital	Not Specified
Consolidated Tier 1 capital	Not Specified	Is estimated financial impact...	Not Specified	Reporting to national CERT/C...	Not Specified
Other reasons for reporting t...	Not Specified	Specify reason for reporting ...	Not Specified	Incident Classification	Significant
Submit to Your National Com...	no	Submit to European Central ...	no	Submit to Your Responsible ...	no
Submit to Your National Data...	no	Submit to Your National NIS ...	yes	Submit to Your National CSIRT	no
RiskAssesment	Not Specified	Assets	Not Specified	NewVulnerability	Not Specified

Figure 150. Fields of a template.

First, the execution of the responder will invoke the AIRE asset to determine if the user has permissions to execute for this phase of the workflow. If not, the user will receive a notification similar to that in Figure 151.

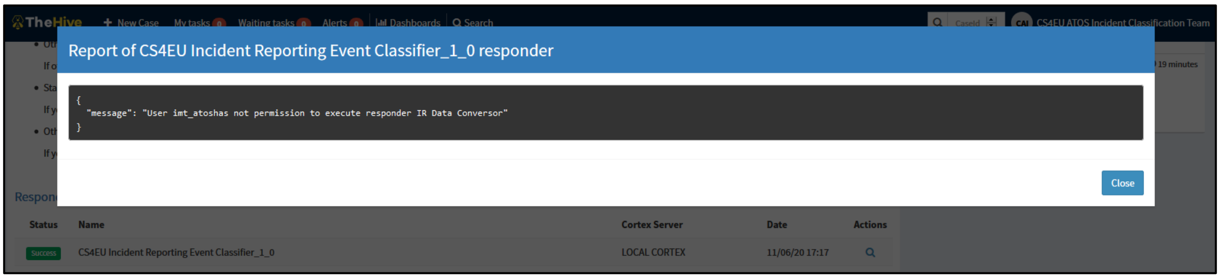

Report of CS4EU Incident Reporting Event Classifier_1_0 responder

```
{
  "message": "User imt_atoshas not permission to execute responder IR Data Conversor"
}
```

Status	Name	Cortex Server	Date	Actions
Success	CS4EU Incident Reporting Event Classifier_1_0	LOCAL CORTEX	11/06/20 17:17	

Figure 151. Output of a responder without permission.

5.9 Early Warning and countermeasures evaluation procedure

Once the Incident Classification task has been closed by the user, if it has been established that the incident has been classified as Significant, it will be automatically sent an **Early Warning** to the Supervisory Authorities configured. This is, to the Channels configured with emails that have been set up as enabled.

Automatically, it is triggered in the Incident Reporting module a countermeasures evaluation procedure (as shown in Figure 32) and a new task in TheHive called "Register countermeasures" (see Figure 153) is created and assigned to the IMT group (in SUNRISE, the end-user with role **CI Incident Response and security operation team)**. Once the operator has received from the CSIRTs some mitigation measures and they have been included associated to the ongoing incident through the menu in the tab Incident Additional Info (as shown in Figure 152), the task can be closed. Automatically, a task to evaluate these countermeasures is created in TheHive (as shown in Figure 153). Finally, once the suitable mitigation actions have been applied and the information in the incident updated, the task can be closed and automatically an alarm will be sent to the Cyber-physical risk assessment module of the CPR tool to re-evaluate risk. No task will appear in TheHive for this step since this invocation is done automatically by the Incident Reporting module of CPR tool without intervention of the user.

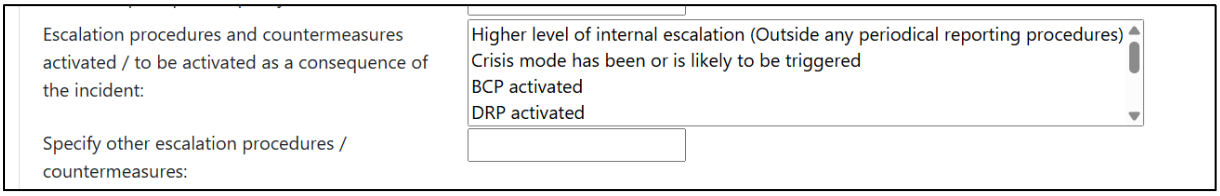

Escalation procedures and countermeasures activated / to be activated as a consequence of the incident:	Higher level of internal escalation (Outside any periodical reporting procedures) Crisis mode has been or is likely to be triggered BCP activated DRP activated
Specify other escalation procedures / countermeasures:	

Figure 152. AIRE dashboard – countermeasures information

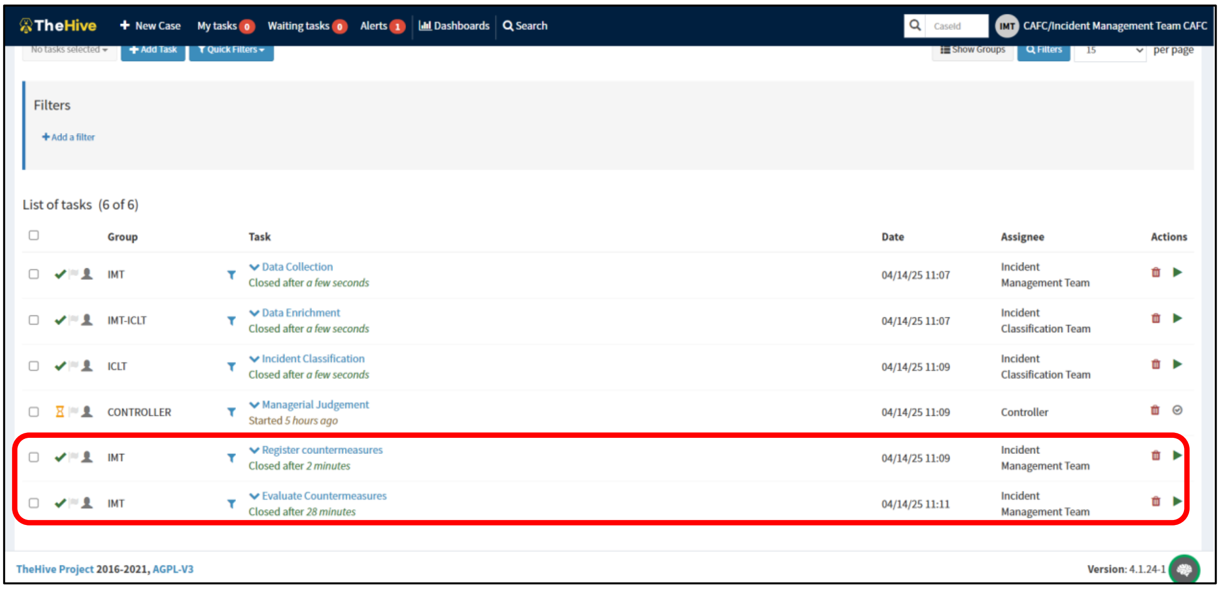

List of tasks (6 of 6)

□	Group	Task	Date	Assignee	Actions
□ ✓	IMT	Data Collection Closed after *a few seconds*	04/14/25 11:07	Incident Management Team	
□ ✓	IMT-ICLT	Data Enrichment Closed after *a few seconds*	04/14/25 11:07	Incident Classification Team	
□ ⌛	ICLT	Incident Classification Closed after *a few seconds*	04/14/25 11:09	Incident Classification Team	
□ ⌛	CONTROLLER	Managerial Judgement Started *5 hours ago*	04/14/25 11:09	Controller	
□ ✓	IMT	Register countermeasures Closed after *2 minutes*	04/14/25 11:09	Incident Management Team	
□ ✓	IMT	Evaluate Countermeasures Closed after *28 minutes*	04/14/25 11:11	Incident Management Team	

Figure 153. Tasks associated to countermeasures management.

5.10 Managerial Judgement Tasks

There are two points in the mandatory incident reporting workflow (as shown in Figure 31) where it is necessary the human intervention of a managerial role to confirm if the incident process must continue or not. First, to confirm the classification of the incident as Significant and proceed with the preparation of the required reports; and second to confirm the report generated contains all the required information and it is ready to be sent to the Supervisory Authorities.

- **Managerial Judgement Task:** This task is automatically created and assigned to the Controller (belonging to the **CI legal and compliance team**) when the *Incident Classification* task is closed, as shown in Figure 154 and Figure 155.

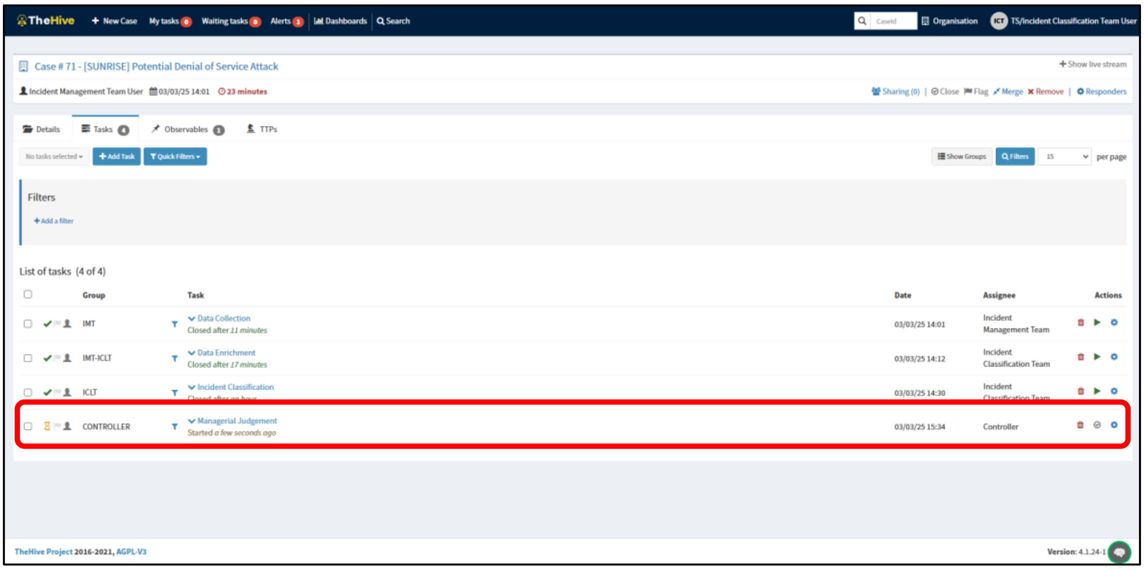

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 23 minutes

Filters

List of tasks (4 of 4)

Group	Task	Date	Assignee
IMT	Data Collection Closed after *11 minutes*	03/03/25 14:01	Incident Management Team
IMT-ICLT	Data Enrichment Closed after *17 minutes*	03/03/25 14:12	Incident Classification Team
ICLT	Incident Classification Closed after an hour	03/03/25 14:30	Incident Classification Team
CONTROLLER	Managerial Judgement Started *a few seconds ago*	03/03/25 15:34	Controller

Figure 154. Automatic creation of Managerial Judgement Task.

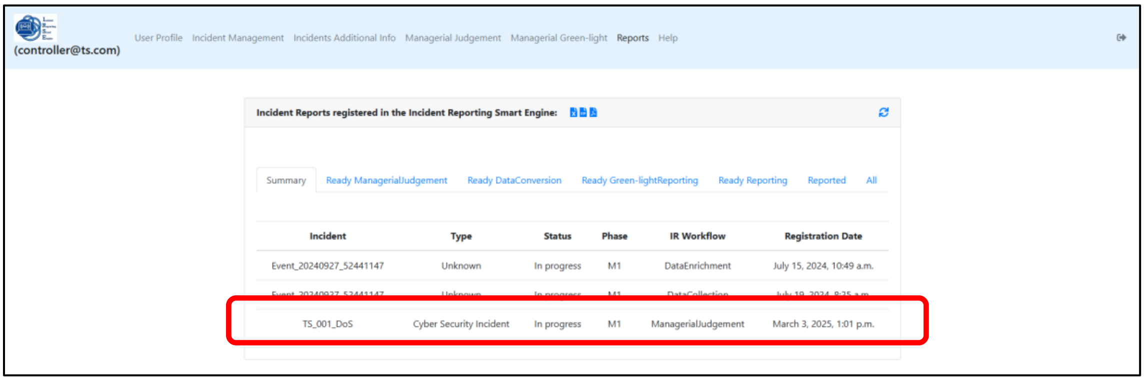

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Type	Status	Phase	IR Workflow	Registration Date
Event_20240927_52441147	Unknown	In progress	M1	DataEnrichment	July 15, 2024, 10:49 a.m.
Event_20240927_52441147	Unknown	In progress	M1	DataCollection	July 19, 2024, 8:25 a.m.
TS_001_DoS	Cyber Security Incident	In progress	M1	ManagerialJudgement	March 3, 2025, 1:01 p.m.

Figure 155. Automatic creation of Managerial Judgement IR Workflow.

Under the menu *Managerial Judgement*, the controller will see the report with the impact classification, as shown in Figure 156.

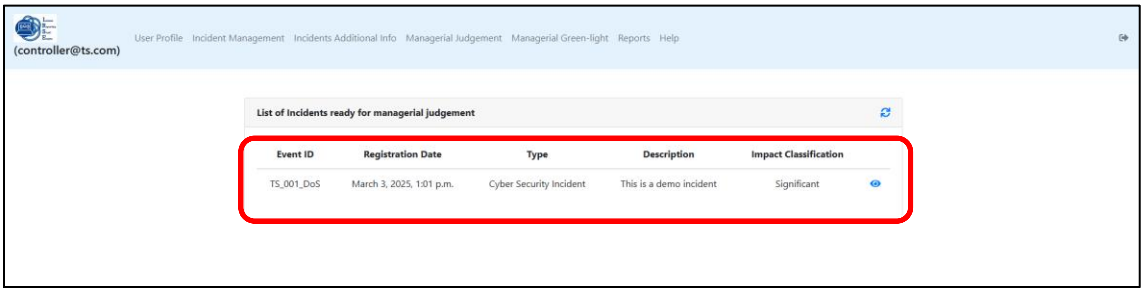

List of Incidents ready for managerial judgement

Event ID	Registration Date	Type	Description	Impact Classification
TS_001_DoS	March 3, 2025, 1:01 p.m.	Cyber Security Incident	This is a demo incident	Significant

Figure 156. List of Incident ready for managerial judgement.

The *Detail* button (the “Eye”) shows the event severity classification, and the suggested mandatory reporting based on the criteria of the regulations enabled, as shown in Figure 157.

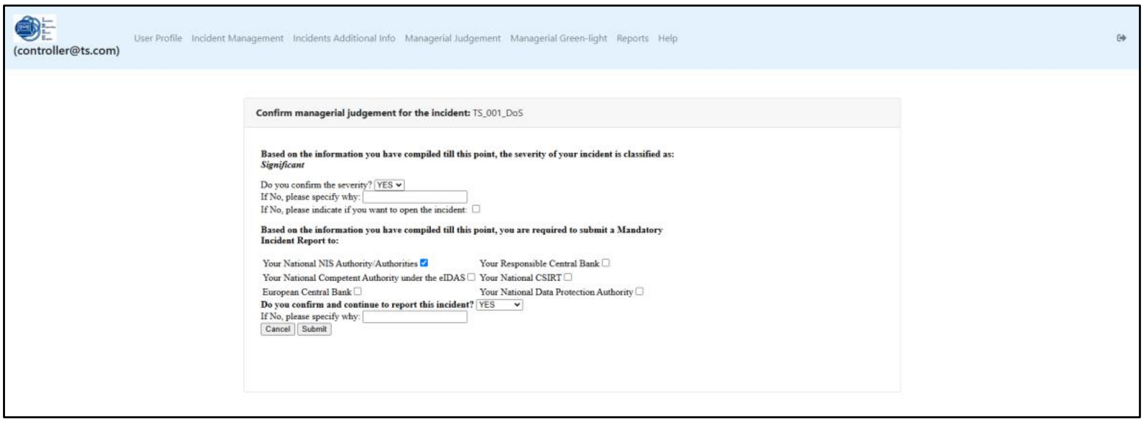

Confirm managerial judgement for the incident: TS_001_DoS

Based on the information you have compiled till this point, the severity of your incident is classified as: ***Significant***

Do you confirm the severity? YES
If No, please specify why:
If No, please indicate if you want to open the incident: □

Based on the information you have compiled till this point, you are required to submit a Mandatory Incident Report to:

Your National NIS Authority/Authorities ☑ Your Responsible Central Bank □
Your National Competent Authority under the eIDAS □ Your National CSIRT □
European Central Bank □ Your National Data Protection Authority □
Do you confirm and continue to report this incident? YES
If No, please specify why:
Cancel Submit

Figure 157. Form for a managerial judgement.

NOTE: If the user with Controller role does not confirm the classification and select the option to open the incident, the incident reporting process will return to the “Data Enrichment” phase. If the controller does not confirm to proceed with the reporting, the incident reporting process will finish, and the reports will be closed.

Automatically, the task called “*Managerial Judgement*” created in TheHive will be closed by AIRE and the workflow of the report will be moved to the following step (depending on the managerial judgement). Then, a new task, “*Data Conversion*”, will be created and assigned to the Incident Reporting Team (in SUNRISE, this is the people of the CI legal and compliance team responsible for preparing the documents to be sent to the Supervisor Authorities) and the incident will be ready for report preparation, as shown in Figure 158. This event is registered in the *Summary* of the *Reports* tab, as shown in Figure 159.

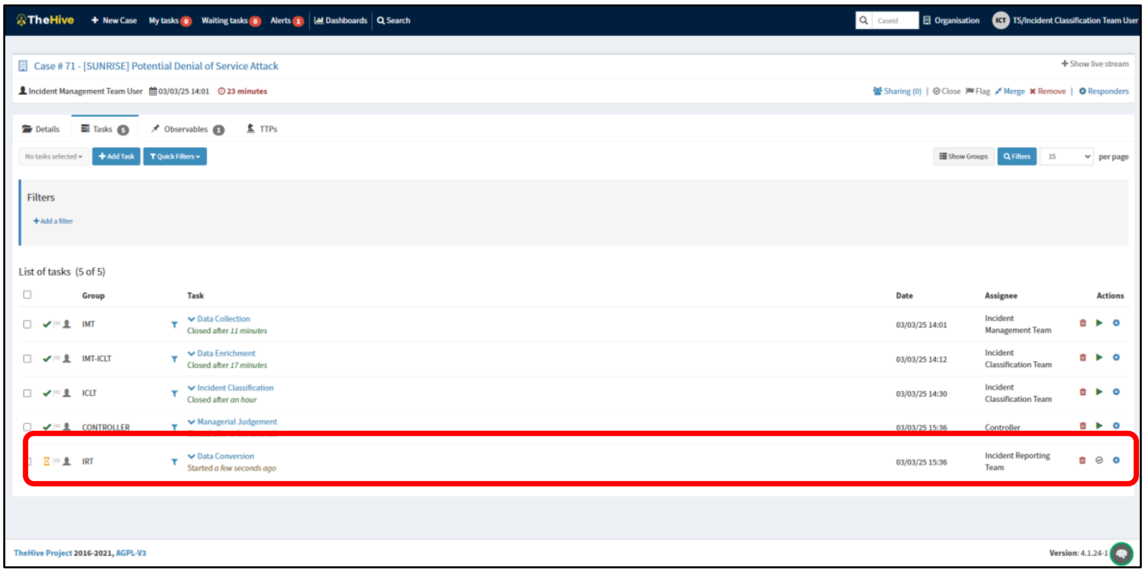

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 23 minutes

Filters

List of tasks (5 of 5)

Group	Task	Date	Assignee
IMT	Data Collection Closed after *11 minutes*	03/03/25 14:01	Incident Management Team
IMT-ICLT	Data Enrichment Closed after *17 minutes*	03/03/25 14:12	Incident Classification Team
ICLT	Incident Classification Closed after *an hour*	03/03/25 14:30	Incident Classification Team
CONTROLLER	Managerial Judgement	03/03/25 15:36	Controller
IRT	Data Conversion Started *a few seconds ago*	03/03/25 15:36	Incident Reporting Team

Figure 158. New Data Conversion Task.

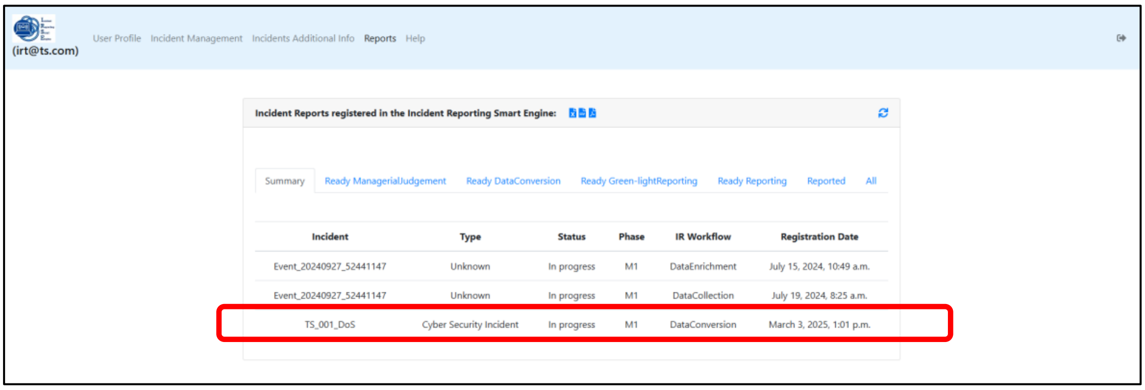

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Type	Status	Phase	IR Workflow	Registration Date
Event_20240927_52441147	Unknown	In progress	M1	DataEnrichment	July 15, 2024, 10:49 a.m.
Event_20240927_52441147	Unknown	In progress	M1	DataCollection	July 19, 2024, 8:25 a.m.
TS_001_DoS	Cyber Security Incident	In progress	M1	DataConversion	March 3, 2025, 1:01 p.m.

Figure 159. Register of Data Conversion Task.

- **Managerial Green-light:** This is the step of the managerial judgement when the report is already ready but need to be confirmed before sending them. At this step, the “Managerial Green-light” tab of the AIRE dashboard, which appears when login to the dashboard with a user with Controller role, will display the list of ready reports, as shown Figure 160.

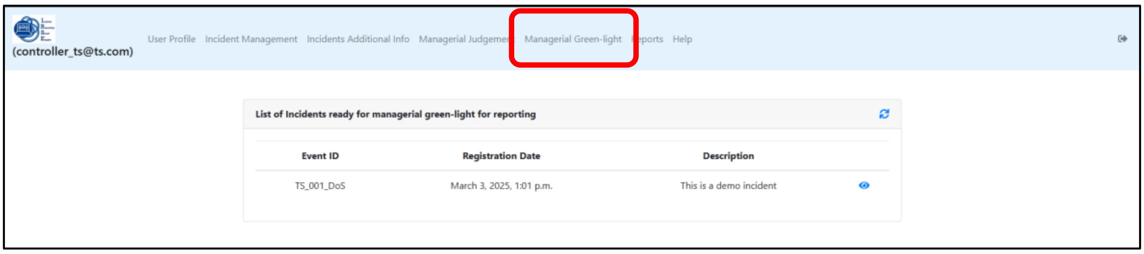

List of Incidents ready for managerial green-light for reporting

Event ID	Registration Date	Description
TS_001_DoS	March 3, 2025, 1:01 p.m.	This is a demo incident

Figure 160. List of Incident ready for managerial green-light for reporting.

Once the details are selected (eye icon), the managerial judgement form will appear to confirm if you wish to proceed with the reporting, as shown in Figure 161.

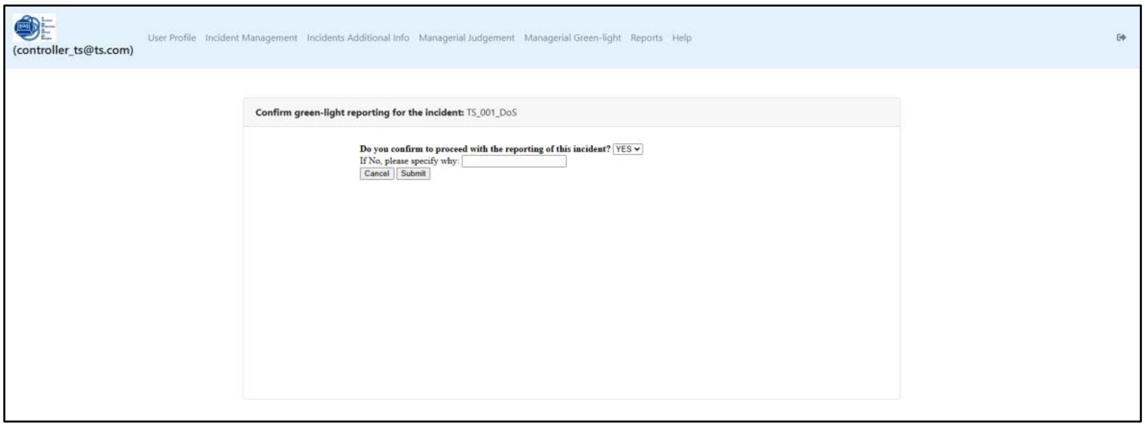

Confirm green-light reporting for the incident: TS_001_DoS

Do you confirm to proceed with the reporting of this incident? YES
If No, please specify why:

Figure 161. Configuration of green-light reporting form.

The *Submit* button automatically closes the task *Green-light for Reporting*, showing a confirmation, as shown in Figure 162. It also creates new task, *Reporting & Release*, which will be assigned to IRT user, as shown in Figure 163.

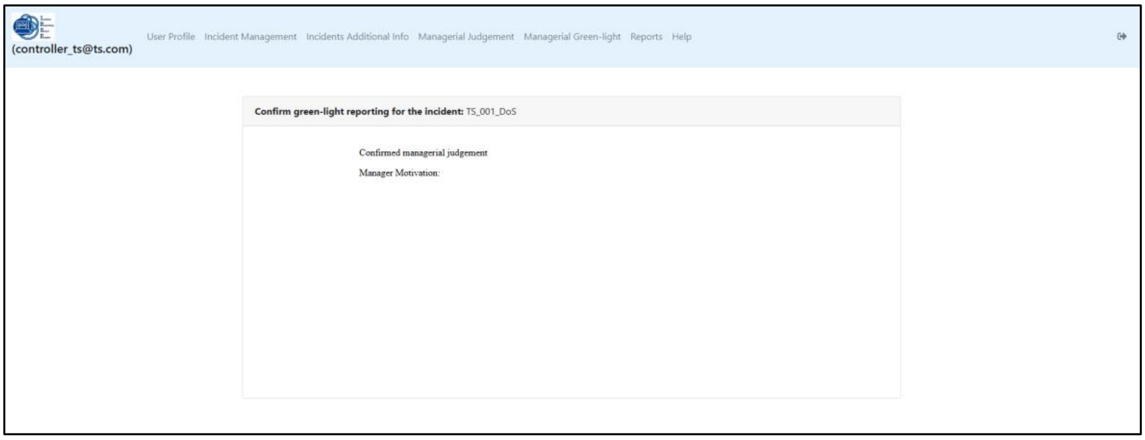

Confirm green-light reporting for the incident: TS_001_DoS

Confirmed managerial judgement

Manager Motivation:

Figure 162. Confirmation of task closed.

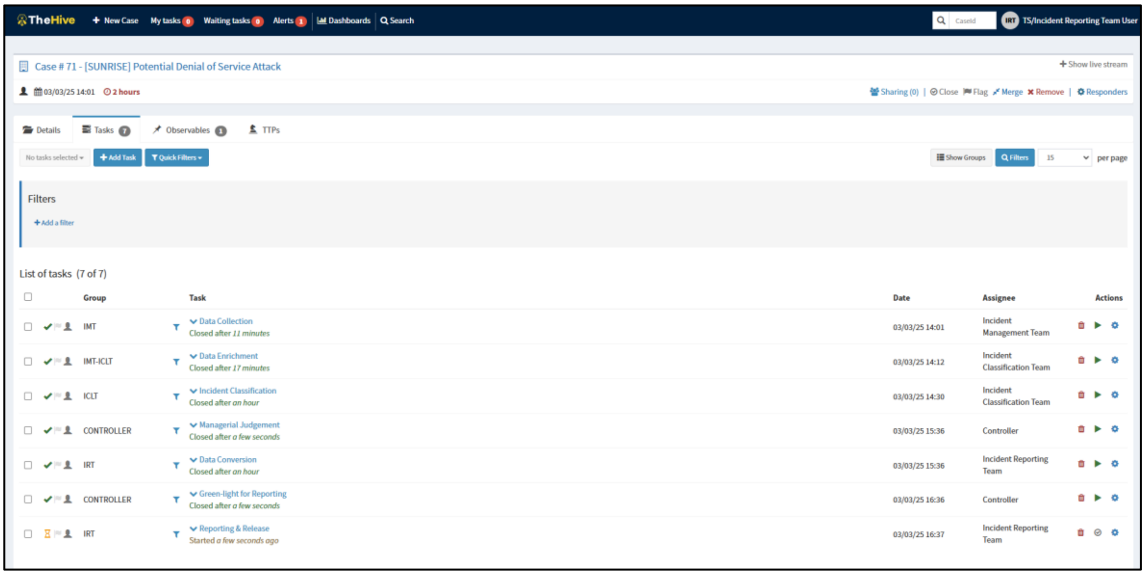

Case # 71 - [SUNRISE] Potential Denial of Service Attack

03/03/25 14:01 2 hours

List of tasks (7 of 7)

	Group	Task	Date	Assignee	Actions
□ ✓	IMT	Data Collection Closed after *11 minutes*	03/03/25 14:01	Incident Management Team	
□ ✓	IMT-ICLT	Data Enrichment Closed after *17 minutes*	03/03/25 14:12	Incident Classification Team	
□ ✓	ICLT	Incident Classification Closed after *an hour*	03/03/25 14:30	Incident Classification Team	
□ ✓	CONTROLLER	Managerial Judgement Closed after *a few seconds*	03/03/25 15:36	Controller	
□ ✓	IRT	Data Conversion Closed after *an hour*	03/03/25 15:36	Incident Reporting Team	
□ ✓	CONTROLLER	Green-light for Reporting Closed after *a few seconds*	03/03/25 16:36	Controller	
□ ⌛	IRT	Reporting & Release Started *a few seconds ago*	03/03/25 16:37	Incident Reporting Team	

Figure 163. New Reporting and Release Task.

The event is registered on the Summary of the Reports, Figure 164.

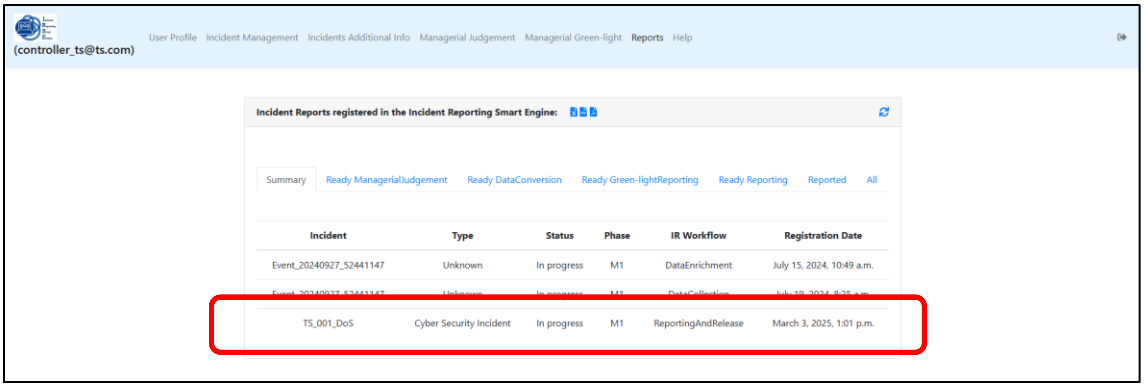

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Type	Status	Phase	IR Workflow	Registration Date
Event_20240927_52441147	Unknown	In progress	M1	DataEnrichment	July 15, 2024, 10:49 a.m.
Event_20240927_52441147	Unknown	In progress	M1	DataCollection	July 19, 2024, 8:25 a.m.
TS_001_DoS	Cyber Security Incident	In progress	M1	ReportingAndRelease	March 3, 2025, 1:01 p.m.

Figure 164. Reporting And Release register in Summary of Reports.

NOTE: If the Controller does not confirm proceeding with the actual reporting, the incident reporting process will finish, and the reports will be closed.

5.11 Generation of Incident Reports

Once the Managerial Judgement task has been closed and it has been created in TheHive the task ***Data Conversion***, it is the moment of generating the documents (PDF, doc, Excel…) that will be sent to the Supervisory Authorities.

At this step, if there is additional information required for reporting that needs to be completed, it must be included in the template or through AIRE dashboard by the Incident Reporting Team user (in SUNRISE, end-user with role ***CI Legal and Compliance team***). Then, the responder *Incident Reporting Data Converter* must be invoked to generate the files associated to the regulations enabled and confirmed by the Controller in the managerial judgement, as shown in Figure 165. These files will follow different format (pdf, word, excel) depending on the templates configured during the configuration of the Incident Reporting module of the CPR tool.

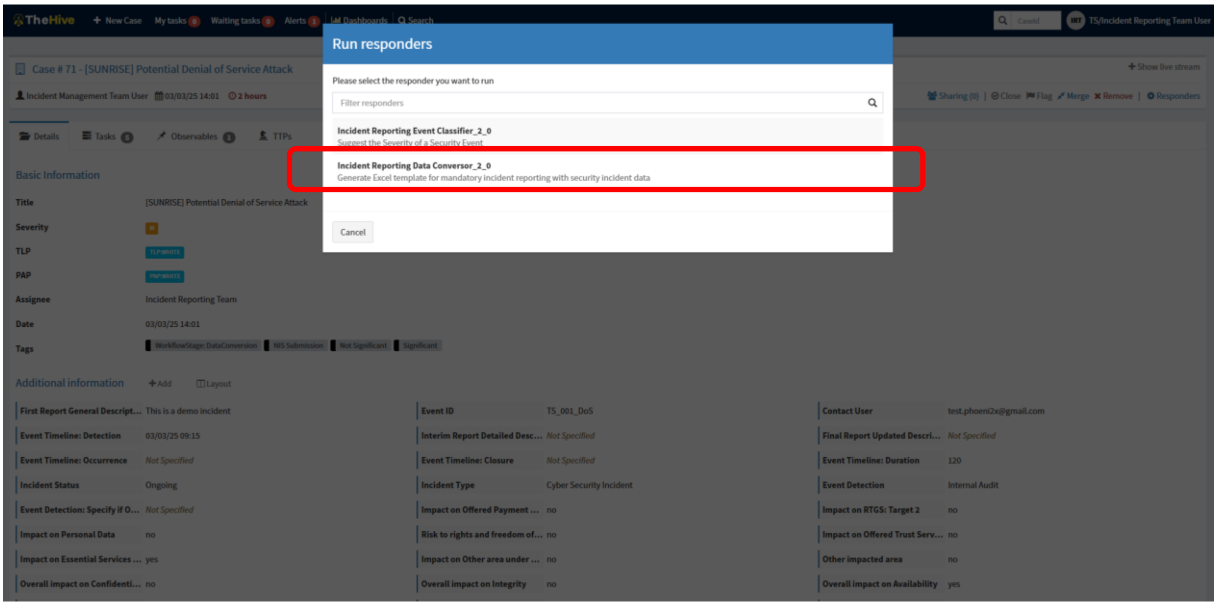

Run responders

Please select the responder you want to run

Filter responders

Incident Reporting Event Classifier_2_0
Suggest the Severity of a Security Event

Incident Reporting Data Conversor_2_0
Generate Excel template for mandatory incident reporting with security incident data

Cancel

Case # 71 - [SUNRISE] Potential Denial of Service Attack

Incident Management Team User 03/03/25 14:01 2 hours

Basic Information

Title	[SUNRISE] Potential Denial of Service Attack
Severity	M
TLP	TLP:WHITE
PAP	PAP:WHITE
Assignee	Incident Reporting Team
Date	03/03/25 14:01
Tags	WorkflowStage: DataConversion NIS Submission Not Significant Significant

Additional information

First Report General Descript…	This is a demo incident	**Event ID**	TS_001_DoS	**Contact User**	test.phoeni2x@gmail.com
Event Timeline: Detection	03/03/25 09:15	**Interim Report Detailed Desc…**	*Not Specified*	**Final Report Updated Descri…**	*Not Specified*
Event Timeline: Occurrence	*Not Specified*	**Event Timeline: Closure**	*Not Specified*	**Event Timeline: Duration**	120
Incident Status	Ongoing	**Incident Type**	Cyber Security Incident	**Event Detection**	Internal Audit
Event Detection: Specify if O…	*Not Specified*	**Impact on Offered Payment …**	no	**Impact on RTGS: Target 2**	no
Impact on Personal Data	no	**Risk to rights and freedom of…**	no	**Impact on Offered Trust Serv…**	no
Impact on Essential Services …	yes	**Impact on Other area under …**	no	**Other impacted area**	no
Overall impact on Confidenti…	no	**Overall impact on Integrity**	no	**Overall impact on Availability**	yes

Figure 165. Run responder invocation.

The result of the execution is shown at the end of the page, as shown in Figure 166. The action button opens a text box with the result of the responder execution, as shown in Figure 167.

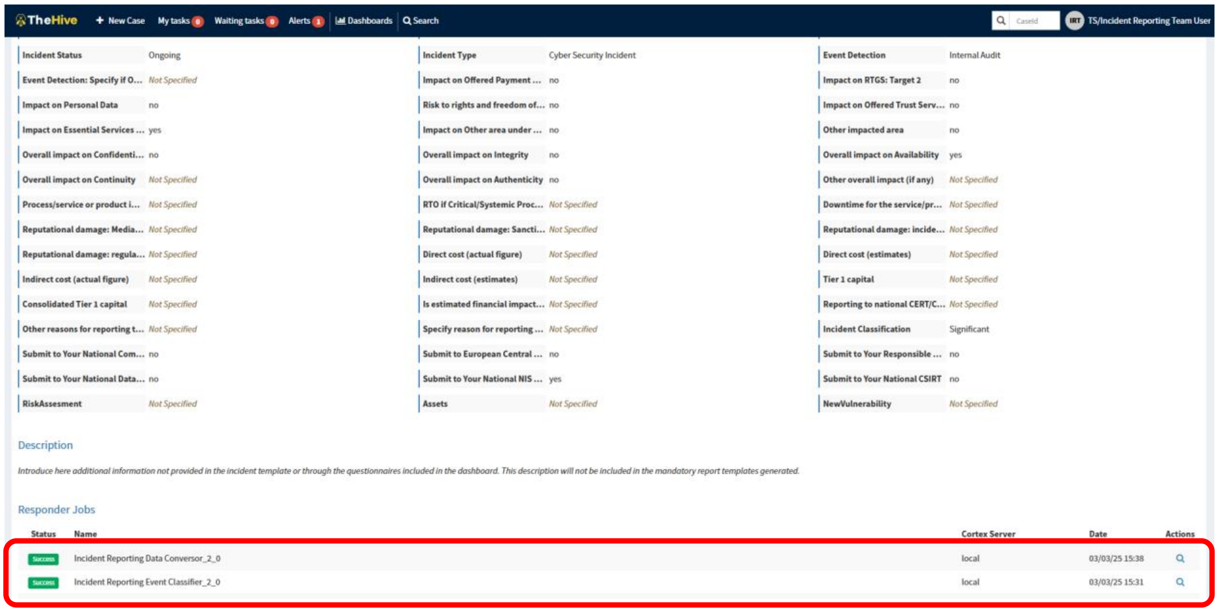

Incident Status	Ongoing	**Incident Type**	Cyber Security Incident	**Event Detection**	Internal Audit
Event Detection: Specify if O…	*Not Specified*	**Impact on Offered Payment …**	no	**Impact on RTGS: Target 2**	no
Impact on Personal Data	no	**Risk to rights and freedom of…**	no	**Impact on Offered Trust Serv…**	no
Impact on Essential Services …	yes	**Impact on Other area under …**	no	**Other impacted area**	no
Overall impact on Confidenti…	no	**Overall impact on Integrity**	no	**Overall impact on Availability**	yes
Overall impact on Continuity	*Not Specified*	**Overall impact on Authenticity**	no	**Other overall impact (if any)**	*Not Specified*
Process/service or product i…	*Not Specified*	**RTO if Critical/Systemic Proc…**	*Not Specified*	**Downtime for the service/pr…**	*Not Specified*
Reputational damage: Media…	*Not Specified*	**Reputational damage: Sancti…**	*Not Specified*	**Reputational damage: incide…**	*Not Specified*
Reputational damage: regula…	*Not Specified*	**Direct cost (actual figure)**	*Not Specified*	**Direct cost (estimates)**	*Not Specified*
Indirect cost (actual figure)	*Not Specified*	**Indirect cost (estimates)**	*Not Specified*	**Tier 1 capital**	*Not Specified*
Consolidated Tier 1 capital	*Not Specified*	**Is estimated financial impact…**	*Not Specified*	**Reporting to national CERT/C…**	*Not Specified*
Other reasons for reporting t…	*Not Specified*	**Specify reason for reporting …**	*Not Specified*	**Incident Classification**	Significant
Submit to Your National Com…	no	**Submit to European Central …**	no	**Submit to Your Responsible …**	no
Submit to Your National Data…	no	**Submit to Your National NIS …**	yes	**Submit to Your National CSIRT**	no
RiskAssesment	*Not Specified*	**Assets**	*Not Specified*	**NewVulnerability**	*Not Specified*

Description

Introduce here additional information not provided in the incident template or through the questionnaires included in the dashboard. This description will not be included in the mandatory report templates generated.

Responder Jobs

Status	Name	Cortex Server	Date	Actions
Success	Incident Reporting Data Conversor_2_0	local	03/03/25 15:38	🔍
Success	Incident Reporting Event Classifier_2_0	local	03/03/25 15:31	🔍

Figure 166. Status of Responders Jobs.

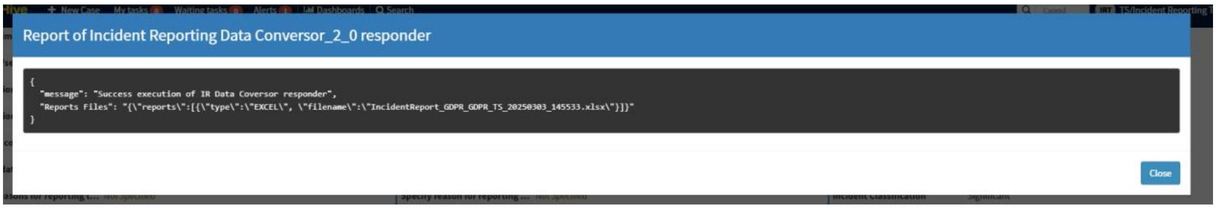
Report of Incident Reporting Data Conversor_2_0 responder

```
{
  "message": "Success execution of IR Data Coversor responder",
  "Reports Files": "{\"reports\":[{\"type\":\"EXCEL\", \"filename\":\"IncidentReport_GDPR_GDPR_TS_20250303_145533.xlsx\"}]}"
}
```

Figure 167. Result of a Responder Job.

An email is sent to the address configured for the user who executed the responder job with the file generated attached (an example is shown in Figure 168).

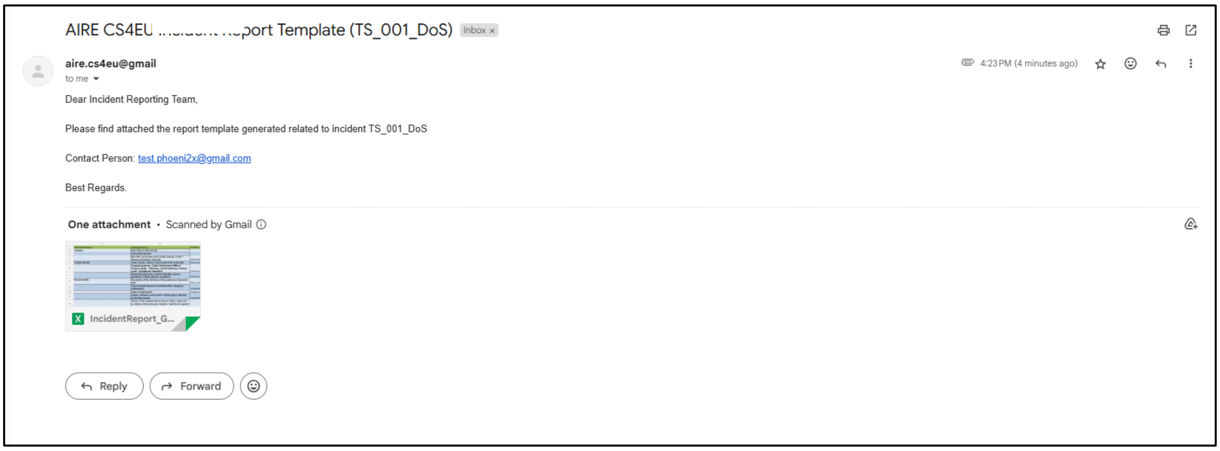
AIRE CS4EU Incident Report Template (TS_001_DoS)

aire.cs4eu@gmail
to me
4:23 PM (4 minutes ago)

Dear Incident Reporting Team,

Please find attached the report template generated related to incident TS_001_DoS

Contact Person: test.phoeni2x@gmail.com

Best Regards.

One attachment • Scanned by Gmail

IncidentReport_G...

Figure 168. Email sent with the report generated.

All the reports are also available in the dashboard under *Reports* tab in *Ready Data Conversion* tab*,* as shown in Figure 169.

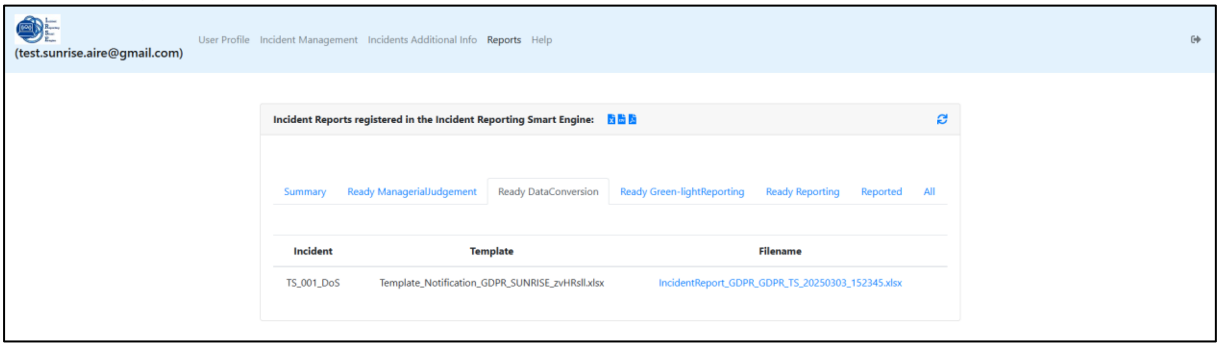
(test.sunrise.aire@gmail.com)

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Template	Filename
TS_001_DoS	Template_Notification_GDPR_SUNRISE_zvHRsll.xlsx	IncidentReport_GDPR_GDPR_TS_20250303_152345.xlsx

Figure 169. List of available reports.

If the file must be modified to change or add information, it can be modified, and the latest version of the documents can be uploaded to the platform from the Reports section. This upload functionality is available from the menu “All” under the “Reports” menu. The name of the file selected to upload to the platform must be the same as one of the existing files, as shown in Figure 170.

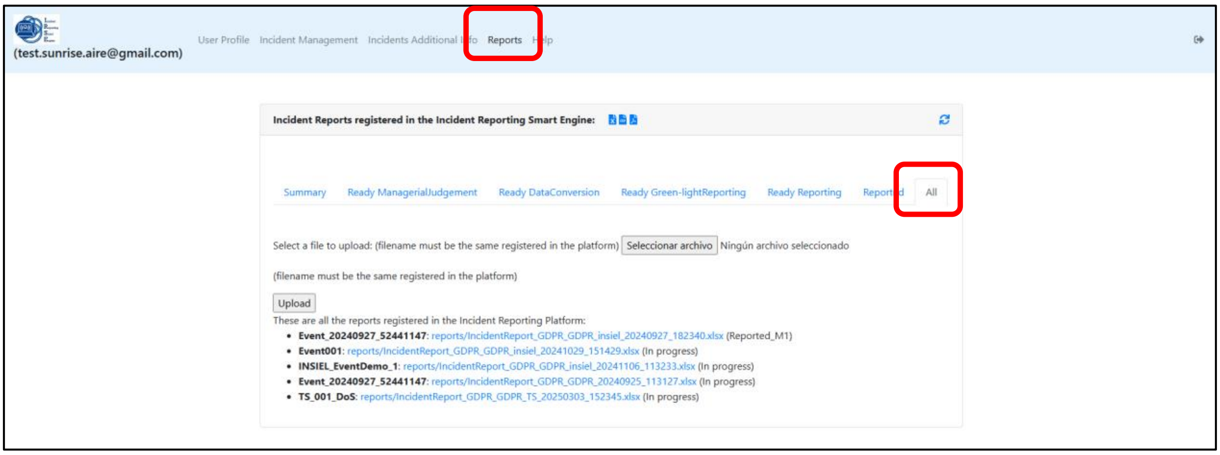

Incident Reports registered in the Incident Reporting Smart Engine:

Select a file to upload: (filename must be the same registered in the platform) Ningún archivo seleccionado

(filename must be the same registered in the platform)

These are all the reports registered in the Incident Reporting Platform:

- **Event_20240927_52441147**: reports/IncidentReport_GDPR_GDPR_insiel_20240927_182340.xlsx (Reported_M1)
- **Event001**: reports/IncidentReport_GDPR_GDPR_insiel_20241029_151429.xlsx (In progress)
- **INSIEL_EventDemo_1**: reports/IncidentReport_GDPR_GDPR_insiel_20241106_113233.xlsx (In progress)
- **Event_20240927_52441147**: reports/IncidentReport_GDPR_GDPR_20240925_113127.xlsx (In progress)
- **TS_001_DoS**: reports/IncidentReport_GDPR_GDPR_TS_20250303_152345.xlsx (In progress)

Figure 170. Upload modified reports menu.

The new file will be registered with the same name but will automatically add a suffix. This way, the users will always be able see the latest version of the reports and can identify which ones were generated automatically by the platform (they end with the timestamp <yyyymmdd_HHMMSS>) or have been modified (they end with <yyyymmdd_HHMMSS> followed by _<suffix>), as shown in Figure 171.

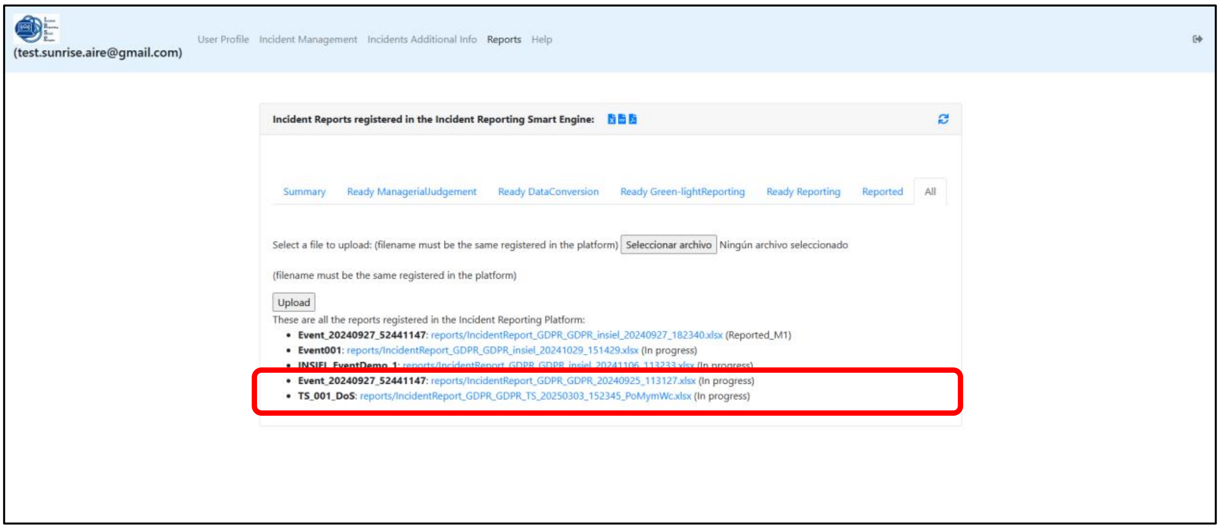

Incident Reports registered in the Incident Reporting Smart Engine:

Select a file to upload: (filename must be the same registered in the platform) Ningún archivo seleccionado

(filename must be the same registered in the platform)

These are all the reports registered in the Incident Reporting Platform:

- **Event_20240927_52441147**: reports/IncidentReport_GDPR_GDPR_insiel_20240927_182340.xlsx (Reported_M1)
- **Event001**: reports/IncidentReport_GDPR_GDPR_insiel_20241029_151429.xlsx (In progress)
- **INSIEL_EventDemo_1**: reports/IncidentReport_GDPR_GDPR_insiel_20241106_113233.xlsx (In progress)
- **Event_20240927_52441147**: reports/IncidentReport_GDPR_GDPR_20240925_113127.xlsx (In progress)
- **TS_001_DoS**: reports/IncidentReport_GDPR_GDPR_TS_20250303_152345_PoMymWc.xlsx (In progress)

Figure 171. Ready Green-light Reporting list.

Once the report has been completed and reviewed by the CI Legal and Compliance Team (IRT group in AIRE), the user will close the associated task *Data Conversion* and a new task, *Green-light for Reporting*, will be assigned to the controller user, as shown in Figure 172.

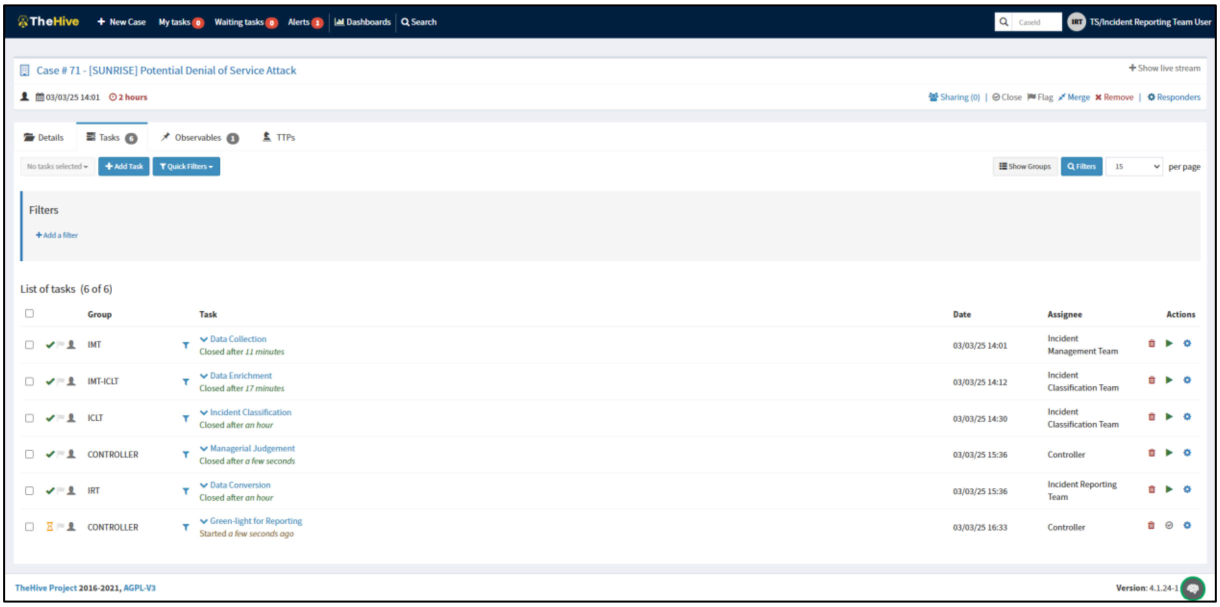

Case # 71 - [SUNRISE] Potential Denial of Service Attack

03/03/25 14:01 2 hours

List of tasks (6 of 6)

Group	Task	Date	Assignee
IMT	Data Collection Closed after *11 minutes*	03/03/25 14:01	Incident Management Team
IMT-ICLT	Data Enrichment Closed after *17 minutes*	03/03/25 14:12	Incident Classification Team
ICLT	Incident Classification Closed after *an hour*	03/03/25 14:30	Incident Classification Team
CONTROLLER	Managerial Judgement Closed after *a few seconds*	03/03/25 15:36	Controller
IRT	Data Conversion Closed after *an hour*	03/03/25 15:36	Incident Reporting Team
CONTROLLER	Green-light for Reporting Started *a few seconds ago*	03/03/25 16:33	Controller

Figure 172. New Green-light for Reporting Task.

A new register is added into the *Summary* tab under the *Reports* tab, as shown in Figure 173.

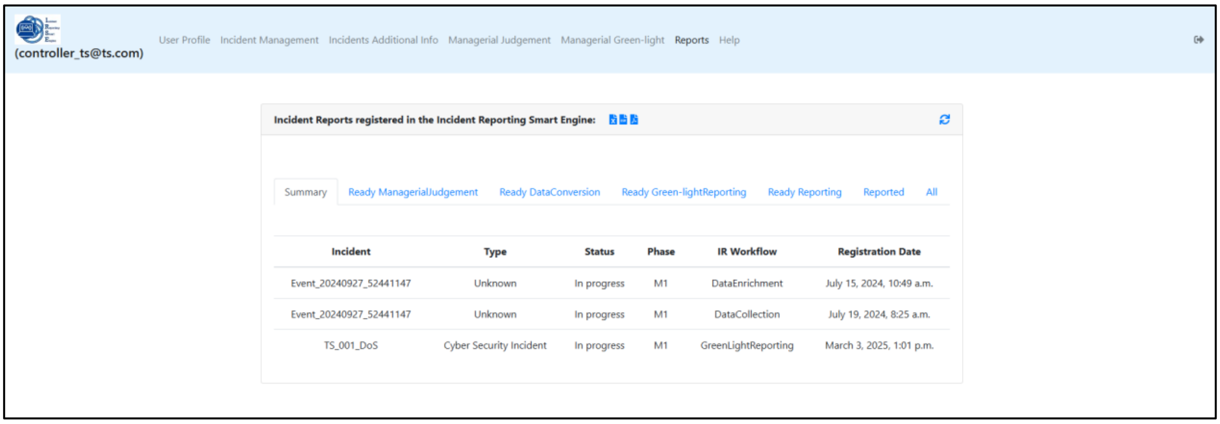

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Type	Status	Phase	IR Workflow	Registration Date
Event_20240927_52441147	Unknown	In progress	M1	DataEnrichment	July 15, 2024, 10:49 a.m.
Event_20240927_52441147	Unknown	In progress	M1	DataCollection	July 19, 2024, 8:25 a.m.
TS_001_DoS	Cyber Security Incident	In progress	M1	GreenLightReporting	March 3, 2025, 1:01 p.m.

Figure 173. Register of Green Light Report.

In the *Ready Green-light Reporting* tab, under *Reports* tab, there is a list of all ready reports, as shown in Figure 174.

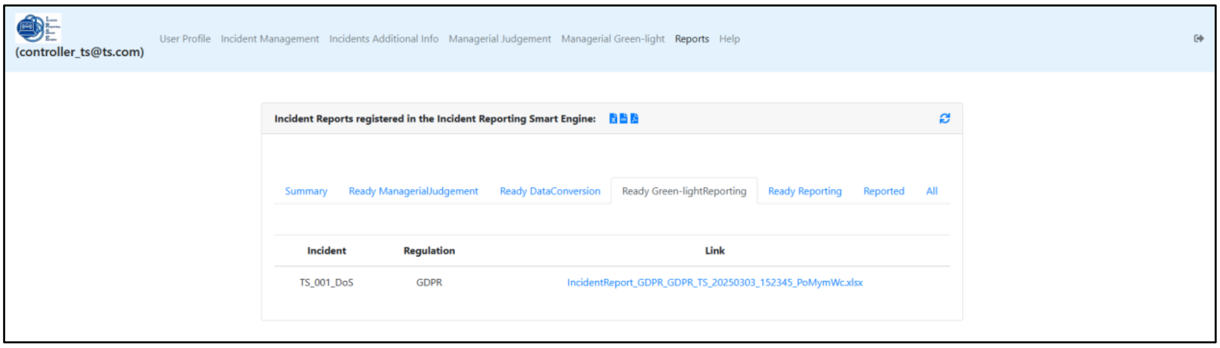

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Regulation	Link
TS_001_DoS	GDPR	IncidentReport_GDPR_GDPR_TS_20250303_152345_PoMymWc.xlsx

Figure 174. List of Ready Green-light Reporting.

5.12 Close incident reporting process

Once it has been received the approval for the Green-Light Reporting task, it is automatically created a task *Reporting & Release* in TheHive. When this task is closed because the notification has been sent, depending on the regulation and the phases included in it (only one or a first report, and intermediate and a final one, for example), the report will appear in the dashboard under *Reports* tab in the status *Reported_M1* (as shown in Figure 175), changing the phase to *M2* and the workflow to a new *Enrichment*, or the incident is closed and it will appear as *Reported*.

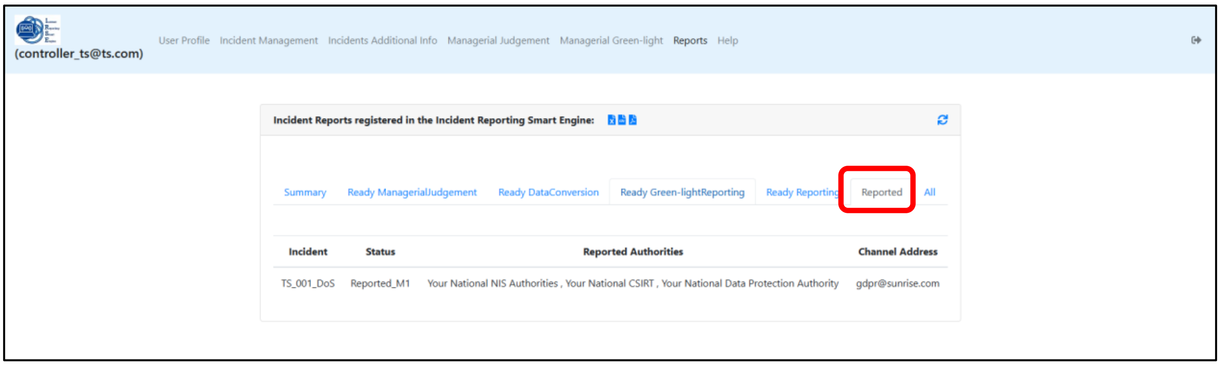

Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Status	Reported Authorities	Channel Address
TS_001_DoS	Reported_M1	Your National NIS Authorities , Your National CSIRT , Your National Data Protection Authority	gdpr@sunrise.com

Figure 175. Reports already reported and closed.

If the regulation has different iterations, a new task “Data Enrichment” will be opened in TheHive to enrich the information about the incident for the Interim Report (M2). The cycle will be repeated to generate the Interim and Final reports depending on the regulations selected as active for the entity and the last phase configured.

In case there is a delay in the reporting process and the reports have not been reported and released on the time configured for some regulations, a notification by email will be sent to the email configured as Contact User in the incident template.

A notification also be sent to the Contact User in case some exception is detected regarding the Incident reporting workflow. For example, if a user without permissions is closing a task which is not assigned to that profile/role.

5.13 REST API

The process of incident management and reporting to authorities can be automated through the AIRE APIs, which allow starting, updating, and closing the enforcement processes. Furthermore, the incident management submodule, TheHive, can receive alerts from SIEMs within the system and IoCs from MISP instances, which contain relevant information about events occurred in other related companies.

Using the REST API, other systems can interact with AIRE’s security incident reporting service, advising of finished tasks, or demanding the end of pending tasks. Users can also consult security incident information and classify the incidents, which must be validated by managerial judgement. Furthermore, they can report to the competent authority. The action that a user can perform depends on its role and the workflow stage of the item. The **aire-workflow-enforcement** service supports the creation of new incidents and assignment of tasks to the different users, depending on their role. The API contains the following functions:

Table 2: aire-workflow-enforcement REST API

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.

HTTP Method	URI	Description
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 for an incident.
POST	/aire/green_light	Submit green-light managerial judgement

The **aire-reports-generator service** transforms the information on security incidents to the different report templates. Then, they are sent to the Competent Authorities based on the regulations. The following table summarizes the API offered by the aire-reports-generator service:

Table 3. aire-reports-generator service REST API

HTTP Method	URI	Description
GET	/aire/generateReports/{incidentId}	Generate report templates for a specific incident.
GET	/aire/earlyNotification/{incidentId}	Send an early notification to the Supervisory Authorities

The **aire-thehive-plugin** service is a middle layer that uncouples the incident management and response tool of organizations from the AIRE engine. It catches actions executed by users inside TheHive and calls the associated actions from AIRE engine, Furthermore, it supports a REST API endpoint to execute actions in TheHive or responders. Such as, checking the user authorization for an action or launching DataConversor Responder, which generates a report of the incident.

Table 4: aire-thehive-plugin service REST API

HTTP Method	URI	Description
POST	/aire/webhook-collector	TheHive Webhook Collector
POST	/aire/executeIR-action	Execute action requests on TheHive
POST	/aire/checkAuthorization	Check authorization for a TheHive Responder execution
GET	/aire/generateReport/{caseId}	Generate report templates for a specific incident.
GET	/aire/earlyNotification/{caseId}	Trigger an early notification to Supervisory Authority for a specific incident.

TheHive features its APIs[8] to control the distinct parts: Organizations, Alerts, Cases, Tasks, etc. This documentation focuses on APIs that are used in automatic mode by other components, such as create new alerts, update a case, or close a task. For this reason, the list of functions is not exhaustive.

- Cases:

Table 5: The Hive REST API for Cases

HTTP Method	URI	Description
POST	/api/case	Create a Case
PATCH	/api/case/{id}	Update a Case
DELETE	/api/case/{id}?force=1	Permanently delete a Case
POST	/api/v0/case/{id1}/_merge/{id2}	Merge two Cases in a single Case
POST	/api/v0/query	List alerts merged in a Case; case ID passed in Request Body

- Alerts:

Table 6: The Hive REST API for Alerts

HTTP Method	URI	Description
POST	/api/v1/query?name=alerts	List of Alerts
POST	/api/alert	Create an Alert
DELETE	/api/alert/{id}?force=1	Delete an Alert
PATCH	/api/alert/{id}	Update an Alert
POST	/api/alert/{id1}/merge/{id2}	Merge an Alert into an existing Case
POST	/api/alert/{id}/createCase	Promote an Alert as a new Case

- Tasks:

Table 7: The Hive REST API for Tasks

HTTP Method	URI	Description
POST	/api/v0/query	List Tasks of a case; case ID passed in Request Body
POST	/api/case/{id}task	Create a Task
PATCH	/api/case/task/{id}	Update a Task
GET	/api/case/task/{id}	Get Task of a case
POST	/api/v0/query	List all waiting Tasks

5.14 OTHER APIs

TheHive can receive security events related to the current infrastructure from other systems, such as alerts from Wazuh SIEM or IoC from MISP instances. It registers the security events and displays them in a dashboard, where users can monitor the system and transform the alerts into incidents with a button on the view.

On the one hand, TheHive can receive the alerts generated and sent by Wazuh to a Kafka broker[9]. Wazuh must send the alerts to the topic *wazuh-alerts*. Then, these alerts are transformed from wazuh format to TheHive alert format and registered in its API.

On the other hand, TheHive can monitor the IoC from MISP instances and shows them in the alert dashboard (see an example in Figure 139). To enable this feature, it is necessary to configure the tool, adding in the application.config file the required information. An example of the configuration is provided in the deployment section 2.3.7.

6 References (Annex I)

[1] https://www.elastic.co/guide/en/elasticsearch/reference/current/getting-started-index-lifecycle-management.html

[2] https://regex101.com

[3] https://en.wikipedia.org/wiki/ISO_8601

[4] **Corte Analyzers Repository**. https://github.com/TheHive-Project/Cortex-Analyzers (last access: 14/04/2025)

[5] https://www.elastic.co/guide/en/beats/filebeat/7.17/index.html

[6] https://docs.oracle.com/javase/7/docs/api/java/text/SimpleDateFormat.html

[7] TheHive http://thehive-project.org/

[8] http://docs.thehive-project.org/thehive/api/

[9] https://kafka.apache.org/