



SUNRISE

Strategies and Technologies for **United** and **Resilient** Critical Infrastructures
and Vital **S**ervices in Pandemic-Stricken **E**urope

D6.2 Cyber-physical resilience tool and training guide V1

Document Identification			
Status	Final	Due Date	30/09/2023
Version	1.0	Submission Date	30/09/2023

Related WP	WP6	Document Reference	D6.2
Related Deliverable(s)	D6.1	Dissemination Level (*)	PU
Lead Participant	XLB	Lead Author	Tomaž Martinčič (XLB)
Contributors	ATS, INS, CAF, PIL, SZ, TS	Reviewers	Olga Segou (INT)
			Milan Tarman (ICS)

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.**

This document and its content are the property of the SUNRISE Consortium. The content of all or parts of this document can be used and distributed provided that the SUNRISE project and the document are properly referenced.

Each SUNRISE Partner may use this document in conformity with the SUNRISE Consortium Grant Agreement provisions.

(*) 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
Tomaž Martinčič	XLB
Pablo de Juan	ATS
Miguel Martín	ATS
Susana González	ATS
Gilda De Marco	INS
Daniele Fabro	CAF
Blaž Jemenšek	PIL
Tomaž Ramšak	SZ
Andreja Markun	TS

Document History			
Version	Date	Change editors	Changes
0.1	24/07/2023	Tomaž Martinčič (XLB)	First version of the document (ToC)
0.2	18/08/2023	Tomaž Martinčič (XLB)	Initial inputs to sections 2.2., 3.1, and Annex I.
0.3	24/08/2023	Tomaž Martinčič (XLB)	Refining the initial inputs from v0.2.
0.4	28/08/2023	Tomaž Martinčič (XLB)	Adding deployment sections and improving the overall quality of the document.
0.5	05/09/2023	Pablo de Juan (ATS)	Input for the remaining sections.
0.6	18/09/2023	Tomaž Martinčič (XLB)	Addressing comments from the internal review.
0.7	19/09/2023	Pablo de Juan (ATS)	Refinement of Annex I.
0.8	20/09/2023	Tomaž Martinčič (XLB)	Merging Annex I updates and fixing formatting.
0.9	25/09/2023	Tomaž Martinčič (XLB)	Added missing references.
0.10	28/09/2023	Tomaž Martinčič (XLB)	Addressing comments from PC review.
0.11	29/09/2023	Antonio Álvarez (ATS)	Final document layout
1.0	29/09/2023	Antonio Álvarez (ATS)	FINAL VERSION TO BE SUBMITTED

Quality Control		
Role	Who (Partner short name)	Approval Date
Deliverable leader	Tomaž Martinčič (XLB)	25/09/2023
Quality manager	Juan Andrés Alonso (ATS)	30/09/2023
Project Coordinator	Antonio Álvarez (ATS)	30/09/2023

Table of Contents

Document Information.....	2
Table of Contents	3
List of Tables.....	5
List of Figures.....	6
List of Acronyms	9
Executive Summary	10
1 Introduction.....	11
1.1 Purpose of the document	11
1.2 Relation to other project work.....	11
1.3 Structure of the document	11
2 Cyber-Physical Resilience Tool	12
2.1 General Context	12
2.2 Architecture	12
2.2.1 Anomaly Detection.....	13
2.2.2 Threat Intelligence.....	15
2.2.3 Risk Assessment	16
2.2.4 Incident Reporting.....	18
2.2.5 Dashboard	19
2.3 Deployment.....	20
2.3.1 Pre-requisites	20
2.3.2 Anomaly Detection Module	20
2.3.3 Threat Intelligence Module	21
2.3.4 Risk Assessment Module	22
2.3.5 Incident Reporting Module	23
3 Cyber-Physical Resilience Tool Methods Tested in Labs	27
3.1 Monitoring System (LOMOS)	27
3.1.1 Data	27
3.1.2 Methods	27
3.1.3 Methods evaluation	28
3.1.4 Throughput.....	30
3.2 Threat Intelligence (TINTED)	31
3.3 Risk Assessment (CERCA)	31
3.4 Incident Reporting (AIRE).....	31
4 Management and What-If Analysis	32
5 Pilot trials execution (feasibility analysis).....	33

5.1 Proofs of concepts	33
5.1.1 Italy: Public Administration	33
5.1.2 Italy: Water.....	33
5.1.3 Slovenia: Telecommunication	34
5.1.4 Slovenia: Transport.....	34
6 Conclusions.....	35
References.....	36
Annex I Training guide and user manual.....	37
I.I Anomaly Detection	37
I.I.I Setting up data sources	37
I.I.II Training a log parser	37
I.I.III Inspecting the log parsing results.....	40
I.I.IV Training an anomaly detection model	41
I.I.V Inspecting the training results.....	43
I.I.VI Setting up live inference.....	43
I.I.VII Inspecting live inference results.....	45
I.II Threat Intelligence	46
I.II.I WEB-GUI.....	46
I.II.II API.....	50
I.III Risk Assessment.....	54
I.IV Incident Reporting.....	58
I.IV.I WEB-GUI.....	59
I.IV.I.I AIRE GUI – ADMINISTRATOR	59
I.IV.I.II AIRE GUI – USER	69
I.IV.II INPUT METHODS	96
I.IV.II.I REST API.....	96
I.IV.II.II OTHER APIs.....	98

List of Tables

<i>Table 1: Empirical evaluation of log-based anomaly detection methods.</i>	29
<i>Table 2: Hardware resources used in the throughput experiments.</i>	30
<i>Table 3: Train log template parsers on 1 million logs from the BGL dataset execution time measurements.</i>	30
<i>Table 4: Train anomaly detection model on 1 million logs from the BGL dataset execution time measurements.</i>	30
<i>Table 5: Anomaly detection inference on 1 million logs from the BGL datasets execution time measurements.</i>	31
<i>Table 6: aire-workflow-enforcement REST API.</i>	96
<i>Table 7. aire-reports-generator service REST API.</i>	96
<i>Table 8: aire-thehive-plugin service REST API</i>	97
<i>Table 9: The Hive REST API for Cases.</i>	97
<i>Table 10: The Hive REST API for Alerts</i>	97
<i>Table 11: The Hive REST API for Tasks</i>	97

List of Figures

Figure 1. Overall architecture of the CPR tool	12
Figure 2. Anomaly Detection Module overview.	13
Figure 3. LOMOS architecture.	14
Figure 4. TINTED architecture.....	15
Figure 5. CERCA architecture.....	17
Figure 6. AIRE architecture.....	18
Figure 7. A sample of raw BGL logs.	37
Figure 8. Regular expressions for extracting messages and timestamps from raw logs and the timestamp format.	37
Figure 9. Regular expression for extracting log message from a raw log.	38
Figure 10. Regular expression for extracting timestamp from a raw log.	38
Figure 11. Drain parameters.	38
Figure 12. Custom masks for complex parameters.	39
Figure 13. Log parser training Elasticsearch connection details and credentials.....	39
Figure 14. Log parser training source index.	39
Figure 15. Log parser training period selection and addition filters.	40
Figure 16. Extracted log templates overview.	41
Figure 17. Log template details.	41
Figure 18. Model training Elasticsearch connection details, credentials, and filters.	42
Figure 19. Training data intervals.	42
Figure 20. Anomaly detection model training hyper-parameters.	43
Figure 21. Training anomaly detection loss chart displayed in MLflow.	43
Figure 22. Pretrained parser MLflow experiment id.....	44
Figure 23. Set the inference schedule period.....	44
Figure 24. Anomaly detection inference Elasticsearch endpoint details and credentials.....	44
Figure 25. Elasticsearch index configuration and data filters.	44
Figure 26. Inference model configuration.	45
Figure 27. Histogram of logs through time (e.g., per day).	45
Figure 28. Average anomaly score.	45
Figure 29. Number of logs with high anomaly score (e.g., above 0.7).....	46
Figure 30. Table of logs with anomaly scores.	46
Figure 31. TINTED Login.	46
Figure 32. System Configuration.	47
Figure 33. Sharing Configuration.....	47
Figure 34. Events on the main page (currently empty).	47
Figure 35. Information Sharing Form.	48
Figure 36. MISP instance.	48
Figure 37. Individual Event details – MISP.	48
Figure 38. TINTED Information Received dashboard.....	49
Figure 39. Individual Event details – TINTED.....	49
Figure 40. Administrator Management Dashboard I.	49
Figure 41. Administrator Management Dashboard II.	50
Figure 42. Selection of multiple events in MISP instance.	50
Figure 43. Download of events in MISP JSON format.....	51
Figure 44. HTTP request to download MISP events.....	51
Figure 45. Orchestrator HTTP request format.....	51
Figure 46. Sharing Agreement field.....	52
Figure 47. MISP instance field.	52

Figure 48. User policies for privacy sharing.....	52
Figure 49. TIE architecture.	53
Figure 50. MISP events enriched with TIE’s threat score.....	54
Figure 51. Vulnerability object.	54
Figure 52. Login form	54
Figure 53. User profile menu	55
Figure 54. Legal entities configuration menu.....	55
Figure 55. Data processing activities configuration	55
Figure 56. Asset view dashboard.....	56
Figure 57. Model configuration dashboard.....	56
Figure 58. Risk model selection dashboard.....	57
Figure 59. Model configuration dashboard updated with risk model.....	57
Figure 60. Questionnaire for risk model.....	58
Figure 61. Risk Report summary.....	58
Figure 62. AIRE naviagtion bar.....	59
Figure 63. Entities Configuration List	59
Figure 64. New entity form.....	60
Figure 65. Regulation for entities.....	60
Figure 66. List of registered users.....	61
Figure 67. List of configured contacts	61
Figure 68. List of regulations registered.....	62
Figure 69. New regulation form	62
Figure 70. List of timers.....	63
Figure 71. Timer edition form.....	63
Figure 72. Example of email sent by AIRE	64
Figure 73. List of Recipient Configurations.....	64
Figure 74. Recipient Edition Form.....	64
Figure 75. List of communication channels	65
Figure 76. Channel Edition.....	65
Figure 77. List of Templates.....	66
Figure 78. Template Edition Form	66
Figure 79. Reported Authorities Configuration	67
Figure 80. Reported Authorities Editor.....	67
Figure 81. List of Criterias.....	68
Figure 82. Criteria Edition.....	68
Figure 83. Help menu	68
Figure 84. TheHive login interface.....	69
Figure 85. TheHive embedded in AIRE interface	69
Figure 86. Template selection for a new Case.....	70
Figure 87. New case form with Incident Respot template	70
Figure 88. New incident in the list of cases	70
Figure 89. Details of a case.....	71
Figure 90. Tasks of associated with a case.....	71
Figure 91. List of Reports.....	72
Figure 92. Task list of a case.....	72
Figure 93. List of incidents registered.....	73
Figure 94. Incident Additional Information Edition	73
Figure 95. Security Event Lifecycle.....	74
Figure 96. List of Essential Services	74
Figure 97. Essential Services edition.....	74
Figure 98. List of Trust Services Assets registered.....	75

Figure 99. Trust Services Asset edition 75

Figure 100. List of Trust Services registered 75

Figure 101. Trust Services affected edition 76

Figure 102. List of Impacted Processes..... 76

Figure 103. Processes Affected edition..... 76

Figure 104. List of Personal Data Breaches registered..... 77

Figure 105. Personal Data Breach Edition..... 77

Figure 106. List of observables 78

Figure 107. Create new observable form 78

Figure 108. Run analyzers option 78

Figure 109. Selection of the analyzers to run 79

Figure 110. Automated creation of Incident Classification task..... 79

Figure 111. Detail of a task..... 80

Figure 112. List of Incident Reports registered..... 80

Figure 113. Responders in an Incident Detail View 81

Figure 114. Selection of one responder to run for an incident case 81

Figure 115. List of Responder Jobs with status..... 82

Figure 116. Output example for after run a responder 82

Figure 117. Tags of an Incident Case..... 83

Figure 118. Fields of a template 83

Figure 119. Output of a responder without permission 83

Figure 120. Automatic creation of Managerial Judgement Task 84

Figure 121. Automatic creation of Managerial Judgement IR Workflow..... 84

Figure 122. List of Incident ready for managerial judgement 85

Figure 123. Form for a managerial judgement 85

Figure 124. New Data Conversion Task..... 86

Figure 125. Register of Data Conversion Task..... 86

Figure 126. Run responder invokation 87

Figure 127. Status of Responders Jobs 87

Figure 128. Result of a Responder Job..... 87

Figure 129. Email to Incident Reporting Team user 88

Figure 130. List of available reports 88

Figure 131. Upload modified reports menu 89

Figure 132. Ready Green-light Reporting list 89

Figure 133. New Green-light for Reporting Task..... 90

Figure 134. Register of Green Light Report 90

Figure 135. List of Ready Green-light Reporting 91

Figure 136. List of Incident ready for managerial green-light for reporting 91

Figure 137. Configuration of green-light reporting form. 91

Figure 138. Confirmation of task closed..... 92

Figure 139. New Reporting and Release Task 92

Figure 140. Reporting And Release register in Summary of Reports 93

Figure 141. Reporting and Release actions 93

Figure 142. New Date Enrichment Task 94

Figure 143. Consult of Security Event Lifecycle..... 94

Figure 144. Security Event Lifecycle..... 95

Figure 145. Email to Contact User reminding of pending reports 95

Figure 146. Notification of closing task without permission 95

List of Acronyms

Abbreviation / acronym	Description
API	Application Programming Interface
BERT	Bidirectional Encoder Representations from Transformers
BGL	BlueGene/L supercomputer
BPMN	Business Process Model and Notation
CERCA	CybEr Risk assessment CAculator
CI	Critical Infrastructure
CPR	Cyber Physical Resilience
CPU	Central Processing Unit
CTI	Cyber Threat Intelligence
CUDA	Compute Unified Device Architecture
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
DB	Database
EC	European Commission
EC2	(Amazon) Elastic Compute Cloud
FN	False Negative
FP	False Positive
FTP	File Transfer Protocol
GB	Gigabyte
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HDD	Hard Disk Drive
HDFS	Hadoop Distributed File System
HTTP	Hypertext Transfer Protocol
ICT	Information and Communications Technology
IP	Internet Protocol (address)
JSON	JavaScript Object Notation
MISP	Malware Information Sharing Platform
REST	Representational State Transfer
SIEM	Security Information and Event Management
SQL	Structured Query Language
SSD	Solid State Drive
TIE	Threat Intelligence Engine
TP	True Positive
VM	Virtual Machine
VPN	Virtual Private Network
WP	Work Package

Executive Summary

This deliverable publicly presents the concept of the cyber-physical resilience (CPR) tool, which was initially presented in deliverable D6.1. The CPR tool is the main output of Work Package 6 (WP6). The goal of the tool is to improve critical infrastructure (CI) cyber-physical security capabilities under extreme situations such as a pandemic. The tool will be validated by multiple CI deployments from different sectors.

The deliverable starts with a description of CPR tool architecture and deployment. There are four modules in the tool. Namely, the anomaly detection, threat intelligence, risk assessment, and incident reporting modules. The modules will be integrated to improve CI cyber-physical security.

The anomaly detection module underwent testing in a lab environment with publicly available data relevant to the use cases and commonly used in the domain of log-based anomaly detection literature. The most relevant indicators are its speed and predictive accuracy. Log-based anomaly detection system showed a high throughput, processing over two thousand log messages per second. Regarding predictive accuracy on the benchmark datasets, it showed a very high recall rate, exceeding 0.9 in some cases and staying above 0.8 in others. However, precision proved to be more sensitive to the datasets and their split methods. In the optimal scenario, our system achieved a precision score of 0.93. The remaining modules were not evaluated during the first year of the project, marked with lower priority as they are not artificial intelligence (AI) tools which need to be measured with different metrics. In the following, as the integration of the different modules within the CPR tool will be progressing, the evaluation will be extended until completing the whole tool, being able to test the whole performance.

Next, possible legal issues are presented and addressed. There are no known problems in the current state of the pilots. This section will be revised in the following iterations of the Cyber-physical resilience tool and training guide.

Following that, the four pilot trials are presented. The pilot trials are focusing on CI and will be used to validate the tool with data from CI partners in the SUNRISE project. Two of the pilots cover public administration and water use cases in Italy. The other two pilots cover telecommunication and transport use cases in Slovenia.

Finally, the user manual for the modules is presented. It covers a step-by-step guide to set up and manage the modules. The manual is enhanced with screenshots from actual demo deployments.

The main outputs of this deliverable are testing and user manual. This is the first version of the cyber-physical resilience tool and training guide, which will be further improved in the following iterations D6.3 and D6.5. Deliverables D6.4 and D6.6 will present a cyber-physical resilience pilot report.

1 Introduction

1.1 Purpose of the document

The purpose of D6.2 is to publicly present the cyber-physical resilience (CPR) tool, initially presented in D6.1 - *Cyber-physical resilience conceptualization* [1]. The deliverable presents the tool's architecture and deployment process. Next, it showcases the outcomes of the tests conducted on the submodules in a lab environment on public datasets. The modules were tested with regard to predictive performance and throughput, both of which hold significant importance in the assessment of cybersecurity tools relevant to critical infrastructure (CI). After that, possible legal issues and pilot trials are discussed. Finally, it presents a user manual and training guide.

1.2 Relation to other project work

This deliverable reports the results of ongoing tasks *T6.1 Cyber-physical security risk assessment*, *T6.2 AI-powered log monitoring*, and *T6.3 Incident response and threat intelligence sharing*. It provides results of lab environment tests and usage guidelines for the initial cyber-physical resilience tool modules presented in D6.1. This includes anomaly detection (LOMOS), threat intelligence (TINTED), risk assessment (CERCA), and incident reporting (AIRE). The tool follows requirements from D3.1 – *Requirements and designs V1* [2]. Deliverable D6.2 is the initial version of the *Cyber-physical resilience tool and training guide*, which will be iterated in D6.3 and D6.5.

1.3 Structure of the document

This document is structured in 6 major chapters and annex.

Chapter 1 presents the purpose of this (D6.2) deliverable, its relation to other project work, and its structure.

Chapter 2 presents the overview of the Cyber-Physical Resilience Tool. The focus is on the architecture and deployment of the four modules.

Chapter 3 presents the evaluation results of lab experiments. The anomaly detection module was tested in the scope of predictive accuracy and processing speed to measure the performance of its models. On the other hand, the rest of the modules were not tested in the first year as they do not incorporate AI algorithms.

Chapter 4 is dedicated to describing possible issues regarding pilots due to legal restrictions related to CI.

Chapter 5 presents four CI pilot trials related to public administration, water, telecommunication, and transport.

Chapter 6 outlines the main findings of this deliverable and provides conclusions.

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

2 Cyber-Physical Resilience Tool

2.1 General Context

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

- ▶ Anomaly detection module utilizes machine learning for log-based anomaly detection.
- ▶ Threat intelligence module integrates pertinent threat intelligence information from each CI sector.
- ▶ Risk assessment module factors in diverse indicators for conducting risk evaluations.
- ▶ Incident reporting module is responsible for automating the reporting process to the relevant authorities in the event of a security incident.

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

2.2 Architecture

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

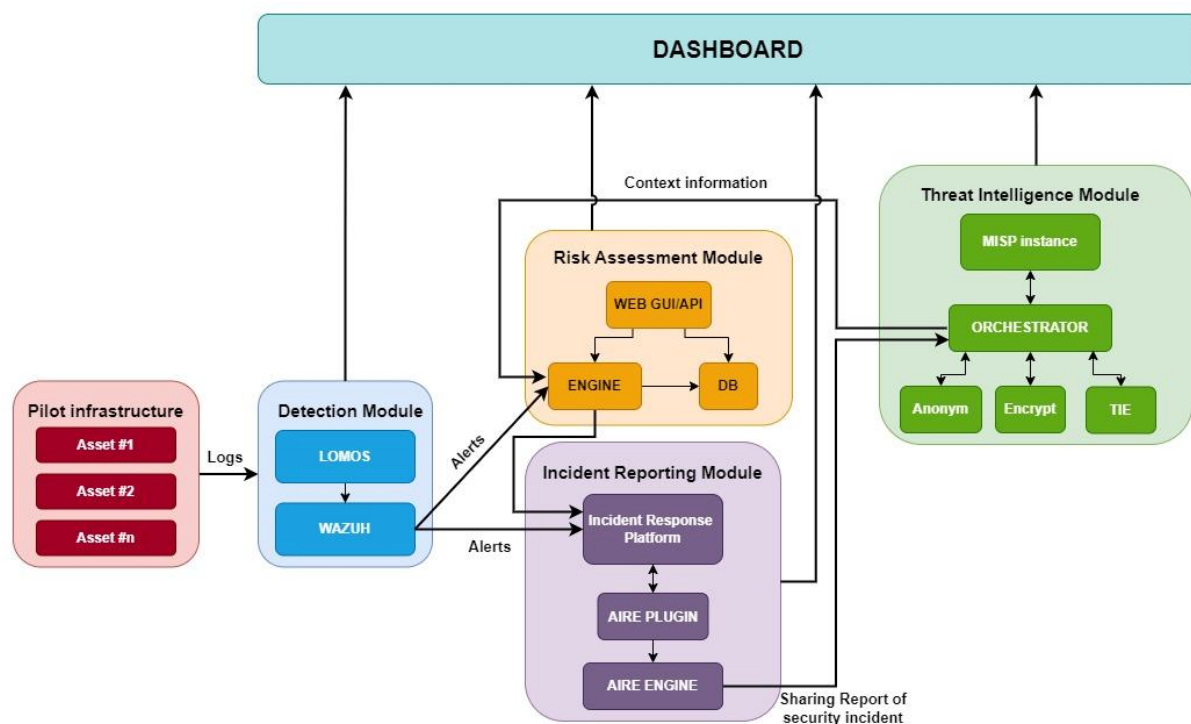


Figure 1. Overall architecture of the CPR tool

The Detection module oversees CI assets and generates security events and low-level alerts. It comprises two main components: an anomaly detector named LOMOS, which uses application logs and raw data to train an AI in recognizing normal system behavior and triggering alerts for deviations, and a SIEM called Wazuh¹, which identifies event sequences related to known threats and raises alarms correlating LOMOS' output. These alerts and events are sent to the Risk Assessment and Incident Reporting Modules.

¹ <https://wazuh.com>

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 incidents based on real-time cyber environment data and the financial impact of protecting different digital assets. It is also sent to the Incident Reporting Module, which evaluates events and alarms from the SIEM to determine security incident presence and severity for potential reporting to authorities. This module also 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. It employs a common MISP² instance to share relevant attack and threat information. After a security incident and notifying authorities, 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.

2.2.1 Anomaly Detection

The main part of the anomaly detection module is LOMOS, which is a self-supervised machine-learning system for log-based anomaly detection [1]. Conceptually the LOMOS workflow consists of four main parts, see Figure 2. The first one is the log parser, which extracts events (log templates) from raw log messages. The second part is the anomaly detector trainer, which trains the model on historical log messages in a self-supervised fashion. The third part is the anomaly detector, which uses the model from the previous step, and new data pre-processed with the log parser, to compute anomaly scores on per per-log basis. The anomaly scores are displayed in the dashboard. A single LOMOS deployment can process multiple log sources with different models.

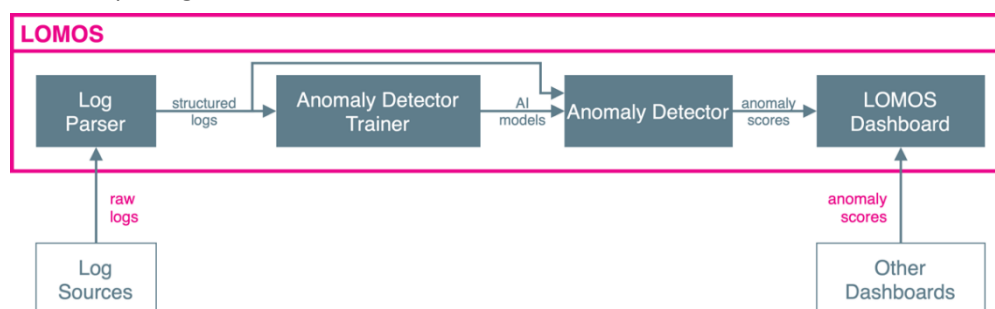


Figure 2. Anomaly Detection Module overview.

Log parser is capable of automatic event extraction from raw log messages. The events are often also called log templates. In LOMOS log templates are extracted with Drain [3]. It is a state-of-the-art online log parsing method, which can parse logs in a streaming and timely manner. Drain uses a fixed depth parse tree, which encodes specially designed parsing rules. The parser can be incrementally trained to add new log templates at any time in the future. However, the anomaly detection model usually needs retraining in such a case.

Log templates are input for the **anomaly detector trainer**. Currently, the LogBERT [4] model is implemented and extended. LogBERT is appropriate method because it is a self-supervised machine-learning method for window log-based anomaly detection. Therefore, it does not need labelled data for training. The log templates, more specifically, their IDs, are grouped into time or index-based windows. Sequences of IDs are input into the BERT-like [5] transformer model. The trained model has a fixed vocabulary, therefore new log templates cannot be added without retraining the model. Because the model is trained in a self-supervised fashion on normal data, it does not need per-log labelling for training. The model is expected to be robust enough so that even a small number of

² <https://www.misp-project.org>

anomalies in training data should not impact its performance significantly. However, the cleaner the training data is, the better the model will be. GPUs are needed to train a model like LogBERT.

The **anomaly detector** employs a model obtained from the anomaly detector trainer. New raw data is first processed with the trained log parser and then passed to the anomaly detector which pulls the trained model from the model registry. Anomaly detection can be executed either on demand or scheduled on a periodical basis. **We extended the idea from the LogBERT method to assign anomaly scores on the per-log level.** When it comes to inference, CPUs are suitable for the majority of use cases. However, GPUs are also supported for inference in case of a need for very high throughput.

The **LOMOS dashboards** provide interactive visualizations of the extracted log templates and logs with anomaly scores. It offers also control panels for model learning and inference. The dashboard guides administrators through the process of training a model and running inference. First, the log parser has to be trained. This is done on a selected period from historical data. The administrator can explore extracted log templates and fine-tune hyper-parameters if needed. In the next step, the administrator trains the model. This is the most hardware resources-consuming step. Afterward, the model is ready to be deployed. The administrator can execute one-time or periodic inference. The latter one is executed based on a given interval. The administrator can monitor logs and anomaly scores in the dashboard. The administrator can also set rules in the dashboard, to get alerts in case of high anomaly scores. The notification about the alert can be sent to an email, Slack, or Microsoft Teams.

The architecture stays the same as presented in D6.1. The architecture diagram is presented in Figure 3. It consists of five main parts. The workflow presented in Figure 2 is distributed through different components. Log data storage is a log source in the workflow diagram. The log parser is contained in the LOMOS Parsing component. LOMOS Anomaly detection is used for anomaly detection trainer and anomaly detector. The jobs are executed by Celery CPU and GPU workers. The results are then presented in the Dashboard. All the components are dockerized, and as such, easy to deploy.

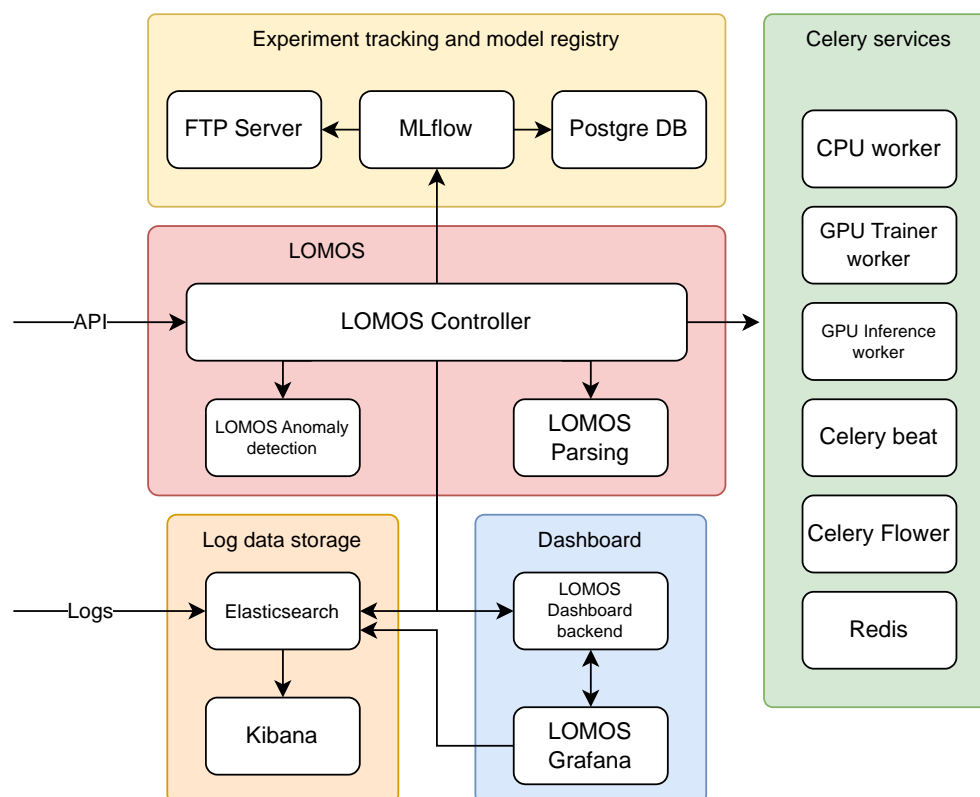


Figure 3. LOMOS architecture.

The main part is **LOMOS** itself, which consists of the LOMOS Controller, LOMOS parsing, and LOMOS anomaly detection modules developed in Python. LOMOS Controller is the main component, which

orchestrates the work. It can execute log parsing and anomaly detection jobs on Celery CPU and GPU workers. The LOMOS Controller exposes REST API, through which it gets requests for training or inference jobs. It is also used to view and manage period jobs.

Celery services are another important part of the system. Celery is used to distribute learning and inference jobs on CPU and GPU workers. Celery includes a web dashboard called Celery Flower, a scheduler Celery Beat, and Redis which is used for communication between the components. Celery and Redis are free and open source.

The third part is dedicated to **experiment tracking and model registry**. It consists of three services. Namely, Postgre DB for storing all the tabular data, FTP server for storing artifacts, and MLflow as a dashboard and access REST API. All of them are free and open source. MLflow offers official Python API which enables developers to seamlessly integrate it in systems. It is used to log training parameters and metrics like loss and other evaluation metrics. Artefacts like weights of the trained model or tabular and visual outputs can also be stored. All the values and artifacts can be retrieved by the experiment's unique identifier.

Logs are stored in the **Log data storage** part, in Elasticsearch indices. Kibana is deployed for managing and exploring Elasticsearch. Elasticsearch is well suited for handling log data, as it can store documents in JSON format. Elasticsearch supports automatic index rollover based on time or index size. The data can be moved through different phases, and later stored in a format which optimized for low storage requirements, rather than for query speed. Old data can also be automatically deleted to keep space for new logs. There is an official Python client for Elasticsearch. The original Elastic is not free anymore for commercial use, so we use the alternative OpenSearch which is licensed under Apache 2.0 and as such free and open source.

The final part is **Dashboard**. The actual dashboard is developed in Grafana, which is a free and open-source analytics and interactive visualization web application. By default, it does not support input forms. We extended it and developed the LOMOS Dashboard backend in Python. Now we support LOMOS configuration forms in the dashboard. This enables users to configure and run the training and inference processes through a web dashboard. The LOMOS dashboard backend passes the requests to the LOMOS controller which executes the jobs.

2.2.2 Threat Intelligence

The threat intelligence module (TINTED) allows the secure exchange of CTI information and provides additional context for the calculation of cyber risk. It is divided into several layers such as the authentication layer, the transport layer, and the privacy and enrichment layer, as shown in Figure 4.

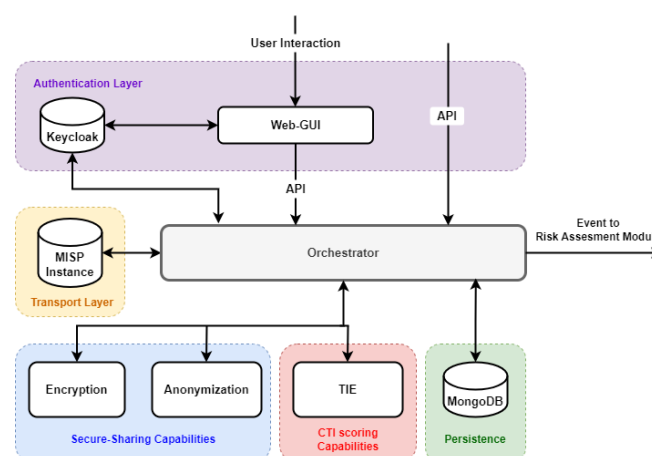


Figure 4. TINTED architecture.

The **MISP Instance** serves as the transport layer of the platform, facilitating the exchange of Cyber Threat Intelligence (CTI) information between different users and organizations. MISP's fundamental

philosophy revolves around the idea of widespread information sharing. However, this approach can introduce risks, as the shared CTI might contain sensitive organizational information that attackers could exploit. TINTED addresses these concerns by offering enhanced sharing capabilities, introducing greater granularity to sharing options. With TINTED, it becomes possible to control who can access specific information and the duration of their access.

Keycloak, an open-source platform, facilitates the efficient and secure implementation of user authentication and authorization processes. Leveraging Keycloak's functionalities, TINTED ensures that user interactions with the platform remain secure and tailored to their appropriate level of access.

MongoDB, as the persistence mechanism within the tool, oversees data storage. This database handles a variety of information, including configuration details for the tool, such as the MISP instance it interacts with and user-specific settings (e.g., user passphrases); user information and data vital for the platform's smooth operation, such as events submitted by specific users; and infrastructure-related information necessary for the proper functioning of the TIE (Threat Intelligence Engine) module.

At the core of TINTED's architecture is the **Orchestrator**. This central component is responsible for facilitating communication among different modules and capabilities of the tool. Additionally, it provides users with an API (Application Programming Interface) to interact with the tool's features and functionalities.

The **WEB-GUI** (Graphical User Interface) acts as an intermediary between users and the Orchestrator API. In addition to simplifying tool usage, the WEB-GUI integrates authentication and authorization mechanisms through Keycloak. This integration ensures secure access to the tool's features and data.

Regarding the privacy-preserving modules, TINTED has the **Encryption module** that introduces the capability to encrypt MISP attribute values and the **Anonymization functionality** which enables the anonymization of specific attribute types.

The **Threat Intelligence Engine** (TIE) constitutes a critical component with two subcomponents: the ZMQClient and the HeuristicEngine. The ZMQClient subscribes to a ZMQ queue within the MISP instance, receiving notifications whenever an event is uploaded. It then forwards these events to the HeuristicEngine, which calculates a corresponding score for each event based on various heuristics.

2.2.3 Risk Assessment

The risk assessment module (CERCA) consists of different submodules, see Figure 5, which finally produce a report that is sent to the incident reporting module, where they are evaluated to determine whether a security incident exists and whether it is critical enough to be reported to the authorities.

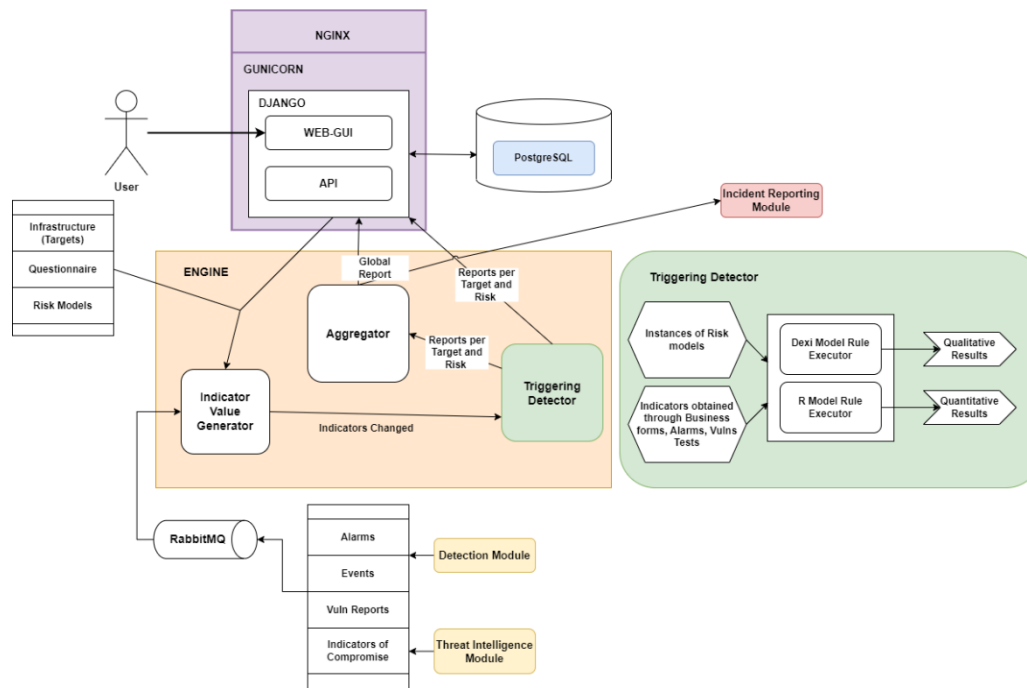


Figure 5. CERCA architecture.

The **WEB-GUI/API** submodule plays a crucial role in facilitating tool configuration. Users can configure the tool by completing various forms that provide details about the entity's infrastructure and pertinent business information.

The **Indicator Value Generator** submodule operates in tandem with the configuration generated by users through the WEB-GUI/API submodule. It also leverages information from the RabbitMQ queue, which includes alarms, events, vulnerability reports, and indicators of compromise sourced from the threat intelligence module. With this amalgamation of data, the submodule either generates new indicators or modifies existing ones, depending on whether it's the initial iteration or not. Once this stage concludes, the generated indicators are dispatched to the Triggering Detector.

The **Triggering Detector** submodule comes into play when changes occur in the indicators or risk models. It triggers the activation of the Risk Model Executors.

The **Risk Model Executors** submodule is responsible for executing algorithms implemented as Dexi and R scripts. It receives three primary inputs: the indicators, the information about the assets with its corresponding mapping to the CIA triad and an instance of the risk model. As an outcome, it generates reports per target and risk. These reports serve as output and are subsequently relayed to the Aggregator.

The **Aggregator** submodule plays a pivotal role in consolidating information. It compiles the generated reports into a comprehensive global report, which is then forwarded to the Dashboard for visualization purposes.

PostgreSQL, serving as the underlying database, is instrumental for data storage. It houses configuration details and past results that have been generated by the system.

Concerning the static information that is used as input, the questionnaire serves as a pivotal input, offering business indicator values that pertain to the infrastructure profile and environment. Designed to encompass the broader aspects of the business and the company's ICT culture, the questionnaire contains model-specific inquiries tailored to the targeted vulnerabilities.

Additionally, the configuration of targets supplies critical data, including IP addresses, ports, and an estimation of potential economic loss should an incident compromise the organization's

confidentiality, integrity, or availability. Complementing this, risk models provide an interpretive framework for risk patterns, considering specific risks within predefined assumptions.

On the other hand, the dynamic information stream consists of various elements that offer real-time insights. Alarms, sourced from the monitored infrastructure, provide immediate notifications about anomalies or potential threats. Events, another integral input, capture significant occurrences within the monitored environment.

Moreover, CERCA assimilates vulnerability reports arising from rigorous penetration tests and vulnerability scans targeted at the platform's defined objectives. Lastly, indicators of compromise contribute a crucial layer of threat intelligence, originating from the MISP instance and meticulously processed by the Threat Intelligence Engine (TIE) module.

Finally, the outputs generated by the CERCA system are a culmination of the processed inputs, yielding a comprehensive assessment of risks and vulnerabilities.

The Global Report encapsulates a holistic evaluation of risk across the entire infrastructure, providing a bird's-eye view of the organization's threat landscape.

For a more granular perspective, the system generates Reports per Target and Risk. These individualized reports are tailored to specific targets, factoring in the specific risk considerations under scrutiny.

Furthermore, the system offers Reports per Model, tailoring insights to the selected risk models, thereby presenting a nuanced understanding of the organization's vulnerabilities.

Lastly, Reports per Target cater to each target being assessed, offering a detailed overview of the risk associated with each one.

2.2.4 Incident Reporting

The incident reporting tool known as AIRE is a central component within this domain. Designed with modularity in mind, AIRE's asset structure consists of various services, see Figure 6, that enable flexibility and adaptability to accommodate diverse regulations and potential regulatory changes.

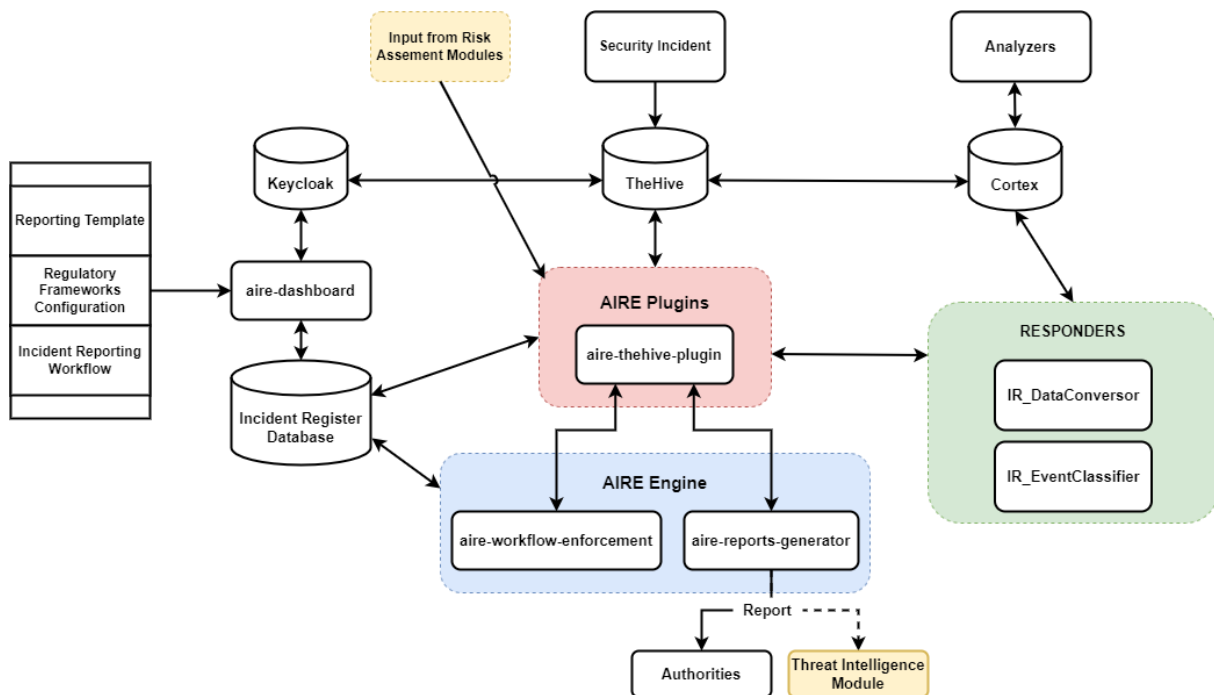


Figure 6. AIRE architecture.

The AIRE engine is comprised of two Springboot microservices: **aire-reports-generators** and **aire-workflow-enforcement**. Acting as an intermediary, the **aire-thehive-plugin** bridges the gap between

the AIRE engine and the open-source Security Incident Response Platform, TheHive. The overarching architecture also integrates an Incident Register Database for comprehensive data storage, and the aire-dashboard web application provides a user-friendly interface to configure and interact with the asset.

The **aire-workflow-enforcement** service plays a pivotal role by orchestrating the incident reporting workflow. This workflow adheres to a Business Process Model and Notation (BPMN) framework. This service efficiently manages and enforces various stages within the common incident reporting process. It encompasses tasks ranging from incident registration and related data gathering to the creation of reports intended for supervisory authorities. Throughout these phases, distinct user roles are involved, and a 4-eye principle is applied, ensuring managerial oversight to prevent inadvertent reporting. The integration of TheHive into this workflow is facilitated through the aire-thehive-plugin service and REST APIs.

Conversely, the **aire-reports-generator** service focuses on managing registered security incident data and tailoring it to different report templates mandated by Competent Authorities. By accessing the Incident Register database, this service adapts the stored information to generate output report files. To facilitate this, the Apache POI library is utilized for Microsoft document formats (Excel and Word), while the Apache PDFBox library is employed for PDF documents, both of which operate under the Apache License v2.0.

2.2.5 Dashboard

The CPR tool is designed to consolidate the diverse modules into a central dashboard, presenting pertinent information for each monitored asset within the Critical Infrastructure. This dashboard offers an initial view divided into distinct sections, each aligned with a module. These sections encompass real-time charts and tables that succinctly summarize CI metrics and status. Further granularity is provided through dedicated component views, granting users access to detailed information, configuration options, and actionable insights. Although customization is anticipated based on pilot requirements, a proposed default view addresses common use cases.

At the dashboard's top, the CERCA section takes precedence. CERCA, focused on risk assessment, features a listing of monitored assets accompanied by risk scores. These scores, color-coded to denote severity, offer a snapshot of asset conditions, with direct links to comprehensive risk reports for individual assets.

Following this, a section is dedicated to the LOMOS module. It encompasses a general chart illustrating anomaly events per minute within the past day. Additionally, a listing of recent anomaly events is presented, each item offering direct access to detailed event information.

The subsequent section compiles information from the Security Information and Event Management (SIEM) module. The default SIEM is Wazuh for the WP6 tool, adaptable to pilot-specific SIEM deployments. This section potentially features a chart displaying alarms per asset over the past month. A list of recent alarms, linked to their detailed views, complements this visual data.

Transitioning to the TINTED module, a chart showcases event types received within the last day (extended if the event count is low). A corresponding list displays relevant events received for the current infrastructure, offering direct access to detailed views for each event.

Concluding the dashboard layout, the AIRE module – responsible for regulatory reporting – presents a list of ongoing reports alongside tasks necessary for reporting completion. This list is thoughtfully organized based on impending deadlines, streamlining the reporting process.

The CPR tool's dashboard unifies the diverse modules, providing an encompassing view of CI performance while enabling deeper dives into specific module functionalities, ultimately enhancing decision-making and operational efficacy.

2.3 Deployment

2.3.1 Pre-requisites

All the modules are containerized using Docker. Consequently, the deployment environment must have Docker (version 20.10.14 or newer) and Docker-Compose (version 1.29.1 or higher) to facilitate component deployment.

- ▶ Installation of docker: <https://docs.docker.com/install/linux/docker-ce/ubuntu>
- ▶ Installation of docker-compose: <https://docs.docker.com/compose/install>

Testing of the application has been conducted on Ubuntu 20.04 LTS and Ubuntu 18.04 LTS operating systems.

2.3.2 Anomaly Detection Module

2.3.2.1 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.

2.3.2.2 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. Example of the nonsensitive parameters configuration file:

```
# 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 actual LOMOS – the controller, workers, celery services, and dashboard. The second one is model registry – Mlflow with FTP server and Postgre 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.

2.3.3 Threat Intelligence Module

2.3.3.1 Hardware requirements

The suggested hardware requirements are as follows:

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

2.3.3.2 Installation procedure

The configuration details for the tool are found within the .env file. It is crucial to verify the accuracy of this information prior to deploying the tool.

In the overarching configuration segment, the DEBUGGING_FLAG and TUNNELING_FLAG indicators must both be set to 0. These flags determine whether the tool is being deployed on a production server. The subsequent part focuses on the individual configurations of specific tool modules. This involves defining container names and the associated ports they utilize. The TIE configuration segment is particularly noteworthy, as it holds environment variables essential for proper tool functioning:

- ▶ MISP_URL: This refers to the URL of the MISP instance, particularly relevant if the MISP instance is located outside of the docker deployment.
- ▶ MISP_SERVICE_NAME: In cases where the MISP instance is part of the same docker deployment, this variable designates the service name of the MISP instance container.

- ▶ **MISP_API_KEY:** This API key is indispensable for interactions with the MISP instance.
- ▶ **ZMQ_CONTAINER_NAME:** The `zmq_client` module connects to the ZMQ Server, which is situated within the MISP instance. Thus, this variable contains the name of the MISP Instance container.
- ▶ **ZMQ_CONTAINER_PORT:** This variable stores the port at which the zmq server is active, enabling the `zmq_client` to subscribe to the queue.

For the configuration of Keycloak, the appropriate adjustments should be made within the `docker/app/controller/keycloak_controller.py` file.

Ultimately, deploying the tool is straightforward and accomplished with a single command:

```
docker compose --env-file .env up -d --build
```

2.3.4 Risk Assessment Module

2.3.4.1 Hardware requirements

The suggested hardware requirements are as follows:

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

2.3.4.2 Installation procedure

To appropriately set up and deploy the Risk Assessment module, the following applications must be installed [6]:

- ▶ The Risk Assessment Engine, which serves as the central element of the tool. It's 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's 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, serving as the storage element of the tool. It's 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, responsible for facilitating communication among various components using a RabbitMQ server.

A Docker container is established 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 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 local SQLite3 database is stored in the file `"rae_dashboard/db.sqlite3"`, and this file is fully compatible with SQLite tools.

In essence, each of the four 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.

Before creating the Docker containers, the Python script `"parser_configurator.py"` (located in the root of the sources) needs to be executed. This script copies specific configuration files, environment files

containing necessary variables, and cryptographic certificates utilized by the applications for secure communication. This preparation is essential for subsequent dockerization:

```
sudo python3 parser_configurator.py
```

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
```

2.3.5 Incident Reporting Module

AIRE component requires a user account with sudo privileges.

2.3.5.1 Aire-reports-generator service

To the deploy the module we need to follow these steps:

- ▶ Run the following commands in the directory containing the *docker-compose.yml* file to construct and deploy the Docker image:

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

Below is the content of the *docker-compose.yml* file:

```
services:
  dashboard:
    build:
      context: ../dashboard/aire_dashboard
      dockerfile: ./Dockerfile
      args:
        DOCKER_DJANGO_DEBUG: "False"
    command: gunicorn config.wsgi:application -c
      ./config/gunicorn_configuration.py
    volumes:
      - ../cs4e-media:/usr/src/app/aire_dashboard/media
    expose: # exposed internally to other Docker services
      - 5602
    env_file: ../dashboard/aire_dashboard/.envs/.env
    container_name: aire_dashboard
```

```

restart: unless-stopped
depends_on:
  - db
networks:
  - "aire-net"
db:
  image: postgres:12.4-alpine
  #image: postgres:11.2-alpine
  volumes:
    - postgres_data:/var/lib/postgresql/data/
  env_file: ../dashboard/aire_dashboard/.envs/.env
  container_name: aire_database
  restart: unless-stopped
  ports:
    - 5432:5432
  networks:
    - "aire-net"
aire-reports-generator:
  build:
    context: ./aire-reports-generator
    dockerfile: ./Dockerfile
  #image: aire-reports-generator:0.1
  #ports:
  # - 8083:8083
  expose: # exposed internally to other Docker services
    - 8083
  volumes:
    - ../cs4e-media/templates:/opt/aire/config/templates
    - ../cs4e-media/reports:/opt/aire/reports
  container_name: aire-reports-generator
  restart: unless-stopped
  depends_on:
    - db
    - dashboard
  networks:
    - "aire-net"
volumes:
  postgres_data:
  static_volume:
networks:
  aire-net:
    external: true

```

- For log verification (with the default file location at `/opt/aire/log/aire-reports-generator.log`) employ the command:

```
sudo docker logs -f aire-reports-generator
```

For altering the configuration file, establish a connection with the container:

- Execute:

```
sudo docker exec -u 0 -it aire-reports-generator bash
```


Subsequently, modify the file located at `/opt/aire/config/application.properties`. Upon completion, restart the container:

► Use:

```
sudo docker restart aire-reports-generator
```

2.3.5.2 Aire-thehive-plugin service

Execute the following commands from the folder where the `docker-compose.yml` file is to build and deploy the docker image:

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

This is the `docker-compose.yml` file:

```
version: "3.3"

# every container should be in the same network
services:
  dashboard:
    build:
      context: ../dashboard/aire_dashboard
      dockerfile: ./Dockerfile
    args:
      DOCKER_DJANGO_DEBUG: "False"
    command: gunicorn config.wsgi:application -c
      ./config/gunicorn_configuration.py
    volumes:
      - ../cs4e-media:/usr/src/app/aire_dashboard/media
    expose: # exposed internally to other Docker services
      - 5602
    env_file: ../dashboard/aire_dashboard/.envs/.env
    container_name: aire_dashboard
    restart: unless-stopped
    depends_on:
      - db
    networks:
      - "aire-net"
  db:
    image: postgres:12.4-alpine
    #image: postgres:11.2-alpine
    volumes:
      - postgres_data:/var/lib/postgresql/data/
    env_file: ../dashboard/aire_dashboard/.envs/.env
    container_name: aire_database
    restart: unless-stopped
    ports:
      - 5432:5432
    networks:
      - "aire-net"
  aire-thehive-plugin:
    build:
      context: ./aire-thehive-plugin
      dockerfile: ./Dockerfile
    expose: # exposed internally to other Docker services
      - 8081
    container_name: aire-thehive-plugin
    restart: unless-stopped
```

```
depends_on:
  - db
  - dashboard
networks:
  - "aire-net"
volumes:
  postgres_data:
  static_volume:
networks:
  aire-net:
    external: true
```

To check the logs (by default file */opt/aire/log/aire-thehive-plugin.log*):

```
sudo docker logs -f aire-thehive-plugin
```

To change the configuration file, connect to the container:

```
sudo docker exec -u 0 -it aire-thehive-plugin bash
```

and edit file */opt/aire/thehive-plugin/config/application.properties*

3 Cyber-Physical Resilience Tool Methods Tested in Labs

3.1 Monitoring System (LOMOS)

In this section, we assess the performance of log-based anomaly detection methods. First, we present the data, the methods, and at the end, the results of the empirical evaluation. The evaluation is performed in the sense of both, predictive accuracy, and throughput.

3.1.1 Data

We evaluate the methods on two public datasets, namely BGL [7] and HDFS [8], labelled by the experts responsible for the system maintenance. One of the datasets contains logs from the BlueGene/L supercomputer, while the other contains logs from multiple hundred EC2 nodes in the Hadoop Distributed File System cluster. The datasets were captured from the actual systems running in production. They capture both, software, and hardware problems. As such, they more than adequately represent behavior captured in logs in complex systems and can serve as an indication of the expected performance of the methods in our pilots. The datasets are well accepted in the log-based anomaly detection research community, as they are used as benchmark datasets in the field.

At the time of the BGL dataset release (2007), this was the largest supercomputer in the world, consisting of 131,072 CPUs and 32,768 GB of RAM. The collection of logs started on the 3rd of June 2005 and lasted for 215 days. In that time 4,747,963 log messages were collected, out of which 348,460 were anomalous. The anomalies are on a per-log level and are classified into 41 categories, including software and hardware problems. The logs were labelled by the experts responsible for the BlueGene/L supercomputer.

The HDFS dataset contains logs from the Hadoop Distributed File System cluster running on over 200 EC2 nodes. It contains 12,577,685 logs captured in two days in November 2009. The logs belong to sessions based on the block ID which is part of the log messages. The labels are binary, where the session is normal or anomalous. The authors of the dataset inspected the HDFS code and consulted with local Hadoop experts to label the data. The dataset contains 575,061 sessions, out of which 16,838 are anomalous.

3.1.2 Methods

Most of the log-based anomaly detection methods follow a typical workflow, where (i) log messages are first represented in a normalized form, known as log templates or events. The log templates (ii) are then modeled as a sequence and (iii) are used for anomaly detection. Some methods skip the log template parsing part and use a pre-trained language model to semantically embed a log message.

Log-based anomaly detection methods mostly work with sequences of log messages. The sequences are generated based either on session windows or sliding windows. The latter session type can be created according to time (e.g., 5 minutes of logs) or a predetermined number of logs (e.g., 100 logs). Choosing the appropriate method for grouping logs into sequences depends on the data characteristics. If logs have session identifiers, it makes sense to use them. In this way, we ensure the logs are related to each other. Otherwise, the sliding window method should be used. The maximum window size is mostly limited by hardware limits. Another problem in very large sequences could be too much noise due to unrelated logs.

We use a state-of-the-art method for log template extraction named **Drain** [3]. Drain is a highly effective tree-based fixed-depth online log parsing method. Log template extraction is a very important step in numerous log-based anomaly detection methods. Logs are first preprocessed with provided general expressions based on domain knowledge. Even though they are not essential, they can improve the log template parsing results. At least general regular expressions for masking IP addresses, timestamps, and similar should be included. Next, new log templates are compared to the

existing log templates. The comparison of log messages is divided into three steps. In the first step, log messages are compared by their length (number of words). In the second step, log templates are compared by the initial words through the Drain's fixed-depth tree structure. The depth of the tree, and with it the number of initial words taken for comparison, is configurable. The assumption here is, that the parameters do not occur at the beginning of log messages. If no exact match is found, a new log template is formed. In the case, there are matches, the new log message is compared word by word with all the log templates in that leaf node. If the highest similarity score is above the threshold, a log template is found for the given log message. If not, a new log template is formed. New log templates are formed only in train mode. If DRAIN is in inference mode, the unrecognized log messages are marked as an unknown log template.

For a baseline, we use an **unknown log template detector**, which is a useful model for anomaly detection in some of the use cases. The idea is to classify the before unseen events as anomalous. A system administrator is notified about it and can respond by either marking the log template as normal or adequately addressing the issue spotted by the detector. The method can handle per-log anomaly detection and per-sequence. In the latter one, we count unknown log templates. If it exceeds the threshold, the sequence is marked as anomalous. In the utmost case, a sequence is marked as anomalous if any of the log templates is unknown.

LOMOS primarily employs the **LogBERT** method [4] for log-based anomaly detection. LogBERT is a self-supervised approach that draws the main architecture and training ideas from BERT [5]. While BERT uses subword tokens, LogBERT employs log keys (log template IDs) as tokens. This adaption involves utilizing masked language modeling training tasks to predict masked log keys. Since this task is a classification into a fixed set of classes, cross-entropy is used as a loss function. The second training objective is a volume of hypersphere minimization, where the loss function ensures embeddings are close together in the vector space. The loss for the second task is defined as the average Euclidean norm of the distance to an average embedding. The losses are weighted and added together. In the inference mode, the anomaly score is computed as the ratio of missed predictions. Random log keys are masked and predicted with the trained model. Top g log keys in the sense of computed probability form a candidate set. If the actual log key is not in the candidate set, the prediction is marked as missed. If the ratio of misses is above the threshold, the sequence is predicted as anomalous.

We **extended LogBERT** to be able to predict anomalies on a per-log basis. The main change is made in the data collator for the inference. Instead of masking log keys at random, we mask only the last log in a sequence. The anomaly score is computed as the relative position of the masked log key in the predicted distribution. If the first predicted log key is the actual one, the anomaly score is 0. If it is the last, the anomaly score is 1.

3.1.3 Methods evaluation

Log-based anomaly detection is a classification problem with imbalanced classes, where the anomalous data is greatly outnumbered by the normal data. Although some of the datasets classify labels into multiple categories, we convert them into binary. The logs are normal or anomalous. To measure the models' predictive performance, we use labelled datasets and metrics that cope well with the imbalanced predictions, namely F1 score, precision, and recall. First, let us define TP as a true positive, FP as a false positive, and FN as a false negative. The precision is defined as

$$precision = \frac{TP}{TP + FP}, \quad (1)$$

and recall as

$$recall = \frac{TP}{TP + FN}. \quad (2)$$

The F1 score is the harmonic mean of precision and recall

$$F1 = 2 \frac{precision \cdot recall}{precision + recall}. \quad (3)$$

We generated the sequences from BGL datasets with a sliding window with a size of 100 logs. For HDFS, we used block ID to generate the sequences based on the sessions. When evaluating the per-sequence anomaly detection methods on datasets with per-log labels, the sequence is anomalous if any of the logs in the sequence is anomalous. Once the sequences have been generated, they're divided into training and testing sets. We keep 40% of the data in the training set and leave out 60% for the test set. There are two prevalent strategies in the literature for dividing the data into train and test sets, namely random and chronological data split. In the random split method, sequences are divided without any particular order. On the other hand, the chronological split method ensures the original sequence order is retained. In this case, the first part of the dataset, according to the timestamps, is used for training sequences, while the latter part is reserved for testing sequences.

In Table 1 we present the results of all three methods described in the previous section. The methods are evaluated on BGL and HDFS datasets. We report the results in terms of F1, precision, and recall. There are two values in each cell, the first for random split, and the second for chronological split of the data.

Table 1: Empirical evaluation of log-based anomaly detection methods.

Method	BGL			HDFS		
	(random/chronological split)			(random/chronological split)		
	F1	Precision	Recall	F1	Precision	Recall
Unknown log template	0.94 / 0.44	0.98 / 0.29	0.91 / 0.91	0.26 / 0.15	1.00 / 1.00	0.15 / 0.08
LogBERT	0.92 / 0.48	0.93 / 0.32	0.91 / 0.97	0.81 / 0.77	0.77 / 0.71	0.86 / 0.82
LogBERT - per log	0.89 / 0.19	0.83 / 0.10	0.95 / 0.99	Labels on sequence level		

The **unknown log template** method works notably well on the BGL dataset, particularly evident with a random data split. In the BGL dataset, log templates mostly have consistent labels. Meaning, that there are not a lot of the log templates that are normal and anomalous, depending on the context they appear in. This leads to a very high recall for the unknown log template predictor. Precision is high only with the random split. In the case of chronological, there are a lot of normal log templates not seen in the training process, that are therefore falsely marked as anomalous. The reason for a very low recall in the HDFS dataset is that most of the anomalous sequences do not contain any of the unknown log templates. Those that do, are practically always anomalous, which is the reason for a high precision in the HDFS dataset.

The **LogBERT** method performs well on all three datasets in both, random and chronological split settings. The only case where it performs not so well is in the BGL dataset with the chronological split. The reason for this is low precision. Specifically, in the BGL dataset, there are a lot of log templates appearing later in the dataset, hence not seen in the training phase. They are marked as unknown and as such have a very low probability of being predicted in the candidate set. Therefore, those unseen log templates get marked as anomalous, even if normal. This effect is not so strong in the HDFS dataset.

The **per log LogBERT** performs a bit worse than the normal LogBERT in the sense of precision and therefore also the F1 score. The reason for this is a much harder task, to classify each single log message as normal or anomalous. LogBERT has to classify if any of the log templates in a sequence of multiple (e.g., 100) logs is anomalous. However, by evaluating each log in the stream, it is capable of detecting more anomalies. For this reason, it has a higher recall than LogBERT.

3.1.4 Throughput

Classification performance is only one of the relevant metrics for log-based anomaly detection. Another important performance indicator is speed. In this section, we describe the results of throughput measurement experiments. We compare the throughput on two different machines. The first one, called Turbo, is highly suitable for machine learning with four CUDA-supported GPUs. The second machine is a virtual machine (VM) without GPU support. Hardware specifications of the machines are listed in Table 2. In the experiments, we test Turbo and hybrid deployments. In the first setting, everything is deployed on Turbo, while in the hybrid deployment, only the GPU worker is on Turbo and everything else is on the VM.

Table 2: Hardware resources used in the throughput experiments.

	Turbo	VM
CPU	48 (AMD Ryzen Threadripper 3960X 24-Core)	8 (Intel Xeon E3-12xx v)
RAM	128 GB	16 GB
DISK	M.2 SSD	HDD
GPUs	4x RTX 2080 Ti (11 GB)	None

For the experiments, we use the first million logs from the BGL dataset. The log-based anomaly detection can be divided into training and inference modes. First focusing on training, the training is split into two tasks. Namely, log template parser training and anomaly detection model training. We measure the speeds separately for each of the tasks. While the parser trains, it also parses the logs and writes the extracted log templates into the database. The trained parser is pushed to the model registry. As seen in Table 3 the training is much faster on Turbo, even though GPUs are not used in this step. The reason for this is faster CPU and disk.

Table 3: Train log template parsers on 1 million logs from the BGL dataset execution time measurements.

Turbo	Hybrid
184 s	628 s

The second part of training is anomaly detection model training. In this step, we use one or four GPUs. The training is done for five epochs with a batch size of 64 sequences having a window size of 100 logs. The results are presented in Table 4. The times are similar for Turbo and Hybrid deployments because the most time-consuming part is the actual training, which is executed on Turbo in all deployment configurations. Having multiple GPUs significantly speeds up the training process.

Table 4: Train anomaly detection model on 1 million logs from the BGL dataset execution time measurements.

Turbo 4	Turbo 1	Hybrid 4	Hybrid 1
1,126 s	3,418 s	1,125 s	3,403 s

Training is done only on demand, initially for the system setup and when a data drift is detected. While the inference is run periodically (e.g., every 5 minutes). For this reason, the inference throughput is usually seen as more important. The inference can be executed either on GPU or even on CPU. The throughput is much higher when running the inference on GPUs, however, the CPUs are often fast enough. The inference consists of extracting log templates from new log messages and combining

them into sequences that are passed through the anomaly detection models. We again use a window size of 100 logs for log sequences and increase the batch size to 256 for the inference. We tested seven different setups. Four of them with GPU and three with CPU workers. The GPU options have one or four GPUs. With the CPU-only deployments, we have a full Turbo deployment CPU inference worker and hybrid deployments with a CPU worker either on the VM or on Turbo.

The training speeds are reported in Table 5. There are significant differences in the speed of different deployments. The biggest bottlenecks are writing and updating database record speeds on the VM, due to HDD. GPUs also have a great impact on computation speed. If we extrapolate the numbers, the slowest option – Hybrid VM-CPU, could process more than 7.5 million logs per day, while the fastest – Turbo 4 more than 185 million logs per day.

Table 5: Anomaly detection inference on 1 million logs from the BGL datasets execution time measurements.

Turbo 4	Turbo 1	Turbo CPU	Hybrid 4	Hybrid 1	Hybrid VM-CPU	Hybrid Turbo-CPU
465 s	1,105 s	5,000 s	830 s	1,652 s	11,092 s	3,745 s

3.2 Threat Intelligence (TINTED)

The threat intelligence module uses MISP underneath. This platform allows users to share and store security events that contain IoCs. The way MISP is envisioned, therefore TINTED, is to facilitate the information of threat intelligence data, which should not be flooded into the platform. That way, the measurement of different metrics such as throughput, precision, recall, and F1 score does not fit with the nature of the tool because it does not make use of AI technology. D6.3 will contain an integrated version of the tool, where functional tests will be carried out to demonstrate the coverage over several use cases, tailored to the nature of the module.

3.3 Risk Assessment (CERCA)

The risk assessment module takes into consideration static and dynamic data. Nonetheless, the dynamic data (SIEM alerts) that it receives is thought to be a small quantity to avoid an infinite loop of production of risk score. Hence, the calculation of the metrics for the log anomaly module does not apply in this case as no AI model is used. D6.3 will feature a unified tool version in which we will conduct functional tests, that will be adapted to the corresponding module, to showcase its effectiveness across various usage scenarios.

3.4 Incident Reporting (AIRE)

The incident reporting module considers if security alerts and events must be elevated to the category of the incident and provides automated processing for reducing workload during the process of reporting to authorities. By design, this tool is not thought to operate with high volumes of data apart from the fact that it does not include any AI model, therefore the measurement of throughput, precision, recall, or F1 score would not make sense. D6.3 will include a consolidated tool version in which we will perform functional tests to demonstrate its efficacy across diverse usage scenarios.

4 Management and What-If Analysis

Section 4 is specifically designated to outline potential legal concerns. This involves scenarios such as data sharing restrictions for CI due to legal regulations or limitations on permissible testing activities.

At this stage, no problems have arisen, although they may appear during the testing phase. This section will be revised in the next versions of this document (D6.3 and D6.5) and is closely monitored in WP9 – Management: Project Coordination.

5 Pilot trials execution (feasibility analysis)

5.1 Proofs of concepts

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 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.
- ▶ Phase 2 (M34): The tools will be piloted in the operational environment of INS. The tools' 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 WP6 tools 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 gained crucial importance 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 the purpose of 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 these 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 analyzed.
- ▶ Phase 2 (M34): CAFC will carry out a trial of the tools in their actual operational environment. This will involve monitoring the tools' 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 WP6 tools on its VPN system for detecting suspicious network behavior. 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 analyzed 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's already in operation within TS.
- ▶ Phase 2 (M34): In this phase, TS will put the tools to the test within its operational setup. The tools' 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 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 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 various 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 with the following stages:

- ▶ Phase 0 (M12): The initial Proof of Concept will illustrate the tools' capability to identify potential threats within the analyzed 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).
- ▶ 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.

In summary, SZ-SZI, acting as Slovenian transport operators, intends to assess 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.

6 Conclusions

This deliverable publicly presents the SUNRISE WP6 Cyber-Physical Resilience (CPR) tool modules and architecture initially presented in deliverable D6.1 with the sensitive dissemination level. The tool consists of four main modules. Namely, the anomaly detection module, threat intelligence module, risk assessment module, and incident reporting module.

The modules were tested in lab environments regarding different metrics relevant to the tools and use cases. For example, the F1 score, precision, and recall were used to measure predictive performance and the processing speed to have a module throughput estimation. The experiments were done with the relevant public datasets that are commonly used in the domain literature. The log-based anomaly detection system was able to process more than two thousand log messages per second. This is more than enough for the expected use cases. The system also showed a high recall, exceeding 0.9 in some cases and staying above 0.8 in others. Meaning it was able to detect most of the anomalies. Precision proved to be more sensitive to the datasets and their split methods. Depending on the setting, it was able to achieve 0.93 in the optimal scenarios but dropped to 0.31 or even a bit lower when evaluating on a per-log basis.

Next, the deliverable focuses on the piloting activities, where the tool will be tested on data coming from real use cases. The tool will be tested in four CI pilots from Italy and Slovenia, including public administration, water, telecommunication, and transport use cases. An iterative, phased approach will be maintained throughout the lifetime of the project, ensuring proper integration and testing of all the related tools.

The D6.2 is the first version of the Cyber-physical resilience tool and training guide deliverable. The content will be iterated to the next versions in D6.3 and D6.5. The pilot reports will be presented in D6.4, and finally in D6.6, illustrating the results of the iterative process adopted by the project, after extensive end-user testing.

References

- [1] **SUNRISE. D6.1** - Cyber-physical resilience conceptualization. Pablo de Juan. 2023.
- [2] **SUNRISE. D3.1** - Requirements and designs V1. George Tsakirakis. 2023.
- [3] **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.
- [4] **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.
- [5] **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.
- [6] **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.
- [7] **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.
- [8] **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.
- [9] **CyberSec4Europe - D3.21** Framework to design and implement adaptive security systems. L. Pasquale and A. Hassan. 2022.

Annex I Training guide and user manual

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

I.I 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, log-based anomaly detection inference.

I.I.I Setting up data sources

Elasticsearch is often used as a central database for logs from different components of simple or complex systems. It is highly suitable for handling large 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 Elastic³. 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 endpoint⁴.

I.I.II Training a log parser

LOMOS is capable of carrying out some preprocessing steps on the logs. Sometimes logs are stored in a raw semi-structured format. In such a case, log messages and timestamps must be extracted first. Readers may find a sample of raw BGL logs in Figure 2.

```

- 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-hammer 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 preprocessing step as in Figure 8.

Preprocessing

Match message regex

```
^(?:(?:\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

```
^(?:(?:\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 regex101⁵ tool (Figure 9 and Figure 10). Based only 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. In the case that the data is already structured

³ <https://www.elastic.co/guide/en/beats/filebeat/7.17/index.html>

⁴ <https://www.elastic.co/guide/en/elasticsearch/reference/current/getting-started-index-lifecycle-management.html>

⁵ <https://regex101.com>

in the Elasticsearch index, we can put a regular expression that matches the whole line “^(.+)\$” and leave the other two fields empty.

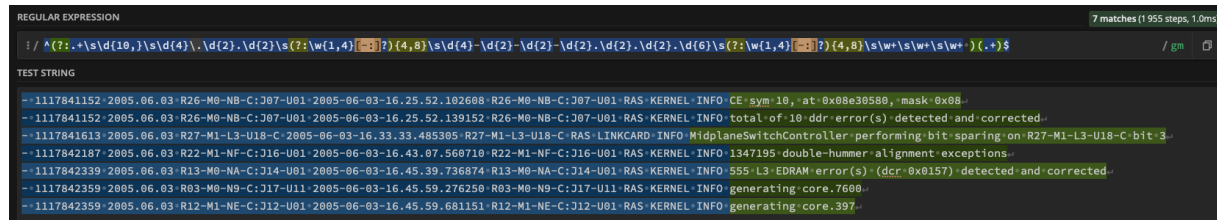


Figure 9. Regular expression for extracting log message from a raw log.

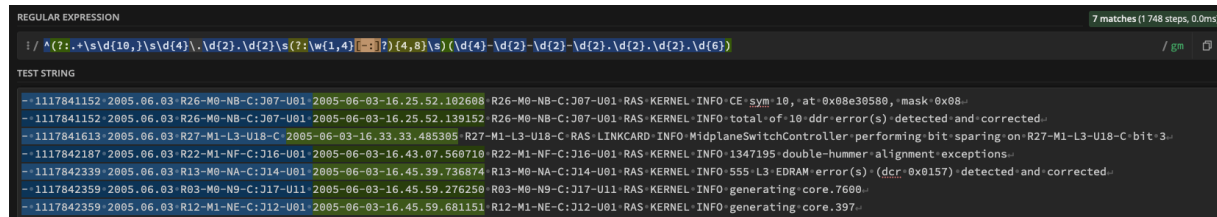


Figure 10. Regular expression for extracting timestamp from a raw log.

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

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

Number of children

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 like 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])|)$`

Remove this mask

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])|)$`

Remove this mask

Add new mask

Figure 12. Custom masks for complex parameters.

The next section is used for setting the connection to Elasticsearch. We have to 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.

After that, we have to 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.

I.I.III 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 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.

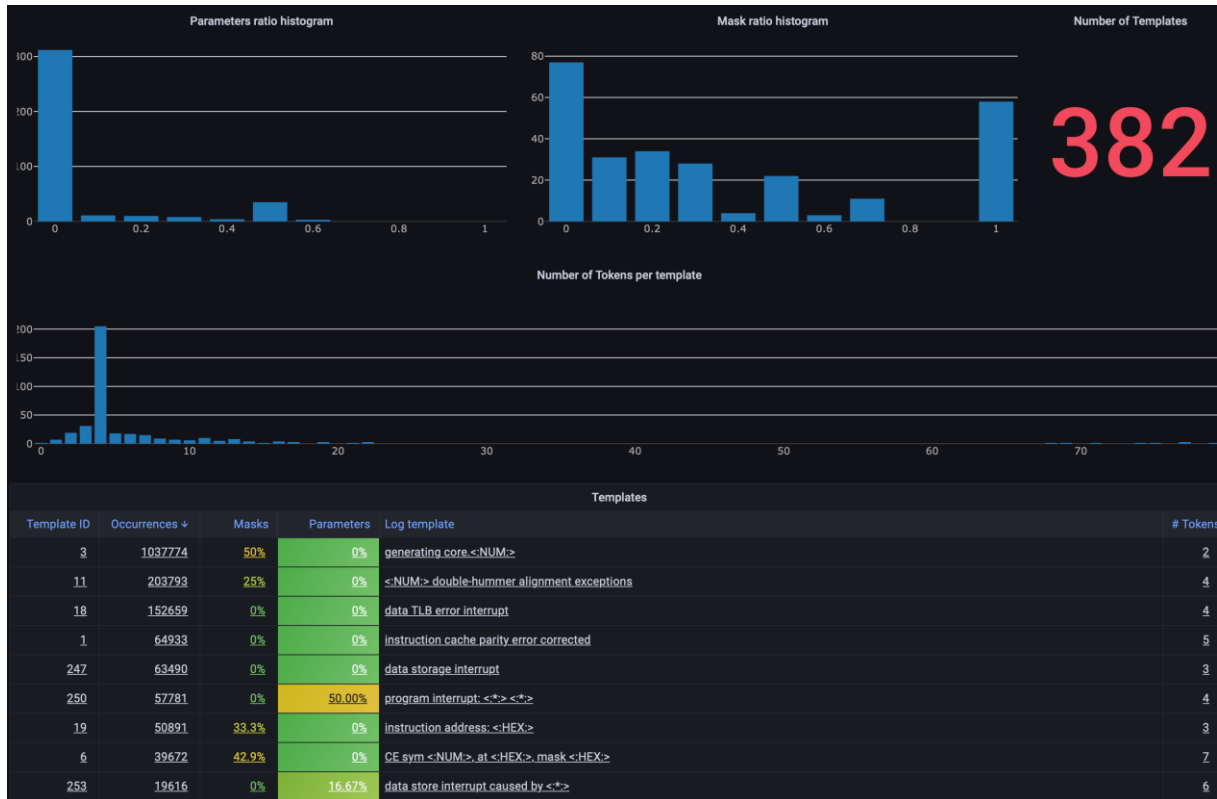


Figure 16. Extracted log templates overview.

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.

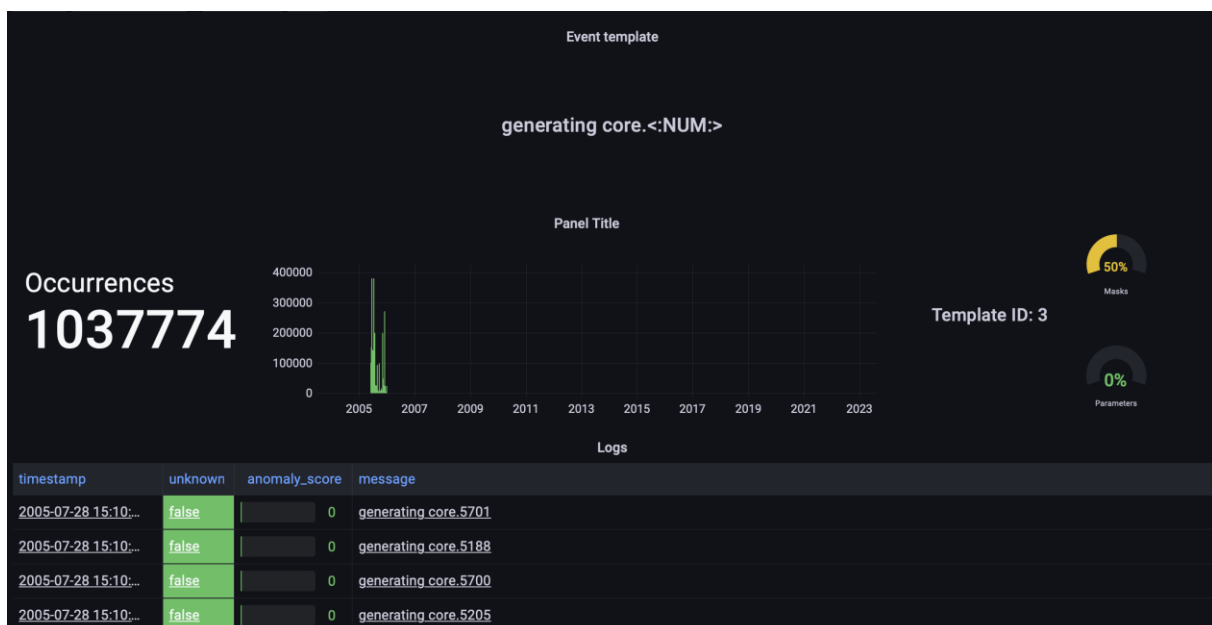


Figure 17. Log template details.

I.I.IV 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 have to 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 skip (potentially) 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

Remove this interval

Add new interval

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 have to 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. After that, 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.

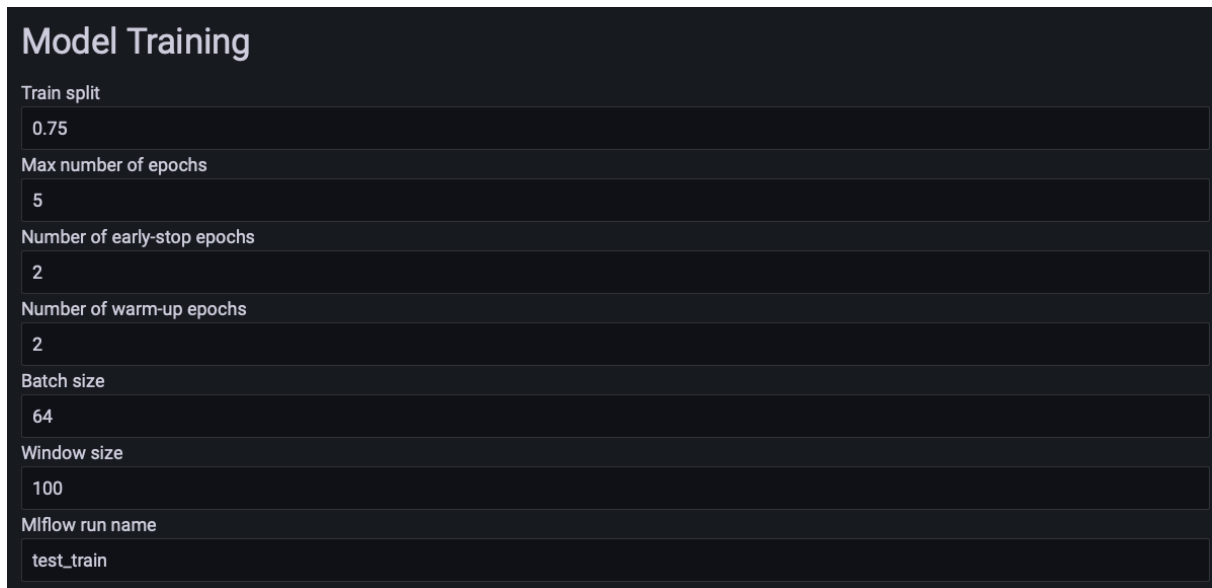


Figure 20. Anomaly detection model training hyper-parameters.

I.I.V Inspecting the training results

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




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

I.I.VI Setting up live inference

Once the log template extraction and anomaly detection training phase are concluded, we can use the model for inference on 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

[Load parser config](#)

Figure 22. Pretrained parser MLflow experiment id.

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

Task period (seconds)

Figure 23. Set the inference schedule period.

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

Elasticsearch host

Elasticsearch port

Elasticsearch username

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. After that, we have to select the period of data that we will pass through the anomaly detector. We can again use timestamps, 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

Elasticsearch log column name

Elasticsearch start

Elasticsearch end

Elasticsearch interval field name

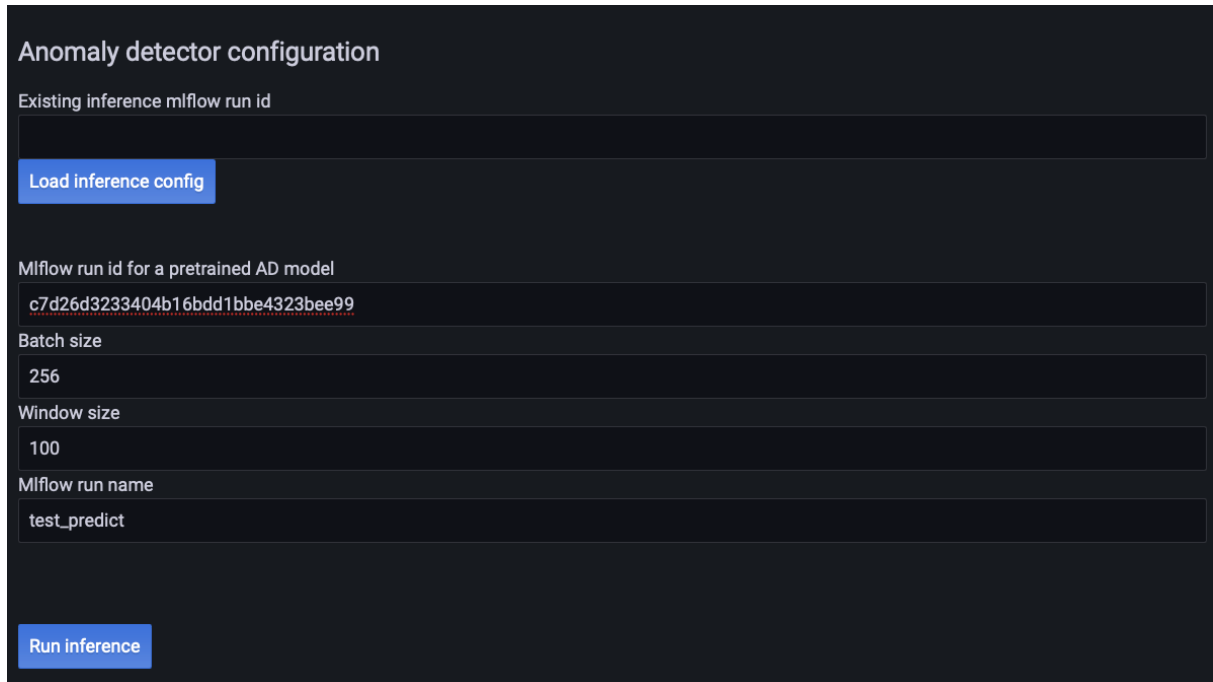
Elasticsearch sorting field name

Elasticsearch sorting field type

Elasticsearch new indices prefix

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

Load inference config

MLflow run id for a pretrained AD model

c7d26d3233404b16bdd1bbe4323bee99

Batch size

256

Window size

100

MLflow run name

test_predict

Run inference

Figure 26. Inference model configuration.

I.I.VII 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 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.

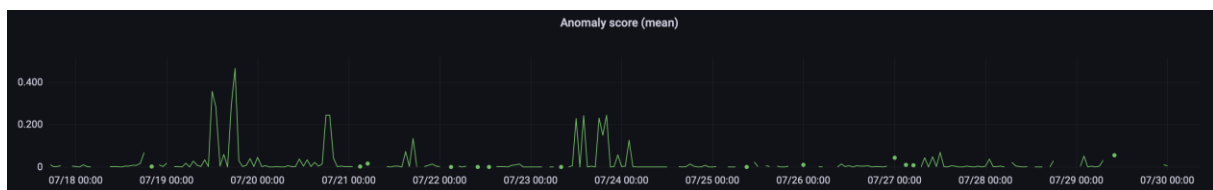


Figure 28. Average anomaly score.

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_YkBlZ5p4GauCjCn
2005-07-13T06:20:38.769000+00-	false	0	6	CE_sym_33_at_0x00dbd200_mask_0x80	ggh7_YkBlZ5p4GauCjCn
2005-07-13T06:20:38.921000+00-	false	0	6	CE_sym_33_at_0x00dbd200_mask_0x80	gwh7_YkBlZ5p4GauCjCn
2005-07-13T06:20:39.072000+00-	false	0	6	CE_sym_33_at_0x00dbd200_mask_0x80	hAh7_YkBlZ5p4GauCjCn
2005-07-13T06:20:39.225000+00-	false	0	6	CE_sym_33_at_0x00dbd200_mask_0x80	hQh7_YkBlZ5p4GauCjCn
2005-07-13T06:20:39.378000+00-	false	0	6	CE_sym_33_at_0x00dbd200_mask_0x80	hgh7_YkBlZ5p4GauCjCn
2005-07-13T06:20:39.530000+00-	false	0.00	7	total of 21 ddr error(s) detected and corrected	hwh7_YkBlZ5p4GauCjCn
2005-07-13T06:27:32.845000+00-	true	0.99	0	ciod: Error loading /home/gloali/src/ddcMD/ddcMD1.1.14/bin/ddcMDbgIV: invalid or missing program image. Exec format error	iAh7_YkBlZ5p4GauCjCn
2005-07-13T06:27:32.881000+00-	true	0.99	0	ciod: Error loading /home/gloali/src/ddcMD/ddcMD1.1.14/bin/ddcMDbgIV: invalid or missing program image. Exec format error	iQh7_YkBlZ5p4GauCjCn
2005-07-13T06:27:32.918000+00-	true	0.99	0	ciod: Error loading /home/gloali/src/ddcMD/ddcMD1.1.14/bin/ddcMDbgIV: invalid or missing program image. Exec format error	igh7_YkBlZ5p4GauCjCn

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. Also, 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 can react appropriately to address the issues found by the system.

I.II Threat Intelligence

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

I.II.I WEB-GUI

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

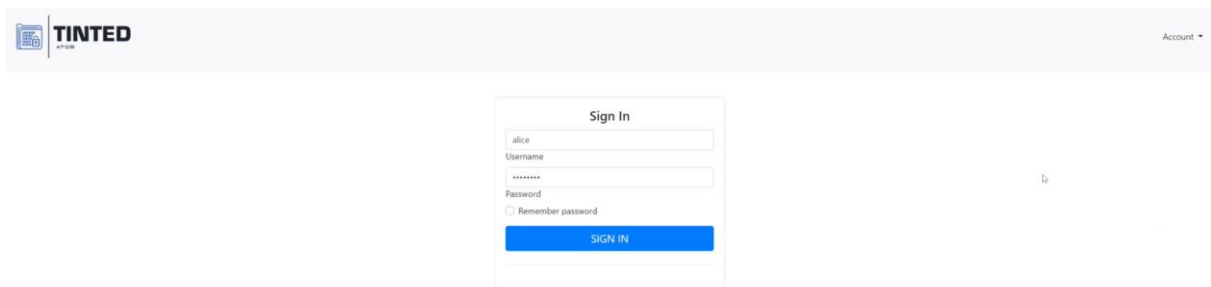


Figure 31. TINTED Login.

Supposing that we have logged in as *Alice* for the first time, then we have to configure the platform. We have to indicate 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).

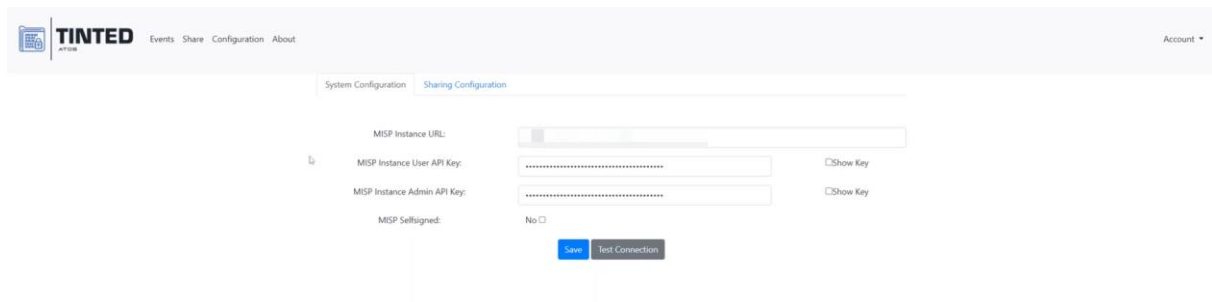


Figure 32. System Configuration.




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).



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

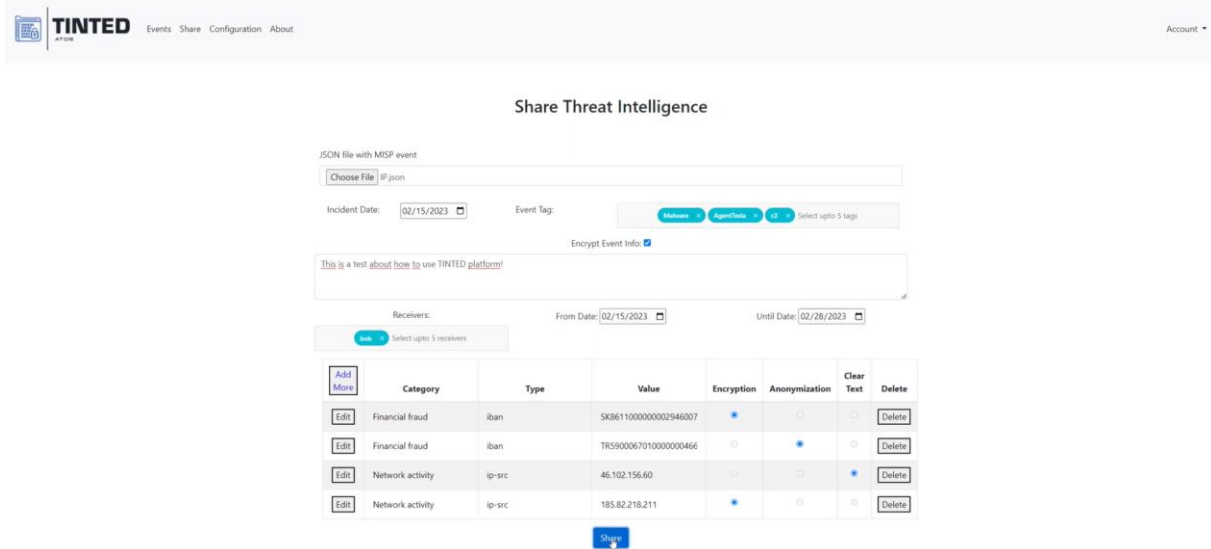
Moving onto the Share webpage, we can see the fields that are available, illustrated in Figure 35.

For the purpose of this guideline, we are going to show how it looks like 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 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.



JSON file with MISP event

Choose File | ip.json

Incident Date: 02/15/2023 Event Tag: Malware Agentless etc. Select upto 5 tags

Encrypt Event Info:

This is a test about how to use TINTED platform!

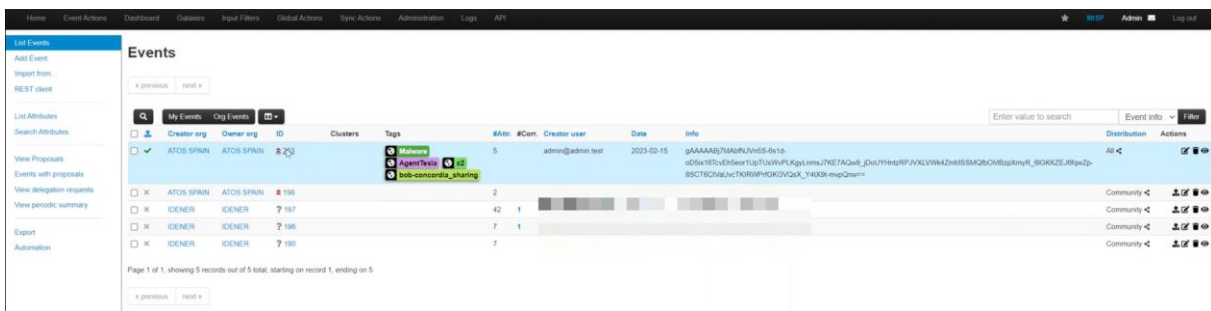
Receivers: From Date: 02/15/2023 Until Date: 02/28/2023

	Category	Type	Value	Encryption	Anonymization	Clear Text	Delete
<input type="checkbox"/>	Financial fraud	iban	SK861100000002946007	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Delete
<input type="checkbox"/>	Financial fraud	iban	TR5900067010000000466	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Delete
<input type="checkbox"/>	Network activity	ip-src	46.102.156.60	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Delete
<input type="checkbox"/>	Network activity	ip-src	185.82.218.211	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Delete

Ship

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.

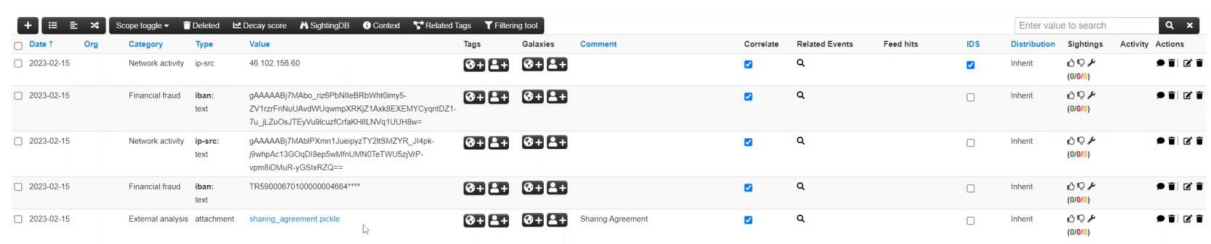


Home Event Actions Dashboard Galaxies Input Filters Global Actions Sync Actions Administration Logs API

Events

My Events	Org Events	ID	Clusters	Tags	#Attr	#Corr	Creator user	Date	Info	Distribution	Actions
<input checked="" type="checkbox"/>	<input type="checkbox"/>	ATOS SPAIN	ATOS SPAIN	Malware Agentless bob-concordia_sharing	5	1	admin@admin.test	2023-02-15	gAAAAABjTAMhNjMhS-5s15- oD5x1M7cE85eur1UpTLWwFLKqLemsJNKE7AQwL_DuYHnRzRPJXKLWmZmM5SAQzOMBzpmRt_80KQZEJfWzZp- 8SCT8CNALvcTKRWRPKGvGxL_Y4KX8-mpQsm=	Community	🔍 🗑️ 🔄

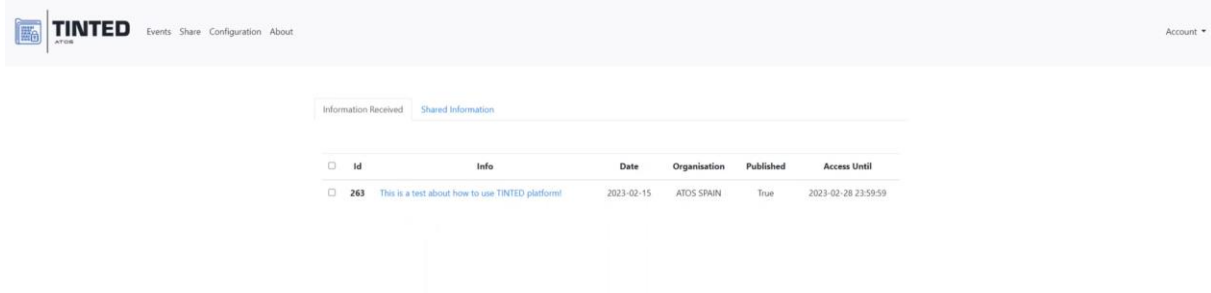
Figure 36. MISP instance.



Date	Org	Category	Type	Value	Tags	Galaxies	Comment	Correlate	Related Events	Feed hits	ID5	Distribution	Sightings	Activity	Actions	
2023-02-15		Network activity	ip-src	46.102.156.60				<input checked="" type="checkbox"/>	Q		<input checked="" type="checkbox"/>	Inherit	(0/0)		🔍 🗑️ 🔄	
2023-02-15		Financial fraud	iban:	gAAAAABjTAMhNjMhS-5s15- ZV1zrFhLJAsdHJQyemjX9KZL4a6SEKfYCygdDZ1- 7uJL2u0sTEyW8RuzdCf9aGHILNvq1ULHlBw=				<input checked="" type="checkbox"/>	Q			<input type="checkbox"/>	Inherit	(0/0)		🔍 🗑️ 🔄
2023-02-15		Network activity	ip-src:	gAAAAABjTAMhNjMhS-5s15- jWpA413GQdD8e9fswfHLMNGTzTWL5qVp- vpm8DMuR-yGSNRZQ=				<input checked="" type="checkbox"/>	Q			<input type="checkbox"/>	Inherit	(0/0)		🔍 🗑️ 🔄
2023-02-15		Financial fraud	iban:	TR5900067010000000466****				<input checked="" type="checkbox"/>	Q		<input type="checkbox"/>	Inherit	(0/0)		🔍 🗑️ 🔄	
2023-02-15		External analysis	attachment	sharing_agreement.pickle			Sharing Agreement	<input checked="" type="checkbox"/>	Q		<input type="checkbox"/>	Inherit	(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 log in into the platform with his credentials (as it was done in Figure 31). After that, he will see that a new event has been received in the dashboard (Figure 38).

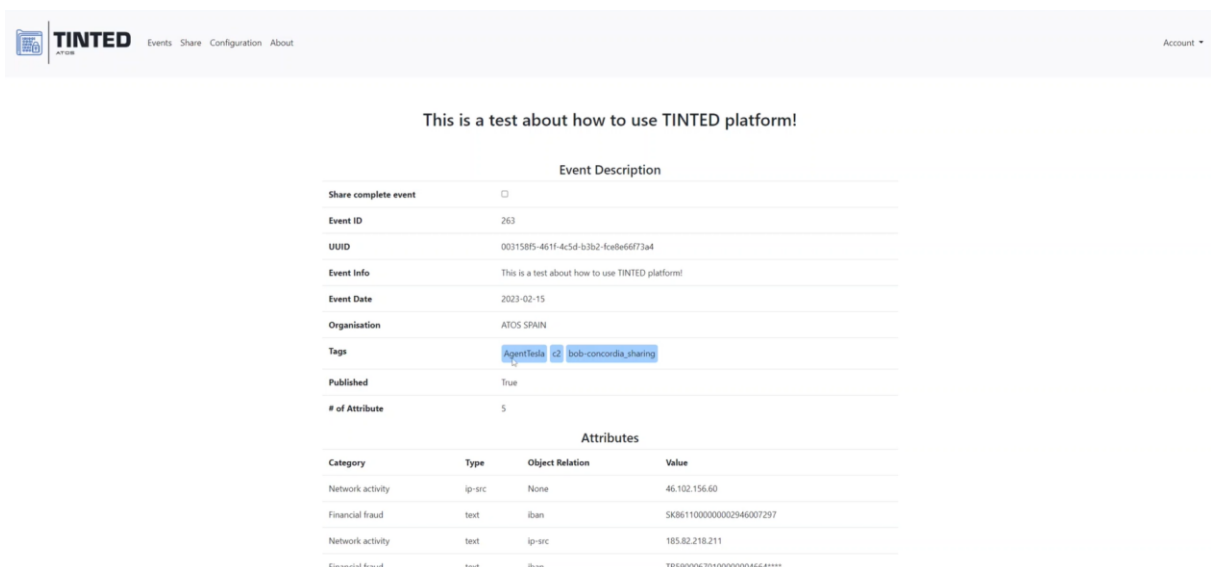


The screenshot shows the TINTED Information Received dashboard. At the top, there is a navigation bar with the TINTED logo and links for Events, Share, Configuration, and About. On the right, there is an Account dropdown menu. Below the navigation bar, there are two tabs: 'Information Received' (selected) and 'Shared Information'. A table lists the received information with columns: Id, Info, Date, Organisation, Published, and Access Until.

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 are able to get the information after the decryption process (Figure 39).

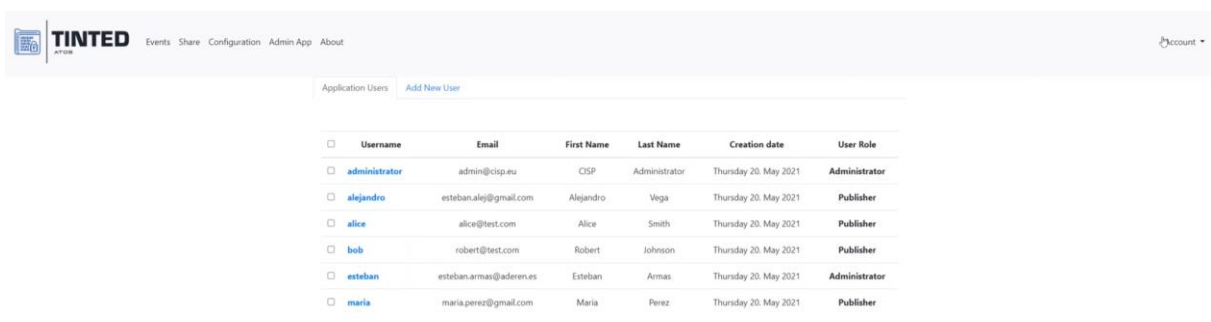


The screenshot shows the TINTED Individual Event details page. At the top, there is a navigation bar with the TINTED logo and links for Events, Share, Configuration, Admin App, and About. On the right, there is an Account dropdown menu. The main content area displays the event details for the event with ID 263. The event description is 'This is a test about how to use TINTED platform!'. Below the description, there is a table of attributes.

Category	Type	Object Relation	Value
Network activity	ip-src	None	46.102.156.60
Financial fraud	text	iban	SK861100000002946007297
Network activity	text	ip-src	185.82.218.211
Financial fraud	text	iban	TR5900067010000004664****

Figure 39. Individual Event details – TINTED.

In case 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.



The screenshot shows the TINTED Administrator Management Dashboard I. At the top, there is a navigation bar with the TINTED logo and links for Events, Share, Configuration, Admin App, and About. On the right, there is an Account dropdown menu. Below the navigation bar, there is a link to 'Add New User'. A table lists the registered users with columns: Username, Email, First Name, Last Name, Creation date, and User Role.

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@ader.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.

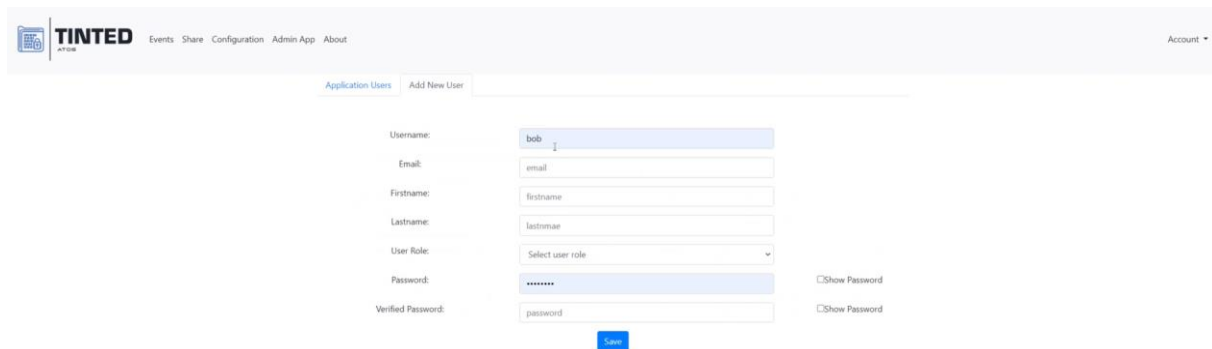


Figure 41. Administrator Management Dashboard II.

I.II.II API

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, like it is shown in Figure 42 and Figure 43, or through an API request to MISP, see Figure 44.

Events

◀ previous 1 2 3 4 5 next ▶

My Events	Org Events	ID	Clusters	Tags	#Attr.	#Corr.	Creator user	Date	Info	Distribution	Actions
<input type="checkbox"/>	<input type="checkbox"/>	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	📄 🗑️ 🔄
<input type="checkbox"/>	<input type="checkbox"/>	21622		Vulnerability TIE:Score-Analysis=2.76	9	1		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	📄 🗑️ 🔄
<input checked="" type="checkbox"/>	<input type="checkbox"/>	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	📄 🗑️ 🔄
<input checked="" type="checkbox"/>	<input type="checkbox"/>	21619		Vulnerability TIE:Score-Analysis=2.67	7	2		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 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	📄 🗑️ 🔄
<input checked="" type="checkbox"/>	<input type="checkbox"/>	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.

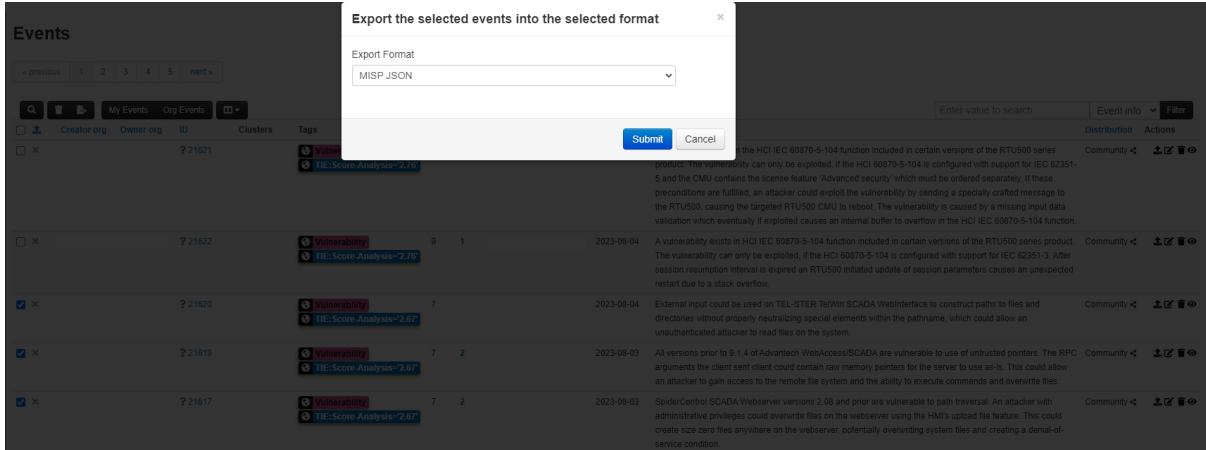


Figure 43. Download of events in MISP JSON format.

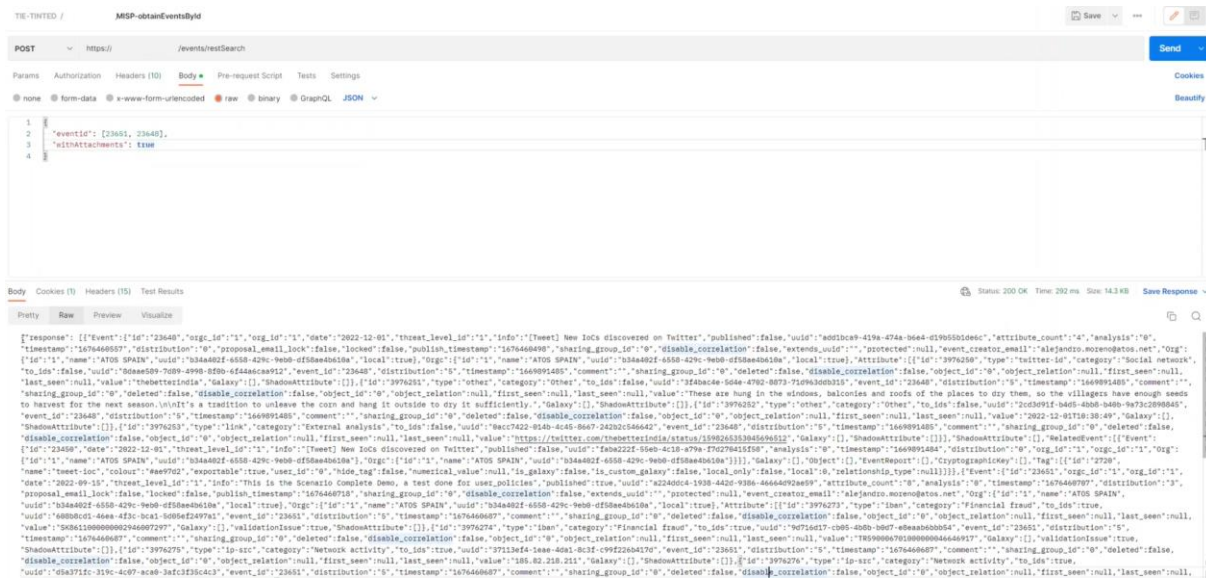


Figure 44. HTTP request to download MISP events.

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

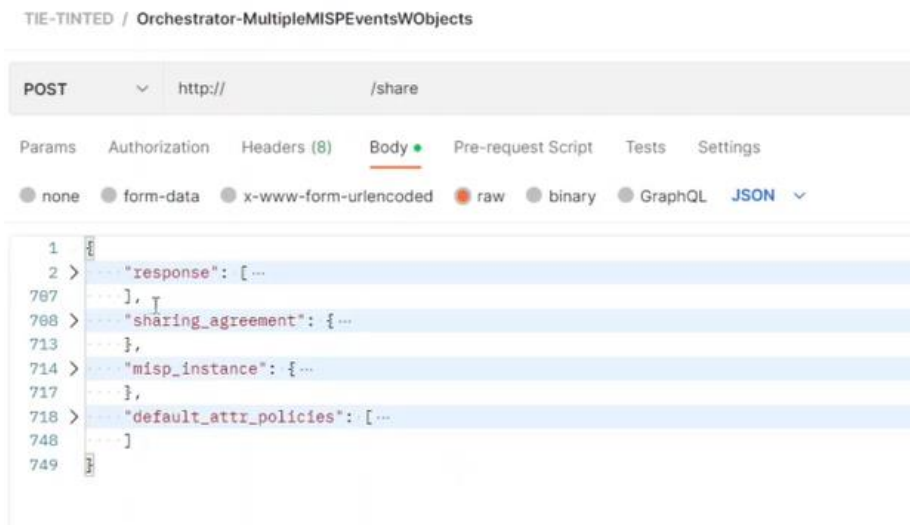


Figure 45. Orchestrator HTTP request format.

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.

```

1 {
2   "response": [
3     {
4       "Event": {
171     }
172   },
173   {
174     "Event": {
705     }
706   }
707 ],
708   "sharing_agreement": {
709     "from_user": "d491e08d-6cec-4560-9187-cd3f06345f69",
710     "to_user": "449340a2-712a-456e-8b18-42eeafedc8d4",
711     "from_date": 1676461798.0,
712     "to_date": 1694779352.0
713   },
714   "misp_instance": {
715     "url": "https://misp.example.com",
716     "api_key": "1234567890"
717   },
718   "default_attr_policies": [
748 ]
749 }

```

Figure 46. Sharing Agreement field.

```

1 {
2   "response": [
3     {
4       "Event": {
171     }
172   },
173   {
174     "Event": {
705     }
706   }
707 ],
708   "sharing_agreement": {
709     "from_user": "d491e08d-6cec-4560-9187-cd3f06345f69",
710     "to_user": "449340a2-712a-456e-8b18-42eeafedc8d4",
711     "from_date": 1676461798.0,
712     "to_date": 1694779352.0
713   },
714   "misp_instance": {
715     "url": "https://misp.example.com",
716     "api_key": "1234567890"
717   },
718   "default_attr_policies": [
748 ]
749 }

```

Figure 47. MISP instance field.

Concerning the user policies, Figure 48 provides an illustration of them.

```

"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's noteworthy 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 already have 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.

Finally, the last module that TINTED has is the Threat Intelligence Engine (TIE). It acquires events from the MISP instance and generates a threat score. This score can be divided into two parts: the initial segment is public, as it is founded on open-source data, and it gauges the threat based on diverse metrics like timeliness, trending, and comprehensiveness. The second segment, which is private, encompasses the aforementioned metrics along with the relevance heuristic that factors in CI infrastructure data. To shield against potential information leaks, this segment is encrypted, particularly due to the criticality of its assessment of the infrastructure's vulnerability against threats. Exposure of this information could have severe consequences, enabling attackers to exploit vulnerabilities and target the entity. By partitioning the score, we adhere to the principle of sharing data through the public metrics while simultaneously safeguarding the organization's vital data by encrypting it. Concurrently, we augment the received event's contextual information, a highly valuable enhancement. This entire procedure is referred to as CTI enrichment, described in Figure 49.

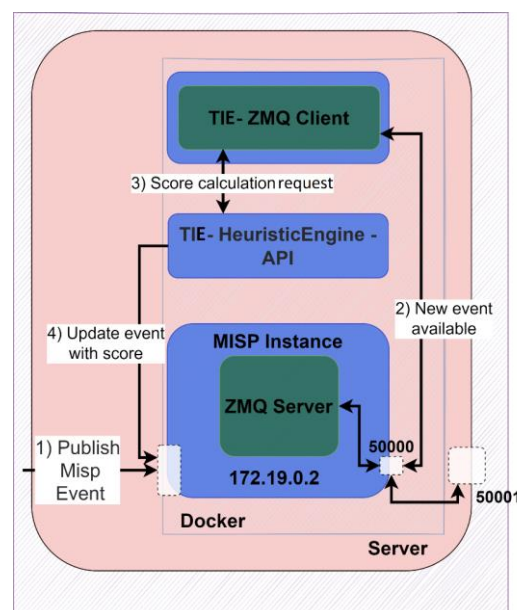


Figure 49. TIE architecture.

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.

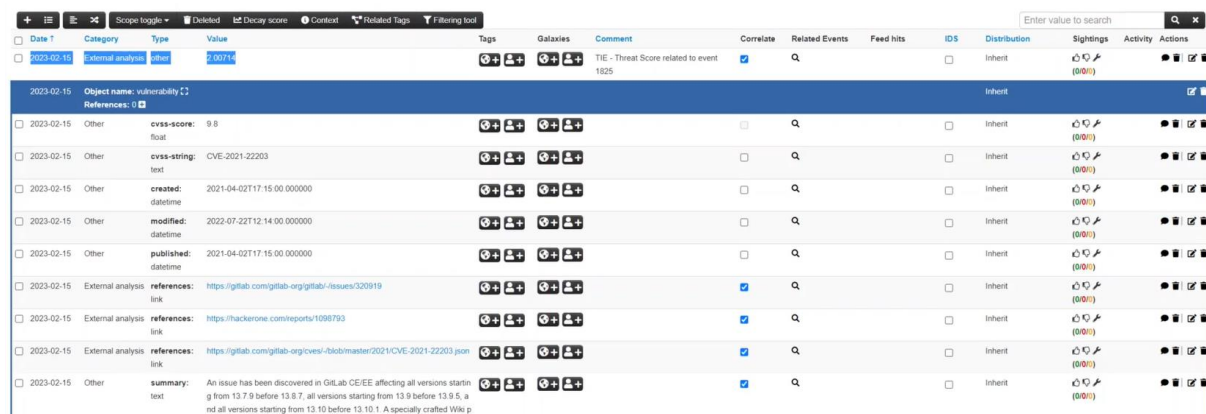
TIE does not encompass all MISP objects; instead, it concentrates on four specific objects considered highly valuable in the realm of threat intelligence: **Vulnerability**, **Domain-IP**, **BTC-Address**, and **File**. These four objects comprise a variety of attributes, including required and optional 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	#Cont	Creator user	Date	Info	Distribution	Actions
X ATOS	ATOS	1828		Ransomware, TIE-Score-Analysis=7.42	3	0	admin@admin.test	2023-02-15	Bitcoin address reported due to illicit activity	Organisation-C	🔍 🗑️ 📄
X ATOS	ATOS	1827		Paylist, Malware, TIE-Score-Analysis=9.1	3	0	admin@admin.test	2023-02-15	File linked to Agent Testa malware family	Organisation-C	🔍 🗑️ 📄
X ATOS	ATOS	1826		Malware, Paylist, malware_classification=malware-category=Stinner, TIE-Score-Analysis=2.18	3	0	admin@admin.test	2023-02-15	IP address linked to Agent Testa malware family	Organisation-C	🔍 🗑️ 📄
X ATOS	ATOS	1825		cert_incident_classification=vulnerability, cyber-threat-framework-Engagement=exploit-vulnerability, TIE-Score-Analysis=2.01	10	0	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-C	🔍 🗑️ 📄

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



Date	Category	Type	Value	Tags	Galaxies	Comment	Correlate	Related Events	Feed hits	IDS	Distribution	Signings	Activity	Actions
2023-02-15	External analysis	other	2.00714			TIE - Threat Score related to event 1825	🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	Other	cvss-score	9.8				🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	Other	cvss-string	CVE-2021-22203				🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	Other	created	2021-04-02T17:15:00.000000				🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	Other	modified	2022-07-22T12:14:00.000000				🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	Other	published	2021-04-02T17:15:00.000000				🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	External analysis	references	https://gitlab.com/gitlab-org/gitlab/-/issues/320919				🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	External analysis	references	https://hackerone.com/reports/1098793				🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	External analysis	references	https://gitlab.com/gitlab-org/cves/-/blob/master/2021/CVE-2021-22203.json				🔍	🔍			Inherit	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄
2023-02-15	Other	summary	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	🔍 🗑️ 📄	🔍 🗑️ 📄	🔍 🗑️ 📄

Figure 51. Vulnerability object.

I.III Risk Assessment

The first step to access the risk assessment module is to enter valid credentials to log in as shown in Figure 52.

LOGIN

Username:

Password:

Figure 52. Login form

Once we have logged in, the user will see a dashboard (Figure 53) with his personal information and a button to update any of the fields.

CERCA User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report

[Launch Risk Assessment](#)

User Profile

Username: sunrise
 Name: sunrisedemo
 Last name: demo
 email: user@sunrise
 Legal Entity: SUNRISE demo entity A (IT Department)
 Current Data Processing Activity: data_share

CERCA v1.3

Figure 53. User profile menu

The following two screenshots, Figure 54 and Figure 55, show information from the user profile with more depth. Regarding Figure 55 we can also analyse the assets that we have served as input for the risk assessment module.

CERCA User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report

[Launch Risk Assessment](#)

Legal Entities Configuration

Name	Department	Description	Responsible
SUNRISE demo entity A	IT Department	SUNRISE demo entity A	SUNRISE admin (admin@sunrise)

CERCA v1.3

Figure 54. Legal entities configuration menu

CERCA User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report

[Launch Risk Assessment](#)

Data Processing Activities Configuration

Name	Controllers	Processors	Subprocessors	Third-Parties	Critical
data_share	SUNRISE demo entity A (IT Department)	SUNRISE demo entity A (IT Department)	SUNRISE Institute Demo (Information Security)	SUNRISE Research Centre	False View Assets

CERCA v1.3

Figure 55. Data processing activities configuration

Figure 56 shows the menu where all the assets are displayed following the CIA triad (confidentiality, integrity, and availability).

CERCA User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report

[Launch Risk Assessment](#)

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

Name	Category	Legal Entity	Data Categories	Availability	Confidentiality	Integrity
Asset	physical	Critical Infrastructure	personal,confidential	4	6	2

CERCA v1.3

Figure 56. Asset view dashboard

In Figure 57, a list of threats, risks and security measures is displayed. We can also select a risk model or clear the current selection.

CERCA User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report

[Launch Risk Assessment](#)

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

Identified Threats [Identified Risks](#) [Identified Security Measures](#) [Risk Models](#)

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.

CERCA v1.3

Figure 57. Model configuration dashboard

After clicking the “Select risk models” hyperlink, we are redirected to another menu, which is shown in Figure 58. The tool can suggest a risk model, but the user is free to choose whatever risk model he thinks that might suit better to his infrastructure.

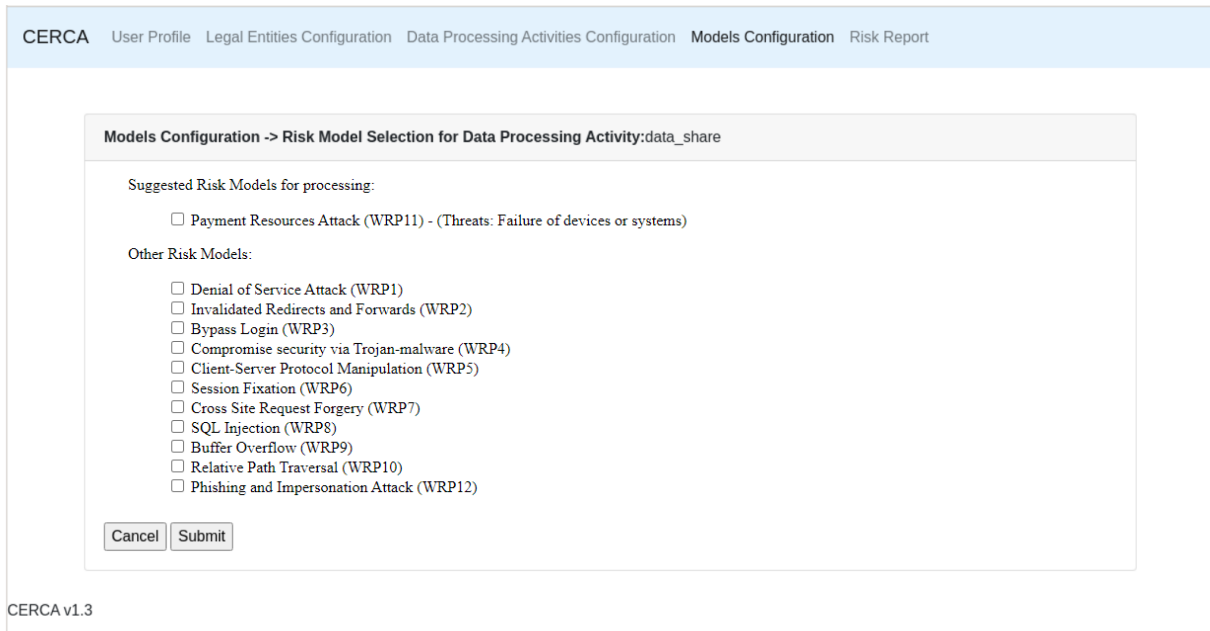


Figure 58. Risk model selection dashboard

After selecting one of the risk models, we can observe that the models' configuration menu (Figure 59) is updated with the new information.

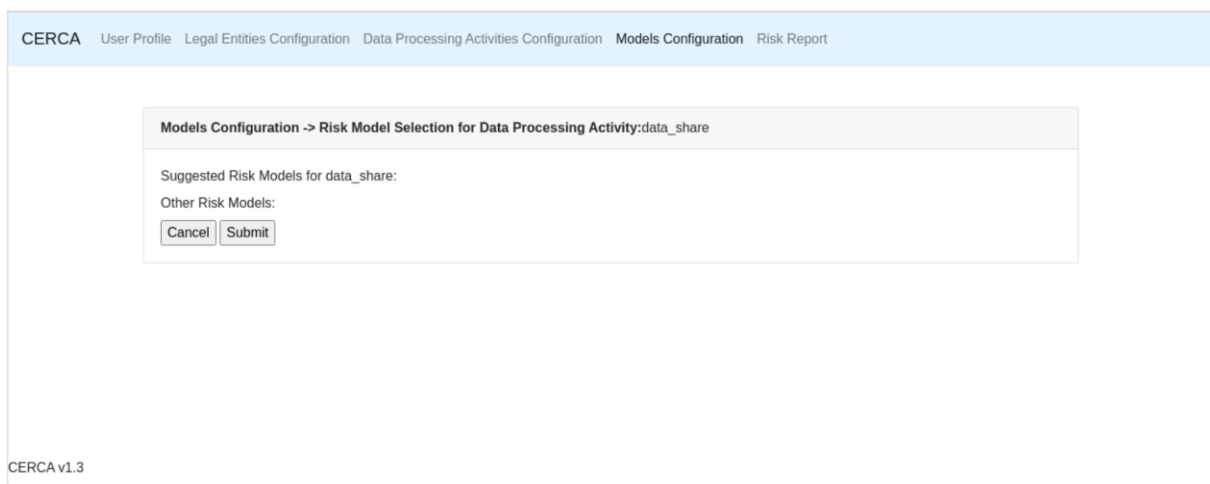


Figure 59. Model configuration dashboard updated with risk model

One of the most important inputs that CERCA needs to process the cyber risk calculation is the questionnaire. In Figure 60 we can observe an example of it.

CERCA User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report

Legal Entities Configuration -> Legal Entity Questionnaire: SUNRISE demo entity A

Q1/39 - Where are your company Head Offices located?

Use of the information on the company business profile

North America
 South & Central America
 Asia
 Europe
 Other

Q2/39 - 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

Q3/39 - Indicate the sensitivity level of the information your company maintains and processes, on average.

Figure 60. Questionnaire for risk model

Finally, after filling the necessary information, we would be able to get the analysis performed by the tool. Figure 61 shows that the risk report works either in a qualitative and quantitative way.

CERCA User Profile Legal Entities Configuration Data Processing Activities Configuration Models Configuration Risk Report

Risk Report: SUNRISE demo entity A

Overall cyber-risk status:

Average value **MEDIUM**

Risk Model:	WRP101: Malware Attack	LOW
Risk WRP101-R1:	Malware attack with loss of Availability	VERY LOW
Risk WRP101-R2:	Malware attack with loss of Confidentiality	LOW
Risk WRP101-R3:	Malware attack with loss of Integrity	VERY LOW
Risk Model:	WRP102: Denial of service Attack	MEDIUM
Risk WRP102-R1:	Denial of service attack with loss of Availability	LOW
Risk WRP102-R2:	Denial of service attack with loss of Confidentiality	MEDIUM

CERCA v1.3

Figure 61. Risk Report summary

I.IV Incident Reporting

This manual is an initial version based on the previous manual of the CyberSec4Europe project [9]. Field names and screenshots are for financial entities and bank regulations because the behavior and functionalities of the module is equal for both cases, and we have not addressed the changes in the graphical interfaces. However, the interfaces and guides will be updated in future releases of the modules and deliverables.

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.

I.IV.I WEB-GUI

The AIRE web allows to create report templates, mapping incident fields with report forms and setting timers, and manage the reports of the incidents; showing the preview, checking that all fields are fulfilled, and finally generating and sending the report.

I.IV.I.I AIRE GUI – ADMINISTRATOR

Before starting to work normally with AIRE, it is necessary to register an organization, users, contact addresses, regulations, etc. To do so, it is necessary to login as admin, then there is the list of tabs for the different configurations, Figure 62:




Figure 62. AIRE navigation bar

1. **Entities:** Show the list of organizations that are registered on AIRE, Figure 63:

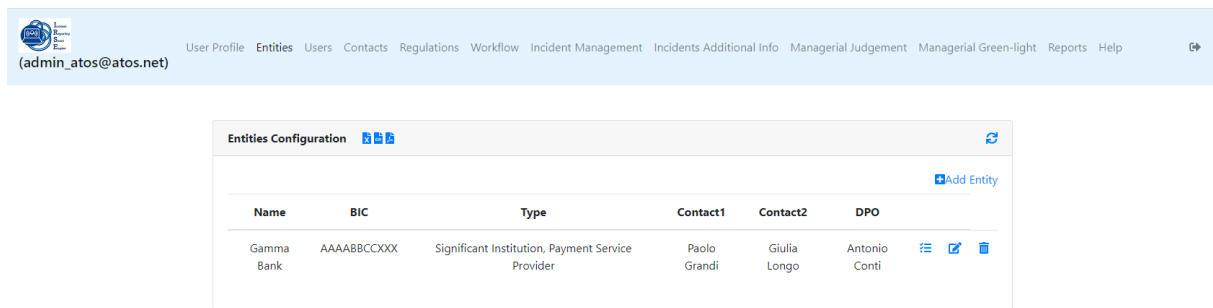


Figure 63. Entities Configuration List

Add Entity or the edit button open a form that allows create or modify a new organization, Figure 64:

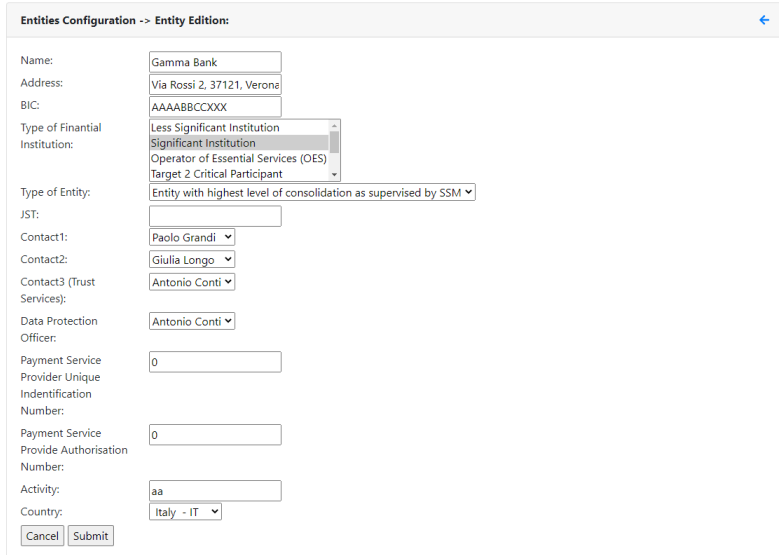


Figure 64. New entity form

Then it necessary to choose the regulation that the entity must meet, only reports associated to enabled regulations will be generated by the platform, Figure 65.

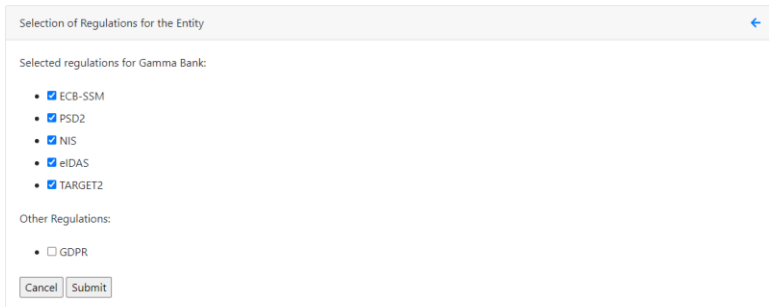



Figure 65. Regulation for entities

2. **Users:** Show the list of registered users with their information: contact information, position, or function. This page allows to create new users, *Add User*, or edit the information of the created users, Figure 66.




Users Configuration						
Name	Surname	Email	Phone	Position	Function	
Incident Management Team	User	imt_atos@atos.net		Specialist	Asset Owner / Incident Management Team	✎ 🗑️
IncidentReporting Team	User	susan.gz@gmail.com	None	Specialist	Incident Reporting Team	✎ 🗑️
Incident Classification Team	User	iclt_atos@atos.net		Specialist	Incident Classification Team	✎ 🗑️
Susana	Gonzalez	admin_atos@atos.net		Specialist	Admin	✎ 🗑️
Controller	User	controller_atos@atos.net		Specialist	Controller	✎ 🗑️

Figure 66. List of registered users

3. **Contacts:** List of all contacts registered in AIRE, Figure 67. There are the following contact types:

- Contact1: primary contact
- Contact2: secondary contact
- Contact3: associated to the trust officer
- DPO: Data Protection Officer

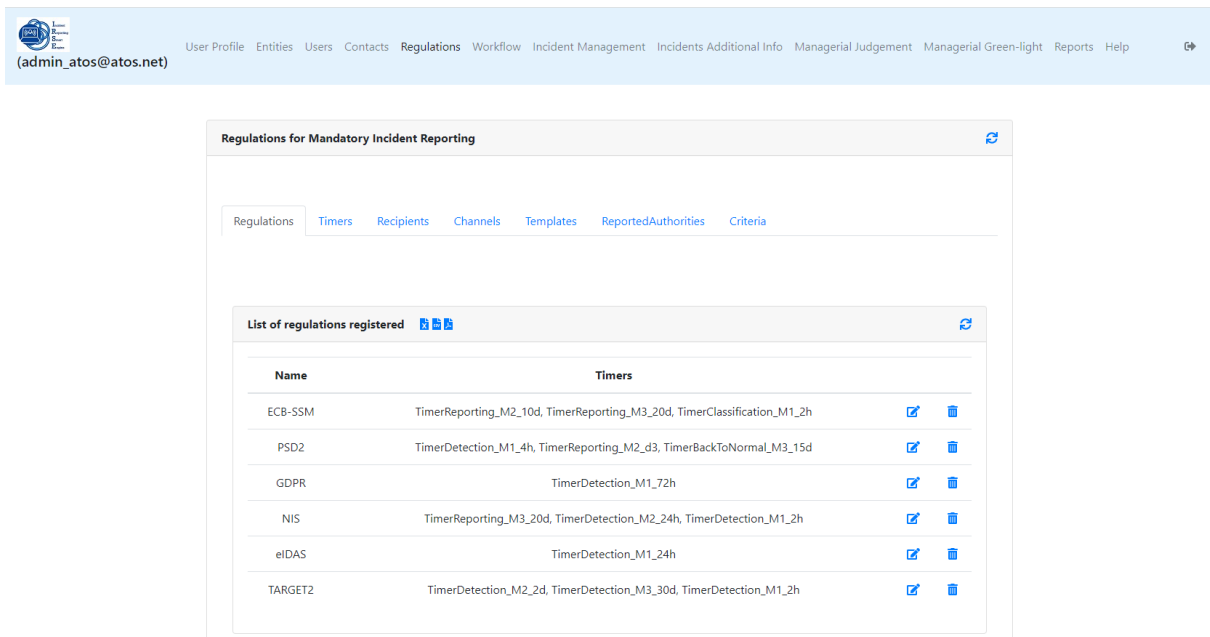


Contacts Configuration						
Name	Surname	Email	Phone	Title	Type	
Paolo	Grandi	Paolograndi@gammabank.com	+393476452952	Controller	Contact1	✎ 🗑️
Giulia	Longo	Giulialongo@gammabank.com	+393459385019	Incident Reporting Team	Contact2	✎ 🗑️
Antonio	Conti	Antonioconti@gammabank.com	+393475930593	Data Protection Officer	DPO	✎ 🗑️
Antonio	Conti	Antonioconti@gammabank.com	+393475930593	Trust Officer	Contact3	✎ 🗑️

Figure 67. List of configured contacts

4. **Regulations:** Regulations for Mandatory Incident Reporting. This section manages the different regulations and has several subsections:

4.1. Regulations: List of registered regulations, Figure 68.



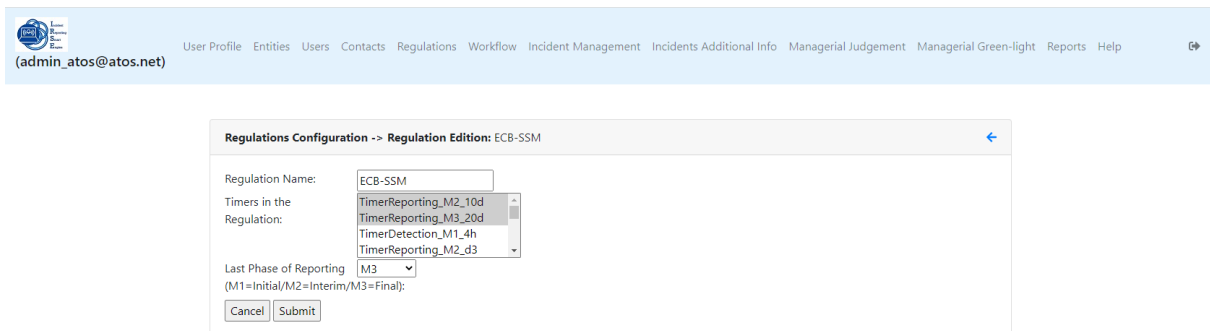
The screenshot shows a web application interface for 'Regulations for Mandatory Incident Reporting'. The user is logged in as 'admin_atos@atos.net'. The main content area displays a 'List of regulations registered' table with the following data:

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 68. List of regulations registered

For each regulation registered in the platform, Figure 68, it is necessary to indicate:

- Last phase of reporting: depending on the reports that have to be disseminated according to a specific directive or regulation, it 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 next point about Timers)



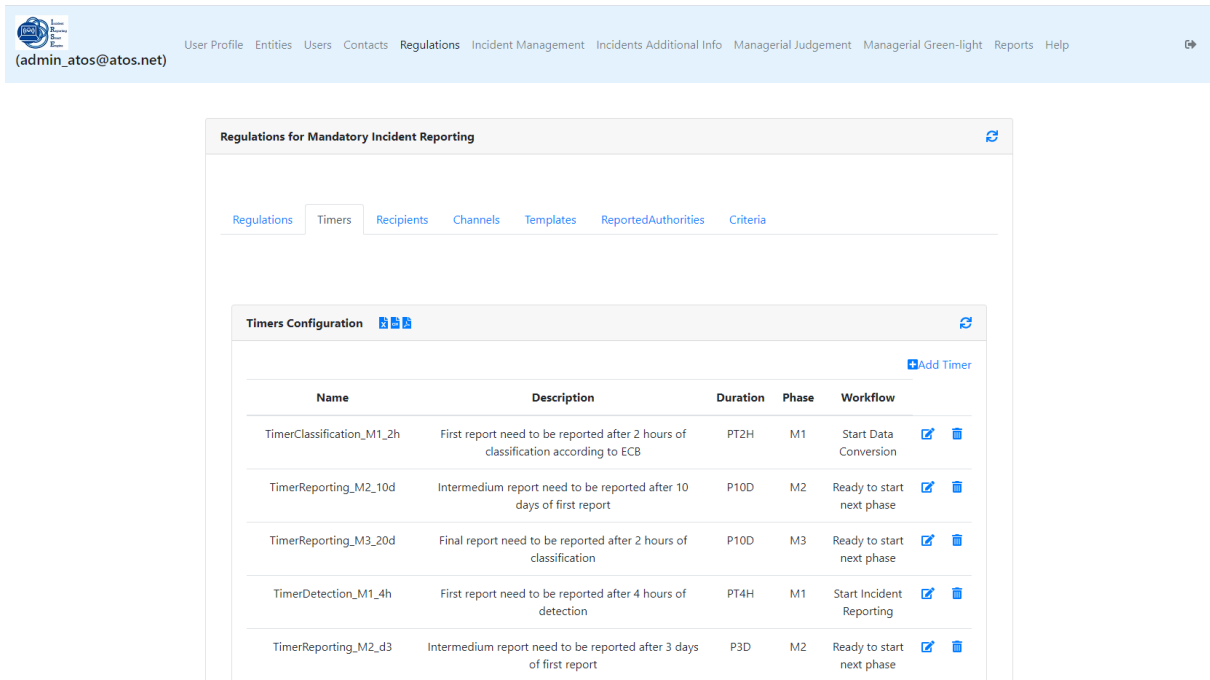
The screenshot shows the 'Regulations Configuration -> Regulation Edition: ECB-SSM' form. The fields are as follows:

- Regulation Name: ECB-SSM
- Timers in the Regulation: A list box containing 'TimerReporting_M2_10d', 'TimerReporting_M3_20d', 'TimerDetection_M1_4h', and 'TimerReporting_M2_d3'.
- Last Phase of Reporting: M3 (with a note: (M1=Initial/M2=Interim/M3=Final))

Buttons for 'Cancel' and 'Submit' are visible at the bottom of the form.

Figure 69. New regulation form

4.2. **Timers:** Indicate when a notification needs to be sent to the incident contact user (the email configured as Contact User in the TheHive template), Figure 68. In case a mandatory report has not been sent to the corresponding Supervisory Authorities in the deadline defined by a specific regulation.

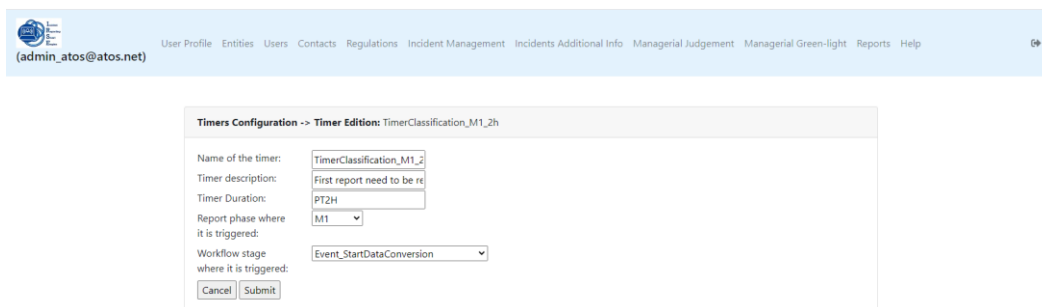


The screenshot shows a web interface for 'Regulations for Mandatory Incident Reporting'. The 'Timers' tab is selected, displaying a table of timer configurations. The table has columns for Name, Description, Duration, Phase, and Workflow. There are five rows of timer configurations listed.

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 70. List of timers

Timer Edition, Figure 71, allows to create new timers or modify existing ones. The *timer duration* field defines the time windows within the report must be sent and uses ISO 8601 durations format⁶. *Report phase* defines the phase in which the report should be sent. And the *Workflow stage* define the event that trigger the current timer.



The screenshot shows the 'Timer Edition' form for 'TimerClassification_M1_2h'. The form contains several input fields and dropdown menus for editing the timer configuration.

Timer Configuration -> Timer Edition: TimerClassification_M1_2h

Name of the timer:

Timer description:

Timer Duration:

Report phase where it is triggered:

Workflow stage where it is triggered:

Figure 71. Timer edition form

⁶ https://en.wikipedia.org/wiki/ISO_8601

AIRE sends an email to the responsible party when a report is not sent on time, Figure 72.

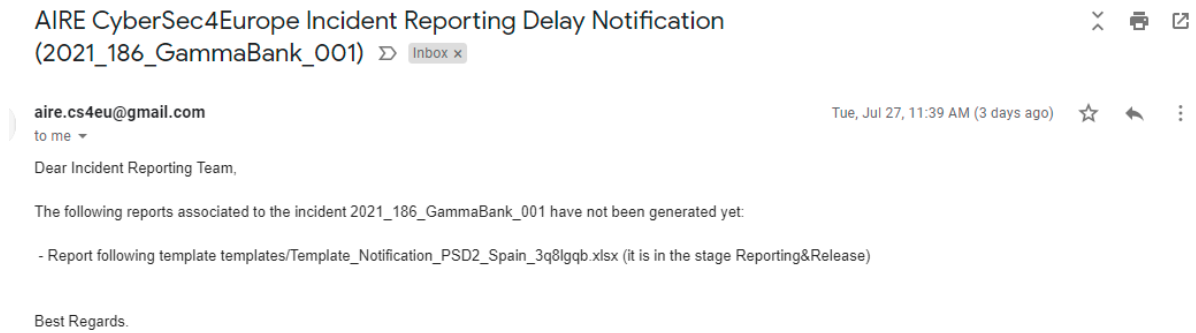


Figure 72. Example of email sent by AIRE

- 4.3. **Recipients:** Associated to each entity and regulation, Figure 73. It will have also a Channel associated, which will be shown once the report has been generated by the platform as a suggestion of channel (e.g., email address) that need to be followed for the reporting.

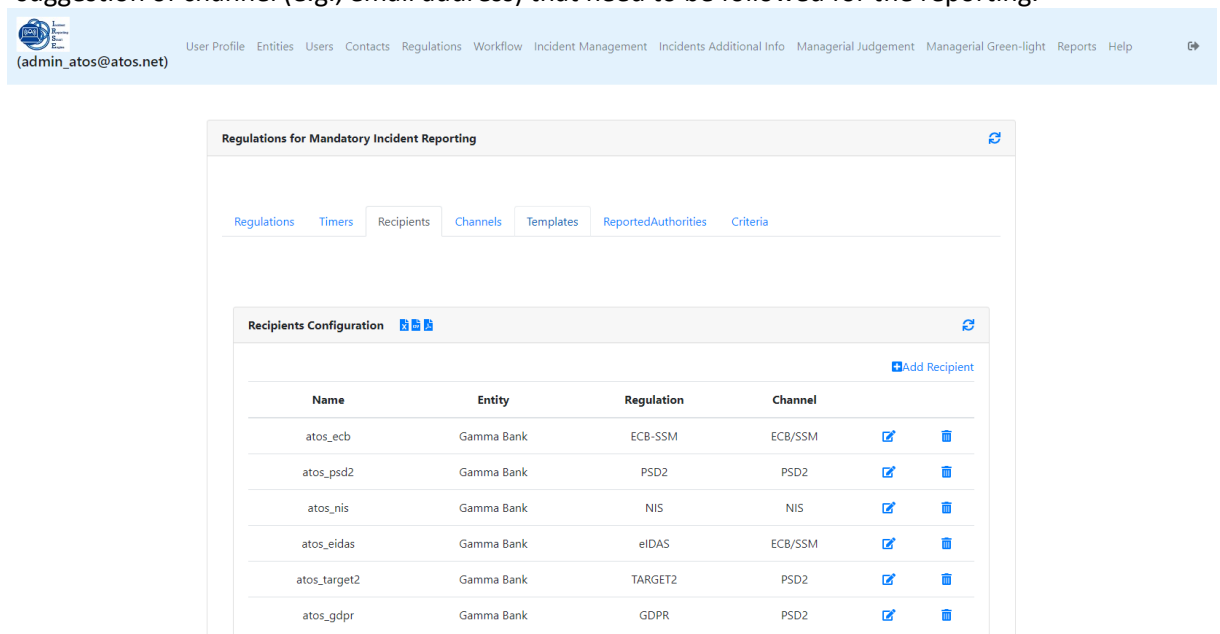


Figure 73. List of Recipient Configurations

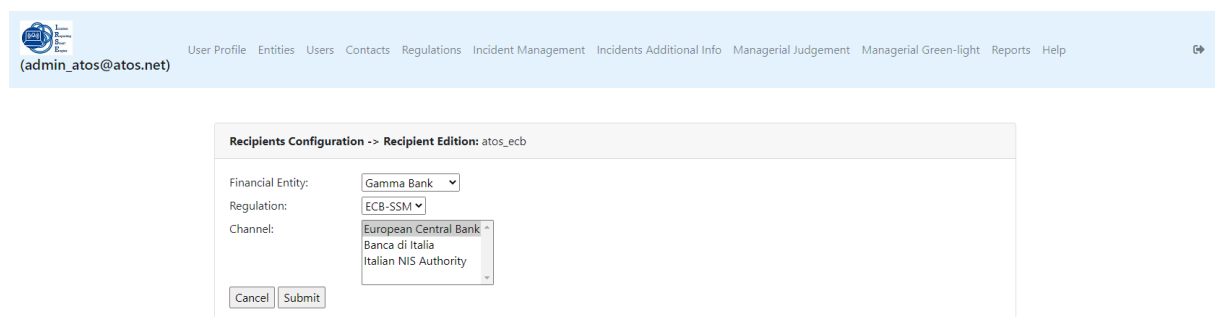
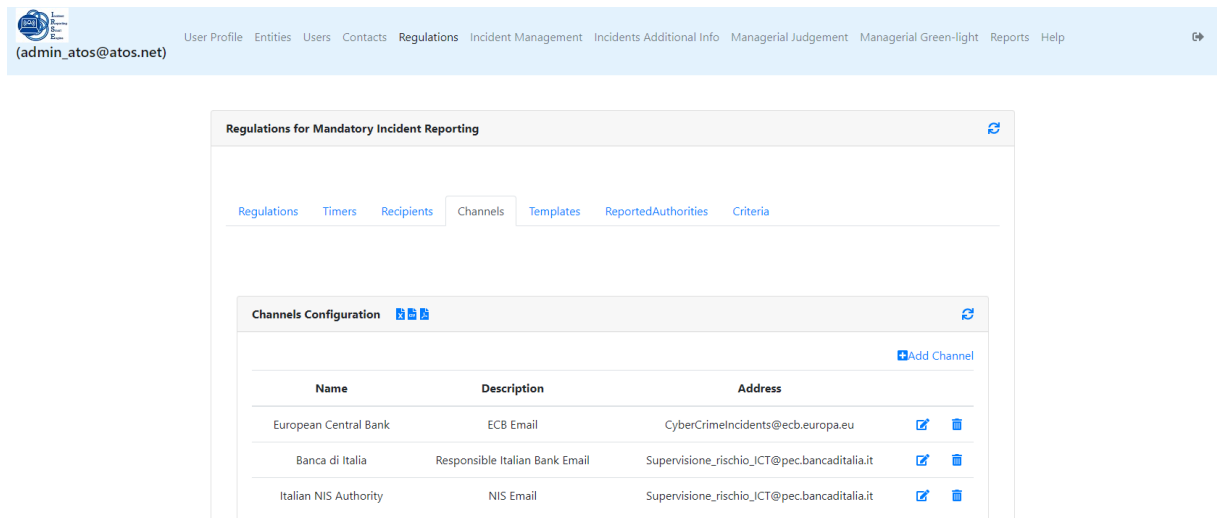


Figure 74. Recipient Edition Form

4.4. **Channels:** The information registered here, Figure 75, will be used in the Recipients and shown to the user as a suggestion when the reports are ready for revision and releasing.

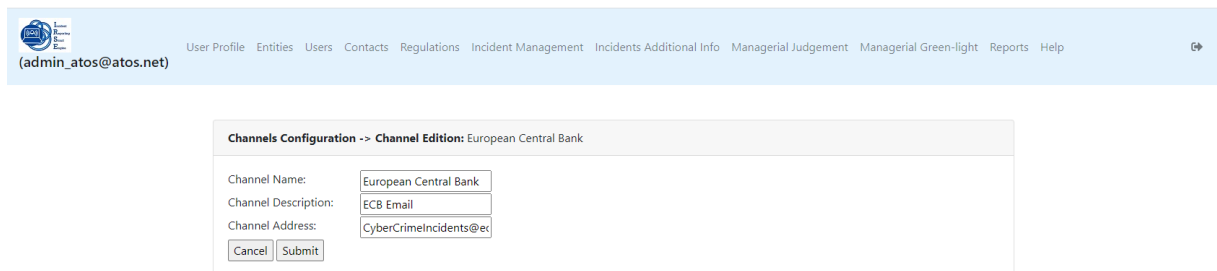


The screenshot shows the 'Regulations for Mandatory Incident Reporting' interface. The 'Channels' tab is selected, displaying the 'Channels Configuration' section. An 'Add Channel' button is visible. The table below lists the registered channels:

Name	Description	Address		
European Central Bank	ECB Email	CyberCrimeIncidents@ecb.europa.eu		
Banca di Italia	Responsible Italian Bank Email	Supervisione_rischio_JCT@pec.bancaditalia.it		
Italian NIS Authority	NIS Email	Supervisione_rischio_JCT@pec.bancaditalia.it		

Figure 75. List of communication channels

Add Channel button allow adding or modifying channels, Figure 76.

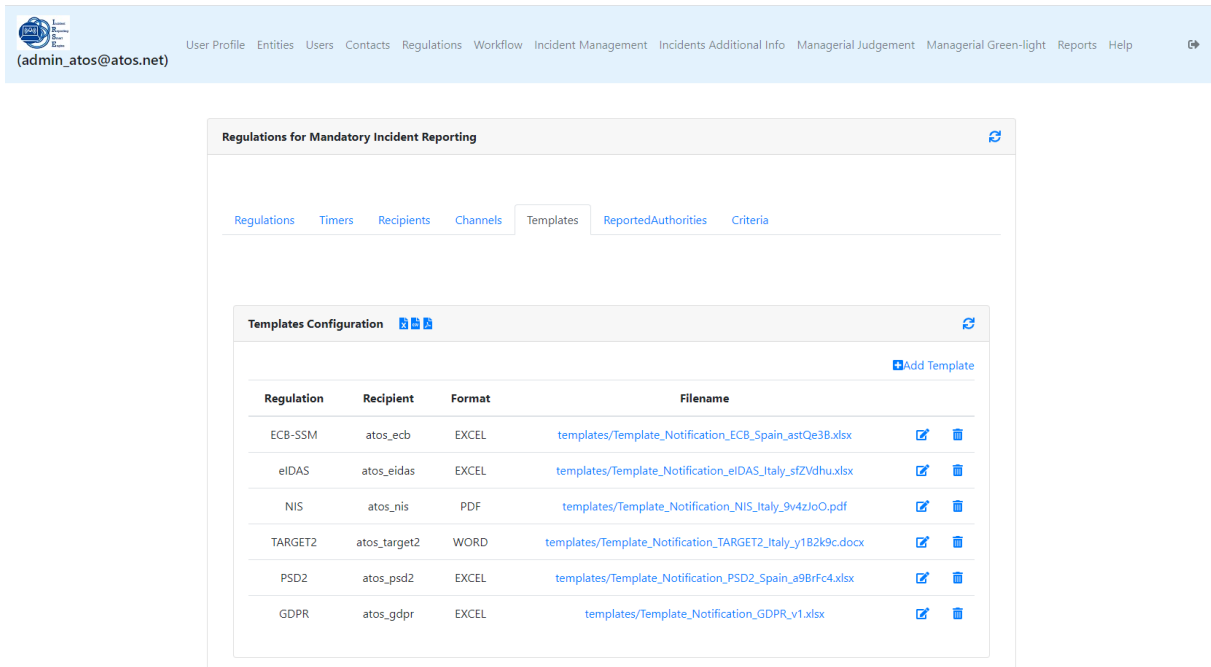


The screenshot shows the 'Channel Edition' form for the 'European Central Bank' channel. The form contains the following fields and buttons:

- Channel Name:
- Channel Description:
- Channel Address:
- Buttons:

Figure 76. Channel Edition

4.5. **Templates:** used by the platform for the generation of the reports will need to be associated to 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 times in the reports. Date and time formats used are the ones defined in SimpleDateFormat⁷.

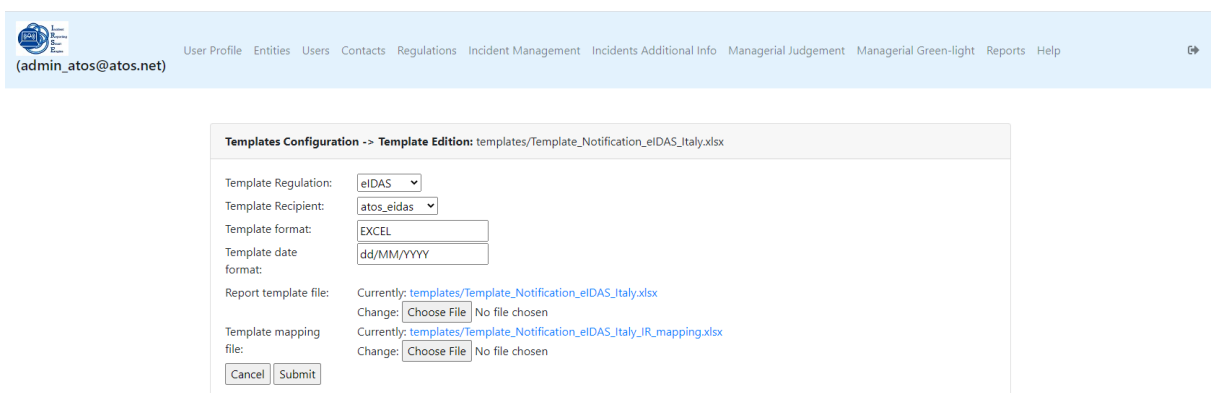


The screenshot shows the 'Regulations for Mandatory Incident Reporting' interface. The 'Templates' tab is selected, displaying a 'Templates Configuration' table. The table lists various regulations and their associated templates.

Regulation	Recipient	Format	Filename
ECB-SSM	atos_ecb	EXCEL	templates/Template_Notification_ECB_Spain_astQe3B.xlsx
eIDAS	atos_eidas	EXCEL	templates/Template_Notification_eIDAS_Italy_sfZVdhu.xlsx
NIS	atos_nis	PDF	templates/Template_Notification_NIS_Italy_9v4zJoO.pdf
TARGET2	atos_target2	WORD	templates/Template_Notification_TARGET2_Italy_y1B2k9c.docx
PSD2	atos_psd2	EXCEL	templates/Template_Notification_PSD2_Spain_a9BrF4.xlsx
GDPR	atos_gdpr	EXCEL	templates/Template_Notification_GDPR_v1.xlsx

Figure 77. 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 allows the extensions *pdf*, *doc*, *docx*, *xls*, and *xlsx*; the extension must be specified on 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 78 shows an example.



The screenshot shows the 'Template Edition Form' for the template 'templates/Template_Notification_eIDAS_Italy.xlsx'. The form includes the following fields:

- Template Regulation: eIDAS
- Template Recipient: atos_eidas
- Template format: EXCEL
- Template date format: dd/MM/YYYY
- Report template file: Currently: templates/Template_Notification_eIDAS_Italy.xlsx; Change: [Choose File] No file chosen
- Template mapping file: Currently: templates/Template_Notification_eIDAS_Italy_IR_mapping.xlsx; Change: [Choose File] No file chosen

Buttons for 'Cancel' and 'Submit' are visible at the bottom.

Figure 78. Template Edition Form

4.6. **Reported Authorities:** In this menu, Figure 79, 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

⁷ <https://docs.oracle.com/javase/7/docs/api/java/text/SimpleDateFormat.html>

authorities that need to be notified in case the criteria and thresholds associated to those specifications are matched.

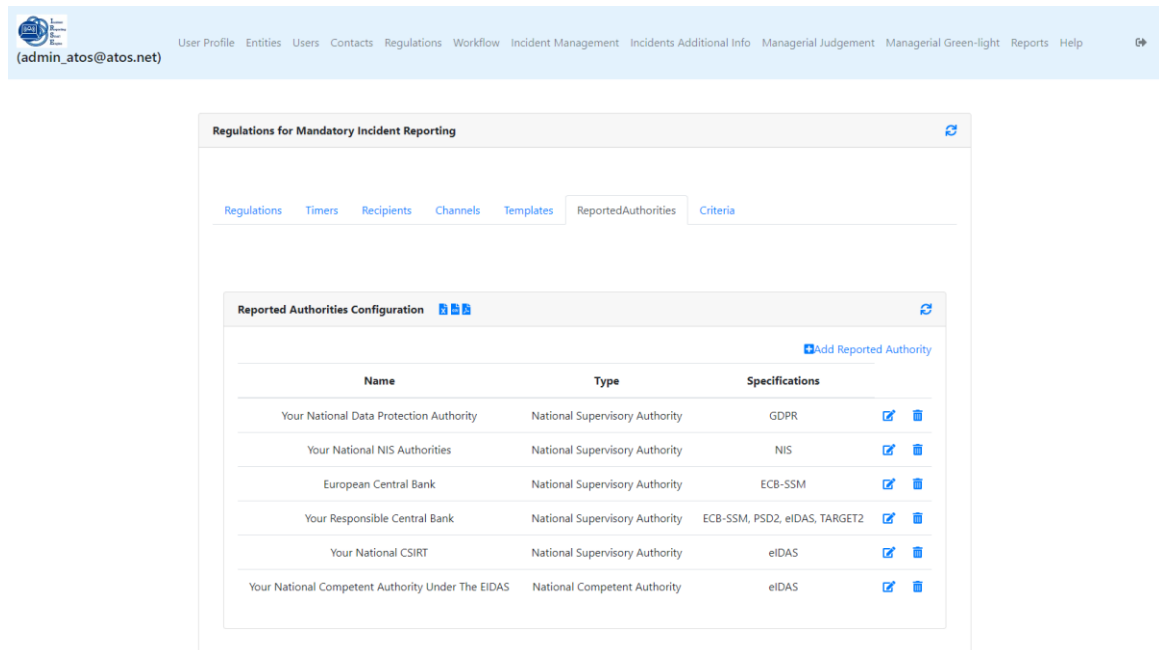


Figure 79. Reported Authorities Configuration

The Reported Authorities Editor, Figure 80, allows create or modify the entries of Reported Authorities.

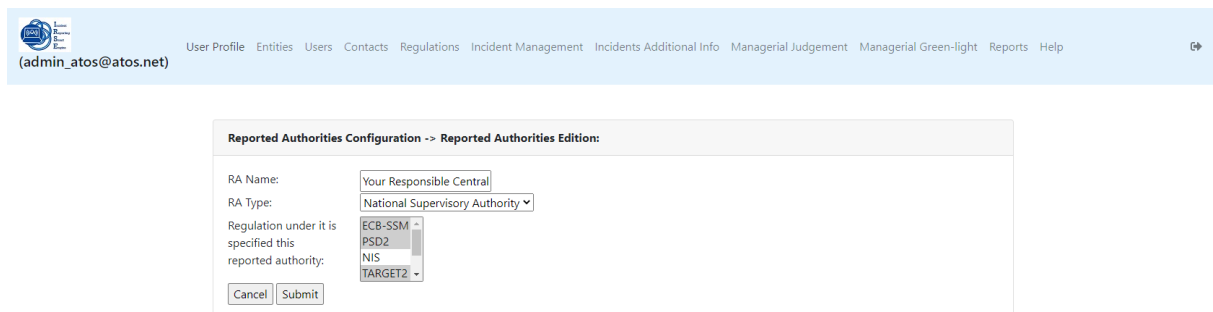
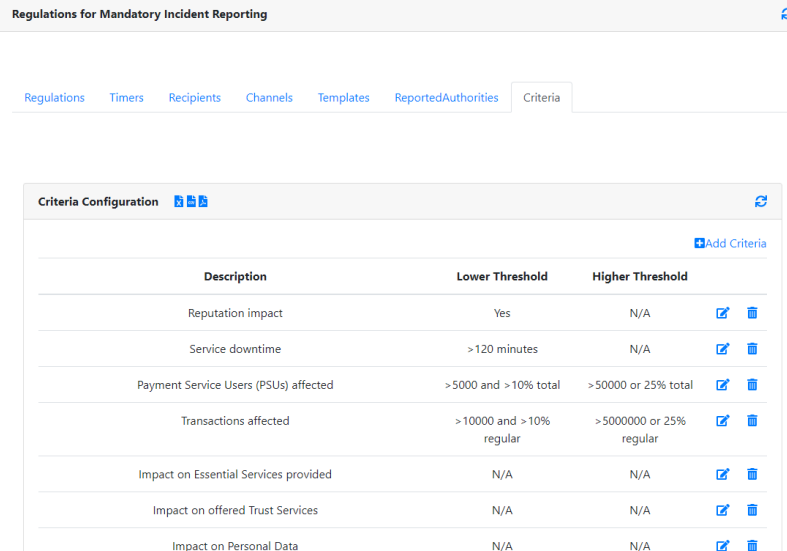


Figure 80. Reported Authorities Editor

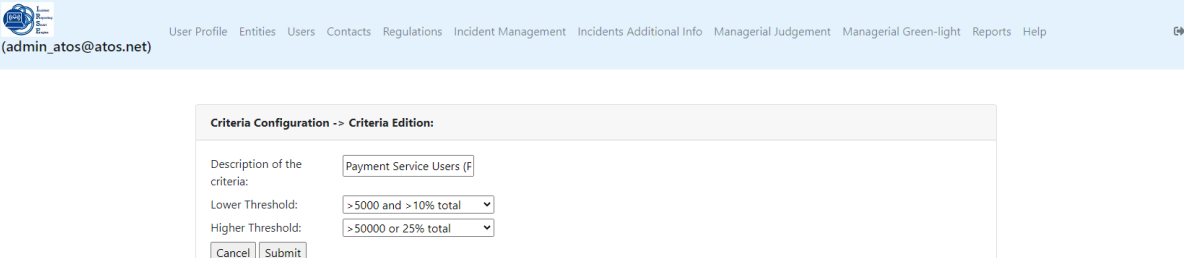
4.7. **Criteria:** Figure 81 shows the criteria supported by the platform for the event classification. The current version of the demonstrator does not support customization of the criteria included in the regulations. These criteria are included here are preloaded and just included for information purposes, but they are not considered in real-time by the responder IR Event Classifier included in the demonstrator. Consequently, if some of them is removed or some is added, the changes will not be considered by the classifier.



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 81. List of Criterias

Figure 82 displays the *Criteria Edition* View where the user can set the lower and higher thresholds.



Criteria Configuration -> Criteria Edition:

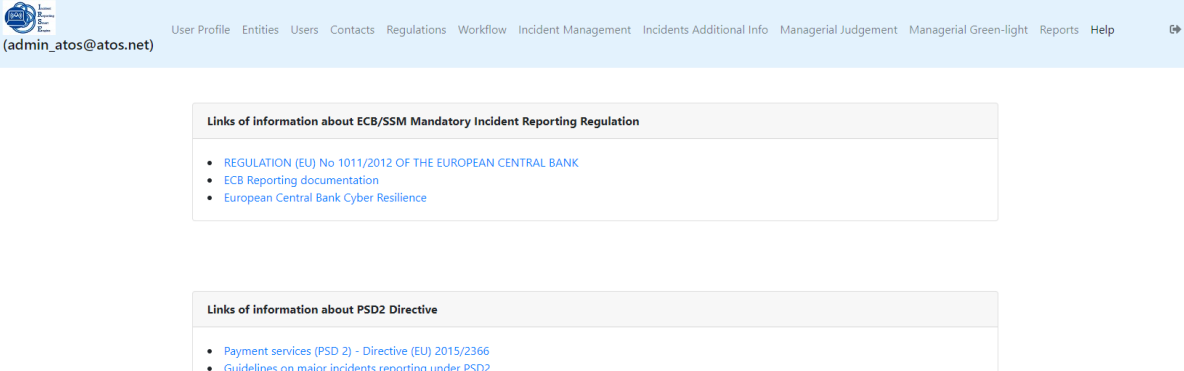
Description of the criteria:

Lower Threshold:

Higher Threshold:

Figure 82. Criteria Edition

More information about regulations can be found under the menu *Help*, Figure 83:



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 83. Help menu

I.IV.I.II AIRE GUI – USER

Once AIRE has been configured by the administrator, users can manage incidents and send reports. This section describes the functionalities that the assets offer.

The incidents are managed by the TheHive⁸ tool. When “Incident Management” is clicked a pop-up window will be shown with the graphical interface provided by TheHive, Figure 84.

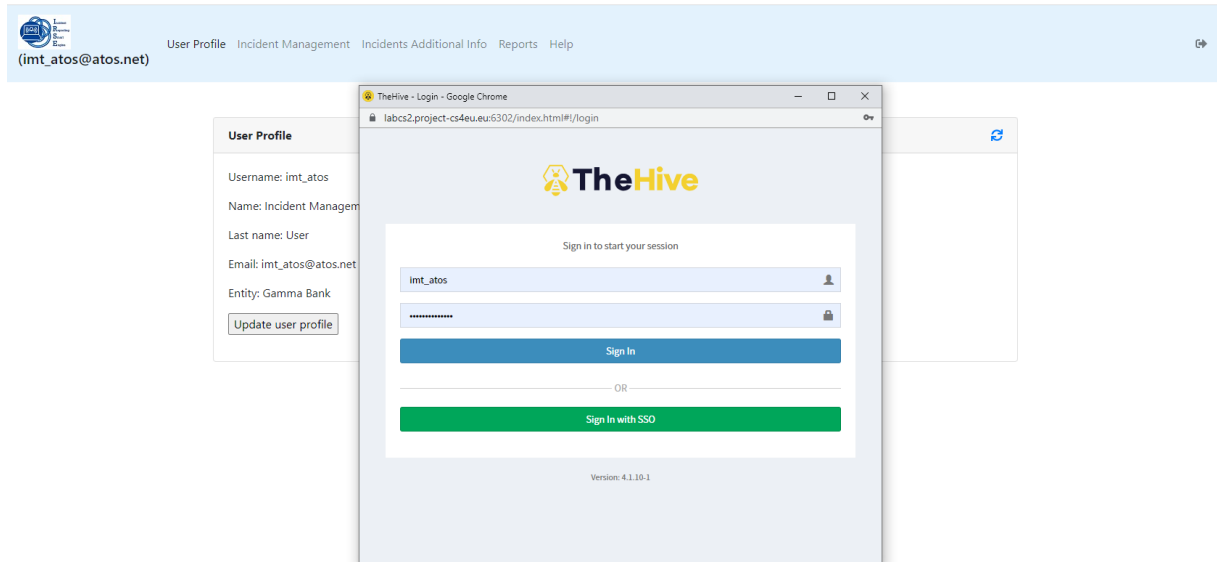


Figure 84. TheHive login interface

After signing in, the TheHive graphical interface is embedded within the AIRE interface, Figure 85. Where the user can perform several actions related with incident management.

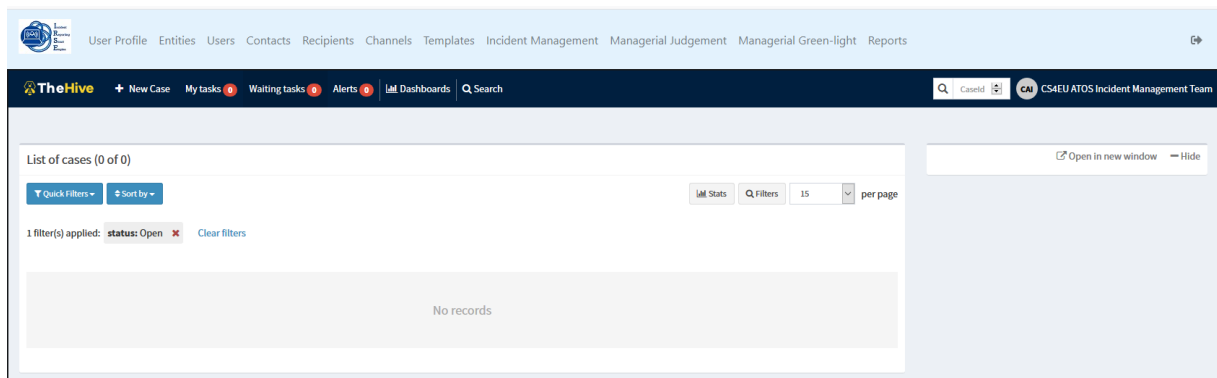


Figure 85. TheHive embedded in AIRE interface

1. **New Case:** Start the process of creating a case for an incident. This may be an empty case or use a predefined template, Figure 86. The template “Incident Report” includes the fields required for the mandatory incident reporting, Figure 87.

⁸ <http://thehive-project.org/>

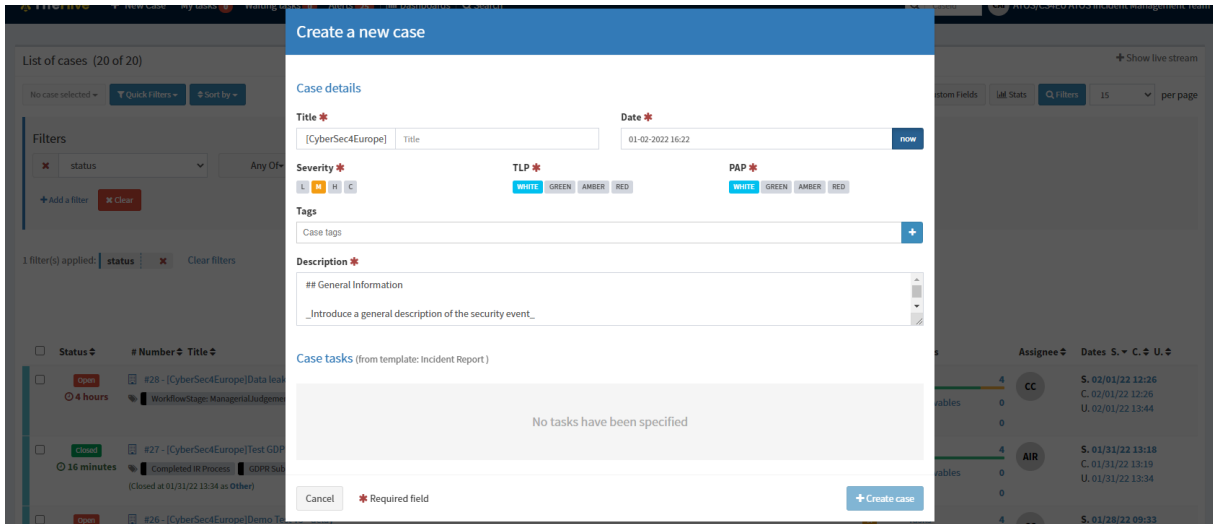


Figure 86. Template selection for a new Case

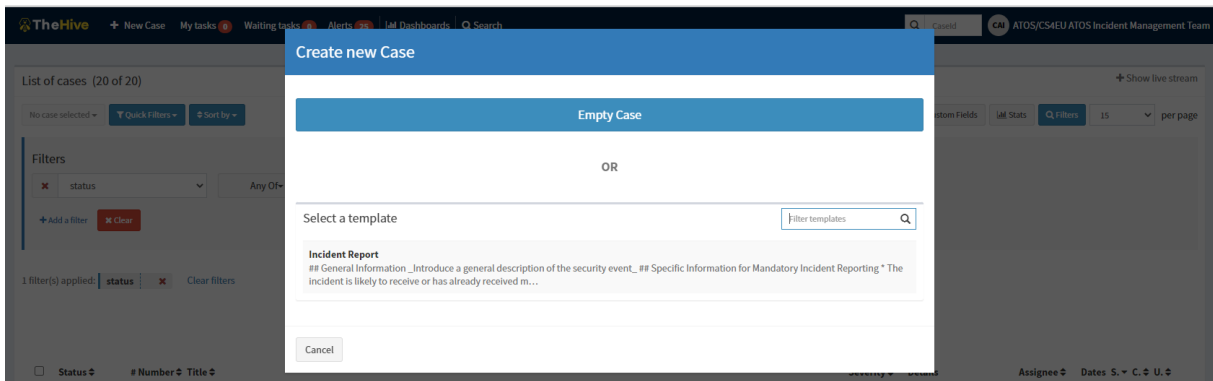


Figure 87. New case form with Incident Respot template

This action creates an empty case in TheHive, Figure 88:

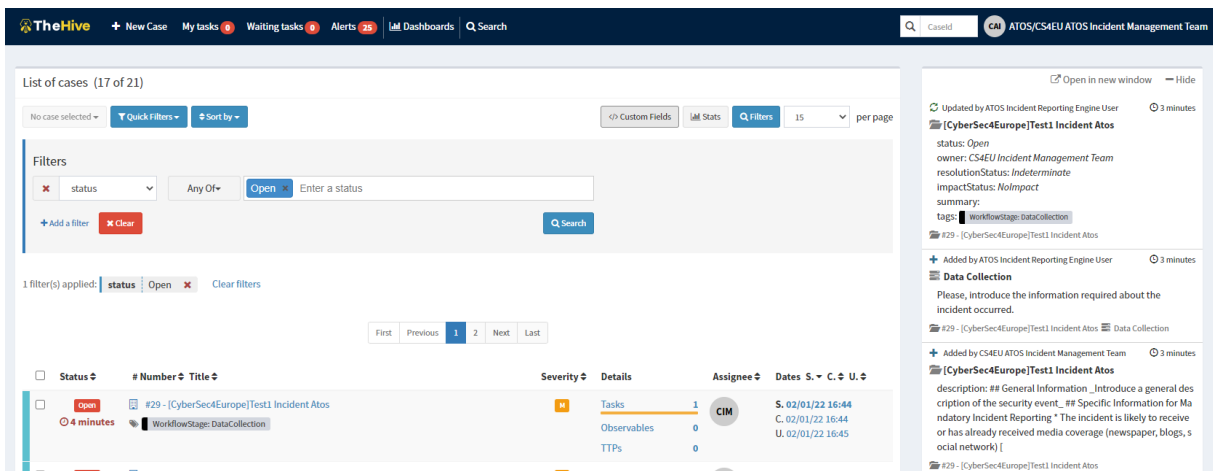
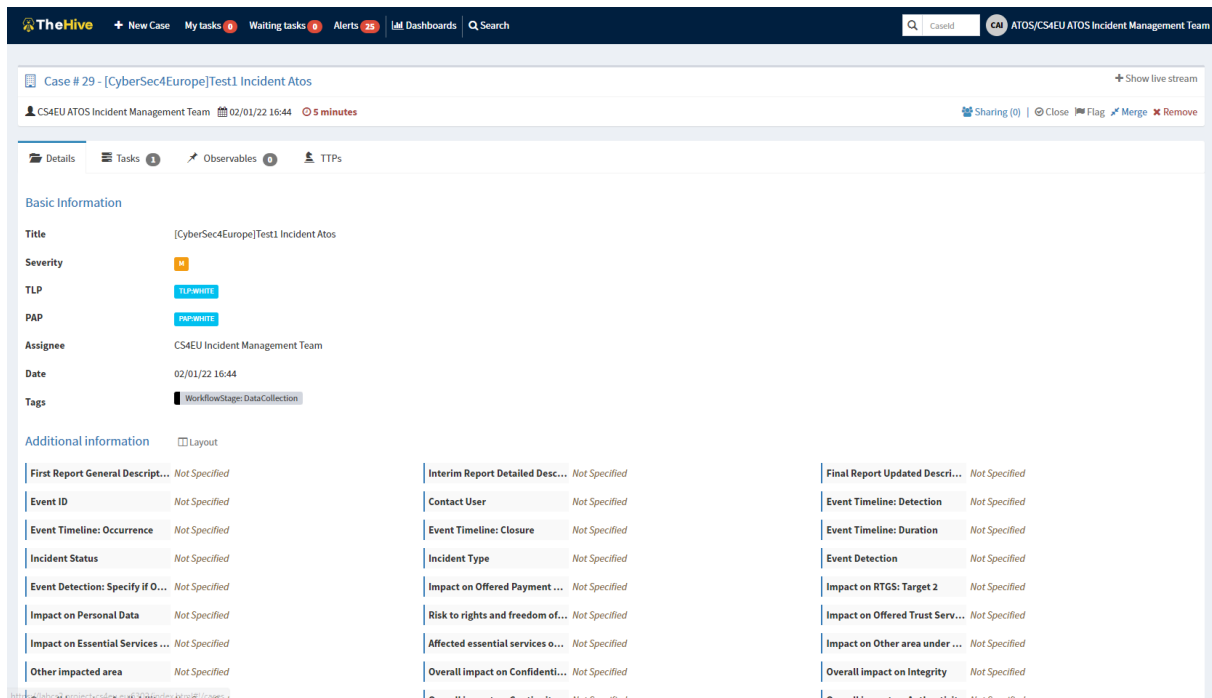


Figure 88. New incident in the list of cases

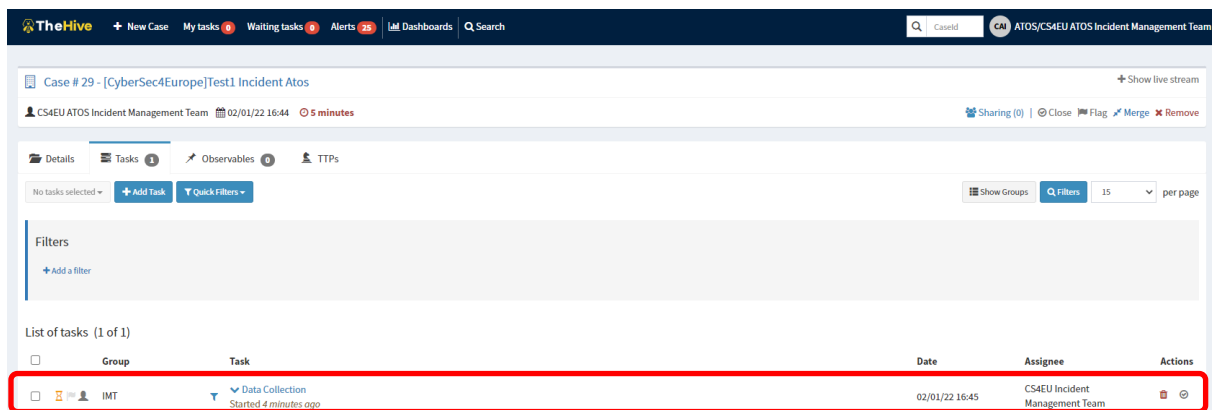
Figure 89 displays the details of the case.



The screenshot shows the 'Details' tab of a case in TheHive. The case title is '[CyberSec4Europe]Test1 Incident Atos'. The severity is 'High' (H), TLP is 'TLP:WHITE', and PGP is 'PGP:WHITE'. The assignee is 'CS4EU Incident Management Team' and the date is '02/01/22 16:44'. A tag 'WorkflowStage: DataCollection' is present. Below this, there is a grid of 'Additional information' fields, including 'First Report General Description', 'Event ID', 'Incident Status', 'Event Detection: Specify if O...', 'Impact on Personal Data', 'Impact on Essential Services', 'Other impacted area', 'Interim Report Detailed Description', 'Contact User', 'Event Timeline: Closure', 'Incident Type', 'Impact on Offered Payment...', 'Risk to rights and freedom of...', 'Affected essential services', 'Overall impact on Confidentiality', 'Final Report Updated Description', 'Event Timeline: Detection', 'Event Timeline: Duration', 'Event Detection', 'Impact on RTGS: Target 2', 'Impact on Offered Trust Services', 'Impact on Other area under...', and 'Overall impact on Integrity'. Most of these fields are currently 'Not Specified'.

Figure 89. Details of a case

In the *Tasks* tab there are the pending task for the case, Figure 90. *Data Collection* task is automatically created and assigned to the *IMT* group.

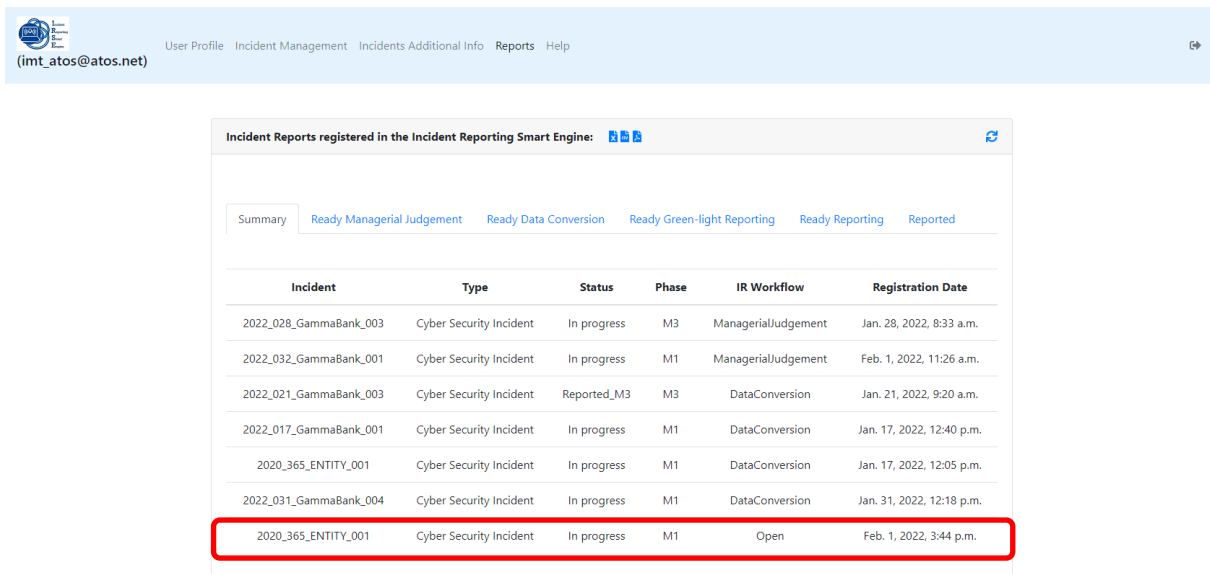


The screenshot shows the 'Tasks' tab for the same case. It displays a list of tasks. One task is visible: 'Data Collection', which is assigned to the 'IMT' group and has a status of 'Started 4 minutes ago'. The task is highlighted with a red box. The interface also shows options to 'Add Task' and 'Quick Filters', and a table with columns for 'Group', 'Task', 'Date', 'Assignee', and 'Actions'.

Group	Task	Date	Assignee	Actions
IMT	Data Collection Started 4 minutes ago	02/01/22 16:45	CS4EU Incident Management Team	[Icons]

Figure 90. Tasks of associated with a case

In the tab *Repos* of the general interface, Figure 91, there will be a new report opened:



Incident Reports registered in the Incident Reporting Smart Engine:

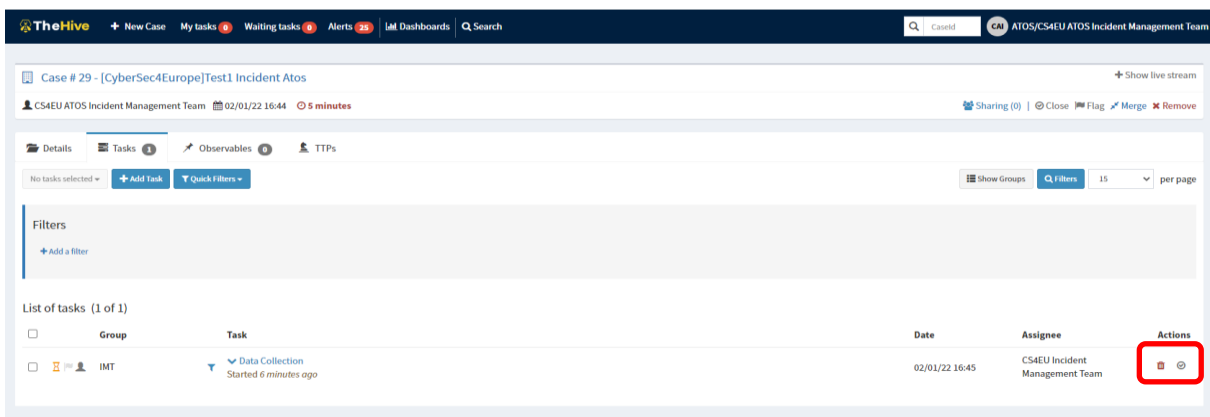
Summary | Ready Managerial Judgement | Ready Data Conversion | Ready Green-light Reporting | Ready Reporting | Reported

Incident	Type	Status	Phase	IR Workflow	Registration Date
2022_028_GammaBank_003	Cyber Security Incident	In progress	M3	ManagerialJudgement	Jan. 28, 2022, 8:33 a.m.
2022_032_GammaBank_001	Cyber Security Incident	In progress	M1	ManagerialJudgement	Feb. 1, 2022, 11:26 a.m.
2022_021_GammaBank_003	Cyber Security Incident	Reported_M3	M3	DataConversion	Jan. 21, 2022, 9:20 a.m.
2022_017_GammaBank_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:40 p.m.
2020_365_ENTITY_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:05 p.m.
2022_031_GammaBank_004	Cyber Security Incident	In progress	M1	DataConversion	Jan. 31, 2022, 12:18 p.m.
2020_365_ENTITY_001	Cyber Security Incident	In progress	M1	Open	Feb. 1, 2022, 3:44 p.m.

Figure 91. List of Reports

NOTE: When *New Case* is registered, since no information has been provided yet, the event ID “2020_365_ENTITY_001” will be assigned by default. Once the information is included through TheHive, it will be reflected also in the dashboard.

2. **Task Actions:** allow to close, resume, or delete the task of a case, Figure 92.



TheHive + New Case My tasks 0 Waiting tasks 0 Alerts 25 IM Dashboards Search CaseID CS4EU/CS4EU ATOS Incident Management Team

Case # 29 - [CyberSec4Europe]Test1 Incident Atos Show live stream

CS4EU ATOS Incident Management Team 02/01/22 16:44 5 minutes Sharing 0 Close Flag Merge Remove

Details Tasks 1 Observables 0 TTPs

No tasks selected + Add Task Quick Filters Show Groups Filters 15 per page

Filters + Add a filter

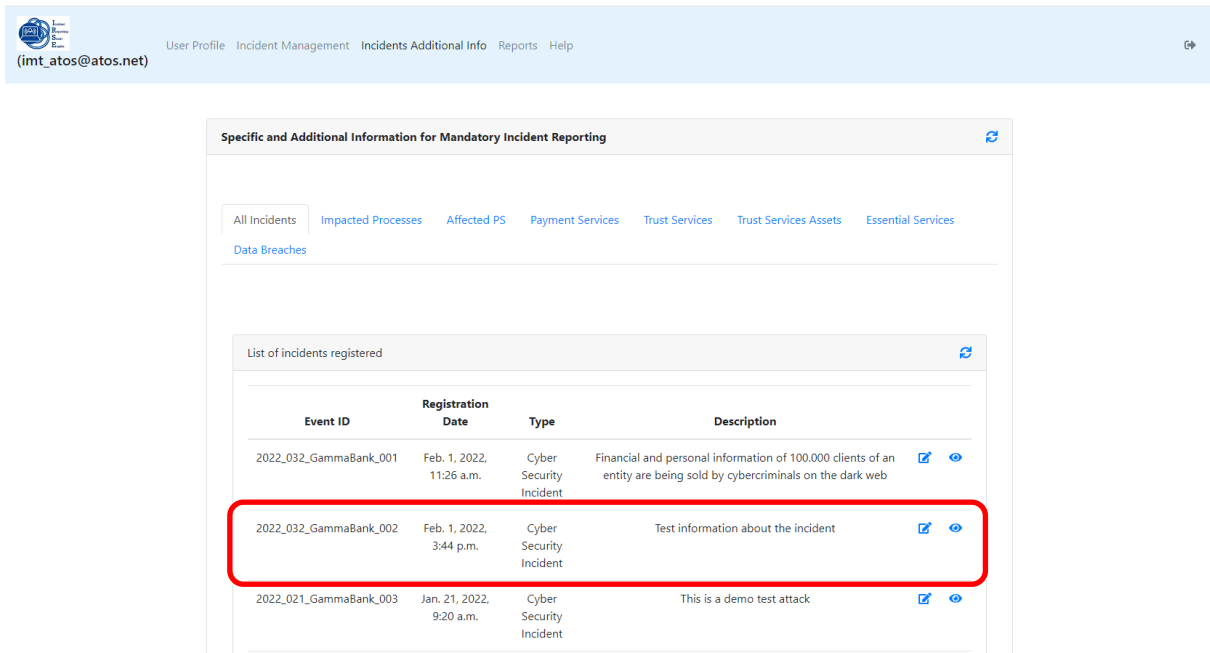
List of tasks (1 of 1)

<input type="checkbox"/>	Group	Task	Date	Assignee	Actions
<input type="checkbox"/>	IMT	Data Collection Started 6 minutes ago	02/01/22 16:45	CS4EU Incident Management Team	<input type="checkbox"/> <input type="checkbox"/>

Figure 92. Task list of a case

When the *Data Collection* task is closed, the associated report changes to *Enrichment*.

3. Incident Additional Info: *All Incidents* tab displays the list of all incidents registered, Figure 93.



Specific and Additional Information for Mandatory Incident Reporting

All Incidents | Impacted Processes | Affected PS | Payment Services | Trust Services | Trust Services Assets | Essential Services

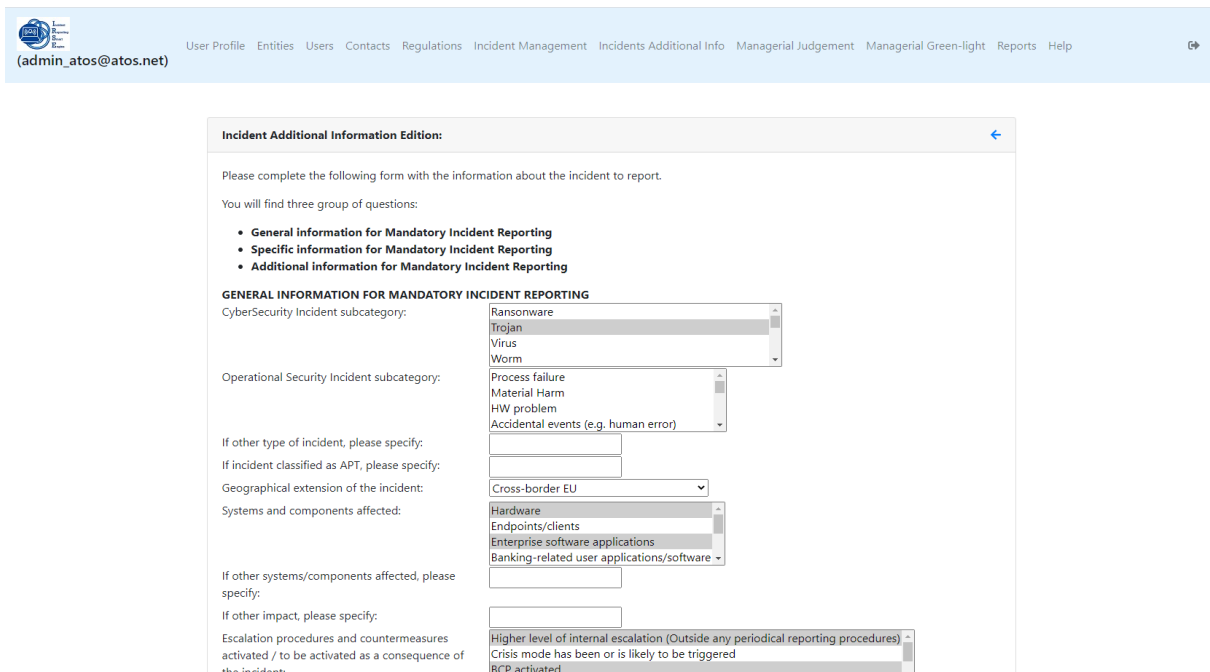
Data Breaches

List of incidents registered

Event ID	Registration Date	Type	Description
2022_032_GammaBank_001	Feb. 1, 2022, 11:26 a.m.	Cyber Security Incident	Financial and personal information of 100.000 clients of an entity are being sold by cybercriminals on the dark web
2022_032_GammaBank_002	Feb. 1, 2022, 3:44 p.m.	Cyber Security Incident	Test information about the incident
2022_021_GammaBank_003	Jan. 21, 2022, 9:20 a.m.	Cyber Security Incident	This is a demo test attack

Figure 93. List of incidents registered

In addition, a user can enrich the data editing the information about a registered incident, Figure 94.



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: Ransomware, 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: [Text input]

If incident classified as APT, please specify: [Text input]

Geographical extension of the incident: Cross-border EU

Systems and components affected: Hardware, Endpoints/clients, Enterprise software applications, Banking-related user applications/software

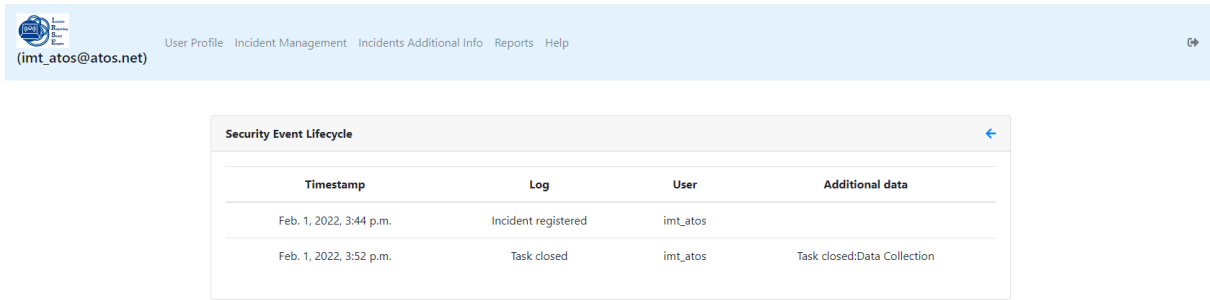
If other systems/components affected, please specify: [Text input]

If other impact, please specify: [Text input]

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, RPD activated

Figure 94. Incident Additional Information Edition

The eye icon, Figure 93, shows the logs registered related to the security event lifecycle, Figure 95 Figure 94:

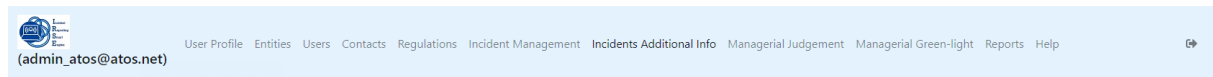
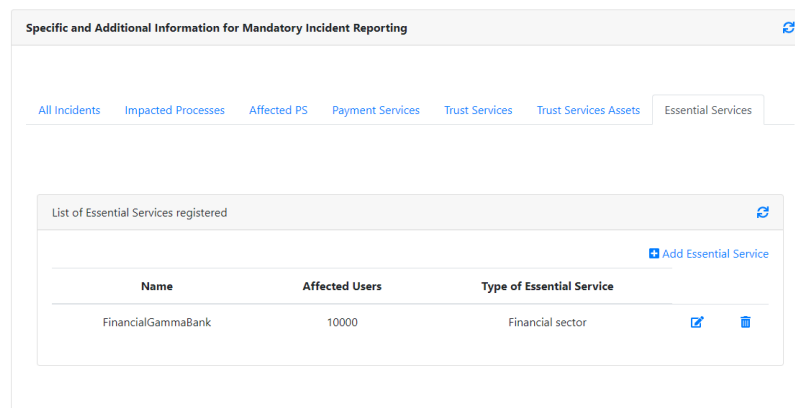


Timestamp	Log	User	Additional data
Feb. 1, 2022, 3:44 p.m.	Incident registered	imt_atos	
Feb. 1, 2022, 3:52 p.m.	Task closed	imt_atos	Task closed:Data Collection

Figure 95. Security Event Lifecycle

The other tabs in the Figure 93 allow to manage the elements that are affected by an incident, such as services, assets, processes, or data. These elements are:

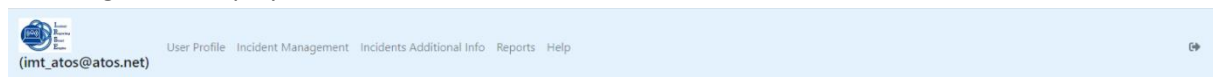
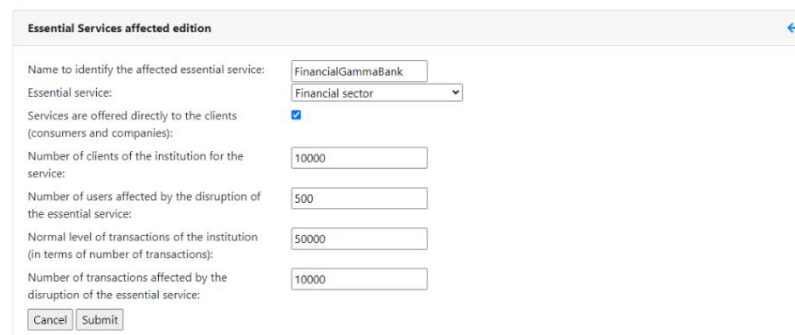
- **Essential Services:** In case the institution is a provider of essential services, they will be defined in this menu. A name needs to be assigned so it can be assigned to the incident.

Name	Affected Users	Type of Essential Service
FinancialGammaBank	10000	Financial sector

Figure 96. List of Essential Services

► Figure 97 displays the edition/creation form for this sub-menu.

Essential Services affected edition

Name to identify the affected essential service:

Essential service:

Services are offered directly to the clients (consumers and companies):

Number of clients of the institution for the service:

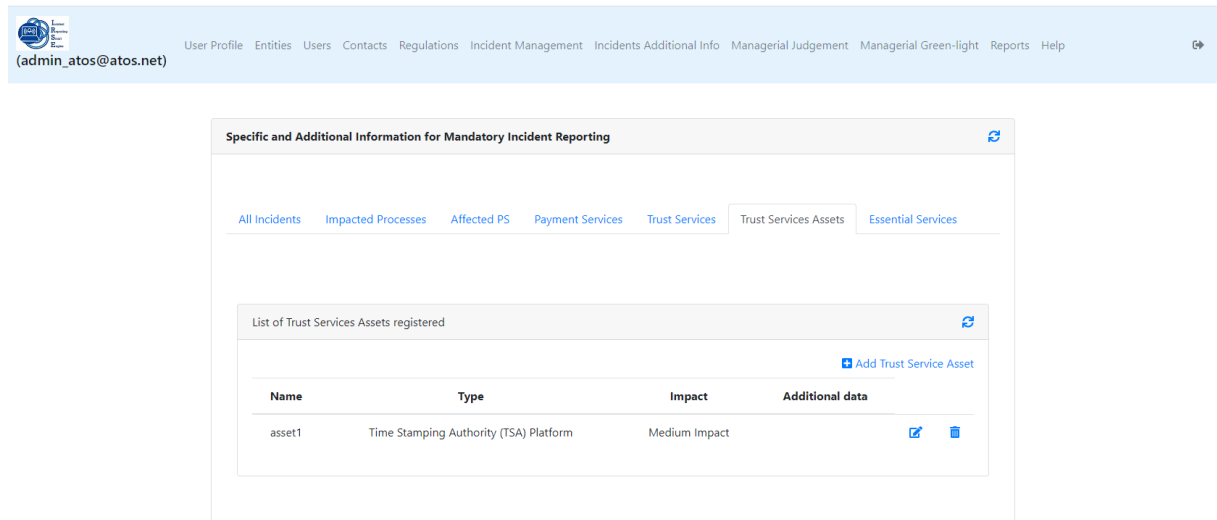
Number of users affected by the disruption of the essential service:

Normal level of transactions of the institution (in terms of number of transactions):

Number of transactions affected by the disruption of the essential service:

Figure 97. Essential Services edition

- Trust Services Assets: External services that support the trust protocols for Critical Infrastructures. Figure 98 displays the list of Trust Services Assets registered.



Specific and Additional Information for Mandatory Incident Reporting

admin_atos@atos.net | User Profile | Entities | Users | Contacts | Regulations | Incident Management | Incidents Additional Info | Managerial Judgement | Managerial Green-light | Reports | Help

All Incidents | Impacted Processes | Affected PS | Payment Services | Trust Services | **Trust Services Assets** | Essential Services

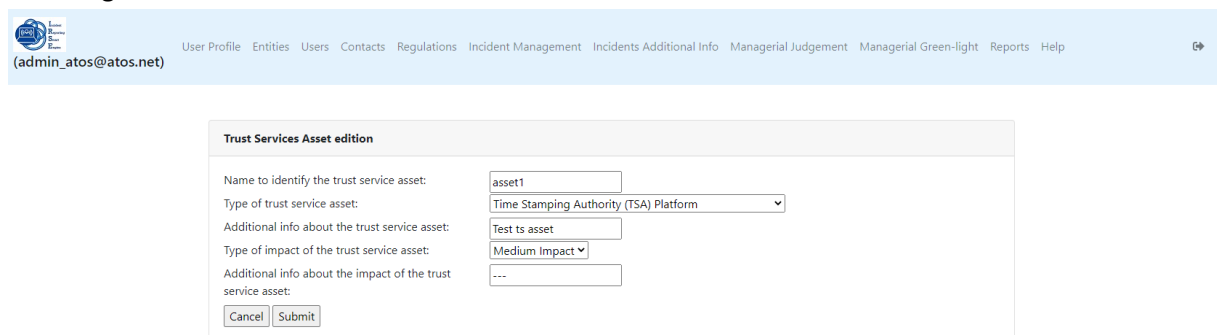
List of Trust Services Assets registered

[Add Trust Service Asset](#)

Name	Type	Impact	Additional data
asset1	Time Stamping Authority (TSA) Platform	Medium Impact	Edit Delete

Figure 98. List of Trust Services Assets registered

- Figure 99 shows the edition form of one Trust Service Asset.



Trust Services Asset edition

Name to identify the trust service asset:

Type of trust service asset:

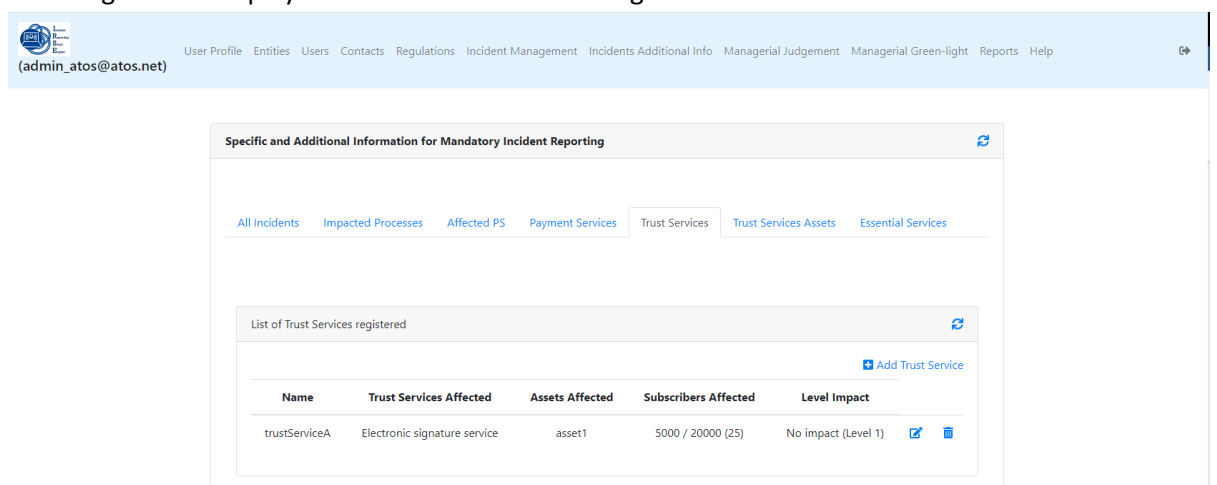
Additional info about the trust service asset:

Type of impact of the trust service asset:

Additional info about the impact of the trust service asset:

Figure 99. Trust Services Asset edition

- Trust Services: Internal services that support the trust protocols for Critical Infrastructures. Figure 100 displays the list of Trust Services registered.



Specific and Additional Information for Mandatory Incident Reporting

admin_atos@atos.net | User Profile | Entities | Users | Contacts | Regulations | Incident Management | Incidents Additional Info | Managerial Judgement | Managerial Green-light | Reports | Help

All Incidents | Impacted Processes | Affected PS | Payment Services | Trust Services | **Trust Services Assets** | Essential Services

List of Trust Services registered

[Add Trust Service](#)

Name	Trust Services Affected	Assets Affected	Subscribers Affected	Level Impact
trustServiceA	Electronic signature service	asset1	5000 / 20000 (25)	No impact (Level 1)

Figure 100. List of Trust Services registered

Figure 101 shows the edition form of one Trust Service Asset.

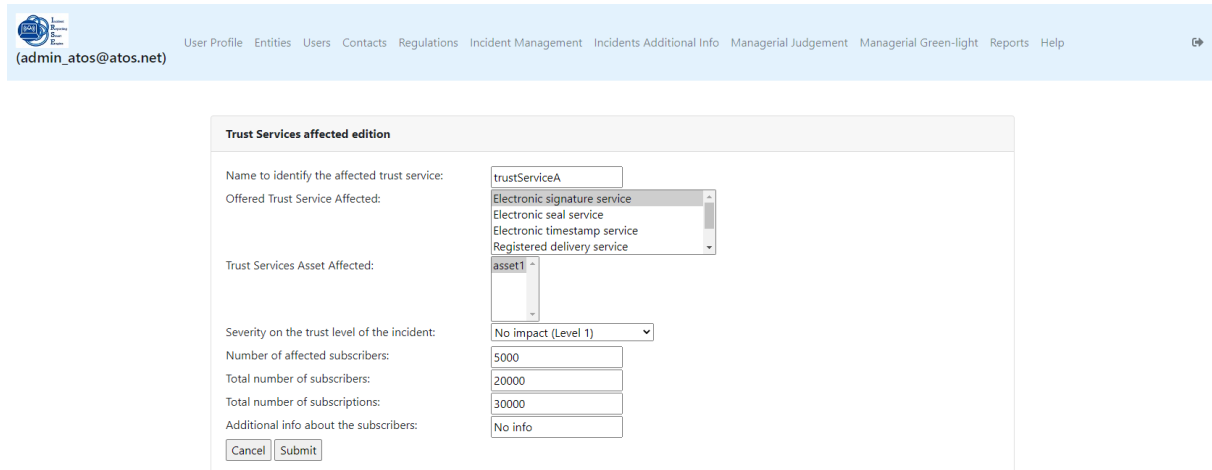
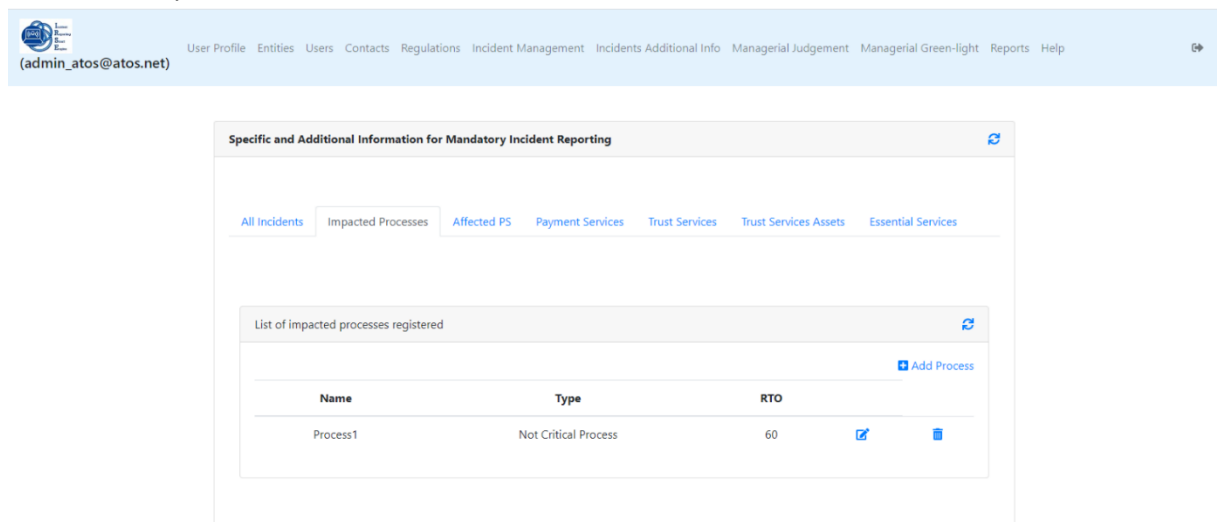


Figure 101. Trust Services affected edition

- **Impacted Processes:** Shows the list of processes impacted by an attack, the impact, and the recovery time.



Name	Type	RTO		
Process1	Not Critical Process	60		

Figure 102. List of Impacted Processes

Figure 103 shows the form to add or create *Processes Affected*.

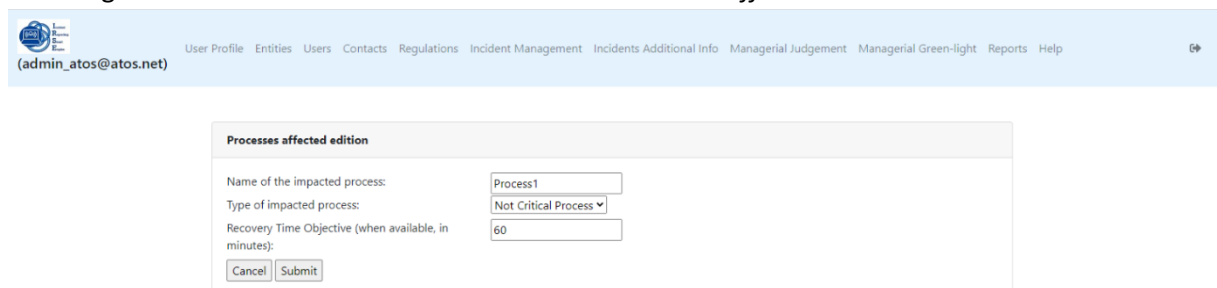
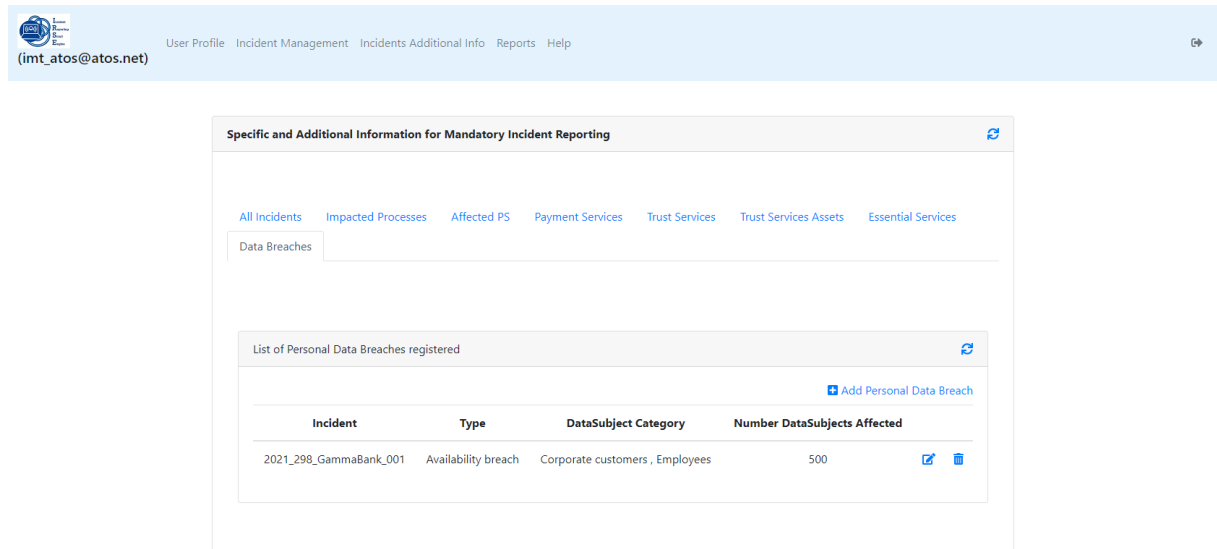


Figure 103. Processes Affected edition

- **Data Breaches:** lists the *Personal Data Breaches* with the type, data category and the number of affected subjects, Figure 104.



Specific and Additional Information for Mandatory Incident Reporting

All Incidents Impacted Processes Affected PS Payment Services Trust Services Trust Services Assets Essential Services

Data Breaches

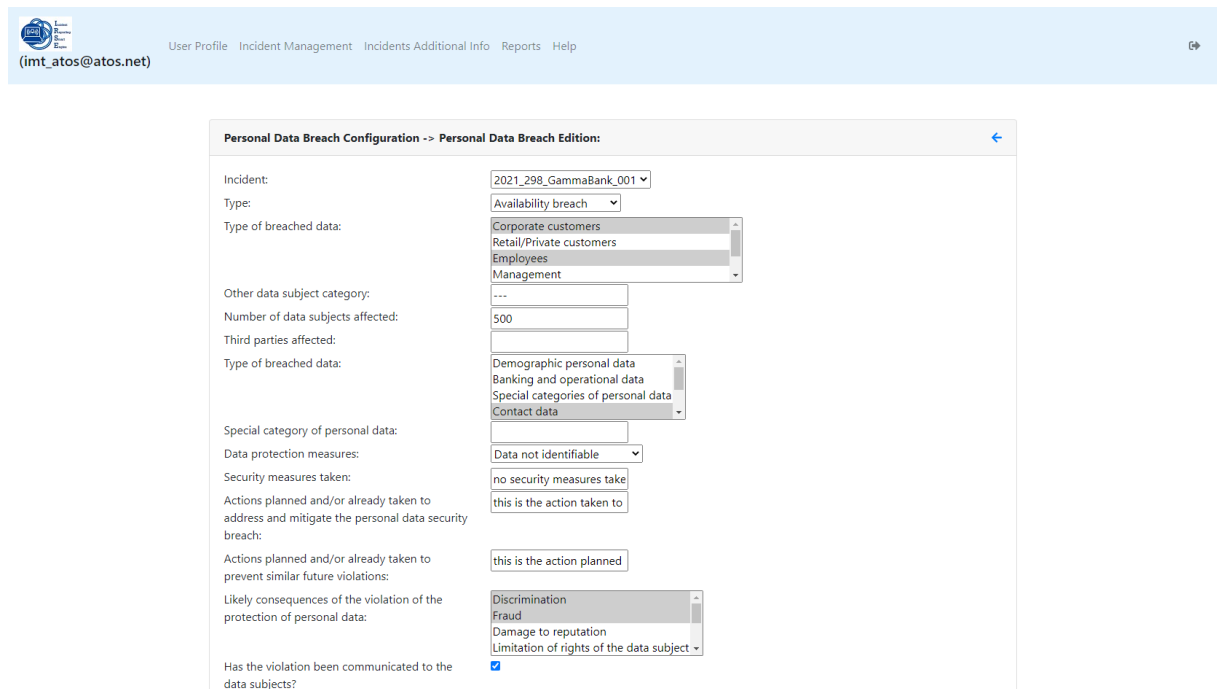
List of Personal Data Breaches registered

[Add Personal Data Breach](#)

Incident	Type	DataSubject Category	Number DataSubjects Affected
2021_298_GammaBank_001	Availability breach	Corporate customers, Employees	500

Figure 104. List of Personal Data Breaches registered

Figure 105 displays the form to add or edit the information of a *Personal Data Breach*.



Personal Data Breach Configuration -> Personal Data Breach Edition:

Incident: 2021_298_GammaBank_001

Type: Availability breach

Type of breached data: Corporate customers, Retail/Private customers, Employees, Management

Other data subject category: ---

Number of data subjects affected: 500

Third parties affected: 500

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 not identifiable

Security measures taken: no security measures take

Actions planned and/or already taken to address and mitigate the personal data security breach: this is the action taken to

Actions planned and/or already taken to prevent similar future violations: this is the action planned

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 105. Personal Data Breach Edition

- ▶ **NOTE:** The association of a data breach with an incident is done by making a selection from the list of active incidents in that menu.

4. **Add observables:** Include information about the incident and run analyzers on them, Figure 106.

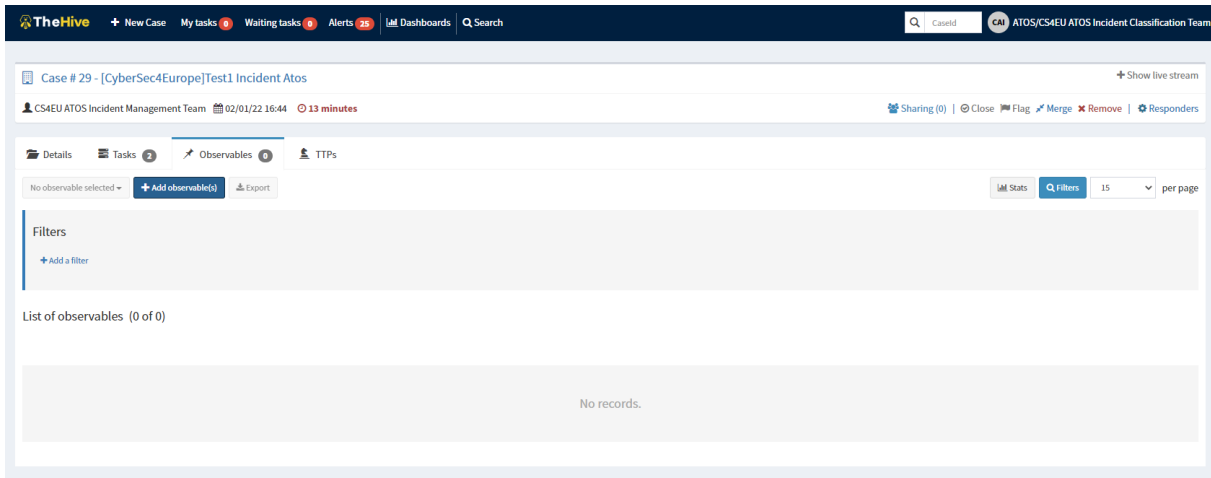


Figure 106. List of observables

The button *Add observable* open a form to create new observables elements, Figure 107.

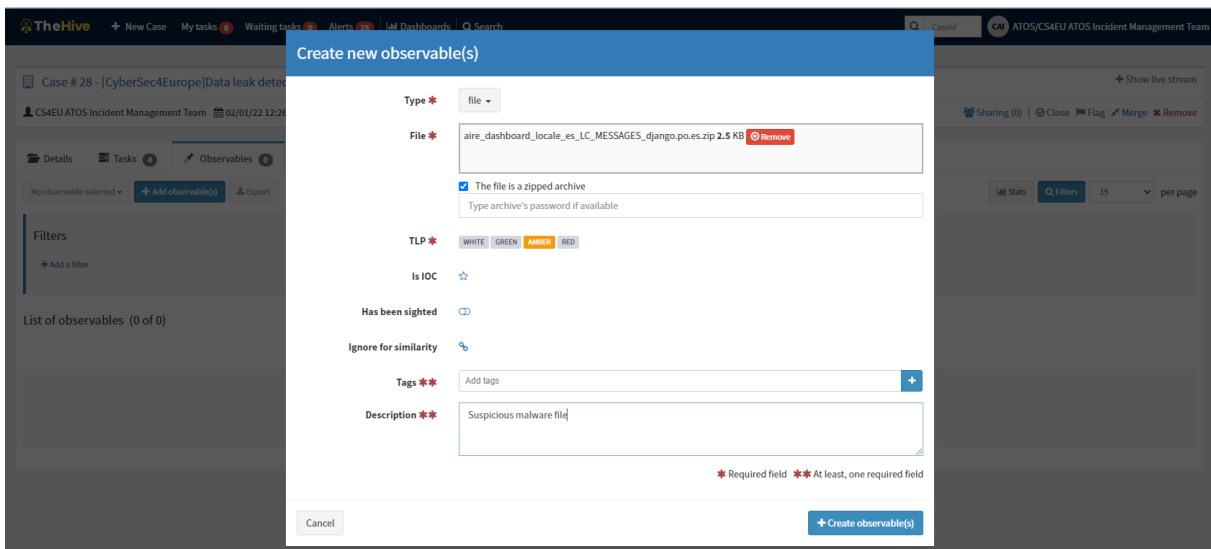


Figure 107. Create new observable form

After selecting some observables, *Run analyzers* allows us to run different analyzers, Figure 108.

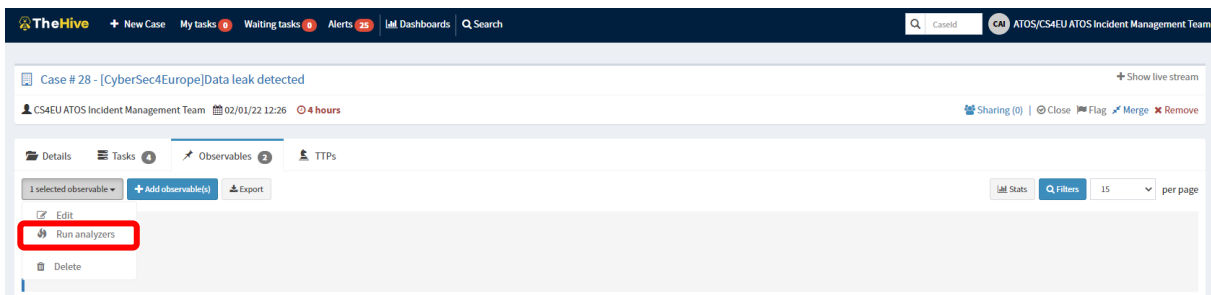


Figure 108. Run analyzers option

Multiple analyzers can be selected to run together, Figure 109.

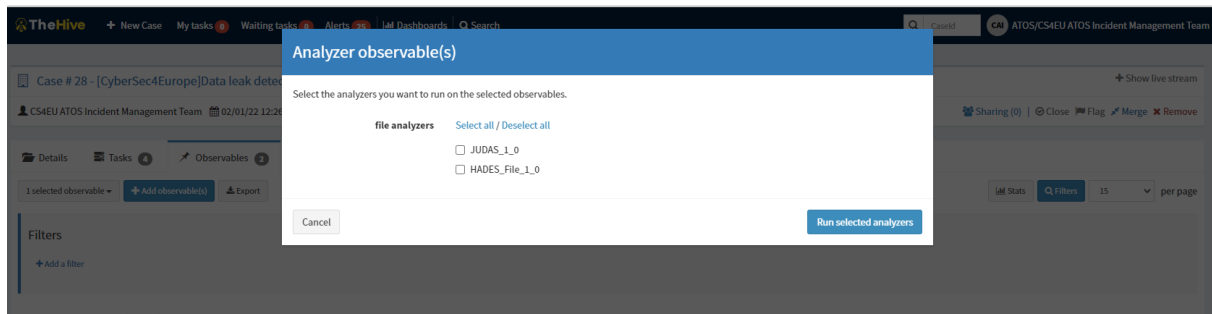


Figure 109. Selection of the analyzers to run

- Incident Classification Task:** is automatically created when the *Data Enrichment* Task is completed, Figure 110.

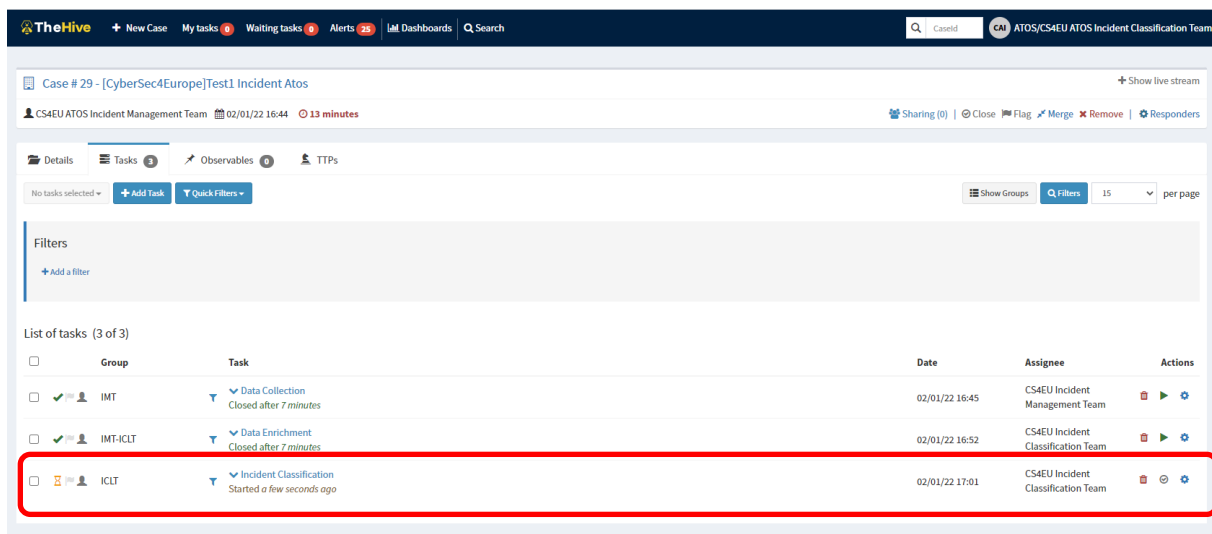
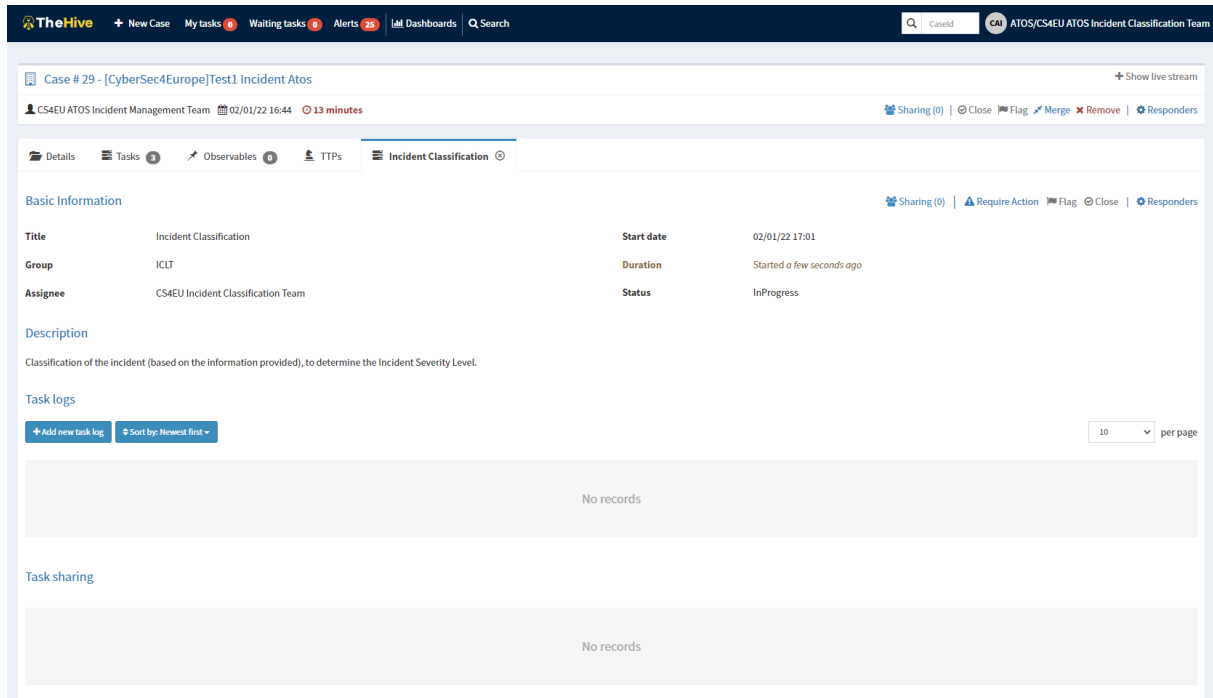


Figure 110. Automated creation of Incident Classification task

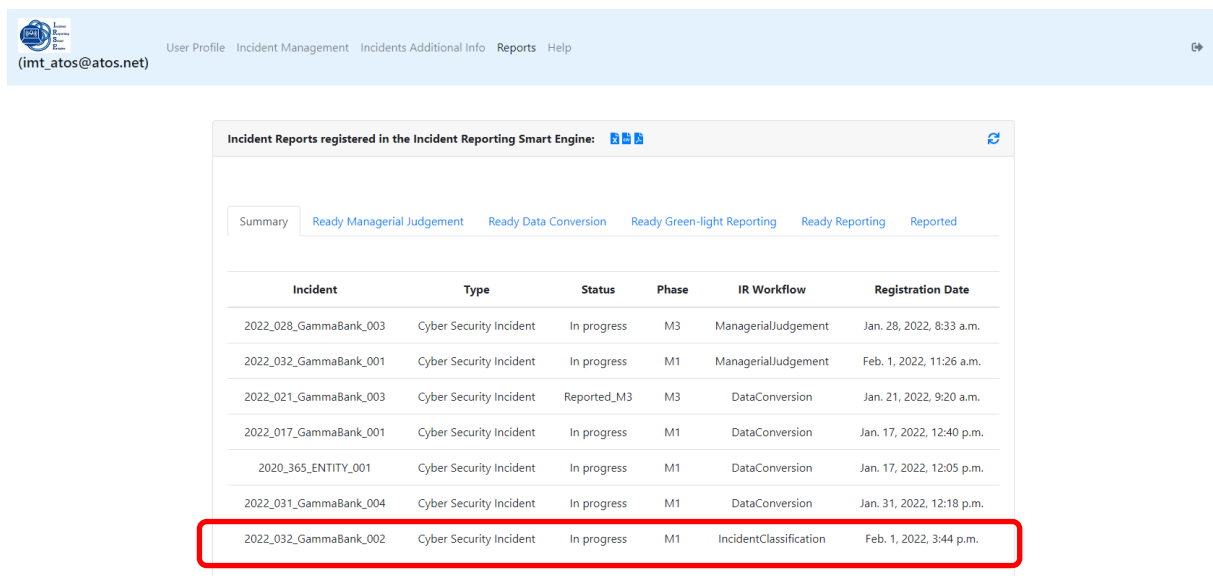
Each task has its own view where details and actions are showed, Figure 111.



The screenshot shows the 'Incident Classification' task details in TheHive. The task title is 'Incident Classification', started on 02/01/22 at 17:01, and is currently 'InProgress'. The assignee is 'CS4EU Incident Classification Team'. The description states: 'Classification of the incident (based on the information provided), to determine the Incident Severity Level.' There are no task logs or task sharing records displayed.

Figure 111. Detail of a task

The report passes to the next IR workflow, Figure 112.



The screenshot shows a list of incident reports registered in the Incident Reporting Smart Engine. The table below lists the reports, with the entry for '2022_032_GammaBank_002' highlighted in red.

Incident	Type	Status	Phase	IR Workflow	Registration Date
2022_028_GammaBank_003	Cyber Security Incident	In progress	M3	ManagerialJudgement	Jan. 28, 2022, 8:33 a.m.
2022_032_GammaBank_001	Cyber Security Incident	In progress	M1	ManagerialJudgement	Feb. 1, 2022, 11:26 a.m.
2022_021_GammaBank_003	Cyber Security Incident	Reported_M3	M3	DataConversion	Jan. 21, 2022, 9:20 a.m.
2022_017_GammaBank_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:40 p.m.
2020_365_ENTITY_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:05 p.m.
2022_031_GammaBank_004	Cyber Security Incident	In progress	M1	DataConversion	Jan. 31, 2022, 12:18 p.m.
2022_032_GammaBank_002	Cyber Security Incident	In progress	M1	IncidentClassification	Feb. 1, 2022, 3:44 p.m.

Figure 112. List of Incident Reports registered

6. **Event Classification Task:** Check the information to do the event classification has been introduced in the template and invoke the *Responder* (located in the upper right corner in TheHive GUI from the page with the Case details), Figure 113.

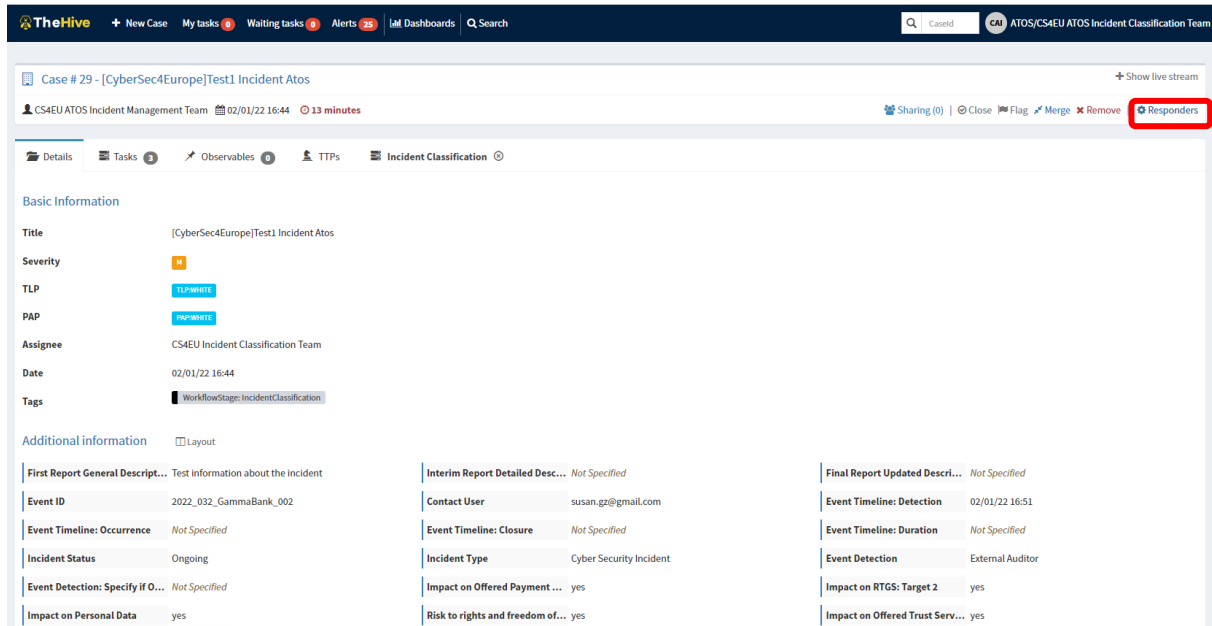


Figure 113. Responders in an Incident Detail View

The Responder is integrated with AIRE asset, Figure 114, so the function of the user who executes it will be checked, and it only will get a result in case it belongs to ICLT (according to the workflow configuration).

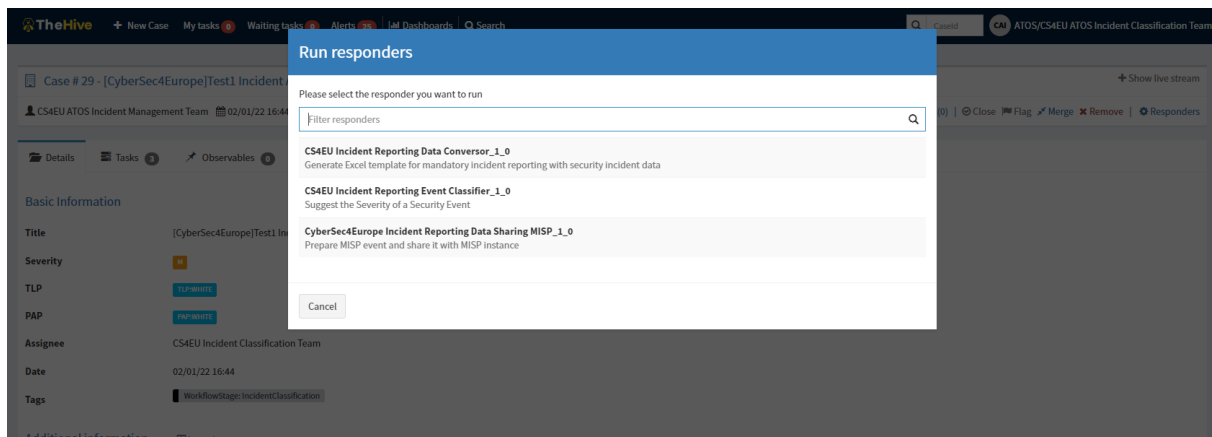
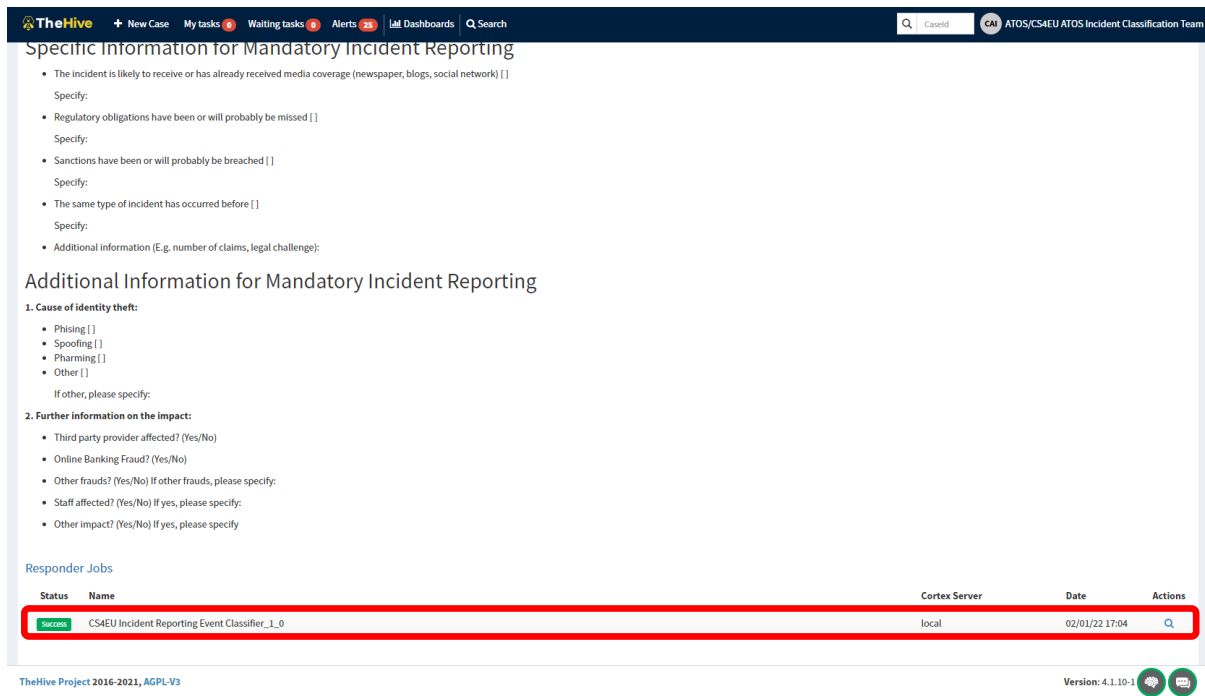


Figure 114. Selection of one responder to run for an incident case

The result of the classification is available at the end of the incident page, Figure 115.



The screenshot shows the TheHive interface for an incident page. The main content area is titled 'Specific Information for Mandatory Incident Reporting' and contains several sections of information, including 'Additional Information for Mandatory Incident Reporting' with sub-sections for 'Cause of identity theft' and 'Further information on the impact'. Below this is a 'Responder Jobs' table. The table has columns for 'Status', 'Name', 'Cortex Server', 'Date', and 'Actions'. One row is highlighted with a red border, showing a 'Success' status for the job 'CS4EU Incident Reporting Event Classifier_1_0' on a 'local' server, completed on '02/01/22 17:04'.

Status	Name	Cortex Server	Date	Actions
Success	CS4EU Incident Reporting Event Classifier_1_0	local	02/01/22 17:04	Q

Figure 115. List of Responder Jobs with status

The action button shows the output, Figure 116.



The screenshot shows a modal window titled 'Report of CS4EU Incident Reporting Event Classifier_1_0 responder'. The window displays a JSON object representing the output of the responder. The JSON includes fields for 'Incident Impact Severity', 'message', and various submission statuses for different systems like 'ICEB-SM', 'PSD2', 'MIS', 'EIDAS', 'TARNO T2', and 'GDPR'.

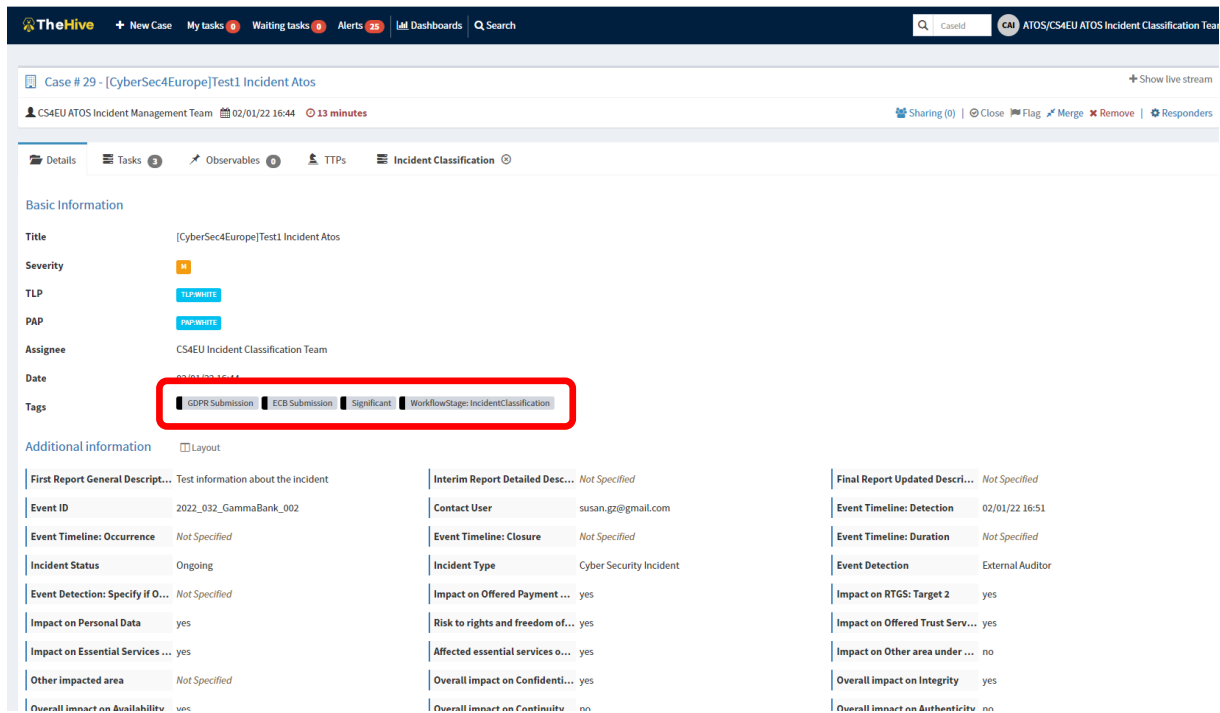
```

{
  "Incident Impact Severity": "Significant",
  "message": "Success execution of IR Event Classifier API",
  "ICEB-SM submission": true,
  "PSD2 submission": false,
  "MIS submission": false,
  "EIDAS submission": false,
  "TARNO T2 submission": false,
  "GDPR submission": true
}

```

Figure 116. Output example for after run a responder

This information will be also automatically updated in the tags of the case, Figure 117, and in the fields of the template, so the ICLT user can check the suggestion and modify it if he/she considers it.



TheHive + New Case My tasks 0 Waiting tasks 0 Alerts 25 Dashboards Search

Case # 29 - [CyberSec4Europe]Test1 Incident Atos

CS4EU ATOS Incident Management Team 02/01/22 16:44 13 minutes

Sharing (0) Close Flag Merge Remove Responders

Details Tasks Observables TTPs Incident Classification

Basic Information

Title [CyberSec4Europe]Test1 Incident Atos

Severity **M**

TLP **TLP:WHITE**

PAP **PAP:WHITE**

Assignee CS4EU Incident Classification Team

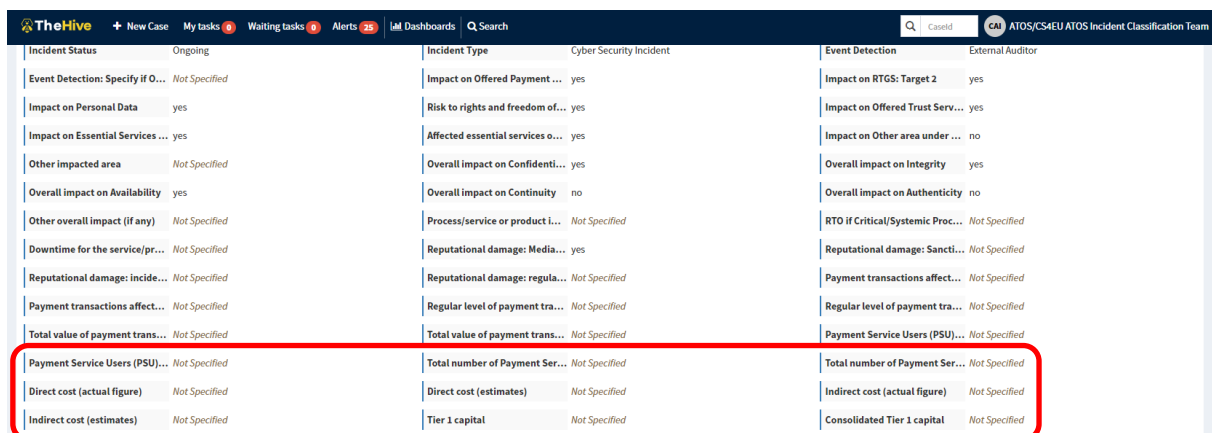
Date 02/01/22 16:44

Tags **GDPR Submission** **ECB Submission** **Significant** **WorkflowStage: IncidentClassification**

Additional information Layout

First Report General Descript...	Test information about the incident	Interim Report Detailed Desc...	Not Specified	Final Report Updated Descri...	Not Specified
Event ID	2022_032_GammaBank_002	Contact User	susan.gz@gmail.com	Event Timeline: Detection	02/01/22 16:51
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	External Auditor
Event Detection: Specify if O...	Not Specified	Impact on Offered Payment ...	yes	Impact on RTGS: Target 2	yes
Impact on Personal Data	yes	Risk to rights and freedom of...	yes	Impact on Offered Trust Serv...	yes
Impact on Essential Services ...	yes	Affected essential services o...	yes	Impact on Other area under ...	no
Other impacted area	Not Specified	Overall impact on Confident...	yes	Overall impact on Integrity	yes
Overall impact on Availability	yes	Overall impact on Continuity	no	Overall impact on Authenticity	no

Figure 117. Tags of an Incident Case



TheHive + New Case My tasks 0 Waiting tasks 0 Alerts 25 Dashboards Search

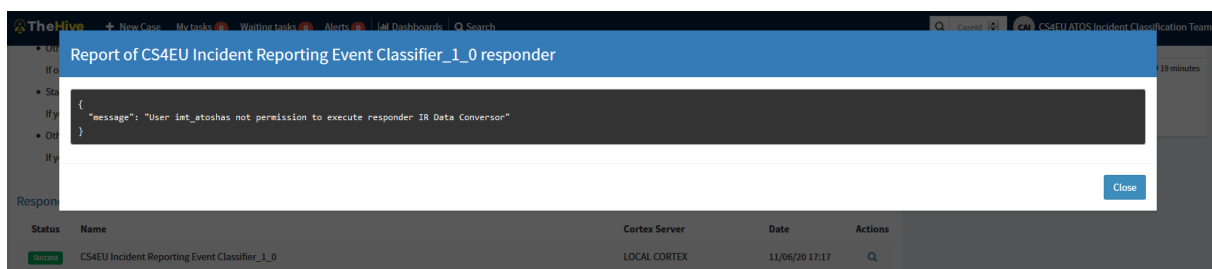
Case # 29 - [CyberSec4Europe]Test1 Incident Atos

CS4EU ATOS Incident Classification Team

Incident Status	Ongoing	Incident Type	Cyber Security Incident	Event Detection	External Auditor
Event Detection: Specify if O...	Not Specified	Impact on Offered Payment ...	yes	Impact on RTGS: Target 2	yes
Impact on Personal Data	yes	Risk to rights and freedom of...	yes	Impact on Offered Trust Serv...	yes
Impact on Essential Services ...	yes	Affected essential services o...	yes	Impact on Other area under ...	no
Other impacted area	Not Specified	Overall impact on Confident...	yes	Overall impact on Integrity	yes
Overall impact on Availability	yes	Overall impact on Continuity	no	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...	yes	Reputational damage: Sancti...	Not Specified
Reputational damage: incide...	Not Specified	Reputational damage: regula...	Not Specified	Payment transactions affect...	Not Specified
Payment transactions affect...	Not Specified	Regular level of payment tra...	Not Specified	Regular level of payment tra...	Not Specified
Total value of payment trans...	Not Specified	Total value of payment trans...	Not Specified	Payment Service Users (PSU)...	Not Specified
Payment Service Users (PSU)...	Not Specified	Total number of Payment Ser...	Not Specified	Total number of Payment Ser...	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

Figure 118. Fields of a template

The execution of the responder will invoke AIRE asset to determine if the user has permissions to execute for this phase of the workflow. If not, it will be shown something like Figure 119.



TheHive + New Case My tasks 0 Waiting tasks 0 Alerts 25 Dashboards Search

Case # 29 - [CyberSec4Europe]Test1 Incident Atos

CS4EU ATOS Incident Classification Team

Report of CS4EU Incident Reporting Event Classifier_1_0 responder

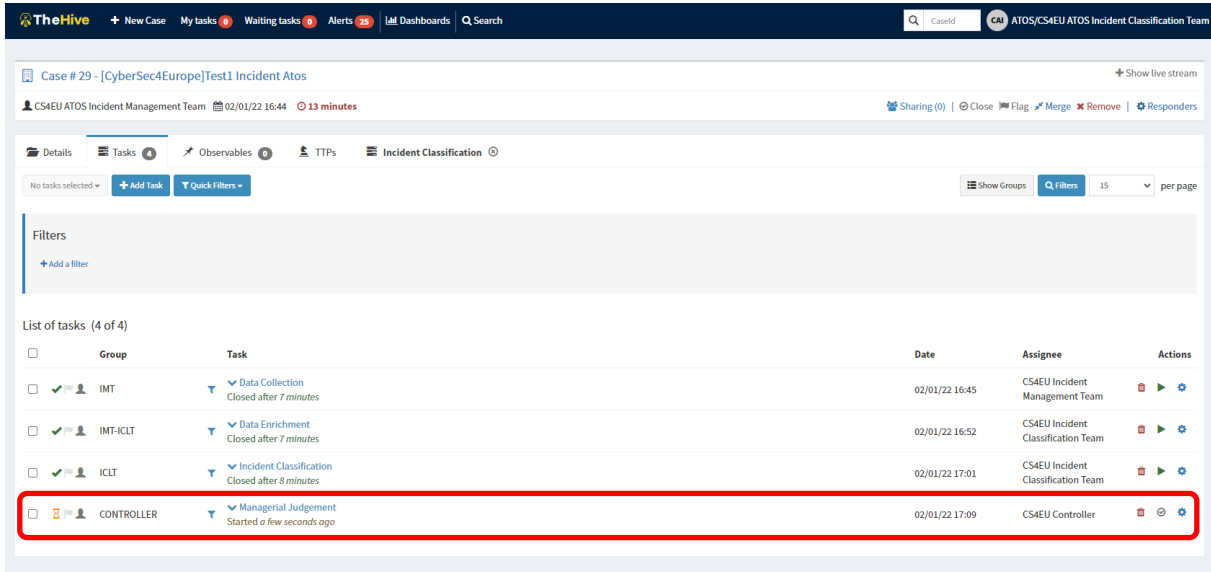
```
{
  "message": "User int_atoshas not permission to execute responder IR Data Converter"
}
```

Close

Status	Name	Cortex Server	Date	Actions
Success	CS4EU Incident Reporting Event Classifier_1_0	LOCAL CORTEX	11/06/20 17:17	Q

Figure 119. Output of a responder without permission

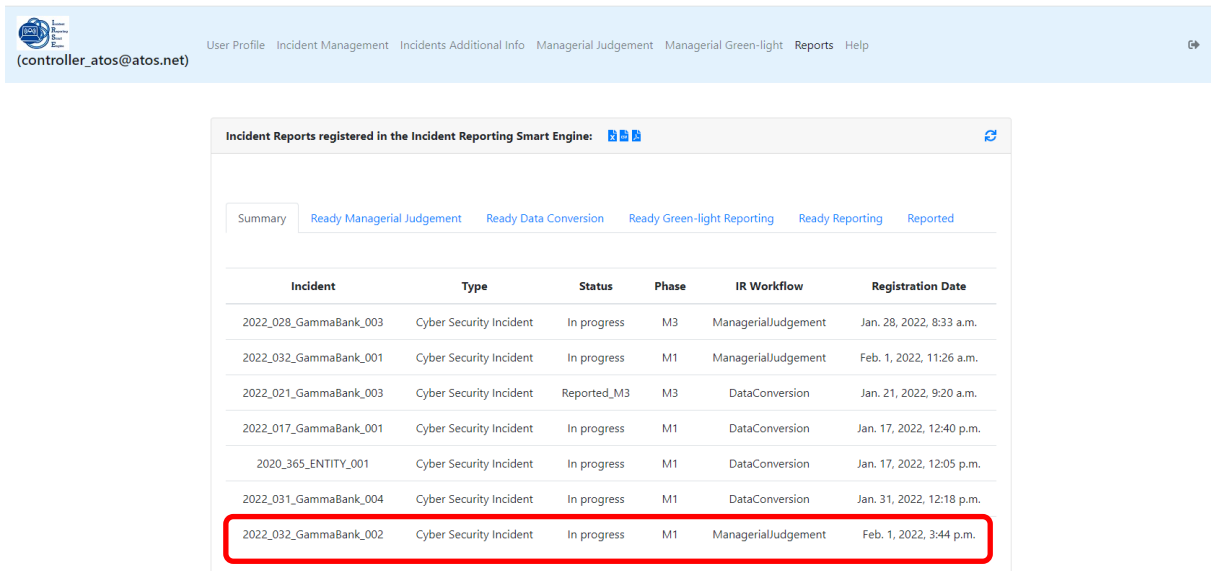
7. **Managerial Judgement Task:** is automatically created and assigned to the Controller when the *Incident Classification* task is closed, Figure 120 and Figure 121.



The screenshot shows the TheHive interface for Case # 29 - [CyberSec4Europe]Test1 Incident Atos. The 'Incident Classification' tab is active, displaying a list of tasks. The 'Managerial Judgement' task is highlighted with a red box. The task details are as follows:

Group	Task	Date	Assignee	Actions
IMT	Data Collection Closed after 7 minutes	02/01/22 16:45	CS4EU Incident Management Team	[Icons]
IMT-ICLT	Data Enrichment Closed after 7 minutes	02/01/22 16:52	CS4EU Incident Classification Team	[Icons]
ICLT	Incident Classification Closed after 8 minutes	02/01/22 17:01	CS4EU Incident Classification Team	[Icons]
CONTROLLER	Managerial Judgement Started a few seconds ago	02/01/22 17:09	CS4EU Controller	[Icons]

Figure 120. Automatic creation of Managerial Judgement Task

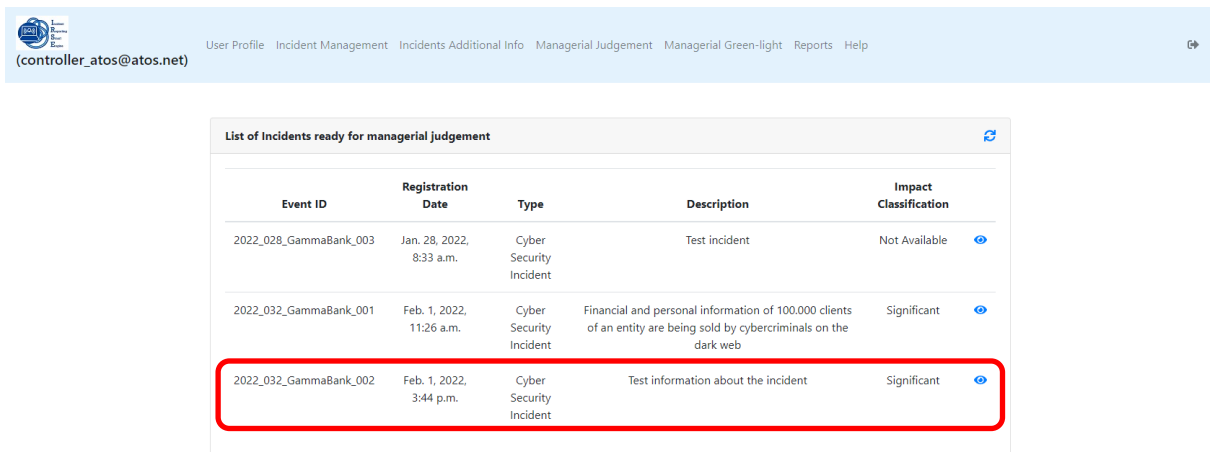


The screenshot shows the Incident Reporting Smart Engine interface. The 'Ready Managerial Judgement' tab is selected, displaying a table of incident reports. The report for incident 2022_032_GammaBank_002 is highlighted with a red box. The table data is as follows:

Incident	Type	Status	Phase	IR Workflow	Registration Date
2022_028_GammaBank_003	Cyber Security Incident	In progress	M3	ManagerialJudgement	Jan. 28, 2022, 8:33 a.m.
2022_032_GammaBank_001	Cyber Security Incident	In progress	M1	ManagerialJudgement	Feb. 1, 2022, 11:26 a.m.
2022_021_GammaBank_003	Cyber Security Incident	Reported_M3	M3	DataConversion	Jan. 21, 2022, 9:20 a.m.
2022_017_GammaBank_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:40 p.m.
2020_365_ENTITY_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:05 p.m.
2022_031_GammaBank_004	Cyber Security Incident	In progress	M1	DataConversion	Jan. 31, 2022, 12:18 p.m.
2022_032_GammaBank_002	Cyber Security Incident	In progress	M1	ManagerialJudgement	Feb. 1, 2022, 3:44 p.m.

Figure 121. Automatic creation of Managerial Judgement IR Workflow

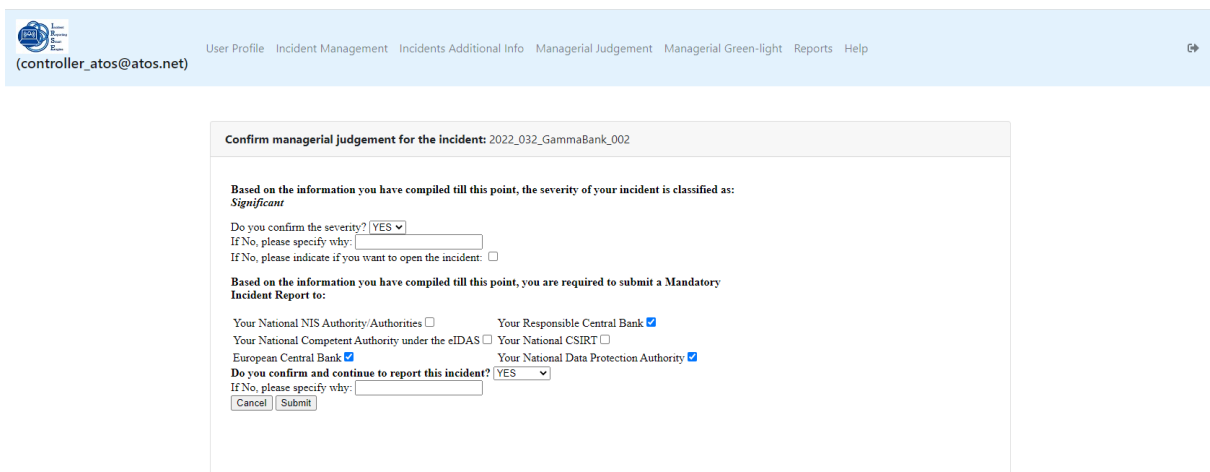
8. Under the menu *Managerial Judgement*, the controller will see the report with the impact classification, Figure 122.



Event ID	Registration Date	Type	Description	Impact Classification
2022_028_GammaBank_003	Jan. 28, 2022, 8:33 a.m.	Cyber Security Incident	Test incident	Not Available
2022_032_GammaBank_001	Feb. 1, 2022, 11:26 a.m.	Cyber Security Incident	Financial and personal information of 100.000 clients of an entity are being sold by cybercriminals on the dark web	Significant
2022_032_GammaBank_002	Feb. 1, 2022, 3:44 p.m.	Cyber Security Incident	Test information about the incident	Significant

Figure 122. 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, Figure 123.



Confirm managerial judgement for the incident: 2022_032_GammaBank_002

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

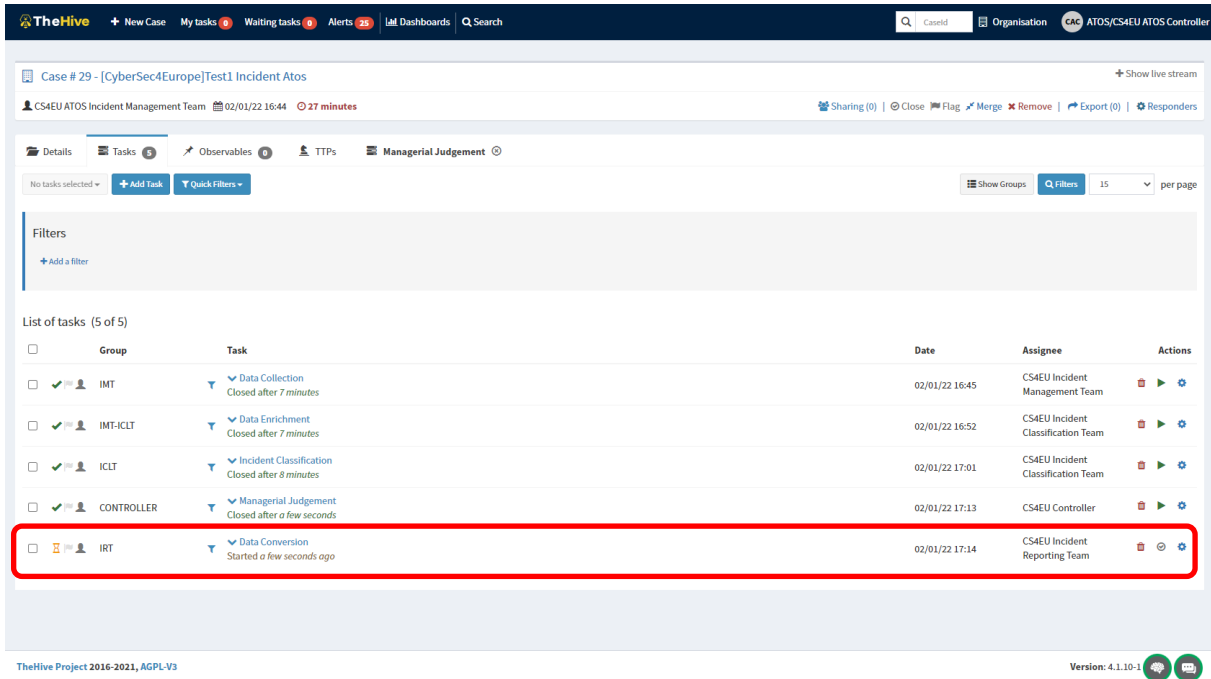
Do you confirm and continue to report this incident? [YES ▼]
 If No, please specify why: _____

[Cancel] [Submit]

Figure 123. Form for a managerial judgement

NOTE: In case the Controller does not confirm the classification and select to open the incident, the incident reporting process will come back to the “Data Enrichment” phase. In this case, the controller does not confirm to proceed with the reporting, the incident reporting process will finish, and the reports will be closed.

Once the managerial judgement is done, go to Incident Management tab to close the task. Automatically, the task called “Managerial Judgement” will be closed by AIRE asset 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 assigned to the Incident Reporting Team and the report will be ready for report preparation, Figure 124.

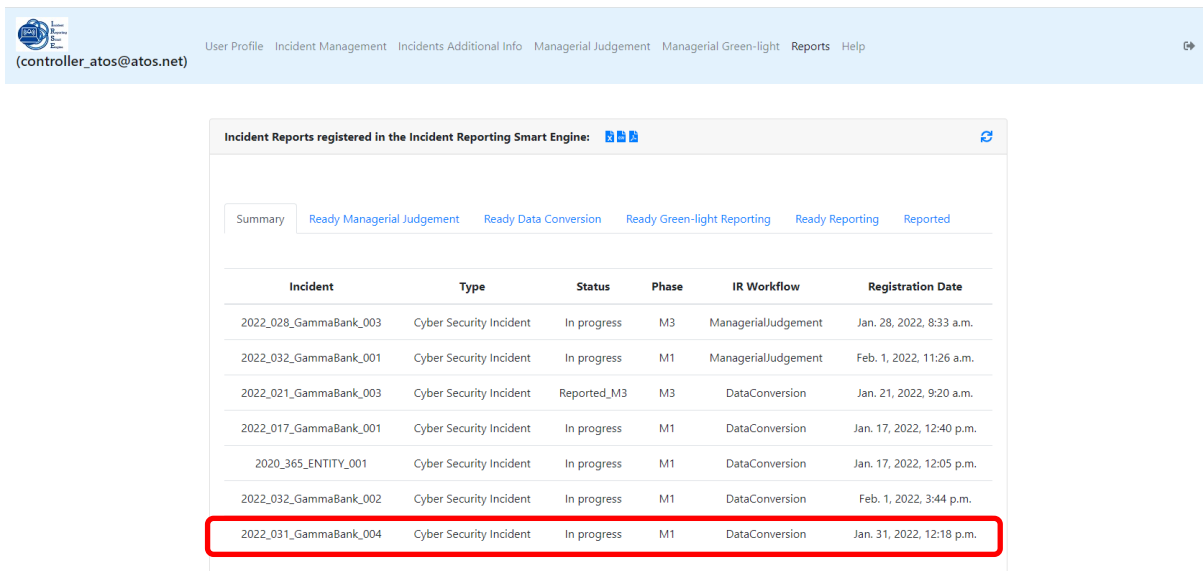


The screenshot shows the 'Tasks' tab in TheHive for Case # 29. The 'Data Conversion' task is highlighted with a red box. The task details are as follows:

Group	Task	Date	Assignee	Actions
IMT	Data Collection Closed after 7 minutes	02/01/22 16:45	CS4EU Incident Management Team	[Close] [Flag] [Merge] [Remove] [Export] [Responders]
IMT-ICLT	Data Enrichment Closed after 7 minutes	02/01/22 16:52	CS4EU Incident Classification Team	[Close] [Flag] [Merge] [Remove] [Export] [Responders]
ICLT	Incident Classification Closed after 8 minutes	02/01/22 17:01	CS4EU Incident Classification Team	[Close] [Flag] [Merge] [Remove] [Export] [Responders]
CONTROLLER	Managerial Judgement Closed after a few seconds	02/01/22 17:13	CS4EU Controller	[Close] [Flag] [Merge] [Remove] [Export] [Responders]
IRT	Data Conversion Started a few seconds ago	02/01/22 17:14	CS4EU Incident Reporting Team	[Close] [Flag] [Merge] [Remove] [Export] [Responders]

Figure 124. New Data Conversion Task

This event is registered in the *Summary* of the *Reports* tab, Figure 125.



The screenshot shows the 'Reports' tab in TheHive. The 'Summary' sub-tab is active. The following table shows the list of incident reports:

Incident	Type	Status	Phase	IR Workflow	Registration Date
2022_028_GammaBank_003	Cyber Security Incident	In progress	M3	ManagerialJudgement	Jan. 28, 2022, 8:33 a.m.
2022_032_GammaBank_001	Cyber Security Incident	In progress	M1	ManagerialJudgement	Feb. 1, 2022, 11:26 a.m.
2022_021_GammaBank_003	Cyber Security Incident	Reported_M3	M3	DataConversion	Jan. 21, 2022, 9:20 a.m.
2022_017_GammaBank_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:40 p.m.
2020_365_ENTITY_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:05 p.m.
2022_032_GammaBank_002	Cyber Security Incident	In progress	M1	DataConversion	Feb. 1, 2022, 3:44 p.m.
2022_031_GammaBank_004	Cyber Security Incident	In progress	M1	DataConversion	Jan. 31, 2022, 12:18 p.m.

Figure 125. Register of Data Conversion Task

9. **Run Responder:** the additional information required for reporting needs to be completed by the Incident Reporting Team user for invoking the responder *CS4EU Incident Reporting Data Converter*, which generates the Excel files associated to the regulations enabled and confirmed by the Controller in the managerial judgement, Figure 126.

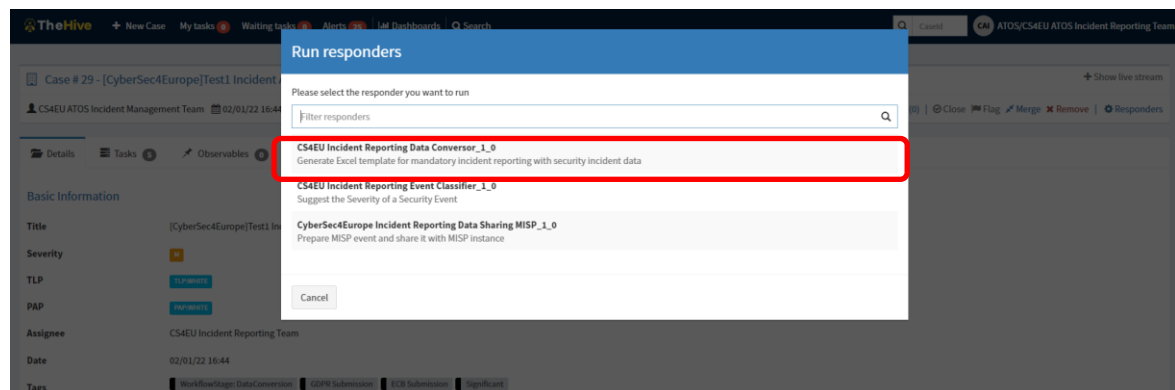


Figure 126. Run responder invocation

The result of the execution is shown at the end of the page, Figure 127.

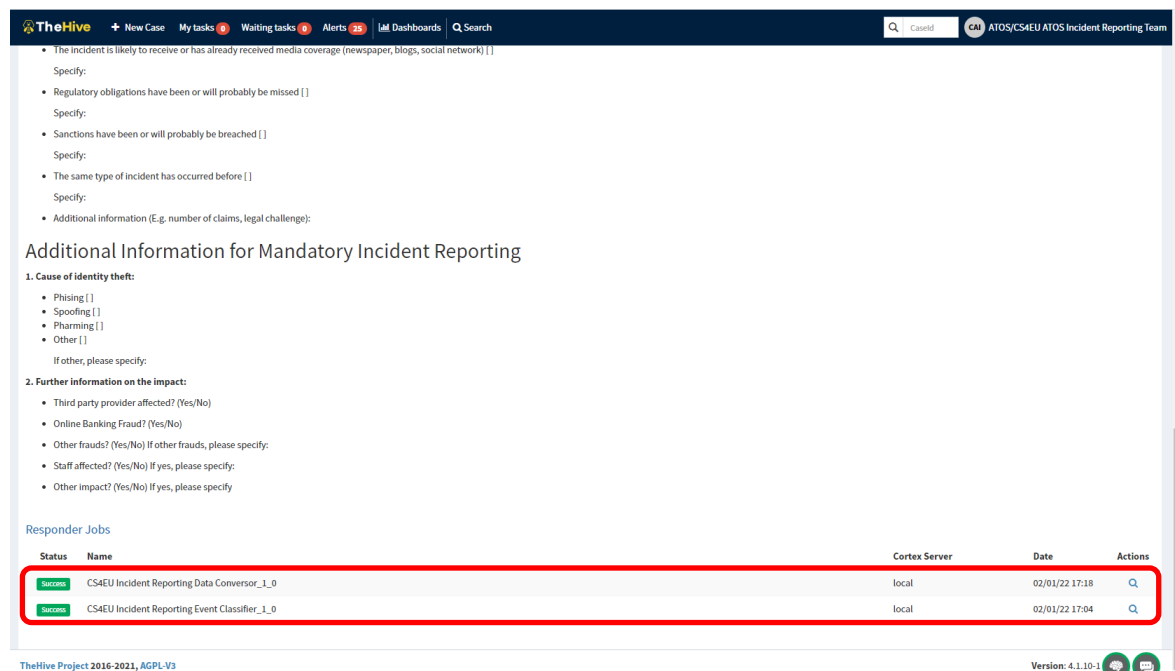


Figure 127. Status of Responders Jobs

The action button opens a text box with the result of the responder execution, Figure 128.

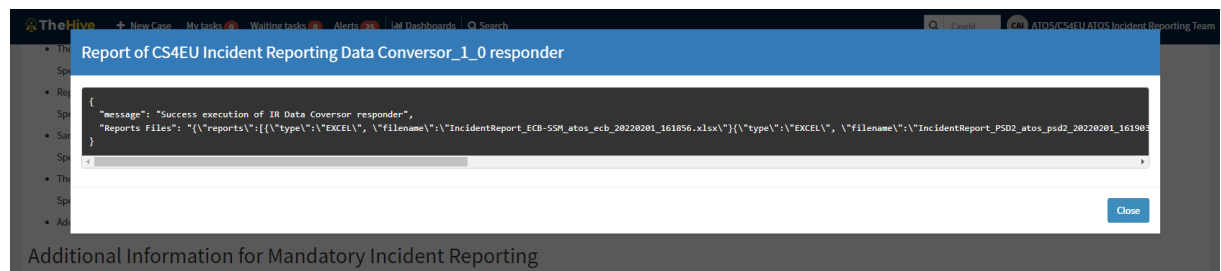


Figure 128. Result of a Responder Job

An email is sent to the Incident Reporting Team user who executed the responder job with the Excel file generated attached, Figure 129.

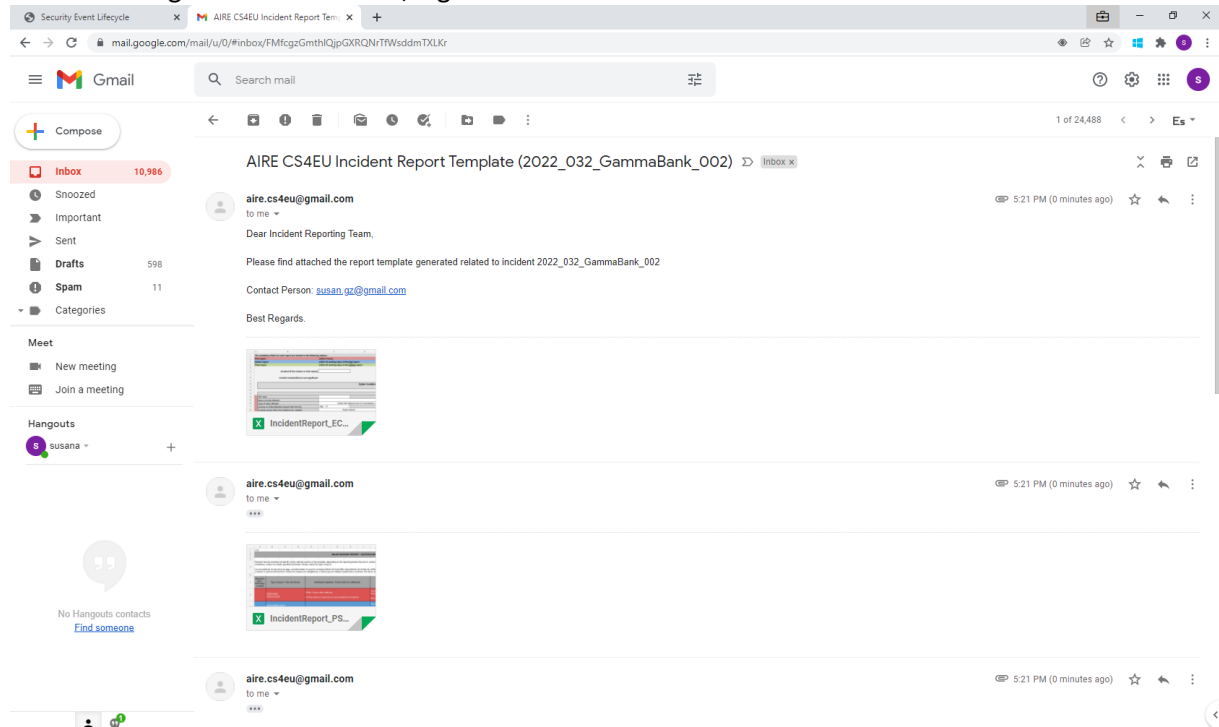


Figure 129. Email to Incident Reporting Team user

All the reports are available in the dashboard under *Reports in Ready Data Conversion*, Figure 130.

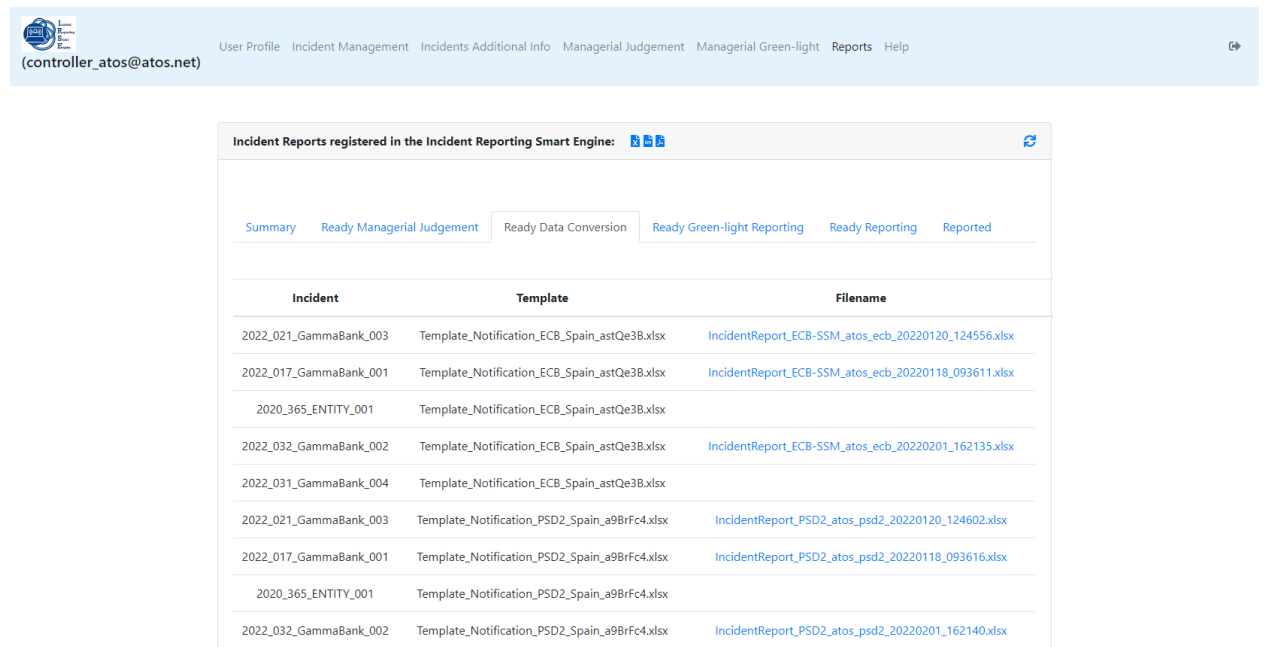


Figure 130. List of available reports

10. Upload modified reports: In case the report is modified, the new version of the documents can be uploaded to the platform from the Reports section. This upload functionality is available from the menu “All” under “Reports” menu. The name of the file selected to upload to the platform must be the same of one already existent, Figure 131.

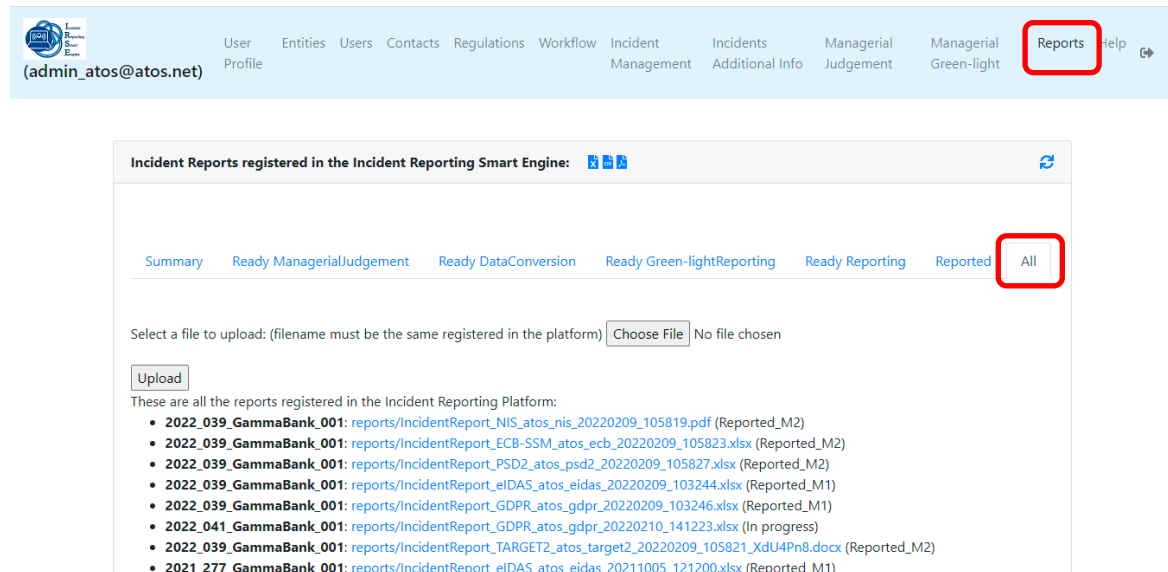


Figure 131. Upload modified reports menu

The new file will be registered with the same name but automatically adding a suffix. In this way, the users will always visualize the last version of the reports but can identify which ones are the generated automatically by the platform (they end with the timestamp <yyyymmdd_HHMMSS>) or have been modified (they end with <yyyymmdd_HHMMSS> followed by _<suffix>), Figure 132

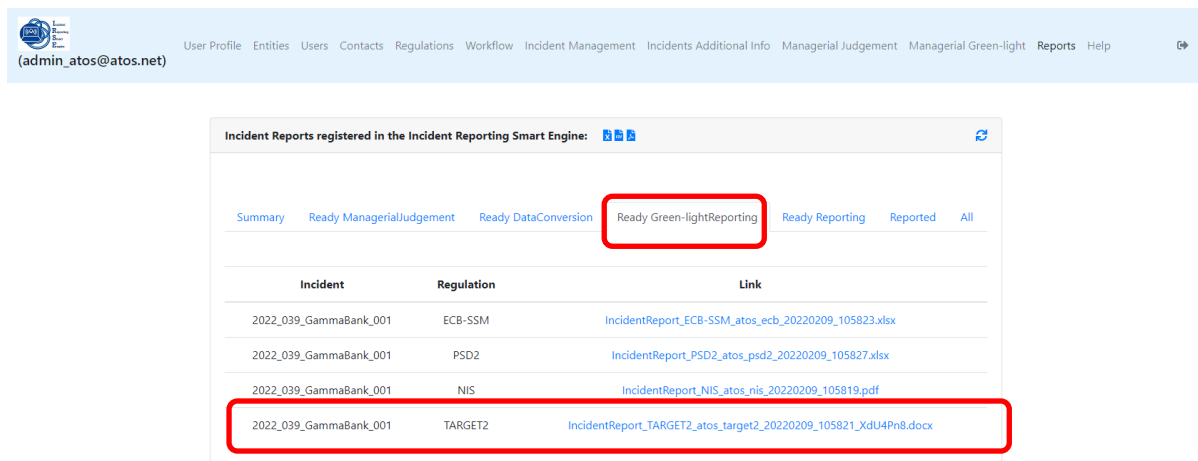
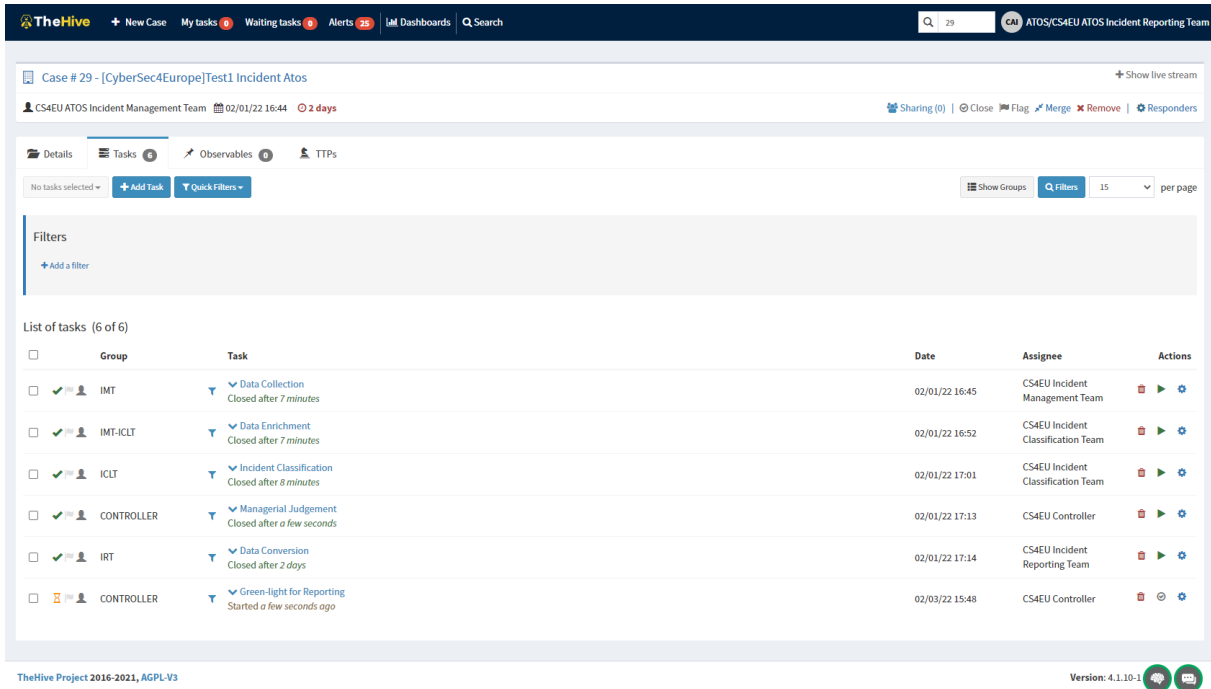


Figure 132. Ready Green-light Reporting list

11. **Close *Data Conversion* Task:** Once the report has been completed and reviewed by the IRT, the user will close the task *Data Conversion* associated and a new task *Green-light for Reporting* will appear assigned to the CONTROLLER, Figure 133.

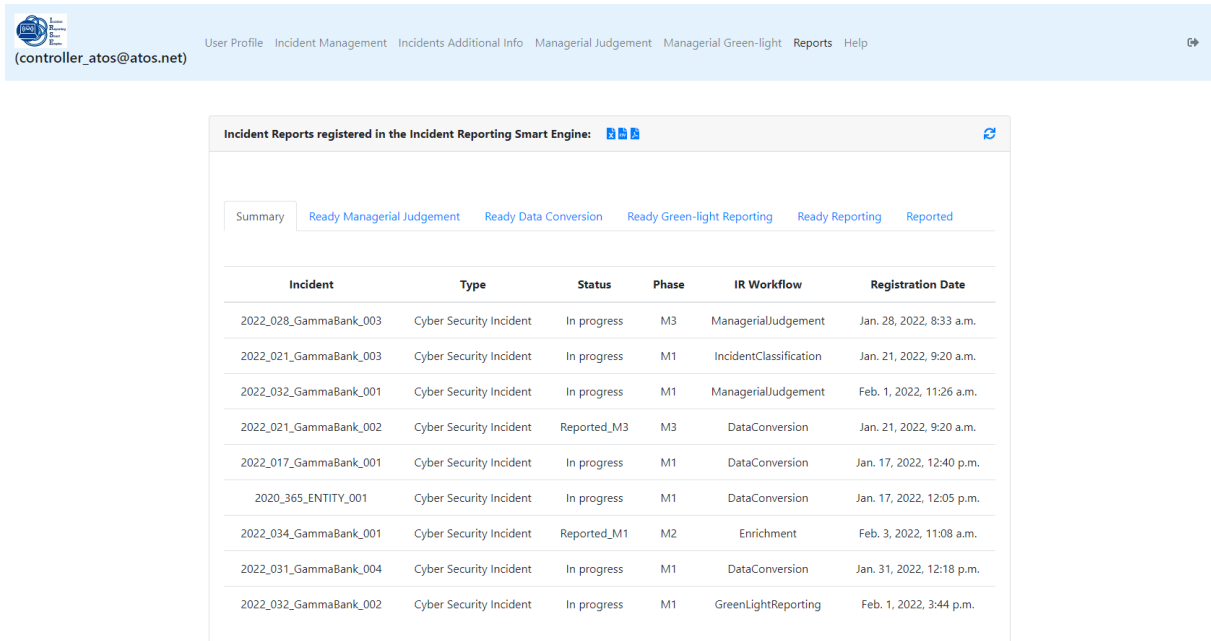


The screenshot shows the TheHive interface for Case # 29 - [CyberSec4Europe]Test1 Incident Atos. The 'Tasks' tab is active, displaying a list of 6 tasks. The 'Data Conversion' task, assigned to the IRT, is marked as closed. A new task, 'Green-light for Reporting', is assigned to the CONTROLLER and is in the 'Started' state.

Group	Task	Date	Assignee	Actions
IMT	Data Collection Closed after 7 minutes	02/01/22 16:45	CS4EU Incident Management Team	[Close] [Flag] [Merge] [Remove] [Responders]
IMT-ICLT	Data Enrichment Closed after 7 minutes	02/01/22 16:52	CS4EU Incident Classification Team	[Close] [Flag] [Merge] [Remove] [Responders]
ICLT	Incident Classification Closed after 8 minutes	02/01/22 17:01	CS4EU Incident Classification Team	[Close] [Flag] [Merge] [Remove] [Responders]
CONTROLLER	Managerial Judgement Closed after a few seconds	02/01/22 17:13	CS4EU Controller	[Close] [Flag] [Merge] [Remove] [Responders]
IRT	Data Conversion Closed after 2 days	02/01/22 17:14	CS4EU Incident Reporting Team	[Close] [Flag] [Merge] [Remove] [Responders]
CONTROLLER	Green-light for Reporting Started a few seconds ago	02/03/22 15:48	CS4EU Controller	[Close] [Flag] [Merge] [Remove] [Responders]

Figure 133. New Green-light for Reporting Task

A new register is added into the *Summary* of the *Reports*, Figure 134.

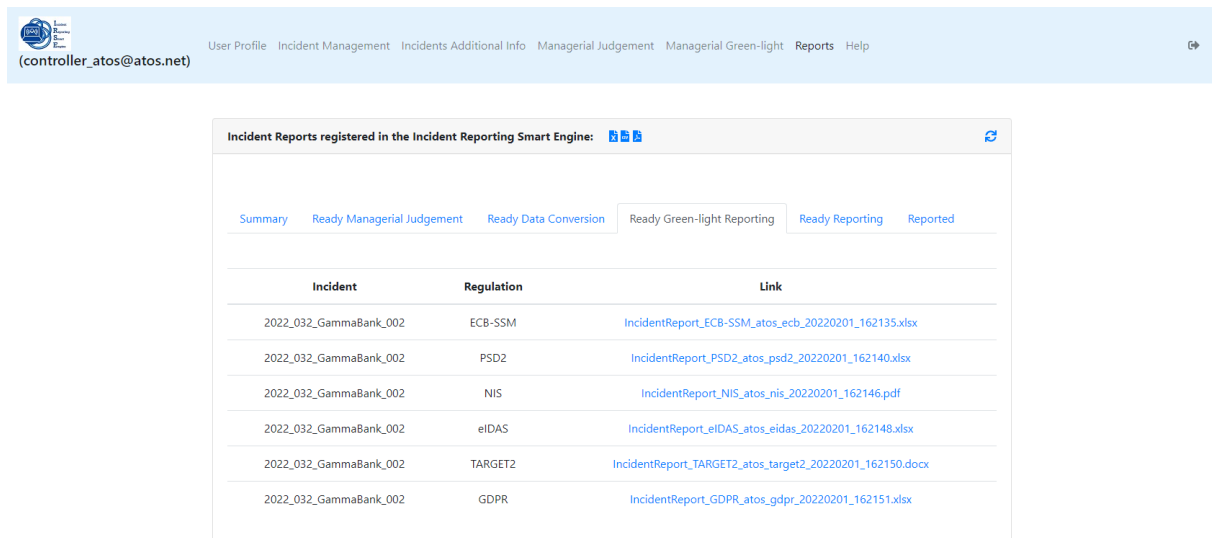


The screenshot shows the 'Reports' section of the interface. A navigation bar includes 'Summary', 'Ready Managerial Judgement', 'Ready Data Conversion', 'Ready Green-light Reporting', 'Ready Reporting', and 'Reported'. The 'Summary' tab is selected, displaying a table of incident reports.

Incident	Type	Status	Phase	IR Workflow	Registration Date
2022_028_GammaBank_003	Cyber Security Incident	In progress	M3	ManagerialJudgement	Jan. 28, 2022, 8:33 a.m.
2022_021_GammaBank_003	Cyber Security Incident	In progress	M1	IncidentClassification	Jan. 21, 2022, 9:20 a.m.
2022_032_GammaBank_001	Cyber Security Incident	In progress	M1	ManagerialJudgement	Feb. 1, 2022, 11:26 a.m.
2022_021_GammaBank_002	Cyber Security Incident	Reported_M3	M3	DataConversion	Jan. 21, 2022, 9:20 a.m.
2022_017_GammaBank_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:40 p.m.
2020_365_ENTITY_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:05 p.m.
2022_034_GammaBank_001	Cyber Security Incident	Reported_M1	M2	Enrichment	Feb. 3, 2022, 11:08 a.m.
2022_031_GammaBank_004	Cyber Security Incident	In progress	M1	DataConversion	Jan. 31, 2022, 12:18 p.m.
2022_032_GammaBank_002	Cyber Security Incident	In progress	M1	GreenLightReporting	Feb. 1, 2022, 3:44 p.m.

Figure 134. Register of Green Light Report

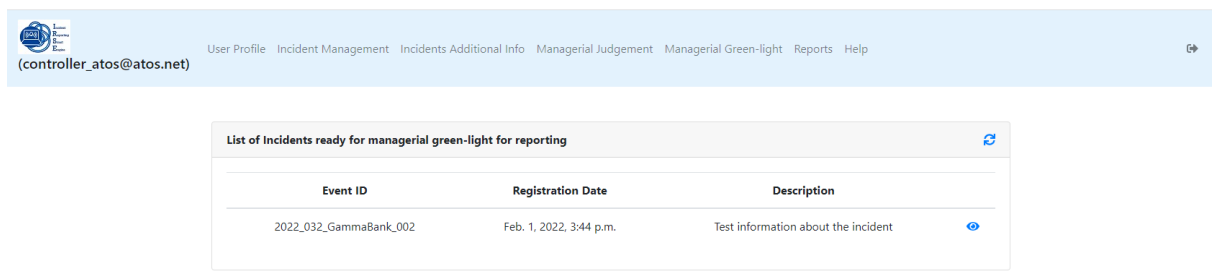
In the tab *Ready Green-light Reporting* inside *Reports* there is a list of all ready reports, Figure 135.



Incident	Regulation	Link
2022_032_GammaBank_002	ECB-SSM	IncidentReport_ECB-SSM_atos_ecb_20220201_162135.xlsx
2022_032_GammaBank_002	PSD2	IncidentReport_PSD2_atos_psd2_20220201_162140.xlsx
2022_032_GammaBank_002	NIS	IncidentReport_NIS_atos_nis_20220201_162146.pdf
2022_032_GammaBank_002	eIDAS	IncidentReport_eIDAS_atos_eidas_20220201_162148.xlsx
2022_032_GammaBank_002	TARGET2	IncidentReport_TARGET2_atos_target2_20220201_162150.docx
2022_032_GammaBank_002	GDPR	IncidentReport_GDPR_atos_gdpr_20220201_162151.xlsx

Figure 135. List of Ready Green-light Reporting

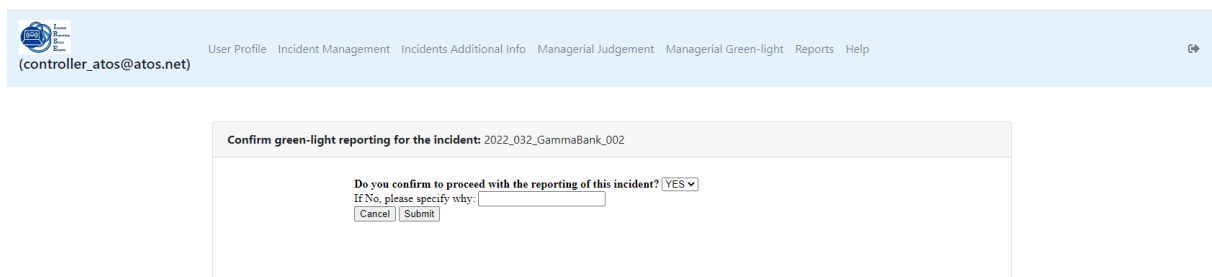
12. **Managerial Green-light:** allows to perform the managerial judgement. The tab displays the list of ready reports, Figure 136.



Event ID	Registration Date	Description
2022_032_GammaBank_002	Feb. 1, 2022, 3:44 p.m.	Test information about the incident

Figure 136. List of Incident ready for managerial green-light for reporting

Selecting in the details (eye) the managerial judgement form will appear to confirm if proceeding with the reporting, Figure 137.



Confirm green-light reporting for the incident: 2022_032_GammaBank_002

Do you confirm to proceed with the reporting of this incident? **YES** ▼

If No, please specify why:

Figure 137. Configuration of green-light reporting form.

Submit button automatically closes the task *Green-light for Reporting*, showing a confirmation, Figure 138.




Figure 138. Confirmation of task closed

It also creates new task, *Reporting & Release*, which will be assigned to IRT user, Figure 139.

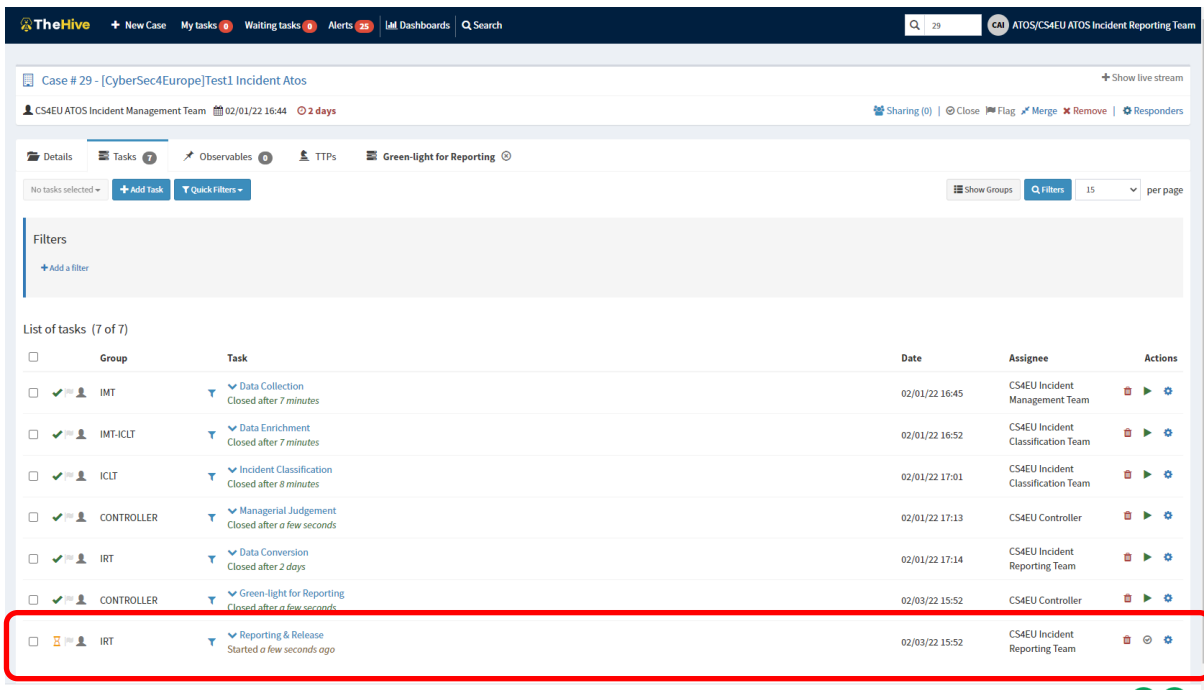
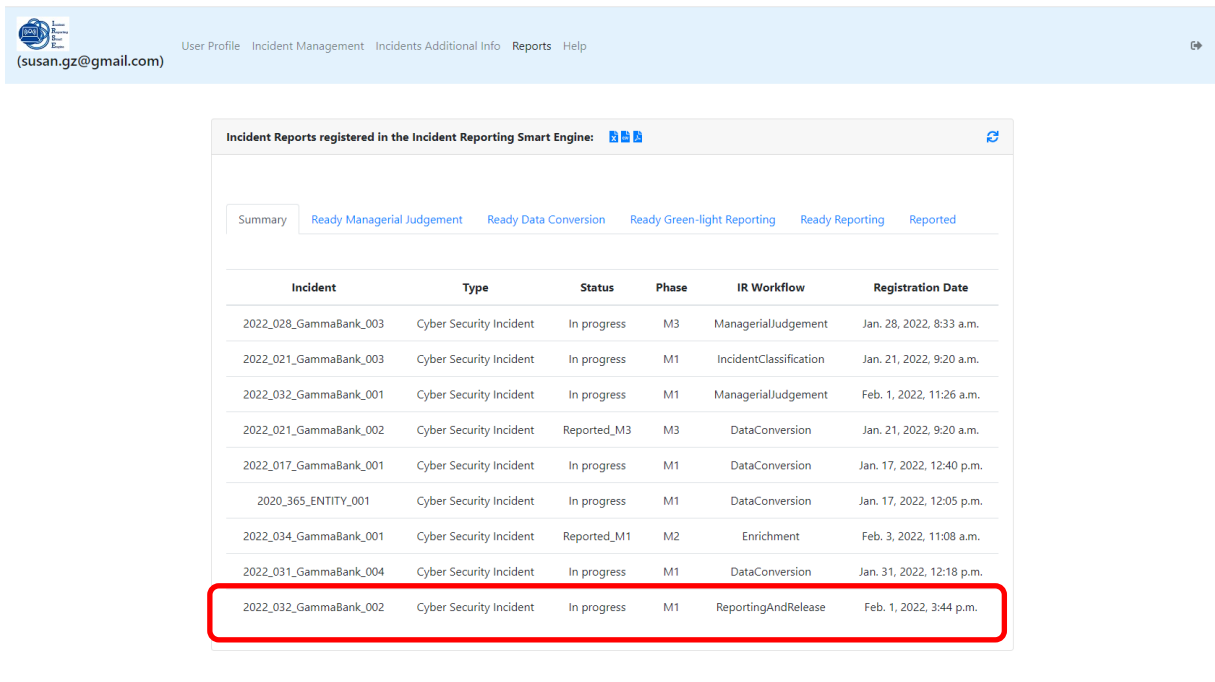


Figure 139. New Reporting and Release Task

The event is registered on the Summary of the Reports, Figure 140.



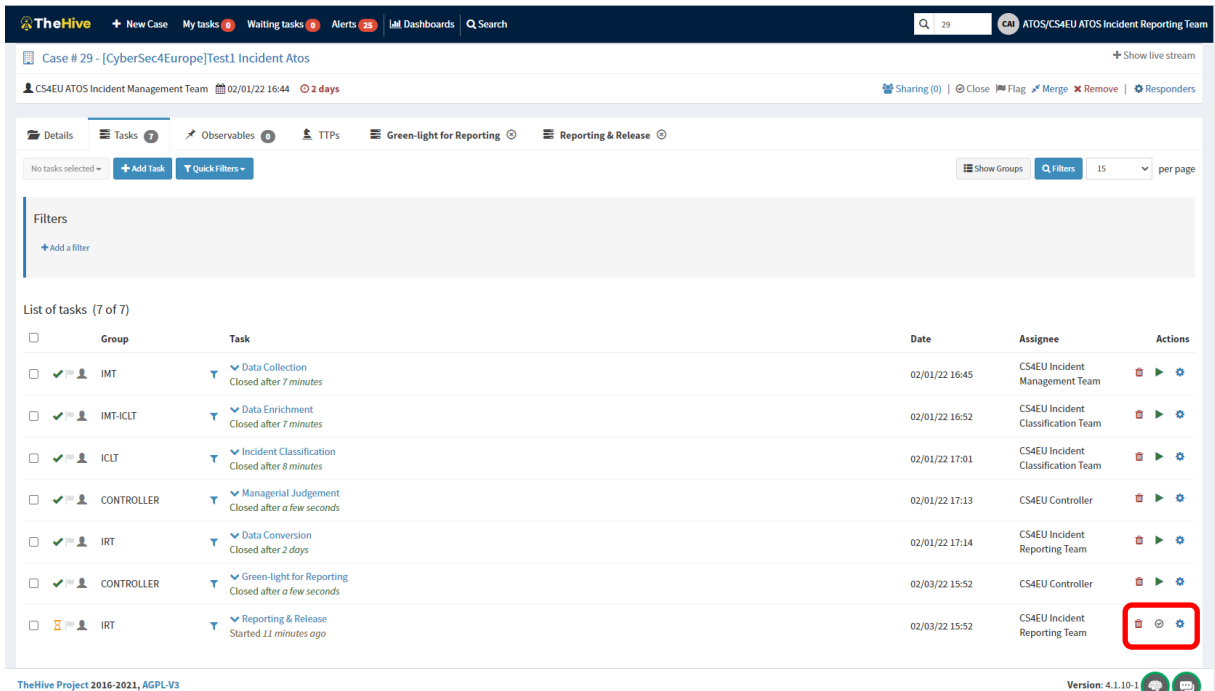
Incident Reports registered in the Incident Reporting Smart Engine:

Incident	Type	Status	Phase	IR Workflow	Registration Date
2022_028_GammaBank_003	Cyber Security Incident	In progress	M3	ManagerialJudgement	Jan. 28, 2022, 8:33 a.m.
2022_021_GammaBank_003	Cyber Security Incident	In progress	M1	IncidentClassification	Jan. 21, 2022, 9:20 a.m.
2022_032_GammaBank_001	Cyber Security Incident	In progress	M1	ManagerialJudgement	Feb. 1, 2022, 11:26 a.m.
2022_021_GammaBank_002	Cyber Security Incident	Reported_M3	M3	DataConversion	Jan. 21, 2022, 9:20 a.m.
2022_017_GammaBank_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:40 p.m.
2020_365_ENTITY_001	Cyber Security Incident	In progress	M1	DataConversion	Jan. 17, 2022, 12:05 p.m.
2022_034_GammaBank_001	Cyber Security Incident	Reported_M1	M2	Enrichment	Feb. 3, 2022, 11:08 a.m.
2022_031_GammaBank_004	Cyber Security Incident	In progress	M1	DataConversion	Jan. 31, 2022, 12:18 p.m.
2022_032_GammaBank_002	Cyber Security Incident	In progress	M1	ReportingAndRelease	Feb. 1, 2022, 3:44 p.m.

Figure 140. Reporting And Release register in Summary of Reports

NOTE: In case the Controller does not confirm proceeding with the actual reporting, the incident reporting process will finish, and the reports will be closed.

13. **Closing the reporting phase:** IRT user can close the task associated *Reporting & Release* from TheHive template. The report will appear in the dashboard under *Reports* tab in the status *Reported_M1*, changing the phase to *M2* and the workflow to a new *Enrichment*, Figure 141.



TheHive Project 2016-2021, AGPL-V3

Version: 4.1.10-1

Group	Task	Date	Assignee	Actions
IMT	Data Collection Closed after 7 minutes	02/01/22 16:45	CS4EU Incident Management Team	[Close] [Flag] [Merge] [Remove] [Responders]
IMT-ICLT	Data Enrichment Closed after 7 minutes	02/01/22 16:52	CS4EU Incident Classification Team	[Close] [Flag] [Merge] [Remove] [Responders]
ICLT	Incident Classification Closed after 8 minutes	02/01/22 17:01	CS4EU Incident Classification Team	[Close] [Flag] [Merge] [Remove] [Responders]
CONTROLLER	Managerial Judgement Closed after a few seconds	02/01/22 17:13	CS4EU Controller	[Close] [Flag] [Merge] [Remove] [Responders]
IRT	Data Conversion Closed after 2 days	02/01/22 17:14	CS4EU Incident Reporting Team	[Close] [Flag] [Merge] [Remove] [Responders]
CONTROLLER	Green-light for Reporting Closed after a few seconds	02/03/22 15:52	CS4EU Controller	[Close] [Flag] [Merge] [Remove] [Responders]
IRT	Reporting & Release Started 11 minutes ago	02/03/22 15:52	CS4EU Incident Reporting Team	[Close] [Flag] [Merge] [Remove] [Responders]

Figure 141. Reporting and Release actions

And a new task “Data Enrichment” will be opened in TheHive to enrich the information about the incident for the Interim Report (M2), Figure 142.

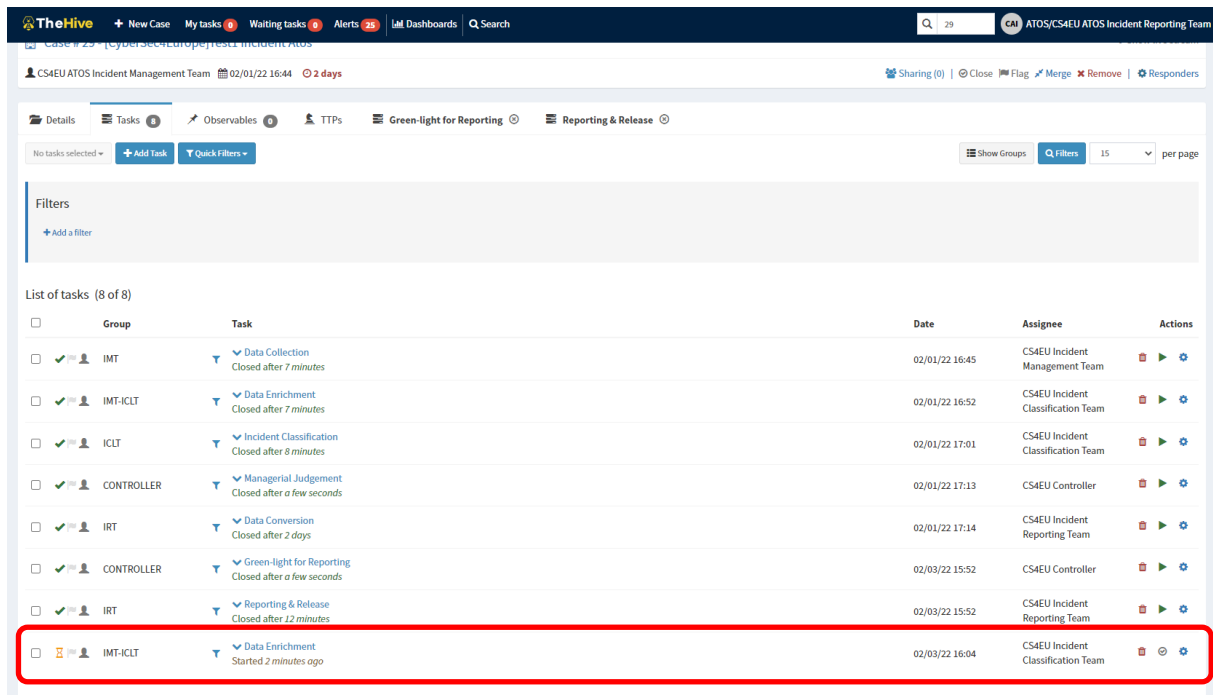


Figure 142. New Date Enrichment Task

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.

14. Security Event Lifecycle: Under *Incidents Additional Info* menu, the *Security Event Lifecycle of All Incidents* can be consulted, Figure 143.

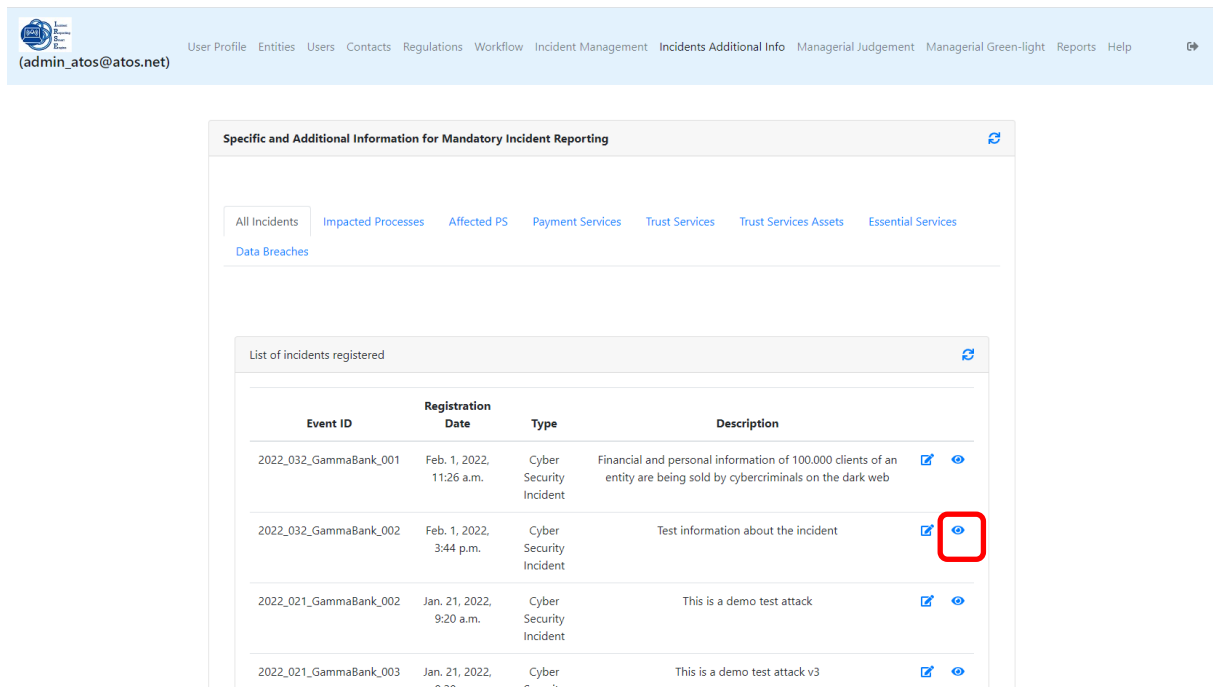
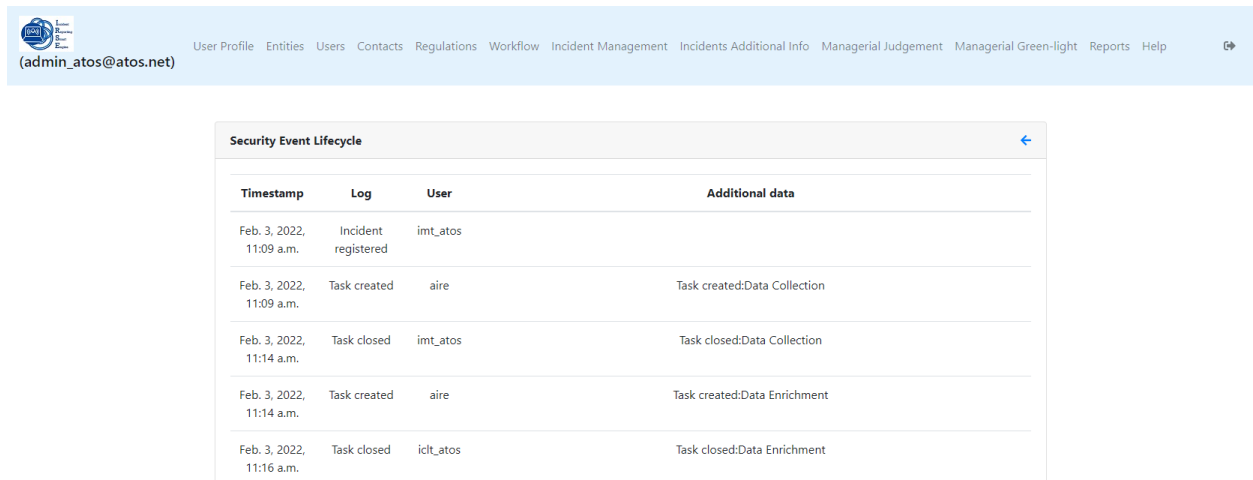


Figure 143. Consult of Security Event Lifecycle

This shows the log with timestamps for all events related with an incident, Figure 144.



Timestamp	Log	User	Additional data
Feb. 3, 2022, 11:09 a.m.	Incident registered	imt_atos	
Feb. 3, 2022, 11:09 a.m.	Task created	aire	Task created:Data Collection
Feb. 3, 2022, 11:14 a.m.	Task closed	imt_atos	Task closed:Data Collection
Feb. 3, 2022, 11:14 a.m.	Task created	aire	Task created:Data Enrichment
Feb. 3, 2022, 11:16 a.m.	Task closed	iclt_atos	Task closed:Data Enrichment

Figure 144. Security Event Lifecycle

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 of the regulations, a notification similar to the one shown in the Figure 145 will be sent to the email configured as Contact User in the incident template.

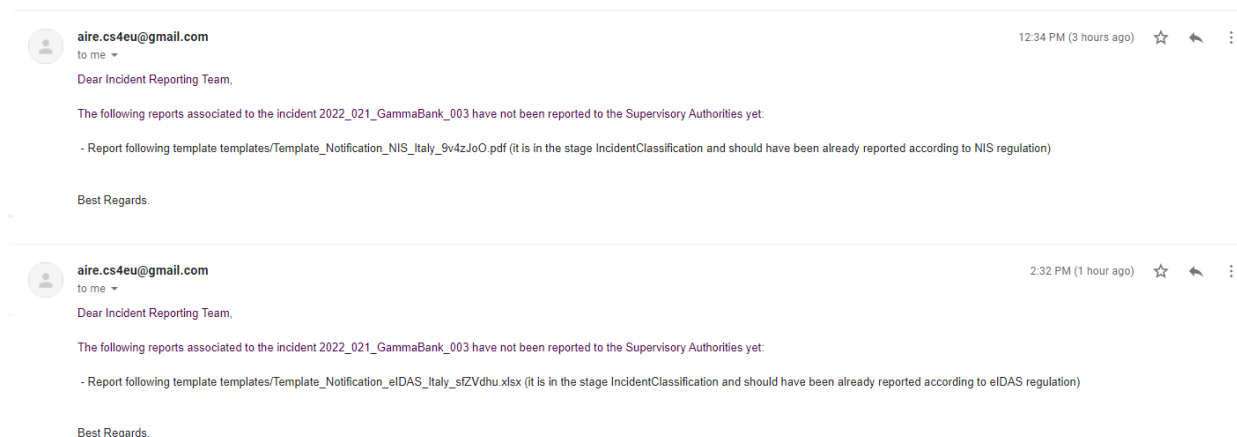


Figure 145. Email to Contact User reminding of pending reports

A notification also be sent to the Contact User in case some exception is detected with respect the Incident reporting workflow. For example, if a user without permissions is closing a task which is not assigned to that profile/role, Figure 146

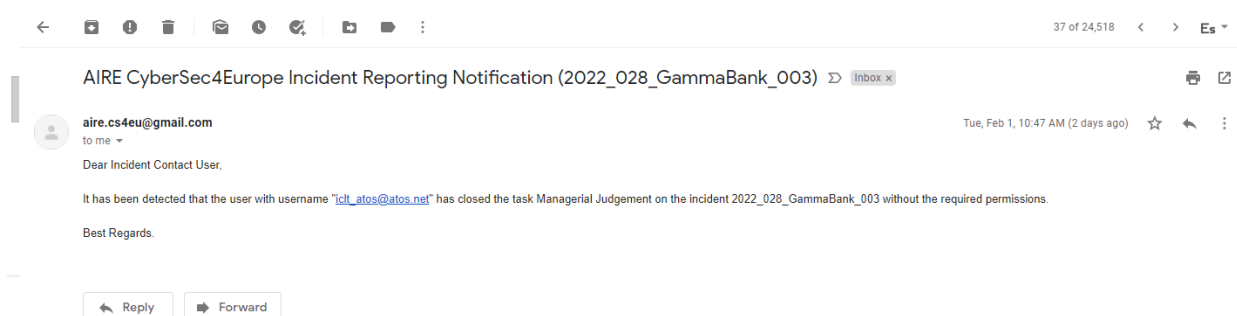


Figure 146. Notification of closing task without permission

I.IV.II INPUT METHODS

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. Also, 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.

I.IV.II.I REST API

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 (CyberSec4Europe - D3.21 Framework to design and implement adaptive security systems. L. Pasquale and A. Hassan. 2022). 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 6: 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.
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 of 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 7. aire-reports-generator service REST API

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

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. Also, 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 8: 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.

TheHive features its own APIs⁹ to control the distinct parts: Organizations, Alerts, Cases, Tasks, etc. This documentation focuses on APIs that are used on automatic mode by other components, such as create new alerts, update a case, or close a task. For this reason, the list of function is not exhaustive.

► Cases:

Table 9: 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 10: 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 11: 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

⁹ <http://docs.thehive-project.org/thehive/api/>

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

I.IV.II.II OTHER APIs

TheHive can receive security events related with the current infrastructure from other systems, for example, 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 one hand, TheHive can receive the alerts generated and sent by Wazuh to a Kafka broker¹⁰. Wazuh has to 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. To enable this feature, it is necessary configure the tool, adding in *application.config* file. It queries the events from the *misp.url.instance* every *1 hour* for the tags *SUNRISE1* and *SUNRISE2*, and then, it creates alarms for the events, which are displayed in the alarm dashboard.

```

misp {
  servers: [
    {
      name = "MISP-NAME"
      url = "https://misp.url.instance"
      auth {
        type = key
        key = "XXXXXX"
      }
      caseTemplate = "Incident Report"
      tags = ["SUNRISE1", "SUNRISE2"]
      #filters
      max-age = 7 days
      max-attributes = 1000
      max-size = 1 MiB
      includedTheHiveOrganisations = ["*"]
      excludedTheHiveOrganisations = []
      #indicate if the tags of the case should be exported to MISP
      event (default: false)
      #exportCaseTags = True
    }
  ]
  # Interval between consecutive MISP event imports in hours(h) or
  # minutes(m) .
  interval: 1 hour
}

```

¹⁰ <https://kafka.apache.org/>