D2.2 System Architecture v1

Deliverable Number	**D2.2**
Deliverable Details: Describes the first version of Resilmesh Architecture	
Deliverable Leading:	**UMU**
Due Date:	**31/05/2024**
Submitted Date:	**25/06/2024**
Author(s)	Jorge Bernal, Pablo Fernandez, David Montoro (UMU), Brian Lee, Xi Lan (TUS), Branka Stojanovic (JR), Jorgeley (SLP), Martin Husak (MUNI)
Reviewer(s):	**RHUL, JR, MUNI; SAB**

Version History

Version	By	Date	Changes
A	Brian lee	25/06/24	SAB review- no changes
A6	Jorge Bernal	03/06/2024	Address internal review comments
A5	All	26/05/2024	Several contributions and refinements and improvements across the document
A4	Brian Lee, Jorge Bernal	22/05/2024	Section 5 and Conclusions
A3	Branka Stojanovic, Jorgeley Olivera, Brian lee, Jorge Bernal, Martin Husak, David Montoro	17/05/2024	Contributions in section 3 and 4
A2	Brian Lee, Jorge Bernal, Pablo Fernandez	06/05/2024	First contributions in Section 2
A1	Jorge Bernal	16/04/2024	ToC, Introduction

Funded by
the European Union

Disclaimer

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Commission. Neither the European Union nor the granting authority can be held responsible for them

Table of Contents

Funded by
the European Union

Funded by the European Union

Funded by the European Union

Executive Summary

This document is the first outcome of Task 2.3 “System Architecture Design”, and it is intended to describe the first High-level Functional Architecture (HLFA) defined in Resilmesh. The architecture identifies several functional components that have been identified considering, 1) the use-cases scenarios described as part of task T2.1 and associated use-case requirements, 2) non-functional requirements defined T2.2, and 3) functional requirements obtained from the technical objectives and goals of the project, being conducted mainly in technical work packages, i.e. WP3, WP4, and WP5.

The architecture supports the aim of Resilmesh to deliver a Cyber Situational Awareness (CSA) based security orchestration and analytics toolset to enable organisations achieve real time defence of essential business functions. Specifically, ResilMesh will help CyS organisations:

- reduce CyS attack surface impact by developing tools to combat complexity (better visibility of assets and services and their dependencies, heterogeneity interoperability and extensibility
- disperse infrastructure, flexible placement of security controls across the CyS infrastructure.
- combat Advanced Persistent Threat (APT) sophistication by developing advanced artificial intelligence (AI) algorithms and tools for early and ongoing attack detection and prediction and improved situation and risk awareness.

The architecture includes not only the structure and functionalities provided by each functional component to comply with the identified requirements, but also the main relationships among the functional components, the services identified per functional components and the interfaces needed.

The architecture definition is used as baseline and reference document, and therefore, will serve as input for technical work packages to develop and integrate the Resilmesh framework. Specifically, the architecture is used in WP3 to guide the development of the Resilmesh platform. The architecture identifies the main approaches to be used in the algorithm and component development of WP4 for threat awareness incident detection as well as for situation assessment and mitigation in WP5.

The overall Resilmesh architecture consists of several planes. Each plane represents a capability that integrates well with other planes. The **infrastructure** plane represents the underlying managed computing and communications cyber systems (CyS) that are protected by Resilmesh. The **Aggregation** Plane is the initial pre-processing step needed to ensure the data coming from the Infrastructure Plane flows properly throughout the pipeline. It collects, aggregates, and normalises data and events from multiple heterogeneous sources. The **Collaboration** mesh provides the connectivity underlay containing the set of supporting functions, protocols and mechanisms required to enable the operation of the ResilMesh system. The **Threat Awareness** layer refers to the set of information processing and analysis functions to manage anomaly

Funded by
the European Union

detection, event correlation and alerting, and attack detection and prediction. The **Situation Assessment** plane contains a set of functional components that collectively provide cyber situational assessment capability to Resilmesh - and which, together with the Threat Awareness plane implements cyber situational awareness in Resilmesh. **The Security Operations Plane** is in charge of deciding and enforcing mitigation actions that should be taken to respond to a detected incident and orchestrate these actions as CoA playbooks and analyse collected information to identify adversaries by tactics and techniques carried out during the incident.

The architecture describes the main planned security workflows supported by the framework, including proactive security posture enforcement through Resilmesh, along with reactive security enforcement workflow aimed to mitigate ongoing attacks through CoA playbooks.

In addition, this document defines several points of extensions of the architecture and open challenges that can be used to extend the ResilMesh framework, through open call specification in WP8. The architecture will be revised after the second software delivery (M24) to consider findings from that work to provide a base for the overall system design and evaluation.

Funded by
the European Union

List of Abbreviations and Acronyms

Abbreviation	Explanation/ Definition
AI	Artificial Intelligence
AID	AI-based Detector
APT	Advanced Persistent Threat (APT)
CSA	Cyber Situational Awareness / Critical Service Awareness
CTI	Cyber Threat Intelligence
AIC	AI-Correlation
CASM	Cyber Attack Surface Management
DN	Data Normalization
ESP	Event Stream Processing
EE	Event Enrichment
EA	Event Aggregator
FL	Federated learning
ISIM	Infrastructure and Services Model
IDS	Intrusion detection systems (IDS)
IPS	Intrusion prevention systems (IPS)
MM	Mitigation Manager
ML	Machine Learning
MB	Message Broker
NSE	Network Status Evaluation
PT	Playbook Tool
PPLF	Privacy-preserving Federated Learning
RM	Resource Orchestrator
RCTI	Robust CTI
SOAPA	Security Orchestration and Analytics Platform Architecture
SOAR	Security Automation, Orchestration and Response
SM	Service Mesh
SIEM	Security Incident and Event Manager
THF	Threat Hunting and Forensics
UEBA	User Entity and Behavioural Analytics
WO	Workflow Orchestrator
XDR	eXtended Detection and Response

Funded by the European Union

Table of Figures

Funded by the European Union

1 Introduction

Motivation & Activities

ResilMesh is devising, designing and implementing a cyber situational awareness-based Security Orchestration and Analytics Platform Architecture (SOAPA) toolset to improve digital infrastructure resilience through fulfilling these objectives:

1. Improving end-to-end data aggregation and security control interoperability in dispersed digital infrastructures
2. Giving CSIRTs better awareness of the service and asset dependencies of their network
3. Helping CSIRTs to build cyber resilience capacity
4. Developing AI based algorithms and tools for early and ongoing attack detection and prediction
5. Developing a situation assessment system to view and forecast network level risk

With these goals in mind, the project has started to build a resilient SOAPA-based platform by combining existing security controls, open-source tools as well as other tools from consortium participants that are being constructed according to high level functional architecture that is presented in this document.

It will develop algorithms and software tools in the project and will integrate these with the platform to form a complete SOAPA system.

The architecture has been designed considering the fact that the associated implementation platform needs to be validated in diverse infrastructure categories, including renewable energy SCADA; smart manufacturing robotics and regional civil infrastructure.

To increase the detection capabilities of the framework, the architecture dedicates different components to deploy AI-based algorithms under different settings centralized and federated, thereby improving the attack detection and prediction for endpoint and network traffic. The architecture deals with the digital infrastructure complexity and heterogeneity by providing tools to give them better awareness of environment dependencies, threats and risk while preserving privacy.

The framework increases the reliability and granularity of shared cyber threat intelligence to improve context for threat hunting and cyber forensics incident response leading to more robust decision making.

Resilmesh Scope and Limitations

The scope of this document is to define the reference high level functional architecture

Funded by
the European Union

Ethics Considerations

We follow the code of conduct specified in D1.6 (Data Management Plan). In essence: ethics considerations related to this deliverable have been tackled in the following ways. Specifically:

- **Personal data and anonymity**: No means of identifying participants beyond Resilmesh Consortium is collected or maintained during or after the completion of this document. No personal data is collected for this document.
- **Confidentiality**: No means of identifying participants nor relevant data are collected or maintained for this document. Data will be protected and kept confidential and will be shared exclusively within the Resilmesh consortium.
- **Usage of the findings**: Findings will be used for Resilmesh innovations and research applications. For example, they will be used to form a requirements analysis for the Resilmesh solution and might be used in future studies and publications.

Participants in this document are Resilmesh Consortium members and they are informed of who to contact for comments, questions or raise complaints. Consent is not required as there is no management of personal data.

Report Structure

The document is structured as follows. Section 1 serves as an introduction to the scope, purpose, and context of the project. Section 2 provides the architecture Foundations and State of the Art analysis. Section 3 defines the architecture and main workflows. Section 4 details the functional components and associated services. Section 5 is intended to define the main open challenges and point of extension. Finally, Section 6 concludes this report.

Funded by
the European Union

2 Architecture foundations and SOTA

Resilience principles in Resilmesh

Resilmesh aims to provide support to enable critical infrastructure providers to improve infrastructure and service resilience across a wide range of dispersed and heterogenous operating environments.

The Resilmesh resilience approach is grounded on several key resiliency best practises described in NIST publication SP 800 160 R2 *Developing Cyber-Resilient System*[1] . SP800 160 v2 R1 sets out to "*address security, safety, and resiliency issues from the perspective of stakeholder requirements and protection needs using established engineering processes*". Most significantly from our point of view is that it identifies fourteen *cyber resiliency techniques* which are described in Appendix D3 of the document. Five of these principles are directly pertinent to CSA to and the table below outlines how these apply- the specific techniques from the standard are indicated in the table below in *italics*.

The table also indicates the primary Resilmesh functional components which implement the different features. These are described in detail further in the document.

Table 1Resilmesh Resiliency Best Practises

Resilmesh Feature	Sub Feature	Description	Functional Component
Context Awareness Ensure the organisation has sufficient overview of its environment to support situational awareness *Contextual Awareness*	Resource Awareness	Maintain information about systems assets, resources, services and their connectivity.	ISIM
	Mission Awareness	Maintain information about enterprise business functions and their dependencies on IS resources and services as well as the threat status of resources.	ISIM
	Threat Awareness	Ensure awareness of threat actors and adversaries. Maintain ongoing observation and analysis of threat event and indicators and enable better prediction of cyber security threats.	SIEM, NSE, RCTI CASM CSA
Security Analytics *Analytic Monitoring*	Sensor Fusion	Combine monitoring data from different sources at different points in the system as well as with externally provided CTI	AD, SIEM EA
	Threat Monitoring	Monitor and analyse system components to look for indication of adversary activity	SIEM, AD NDR
	Forensics and Behavioural Analysis	Analyse indicators of compromise and behaviour and other evidence of adversary presence or activity	THF

[1] https://csrc.nist.gov/pubs/sp/800/160/v2/r1/final

Funded by
the European Union

	Situational Awareness	Correlate and analyse data to monitor the ongoing organisation wide situational awareness	SIEM, AIC, NSE

Adaptive Ensure ResilMesh can be used for a variety of application domains, computing topographies, and mission types. *Coordinated Protection* *Dynamic Positioning* *Adaptive Response*	Heterogeneous Infrastructures	Ensure ResilMesh can support organisations of different types, domains and size.	DN, Collaboration Mesh
	Flexible function allocation and composition	Ensure that security functions can be flexibly located where required across the dispersed IS infrastructure. Enable analytic functions to be pipelined/composed to support flexible aggregation and processing of events.	SM, RO
	Interoperability	Ensure that diverse security function and tools can share information and actions	Collaboration Mesh
	Operation Orchestration	Ensure that critical cyber operations are automated, and course of action playbooks are prepared to ensure collaboration between diverse tools to ensure effective detection, analysis and response.	WO, Playbook Tool
	Risk based response	Ensure that threat response/mitigation actions are chosen to prioritise those IS capabilities required to ensure enterprise critical mission continuity.	MM

Cyber situational Awareness

ResilMesh aims to provide critical infrastructure security teams with a greater cyber resilience capability by improving cyber resilience using cyber situational awareness (CSA). CSA is an adaptation of generic situational awareness known from other fields, such as aviation or military, and is most often defined using the three-level definition as *"the perception of the elements in the (cyber) environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future"* [ENDSLEY, JIRSÍK].

The three levels of CSA are displayed in Figure 1. Each level depends on the level below, so that the projection of the projection of the environment requires a solid comprehension of it, which is fully dependent on the perception. Achieving the higher

Funded by the European Union

levels of situational awareness allows for making the correct decisions and taking the appropriate actions.

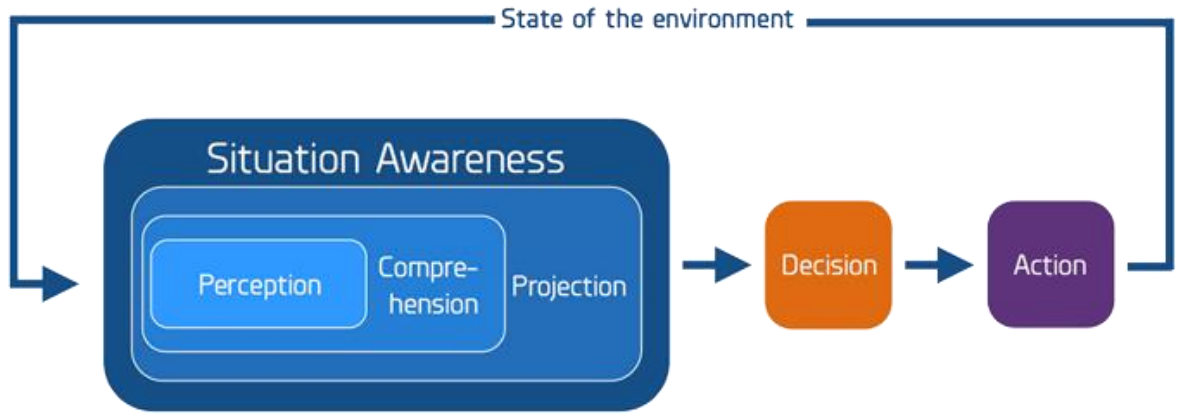

Figure 1 - Situation Awareness[2]

Taking the actions and other changes in the environment make this a perpetual cycle; the changes in the environment call for going through the perception again, followed by comprehension and projection and other decisions and actions.

Perception relates to i) identification and enumeration of the assets and components that constitute the system as well as their dependencies and security posture (*critical infrastructure structure awareness*) and ii) being aware of the threats within and without the system that may impact the system (*threat awareness*). Perception is usually supported by asset management systems, attack surface mapping, vulnerability scanners, intrusion detection systems, cyber threat intelligence, and other sensors. However, the number of heterogeneous sensors and data types, as well as the quantity of data make the perception a challenging task. [HUSÁK, GUTZWILLER]

Comprehension relates to understanding the risks that threats may pose to an organisation by combining perception components with knowledge of the organisation's critical missions and assets. Comprehension is usually facilitated by various methods of data visualisations, correlation, and analytics. Achieving the level of comprehension requires timely perception of the environment and frequent actualization when the elements or the environment changes. Similar to perception, comprehension of the cyber environment is hindered by large volumes of heterogeneous data in practice. Moreover, only a small fragment of the data is usually relevant for investigating cyber incidents, which makes the cyber environment especially prone to "information overload, meaning underload" and poses additional challenges to resolving the task. [HUSÁK, GUTZWILLER]

Projection refers to the prediction of future events based on knowledge of current and past threats and state of the environment. There are three types of projection that we may consider [4]:

[2] https://stan-institute.com/en/news/situational-awareness/

Funded by
the European Union

- *Attack Projection* of *Attack Step Prediction* answers the questions "What is an adversary going to do next?" The existing approaches build upon alert correlation and multi-step attack detection, often intertwining with attacker intention recognition. Existing research shows promising results for real-time predictions, but with only short time left for attack mitigation.
- *Attack Prediction* or *Intrusion Prediction* answers the question "What type of attack will occur when and where?" This is a more general question than the previous one and is rather approached by network entity reputation databases and predictive blocklisting. Analogous techniques could be used to predict the emergence of novel vulnerabilities.
- *Network security situation forecasting (NSSA)* or *Attack Forecasting* answers the question "How is the overall situation going to evolve?" In this case, the predictions are not made for specific attackers or targets, but for their overall numbers throughout the network. Existing NSSA methods based on time series allow for forecasting the cyber-attack rates and estimate the workload of the security teams.

Nevertheless, despite numerous promising research results, it is imperative though to remind that successful projection heavily depends on well executed perception and comprehension.

The CSA system enables an organisation to assess the overall risk situation and to take mitigation actions when and where necessary. A cyber situational awareness capability is critical to enable organisations implement mission based cyber resilience solutions. Several technical solutions and toolsets for supporting CSA are available as open-source (e.g., CRUSOE [CRUSOE]) or commercial products. However, situational awareness is not solely a technical problem and requires the coordination of technical, operational, and management perspectives to be used properly [AHMAD].

In the context of Resilmesh, CSA refers specifically to the perception of the enterprise's security posture and its threat environment, the comprehension of both as risks, and the projection of their status, which enables prompt mitigation of threats and their materialisations.

Collaboration Mesh

The *Collaboration Mesh* refers to the principles and mechanisms employed in Resilmesh to **improve interoperability and integration** of different - often siloed - security controls and tools.

It is predicated on two main approaches

1. A **connectivity underlay** containing the set of supporting functions, protocols and mechanisms required to enable the operation of the ResilMesh system across dispersed operation environments, including support for creation of data processing pipeline to support event and anomaly detection flows through the system. It will use, to the greatest extent possible, ***container orchestration capability*** to implement services and APIs for the upper layers to compose, deploy, modify and migrate the security microservice pipelines across edge and

Funded by the European Union

cloud. The mesh will also provide ***data streaming backplane*** to transfer real-time data as needed and will provide a pub/sub broker to transfer and route messages. Moreover, it will provide a ***service mesh capability*** for critical Resilmesh deployments where extended app resilience, observability, and security are required.

2. **Composability and interoperability of security controls and tools**: Resilmesh adopts and adapts principles from existing industry best practise guidelines, in particular the Gartner Cybersecurity Mesh[3] and the Open Cybersecurity Alliance Open XDR initiative[4]. These efforts[5] both aim "*to facilitate interactions between security products, using a mix of open standards and interfaces, proprietary APIs, and point integrations*".
3. ResilMesh will provide an ***operations orchestration*** framework to enable integration of several tools in the security operations centre into a single plane and will enable automation of several critical processes, including CTI processing, event management, etc. It will be implemented using existing open-source components. The framework will use ***existing interoperability standards*** to provide a uniform message passing syntax and semantics between the SOC applications e.g. STIX for CTI sharing and OpenC2 Command and Control language to provide a set of common commands to trigger Course of Action (CoA) playbooks based on OCA CACAO[6] to enable sharing of playbooks.

Distributed AI-based Anomaly detection

Anomaly detection is a crucial component for identifying potential attacks based on identifying deviations from normal network traffic behaviour. Traditional methods for detecting such anomalies often rely on signature-based mechanisms, which have been demonstrated to be ineffective for detecting zero-day attacks [Khraisat]. However, the advent of machine and deep learning applied to attack detection has significantly enhanced the ability to analyse and interpret complex data patterns associated with cybersecurity threats. The implementation of deep learning models typically necessitates the availability of substantial computational resources and access to centralised, extensive datasets, which can be challenging due to concerns regarding data privacy or data governance regulations. In this context, the use of distributed artificial intelligence (AI) through federated learning presents a promising solution.

Federated Learning [McMahan] allows multiple decentralised entities, such as servers across various network segments, to collaboratively train a global model while maintaining the confidentiality of all training data, thus negating the need for direct data exchange. This method is particularly advantageous for cybersecurity anomaly detection, in addition to the aforementioned advantages pertaining to privacy and

[3] https://www.gartner.com/en/information-technology/glossary/cybersecurity-mesh
[4] https://github.com/opencybersecurityalliance/oxa
[5] The Cybermesh and the OpenXDR are described in greater detail in D3.1 "Resilmesh Platform Reference Implementation"
[6] CACAO https://docs.oasis-open.org/cacao/security-playbooks/v2.0/security-playbooks-v2.0.html

Funded by
the European Union

training computation decentralisation. Furthermore, it improves attack detection time since the final detection model after federated training is already at the locations where it will be deployed and used.

Furthermore, the distributed application of AI through federated learning not only enhances the generalisability and robustness of anomaly detection models but also addresses the scalability demands of cybersecurity frameworks. Training models across diverse environments improves their performance against new and sophisticated types of attacks [Singh], a crucial feature for cybersecurity where anomalies are unpredictable. The decentralised nature of federated learning also reduces risks such as model poisoning and security vulnerabilities inherent in centralised data collection systems. However, model poisoning and membership inference attacks are still present and need to be mitigated, for instance, by using Differential Privacy techniques [Ruzafa]. Furthermore, Federated Learning can scale to thousands of nodes [Kairouz], enabling substantial computational power to be distributed across nodes without the need for extensive data transfer or central processing. This scalability is of paramount importance for the real-time monitoring and mitigation of threats within network traffic, ensuring continuous and effective anomaly detection.

Security Operations and Analytics Platform Architecture

The ResilMesh software architecture is constructed according to the Security Operations and Analytics Platform Architecture (SOAPA)[7] concept. SOAPA is an emerging SOC architecture that essentially combines SIEM with an extended set of capabilities as below:

- **Security Incident and Event Manager** (**SIEM**)- The SIEM is primarily responsible for logging events and alerts and using rule-based correlation to highlight any deviations or alarms. The Resilmesh SIEM is implemented by Wazuh and the ELK stack
- **eXtended Detection and Response (XDR**): is an extension and centralisation of the Endpoint Detection and Response (EDR) concept. XDR is a consolidation of tools and data that provides extended visibility, analysis, and response across endpoints, workloads, users, and networks. XDR functionality in Resilmesh is primarily provided by the Wazuh XDR tools as well as the AI correlation (AIC) components.
- **Network security analytics**: SIEM's log analysis and XDR host behaviour monitoring are complemented by flow and packet analysis in SOAPA. This is implemented in Resilmesh by the NSE, NDR and AIC functional components.
- **User Entity and Behavioural Analytics (UEBA)**: is a category of security solutions that use innovative analytics technology, including machine learning and deep learning, to discover abnormal and risky behaviour by users, machines

[7] Goodbye SIEM, Hello SOAPA https://bit.ly/3DMQKcB

Funded by the European Union

and other entities on the corporate network. UEBA will be addressed in Resilmesh with the use of the Anomaly Detection (AID) function component whether standalone or embedded in other functions such as NDR.

- **Cybersecurity Asset Management**: consolidates, standardises, removes duplicates, and correlates asset information from various data sources, providing a comprehensive cyber asset inventory. This approach identifies security vulnerabilities and facilitates automated remediation, lowering the attack surface and simplifying workflows. In Resilmesh asset management is realised by the CASM and ISIM functional components.
- **Risk Assessment**: triages alerts and estimate network risk based on security posture and asset criticality. This is realised in Resilmesh through the CASM, CSA and NSE functional components
- **Security orchestration and Automated Response (SOAR):** automates repetitive and routine security tasks and processes and orchestrates security processes and workflows across various security tools and technologies, and facilitates incident response by providing capabilities to investigate, triage, and respond to security incidents.

These various services are tightly integrated into a SOAP security operations stack consisting of:

- **Common data services** in security operations manage a growing volume of diverse data types, amounting to terabytes per day. This data is ingested, processed, and prepared for analysis within SOAPA which centralises these functions and allows analytics engines to focus solely on analysis tasks.
- The **software services layer** within SOAPA, akin to traditional middleware, handles the delivery of data elements to analytics engines in appropriate formats and contexts.
- The **analytics layer** within SOAPA is where data is transformed into actionable insights using tools like threat intelligence platforms, behavioural analytics, and SIEM (Security Information and Event Management) systems.
- Finally, the **security operations layer** within SOAPA executes actions based on the analysed data, such as system quarantining, security control modifications, or software patch installations, among other security tasks necessary for effective response and mitigation.

These are organised architecturally in a generic SOAPA architecture[8] as shown in

[8] https://www.csoonline.com/article/566687/security-operations-activities-to-watch-in-2019.html

Funded by the European Union

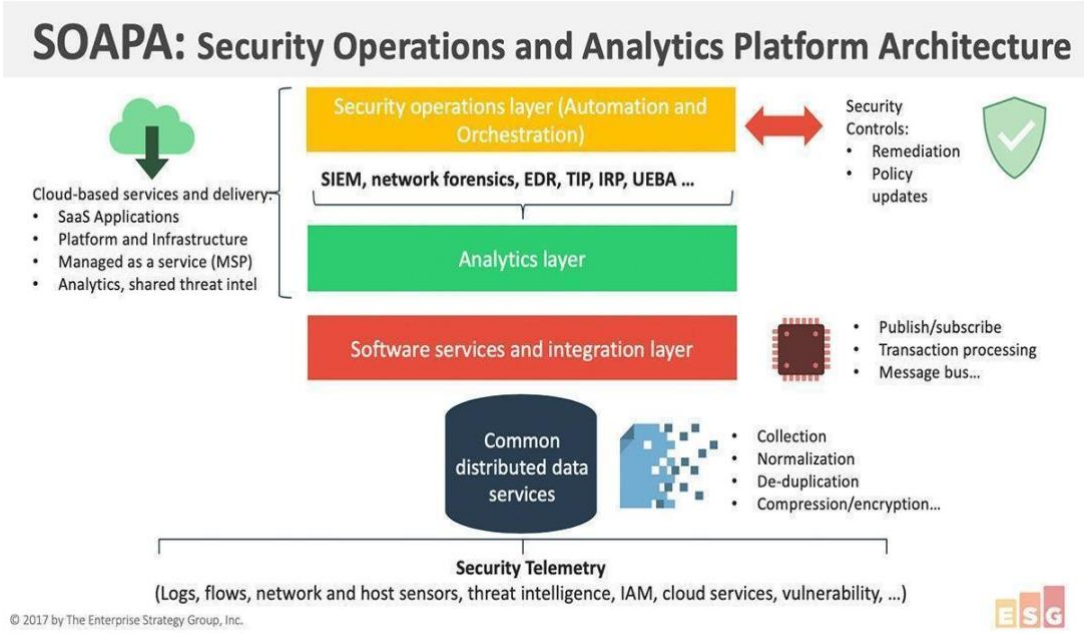

Figure 2 - SOAPA

It maps the ResilMesh CSA + SOAR functions into the SOAPA concept and layers:

- Security Operations Layer – contains parts of the collaboration mesh plus SOAR functions.
- Analytics Layer – contains threat awareness and situation assessment
- Software Services Layers – contains parts of the collaboration mesh function
- Distributed Data Services – contains the aggregation functions

This is shown in the figure below:

Funded by
the European Union

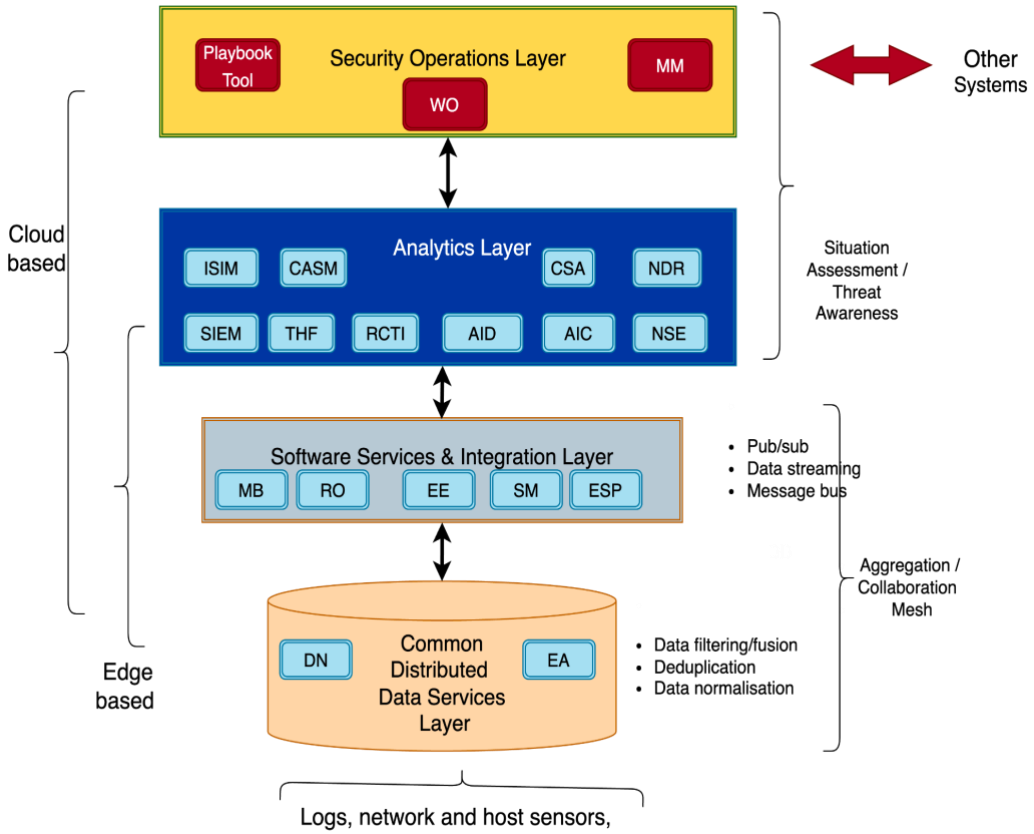

Figure 3 - Resilmesh SOAPA

Funded by the European Union

3 High Level Architecture design

This section describes the Resilmesh High Level Functional Architecture design. The section dedicates different subsections to describe the planes and the associated functional components that fall into each plane. In addition, this section describes the main workflows envisioned in the project to satisfy the project goals.

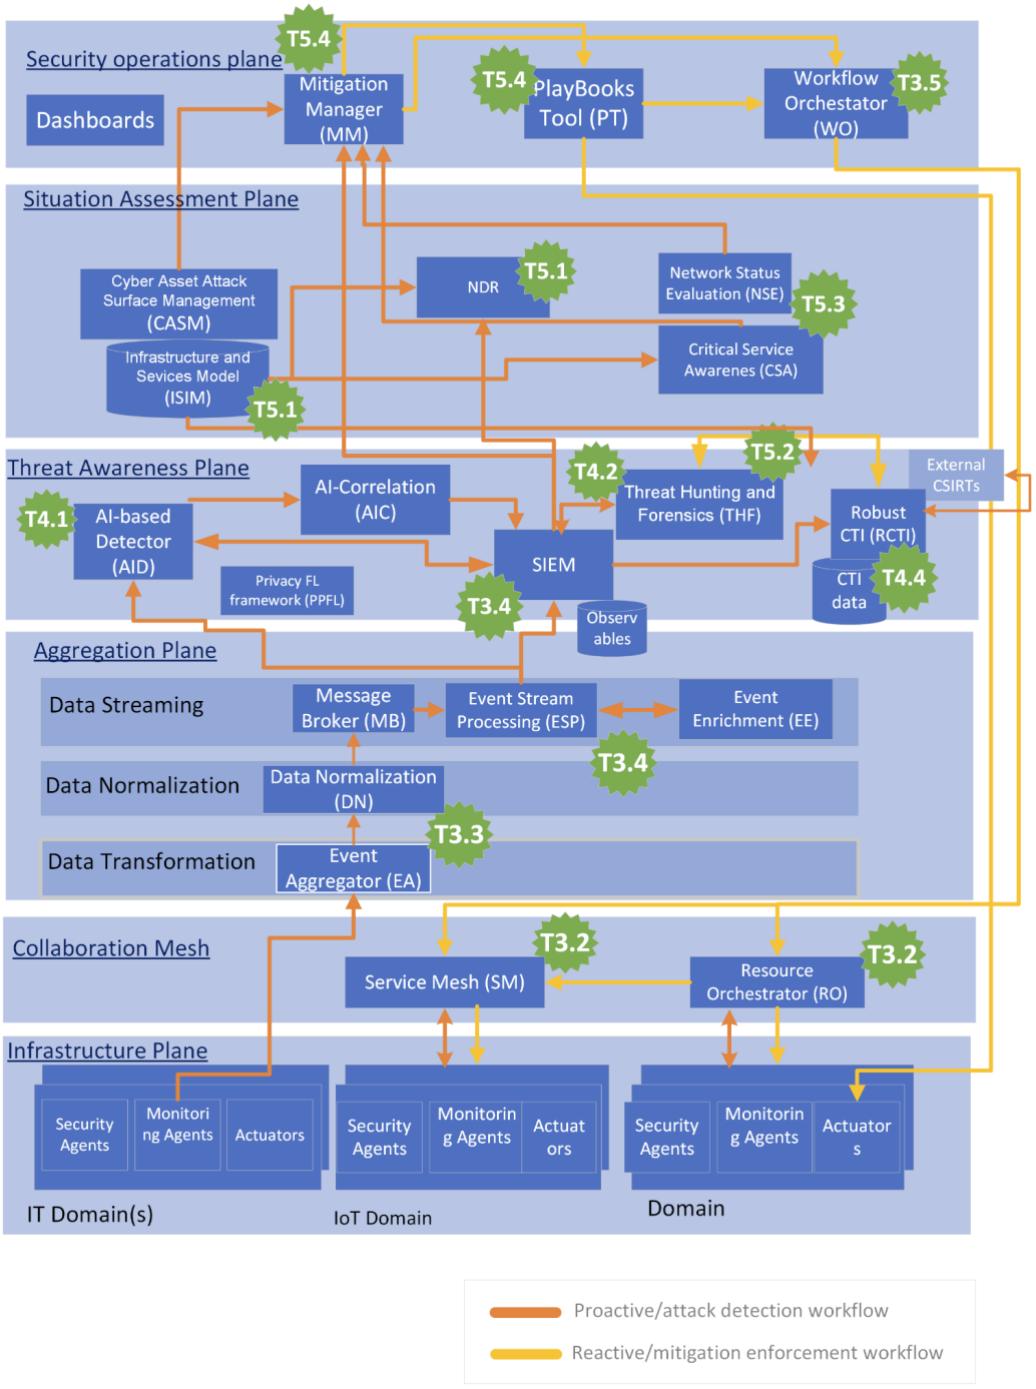

Figure 4 High Level Architecture

Funded by the European Union

The Infrastructure Plane

The network communication between components is handled in this plane. It includes both physical and virtual network elements responsible for tasks such as forwarding the traffic according to commands and rules received by the Network Controllers through the southbound API. In virtualized scenarios, this infrastructure includes the cloud computing technologies (i.e., computing, storage and networking) to deliver virtual Infrastructure-as-a-Service (IaaS). Special controllers are used to control devices in more heterogeneous scenarios that include IoT.

Computing and communications Cyber Systems (CyS) domains include a wide range of civil and critical infrastructures that have very varying technologies, topology, and application requirements. Topologies can be widely dispersed (water and energy infrastructures), concentrated in a few locations (manufacturing, health) or widespread (communications infrastructure). CyS resources are a mixed of constrained (IoT/edge) and powerful (cloud) computing devices and maybe a single technology (IT or OT) or a mix of both. ResilMesh use case pilots have thus been chosen to demonstrate the applicability of the ResilMesh approach across these dimensions e.g. across the three topologies types, dispersed (renewable energy), concentrated (flexible manufacturing) and widespread (regional infrastructure). In Resilmesh this Infrastructure plane represents the infrastructure needed for the three use cases: renewable energy, and flexible manufacturing and civic regional infrastructure.

- The **manufacturing environment** considers several industrial robots, powered by the robotics “de factor standard” of Robot Operating System (ROS) - middleware. In such a manufacturing environment, IT and OT operations are interconnected according to the well-known layered Purdue Enterprise Reference Architecture and Purdue reference model for ICS, SCADA and OT systems. The manufacturing environment is structured in at least 5 levels, including, corporate Level 4 IT networks, Level 3 Operations, Level 2 - Central control Stations, Level 1 Control Network and controllers and Level 0 comprising physical actuators. Hardware components will include widely used and popular industrial grade robotics platforms, industrial grade networking equipment and engineering infrastructure. Software components will include firmware versions by robot manufacturers, drivers (official manufacturer ROS drivers) and available ROS versions.
- The **renewable energy Infrastructure** is centralised platform that collects data from PV plant SCADA distributed all over Italy. The PV plant O&M use case addresses the integration of different, complementary technologies including SCADA technologies, cybersecurity technologies and protocols (to ensure data integrity and availability and avoid any malicious interference that might alter the behaviour of the plants and), telecommunication protocols.
- The **civil region infrastructure** is a physical and virtualized and IT system and network composed of several data centres and information systems that are deployed in a distributed and dispersed subnetworks managed by the civil government.

Funded by
the European Union

Security probes or **sensors** are deployed in the infrastructure to support monitoring the managed system. For instance, Packetbeat [9]agents is an analyser that gathers and sends data from hosts and containers, FileBeat[10] agents for collecting logs from security devices, cloud, containers, hosts. In addition, special purpose probes for OT can be shipped for gather logs, and data from the OT network protocols and system logs.

In addition, this infrastructure layer will host the set of **actuators** that will help to configure dynamically and on-demand the security command and controls and enforce the remediation(s) and mitigations as requested by the Resilmesh framework. These actuators and security agents can be managed through diverse protocols such as OpenC2.

Aggregation Plane

The aggregation plane is the initial pre-processing step needed to ensure the data coming from the Infrastructure Plane flows properly throughout the pipeline. It collects, aggregates, and normalises data and events from multiple heterogeneous sources including logs, IDS, network sources, AI models etc. It contains a rich set of data/event (including network traffic) filtering, fusion, logging, storage and forwarding functions.

It is comprised of three functional sub-layers:

- **Data Transformation** - this receives the data from on-device agents and other sensors and filters and aggregates the incoming raw data. It contains the *Event Aggregation* (EA) functional component.
- **Data Normalisation** - this layer adapts the aggregated data to a common data format scheme for processing in Elasticsearch - the Elasticsearch Common Schema (ECS). It contains the *Data Normalisation* (DN) functional component.
- **Data Streaming** - this layer distributes the normalised data to the upper layers and also performs further processing on the data streams, if required. It contains the *Message Broker (MB), Event Stream Processing (ESP) and Event Enrichment (EE*) functions.

Event Aggregation (EA): This functional component ingests event logs from multiple sources and transforms the events via various operations (routing / logging / fusion / filtering / augmentation / reduction / monitoring) and then outputs the data to one or more destinations via the streaming layer. Outputted events are transferred to the DN function

Data Normalisation (DN): transforms data from multiple disparate formats coming from different sources to a single common format that can then be used for analytics, visualisation, reporting, etc

[9] https://www.elastic.co/es/beats/packetbeat
[10] https://www.elastic.co/es/beats/filebeat

Funded by
the European Union

Message Broker (MB): This is an intermediary function that applications and services use to communicate with each other to exchange information. Message brokers can be used to route and deliver messages to the required destinations.

Event Stream Processing: This is an optional component that may be deployed for specific purposes. It provides a capability for real-time processing of continuous data streams through.

- Stream Processing (SP): involves the real-time handling of data, where computation occurs directly as data is generated or received. Most data are produced incrementally over time as a sequence of events.
- Complex Event Processing (CEP): is a generalisation of traditional stream processing for aggregating, processing, and analysing data streams in order make high-level inferences about complex events within the business domain using models of causality and conceptual hierarchies. CEP is often used for tasks such as event correlation.

Event Enrichment
This function enriches the events with contextual information from the enrichment API provided by partner Silent Push. Silent Push scans, clusters, scores and enriches the global IPv4 range in a first-party database that outputs Indicators of Future Attack (IOFA) – domain, IP and URL data that explains the relationship between billions of observable data points across the internet.

The Collaboration Mesh Plane

This is the Collaboration Mesh **connectivity underlay** as described previously in the *Architecture Foundations* section. Note that the SM is an optional component and is deployed only for specific scenarios. Also, while a micro-service based architecture (with container RO) is the preferred application software architecture - other approaches such a bare-metal may be used when required.

The Threat awareness plane

The Threat Awareness layer of the ResilMesh platform contains a set of information processing and analysis functions to manage anomaly detection, event correlation and alerting, and attack detection and prediction. The layer functions are implemented as a combination of new development and reuse of existing components.

The Threat Awareness layer includes an **AI-based Anomaly Detection (AID)** module to detect anomalies for any type of IT or OT application and/or network protocol. Models can be placed at the endpoint/edge or in the cloud as required. While ResilMesh supports the use of any ML or AI techniques to develop models, the focus is on the use of deep learning techniques. It will evaluate different multi-view deep learning approaches, such as multi-view fusion-based methods and multi-view alignment-based methods, to deal with the intrinsic heterogeneity of mixed technology domains. ResilMesh will evaluate the use of distributed edge/endpoint-based anomaly detectors

Funded by the European Union

and will consider the use of stacked or hierarchical deep learning techniques for both feature level fusion and/or decision (model) level fusion. This layer also includes a federated learning platform - the **Privacy FL Framework (PPFL)** - to support federated training of the different types of deep learning algorithms mentioned above. The platform will support both horizontal and vertical federated learning, as well as a combination of both. The final output of these detectors will be events sent to the AIC and/or SIEM correlation function, or as features to be passed on to other models.

The **AI Correlator (AIC)** functional component is intended to provide AI methods to correlate security events includes with application to i) classify security events for *event detection ,event grouping ,* and *event pattern extraction ,* ii) *Intrusion detection* which deals with multi-stage and targeted attacks or *anomaly detection* to notify the security administrator about misuses and deviations from normal behaviour, respectively and iii) *Intrusion/attack projection* based on incoming events, which allows early detection of intruder targets. AI correlation may be embedded within other components e.g. such as NSE or NDR or may exist as a stand-alone component.

The **Security Incident and Event Manager (SIEM)** is a central component in Resilmesh and provides a number of capabilities including event logging (through Elasticsearch), event correlation as well as a range of XDR capabilities.

The **THF (TTP-based Hunting and Forensics)** TTP-based threat hunting focuses on identifying adversaries by their tactics, techniques, and procedures—the "how" of their operations rather than the "what" or "when." This approach leverages the ATT&CK framework developed by MITRE. THF supports the use of TTP-based hunting techniques for cyber-attack investigation. THF will have an UI where the user can carry out the mentioned hunting and analysis features.

The Situation Assessment plane

The **Situation Assessment** plane contains a set of functional components that collectively provide cyber situational assessment capability to Resilmesh - and which, together with the Threat Awareness plane implements cyber situational awareness in Resilmesh. These components are:

The **Infrastructure and Service Information Model (ISIM)** captures and represents all the entities of interest in the environment including devices, networks applications (services) users and data. The information model interconnects the pieces of information on the assets in the environment. When deploying the overall system in a new environment there is a need to fill the database with data from external sources. For each category of assets, their enumeration will be collected from existing databases, repositories, or service, or collected via a set of custom tools. The ISIM provides a GraphQL[11] interface to allow applications to query asset status.

The **Cyber Asset Attack Surface Management (CASM)** tool ingests data from the ISIM. The principal role of the CASM is to monitor the organisation's internal and external

[11] https://graphql.org/

Funded by the European Union

attack surface and security posture. It provides an interactive query capability to allow operators to determine cyber security posture based on the relationship between assets - based on ISIM. It also carries out domain (and sub-domain) enumeration to discover so-called 'shadow-IT' and to identify and manages threats discovered in internet-facing assets using independent scans of the organisation attack surface. CASM can detect changes of asset status and trigger required actions or alerts to the Mitigation Manager (MM) as required.

The **Network Detection and Response (NDR)** component continuously monitors and analyses raw enterprise network traffic to establish a baseline of normal behaviour. Any deviations from this baseline are flagged as potentially threatening, prompting alerts for security teams to investigate and respond to potential threats within their environment. NDR ingests data from several sources including ISIM (information about assets) and the SIEM security events and alerts.) and it triggers action invocation in the Mitigation Manager when deviations from expected behaviour are detected.

The **Critical Service Awareness (CSA)** component will provide hierarchical risk assessment to aggregate infrastructure risk into a risk for the critical service or business mission (CS/M). It will implement tools to assess such risks. To achieve this, the component will use data stored in ISIM (information model) and forward the outputs to the NSE (Network Situation Evaluation).

The **Network Situation Evaluation (NSE)** functional component provides a risk assessment of the overall network based on input from other functions and can also project the attack intensity for the network. It provides visualisation of both current and future network risk status. It uses inputs from a number of tools to do this including the ISIM, SIEM, CSA and NDR.

The Security operations plane

The Security Operations Plane contains a set of functional components aiming to decide the most suitable mitigation actions that should be taken to respond to a detected incident, orchestrate these actions as CoA playbooks and analyse collected information to identify adversaries by tactics and techniques carried out during the incident.

Mitigation Manager (MM) Functional Component is responsible for deciding which mitigation actions, if any, are taken to deal with an incident. MM relies on CSS and Network Situation Assessment (NSA) to get information about mission, risk and network status, which are considered factors in the mitigation decision process. Mitigation Manager should define versatile and dynamic-based algorithm(s) to decide which are the best mitigation actions to enforce. Mitigation Manager's decision algorithm could be a rule-based inference engine or even it could be based on AI.

Once mitigation actions have been decided, MM interacts with **Playbooks Tool (PT)** through a REST API to trigger the orchestration of the selected mitigation playbook(s)

Funded by
the European Union

to counter the incident. Then, playbooks Tool (such as Shuffle) launches CoA playbooks that will execute workflows which contain a set of several actions to mitigate incidents, like OpenC2 commands. Some of the mitigation actions could be network filtering, quarantining hosts or updating outdated software. Workflow actions also can enrich SIEM adding new rules (like yara or sigma rules) or execute other tools like Security Orchestrator that can enforce security or privacy policies.

Playbooks Tool should also be user-friendly as it should enable the user to create automated and exportable workflows and design associated PT apps to integrate tools like Security Orchestrator as a workflow action.

If a detected anomaly is confirmed as malicious activity, this is investigated, gathering information to understand the threat and identify potential adversaries. Once the gathered information is sufficient, mitigation actions can be deployed by PT to counter the attack.

To support PT, the **Workflow Orchestrator (WO)** will orchestrate and automate complex actions taken as a CoA playbook step. WO will perform part of the playbook actions (those which are difficult to implement in Shuffle PT) when requested by Playbooks Tool within a CoA playbook execution. WO then uses a subset of tools from Target Actuators to carry out the requested actions. PT can request an action to the WO via a REST/gRPC API, and then Workflow Orchestrator can provide information about the progress and outcome of the performed actions to the Mitigation Engine, establishing a constant feedback loop.

Main workflows

Attack/incident detection flow

There is not a single control flow for attack detection but rather a main flow (SIEM flow) and a number of alternate that may also occur in conjunction with main flow. These are now described.

Funded by
the European Union

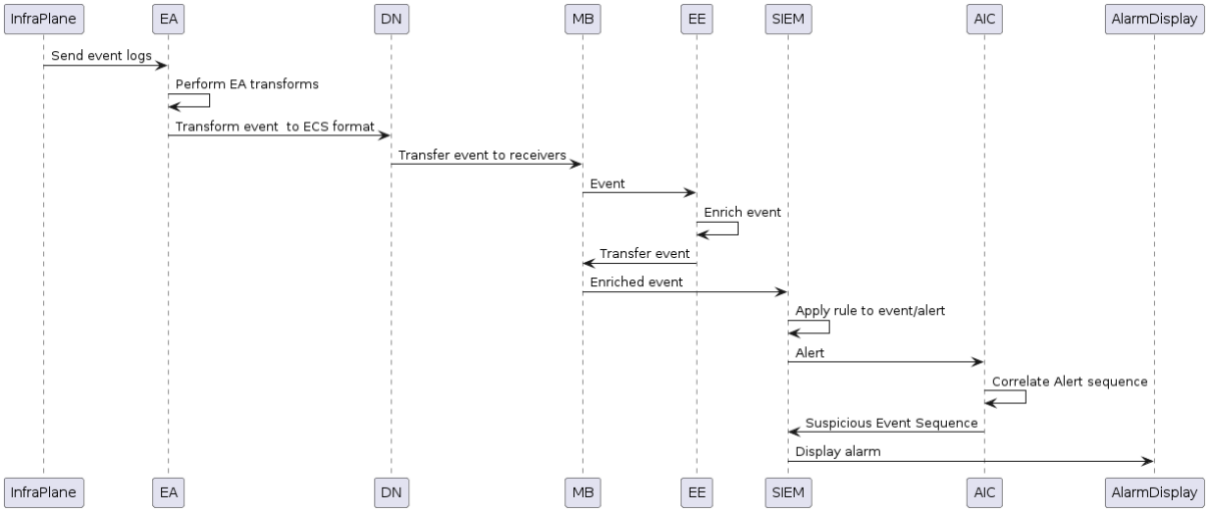

Figure 3 - SIEM Main Flow

1. Events are emitted from endpoint logs and/or security sensors such as Intrusion Detection Systems (IDS) , firewall logs etc. These are collected by the EA function where various data aggregation and transformation are performed on them. The events are then normalised to a common format by the DN function. Events are then passed to MB,
2. The MB is a publish/subscribe and pushes the events forward to subscribing functional components. These components may in turn perform some processing and then return the processed event to the MB for distribution to the next stage.
3. Events are enriched by the EE functional component which calculates several risk metrics based on screening IP addresses or domain name within the event.
4. The event is then usually forwarded to the SIEM for storage and checking for evidence of anomalous or malicious activity. This checking is done by a rule based detection system within the Wazuh manager. Events from endpoint logs are matched against various rules for this purpose and an alert is generated if anything suspicious is noted. Alerts from endpoints or NIDS e.g. Suricata are managed in the same way i.e. these alerts are match against a rule
5. Alerts are given a priority score and forwarded to the alarm manager dashboard (in Kibana). This dashboard is the central point in the SOC for triaging all alerts.
6. The AIC may optionally be deployed in an alternate flow after step 4 i.e. instead of invoking step 5, events may be forwarded to the AIC[12] for further correlation and analysis. The AIC is used to examine and detect suspicious event sequences which then raise an alarm for the analyst. Non suspicious event sequences are suppressed and not presented to the analyst.

In addition to the main flow there are a number of supplementary flows which interwork with it. We begin with the **AID flow**

[12] Note the AIC may have different uses - see the functional component description - here it is used for reducing the number of events shown to the SOC operator;

Funded by
the European Union

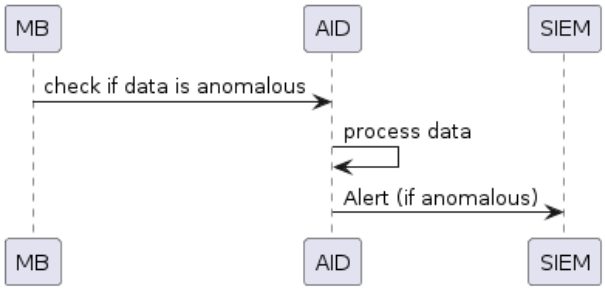

Figure 4 -AI-based Anomaly Detection

1. In this flow data is transferred to the (after step 3 in the main flow) AID to detect suspicious or anomalous behaviour. This data may be events (e.g. Windows events) or it may be pre-processed data for the detector - for example network traffic flow data may be pre-processed to extract traffic feature which are then sent to the AID for analysis. This prep-processing takes place in the infrastructure plane and is not shown here. Also the AID functional component may be centrally located (as depicted in the flow above) or it may be placed at the edge. Moreover, the AID may occur as a standalone device e.g. NIDS or it may be embedded in another functional component such as NDR.
2. The AID will analyse the received data, then generate an alert if it detects a suspicious activity and will forward this alert to the SIEM as shown.

Another supplementary flow is the **ESP flow.** In this flow events are forwarded after step 3 in the main flow to the ESP. Subsequent steps depend on the particular ESP service invoked.

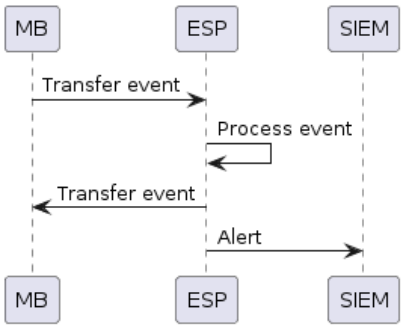

Figure 5 - ESP Flow

1. If the ESP Stream Processing service is used then low level processing is applied to the data stream e.g. aggregating a number of events before storing them. In this case the next step in the flow is to route the processed data back to the MB for distribution to the next stage.
2. If the ESP Complex Event Processing is used the a high level correlation of events takes place with the usual outcome to be generation of alert which is then transferred to the SIEM.

Cyber Threat Intelligence Sharing is another supplementary flow.

Funded by
the European Union

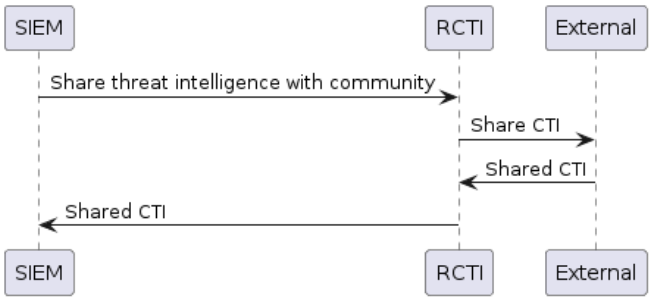

Figure 6 - CTI Flow

1. The SIEM (or other functions) may share threat intelligence with other partners.
2. The system may also receive information from other partner

Threat hunting is another supplementary flow that , while not normally part of the attack detection flow, may be used to support detection.

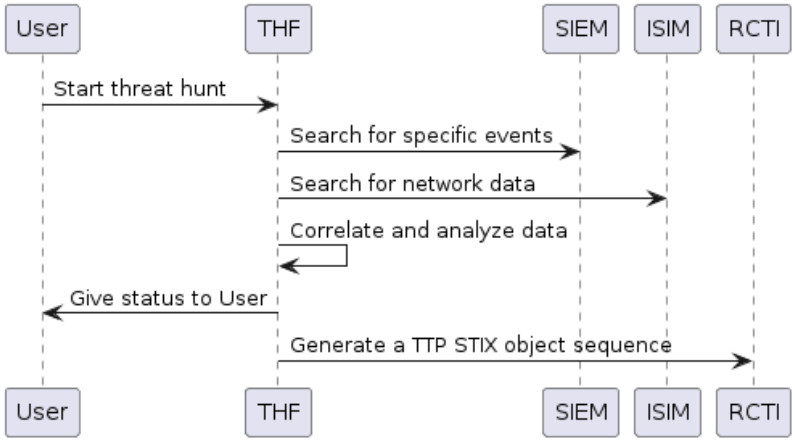

Figure 7 - Threat hunting Flow

1. The user (threat hunter) initiates a threat hint based on a specific TTP based hypothesis.
2. The THF fetches network and TTP data based on the provided hypothesis.
3. Depending on the specific hypothesis and user request different analyses may be carried out .
4. Hunt progress and status is displayed to the User.
5. Steps 1-4 are iterative and may be repeated a number of times.
6. In some cases the output may be shared as threat intelligence as STIX behaviour sequence object.

Funded by
the European Union

Situation Assessment flow

The SA flow also has a number of alternate sub flows detailed below.

ISIM Sub flow

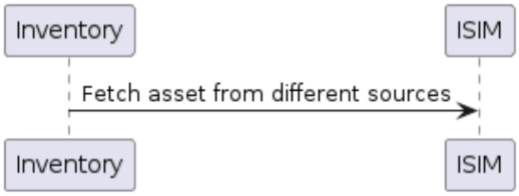

Figure 8 - ISIM Sub flow, Uploading Asset Information to ISIM

The ISIM stores information on network assets and is updated form a variety of different sources.

NDR Sub flow

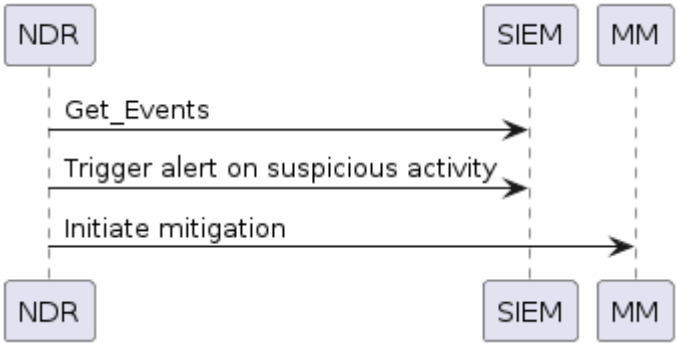

Figure 9 - NDR alerting flow

NDR collects information from network sources (via the SIEM) and checks for anomalous behaviours. NDR contains embedded analysis components to carry out such checks , including possibly the AID. If it detects suspicious activity it generates an alert and forwards it to the SIEM for triage. It also initiates mitigation actions if required.

Cyber ASM Sub flow

CASM monitors both internal and external assets attack surfaces.

1. It scans the internet to discover ‘shadow IT’ and also checks the attack surface using a variety of network sources e.g. such as Shodan.
2. It updates the ISIM with discovered information.
3. It also performs queries against the ISIM to discover potential deviations from the desired security posture.

Funded by
the European Union

4. It may initiate one or more actions to the XDR towards the endpoint e.g. it may initiate a vulnerability scan.
5. It notifies user of discovered deviations and may also initiate mitigation actions in some cases

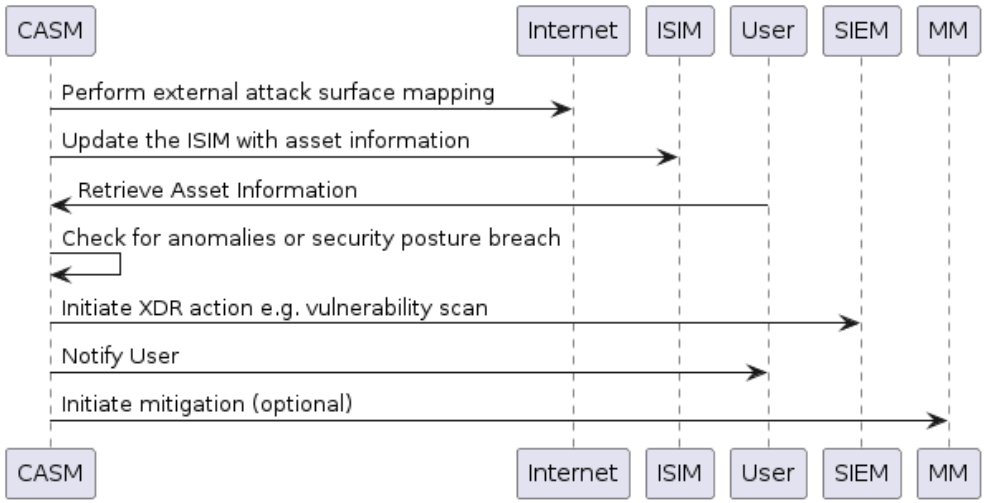

Figure 10 - CASM flow

NSE flow

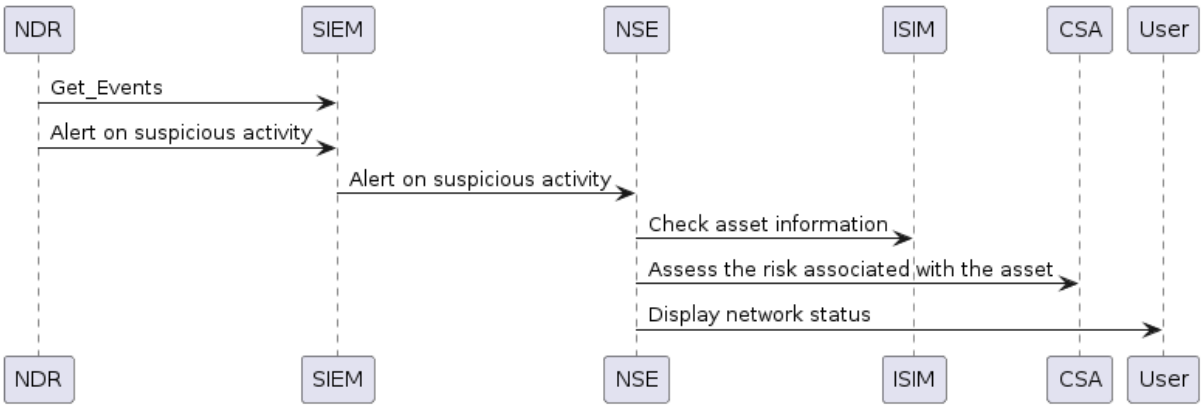

Figure 11 - NSE Flow

The NSE maintains an overview of network risk status.

1. It receives notification of suspicious events from the SIEM and fetches asset information from the ISIM.
2. It checks the asset criticality with the CSA and calculates associates risk if any .
3. Finally it updates the network status.

Funded by
the European Union

Reactive/mitigation flow

Reactive / mitigation enforcement workflow consists of the set of steps and actions performed by a set of functional components to respond to a detected threat by enforcing CoA playbooks.

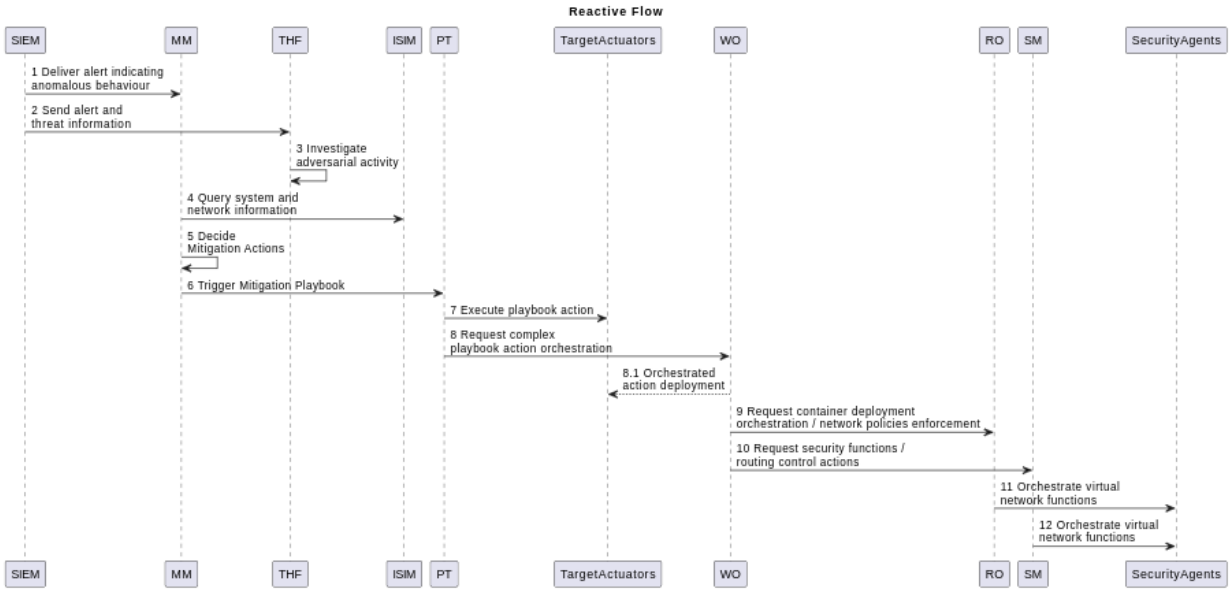

Figure 12 - Reactive flow

1. Firstly, SIEM collects logs from the monitored endpoints to analyse systems' behaviour. SIEM contains a predefined rule set and additional rules created by administrators. When some of these rules are triggered by the information contained in the received logs, SIEM will fire an alert and deliver it to the MM to start a potential threat mitigation process.
2. SIEM will also send the alert and threat information to THF, which based on the alert, gathers contextual information to fully understand the threat and tries to identify the adversaries which are carrying out these activities.

3. Then, THF will investigate threat details and it will provide capabilities for generating reports and visualising data trends. Using its correlation function, THF will link related activities across time and terrain to construct a coherent narrative of potential adversarial actions. All of these investigation and hunting processes will feed the hunting database with the fetched data and the outcomes of the correlation analysis and will also support further investigations, trend analysis, and threat intelligence enrichment.

4. Before starting the MM decision process, it needs to obtain situational information about network and system status. To make this, MM will query ISIM to fetch this information. ISIM is made of an information model, using the CRUSOE data mode. The information model is to be materialised as a database that will serve as a data

Funded by
the European Union

repository for other components. CRUSOE will also carry out threat risk assessment, whose output is needed by MM to decide mitigation actions.

5. MM will have to decide the most suitable mitigation actions to counter the current threat. The decided mitigation playbook will be the result of its decision algorithm execution. This takes into account the dynamic status of system and network (via queried ISIM information), threat risk assessment and mitigation execution cost.

6. Once mitigation playbook selection is finished, MM triggers its execution by requesting it to PT. MM indicates the CoA playbook to execute to the PT.

7. PT will start the execution of the requested mitigation playbook. Most of the playbook actions will lead PT into deploying actions in some target actuators. Among the options, this can be achieved by creating and sending OpenC2 commands to the actuators.

8. During PT playbook execution, some of the actions could not be suitable to be performed directly from PT because of their complexity. To handle this situation, PT will contact WO for requesting the orchestration of the playbook complex actions. Then, as shown in 8.1, WO will deploy the orchestrated actions into the proper target actuators.

9. Playbook mitigation actions can also provoke the orchestration of virtual network functions to ensure that the needed security actions are taken. WO will request container deployment orchestration and/or network policies enforcement to RO that will create a virtual network function (security agent) to accomplish the requested security actions.

10. In a similar fashion, WO can request security functions or routing control actions to SM which also creates virtual network functions to orchestrate these actions.

11. RO can manage container lifecycle, including deployment, scaling, updating and terminating. Moreover, it defines and enforces network policies to control the communications among containers within the cluster or externally. RO supports integrating with SM to facilitate advanced network management.

12. By using SM, the WO can invoke security functions including end-to-end encryption (mutual TLS), authentication and authorization policies. Furthermore, WO is able to orchestrate routing control actions like configuring traffic shifting and mirroring.

This section has defined the main high level functional architecture, including their planes and functional components. Furthermore, the section has delved into and the main associated interactions between functional components and general workflows expected in Resilmesh to fulfil the project requirements.

Funded by
the European Union

4 Functional component descriptions

This section aims to detail the functional components that have been introduced previously in section 3. Each functional component is described using the same template that firstly defines a description of the Function associated to the functional component. Then, each functional component describes the main services that it will feature to implement the function in the framework. Each service defines the capabilities it offers, the type of service, who are the consumers services, pre and post conditions as well as the main envisioned interfaces.

Resource Orchestration

Function

The Resource Orchestrator manages containerised workloads and services through declarative configurations. It dynamically manages the lifecycle of containers, including deployment, scaling, and networking, and self-healing ensuring high availability and reliability of services. It automatically allocates and manages computing resources like CPU, memory, and storage across a cluster based on resource usage and demand, optimising performance and reducing costs. The Resource Orchestrator uses declarative configurations to streamline the deployment processes and manages the secrets, allowing applications to adapt to different environments without code changes.

Provided services;

Orchestration Services

A. Description

Cluster Management

- The Resource Orchestrator has high flexibility to scale up or down automatically based on the workload requirements to optimise the usage of resources. It also maintains cluster availability over updates and failures with minimal downtime.

Container Orchestration

- The Resource Orchestrator can manage container lifecycle, including deployment, scaling, updating and terminating. It also continuously detects containers´ health and makes them recover from the failures, such as restarting or rescheduling failed containers on other working nodes. The Resource Orchestrator has integrated service discovery and load balancing abilities to allow containers to communicate with each other and distribute traffic efficiently. The Resource Orchestrator supports automated provisioning of resources, including ephemeral volumes and persistent storages.

Network Management

Funded by
the European Union

- Resource Orchestrator defines and enforces network policies to control the communications among containers within the cluster or externally. It also provides network isolation between different applications to enhance security and reduce interference. The Resource Orchestrator supports integrating with service mesh to facilitate advanced network management.

Security

- The Resource Orchestrator provides robust mechanisms like access control policies to restrict access to the orchestration platform, only authorized users and services can perform operations. The Resource Orchestrator can securely store and manage secrets of containers without exposing credentials in configurations or codes.

B. Capabilities

Enables deployment of Resilmesh functions in a flexible and dynamic manner.

C. Type

External.

D. Consumers

- SOAPA layers

E. Pre-conditions to consume the service

-

F. Interfaces

We do not give a detailed interface breakdown but instead point to the Kubernetes and Docker Swarm documentation.

nRO_services	**Detailed Description**	These interfaces enable the applications to be orchestrated as required for different deployments.
	From provider	RO
	To Consumer	RO
	Technology	Rest API
	API Documentation	https://kubernetes.io/docs/concepts/cluster-administration/ https://docs.docker.com/engine/swarm/
	Partners involved	TUS,

Service Mesh

Function

Service Mesh Platform is a dedicated infrastructure layer that enables, secures and monitors service-to-service communications within distributed applications. The

Funded by the European Union

Service Mesh Platform supports automatic service discovery within the network, so that services can identify and locate each other in a highly dynamic environment where services frequently change due to deployments, scaling, and failures. It has traffic management functionalities, including load balancing that distributes network traffic across multiple backend services to ensure reliable and efficient data handling, and fine-grained traffic control such as detailed telemetry, encryption settings, and specific routing rules. The Service Mesh Platform enhances the security posture of the overall application infrastructure through fine-grained access control policies, transparent TLS encryption, network segmentation, and Authentication, Authorization and Audit (AAA) tools. It offers detailed insights into the behaviour of services, including monitoring, logging, and tracing capabilities to help diagnose and resolve issues quickly.

Provided services;

Mesh Services

A. Description

A service mesh is an infrastructure layer designed for managing interactions between services/microservices. It helps your microservices run smoothly, securely, and stable (while telling you what is going on with them. It handles things such as discovery, load balancing, failure recovery, metrics, monitoring, rate limiting, access control, and authentication.

The following services, amongst others, are provided by the platform:

- Service Discovery
- Security
- Observability
- Routing Control

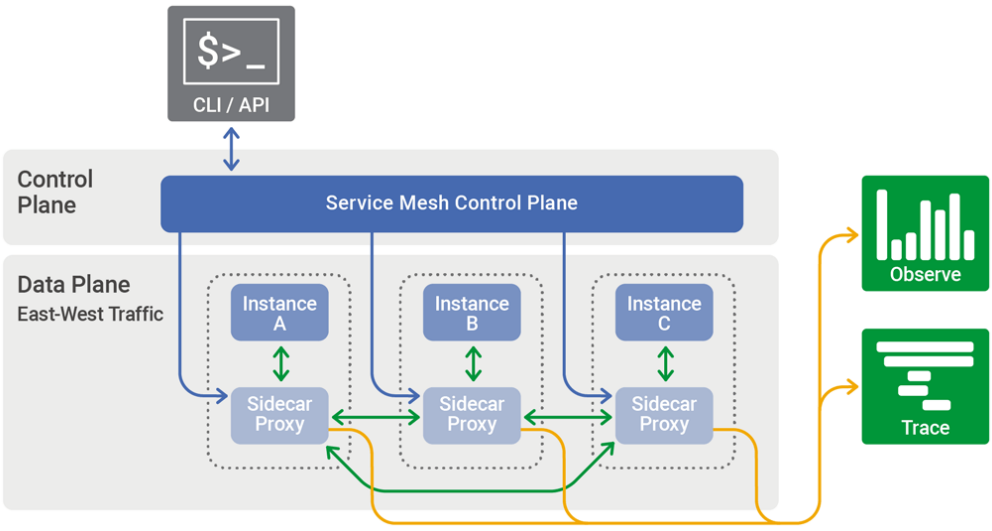

Figure 13 -Service Mesh

Image Source: https://www.nginx.com/blog/what-is-a-service-mesh/

B. Capabilities

Funded by the European Union

Provides resilience through simplifying observability, traffic, security, and policy management

C. Type

External.

D. Consumers

- SOAPA layers

E. Pre-conditions to consume the service

A service mesh capability such as Istio or NATS must be provided

F. Interfaces

These interfaces are generic but reference is given to the NATS implementation for illustration purposes. Alternative implementations include Istio, Linkerd etc.

nSM-Discovery	**Detailed Description**	This interface allows the Resilmesh functions to discover where the other services are via a service registry
	From provider	SM
	To Consumer	Resilmesh Functions
	Technology	Rest API
	API Documentation	https://nats.io/tags/service-mesh/ https://nats.io/blog/nats-to-implement-servicemesh-part1-service-discovery/
	Partners involved	TUS,

nSM_Security	**Detailed Description**	This interface allows the Resilmesh functions invoke security functions including end-to-end encryption (mutual TLS), authentication, authorization policies as well as service-to-service access control among the services.
	From provider	Service Mesh
	To Consumer	Resilmesh Functions
	Technology	REST API
	API Documentation	https://nats.io/tags/service-mesh/ https://nats.io/blog/nats-to-implement-service-mesh-functionality-part2-security/
	Partners involved	TUS

Funded by the European Union

nSM_Observability	**Detailed Description**	This interface allows the Resilmesh to invoke observability such as metrics, tracing, and alerting functions
	From provider	Service Mesh
	To Consumer	Resilmesh Monitor
	Technology	REST API
	API Documentation	https://nats.io/tags/service-mesh/ https://dale-bingham-soteriasoftware.medium.com/using-nats-to-implement-service-mesh-functionality-part-3-metrics-tracing-alert-observability-f77cf5ab7db1
	Partners involved	TUS

nSM_Routing Control	**Detailed Description**	This interface allows the mesh to configure traffic shifting and mirroring
	From provider	Service Mesh
	To Consumer	Resilmesh Functions
	Technology	REST API and/or Fabric
	API Documentation	https://nats.io/tags/service-mesh/
	Partners involved	TUS

nSM_Loadbalancing	**Detailed Description**	This interface allows the service mesh to perform load balancing easily when 2 or more services are setup as replicas/copies
	From provider	Service Mesh
	To Consumer	Resilmesh Functions
	Technology	API REST
	API Documentation	https://nats.io/tags/service-mesh/
	Partners involved	TUS,

Funded by
the European Union

Event Aggregation

Function

This functional component ingests event logs from multiple sources and transforms the events via various operations (routing / logging / fusion / filtering / augmentation / reduction / monitoring) and then outputs the data to one or more destinations via the streaming layer. Aggregator transformation components may be linked in data pipelines giving rise to arbitrarily complex processing topologies. An example of such a pipeline is shown below for the *Vector.dev* aggregator.

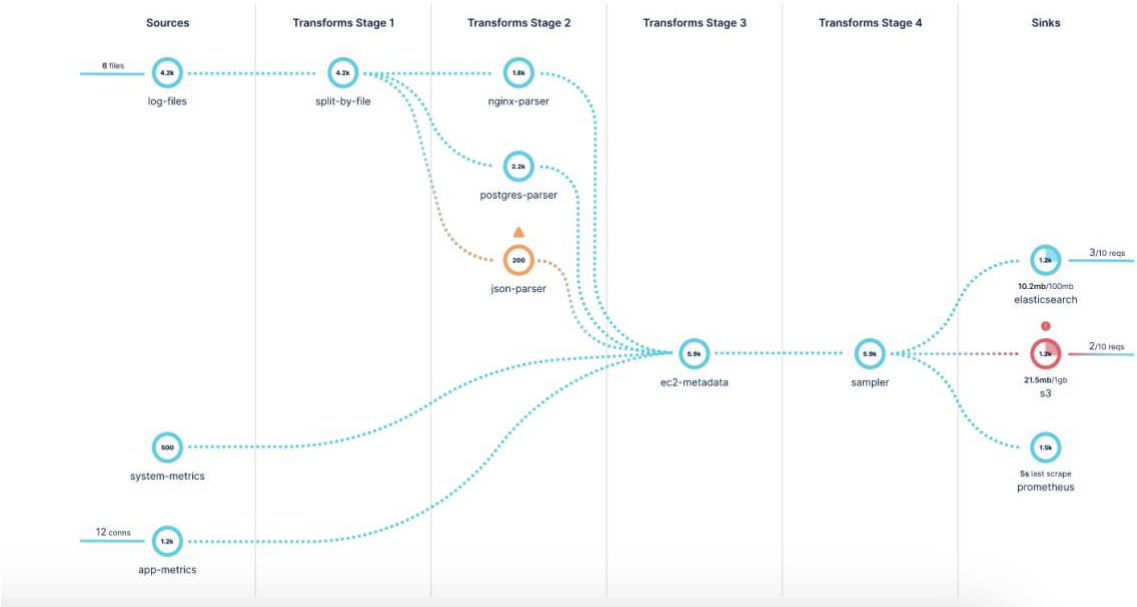

Figure 14 -Orchestration[13]

Provided services

The EA has three services as described below.

Event Aggregation

A. Description

[13] https://vector.dev/docs/about/under-the-hood/architecture/pipeline-model/

Funded by
the European Union

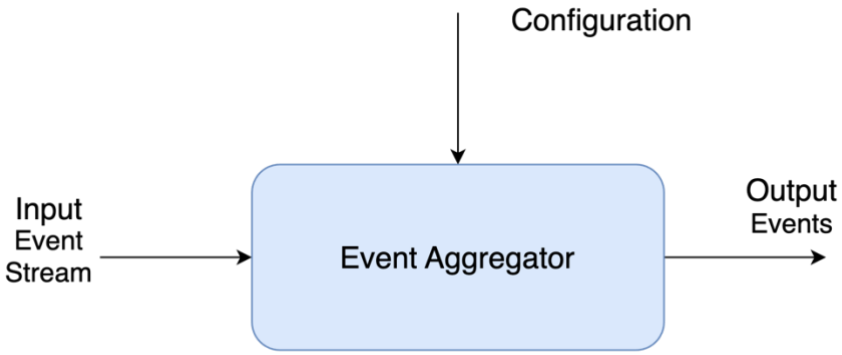

Figure 15 -Event Aggregation

B. Capabilities

Create

C. Type

Internal / External.

D. Consumers

- Decision engine

E. Pre-conditions to consume the service

Security(s) policies are defined.

F. Interfaces

nEA_Input	**Detailed Description**	This interface allows ingestion of events to the EA
	From provider	Log collection agents
	To Consumer	EA
	Technology	Rest API
	API Documentation	NA
	Partners involved	SLP

nEA_Output	**Detailed Description**	This interface allows output of events to the CEP
	From provider	EA
	To Consumer	DN. Other EA
	Technology	Rest API
	API Documentation	NA

Funded by
the European Union

	Partners involved	SLP

nEA_Configure	**Detailed Description**	This interface allows definition of EA pipeline configurations and addition of own transform /
	From provider	EA
	To Consumer	Resilmesh Applications
	Technology	Rest API or File
	API Documentation	NA
	Partners involved	SLP

Data Normalisation

Functions

Data Normalisation is the processing step used to map the events to a common format schema, it's the transformation of heterogeneous data coming from different sources into a single unified schema, so that they can be properly consumed by the different components of the framework, since they will know what to expect from the consumed data. For such, we will use ECS, which has several fields that can be used for the majority of use cases.

Provided services

- **Standardisation**
  The events across the framework's pipeline will have a standard format schema

Capabilities
Common event schema
Type
External
Consumers
Threat Awareness Plane
Pre-conditions to consume the service
-
Interfaces

nDN_services	**Detailed Description**	this interface provides the standardisation of the events

Funded by
the European Union

	From provider	DN
	To Consumer	DN
	Technology	ECS
	API Documentation	https://www.elastic.co/guide/en/ecs/current/index.html
	Partners involved	SLP, GMV, TUS, UMU

Message Broker (MB)

Functions

Guarantee the correct flow of the events being processed by the Resilmesh Framework by queueing them according to their state (normalised, enriched, etc). Essentially, every component of the framework will rely on the Message Broker for delivering their processed events, so that other components will do the same and so on, until the whole pipeline is finished and the event is ready for the Analytics Layer.

Provided services

Load Balancing

The Message Broker has the capability of reliably queueing events in such a manner that they can be consumed by any number of components. If the component is replicated, as in a Docker Compose or Kubernetes architecture, the events will be load balanced/distributed among these replicas.

Horizontal Scalability

The Message Broker is an essential piece of technology that aligns with the capability of horizontally scaling the components, since the events will be distributed among the consumers/subscribers, if the load of the framework increases, we increase the components replicas.

Capabilities

Leverages the resilience of the framework

Type

External

Consumers

Threat Awareness Plane

Pre-conditions to consume the service

NATS Clients, see: https://nats.io/download/#clients

Interfaces

Funded by
the European Union

nMB_services	**Detailed Description**	NATS.io implements the client-server paradigm, any client can connect to the server by passing the connection string in the format "nats://<IP>:<PORT>"
	From provider	any
	To Consumer	any
	Technology	NATS protocol over TCP
	API Documentation	https://docs.nats.io/
	Partners involved	SLP, GMV, TUS, UMU

Event Stream Processing (ESP)

Function

This component provides a capability for real-time processing of continuous data streams.

- *Stream processing* (SP) involves the real-time handling of data, where computation occurs directly as data is generated or received. Most data is produced incrementally over time as a sequence of events. In stream processing, applications maintain a constant presence for executing logic, performing analytics, and running queries, with data continuously passing through them. When an event is received from the stream, a stream processing application responds accordingly, potentially initiating an action, modifying an aggregate or statistic, or storing the event for future use.
- *Complex event processing (CEP)* is a generalisation of traditional stream processing for aggregating, processing, and analysing data streams in order to gain real-time insights from events as they occur. However whereas traditional stream processing is concerned with finding low-level patterns in data, such as the number of mouse clicks within a fifteen-minute window CEP can make high-level inferences about complex events within the business domain using models of causality and conceptual hierarchies. CEP is therefore suitable for tasks such as event correlation.

Provided services

Complex Event Processing

A. Description

This function aggregates a lot of different information that identifies and analyses cause-and-effect relationships among events in real time by querying data before storing it within a database or, in some cases, without it ever being stored. Events can be received from different sources.

Funded by the European Union

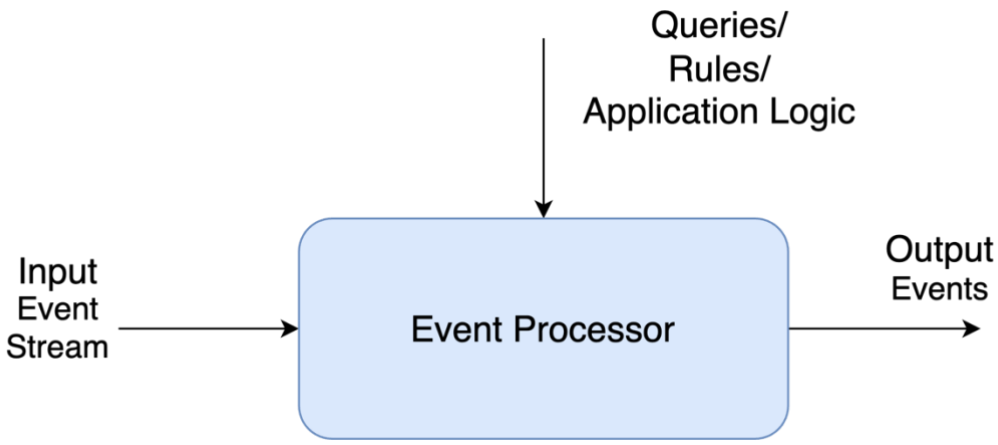

Figure 16 - Complex Event Processing

B. Capabilities

Correlation of Alerts to detect unusual patterns or events

C. Type

External

D. Consumers

Resilmesh applications

E. Preconditions to consume the service

F. Interfaces

nESP_CEPInput	**Detailed Description**	This interface allows ingestion of events to the CEP
	From provider	CEP
	To Consumer	Resilmesh Applications e.g. Elasticsearch and others
	Technology	Rest API
	API Documentation	NA
	Partners involved	TUS,

nESP_CEPOutput	**Detailed Description**	This interface allows output of events to the CEP
	From provider	CEP
	To Consumer	Resilmesh Applications e.g. Elasticsearch and others
	Technology	Rest API

Funded by the European Union

	API Documentation	NA
	Partners involved	TUS

nESP_CEPConfigure	**Detailed Description**	This interface allows definition of CEP applications processing logic i,e, queries and rules
	From provider	CEP
	To Consumer	Resilmesh Applications
	Technology	Rest API
	API Documentation	NA
	Partners involved	TUS

-

Stream Processing

A. Description

The stream processing services enables the creation of individual 'stream processors' to operate on events in the stream and also allows the creation of a stream processor topology to carry out composite stream processing on the events. An example is Kafka streams that are built on top of the Kafka messaging broker. Events are ingested and output via the message broker API's.

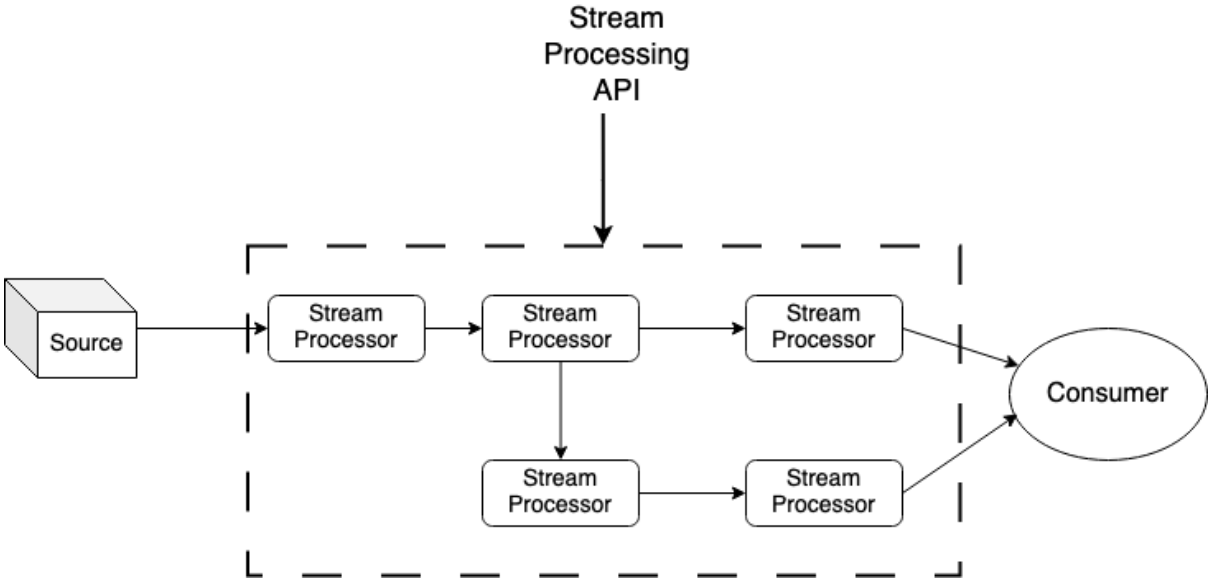

Figure 17 - Stream processing

B. Capabilities

Creates and deploys event stream processing applications

C. Type

Funded by
the European Union

Internal / External.

D. Consumers

- Decision engine

E. Pre-conditions to consume the service

Security(s) policies are defined.

F. Interfaces

nESP_Streamprocessc	**Detailed Description**	This interface allows the creation and deployment of stream processing applications
	From provider	Stream Processor
	To Consumer	Resilmesh Applications e.g. Elasticsearch and others
	Technology	Rest API
	API Documentation	https://kafka.apache.org/20/documentation/streams/developer-guide/dsl-api.html
	Partners involved	SLP,

Event Enrichment (EE)

Function

This function will enrich the events with contextual information from the Silent Push API. Silent Push scans, clusters, scores and enriches the global IPv4 range in a first-party database that outputs Indicators Of Future Attack (IOFA) – domain, IP and URL data that explains the relationship between billions of observable data points across the internet. The enriched data will make events more relevant through contextual information, depending on the type of the event (IP, Domain, etc) such as: whois, DNS records, ASN, Name Server, Certificates, Subnet information etc

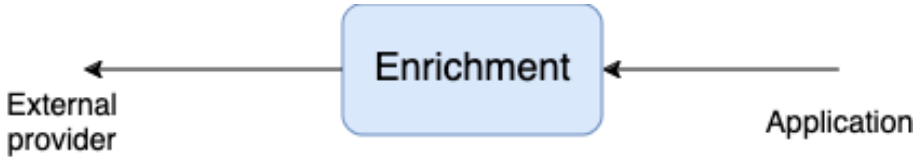

Figure 18 - Event Enrichment

Provided services

Event Enrichment

A. Description

This service enriches the incoming events with a callout to the Silent Push API.

Funded by the European Union

B. Capabilities
enriches events to indicates possible suspicious IP/ASN/URL.
C. Type
Internal / External.
D. Consumers
- Various applications

E. Pre-conditions to consume the service
License agreements defined with SLP
F. Interfaces

nSLP-Enrich	**Detailed Description**	This interface from SLP enriches events with various risk and reputation scores to indicate possibly risky indicators of attack
	From provider	SLP
	To Consumer	Resilmesh Applications
	Technology	REST API
	API Documentation	https://docs.silentpush.com/
	Partners involved	SLP

nRCTI-Enrich	**Detailed Description**	This internal interface allows Resilmesh applications to call out enrich events.
	From provider	SLP
	To Consumer	Resilmesh Applications
	Technology	REST API
	API Documentation	NA
	Partners involved	SLP

Security Incident and Event Manager (SIEM)

Function

Security Information and Event Management (SIEM) operates within the realm of computer security, integrating software solutions and services that merge security information management with security event management. SIEM serves as the central element in a standard Security Operations Centre (SOC), which serves as the focal point for addressing security concerns across an organisation. It conducts immediate analysis of security alerts generated by both applications and network hardware. SIEM solutions are offered by vendors as software packages, appliances,

Funded by the European Union

or managed services, serving the dual purpose of logging security information and generating compliance reports.

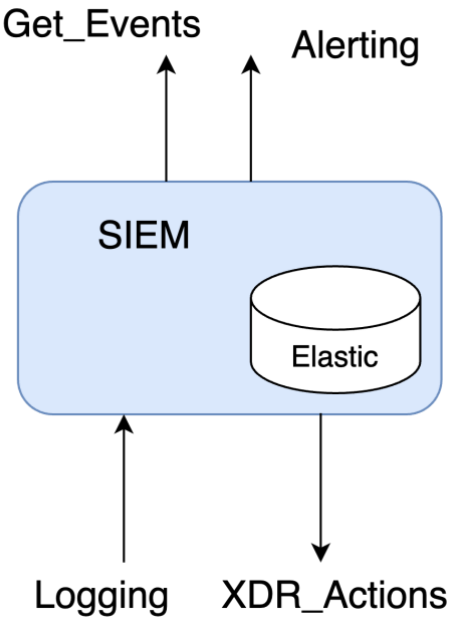

Figure 19 - SIEM Functional Architecture

The Resilmesh SIEM is based the open source SIEM/XDR application, Wazuh[14],and also uses **ElasticSearch**[15] as the event store

Provided services

Event Logging

A. **Description**
   Wazuh collects, analyses, and stores logs from endpoints, network devices, and applications. The Wazuh agent, running on a monitored endpoint collects and forwards system and application logs to the Wazuh server for analysis.
B. **Capabilities**
   Security log analysis; Vulnerability Detection
C. **Type**
   Internal
D. **Consumers**
   Analytics applications
E. **Preconditions for service**
   Wazuh installed
F. **Interfaces**

[14] https://wazuh.com/platform/overview/
[15] https://github.com/elastic/elasticsearch

Funded by
the European Union

nSIEM-LogCollection	**Detailed Description**	This interface ingests logs from the endpoints and stores on Wazuh or r Elastic
	From provider	End Points
	To Consumer	SIEM
	Technology	Rest API
	API Documentation	https://documentation.wazuh.com/current/getting-started/use-cases/log-analysis.html#log-data-collection
	Partners involved	TUS, SLP,

Event Correlation and Alerting

A. Description

Wazuh ruleset detects security events and anomalies in log data. These rules are written in a specific format and they trigger alerts when certain conditions are met. The rules are defined based on certain criteria like log fields, values, or patterns to match specific log entries that may indicate security threats. Wazuh provides a wide range of pre-built rules covering common security use cases. Additionally, administrators can create custom rules tailored to their specific environment and security requirements.

B. Capabilities

Intrusion detection

C. Type

Internal

D. Consumers

Alarm Dashboard (Kibana) ; Analytics functions; ; Mitigation Manager

E. Preconditions for service

Wazuh installed

F. Interfaces

nSIEM-Alerting	**Detailed Description**	This interface output alerts that indicate suspect or anomalous behaviour.
	From provider	SIEM
	To Consumer	Mitigation Manager, Alarm presentation in Elastic, Analytics functions
	Technology	Rest API
	API Documentation	https://documentation.wazuh.com/current/getting-started/use-cases/log-analysis.html#rules-and-decoders

Funded by
the European Union

	Partners involved	TUS, SLP,

XDR Actions

A. Description

Wazuh also provides an Extended Detection and Response (XDR) platform with a comprehensive security solution that detects, analyses, and responds to threats across multiple IT infrastructure layers. Wazuh collects telemetry from endpoints, network devices, cloud workloads, third-party APIs, and other sources for unified security monitoring and protection.

B. Capabilities

Threat Hunting; File Integrity Monitoring; Behavioural Analysis; Endpoint mitigation actions.

C. Type

Internal / External.

D. Consumers

CASM;

E. Pre-conditions to consume the service

Wazuh installed.

F. Interfaces

nSIEM_XDR_Action	**Detailed Description**	This interface allows the WAZUH XDR manager to invoke detection and response on the endpoint.
	From provider	SIEM
	To Consumer	EDR agent
	Technology	REST API, or Fabric
	API Documentation	https://wazuh.com/platform/xdr/
	Partners involved	TUS, SLP

Get_Events

A. Description

Elasticsearch is a distributed search and analytics engine optimised for speed and relevance on production-scale workloads.

B. Capabilities

Event storage and retrieval

C. Type

Internal

D. Consumers

CASM;NDR TFIR, NSE

E. Pre-conditions to consume the service

Funded by the European Union

Elastic installed.

F. Interfaces

nSIEM_Get_Events	**Detailed Description**	This interface allows the applications such as NDr and NSE to retrieve events from SIEM logs
	From provider	SIEM
	To Consumer	CASM NDR TFIR, NSE
	Technology	REST API, or Fabric
	API Documentation	https://github.com/elastic/elasticsearch
	Partners involved	TUS, All

AI-based detector (AID)

Function

The AI-based detector (AID) functional block is a system component for anomaly detection across a spectrum of IT and OT applications and network protocols. Leveraging machine learning (ML) and artificial intelligence (AI) techniques, the system offers anomaly detection models that can be deployed on endpoints, edges, or in the cloud based on specific requirements.

This component will include multi-view deep learning approaches, including fusion-based and alignment-based methods, to effectively handle the inherent heterogeneity present in mixed technology domains. For instance, a multi-view anomaly detector within ResilMesh may combine data from disparate sources such as an OT PLC (view), Modbus network data (view), and IT IDS data (view) to detect potential cross-node attacks.

The implementation extends to distributed, edge/endpoint-based anomaly detectors. The output generated by these detectors comprises events forwarded to the correlator/SIEM (Security Information and Event Management) or features transmitted to other models, ensuring a comprehensive and integrated approach to anomaly detection within diverse technological environments.

It has the following main features:

- Detect suspicious events and attacks at edge and potentially cloud
- Support multi-view anomaly detection for blending heterogeneous data sources, crucial for mixed IT/OT critical infrastructures
- Implement hierarchical feature fusion (e.g., Autoencoders) and decision fusion (e.g., ensemble learning) models for edge-based anomaly detection

Funded by the European Union

Provided services

AI based anomaly detection in OT environment

A. Description

This service provides AI based anomaly detection.

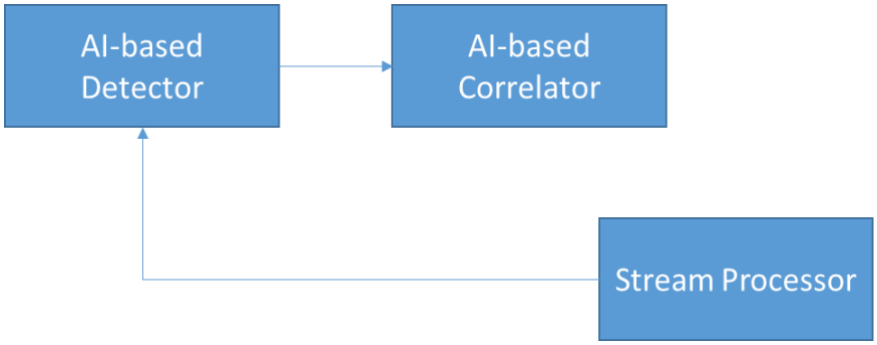

Figure 20 -AID Architecture

B. Capabilities

Anomaly detection, multi-view data fusion.

C. Type

Internal.

D. Consumers

- AIC
- Applications

E. Pre-conditions to consume the service

Pre-processed and streamed features.

F. Interfaces

nAID_IngestFeatures	**Detailed Description**	This interface ingests pre-processed and streamed features.
	From provider	Event stream processor (ESP)
	To Consumer	AID
	Technology	REST API or Method call
	API Documentation	NA
	Partners involved	JR, MONT, UMU

Funded by the European Union

nAID_DetectionEvent	**Detailed Description**	This interface outputs the result of the anomaly detection process in the form of security events and alerts.
	From provider	AID
	To Consumer	AIC, Application
	Technology	REST API, or Method call
	API Documentation	NA
	Partners involved	JR, TUS, MONT, UMU, Application Developer

Privacy preserving model training (PPFL)

Function

The privacy preserving model training is responsible for training the AI-based anomaly detection model developed in T4.1 using Federated Learning along with privacy-preserving algorithms based on data perturbation mechanisms and Privacy-Enhancing Technologies (PETs) to prevent membership inference and model poisoning attacks during the federated training process.

It interfaces with the Message Broker and with the SIEM to receive the information needed for training and will make available to the AI-based Detector the final model obtained at the end of the final round to be used for real-time evaluation.

For the Federated Learning training setting, several aggregation algorithms (e.g., FedAvg, FedProx) will be analysed and evaluated, as well as different privacy-preserving mechanisms (Differential Privacy based on noise-adding mechanisms such as Laplace or Gaussian-based), so that the trade-off between accuracy and privacy is maximised.

The actors of the FL scenario, i.e., aggregators and agents, will be served as Virtual Network Functions (VNFs) so that they can be orchestrated (deployed and configured dynamically) based on policies, where specific configuration parameters, such as the number of training rounds to be executed or the aggregation algorithm to be used, can be specified.

Provided services

Privacy-preserving training

A. Description

The Federated Learning process workflow that would be used to train the anomaly detection model can be consulted in the following figure. As can be seen, a fixed number of training rounds are executed. During each round, the agents will train their

Funded by the European Union

local models upon local training data, apply a certain noise-adding mechanism, and share the resulting protected model updates with a central aggregator that will mix them using a certain aggregation function (e.g., FedAvg). At the end of the final round, the final model is shared with the AI-based Detector so it can be used for real-time anomaly detection.

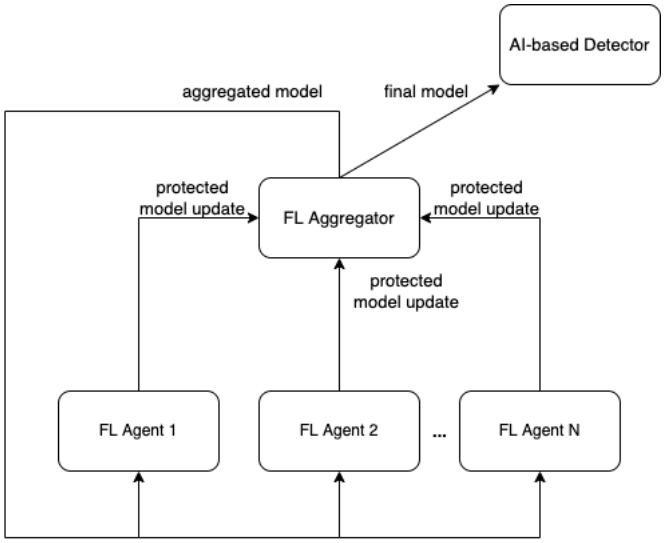

Figure 21 - Privacy Preserving Federated Learning

B. Capabilities

Performs federated learning based on the deployed agents

C. Type

Internal / External.

D. Consumers

- AI-based Detector

E. Pre-conditions to consume the service

Provide training data for each agent in the Federated Learning scheme.

F. Interfaces

AI-based-Detector-interface-1	**Detailed Description**	This interface allows interacting with the AI-based detector to receive the model architecture to be trained
	From provider	AI-based Detector
	To Consumer	Privacy-preserving model training
	Technology	Rest API or Fabric
	API Documentation	NA
	Partners involved	UMU, JR

Funded by
the European Union

AI-based-Detector-Interface-2	**Detailed Description**	This interface allows interacting with the AI-based detector to send the final model trained through Federated Learning
	From provider	Privacy-preserving model training
	To Consumer	AI-based Detector
	Technology	REST API or Fabric
	API Documentation	NA
	Partners involved	UMU, JR

AI Correlation (AIC)

Function

Event correlation is the process of finding the relationships between events. Correlation creates context between individual events and information previously collected in real-time, and also normalises it for subsequent processing. The primary purpose of alert correlation is to identify the most significant events in the security dataset. Security event correlation should increase the quality of information about events while decreasing their number and interpreting multiple alarms.

The main directions for **application** of AI methods to correlate security events includes:

- classify security events for *event detection, event grouping,* and *event pattern extraction*
- *Intrusion detection* which deals with multi-stage and targeted attacks or *anomaly detection* to notify the security administrator about misuses and deviations from normal behaviour, respectively.
- I*ntrusion/attack projection* based on incoming events, which allows early detection of intruder targets.

There are three main areas of event correlation **methods**:

- *Similarity-based* methods are based on the idea that similar events can have the same root cause or the same type, and the found links depend on the inherent similarity between attributes of each event. Similar alerts are usually aggregated into a composite so-called *meta/hyper-alert*.
- *Causal-based* methods focus on the causal structure of an event sequence, when previous steps determine the ones that follow.
- *Data mining* is a process of discovering significant patterns, especially in a large amount of data

Different methods are appropriate for different applications.

Event correlation **AI-models** compromise the following approaches:

- *Rule-based correlation models* – similarity rules , causal rules , composite rules and rule mining models

Funded by the European Union

- *Semantic correlation models* – signature language-based , event embedding and ontology learning models.
- *Graphical correlation models* – knowledge provenance graphs and probabilistic graphical models.
- *Machine learning correlation models* – shallow and deep learning models.

AIC can be decomposed to three sub functional components

- **AI Attack Projection (AIA)-** This component uses, primarily causal-based, correlation methods to predict what an attacker in an (already observed) multi-step attack is likely to do next.
- **AI Pruning (AIP)**- This component uses, primarily similarity based, methods to reduce the flow of information to be consumed e.g. to be presented to the SOC threat analysts.
- **AI-Root Cause Analysis (AI-RCA)** - This component uses primarily similarity based methods to help identify root cause of an anomaly, very often as part of a set RCA approaches.

AIC is most often included as a component in a higher level application.

AIC is complementary to non AI based rule based correlation.

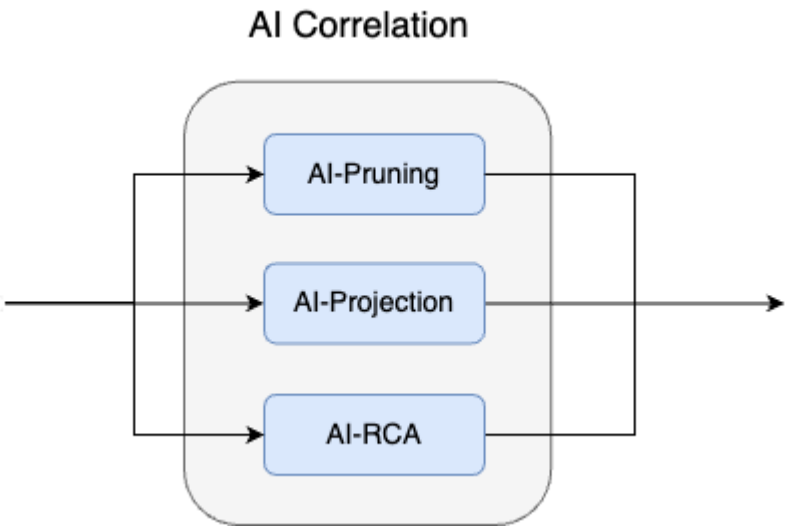

Figure 22 - AI Correlation

Provided services

AI Correlation

A. Description

This service provides generic AI correlation.

B. Capabilities

Alert grouping, prediction

C. Type

Internal.

D. Consumers

Funded by
the European Union

- Applications

E. Pre-conditions to consume the service

-.

F. Interfaces

nAIC_IngestEvents	**Detailed Description**	This interface ingests security events and alerts to the correlation process.
	From provider	AIC
	To Consumer	Alert Stream
	Technology	REST API or Method call
	API Documentation	NA
	Partners involved	TUS, MONT,

nAIC_CorrelationResult	**Detailed Description**	This interface outputs the result of the correlation process. The output depends on the type of correlation method and AI approach used. Typically it will be a meta-alert of some type or an event prediction.
	From provider	AIC
	To Consumer	Application
	Technology	REST API, or Method call
	API Documentation	NA
	Partners involved	Application Developer

TTP-based Threat Hunting and Forensics (THF)

TTP-based threat hunting focuses on identifying adversaries by their tactics, techniques, and procedures—the "how" of their operations rather than the "what" or "when." This approach leverages the ATT&CK framework (Adversarial Tactics, Techniques, and Common Knowledge) developed by MITRE to categorize and describe a comprehensive matrix of known adversary behaviours and methods. THF supports the use of TTP-based hunting techniques for cyber-attack investigation. It contains a number of sub-functions that support steps of the hunting process.

These are:

TTP Analytics

Funded by the European Union

This function allows the hunter to define the *analytic*[16] requires to collect TTP related information, and specify the data source(s) from which the data must be fetched. It provides a data model to represent entities of interest in the target domain e.g. processes, files etc. The analytic is expressed with an abstract syntax and is translated to the format (query language) for each source. The query is saved for later use to test the analytic against the data source and also for use during the hunt phase.

Analysis & Hunting

This function begins with tuning analytics for initial threat detection. If anomalies or outliers are detected, they're evaluated to confirm potential adversarial activity. Confirmed malicious activities are then investigated, gathering contextual information to fully understand the threat. Should the information be sufficient, countermeasures are imposed to disrupt the adversary. Otherwise, the hits continue to be triaged until sufficient understanding is achieved to move on to subsequent analytics. It will provide capabilities for generating reports and visualising data trends, such as heat maps or graphs, to represent the frequency or distribution of specific TTPs, attacks, or vulnerabilities over time.

Correlation

This function will develop rules or algorithms to correlate different data points and events. It will link related activities across time and terrain to construct a coherent narrative of potential adversarial actions. For example, correlating a suspicious file download on one endpoint with lateral movement attempts to another part of the network.

Hunting Database

This function will store the fetched data for the hunt, the outcomes of the correlation analysis and will also support further investigations, trend analysis, and threat intelligence enrichment.

[16] MITRE defines an ATT&CK analytic as "to the processes, techniques, and tools used to analyse data for signs of adversarial tactics, techniques, and procedures (TTPs)".

Funded by
the European Union

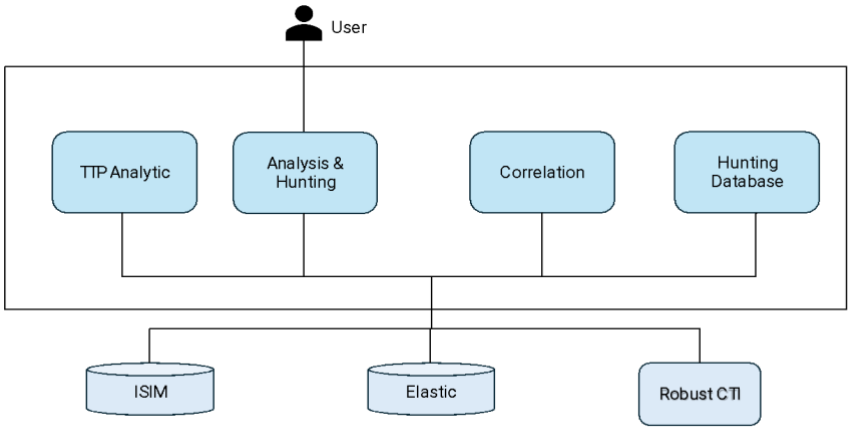

Figure 23 - Threat Hunting and Forensics

User Interaction

The user can carry out the following operations via the UI:

- Define a TTP of interest with associated data sources and analytic
- Translate the abstract analytic to the desired query language of the data sources
- Query the data sources to discover occurrences of the TTP
- Carry out investigation and analyses to discover trends and correlations between indicating TTP related malicious activities

Provided services

THF does not provide any API based services but does consume a number of other services

A. Interfaces

nISIM_AssetSearch	**Detailed Description**	This interface allows the CASM to retrieve asset details from ISIM
	From provider	THF
	To Consumer	Infrastructure and Services Information Model (ISIM)
	Technology	REST API,
	API Documentation	NA
	Partners involved	TUS, MUNI

Funded by the European Union

nSIEM_Get_Events	**Detailed Description**	This interface allows the THF to retrieve alert and threat data from the SIEM (elastic search)
	From provider	THF
	To Consumer	SIEM
	Technology	REST API
	API Documentation	https://www.elastic.co/webinars/security-and-threat-detection-with-the-elastic-stack
	Partners involved	TUS

nRCTI_ECS2STIX	**Detailed Description**	This interface allows the THF to convert behavioural observations to STIX format
	From provider	THF
	To Consumer	SIEM
	Technology	REST API
	API Documentation	NA
	Partners involved	TUS

Robust Cyber Threat Intelligence (RCTI)

Function

This function contains two sub-functions

- **CTI sharing** internally and externally. This is implemented in the first instance using MISP and may also be extended to sharing via STIX2
- **Indicator of Behaviour (IoB) creation** – IoB are activities or events that indicate possible anomalous behaviour and are seen as the next step in improving threat detection and response. These could include anomalies in network traffic, unexpected file modifications, or irregular user activities. This function builds on work carried out in the Open Cyber Security Alliance (OCA) that aims to create a standardised approach for representing cyber threat actor behaviours in a shareable format. focus on patterns of behaviour associated with malicious cyber activity. By understanding the behaviour patterns innovative solutions can be developed to enable shared behaviour sets.

Funded by the European Union

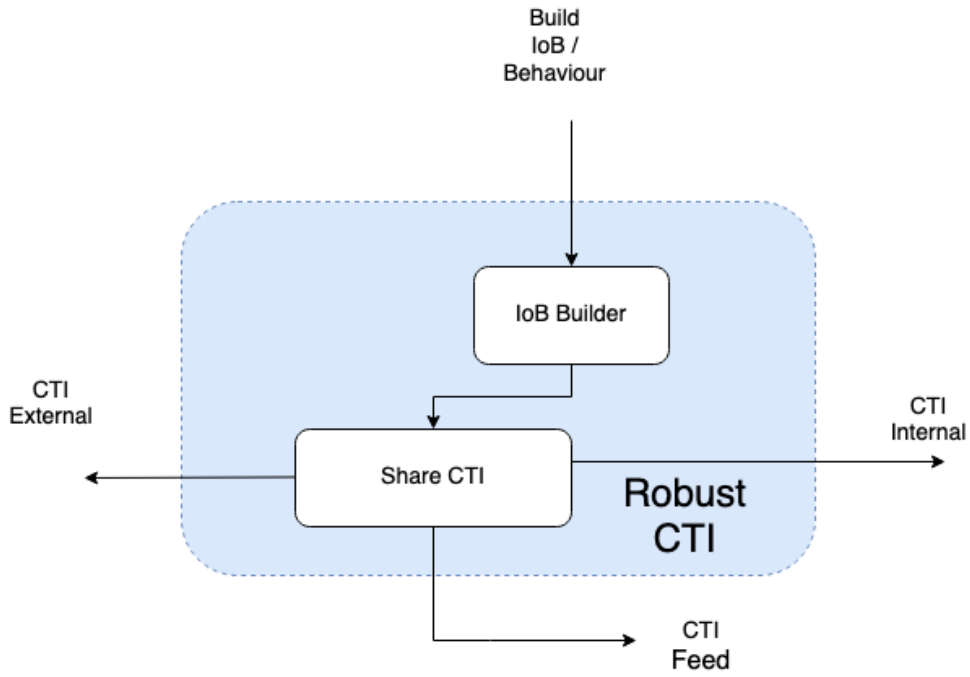

Figure 24 -Robust CTI Architecture

Provided services

CTI Sharing

A. Description

This service shares CTI internally and externally

B. Capabilities

store and share CTI

C. Type

Internal / External.

D. Consumers

- Various applications

E. Pre-conditions to consume the service

Sharing agreements defined with external partners.

F. Interfaces

nRCTI-ExtCTI	**Detailed Description**	This interface allows Resilmesh to share CTi with external organisations via MISp (or other CTI platform). Appropriate security arrangements must be defined. This is implemented in the first place for MISP
	From provider	External CTI ProviderI
	To Consumer	RCTI
	Technology	Rest API
	API Documentation	https://www.circl.lu/doc/misp/automation/ ,

Funded by
the European Union

	Partners involved	SLP

nRCTI-IntCTI	**Detailed Description**	This interface allows Resilmesh applications to push and pull CTI . This is implemented in the first place for MISP using PyMISP
	From provider	RCTI
	To Consumer	Resilmesh Application
	Technology	Rest API
	API Documentation	https://www.circl.lu/doc/misp/pymisp/
	Partners involved	TUS

nRCTI-IntFeed	**Detailed Description**	This interface enables local feeds into MISP. Feeds are remote or local resources containing indicators that can be automatically imported into MISP at regular intervals
	From provider	Resilmesh Applications
	To Consumer	RCTI
	Technology	REST API
	API Documentation	https://www.circl.lu/doc/misp/managing-feeds/
	Partners involved	TUS

IoB construction

A. Description

This service allows the construction of STIX2.1 objects to describe a chain of suspected adversarial behaviours.

B. Capabilities

generate STIX IoB objects

C. Type

Internal / External.

D. Consumers

- THF

E. Pre-conditions to consume the service

-

F. Interfaces

Funded by the European Union

nRCTI_buildIoB	**Detailed Description**	This interface enables the construction of STIX IoB behavioural object sets.
	From provider	RCTI
	To Consumer	THF
	Technology	REST API
	API Documentation	https://docs.silentpush.com/
	Partners involved	TUS

Infrastructure and Service Information Model (ISIM)

Function

The information model captures and represents all the entities of interest in the environment. The environment consists of many components of various types, which has to be reflected. The information model interconnects the pieces of information on the assets in the environment.

The information model is to be materialised as a database that will serve as a data repository for other components of T5.1 (e.g., CASM, CSA, NSE). The other components will have access to the database for reading and editing. However, there is also a need to fill the database with data from external sources, namely when deploying the overall system in a new environment. Thus, the component needs to implement a connector to an external service or repository and store the data about the infrastructure to ISIM in bulk.

Following the Cyber Defence Matrix [CDM], all assets of the following categories are considered:

- Devices,
- Network,
- Applications,
- Data,
- Users.

For each category of assets, their enumeration will be collected from existing databases, repositories, or service, or collected via a set of custom tools. Data collection is split between asset management tasks (ISIM) and attack surface discovery tasks (CASM). ISIM collects the data on assets, primarily from asset inventories like the NetBox tool [NetBox]. Widely used asset management systems and asset descriptions in non-IT domain (e.g., IoT, OT) will be provided by partners during the development.

CRUSOE data model [KOMÁRKOVÁ] serves as a foundation for this component, yet still, it will be updated to better fit the ResilMesh use cases. CRUSOE was intended to store data coming from network monitoring tools. The list of changes would go as follows:

Funded by the European Union

- Devices - referred to as *Node* and *IP* in CRUSOE, can me merged into a *Device* entity,
- Networks - CRUSOE nowadays only considers the hierarchy or subnets, but not separate networks, the change will be trivial in this regard,
- Applications - CRUSOE models only network services associated with the IP, which will be generalised,
- Data - not modelled in CRUSOE, adding them and connecting them to devices would not be complicated,
- Users - already modelled in CRUSOE, although not yet used in practice.

Moreover, ISIM will store information about critical services and business processes/missions and their mapping to assets to keep track of which assets support each critical service or missions. These mappings will be provided by CSA and, if possible, stored alongside other data in the ISIM database. In case this is not feasible, there will be a separate repository of mappings in the CSA component. Models used in the CRUSOE Decide component will be used for start.

Provided services

Three provided services are assumed: 1) DB that stores the data according to the information model, 2) data collectors or connectors to external systems, and 3) API for the use by NSE and CSA.

Information model and Data Base

A. Description

The component consists of two parts - database and a data model. The database is a functional component of ResilMesh, the data model defines the structure of the data in it and provides a common language for this and other components of WP5.

To start with, the CRUSOE data model and a graph database (preferably Neo4J) will be used.

B. Capabilities

1) Persistent storage of data about assets in the environment.
2) Common language for describing and categorising the assets.

C. Type

Internal.

D. Consumers

- NSE
- CSA

E. Pre-conditions to consume the service

ISIM and CSA data models are defined.

F. Interfaces

Only an admin interface of the database for its management and/or orchestration.

Funded by the European Union

Data collector / connector

A. Description

A data collector or connector capable of filling the ISIM database with data from external sources, there might be multiple connectors for multiple types of data sources.

Primarily, there will be an option to insert a bulk of asset descriptions in CSV or JSON format directly - for all asset types and domains. Structure of such CSV/JSON is to be defined.

Optionally, for selected asset types and scenarios, we may use existing data sources and write connectors that extract data from them, directly. For example, a widely used asset management system in IT is Netbox - an ISIM-Netbox connector would get the relevant data from Netbox (list of assets - devices, networks, …) and store them in ISIM DB. Similar connectors could be implemented to get lists of assets of other types (users, data, …) or in other domains (IoT, OT). Such a connector is run manually or optionally scheduled to update the list of assets.
if no such service exists or is not considered for an asset type or domain, ISIM falls back to textual input described above.

B. Capabilities

Reading the description of the environment from an external source (service, repository) and storing it into the ISIM database, in bulks of multiple assets of the same or several types.

C. Type

Internal / External.

D. Consumers

- ISIM

E. Pre-conditions to consume the service

External data sources are up and running or the data on assets are available elsewhere, e.g., in a CSV or JSON file. ISIM DB is up ad running.

F. Interfaces

Just input.

ISIM API

A. Description

API that allows the user/NSE/CSA to access the contents of the ISIM DB. The API implements queries expected from NSE/CSA and ensures their correct execution, conforming to the data model, etc.

B. Capabilities

Translating queries from NSE/CSA/user into valid ISIM DB queries and ensuring their proper execution. List of queries is to be specified by NSE/CSA.

C. Type

Internal

Funded by the European Union

D. Consumers

- NSE
- CSA

E. Pre-conditions to consume the service

ISIM DB is up and running

F. Interfaces

REST API or GraphQL interface that allows other components of WP5 to access ISIM data.

nISIM_AssetSearch	**Detailed Description**	This interface will allow other components in WP5 (namely NSE and CSA, in some cases even the user directly) to view and manipulate with the content of ISIM.
	From provider	ISIM
	To Consumer	NSE, CSA
	Technology	Rest API or GrapQL
	API Documentation	TBD
	Partners involved	MUNI, TUS

Cyber Asset Attack Surface Management (CASM)

Function

The Cyber Asset Attack Surface Management (CASM) provides contextual awareness about all these assets in the organisation. It has the following main features;

1. It provides status and posture information for every asset – internal and external - in the enterprise across all technology types and it continuously monitors assets for any change of status to identify risky services running the organisations network for remediation and attack surface area reduction.
2. It enables enterprise security teams to check how assets comply with the organisation security policy i.e. it can be used to detect deviations from policy by examining relationships between objects via queries to the underlying database. Security teams can ask questions such as “Which users do not have multi factor access enabled on AWS?” or “Show me which devices do not have virus checkers installed” or “Show me all assets with highest vulnerabilities”. Queries may be saved and shared.
3. It provides a capability to take action via alerting or enforcement. Users may e.g. open a trouble ticket to remedy some problem or send an email or deploy security controls etc. Actions may be manually or automatically invoked.

The function works for both IT and OT assets. It enables users to add a customised dashboard to manage their own activities.

Funded by the European Union

It has the following subfunctions:

External ASM (EASM)

This function is used to map the external attack surface of an organisation i.e. it identifies and manages threats discovered in internet facing assets using independent scans of the organisation attack surface. This includes aspects of so-called 'shadow IT' which is defined as "all hardware, software or any other solutions used by employees inside the organisation which have not been approved by the IT department". EASM discover and enumerates internet facing assets using a number of DNS enumeration techniques such as

- Brute-force of subdomains using a domain name wordlists and alteration wordlists
- Identify subdomains by reading SSL/TLS certificates, performing DNS zone transfers or checking certificate transparency logs
- Recursive subdomain discovery on identified domains

The enumeration results are stored in ISIM and the tool provides continuous monitoring by performing regular repeated enumerations and comparing and highlighting differences between enumerations. It can calculate and highlight risks associated with different assets and initiate remedial actions including decommissioning or isolating assets that don't need to be Internet facing.

Internal ASM

This function provides visibility to the organisation's internal assets that are collected in the ISIM database. It leverages the ISIM graph query language to enable security personnel to check asset status as well as relationships between assets. It will:

- discover exposures including vulnerabilities, expired certificates, etc.
- ensure compliance with regulations,
- continuously monitor all assets,
- notify security teams of changes or omissions,
- trigger actions to isolate or remediate exposures and vulnerabilities.

Alerting

This function triggers alerts and mitigation actions in response to input from ASM modules

UI

This function provides the main dashboard and user interface functionality. This will entail the use of (sub) function specific visualisations and UI screen as well as use of the ISIM console and visualisation functions. It provides the follow services to the operator/analyst via the UI:

- discover and visualise the internal and Internet facing assets
- enumerate the discovered Internet facing assets to uncover vulnerabilities and estimate risk
- discover the status, compliance and risk situation of internal assets
- initiate automated continuous monitoring of both internal and Internet facing assets
- trigger report or remediating actions.

Funded by
the European Union

Provided services

Attack Surface Management

This service provides internal and attack surface management scans and alerting. It is triggered either by a schedule or on demand by an Operator.

A. Description

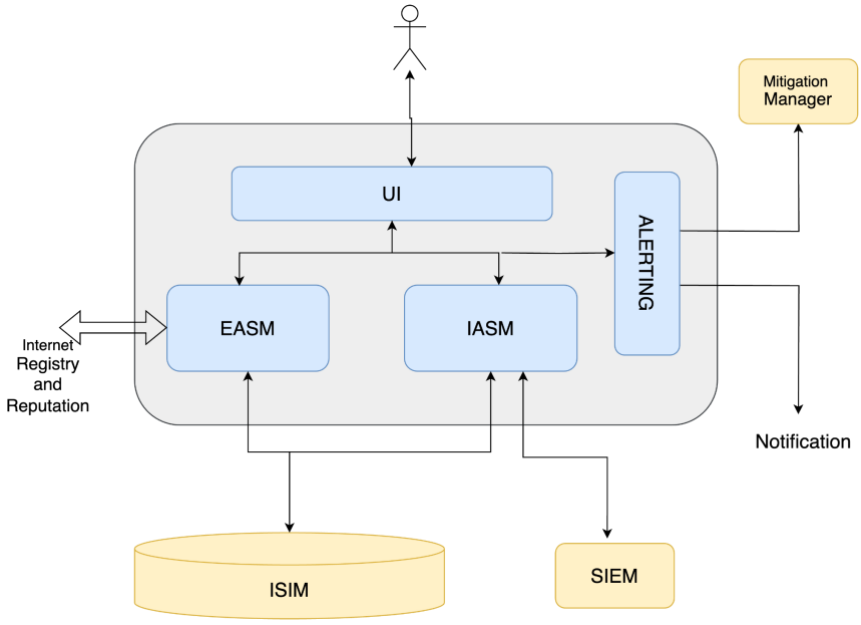

Figure 25 - CASM Architecture

B. Capabilities

Scans internal and external assets and triggers alerts

C. Type

Internal and External.

D. Consumers

Mitigation Manager; Other services e.g. email

E. Pre-conditions to consume the service

Asset connectors exist and external interfaces keys obtained

F. Interfaces

nCASM-ExtScan	**Detailed Description**	This interface enables the EASM to scan the Internet to discover the organisations external facing services and ports and to discover shadow IT
	From provider	CASM
	To Consumer	Internet Registry and Reputation Services
	Technology	Rest API
	API Documentation	NA

Funded by the European Union

	Partners involved	TUS, MUNI

nISIM_AssetSearch	**Detailed Description**	This interface allows the CASM to retrieve asset details from ISIM
	From provider	CASM
	To Consumer	Infrastructure and Services Information Model (ISIM)
	Technology	REST API,
	API Documentation	NA
	Partners involved	TUS, MUNI

nSIEM_XDRaction	**Detailed Description**	This interface allows the CASM to check assets for software vulnerabilities and/or invoke other XDR functions such as file integrity checking etc
	From provider	CASM
	To Consumer	SIEM
	Technology	REST API
	API Documentation	https://wazuh.com/platform/xdr/
	Partners involved	TUS

nCASM_Alerts	**Detailed Description**	This interface allows the CASM to initiate mitigation actions. The specific action and enforcement agent will depend on the asset and context.
	From provider	CASM
	To Consumer	Mitigation Manager
	Technology	REST API and/or Fabric
	API Documentation	NA
	Partners involved	TUS, UMU

nMISC_Notify	**Detailed Description**	This interface allows the orchestrator to notify SOC operators and other relevant personnel of alerting events.
	From provider	CASM
	To Consumer	Email, Slack and other such providers.

Funded by the European Union

	Technology	API REST
	API Documentation	NA
	Partners involved	TUS

nSIEM_Get_Events	**Detailed Description**	This interface allows the applications such as CASM to retrieve events from SIEM logs
	From provider	SIEM
	To Consumer	CASM
	Technology	REST API,
	API Documentation	https://github.com/elastic/elasticsearch
	Partners involved	TUS, All

Critical Service Awareness / Mission Awareness (CSA)

Function

The critical service awareness component will provide hierarchical risk assessment to aggregate infrastructure risk into a risk for the critical service or business mission (CS/M). It will implement tools to assess, visualise, and manage such risks. To achieve this, the component will use data stored in ISIM (information model) and forward the outputs towards NSE (Network Situation Evaluation).

The overall functionality of CSA goes as follows:

1. Information on the assets in cyber infrastructure are stored in ISIM.
2. User defines which assets support which critical service or business mission and relations are there between the assets, services, and missions. The definition is stored in ISIM (or a separate CSA repository, if needed).
3. On demand or periodically, the component will (for each critical service/mission):
    a. Iterate the assets (supporting the service/mission) and risks associated with them,
    b. hierarchically asses the risks by aggregating them bottom-up,
    c. return the overall risk to each critical service or mission,
    d. (optional) save the risks into time series for the use by NSE.
4. The outputs will be used as follows:
    a. formatted for visualisation in the NSE,
    b. critical services and missions will be sorted by the risk level,
    c. (optional) key factors contributing to the overall risk scores are highlighted (root cause analysis - perhaps leave to NSE),
    d. recommendations of tasks to manage the risks will be provided.

Funded by the European Union

CSA will allow the user to check whether a risk to an asset impacts a critical service or mission and which risks impact each critical service or mission. More precisely, it should allow the user to see which risks are the most severe with regard to CS/M.

CSA will rely on the data stored in ISIM, namely the enumeration of assets. However, it will also build a mapping between assets and CS/M. The data will be stored either in ISIM or in a separate repository of CSA, depending on practical issues. The outline of the mapping can be seen to the right of this text. Critical services ("Service") is mapped on assets (in red) via AND/OR notation. The same notation is used to map services and missions ("Supportive Process").

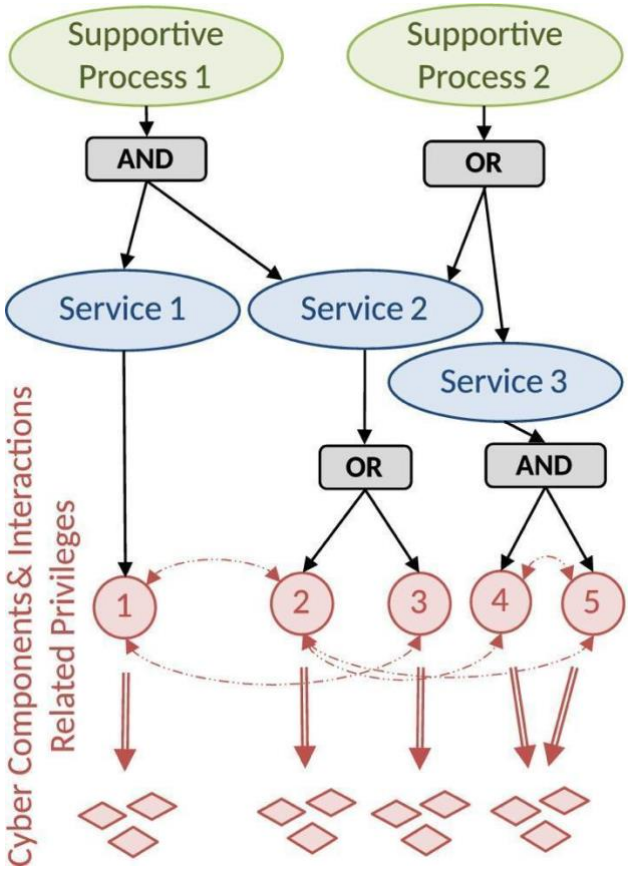

Provided services

Two services to be provided: 1) management of the repository of mappings and 2) risk assessment.

Management of asset-CS/M mappings repository

A. Description

CRUD interface that will allow for accessing and manipulating a repository of mappings.

B. Capabilities

Insert a new mapping

User prepares a JSON-formatted (preferably) document describing the mapping of CS/M to assets. The CSA will receive the JSON, validate it and check if all the assets are defined in ISIM. Then, it will check if the CS/M is already defined and either inserts or updates it.

Permanently store mappings

Resolved by ISIM, unless not technically feasible - then a separate DB will be created.

Read one or more mappings

JSON-formatted descriptions of one or more mappings are returned on demand.

Modify or delete a mapping

For each modification, a validity check will be performed - all the assets need to be
defined beforehand,

C. Type

Funded by
the European Union

Internal

D. Consumers

- ISIM (Information model)
- NSE (Network Situation Evaluation)

E. Pre-conditions to consume the service

ISIM and NSE are up and running, although it should be usable even without NSE (via text-based API)

F. Interfaces

nCSA_UpdateMission	**Detailed Description**	This interface will allow the user (via the dashboard) to read, insert, or manipulate CI/mission descriptions
	From provider	CSA
	To Consumer	NSE, ISIM
	Technology	Rest API or GrapQL
	API Documentation	NA
	Partners involved	MUNI, TUS

CS/M assessment

A. Description

The initial version will use the algorithm described in the research paper by Javorník and Husák [Javorník]. For example, let's assume a mission is supported by two critical services, each supported by an application running on a device. Applications or devices are found to be vulnerable - risks based on CVSS scores of the vulnerability are assigned to them. Risks to the applications and devices are then propagated upwards to the critical service - maximal, average, or otherwise calculated risk scores are estimated. Then, the same is done with missions - risk score is aggregated from the scores assigned to critical services.

Nevertheless, a more elaborated approach inspired by other existing solutions or operational needs of ResilMesh will be elaborated.

B. Capabilities

Assessing the risk for CS/M - each CS/M is assigned a risk score based on its underlying assets.

Funded by the European Union

C. Type
Internal
D. Consumers
- ISIM (Information model)
- NSE (Network Situation Evaluation)

E. Pre-conditions to consume the service
ISIM and NSE are up and running.
F. Interfaces

nCSA_AssessSituation	**Detailed Description**	This interface will allow the user (via the dashboard) or orchestrator (for periodically triggered assessments) to run the risk assessment routine and output its results.
	From provider	CSA
	To Consumer	NSE, ISIM
	Technology	Rest API or GraphQL
	API Documentation	NA
	Partners involved	MUNI, TUS

Network Detection and Response (NDR)

Function

Network Detection and Response (NDR) leverages non-signature-based analytical methods, such as machine learning, to identify suspicious network activities. NDR solutions continuously monitor and analyse raw enterprise network traffic to establish a baseline of normal behaviour. Any deviations from this baseline are flagged as potentially threatening, prompting alerts for security teams to investigate and respond to potential threats within their environment.

It contains the following capabilities:

1. **Traffic Analysis**: Analyse network traffic patterns to identify normal behaviour and potential anomalies. This includes monitoring bandwidth usage, protocols, and communication patterns.
2. **User Behaviour Analysis:** Analyse user activities on the network to detect anomalous behaviour. This can include identifying suspicious login patterns, unauthorized access attempts, or unusual data access patterns.
3. **Anomaly Detection:** Implement algorithms and techniques to detect anomalies in network behaviour. This can include unusual patterns of traffic, unexpected connections, or deviations from established baselines.
4. **Root Cause Analysis (RCA):** RCA aims to identify the root causes of incidents. By understanding the fundamental issues, SOC analysts can address the core problems rather than dealing with symptoms (cascading effect). This can reduce recurring incidents

Funded by the European Union

5. **Explainable AI (XAI):** XAI contributes to the Anomaly Detection model transparency as it provides insights into the model's decision-making process. (aka, explainable when a particular instance is flagged as anomalous). It also helps highlighting contributing features (Feature Importance)
6. **Reporting and Visualization:** Generate reports and visualizations to present network situational awareness information to stakeholders. This can include dashboards, graphs, and charts to communicate key insights.

NDR will build on anomaly detection and root cause correlation functions described elsewhere in this document. These will be included directly as subcomponents/analysis modules and will not be consumed as services.

The NDR tool will be based on the MMT toolset from Montimage[17]. An representative architecture of the MMT-based NDR is shown in the diagram below - the final version may deviate somewhat but will closely follow this structure

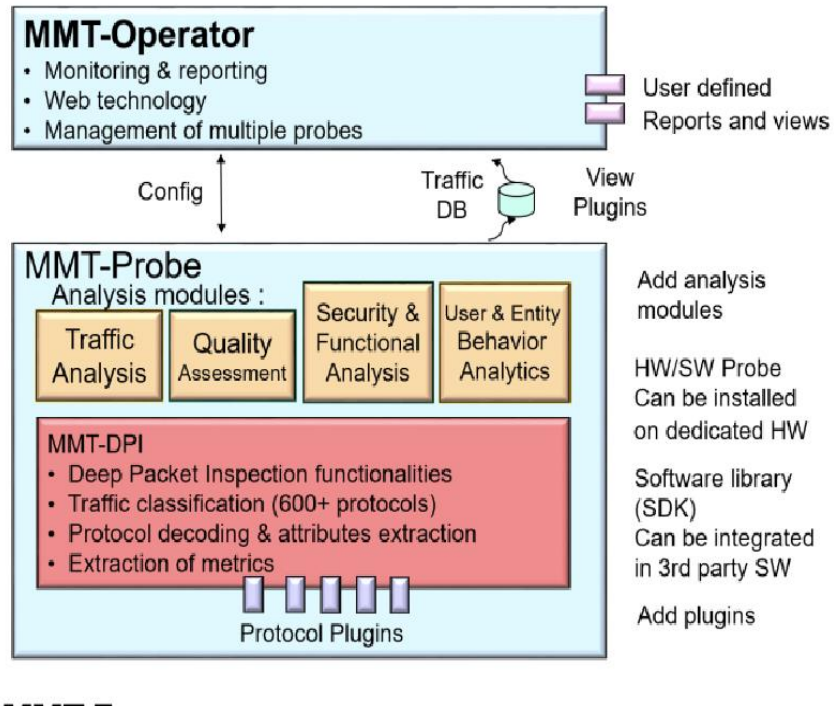

Figure 27 - NDR based on MMT

[17] https://github.com/montimage

Funded by the European Union

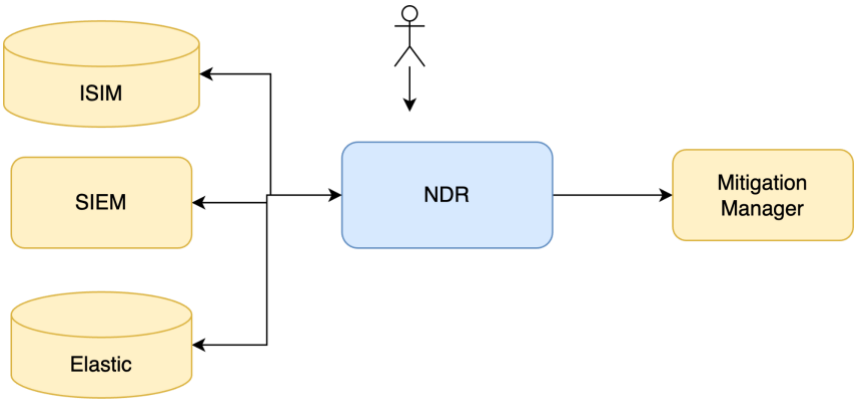

Figure 28 -NDR Functional Architecture

Provided services

NDR does not provide services to other components but does consume services from other components

A. Interfaces

nISIM_AssetSearch	**Detailed Description**	This interface allows the NDR to retrieve asset details from ISIM
	From provider	NDR
	To Consumer	Infrastructure and Services Information Model (ISIM)
	Technology	REST API,
	API Documentation	NA
	Partners involved	MONT, MUNI

nNDR_Alerting	**Detailed Description**	This interface output alerts that indicate suspect or anomalous behaviour.
	From provider	NDR
	To Consumer	NSE
	Technology	REST API,
	API Documentation	NA
	Partners involved	MONT, TUS

Funded by
the European Union

nSIEM-Alerting	**Detailed Description**	This interface outputs alerts that indicate suspect or anomalous behaviour.
	From provider	SIEM
	To Consumer	NDR
	Technology	Rest API
	API Documentation	https://documentation.wazuh.com/current/getting-started/use-cases/log-analysis.html#rules-and-decoders
	Partners involved	TUS, MONT

nSIEM_Get_Events	**Detailed Description**	This interface allows the applications such as NDR to retrieve events from SIEM logs
	From provider	SIEM
	To Consumer	NDR
	Technology	REST API,
	API Documentation	https://github.com/elastic/elasticsearch
	Partners involved	TUS, All

Network Situation Evaluation (NSE)

Function

Network Situation Evaluation provides a risk assessment of the overall network based on input from other functions and can also project the attack intensity for the network.

It provides the following capabilities:

1. calculates the overall network risk based on inputs from different sources including CSA, CASM and NDR
2. *attack intensity prediction* fuses information about the ongoing attacks from diverse sources and estimates an overall attack intensity. The overall intensity is derived from the number and severity of attacks against the whole network. The prediction can then give a warning about incoming increase or decrease of attacks
3. provides a visualisation of both current and future network risk status

Funded by the European Union

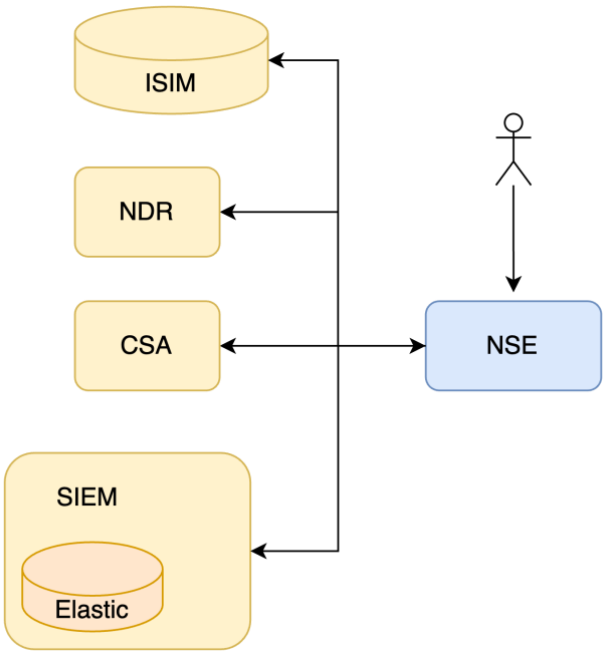

Figure 29 - NSE Functional Architecture

Provided services

NSE does not provide any services of itself but consumes many others as described below.

A. Interfaces

nISIM_AssetSearch	**Detailed Description**	This interface allows the NDR to retrieve asset details from ISIM
	From provider	NSE
	To Consumer	Infrastructure and Services Information Model (ISIM)
	Technology	REST API,
	API Documentation	NA
	Partners involved	MONT, MUNI

nSIEM-Alerting	**Detailed Description**	This interface output alerts that indicate suspect or anomalous behaviour.
	From provider	SIEM
	To Consumer	NSE
	Technology	Rest API

Funded by
the European Union

	API Documentation	https://documentation.wazuh.com/current/getting-started/use-cases/log-analysis.html#rules-and-decoders
	Partners involved	TUS, MONT,

nNDR_Alerting	**Detailed Description**	This interface output alerts that indicate suspect or anomalous behaviour.
	From provider	NDR
	To Consumer	NSE
	Technology	REST API,
	API Documentation	NA
	Partners involved	MONT, TUS

nCSA_AssessSituation	**Detailed Description**	This interface will allow the user (via the dashboard) or orchestrator (for periodically triggered assessments) to run the risk assessment routine and output its results.
	From provider	CSA
	To Consumer	NSE ,
	Technology	Rest API or GraphQL
	API Documentation	NA
	Partners involved	MUNI, TUS

Mitigation Manager (MM)

Function

The mitigation manager functional component is responsible for deciding which mitigation actions, if any, should be taken in response to a detected incident.

It interfaces with the Critical Service Awareness (CSA) to analyse mission, risk and network status projection as factors in the mitigation decision process. The mitigation manager also interacts with the CoA playbooks engine (such as Shuffle) to trigger the orchestration of the decided mitigation playbook(s) in response to the incident. It also interface with the Infrastructure and Service Model ISIM to gather the information model captures and represents all the entities of interest in the environment such as Devices, Network, Applications, Data, Users

Funded by
the European Union

The logic of the mitigation functional component might be based on a rule-based inference engine and in some cases, it might also be based on AI.

The kind of mitigation actions to be enforced as part of the playbooks might include:

- Network Filtering: divert, drop-block, mirror network traffic.
- Set-up Channel protection
- Run Ansible Scripts
- Quarantine host(S)
- Update outdated software
- Restore services.
- Update Indicators in SIEM
- Add new rules in SIEM, e.g. yara rule, sigma rules for detection.
- Share CTI

Provided services

Mitigation Service

A. Description

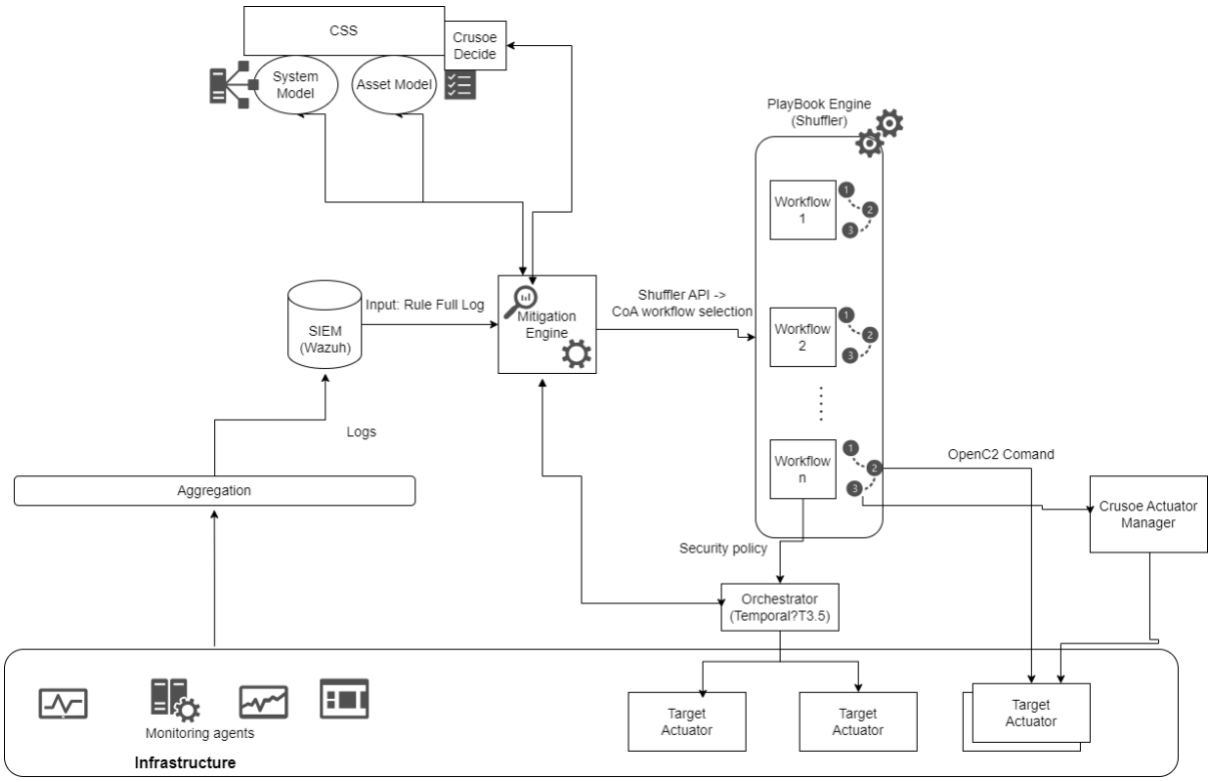

Figure 30 - Mitigation Manager

This service implements the Mitigation Manager function.

Funded by the European Union

B. Capabilities
Create
C. Type
Internal / External.
D. Consumers
- Playbook engine, Orchestrator,

E. Pre-conditions to consume the service
Abstract mitigation playbooks are defined (e.g. in CACAO)

F. Interfaces

Playbook-interface	**Detailed Description**	This interface allows interacting with the Playbook engine in order to trigger the enforcement of 1 or more selected playbooks.
	From provider	Mitigation Service
	To Consumer	Playbook engine
	Technology	Rest API or Fabric
	API Documentation	NA
	Partners involved	UMU

nSIEM_alerting	**Detailed Description**	It allows the Mitigation Engine to receive alerts that will trigger the reaction process to select the best countermeasure to apply
	From provider	SIEM
	To Consumer	MM
	Technology	PUB/SUB
	API Documentation	
	Partners involved	UMU

nCSA_AssessSituation	**Detailed Description**	This interface allows the MM to obtain risk scores from CSA
	From provider	CSA
	To Consumer	MM

Funded by the European Union

	Technology	REST API
	API Documentation	
	Partners involved	MUNI-UMU

nNSE_risk	**Detailed Description**	This interface allows the MM to obtain risk scores from Network Status Evaluation
	From provider	NSE
	To Consumer	MM
	Technology	REST API
	API Documentation	
	Partners involved	TUS

nNDR_Alerting	**Detailed Description**	This interface output alerts that indicate suspect or anomalous behaviour.
	From provider	NDR
	To Consumer	NSE
	Technology	REST API,
	API Documentation	NA
	Partners involved	MONT, TUS

nWO_orchestrate	**Detailed Description**	This interface allows the MM to enforce directly mitigation through Work Flow Orchestrator (WO) component
	From provider	MM
	To Consumer	WO
	Technology	REST API
	API Documentation	
	Partners involved	UMU-MUNI

Funded by
the European Union

nCASM_alerts	**Detailed Description**	This interface allows the MM to get alerts from the CASM component
	From provider	CASM
	To Consumer	Mitigation Manager
	Technology	REST API and/or Fabric
	API Documentation	NA
	Partners involved	TUS, UMU

PlayBooks Tool (PT)

Function

The PT is a CoA playbooks engine intended to execute specific workflows that contains as a set of actions to be enforced as countermeasures to mitigate the attacks. Those workflows can contain several steps to enforce OpenC2 actions and any other step needed in the workflow, such as deploying additional rules in SIEM.

Additionally, the playbook can include in its steps the invocation of Security Orchestrator that, in turn, can enforce certain security/privacy policies, beyond common-basic actions as defined in OpenC2. It might include for instance the deployment and configuration of certain virtual security functions in the network.

The PT should be ideally able also to define graphically standardised workflows (using standard CACAO).

The kind of mitigation actions to be enforced as part of the playbooks might include:

- Network Filtering: divert, drop-block, mirror network traffic.
- Set-up Channel protection
- Run Ansible Scripts
- Quarantining host(S)
- Update outdated software
- Restore services.
- Update Indicators in SIEM
- Add new rules in SIEM, e.g. yara rule, sigma rules for detection.
- Share CTI

Provided services

Playbook Service

A. Description

Funded by the European Union

The Play book tool is highlighted in green in the figure.

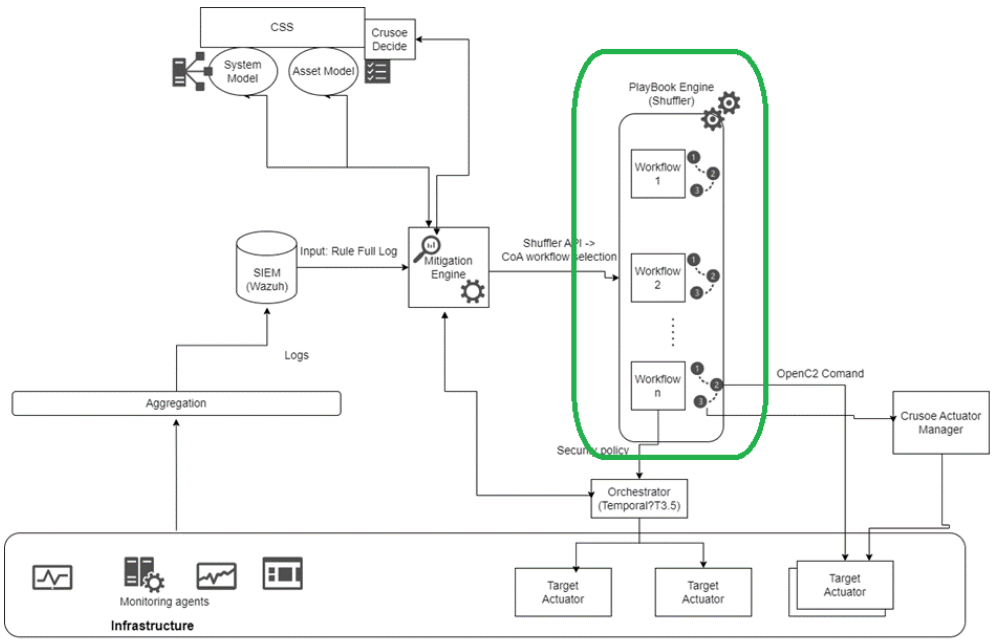

Figure 31 - Playbook Tool

The PT service is the tool that implements the PT functional component functionality.

B. Capabilities

Enforce and define CoA playbooks

C. Type

Internal / External.

D. Consumers

- Mitigator, Orchestrator,

E. Pre-conditions to consume the service

Abstract mitigation playbooks are defined (e.g. in CACAO)

F. Interfaces

nMM_runPlaybook	**Detailed Description**	This interface allows interacting with the Playbook engine in order to trigger the enforcement of 1 or more selected playbooks.
	From provider	Mitigation Service
	To Consumer	Playbook engine
	Technology	Rest API or Fabric

Funded by the European Union

	API Documentation	NA
	Partners involved	UMU

nRO_orchestrate	**Detailed Description**	This interface allows the PT to enforce complex actions through the Resource Orchestrator
	From provider	PT
	To Consumer	RO
	Technology	Rest API, Temporal interface
	API Documentation	
	Partners involved	UMU, MUNI

nActuator_enforce	**Detailed Description**	This interface allows the PT to enforce mitigations actions such as OpenC2
	From provider	PT
	To Consumer	Actuator
	Technology	Rest API, OpenC2 and others
	API Documentation	https://docs.oasis-open.org/openc2/oc2ls/v2.0/oc2ls-v2.0.html
	Partners involved	UMU, MUNI

Workflow Orchestration and Automation (WO)

Function

- Orchestration and automation to support actions taken as part of Courses of Action (CoA) playbooks
- Playbooks are advertised and requested through the *Playbook Engine* interface
- Playbooks can be triggered either on-demand or scheduled periodically
- The *Orchestrator* component can provide the data back to the *Mitigation Engine* creating a constant feedback loop
- Workflow orchestration consists of two parts:
  - high-level, low-code, visual workflow orchestration in Shuffler.io
  - low-level, all-code, complex task-as-a-workflow orchestration in Temporal.io
- The workflow orchestration and automation should be easily extensible, simple to

Funded by
the European Union

use, durable, not hard to maintain, not hard to define, and perform complex tasks. There is no existing orchestration engine that fulfils all of the mentioned requirements; therefore, the orchestration consists of the *Playbook Engine* component and the *Orchestrator* component.

- The *Playbook Engine* component represents the standardised interface for playbook management. The playbooks will be easy to view, edit, and use from a high-level point of view.
- The *Orchestrator* component will perform complex and/or time-demanding parts of the workflow. It will ensure that the execution of the actions defined via a workflow is timely, durable, and complete. Typical complex tasks are e.g.:
    - building an environment for analysing network traffic and performing the analysis,
    - building an environment for disk analysis and performing the analysis.
- It provides information about the result of actions and data obtained during their execution.
- It uses a defined subset of tools from Target Actuators to perform the actions.

Provided services

Orchestration and Automation platform (Shuffler.io + Temporal.io)

A. Description

The Playbook Engine defines use cases for which it needs cooperation from the orchestrator, typically more complex actions that are difficult to implement in Shuffler. In the orchestrator, we implement the actions by breaking them down into individual steps. Based on the request from the Playbook Engine, the Orchestrator will execute the actions on the Target Actuators and monitor them. It will provide information about the progress and outcome of the actions back to the Mitigation Engine.

Funded by the European Union

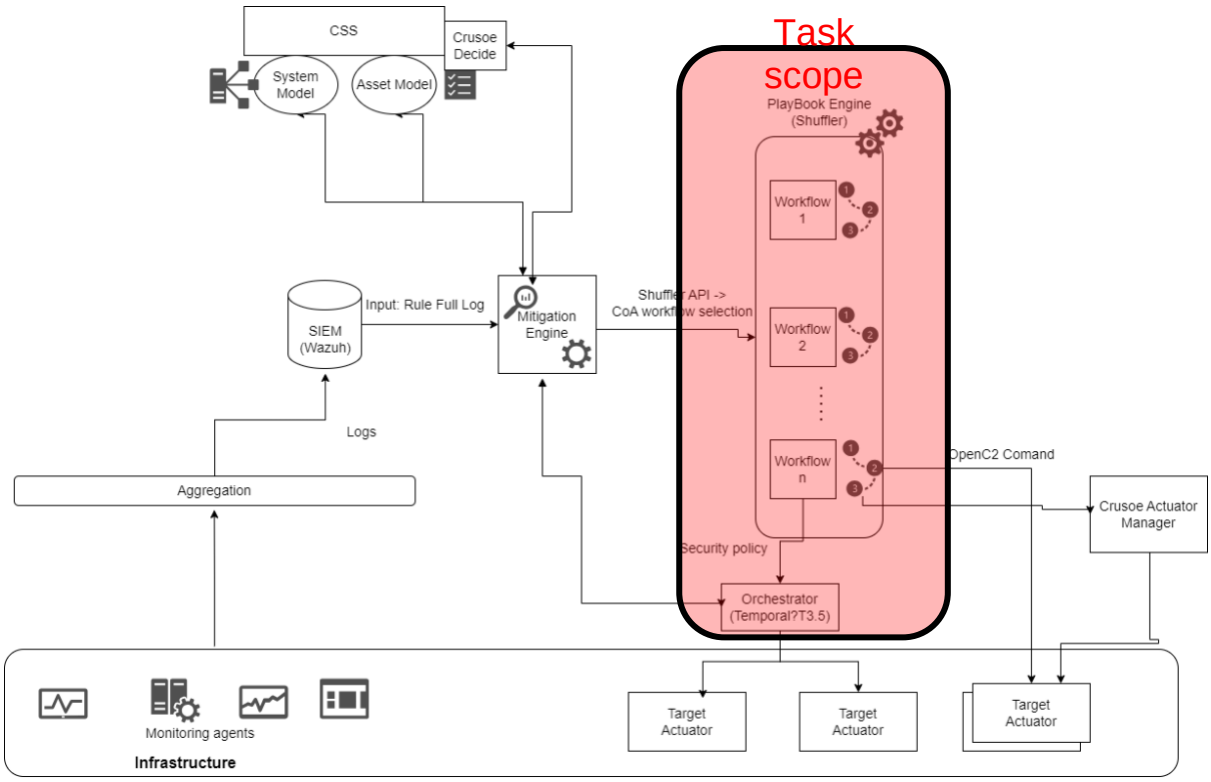

Figure 32 - Orchestration Service

B. Capabilities

Manage playbooks and execute actions

C. Type

Internal / External.

D. Consumers

Mitigation Engine, external clients.

E. Pre-conditions to consume the service

Playbook Engine needs to define a playbook providing the required workflow.

F. Interfaces

Playbook Engine Interface (Shuffler.io)	**Detailed Description**	This interface allows the Mitigation Engine and external clients to send a request for a playbook execution. The playbook will represent a predefined use case and will be accompanied with any necessary parameters and features.
	Interface provider	Playbook Engine

Funded by the European Union

	Interface consumer	Mitigation Engine, external clients
	Technology	REST API, UI
	API Documentation	NA
	Partners involved	

Orchestrator Interface (Temporal)	**Detailed Description**	This interface allows the Playbook Engine to send a request for orchestration of a complex action. The action will belong to a predefined set of supported use cases and will be accompanied by any necessary parameters and features.
	Interface provider	Orchestrator
	Interface consumer	Playbook Engine
	Technology	gRPC API, REST API
	API Documentation	https://docs.temporal.io/
	Partners involved	

Mitigation Engine Interface	**Detailed Description**	This interface allows the Orchestrator to send the feedback on the orchestrated actions back to the Mitigation Engine.
	Interface provider	Mitigation Engine
	Interface Consumer	Orchestrator
	Technology	REST API
	API Documentation	NA
	Partners involved	

Target Actuator Interface	**Detailed Description**	This interface allows the Orchestrator to manage devices required to carry out the orchestrated actions, including operations such as deploying a virtual environment, setting up a monitoring infrastructure, and carrying out data analysis.
	Interface provider	Target Actuators
	Interface consumer	Orchestrator
	Technology	REST API
	API Documentation	NA
	Partners involved	

Funded by the European Union

Reinforcement Learning based automated security testing (RLBAST)

Function

The Reinforcement Learning (RL) security testing is a training method in machine learning to enable intelligent vulnerability identification and improve resilience preparation and capacity building. By doing so, an agent performs actions on a given environment guided by a defined reward function. The agent receives rewards or penalties based on its actions. The goal is to learn the optimal behaviour that maximises the cumulative reward over time across all assets of a given environment.

It has the following main features:

- Performs an automated security testing using reinforcement learning
- Improves resilience preparedness of cyber security teams
- Supports capacity building for security test teams via a guided penetration testing approach.

Provided services

Reinforcement Learning based automated security testing Service

A. Description

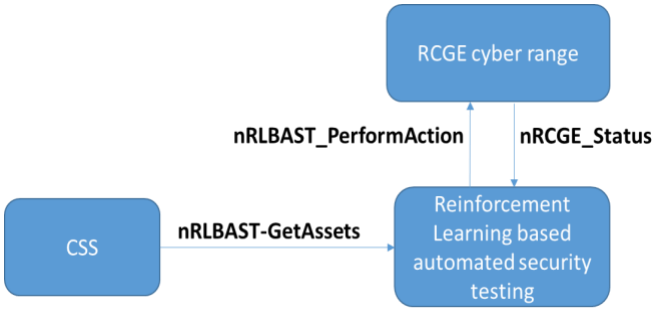

Figure 33 - RLBlast service

B. Capabilities

C. Type

D. Consumers

Cyber range testing environment

E. Pre-conditions to consume the service

Operation modes, actions and reward functions for chosen environment are defined.

F. Interfaces

Funded by the European Union

nRLBAST-GetAssets	**Detailed Description**	This interface allows the RL based automated security testing to obtain information about the assets used in the environment
	From provider	NSA
	To Consumer	RL based automated security testing
	Technology	Rest API
	API Documentation	NA
	Partners involved	TUS

nRLBAST_PerformAction	**Detailed Description**	This interface allows the RL based automated security testing to interact with the components of the environment to test.
	From provider	RL based automated security testing
	To Consumer	RCGE cyber range
	Technology	NA
	API Documentation	NA
	Partners involved	YAMK

nRCGE_Status	**Detailed Description**	This interface allows the RL based automated security testing to obtain some status information from the RCGE.
	From provider	RCGE cyber range
	To Consumer	RL based automated security testing
	Technology	NA
	API Documentation	NA
	Partners involved	YAMK

Funded by
the European Union

5 Architecture Extensions and Open Challenges

This section of the document outlines extension points in Resilmesh that can be leveraged by Open Call parties to add to the Resilmesh functional scope.

We consider four main extension categories as described below. Each category includes examples but suggestions are not limited to just these.

Extension to new domains and systems

This is concerned with extending i) the **detection capabilities** of Resilmesh to i) new OT domains - other than the three addressed directly within the project already or new OT protocols and device types and /or ii) the interoperability of Resilmesh third party security controls and tools . Examples include bit are not limited to:

- **Novel OT/IT Datasets:** This extension refers to the development, deployment and testing of innovative monitoring/probes agents as well as data processors that can gather, in real-time, data from underlying IT/OT infrastructure – other than the three domains already addressed within the project - and calculate novel and meaningful metrics and features not present in current state of the art datasets. The resultant set of data features will generate a novel dataset that can be used afterwards for the Resilmesh AI-detection engine(s) for training. In addition, the proposed monitoring agents and data processors can be deployed in Resilmesh for real-time anomaly detection.

- **Extension of the Asset Management functions** through Integration of new device types to ISIM + new applications based on ISIM/CASM: ISIM is the Resilmesh asset database and is developed in the project for a number of project use case device types. It is intended however to be extended to integrate new areas (e.g. OT domains or cloud/containerised IT systems) and the device types within these domains.

New Analytic Algorithms and Architectures

- **User and Entity Behaviour Analytics:** UEBA shifts the focus of detection from Indicator of Compromise (IoC) approaches to focus on higher level Indicators of Behaviour (IoB). UEBA can apply to both endpoint and network traffic behaviours. One approach here could be to extend the Resilmesh NDR functional component with network behaviour analytics such as those identified in the Network Traffic Analysis category in the Mitre D3FEND taxonomy (https://d3fend.mitre.org/). UEBA analytics for IIoT/OT infrastructure in particular are of interest.

Funded by the European Union

- **Novel edge AI AD architectures**: The deployment of edge-based AI opens many possibilities for experimenting with different algorithms and architectures, taking into consideration the needs of the domain and the data. Some possible approaches might be
    - Use Ensemble methods
    - Distributed deep learning
    - Incremental learning
    - Edge-to-Edge Collaborative Anomaly Detection

  These are suggestions and there can be many other approaches.

Other possible approaches may also be suggested.

Stream Processing of Security Events

This involves the real-time handling of data, where computation occurs directly as data is generated or received.

Data processing pipelines based on platforms such as SPARK , KAFKA Streams, Esper etc can be used to for many purposes including the processing of security data 'on the fly' in real time e.g. for data aggregation etc. They can also enable real-time AI processing of events streaming of events such as the novel Ad architectures outlines above

Complex event processing is a generalisation of traditional stream processing for aggregating, processing, and analysing data streams in order to gain real-time insights from events as they occur. CEP systems can be used to gain critical insights into security incident through event correlation e.g. DDoS attacks, which can help cyber and IT teams detect and prevent attacks.

Stream processing may be facilitated by the use of ***Integration of a data Lakehouse***: *A data Lakehouse is a centralised storage repository capable of accommodating* structured and unstructured data at virtually any scale. It facilitates storing data in its original format and supports a wide range of analytics, including data extraction, visualisation, big data processing, and machine learning. It consolidates data from various sources such as logs, system events, threat intelligence feeds, and others to provide a holistic perspective of security events within an organisation.

Security Operations

This category addresses potential expansion to Resilmesh functional components on the Security Control Plane. The goal is to demonstrate the use of Resilmesh mitigation orchestration and enforcement capabilities in OT domains.
Some possibilities include:

- **Novel mitigation playbooks** for handling response for new attack types. This may require the development accompanying dataset for attack detection.

Funded by the European Union

- **Development of novel software actuators** tailored for enforcing actions in OT networks and systems, such as for instance, enablers for controlling SCADA systems
- **Enhancement of Resilmesh XDR capabilities** through integration with existing EDR (Endpoint Detection and Response) systems to improve the range of attack mitigation and response controls.

Funded by
the European Union

6 Conclusion

This document has described the first version of the Resilmesh Functional architecture, aimed to help infrastructure providers to improve resilience across a wide range of dispersed and heterogenous operating environments.

The architecture has been structured in several planes to cope with the functional requirements defined in D2.1. It leverages on the (SOAPA) model, and it considers best practices by NIST publication SP 800 160 regarding cyber resiliency techniques, such as, Contextual Awareness, Analytic Monitoring, Coordinated Protection, Dynamic Positioning and Adaptive Response.

The architecture is split in some functional planes, including: Infrastructure, Aggregation, Collaboration, Threat Awareness, Situation Assessment, Security Operations. The functional components identified per plane have been thoroughly described, along with the main services and interfaces expected per each of them.
In addition, the document has described the main workflows supported by the architecture, i.e. incident detection workflow and reactive/mitigation workflow to enforce CoA playbook in response to an ongoing detected incident.

Additionally, this document has defined some possible points of extension of the architecture and open challenges that can be used to extend the ResilMesh framework, through open call specification in WP8. This architecture will be revised after second software delivery (M24) to consider findings from that work to provide a base for the overall system design and evaluation.

Funded by
the European Union

References

[ENDSLEY] ENDSLEY, Mica R. Toward a theory of situation awareness in dynamic systems. *Human factors*, 1995, 37.1: 32-64.

[JIRSÍK] HUSÁK, Martin; JIRSÍK, Tomáš; YANG, Shanchieh Jay. SoK: contemporary issues and challenges to enable cyber situational awareness for network security. In: Proceedings of the 15th International Conference on Availability, Reliability and Security. 2020. p. 1-10.

[GUTZWILLER] GUTZWILLER, Robert; DYKSTRA, Josiah; PAYNE, Bryan. Gaps and opportunities in situational awareness for cybersecurity. *Digital Threats: Research and Practice*, 2020, 1.3: 1-6.
[HUSAK] HUSÁK, Martin, et al. Survey of attack projection, prediction, and forecasting in cyber security. IEEE Communications Surveys & Tutorials, 2018, 21.1: 640-660.

[CRUSOE] HUSÁK, Martin, et al. CRUSOE: A toolset for cyber situational awareness and decision support in incident handling. *Computers & Security*, 2022, 115: 102609.

[AHMAD] AHMAD, Atif, et al. How can organizations develop situation awareness for incident response: A case study of management practice. *Computers & Security*, 2021, 101: 102122.

[Khraisat] Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (12 2019). Survey of intrusion detection systems: techniques, The ResilMesh software architecture is constructed according to the Security Operations and Analytics Platform Architecture (SOAPA and challenges. *Cybersecurity*, *2*. doi:10.1186/s42400-019-0038-7

[McMahan] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (20--22 Apr 2017). Communication-Efficient Learning of Deep Networks from Decentralized Data. In A. Singh & J. Zhu (Eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (pp. 1273–1282). Retrieved from https://proceedings.mlr.press/v54/mcmahan17a.html

[Singh] Singh, P., Singh, M. K., Singh, R., & Singh, N. (2022). Federated Learning: Challenges, Methods, and Future Directions. In S. P. Yadav, B. S. Bhati, D. P. Mahato, & S. Kumar (Eds.), *Federated Learning for IoT Applications* (pp. 199–214). doi:10.1007/978-3-030-85559-8_13

[Ruzafa] Ruzafa-Alcázar, P., Fernández-Saura, P., Mármol-Campos, E., González-Vidal, A., Hernández-Ramos, J. L., Bernal-Bernabe, J., & Skarmeta, A. F. (2023). Intrusion Detection Based on Privacy-Preserving Federated Learning for the Industrial IoT. *IEEE Transactions on Industrial Informatics*, *19*(2), 1145–1154. doi:10.1109/TII.2021.3126728

[Kairouz] Kairouz, P. et al. (2021). *Advances and Open Problems in Federated Learning*. Retrieved from http://ieeexplore.ieee.org/document/9464278

[CDM] https://cyberdefensematrix.com/

[NetBox] https://docs.netbox.dev/en/stable/

Funded by
the European Union

[KOMÁRKOVÁ] KOMÁRKOVÁ, Jana, et al. CRUSOE: Data model for cyber situational awareness. In: Proceedings of the 13th International Conference on Availability, Reliability and Security. 2018. p. 1-10.

[1] Javorník, Michal, and Martin Husák. "Mission-centric decision support in cybersecurity via Bayesian Privilege Attack Graph." *Engineering Reports* 4.12 (2022): e12538. https://onlinelibrary.wiley.com/doi/pdf/10.1002/eng2.12538

Funded by
the European Union