

Securing Cyber Resilience Through Situational Awareness: The Resilmesh Approach



Brian Lee, Martin Husák, Ekam Puri Nieto, Jorge Bernal,
Maxime Vinh Hoa La, Branka Stojanovic, and Tengfei Yang

Abstract The threats toward digital services and infrastructures continue to mushroom and advanced persistent attacks continue to proliferate. Increasingly service providers are realizing attack prevention is no longer sufficient to prevent infiltration and they have to adopt a resilience-based approach that will allow them to recover from the almost inevitable service breach. In this chapter, we therefore outline a contributing approach to achieving cyber resilience based on the use of cyber situational awareness and present how this approach is implemented in the Resilmesh system. Given the fact that cyber infrastructures are complex, heterogeneous, and dispersed, and exposed to advanced persistent threats, there is a need to map and understand them well to effectively protect them. Therefore, Resilmesh addresses these challenges by developing a resilient and collaborative event aggregation and processing mesh across dispersed digital infrastructures, giving CSIRTs better awareness of the service and asset dependencies in their network, developing AI-based algorithms and

B. Lee (✉) · T. Yang

Software Research Institute, Technological University of the Shannon, Athlone, Ireland

e-mail: brian.lee@tus.com

T. Yang

e-mail: tengfei.yang@tus.ie

M. Husák

Institute of Computer Science, Masaryk University, Brno, Czech Republic

e-mail: husakm@ics.muni.cz

E. P. Nieto · J. Bernal

University of Murcia, Murcia, Spain

e-mail: ekam.purin@um.es

J. Bernal

e-mail: jorgebernal@um.es

M. V. H. La

Montimage sarl., Paris, France

e-mail: vinhhoa.la@montimage.eu

B. Stojanovic

Joanneum Research Digital, Graz, Austria

e-mail: branka.stojanovic@joanneum.at

tools for early and ongoing attack detection and prediction, and developing a situation assessment system for system risk comprehension, projection, and response.

Keywords Cyber situational awareness • Resilience • Intrusion detection • Attack mitigation • Cybersecurity mesh architecture

1 Introduction

Resilmesh was developed to meet the expected outcomes of the Horizon Europe call HORIZON-CL3-2022-CS-01-01-“Improved monitoring of threats, intrusion detection and response in complex and heterogeneous digital systems and infrastructures”. The need to increase the resilience capability of digital infrastructures and the resilience capacity of the organizations that run them was a key thread in the call.

Resilmesh thus aims to address this demand to improve resilience in infrastructures and organizations in a systematic and structured manner. This chapter describes our approach to the former issue of improving resilience of digital services and infrastructure by improving the capabilities of improving Security Operation Center (SOC) operations. Resilmesh adopts two key principles to achieve this:

- Placing Cyber Situational Awareness (CySA) at the heart of the Resilmesh approach.
- Grounding our approach to resilience on a sound, state-of-the-art resilience engineering framework, devised by NIST [1].

Cyber situational awareness is an approach to cyber defense that entails perceiving the cyber environment, understanding the current security situation, and being able to anticipate how the situation may evolve [2]. It provides a cyber-defense team with powerful capabilities to anticipate and withstand adversarial attacks and lays the basis for recovery of services and critical infrastructure, which is the core of the resilience lifecycle. NIST defines a systematic approach to engineering resilience capabilities into the network and enshrines several types of engineering constructs to develop cyber resilient solutions. Some of these constructs relate to improving CySA and Resilmesh has adopted these practices and developed functions to implement them.

Resilmesh includes a number of novel contributions:

- The system has is specifically designed to address cyber situational awareness aided resilience and this is reflected in the architecture where the functional layers map to clear CySA functions. Moreover, novel CySA functions such as network situation awareness and mission awareness are provided to give operator powerful insight into infrastructure a service situational awareness. Resilmesh also integrates various CySA functions into a consolidated dash board framework to enable network engineers to opt to only see network status but to drill down to specific items to investigate underlying root causes.

- The system also has baked in several industry best practices to improve both interoperability following the Open Cybersecurity Alliance Open XDR architecture [3] and cyber resilience engineering best practices from NIST [1].
- A comprehensive set of integration reference points (IRP's) are provided to enable easy extensibility (at all layers) to adapt Resilmesh to a wide array of digital service domains. For example this new asset types or customized Artificial Intelligence algorithms and functions can be added and as well as new attack mitigations to address new threats.

This chapter is structured into eight sections, including this introduction. Section 2 outlines the background on cyber situational awareness and Sect. 3 introduces the concept of cyber resilience, which are the two building blocks of Resilmesh. Section 4 outlines the general approach taken by Resilmesh, while its key components are discussed in follow-up sections. Section 5 describes the contextual awareness components, Section 6 describes the analytic monitoring, and Sect. 7 describes the adaptive response. Section 8 concludes the chapter.

2 Cyber Situational Awareness

Cyber situational awareness (CySA) is an application of situational awareness into the cyber domain and a frequent research topic. In essence, it is a holistic approach to cyber defense and corresponds to perceiving the cyber environment, understanding the current security situation, and being able to project how the situation will evolve [2]. While the notion of situational awareness is very old, the theoretic foundations were laid during the last decade of the twentieth century. One of the fundamental works and definitions of SA was published by Mica Endsley [4], who proposed the three-level model definition. The definition consists of three ascending components referred to as levels. The definition can be formulated as follows: Situational Awareness is the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future [4]. Figure 1 illustrates how projection builds on comprehension, which builds on perception.

At the first level, the level of perception of the elements in the environment, status attributes and dynamics of relevant attributes of an environment are perceived [4]. The analysis and interpretation of the data is left to the next level. Nevertheless, if the data are biased or faulty, the operator gains a misleading notion of the situation and actual state. Without the correct basic perception of relevant information, the chances of making wrong conclusions increase dramatically [5]. A common problem, especially in CySA, is the data overload [6]. An operator is overwhelmed with large volumes of data from sensors, and omitting or not even encountering relevant pieces of information.

The second level, comprehension, goes beyond the perception of the data and attempts to understand the situation correctly. The comprehension of a situation is




Fig. 1 Three-level model. Adapted from [4]

based on a synthesis of disjoint elements perceived at Level 1 [4]. The elements are combined, interpreted, assigned with significance concerning given goals, and combined with knowledge of an operator to form a holistic picture of an environment. The difference between the levels is analogous to reading individual words (Level 1) and understanding the sentences (Level 2) [2].

The projection of future status is the third and highest level of situational awareness, but also the hardest to reach. In the majority of fields, where situational awareness is of importance, including cyber, the experienced operators rely heavily on future projection. The future projection based on current events and dynamics enables them to anticipate future events and their implications, which allows them to make decision in time [5].

Although CySA shares a lot with general situational awareness, there are several specific features of the cyber environment that need to be considered. First, the cyber environment has no borders. This dynamic world is highly malleable and potentially scale-free, not limited by the law of physics as in other domains. If one wishes to set any boundaries and avoid noting that everything can be everywhere, the physical location of the network or system is usually used as a spatial boundary for CySA [7]. However, even that can have its limitations due to wide use of, for example, cloud computing or network tunnels.

Second, the perception in cyber environment is more complicated. We can no longer rely on physical sensors or human perception, but use specialized hardware or software sensors and rely solely on them, since it is not possible to observe the information directly. Each sensor is a complex system that is prone to errors, measurement artifacts, or false positives, in case of various detection systems. It is also complicated to observe the adversary directly—we may only observe the activity or traces of the adversary, who typically stays hidden in both cyber and physical environment. Related factors to consider concerning the performance is the speed of the events and resources required to process them. The speed of the events is orders of magnitude quicker than in case of physical conflicts. The resources needed for the processing of such large volumes of information are significantly higher in the case of CySA [2].




Fig. 2 Taxonomy of cyber situational awareness tools and components [2]

Finally, contrary to traditional military doctrine, in which defender gains numerous advantages [6], a cyber attacker overtakes all advantages in the case of cyber conflicts. The advantages of a cyber attacker include, but are not limited to, anonymity (it can hide across national sovereignty boundaries), global reach to probe weaknesses, social engineering to exploit human weaknesses, and a possibility to pick a time, place and tools for an attack. A related issue is the disproportion of resources required for launching the attack compared to resources required to successfully detect it [2].

Figure 2 illustrates which cybersecurity tools and techniques belong to are used to reach which level of CySA [2]. Readers interested in further details on the topic of CySA are kindly referred to fundamental book by Kott et al. [6], an older, yet exhaustive literature survey by Franke and Brinielsson [8], or a more recent review by Husák et al. [2]. Nevertheless, CySA is no longer a purely academic term, but is gaining attention and recognition in the industry and commercial cybersecurity products.

3 Engineering Cyber Resilience

Cyber resilience is commonly understood as the ability of a system, or component, to recover its capability or performance after this capability has been diminished as a result of a cyber attack [6]. Resilience is related to but should not be confused with other cyber concepts such as robustness, risk and security. Traditionally, risk management is used to identify vulnerable components of a system and then apply security controls to harden these parts. As attacks have become more sophisticated and underlying technologies have become more complex and diverse it has become increasingly difficult and expensive to harden every entry point. This has led to an

acceptance that prevention is almost impossible and that organizations must adapt to attacks to bounce back from the effects of these attacks. Viewed in this way cyber resilience can be seen as the capacity of a system to:

1. *Anticipate*: Maintain a state of informed preparedness for adversity.
2. *Withstand*: Continue essential mission or business functions despite adversity.
3. *Recover*: Restore mission or business functions during and after adversity.
4. *Adapt*: Modify mission or business functions and/or supporting capabilities in response to predicted changes in the technical, operational, or threat environments.

NIST has created a systems engineering framework to develop cyber resilient solutions for cyber systems [1]. Figure 3 shows how an organization’s risk management strategy defines the priorities and requirements for its cyber resiliency goals and objectives, and how these, in turn, guide the selection and application of desired cyber *resilience techniques* to complete the required cyber resilient solution. Goals and objectives define WHAT capabilities are needed while techniques inform HOW the systems engineering framework can be used to develop cyber resilient systems. A cyber resiliency technique is thus “.. a set or class of practices and technologies intended to achieve one or more goals or objectives by providing capabilities”. The cyber resilience framework defines 14 such techniques, a sample of which are described in Table 1 to illustrate the concept. Techniques are stated at a high level and have one or more implementation approaches, i.e., “a subset of the technologies and processes included in a technique that are defined by how the capabilities are implemented,” e.g., see the privilege restriction implementation approaches in Table 1.

The standard develops the techniques and implementation approaches in greater depth and the reader is referred to the standard for a more complete description.

A subset of these techniques is used as the basis for achieving resilience in Resilmesh.




Fig. 3 Cyber resilience solution. Adapted from [1]

Table 1 Cyber resilience techniques and implementation approaches in Resilmesh

Technique	Description	Implementation approach
Coordinated protection	Ensure that protection mechanisms operate in a coordinated and effective manner.	– Trust based privilege management – Attribute based usage restriction – Dynamic privileges
Privilege restriction	Restrict privileges based on attributes on environmental factors.	Restriction of users and system elements, as well as on environmental factors.
Unpredictability	Make changes randomly or unpredictably.	– Temporal unpredictability
Segmentation	Define and separate system elements based on criticality and trustworthiness etc.	–

4 Resilmesh Approach

4.1 Challenges and Objectives

Resilmesh was developed to meet the expected outcomes of the Horizon Europe call HORIZON-CL3-2022-CS-01-01- “Improved monitoring of threats, intrusion detection and response in complex and heterogeneous digital systems and infrastructures” for which we identified the following challenges:

1. *Complex, heterogeneous and dispersed*: with multiple infrastructure layers with many types of technologies spread over wide geographical areas (cloud/edge/endpoint), creating dependencies between services and assets and thus leading to multiple attack vectors.
2. *Exposed to advanced persistent threats (APTs)*: that seek to exfiltrate information or impede critical aspects of a mission or organization based on multiple attack vector approaches.
3. *Slow adoption of best practices*: such as Zero Trust and Secure System Edge enabling easier access to attackers.

Resilmesh aims to address these challenges by improving cyber resilience through the following:

- Developing a resilient and collaborative event aggregation and processing mesh across dispersed digital infrastructures.
- Giving CSIRTs better awareness of the service and asset dependencies in their network.
- Developing AI-based algorithms and tools for early and ongoing attack detection and prediction (counter APT’s).

- Developing a situation assessment system for system risk comprehension, projection and response.

In order to achieve these goals, Resilmesh adopts the NIST resilience engineering framework [1] as described in the next section.

4.2 Resilience in Resilmesh

In this section, we describe how a subset of NIST cyber resiliency techniques—and their related implementation approaches—form the basis of the Resilmesh resilience approach.

The relevant techniques are found in Table 2. While the techniques define the high-level approach the implementation approaches actually develop the Resilmesh functional components that implement the resilience solution. These are described in detail in subsequent sections, after we first describe the Resilmesh architecture design.

4.3 Situational Awareness in Resilmesh

Resilmesh extends the CySA process described in Fig. 1 and ‘unfolds’ it into an end-to-end processing pipeline with steps from when the event is emitted by the endpoint to its final consumption by situation assessment—see Fig. 4.

This CySA ‘logical architecture’ is crystallized by the Resilmesh reference architecture, as shown in Fig. 5. This is realized as a number of distinct management planes:

- *Infrastructure Plane*: the network and system being monitored.
- *Aggregation Plane*: collects, aggregates, normalizes, and streams data and events from multiple heterogeneous sources including logs, Intrusion Detection Systems, and other network sensors.
- *Collaboration Mesh*: integration and interoperation to enable the operation of the system across the dispersed digital infrastructure.
- *Threat Awareness Plane*: suite of analytics functions to manage event correlation and alarming, attack detection and prediction, CTI sharing, and threat hunting.
- *Situation Assessment Plane*: captures dependencies between services onto the IT/OT resources that realize them; visualizes the current network risk status; forecasts the near-term situation evolution.
- *Security Operations Plane*: automates and orchestrates response and mitigation actions.

Table 2 Cyber resilience techniques and implementation approaches

Technique	Implementation approach	Description
Contextual awareness	Resource awareness	Maintain information about systems assets, resources, services and their connectivity.
	Mission awareness	Maintain information about enterprise business functions and their dependencies on IS resources and services as well as the threat status of resources.
	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.
	Risk awareness	Ensure awareness of the risk of the end point and networking components. Enable a hierarchical risk estimation for the network by aggregating the risk upwards.
Adaptive response	Dynamic reconfiguration	Make changes to systems, system elements and resources to change functionality or behavior without interrupting service.
	Adaptive management	Change how mechanisms are used based on changes in the operational environment as well as changes in the threat environment.
Analytic monitoring	Monitoring and damage assessment	Monitor and analyze behavior and characteristics of components and resources to look for indicators of adversary activity or precursor conditions or indications of other threat events and to detect and assess damage from adversity.
	Sensor and fusion analysis	Fuse and analyze monitoring data and analysis results from different information sources or at different times together with externally provided threat intelligence.
	Forensics and behavioral analysis	Analyze indicators and adversary TTPs, including observed behavior, malware, and other artifacts left behind by adverse events.

The Aggregation plane and Collaboration mesh represent the Resilmesh ‘platform’ and provide a base for the application-oriented upper layers. The threat awareness layer provides the essential basic detection capability and provide hooks for the situation assessment and security operations layer which consequently depend on the threat awareness layer. Resilmesh functions to improve resilience are primarily focused in the upper planes of the architecture (Threat Awareness (Analytic Monitoring), Situation Assessment (Contextual Awareness) and Security Operations (Adaptive Management) as these address improving resilience of enterprise services and infrastructure. However, the project also seeks to improve the resilience of the Resilmesh platform itself through the Collaboration Mesh (Adaptive Management).




Fig. 4 CySA in Resilmesh

5 Contextual Awareness

Contextual awareness corresponds to the awareness of the cyber environments, including the cyber assets, actors, and threats. Resilmesh divides the contextual awareness into three parts: resource awareness focusing on the cyber assets, such as hosts and services, in purely technical context; mission awareness that links cyber assets to business objectives and missions, including the notion of their criticality and dependencies; and risk awareness, that approaches the risks based on cyber threats and malicious behavior.

5.1 Resource Awareness

Resource awareness is paramount to the perception level of CySA as its main goal is to map the cyber environment and enumerate the assets. Resilmesh implements two closely connected components, CASM and ISIM, to conduct periodical network scans and store the results in a comprehensive way to enable resource awareness.

The **Infrastructure and Service Information Model (ISIM)** is the foundational component, which materializes the data model and provides the central point for the other components. It consists of the database, where all the data collected by CASM are stored. The database conforms to an elaborated ISIM data model, which is enforced by data quality checks during every insert operation or on demand.

We aimed to represent all assets from the cyber defense matrix (CDM)¹—devices, applications, networks, data, and users. Networks are represented as subnets with Classless Inter-Domain Routing (CIDR) notations. Moreover, their relationships in

¹ <https://cyberdefensematrix.com/>




Fig. 5 Resilmesh reference architecture

the data model express the domain names of hosts in the network, their IP addresses, and connections. All types are connected with components they are processed by or stored at. Such a representation can express captured indicators and sensitive data, which helps model cyber threats and helps with situation assessment. Subsequently, we added two additional layers to the data model to support a robotics use case. The layers represent hardware (e.g., hardware versions and device groups) and physical locations. The current draft was based on data models and ontologies for industrial robotics [9]. The model aims to support storing information about cloud devices. We added a Uniform Resource Identifier (URI) to the data model. Furthermore we added a labeling attribute to distinguish types of network nodes based on MITRE

D3FEND's Digital Artifact Ontology (DAO).² We store timestamps for some vertices and edges for situation awareness.

ISIM implementation uses Neo4j, one of the most widely known and used graph databases, since the data model is modeled as a graph and because the graph structure allows for easy extensibility of the model and the data. To facilitate the communication between ISIM and other components, two APIs are created, REST API and GraphQL.

The **Cyber Attack Surface Management (CASM)** is the data collection component. The principal role of the CASM is to monitor the organization's internal and external attack surface and security posture. It carries out domain (and sub-domain) enumeration to discover so-called 'shadow-IT', and to identify and manage threats discovered in internet-facing assets using independent scans of the organization's attack surface. CASM can detect changes of asset status and trigger required actions or alerts to the Mitigation Manager (MM) as required.

CASM component is orchestrated with Temporal.io tool,³ which allows for scheduling the scans, ensuring all tasks end successfully, and automating reactions to failures, such as waiting a specified time period and attempting the task again. This allows for continuous, reliable, and almost zero-touch scanning of the infrastructure, providing fresh data for the use by the other components.

We built the CASM around the EasyEASM.⁴ EasyEASM scans the infrastructure using a set of various tools. CASM enumerates the subdomains of a given domain and scans the IP range of the infrastructure to enumerate the active devices and network services. For scanning internal networks and non-Internet-facing assets, we also implemented a scanner based on common Nmap tool. Both tools parse CPE strings describing the software on the discovered hosts. CASM then looks up the vulnerabilities of the given CPE strings. Depending on the needs of deployment environment, for example in the realms of IoT and OT, a specialized tools can be integrated and orchestrated to collect additional data (e.g., specialized IoT or OT scanners).

5.2 Mission Awareness

Mission awareness is an innovative approach that effectively combines cyber situational awareness with business objectives of the organization operating the infrastructures protected by Resilmesh. A mission is decomposed into mission-supportive processes that are mapped to IT services that are consequently mapped to cyber components. The functional requirements are posed on IT services by the mission-supportive processes and on cyber components by the IT services. A satisfying mission configuration is a subset of mission-supportive processes, IT services, and

² <https://d3fend.mitre.org/dao/>


³ <https://temporal.io/>

⁴ <https://github.com/g0ldencybersec/EasyEASM>

cyber components that align with the mission’s functional requirements. Finally, mission resilience is the mission’s ability to continue its operations while maintaining all the required functional and security requirements.

The mission decomposition model is a graph structure with the mission-supportive processes, IT services, and cyber components as nodes [10]. The functional requirements are represented as edges between mission-supportive processes and IT services or between IT services and cyber components. The alternative mission configurations are expressed via the AND/OR nodes [11]. An AND node in a mission configuration poses a requirement for the presence of all of its child nodes. An OR node expresses a mission configuration option since it poses a requirement for the presence of at least one of its child nodes. A mission-supportive process may be satisfied by more mission configurations, i.e., combinations of IT services and cyber components. For example, a mission-supportive process dependent on one of the two IT services produces three satisfying mission configurations, one with the first IT service, one with the second IT service, and one with both. See Fig. 6 for an illustrative example of a mission decomposition model.

Fig. 6 Mission decomposition model [10]: supportive processes of business mission are mapped to abstract services (e.g., information systems), which are then mapped to specific cyber components (e.g., devices and services)



In Resilmesh, the mission awareness is implemented in the **Critical Service Awareness (CSA)** component. This component provides hierarchical risk assessment to aggregate infrastructure risk into a risk for the critical service or business mission. It also uses the data stored in ISIM and forward the outputs to the Network Situation Evaluation (NSE). The critical services or business missions can be represented as a tree structure, in which nodes representing cyber assets (devices or services) or children of nodes representing the missions or its logical component. In addition, AND/OR nodes can be placed between missions and assets to model situations, in which the mission relies on multiple assets simultaneously (AND) or one of interchangeable assets (OR). Tree structure can be provided as JSON files, stored in ISIM, and visualized in the dashboard.

5.3 Risk Awareness

This subsection introduces the Risk Situation Awareness function within the Resilmesh framework, designed to provide a comprehensive understanding of network risk posture and enable timely detection and mitigation of cyber threats. The function integrates two complementary components: **Network Status Evaluation (NSE)** and **Network Detection and Response (NDR)**.

The NSE component is responsible for assessing the overall security status of the network. It interacts with the SIEM to receive threat information and with the ISIM to retrieve systems information, e.g., node information including vulnerability status. It computes risk scores based on asset criticality, threat activity, and vulnerability status, providing a high-level view of risk across the infrastructure. NSE also supports hierarchical analysis, allowing operators to drill down into specific network segments or entities for more detailed risk insights.

The NSE application server, forming the core of the logic layer, manages the computation and visualization of network risk. It constructs a virtual representation of the monitored infrastructure by aggregating individual nodes based on user-defined criteria. For IPv4 networks, aggregation is performed using Classless Inter-Domain Routing (CIDR), which groups IP addresses into supernets for hierarchical analysis. It calculates risk scores for each aggregated node using a defined model that takes into account asset criticality, threat status, and vulnerability information. These scores are then passed to the front end for visualization. While IPv4-based aggregation is fully supported, IPv6 aggregation methods remain under consideration for future development.

The NDR component, built on top of the Montimage AI Platform (MAIP) delivers advanced threat detection and mitigation capabilities by continuously monitoring, analyzing, and responding to network traffic anomalies.

At the core of the NDR system is its AI-based behavioral anomaly detection capability. MAIP provides a suite of pre-trained machine learning models, including

supervised, unsupervised, and deep learning approaches that are optimized for identifying irregularities in network traffic. These models detect deviations from baseline behaviors and known threat signatures, thus allowing for a hybrid detection approach. To ensure transparency and interpretability of detection outcomes, MAIP integrates Explainable AI (XAI) techniques such as SHAP and LIME. By bridging the gap between high-performance automated detection and human decision-making, XAI fosters trust and accelerates incident response. A fundamental strength of the NDR system lies in its Root-Cause Analysis (RCA) module. Rather than stopping at identifying anomalous activity, the system performs in-depth analysis to trace the origin of the issue. It leverages technologies such as graph databases (e.g., Neo4j) to map dependencies between assets, services, and network segments. This allows for pinpointing the underlying causes of alerts, which is crucial for understanding how threats propagate and where mitigation efforts should be focused.

6 Analytic Monitoring

Analytic Monitoring involves adjusting the use of mechanisms in response to changes in both the operational and threat environments. Resilmesh divides analytic monitoring into three parts: *monitoring and damage assessment*, focusing on monitoring and analysis of components and resources to identify threat indicators, *sensor and fusion analysis*, focusing on fusion and analyzes of monitoring data with threat intelligence across sources and timeframes; and *forensics and behavioral analysis*, including indicators and adversary TTPs. These functions are encapsulated within the Threat Awareness Layer of Resilmesh—a key subsystem responsible for anomaly detection, event correlation, alerting, and predictive analysis.

At the core of this layer is the **Security Incident and Event Manager (SIEM)**, a central component built on Wazuh.⁵ The SIEM plays a vital role in collecting, logging, visualizing, and correlating events from across the infrastructure. It supports both rule-based and AI-assisted correlation to highlight anomalies and generate alerts. The SIEM acts as the main conduit for integrating the output of various detection models, forwarding correlated events to security operators and feeding additional analytics pipelines.

The SIEM is in turn supported by a number of other layer functional components as described in the following sessions. Together, these integrated analytic functions—anchored by the Wazuh-based SIEM—enable Resilmesh to deliver responsive, intelligent, and resilient cybersecurity monitoring across heterogeneous, distributed infrastructures.

⁵ <https://wazuh.com/>

6.1 *Monitoring and Damage Assessment*

Monitoring involves observing and analyzing the behavior and characteristics of system components and resources to identify signs of adversary activity, early threat indicators, and to detect and assess any resulting damage. Resilmesh addresses two main components: artificial intelligence (AI) machine learning (ML)-based anomaly detection and correlation.

6.1.1 *Anomaly Detection*

The Anomaly Detection Component (AID) in the Resilmesh framework is designed to detect suspicious events and potential cyberattacks in both IT and OT environments. Utilizing advanced ML and AI techniques, AID delivers robust anomaly detection across diverse applications and network protocols.

This component supports deployment on endpoints, edge devices, or in the cloud, offering flexibility based on operational needs and resource constraints. Resilmesh emphasizes the use of deep learning techniques—such as hierarchical feature extraction using autoencoders and decision-level ensemble models—to ensure accurate and efficient detection at the edge. These methods enable localized processing with minimal latency and reduced energy consumption during both training and inference.

Traditional anomaly detection methods, particularly signature-based systems, are often ineffective in identifying zero-day attacks or novel adversary behaviors. In contrast, ML and deep learning models can identify previously unseen anomalies by learning normal behavior patterns and flagging deviations from them. However, deploying these models at scale poses significant challenges due to the need for centralized datasets, substantial compute power, and compliance with privacy and data governance constraints.

To address these limitations, Resilmesh explores the use of federated learning, a distributed AI paradigm that allows multiple decentralized nodes to collaboratively train a shared model without transferring raw data. This method enhances privacy, as local data remains confined to its origin, while enabling collaborative learning across diverse systems and environments. It also improves detection latency, since the trained model is directly deployed where it is needed—at the edge.

Federated learning offers additional benefits, such as improved robustness and generalizability. By training across a wide variety of local conditions and usage patterns, the resulting model is better equipped to detect sophisticated or previously unseen threats. Moreover, distributing training workloads enhances scalability and reduces the risk of bottlenecks or single points of failure associated with centralized systems.

Nevertheless, federated learning introduces new security concerns, including model poisoning and inference attacks. Resilmesh addresses these by evaluating privacy-preserving techniques such as Differential Privacy, which ensures that the contribution of individual data points remains untraceable within the trained model.

The approach also supports massive scalability—federated learning frameworks can span thousands of nodes, enabling high-throughput, real-time monitoring of critical systems.

By combining localized deep learning, privacy-aware distributed training, and efficient edge deployment, the AID component provides a high-performance, scalable, and secure solution for anomaly detection in modern cyber-physical infrastructures.

6.1.2 Anomaly Correlation

The integration of AI in cybersecurity anomaly detection serves to enhance the correlation of events, thus reducing the number of false positives, and, moreover, facilitating context-aware anomaly detection. AIC is based on the DeepCASE framework [12], which combines contextual learning, deep learning models, and Large Language Models (LLMs) to identify and interpret complex event sequences.

This system begins with the **Context Builder** module, whose function is to extract and encode contextual information from raw event logs. The model incorporates a transformer-based architecture to produce an attention vector for the event context, which has been developed to more effectively identify sequential dependencies and enhance the model's sensitivity to minor alterations in event behavior. This enriched representation of event context forms the foundation for the subsequent analysis stages.

Then, the **Sequence Detector** module leverages transformer networks to detect and classify event sequences. The use of attention vectors ensures that the detector considers both the semantic and temporal relationships between events, thereby enhancing detection accuracy.

Using the attention vectors generated by the Context Builder module and the detection result from the Sequence Detector module, the **Interpreter** module groups similar event trajectories into distinct clusters. These clusters function as reference sets, representing known behavior patterns. The Interpreter enables the system to contextualize incoming event sequences against a repository of known patterns. This, in turn, supports the rapid identification of novel or anomalous activities.

The final module is an **LLM-based explanation** module. Initially, an LLM was employed to extract relevant information from a large corpus using a hybrid rule-based approach. In subsequent phases of development, integration between the LLM and the event clusters has been implemented to generate targeted, context-aware explanations for detected anomalies. This not only assists analysts in their understanding and decision-making processes, but also addresses the requirement for Explainable AI in security operations.

6.2 *Sensor and Fusion Analysis*

Sensor and fusion analysis involves combining and analyzing data from diverse monitoring sources and timeframes, integrating it with external threat intelligence to generate a more comprehensive understanding of potential threats. In complex IT and OT environments, where systems operate across different technologies and protocols, effective fusion of sensor data is critical to uncovering hidden or coordinated attack patterns.

Resilmesh extends the state of the art by incorporating advanced multi-view learning techniques specifically designed to address the challenges of analyzing heterogeneous data sources. These techniques enable the system to integrate and correlate information from various types of monitoring components—such as operational technology (OT) devices, industrial control protocols (e.g., Modbus), and IT intrusion detection systems (IDS)—to identify anomalies that may not be evident when sources are analyzed in isolation.

A core focus of this component is the use of deep learning methods for both feature-level and decision-level fusion. Hierarchical model architectures, including deep stacked networks (DSNs) and autoencoders, are employed to learn and merge features from different views in a structured way. In parallel, ensemble learning methods support decision fusion, enhancing model robustness and improving detection accuracy. This layered approach allows the system to account for the natural hierarchies and distribution patterns commonly found in edge and fog computing environments.

By supporting distributed deployment, fusion analysis can be performed locally at the edge or aggregated in the cloud, depending on the needs of the deployment. This ensures scalable performance and efficient resource usage, while preserving the richness of data-driven insights across domains.

6.3 *Forensics and Behavioral Analysis*

Forensics and behavioral analysis in Resilmesh focuses on Threat Hunting and Digital Forensics for Robotics.

6.3.1 *Threat Hunting*

Threat hunting is a cybersecurity practice that involves the systematic search for hidden threats, anomalous behaviors, and advanced attack techniques within the digital environment of an organization. Although it greatly enhances an organization's capacity to detect and mitigate threats at an early stage, the current approach to threat hunting is largely manual and dependent on SIEM systems and analyst expertise. A promising shift in this domain is presented by recent advances in LLMs. The

role of LLMs in structured threat hunting processes has not yet to be fully explored. To address this issue, this study focuses on integrating LLMs into the 3C (Construct, Conduct, Conclude) methodology developed by the HORIZON Resilimesh project.

Construct the Hunt. The initial phase of the 3C threat hunting process involved the establishment of the objectives of the hunt, which were typically initiated by recent incidents, CTI reports, or shifts in the threat landscape. **Conduct the Hunt.** In the second phase, analysts collect and analyze security data to test the hunt hypothesis. In the event of confirmation, the findings are documented; if inconclusive, the hypothesis may be refined and the search repeated to enhance detection accuracy. **Conclude the Hunt.** The final phase involves the documentation and reporting of findings, and the preservation of the hunt for future reference.

Each phase—Collect, Construct, and Conclude—can benefit from LLMs, especially when deployed as part of agent-based systems that support task automation and decision-making. In the **Construct the Hunt**, LLMs are capable of employing a range of techniques to gather and interpret CTI from both internal and external sources. Depending on the nature of the hunt, LLMs can operate within a conversational, single-agent framework or interact with multiple tools or agents. More research is required to identify the optimal agent architectures for different hunting scenarios and determine how LLM capabilities can be best integrated into existing security workflows. In the **Conduct the Hunt**, LLMs can dynamically configure data collection agents (e.g., Filebeat), establish SIEM ingestion pipelines, and generate detection rules tailored to the hypothesis. By reasoning across this multi-source data, LLMs provide insight into whether the hypothesis is supported. If it is not supported, the LLM has the capacity to revise the hypothesis or refine the search parameters. Given the largely autonomous, tool-rich nature of this phase of development, a multi-agent system is particularly suited to support the diverse tasks involved. **Conclude the Hunt**, it is possible to utilize the generative capabilities of the LLM to a considerable extent. An LLM is able to summarize and document the results and outputs of the search and communicate these findings to all relevant parties, as well as archive the results. It appears that a single-agent approach is adequate for this stage.

When integrated within an agent-based architecture, LLMs can operate autonomously across the 3C phases. This orchestration not only streamlines the threat hunting process, but also enables adaptive, continuous analysis at scale, enhancing the analyst's ability to detect and respond to emerging threats with greater speed and precision.

7 Adaptive Response

Adaptive Response refers to the capability of the system to take action to adjust to changes in its environment through *dynamic reconfiguration* at the collaboration mesh layer and *adaptive management* at the security and operations layer to apply attack mitigations.

7.1 *Dynamic Reconfiguration*

Resilmesh builds resiliency into its aggregation functions through the use of cloud native reconfiguration capabilities based on Service Mesh⁶ capabilities. A service mesh is a specialized infrastructure layer that manages and secures communication between microservices within a virtualised application. It provides features like traffic routing, load balancing, security, and observability without requiring changes to the application code itself. When used with appropriate resource provisioning the service mesh allows a platform service to be effectively reconfigured seamlessly across dispersed digital infrastructures with minimal effort.

The Resilmesh service mesh is implemented using the NATS message broker.⁷

7.2 *Adaptive Management*

The **Mitigation Manager (MM)** functional component—see Fig. 7—is responsible for deciding which mitigation actions, if any, should be taken in response to a detected incident. To achieve this, it follows the Reactive/Mitigation Flow:

1. As the SIEM component collects and analyzes system behavior, such traffic might trigger one or more predefined SIEM rules when it is considered to be anomalous. When this happens, the SIEM generates and fires an alert containing related information. This alert is directly consumed by the Mitigation Manager, and is the initial step in attack mitigation.
2. In order to determine a suitable mitigation for the received alert, the Mitigation Manager will utilize Cyber Situational Awareness information, the purpose of which is to provide critical and up-to-date information about the current status of the network including risk scores. This information is provided by the Infrastructure and Services Model (ISIM) component, which keeps an updated representation of the complete system explained above.
3. With the information gathered in the previous steps, the Mitigation Manager evaluates the collection of actions available for mitigation—the Course of Action (CoA) Playbooks—and selects the one most appropriate for the current situation. The Mitigation Manager (MM) optimizes its decision based on the cost of executing playbooks in order to ensure that the selected mitigation is the least disruptive.
4. Once the decision process has landed on a CoA playbook, the Mitigation Manager requests the execution of the strategy. This process requires the sourcing of playbook parameters used to determine the target/s of a mitigation strategy, the degree

⁶ <https://www.cncf.io/blog/2017/04/26/service-mesh-critical-component-cloud-native-stack/>

⁷ <https://dale-bingham-soteriasoftware.medium.com/using-nats-to-implement-service-mesh-functionality-part-1-service-discovery-5f2e6088843b/>

- to which a dynamic action is performed, a list of additional targets for the mitigation, and so on. These parameters are extracted, inferred, or computed from both the threat alert that triggered the attack mitigation process and the information retrieved from the ISIM.
5. After determining the course of action, the Mitigation Manager requests the execution of a playbook. This mitigation can be applied by two components: Playbook Tools or Workflow Orchestrator. The kinds of mitigation actions to be enforced as part of the playbooks might include: network filtering (diverting/drop-blocking traffic), network traffic mirroring, channel protection setup, Ansible script execution, host quarantining, remote software upgrading, service restoration, SIEM indicator/rule updating (YARA rules, Sigma rules for detection), and CTI Sharing.

For its main decision-making algorithm, the Mitigation Manager uses a **graph-based approach** where an attack is divided into steps, with each edge representing an attacker’s action as progress toward a goal such as exfiltrating data, or causing damage to critical infrastructure. This graph structure has been modeled after the MITRE Attack Flow data model, which defines attack graphs as directed acyclic graphs (DAGs).

Combining attack graphs and workflows, the Mitigation Manager is able to evaluate threat alerts and select optimal mitigations. The complete **mitigation process** consists of the following steps:




Fig. 7 Mitigation Manager architecture and overview

1. **Alert parsing:** When an alert is received, the Mitigation Manager dispatches a task to handle the event. SIEM alerts contain a large amount of fields, and different alert types contain different fields.
2. **Introduction of new attack graphs:** The Mitigation Manager can track multiple concurrent attacks taking place, thanks to the attack graph structure. When evaluating an alert, the attack graph database is first checked to locate any potential new attacks triggered by the alert and added to a data structure known as the “attack front”, a collection of Attack nodes considered to be the next steps taken by adversaries.
3. **Attack front advancement:** After any new attack graphs are included, the attack front is evaluated for ATT&CK Technique matches. When a technique matches a node in the attack front, the node is treated as a “compromised node”—meaning that it might have been exploited by an attacker.
4. **Updating probabilities:** Alongside attack front advancement, all nodes in attack graphs belonging to the attack front receive an update to their risk value. Risk is dynamic and changes based on multiple factors including total attack graph progression, number of conditions to fulfill, and proportion of conditions fulfilled. **Node filtering:** When selecting mitigations to apply, the Mitigation Manager distinguishes between three types of nodes: those that are directly impacted by the threat alert, those that are highly likely to be progressed to, and those that have been “historically risky”.
5. **Workflow filtering:** For each node in the previously mentioned collection, a workflow is selected. Workflows are filtered using the previously described procedure: filtering based on ATT&CK Technique and applicability, and reduced by lowest associated cost. **Workflow execution:** Once a list of workflows has been selected, they are applied using their respective components: PT workflows are forwarded to Shuffle, and WO workflows are forwarded to the selected orchestrating service in question, such as Temporal.

8 Conclusion

In this chapter, we discussed the issues of cyber situational awareness and cyber resilience in IT infrastructure and presented the approach toward achieving those concepts implemented in the Resilmesh system. We reprised the fundamental background cyber situational awareness and cyber resilience to illustrate how the Resilmesh system implements these concepts. Given the fact that current cyber infrastructures are complex, heterogeneous, and dispersed, and exposed to advanced persistent threats, there is a need to map and understand them well to effectively protect them. Therefore, Resilmesh addresses these challenges by developing a resilient and collaborative event aggregation and processing mesh across dispersed digital infrastructures, giving CSIRTs better awareness of the service and asset dependencies in their network, developing AI-based algorithms and tools for early and ongoing attack

detection and prediction, and developing a situation assessment system for system risk comprehension, projection, and response.

Declarations.

This project has been funded by the European Union as part of the Horizon Europe Framework Program (HORIZON), under the grant agreement 101119681, Resilmesh.

References

1. Ross, R., Pilliteri, V., Graubart, R., Bodeau, D., McQuaid, R.: Developing cyber-resilient systems: A systems security engineering approach. NIST Special Publication 800-160, Volume 2 Revision 1 (1), 4–32 (2021)
2. Husák, M., Jirsík, T., Yang, S.J.: 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. ARES '20. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3407023.3407062>
3. Various.: Open Cybersecurity Alliance. Open XDR (2024). <https://github.com/opencybersecurityalliance/oxa>
4. Endsley, M.R.: Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors* 37(1), 32–64 (1995)
5. Endsley, M.R.: Theoretical underpinnings of situation awareness: A critical review. In: Proceedings of the International Conference on Analysis and Measurement of Situation Awareness, p. 24 (1995)
6. Kott, A., Wang, C., Erbacher, R.F.: Cyber Defense and Situational Awareness. *Advances in Information Security*, vol. 62. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11391-3_12
7. Brynielsson, J., Franke, U., Varga, S.: Cyber Situational Awareness Testing. In: Combatting Cybercrime and Cyberterrorism, pp. 209–233. Springer, ??? (2016). https://doi.org/10.1007/978-3-319-38930-1_12
8. Franke, U., Brynielsson, J.: Cyber situational awareness—A systematic review of the literature. *Computers & Security* 46, 18–31 (2014) <https://doi.org/10.1016/j.cose.2014.06.008>
9. IEEE Robotics and Automation Society: Ieee standard ontologies for robotics and automation. *IEEE Standards*, 1–60 (2015) <https://doi.org/10.1109/IEEESTD.2015.7084073>
10. Javorník, M., Husák, M.: Mission-centric decision support in cybersecurity via Bayesian Privilege Attack Graph. *Engineering Reports* 4(12), 12538 (2022) <https://doi.org/10.1002/eng2.12538>
11. Jakobson, G.: In: Kott, A., Wang, C., Erbacher, R.F. (eds.) *Mission Resilience*, pp. 297–322. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11391-3_14
12. Ede, T., Aghakhani, H., Spahn, N., Bortolameotti, R., Cova, M., Continella, A., Steen, M., Peter, A., Kruegel, C., Vigna, G.: DeepCASE: Semi-Supervised Contextual Analysis of Security Events. In: Proceedings of the IEEE Symposium on Security and Privacy (S&P) (2022). IEEE

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits any noncommercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if you modified the licensed material. You do not have permission under this license to share adapted material derived from this chapter or parts of it.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

