CUREX

SECURE AND PRIVATE HEALTH DATA EXCHANGE

DETECTION OF SECURITY EVENTS IN HEALTHCARE INFRASTRUCTURE USING THE CUREX APPROACH

Authors: Gustavo Gonzalez Granadillo (ATOS), Rodrigo Diaz Rodriguez (ATOS), Antonio Jesus Diaz Honrubia (UPM), Alejandro Rodríguez González (UPM), Evmorfia Biliri (S5), Sotiris Koussouris (S5), Christos Bellas (AUTH), Athanasios Naskos (AUTH), Georgia Kougka (AUTH), Anastasios Gounaris (AUTH)

DECEMBER, 2021

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826404

Health Data Assets and Ecosystem

The CUREX project emerged with the aim of providing a full ecosystem that ensures the security, confidentiality, integrity and availability of health data, especially during a data transfer between two different institutions or even inside an institution. Moreover, this ecosystem should be developed guarantying a flexible, scalable and situational awareness-oriented platform.

Towards this objective, CUREX has delivered an Asset Discovery Tool (ADT) that allows the detection of the different devices and system running in the hospital or care centres. It has also designed and developed a Vulnerability Discovery Manager (VDM) that is in charge of analysing the healthcare systems at different layers and provide the necessary information for in-depth analysis from a threat intelligence point of view. Harvesting the power of machine learning and data analytics, CUREX has proposed a Knowledge Extraction & Analytics (KEA) module, able to collect knowledge out of different infrastructure components with the purpose of developing classification models for revealing vulnerabilities and profiling threats. Finally, within CUREX a Threat Intelligence Engine (TIE) has been implemented that provides real-time information about malicious activities detected in the target system, relying on advanced big data analytics.

In order to measure the cybersecurity and privacy risk of any pair of devices that will be involved in a data exchange, the proposed solution needs to perform the following processes (Figure 1):

- Discover the network topology and devices connected to the target network (Asset Discovery);
- Identify vulnerabilities associated to the discovered devices (Vulnerability Management);
- Analysing the network to discover threats and abnormal situations (Knowledge Extraction and Analytics);
- Correlate security events and provide intelligence data related to threats and attacks (Threat Intelligence).

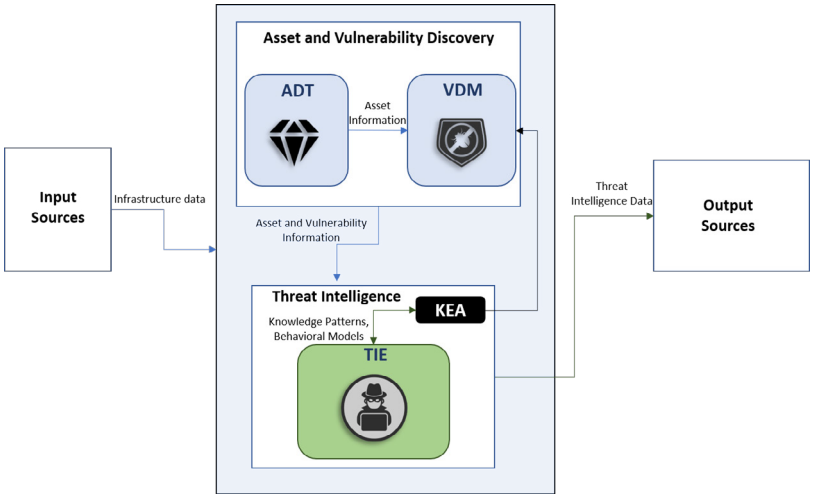

FIGURE 1. CUREX HEALTH DATA ASSETS AND ECOSYSTEM COMPONENTS

This white paper aims to describe different solutions based on activities related to the analysis and identification of security risks in the health data domain. For this purpose, in CUREX we have proposed the creation of an environment that allows the acquisition and representation of knowledge in the healthcare domain, and more specifically of those aspects that are associated to the security and integrity of both the data (patient data) and the computational-based systems (hardware and software) where the patient data is stored, detailing also how these systems are interrelated and what are their main security features. The remainder of this paper details each of the tools and services composing the CUREX health data asset and ecosystem.

Asset Discovery

This service is performed by the Asset Discovery Tool (ADT) [1], aiming to discover the assets (devices and other information about them) that are connected to a health-related IP network. When it is possible, ADT discovers open ports, services and software that are being running in the devices, the operating system, hardware information, etc

ADT has been thought as a distributed application to cope with the problem of network isolation. That way, the tool (or more specifically, a part of it) should be running in each isolated network called ADT Distributed Node. That part of the tool that is contained in each ADT Distributed Node, has been called ADT Distributed Software, will be replicated in each subnetwork and it will be in charge of discovering all the devices that are connected to that part of the network (as seen in Figure 2).

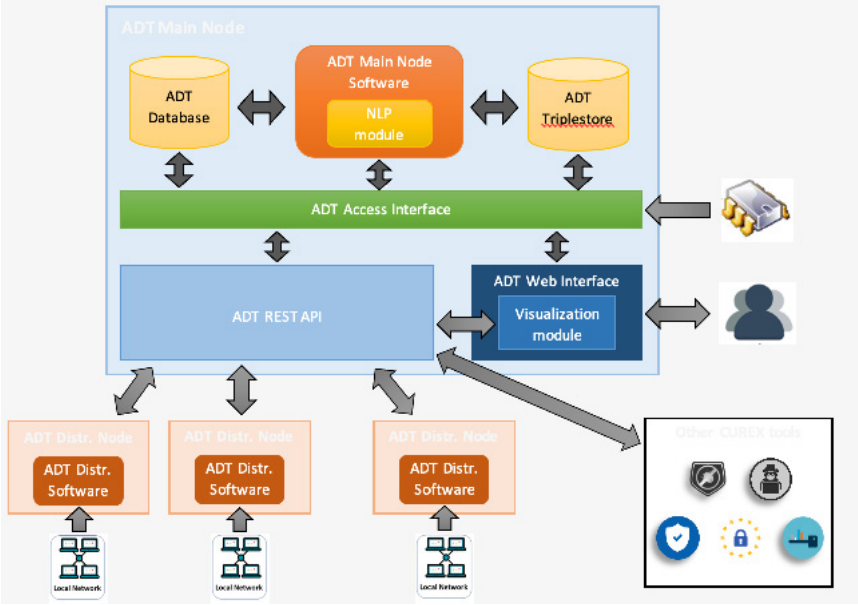

FIGURE 2. ADT GLOBAL ARCHITECTURE

The tool makes use of semantically annotated data by means of ontologies. All the gathered information is analysed and transformed into subject/predicate/object triples and stored in an ADT Main Node using an ADT Triplestore, which will be used to store semantic data. This latter can also be used to store other type of information, such as specific software details about the components installed on a device. Moreover, the information is also stored without semantical representation in the ADT Database, which is a relational database. Users and other applications should only interact with this ADT Main Node.

Information security and communication reliability are provided by the use of a Trusted Platform Module (TPM) that is used to ensure that the hardware and/or the software in the ADT Main Node are not altered.

Natural Language Processing (NLP) technology is also used so that users may provide a semi-structured textual description of the network topology. This information is then incorporated in the ADT Triplestore and the ADT Database as if it had been discovered by a distributed node.

ADT is in charge of discovering network topology and information associated to discovered hosts, but, firstly, an input from the network administrator is required. This user input information should define the networks that are going to be scanned by means of their network IP addresses, network masks, gateways, and a network identifier. Once this information is known, the ADT Distributed Nodes will search all the related information in those VLANs, being their input the responses to the network messages that they send to discover all the assets in it.

In the cases in which this information cannot be automatically discovered, the network administrator is also allowed to supply it. In this case, two kinds of structure for the provided information are available:

1. Structured information: the network administrator can send the information in JSON format emulating what an ADT Distributed Nodes would send after discovering the information. This information could be obtained in two different ways: i) the JSON has been produced by an ADT Distributed Node; 2) the network administrator writes the JSON from scratch. The JSON file must specify the VLAN ID, the field “node_ip_address” can be set to any string that, for instance, describes that this is a user input (i.e., it could say ““node_ip_address”:”user_input””), while the timestamps could be set to the current time. The JSON must also specify the list of devices of that network, with the timestamps (the same as before can be used), the MAC and IP address, the MAC card manufacturer, the operating system, and a list with all the detected open ports. For each port, it includes the number of the port, the service that is inferred to be running, the version of the service, the method for discovering this information, the state (open, closed or filtered), and the protocol of the port (TCP or UDP).

2. Semi-structured information: in this case, the network administrator does not need to build or obtain a JSON document. Instead, he/she can provide data in a plain text file per VLAN in natural language following some guidelines about the order of the elements to be specified. In these files, the first line contains information about the description of the network (its IP address and how the devices are connected to it, wire or Wi-Fi). Each of the following lines correspond to the information of each device, being the device considered as the “ADT Distributed Node” for this VLAN the first one listed. This information includes IP and MAC address (notated following the standards), open ports and services (notated as “SERVICE (PORT)”) and the device operating system. MAC address of each device should be provided as this will be part of the device identifier.

Knowledge Extraction and Analytics

This service is performed by the Knowledge Extraction and Analytics tool (KEA) [2] that harvests knowledge out of systems, sub-components and network communication interfaces that deal with health data sources (assets) in order to construct models and design methods capable of detecting threat patterns. Machine learning and broader data analytics methodologies are utilized, in order to develop classification models for revealing vulnerabilities and profiling threats; these models are then used for runtime threat detection and prediction.

The main input sources of the KEA analytics mechanism are online network traffic metrics obtained by appropriate network monitoring tools and intrusion incidents logs; system profiles provided by the Asset Discovery Tool (ADT); and manually provided audit logs, understood as historical data containing the following event types:

- infrastructure status logs;
- intrusion incidents;
- Suspicious actions; and
- Raw measurements of relevance that can be transformed to artificial events during the preprocessing phase.

Audit logs are contextualized with the help of the Asset Discovery Tools (ADT) tool and fed to KEA as timestamped annotated events. The analytics mechanism consists of three sub-components, (i) the Complex Event Processing-based Threat Detection (CEPTD), (ii) the Machine Learning-based Threat Detection (MLTD) and iii) Outlier Detection (OD). More specifically, the CEPTD uses the Common Attack Patterns Enumeration and Classification (CAPEC) [3] database as a source of predefined threat profiles for the construction of threat detection rules, used to process network traffic metrics and audit logs from Intrusion Detection Systems (IDSs) in order to detect threats. The MLTD utilizes ML algorithms and more specifically, event-based prediction algorithms, which are trained with timestamped annotated events and intrusion incidents obtained by the systems that deal with health data, in order to identify threat patterns, predict prominent threats and report the patterns and the threats to the TIE component. Multiple algorithms are employed; therefore, the final KEA solution is an ensemble one. Finally, the OD encapsulates the outlier detection technique based on unsupervised learning. This mechanism will act as a complementary tool to the deployed CEPTD and MLTD subcomponents in order to detect threats outside the CAPEC sphere.

The output of the KEA component is in the form of patterns describing threats and vulnerabilities and early detected or predicted intrusion incidents or preceding patterns, using the CEPTD, MLTD and outlier detection approach. KEA can be deployed both as an embedded component in the CUREX platform, depicted in Figure 3. In the standalone version of the KEA, the component does not require input from other CUREX components, instead it receives all the necessary input data from other sub-components in the KEA. For example, the data sources of MLTD differ between the two KEA versions, but they are still compatible. Both of the deployments share the same sub-components and functionality, having as only difference the internal and external communication.

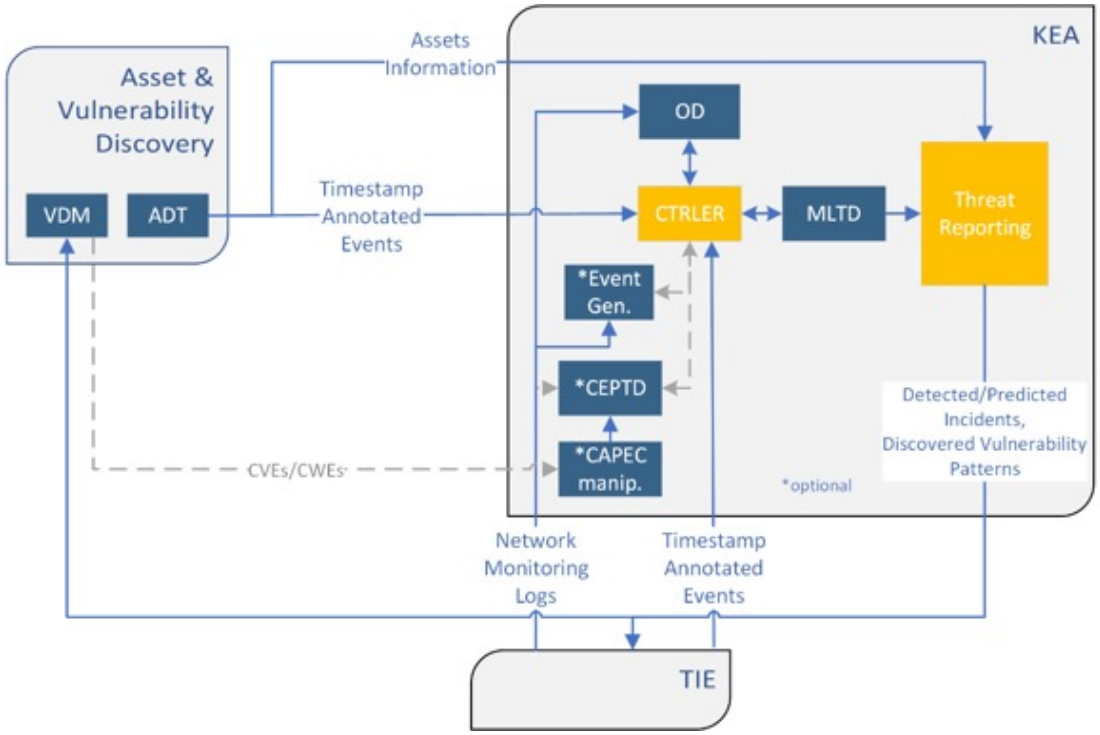

FIGURE 3. KEA REFERENCE ARCHITECTURE (EMBEDDED DEPLOYMENT)

Overall, the input to KEA component is categorized in five types, as follows:

1. Manually generated timestamp annotated events (i.e., system logs manually uploaded on the Asset and Vulnerability Discovery AVD) by IT admins annotated with cybersecurity/privacy incidents).
2. Automatically obtained timestamp annotated events.
3. Network Traffic Monitoring, e.g. pcap files.
4. Information regarding the available assets.
5. Common Vulnerabilities and Exposures (CVEs) [4] and Common Weakness Enumeration (CWEs) [5], which may, but not necessarily, be provided by VDM.

Additionally, the output of KEA consists of three types:

1. Anomalous events reported at runtime.
2. Detected/Predicted Incidents.
3. Discovered patterns of threats/vulnerabilities.

Vulnerability Management

This service is provided by the Vulnerability Discovery Manager (VDM), a domain-specific tool that identifies, analyses and manages vulnerabilities in the target infrastructure. VDM has been designed to support the key needs of the e-health sector (e.g., criticality and availability). The main role of the

- identify vulnerabilities to system resources (discovery),
- classify and assign priorities to detected vulnerabilities (assessment),
- define remediation solutions to mitigate detected vulnerabilities (treatment).
- provide intelligence sharing functionality in order to share the information with other hospitals and healthcare providers (sharing).

VDM has four main components (as depicted in Figure 4): (i) vulnerability scanner, aiming at discovering vulnerabilities in every asset composing the target system; (ii) vulnerability storage, used to store reports and information about the detected vulnerabilities; (iii) vulnerability assessment, aiming at classifying assigning priorities and treatments to all detected vulnerabilities; and (iv) vulnerability reporting and sharing, focusing on generating a vulnerability report that can be shared among organizations.

The VDM tool produces a vulnerability report with valuable information about the target network/hosts and their associated vulnerabilities. The report can be displayed in a dashboard or exported in JSON, PDF, CSV, Word or Excel format.

VDM makes use of the OpenVAS dashboard to display the list of vulnerabilities identified for a particular target, the severity of each vulnerability, a Quality of Detection (QoD) value that indicates the reliability level (0% to 100%) of the executed vulnerability detection or product detection, the affected host (IP address), location (port number and port type), and actions (e.g., exporting options).

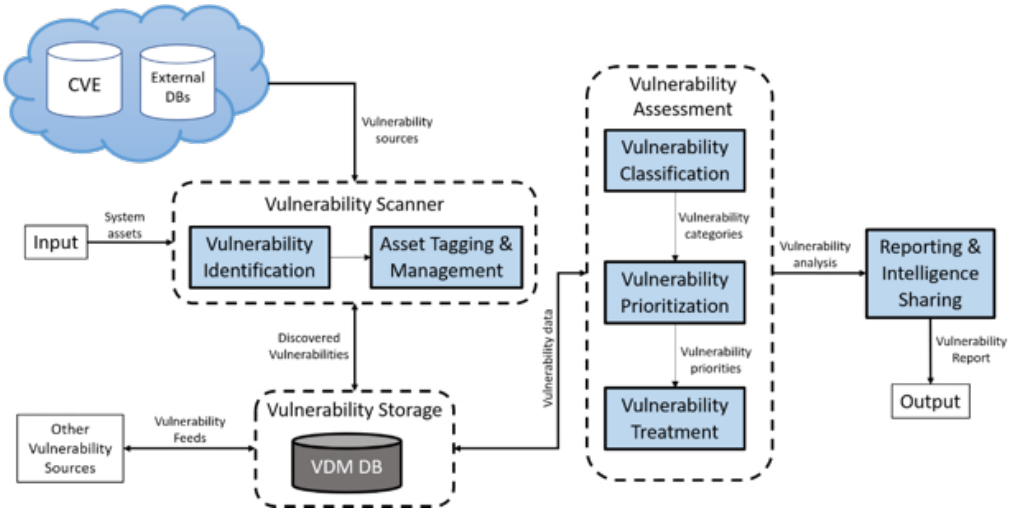

FIGURE 4. VULNERABILITY DISCOVERY MANAGER ARCHITECTURE

By clicking on a particular vulnerability from the list of results, we can obtain specific information of the vulnerability, its impact level, affected software, remediation actions and associated CVE [4]. The output of the vulnerability analysis helps security administrators to discover possible attack spaces, adversarial models and provides useful information to the threat intelligence phase, were all threats are compiled in an online dossier. The result of VDM is provided to the Threat Intelligence Engine, which will perform further and in-depth analysis of the system.

The VDM report contains in the cover page a summary of the findings that includes the overall risk level based on the Common Vulnerability Scoring System (CVSS) [6] (i.e., low, medium, high), the risk ratings: number of findings identified as low, medium, high or info; and information about the scan performed (e.g., start and finish time, duration, number of tests performed, status of the scan, etc.).

The main part of the VDM report details the findings one by one. Each finding has a flag in red, orange, and blue colors indicating its severity. Green flags provide scan coverage information (e.g., open ports pert target host). In addition, the report provides information about the name of the finding, a description, an impact level and recommendations (whenever possible).

Threat Intelligence

The Threat Intelligence Engine (TIE) [7] is a platform designed to help organisations turn the voluminous and heterogenous threat-related data generated daily by internal and external sources into actionable insights that will facilitate them in their effort to perceive, reason, learn and ultimately defend against cyber-threats. TIE constitutes a set of services integrated and calibrated to instantiate a powerful next generation SIEM that:

- Contributes towards bridging the gap between the traditional rule-based approach of SIEM systems and the currently untapped potential of applying machine learning models in real-world security systems, by combining misuse detection with anomaly detection techniques and bringing explainable AI techniques in cybersecurity.
- Extends and improves the correlation techniques applied on detected events by enabling cross-layer correlation of input data.
- Leverages new technologies towards information contextualisation and visualisation, offering more pleasant and intuitive graphical interfaces that go beyond passive consumption of threat information to empowering the user to explore and query the enriched underlying information in a more direct and flexible way that enables forensics workflows and promotes deeper understanding of the underlying system.
- Facilitates information integration and sharing, by leveraging modern technologies and established standards and semantics in the cybersecurity domain in order to enable scalable aggregation, enrichment, normalisation and processing of a wide range of input data and finally threat information sharing through multiple means.

TIE comprises the following main sub-components: (i) an Agent, responsible for the collection of the input data, their normalisation and their transfer to the sub-components responsible for the core processing, i.e. either to the main XL-SIEM engine or to the XAI engine, as shown in Figure 5; (ii) an XAI Engine, which implements the ML-enabled analysis of the input data and further comprises two interrelated modules: the first performs threat detection by applying the developed supervised machine learning models to the input data, and the second is responsible for the interpretability mechanism that generates explanations for the ML model outputs at a model- and at an instance- level; (iii) the core XL-SIEM Engine, which is responsible for the aggregation, filtering and correlation of the events collected by the agents and for the alarm generation; (iv) the TIE database, responsible for the storage of events, threat analysis results, alerts, user configurations and data coming from external sources for enrichment and conceptualisation purposes; and (v) the TIE dashboard, responsible for the data visualisation in a web graphical interface and the presentation of threat-related information in an intuitive manner.

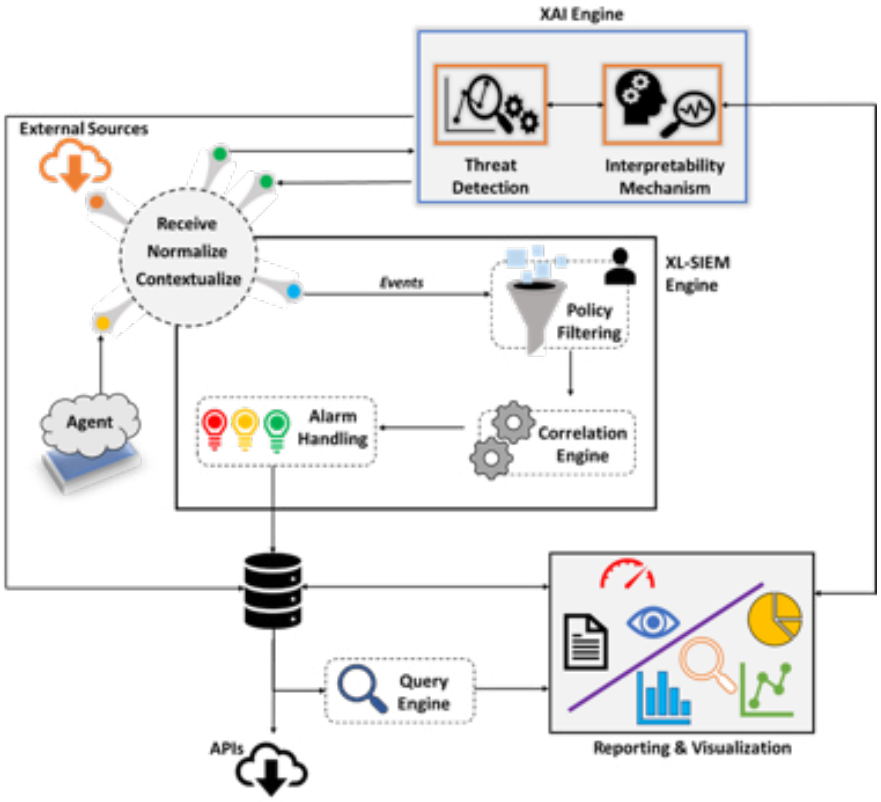

FIGURE 5. TIE ARCHITECTURE

Overall, the main novelties of the TIE are the following:

- It brings explainable AI techniques in cybersecurity and enhances rule-based SIEM systems with machine learning models for threat detection in real-world security problems.
- It provides cross-layer event correlation capabilities to take advantage of data aggregated across all layers of the monitored infrastructure.
- It enhances the contextualisation, visualisation and sharing of threat related information and enables the user to move beyond passive information consumption to flexible exploration and investigation.

Example of Usage

The output of the ADT, after being installed and executed in the target infrastructure, is depicted in Table 1, with information of the hosts composing the network, the IP addresses associated, and the operating system installed in each host.

TABLE 1. ADT OUTPUT

HOST ID	IP ADDRESS	Description	OS
DEVCUREXSAVAC	192.168.40.2	Savac server	Centos 7.7, BBDD Oracle 11.2.0.4.0
PREPACS	192.168.40.3	PACS server	Windows 2016, IIS. Ports DICOM ,1010 1011 and 104
PREPACSSQL	192.168.40.4	SQL server	Windows 2016, SQL Server
DEVCUREXSAVIO5	192.168.40.5	Integration Server	Windows 2012, Apache Tomcat

DEVCUREXSAVIFS	192.168.40.6	Savac Template server	Windows 2012, File server
DEVCURESSAVIFS1	192.168.40.7	Savac Reports server	Windows 2012, File server
PCCUREX01	192.168.41.2	Workstation 1	Windows 10
PCCUREX02	192.168.41.3	Workstation 2	Windows 10
PCCUREX03	192.168.41.4	Workstation 3	Windows 10

For each IP address discovered by the ADT, a vulnerability analysis has been performed by the VDM. Table 2 summarizes the results associated to IP 192.168.40.4. Please note that the same process has been performed in all IP addresses from Table 1. VDM provides not only the CVEs found for this particular host, but also the severity level and potential impact that such vulnerability may cause in the system if successfully exploited.

TABLE 2. VDM OUTPUT

IP ADDRESS	Vulnerability	Priority	Potential Impact
172.21.40.4	CVE-2020-11896	High	Remote Code Execution
	CVE-2020-11897	High	Possible Out-of-Bounds Write
	CVE-2020-11898	High	Possible Exposure of Sensitive Information
	CVE-2020-11901	High	Remote Code Execution
	CVE-2020-11899	Medium	Possible Out-of-bounds Read, and DoS
	CVE-2020-11903	Medium	Possible Exposure of Sensitive Information
	CVE-2020-11914	Low	Possible Out-of-bounds Read

By using ML algorithms, KEA performs an analysis to search for threat patterns that could be potentially associated to the target host. As a result, a list of weaknesses and/or attack patterns is added by the KEA to the list of vulnerabilities discovered by the VDM. Table 3 summarizes these findings.

TABLE 3. KEA OUTPUT

KEA– OD (Detection)

IP ADDRESS	Incident Time	Risk	Attack Type
172.21.40.4	2021-11-05 02:23:57	14.28%	DoS
	2021-11-05 02:24:03	23.51%	
	2021-11-05 02:37:11	13.64 %	

KEA– MLTD 1 (Prediction)

IP ADDRESS	Incident Time	Risk	Timeframe (secs)	Attack Type
172.21.40.4	2021-11-05 02:23:44	4.8%	13	DoS
	2021-11-05 02:23:48	25.3%	10	
	2021-11-05 02:23:57	63.64 %	2	

KEA– MLTD 2 (New Threat Pattern Extraction)

IP ADDRESS	Pattern
172.21.40.4	TCP SYN flood
	Packet Loss
	DoS

Finally, TIE recovers the information of the assets from the target network, their vulnerabilities, and weaknesses, as well as the security events detected in real time and correlates all this information to identify threats and/or attacks or any kind of malicious event affecting the target infrastructure. Table 4 summarizes this information and provides the number of events associated to each security alarm generated by the tool, its severity level, and the IP addresses and ports (source and destination) involved in the incident.

TABLE 4. TIE OUTPUT

Signature	Events	Severity	Source	Destination
Network scan, Nmap scan	3	Low	10.177.84.5:3510	192.168.21.4:ANY
SQL injection attempts detected	2	Medium	172.16.4.235:ANY	192.168.21.4:ANY
Host discovery attempt using PING scan	3	Low	10.177.84.5	192.168.21.4:ANY
Abnormal observation on flow	2	Medium	13.89.190.196:ANY	192.168.21.4:ANY

Conclusions

This paper described the four main aspects to be considered during the analysis and identification of security risks in the health data domain, including: (i) the asset discovery; (ii) the vulnerability management; (iii) the knowledge extraction and analytic; and (iv) the threat intelligence. All these are services composing the CUREX health data asset and ecosystem module aiming to identify potential threats and detect malicious events originated in the target infrastructure.

A high-level architecture of each service is depicted in the document and a description of their main characteristics is briefly presented along with a concrete example of usage to illustrate the applicability of the developed solutions. As a result, it is possible to detect in real time known threats and unknown attack patters related to the assets discovered in the target network. Malicious events are correlated, and a severity level is generated for each correlated alarm, to help security administrators in the decision-making process of prioritize their mitigation measures.

References

[1] CUREX Consortium. Asset Discovery Tool (ADT). CUREX deliverable 3.1 (January 2020)

[2] CUREX Consortium. Knowledge Extraction and Analytics (KEA). CUREX deliverable 3.2 (November 2019)

[3] MITRE. Common Attack Pattern Enumeration and Classification. A community resource for identifying and understanding attacks. Available at: https://capec.mitre.org/

[4] MITRE. Common Vulnerabilities and Exposures (CVE). Available at: https://cve.mitre.org/

[5] FIRST. Common Vulnerability Scoring System SIG (CVSS-SIG). Available at: https://www.first.org/cvss/

[6] CUREX Consortium. Vulnerability Discovery Manager (VDM). CUREX deliverable 3.3 (March 2020)

[7] CUREX Consortium. Threat Intelligence Engine (TIE). CUREX deliverable 3.4 (May 2020)

Partners

Contact us

Project Coordinator: University of Piraeus Research Center (UPRC)

www.curex-project.eu

info@curex-project.eu

Twitter: @CUREX_H2020

LinkedIn: CUREXH2020

Facebook: CUREXH2020

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826404