A Collaborative Analysis System for Cross-Organization Cyber Incident Handling

Giuseppe Settanni, Florian Skopik, Yegor Shovgenya and Roman Fiedler
Digital Safety and Security Department, AIT Austrian Institute of Technology, Donau-City-Strasse 1, 1220 Vienna, Austria
Giuseppe.Settanni@ait.ac.at, Florian.Skopik@ait.ac.at, Yegor.Shovgenya.fl@ait.ac.at, Roman.Fiedler@ait.ac.at

Keywords: Cyber Incident Handling, Security Operation Center, Situational Awareness

Abstract: Information and Communication Technology (ICT) systems are predominant in today's energy, finance, transportation and telecommunications infrastructures. Protecting such Critical Infrastructures (CIs) against modern cyber threats and respond to sophisticated attacks is becoming as complex as essential. A synergistic and coordinated effort between multiple organizations is required in order to tackle this kind of threats. Incidents occurring in interconnected CIs can be effectively handled only if a cooperation plan between different stakeholders is in place. Organizations need to cooperatively exchange security-relevant information in order to obtain a broader knowledge on the current cyber situation of their infrastructures and timely react if necessary. National cyber Security Operation Centers (SOCs), as proposed by the European NIS directive, are being established worldwide to achieve this goal. CI providers are asked to report to the national SOCs about security issues revealed in their networks. National SOCs correlate all the gathered data, analyze it and eventually provide support and mitigation strategies to the affiliated organizations. Although most of these tasks can be automated, human involvement is still necessary to enable SOCs to adequately take decisions on occurring incidents and quickly implement counteractions. In this paper we therefore introduce and evaluate a semi-automated analysis engine for cyber incident handling. The proposed approach, named *CAESAIR* (*Collaborative Analysis Engine for Situational Awareness and Incident Response*), aims at supporting SOC operators in collecting significant security-relevant data from various sources, investigating on reported incidents, correlating them and providing a possible interpretation of the security issues affecting concerned infrastructures.

1 Introduction

Advanced multi-stage cyber campaigns are continuously being put in place aiming at disrupting organizations' business continuity, exfiltrating sensitive data, or disturbing critical services. Such elaborated attacks intrude and leverage multiple systems and networks at the same time, therefore localized detection and isolated reaction are no longer effective. A more comprehensive and distributed approach is therefore required (Tankard, 2011).

In fact, recently issued European directives such as the NIS directive (European Commission, 2013), explicitly demand for the establishment of collaboration networks among the Member States, in order to allow secure sharing of relevant incident information and gain cyber situational awareness both on national and international level. CI providers are requested to report cyber security issues occurring on their networks to cyber SOCs. These centers are intended to collect reports, gather intelligence and provide cyber security support to the affiliated organizations while preserving their privacy.

Our work is being carried out in the context of a European funded research project named ECOSSIAN[1] (European Control System Security Incident Analysis Network) that aims at improving the detection and management of highly sophisticated cyber security incidents and attacks against CIs by implementing a pan-European early warning and situational awareness framework with command and control facilities (Kaufmann et al., 2014).

In a previous work (Settanni et al., 2015) we presented and shortly described a blueprint for our cyber incident analysis system. The system's objective is twofold: i) to gain cyber situational awareness of national CIs by analyzing security intelligence obtained from multiple sources, ii) to timely and securely distribute early warnings and advisories in the case of

[1] www.ecossian.eu

detected threats. Here we narrow down the focus on the acquisition, aggregation and analysis components outlined in that work by introducing a semi-automated system for incident-data analysis. The proposed system gathers security intelligence from multiple trusted sources, combines and correlates relevant information with reported cyber incidents, and derives possible conclusions on the occurring security issues.

An assisted-learning function allows the system to automatically determine similarities between reported issues and every other significant resource contained in the knowledge base in order to ease the analysis phase. Also, this learning process takes into account and adapts itself to operator's feedback. Operators can train the system by accepting or denying every automated association (or derived conclusion), scoring their usefulness, and providing comments about them.

We provide an illustrative scenario that points out the need for a collaborative analysis system for incident handling. The main functional requirements for our approach are derived from this scenario. We illustrate the designed system components highlighting their functionality; we give an overview on the current status of our system prototype implementation, and we evaluate our approach. We finally discuss how security intelligence is handled by the introduced system by demonstrating how it would address the challenges outlined in the presented scenario.

This paper is structured as follows. In Section II we present the ECOSSIAN ecosystem. In Section III a realistic scenario is described. In Section IV we collect the main challenges and requirements derived from the use-case. In Section V we introduce the formal model and the architectural components of our incident analysis system. Section VI reports the implementation details of the main system components. In Section VII we evaluate the performance of our system, while in Section VIII we highlight how our approach can be employed to address the issues raised in the use-case Section, by providing collaborative cyber incident handling. In Section IX we review the state of the art of distributed analysis systems for cyber incident handling. We conclude the paper in Section X with outlook and future work.

2 The ECOSSIAN Ecosystem

In order to discover and properly handle modern sophisticated cyber threats, organizations need to cooperate and exchange cyber security information so to gain knowledge of the current situation of their infrastructures (ENISA, 2013a).

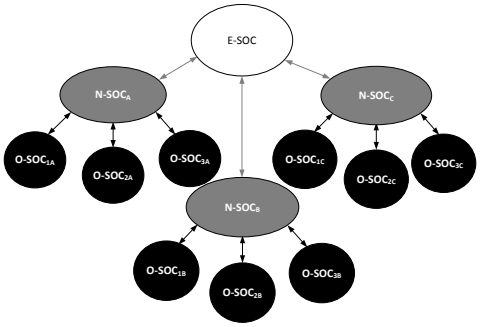

Figure 1: ECOSSIAN ecosystem.

Most organizations deploy in-house security procedures and tools aiming at detecting threats localized within their infrastructures. However, there is a significant amount of cyber threats impossible to be revealed without further information gathered from sources external to the enterprise network. Moreover, security information is most of the time also sensitive information, therefore the sharing infrastructure needs to employ mechanisms to protect the shared data against leakage, misuse and loss of reputation. Organizations tend to cooperate more between each other if a trusted third party guarantees for privacy in the communication (ENISA, 2013b).

To bridge this gap, SOCs such as (German Federal Office for Information Security, 2014), (National Cyber Security Center - Ministry of Security and Justice - Netherlands, 2014) and (Government of Canada - Cyber Security, 2014), are recently being established. SOCs collect security-relevant information generated by a multitude of systems, correlate it and try to comprehend whether the monitored systems are encountering anomalies, or are under attack. Once a threat is identified, a SOC also provides technical support to mitigate the effects it might cause, and proposes strategies for solving the issue.

In the ECOSSIAN project, we propose a Pan-European three-layered approach (Kaufmann et al., 2014) to protect CIs by detecting cyber incidents and timely generating and distributing early warnings to the monitored infrastructures. As depicted in Figure 1 we foresee three types of SOCs: Organization SOC (O-SOC), National SOC (N-SOC), and European SOC (E-SOC). At O-SOC level organizations deploy sensors and tools for threat detection, and report to N-SOCs about incidents that might have cross-organizational relevance. There are several different types of information O-SOCs share with their respective N-SOC. Data generated by sensors at O-SOC level can be automatically forwarded to the N-SOC acquisition module; security relevant information (such as threats, vulnerabilities, observations,

etc.) obtained by analyzing locally detected anomalies, can instead be manually reported by O-SOC operators. N-SOCs are deployed by European Member States joining the ECOSSIAN network, they are responsible for gaining cyber situational awareness on the network on national CIs. Here cyber intelligence is acquired by analyzing information gathered from different data sources such as reporting O-SOCs, federated N-SOCs, publicly available sources, etc. Cyber incident information aggregation, correlation, clustering and analysis are the main functionalities provided at this level. Once the evaluation of the analysis results is concluded, mitigation steps, advisories, or early warnings are sent back to the reporting O-SOCs. At the highest level an E-SOC performs analysis of strategic information shared by the different N-SOCs and distributes advisories to targeted lower level SOCs.

3 Illustrative Scenario

In this section we describe a fictional but realistic scenario of advanced cyber attacks in today's setting. It further demonstrates the need for SOCs and cooperation.

WonderLight is the main energy provider in CountryX, one of the European Member States. The corporate is structured in five different departments located in three different sites on the national territory. Three power plants of different size produce electricity distributed to more than 40% of the population.

Each chief of department receives a well-crafted email, apparently sent by the CEO asking the recipient to click on a link to get access to the last quarterly financial report. Two recipients (Dr. A and Dr. B) click on the link and are directed to an empty web page. Dr. A and Dr. B. close the web browser ignoring the effects of their clicks and keep carrying out their activities. Mrs. C, who received the message as well, suspects about the trustworthiness of the email, but being uncertain of what to do, she deletes the email from the mail client without clicking on the link. The remaining two victims (Mrs. D and Dr. E) are more familiar with cyber-security issues and since they are aware that the financial report would not be ready before three weeks, they immediately realize the email they just received is malicious. They report the security event to the company's IT department in charge, among other tasks, of investigating cyber-security issues.

The IT department examines the reported issue, but due to lack of resources (personnel, time, tools) for a deep inspection, and noticing that the web page is not reachable anymore when clicking on the received link, they simply alert all the company employees about the occurred phishing attempts and instruct them not to click on any link received in similar emails.

Moreover, Dr. A and Dr. B, by clicking on the link, inadvertently downloaded through their web browser a sophisticated malware undetectable by the antivirus solutions installed on their machines. The malware exploits a Windows vulnerability, executes a daemon performing port scanning and communicating with the command and control server through an encrypted connection. The malicious software detects a list of open ports and hence allows the attackers to gain remote privileged access to the infected computers. The attackers remotely connect to these machines and thanks to their privileges are able to browse through the corporate network, explore privacy-sensitive documents, and get access to the Industrial Control Systems (ICSs) managing the energy production at the power plants.

The attackers are skilled and motivated activists aiming at boycotting and damaging WonderLight's and other European energy providers' reputation because of their questionable environmental policy. After three weeks of undetected intrusion, the attackers have gained a good knowledge of WonderLight's infrastructure. They start maneuvering numerous Programmable Logic Controllers (PLCs) employed in the ICSs at the main power plant, and gradually decrease the energy production until shutting down entire portions of the electrical network. This causes hours of blackouts in several regions of the country until the business continuity plan of WonderLight is activated, the intrusion is detected and the normal energy production is restored. Such an event implies a considerable loss of revenue and reputation for WonderLight.

If similar APTs target at the same time every other energy provider in the country, the effects of such distributed attack can be dramatic (Reichl et al., 2013).

Although modern European States rely on CIs employing state-of-the art intrusion detection systems (IDs) and off-the-shelf industrial control systems, they are not able to prevent the described attack scenario because *they lack of dedicated organization and especially national cyber SOCs, the CIs do not cooperate with one another, and they rarely examine cyber threat intelligence generated by third parties*.

Starting from the presented scenario we derive in the following section the main functional requirements for the analysis system, deployed at the N-SOC in the ECOSSIAN ecosystem, to tackle the outlined security challenges.

4 Functional Requirements

Diverse aspects are to be considered when designing an incident analysis system for a national SOC in the ECOSSIAN ecosystem.

In order to have a comprehensive picture of the cyber security situation of the CIs in a country, it is necessary for the N-SOC to collect relevant information from multiple sources.

Publicly available information should be collected and taken into account to obtain a broad knowledge on the most recent security events. Information such as latest available updates and patches, discovered vulnerabilities, released security bulletins should be considered while handling cyber incidents on a national level.

At organization's level state-of-the-art intrusion, threat and anomaly detection systems are required to be in place aiming at revealing malfunctions, attacks and deviations from the normal systems' behavior. The outcomes generated by such tools (alarms, alerts, logs, etc.) are analyzed and processed by O-SOC operators and if necessary forwarded to the N-SOC. In ECOSSIAN we foresee two types of information to be transmitted from O-SOCs to N-SOC. Manual reports can be filled out by O-SOC operators describing observed security incident and then sent to the N-SOC through a secure and privacy-preserving sharing network; moreover data automatically generated by sensors[2] deployed in the infrastructure (including IDS alerts, firewall logs, SIEM alarms, and ECOSSIAN sensor readings) can automatically be collected by the N-SOC acquisition module (Settanni et al., 2015).

In order to regulate the information sharing process between O-SOCs and the respective N-SOC, reporting policies must be in place. Source and destination entities, formats, timing, frequency, content and every other detail of the sharing processes has to be specified within these policies. Only adhering to this commonly agreed set of rules different organizations can cooperate effectively in a trusted manner.

Another crucial aspect to take into account is privacy. Organizations will be prone to cooperate with one another and share security data only if confidentiality is guaranteed. Data protection is fundamental for the design of a secure incident sharing network and the implementation of an efficient incident analysis system. Access control, anonymization and encryption functionalities must be therefore enabled.

The process of cyber incident handling can massively be supported by automated procedures and software tools for aggregating, classifying and analyzing frequently generated and high volume incident data. However, not all the analysis tasks can be carried out without human involvement; complex issues and manually reported incidents, received less frequently than sensor readings, need to be thoroughly examined by N-SOC security experts.

The cooperative analysis system we propose in this paper has the purpose of facilitating the acquisition of national cyber situational awareness, and assisting the decision making process by bringing significant intelligence to the attention of N-SOC operators. To achieve this goal sophisticated visualization tools demonstrating the current security situation of the monitored national CIs are required.

[2]We assume here that EOCSSIAN sensors are deployed at each organization's infrastructure and are configured to generate and periodically transmit (after approval) interpreted security data representing the health status of services running at the specific infrastructure they monitor. This data does not need to be further interpreted by the N-SOC.

5 Cooperative Analysis: The Linking Model

In this section we introduce *CAESAIR*: a *Cooperative Analysis Engine for Situational Awareness & Incident Response*. First a theoretical description of the model is given, then details on the architectural components and their functionalities are provided.

5.1 Information Entities

In the following we define the main information entities comprising the data that CAESAIR collects and generates while processing reported security issues.

We define as *Resource* any relevant document collected and stored in the system, such as an incident report from an O-SOC, a security advisory, a forum post or an email message. Resources are not changed over the course of their processing at N-SOC. Both operators and the analysis engine can attribute a Resource to clusters and classes of Resources. The set of existing classes is defined by the N-SOC personnel, while the clusters are discovered by the analysis system during the knowledge base evaluation.

An *Artifact* identifies a certain concept and an unlimited number of its text representations (phrases or regular expressions). The representations are used by the system to detect the concept in free text, e.g. the terms "Windows 7" and "ms-win7" identify the same Artifact *MS Windows 7*. Two Artifacts may build multiple one-sided "is-a" relations, such as *Linux distribution* to *Operating system*, or *Fedora* and *Debian* to *Linux distribution*.

The frequency of Artifacts' occurrences is one of the basic metrics used for estimating similarity between Resources. The more concepts are reflected as Artifacts, the more information is available about each Resource.

A *Tag* is a text label that may be attached to a Resource or Artifact, showing their connection to a certain concept that is not explicitly mentioned in them. For example, reports describing a highly targeted spear phishing attack may be tagged by a N-SOC operator as "suspected APT" and "social engineering", although the terms APT or social engineering do not occur in them. Tags also help the operators to group Resources and Artifacts in an intuitive and flexible way.

5.2 Information Importing

Resources are added to CAESAIR either manually by the O/N/E-SOC personnel or automatically via preconfigured import interfaces, e.g. web crawlers or remote database APIs.

When a new document is acquired by the analysis system its text is first indexed to the search engine; then a Resource object is created in the system, referencing the search index entry. The Resource's text is therefore scanned for occurrences of Artifacts known to the system, or possible new Artifacts. Based on detected Artifacts and the original text itself, the system attempts to attribute the Resource to one of the known classes and clusters. Finally, the new Resource is forwarded for evaluation by an operator, who can confirm or reject the system's suggestions.

Artifacts may be created by the system according to predefined rules, or manually by the O/N/E-SOC operator. All stored Resources will be regularly scanned for occurrences of newly added Artifacts; if an Artifact is deleted, records of its occurrences are also removed.

Furthermore, the system may create new Artifacts if it discovers a phrase matching a certain rule defined by the operator, such as "two words beginning with capital letters". This behavior differs from detecting representations of a single Artifact based on a regular expression: in that case the text match was marked as an occurrence of the existing artifact, and here we create a new Artifact and set its first representation string equal to one that matched our rule.

5.3 Resource Linking Model Definition

CAESAIR's analysis process aims at identifying all existing Resources *linked* to the Resource under examination or to a given text, deriving possible correlations between Resources, and revealing patterns in distribution of Resources over time and locations.

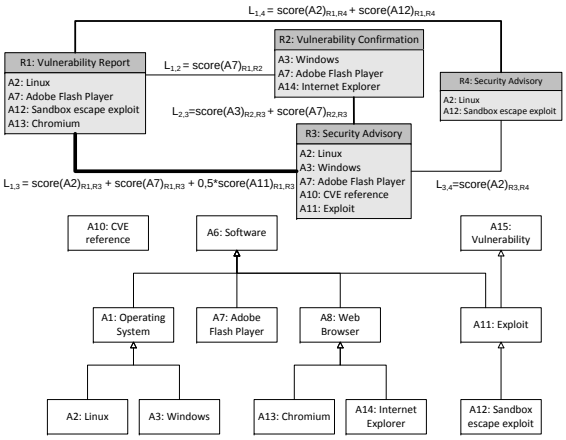

Figure 2: Top: Example of a Resource Linking Diagram. Bottom: Example of an Artifact Relation Diagram

The resource linking process is based on interrelations between Resources, which represent documents containing text, and Artifacts, which represent concepts and their text representations.

If a Resource's text contains at least one representation of a certain Artifact, this Artifact is considered related to the Resource. Each such relation is further referred to as *Occurrence* of an Artifact in a Resource. A Resource may or may not have any arbitrary number of Occurrences of any Artifacts, as long as each Occurrence has at least one unique representation in the Resource's text (it may be a phrase, word or even a single character). Figure 2 depicts an example of Artifacts included in different Resources (upper part of the figure), and how Artifacts can be related to one-another (lower part of the figure).

Given a certain Resource r_1, containing the artifacts a_1, a_2, a_3 and a_4, the algorithm calculates its linkage to the sample Resource r_0 as follows.

Let R be the set of all Resources, and R_{0m} the set of Resources having at least one Artifact in common with r_0. Let us assume $R_{0m} = \{r_2, r_3, r_4, r_5\}$. Let $A_{0,1}$ be a set of Artifacts present both in r_0 and r_1. Let us assume it contains the Artifacts a_1, a_2 and $a_3 : A_{0,1} = \{a_1, a_2, a_3\}$

Now for each artifact a_i in $A_{0,1}$, we define the rating score as:

$$score_{r_{01}}^{a_i} = f_s(TFIDF(a_i, r_1, R_{1m}), TFIDF(a_i, r_0, R_{0m}), freq(a_i, R)) \quad (1)$$

where:

$TFIDF(a,r,Rm)$	function determining the Term Frequency-Inverse Document Frequency weight of an Artifact a in the resource r considering the set of resources havinig mutual artifacts R_m
$freq(a,R)$	function returning the sum of boolean frequencies of the Artifact a in the set R of all Resources
f_s	customizable scoring function

We then determine the linkage between r_1 and r_0 as the sum of scores for each Artifact in $A_{0,1}$:

$$link_{r_{01}} = \sum_{i=0}^{N} score_{r_{01}}^{a_i} + fb \quad (2)$$

where N is the number of Artifacts that occur in $A_{0,1}$, and fb is an optional value representing the operator's feedback on the goodness of the link (its value can be grater or smaller than zero).

Figure 2 also shows how four Resources are linked to one-another. The thickness of the lines connecting the Resources indicates how significant the Resources are to one-another and is proportional to the calculated *link*.

5.4 Knowledge Base Querying

The knowledge base contains the full and up-to-date information available at the N-SOC; search and filtering functions make this information available to the N-SOC personnel and authorized parties.

Two ways of searching are provided:

1. By text: to find Resources containing the specified text. In this case the results are rated by a customizable combination of:
   - number of matches in a single Resource;
   - number of matches in Resource's Tags or related Artifacts;
   - proximity between the query and found matches.
2. By Resource: to find Resources that are *linked* to the specified sample Resource. In this case the search results (Resources) are rated by:
   - text match criteria mentioned in the previous paragraph;
   - number of Artifacts that are present in both the search match and the sample;
   - cross-combinations of matching Artifacts in other Resources: a Resource containing matching Artifacts that do not appear in other Resources may be rated higher;
   - number of Artifacts that do not directly match with those from the sample, but have a common parent node in the is-a hierarchy.

Let us assume there are Artifacts representing some operating systems named *MS Windows*, *FreeBSD* and *Debian*, whereby *FreeBSD* and *Debian* have an "is-a" relation to an Artifact representing *UNIX-like* operating systems. All of the artifacts also have an "is-a" relation with the Artifact "Operating system". Now, if the Resource given as search sample contains the Artifact "Debian", then a Resource containing "MS Windows" will be rated lower than a Resource having "FreeBSD". A Resource that has no common Artifacts with the sample will, as expected, not receive any rating for this, but can still be rated higher than others if it better satisfies the text match criteria.

Filtering allows, instead, to select groups of Resources that exactly satisfy given binary criteria such as presence of specific Artifacts or words in the Resource's text, or being part of a specific class or cluster.

5.5 Statistical Analysis

The population of Resources in CAESAIR's knowledge base comprises all information the N-SOC has about the current security situation. To keep track of this situation, CAESAIR regularly evaluates the knowledge base. First it computes statistics describing a predefined set of Resources' properties, such as frequency of Windows software vulnerability reports, top 100 incident reports sorted by damage level, as submitted by the victim, etc. Moreover it determines Resource clusters (based on quantitative properties of Resources), and then it looks for patterns or anomalies in relations between the Resources. For instance, an unprecedented increase in attack reports submitted by energy sector companies, which the N-SOC personnel noticed anyway, can appear correlated with phishing campaigns recently reported by another federated N-SOC.

6 CAESAIR: System Components and Implementation

In this section we describe the CAESAIR system components and provide details on their current software implementation.

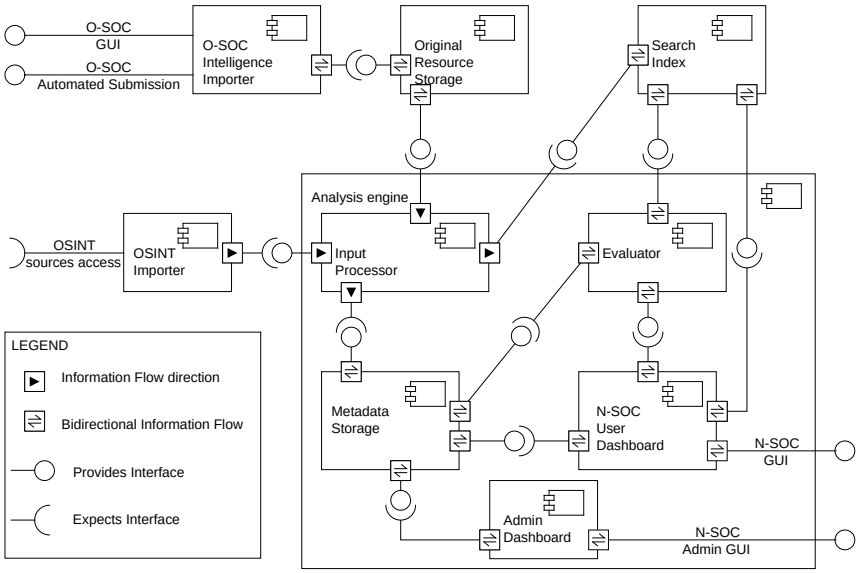

Figure 3: CAESAIR system components.

6.1 System Components

As shown in Figure 3, CAESAIR consists of the following components:

- Importers for both open intelligence (OSINT) and the data generated within ECOSSIAN;
- Original resource (document) storage;
- Search index;
- Analysis engine, comprising:
  - Incoming data processor;
  - Metadata storage;
  - Evaluator;
  - Dashboards for the N-SOC personnel.

CAESAIR operates based on incoming Resources, the basic units of information for the system. The **importers of intelligence data** acquire new documents actively (e.g. by crawling given web resources, databases) or passively from open sources or O-SOCs through a dedicated interface provided by importers, either graphical or non-graphical.

Importers validate, sanitize each document's contents and then forward it to the **original resource storage**, where it will be kept with read-only access for future reference.

A copy of each Resource is saved in the **search index**, from where it can be retrieved by Resource ID, or via full-text search over all properties of a Resource.

The **input processor** collects Resources from the original storage and checks them for occurrences of known Artifacts, or for possible new Artifacts. All detected occurrences and new Artifacts are saved to the **Metadata storage**, whereas the new Resources are forwarded to the search index.

Metadata storage holds all the data produced within CAESAIR, such as: Artifacts known to the system, relations between them and Resources, attribution of Resources to clusters and classes, comments that N-SOC operators add to Artifacts or Resources.

Evaluator is the core component that helps drawing conclusions from the accumulated data: it periodically, in configurable intervals, queries both the Metadata storage and the search index to keep track of interdependencies between Artifacts and Resources. Using the resource linking model discussed in Section V, it also estimates similarities between Resources, attributes them to clusters and suggests classification for them. On request from an N-SOC user the Evaluator also answers the question: "what Resources known to the N-SOC are similar to the given one?". By running the *Linking Model* (see previous section) the Evaluator identifies and prioritizes the Resources with highest pertinence to the Resource currently analyzed.

The **N-SOC dashboard** is the primary way for N-SOC personnel to use the system. It provides a graphical and a programming interface for querying the analysis system and giving feedback on the queries, searching for Resources and monitoring patterns in the accumulated data.

When an N-SOC operator is handling an incoming incident report, the graphical interface will use the Evaluator's search and filtering API to display the list of Resources that seem to be relevant to the current context. The operator may narrow down the results set by setting additional filtering criteria (e.g. only display the Resources added in March 2015 and from two specific sources).

Finally, the **administrator dashboard** allows manually tuning the system's parameters and directly accessing the Metadata storage. It is supposed to be used for maintenance rather than normal work-flow.

6.2 Implementation

CAESAIR's storage components (Metadata and Original Resource Storage) are based on *PostgreSQL*[3] database servers. For the search index we employ *Elasticsearch*[4], an open-source search engine that is optimized for frequently updated dataset and offers a flexible API for querying and indexing data.

The logic components (*Importers*, *Input Processor*, *Evaluator* and *Dashboards*) run by the N-SOC, are developed as Python modules.

O-SOC intelligence importer will function as a service available to O-SOCs via an online ticketing system (*Request Tracker*[5]) hosted at the N-SOC. The service will also provide an API, allowing O-SOCs to interact with CAESAIR using custom software. The open source intelligence importer incorporates Python modules for retrieving and parsing data in various formats from different sources (such as a website, database, file system). It runs on-demand or permanently, pushing new data to the Analysis engine as they appear in the targeted source. The Input Processor is triggered by the Open Intelligence Importer whenever new documents are available; it also actively checks the Original Resource Storage for updates.

The Evaluator runs periodically to recalculate the statistics (information on clusters, classes, patterns detected in Resources and Artifacts) and save it to the Metadata Storage.

The N-SOC user and administrator dashboards, implemented as web applications with Python, serve Javascript-powered client GUI and querying the Metadata storage, the Search index and the Evaluator. They also validate and forward user commands and feedback to the Evaluator, and update the Metadata storage as needed.

[3]http://www.postgresql.org/

[4]https://www.elastic.co/

[5]https://www.bestpractical.com/rt/

Figure 4 depicts a screen-shot of the main view of the N-SOC user dashboard. The left panel lists the Resources to be analyzed. On the right panel a set of Resources, linked to the one currently selected, are displayed.

7 System Evaluation

In order to assess the performance and the soundness of the implemented system we defined relevant quantitative and qualitative metrics and we ran a number of tests to extract measures under different functional conditions. The evaluation setup, the defined performance metrics, the obtained measures and the derived conclusions are reported in this section.

7.1 Evaluation Setup

All the tests have been executed on a machine with Intel(R) Xeon(R) CPU E5-1620 v2 at 3.70GHz 8 cores and 16GB memory, running Ubuntu 12.04.5 LTS operating system.

To our best knowledge no publicly available repository of incident reports exists, hence to measure the performance of our system we imported in the Resource knowledge base elements from *Common Vulnerabilities and Exposures* (CVE)[6] database with entries registered between 2013 and 2015, including 16 thousand elements.

For the sake of simplicity in this work we limited the Artifacts extraction to the identification of only product and platform names within the Resources. We therefore imported the complete Common Platform Enumeration (CPE)[7] repository (counting around 100 thousand entries), and generated an artifact (with its different representation) for each entry of this archive.

7.2 Quantitative Evaluation

To quantitatively evaluate our system we observed its scalability, in terms of resources required to perform its operations, considering a variable number of existing resources and artifacts in the knowledge base. In close cooperation with the ECOSSIAN partners we specifically took care about a realistic setup. We focused on two main timing metrics:

- t_{Aext}: average time required to extract Artifacts (CPE entries) relevant to a Resource r_i

[6]https://cve.mitre.org/

[7]https://cpe.mitre.org/

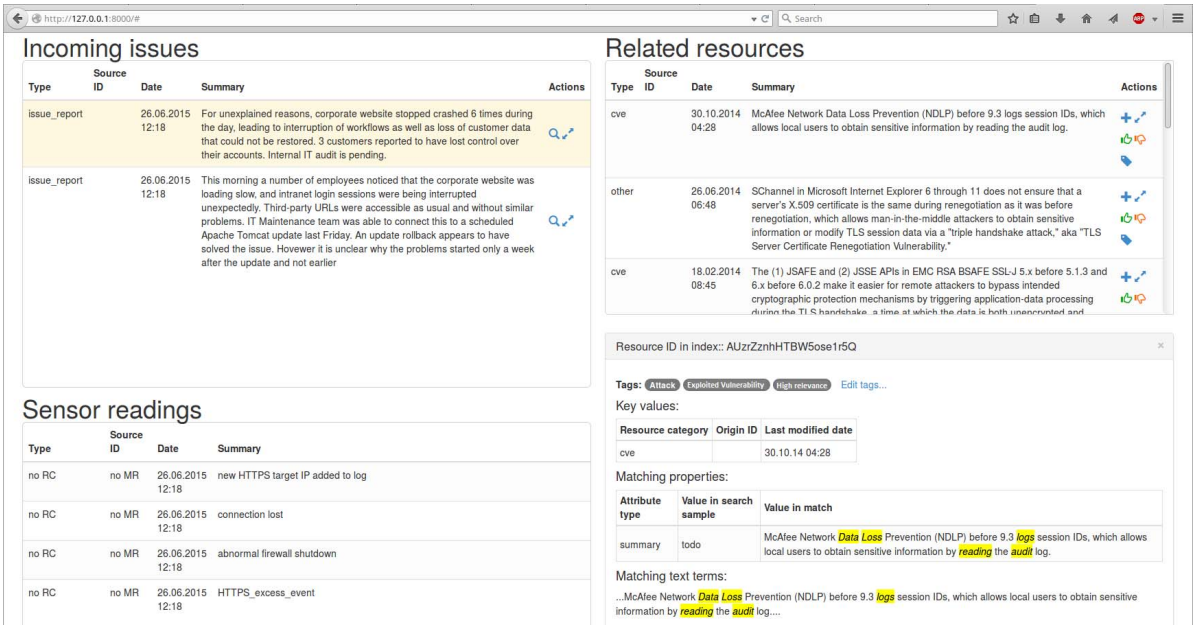

Figure 4: Screenshot of one view of CAESAIR's dashboard.

- t_{Link}: average time to link a Resource r_i to any other relevant Resource r_j in the set R_{im}[8]

Two Resources having at least a common Artifact are linked according to the method described in Section 5.3. The scoring function f_s adopted in our tests is the sum. We did not consider any operator's feedback not to bias the evaluation of the automated linking. Hence given an Artifact a_x present in 2 resources r_i and r_j:

$$score_{r_{ij}}^{a_x} = TFIDF(a_x, r_i, R_{im}) + TFIDF(a_x, r_j, R_{jm}) + freq(a_x, R) \quad (3)$$

and

$$link_{r_{ij}} = \sum_{x=0}^{N} score_{r_{ij}}^{a_x} \quad (4)$$

Where N is the number of Artifacts present both in r_i and in r_j.

During our tests we imported into the system, through the CAESAIR importing component, a variable number of randomly selected CVE entries. We also let the number of imported CPE entries variate, so to determine how the number of existing Resources and Artifacts influences the system's performance.

Figure 5 shows how the average Artifacts extraction time t_{Aext} changes when considering a variable number of Artifacts (CPEs), and a variable number

[8] R_{im} is the set of Resources having at least one Artifact in common with r_i

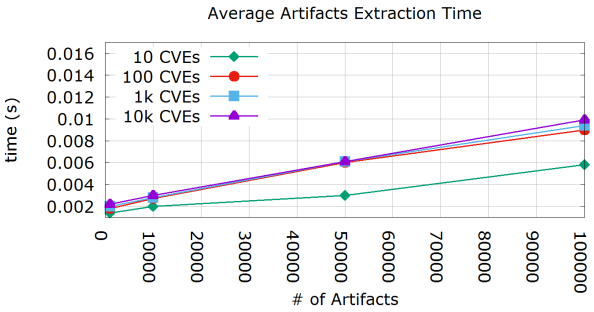

Figure 5: Average Artifacts extraction time vs number of CVE Resources vs number of Artifacts (CPE).

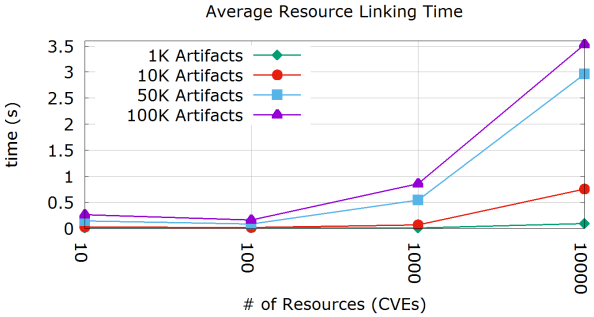

Figure 6: Average Linking time vs number of CVE Resources vs number of Artifacts (CPE).

of Resources (CVEs). Analyzing this graph we can observe that the Artifacts extraction function scales with the number of Artifacts present in the knowledge base. In fact even considering 100 thousand CPEs the average time required to extract all the Artifacts from a Resource is around 10 milliseconds. We can notice

this is true for every number of Resources we considered.

Figure 6 shows how the average Resource linking time t_{Link} changes when considering a variable number of Resources (CVEs), and a variable number of Artifacts (CPEs). Observing this graph we can conclude that, as predicted, the linking method requires longer time when the number of Resources to analyze grows. Moreover the average Resource linking time grows with the number of existing Artifacts. This is due to the fact that the more Artifacts and Resources are present in the knowledge base, the more queries the system has to execute to obtain the Resources related to the one under analysis. Looking at graph in the case of 10 thousand Resources and 100 thousand Artifacts, we can observe how, given a Resource, the linking function needs around 3.5 seconds to obtain all the Resources in the knowledge base having at least one mutual Artifact. This is a relatively short time considering that it corresponds to the time the operator will need to wait after clicking on a Resource to obtain all the related Resources. We foresee that this time could get smaller if we would adopt a less stringent linking method not based on perfect matching of all the Artifacts Resources contain, but on a fuzzy search of only the most relevant Artifacts of the Resource.

7.3 Qualitative Evaluation

In order to perform a qualitative assessment of the designed Resource linking method we adopted the Delphi method (Sackman, 1974). A group of 3 experts with cyber security background and periodically involved in cyber incident handling activities, have been selected and asked to judge the results of our Resource linking method. Statistics on this manual assessment allow a preliminary qualitative evaluation of our approach.

We let the system extract Artifacts (from the 100 thousands imported CPEs) and generate links for a set of 10 thousand Resources (CVEs). We considered a set of 25 different randomly selected linked Resources (i.e., Resources with at least 1 link) and assigned them to each expert. The experts graded the relevance of all the links the system created for each analyzed Resource, with a score S between 0 (not relevant at all) and 5 (very relevant). The average $\bar{S}$ and the median $\tilde{S}$ of the scores obtained by each expert are reported in Table1. The table reports small discrepancies between the evaluation results of the 3 experts. Indeed the average scores variate in a range between 2.8 and 3.28. Moreover, given the experts' evaluation, we can derive from a preliminary and rough analysis that the linking method provides quite accurate matches between existing related Resources. The qualitative assessment proposed here will be further validated in the future by the design and the adoption of an analytical evaluation approach.

	$\bar{S}$	$\tilde{S}$
Expert 1	2.96	3
Expert 2	3.28	4
Expert 3	2.8	3
Average	3.01	3.33

Table 1: Average results of 3 experts' qualitative evaluation. Analysis based on the links obtained on 25 randomly selected Resources.

8 Illustrative Application

In this section we validate the presented approach. We focus on the advantages derived from its deployment in the ECOSSIAN ecosystem to tackle the security issues outlined in the scenario in Section 3.

Considering the scenario, if WonderLight (the targeted company) would be part of ECOSSIAN and it would therefore have a highly qualified O-SOC in place (required to join the ECOSSIAN ecosystem), the O-SOC operators would receive notification by the employees recognizing the suspicious mails, and using the ECOSSIAN platform they would report a security incident to the N-SOC including a copy of the suspicious email.

Let us assume the attack was performed on international scale, and targeted many European small and medium energy providers employing a particular series of PLCs produced by a specific vendor. In this case the CountryX's N-SOC receives security issues not only from the WonderLight's O-SOC, but also from other affected energy organizations in CountryX. The N-SOC uses CAESAIR to optimally prioritize and analyze the acquired reports, and correlating them it derives that a common threat is targeting local energy providers. The N-SOC operators generate hence a security alert and distributes it to all O-SOCs belonging to the energy sector, asking to those which didn't report any issue yet, to inform the N-SOC if they encountered any similar situation.

Some hours later, the CountryY's N-SOC learns, through open source intelligence feeds, about a *zero-day exploit* that takes advantages of a newly discovered shell vulnerability and gains privileged remote access to a system. This N-SOC informs the E-SOC which opportunely forwards this security information to all the other N-SOCs.

In order to check if the suspicious emails were connected to the newly discovered zero-day exploit, the CountryX's N-SOC asks the energy O-SOCs to configure the ECOSSIAN sensors to provide information about IP addresses and ports currently active in their infrastructure. The N-SOC obtains automated data from the ECOSSIAN sensors, and employing CAESAIR is able to discover whether the attacker gained already remote access to the different organizations' infrastructure. Affected organizations would in fact have same suspicious ports and IP addresses in use.

Our system provides a user-friendly graphical interface that allows the operator to quickly browse through relevant resources, conveniently ranked by the analysis engine, and investigate on reported incidents by querying the knowledge base. N-SOC operators are therefore able to easily assess the cyber security situation of the CountryX by analyzing the derived intelligence, and will inform the affected organizations (by distributing tailored advisories to the respective O-SOCs) on how to mitigate the detected threat.

Indicators of Compromise (IoCs) will be generated by the CountryX's N-SOC and they will be sent (within an incident report informing about the detected cyber threat) to the E-SOC. The E-SOC will possibly receive multiple reports related to the same issue from different affected N-SOCs. After acknowledging the incident the E-SOC will distribute the respective IoCs and inform interested N-SOCs about it. Receiving N-SOCs will investigate through the O-SOCs if any organization in their country was affected by the security issue, and will give their feedback to the E-SOC. The information collected at E-SOC will be analysed and will allow to obtain cyber situational awareness at European level.

This use-case demonstrates how the ECOSSIAN ecosystem would efficiently address one of the possible security issues targeting multiple organizations located in different European countries. By cooperatively employing CAESAIR's outcomes the N-SOCs operators are able to quickly assess the severity of distributed threats, contrive possible mitigation steps, allow operators to timely react to the security incident and restrain its impact on their infrastructures.

9 Related Work

Cyber incident handling is becoming a crucial task for SOCs responsible for the protection of modern CIs. Increasing amounts of data needs to be properly collected, processed and analyzed to obtain knowledge on the current security situation and derive proper countermeasures in case of threats.

Regulatory bodies have recently started issuing directives (European Commission, 2013), frameworks (European Commission, 2009), and guidelines (ENISA, 2013c), in order to raise the awareness on cyber threats among CI organizations, regulate on the establishment of required countermeasures, and define roles and responsibilities of the entities involved in handling cyber incidents on national and international level.

Moreover, several solutions have been proposed which aim at efficiently process and share threat information.

In (IBM, 2013) *IBM* presents *X-Force*, a system for automated threat analysis based on dynamic Internet data to gain insight and context in security incidents.

FIDO, an open source system for automatically analyzing security events and responding to security incidents is introduced by *NETFLIX* in (NETFLIX, 2015). *FIDO* is an orchestration layer that automates the incident response process by evaluating, assessing and responding to malware and other detected threats.

A more sophisticated solution is proposed by *Airbus Defence & Space* in (AIRBUS, 2014); *Cymerius* is a decision-making engine which operates within a SOC and permits real-time management of IT system security by taking into consideration the organizational aspects of the monitored infrastructure.

ENISA, in cooperation with a group of four European CERTs, is currently working on an *Incident Handling Automation Project (IHAP)* (ENISA, 2015) which aims at improving the incident handling process by increasing automation and providing an easy-to-set-up and deploy solution for incident response process. In IHAP CERT teams use a unified *Data Harmonization Ontology* to enhance the actionable reporting and analysis of the collected abuse information.

Since 2010 an open community is developing *CRITS (Collaborative Research Into Threats)*[9], an open source malware and threat repository which combines an analytic engine with a cyber threat database. CRITS serves as a repository for attack data and malware, and provides analysts with a platform for conducting malware analysis and correlation of hierarchically structured cyber threat information[10].

All the aforementioned approaches and tools are based on automated processing of incident data and do not foresee any human involvement while handling

[9] https://crits.github.io/

[10] Open source structured data exchange formats such as STIX, CybOX and TAXII are used within CRITS.

security incidents. The system we presented in this paper differs by massively supporting SOCs' analysis tasks with automated information aggregation and correlation, but at the same time involves security analysts in the decision making phase. Since CAESAIR is an assisted-learning system, it also allows analysts to provide feedback on the soundness of the conclusions derived, and therefore train the engine to dynamically adapt to evolving security issues.

10 Conclusion and Future Work

In this paper we introduced CAESAIR, a model for the analysis of cross-organizational cyber incidents that aims at supporting National Security Operation Centers in gaining cyber situational awareness and responding to detected security issues.

Thanks to its self-learning search engine, CAESAIR improves its incident querying capabilities and hence allows N-SOCs to achieve high incident handling rates.

The system performance evaluation proved the feasibility and the scalability of our approach by demonstrating how the main functional tasks provided by CAESAIR can be implemented and executed on real data at a high rate. We eventually provided a preliminary qualitative assessment of the CAESAIR Resource linking function adopting the Delphi method.

Future work include the investigation of advanced Resource linking procedures (e.g., based on Bayesian and neural networks), the development and integration of functionalities for allowing assisted learning through operator's feedback, code optimization in order to guarantee improved scalability, and the design as well as the implementation of an analytical evaluation method for the assessment of the different foreseen Resource linking procedures.

Acknowledgements

This work was partly funded by the European Union through the FP7 project ECOSSIAN (607577) and the Austrian FFG research program KIRAS in course of the project CIIS (840842)

REFERENCES

AIRBUS (2014). CYMERIUS - Security Management Tool. http://www.defenceandsecurity-airbusds.com/web/guest/1299.

ENISA (2013a). Cybersecurity cooperation: Defending the digital frontline.

ENISA (2013b). Detect, share, protect. Technical report, EU Agency for Network and Information Security.

ENISA (2013c). Technical Guideline on Incident Reporting.

ENISA (2015). Incident Handling Automation Project.

European Commission (2009). Article 13a of the Regulatory Framework for Electronic Communications in the European Union.

European Commission (2013). Commission proposal for a directive concerning measures to ensure a high common level of network and information security across the union.

German Federal Office for Information Security (2014). German it crisis management webpage. https://www.bsi.bund.de/EN/Topics/IT-Crisis-Management/itcrisismanagement_node.html.

Government of Canada - Cyber Security (2014). Canadian cyber incidents response center homepage. http://www.publicsafety.gc.ca/cnt/ntnl-scrt/cbr-scrt/ccirc-ccric-eng.aspx.

IBM (2013). Combat the latest security attacks with global threat intelligence.

Kaufmann, H., Hutter, R., Skopik, F., and Mantere, M. (2014). A structural design for a pan-european early warning system for critical infrastructures". In *Elektrotechnik und Informationstechnik*. Springer.

National Cyber Security Center - Ministry of Security and Justice - Netherlands (2014). Dutch cyber security center homepage. https://www.ncsc.nl/english/.

NETFLIX (2015). Introducing FIDO: Automated Security Incident Response.

Reichl, J., Schmidthaler, M., and Schneider, F. (2013). The value of supply security: The costs of power outages... *Energy Economics*, 36:256–261.

Sackman, H. (1974). Delphi assessment: Expert opinion, forecasting and group process. An Experiment in Probabilistic Forecasting.

Settanni, G., Skopik, F., Fiedler, R., and Shovgenya, Y. (2015). A blueprint for a pan-european cyber incident analysis system". In *Proceedings of 3rd International Symposium for ICS and SCADA Cyber Security Research*, pages 84–88.

Tankard, C. (2011). Advanced persistent threats and how to monitor and deter them. *Network Security*, 2011(8):16–19.