Enriching Threat Intelligence Platforms Capabilities

Mario Faiella[1], Gustavo Gonzalez-Granadillo[1], Ibéria Medeiros[2],
Rui Azevedo[2] and Susana Gonzalez-Zarzosa[1]
[1]*Atos Research & Innovation, Cybersecurity Laboratory, Spain*
[2]*LASIGE, Faculty of Sciences, University of Lisboa, Portugal*
{*mario-ferdinando.faiella, gustavo.gonzalez, susana.gzarzosa*}*@atos.net, imedeiros@di.fc.ul.pt, razevedo@lasige.di.fc.ul.pt*

Keywords: Threat Intelligence Platforms, Open Source Intelligence (OSINT), Data Enrichment, MISP, Threat Score.

Abstract: One of the weakest points in actual security detection and monitoring systems is the data retrieval from Open Source Intelligence (OSINT), as well as how this kind of information should be processed and normalized, considering their unstructured nature. This cybersecurity related information (e.g., Indicator of Compromise - IoC) is obtained from diverse and different sources and collected by Threat Intelligence Platforms (TIPs). In order to improve its quality, such information should be correlated with real-time data coming from the monitored infrastructure, before being further analyzed and shared. In this way, it could be prioritized, allowing a faster incident detection and response. This paper presents an *Enriched Threat Intelligence Platform* as a way to extend import, quality assessment processes, and information sharing capabilities in current TIPs. The platform receives structured cyber threat information from multiple sources, and performs the correlation among them with both static and dynamic data coming from the monitored infrastructure. This allows the evaluation of a threat score through heuristic-based analysis, used for enriching the information received from OSINT and other sources. The final result, expressed in a well defined format, is sent to external entities, which is further used for monitoring and detecting incidents (e.g., SIEMs), or for more in-depth analysis, and shared with trusted organizations.

1 INTRODUCTION

The number and the impact of cyber attacks has drastically increased during the last years, as revealed by reports written by governments and companies, especially in terms of how much these threats could harm them from an economical point of view. The Council of the Economic Advisers of the United States[1] estimated that malicious cyber activity had an economic impact in the U.S. economy between 57 billion and 109 billion dollars in 2016 (CEA, 2018). Cybersecurity Ventures[2] identified cyber crime as the *"greatest threat to every company in the world"*, predicting that it will cost the world more than six trillion dollars annually by 2021 (Ventures, 2017). Moreover, the global management consulting firm Accenture[3], during a study conducted in 2017 (Accenture, 2017), affirmed that cyber crime, on an annual average, is costing organizations 11.7 million dollars, more or less 23 percent more than the previous year. These successful incursions potentially allow groups of attackers to acquire valuable intellectual properties and secrets. With the aim of facing these menaces, it is crucial to have timely access to relevant, accurate information about them, for protecting precious internal and sensitive data as well as critical assets.

Collecting and processing Open Source Intelligence (OSINT) information is becoming a fundamental approach for obtaining cybersecurity threat awareness. Recently, the research community has demonstrated that useful information and Indicators of Compromise (IoC) can be obtained from OSINT (Liao et al., 2016; Sabottke et al., 2015). Besides the research oriented efforts, all Security Operation Centre (SOC) analysts get updated about new threats against their IT infrastructures by collecting and analyzing cybersecurity OSINT data. Nevertheless, skimming through various news feeds is a time-consuming task for any security analyst.

Furthermore, analysts are not guaranteed to find news relevant to the IT infrastructure they oversee. Tools are therefore required, not only to collect OSINT, but also to process it, aiming at enhancing the

[1]https://whitehouse.gov/cea/
[2]https://cybersecurityventures.com/
[3]https://www.accenture.com/

quality of the information carried on OSINT to SOC analysts, for instance, to benefit from the potential they have. In addition, such tools must filter only the relevant parts for the SOC analysts, thus decreasing the amount of information and consequently, the time required to analyze it and act upon. When appropriate, the filtered information must be further processed to extract IoCs.

Moreover, a proper quality assessment is needed, to check if gathered data can be considered as valuable Threat Intelligence (denoted as TI). Sillaber et al. (Sillaber et al., 2016) identified TI quality evaluation as one of the main challenges in actual cybersecurity information sharing scenarios, mainly caused by the limitation of existing TI sharing tools, as well as the lack of suitable and globally recognized standards and ontologies (Mavroeidis and Bromander, 2017). These assessment processes can provide more insights for inferring the impact that some cyber attacks could have with respect to internal assets and resources, prioritizing threat detection and incident response.

In addition, the ability to share OSINT information is often not enough. TI must be expressed, and then, shared using specific standards, allowing involved parties to speed up processing and analysis phases of received information, achieving interoperability among them.

In this paper we propose an *Enriched Threat Intelligence Platform*, (hereinafter denoted as `ETIP`), aiming at extending import and information sharing capabilities of internal detection and monitoring systems (e.g., SIEMs) improving also quality assessment of received cybersecurity events.

The final objective is to integrate the relevant security data coming from public sources (e.g., social networks), after going through a quality information enhancing process, with data gathered from the infrastructure through specific detection and monitoring systems (e.g., SIEMs, IDS, IPS), to anticipate and improve threat detection and incident response. This integration has been defined as a crucial activity in order to produce real and valuable TI (Skopik et al., 2016).

In this context, on one hand, it arises the need of a component that relates and aggregates collected OSINT data, generating thus new enriched data. On the other hand, it also requires a component that considers potential security issues in the monitored infrastructure, to be correlated with the received OSINT data, providing a threat score for the latter that helps to identify its relevance and priority.

This threat score will complement the usage of static information about the monitored infrastructure with dynamic and real-time threat information reported from inside the own monitored infrastructure in the way of IoCs. This dynamic evaluation is based on heuristic analysis which allows determining the priority of the incoming OSINT data, by assigning a threat score to it. The produced object integrating the information received from OSINT data sources through its calculated threat score is sent directly to other security systems and tools (e.g., SIEMs) for visualization, storage, processing, or feedback and, optionally, could also be shared with external trusted organizations.

The remainder of this paper is structured as follows: Section 2 presents related works, while Section 3 introduces and compares threat intelligence platforms. Section 4 describes the architecture of our proposed Enriched Threat Intelligence Platform (ETIP). Section 5 details the Threat Score Evaluation process and Section 6 illustrates the applicability of our approach with a use case scenario and preliminary results. Finally, conclusions and perspective for future work are presented in Section 7.

2 RELATED WORK

Several standard formats have been proposed to facilitate cyber intelligence sharing among platforms. Examples of such formats are the Open Indicators of Compromise (OpenIoC[4]), Structured Threat Information eXpression (STIX[5]), Trusted Automated eXchange of Indicator Information (TAXII[6]).

Few studies of existing threat intelligence platforms (TIPs) have been identified. Tounsi and Rais (Tounsi and Rais, 2018) provides a survey about open source threat intelligence platforms, including the Malware Information Sharing Platform (MISP)[7], the Collective Intelligence Framework (CIF)[8], the Collaborative Research Into Threats (CRITs)[9], and Soltra Edge[10]. Sauerwein et al. (Sauerwein et al., 2017), provide an exploratory study of software vendors and research perspectives of threat intelligence sharing platform, and conclude that the market for threat intelligence sharing is still developing. Moreover, also

[4]https://www.darknet.org.uk/2016/06/openioc-sharing-threat-intelligence/

[5]https://oasis-open.github.io/cti-documentation/stix/intro

[6]https://oasis-open.github.io/cti-documentation/taxii/intro

[7]http://www.misp-project.org

[8]https://csirtgadgets.com/collective-intelligence-framework

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

[10]https://www.soltra.com/en/

ENISA provides an updated report about opportunities and limitations of actual TIPs (ENISA, 2017), suggesting various guidelines that should be followed for overcoming them.

Owen (Owen, 2015) proposes Moat, a powerful tool that covers known bad actors and consume data from multiple sources such as vulnerability systems and port scanners. Moat has been integrated with SIEMs using STIX and XML formats for sharing purposes but it is not yet defined for other well-known standards such as TAXII.

Some commercial SIEMs (e.g., LogRhythm[11]) have added security intelligence to its SIEMs and analytic platforms. Their approach uses rich context enabled by threat intelligence from STIX/TAXII-compliant providers, commercial and open-source feeds, as well as internal honeypots. As a result, the platform uses these data to reduce false-positives, detect hidden threats, and prioritize concerning alarms.

To the best of our knowledge, more research is needed about threat intelligence sharing platforms, and their integration with other security tools. Our approach suggests the use of a platform for collecting and aggregating cyber security related information from OSINT, relying on MISP for storing and managing the resultant IoCs, which will be further enriched with a threat score, for prioritizing possible defence actions. The outcome of this platform will feed systems, like SIEMs and IDS, with actionable information that will improve the detection of cyber threats, and could also be shared, in an automated way, with internal SOCs and CSIRTs, as well as with other trusted organizations.

3 THREAT INTELLIGENCE PLATFORMS

Many companies started relying on Threat Intelligence Platforms (TIPs) for overcoming gaps and limitations of actual detection and monitoring systems, especially SIEMs (ThreatConnect, 2018). They are in charge of retrieving structured and unstructured data from diverse external sources, and perform various complex operations, such as filtering, aggregation, normalization, detection, analysis and enrichment, as well as the injection of results into SIEMs. However, their implementation and usage are still in their infancy and, as stated in (Sauerwein et al., 2017), many drawbacks have to be addressed, for instance, in terms of dynamic trust assessment of external sources and advanced analysis capabilities, where manual work is still needed, especially for making the retrieved information effectively actionable.

TIPs are ideal tools for data collection, storage, sharing, and for integration with external entities, that could be other security platforms and tools, as well as specific groups for handling incident response and threat management (e.g., SOC, CSIRTs). Several TIPs are available in the market (most of them under commercial license). In terms of open-source solutions, we have identified the following:

1. The Malware Information Sharing Platform (MISP),
2. The Collective Intelligence Framework (CIF),
3. The Collaborative Research Into Threats (CRITs), and
4. SoltraEdge (SE), but only a limited version is available with this kind of license.

The comparison among them is summarized in Table 1, and has been performed taking into account the following criteria, mainly based on the study conducted by Tounsi et al. (Tounsi and Rais, 2018), together with some personal considerations, especially about hardware requirements:

- **Import/Export Format:** MISP and CRITs are able to work with a great number of formats (e.g., PDF, doc, xls, txt, JSON, XML, STIX). MISP supports an ad-hoc standard for representing Threat Intelligence (a customized JSON[12] format), and basic built-in capabilities for STIX v.2.0[13] conversion. It also allows adding modules for ad-hoc importing/exporting without modifying the core functionalities. CIF is not as flexible as the previous two TIPs, especially if specific standards (e.g., STIX) will be considered, while the free and limited version of SoltraEdge presents some limitations when dealing with non-STIX data.
- **Integration with/Export to Standard Security Tools:** MISP allows an easy interaction with Intrusion Detection Systems (IDSs) and SIEMs, and contains flexible REST APIs for integrating internal solutions with the platform. CIF is also a viable platform to integrate with IDS and SIEMs, although less flexible than MISP. CRITs is a huge repository of TI, not specifically built for interacting with systems such as SIEMs and IDS, however its flexibility allows building ad-hoc solutions for these purposes. The free SoltraEdge ver-

[11] https://logrhythm.com

[12] JavaScript Object Notation, https://www.json.org

[13] https://oasis-open.github.io/cti-documentation/stix/intro

sion has many limitations, especially in terms of API support for interacting with other platforms.

- **Support of Collaboration:** MISP allows centralized support, by sharing the same instance among a trusted community; and decentralized support, when multiple instances interact in a peer-to-peer way. CIF allows the usage of a private instance and the implementation of a shared instance through a centralized service. CRITs and SoltraEdge allow the usage of a private instance, or a shared one in the context of a trusted community. However, CRITs has very poor built-in sharing capabilities.
- **Data Exchange Standards:** MISP and CRITs are able to deal with many different standards, including STIX and TAXII. The limited version of Soltra-Edge has very poor capabilities to deal with standards different to STIX and TAXII. CIF, has been designed to work with other CIF instances using private solutions for meeting high performance requirements with partial or no support on standards such as STIX and TAXII.
- **Analysis Capabilities:** high analysis capabilities is a weakness for all current TIPs (Sauerwein et al., 2017). CRITs and SoltraEdge are huge central repositories for collaborating analysis rather than pure sharing platforms. They have better built-in analysis capabilities compared to MISP and CIF. However, this advantage is partially lost with the limited version of SoltraEdge.

Table 1: Comparison of Threat Intelligence Platforms.

Evaluated Criteria	MISP	CIF	CRITs	SE
Import/Export Format	●	○	●	–
Integration Capabilities	●	●	○	○
Data Exchange Std.	●	○	○	○
Support of Collaboration	●	●	○	○
Analysis Capabilities	○	○	●	○
Graph Generation	○	○	●	○
License	●	●	●	○
Hardware Requirements	●	–	●	●

–Low/Basic ○ Medium/Average ● High/Advanced

- **Graph Generation:** visualization capabilities are strictly related to the analysis of the aforementioned features, and the same consideration can be deducted. More generically, this is another limitation of current TIPs (Sauerwein et al., 2017).
- **License:** all TIPs considered in this analysis (including the limited version of SoltraEdge) are released with open source licenses.
- **Hardware Requirements:** MISP, CRITs and SoltraEdge have very similar requirements in terms of RAM and hard disk size needed. CIF, instead, has higher requirements, especially in terms of processing capabilities.

Due to its ability to be integrated with SIEMs and IDSs; its high flexibility features for integrating internal and custom solutions; the support of specific data exchange standard, such as STIX, as well as good built-in information sharing capabilities; the availability of a very detailed on-line documentation[14]; and a huge and responsive on-line community, in case of development issues; the selected threat intelligence platform is the MISP.

4 ENRICHED THREAT INTELLIGENCE PLATFORM ARCHITECTURE

The Enriched Threat Intelligence Platform (ETIP) is composed of two main modules: (i) the Composed IoC Module, and (ii) the Context Aware Intelligence Sharing Module as shown in Figure 1.

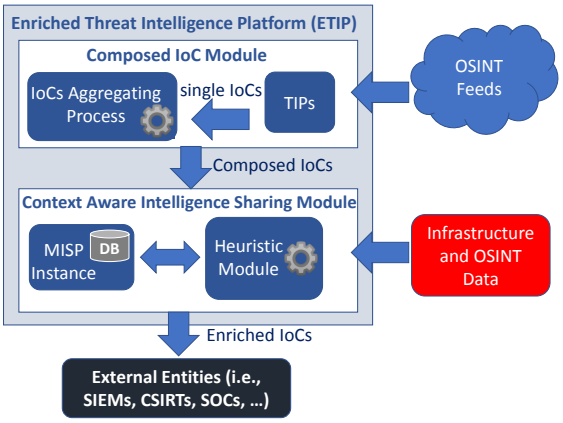

Figure 1: Enriched Threat Intelligence Platform architecture.

While the first module collects several security events (i.e., IoCs) provided from different OSINT feeds, and then interrelates them, generating thus IoCs with more information – *composed IoCs* –, the second module receives these composed IoCs and correlates them with information provided by security tools deployed in the organization network infrastructure. Applying a heuristic analysis to these data, the resulting IoC can be further enriched, providing more insights about how much the incoming information could be considered as real intelligence by the enterprise.

[14]https://www.circl.lu/doc/misp/book.pdf

4.1 Composed IoC Module

We propose a component to generate composed IoCs, i.e., for collecting different IoCs from different TIPs, correlating them, and then, generating new composed IoCs. The architecture of the component is represented in Figure 2. The main parts are the following:

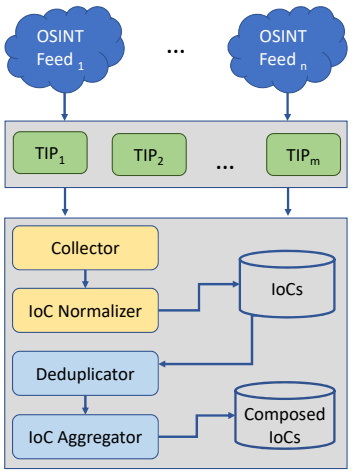

Figure 2: Composed IoC Module Architecture.

- *OSINT Feeds.* Open source intelligence information regarding to security events, such as cyber-attacks, exploitation of software vulnerabilities, malware domains, IP blacklists.
- *TIPs.* Different TIPs are used in parallel to collect several OSINT data provided from diverse feeds, which take advantage of the enrichment capabilities they offer, such as improving OSINT threat intelligence with external data not included in OSINT feeds (e.g., asn source, whois).
- *Collector.* The output of the different TIPs, as a form of IoCs, is channelled to a collector module. The TIP's IoCs are seen as OSINT feeds but in an IoC format (e.g., STIX, MISP format).
- *IoC Normalizer.* Since IoCs might be collected in different formats (depending on the format adopted by TIPs), it is necessary to normalize them in a single and common format (e.g., MISP format). After this process, they are stored in a database to be processed by the component.
- *Deduplicator.* IoCs received from different TIPs can be equals, since TIPs can be configured with the same OSINT feeds. The deduplicator module analyzes the received IoCs with those already exist in the database with the aim to identify duplicated IoCs and remove them before being processed by the IoC aggregator module.
- *IoC Aggregator.* Aggregates different but related IoCs, and generates new ones. The process consists on identifying IoCs that contain relevant interrelated information, aggregating them in a same set, and then, merging that information into a single IoC, creating a new IoC that we call *composed IoC*. These new IoCs are stored in the database for later be used by the context aware threat intelligence component (see Section 4.2).

4.1.1 Implementation

The proposed component was implemented using the MISP platform, for which the deduplicator and aggregator modules were developed in Python 3, and integrated in MISP. This latter acts as the collector and IoC normalizer to receive and normalize the data that can be provided from different TIPs, and then uses the developed modules to process the received data.

The component offers to the end user the possibility of configuration based on two criteria: (1) the trust level of the IoC assigned by the MISP community, where, for example, an IoC with level 2 means that IoC has the trustiest level of confidence and its information is relevant; (2) the interrelation type between IoCs which will be considered by the IoC aggregator module. This interrelation can be based on IoC as a whole or their attributes, which the former only allows interrelations between IoCs that belong to the same threat category, whereas the latter permits a most deep analysis and connection between IoCs, allowing to generate new data provided by IoCs of different categories (e.g., IoCs belonging to the MISP *network* category and from type of *vulnerability*).

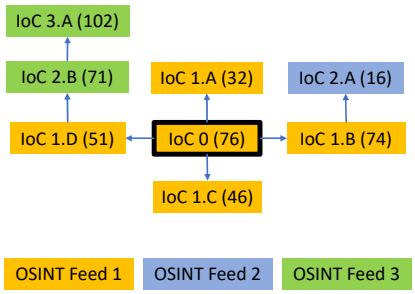

Figure 3: Schematic for the creation of a composed IoC.

The component considers dividing the OSINT feeds in two categories: (1) low level feeds, which consist mainly of IP addresses and URLs; and (2) high level feeds, which contain a more advanced analysis with information about network artifacts, campaigns, etc.; that feed TIPs (e.g., CRIT, MISP). It performs queries to the database to identify new entries and other entries that have matches, and then merge them

forming a new IoC and inject it into the database, which is labelled with a tag that allows identifying it as a rich IoC and avoids the creation of loops.

Figure 3 exemplifies a composed IoC. In the figure, starting from an IoC that contained 76 elements, we were able to identify 7 other IoCs, originating from 3 distinct OSINT feeds, that were correlated. The merging of these IoCs allowed the creation of a new IoC containing 468 elements.

4.2 Context Aware Intelligence Sharing Module

The context aware intelligence sharing module is able to correlate static and real-time information (e.g., Indicators of Compromise), related to the monitored infrastructure, with data coming from external OSINT sources through OSINT data fusion and analysis tools, to check the relevance and accuracy of the data. Furthermore, the module is also able to share both the original and the enriched information with external entities, in an automated way.

The proposed module architecture, depicted in Figure 4, is composed of two main elements: (i) a MISP Instance, in charge of gathering data from both OSINT-based sources and the monitored infrastructures, as well as sending the *enriched IoCs* to internal components, systems and tools (e.g., SIEMs) or sharing them with trusted organizations; and (ii) the Heuristic Component, in charge of performing the heuristic analysis, with the final aim of computing a Threat Score, enriching the data coming from OSINT-based sources, and sending it back to the MISP Instance.

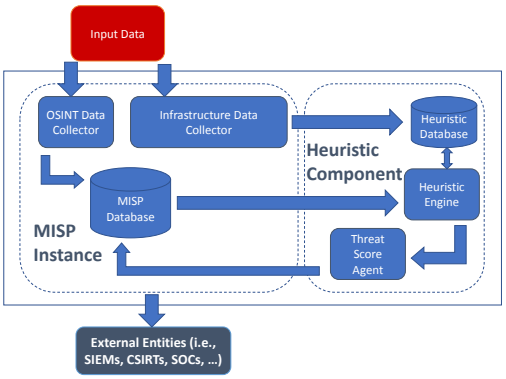

Figure 4: Context Aware Intelligence Sharing Module.

4.2.1 MISP Instance

From Figure 4, the OSINT and Infrastructure data collectors are responsible of capturing useful data from OSINT, IoCs and the infrastructure in order to evaluate a set of pre-defined heuristics and to compute a threat score. Collected OSINT data are stored in the MISP Database, whereas collected infrastructure data are stored in the Heuristic Component Database.

The integration between security tools, as well as internal SOC and CSIRTs, and the context aware intelligence sharing module is possible thanks to the adoption of MISP. The objective is to use, as much as possible, the built-in sharing capabilities of the platform when this interaction takes place, such as a zeroMQ publish-subscribe model[15]. MISP comes with so-called "MISP-modules", used both for ad-hoc import and export of threat information. If required, new modules could be created from scratch and integrated with the MISP Instance, without modifying the core functionalities of the platform.

The heuristic component receives all data coming from the monitored infrastructures, through MISP. Data could be dynamic (e.g., IoC detected in the infrastructures) or static and generic information about a specific infrastructure (e.g., used sensors, operating systems, specific lists of IP addresses). These data are stored in the MISP database, represented through the JSON format (e.g., STIX, MISP events), or through simple documents related to generic information. Since its usage is of great interest to the heuristic component, data could be also stored in a different way, using for instance a private non-relational database such as MongoDB[16], which simplifies the information retrieval by the heuristic engine and allows for a full control of the analysis performed by the tool.

The adoption of MISP makes it possible to automatically share data with external entities thanks to its built-in information sharing capabilities. For those cases in which the external entity is using a MISP instance, the sharing process is performed by simply synchronizing both instances. Otherwise, MISP comes with a list of REST APIs, which are accessed from any internal and external services with different levels of access rights, to directly interact with its database, to push/pull cyber-security related events.

4.2.2 Heuristic Component

The heuristic component receives information coming from multiple sources (e.g., OSINT data, infrastructure, IoCs, etc.) to be used in the Threat Score (TS) analysis performed by the heuristics engine. This latter considers a set of conditions that are evaluated for every single feature. A score is assigned to

[15] http://zeromq.org/

[16] https://www.mongodb.com/

every feature (i.e., individual score). The sum of all individual scores results into the Threat Score associated to the data being analyzed.

The database from the heuristic component stores the information of the infrastructure and the OSINT data collector. The Threat Score Agent is responsible for the generation of the resulting enriched Indicator of Compromise (eIoC), including the Threat Score for the security information received from OSINT data sources. The eIoCs shared by this component includes the same information received from OSINT, as well as the associated Threat Score and the features considered in the evaluation.

5 THREAT SCORE (TS) EVALUATION

The threat score evaluation is part of the heuristic component that uses a threat score function (detailed in Section 5.1) to compute the relevance of the received data. The process performs an analysis methodology composed of the following steps:

1. **Source Identification:** during this phase we search and identify all possible sources of information. Examples of these sources are: security logs, databases, report data, OSINT data sources, IoCs, etc.
2. **Heuristics Identification:** different features (e.g., heuristics) are identified from the input data. Such features provide relevant information about the infrastructure (e.g., vulnerabilities, events, faults, errors, etc.) useful in the threat analysis and classification process. Examples of heuristics are: CVE, IP source, IP destination, port source, port destination, timestamp, etc.
3. **Threshold Definition:** for each heuristic, minimum and maximum possible values are defined based on characteristics associated to the instance. We checked, for instance, if the input data contains or not a CVE for the detected threat. A threshold (e.g., 0-5) is assigned to cover all possible results.
4. **Score Computation:** for each possible instance of the identified heuristic, a score value is assigned based on expert knowledge. All individual scores are then aggregated and a final score is computed. The resulting value will indicate the priority and relevance of the security information coming from OSINT data sources and the monitored infrastructure.
5. **Training Period:** a set of preliminary tests need to be performed during a training process to evaluate the performance of the engine. The tests include real data to analyze the score obtained individually (for each heuristic) and globally (for the whole event) which help to analyze false positive and negative rates.
6. **Engine Calibration:** in order to minimize deviations (e.g., reduce number of false positive, false negative) the engine must be calibrated by analyzing the obtained results, adding other heuristics and/or modifying the assigned values to current attributes.
7. **Final Tests:** Once the engine is calibrated, we can repeat previous tests or add new ones in order to evaluate the performance of the tool.

5.1 Threat Score Function

The heuristics-based threat score is composed of a set of individual scores as a complement of other prediction tools to indicate the priority and relevance of incoming security information received from OSINT and infrastructure data sources. There exists a large number of different aggregation operators (e.g., Arithmetic Mean, Geometric Mean, Harmonic Mean, Weighted Mean, Ordered Weighted Averaging) (Ravana and Moffat, 2009; Torra and Narukawa, 2007) that can be used for the computation of the threat score. They differ on the assumptions about the data (data types) and the type of information that we can incorporate in the model.

From the aforementioned aggregation operators, the Weighted Mean (WM) is the selected function to compute the threat score, due to the following advantages: (i) simple and straightforward function; (ii) avoids indeterminate results and/or null values; (iii) can be used if one or more individual scores are zero; and (iv) individual scores are assumed to have different weights depending on the source and the relevance of the information.

The proposed Threat Score (TS) is defined as the sum of all individual heuristic values (X_i) times its corresponding weight factor (P_i). This latter considers multiple criteria (e.g., relevance, accuracy, timeliness, variety). The sum is then affected to the completeness parameter (C_p), as shown in Equation 1.

$$TS = C_p \times \left(\sum_{i=1}^{t} X_i \times P_i \right) \quad (1)$$

The resulting TS ranges from zero to five ($0 \leq TS \leq 5$), the higher the TS value, the more reliable the IoC. Thus, a TS with a value between zero and one

($0 \leq TS \leq 1$) indicates a Very Low level of priority; a TS with a value between one and two ($1 \leq TS \leq 2$) indicates a Low level of priority; a TS with a value between two and three ($2 \leq TS \leq 3$) indicates a Medium level of priority; a TS with a value between three and four ($3 \leq TS \leq 4$) indicates a High level of priority; and a TS with a value between four and five ($4 \leq TS \leq 5$) indicates a Very High level of priority.

5.2 Heuristic Value

The first part of the (TS) function refers to the value assigned to a given heuristic (X_i) based on the type of information processed during the evaluation. Considering, for instance, that one of the features to be evaluated is the presence of a Common Vulnerability Exposure (CVE)[17] identified in the input data, the engine will check if the word 'CVE' appears in the input data in order to retrieve the complete CVE number (i.e., CVE-AAAA-NNNN).

If a CVE is found, the engine checks for its associated Common Vulnerability Scoring System (CVSS)[18]. More specifically, the engine will search for its associated base score, which considers access vector, access complexity, authentication, and impact related information based on availability, confidentiality and integrity. Depending on the CVSS score, the vulnerability is labeled as none, low, medium, high or critical, as shown in Table 2.

Table 2: Common Vulnerability Score System (CVSS) v3 Ratings.

Severity	None	Low	Med	High	Critical
Lower Bound	0.0	0.1	4.0	7.0	9.0
Upper Bound	0.0	3.9	6.9	8.9	10.0

Source: https://www.first.org/cvss/specification-document

Each evaluated feature is assigned an individual score based on the defined threshold (e.g., from 0 to 5) that will indicate the level of impact of the feature with respect to the event. We define the variable "Score_CVE" that will compute the individual score value assigned to the presence of a CVE in the input data based on the conditions described in Table 3.

Other features (e.g., source/Destination IP, creation and validity timestamps, etc.) may use positive and/or negative values in the assignment process. Such individual values are then tuned in the training and calibration processes so that the final threat score reduces the number of false positives and negatives.

[17] https://www.mitre.org

[18] https://www.first.org/cvss/specification-document

Table 3: Examples of Individual Threat Score.

Criteria	Condition	Score
No CVE	If CVE == ''	0
CVE exists with CVSS 'none' or 0.0	If CVE $\neq$ '' & CVSS = 'none' | CVSS = 0.0	1
CVE exists with CVSS 'low' or less than 4.0	If CVE $\neq$ '' & CVSS = 'low' | CVSS $\leq$ 4.0	2
CVE exists with CVSS 'medium' or less than 7.0	If CVE $\neq$ '' & CVSS = 'med' | CVSS $\leq$ 7.0	3
CVE exists with CVSS 'high' or less than 9.0	If CVE $\neq$ '' & CVSS = 'high' | CVSS $\leq$ 9.0	4
CVE exists with CVSS 'critical' or less than 10.0	If CVE $\neq$ '' & CVSS = 'critical' | CVSS $\leq$ 10.0	5

5.3 Weighting Criteria

The second part of the (TS) function corresponds to the weighting criteria (P_i). According to Henry Dalziel (Dalziel, 2014), Threat Intelligence refers to specific information that must meet three specific criteria: (i) it must be relevant, for the entity who receives it, (ii) actionable and (iii) valuable, from a business perspective. In (ENISA, 2015) the concept of "actionable information" is defined by the European Union Agency for Network and Information Security (ENISA), from an organization point of view as the information that can be used immediately for specific and strategical decision making. Considering (ThreatConnect, 2018) and (ENISA, 2015), in order to be "actionable", information must meet the following criteria:

Relevance: it must have some impacts on specific receiver's assets, such as networks, software and hardware. That is, indicators of compromise will usually be considered relevant when a threat could affect the monitored infrastructure. In order to determine the relevance, it is crucial to determine types of threats targeting your assets/systems, considering real-time information (e.g., IoC), from many internal sources, because they are able to provide dynamic and continuous information about current internal monitoring operation, together with a global view of the infrastructure status.

In our analysis, this criterion evaluates if the information associated to a given attribute is useful to identify a threat. Relevance is computed as shown in Table 4.

Table 4: Relevance Criteria.

Relevance	Score
Attribute with no data	0
Optional Attribute	1
Attribute does not identify threat but helps in the analysis	5
Attribute is useful to identify threat	7
Mandatory attribute to identify threat	10

Timeliness: threat intelligence is more reliable when it allows detecting attacker's activity, especially during the same intrusion, to monitor how it evolves during time. Moreover, information about events older than a few hours are, most of the times, irrelevant and non-actionable due to the dynamic nature of some threat's characteristics, considering that some of them are discovered and analyzed months after the initial compromise.

In our analysis, this criterion evaluates if a detected event is related to an already detected one, by the infrastructure or by the OSINT-based components, and if for instance, such events refer to the same threat, but with a different level of intrusion, providing new or updated information. Timeliness is computed as shown in Table 5.

Table 5: Timeliness Criteria.

Timeliness	Score
Attribute with no data	0
Attribute has never been seen	1
Attribute has been seen with the same value	5
Attribute has been seen with a different value	10

Accuracy: the receiver side should be able to process the received data as soon as possible. It depends mainly on three factors, which are the confident of the source from which data is retrieved, the trust level placed in those sources (which, in turn, could depend on factors such as false positives/false negatives rates) and the local dynamic context of the receiver. The latter is crucial in order to avoid inaccurate results and efforts when dealing with incident response.

In our analysis, information coming from OSINT-based components will be compared to the information coming from the infrastructure, if there is a match of one or more attributes, a score will be computed. Accuracy is computed as shown in Table 6.

Table 6: Accuracy Criteria.

Accuracy	Score
Attribute with no data	0
Attribute has some data with no match	1
There is a match of one source and the infrastructure	5
There is a match of two sources and the infrastructure	10

Variety: detection and prevention should not rely on a single technique or tool. It is crucial to use a combination of systems, tools (e.g., IDS, IPS and Firewalls) and sources (e.g., OSINT), especially when they are able to detect the threat at different levels of intrusions (kill chain phases).

In our analysis, this criterion evaluates the sources from where the information is originated or detected e.g., infrastructure, OSINT-based components. Variety is computed as shown in Table 7.

Table 7: Variety Criteria.

Variety	Score
Attribute with no data	0
Data come from one source	1
Data come from two sources	5
Data come from all sources	10

Ingestibility: received information must be easy to ingest into internal data management systems for further processing and analysis phases. This is achievable using specific standards for representing this data, allowing the receiver to process data as fast as possible, helping also security analysts, as well as through the usage of specific transfer protocols for sharing the related intelligence.

Ingestibility is not considered in our analysis since we are assuming that all received data is expressed in a structured way and uses a specific standard format to be processed in the system. The data collection will be handled directly by the MISP instance. This criterion would have been meaningful in case of reception of unstructured information, but this scenario is not considered by the context aware OSINT platform. The analysis will focus on other criteria with the possibility of adding new ones in the future.

Completeness: Threat intelligence should provide valuable and complete information to the final receiver, evaluated from the local cyber context point of view of the latter. Sometimes, sources are incomplete when considered alone, but their provided data

become actionable once combined or processed with other internal data available to the destination or received from other external sources.

In our analysis, this criterion is used as an overall assessment of the heuristic and not for individual score evaluation of the attributes. Each heuristic is composed of one or more attributes (e.g., CVE is composed of six attributes: (i) no_cve, (ii) cvss_none, (iii) cvss_low, (iv) cvss_medium, (v) cvss_high, (vi) cvss_critical. Completeness is measured as the number of attributes with a non-empty value over the total number of attributes, as shown in Equation 2.

$$C_p = \frac{Non\_Empty\_Attributes}{Total\_Attributes} \quad (2)$$

6 PRELIMINARY RESULTS

In this section we will illustrate the advantages of our approach, and the benefits our platform will bring in terms of prioritizing threat information received from external sources (e.g., OSINT). We will use the Context Aware Intelligence Sharing Module to evaluate relevance, accuracy, and other features on the information received from the Composed IoC Module. For the threat score evaluation we will consider both: the data coming from the infrastructure through security systems and tools (e.g., SIEMs, IDS), as well as, IoCs obtained by open source and public feeds. The resulting threat score will be inserted in the information associated to the analized IoC. The higher the threat score value, the higher the priority of the associated information when handled by incident response teams and security analysts.

The Composed IoC module will provide both: single and composed IoCs to the Context Aware Intelligence Sharing module. We expect that the composed IoCs will have an associated threat score higher than the threat score of each single IoCs they contain. The entire process considered in this use case is characterized by four sequential phases: **Collecting** and **Aggregating** phases, performed by the Composed IoC module, followed by the **Sharing** phase, which involves both modules, and the **TS Evaluation** phase, completely handled by the Context Aware Intelligence Sharing module.

Collecting Phase: to collect OSINT data we configured a MISP instance with 34 OSINT feeds from higher value information (e.g., CVE vulnerabilities) and low value information (e.g., IP blacklists). These feeds are provided by diverse public free entities and reach MISP in different formats, such as csv and txt files. OSINT data are normalized in a single format, namely the MISP format, and then stored as IoCs in the MISP database. Afterwards, the deduplicator module we developed is executed to load the IoCs and search for duplicates in order to delete them. This task allows improving MISP in two forms: identify duplicated IoCs and reduce the quantity of data stored, and therefore, increasing the MISP performance.

Aggregating Phase: After removing the duplicated IoCs, the aggregator component analyses the resulting IoCs to look for connections among them. For the connections found, the aggregator puts the involved IoCs in the same group of IoCs, since they are related to the same malicious threat. At the end, we have several and different groups of IoCs forming clusters, each one for a particular threat category. At the point of view of a threat category, a cluster can contain IoCs correlated between them and related with a same (sub-)threat (or attack) and possibly with other valuable malicious information that can be provided in a same IoC. This means that a cluster can contain sub-clusters of IoCs regarding to different attacks. Such sub-clusters can well characterize, this point of view, attacks that have been executed, for which individual IoCs could not allow their identification. Finally, each sub-cluster is represented as one IoC, i.e., all its IoCs are merged in a single one, generating a composed IoC, and then, they are stored in the MISP database.

Sharing Phase: the final outcome of the composed IoC module is sent to the context aware intelligence sharing module for proceeding with the final computation of the threat score. This integration is achieved easily thanks to the adoption of MISP.

More precisely, two different MISP instances are used, one by each module. For simplicity, and for facilitating reader comprehension, we will refer to them as $MISP_A$ and $MISP_B$, respectively. These instances have been synchronized between each other to allow a real-time and completely automated information sharing, following the guidelines provided in the MISP book[19] for setting up a MISP synchronization server on $MISP_A$. This server has been associated to a specific user with synchronization privileges, which is replicated in both instances. Injecting, and publishing IoCs in $MISP_A$ on behalf of this user, triggers a push operation of one or more IoCs directly on $MISP_B$, completing the one-way information sharing needed for this use case scenario.

When the synchronization server is set up, the sync user authentication key must be specified. This

[19] https://www.circl.lu/doc/misp/book.pdf

information is provided by $MISP_B$, when the user is created.

TS Evaluation Phase: to perform the computation of the threat score, $MISP_B$ needs to be extended with new functionalities. For this specific use case, we developed a new MISP export module, integrating it with the core software of the platform, following the guidelines[20] provided by the MISP community and developers. In this way, the module is available directly from the MISP UI, where it can be triggered manually by the user for a specific IoC, to retrieve and send the IoC directly to the Heuristic Module of the platform (Figure 4), where the threat score function has been implemented, and added to the original IoC as a new MISP attribute.

Aiming at correlating IoCs with infrastructure data, as well as with other useful information about cyber events from open source and public feeds, a MongoDB[21] database is used, and the information is stored as JSON documents. The final IoC (i.e., enriched IoC) could be shared with specific security tools or internal SOCs and CSIRTs, with the additional threat score used for determining the priority of the contained data, in case of some defence activities would be needed.

In order to provide practical examples of the functionality of our platform, eleven samples of composed IoCs have been considered, and the entire process previously described has been executed in each of them. As a result, the threat score (TS) is computed for every single IoC ($sIoC$) integrating the composed ones, making it possible to compare with the threat score computed for the composed IoCs ($cIoC$).

For the heuristic analysis, a subset of nine MISP attributes[22] has been selected, composed of the ones which are more relevant according to the monitored infrastructure (i.e., vulnerability, filename, ip-src, ip-dst, hostname, domain, url, link, and md5). This does not mean that other attributes are discarded, they simply have a higher importance when specific criteria are evaluated, especially for relevance and completeness.

Results are summarized in Table 8. Each row is associated to a sample of composed IoCs (e.g., S_1, ..., S_{11}), specifying the number of single IoCs ($sIoC$) composing them, their individual Threat Score (TS), and the global TS associated to the composed IoC ($cIoC$).

More information about our platform can be found in https://caisplatform.wixsite.com/english.

[20]https://github.com/MISP/misp-modules

[21]https://www.mongodb.com/

[22]https://www.circl.lu/doc/misp/categories-and-types/

Table 8: Threat Score Results.

Samples	N. of sIoCs	individual TS	Global TS
S_1	5	1.86, 2,55, 1.80, 0.71, 1.94	3.18
S_2	3	1.43, 2.32, 1,58	2.53
S_3	3	2.48, 1.54, 1.09	2.87
S_4	6	1.18, 1.40, 1.54, 0.64, 1.41, 2.03	3.07
S_5	2	2.84, 1.66	3.22
S_6	17	1.39, 2.22, 2.21, 1.99, 1.87, 0.70, 1.66, 1.10, 0.56, 0.96, 0.94, 0.56, 1.58, 2.66, 2.27, 1.36, 1.08	3.98
S_7	7	2.09, 3.27, 1.89, 0.89, 2.88, 1.93, 1.66	2.84
S_8	4	3.06, 2.68, 2.11, 1.55	3.11
S_9	11	1.66, 1.21, 2.35, 1.92, 1.33, 1.29, 1.6, 0.90, 0.88, 1.02, 0.56	4.13
S_{10}	2	2.43, 2.31	2.54
S_{11}	2	0.99, 0.55	1.29

As depicted in Table 8, in most of the cases, the TS of the composed IoCs is higher than the TS for each of the related single ones. This improvement is strictly dependent on two main factors:

1. The number of attributes present in the IoC. The higher this number, the higher the probability of increasing the overall quality when the aggregation is performed; and
2. The quality of the single IoCs. The higher the quality of the information found in the attributes present in the sIoCs, the lower the probability to increase the overall quality when aggregating several IoCs.

For the first factor, we have seen samples S_6 and S_9 with a high number of single IoCs (17 and 11 IoCs respectively), for which the TS of their $cIoC$ has greatly improved compared to the one of each $sIoC$. For the second factor, we have seen samples S_2, S_3, S_7, and S_{10}, in which the aggregation process is not able to add a relevant level of quality to the final IoC. In these cases, the quality of the information identified in the $sIoC$ samples results in high TS values. Although in most of the cases, the TS value of the $cIoC$ is higher than the one associated to each $sIoC$, the improvement

is low, having in some cases a lower TS value in the *cIoC* compared to one of the *sIoC* (i.e., S_7).

It is important to note that among all the criteria used in the TS computation, the completeness (i.e., C_p) is the criterion that affects the most the final result. Whereas, all other criteria are adding individual values to the TS, the completeness criterion is multiplying to the overall addition, affecting to a higher level the TS results of single or composed IoCs.

7 CONCLUSION

This paper presents *ETIP*, an enriching threat intelligence platform, as an extended import, quality assessment processes and information sharing capabilities in current TIPs. The proposed platform gathers and processes structured information from external sources (e.g., OSINT sources) and from the monitored infrastructure. The platform is composed of two main modules: (i) a Composed IoC Module, in charge of collecting, normalizing, processing and aggregating IoCs from OSINT feeds; and (ii) a Context Aware Intelligence Sharing Module, able to correlate, assess and share static and real time information with data obtained from multiple OSINT sources.

The ETIP platform computes a Threat Score (TS) associated to each IoC before sharing it with both internal monitoring systems and tools (e.g., SIEMs) and trusted external parties. Enriched IoCs will contain a threat score that will enable SOC analysts to prioritize the analysis of incidents. The Threat Score evaluates heuristics with two types of weights: (i) individual weights assigned to every attribute (e.g., relevance, accuracy, variety, etc.); and (ii) global weight (i.e., completeness criterion) assigned to the heuristic. The higher the TS value, the more reliable the IoC. Thus, as the TS value approaches to zero, the IoC can be considered as poor, incomplete and/or not reliable with a very low priority level.

Future work will concentrate in developing new attributes to enrich the threat score analysis, improving the quality of the refined threat intelligence to be shared, providing not only the final threat score, but also detailed information about each single criterion used in the evaluation of the score itself, which in turn helps to improve threat detection and incident response.

ACKNOWLEDGMENT

The research in this paper has received funding from the EC through funding of DiSIEM project, ref. project H2020-700692, NeCS project, ref. project H2020-675320 and LASIGE Research Unit, ref. UID/CEC/00408/2019.

REFERENCES

Accenture (2017). Cost of cyber crime study. Online.

CEA (2018). The cost of malicious cyber activity to the u.s. economy. Online Report.

Dalziel, H. (2014). How to define and build an effective cyber threat intelligence capability. In *Syngress, eBook*.

ENISA (2015). Actionable information for security incident response. In *Online Technical Paper*.

ENISA (2017). Exploring the opportunities and limitations of current threat intelligence platforms.

Liao, X., Yuan, K., Wang, X., Li, Z., Xing, L., and Beyah., R. (2016). Acing the ioc game: Toward automatic discovery and analysis of open-source cyber threat intelligence. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 755–766.

Mavroeidis, V. and Bromander, S. (2017). Cyber threat intelligence model: An evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence. In *Intelligence and Security Informatics Conference*, pages 91–98. IEEE.

Owen, T. (2015). Threat intelligence & siem. In *Masters Research Project, Lewis University*.

Ravana, S. D. and Moffat, A. (2009). Score aggregation techniques in retrieval experimentation. In *Twentieth Australasian Database Conference*.

Sabottke, C., Suciu, O., and Dumitras, T. (2015). Vulnerability disclosure in the age of social media: exploiting twitter for predicting real-world exploits. In *In 24th USENIX Security Symposium*, pages 1041–1056.

Sauerwein, C., Sillaber, C., Mussmann, A., and Breu, R. (2017). Threat intelligence sharing platforms: An exploratory study of software vendors and research perspectives. In *Conference on Wirtschaftsinformatik*.

Sillaber, C., Sauerwein, C., Mussmann, A., and Breu, R. (2016). Data quality challenges and future research directions in threat intelligence sharing practice. In *ACM on Workshop on Information Sharing and Collaborative Security*, pages 65–70. ACM.

Skopik, F., Settanni, G., and Fiedler, R. (2016). A problem shared is a problem halved: A survey on the dimensions of collective cyber defense through security information sharing. *computers & security*, 60:154–176.

ThreatConnect (Accessed February 2018). Threat intelligence platforms. everything you've ever wanted to know but didn't know to ask. In *Ebook*.

Torra, V. and Narukawa, Y. (2007). Modeling decisions: Information fusion and aggregation operators. In *Springer-Verlag Berlin Heidelberg*.

Tounsi, W. and Rais, H. (2018). A survey on technical threat intelligence in the age of sophisticated cyber attacks. In *Computers & Security*, volume 72, pages 212–233.

Ventures, C. (2017). 2017 cybercrime report. Online.