Automated Solution for Enrichment and Quality IoC Creation from OSINT

Rui Azevedo, Ibéria Medeiros, and Alysson Bessani

LASIGE, Faculdade de Ciências da Universidade de Lisboa, Portugal
razevedo@lasige.di.fc.ul.pt, imedeiros@di.fc.ul.pt, anbessani@fc.ul.pt

Abstract. Cyber-security has become a top priority for most organizations, as the impact costs of cyber-attacks has risen to the billions of dollars. Organizations, to protect themselves, are resorting to security information and event management (SIEM) systems to monitor their infrastructures while investing in cyber threat intelligence (CTI) to provide them forewarning about the risks they face, as well as to accelerate their response times in the detection of attacks. One path to obtain CTI is the collection of open source intelligence (OSINT) via threat intelligence platforms (TIP) and their representation as indicators of compromise (IoC). However, most of TIPs provide threat information with little to no processing. This situation increases the pressure on security analysts who, already faced with the arduous task of sorting the alerts originating from their networks, must also sort this additional flow of data to find relevant intelligence. This paper proposes an approach to generate *threat intelligence of quality* based on collected OSINT feeds that can later be used in defensive infrastructures, such as SIEMs. The approach, implemented in a platform and assessed with 34 OSINT feeds, was able to create *enriched IoCs* that allowed e identification of cyber-attacks previously not possible by analyzing the IoCs individually.

Keywords: cyber-security · open source intelligence (OSINT) · threat intelligence platforms (TIP) · indicators of compromise (IoC) · security

1 Introduction

The estimated cost of cyber-crime is expected to reach $2,1 trillion by 2019 [7] placing cyber-security at the utmost priority level for most organizations. A reason for this is the changing threat landscape that also impacted the defensive posture assumed by companies, who moved from a perimeter defense approach to an *in depth defense* architecture, to improve the chances of avoiding an attack [5]. Central to this approach are the Security Information and Event Management (SIEM) systems, whose ability to compile and correlate the vast amount of information produced by the different sensors (e.g., firewalls) that monitor the networks activity vastly increase the capacity to detect an attack. However, Advanced Persistent Threats (APT) [2] have emerged and propagated as being an adversary that aims to steal data instead of causing damage to the organization

network and maintain its access undetected for longer periods of time [4]. Faced to this new type of adversary, SIEMs have been found wanting, as they struggle to react to the specific features of these new and sophisticated attacks.

Cyber threat intelligence (CTI) has emerged as a key asset to break such an adversary, detecting its activities by improving the defensive posture via the anticipation of future attacks. One way to obtain CTI is by accessing open source intelligence (OSINT) feeds with information about cyberspace threat activities, under the form of security events.

These feeds are typically collected through *threat intelligence platforms* (TIP) under the form of *indicators of compromise* (IoC), an information artifact that aggregates data on malicious activity in a system or a within a network. Despite the increase of such technology, there are still limitations [6], namely the harnessing of the volume of threat information that is produced and shared in the diverse formats of IoCs. Moreover, most of TIPs are providing threat information with little to no processing which, associated to the crescent volume of data that they collect, makes finding relevant and quality intelligence from them a hard task for a security analyst. This fact has increased the pressure on analysts, who are already faced with the arduous task of sorting the multitude of SIEM alerts, by forcing them to also sort these additional data. Therefore, an increase in the quality of the intelligence that is offered is needed to make it more useful, by providing context and prioritization that currently are lacking [8].

This paper proposes an approach to generate *threat intelligence of quality* based on collected OSINT feeds which can be later used in defensive infrastructures, such as SIEMs. This improved intelligence translates in new *enriched IoCs* that are obtained by correlating and combining IoCs from different OSINT feeds, aggregating them into clusters, and then representing the most relevant threat information of clusters in a single, enriched IoC. The approach was implemented in a platform and assessed with 34 OSINT feeds. The platform was able to create enriched and quality IoCs that should allow the identification of cyber-attacks previously not possible by analyzing those IoCs individually.

The contributions of the paper are: (1) an approach using OSINT to improve cyber-security; (2) two methods to correlate and aggregate TI; (3) a platform to generate enriched and quality IoCs by using these methods; (4) an experimental evaluation that shows the ability of this platform to create IoCs.

2 Background and Related Work

This section provides some context and related work on CTI, focusing on the platforms used, their challenges and on how quality is defined in this context.

2.1 A New Threat Landscape

The increased dependency in information technology has changed the threat landscape faced by companies, since currently, their most valuable assets exist in the form of bits within their network. This reality leads to the development

of new terms and concepts to define the threats that are faced and how they operate. One of them is the Advanced Persistent Threats (APT) that characterizes a novel type of adversary that, under patronage of a state or a private entity, uses diverse and significant resources to attack a strategically valuable target. This new actor can be distinguished from other criminal enterprises by four key features: specific and clear objectives; highly organized and well resourced; longer duration of the attacks; high degree of stealth [4]. The attack vectors have also changed focusing on the use of polymorphic (software capable of hiding its signature) and composite threats, which exploit both syntactic and semantic attacks against, respectively, technical and social vulnerabilities [14], allowing the execution of new and sophisticated attacks.

2.2 Threat Intelligence

To face this new landscape, companies are moving toward a proactive approach where threat intelligence (TI) plays a central role to inform their decisions.

Gartner defines TI as "evidence-based knowledge, including context, indicators, (...) about hazard to assets that can be used to inform decisions regarding the subjects response to that hazard" [15]. TI is the product of the threat intelligence life cycle that allows the transformation of raw data from the operational environment into intelligence that can be acted upon [10]. The main goal of this life cycle is to produce good quality intelligence. However, it presents a challenge as there is no definitive way to measure this value and the distinction between good and bad TI depends on the perception of the analyst that receives it.

TI quality is based on four characteristics [13]: *timeliness*, which indicates the time between creation and usage; *relevance*, which indicates if the intelligence is related to actual risks to the consumer; *accuracy*, which measures if the intelligence improves the response to incidents; *completeness*, which measures if the intelligence allows by itself the identification of an incident. Timeliness can be easily quantifiable, while relevance is determinable but it is closely tied to the recipient, so even within the same organization, different analysts may place different value to the same TI object. Accuracy and completeness can only be quantified after the fact and will always be estimated.

2.3 Threat Intelligence Platforms

As stated before, the objective of TI is the creation of a product that can be acted upon, which implies that it must reach the target population who require it. The increased need to share this type of product and the volume of information shared led to the traditional models of sharing information (e.g., mails) to become obsolete, as they were both inconsistent and non-scalable [1]. The recent solution for these issues came in the form of Threat Intelligence Platforms (TIP), such as Threatconnect or MISP, that promote the exchange of information, enable automation, and facilitate the generation, refinement and vetting of data [12].

There is no, to the best of our knowledge, architecture that illustrates the way TIPs work. Therefore, Fig.1 presents a proposal for such architecture based on

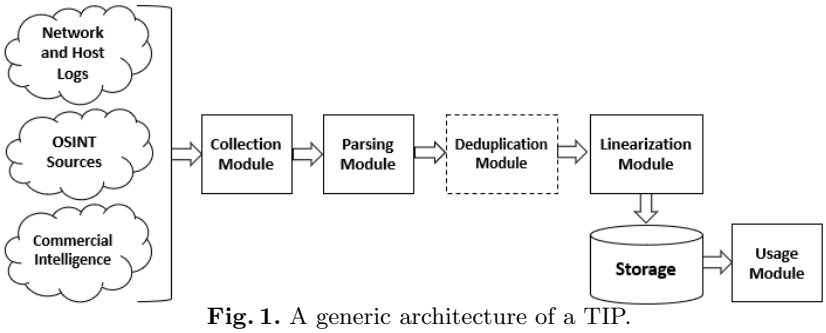

Fig. 1. A generic architecture of a TIP.

the documentation of platforms we have studied. The information, obtained from different sources, such as commercial, OSINT feeds, network logs, and forensic analysis, is collected for processing by a parser module. Some TIPs (such as IntelMQ) have the possibility of eliminating duplicated entries before linearizing them for long term storage. This stored information can then, depending on the platform, be used to produce rules to be deployed in security equipments (e.g., IDS) or undergo further analysis.

Alongside the evolution of TIPs, shared and standard formats have also evolved to represent TI processed by TIPs, i.e., to represent the IoCs. For instance, OpenIOC and STIX/TAXII are two of these standard formats, which have the capability of representing the TI parsed and normalized, respectively, on JSON or XML, identifying clearly each piece of information. This standardization facilitates the sharing of information among diverse entities, as well as its conversion into rules that may be applied by network ,defense mechanisms.

Despite the efforts put on the improvement of TIPs, a study by ENISA [6] found that these technologies still present many challenges that must be overcome, such as the definition of common standards for intelligence sharing, moving from data collection to analysis, reducing the need for human participation in the TI life cycle, and improving the quality of the TI that is provided. A solution to part of these issues lies in augmenting the automation capacities of TIPs to enable the processing of TI to increase its quality [8].

3 IoC Enrichment and Quality Platform

The proposed platform collects, aggregates, and correlates IoCs from OSINT sources, which is tightly related to *cyber-security responses*: get relevant and quality security TI insights to be used by organizations to improve their response times and to try to anticipate the attacker's actions. The improved intelligence translates in new enriched IoCs that can be used by defense mechanisms belonging to an organization's network infrastructure.

Given the complexity of TIPs, the proposed approach addresses three challenges (presented next), identified as TIPs limitations in [6], with the aim of obtaining quality TI in an automated and efficient way. Therefore, in a first

instance this section presents which are these challenges and how they are addressed by the proposed approach, justifying thus the options taken on it, and then it presents the approach itself and the platform architecture.

3.1 Threat Intelligence Quantity and Quality

Reducing the quantity of information that reaches a security analyst. TI is usually dispersed through multiple sources [9], which implies that to maximize the knowledge on a specific issue, an analyst must identify and access these, making his task harder and creating the risk of useful information being overlooked.

Our approach addresses this issue by collecting different sources dispersed through the Internet, aggregating them in clusters, and then representing each cluster by a single IoC. In doing so, we also reduce the volume of information stored by TIPs.

Increasing the quality of intelligence that arrives to an analyst. As stated previously, the quality of TI derives from four factors (see Section 2.2). To guarantee an increase in the quality, these factors must be addressed by working on both (1) the TIP configuration and (2) on its internal processing capabilities. (1) is dependent of the clear definition of the goal of the platform, meaning the purpose that the intelligence obtained will be applied on must be well-defined and a specific scope must be established. This will allow the selection of sources that focus on those targets, reducing the collection of irrelevant information. The factors of accuracy and timeliness are also dependent on sources selection, which have to be vetted on the quality of TI they provide (e.g., how recent are the indicators they share?). Regarding (2), a platform must guarantee that it will be able to process the information it receives to extract the most relevant details for the analysts.

The approach aims to address this, first by using sources containing high level detail TI, i.e., different from blacklisted IPs, secondly by employing correlation techniques capable of interconnecting different security events related to the same threat, and then generate new data resulting from this processing.

Facilitate the automation of the generation of improved intelligence. Any proposed platform must be designed having in mind that it should be able to work with minimal human intervention, and that it should be possible to integrate its products into other technological platforms.

The proposed approach comprises a set of modules interacting between them autonomously and automatically to generate new enriched TI that can be integrated in defense equipment and mechanisms.

3.2 Platform Overview

To achieve the challenges just described, we propose a platform capable of generating enriched IoCs, i.e., that collects different IoCs obtained from OSINT

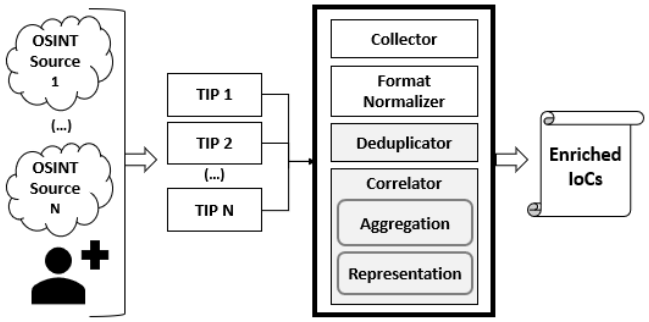

Fig. 2. The enriched and quality IoC platform architecture.

sources, aggregates them in clusters, correlates IoCs within clusters, and generates new IoCs that represent the most relevant threat information of the clusters in a concise form, offering a better intelligence product than the individual IoCs. Fig. 2 shows the architecture of the platform with the following main components:

- *TIP Sources.* Different TIPs are placed in parallel collecting feeds from different OSINT sources and inputs from different communities. This allows the system to take advantage of the different enrichment capabilities these platforms offer as well as from the different communities that use them.
- *Collector.* The output of the different TIPs is channeled to the collector, which stores IoCs temporarily until they are processed.
- *Format Normalizer.* Converts the TIPs IoCs into a single format to allow them to be processed by the platform.
- *Deduplicator.* The deduplicator analyzes received IoCs to remove duplicates before sending them to the correlator module.
- *Correlator.* Correlates different IoCs and generates new ones. Processing consists of querying the platform database to identify IoCs that contain relevant related information fields, and then merging that information into a single IoC, creating a new rich IoC that is stored for later can be used in a variety of purposes.

Given the relevance of the *deduplicator* and the *correlator* components, the next sections provide further detail on them.

3.3 Deduplicator

As received IoCs may already be present in the platform, the deduplicator module aims to eliminate duplicates and reduce the necessary storage space. This module analyses incoming IoCs, comparing the elements that compose them with those present in other IoCs already in the platform. If the module detects that all the information presented in the incoming IoC is already present in one of the existing IoCs, or that one of the existing IoCs is fully represented in the incoming one, then it will store the most complete, or all things being equal, the oldest IoC, discarding the superfluous one.

3.4 IoC Correlation

The correlator is the core of the platform, performing the aggregation, correlation and representation tasks. The process starts by searching for correlations between the different IoCs and, once correlations have been identified, generating new ones. It performs its function in two steps: (1) *aggregation*, it queries the database to identify IoCs that contain relevant related information fields determining clusters of related IoCs; (2) *representation*, for each resulting cluster, it merges the information contained in different IoCs into a single one, eliminating duplicated attributes, and stores the new enriched IoC for later use.

To establish the correlations the platform can use one of two correlator methods we defined, named *naive* and *cluster*, that would allow the identification of groups of correlated IoCs, generating the clusters of IoCs. However, before performing any method, an initial filtering stage is applied to identify only IoCs that respect specific rules, i.e., by eliminating events that will bring no added value (such as blacked IP lists), this allows to create a *subset of IoCs of interest*. In the naive method only direct correlations are identified, in the sense that the enriched IoC is built from a central IoC and all those IoCs that share one or more attributes with it. On the other hand, the cluster method creates a graph with all events in the IoCs set, where each IoC is a node and the edges represent shared attributes between IoCs, and clusters interconnecting IoCs are identified as a source for a new enriched IoC. Algorithm 1 presents the algorithm of both methods, and Fig. 3 shows the resulting enriched IoCs for the cluster method. In part A of the figure is shown the formed clusters (the black points), and in part B an example of a enriched IoC is illustrated, where the nodes represent the different sources and the edges their interconnections.

Algorithm 1: Correlator Algorithm

```
Input: The subset of IoCs of interest
Output: A list of enriched IoCs
ListsOfEventsToMerge = [ ];
if Naive approach used then
    foreach event of subset do
        if event share attributes with related_event of subset then
            Create list L with event and all related_events;
            Add L to ListsOfEventsToMerge;
if Cluster approach used then
    Create graph G;
    foreach event of subset do
        Add event to G as a node;
        if event share attributes with related_event of subset then
            Add edge to G connecting event to related_event;
    Identify all clusters in G;
    Add all nodes in a cluster as a list to ListsOfEventsToMerge;
foreach list in ListsOfEventsToMerge do
    Create AttributesToMerge=[ ];
    foreach event in list do
        Collect all attributes in event;
        if attribute not in AttributesToMerge then
            Add attribute to AttributesToMerge;
    Create new event containing attributes in AttributesToMerge;
```

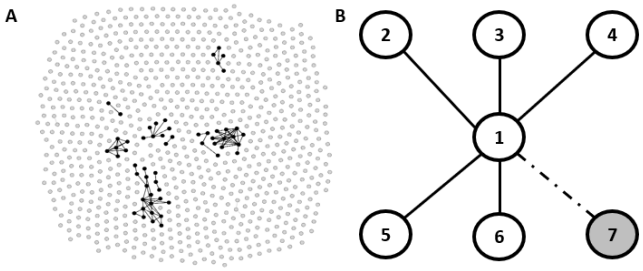

Fig. 3. A - Graphical representation of the nodes in clusters outlined. B - Schematic representation of a enriched IoC.

4 Implementation

To implement the proposed platform, we developed the deduplicator and correlator modules in Python 3, and we used the MISP [11] platform to receive and normalize the data coming from TIPs, and then to process that data by using our modules. MISP was installed in a Docker Container.

For the development of deduplicator and correlator modules, we made use of third party libraries such as pymysql, pymisp and networkx to provide basic functions to query the database, create graphs and identify their subgraphs and, to create the enriched IoCs.

The platform can be configured following a selection criteria that gives the opportunity to the user to choose the trust level and similarity to generate the enriched IoCs, namely: (1) *Node selection:* option to select the level of trust for the event to be included and to indicate if events tagged as black lists should be excluded; (2) *Edge selection:* option to select different rules in the definition of which similarities should be taken into consideration, such as accept all relationships, exclude all relationships based on attributes belonging to the MISP 'Network' category or only include relationships based on attributes belonging to the MISP categories 'Attribution', 'Targeting data' or of the type 'vulnerability'.

5 Evaluation

This section presents the evaluation of our platform on generating enriched IoCs. Section 5.1 describes how we set up and configured the platform and Section 5.2 presents and discusses the obtained results.

5.1 Setup

The platform was built by mounting a MISP Docker Container on a Dell PowerEdge R420 server with a Intel Xeon E5-2407 CPU 2.2GHz/10 MB cache processor, 32 GB RAM and a 300 GB hard disk. It was configured to collect security events from 34 OSINT feeds provided by 12 distinct organizations, which allowed the collection of 1174 events during 35 days. Table 1 presents the distribution of

Table 1. Distribution of OSINT events by Organization and Trust Level.

Organization Name	Trust Level			Events	Avg. Attr. by Event
	0	1	2		
CERT-RLP	0	0	1	1	62
CIRCL	182	138	484	804	126
CUDESO	0	1	121	122	40
CiviCERT	0	0	1	1	79
Crimeware	0	0	1	1	26
CthulhuSPRL.be	7	1	209	217	422
ESET	0	0	1	1	263
FOXIT-CERT	0	0	1	1	180
INCIBE	0	1	0	1	64
NCSC-NL	0	1	1	2	125
clearskysec.com	1	0	0	1	110
inThreat	0	20	2	22	6232
TOTAL	190	162	822	1174	286

events by the organizations (column 'Events') and the event trust level based on the analysis indicator present in the event, where level 2 is the maximum trust. The last column indicates the average number of attributes of the events.

For the assessment we made two assumptions: (A1) the level of analysis associated to an event directly correlates to the level of trust that can be placed on the information contained; (A2) all black list events are correctly tagged.

We focused on the results obtained when evaluating only fully analyzed events that did not contain blacklisted IPs (representing a subset of 820 IoCs from our collected events), which we considered to be the minimum requirements for the selection process in the context of the creation of improved IoCs. The rationale behind the first criteria is, as per our assumption A1, that events marked as fully analyzed will only contain information that was confirmed by a human actor thus providing some guarantee of their quality. As to the second criteria, it was deemed that events containing black lists will reduce the quality of the end product as their size and indiscriminate nature will dilute the relevance of the information in the enriched IoC, as well as increasing the chance of unrelated incidents being added during the aggregation stage of the processing (while we will not present these results here, this was experimentally confirmed).

5.2 IoCs Generation

To test the platform, we ran multiple simulations (both with and without the creation of the resulting enriched IoCs), applying the different filters and correlate methods to the creation of the enriched IoCs, that allowed us to observe the platform's performance under these different constrains.

Table 2. Number of potential enriched IoCs obtained when applying different filters.

Filter	Naive Approach	Cluster Approach
All connections	423 [2:44]	69 [2:276]
No network connections	351 [2:43]	65 [2:205]
Only attacker, target or vulnerability	54 [2:7]	11 [2:17]

In Table 2, we present the results, in terms of potential new enriched IoCs, of using the naive and the cluster methods on our collected events with the application of the different filters to the connections, and with the indication of the minimum and maximum number of contributing IoCs (inside brackets). The results obtained from these experiments show that the naive method produces a higher number of possible enriched IoCs in all implementations, however, the maximum number of IoCs proposed for the inclusion in an enriched IoC is lower than when the cluster method is used. This result was to be expected as the naive method will present all subsets of the different proposals identified via our cluster method. From these results, it is also possible to confirm that a more targeted approach to filter the correlations allows the identification of fewer and smaller clusters of correlated IoCs.

An analysis of a sample of the produced IoCs showed that the best intelligence products were obtained from the application of the cluster approach with a targeted filter using only connections from attacks that explored the same vulnerability, were performed by the same actor or had the same target.

Table 3. Resulting enriched IoCs by using cluster approach.

Enric IoC	Characteristics
1	Composed of two distinct IoCs from the same organization connected by a WhoIs registrant email. It presents the connection between a specific malware campaign and an email used in the registration of multiple domains involved in attacks.
2	Composed of two distinct IoCs from two distinct organizations that share 74 attributes and that both focus on reviews of information collected on an attacking group called 'Packrat'.
3	Composed of two distinct IoCs from the same organization that are related by a WhoIs registrant email and a domain belonging to that actor, connecting an event with domains used by that actor with the 'Multichair' campaign.
4	Composed of three distinct IoCs from two organizations that present different attacks where the same vulnerability was used CVE-2017-11882.
5	Composed of three distinct IoCs from two organizations connected by the use of the vulnerability CVE-2017-0262. This new IoC connects the activities of the APT 'Sofacy' with a campaign against financial institutions in Ukraine and an analysis of EPS Processing zero-days attacks.
6	Composed of four distinct IoCs from two organizations it provides a connection between two previously unconnected APTs BlackVine and Cyber Kraken (aka Threat Group 3390/Emissary Panda), both active in 2015, who both used the ScanBox framework.
7	Composed of five distinct IoCs from one organization, it connects four different studies (of which only two are also connected when applying our filter) on the APT Sofacy (aka APT 28/Fancy Bear) via one central IoC dedicated to the use of a specific provider.
8	Composed of six distinct IoCs from two organizations, it is composed of a central cluster of three interconnected IoCs centered around the Turla (aka Snake/Uroburos) attacks and the APT Sofacy, to which the three other events are connected independently, bringing further information on the Turla attack, introducing an expansion of the Snake attack and, also, a framework called Cobra that was used by attackers also using Uroburos.
9	Composed of seven distinct IoCs from two organizations, it is created around an IoC on the Neutrino Exploit Kit connecting via one of two vulnerabilities, CVE-2014-6332 and CVE-2013-2551, to multiple unconnected attacks and to another exploit kit.
10	Composed of eleven distinct IoCs from two organizations, it contains a group of ten interconnected events that focus on the activities of the APT Sofacy and allows the connection of an IoC on Operation Pawn Storm to them.
11	Composed of seventeen distinct IoCs from three organizations, it contains a central core, composed of nine IoCs that are connected by the CVE-2012-0158, that branches out to three other groups of IoCs on campaigns and tools.

Table 4. Details of the components of the *enriched IoC 9*.

ID	IoC Name	# Att.	Creation Date
1	OSINT Neutrino Exploit Kit - One Flash File to Rule Them All by SpiderLabs	38	28-12-2015
2	OSINT DTL-12012015-01: Hong Kong SWC attack from Dragon Threat Labs	40	11-01-2015
3	OSINT Operation Double Tap from FireEye	30	21-11-2014
4	OSINT Infected Korean Website Installs Banking Malware by Cyphort	30	28-09-2015
5	OSINT Evilgrab Delivered by Watering Hole Attack on President of Myanmars Website by Palo Alto Unit 42	30	11-06-2015
6	OSINT Evolution of the Nuclear Exploit Kit by Cisco Talos group	225	09-10-2014
7	OSINT - powershell used for spreading trojan.laziok through google docs	6	22-04-2016

The application of the cluster method with the targeted filter allowed the identification of 11 possible enriched IoCs. Table 3 presents these IoCs and their characteristics. These 11 enriched IoCs show indications of containing higher quality information than the individual original IoCs, either by bringing together further technical information or by allowing to look at the patterns of attacks and tools that emerge around a specific vulnerability.

The *enriched IoC 9* (presented in Table 4) offers a good example of the potential offered by enriched IoCs. The cluster that originated it is composed of IoCs only correlated by a connection to a vulnerability, either CVE-2014-6332 or CVE-2013-2551, with the other technical information they contain reinforcing the information existing in the other components. This is particularly relevant in the case of IoC 7, since it only contains 6 attributes of which one is the shared one and two others that refer to external analysis, it is to be expected that by itself it would not allow the identification of any threat. Furthermore, this enriched IoC also shows the evolution of this vulnerability and cases where it was used in attacks. Fig. 3–B is a schematic representation of the enriched IoC 9 with the color of the nodes representing their source organization [white: CthulhuSPRL.be, grey: CIRCL] and the type of connector their shared attribute [full: CVE-2014-6332, dotted: CVE-2014-2551]

The results show that the platform may serve a purpose at both technical and strategic levels of intelligence, allowing the production of intelligence that may cover all levels of the pyramid of pain [3], which defines the sophistication levels to deploy attacks (and to detect them).

6 Conclusion

This paper presented a solution for the issue with the overload of information that reaches organizations and their defensive structure, aiming to improve their protections by using threat intelligence (TI) of quality. Focusing on collected OSINT feeds, the solution correlates events, and applies a clustering technique to identify groups of indicators of compromise (IoCs) to create enriched TI. These Enriched IoCs offer added value by grouping information that would otherwise be dispersed in multiple different sources. The solution was implemented in a platform and validated with 34 OSINT feeds. The resultant Enriched IoCs may allow the characterization of cyber-attacks which were not possible by analyzing the original IoCs individually.

Acknowledgements This work is supported by the European Commission through the H2020 programme under grant agreement 700692 (DiSIEM), and LASIGE Research Unit, ref. UID/CEC/00408/2013.

References

1. Bank of England: CBEST intelligence-led testing: Understanding cyber threat intelligence operations (2016), https://www.bankofengland.co.uk/-/media/boe/files/financial-stabilit/financial-sector-continuity/understanding-cyber-threat-intelligence-operations.pdf
2. Bejtlich, R.: What apt is (and what it isn't). Information Security pp. 20–24 (July/August 2010)
3. Bianco D.: The pyramid of pain (Mar 2013), http://detect-respond.blogspot.pt/2013/03/the-pyramid-of-pain.html
4. Chen, P., Desmet, L., Huygens, C.: A study on advanced persistent threats. In: Proceedings of 15th IFIP TC 6/TC 11 International Conference. pp. 63–72 (Oct 2014)
5. Cole, E.: Advanced persistent threat - understanding the danger and how to protect your organization (2013)
6. ENISA: Exploring the opportunities and limitations of current threat intelligence platforms (version 1.0) (Mar 2018), https://www.enisa.europa.eu/publications/exploring-the-opportunities-and-limitations-of-current-threat-intelligence-platforms
7. Internet Society: Global internet report 2016 (Oct 2016), www.internetsociety.org/globalinternetreport/2016/
8. Itay Kozuch: Cyber threat intelligence:how to turn quantity into quality (Apr 2018), https://www.peerlyst.com/posts/cyber-threat-intelligence-how-to-turn-quantity-into-quality-itay-kozuch
9. Liao, X., Yuan, K., Wang, X., Li, Z., Xing, L., Beyah, R.: Acing the ioc game: Toward automatic discovery and analysis of open-source cyber threat intelligence. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 755–766 (Oct 2016)
10. Martin E. Dempsey: Jp2-0, joint intelligence (2013), http://www.dtic.mil/doctrine/new_pubs/jp2_0.pdf
11. MISP Project: Tmisp - open source threat intelligence platform & open standards for threat information sharing, http://www.misp-project.org
12. Sauerwein, C., Sillaber, C., Mussmann, A., Breu, R.: Threat intelligence sharing platforms: An exploratory study of software vendors and research perspectives. In: Towards Thought Leadership in Digital Transformation: 13. Internationale Tagung Wirtschaftsinformatik. pp. 12–15 (Feb 2017)
13. Sergio Caltagirone: The 4 qualities of good threat intelligence (Jul 2015), http://www.activeresponse.org/the-4-qualities-of-good-threat-intelligence/
14. Tounsi, W., Rais, H.: A survey on technical threat intelligence in the age of sophisticated cyber attacks. Computers & Security **72**(C), 212–233 (Jan 2018)
15. Webroot: Threat intelligence: What is it, and can it protect you from today advanced cyber-attacks? (2014), https://www.gartner.com/imagesrv/media-products/pdf/webroot/issue1_webroot.pdf