An Empirical Evaluation of Cyber Threat Intelligence Sharing in the ECHO Early Warning System

Ioannis Chalkias[1], Cagatay Yucel[1], Dimitrios Mallis[1], Jyri Rajamaki[2(✉)], Fabrizio De Vecchis[3], Peter Hagstrom[3], and Vasilis Katos[1]

[1] Bournemouth University, Fern Barrow, Poole, Dorset BH12 5BB, UK
{ichalkias,cyucel,dmallis,vkatos}@bournemouth.ac.uk
[2] Laurea University of Applied Sciences, Espoo, Finland
jyri.rajamaki@laurea.fi
[3] RHEA Group, Avenue Einstein 8, 1300 Wavre, Belgium
{f.devecchis,p.hagstrom}@rheagroup.com

Abstract. This paper reports on the information sharing practices of cyber competency centres representing different sectors and constituencies. The cyber competency centres participated in the form of CSIRTs employed the ECHO Early Warning System. Through a structured tabletop exercise, over 10 CSIRTS were engaged and a number of features were captured and monitored. A key research question was to determine the factors that can potentially hinder or amplify Cyber Threat Intelligence information sharing. The exercise imitated real attack scenarios using state-of-the-art tactics techniques and procedures as observed by real-world APT groups and daily incidents. The findings revealed differences in terms of timeliness, response time and handling tickets with different Traffic Light Protocol classifications, duration of handling a ticket and intention to disclose.

Keywords: cybersecurity tabletop exercise · extra-constituency information sharing

1 Introduction

As attacks become frequent and with high impact, an Early Warning System (EWS) for Cyber Threat Intelligence (CTI) aims at serving as a security operations support tool enabling the members of the network to coordinate and share information in near real-time in order to develop and maintain efficient incident handling capabilities. At the same time EWS, stakeholders must retain their fully independent management of cyber-sensitive intelligence and related data management.

The design and development of the proposed EWS followed the main concepts and practices of information sharing and trust models from the cyber domain. The requirements elicitation and analysis process included the requirements for information sharing

A preliminary version of this paper appeared in DIGILIENCE 2020 [17].

T. Tagarev and N. Stoianov (Eds.): DIGILIENCE 2020, CCIS 1790, pp. 23–40, 2024.
https://doi.org/10.1007/978-3-031-44440-1_3

within and between partners across organisational boundaries as derived from a multi-sector analysis. This paper reports on the empirical evaluation of the proposed EWS which was conducted through a table-top exercise (TTX). To the best of our knowledge, the literature currently lacks publicly available datasets capturing the communication dynamics and exchange between different CSIRTS, across different sectors. In this work we attempt to create such a dataset and use it in order to evaluate both an information sharing system (namely the Early Warning System of the ECHO pilot project) as well as the team and communication dynamics within and among the participating CSIRTS.

The rest of the paper is structured as follows. In Sect. 2 we outline how tabletop exercises provide the means to evaluate CTI sharing approaches. Section 3 outlines the particular features of ECHO's Early Warning System (EWS). In Sects. 4 and 5 the empirical analysis and evaluation through TTX exercise are presented. Section 6 summarises the findings and presents areas for future research.

2 CTI Sharing and Tabletop Exercises

The literature review indicates that there is no pervasive or widely agreed upon definition of "cyber security information sharing". As such, the structures of the information sharing models can be sector-specific and created in different environments. There is a need for a common early warning solution. The fight against hybrid threats means not only preventing cyber attacks, but also identifying, tracing and prosecuting a criminal/criminal group [19]. This means an even deeper integration of government systems in the future as the term "warning" includes also preventive functions.

Relevant information from the site of a major hybrid incident should be directly shared with the national CERTS and follow a coordinated response. Combining pieces of information to ensure the correct and reliable information to be shared is deemed to be significant in establishing cyber capacity. The shared information should be in a form that is unambiguous and accessible to the involved parties. In the future cyber defence operations are expected to be more integrated and automated according to local capabilities, authorities and mission needs [15]. A shared common operational picture means that real-time communication link from the local level to nation and EU level exist. A common cyber situational awareness is needed for both operating Cyber Physical Systems (CPS), and for emergency and crisis management [4]. There should be the connection between cyber situational awareness and emergency management [13].

Moreover, it is important to take into account how national Cyber Security Centres cooperate with other organisations within critical infrastructure on a national level. The states departments of the United States work closely together in the fight against cyber security threats. The organisations of public administration of the Member States in European Union cooperate on a formal basis as set out in the NIS directive [1] and the Cyber Security Act [2].

It could be argued that cooperation outside the EU borders can be hindered by the fundamental differences of administrative functions between European Union and the other country/coalition. However, as Ilves et al. [12] mentioned, there are no crucial barriers to increase collaboration concerning early warning solutions between US, NATO and EU. US's Cyber security sharing act and Europe's directive on Network and Information Security (NIS) have similar goals. In addition to this, EU and NATO signed

a technical arrangement in 2016 to increase information sharing between the NATO Computer Incident Response Capability and EU Computer Emergency Response Team [10]. Public safety actors like European law enforcement agencies need common shared situational picture for the cross-boarding tasks in a way that operational cooperation will be based on reliable platform [7].

As part of cyber crisis management, Tabletop Exercises (TTXs) for cyber security were established around the early 2000s to provide the response-plan developers with insights about the efficiency and the effectiveness of the proposed action plan in case of a malicious cyber incident. In [20], the authors discuss in detail the very first government/multi-sector joint event known as the "Dark Screen" that took place in San Antonio, Texas. The aim was to assess the preparedness and the abilities of the districts of San Antonio and Bexar against cyber terrorist threats. Dark Screen was a TTX separated into 3 stages, each of them with different objectives; these included the successful interconnection among the participating organisations, penetration and incident response testing, and information sharing observation. Moreover, examples of sector based tabletop exercises are presented in the same document, where the respective sectors make up national critical infrastructures. In all cases, information sharing among the various sectors and the related organisations has been repeatedly identified as a critical and necessary measure towards cyber resilience enhancement.

A TTX must be designed by experienced facilitators and includes one or more realistic, but fictional, scenarios tailored to the participants of the exercise; these scenarios include cyber attack simulations with well-defined objectives that aim at exploiting loopholes in systems, response plans, behaviours, but also highlight the importance of individual roles and responsibilities, and timely decision-making activities, such as attack mitigation and information sharing, against potential cyber threats [14]. Upon the completion of each TTX an "after-action" report is deemed necessary in order to describe if the original goals were met, the overall experience of the participants, the lessons learned and the associated proposed solutions. Such an example is given by the EU Agency for cyber security (ENISA) [5] who recently conducted a TTX in an attempt to assess and evaluate the EU's crisis plans and mitigation mechanisms in case of malevolent cyber incidents related to the EU elections. Similarly, the U.S. Department of Homeland Security (DHS) Cybersecurity and Infrastructure Security Agency (CISA) designed the Elections Cyber Tabletop Exercise Package (ECTEP) Situation Manual (SitMan), based on the TTX format, as part of a strategic initiative to strengthen the cyber training capability of stakeholders [6].

As EWS and CTI play a supportive role on security operations [16], the approach of conducting tabletop exercises has been chosen for the evaluations of EWS, throughout its development.

3 Early Warning System Features and Requirements

At the kernel of information sharing lies the intelligence data item (IDI). In the context of ECHO, an intelligence data item is defined as any piece of data that potentially contains actionable information relating to cyber security [8]. Appreciating the enormous value of information and its potential, an information sharing framework is required in

order to appropriately manage the life-cycle of the corresponding data items, from their generation, processing, dissemination all the way to their destruction. This approach is expected to facilitate the creation of a community of a large pool of stakeholders who will engage in joint intelligence activities and reliably share information and collaborate in handling security incidents in an effective and timely manner. As such, establishing and ensuring trust is a key factor for the successful adoption of the EWS [18].

3.1 Characteristics of Intelligence Data Items

At a first level of discrimination, IDIs can be structured, semi-structured or unstructured. Typically unstructured data refer to primary sources of information that are normally processed by automated or human means for extracting the necessary information. This process would generate structured IDIs that would allow automated processing. It should be noted though that there can be primary sources ingested into the EWS that are structured (e.g. log files).

IDIs can also be distinguished as reference information or operational information. Reference information refers to the IDIs that contribute in achieving situational awareness, allowing the beneficiary to make informed judgements on the cyber risks of the organisation. Operational information relates to those IDIs that support the actual decision making, handling incidents and so forth.

The IDIs should be accompanied with metadata that will contextualise the contained information but also enable the EWS to implement and enforce authorisation and access control mechanisms. Common identifiers and enumerations should be used whenever possible.

Table 1 presents an initial list of the categories of information and their expressions as IDIs. IDIs that potentially contain Personal Information will need to also meet the privacy requirements. These categories were further expanded and refined following the requirements elicitation and specification.

Information Sharing Model Assumptions

Against all the above, the proposed ECHO information sharing model is based on the following assumptions or premises:

- There will be a clear and concise governance model for the intelligence data items, where each item will be described by a comprehensive list of contextual information (metadata) to allow fine-grained decision making on the management and handling of the data.
- There will be a clear process for on-boarding and off-boarding of participating organisations.
- It is expected that it would be easier for organisations being in the same sector or having similar goals and purpose to form easier clusters for sharing threat intelligence information, as they are more likely to have established and mature exchange arrangements; therefore they are more likely to reach consensus. On the contrary, organisations that operate in orthogonal industries (i.e. where their respective industries have virtually nothing in common) is expected that would be less forthcoming in sharing information.

Table 1. Intelligence data items

Information category	IDI	structured/ unstructured	reference/ operational	Personal Information
Technical threat indicator	IOC (email, IP address, file hash, mutex, domain)	S	R	
Intrusion attempt	Threat Actor	S	O	X
	IOC (atomic, composite, behavioural)	S	O	
Security alert	Ticket	Semi	O	
	Readiness level	S	R/O	
Vulnerability information	CVE	S	R/O	
	CVSS	S	R/O	
	Threat identification	Semi	O	
	Geopolitical	U	R	
	Exploitability	S	R	
Vulnerability report	Vulnerability scanning report	S	R	
Incident report	Report	U	O	?
TTP	ATT&CK	S	R/O	
	STIX object	S	R/O	
Remediation actions	Operating procedure	U	O	
	Playbook	U	O	
Asset	CPE to describe system platforms	S	R/O	
	CCE (common configuration enumeration)	S	R/O	
Discussion	Discussion item	U	R/O	?
Blog post	Reference	U	R/O	
Poll	Poll item	U	R/O	
Raw data	Log-file	S	O	
	Netflow	S	O	X
	Packet capture	Semi	O	X
	RAM image dump	Semi	O	X
	Malware sample	Semi	O	?
	VM Image	U	O	X
	File	U	R/O	X
	Email	U	R/O	

- Stakeholders and participants are expected to join predefined and ad-hoc groups.
- Trust will be delivered through technical, organisational and human means.
- Due to the nature and diversity of sectors, in order for information sharing to provide a meaningful and accurate services, the scope of the data items should be extended to encompass Cyber Physical Systems; indicatively, this can consider the practices found in the Maritime Sector where there is a clear distinction between cyber (e.g. IT networks) and Physical (e.g. Operational Technology networks) highlighting the existence and inter-dependencies between the physical and cyber plane.
- Translation and normalisation services will allow the standardisation of intelligence data items. The underlying taxonomies and schemas should cater for the verticals by including optional fields.
- Existing standards for information processing and sharing will be adopted wherever possible.

3.2 Information Sharing Architecture

Information sharing is highly dependent upon and influenced by the regulatory frameworks as well as the cultural norms both within a sector and the organization itself. In academia for example, barriers to sharing are expected to be lower than the other sectors, due to the culture of freedom of academic expression and an academic citizen mentality of peer review and dissemination of research output. On the other hand, in critical infrastructure type of sectors such as Energy, or in banking [3], information sharing is more intensely regulated, and this also is reflected in the respective organisational cultures. This creates a tessellation of regulatory frameworks and cultural antecedents on the following levels:

- Intra-Organisational, influenced by specific internal policies and procedures.
- Intra-Sector, imposed by the respective sector.
- National-governmental, governed by the respective strategic decisions on a national level.
- Transnational, through the international agreements, treaties and EU legislation and directives, in the case of the organisation operating within the EU. This may include frameworks for information sharing with Law Enforcement entities.

The above are also complemented by horizontal legislation such as the GDPR that cuts across all sectors. Provided that:

- The jurisdictional and operational environment of the proposed system relates to the EU initiative on establishing a network of competency centres, and
- the EWS is intended to support information sharing among and between a multitude of sectors such as Healthcare, Energy and Maritime,

A modified hybrid model architecture is recommended as this appears to best fit the requirements following the cross-case analysis. In essence, the hybrid approach will allow to maintain a basic form of hierarchy, and at the same time it will allow the connection of different hubs, forming a higher peer to peer level. This is also in accordance to how CERTs operate and share information, which is done on a peer to peer basis but also within their level of operation (e.g. national, organisational, etc.). Allowing some degree of centralisation will also enable centralised decision making and support the emergence of Coordination Centres. A hub could represent a variety of communities, such as a specific sector, an interest group or a national point. It is recommended that each hub will refer to organisations of common characteristics, goals or sector, simplifying its management, internal governance and deployment complexity.

This would be inline with the EWS architecture supporting tenants allowing also seamless integration through the sharing API capability that will connect EWS instances.

From a governance perspective, the immediate consequence of this would be to have trust realms, two tiers of cross organisational boundaries, as shown in Fig. 1.

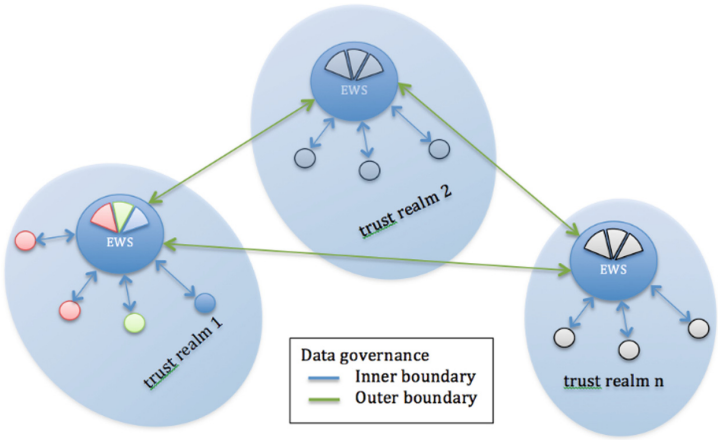

Fig. 1. Information sharing architecture

In this example, three trust realms are presented. Each realm can correspond to any type of organisational cluster, e.g. realm 1 could be academic CERTs, realm 2 national cyber security competency centres and realms 3 maritime sector. Every trust realm can have more than one EWS instances, for scalability and resilience purposes. The governance model could refer to policies and security certification requirements for deploying an EWS instance.

The first step for an organisation or individual joining the EWS ecosystem is to complete the on-boarding process. Upon successful application, the organisation is allocated a tenant slice. This will host all information provided by the participating organisation. Organisation boundaries can be crossed within a given trust realm and these are specified through the inner boundary data governance. It is expected that these will be the first to be formed, upon the emergence of the EWS.

Inter-realm information sharing is controlled by the outer boundary data governance models. These are expected to be more complex and diverse and will require a longer maturity period. It should be noted that not all trust realms will necessarily connect to each other; such configurations imply that some realms will emerge to be more authoritative and trustworthy than others, but should also indicate that transitive trust should not be guaranteed or offered.

IDIs containing personal information will be go through anonymization and redaction layers prior to leaving a tenant's area. For structured IDIs, automated processes would seamlessly and efficiently implement the underlying privacy policy. Information classification schemes will be enforced at the organisational boundaries (coarse grained access control) as well as internally (fine grained).

As the organisation participation and connectivity between the hubs increase, the value of the network is expected also to increase, in accordance to Metcalfe's Law. However, as this increase is very likely to result to generation of large volumes of data, the perceived usefulness is expected to decrease. In order to compensate for this, information sharing should not only be limited by access control criteria, but additional contextual features to enable effective filtering of non-relevant information (noise). A representative feature for this task is asset information. For example, by using the Common Platform Enumeration (CPE) convention, an organisation can describe their assets in a standardised way. By doing this it would be possible to quickly filter out attacks and vulnerabilities that are not applicable to a particular organisation's attack surface.

3.3 Features of the Information Sharing System

A modular approach for the EWS is considered. The core EWS should be comprised of a ticketing system supporting distributed workflow among a number of different partners and organisations. The EWS should allow the enrichment and contextualisation of the introduced and ingress information [11]. As such, a standard description and an expandable information taxonomy should be considered.

An initial list of features of the perspective EWS is presented below:

– **A suitable confidentiality model, such as the traffic light protocol.** All intelligence items will need to be assigned with a designation to ensure that the sensitive information is shared with the appropriate audience. The Traffic Light Protocol (TLP) is

recommended because it is less formal, does not really require NDAs, etc., it is more of a "gentlemen's agreement" and allows a faster communication of incident data. TLP will of course run in conjunction with the standard system's access control mechanisms, such as RBAC. For the EWS system in particular and upon a joint decision, FIRST's TLP definition is adopted to support future interoperability and standardisation with all pilots. Moreover, the confidentiality model – due to the nature of the EWS – should include introduction of information by protecting source attribution (Chatham House rule), in order to facilitate the submission of any information that can be vital when handling security incidents. A direct consequence of this is the consideration of the reliability of the data, defined further below.

- **An access control scheme, capable of making fine-grained access control decisions.** The audience accessing intelligence items shall be controlled through access classifiers such as organisations, groups, and roles.
- **Support of multiple taxonomies and standards for intelligence sharing.** This will allow the hosting of organisations belonging in different sectors.
- **Capabilities for a structured sharing of intelligence data e.g. use of Structured Threat Information eXpression (STIX).**
- **The system should facilitate the exchange of intelligence between CERTS/CSIRTS and LEAs.** Terminologies used in the two communities are sometimes different. ENISA recommends using the 'Common Taxonomy for Law Enforcement and The National Network of CSIRTs' [7].
- **Common data and document formats support.** Use of common formats e.g. Word, PDF, and CSV facilitate intelligence sharing where the use of specialised formats is not an option.
- **Capability to evaluate the reliability of the source of an intelligence data item.** All information sources should be assessed for reliability based on a technical assessment of their capability, or in the case of human intelligence source, their history.
- **Assessment of the credibility of an intelligence data item based on likelihood and levels of corroboration by other sources.** An EWS allowing a quick turnaround and fast decision making requires that the ingress information is trusted. The system should have mechanisms to assess the credibility of the information and include fake news protection mechanisms.
- **A shared workflow management system for incident handling.** This is one of the main purposes and core functionalities of the EWS, allowing also to monitor the effectiveness and efficiency of the system.
- **Trust-boosting security technologies.** Supporting the creation of closed communities and encrypted peer to peer communication.
- **Data redaction capabilities, for privacy compliance.** The system will need to redact personal information for data items marked to contain PI when exporting them to other EWS instances based on a privacy protection policy. For structured data, this can be done automatically. For unstructured data, this can be done semi-automatically, but may require human inspection and approval.
- **Attribution capabilities, identification of the origins of the source of information.** For traceability, disseminated information shall contain appropriate origin describing meta-data.

- **Anonymous sharing of information.** Despite the attribution requirements, it is advised that the system would still allow anonymous information, however, these items will need to be clearly marked as anonymous and is expected to have an impact on the reliability of the information.
- **Customisable exchange of intelligence data.** Customisation may be in accordance with internal (originating organisation) or external requirements.
- **Predefined criteria for data dissemination.** This relates to both the originator of the information (e.g. the criteria a set in accordance with audience, trust realms etc.) and the consumer of the information (e.g. data versions and revisions, severity, etc.)
- **Data normalisation.** The system shall normalise all ingress data under a common format, or data model. This will enable compatibility, interoperability and other functions (correlation).
- **A flexible data model.** Expansion of the data model is a prerequisite to allow EWS to grow across different domains and verticals. The system can allow custom creation of tags and the enrichment of existing IDIs. This could be automatic or manual. For example, an IDI may be enriched by external information from OSINT activities.
- **Correlation capabilities.** At a minimum level, the system should automatically link newly imported IDIs with existing IDIs.
- **Data items curation.** The system shall curate and de-duplicate IDIs imported from different sources and datasets. This is for ensuring that the integrity and accuracy of analytics is offered.
- **Advanced data analytics.** Situational awareness will be considerably supported from data analytics techniques (e.g. clustering and classification). This could include production of trends over time related data to support predictive analytics.
- **Visual analytics.** The system should provide visual analytics through a dynamic, interactive UI.
- **Pivoting capabilities.** In order to support the analytics processes and allow complex correlations and analytics, the system should offer pivoting capabilities over data.
- **Data exporting formats.** The system shall support exporting of data in different formats e.g. STIX, OpenIOC, CSV, Yara, sigma, etc.
- **Filtering capabilities.** The system should support filtering of information across a number of parameters and features. This also includes both whitelisting, blacklisting, to filter out benign activity and to pin down suspicious/malicious events.
- **Triaging.** The system should provide a high level overview of the data so that the analyst can quickly get a "gist" of what they contain. For example, for numerical data, the basic statistical information should be presented.
- **Alerting and communication.** This feature is required to improve the response times to incidents. This involves capabilities to match asset configuration with vulnerability information (for example describing assets as CPE and pairing with CVE and CVSS items) and sending a message to a designated contact point if a criticality level of an event exceeds some threshold. For example, this can be done if an asset described through a configuration is detected to be vulnerable to an exploit with a CVSS score.
- **Intelligence report generation.** The information shared should be available to the stakeholders in an appropriate format and level of detail.

4 Tabletop Exercise Approach

The tabletop exercise (TTX) was conducted as part of the evaluation of the development of the EWS platform. During the TTX the developers were given the chance not only to evaluate the technical features of the platform but also to test the assumptions mentioned above and ECHO's information-sharing architecture. It can be segmented in three stages: preparation, exercise conduct and post-analysis. The preparation of the exercise was based on the evaluation objectives, set from the ECHO consortium. For this TTX the defined objectives were the following:

- Assess the information-sharing policies of the EWS
- Evaluate the development of the EWS
- Detect and identify bugs and possible improvement
- Cyber attribution of the attacks

The next step of the preparation stage was to align the background of the exercise with current events or a major event that increases the alert status of a CSIRT or requires different incident management. In the case of the TTX, a mixture of the current Pandemic and a fictional EU-elections was used. An event like the EU-elections (and the extended period that is affected by it) can be exploited by the perpetrators of the cyber realm and this fact would create an environment that the participants would be prompted to optimise the information-sharing decisions and also attribute the incidents appropriately. To complete the backstory of the exercise, a scenario was created, according to which members of the consortium were the developers of an e-voting application that would allow the European citizens to vote electronically; requesting from the remaining members to contribute to the development by taking part in the evaluation and testing of the tool.

In the process of creating a coordinated framework that would coordinate the efforts of the participants, TTX also adopted the basic principles of the "Blueprint for Coordinated response to large scale cross border cyber security incidents and crises" [9]. Following the guidelines for increased technical and situational awareness, two entities were created and added to the platform with the purpose of providing appropriate information to the participants.

The remainder of the preparation process was focused on the preparation of the incidents that were developed under selected taxonomies and stages as follows.

The first taxonomy reflected on the identity of the receiver of the incident:

- Incidents addressed from the whole consortium
- Inter-sector incidents
- Sector-specific incidents
- Organisation-specific incidents
- The second taxonomy reflected on the content of the incidents
- EU-Elections focused content
- Covid-19 focused content
- Miscellaneous content

The incidents would involve the use of commonly known malware, recently released vulnerabilities, fake news dissemination and alerts for unauthorised activities. The actions were related and attributed, mostly, to APT-28, other APT groups and other

criminal actors. The intentions behind the choice of criminal actors were aiming to highlight that fact that cyber perpetrators could take advantage of a period when the attention is focused to a major event and target other areas that would appear of a lesser priority at that time.

The second stage of TTX, the exercise conduct, lasted three hours during which the participants were exposed to the incidents designed and delivered by the organisers. The participating teams were equally shared between two instance servers. The incidents were delivered in three waves with alternating volumes of information flow that would simulate a notion of realistic flow of incidents that are reported to a CSIRT.

Along with the injections, the teams were receiving hourlies, from two entities run by the organisers, as described in the preparation stage. The hourlies were delivered in three sets and offered the teams the opportunity to increase their awareness in the following topics:

- Voting applications
- E-voting
- APT groups
- APT-28
- Newly exposed vulnerabilities

The development of the scenarios of the injects aimed to trigger the exchange of CTI between the members of the EWS constituencies; the interaction, for each wave, between the teams and the related incident are displayed in Fig. 2.

During the three hours of the exercise the participants processed the information they received from the incoming reports and hourlies and proceeded on exchanging tickets internally and externally. After the conclusion of the exercise, the participants completed an online questionnaire focusing on the following aspects

- Profiling of the participants
- Information-sharing decisions
- Bug reporting
- Technical aspects of the tickets
- Evaluation of the exercise

The outcome of the questionnaire combined with the analysis of the log files of the instance servers are discussed in the following section.

5 Analysis

The two sources of the primary data used for the analysis were the logs of the EWS system and the questionnaire provided to the team members at the end of the exercise. The analysis includes merging these two sources to produce results for the exercise. The analysis is composed of three parts: i) the activity of the exercises including the sharing activities of the teams and constituencies; ii) the analysis on team dynamics which includes the following features: the number of team members, experience of team members and the correlation of those with the alive time of the tickets and the distribution of tickets; and iii) the characteristics retrieved from the nature of the tickets such as the assigned TLP levels and their distribution in general.

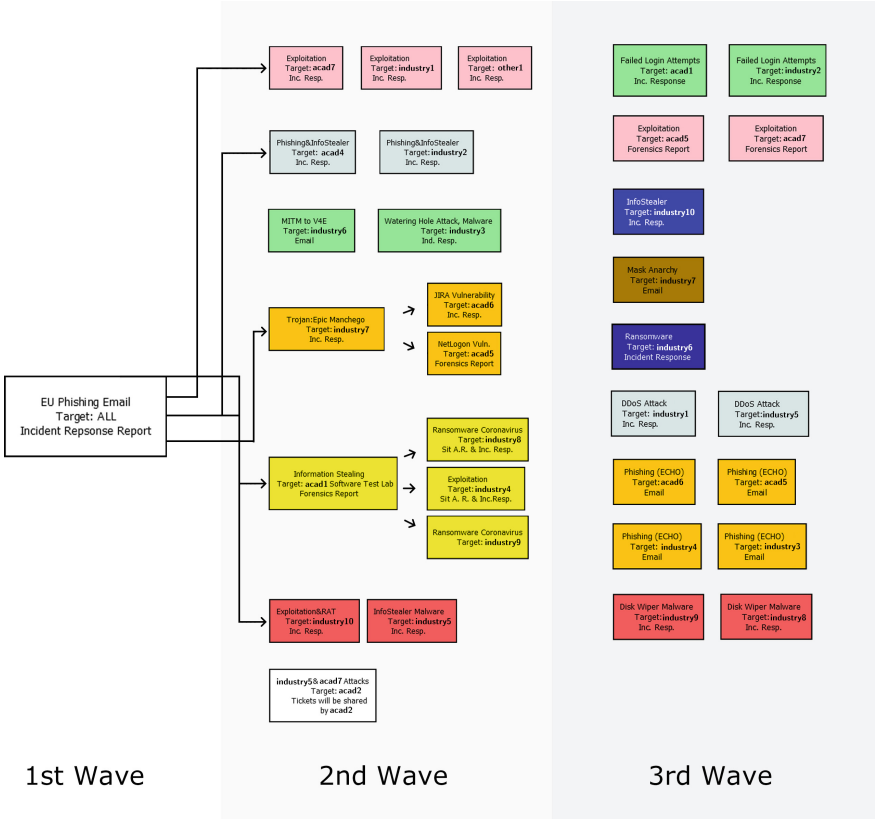

Fig. 2. Staging and timeline of the TTX

5.1 Exercise Activity Analysis

As explained in the description of the exercise, the teams are separated into logical constituencies defined by the election scenario. These constituencies consist of sectors of Defense, (European) Elections, IT and Healthcare. In addition to these 4 sectors, an ECCC constituency which represents member states from EU and a MEDIAROOM constituency were defined to distribute information to all of the audience from the participants. The MEDIAROOM constituency was created and utilised to share hourlies publicly by the organisers of the exercise.

Above all of these logical constituencies, the teams are coming from the industry or academia, therefore research and private sector constituencies were also utilised to create distribution channels between universities and private sector representatives respectively.

As can be seen from the Fig. 3, Elections Constituency which is the main theme of the scenario retrieves the most tickets during the exercise. Since the theme also includes European Elections process, the second high attraction constituency is the ECCC. The third one is the MEDIAROOM since this constituency is utilised for the distribution of the hourlies as aforementioned in the previous chapters.

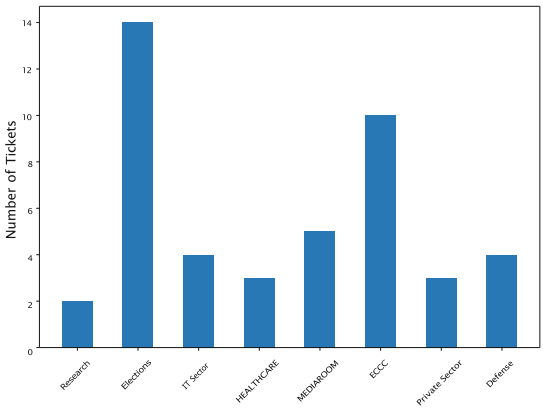

Fig. 3. Number of tickets shared by each constituency

The number of tickets shared by each organisation is given in Fig. 4. All the tickets are separated into two colours - internal (orange) and external (blue). As you can see from the figure, there are some organisations that didn't share any internal or external tickets. The organisation acad2 was included in the organisation committee, acad4 had technical issues and therefore they couldn't share any tickets.

On the day of the exercise, as explained above, incidents that are related to the elections theme of the scenario and incidents crafted for specific partner targets were released. It was seen that the majority of the partners shared these incidents via tickets after discussing them using the internal tickets while partners industry2, acad5 and industry5 only shared tickets internally.

As it can be seen from Fig. 4, the highest number of tickets are shared by the team acad3 although they received approximately the same amount of injects with all the partners. After the inspection of logs, this has been explained by the workflow decided by the team acad3. acad3 had decided to create several templates and workflows regarding the incident management; for each incident there are several tickets created for all the subtasks of the incident response. Therefore, most of the tickets were shared internally between these sub-teams of the incident response team of acad3 and after the closure of the event, a ticket is decided to be shared or not. To this end, we defined the metric extrovert to measure the percentage of tickets our unit of enquiry (team/constituency category) that were shared externally:

$$extrovert = \frac{\#externally\_shared\_tickets}{total\_tickets} \tag{1}$$

Following an analysis on a sample of 70 tickets, the academic constituency was (significantly) faster in closing the tickets than the industrial partners. In terms of making the tickets externally, there were no significant differences.

Figure 5 is an example analysis of the ticket alive time and also shows the breakdown the steps a ticket is going through the incident response workflow. Orange markers show the ticket operation such as assigning to a specific handler or updating the status of a

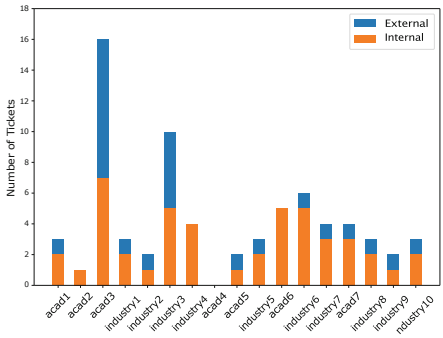

	Constituency category	
	Academic	Industry
tickets	30	40
$\bar{x}_{time}$	13.08	31.60
S_{time}	15.60	29.41
$\bar{x}_e$	0.543	0.612
S_e	0.41	0.40
t-test[prob]	mean time	-3.397[0.0011]
	extrovert	-0.700[0.4862]NS

Fig. 4. Tickets produced and shared by partners (Color figure online)

ticket. Green and blue markers represent attaching a document and referencing a resource to a ticket on the system respectively. Red markers show that a comment is added.

5.2 Team Analysis and Characteristics

Figure 6 shows the results of a regression on the number of team members against the produced number of tickets. We observe that the relationship between these two is exponential, showing that the larger the team size, the exponentially more the number of tickets are produced. However, there was no significant relation in the average lifetime of a ticket and the number of team members ($p = 0.778$).

5.3 Ticket Analysis

All the tickets that are created and shared within the system has a TLP level as described in Sect. 4. White tickets are for the general audience and can be considered open source while the red level being the disclosed information only distributed to the respondent of the incident. As can be seen in Fig. 7, the average alive time increases monotonically following the increase of confidentiality with a high correlation coefficient of $r = 0.9463$. This could be attributed to the nature of the cyber security incidents as the white and green level tickets are considered as public or generally publishable information and do not necessarily need an action while amber and red tickets are more organisation specific and need more time to analyse, plan and resolve. The pairwise comparisons between partners on their TLP assignments is also shown. While in many instances the sample size may not be considered large enough, the particular comparison of acad3 and industry3 which is significant is noteworthy, as these two partners have the highest number of tickets.

Figure 8 shows the results of a hierarchical clustering using Ward's method on the alive times (as shown with the example partner in Fig. 5) across all partners. Setting the cutoff at 4, we observe three clusters. Interestingly, while there are two small clusters with pure academic or industry members, the bigger cluster (in red) displays itself some visible segmentation between academia and industry, suggesting that these two constituency categories have emergent and esoteric behaviours and dynamics.

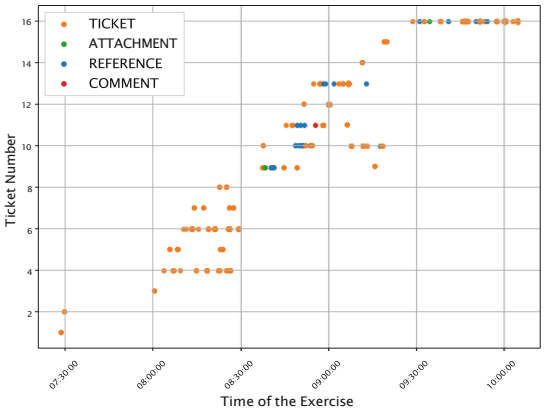

Fig. 5. Example life time of a ticket from the constituency acad3 (Color figure online)

Model Info			**Model Fit**	
Observations:17			F(1,15)=52.65, p=0.000	
Dependent variable: ltickets			R^2 = 0.778	
Type: OLS			Adj. R^2 = 0.763	
Coefficients				
	coef.	std. error	t	p
intercept	0.6174	0.120	5.140	0.000
members	0.2656	0.037	7.256	0.000

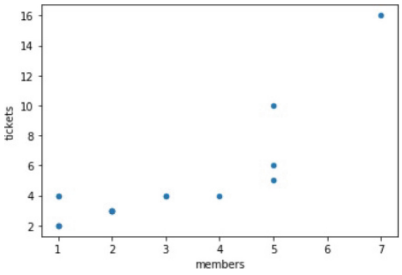

Fig. 6. Regression results (members~log(tickets))

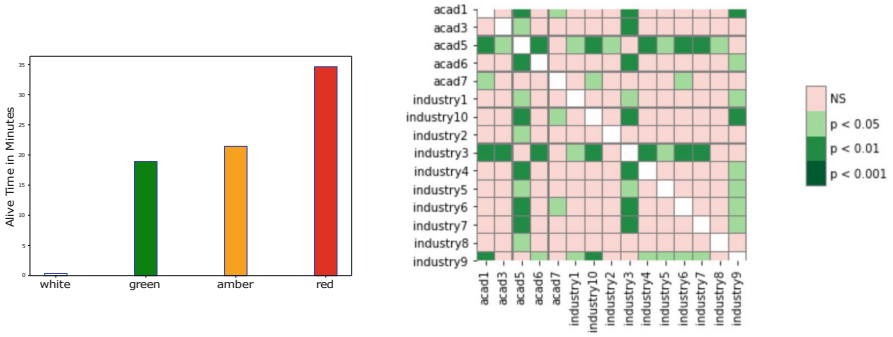

Fig. 7. TLP time and pairwise comparisons (Color figure online)

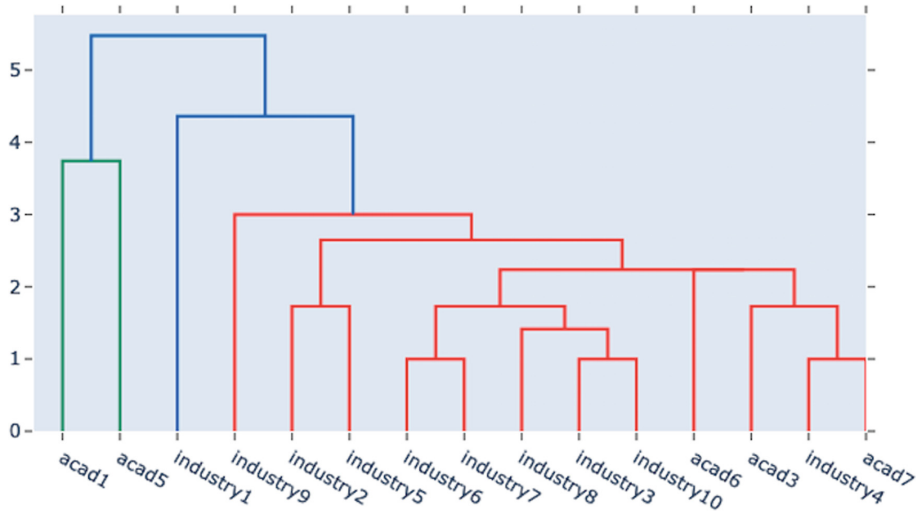

Fig. 8. Hierarchical clustering on ticket lifetime

6 Conclusions and Future Directions

In this paper we empirically evaluated the deployment of ECHO's Early Warning System through a structured tabletop exercise. We observed that the key features that were significant discriminators of the behaviours of the various constituencies were the TLP level and the size of the respective team. Specifically, the higher the confidentiality the ticket was deemed, the more time it took to resolve it. With regards to the size of the team, we noticed that the number of tickets grow exponentially to the size.

We also observed emerging clusters between academic and industry partners; this observation can be partially explained by the statistically significant difference between the average ticket alive times, where academic CSIRTs seems to be more than twice as fast in closing tickets than their industry counterparts.

The above findings can support and guide the development of additional functionality of an EWS. For example, the large alive times of TLP:RED tickets are non surprising, but warrant the investment in developing plugins to redact confidential data in a semi-automated manner.

The analysis exercise showed that there could be more metrics and features to be synthesized in order to assess the communication dynamics of the CSIRTS and the EWS system itself. In this paper we introduced the metric extrovert which shows the degree of a CSIRT sharing their tickets with their peers; however with the current dataset there were no significant differences observed with the average alive time.

For future work, a more expanded dataset would allow the application and exploration of this domain through more advanced algorithms.

Acknowledgement. This work has received funding from the European Union's Horizon 2020 research and innovation program under the grant agreement no 830943 (ECHO).

References

1. Directive (EU) 2016/1148 of the European Parliament and of the Council of 6 July 2016: Concerning measures for a high common level of security of network and information systems across the union. OJ **L 194**, 1–301 (2016)
2. Regulation (EU) 2019/881 of the European Parliament and of the Council of 17 April 2019 on ENISA (The European Union Agency for Cybersecurity) and on information and communications technology cybersecurity certification and repealing regulation (EU) no 526/2013 (cybersecurity act). OJ **L 151**, 15–69 (2019)
3. Bianchi, G., Dargahi, T., Caponi, A., Conti, M.: Intelligent conditional collaborative private data sharing. Futur. Gener. Comput. Syst. **96**, 1–10 (2019). https://doi.org/10.1016/j.future.2019.01.001
4. Burger, E.W., Goodman, M.D., Kampanakis, P., Zhu, K.A.: Taxonomy model for cyber threat intelligence information exchange technologies. In: Proceedings of the 2014 ACM Workshop on Information Sharing and Collaborative Security, WISCS 2014, pp. 51–60. Association for Computing Machinery, New York, NY, USA (2014). https://doi.org/10.1145/2663876.2663883
5. Cybersecurity and Infrastructure Security Agency: Elections Cyber Table-Top Exercise Package (2020)
6. ENISA: EUELEX 19 After Action Report, May 2019
7. ENISA, Europol/EC3: Common Taxonomy for Law Enforcement and The National Network of CSIRTs, p. 16 (2017). https://www.europol.europa.eu/sites/default/files/documents/commontaxonomyforlawenforcementandcsirtsv1.3.pdf
8. ENISA, C., Polska/Nask, C.: Actionable Information for Security Incident Response, November 2014
9. European Commission: Blueprint for Coordinated Response to Large Scale cross-border Cybersecurity Incidents and Crises (2017)
10. European Council: EU Cyber Defence Policy Framework. Technical report, Brussels, November 2018. https://www.consilium.europa.eu/media/37024/st14413-en18.pdf
11. Faiella, M., Gonzalez-Granadillo, G.: Enriching Threat Intelligence Platforms Capabilities (2016)
12. Ilves, L., Evans, T., Cilluffo, F., Nadeau, A.: European Union and NATO global cybersecurity challenges: a way forward. PRISM **6**(2), 126–141 (2016)
13. Kaufmann, H., Hutter, R., Skopik, F., Mantere, M.: A structural design for a pan-European early warning system for critical infrastructures. Elektrotechnik und Informationstechnik **132**(2), 117–121 (2015). https://doi.org/10.1007/s00502-015-0286-5
14. Kick, J.: Cyber exercise playbook. Cyber Exerc. Playbook **7013**(November), 1–40 (2014)
15. Li, V.G., et al.: Reading the tea leaves: a comparative analysis of threat intelligence, pp. 851–868 (2019)
16. Mavroeidis, V., Bromander, S.: Cyber threat intelligence model: An evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence. In: Proceedings - 2017 European Intelligence and Security Informatics Conference, EISIC 2017, pp. 91–98, January 2017. https://doi.org/10.1109/EISIC.2017.20
17. Rajamaki, J., Katos, V.: Information sharing models for early warning systems of cybersecurity intelligence. Inf. Secur. Int. J. **46**(2), 198–214 (2020)
18. Schaberreiter, T., et al.: A quantitative evaluation of trust in the quality of cyber threat intelligence sources. In: ACM International Conference Proceeding Series (2019). https://doi.org/10.1145/3339252.3342112

19. Wang, K., Du, M., Sun, Y., Vinel, A., Zhang, Y.: Attack detection and distributed forensics in machine-to-machine networks. IEEE Network **30**(6), 49–55 (2016). https://doi.org/10.1109/MNET.2016.1600113NM
20. White, G.B., Dietrich, G., Goles, T.: Cyber security exercises: testing an organization's ability to prevent, detect, and respond to cyber security events. Proc. Hawaii Int. Conf. Syst. Sci. **37**(C), 2635–2644 (2004). https://doi.org/10.1109/hicss.2004.1265411