AUTHOR'S ACCEPTED VERSION – DOI: [pending, paper in press]
International Conference on Cybersecurity, Situational Awareness and Social Media (Cyber Science 2022)

Practical Cyber Threat Intelligence in the UK Energy Sector

Alan Paice[1] and Sean McKeown[2] [0000-0001-7231-1682]

[1] EDF, UK
alan.paice@edf-energy.com
[2] Edinburgh Napier University, Edinburgh, Scotland
s.mckeown@napier.ac.uk

Abstract. The UK energy sector is a prime target for cyber-attacks by foreign states, criminals, 'hacktivist' groups and terrorists. As Critical National Infrastructure (CNI), the industry needs to understand the threats it faces to mitigate risks and make efficient use of limited resources. Cyber Threat Intelligence (CTI) sharing is one means of achieving this, by leveraging sector wide knowledge to combat ongoing mutual threats. However, being unable to segregate intelligence or to control what is disseminated to which parties, and by which means, has impeded industry cooperation thus far.

The purpose of this study is to investigate the barriers to sharing and to add to the body of knowledge of CTI in the UK energy sector, while providing some level of assurance that existing tooling is fit for purpose. We achieve these aims by conducting a multivocal literature review and by experimentation using a simulated Malware Information Sharing Platform (MISP) community in a virtual environment.

This work demonstrates that trust can be placed in the open-source MISP platform, with the caveat that the sharing models and tooling limitations are understood, while also taking care to create appropriate deployment taxonomies and sharing rules. It is hoped that some of the identified barriers are partially alleviated, helping to lay the foundations for a UK Energy sector CTI sharing community.

Keywords: Cyber Threat Intelligence · CTI · Information Sharing · Cybersecurity · Situational Awareness

1 Introduction

The UK energy sector is classed as Critical National Infrastructure (CNI) [39] and is therefore vital to UK national security, making it a high-profile target for cyber adversaries. Over the last decade a persistent threat of cyber espionage from hostile state actors towards the UK energy sector has been observed. Links between hostile state actors (HSAs) and cyber criminals are also growing. Hostile state actors have reportedly tasked hacking groups with malicious cyber activity such as data theft on their

behalf. The growing sophistication of the threat against the UK energy sector has led to a rethink to siloed defence of critical assets. With an inside view of the sector, it was noted that there is currently a duplication of effort and time when disseminating new threats which could allow an adversary to attack several organisations in turn, probing their defences and potentially exploiting a shared weakness. These issues indicate that there may be the potential to improve the collective defence of the industry as a whole.

By engaging in Cyber Threat Intelligence (CTI) sharing activities organisations can inform each other of cyber incidents in near real-time, allowing for the timely deployment of countermeasures. Reporting organisations could then expect that service in return. Although it appears to be a logical step to share threats discovered in an organisation in the same sector with others, the practice has not yet seen widespread adoption in the sector.

Information and intelligence on threats to the UK energy sector are obtained from open source, proprietary vendors via alerts sent by the National Cyber Security Centre (NCSC), and other agencies, as well as internal monitoring performed by each organisation. While this reporting goes some way to keep the sector informed, each separate part of the UK energy sector may not know which threats their competitors and peers are facing in real-time.

There are many reasons why CTI is not shared and many have attempted surveys, literature reviews and assorted studies. This paper examines some of the key reasons the UK energy sector has been slow to adopt CTI sharing and through use of an experimental case study, constituting several sharing models and taxonomies, we aim to provide assurance that these potential barriers can be overcome.

The main contributions of this paper are:

- An exploration of the barriers to entry for CTI sharing in the energy sector, and what has stopped its adoption in the past.
- Empirical testing of a trust model methodology to provide assurance for CTI deployment in the UK energy sector.

We perform a multivocal literature review to identify barriers, while making use of the open-source Malware Information Sharing Platform (MISP)[1] to simulate and empirically verify a variety of sharing models and intelligence tagging taxonomies identified in the literature.

2 Background

Sharing CTI, knowledge and information between peers and experts is seen as a critical countermeasure to the growing cyber threat [38]. Sharing can allow for a tremendous situational awareness and faster responses to emerging threats [44]. There are, however, potential barriers to sharing: standardisation, competition and trust have been identified among the reasons why organisations struggle to share [44]. The ob-

[1] https://www.misp-project.org/

jective of CTI sharing is the exchange of information and intelligence across traditional boundaries [7]. Chandel et al. [9] claim that CTI sharing is underdeveloped and limited by many technical barriers. They also argue that poorly defined CTI community standards lead to opportunistic and non-sharing behaviours, such as free-riding, where participation is limited to consuming the CTI and not contributing to the community [4]. There appear to be reasons deeper than organisations just wanting to consume CTI as it is seen that some entities are happy to pay for CTI from a central source [20]. Tounsi el al. [38] describe the benefits of sharing as 'undeniable', though other work claims that we currently lack empirical evidence to support such positive claims [43].

2.1 Barriers to CTI Sharing

We will first explore barriers to adopting mutual CTI sharing, before discussing the existing sharing ecosystem in the UK Energy Sector in Section 2.2.

Trust: Many of the themes in this section fundamentally reduce to some facet of trust in sharing, such as: confidence in the sharing mechanisms, legal oversight, and quality of intelligence. If there were total trust in a system of sharing, then many of the issues explored in the literature would be solved, and sharing would be seen as a totally normal and worthwhile endeavour. Trust in CTI sharing is a significant area of research and is ongoing [1,41]. A key theme which has been identified is that one of the challenges to solving CTI sharing is the establishment of a trust relationship between those entering a sharing relationship. Trusted relationships can take a long time to build and constant effort to sustain, i.e., they are hard to gain, and easy to lose [25]. Additionally, such relationships are difficult to build for untrusted participants [38].

Wagner et al. [41] examine CTI sharing platforms, noting that many of them establish trust manually, arguing that platforms such as Malware Information Sharing Platform (MISP) need to be found through traditional trust establishment techniques such as face to face meetings and between a closed circle of trusted members. The authors also suggest that this limits the usefulness of this small trust circle as much of the sharing participation is in private. Indeed, this behaviour may be a reflection of the human condition, with Tounsi et al. [38] suggesting that this may be an instinctual response to the unknown. Face to face trust is emphasised in the European FI-ISAC (Financial Institutes – Information Sharing and Analysis Centre), where stakeholders are required to attend meetings, being excluded if they fail to attend three successive meetings [15], building a degree of trust and investment in the relationships which are formed. In regulated sectors, such as the UK energy sector, which hold sensitive information, it could be argued, for the purposes of safeguarding sensitive and classified information, that not trusting without due diligence is a sensible approach. Such considerations are compounded due to some CTI sharing platforms conducting insufficient peer vetting when sharing information about vulnerabilities [41], potentially causing some organisations to have difficulty placing trust in them.

Wagner et al. [41] attempt to address the issue of trust in sharing approaches with a trust taxonomy. The taxonomy attempts to associate a trust level with a source of CTI

for its full life cycle. Peers rate the quality, timeliness and other criteria to generate a score. Sharing activities would reveal the presence of those peers who just want to consume (in which case another sharing model would be more appropriate for them). This approach is considered to be most appropriate for industries which widely share similar problems and experiences across peers, fostering a sense of mutual defence in a closed 'trusted community' [28,46].

Reputation: Reputational damage is one of the reasons a business may be reticent to make contributions, at least without some degree of anonymity. Bad press from a breach could damage the organisation's reputation [38] and make them and others more cautious about sharing openly. For organisations to participate in effective CTI sharing, they must build up a reputational capital and earn credibility, which may limit the participation of newer members. In a similar manner to trustworthiness, the reputation of stakeholders is gained over time and is damaged easily [42], but may be even more difficult to re-accumulate. Some form of anonymisation and unattributed information could help solve this worry [17], however there is currently no complete solution to the problem [10,20,32,34]. An additional concern is that the sharing of raw data could expose details of the victim's infrastructure and encourage other threat actors to attack based on the information presented [41].

Legal and Privacy Issues: The framework of obligations and information exchange required for CTI sharing invokes a complex regulatory landscape. Organisations find data protection and privacy laws as one of the biggest concerns for CTI sharing [25]. Of all the challenges examined, legal and privacy issues could stop organisations sharing at all. This is understandable, as the law is continuously changing and the risk appetite of legal action could be too much for some companies to manage. One example is the EU's adoption of the General Data Protection Regulation (GDPR), prompting large scale international alignment for the sake of harmony, the complexities of which are discussed by Sullivan and Burger [37]. One point of note is the introduction of several new categories of personal data, IP addresses, which could cause issues for sharing CTI, though in many cases this data may be processed for the purposes of legitimate public interest [25]. However, there are potential legal challenges which could be mounted [8], with issues pertaining to privacy being thematically similar to those of reputation and trust [44], such as with pseudo anonymisation of data.

Competition and Conflicts of Interest: The cooperation between competing firms where they are seeking a competitive edge, to protect commercial interests and intellectual property has been dubbed 'coopetition' [46]. While organisations work together in mutual interest to lower costs, blame culture and reluctance to admit fault can make it difficult to participate in information sharing [29]. There is also some concern that collusion attacks could cause organisations to be forced out of trusted communities, damaging reputations, such as by collectively scoring their contributions poorly [42].

Technical and Financial: Some organisations may be ready to share their threat intelligence but feel there are insufficient CTI sharing models (which are discussed in

Section 3.2) and collaboration platforms that cater to their particular needs, creating a barrier to entry [41,42]. This is particularly the case in sensitive or critical industries, such as the UK energy sector. 'Sharing security artefacts between industry peers is a technically complicated, slow, untrusted, and an overly bureaucratic task' [1]. Interoperability and automation have been highlighted as issues to CTI sharing, especially in peer to peer models [25]. Economic considerations are also at play. CTI sharing can be seen as being expensive [31], a drain on resources [44], or as a means of eroding competitive economic advantages [1].

Quality of Intelligence and Sources Incorrect or bad quality CTI can cause resources to be expended unnecessarily. One issue is that there are limited tools for formally validating report structures, such as the commonly used STIX [26] format. As a result, shared CTI data often include incomplete or incorrect information [27]. Additionally, there is concern around the quality and validation of indicators of compromise [9], particularly when pertaining to relevance, timeliness, accuracy, comparability, coherence and clarity [38]. It is also argued that many platforms are good at providing a quantity of data, rather than quality intelligence that can be actionable [33]. Abu et al. [3] state that there can be a problem with data overload and that 70% of feeds are 'sketchy and not dependable in terms of quality', they, however, offer no quantifiable scale to measure 'sketchiness'. A Quality of Indicators (QoI) model is proposed by Al-Ibrahim et al. [4] for the assessment of the level of contribution by participants in CTI sharing by measuring the quality rather than the quantity of participation. In this model the QoI and intelligence are sent to an assessor and given a score. The QoI model uses machine learning which contrasts with the system model [41] where industry peers decide on the quality. Currently there has been no research to compare which model could confirm the quality of CTI being shared.

2.2 Current CTI Sharing in the UK Energy Sector

Currently, the UK Energy Sector looks to the European Energy Information Sharing & Analysis Centre EE-ISAC [11] for guidance on CTI, as well as UK government sources [40]. Academic studies on CTI sharing in the energy sector are not numerous. There has been some work by governments and international organisations such as the EU. The Cyber Security in the Energy Sector Recommendations [13] gives a broad outline of best practices and broad guidance to share CTI via an energy Information Sharing and Analysis Centre (ISAC). The study also recommends a united interface for sharing with international allies. The study also finds that there is no EU supported pan-European trusted platform (e.g., ISAC) for exchanging CTI the energy sector.

The UK's Cyber Security Information Sharing Partnership (CiSP) and the European Network and Information Security Agency (ENISA) both help with the sharing of cyber threat information, allowing organisations to better detect campaigns that target particular industry sectors [17]. The Cyber Security Information Sharing Partnership (CiSP) is a joint industry and government initiative set up to exchange CTI as fast as possible as it was seen that the time it was taking to disseminate products was taking

too long. The goal of the project was to increase situational awareness and to reduce the impact on UK business of the increasing cyber threat to UK industry, especially that of critical national infrastructure such as the energy sector [36,44]. Launched in March 2013, CiSP now sits under the management of the NCSC, a part of the Government Communications Headquarters (GCHQ) [36,40]. Membership of CiSP provides the ability to securely engage with other government departments and industry peers and partners to seek advice and learn from each other. Discussion on CTI matters is encouraged at all levels, and the collaborative environment helps to provide an earlier warning of threats that had been seen before. CiSP also helps to improves the members ability to protect their assets and provides free access to network monitoring reports [30]. However, these approaches rely on a top-down sharing, rather than industry peer collaboration, potentially limiting its effectiveness [38].

2.3 Evaluation of Sharing Methods and Platforms

Compounding the problems noted above, particularly in relation to issues of trust, the robustness and effectiveness of the platforms which implement CTI sharing is unclear. Sauerwein et al. [32] note that there is sparse scientific analysis of the state of the art threat intelligence sharing platforms, with very little empirical research exploring this space [45]. Few overviews and comparisons of sharing platforms are available, and many of them are incomplete, sufficiently transparent or are outdated [7], with attempts to study tools being hindered by the lack of detailed information on proprietary platforms [7], or clear bias towards owned commercial products (e.g. [5]). Modern studies which evaluate tools either do so via a literature based evaluation [7], or by surveying organisations [10], rather than direct empirical testing.

The work presented in this paper aims to address some of the barriers and issues identified in this section by demonstrating, empirically, that existing opensource tooling can remove, or mitigate, some of these barriers, and provide a level of assurance that CTI can be shared effectively, and securely, in the energy sector and other sensitive industries.

3 Methodology

In order to explore the barriers preventing CTI sharing, and to develop an empirical understanding of whether the MISP platform is suitable for the energy sector, two main evaluations were performed in a simulated environment:

1. **Sharing Model Evaluation:** Exploring four main abstract sharing models: Source/Subscriber; Hub and Spoke; Peer to Peer; and Hybrid (described in Section 3.2 and Table 1). Models were evaluated in order to determine their suitability for the energy sector, and whether they address barriers identified in the literature.

2. **Taxonomy and Tag Evaluation:** To determine whether the use of tagging and application of taxonomies (described in Section 3.3) can be used to explicitly en-

hance *privacy* and *trust*, both of which have been identified as critical barriers to entry. The same approach is taken for the sharing model evaluation, however, there are scenarios where CTI is classified only for use in UK organisations, limiting the propagation of specific documents, which should take precedent over the deployed sharing model.

3.1 Experimental Configuration

The CTI sharing tool used in this work, MISP, was chosen as it is open-source software, which is freely available, mitigating some financial barriers, while also seeing more usage in the industry as the emphasis shifts to attacker Tactics, Techniques and Procedures (TTPs), rather than simply Indicators of Compromise (IoCs) aggregation [19]. The aim is not to test MISP as a tool, as such, but rather to demonstrate that identified barriers could be overcome with a little assurance and some empirical experimentation.

Ten intelligence reports were obtained from open-source threat intelligence providers, all of which are marked for free distribution. A large sample was not required, as scaling is not tested, only the sharing behaviour of the platform. Reports were ingested manually, with the associated JSON file for MISP being made available on Github[2], presented in the STIX format [26].

The structure of a cyber event on MISP (e.g., attack, malware identification), can be split into three phases: event creation; populating of attributes and attachments; and publishing/sharing. Each event can contain multiple reports pertaining to the same attack vector, malware, etc. Manually populating these events allowed for repeated, controlled, experimentation. Repeat experimentation was facilitated via use of Virtual Machines (VMWare Fusion Pro 11.1.0) and snapshots, such that each run could be repeated without introducing ordering effects, re-importing data each time. This was necessary as MISP creates a 128-bit Universal Unique Identifier for each event (UUID)[18], which prevents deleted items from being re-imported.

Multiple MISP instances, based on the master MISP OVA files (version 2.4.130) were run (one for each entity in the sharing models in Section 3.2) to allow for the propagation of events to be confirmed. The MISP Hardware Sizer[3] was used to determine that 1 virtual CPU and 2GiB of RAM would suffice for each instance.

3.2 CTI Sharing Models

The four sharing models discussed in the literature are outlined here. For convenience a summary of the models is depicted in Table 1, with a brief discussion of each presented afterwards.

[2] https://github.com/smck1/Energy_CTI_Experimental_Files
[3] https://www.misp-project.org/MISP-sizer/

Table 1. CTI Sharing Model Overview

Model Name	Description	Use case
Source and Subscriber	— Single central source shares with subscribers — Subscribers do not share with central source — The peers do not share with each other directly	A subscription to a government RSS feed or email list that is giving a subscriber regular update on threats.
Hub and Spoke	— Peers share with a central hub — The central hub shares with peers — The peers do not share with each other directly	A central intelligence repository such as a government agency who wish to produce intelligence for consumption but wish to have feedback and intelligence fed back to the central source in a 1:1 relationship.
Peer to Peer	— Each Peer shares with each other — No need for any intermediary hubs or central sources	Peer to Peer is post to all. A detected threat could be shared rapidly to all members in the mesh.
Hybrid	Any combination of: — Peer to peer — Hub and Spoke	A sector providing feeds to peers in that industry such as energy. The peers could share their own CTI with each other and not have to report back to the central source if they so choose.

Source/Subscriber: Also known as the centralised model [20], the simplest of the four models, where CTI is consumed by the subscribers, but not produced by them. Subscribers require a great deal of trust in the source to participate in this model [38]. In many cases will come from a national source [2].

Hub and Spoke: This model places a central clearing house for CTI sharing that intelligence or information with the spokes. The spokes consume CTI from the hub and can share back CTI with the central hub in a 1:1 relationship. The model allows the group of intelligence producers and consumers to share information [34]. Where private companies direct the source, this can be seen as controversial as they may be more interested in profit or competition with other providers [35].

Peer to Peer: Also known as the decentralised model [20], this approach may be the most effective in the sharing of CTI between peers in similar sectors and

removes the problem of needing to trust a central repository [1]. A robust system of trust and regulatory conformity needs to be established between each of the peers [42]. If implemented successfully, this model could be compelling for sharing thematically similar problems and solutions to those peers that are qualified in that field of interest [34]. The model could help produce timely and actionable products much faster than waiting for a central entity to make those decisions on what the end peer receives.

Hybrid: A combination of the Peer to Peer and Hub and Spoke models [42]. Noor et al. [24] state that the hybrid model enables the best elements of both models, with Peer to Peer elements effectively collecting strategic CTI, while Hub and Spoke behaviour adds value to the raw CTI. The Hybrid model allows CTI from a central source to be shared with their peers [35], however, this could cause some issues for some proprietary platforms where the licensing only allows exclusive use by the subscriber and concerns around classification and source protection. This model could be useful for the energy sector, where feeds come from central government sources, with members being vetted before joining a community of sharers [6], while also facilitating timely peer sharing.

3.3 Tags and Taxonomies

- OFFICIAL – SENSITIVE: SNI *(Sensitive Nuclear Information)*
- LEGALLY PRIVILEGED
- PROTECT – PRIVATE
- PROTECT – COMMERCIAL & CONTRACTS
- UK SUBJECT TO EXPORT CONTROL
- NOT PROTECTIVELY MARKED
- EXPORT CONTROLLED (EC)
- SENSITIVE BUSINESS INFORMATION (SBI)

Fig. 1. EDF UK Protective Marking Taxonomies

Many terms are used in the literature to define ordered classifications systems used in threat intelligence[4], to facilitate understanding in both technical and non-technical consumers. While there is currently no consensus on concepts and definitions related to CTI taxonomies [12,22], all that is required for our definition is that a taxonomy groups objects and describes their relationships, has a set vocabulary, and maps this knowledge in a readable format [14].

A comparison of every such taxonomy would not be feasible in this study, as the research is ongoing, and taxonomies are continually evolving with the everchanging nature of the threats [14,22,41]. As such, we will focus on terminology used in the

[4] e.g., 'information exchange standard', 'ontology', 'taxonomy', 'data type', 'data/field format standard', 'data/Field representation format', 'classification', 'semantic vocabulary', 'field' and 'knowledge map', 'Machine tag' [14].

UK energy sector, which uses the Traffic Light Protocol (TLP), UK government protective marking and custom company protective marking. An example given from EDF is depicted in Figure 1.

The Traffic Light Protocol was created by the UK's Centre for Protection of National Infrastructure (CPNI), a UK government agency. The TLP design encourages the sharing of sensitive information and helps establish trust within the sharing of information and intelligence. Its purpose is to ensure that sensitive information is shared with the appropriate audience. It is not, however, a classification marking scheme on its own, more that it is used as an indicator to reflect how sensitive the information or intelligence is to aid in collaboration. The current standard is defined by the Forum of Incident Response and Security Teams (FIRST) Standards Definitions and Usage Guidance [16]. Restrictions are colour coded: TLP:RED for non-disclosure; TLP:AMBER for limited disclosure to participant's organisations; TLP:GREEN for limited disclosure to a restricted community; and TLP:WHITE for unrestricted disclosure [17,23].

In order to evaluate sharing taxonomies, and determine an order of primacy (sharing model vs. document sharing restrictions), the EDF taxonomy depicted in Figure 1 was converted to the JSON format (Figure 2) and imported into MISP (Figure 3).


```
{
    "namespace": "UK-Energy-Sector ",
    "expanded": "UK Energy Sector",
    "description": "Protective markings used in a typical UK Energy sector organisation",
    "version": 3,
    "predicates": [{
            "value": "OFFICIAL – SENSITIVE : SNI",
            "expanded": "OFFICIAL – SENSITIVE : Sensitive Nuclear Information",
            "description": "Information relating to activities carried out on or in relat
        },
        {
            "value": "Sensitive Business Information",
            "expanded": "Sensitive Business Information (SBI)",
            "description": "Data that covers commercial, legal and personal data. This in
        },
        {
            "value": "PROTECT – PRIVATE",
            "expanded": "PROTECT – PRIVATE",
            "description": "Private data"
        },
        {
            "value": "PROTECT – COMMERCIAL & CONTRACTS",
            "expanded": "PROTECT – COMMERCIAL & CONTRACTS",
            "description": "Commerical contracts"
        },
        {
            "value": "LEGALLY PRIVILEGED",
            "expanded": "LEGALLY PRIVILEGED",
            "description": "Legal documents"
        },
        {
            "value": "Export Controlled (EC)",
            "expanded": "Export Controlled (EC)",
            "description": "Subject to export controls "
        },
```


Fig. 2. Snippet of the custom taxonomy created for this experiment, in raw JSON format.

3.4 Sharing in MISP

Your Organization Only: Intended for internal dissemination only. Events with this setting will not be shared outside of the instance. Upon Push: do not push. Upon Pull: pull only for internal users.

This Community Only: Users that are part of the local MISP community will be able to see the event. This includes the internal organisation, organisations on this MISP server and organisations running MISP servers that synchronise with this server. Upon Push: do not push. Upon Pull: pull and downgrade to 'Your Organization Only'.

Connected Communities: Extends 'This Community Only' to servers connected to synchronising servers (i.e., extending to two hops away from the originating instance). Any other organisations connected to linked instances that are two hops away from this own will be restricted from seeing the event. Upon Push: downgrade to 'This Community Only' and push. Upon Pull: pull and downgrade to 'This Community Only'.

All communities: This will share the event with all MISP communities, allowing the event to be freely propagated from one instance to the next. Upon Push: push. Upon Pull: pull.

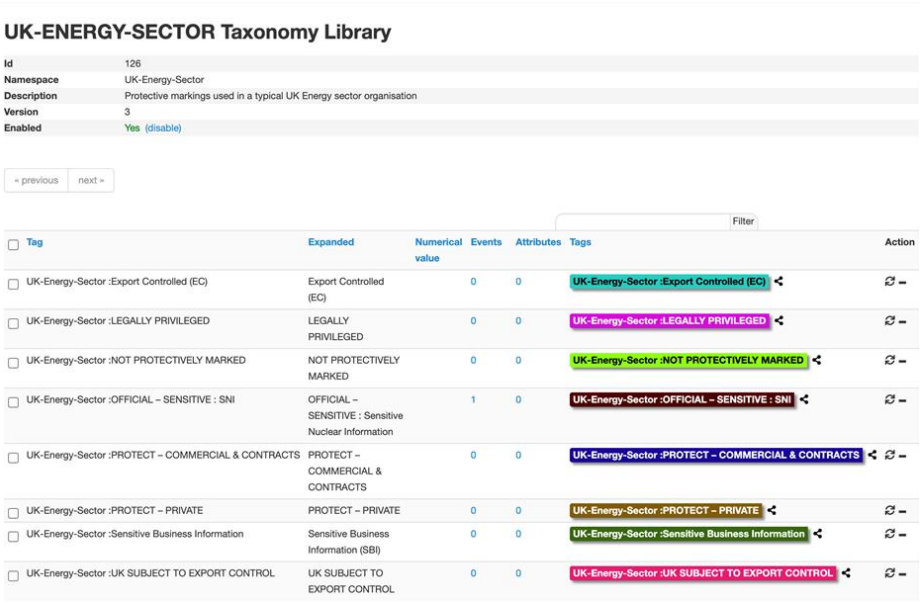

UK-ENERGY-SECTOR Taxonomy Library

Id	126
Namespace	UK-Energy-Sector
Description	Protective markings used in a typical UK Energy sector organisation
Version	3
Enabled	Yes (disable)

« previous | next »

Filter

Tag	Expanded	Numerical value	Events	Attributes	Tags	Action
UK-Energy-Sector :Export Controlled (EC)	Export Controlled (EC)		0	0	UK-Energy-Sector :Export Controlled (EC)	
UK-Energy-Sector :LEGALLY PRIVILEGED	LEGALLY PRIVILEGED		0	0	UK-Energy-Sector :LEGALLY PRIVILEGED	
UK-Energy-Sector :NOT PROTECTIVELY MARKED	NOT PROTECTIVELY MARKED		0	0	UK-Energy-Sector :NOT PROTECTIVELY MARKED	
UK-Energy-Sector :OFFICIAL – SENSITIVE : SNI	OFFICIAL – SENSITIVE : Sensitive Nuclear Information		1	0	UK-Energy-Sector :OFFICIAL – SENSITIVE : SNI	
UK-Energy-Sector :PROTECT – COMMERCIAL & CONTRACTS	PROTECT – COMMERCIAL & CONTRACTS		0	0	UK-Energy-Sector :PROTECT – COMMERCIAL & CONTRACTS	
UK-Energy-Sector :PROTECT – PRIVATE	PROTECT – PRIVATE		0	0	UK-Energy-Sector :PROTECT – PRIVATE	
UK-Energy-Sector :Sensitive Business Information	Sensitive Business Information (SBI)		0	0	UK-Energy-Sector :Sensitive Business Information	
UK-Energy-Sector :UK SUBJECT TO EXPORT CONTROL	UK SUBJECT TO EXPORT CONTROL		0	0	UK-Energy-Sector :UK SUBJECT TO EXPORT CONTROL	

Fig. 3. Snippet of the custom taxonomy created for this experiment, as represented in MISP.

As this work focuses on dissemination to external organisations, experiments use both the 'Connected Communities' and 'All Communities' models.

Experiments were repeated 25 times to assure consistent behaviour and were carried out by a trained intelligence analyst. Tested configurations use Push unless otherwise specified. Five Peers were simulated: The local EDF UK organisation, and four external peers, with names being chosen to demonstrate potential example peers, but otherwise referred to in the results as v1 (EDF UK) and v2 – v5 (external peers).

Discussion pertaining to the verification of sharing models is presented in Section 4, while Section 5 discusses the validation of tags and taxonomies.

4 Verifying Sharing Models

4.1 Source and Subscriber

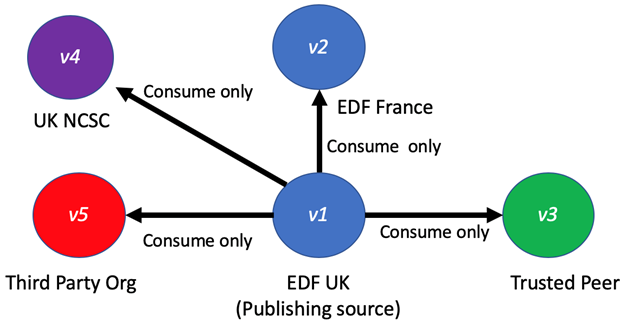

(a) Source and Subscriber model.

$$\text{Shared from} \begin{array}{c} \\ v1 \\ v2 \\ v3 \\ v4 \\ v5 \end{array} \begin{array}{c} \begin{array}{ccccc} v1 & v2 & v3 & v4 & v5 \end{array} \\ \begin{pmatrix} 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \end{array}$$

(b) 'Connected' and All Communities results (identical).

Fig. 4. Source and Subscriber sharing model results.

Beginning with the simplest case, a single source disseminates information to consumers (Figure 4a), allowing for verification that this works as intended, and that sharing is uni-directional.

The directed adjacency matrix, depicting when data is propagated to each party, is provided in Figure 4b. A '1' indicates event propagation from the party vy (left) to vx (top), while a '0' indicates that the event was not shared. By convention we ignore instances of a party sharing with itself, such that (v1,v1) (EDF UK, the source sharing with itself) is 0.

In this case both Connected and All Communities produced the same results, which were both in-line with expectation. All subscribers received the events shared with them. Once the events were published the subscribers received the events. This model does not allow flow back to the source and that was replicated with this model. However, it should be noted that server misconfiguration can cause this model to fail, which could cause serious data breaches.

4.2 Hub and Spoke

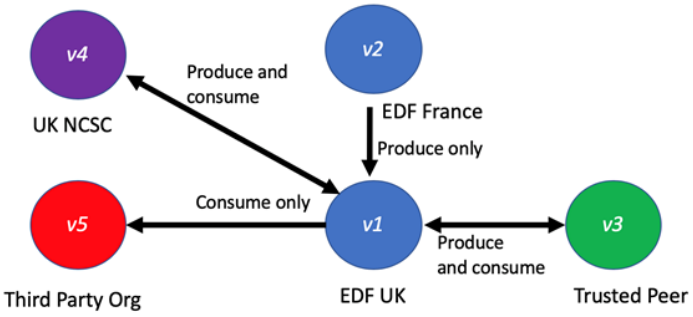

(a) Hub and Spoke model.

Shared from \	$v1$	$v2$	$v3$	$v4$	$v5$
$v1$	0	0	1	1	1
$v2$	1	0	0	0	0
$v3$	1	0	0	0	0
$v4$	1	0	0	0	0
$v5$	0	0	0	0	0

Shared from \	$v1$	$v2$	$v3$	$v4$	$v5$
$v1$	0	0	1	1	1
$v2$	1	0	**0**	**0**	**0**
$v3$	1	0	0	**0**	**0**
$v4$	1	0	**0**	0	**0**
$v5$	0	0	0	0	0

(b) (left) 'Connected Communities' results.
(c) (right) All Communities results. Unexpected results are marked in red.

Fig. 5. Hub and Spoke sharing model results.

In Hub and Spoke there is again a central source, but it can receive events from the consumers, though they do not share with each other directly (Figure 5a). One of the behaviours discovered here was that the event was not automatically shared beyond the first hop as the event was downgraded to share with 'This Community Only'. This was predicted for the Connected Communities setting (Figure 5b) (though sharing had to be triggered manually by the administrator after the first hop), but the same behaviour was also present for the All Communities setting, which was unexpected. Note the red/underlined items in Figure 5c which denote sharing that defied expected behaviour. In this case propagation required re-publishing to reach the second hop, which was not expected behaviour for All Communities. Fortunately, the issue results in under-sharing, rather than over-sharing, meaning that no critical information can be accidentally disclosed.

4.3 Peer to Peer

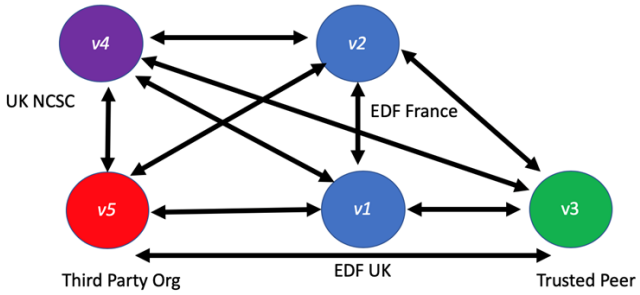

(a) Peer to Peer model.

$$\text{Shared from} \begin{array}{c} \\ v1 \\ v2 \\ v3 \\ v4 \\ v5 \end{array} \begin{array}{c} \begin{array}{ccccc} v1 & v2 & v3 & v4 & v5 \end{array} \\ \begin{pmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 \end{pmatrix} \end{array}$$

(b) 'Connected' and All Communities results (identical).

Fig. 6. Peer to Peer sharing model results.

Peer to Peer connects peers directly to each other in a distributed fashion, with no central source (Figure 6a). As per expectation, regardless of whether All or Connect Communities settings were used, peers propagated to each other without fail (Figure 6b). Each of the four external nodes was updated immediately upon publication by the sharing node. Due to the downgrading (to 'This Community') that happens at each hop, events did not propagate more than a single hop, preventing unintentional disclosure.

4.4 Hybrid

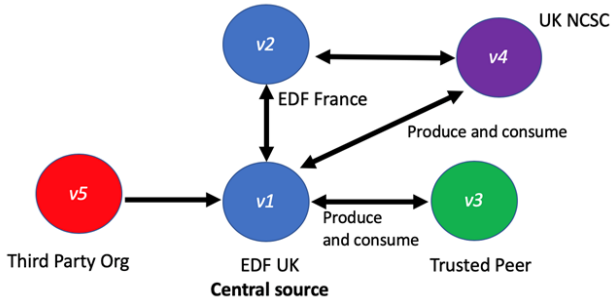

(a) Hybrid model - Firstly in all PUSH configuration, with a second run where PULL is enabled from *v*5 to *v*1.

Shared from	v1	v2	v3	v4	v5
v1	0	1	1	1	0
v2	1	0	0	1	0
v3	1	0	0	0	0
v4	1	1	0	0	0
v5	1	0	0	0	0

Shared from	v1	v2	v3	v4	v5
v1	0	1	1	1	0
v2	1	0	**0**	1	0
v3	1	**0**	0	**0**	0
v4	1	1	**0**	0	0
v5	1	**0**	**0**	**0**	0

(b) (left) 'Connected Communities' results. PUSH and PULL configurations.
(c) (right) All Communities PUSH results. Unexpected results are marked in red.

Shared from	v1	v2	v3	v4	v5
v1	0	1	1	1	0
v2	1	0	0	1	0
v3	1	0	0	0	0
v4	1	1	0	0	0
v5	1	**1**	**1**	**1**	0

(d) All Communities PULL results. Differences to PUSH are highlighted in bold.

Fig. 7. Hybrid sharing model results.

The Hybrid model is the most flexible, facilitating peer sharing, but optionally making use of a centralised source, with optional upstream sharing. Once again the Connected Communities (Figure 7b) behaviour was as expected, with downgrading meaning that events are only propagated a single hop. However, in the default Push configuration, once again the All Communities setting resulted in some unexpected behaviour (Figure 7c, with unexpected results in red/underlined). While unexpected a priori, these results are consistent with the findings for the Hub and Spoke model.

In order to explore the propagation issues, this setup was tested using a Pull based configuration to EDF UK (v1) from the Third Party Organisation (v5). Once the EDF UK instance pulled the events from the Third party, the other instances shared with each other, demonstrating the predicted behaviour (Figure 7d).

MISP's behaviour in PUSH mode was then considered to be a bug at this point, with results being reported to the project[5]. The MISP project responded with an update that the description of delegation behaviour in their documentation did not align with the actual behaviour, such that it is actually intended, but was initially misrepresented.

5 Verifying Taxonomies - Events and Tags

The testing in Section 4 demonstrated that many models in MISP require manual intervention in order for events to pass beyond the neighbouring nodes. As such, the Peer to Peer model was used to test tagging behaviours as we only need to focus on whether the event is initially shared or not, as opposed to its subsequent tertiary prop-

[5] https://github.com/MISP/misp-book/issues/202

agation. Tags correspond to the levels from the EDF energy sector taxonomy in Figure 1. Rules for propagation are set by the node, and individual CTI documents are tagged with these properties at creation time by analysts. Additional rules are set to Allow or Block specific organisations, with examples rules, in a Boolean combination, depicted in Figure 8. It should be noted that MISP was found to be very sensitive to formatting of the imported custom taxonomies, and care should be taken when creating and validating the JSON input (Figure 2).

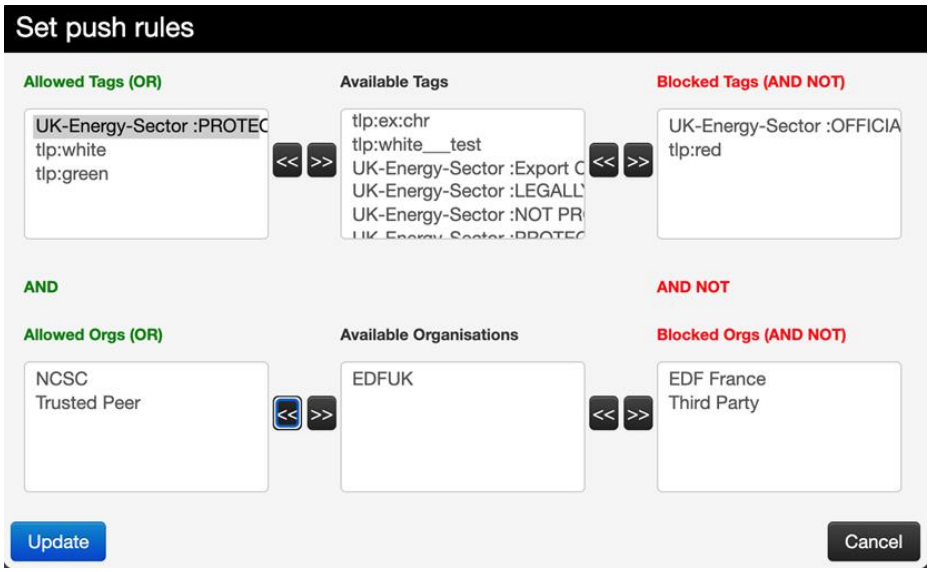

Fig. 8. Example MISP tag configuration - Allowed/Block propagation rules using Boolean combinations.

Several mixes of rules were explored in order to determine the order of precedence, and whether or not unintentional sharing could be triggered, resulting in information breaches.

Allow All / Block All: Sharing all tags and blocking all tags in order to ascertain whether basic rules work as intended.

Combinations of Allow and Block: A mix of allow and block, again to verify basic functionality.

Contradicting Tags: Opposing tags - verifying, for example, that an event tagged with TLP:GREEN (marked as Allow) and OFFICIAL – SENSITIVE: SNI (marked as Block) would ideally be blocked from sharing. In reality similar occurrences would likely be organisations (or groups thereof) specific rules. This testing also determines if multiple Allow tags override a single Block tag.

Across all combinations of tags, the ideal behaviour was observed. Events were Allowed or Blocked, and in cases where multiple tags were assigned, Block always takes precedence over Allow. No number of Allow tags override a single Block tag. For CTI sharing, particularly in critical infrastructure, this is optimal behaviour as it avoids potentially dangerous over-sharing of sensitive information. This behaviour holds true when setting rules for specific organisations, where Blocked organisations override event tags.

6 Discussion

The various abstract models presented in the literature can be replicated in free, open, platforms such as MISP. The experiments have verified that such tools, when configured appropriately, are fit-for-purpose, with the caveat that some behaviour should be verified for deployment (as with any piece of software). In the case of the experiments in this paper, the unexpected behaviours were simply miscommunications in the documentation, however misconfiguration was also verified as being an avenue leading to unwanted disclosure. This is particularly critical in the energy sector, where sharing of sensitive information can be an offence of criminal negligence, not to mention the reputational damage and financial consequences. Just as it is essential for a human analyst to be involved in creating the intelligence [42], similar care must be taken when maintaining CTI sharing systems. In particular, the creation of sharing taxonomies must be done carefully, with MISP being sensitive to formatting errors. Applying organisation based filters appears to be a good way to mitigate some of these problems. While relatively minor in practice, the issues discovered here could cause serious consequences, meaning that the authors recommend more public testing of platforms is pursued going forward, in addition to testing carried out before deploying a new tool, or version, in-house.

The exposition of the various sharing models in this work should be useful for implementation going forward, with the Hybrid model being flexible enough to allow peer participation and nuanced dissemination controls. Assuming an understanding of the sharing implementation, and correct deployment, this testing should provide some level of assurance that tools such as MISP are suitable to engage in inter-organisation CTI sharing, and to mitigate risks of trust and over-sharing resulting in reputational damage.

It should also be noted that there is a degree of anonymity built-in when sharing events, as no evidence of the source of the event after the first hop, meaning that only direct trusted peers would have direct source identification, unless specified in the documents themselves. This in itself may reduce the perception of appearing vulnerable, and influence the rate of sharing and cooperation which is usually tempered by the competitive barriers. Ultimately, as the NCSC put in their annual report in 2019: ‘improving the cybersecurity of the UK is far from a solo effort.’ [21], and the fewer barriers to wider, meaningful, participation in CTI sharing, the better.

7 Conclusions

A number of barriers to Cyber Threat Intelligence (CTI) sharing in the UK energy sector have been discovered, with a focus on assurance to alleviate issues caused by the barriers of trust; reputation; competition; and technical/financial constraints. A simulated Threat intelligence sharing infrastructure was created in a virtual environment using the open-source Malware Information Sharing Platform (MISP). Four CTI sharing models from the literature were tested using real CTI data, processed by a trained CTI analyst to emulate a live system as closely as possible.

While the testing discovered some deployment issues with the platform, they provide a level of assurance that the segregation and dissemination of CTI data, for some set of criteria, should be possibly when carefully crafted taxonomies, tags, and forwarding rules, are maintained by trained professionals. In particular, safeguards on the platform would prevent critical events such as accidental dissemination of classified or sensitive information, which can be demonstrated to regulators. We do, however, recommend that there is an increased rate of publicly disseminated testing and assurance information made available in the CTI domain in future, as it will help avoid potential issues and bugs, and therefore dissemination catastrophes.

The COVID19 global pandemic has demonstrated that UK energy sector organisations cannot stand alone in ever-increasing complex threats to critical national infrastructure. MISP allows for a lower barrier of entry in terms of cost and ease of getting CTI shared quickly. With adequate testing, fears can be assuaged (such as those of over-sharing), and an appropriate sharing model for the industry can be adopted to allow for high-levels of participation. The use of MISP in the UK energy sector would enable the threats to the industry to be shared quickly and securely, and this study can be used as a foundation to encourage wider CTI sharing in the sector, ideally in the near future.

7.1 Future Work

Future work could increase the interactions between academics; intelligence practitioners; and energy sector companies, in order to create a deeper understanding of the field and add to the body of knowledge. The energy sector is highly regulated, and higher levels of assurance can be achieved by working together to test and formalise safe and secure processes. Automating the testing of CTI processing would allow for frictionless deployment while increasing assurance. The use of directed adjacency graphs could help model large sharing networks and predict future behaviour of sharing networks, which could accurately predict results for a system working as it should and even when it is not as observed during experimentation. Other pragmatic concerns could revolve around working with live feeds in order to better understand the characteristics of large scale CTI networks, and to include multiple levels of internal dissemination in the modelling.

References

1. Sharing is Caring: Collaborative Analysis and Real-time Enquiry for Security Analytics. 2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData) (2018). https://doi.org/10.1109/Cybermatics 2018.2018.00240
2. CiSP - NCSC.GOV.UK (2020), https://www.ncsc.gov.uk/section/keep-up-todate/cisp
3. Abu, M.S., Selamat, S.R., Ariffin, A., Yusof, R.: Cyber threat intelligence – Issue and challenges. Indonesian Journal of Electrical Engineering and Computer Science **10**(1), 371–379 (apr 2018). https://doi.org/10.11591/ijeecs.v10.i1.pp371-379
4. Al-Ibrahim, O., Mohaisen, A., Kamhoua, C., Njilla, L.: Beyond Free Riding: Quality of Indicators for Assessing Participation in Information Sharing for Threat Intelligence (2017). https://doi.org/10.1145/1235.
5. ANOMALI: The Definitive Guide to Sharing Threat Intelligence. Tech. rep. (2019), https://www.anomali.com/resources/whitepapers/the-definitive-guide-tosharing-threat-intelligence
6. Bakis, B.J., Wang, E.D.: Building a National Cyber Information-Sharing Ecosystem. Tech. rep., MITRE (2017), https://www.mitre.org/publications/technicalpapers/building-a-national-cyber-information-sharing-ecosystem
7. Bauer, S., Fischer, D., Sauerwein, C., Latzel, S., Stelzer, D., Breu, R.: Towards an Evaluation Framework for Threat Intelligence Sharing Platforms. In: Proceedings of the 53rd Hawaii International Conference on System Sciences. Hawaii International Conference on System Sciences (2020). https://doi.org/10.24251/hicss.2020.239
8. Borden, R.M., Mooney, J.A., Taylor, M., Sharkey, M.: Threat Information Sharing and GDPR: A Lawful Activity that Protects Personal Data. Tech. rep., FS-ISAC (2018)
9. Chandel, S., Yan, M., Chen, S., Jiang, H., Ni, T.Y.: Threat Intelligence Sharing Community: A Countermeasure Against Advanced Persistent Threat. In: Proceedings - 2nd International Conference on Multimedia Information Processing and Retrieval, MIPR 2019. pp. 353–359. Institute of Electrical and Electronics Engineers Inc. (apr 2019). https://doi.org/10.1109/MIPR.2019.00070
10. Chantzios, T., Koloveas, P., Skiadopoulos, S., Kolokotronis, N., Tryfonopoulos, C., Bilali, V.G., Kavallieros, D.: The Quest for the Appropriate Cyber-threat Intelligence Sharing Platform. In: Proceedings of the 8th International Conference on Data Science, Technology and Applications. pp. 369–376. SCITEPRESS - Science and Technology Publications (2019). https://doi.org/10.5220/0007978103690376.
11. EE-ISAC - European Energy - Information Sharing & Analysis Centre: EEISAC - European Energy - Information Sharing & Analysis Centre (2020), https://www.ee-isac.eu/
12. ENISA: Cybersecurity Incident Taxonomy. Tech. rep. (2018), https://ec.europa.eu/informationsociety/newsroom/image/document/2018-30/cybersecurity incident taxonomy 00CD828C-F851-AFC4-0B1B416696B5F710 53646.pdf
13. EU Commission: Cyber Security in the Energy Sector Recommendations for the European Commission on a European Strategic Framework and Potential Future Legislative Acts for the Energy Sector. Tech. rep. (2017)
14. European Union Agency for Network and Information Security (ENISA): A good practice guide of using taxonomies in incident prevention and detection — ENISA (2017), https://www.enisa.europa.eu/publications/using-taxonomiesin-incident-prevention-detection

15. Financial Institutes – Information Sharing and Analysis Centre: European Financial Institutes – Information Sharing and Analysis Centre, A Public-Private Partnership — ENISA (2020), https://www.enisa.europa.eu/topics/cross-cooperationfor-csirts/finance/european-fi-isac-a-public-private-partnership
16. FIRST: Traffic Light Protocol (TLP) (2020), https://www.first.org/tlp/
17. Johnson, C., Badger, L., Waltermire, D.: Guide to Cyber Threat Information Sharing. Special Publication - Council for Agricultural Science and Technology (2016). https://doi.org/10.6028/nist.sp.800-150
18. Leach, P., Mealling, M., Salz, R.: RFC 4122 - A Universally Unique IDentifier (UUID) URN Namespace (2005), https://tools.ietf.org/html/rfc4122#section4.1.1
19. Lee, R.M.: 2020 SANS Cyber Threat Intelligence (CTI) Survey (2020), https://www.sans.org/reading-room/whitepapers/analyst/2020-cyber-threatintelligence-cti-survey-39395
20. Leszczyna, R., Wro´bel, M.R.: Threat intelligence platform for the energy sector. Software - Practice and Experience **49**(8), 1225–1254 (aug 2019). https://doi.org/10.1002/spe.2705
21. Levy, I., Maddy, S.: Active Cyber Defence (ACD) - The Second Year - NCSC.GOV.UK (2019), https://www.ncsc.gov.uk/report/active-cyber-defencereport-2019
22. 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. vol. 2017-Janua, pp. 91–98. Institute of Electrical and Electronics Engineers Inc. (dec 2017). https://doi.org/10.1109/EISIC.2017.20
23. Mutemwa, M., Mtsweni, J., Mkhonto, N.: Developing a cyber threat intelligence sharing platform for South African organisations. In: 2017 Conference on Information Communication Technology and Society, ICTAS 2017 - Proceedings. Institute of Electrical and Electronics Engineers Inc. (may 2017). https://doi.org/10.1109/ICTAS.2017.7920657
24. Noor, U., Anwar, Z., Altmann, J., Rashid, Z.: Customer-oriented ranking of cyber threat intelligence service providers. Electronic Commerce Research and Applications **41**, 100976 (may 2020). https://doi.org/10.1016/j.elerap.2020.100976
25. Nweke, L.O., Wolthusen, S.: Legal Issues Related to Cyber Threat Information Sharing Among Private Entities for Critical Infrastructure Protection. In: 2020 12th International Conference on Cyber Conflict (CyCon). pp. 63–78. IEEE (may 2020). https://doi.org/10.23919/CyCon49761.2020.9131721.
26. OASIS-OPEN: STIX 2.1Standard (2020), https://docs.oasisopen.org/cti/stix/v2.1/cs01/stix-v2.1-cs01.pdf
27. Qamar, S., Anwar, Z., Rahman, M.A., Al-Shaer, E., Chu, B.T.: Data-driven analytics for cyber-threat intelligence and information sharing. Computers and Security **67**, 35–58 (jun 2017). https://doi.org/10.1016/j.cose.2017.02.005
28. Rashid, Z., Noor, U., Altmann, J.: Network externalities in cybersecurity information sharing ecosystems. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). vol. 11113 LNCS, pp. 116–125. Springer Verlag (sep 2019). https://doi.org/10.1007/978-3-030-13342-9 10
29. Ring, T.: Threat intelligence: Why people don't share. Computer Fraud and Security **2014**(3), 5–9 (Mar 2014). https://doi.org/10.1016/S1361-3723(14)70469-5
30. Rosemont, H.: Public-Private Security Cooperation From Cyber to Financial Crime. Tech. rep. (2016), www.rusi.org
31. Rowley, L.: The value of threat intelligence. Computer Fraud & Security **2019**(10), 20 (oct 2019). https://doi.org/10.1016/s1361-3723(19)30109-5

32. Sauerwein, C., Sillaber, C., Mussmann, A., Breu, R.: Threat Intelligence Sharing Platforms: An Exploratory Study of Software Vendors and Research Perspectives. Tech. rep. (2017), https://aisel.aisnet.org/wi2017/track08/paper/3/
33. Shin, B., Lowry, P.B.: A review and theoretical explanation of the 'CyberthreatIntelligence (CTI) capability' that needs to be fostered in information security practitioners and how this can be accomplished (may 2020). https://doi.org/10.1016/j.cose.2020.101761
34. Skopik, F., Settanni, G., Fiedler, R.: A problem shared is a problem halved: A survey on the dimensions of collective cyber defense through security information sharing. Computers and Security **60**, 154–176 (jul 2016). https://doi.org/10.1016/j.cose.2016.04.003
35. Skopik, F., Settanni, G., Fiedler, R.: Cyber Threat Intelligence Sharing through National and Sector-Oriented Communities. In: Collaborative Cyber Threat Intelligence: Detecting and Responding to Advanced Cyber Attacks at the National Level, pp. 129–186. CRC Press (jan 2018). https://doi.org/10.4324/9781315397900
36. Stoddart, K.: UK cyber security and critical national infrastructure protection; UK cyber security and critical national infrastructure protection. Tech. rep. (2016). https://doi.org/10.1111/1468-2346.12706, http://www.bbc.co.uk/news/world-uscanada-34641382.
37. Sullivan, C., Burger, E.: "In the public interest": The privacy implications of international business-to-business sharing of cyber-threat intelligence. Computer Law and Security Review **33**(1), 14–29 (feb 2017). https://doi.org/10.1016/j.clsr.2016.11.015
38. Tounsi, W., Rais, H.: A survey on technical threat intelligence in the age of sophisticated cyber attacks (jan 2018). https://doi.org/10.1016/j.cose.2017.09.001
39. UK Government: Cyber Threat Intelligence in Government: A Guide for Decision Makers & Analysts. Tech. rep. (2019), https://hodigital.blog.gov.uk/wpcontent/uploads/sites/161/2020/03/Cyber-Threat-Intelligence-A-Guide-ForDecision-Makers-and-Analysts-v2.0.pdf
40. UK Government: Detecting the Unknown: A Guide to Threat Hunting. Tech. rep. (2019), https://hodigital.blog.gov.uk/wp-content/uploads/sites/161/2020/03/Detecting-the-Unknown-A-Guide-to-ThreatHunting-v2.0.pdf
41. Wagner, T., E, P., K, M., Abdallah, A.: A Novel Trust Taxonomy for Shared Cyber Threat Intelligence (2018), https://www.hindawi.com/journals/scn/2018/9634507/
42. Wagner, T.D., Mahbub, K., Palomar, E., Abdallah, A.E.: Cyber threat intelligence sharing: Survey and research directions. Computers and Security **87**, 101589 (nov 2019). https://doi.org/10.1016/j.cose.2019.101589
43. Zibak, A., Simpson, A.: Can we evaluate the impact of cyber security information sharing? In: 2018 International Conference on Cyber Situational Awareness, Data Analytics and Assessment, CyberSA 2018. Institute of Electrical and Electronics Engineers Inc. (nov 2018). https://doi.org/10.1109/CyberSA.2018.8551462
44. Zibak, A., Simpson, A.: Cyber Threat Information Sharing: Perceived Benefits and Barriers (2019). https://doi.org/10.1145/3339252.3340528.
45. Zibak, A., Simpson, A.: Cyber threat information sharing: Perceived benefits and barriers. In: Proceedings of the 14th international conference on availability, reliability and security. pp. 1–9 (2019)
46. Zrahia, A.: Threat intelligence sharing between cybersecurity vendors: Network, dyadic, and agent views. Journal of Cybersecurity **4**(1) (jan 2018). https://doi.org/10.1093/cybsec/tyy008.