The limitations of national cyber security sensor networks debunked: Why the human factor matters

Florian Skopik, AIT Austrian Institute of Technology, Austria, florian.skopik@ait.ac.at

Abstract: Organizations recently started to exchange security relevant information on cyber incidents to timely mitigate the effects of newly discovered malware and other forms of cyber attacks. Moreover, state actors assume their role as information brokers through national cyber security centers and distribute warnings on new attack vectors and vital recommendations on how to mitigate them. Although many of these initiatives are effective to some degree, they also suffer from considerable limitations. When going beyond pure technical indicators, extensive human involvement is required to manually review, vet, enrich, analyze and distribute security information until relevant information reaches a decision maker. Recent research therefore proposes the automatic collection, analysis and preparation of security data to effectively overcome limiting scalability factors. While this seems to work at an organizational level, the elevation of these approaches to a cross-organizational and even national level is not straight forward. In this paper we investigate where and why the human factor seems irreplaceable. We shed light on the limitations of autonomous cyber security sensor networks at the national level and outline important research areas that need further attention in order to address the remaining issues.

Keywords: cyber security sensors, national cyber security, cyber situational awareness, national cyber security centers

1. Introduction

Recent legal and regulatory advancements, such as the EU NIS directive (European Commission, 2016) and the US CISA (US Congress, 2015) support the development towards a more connected cyber security community, especially a more open culture of exchanging information on security incidents. These initiatives foresee that organizations, especially critical infrastructure providers, report incidents and critical situations to the authorities, essentially cyber security centers or national CERTs/CSIRTs, which take these reports to create common cyber operational pictures (CCOPs) (Pahi et al., 2017). These CCOPs are the foundational basis to establish cyber situational awareness (Franke and Brynielsson, 2014) and aid decision making at the different levels of organizations and nation states. However, the whole process of data collection and approval on an organization's side, as well as the data review, interpretation, aggregation and analysis on the national side is error-prone, involves large amounts of human intelligence, introduces significant lags and therefore does not scale. Thus, recent research has proposed numerous models for cyber security sensor networks to overcome these issues caused by manual reporting (Swart et al., 2016) (Coldebella and White, 2010).

The vision of autonomously reporting sensors is that the authorities can discover within minutes how widely distributed a certain malware is, how vulnerable organizations or industry sectors are in general, and thus, derive the general threat level of a whole society within a nation state at any point in time. That is the theory. Unfortunately, there are numerous limiting factors to this vision. Set aside all the issues regarding data privacy, legal barriers and complex governance aspects, the single factor that seems to render this vision unachievable is still the human in the loop. In order to decrease the dependency on human skills in the whole national security eco-system, it is of paramount importance to understand where and why these human factors are strictly required. The grand vision of many authorities is that a "black box", deployed at the perimeter of each and every organization's network collects traffic data and reports suspicious behavior and sightings of malicious activities to a central entity. What sounds like a rather simple autonomously working system, means however a very complex, human-driven effort towards collaborative cyber security.

This paper deals in with the human factor of cyber security sensor networks. The actual contributions are:

- We outline the **foundational concepts of a cyber security sensor network** to better understand where the human is indispensable in the whole process.
- We define **a process model to design and run a cyber security sensor network**, which was derived together with national authorities in course of a research project.

- We **critically review** this model with respect to **the degree of achievable automation** in each step and discover clear limitations due to required human involvement. We also highlight future research directions to relax the situation.

2. Related Work

Indicators of compromise are a means to validate the exploitation of a vulnerability (Rid and Buchanan, 2015). They are used to look for traces that a system has been hacked, modified or exploited in some malicious way. Therefore, vendors of malware scanning solutions distribute IoCs to their deployments so that customers can automatically verify infections. Eventually, this makes malware scanners the simplest form of sensor nodes which are supplied with new signatures on a regular basis. Signatures to identify malicious domains and IP addresses may also be developed by analyzing DNS traffic (Passive DNS). These types of sensors are de-facto state of the art in more mature organizations and can be connected to security information and event management (SIEM) solutions (Miller et al., 2010) to evaluate their results and get an overview of the current threat situation. However, this knowledge resides mostly within organizations only and is just useful to them since only they know their specific processes and are capable of interpreting the results correctly.

National authorities may receive manual reports from organizations which may be based on automatically collected data. This reporting should tremendously increase the awareness of national cyber security centers and CERTs/CSIRTs as intended by the EU's NIS directive (European Commission, 2016), US CISA (US Congress, 2015). Additionally, information sharing across organizations (Skopik et al., 2016) takes mostly place within industry sectors, which run similar services deployed on similar technologies, and thus, fighting with the same security issues. However, this information sharing processes are mostly initiated on demand and performed manually instead of automatically; e.g., manual exchange of indicators in MISP (Wagner et al., 2016). What is missing, is a means for a near real-time evaluation of the current situation, e.g., in case of spreading malware or serious and widely distributed vulnerabilities. Getting to know, who is affected and how serious the problem is, requires tremendous human effort. So, an automatic evaluation and forwarding would be desirable for the national authorities and considerably relax the situation in the beginning of a new attack wave. Cyber security sensor networks, as proposed by several national cyber security strategies (Luiijf et al., 2013) could be of great help, and some real-world examples are already deployed, e.g., in France, Finland (Rantapelkonen et al., 2013), and Switzerland (Cavelty, 2014) for exactly that purpose.

Eventually, collecting information about cyber attacks, incidents and threats in a timely manner is essential to gather cyber situational awareness (Franke and Brynielsson, 2014), and a prerequisite of justified decision making (Stotz and Sudit, 2007).

3. The Foundations of Cyber Security Sensor Networks

3.1 Motivation for Sensor Networks

Especially, at the beginning of a nation-wide cyber security incident, such as spreading malware in critical infrastructures (Chen and Bridges, 2017) or the recent discovery of a wide-spread vulnerability (Durumeric et al., 2014), information of national CERTs or National Cyber Security Centers about the situation of privately-owned organizations are scarce at best. Most urgent questions in such situations for which timely answers may be essential for the survival of a nation state's industry, include:

- Who is affected?
- Who needs help?

Unfortunately, affected organizations tend to report quite late or not at all (Choo, 2011). However, the obligations of the EU's NIS directive, as well as the US CISA might ease the situation here, depending on how strict the thresholds for reporting obligations are set. Moreover, there are many exemptions for numerous kinds of organizations which do not need to report incidents with significant impact. Even if an organization reports cyber security issues to the authorities, this information is just useful if the authorities get to know the potential impact, and an estimate whether the reporting organization can deal with the issues alone or needs external help. For instance, in the WannaCry case early 2017 (Chen and Bridges, 2017), many organizations were down but busy with restoring their data from recent backups, while others were still debating whether this small payment to an anonymous bitcoin account would unlock their data. The latter didn't have disaster recovery

plans at this time. So, generally speaking, it is unclear how much impact an attack wave has on a larger scale, i.e., across different organizations, and how widespread a certain malware (or exploitable vulnerability) is.

Up to now, in such cases, national CERTs and CSCs have simply asked organizations to periodically send reports stating their operational status to estalbish a clear picture. This situation is unsatisfactory and demands a more automated approach.

3.2 Illustrative Application of Cyber Security Sensors

The demand for (semi-)automatically collecting cyber security relevant information (e.g., through sensors) is not exactly new and has been enforced by secret services since years (Coldebella and White, 2010). While with a high degree of automation for data collection, aggregation, analysis and interpretation a high scalability factor seems achievable, the setup and application of cyber security sensors across organizations is highly non-trivial. A comparatively primitive, but yet non-trivial, application case of cyber security sensors is the registration of IoC sightings. An IoC (Rid and Buchanan, 2015) is a unique data particle that verifies the presence of a malware or the exploitation of a service. This is for instance the presence of a certain file (name, hashsum etc.), a specific process, log line in DNS records, specific network events etc. Numerous initiatives to detect the presence of these IoCs automatically on a national level have been undertaken. They use distributed sensors, deployed on the Internet as well as within organizations, to constitute a sensor network. For the WannaCry ransomware campaign, the basic questions that required timely answers through consulting such a sensor network at an early stage were (answers in parentheses):

- **What** needs to be detected? (certain patterns in SMB traffic; specific strings in memory)
- **Where** can be detected? (between network segments (firewalls); at all potentially vulnerable hosts)
- **How** can be **detected** and captured? (analysis of netflows, deep packet inspection, host scans)
- **How** (quickly) does the result need to be **forwarded**? (encrypted and anonymized on detection)
- **Why** should it be detected and captured? (estimation of how widespread WannaCry is)

3.3 Design Considerations for Cyber Security Sensor Networks

With respect to the questions raised above, multiple vital aspects need to be considered during the design phase of a sensor network. **WHAT** needs to be detected, might not be straight forward to answer; either simple IoCs are known and can be distributed to sensors (e.g., via detection rules that are pushed to sensor nodes), or complex contextual data needs to be captured for further analysis (e.g., type and degree of dependency of critical business services on underlying technologies). The latter is specifically challenging due to the diversity and large volumes of data that needs handling. **WHERE** data can be detected does not only depend on the type of data, but also on the type and structure of the monitored infrastructure. Newer concepts, such as “Bring-Your-Own-Device” and end-to-end encryption has caused the network perimeter to diminish. Additionally, software-defined networks (SDN) and virtualization techniques rather require host-based detection, which is however much trickier to implement and to enforce. This last point is closely connected to the question **HOW** data is being **detected** and captured. Especially host-based detection raises concerns with respect to scalability, performance, data privacy, and accountability in case something goes wrong – and of course: security! Once data is being captured – on the host or on the network – **HOW** data is being securely **forwarded** is similarly important and includes aspects of encryption, transport priorities, aggregation and buffering mechanisms. Eventually, the last question is **WHY** data should be detected and captured. If there is no clear contribution of sensor readings to a higher level analysis for both the nation state and the privately-owned companies that report, then the acceptance of this technology within organizations will be limited.

4. From Simple Data Acquisition to Cyber Situational Awareness

4.1 On common cyber operational pictures (CCOPs)

The transformation of vital information into a common (cyber) operational picture (CCOP) is key to justified decision making. However, an important lesson to learn is that despite the “common” in the name, there is no “one fits all” solution in the cyber domain. A CCOP is always layer- and application-specific. Stakeholders on different layers need to be supplied with information that is specific to their tasks, functions and decisions to make. For instance, a server administrator needs different information and makes different decisions than a chief information security officer above him. The information required by business management on the strategic layer is different from the information required by decision makers in political functions. Moreover, independent

from the actual layers and roles, different applications of CCOPs require different data. For example, incident response focuses on other aspects than proactive risk assessments.

4.2 The dilemma of layer-adjusted CCOPs

Based on the explanation above, the actual dilemma is to keep a balance, on the one side, to supply stakeholders with information shaped to their needs, and on the other side not influence their decision making by filtering information which is below a preset relevancy threshold. The same information may trigger different actions depending on how it is visualized in dashboards. Colors, font sizes, and orders of table entries have great influence. The application of filters is even more critical.

4.3 A rigorous model to establish cyber security sensor networks

To tackle the dilemma of creating layer-specific CCOPs, we came up with a rigorous design process that help us to design application-specific cyber security sensor networks, which go far beyond simple IoC validation and thus incorporates both, carefully selected core data and contextual data. In this model, we run through five consecutive steps, each dealing with a specific question, depicted in Figure 1.

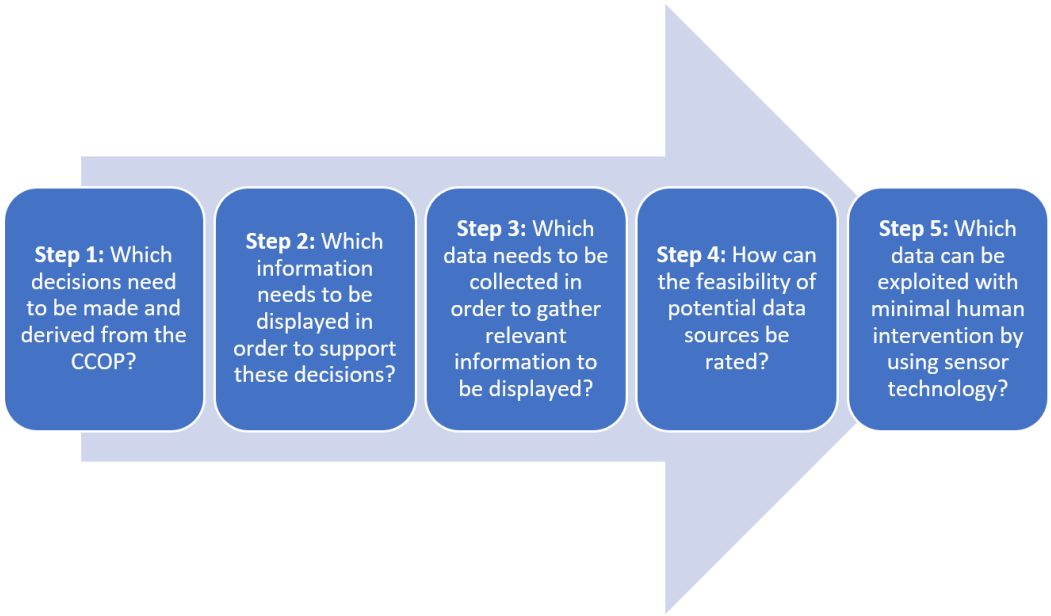

Figure 1: CCOP Design Process and Guiding Questions

Finding answers to the questions in these five steps helps us to set the corner stones of an application-specific sensor network. However, many of these answers are quite tricky to answer and conflict with the need for a high degree of automation. Eventually, deploying, running and maintaining a cyber security sensor network following this model requires considerable human involvement, which is the root cause of serious limitations for their fast and cost-efficient adoption as outlined in the next section.

5. Major Limitations and Critical Factors

In this section we are going to dissect the introduced top-down model of Section 4 and investigate the multiple factors that need to be considered during the design of a cyber security sensor network. Automation in the various steps is the key to a fast, scalable and efficient sensor network. Therefore, we specifically focus on steps in the model that strictly require human interaction.

5.1 Step 1: Which decisions need to be made?

Eventually, sensor data is retrieved, analyzed, interpreted and transformed into common cyber operational pictures (CCOPs) to aid some sort of decision making. The decisions in national CERTs and security centers might be manifold at an operational and strategic level. We distinguish decisions depending on their time horizon

(Table 1 based on (Pahi and Skopik, 2016)). All these decisions require the human in the loop due to usually complex cyber security situations at the national level – either to make or to approve the decision.

Table 1: Decisions at the national level supported by sensor networks

Short-term (hours)	Send out warnings to potentially affected organizations Provide immediate help (incident response, disaster recovery) Provide recommendations Enforce information sharing within a sector
Medium-term (days)	Create task force to overcome a crisis Coordinate actions together with vendors or service providers Assist in disaster recovery across organizations Update best practice guidelines and distribute updates Start prosecution
Long-term (months)	Financial support for education of experts Provide trainings Coordinate periodic external audits Adapt laws and regulations (e.g., thresholds in obligations to report)

5.2 Step 2: Which information need to be displayed?

The information to be visualized in CCOPs is always application-specific and stakeholder-specific. In other words, decision makers need to be supplied with information for a particular situation. The application cases, the decision in context of these application cases and the type of relevant information need to be defined in advance, otherwise decision makers run the risk of getting overloaded with tons of unspecific data. That is a major difference to the application of COPs in the physical world.

For example, let us consider the goal of a CCOP is to support justified risk assessment of cyber crime using DDoS (Radunovic, 2013). First information of interest is about the current threat level. For instance, how much does renting a DDoS attack with, say, 5 GBit/s for a duration of 1 week, in the Darknet cost? And regarding the motivation: Is it in someone's interest to harm an organization, industrial sector or nation state? The latter is often connected to evaluating the current political situation and attitude of activists and hacktivists. Besides the current threat situation, historic data may be of great help to estimate trends. Examples are questions on how many organizations where victims of such an attack in the last year; or how well organizations of a certain sector are prepared, e.g., with respect to their business continuity planning. Also, which mitigation controls are already in place, such as backup sites, or traffic scrubbing contracts, is essential to draw the right picture. And in case despite all preparation an attack happens, it is important to identify the properties of attacks early, e.g., coming from a known or unknown bot net, the modus operandi of the attacker, the type and criticality of affected systems, and the ransom demanded. Looking at this single case of CyberCrime via DDoS (which is not even exhaustively presented here), it becomes obvious that collecting network data alone will not help us much to understand cyber attacks, but carefully collecting contextual data to aid the interpretation of collected network data and eventually to draw the right conclusions is the key to an effective cyber security sensor network. Unfortunately, coming up with these application cases requires a sharp human mind and cannot be done automatically.

5.3 Step 3: Which data need to be collected?

Now, after we know what information needs to be visualized for applicable CCOPs, we need to decide what data is required to derive this information. Here, we distinguish between two essential classes: (i) core data and (ii) contextual data (Pahi et al., 2017). Many works in the threat intelligence domain focus only on core data, e.g. STIX (Barnum, 2012), which is comparatively easy to model, gather and process automatically. The core data includes indicators (IoCs), observables (aka sightings of these IoCs in real infrastructures), information about previous incidents and targets, applied tactics, techniques and procedures (TTPs), and so on. However, the interpretation of these core data requires contextual data for decision making. For instance, just because some IoC was spotted, it does not mean its occurrence is critical. Without knowledge of the layout of the target network, criticality of hosted services and supported processes, as well as capabilities of the people handling a problem, a reliable decision on the next steps is hardly possible. Thus, to interpret core data, we require

contextual data, including information on organizations, their services to the public and the dependency of their business processes to the underlying technology. Furthermore, essential industry know-how and lessons learned from previous incidents increase the quality of decision making. A rule of thumb is that core data can be gathered as is and thus exploited automatically, while contextual data is usually too diverse, unstructured, noisy and often not explicitly documented in a machine-readable format so that it requires human processing. More insights on security-related data and its classification is discussed in Pahi et al. (2017).

5.4 Step 4: How can potential data sources be rated?

A key question always is where to gather the data from. Often there are multiple sources, under control of different parties, which however emit the same or similar data. In that cases we need a sound approach to decide which source should be consulted. There are numerous metrics to consider, based on the specific application case of the CCOP, including timeliness, relevance, control, completeness, trust, availability, sensitivity, and accuracy. For instance, one source might be manually vetted, and cross-checked and thus delivers high quality trust-worthy results, while another one misses these features but delivers more up-to-date information. Other questions might be centered on who is in control of a source, can it be trusted or easily be manipulated. Which sources to pick is especially tricky in case of conflicting data from different sources. Here, a human in the loop who manually makes a justified decision on which source to pick is required. Since sources are mostly volatile in nature, verifying the quality of sources is a re-occurring task.

5.5 Step 5: Which data can be exploited using sensors?

The last step is the one with the highest potential for automation. Once it was determined what needs to be collected and for what reasons, the pure data acquisition takes place. The low-hanging fruits are all technical information about an organization that can be gathered from the outside, such as (i) the IP addresses an organization uses, which is highly relevant for botnet detection; (ii) the core services an organization offers to identify dependencies to other organizations; and (iii) public keys and certificates, such as S/MIME, SSL/TLS used for external communication. It becomes trickier when we look into an organization and collect information from within – either captured from the network or directly on hosts. These data include, but is not limited to, (i) assets and configurations, e.g., collected via SNMP to match against known vulnerabilities and weaknesses); (ii) data about system usage and behavior, e.g., usual bandwidth consumption profiles, degree of statistical anomalies in data flows; (iii) attacked services, exploited vulnerabilities and used attack vectors (if known at all); (iv) the results of periodically performed malware scans and internal audits (which however require contextual data for interpretation); and (iv) simple IoC sightings within the network, e.g. suspicious files, IP/mail addresses, scheduled reboots etc. incl. sandbox analysis results.

However, particularly interesting is the information which can be inferred from this simple data. Just to provide some examples: One could derive an organization's "patch mentality" by just measuring the time span between the release of a new patch and its deployment at the organization's assets. Another example is mining of operational best practices, e.g., how frequently passwords are changed, or what types of roles to restrict system access is applied. One could even determine if there were successful malware attacks, although a patch was available to fix the exploited vulnerability. Anyway, the key question is which – if any – of these details is useful to aid the decision making on a national level as outlined in Step 1 of this process.

6. Conclusion and Future Work

There is an infinite amount of data that can be gathered from hosts and networks, but the essential question is which of this data is relevant to support decision making processes at the national level. An accurate overview of the status of critical infrastructures helps to determine and coordinate mitigation actions across organizations, while the actual mitigation actions are mostly carried out within the organizations. For instance, the distribution of detailed recommendations to handle a recently discovered vulnerability or the distribution of an unofficial emergency patch are good examples here. However, to determine which information is of importance to decision makers at diverse layers and in various roles, complex human knowledge and skills are required for designing, deploying, operating and maintaining sensor networks. It is obvious that just collecting huge amounts of data on a broad scale does not aid decision making and is thus meaningless without a rigorous justification for doing so. The application of sensors also requires a concept to gather the context in which sensor values are collected. If a sensor reports suspicious traffic or the sighting of an IoC, the analysts at the national layer also need background information of the affected organization, its infrastructure, assets and services.

As more human skills are required, the less scalable the whole system becomes. It is therefore of paramount importance for future research to determine ways to relieve the human from reoccurring tasks and allow him to focus scarce and valuable resources on other tasks. Future research therefore needs to deal with:

- **Data retrieval:** Automatizing data retrieval from sources that emit data in a wide range of different formats is a key to increase the scalability. Besides getting along with numerous interface styles and protocols, the automatic generation of parsers to transform data in any format (STIX, OpenIoC, logs etc.) into one consistent format for later reasoning is particularly of interest here.
- **Natural language processing:** Besides the challenges of harmonizing different syntaxes, the much bigger challenge is to automatically process and understand free text messages, e.g., incident reports, threat assessment reports, whitepapers, blog entries, Internet forums etc. Many valuable sources are available only in an unstructured text format, which is hard to digest automatically. However, accounting for high-level TTPs described in these sources is much more effective than an analysis based on technical indicators only. Either a machine-readable representation of these high-level TTPs or some smart algorithms to process this natural language texts directly are desirable.
- **Autonomous security database management and lookup:** Maintaining, querying and even cross-connecting (public) databases for targeted lookups when investigating incidents or assessing threat levels, e.g., when a suspicious IP address, url name or file is found in a company network, is a key requirement for automated threat assessment. Particular examples for such lookup databases are VirusTotal, ThreatMiner, ThreatCrowd and DNSDumpster.
- **Information fusion and semantic reasoning:** Cross-connecting aforementioned sources, such as malware domains with file hashes, or CVE entries with information on exploits, is key to avoid tedious manual search activities and free the analyst's time for actual analysis instead of data collection activities. This will however require at least some semantic understanding of the information delivered by the sources.
- **Decision making support systems:** Once the analysis is largely performed automatically, a human decision maker would only need to review the results and make a decision appropriate for a given situation. Re-occurring decisions, e.g., the triage in incident response, may, however, be automatized by using self-learning systems which monitor human decisions and comprehend which factors lead to certain decisions.

Eventually, nation states need to ensure transparency regarding the application of cyber security sensor networks. If organizations do not know what the authorities are looking for, if there is no clear benefit for the monitored organizations and no reasons for collecting specific types of data, the acceptance of this technology will be extremely limited and thus its effectiveness suffer. In the best case, organizations and the nation state build a public-private-partnership were both sides benefit equally.

Acknowledgement:

This work was partly funded by the Austrian FFG research program KIRAS in course of the projects ACCSA (860649) and CRISCROSS (860678).

References

Barnum, S. (2012), 'Standardizing cyber threat intelligence information with the structured threat information expression (stix)', MITRE Corporation 11, 1–22.

Cavelty, M. D. (2014), Reporting and analysis center for information assurance (melani) (phase 2: 2004–2010), in 'Cybersecurity in Switzerland', Springer, pp. 39–55.

Chen, Q. and Bridges, R. A. (2017), 'Automated behavioral analysis of malware a case study of wannacry ransomware', arXiv preprint arXiv:1709.08753.

Choo, K.-K. R. (2011), 'The cyber threat landscape: Challenges and future research directions', Computers & Security 30(8), pp. 719–731.

Coldebella, G. P. and White, B. M. (2010), 'Foundational questions regarding the federal role in cybersecurity', Journal of National Security Law & Policy 4, 233.

Durumeric, Z., Kasten, J., Adrian, D., Halderman, J. A., Bailey, M., Li, F., Weaver, N., Amann, J., Beekman, J., Payer, M. et al. (2014), The matter of heartbleed, in 'Proceedings of the 2014 Conference on Internet Measurement Conference', ACM, pp. 475–488.

European Commission (2016), 'The directive on security of network and information systems (nis directive)', viewed 17 November 2018, https://ec.europa.eu/digital-single-market/en/network-andinformation-security-nis-directive.

Franke, U. and Brynielsson, J. (2014), 'Cyber situational awareness–a systematic review of the literature', Computers & Security 46, pp. 18–31.

Luiijf, E., Besseling, K. and De Graaf, P. (2013), 'Nineteen national cyber security strategies', International Journal of Critical Infrastructures 9(1-2), pp. 3–31.

Miller, D. R., Harris, S., Harper, A., VanDyke, S. and Blask, C. (2010), Security Information and Event Management (SIEM) Implementation (Network Pro Library), McGraw Hill.

Pahi, T., Leitner, M. and Skopik, F. (2017), 'Preparation, modelling, and visualisation of cyber common operating pictures for national cyber security centres', Journal of Information Warfare 16(4), pp. 26–40.

Pahi, T. and Skopik, F. (2016), 'A public-private-partnership model for national cyber situational awareness', Intl. Journal on Cyber Situational Awareness 1.

Radunovic, V. J. (2013), Ddos-available weapon of mass disruption, in 'Telecommunications Forum (TELFOR), 2013 21st', IEEE, pp. 5–8.

Rantapelkonen, J., Salminen, M. et al. (2013), 'The fog of cyber defence', Julkaisusarja.

Rid, T. and Buchanan, B. (2015), 'Attributing cyber attacks', Journal of Strategic Studies 38(1-2), pp. 4–37.

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

Stotz, A. and Sudit, M. (2007), Information fusion engine for real-time decision-making (inferd): A perceptual system for cyber attack tracking, in 'Information Fusion, 2007 10th International Conference on', IEEE, pp. 1–8.

Swart, I., Irwin, B. and Grobler, M. M. (2016), 'Adaptation of the jdl model for multi-sensor national cyber security data fusion', International Journal of Cyber Warfare and Terrorism (IJCWT) 6(3), pp. 17–30.

US Congress (2015), 'S.754 - to improve cybersecurity in the united states through enhanced sharing of information about cybersecurity threats, and for other purposes', viewed 17 November 2018, https://www.congress.gov/bill/114thcongress/senate-bill/754.

Wagner, C., Dulaunoy, A., Wagener, G. and Iklody, A. (2016), Misp: The design and implementation of a collaborative threat intelligence sharing platform, in 'Proceedings of the 2016 ACM Workshop on Information Sharing and Collaborative Security', ACM, pp. 49–56.