Journal of Information Warfare

Published by
Peregrine Technical Solutions, LLC
Yorktown, Virginia, USA

Print Version
ISSN 1445-3312

Online Version
ISSN 1445-3347

National Cyber Security Sensor Networks and the Human in the Loop

F Skopik

Center for Digital Safety & Security
AIT Austrian Institute of Technology
Vienna, Austria

E-mail: florian.skopik@ait.ac.at

Abstract: *Organisations recently started to exchange security relevant information on cyber incidents to timely mitigate the effects of newly discovered malware and other forms of cyberattacks. Moreover, state actors take over their role as information brokers through national cyber security centres 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, analyse, 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 organisational level, the elevation of these approaches to a cross-organisational and even national level is not straightforward. This paper investigates where and why the human factor seems irreplaceable and sheds light on the limitations of autonomous cyber security sensor networks at the national level.*

Keywords: *Cyber Security Sensors, National Cyber Security, Cyber Situational Awareness, National Cyber Security Centres*

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 of a more connected cyber security community, especially a more open culture of exchanging information on security incidents. These initiatives foresee that organisations, especially critical infrastructure providers, report incidents and critical situations to the authorities, essentially cyber security centres or national Computer Emergency Response Teams (CERTs)/Computer Security Incident Response Teams (CSIRTs), which take these reports to create Common Cyber Operational Pictures (CCOPs) (Pahi, Leitner & Skopik 2017). These CCOPs are the foundational basis for establishing cyber situational awareness (Franke & Brynielsson 2014) and aiding decision making at the different levels of organisations and nation states. However, the whole process of data collection and approval on an organisation’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, Irwin & Grobler 2016) (Coldebella & 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 organisations or industry sectors are in general, and can, thus, determine 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. If all the issues regarding data privacy, legal barriers, and complex governance aspects are set aside, the single factor that seems to render this vision unachievable is still the human in the loop (Skopik 2019). 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 organisation's network, collects traffic data and reports suspicious behaviour 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 addresses the human factor of cyber security sensor networks. The actual contributions are

- An outline of the **foundational concepts of a cyber security sensor network** to better understand where the human is indispensable in the whole process.
- A definition of **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.
- A **critical review** of this model with respect to **the degree of achievable automation** in each step and summary of clear limitations due to required human involvement.

Related Work

Indicators of compromise are data found in system log entries or files that identify potentially malicious activity on a system or network. They are a means to validate the exploitation of a vulnerability (Rid & 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 Indicators of Compromise (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 analysing DNS traffic (Passive DNS). These types of sensors are de-facto state of the art in more mature organisations and can be connected to Security Information and Event Management (SIEM) solutions (Miller *et al.* 2010) to evaluate their results and to get an overview of the current threat situation. However, this knowledge resides mostly within organisations only and is just useful to them since only they know their specific processes and can interpret the results correctly.

National authorities may receive manual reports from organisations which may be based on automatically collected data. This reporting should tremendously increase the awareness of national cyber security centres and CERTs/CSIRTs as intended by the EU's NIS directive (European Commission 2016) and US CISA (US Congress 2015). Additionally, information sharing across organisations (Skopik, Settanni & Fiedler 2016) takes place mostly within industry sectors, which run similar services deployed on similar technologies, and, thus, are fighting with the same security issues. However, these information sharing processes are mostly initiated on demand and are performed manually instead of automatically, for example, manual exchange of indicators in Malware Information Sharing Platform (MISP) (Wagner *et al.* 2016). What is missing is a means for a near real-time evaluation of the current situation, such as in the case of spreading malware or serious and widely distributed vulnerabilities. Getting to know whom is affected and how serious the problem is requires tremendous human effort. So, an automatic evaluation and forwarding system would be desirable for the national authorities and would 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, Besseling & De Graaf 2013) could be of great help, and some real-world examples are already deployed, such as in France, Finland (Rantapelkonen & Salminen 2013), and Switzerland (Cavelty 2014) for exactly that purpose.

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

The Foundations of Cyber Security Sensor Networks

Motivation for sensor networks

Especially at the beginning of a nation-wide cyber security incident, such as spreading malware in critical infrastructures (Chen & 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 organisations 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 organisations tend to report quite late or not at all (Choo 2011). However, the obligations of the EU's NIS directive, as well as of 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 organisations which do not need to report incidents with significant impact. Even if an organisation reports cyber security issues to the authorities, this information is just useful if the authorities get to know the potential impact and can estimate whether the reporting organisation can deal with the issues alone or needs external help. For instance, in the WannaCry case in early 2017 (Chen & Bridges 2017), many organisations 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. Therefore, it is unclear how much impact an attack wave has on a larger scale, such as across different organisations, and how widespread a certain malware (or exploitable vulnerability) is.

Up to now, in such cases, national CERTs and Cyber Security Centers (CSCs) have simply asked organisations to periodically send reports stating their operational status to establish a clear picture. This situation is unsatisfactory and demands a more automated approach.

Illustrative application of cyber security sensors

The demand for (semi-) automatically collecting cyber security relevant information (such as through sensors) is not exactly new and has been enforced by secret services for years (Coldebella & 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 organisations is highly non-trivial. A comparatively primitive, yet hardly achievable, application case of cyber security sensors is the registration of Indicator of Compromise (IoC) sightings. An IoC (Rid & 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 or hashsum, for example), a specific process, a log line in DNS records, or specific network events. 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 organisations, 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 it be detected? (between network segments [firewalls]; at all potentially vulnerable hosts)
- How can it be detected and captured? (analysis of netflows, deep packet inspection, host scans)
- How (quickly) does the result need to be forwarded? (encrypted and anonymised on detection)
- Why should it be detected and captured? (estimation of how widespread WannaCry is)

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. The question of **WHAT** needs to be detected might not be so straightforward to answer; either simple IoCs are known and can be distributed to sensors (such as via detection rules that are pushed to sensor nodes), or complex contextual data needs to be captured for further analysis (for example, 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 need 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 have caused the network perimeter to diminish. Additionally, software-defined networks (SDN) and virtualisation techniques rather require host-based detection, which is much trickier to implement and to enforce. This last point is closely connected to the question of **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 organisations will be limited.

Creating 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.

Creating layer-specific CCOPs requires a rigorous design process for the underlying cyber security sensor networks. The model introduced in **Figure 1**, below, runs through five consecutive steps, each dealing with a specific question. Finding answers to the questions in these five steps allows organisations to set the cornerstones of an application-specific sensor network. However, many of these questions 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.

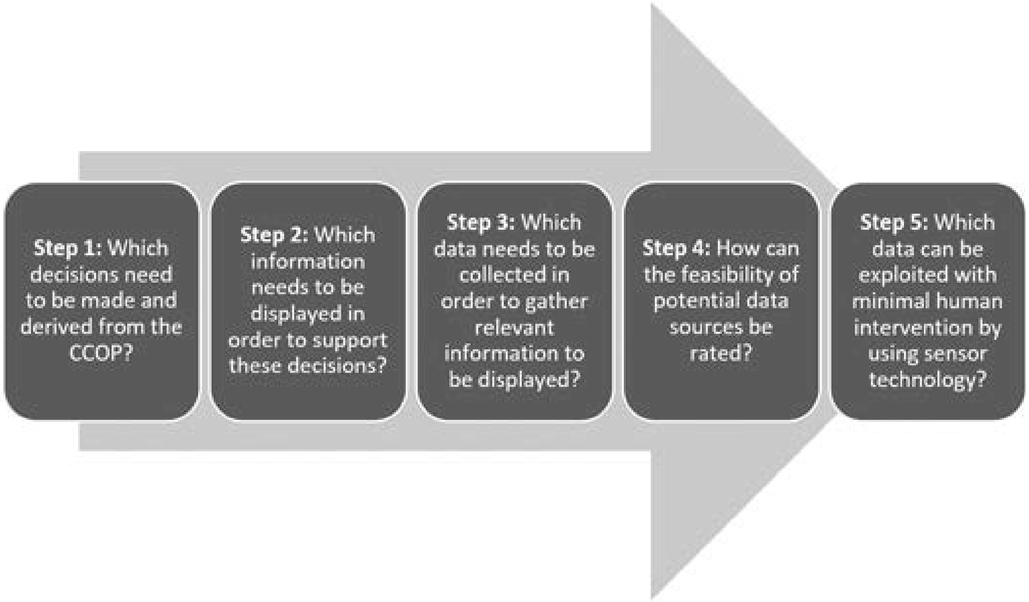
Step 1: Which decisions need to be made and derived from the CCOP?
Step 2: Which information needs to be displayed in order to support these decisions?
Step 3: Which data needs to be collected in order to gather relevant information to be displayed?
Step 4: How can the feasibility of potential data sources be rated?
Step 5: Which data can be exploited with minimal human intervention by using sensor technology?

Figure 1: CCOP design process and guiding questions

Major Limitations and Critical Factors

This section takes a closer look into the introduced top-down model of the previous section and investigates 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, the following subsections specifically focus on steps in the model that require human attention.

Step 1: Which decisions need to be made?

Eventually, sensor data is retrieved, analysed, interpreted, and transformed into Common Cyber Operational Pictures (CCOPs) to aid in some sort of decision making. The decisions in national CERTs and security centres might be manifold at an operational and strategic level. Decisions depend on their time horizon (**Table 1**, below, is based on Pahi & Skopik 2016). Due to usually complex cyber security situations at the national level, all these decisions require the human in the loop—either to make or to approve complex decisions.

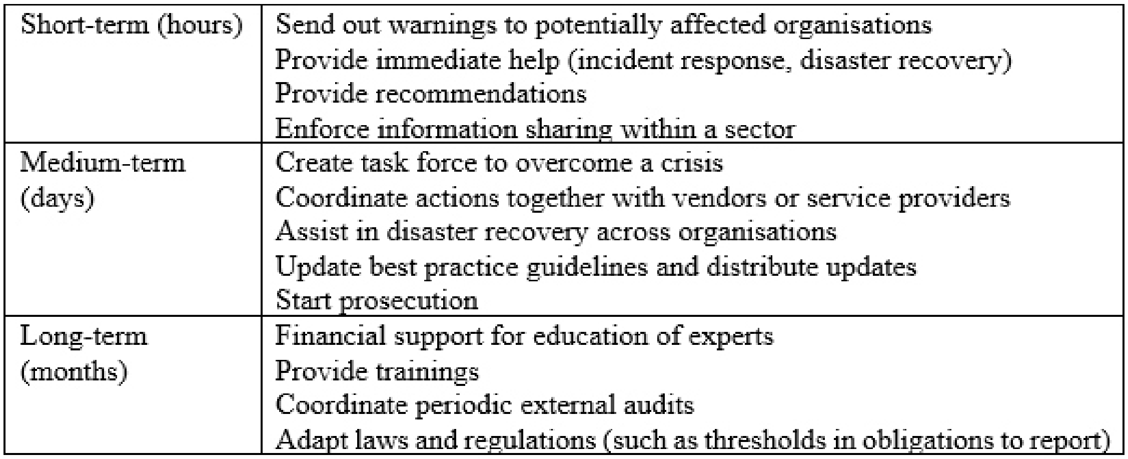

Short-term (hours)	Send out warnings to potentially affected organisations Provide immediate help (incident response, disaster recovery) Provide recommendations Enforce information sharing within a sector
Medium-term (days)	Create task force to overcome a crisis Coordinate actions together with vendors or service providers Assist in disaster recovery across organisations 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 (such as thresholds in obligations to report)

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

Step 2: Which information needs to be displayed?

The information to be visualised in CCOPs is always application-specific and stakeholder-specific. In other words, decision makers need to be supplied with information for a specific 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 unspecific data. This is a major difference to the application of COPs in the physical world.

For example, consider the goal of a CCOP is to support justified risk assessment of cyber crime using DDoS (Radunovic 2013). The first information of interest is about the current threat level. For instance, how much does renting a DDoS attack with, for instance, 5 GBit/s for a duration of one week, in the Darknet cost? And regarding the motivation: is it in someone's interest to harm an organisation, 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 organisations were victims of such an attack in the last year; or how well organisations of a certain sector are prepared, such as, with respect to their business continuity planning. Also, considering which mitigation controls are already in place, such as backup sites or traffic scrubbing contracts, is essential to draw the right picture. And despite all preparation, if an attack happens, it is important to identify the properties of attacks early, such as whether they were coming from a known or unknown botnet, 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 yield a better understanding of cyberattacks, but carefully collecting contextual data to aid the interpretation of collected network data and eventually to draw the right conclusions is 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.

Step 3: Which data need to be collected?

After knowing what information needs to be visualised for applicable CCOPs, experts decide what data is required to derive this information. Here, two essential classes need to be distinguished: (1) core data and (2) contextual data (Pahi, Leitner & Skopik 2017). Many works in the threat intelligence domain focus only on core data, such as STIX (Barnum 2012), which is comparatively easy to model, gather, and process automatically. The core data includes indicators (IoCs); observables (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, the fact that an IoC was spotted 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, contextual data is required, including information on organisations, 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 are discussed in Pahi, Leitner & Skopik (2017).

Step 4: How can potential data sources be rated?

A key question 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 these cases, a sound approach is needed 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 deliver high quality, trustworthy results, while another one misses these features but delivers more up-to-date information. Other questions might be centred on who is in control of a source, whether it can be trusted or easily manipulated. Which sources to pick is especially tricky in cases with 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.

Step 5: Which data can be exploited using sensors?

The last step is the one with the highest potential for automation. Once it is 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 organisation that can be gathered from the outside, such as (1) the IP addresses an organisation uses, which is highly relevant for botnet detection; (2) the core services an organisation offers to identify dependencies to other organisations; and (3) public keys and certificates, such as S/MIME, SSL/TLS used for external communication. It becomes trickier when looking into an organisation and collecting information from within—either captured from the network or directly on hosts. These data include, but are not limited to, (1) assets and configurations, such as those collected via SNMP to match against known vulnerabilities and weaknesses; (2) data about system usage and behaviour, such as usual bandwidth consumption profiles, degree of statistical anomalies in data flows; (3) attacked services, exploited vulnerabilities, and used attack vectors (if known at all); (4) the results of periodically performed malware scans and internal audits (which however require contextual data for interpretation); and (5) simple IoC sightings within the network, such as suspicious files, IP/mail addresses, scheduled reboots, including sandbox analysis results.

Particularly interesting is the information which can be inferred from this simple data. For instance, one could derive an organisation's 'patch mentality' by just measuring the time span between the release of a new patch and its deployment at the organisation's assets. Another example is mining of operational best practices—for example, how frequently passwords are changed or what types of roles to restrict system access are 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.

Roles, Responsibilities, and Interactions

The envisioned cyber security sensor network comprises numerous types of stakeholders with individual roles and pre-modelled interaction patterns. Even for a rather conservative IoC-based sensor network, the human in the loop is inevitable for numerous vital operational tasks, as outlined in this section.

Overview of the Stakeholders of the Sensor Network

Figure 2, below, shows the stakeholder structure to support the information sharing process between organisations, sector-CERTs/CSIRTS, and national authorities. The process starts at the top, where a national Cyber Security Centre (CSC) maintains cyber situational awareness for high-level decision making (Franke & Brynielsson 2014). To fulfil this task, the CSC gains access to various non-public threat information sources, such as confidential repositories and lists from secret services and law enforcement. Carefully selected subsets of these threat information sources are then forwarded, preferably as Indicators of Compromise (IoCs) (Obrst, Chase & Markeloff 2012), to the sector-CERTs, where they feed an MISP server (Wagner *et al.* 2016) with this (partly confidential) information. The sector-CERTs may additionally subscribe to various other useful public sources (such as vendor lists). This preselection process of relevant threat information sources reduces the work for the individual organisations and ensures the delivery of high-quality information shaped to the needs of an industrial sector (for example, fitting to the commonly used assets in a certain domain). The preselected events are then consumed by Sensor Network Nodes (SNNs) distributed throughout the organisations. An SNN searches for the received IoCs in its associated network and reports found sightings back to the sector-CERT's MISP server. The CERT aggregates received feedback on a continuous basis and creates Common Cyber Operational Pictures (CCOPs) (Pahi, Leitner & Skopik 2017)—some of them for specific industry sectors. Expert circles and national decision makers use these CCOPs as the basis to assess the overall cyber security situation within an industry sector. This is an essential prerequisite to take timely counteractions in case of large-scale coordinated cyberattacks.

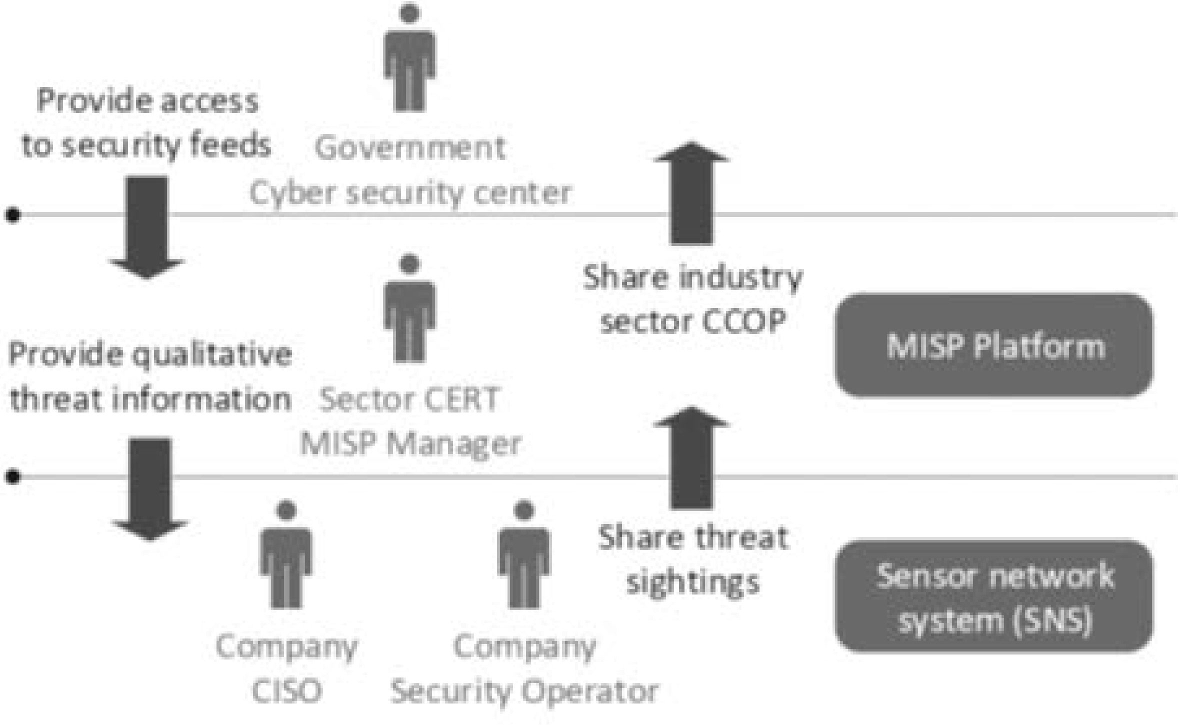
Provide access to security feeds
Government Cyber security center
Share industry sector CCOP
MISP Platform
Provide qualitative threat information
Sector CERT MISP Manager
Share threat sightings
Sensor network system (SNS)
Company CISO
Company Security Operator

Figure 2: Human roles and their interactions in sensor network operations

The governmental Cyber Security Center

A Cyber Security Center (CSC) is a governmental institution which has the responsibility to periodically assess the overall cyber security situation of critical industry sectors (European Commission 2016). For that purpose, the CSC collects information about the security status of critical infrastructures to gain knowledge whether an infrastructure is the victim of any serious attack.

The distribution of IoCs to organisations and their validation (for example, the discovery of IoCs in distributed networks) is an essential tool to assess the current situation and anticipate further trends. The CSC, therefore, asks specific sector-CERTs for support to send them CCOPs for their respective sectors and, eventually, generates an overall picture across all sectors.

In short, the tasks of the CSC are as follows:

- Provide access to non-public feeds: the CSC exchanges threat information with vendors and institutions from other countries. Because of this cooperation and because the CSC is close to the government, it gets access to non-public threat information. The CSC revises and refines gathered information and provides it in an applicable data format to the sector-CERTs;
- Receive CCOPs for each industry sector: each sector-CERT regularly creates a CCOP for its respective industry sector, which is passed on, combined with additional threat information to the CSC. This information is the basis for national decision making;
- Derive ongoing attacks against specific sectors: after one sector-CERT passes on a new CCOP or threat information to the CSC, the CSC assesses the collected information and compares it with gathered information from other industry sectors. The CSC can then determine whether there is an attack especially targeted against an industry sector (or individual nation state).
- Advise CERTs about attacks on other industry sectors: based on the evaluated threat level and based on the received CCOPs, the CSC informs sector-CERTs about emerging attacks in other domains. As a result, the CSC is the connecting link between multiple sector-CERTs and other involved organisations;
- Create a CCOP for a national high-level overview: with the collected information from all previous steps, the CSC creates an aggregated national CCOP for national decision making;
- Evaluate threat level based on the high-level CCOPs: in a final step the CSC rates the overall threat level and identifies vulnerable industry sectors, such as their above-average attack surface or commonly outdated hardware and software assets.

Sector-CERTs

A sector-CERT is an independent organisation of information technology specialists that advises in case of cyber security incidents in their industry sector. The sector-CERT collects information about recent cyberattacks and provides general recommendations, as well as advice to individual organisations, to help fend off attacks. A popular method for sharing threat information is a mailing list which is run by a CERT; but in order to collect large amounts of specific information (‘IoC sighting’) in very short time-frames (such as within hours), an automated threat sharing mechanism offers numerous advantages. That is why the sector-CERT operates an automated threat sharing service, such as an MISP server.

The tasks of a sector-CERT in the scenario depicted in **Figure 2**, above, are to

- Receive feeds from CSC: the sector-CERT receives threat information and enters it in its

MISP service;
- Operate MISP service: the MISP service holds IoCs to be scanned for with the cyber security sensor network and is maintained and operated by the sector-CERT. On this platform, the threats combined with their indicators and the reported sightings from the organisations are stored;
- Determine qualitative threat feeds for industry sectors: one of the sector-CERT's core tasks is to find and to subscribe to high-quality and applicable (mostly public non-commercial) threat information sources, also called feeds. After the subscription to this service, the MISP starts to collect threat information from these feeds regularly.
- Provide and manage API access for the sensor network system: the sector-CERT uses the MISP service to provide an interface, so that the organisations (more specifically, sensor nodes deployed in their infrastructures) can access the collected threat information. However, access is limited, and organisations must register with the threat sharing program first. Once registered, organisations share their public keys with the sector-CERT and get credentials, such as a client-side certificate to use the interface and to request threat information;
- Retrieve sightings from organisations: after an organisation has requested threat indicators from the MISP service and has found one of these indicators in its infrastructure, it reports a sighting back to the MISP service. These sightings are then merged with their associated threat indicator and are stored in the MISP's database. In a following step, the sector-CERT aggregates the received sightings to maintain overview about discovered instances;
- Identify wide-spread malicious activities and inform companies: with the help of the collected sightings, the sector-CERT can re-evaluate a large-scale threat situation. Once it has identified a dangerous trend, it first informs the companies participating in the threat-sharing program of the found threat and raises its threat level. Informed companies can then reassess the risk based on the warning and improve planned countermeasures if necessary;
- Derive conclusions from sightings and create a sector CCOP: another task of the sector-CERT is to create a CCOP by aggregating sightings and feedback from the organisations that consumed IoCs;
- Provide sector CCOPs to the CSC: at regular intervals, on request, and in the event of an immediate threat, the sector-CERT sends the current sector-specific CCOP to the CSC. In the event of a major attack, the CSC can use additional information channels to best support the CERT sector with additional information about the threat.

Organisations

An important stakeholder within each organisation is the security operator (and the whole IT team respectively) who keeps (in accordance with the organisation's Chief Information Security Officer [CISO]) the organisation's infrastructure secure. Note, this abstract role resides at a lower, more technical, level within the organisation.

The associated tasks are as follows:

- Manage the Sensor Network Node (SNN): the security operator's main task is to maintain the SNN and to keep it operational and secure. He or she is responsible for the correct positioning, implementation, and configuration, as well as maintenance of the SNNs. He or she is in control of the prioritisation of the scanning tasks and ensures the security of the communication channels between the SNNs and its sensors;

- Manage sensors: the security operator is responsible for registering all to-be-monitored devices at the SNN and equipping them with appropriate sensors. Therefore, he or she must either install a sensor software (agents/daemons) on the devices or gain access to the relevant data via external means (for example, via network taps) and then register them to the SSN and update their configurations accordingly. After this registration is completed, the security operator must determine which types of IoCs can be searched for on a device; then, the corresponding scanning tools need to be installed and configured. Once these two initial steps are complete, the sensor is functional on its own and the SNN can send scan jobs to the sensor;
- Evaluate sightings and create an organisation’s CCOP: based on the collected data of the SNN, the security operator evaluates the current threat status. Based on his assessment of the situation, he or she prepares a CCOP and transmits it to the CISO, which he or she then accounts for in future risk analysis activities;
- Receive action plan from CISO: the security operator receives an action plan comprising upcoming tasks, based on anticipated strategic risks analysed by the CISO. These tasks are to be realised by the security operator in order to implement the strategic security goals of an organisation. Specifically, these strategic goals determine protection levels of assets and, therefore, also scanning intervals, depths, and invested effort into discovering threats in certain network segments or categorised assets. Therefore, the organisation itself determines where to look for certain distributed IoCs.

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 to coordinate mitigation actions across organisations, while the actual mitigation actions are mostly carried out within the organisations. The distribution of detailed recommendations to handle a recently discovered vulnerability and the distribution of an unofficial emergency patch are good examples. 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 regarding the affected organisation, its infrastructure, assets, and services. Eventually, the human in the loop is key to creating appropriate situational pictures.

Acknowledgement

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

References

Barnum, S 2012, ‘Standardizing cyber threat intelligence information with the Structured Threat Information eXpression (STIX)’, MITRE Corporation 11, pp. 1-22.

Cavelty, MD 2014, ‘Reporting and analysis center for information assurance (melani) (phase 2: 2004–2010)’ *Cybersecurity in Switzerland*, SpringerBriefs in Cybersecurity, Springer, CH, pp. 39-55.

Chen, Q & Bridges, RA 2017, ‘Automated behavioral analysis of malware: A case study of wannacry ransomware’, arXiv preprint arXiv:1709.08753.

Choo, KKR 2011, ‘The cyber threat landscape: Challenges and future research directions’, *Computers & Security*, vol. 30, no. 8, pp. 719-31.

Coldebella, GP & White, BM 2010, ‘Foundational questions regarding the federal role in cybersecurity’, *Journal of National Security Law & Policy*, vol. 4, p. 233-45.

Durumeric, Z, Kasten, J, Adrian, D, Halderman, JA, Bailey, M, Li, F, Weaver, N, Amann, J, Beekman, J, Payer, M & Paxson, V 2014, ‘The matter of heartbleed’, *Proceedings of the 2014 Conference on Internet Measurement Conference*, ACM, pp. 475-88.

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 & Brynielsson, J 2014, ‘Cyber situational awareness—A systematic review of the literature’, *Computers & Security*, vol. 46, pp. 18-31.

Luiijf, E, Besseling, K & De Graaf, P 2013, ‘Nineteen national cyber security strategies’, *International Journal of Critical Infrastructures*, vol. 9, nos. 1-2, pp. 3-31.

Miller, DR, Harris, S, Harper, A, VanDyke, S & Blask, C 2010, *Security Information and Event Management (SIEM) Implementation (Network Pro Library)*, McGraw Hill, New York, NY, USA.

Obrst, L, Chase, P & Markeloff, R 2012, ‘Developing an ontology of the cyber security domain’, *STIDS*, pp. 49-56.

Pahi, T, Leitner, M & Skopik, F 2017, ‘Preparation, modelling, and visualisation of cyber common operating pictures for national cyber security centres’, *Journal of Information Warfare*, vol. 16, no. 4, pp. 26-40.

Pahi, T & Skopik, F 2016, ‘A public-private-partnership model for national cyber situational

awareness', International Journal on Cyber Situational Awareness, vol. 1, no. 1, pp. 31-53.
Radunovic, V J 2013, 'DDoS-Available weapon of mass disruption', *Proceedings of the 21st Telecommunications Forum—TELFOR 2013*, IEEE, pp. 5-8.

Rantapelkonen, J & Salminen, M (eds.) 2013, *The fog of cyber defence*, Series 2 Article Collection No. 10, National Defence University, Department of Leadership and Military Pedagogy, Juvenes Print Oy, Tampere, FI.

Rid, T & Buchanan, B 2015, 'Attributing cyber attacks', *Journal of Strategic Studies*, vol. 38, nos. 1-2, pp. 4-37.

Skopik, F 2019, 'The limitations of national cyber security sensor networks debunked: Why the human factor matters', *14th International Conference on Cyber Warfare and Security*, ACPI, pp. 1-8.

Skopik, F, Settanni, G & 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*, vol. 60, pp. 154-76.

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

Swart, I, Irwin, B & Grobler, MM 2016, 'Adaptation of the JDL model for multi-sensor national cyber security data fusion', *International Journal of Cyber Warfare and Terrorism*, vol. 6, no. 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 & Iklody, A 2016, 'MISP: The design and implementation of a collaborative threat intelligence sharing platform', *Proceedings of the 2016 ACM Workshop on Information Sharing and Collaborative Security*, ACM, pp. 49-56.