Contextualised and actionable information sharing within the cyber-security community

Frédéric Garnier

CERT-EU [**]

Abstract. In today's digital world, attacks proliferate and targeted organisations imagine new strategies to better detect, prevent or respond to threats. Information exchange on cyber security and especially cyber threats is developing fast. Information sharing communities take shape, by sector, by country or at international level, usually on a voluntary basis in trusted circles. Security firms understand and stimulate this move by developing new products and services. Organisations foresee benefits from leveraging information sharing. However, as threat information sharing networks emerge and develop, it is necessary to consider how those networks should best be organised and what performance they should deliver on the consuming end. Indeed, nodes constituting networks should have a minimum of functional characteristics to best connect and interact with each other and create added value. This also implies that information exchanges within networks should meet minimum quality criteria, especially in terms of threat contextualisation. Faced with an extremely dynamic threat landscape, the challenge is to automate information sharing and make information delivered on the consuming end immediately actionable.

The purpose of this paper is to introduce and describe a model for a cyber-threat intelligence network as a means for organisations to develop accurate threat situation awareness and better detect or prevent targeted attacks. This is a network of organisations inter-connecting technical platforms for automated exchanges of structured and actionable threat information. The paper proposes a scheme for information flows within the network and a functional model for the nodes (organisation and technical platform) constituting the network. The concept of a cyber-threat intelligence fusion node is developed. Finally, minimum criteria for making such a network efficient are proposed: contextualisation of information, automation of exchanges and structured data packages.

1 Introduction

In today's digital world any organisation having a footprint on the Internet is susceptible of being targeted by attacks. Organisations develop strategies to deter, prevent, detect or respond to such attacks. The diversity

[**]. CERT-EU is the computer emergency response team for the EU institutions, bodies and agencies. See `http://cert.europa.eu`

and dynamic of threats coupled with the intrinsic vulnerability of Internet based technologies make zero risk impossible. In other words, a given organisation will always be exposed to successful attacks. Defenders might schematically action two levers to minimize risks. They might structurally reduce the vulnerabilities, weaknesses and exposure of their most valuable assets. This reduction will be limited by financial, technical and cultural constraints. They might also aim at understanding and mitigating threats they are more specifically subject to. Indeed, not all threats are equally relevant for a given organisation or industry sector and there is always a relation between an attacker and in its victim. An organisation will be more concerned by threats targeting its sector, supply chain or geographical area. It will handle as a priority threats causing more damages, being the most intense or the most persistent, having specific motivation (cyber-crime, espionage or hacktivism). Monitoring these characteristics allows organisations to identify which threats are the most pertinent for them at any specific time.

The activity of monitoring, understanding, characterising and mitigating threats is usually called cyber-threat intelligence (CTI). Because no one can monitor the threat landscape in isolation, organisations engage in threat information sharing. Information sharing groups take form, interact and sometimes overlap. Gradually, a global threat intelligence network is taking shape. It includes organisations of various countries and sectors, mostly on a voluntary basis. It consists of several sub-networks, is decentralised and composed of interconnected nodes. It receives data from collecting sensors and releases data to consuming sensors. The time has come to start modelling this network, in order to describe its expected characteristics and specify minimum performance criteria. The aim is to enable threat mitigation via the proper collection, sharing and consuming of threat actionable data.

This paper provides a synthesis of threat intelligence essentials. On this basis, a model is proposed for the threat intelligence network and its interconnected atomic elements (aka ‘nodes’). Finally, minimum characteristics for contextualised and actionable threat data are proposed.

Section 2 recalls the main concepts of cyber-threat intelligence and provides references to some key contributions in this field. Section 3 introduces the model for a cyber-threat intelligence network. Section 4 deals with the functional architecture of the most important element of the network — the CTI fusion node. Section 5 describes the minimal context information that shall be exchanged. Section 6 introduces criteria for actionable threat information.

Sections 2, 3 and 4 are mainly intended for readers interested in the concept of cyber-threat intelligence networks and fusion nodes. Sections 5 and 6 are for readers interested in minimum criteria for contextualised and actionable information.

2 Cyber threat intelligence (CTI)

The axiom underlying cyber-threat intelligence practices is that organisations have a better chance of defending themselves against attacks if they understand:

- Who is attacking or potentially to targeting them,
- How the adversary is realising attacks,
- What is being targeted,
- Where attacks are taking place,
- When the attackers are active.

Cyber-threat intelligence essentials are about profiling the malicious actors of the Internet (their motivation, capabilities, and historical activities), understanding which techniques, tactics and procedures (TTPs) are being used to better detect or counter them, and monitoring past or current campaigns to assess proximity, imminence or likelihood of attacks. Some notable contributions in this field are available in references [1,2,3].

This intelligence must rely on facts and technical observations. Defenders collect technical data related to attacks from different sensors and investigation tools. Indeed, malicious activities on the Internet leave traces:

- Command & Control IP addresses and domain names,
- Malicious URLs,
- Simple Mail Transport Protocol (SMTP) headers, email addresses, subject lines, and contents of emails used in phishing attacks,
- Malware samples and artifacts,
- Exploit code,
- Packet captures of attack traffic,
- NetFlow data.

A set of observations related to a suspicious or malicious activity is usually called an “indicator”. An indicator packs these observations (aka observables) with associated context: when and where it was seen, at which stage of the attack sequence was it observed, what level of certainty one has that it is related to malicious activities.

Security actors cooperate on cyber-threat intelligence. Information exchanges develop between cyber-security firms, independent experts, research institutes, computer emergency response teams (CERT), law enforcement authorities and organisations wishing to improve their cybersecurity posture. Such exchanges help sharing work (no single organisation can monitor everything), sharing expertise (defeating advanced intrusion techniques requires specialisation of researchers on the defender side) and improving situation awareness. Ultimately, such exchanges must deliver valuable and actionable output to organisations defending their assets against attacks. In this context, it is important to formulate objectives that should be achieved by threat intelligence activities:

- **Technical defence** — ability to detect, prevent or respond to single instances of malicious activities (identify and block a spear-phishing attempt, prevent the execution of an exploit, block the navigation to a temporarily infected legitimate website, eradicate malware implants on a set of infected host, etc.)
- **Tactical defence** — ability to detect, analyse and defeat a campaign of attacks lasting several weeks or months and leveraging special techniques, tactics and procedures (raise awareness concerning social engineering, block delivery and command and control infrastructure, deploy relevant patches across the defender's infrastructure, etc.)
- **Strategic defence** — ability to recognise the malicious activities of a group of malicious actors over several months or years, deter them or make the cost significantly higher for the attacker.

Figure 1 illustrates interactions and objectives of cyber-defense levels.

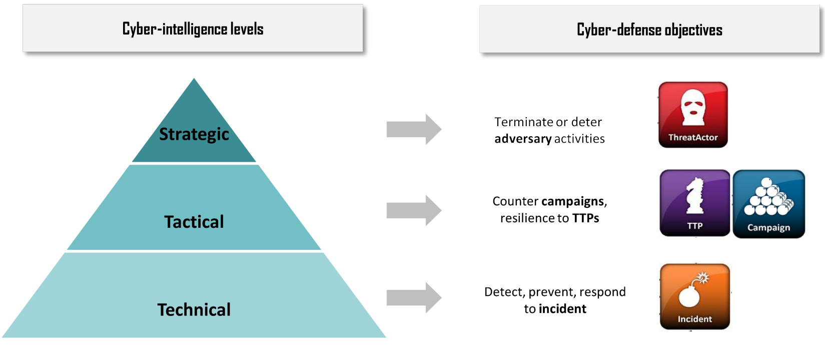
Cyber-intelligence levels
Cyber-defense objectives
Strategic
Tactical
Technical
Terminate or deter adversary activities
ThreatActor
Counter campaigns, resilience to TTPs
TTP
Campaign
Detect, prevent, respond to incident
Incident

Fig. 1. Intelligence and layered defense

Upper layers of defence depend on the underlying one(s), and vice-versa. It is illusory to think strategic defence, without solid technical and tactical defences. These layers enable collection of technical information on the attackers / campaigns / TTPs and support the development of solid strategic defence plans. On the other hand, focusing on technical defence only may just consume resources on trying to prevent multiple attacks without prioritising, deriving lessons learnt and setting plans to defeat or deter the most dangerous adversaries.

3 Networked cyber threat intelligence

No single organisation can in isolation appropriately monitor, understand and characterise dynamic threats. Organisations want to benefit from detections, investigations, analyses and context enrichments shared by others. In this paper, “organisation” refers, on the one hand, to any entity owning an IT infrastructure with a footprint in the cyber-space and wishing to defend itself against attacks (government institutions or agencies, NGOs, companies, etc.), and on the other hand, to professionals and firms delivering cyber-security services.

Cyber-defence evolves toward communities of organisations willing to improve their security posture (or the posture of their constituents / customers) leveraging cyber information sharing. Sharing communities are formed according to diverse criteria, such as:

— Industry sector (e.g. Information Sharing and Analysis Centers — ISACs),
— Country or group of countries (e.g. national CERTs and their constituency, European Governmental CERTs, etc.),
— Other mutual interests (e.g. a firm in its supply chain, a private CERT and its constituency, etc.).

Threat information circulates within and across sharing groups and can be of diverse nature (cf. section 2). Single organisations often participate in several sharing groups. Hence, threat information sharing groups form a complex eco-system. To ensure information sharing delivers the expected added value, these sharing groups can be thought of as networks and basic network engineering techniques can be applied to model this eco-system.

The present section introduces a high level model for networked CTI:

— What circulates in the network (e.g. data flows),
— Where data are introduced and consumed in the network (entry- and exit- points),
— How the global CTI network and sub-networks (e.g. sharing groups) are structured.

3.1 Data flows

Different levels of threat intelligence circulate in these networks. We focus here on the following flows:

- **Technical** data flows — Host or network based detections, indicators of compromise with minimal context (see section 5). Technical threat intelligence can typically be used automatically and immediately in IT security devices (host or network-based IDS, host scanners, etc.). Examples:
  - Indicators for a spear-phishing email (email source, malicious URL, filename | hash value of attachment, date and time),
  - Malware sample with hash value and embedded C&C,
  - Malicious servers delivering a malware,
  - Infected domains redirecting to malware delivery servers,
  - Etc.
- **Tactical** data flows — Investigation findings on special techniques, tactics and procedures or campaigns. Tactical threat intelligence may usually not be used automatically and directly in IT security devices. It doesn't lead to immediate technical action. Consuming of tactical threat intelligence can help review security controls, better organise defence-in-depth, raise awareness, or prioritise hunting of threats. Examples:
  - Lateral movement techniques (e.g. Pass-The-Hash),
  - Techniques for tracking individual victims (e.g. implant of persistent cookies),
  - Techniques to evade anti-malware capacities,
  - Campaign of attacks leveraging watering holes techniques to infect victims from a special industry sector or a geographic area,
  - Etc.
- **Strategic** data flows — Actor's profiling, objectives, current and past activities and possible weaknesses. Consuming of strategic threat intelligence can support policy making. Examples:
  - Historical activities of a specific threat actor,
  - Threat landscape within a given industry sector,
  - Etc.

Technical, tactical and strategic exchanges enrich and complement each other. Not all organisations are capable of generating information for the tactical and strategic levels because this requires experience and advanced capabilities in terms of detection and investigation. However any organisation can achieve a minimal level of maturity, become able to detect single instances of attacks and hence generate useful information flow at technical level. The active participation of as many organisations as possible at technical level is essential because the variety of technical detections is a key success factor for investigation and intelligence at tactical and strategic levels.

Figure 2 illustrates the levels of information sharing.

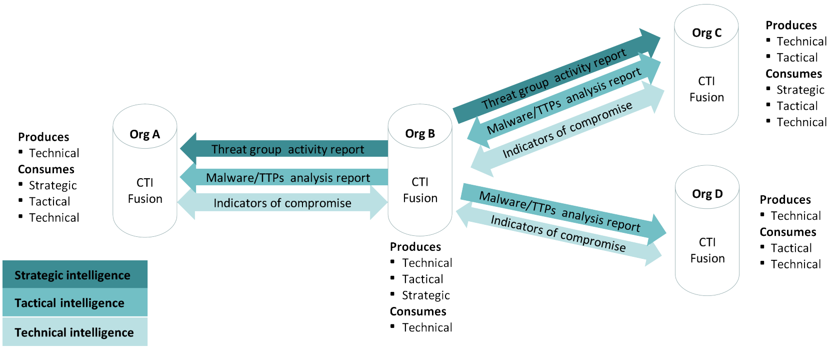

Fig. 2. Levels of cyber-intelligence exchanges

In this example, different levels of maturity cooperate and provide added value in the network:

- **High maturity** — Organisation B produces and consumes any levels of threat intelligence, from technical to strategic. It is able produce a comprehensive threat landscape description for given sectors or geographical areas. It plays a central role in this network. However its capacity to produce tactical and strategic intelligence depends on the contributions from others.
- **Intermediate maturity** — Organisation C produces threat intelligence up to the tactical level and consumes up to the strategic. This organisation provides a good tactical contribution to the network by releasing or enriching investigations on TTP or malware analysis.
- **Minimal maturity** — Organisation D produces only technical and consumes up to tactical. Organisation D can receive strategic threat analysis, but is not in position to make policy decision to influence the threat landscape. Organisation A produces only technical and consumes up to strategic. Such organisations have no specific investigation capabilities, but share interesting technical detections with the community.

Beyond the maturity level, there are other factors that limit active participation of all actors to the tactical and strategic levels (e.g. geopolitical context, economic competition, etc.). If they cannot or do not want to participate to tactical or strategic exchanges, organisations participating to the sharing network should whenever possible share indicators of compromise and hence contribute to technical threat intelligence data flows.

3.2 Entry and exit points

Indicators of compromise exchanged at technical level are generated based on detections by sensors or and are aimed to be consumed by other sensors on the other end of the data flow. The technical threat intelligence network shall support sensor-to-sensor information flows. In this model, sensor means any device or collection of devices capturing and handling network or application-level data flows in view of detecting, preventing or responding to attacks (suspected or actual).

At the beginning of the sensor-to-sensor chain, originating sensors are those via which attacks are observed. They support the production of initial data. At the end of the chain, consuming sensors make use of data as feeds for prevention or detection. These data are those that have been produced from the originating sensors and have been handled and enriched throughout the sensor-to-sensor chain. The same sensor can be either an originating or consuming sensor depending its role in a given sensor-to-sensor data flow. Based on data used as feeds, a consuming sensor may detect attacks and allow collection of more data on the specific threat. This data is returned into the sensor-to-sensor chain and the sensor becomes an originating sensor for this new data flow.

In between originating and consuming sensors, a series of CTI fusion nodes collect, handle and share information. CTI fusion nodes exchange data between each other. There can be one or more CTI fusion nodes between originating and consuming sensors. Most of the fusion nodes also interact with originating and/or consuming sensors.

Figure 3 provides a high level illustration of information flows from sensors to sensors via fusion nodes.

Across the network, end-to-end information flows remotely connect sensors to sensors, via CTI fusion nodes. At each CTI fusion node, handling can be made on data. Through the network, the high level workflow is:

1. Production — Original technical data are collected via originating sensors (IDS, SIEM, forensic tools, etc) on the infrastructure of a

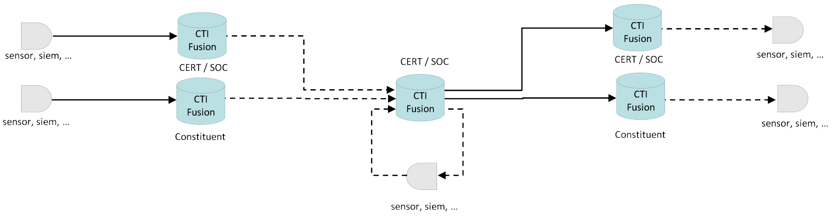

Fig. 3. Information Flow through Fusion Nodes

victim of an attack. This collection is done in the context of incident response, or during the monitoring of malicious activities (intrusion attempts, etc). These sensors are under the responsibility of a first CTI fusion node operated directly by the victim (e.g. an internal SOC) or by a CERT to which the victim reports. In the latter case the victim is a constituent of the CERT. This collection results in the production of information (e.g. indicators of compromise) to be transported through the network.

2. Collect-Handle-Share — Information is shared from the first CTI fusion node with other CTI fusion nodes in accordance with authorisations/restrictions set by the victim organisation. Each CTI fusion node should have a clearly established information sharing policy which regulates what can be shared with whom and when. At each hop from one CTI fusion node to another, information is handled, possibly enriched and consolidated with information coming from either local sensors or other CTI fusion nodes. To perform fusion tasks, it is essential that a node receives a minimal set of information on the threat context (when, where, how, etc.), i.e. contextualisation. Once this handling is completed, information can be further shared with other CTI fusion nodes and/or directly with consuming sensors.

3. Consuming — At the end of the chain, information serves as feed for IT security devices (aka sensors) of a consuming organisation. An organisation may operate its own CTI fusion capability and then merge technical data from multiple sources before consuming. Or the organisation relies on the CTI fusion node of a parent CERT or SOC and will consume technical data released by it. In either case the consuming organisation needs to obtain data that can unambiguously and directly be used to feed its sensors, i.e. actionable information.

3.3 Local sharing networks

The CTI network can take different forms. It is best to consider it as a network of networks: a global network made of local networks (or clusters of local networks). Indeed groups of organisations organise their threat intelligence information sharing and form local networks. Nodes within these local networks are connected to each other. Some nodes belong to different local networks and support connectivity between these local networks. Most also connect to the global network.

Local networks can be organised in different manners. Two typical models are presented below:

- *Hub and spoke* (figure 4) — A central CTI fusion node in the local network (A) is the prime interface between the global network and other organisations in the local network. Within this local network, the central CTI fusion node pools and shares information. This is typically the model for a CERT and its constituency.

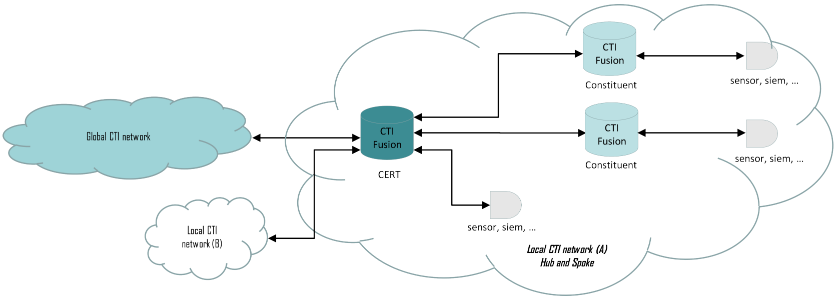

Fig. 4. CTI local network — Hub and spoke

- *Peers to peers* (figure 5) — Several CTI nodes form trusted groups and share information on peer-to-peer basis. Each CTI node may participate to other sharing groups and is integrated into the global CTI network.

4 CTI fusion nodes

In this section we will focus on the core component of the sensor-to-sensor chain, the CTI fusion node. Any organisation engaging in CTI networking operates some kind of CTI fusion node for handling threat data. Technical features (volume of data, number of connections with other nodes, etc.) can vary depending on the position of the operating

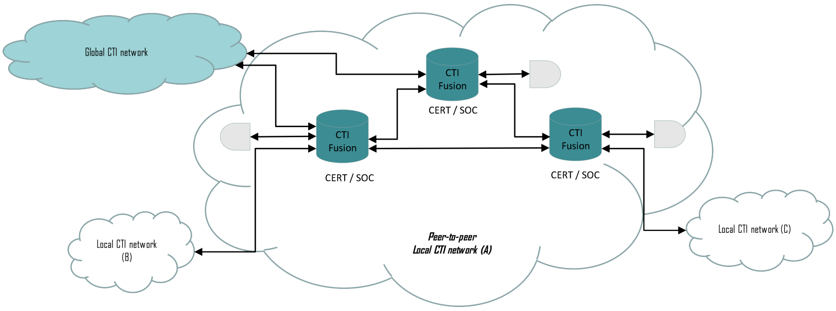

Fig. 5. CTI local network — Peer to peer

organisation in the network, its capacity and maturity. However, functions of a CTI fusion node always include the collection of information from originating sensors or other nodes, the handling this information and its sharing with consuming sensors or other nodes. Additionally, the functioning of the sensor-to-sensor chain supposes that individual nodes meet minimal performance characteristics (accuracy, freshness, completeness, etc.). Indeed, network nodes must not create data quality degradation. This section introduces a standard functional model for a CTI fusion node.

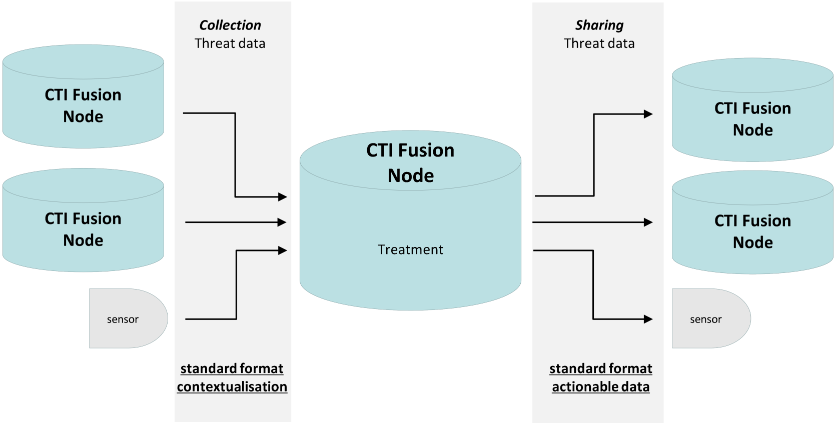

Fig. 6. CTI Fusion Node

4.1 Collection

CTI fusion nodes collect information from internal and external sources. From manual to fully automated collection, different collection modes are possible. For a given organisation, internal sources consist in individual sensors or clusters of sensors orchestrated by devices such as security incident and event managers (SIEM). External sources are other CTI fusion nodes, operated by diverse organisations. For a CERT or SOC, usual external sources are peers (e.g. other CERTs/SOCs within the same sector or cross sectors), partners (other categories of cyber-security actors with which the organisation has established partnership and information exchange agreements), and open sources. Information is shared on a bilateral basis or within information sharing groups (e.g. ISAC). These sources should operate functionally equivalent CTI fusion nodes.

4.2 Handling

CTI fusion nodes realise diverse handling operations on data that they ingest, aiming at:

1. Operational exploitation (consuming edge) of data within the organisation operating the fusion node,
2. Prepare sharing of data with other CTI fusion nodes,
3. Threat situation awareness.

It is crucial that a CTI fusion node meets minimum performance characteristics: it shall not reduce the quality level of data received and should enrich the data whenever possible. The functional modules of a CTI fusion node are described in the subsequent sub-sections.

Import control This function ensures that data ingested into a CTI fusion node meets minimum quality standards. It makes use of the technical checks capacity (see 4.2) before external data are actually imported in the CTI fusion node. It typically checks that data are properly contextualised and reasonably fresh, that they are appropriately structured and will not generate “noise” in the CTI operational picture of the receiving organisation. For example, this capability controls that:

1. Information originating from public sources will not be imported twice from different channels,
2. Incoming data is properly contextualised (e.g. timing, sighting, kill chain),
3. Data obtained from CTI servers / sensors are refreshed appropriately,
4. Minimal metadata are contained in the incoming data package (e.g. producer, traffic light protocol label, title, description, etc.),

5. New sources are tested and meet minimal performances for the CTI node before the plug-in is considered operational.

Reliability This function indexes information with a reliability metric. Each organisation operating a CTI fusion node should characterise the sources of information it uses in terms of reliability. This is essential to maintain trust through the sensor-to-sensor value chain. A possible model is based on military intelligence notation:

1. Source reliability: A — Completely reliable source, B — Usually reliable, C — Fairly reliable, D — Not usually reliable, E — Unreliable, F — Reliability cannot be judged. This criteria should be managed dynamically (i.e. the note of a source should change if it show variations in the quality of information provided). Observed degradation of reliability should be lead to revising import control settings appropriately (e.g. ban the source).
2. Information reliability: 1 — Confirmed; 2 — Probably true, 3 — Possibly true, 4 — Doubtfully true, 5 — Improbable, 6 — Cannot be judged. A rating should be provided by the producer, it can be modified by the receiving organisation based on its own analysis or by correlating the same information coming from distinct independent sources.

Technical checks This function verifies that technical data ingested in the CTI fusion node are actually indicators of malicious activity and are actionable. The goal is to focus on the most pertinent and fresh technical data, limit errors and reduce false-positives.

Example of technical checks classes are listed in the following table.

Table 1: Technical Checks — Examples

Class	Example of checks	Applicable to
False Positive	• Reported as false positive by a partner • Reported as false positive by a constituent	IPs, domains, email addresses, etc.
White listed	• Legitimate and owned by a partner • Legitimate and owned by a constituent • Good reputation / high ranking	IPs, domains, email addresses, etc.
	• Known hash	Hash values (MD5, SHA1, SHA256, etc.)

Class	Example of checks	Applicable to
Invalid	• Invalid syntax • Invalid hash	IPs, domains, email addresses, etc. Hash values (MD5, SHA1, SHA256, etc.)
Time To Live (TTL)	• Time to live expired	IPs — TTL : Hours — Days (e.g. 72 hours) Domains — TTL : Weeks — Months (e.g. 6 months) URLs — TTL : Weeks — Months (e.g. 6 months) Hash values — TTL: Years
Not actionable	• Too generic • Valid user agent	Pattern in traffic User agent

These checks should be run:

1. While data are being imported (import control),
2. Before they are shared with other CTI fusion nodes or released to sensors (sharing control),
3. On an ongoing basis, data should also be periodically re-checked (e.g. overnight tests over the database).

Indeed, the malicious nature of technical CTI data is intrinsically dynamic. Attackers use dynamic combinations of legitimate, dedicated, randomly generated set domains / hostnames, IP addresses, email or social media accounts to build their multi-stages malware delivery strategies or evolving C&C infrastructures. When technical checks detect that CTI data correspond to domains, IP addresses of the organisation, its constituency or its partners, results can be used to alert them and avoid adverse impact in terms of possible extensive blacklisting.

Contextualisation This function manages the context information associated with raw CTI data. Context information is essential to support correlation (see section 4.2) and to enable appropriate use of CTI data in consuming sensors. Minimal contextualisation includes:

1. Timing: when a threat observation took place, when the threat vector or payload started to be active, if it still in use.
2. Sighting and targeting: where the observation was made and if specific industry sectors or geographical areas might be targeted.
3. Kill chain: at which stage of the intrusion kill chain the observation was made.

Details on contextualisation are addressed in section 5. In a nutshell, this purpose is to make sure that timing, sighting, source and kill chain information exist, are complete, accurate and maintained / enriched over time. When different sources provide information in various manners, the contextualisation function of the CTI nodes makes sure that they are represented in an homogeneous way in the node and that possible gaps in contextualisation are appropriately addressed (e.g. setting 'conservative' default values, degradation of the reliability ranking of the source, local 'ban' of the data within the data base, etc.).

Correlations This function detects correlations between data in the fusion node. A CTI fusion node collects information from diverse sources. There can be overlaps between incoming CTI data packages. If correlation is not correctly handled, the risk is redundancy and confusion. With proper correlation, this is a source of enrichment. Correlations can support inter alia: situation awareness (e.g. understanding which threat is targeting whom and when), TTP characterisation (e.g. aggregating diverse pieces of knowledge for a given modus operandi), adversaries profiling (e.g. recognising the signature of an adversary, its objectives and historical activities), malicious campaigns monitoring (e.g. first signs, intensification, going below the radar, resuming, etc.). One can schematically distinguish two levels of correlations:

1. Technical correlations — Indicators that share similar observables are "correlated". Observables (IP addresses, hostnames, email addresses, etc) play the role of pivot to link different indicators.
2. Advanced correlations — Campaigns, threat actors or techniques / tactics / procedures can be associated based on technical correlations.

Correlation is however a complex issue. Research has just begun and spectacular developments are possible. More advanced functional description of threat correlations can be developed. One should however always bear in mind that deception is possible and that the maliciousness of IP addresses, hostnames or other observables is very volatile.

Course Of Action This function manages recommended and actual actions made with threat data. Any CTI fusion node should be able to release information directly usable in sensors. However, not all observables may be used for detection, prevention or investigation. An important capability is to generate relevant rules from observables and specify recommended actions that may be performed with such observables or rules. Typical categories of actions to be performed with threat observables are:

1. Detection (via host-based or network-based intrusion detection systems, etc.).
2. Prevention / Denial (via intrusion prevention systems, proxies, mailguards, firewalls, etc.).
3. Investigation (via security information and event managers, log analysers, etc.).
4. Intelligence / awareness raising (via CTI fusion nodes, etc).

The course of action may depend on a consuming organisation's specific policy and constraints. Recommended course of action will guide them and limit the risk of exploitation mistakes (e.g. blacklisting a legitimate website that was temporarily used to redirect to a malicious domain, blocking a C&C domain when it would be preferable to observe the adversary behaviour before initiating the response, etc). Finally, proactive researches of infections by organisations might be time consuming when dealing with advanced and very stealthy threats. Course of action functionalities should help prioritisation and decide which threats should be hunted first.

Situation awareness This function elaborates and maintains pertinent threat situation awareness. A CTI fusion node provides situation awareness for the community it serves. The networking of CTI fusion nodes augments this capacity. CTI fusion nodes assemble from local and remote sensors discrete instances of attacks (incidents), infrastructures, vectors, targeting, etc. By merging data, CTI fusion nodes can realise dynamic threat characterisation. Typical drivers for threat situation awareness are:

1. Know threats that are currently targeting or susceptible to target my organisation, my sector, my constituents, my supply chain or my partners.
2. Know characteristics of most critical threats and recommended / actual strategies to mitigate them.
3. Know most active partners for monitoring, characterising, sharing given threats.

For situation awareness functions, CTI fusion nodes interact with other tools used by the community it serves.

Analysis Given the complexity of characterising threat components, especially techniques, tactics and procedures, a CTI fusion node incorporates analytic tools to support operator's work. The objective here again is to create added value in the sensor-to-sensor chain. Each CTI node fuses information depending on its capacity. Organisations engaging in CTI networks produce their own analytics and share with others. Analytics complete or confirm each other and enable to derive better description on the adversaries' TTP, historical activities / targeting and ongoing campaigns.

Taxonomy This function ensures that the different threat components are appropriately categorised and assembled. CTI fusion nodes must understand each other. It is vital that they adopt common taxonomies and represent threats the same way. The current recommended taxonomy for CTI information is STIX (see reference [3]). Additionally, there is a need to manage an object which is currently not part of STIX, "organisation": an organisation is a producer of threat information, a victim of attacks, a constituent or a partner. The main threat components to be handled in a CTI fusion node are hence:

1. Observables,
2. Indicators,
3. Incidents,
4. Campaigns,
5. Threat actors,
6. TTPs,
7. Exploit target,
8. Courses of actions,
9. Organisations (producers, partners, constituents, victims).

Import/Export formats This function enables read & write for relevant technical data / file formats. To integrate in sharing networks CTI fusion nodes should be able to read and produce cyber threat information in different formats. The adoption of a common and universal format (such as STIX's XML) is a goal for many organisations but it cannot be mandated. A pragmatic approach for a CTI fusion node is to understand with whom it exchanges information and implement the relevant format converters. Any structured format should be acceptable as long as the taxonomy of represented objects is consistent with the common taxonomy (e.g. STIX threat components).

Sharing policy This function manages the sharing rules. An essential function of CTI fusion nodes is the implementation of rules determining which information can be shared with whom and when. There are legal and policy constraints related to information sharing that have to be addressed by each organisation and within each sharing network. A CTI fusion node will properly integrate in a sharing network if its sharing policy is well defined and is compatible with those of other organisations in the network. Examples of parameters that might be considered to define an information sharing policy are:

1. Anonymisation — Threat intelligence activities involve the handling of personal data, with different regulations associated to such handlings. Additionally, the security and interests of victims (reputation, competitiveness...) must always be preserved. Therefore victim's anonymisation is fundamental in information exchanges. A correct data structure inside the CTI fusion nodes allow identity to be handled locally only and not be exposed to sharing.
2. Intellectual property — Threat information may be obtained from sources (e.g. commercial feeds) that impose restriction in terms of further sharing.
3. Traffic Light Protocol — CTI fusion nodes handle information that they own (i.e. generated from their own sensors) and information that is owned by others. Different classification, need-to-know and right-to-share systems exist. The Traffic Light Protocol is a widely used right-to-share system. It has however to coexist with regulations and policies depending on the sector, country or group of countries.
4. Recipient — The organisation operating the CTI fusion node should categorise possible sharing partners since not all information may be shared with all. Typically, the TLP refers to peers, constituents, customers, membership, etc.
5. Producer. The organisation operating the CTI fusion node should categorise possible sources of information as this will influence the recipients of sharing.
6. Targeted domain — To limit the production of noise, CTI fusion nodes may decide to share with some partners threat that have been detected in a given domain only.
7. Threat level — To limit the production of noise, CTI fusion nodes may decide to share with some partners only information related to a significant threat level.

Figure 7 illustrates how the different parameters can be combined to form sharing rules.

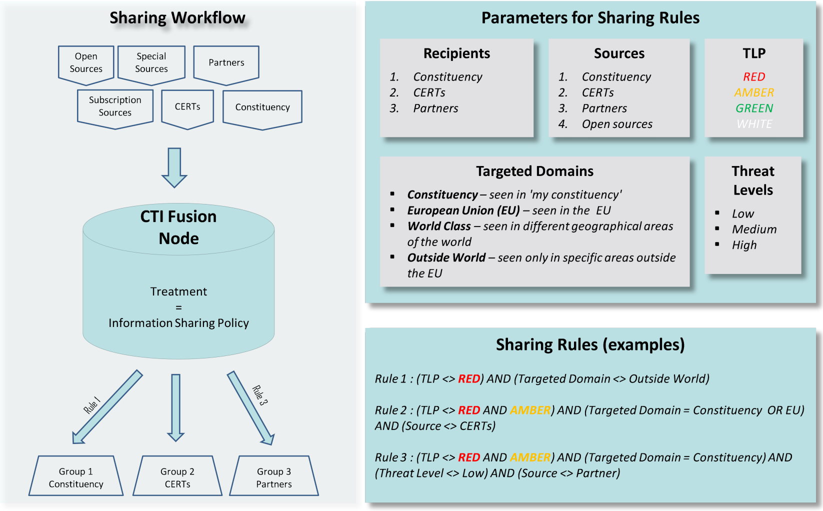

Parameters for Sharing Rules

Recipients

1. *Constituency*
2. *CERTs*
3. *Partners*

Sources

1. *Constituency*
2. *CERTs*
3. *Partners*
4. *Open sources*

TLP

RED
AMBER
GREEN
WHITE

Targeted Domains

- ***Constituency*** *– seen in 'my constituency'*
- ***European Union (EU)*** *– seen in the EU*
- ***World Class*** *– seen in different geographical areas of the world*
- ***Outside World*** *– seen only in specific areas outside the EU*

Threat Levels

- *Low*
- *Medium*
- *High*

Sharing Rules (examples)

*Rule 1 : (TLP <> **RED**) AND (Targeted Domain <> Outside World)*

*Rule 2 : (TLP <> **RED** AND AMBER) AND (Targeted Domain = Constituency OR EU) AND (Source <> CERTs)*

*Rule 3 : (TLP <> **RED** AND AMBER) AND (Targeted Domain = Constituency) AND (Threat Level <> Low) AND (Source <> Partner)*

Fig. 7. Example of sharing rules construction

Sharing control This function controls the information sharing workflow. A CTI fusion node must establish controls to ensure that the sharing policy is being strictly respected and that data leakage is prevented. Interconnection of CTI fusion nodes must remain a lever for the community of defenders to improve collectively their security posture. It must not create adverse impacts for members of the community (i.e. leakage of sensitive information) or flood the CTI network with noise (i.e. redundant, unconfirmed, low quality / low pertinent information).

Production A CTI fusion node interacts primarily with: (a) other functionally similar nodes, (b) sensors. A CTI fusion node shall therefore produce data that can be used by other nodes in the CTI sharing networks and serve as feeds in sensors (local or remote). In the first case, produced data shall correspond to the expectations of another CTI fusion node in terms of structure and content. The structure should whenever possible comply with most popular standards (e.g. STIX). The content shall include CTI raw and context data to enable receiving CTI node to do their fusion work. These minimum context data will be detailed in the next section of this paper.

4.3 Summary

Figure 8 illustrates the standard functional architecture of a CTI fusion node.

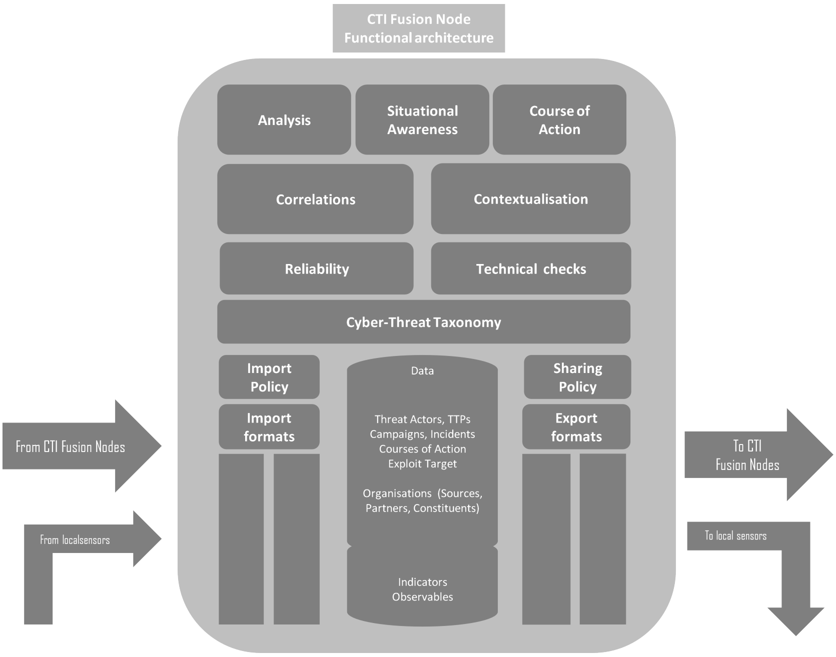

Fig. 8. CTI Fusion Node

The following table summarises the fundamentals of CTI fusion nodes functions.

Table 2: CTI fusion node — functional basics

Category	Function name	Function essentials
Collection		• Internal sources • External sources • Fusion nodes • Sensors

Category	Function name	Function essentials
Handling	Import control	• Completeness • No redundancy • Granularity
	Reliability	• Source reliability • Information accuracy • Confidence rating • Maintain trust through the sensor-to-sensor value chain
	Technical checks	• Maliciousness • Actionable • Freshness
	Contextualisation	• Timing • Sighting / Targeting • Kill chain
	Correlations	• Fact-based (Observable based) • Situation awareness • TTP characterisation • Adversaries profiling • Campaigns monitoring
	Course Of Action	• Recommended use of CTI data • Threat handling prioritisation • Compliance with security policy • Avoid exploitation mistakes • Reduce false positives
	Situation awareness	• Active threats • Active reporters, investigators, monitors • Current defensive tactics
	Analysis	• Create added value in the CTI fusion node network • Complete and confirm analytics • Increase TTPs, Campaigns and Adversaries understanding
	Taxonomy	• Common threat taxonomy • Interoperability between CTI fusion nodes
	Import/Share formats	• Read / write different formats • Converters

Category	Function name	Function essentials
		• Adapt to other interacting CTI fusion nodes
	Sharing policy	• Combine need-to-know and right-to-share • Implement TLP and other relevant classification systems • Victim's identity anonymisation • Compatible sharing policies in CTI networks
	Sharing control	• Implement sharing controls iaw sharing policies • Prevent data leakage
	Production	• Data content adapted to intended use (other CTI fusion nodes or sensors) • Data structure supporting automation
Sharing		• Constituents / Customers • Peers • Partners • Public • Sensors

5 Contextualised data

Individual performances of each CTI fusion node are essential for the well-functioning of the sensor-to-sensor chain. Fusion nodes must receive and produce data of sufficient quality. This assertion is valid whether data is going to be used by other fusion nodes or is going to be consumed locally by sensors owned by the operating organisation. When they handle data, fusion nodes should not cause data quality deprecation and should create added value for the next fusion node. The quality of CTI data relies on the existence of context information. Without minimal context (what, where, when, how), "raw" threat data is not properly actionable and might even be counter-productive.

The causes of lack of context are multiple. CTI fusion nodes collect data from diverse sources and usually the volume of data is huge. There are cases where accuracy of data is not guaranteed (poor quality sources or poor quality data in certain circumstances when a threat is not well understood yet). Some sources only provide a limited context: raw data is provided without clear timing (when was this detected, when did this start to be malicious, etc.), no clear sighting or targeting (where was this threat seen, what sector / location was it targeting), no clear scope (what is this threat actually causing in terms of damaging consequences, is it about denial of services, data leakage, ICS disruption / tampering, etc.). In many cases, it is simply because the preservation of the interests of the victim dictates too many restrictions in terms of context sharing. It can also be because the source does not consider the expectations of the sensor-to-sensor model, which implies that shared information will ultimately be used in a sensor.

The damaging consequences are multiple:

- **Noise** — Non contextualised data will cause important data (e.g. targeting my industry sector or seen in my supply chain) to be lost in noise (e.g. low priority or not related to pertinent threat).
- **Difficult prioritisation** — Without context, it is difficult to determine which threat should be handled first, especially when resources of the defender are scarce. Some sophisticated attacks cannot be detected by automated sensors and will need some proactive and time consuming research by the possible victim.
- **False positive** — Without context, threat data are difficult to exploit appropriately in sensors. Data can be outdated or too generic and therefore create false positives.
- **Not actionable** — When there is no indication on how the observable was used in the intrusion kill chain, it is difficult to make the right decision in terms of handling (monitoring, blocking or investigating possible intrusion).

The structure of technical CTI data is composed of raw data and context:

- **Raw data** — Atomic threat data are labelled as observables in the STIX model. They are technical parameters used to create rules for detection, prevention and investigation. Example of raw data types are IP addresses, domains/hostnames, filenames, hash values (md5, sha1, sha256), URI/URLs, email-src, email-dest, email-subject, etc. There are still too many cases where set of raw data are provided in files with no or only vague references to malwares, a campaigns, a threat groups and without a minimum level of technical context (timing, sighting, kill chain).

- **Context** — There are different expectations and understandings for contextualisation. The definition proposed in this paper intends to support minimal performances expected for CTI fusion nodes in the sensor-to-sensor chain. Four components of technical context are recommended to be shared along with raw data:
  - Timing (WHEN),
  - Targeting (WHERE),
  - Kill chain (HOW),
  - Scoping (WHAT).

The cases depicted in Figures 9, 10 and 11 provide examples of technical threat data shared with more or less context.

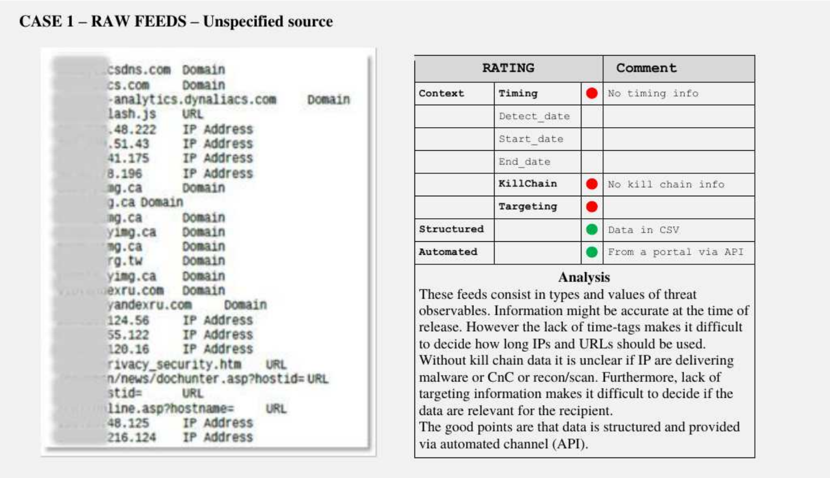

CASE 1 – RAW FEEDS – Unspecified source

[redacted]

```
csdns.com  Domain
cs.com     Domain
-analytics.dynaliacs.com    Domain
lash.js    URL
.48.222    IP Address
.51.43     IP Address
41.175     IP Address
8.196      IP Address
mg.ca      Domain
g.ca Domain
mg.ca      Domain
yimg.ca    Domain
mg.ca      Domain
rg.tw      Domain
yimg.ca    Domain
exru.com   Domain
yandexru.com     Domain
124.56     IP Address
55.122     IP Address
120.16     IP Address
rivacy_security.htm    URL
n/news/dochunter.asp?hostid= URL
stid=      URL
line.asp?hostname=    URL
48.125     IP Address
216.124    IP Address
```

RATING			Comment
Context	Timing	🔴	No timing info
	Detect_date		
	Start_date		
	End_date		
	KillChain	🔴	No kill chain info
	Targeting	🔴	
Structured		🟢	Data in CSV
Automated		🟢	From a portal via API

Analysis

These feeds consist in types and values of threat observables. Information might be accurate at the time of release. However the lack of time-tags makes it difficult to decide how long IPs and URLs should be used. Without kill chain data it is unclear if IP are delivering malware or CnC or recon/scan. Furthermore, lack of targeting information makes it difficult to decide if the data are relevant for the recipient.

The good points are that data is structured and provided via automated channel (API).

Fig. 9. Case 1 — Raw feeds

5.1 Timing

The purpose of time tagging threat data is to understand **when** the threat was detected. Association of timing context to raw threat data serve several objectives for a local CTI fusion node and throughout the sensor-to-sensor CTI network:

- **Manage ageing data** — Given the dynamic nature of threats, CTI fusion nodes continuously check that raw threat data are not outdated. CTI fusion nodes may not ingest outdated data and should refrain from sharing it further. For sensors under their

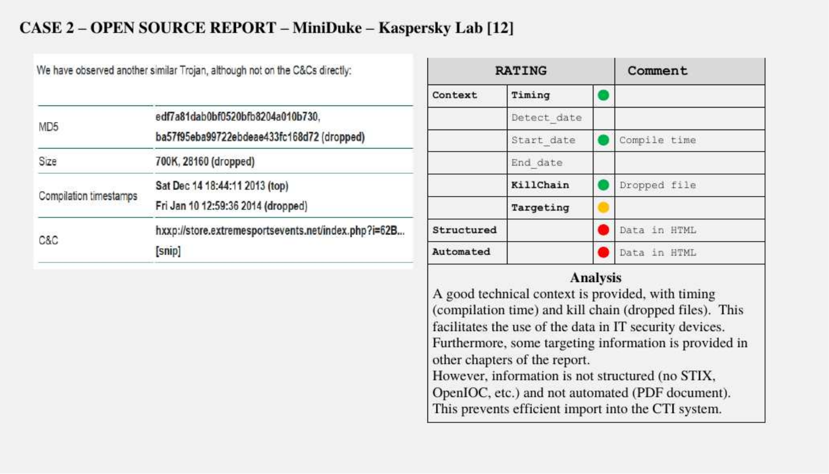

CASE 2 – OPEN SOURCE REPORT – MiniDuke – Kaspersky Lab [12]

We have observed another similar Trojan, although not on the C&Cs directly:

MD5	edf7a81dab0bf0520bfb8204a010b730, ba57f95eba99722ebdeae433fc168d72 (dropped)
Size	700K, 28160 (dropped)
Compilation timestamps	Sat Dec 14 18:44:11 2013 (top) Fri Jan 10 12:59:36 2014 (dropped)
C&C	hxxp://store.extremesportsevents.net/index.php?i=62B... [snip]

RATING			Comment
Context	Timing	🟢	
	Detect_date		
	Start_date	🟢	Compile time
	End_date		
	KillChain	🟢	Dropped file
	Targeting	🟡	
Structured		🔴	Data in HTML
Automated		🔴	Data in HTML

Analysis

A good technical context is provided, with timing (compilation time) and kill chain (dropped files). This facilitates the use of the data in IT security devices. Furthermore, some targeting information is provided in other chapters of the report.
However, information is not structured (no STIX, OpenIOC, etc.) and not automated (PDF document).
This prevents efficient import into the CTI system.

Fig. 10. Case 2 — Open Source Report

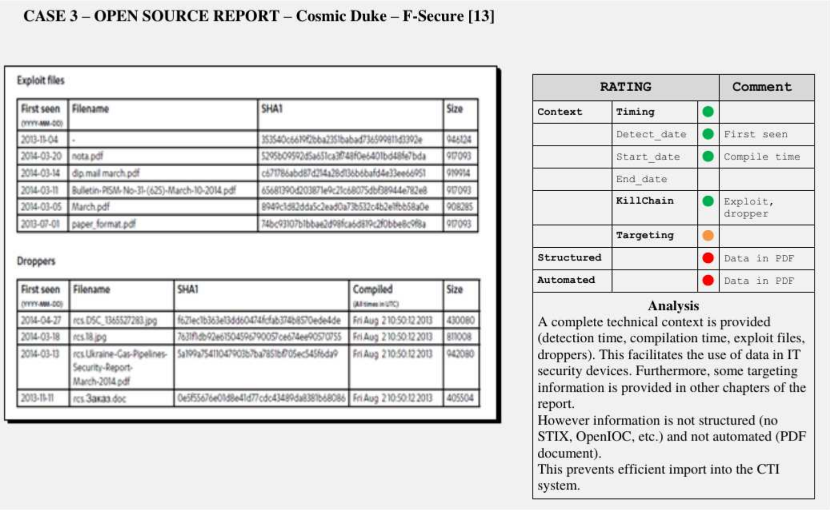

CASE 3 – OPEN SOURCE REPORT – Cosmic Duke – F-Secure [13]

Exploit files

First seen (YYYY-MM-DD)	Filename	SHA1	Size
2013-11-04	-	353540c6619f2bba2351babad736599811d3392e	946124
2014-03-20	nota.pdf	5295b09592d5a651ca3f748f0e6401bd48fe7bda	917093
2014-03-14	dip.mail march.pdf	c671786abd87d214a28d136b6bafd4e33ee66951	919914
2014-03-11	Bulletin-PISM-No-31-(625)-March-10-2014.pdf	65681390d203871e9c21c68075dbf38944e782e8	917093
2014-03-05	March.pdf	8949c1d82dda5c2ead0a73b532c4b2e1fbb58a0e	908285
2013-07-01	paper_format.pdf	74bc93107b1bbae2d98fca6d819c2f0bbe8c9f8a	917093

Droppers

First seen (YYYY-MM-DD)	Filename	SHA1	Compiled (All times in UTC)	Size
2014-04-27	rcs.DSC_1365527283.jpg	f621ec1b363e13dd60474fcfab374b8570ede4de	Fri Aug 2 10:50:12 2013	430080
2014-03-18	rcs.18.jpg	7631f1db92e61504596790057ce674ee90570755	Fri Aug 2 10:50:12 2013	811008
2014-03-13	rcs.Ukraine-Gas-Pipelines-Security-Report-March-2014.pdf	5a199a75411047903b7ba7851bf705ec545f6da9	Fri Aug 2 10:50:12 2013	942080
2013-11-11	rcs.Заказ.doc	0e5f55676e01d8e41d77cdc43489da8381b68086	Fri Aug 2 10:50:12 2013	405504

RATING			Comment
Context	Timing	🟢	
	Detect_date	🟢	First seen
	Start_date	🟢	Compile time
	End_date		
	KillChain	🟢	Exploit, dropper
	Targeting	🟠	
Structured		🔴	Data in PDF
Automated		🔴	Data in PDF

Analysis

A complete technical context is provided (detection time, compilation time, exploit files, droppers). This facilitates the use of data in IT security devices. Furthermore, some targeting information is provided in other chapters of the report.
However information is not structured (no STIX, OpenIOC, etc.) and not automated (PDF document).
This prevents efficient import into the CTI system.

Fig. 11. Case 3 — Open Source Report

control, depending on the age, they will recommend use of data for detection.

— **Narrow investigations** — Forensics investigations and logs management require time tagging of threat data to narrow research.

— **Counter-measures time window** — Time tags help setting time windows for blocking or monitoring hostnames / domains / IP addresses. Time tags are here important to reduce false positives and prevent blocking legitimate sites where not necessary.

— **Help prioritising** — Infrastructures, intrusion vectors and malicious payloads change all the time. Time tagging of data is essential to make sure that detection and prevention measures will focus on the most recent variants of malware and techniques/tactics/procedures.

Timing parameters

— **Detect date** — when threat data was first seen.

— **Start date** — when threat data was created or started to be malicious.

— **End date (optional)** — when threat data stopped being malicious.

Examples of time tags are given in Table 3.

Date	**Detect date**	**Start date**	**End date**
Definition	Time of detection of the indicator / observable	Start of malicious activities related to the indicator / observable	End of malicious activities related to the indicator / observable
Example types	filename, hash	domain, hostname, url	domain, hostname, url
Example time	Compilation time	Registration date Infection date (watering hole)	Infected legitimate domain being cleaned-up Malicious domain going offline or stopped to be used
Status	Mandatory	Mandatory	Optional

Table 3. Time tag examples

5.2 Targeting and Sighting

The purpose of targeting and sighting is to indicate **where** (location or sector) the threat instance was detected. Different CTI fusion nodes might each detect discrete instances of the same attack (e.g. points of malware delivery). Assembling the pieces of the puzzle supposes that each piece is ‘geo-tagged’ or ‘sector-tagged’. Once assembled, the network of CTI fusion nodes obtain a better threat targeting picture and hence can measure the danger for a given organisation, industry sector or country. Here are two fundamental activities that can be supported by this kind of context data:

- **Threat proximity metrics** — Threats are global and attacks can either target specific sectors and countries or be opportunistic. In either case, given the volume of malicious activities, organisations can strongly benefit from a proximity indicator: how close to me, my sector, my country, or my supply chain is this threat. This can support triage and prioritisation. Indeed, when proactive researches require resources, it is critical to prioritise handling of threats in sensors.
- **Support CTI fusion** — CTI fusion and analysis require information on where the threat was detected and what it was targeting. This supports objectives, such as threat actor profiling (what are the motivations and typical sectors targeted by a given group) or campaign scoping (which sectors or countries are being targeted). For a given malware, collecting several indicators with targeting / sighting data helps understanding who has been targeted and to which partners alerts should be released.

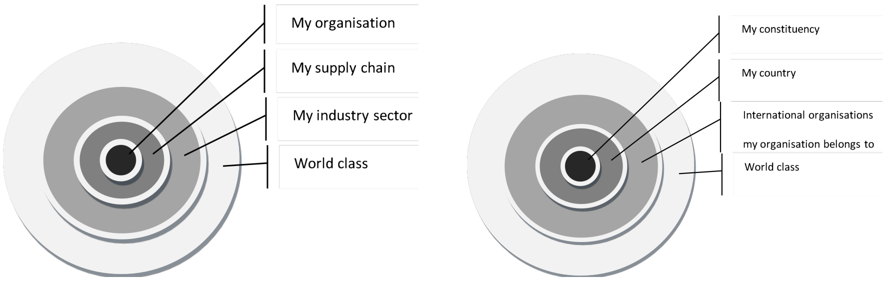

(a) Private organisation perspective (b) National CERT perspective

Fig. 12. Targeting and prioritisation (examples)

The main obstacle to sharing targeting / sighting data is the preservation of interests of the victim (security, reputation, competitiveness...). The model described here introduces a solution between two extreme practices which are equally undesirable:

- **Minimalist** — No sharing of targeting / sighting information — This practice undermines data fusion and action within the CTI network and should be avoided.
- **Maximalist** — Sharing all about the identity of the target — This practice may adversely impact the victim and is usually avoided. This is however possible in certain circumstances (e.g. when the victim decides to publicised at attack).

Targeting	Level	Description	Status
Geo-location	Continent	Continent(s) where a threat instance or indicator has been detected. E.g. North America, Latin America, Europe, Middle East, Africa, Oceania, Asia	Mandatory
	Country	Country(ies) where a threat instance or indicator has been detected.	Whenever possible
	Organisation	Organisation(s) where a threat instance or indicator has been detected.	Optional
Sector		’Aerospace’, ’Banking’, ’Biomedical’, ’Chemical’, ’Defense’, ’Diplomacy’, ’Education’, ’Electricity’, ’Electronic’, ’Energy’, ’Government-Administration’, etc.	Whenever possible

Table 4. Targeting and sighting

— **Intermediate** — The intermediate option is to specify the continent, country and/or sector if the targeted organisation cannot be revealed. This should be the preferred option in most cases.

Targeting / Sighting parameters (see Table 4) :

- **Geo-location**
  - **Continent level**, and/or
  - **Country level**, and/or
  - **Organisation level**

  and/or
- **Sector**

When several instances of the same threat have been detected, multiple targeting / sighting values are possible.

5.3 Kill chain

Kill chain parameter enables to understand **how** the threat materialises. This is necessary to understand the position of threat data in the intrusion kill chain. Examples: Was this IP address used to perform some recon/scan activities or is it a command and control server? Is the hostname a legitimate website temporarily infected to be a redirect / watering hole or is it where a malicious payload is being delivered? In the CTI network and at sensor level, the objectives of the kill chain context are:

- **Understand techniques / tactics / procedures** — The position of the indicator in the kill chain (e.g. delivery, installation, CnC, etc.) allows to understand the intrusion sequence, and therefore the techniques used by the attacker.
- **Act with threat-data** — The kill chain data enables the defender to make the right decision in terms of how to best use the data in its IT security devices (monitoring, blocking, scanning, researching in logs, etc.). It is essential to avoid inappropriate action for the defender community (e.g. make the adversary aware that you know) and prevent adverse business impact for the defender (e.g. blocking a legitimate website that is temporarily infected).

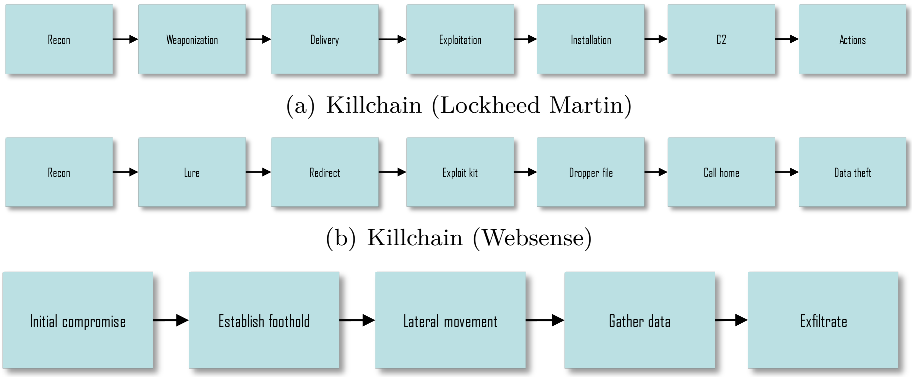

(a) Killchain (Lockheed Martin)

(b) Killchain (Websense)

(c) APT lifecycle (Hacker Intell Initiative)

There are different kill chain models, the most popular being the one developed by Lockheed Martin (reference [1]). Detailed explanations of these models are already widely available. Figures 13(a), 13(b), and 13(c) are therefore just a summary representation.

5.4 Scoping

The purpose of scoping threat data is to understand **what** kind of threat we are dealing with. This aspect is the most complex element of contextualisation. Different approaches are possible, from some basic scoping indication to the most complete and complex description. The STIX model offers a complete tool box in order to describe extensively the threat context.

It is up to each organisation or information sharing group to decide what level of contextualisation it wants or is able to share. In section 3 of this paper, the concept of multi-layered information sharing (technical, tactical, strategic) supposes that the description of the “what” can be minimalist for technical information exchanges, while it should be more developed for tactical and strategic exchanges.

This aspect of the contextualisation is mentioned here as background information only. We will not enter into further detailed specifications here.

5.5 Minimal and extended context

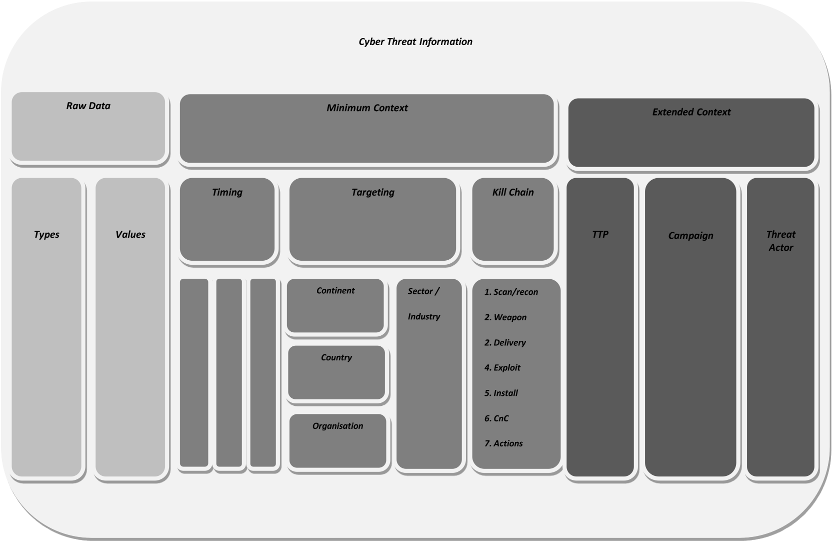

Fig. 13. Minimum and extended context

In a nutshell, the model that has been described here implies that within the CTI network, technical information must always be exchanged with a minimum context in terms of timing, targeting/sighting and kill chain.

For those organisations or communities which are mature enough in terms of CTI, extended context information may be shared in order to provide for better tactical and strategic threat situation awareness.

Figure 13 summarizes this minimum vs extended context data.

6 Actionable data

Information is actionable when it can be routed directly as useful feeds to IT security devices on the consuming side of the sensor-to-sensor model. Typical functional families of CTI consuming tools or activities include:

- **Detection tools** — Network-based intrusion detection systems (N-IDS), host-based intrusion detection systems (H-IDS), etc.
- **Prevention tools** — Network-based intrusion prevention systems (H-IPS), firewalls, proxies, mailguards, patch or vulnerability management systems, etc.
- **Investigation tools** — Security incident and event management (SIEM) systems, log analysers, host or network scanners, etc.
- **Intelligence and awareness** — Analysis of sectorial or geographical targeting, motivation of threat actors, specific techniques, tactics and procedures, etc.

Organisations on the consuming edge of CTI networks typically face constraints. They must be taken into account so that minimum performances for CTI fusion nodes can be identified. Consuming organisations have limited resources, they should be able to triage and prioritise threat data, especially when the volume of incoming data becomes too large and some manual handling is required. Consuming organisations use specific security tools and should be able to decide which set of data are most appropriate for the tools they use. Organisations adopt diverse security policies (more or less protectionist), and should be able to decide on the most appropriate strategies to take regarding threats.

Typical objectives for consuming organisations:

- **Prioritization** — Consuming organisations should receive sufficient information in order to be able to make decisions on the priority of threats. Indeed, for more and more advanced persistent threats, if only fully automated IT security devices are used, there is only a limited chance that infections are detected. Proactive research involving human judgement is often required against advanced threats.
- **Time-To-Live** — Threats are dynamic. Some indicators are actually useful to detect threats only for a limited period of time. Time-to-live consist in focusing on threat indicators that are actually alive.
- **Automation** — Information should be directly usable in tools in use by the consuming organisation. Indeed, dynamic techniques

implemented by attackers require fast reaction and automation in the defense chain.

- **False-positive reduction** — Use of threat data should not generate (too many) false positives, as this creates confusion on the side of consuming organisations.
- **Exploitation mistakes reduction** — The risk of threat data consumption generating adverse impact on organisations should be minimised (e.g. blocking legitimate websites, damage to the reputation of the organisation, etc.).

Figure 14 illustrates how context information helps prioritisation, decision and acting regarding threat data.

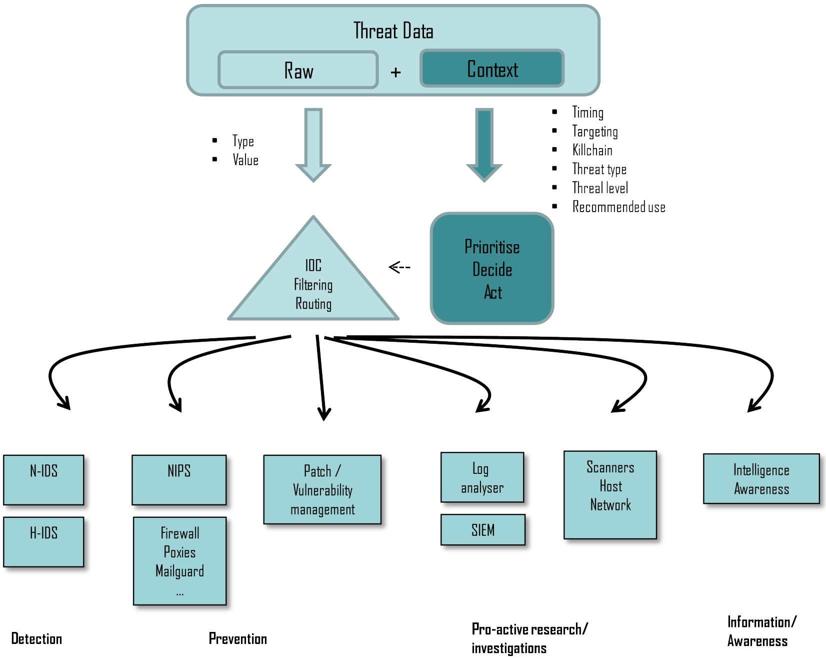

Fig. 14. Threat data on the consuming side

Table 5 summarises how context information can help consuming organisations in terms of prioritisation, automation, and reduction of false positives and mistakes.

Context	Level	Usage	Rationale
Timing	WHEN	• Prioritisation • Time To Live (TTL) • False positive reduction	• Discriminate recent attacks from old ones • Don't use outdated data • Indicate when a legitimate website stopped being infected
Targeting	WHERE	• Prioritisation	• Proximity metrics—discriminate "close" and "far-away" attacks
Kill chain	HOW	• Automation • Exploitation mistakes reduction • False positive reduction	• Kill chain indicates how to use the observables so that they can be routed directly in the right consuming sensor. • Organisation understands how the observable was used by the adversary and makes right handling decision (e.g. avoid blocking legitimate website)
Threat type	WHAT	• Prioritisation	• Organisation decides to handle a threat first depending on its type
Threat level	WHAT	• Prioritisation	• Organisation decides to handle a threat first depending on its level

Table 5. Contextualisation and consuming

7 Conclusion

Developing threat information sharing and automated handling are two priorities in a very dynamic threat landscape. A global cyber-threat intelligence sharing network is taking shape. It enables end-to-end, sensor-to-sensor threat information work-flows. The nodes constituting the network realise cyber-threat intelligence fusion, context enrichment and further sharing. This helps organisations to better respond to threats. The present paper demonstrates why contextualisation of information shared in the network is vital for drawing an accurate threat situation picture and supporting organisations on the consuming end.

A high-level functional model for a CTI fusion node has been introduced in this paper. A first implementation of such a CTI fusion node has been developed and is being used within CERT-EU [1]. Some of the key concepts exposed in this paper have hence been tested and validated. This should allow further refinement and complement of the model in the future. The aim is to help the development of automated information sharing networks and facilitate the integration of new participants. It is expected that the barrier to join automated information exchanges will be lowered for many organisations

Additional research will have to be performed and more experience needs to be gained to increase the performance of local and global information sharing networks. It will be necessary that as many organisations as possible will participate actively in this effort.

References

1. Eric M. Hutchins, Michael J. Cloppert, Rohan M. Amin, Ph. D. – Lockheed Martin Corporation. Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains.
2. Sergio Caltagirone, Andrew Pendergast, Christopher Betz. The Diamond Model of Intrusion Analysis.
3. STIX$^{\text{TM}}$. Structured Threat Information eXpression. `https://stix.mitre.org`
4. TAXII$^{\text{TM}}$. Trusted Automated eXchange of Indicator Information. `http://taxii.mitre.org`
5. Wikipedia. Intelligence source and information reliability `http://en.wikipedia.org/wiki/Intelligence_source_and_information_reliability`
6. NIST Special Publication 800-150 (Draft). Guide to Cyber Threat Information Sharing (Draft)

[1] CERT-EU is the computer emergency response team for the EU institutions, bodies and agencies. See `http://cert.europa.eu`

7. 2014 Threat Report. M Trends. Beyond the Breach (Mandiant).

8. 2014 Threat Report. Websense.

9. CrowdStrike Global Threat Report. 2013 Year in Review.

10. McAfee. The Economic Impact of Cybercrime and Cyber Espionage. July 2013.

11. Imperva. The Non-Advanced Persistent Threat. `http://www.imperva.com/docs/HII_The_Non-Advanced_Persistent_Threat.pdf`

12. Kaspersky Lab. MINIDUKE is back. July 2014. `https://securelist.com/blog/incidents/64107/miniduke-is-back-nemesis-gemina-and-the-botgen-studio`

13. F-Secure. COSMICDUKE — Cosmu with a twist of MiniDuke — TLP:WHITE. July 2014