National Cyber Security Centre
Ministry of Security and Justice

Indicators of Compromise

Effectively apply threat information

Factsheet FS-2016-02 | version 1.0 | 1 June 2017

If you are responsible for securing the network of your organisation, you will often hear the term: IoC, or Indicator of Compromise. In short, an IoC is an indicator that makes it possible to detect the presence of a specific threat within your network.

When receiving an IoC, a great number of organisations wonder what they should do with this information. How do I process an IoC? What will I find? And what should I do when it turns out my own organisation was hit as well? These are all valid questions and will be answered in this factsheet.

Background

Incidents within a network often go unnoticed for a long time. Several malware campaigns like *Carbanak*[1] and *SYNful Knock*[2] show that attacks can sometimes go undetected for years. All this time, attackers are able to extract sensitive information from networks of compromised organisations without those organisations being aware.

Target audience

Information security professionals

This factsheet was written in collaboration with:

The Dutch Tax Authority
ING Bank NV

[1] https://securelist.com/blog/research/68732/the-great-bank-robbery-the-carbanak-apt/
[2] https://www.fireeye.com/blog/threat-research/2015/09/synful_knock_-_acis.html

Using Indicators of Compromise, insights in your own incidents become shareable with other organisations. Over the past few years, organisations have found that detecting digital compromise can be difficult. An incident at one organisation can be only one of multiple, similar incidents at other organisations. Information regarding an incident at one organisation can lead to detection and possibly prevention within other organisations.

Indicators of Compromise

An Indicator of Compromise (IoC) is information that can help with identifying specific malicious behaviour on a system or within a network.[3]

In practice, IoCs are often IP addresses or domain names. If these IP addresses or domain names are encountered within a network, this may indicate an infection. For example, an organisation determines that they are a victim of an attack where the attackers have uploaded sensitive information to an external server. To help other organisations to detect a similar attack, the victim organisation shares the properties of the attack (IoCs) known thus far, such as the IP address and domain name of the external server. Based on this information, other organisations can investigate if there are similar, malicious activities taking place within their own networks. If this is the case, this might be an indication that these other organisations are dealing with the same attack.

During an investigation, multiple IoCs related to an incident can become available. With these IoCs, organisations can extend the characteristics further so that they get a better picture of the incident. As an example, after some investigation it may be discovered that an attack started with an email with a specific subject, a known vulnerability was exploited, specific files were used and a specific, repeating pattern was seen in communication with an external server. Based on this information, the set of IoCs for this attack might be as follows:

```
IP address: 192.168.23.29
Domain: cc.evilserver.domain
Email subject: "Something interesting!"
Vulnerability: CVE-4242-4242
MD5 hash file:
24b6dc71b26debdfb8c7cfbd3a55abbd
URL pattern: abc\d\d\d(.*?)\.php
```

With IoCs, organisations have more possibilities to detect similar attacks.

[3] There are standards where the definition of an IoC differs or where it has a different name.

The importance of contextual information

It is important to provide as much contextual information as possible to IoCs. The following questions can help with acquiring relevant contextual information:

- **Who** was hit by this attack? Did the attack target a specific organisation, a specific sector or were multiple organisations/sectors hit?
- **Who** is behind this attack, and what is the sophistication level of this attacker? For example, is the attacker a script kiddie or a state sponsored actor?
- **What** happened and **what** is the damage done? Was the attack aimed at making a system temporarily unavailable (Denial-of-Service) or were large amounts of sensitive documents exfiltrated?
- **Where** in the network did the attack take place? Was the attack aimed at, for example, workstations within the organisation or did the attack target ICS/SCADA systems?
- **When** did the attack take place? This information helps searching in logs and can also help answer if there is indeed a problem. For example, the domain name of a legitimate website that spread malware for five minutes might be a valid IoC, but only hits within these five minutes are relevant.
- **Why** did this attack take place? The why question can help determine how likely it is that other organisations are also hit.

Always provide as much contextual information as possible with an IoC. In practise, this will often be limited. Sometimes organisations are hesitant to share information, because they might not want to share details surrounding an attack or the fact that they are a victim.

Confidentiality of information

Organisations that trust each other tend to share information with each other more easily. When an organisation shares information and does not want this information to be spread further by the recipients, then this organisation must be able to trust the recipients with regard to how they treat this information.

With the Traffic Light Protocol (TLP), the supplier of the information can indicate with whom the recipient(s) may share this information. When you receive information with a TLP classification, you will know what you may and may not do with this information. The term Traffic Light refers to the different colours of a traffic light. An organisation can share information under TLP Red, TLP Amber, TLP Green or TLP White. The following table has a short summary of the TLP categories.

TLP categories

TLP-RED	"For your eyes only". Only to be used by you and not to be spread to other people, even within your own organisation.
TLP-AMBER	To be used and shared with co-workers within your organisation on a need-to-know basis and with clients or customers who need to know this information to protect themselves or prevent further damage.[4]
TLP-GREEN	Used for information that is not very sensitive and can be shared with partners and peers, but not via publicly accessible channels (e.g. websites).
TLP-WHITE	Public information that can be shared freely, taking into account standard copyright rules.

When making your own IoCs available, decide if and how much further dissemination you want to allow and choose the TLP category that matches. Keep in mind that when you classify information as TLP Red, the limitations are so strict that the question arises whether it is still useful to share this information in the first place. When processing IoCs you receive from others, respect the TLP category that the sender has given this information.

When you have outsourced (a part of) your infrastructure to a third party, keep in mind that possible limitations regarding the TLP category might mean that you cannot always share IoCs that you receive with this third party. When in doubt, we advise you to ask the party that shared the information with you.

Applying IoCs

To apply IoCs within your organisation, you will need, depending on the size of your organisation, one or more people with knowledge of the subject matter. You will also need to have logging enabled on central systems within your organisation. You can use specific tooling to search for hits on IoCs within your organisation.

How to implement this within your organisation is out of scope for this factsheet.[5]

Often there are systems within an organisation that can help in the search for hits on IoCs. For inspiration:

- **Proxy servers** register the websites that users visit intentionally or unintentionally. Domain names and URLs can be found in the logs of these systems.
- **DNS servers** answer DNS requests that systems within the organisation perform. Logging on DNS servers is essential when looking for malicious IP addresses, domain names and DNS servers.
- **Mail servers** are used for receiving or sending email messages. You can use logging on your mail servers to see if your organisation received specific malicious email messages by searching for specific subjects, attachments or senders.
- **Firewalls** monitor all kinds of network flows within the network and can allow or block traffic based on rules. Advanced firewalls also look at other traffic characteristics besides IP addresses and port numbers. Logging of firewalls can therefore be of great value.
- **Intrusion Detection Systems** (IDS) and **Intrusion Prevention Systems** (IPS) are meant to detect or block attacks on a network. It is useful to see if an IDS or IPS has seen an IoC before. If not, it is wise to add detection rules to these systems so that this IoC will be detected or blocked in the future.
- **Antivirus software** is aimed at the individual systems within an organisation. Such software monitors files and processes on a system. By supplying antivirus software with information on malicious files and processes, it is possible to detect the presence of such files and processes on the various systems within an organisation.
- **Security Information and Event Management** (SIEM) is a solution that is especially suited for this task because as a central system, it contains logging from a variety of systems and applications.

Start with IoCs that can be deployed on systems where information on (parts of) the internal network passes through or is stored. Examples are SIEM solutions, mail servers or proxy servers. This way, an IoC can be deployed quickly to monitor for many different systems within the network. Sometimes, the only available IoCs are those with which individual systems can be investigated. In these cases, you can often use contextual information of the IoCs to deduct for which type of system this

[4] The TLP-Amber definition has changed recently. The old definition allowed for information to be shared only within your own organisation. The new definition allows for sharing, when necessary, with clients and customers. To prevent breaches of trust, ask the sender of the information which TLP-Amber definition is being used if unsure. Also see: https://www.first.org/tlp.

[5] For more information on implementing detection solutions, see our whitepaper Handreiking voor implementatie van detectie-oplossingen (Dutch): https://www.ncsc.nl/actueel/whitepapers/handreiking-voor-implementatie-van-detectie-oplossingen.html.

IoC can be used (workstations, mail server, web server, etc). This often narrows the search space significantly.

An IoC generated a hit, now what?

When one or more IoCs generate a hit, you will have to determine where you want to take action. To determine whether a hit requires action, you will have to look deeper to gain a clear situational picture.

Take a close look at the IoC that generated a hit. A hit generated by an IoC does not always mean that malicious activity has taken place. It is also possible that this hit is a false positive. A false positive is a hit on an IoC where the behaviour that triggered the hit is not malicious.

Some types of IoC are more sensitive to false positives than others. A good example is an IP address. Sometimes an IP address is used exclusively for malicious activities, but other times an IP address is used for shared webhosting. In the latter case, it is possible that a single IP address houses multiple websites where only one website is malicious. When you see a hit on this IP address, it does not automatically mean that malicious activities took place. It can also be a connection to one of the non-malicious websites on this IP address. On the other hand, a hit on a hash value of a malicious file has a much lower chance of being a false positive. For this reason, it might be necessary to obtain additional contextual information, for example by obtaining information from the DNS servers indicating which domain name was queried that lead to a hit on an IoC of an IP address.

If you encounter a hit on an IoC, find out which system within the network generated this hit. Where you will have to search depends on the type of IoC that generated a hit. As soon as you know which system within the network generated the hit, you can take action, such as isolating the system from the network. You can also consider taking additional actions, such as forensic investigation. The actions will differ per organisation and per case. For example, an infection of a visitor's system that is connected to the public Wi-Fi network will be of less importance than an infection on the internal mail server, and therefore will most likely call for different actions to be taken.

How can I create an IoC?

First, assess which information you have on the incident.

Choose a starting point for the gathering of information, as closely as possible to the source of the infection. If you only have metadata as information (subject of an email, a pattern in a requested URL, etc), you will have to conduct searches within systems to obtain the actual source of the infection within your organisation.

If the information you have is file-based, obtain unique characteristics of these files that you can apply and share. When you encounter malicious files, you can verify this, for example, by letting your antivirus software scan the files. In most cases, the antivirus software will recognize the files as malicious. For files, there are several characteristics that you can use and share as an IoC. For inspiration:

- The **hash value** of the file (MD5/SHA1/SHA-256). This hash value is characteristic for the file so that other parties can use this hash value to detect the same malicious file within their organisation.
- The **location** and **name** of the file. Often malware copies itself to a specific location on the system and renames the copy of itself so it can start again when a restart of the system takes place.
- **Distinctive patterns** within a malicious file. Often, different variants of the same malicious file are being used by attackers. These files are in essence all the same malware. Each file differs in small points from the other files. Attackers do this to prevent detection based, for example, on hash values of the files. These files often share the same patterns in their contents. With specific tooling, such as Yara[6], it is possible to write rules to detect malicious files based on patterns in the contents of the files.
- **Specific registry keys** that the malware creates or queries. Some Windows malware will want to make changes in the registry, for example, to change security settings on infected systems or to configure the system so that the malware is executed at every restart of the system. These registry keys can be unique enough to serve as a valid IoC.

If the information you have is network-based, obtain unique characteristics of the malicious network traffic. For example, you can search through the logs of the proxy server or DNS server for traces of malicious activities. You can also execute the malware in a controlled environment to monitor network traffic. This might result in sufficient information from which to create IoCs. At the network-level, there are several characteristics that you can use and share as an IoC. For inspiration:

- **Domain names**, **IP addresses** or **URLs** to which the malware connects or from where the malware is downloaded.
- The **User Agent HTTP header** that is used by some types of malware when making an HTTP request. Attackers do this to pretend to be a browser to try to prevent standing out from the legitimate network traffic within an organisation.

[6] https://plusvic.github.io/yara/

Sometimes, this User Agent header is so unique that it differs from the User Agent header that legitimate browsers use. This can be a good IoC to monitor within the organisation.

- **Distinctive patterns** in network traffic. Just like files, network traffic can also contain patterns that can be used to create a good IoC. Many malware families communicate in such a unique manner with their Command & Control servers that this is an excellent way to monitor for malicious activities within the network. To do this, specific tooling with the accompanying rulesets is required, such as Snort[7], Suricata[8] or Bro[9].

If you have created your own IoCs, make sure that these IoCs do not generate false positives. If you use filenames as IoCs, check non-infected systems to see whether these files are indeed not present on these systems. If you are using domain names or IP addresses as IoCs, check the logging of your systems. It is important to be sure that these IoCs only result in hits when there are actual infections.

Final note

Processing IoCs within your organisation is a process for which you need to implement specific workflows and tooling. Furthermore, it will need an investment of time. This can be a pre-existing process within your organisation, or it can be a completely new process in the information security within your organisation. It is a process that helps with, but does not guarantee, the detection of infections.

When your organisation has adopted this process to a sufficient degree, it can be a powerful method of protecting yourself against different threats. It also helps to share information on threats with other organisations, by which you might help prevent incidents at other organisations.

[7] https://www.snort.org/
[8] http://suricata-ids.org/
[9] https://www.bro.org/

Publication
National Cyber Security Centre (NCSC)
P.O. Box 117, 2501 CC The Hague
Turfmarkt 147, 2511 DP The Hague
+31 (70) 751 5555

More information
www.ncsc.nl
info@ncsc.nl
@ncsc_nl

FS-2017-01 | version 1.0 | 1 June 2017