Cybersecurity Automation and Threat Intelligence Sharing Best Practices

April 2021

JOHNS HOPKINS APPLIED PHYSICS LABORATORY

APPLYING "LOW-REGRET" METHODOLOGY FOR CYBER THREAT INTELLIGENCE TRIAGE

Rapidly Sharing Actionable Intelligence for Network Defense

Charles Frick

Many organizations choose to participate in receiving shared Cyber Threat Intelligence (CTI) via subscribing to various feeds of information that include many forms of data, including Indicators of Compromise (IOCs). This information is often focused on a particular community through participation in an Information Sharing and Analysis Center or Organization (ISAC/ISAO).

Unfortunately, there is a significant challenge in finding actionable IOCs within these streams of data that provide any benefit towards network defense, leaving much of these data unused or not used until it is no longer valuable, as cyber attackers retire their use of specific IOCs rapidly.

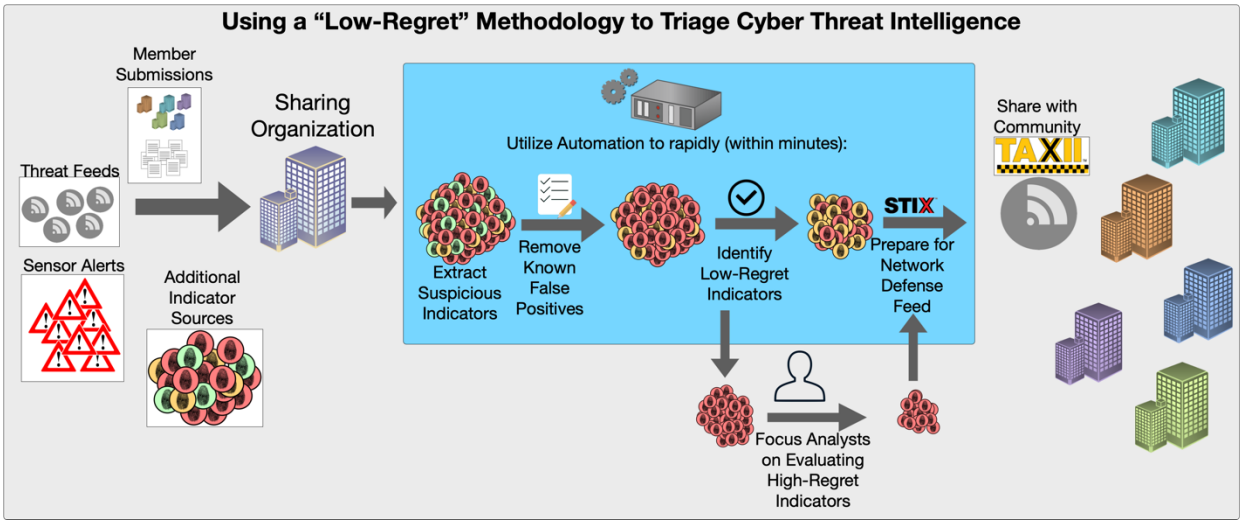

Figure 1 Process for Application of "low-regret" methodology for CTI triage

Through multiple research and pilot efforts, the Johns Hopkins University Applied Physics Laboratory (APL) has successfully deployed threat feeds that within minutes of receipt extract, identify, and share actionable IOCs to a sharing organization such as an ISAC or ISAO. Figure 1 provides a high-level visualization of this process. In this paper, the methodology and process are provided in more detail to help other organizations leverage these capabilities for their communities' network defense needs.

"Low-regret" methodology

What does it mean to employ a "low-regret" methodology towards network defense? In short, it means to use a benefit vs. regret assessment to make decisions about implementing automated actions. This leads organizations to focus on ***when*** to take an action in an automated manner ***instead of whether*** the action should be automated. With respect to automated responses based on cyber threat intelligence, the definition of regret can be simply defined as:

- **"Low-Regret":** Taking automated action against this intelligence is extremely unlikely to disrupt operations, regardless of whether or not the intelligence assessment is correct.
- **"High-Regret":** Taking automated action against this intelligence may have impact to operations.

More detail on the "low-regret" methodology is freely available via the APL GitHub page: https://github.com/JHUAPL/Low-Regret-Methodology.

Applying a "low-regret" methodology to triage threat intelligence

Applying the "low-regret" methodology to CTI triage, as depicted in this paper, revolves around the ISAC/ISAO concept. Malicious cyber campaigns, such as ransomware, often target a specific industry or community within an industrial sector. This section of the paper provides details on how a sharing organization such as an ISAC/ISAO can use the methodology to develop automation to rapidly identify and share "low-regret" IOCs to their community.

Extract suspicious indicators

Information sharing organizations receive a large quantity of IOCs from many sources (other threat feeds, system alerts, member submissions, etc.). The key first step to the process is to extract the suspicious IOCs from the deluge of intelligence received daily. This is not merely regular expression (REGEX) extraction of everything that looks like an indicator. Any information sharing organization should utilize a process for identifying potentially malicious indicators based on the context in which they were received. This can be through tools such as malicious signatures, tags, or other processes utilized within the sharing organization to identify data that have the appearance of being related to malicious cyber activity. Without this step, the amount of data to parse is simply too vast to provide actionable intelligence to a community that is under constant cyber attack.

The central tenet within the application of "low-regret" towards this extraction is found in the words ***"potentially malicious indicators."*** Many information sharing mechanisms fail to provide data in an actionable timeframe due to a desire to irrefutably prove that an

IOC is malicious before sharing it. Unfortunately, the time to make this determination is often greater than the time the IOC is in active use by a cyber attacker. Thus, the process for identifying potentially malicious indicators must be automated and utilize repeatable, codified steps in making the determination of potential. It does not mean to ignore all other data, as this is a key function of the additional intelligence processing capabilities within an ISAC/ISAO.

Remove known false positives

Signatures are not perfect and there is always the potential for inconsistencies in the source information. Certain threats target popular or business-critical sites, which means that certain IOCs associated with malicious behavior may actually be high-regret. Because the initial extraction is fully automated, there is a need to implement automation to filter out potentially malicious IOCs that would most likely impact operations if they were automatically blocked. These are considered “high-regret” and must be assessed by other means.

This is where the power of a community-focused sharing organization such as an ISAC/ISAO can significantly improve the process. Some IOCs related to malicious activity are easily identifiable as things that would impact operations such as Internet search providers, popular Domain Name Servers, etc. However, an ISAC/ISAO can also maintain an “allow list” for domains, IP addresses, etc. that are known to be critical resources for their community, even if their impact to that community is not well-known globally.

Identify “low-regret” indicators

Once the known false positives have been removed, automation is ready to apply scoring to IOCs in order to identify those that are highly unlikely to impact operations based on criteria that have been defined according to the organization’s policies and risk tolerance. The key to this step is to understand attributes that malicious IOCs of this type tend to have in common, but are not shared by authorized IOCs. After those attributes are identified, the next step is to figure out where and how to access the information about that attribute to make this determination for an IOC. That information is next made accessible by automation, then the low regret score can be assessed via several rapid queries based on indicator type. These queries may include but are not limited to:

- Domain age: A newly registered domain is less likely to be a critical asset
- Number of mapped domains: IP IOCs that historically resolve to one or two domains are less likely to migrate towards critical assets across the Internet even when they reside on shared infrastructure

- Known malicious behavior: Malicious files often have particular traits flagged by analysis software that are not shared by legitimate files.
- Analyst Vetted IOCs: If analysts at the sharing organization or a trusted source of threat bulletins have provided high confidence that an IOC is malicious and that IOC appears in the process, it should be flagged as "low-regret" as it is highly likely to be malicious

IOCs that fail to meet these checks are not to be ignored. Rather, they are the (now much smaller) pool of intelligence that the threat intelligence analysts will study. If the analysts determine an IOC does threaten malicious activity, the IOC can be sent through the triage process again with a tag of "analyst-vetted."

Prepare indicators for a network defense feed

Once the automation has identified the "low-regret" IOCs, it can then rapidly transcribe each IOC into a shareable format, such as the Structured Threat Information eXpression (STIX) standard. The information used to determine the IOC as "low-regret" should also be included in the machine readable data object for the IOC so that receivers do not need to repeat the steps conducted by the information sharing organization.

Share with community

Once the IOCs have been properly formatted, the automation can then share the data through machine-speed transfer mechanisms, such as the Trusted Automated eXchange of Intelligence Information (TAXII) protocol or other accepted machine-speed transfer mechanisms employed by the community. Regardless of the sharing mechanism employed, it is critical to ensure recipients receive all the relevant context with the IOCs as they are shared.[1]

Conclusion

The utilization of a "low-regret" methodology has the capability to extract operational value of much CTI that is currently ignored by network defense operations. It is not a panacea, it will not capture the insights derived from active analysis of CTI, but it can provide actionable data that community members can use in their security operations to disrupt malicious campaigns against their networks.

[1] Watson, K., "Cyber Threat Intelligence (CTI) Sharing Infrastructures", Feb 2021.

Acknowledgement

This material is based upon work supported by the U.S. Department of Homeland Security / Cybersecurity and Infrastructure Security Agency under Grant Award Number DHS-19-CISA-128-SLT-001 (State, Local, Tribal, and Territorial Indicators of Compromise Automation Pilot).

Disclaimer

The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security / Cybersecurity and Infrastructure Security Agency.

✉ Charles Frick
Charles.Frick@jhuapl.edu