Cybersecurity Automation and Threat Intelligence Sharing Best Practices

April 2021

JOHNS HOPKINS APPLIED PHYSICS LABORATORY

APPLYING "LOW-REGRET" METHODOLOGY FOR RESPONSE TO INDICATORS

Rapidly mitigating IOCs at scale

Charles Frick

Analysis and response to cyber Indicators of Compromise (IOCs) is so resource consuming that many cybersecurity teams do not even attempt to use them in operations. This paper showcases how to apply a "low-regret" methodology for rapid evaluation and response to these IOCs via Security Orchestration, Automation, and Response (SOAR) tools. Using this methodology depicted in Figure 1, organizations have been able to add IOC mitigation into security operations in a value-added and sustainable manner.

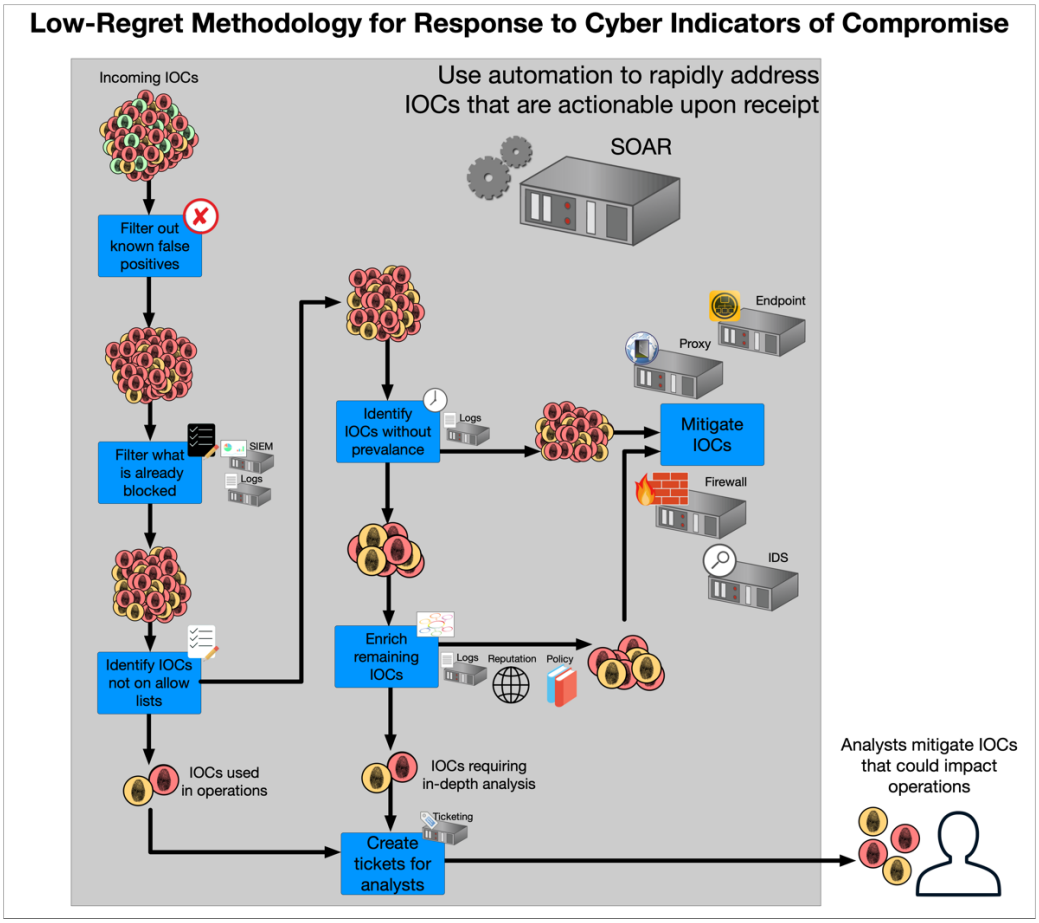

Figure 1 Sample process for applying Low-Regret to IOC Response

"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 Johns Hopkins University Applied Physics Laboratory (JHU/APL) GitHub page: https://github.com/JHUAPL/Low-Regret-Methodology.

Applying a "low-regret" methodology to indicator response

Many organizations remain hesitant to use automation to respond to IOCs due to concern over adversely impacting operations by blocking access to business-critical resources. This has often led to placing an analyst "in the loop" to review the IOC before deciding whether or not to block it. However, the decision logic for much of this process is rigidly defined and is repetitively applied, normally when an analyst "has the time" to address the IOCs amongst their other duties and tasks. By defining tailored, orchestrated automation workflows that account for organizational policies and risk tolerance, tools such as SOAR can process a majority of the IOCs in the background while placing the operator "on the loop" to review overall process performance and intervene when needed.

Automated filtering of indicators

The first step in automated processing of IOCs is filtering of the incoming IOCs to reduce the set to only those IOCs that are actionable. Through local policy rules, many organizations have defined traits for indicators that are indicative of potential malicious intent and automation can conduct a query against these traits to remove IOCs that would be considered false positives.

The next automated check is to determine if an IOC is already blocked by the security controls already in place. Many security tools receive automated updates for IOCs that have been determined to be "known bad." A human analyst should never have to

conduct a check against IOCs to see if they are already blocked. Those data are readily available via logs and API queries to the existing security tools.

Additionally, the "low-regret" methodology has great utility to an organization by utilizing an "allow list" of IP addresses, websites, domains, and files that are known to be used by operations. If an IOC flagging these resources is received and not determined to be a false positive, it is "high-regret" and should be removed from the automated process and sent to an analyst for review.

It is important to note that many of these steps can also be augmented through the participation in an Information Sharing and Analysis Center/Organization (ISAC/ISAO) that provides a network defense feed of indicators that employs "low-regret" methodologies for triage. In a related paper[1], JHU/APL provides detail on this application of the methodology.

Automated local enrichment

In order to account for organizational policies and risk tolerance, IOCs often require local enrichment in order to decide whether or not an IOC should be blocked. Automating this process often creates some concern as resources such as bandwidth and licensing can be heavily utilized when being applied automatically throughout the day.

One key aspect of "low-regret" implementations is to re-think the order of enrichment steps for IOCs. First, a query for network prevalence against the IOC can identify whether any tool or user within the organization has ever accessed a particular IP address, website, or file. If the IOC is not a false positive, and nothing in the organization has used it before, it is not likely to impact operations and can be automatically blocked. In previous pilot work,[2] JHU/APL has found up to 99% of incoming IOCs that are not false positives can meet this threshold.

For the remaining IOCs, automation can further conduct the basic queries against logs, reputation engines, and organizational policy to determine which IOCs can meet the threshold for blocking and extract the ones requiring additional analysis for manual review.

Automated mitigation

Once a set of "low-regret" IOCs has been identified, automated workflows can also execute the blocking of the IOCs through a variety of tools such as firewalls, proxies, or endpoint detection and response, provided that these tools properly support

[1] Frick, C. "Applying Low-Regret methodology for cyber threat intelligence triage", April 2021.
[2] Frick, Charles K. "SLTT Pilot Shareable Workflows." IACD, Integrated Adaptive Cyber Defense, Dec. 2020, www.iacdautomate.org/sltt-pilot-shareable-workflows.

automation.[3] When designing automated workflows for these mitigations, it is highly recommended to utilize a modular design that easily enables undoing of a mitigation if at a later date it is determined that access to a specific IOC-identified resource is now needed by operations. By maintaining an audit and review capability for the workflows, organizations can easily track the automated mitigations at scale on a regular basis.

Identifying tasks requiring human interaction and controls

The IOCs that do not meet the thresholds for automated mitigation, and thus require decisions and mitigations by a human analyst, comprise a significantly smaller subset of the incoming IOCs. By integrating the workflow with existing case management tools, such as ticketing, this process can be accelerated and easily tracked to ensure the organization's exposure to potential threat is reduced.

Conclusion

Manual review of large IOC sets causes significant impact for network defense operations. Employing the "low-regret" methodology to know when mitigations can occur faster, even before complete attribution to a threat, provides significant improvement in security posture while reducing workload on human analysts within a security team. Operational pilots have shown that organizations using this type of methodology can successfully integrate IOC mitigation into their security operations. Metrics have shown that it was more efficient for the organizations to spend analyst time reviewing the performance of automated mitigation and addressing any issues than it was to review each IOC prior to mitigation. This process will not serve as a panacea for security operations, but by allowing automation to augment security teams, an organization can greatly improve its ability to address the constantly increasing speed and scale of cyber threats.

[3] Watson, K. "Enabling automation in security operations – assessing automation potential of products and services", February 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