DFLABS
CYBER INCIDENTS UNDER CONTROL

Automate or Die without Dying

How to safely automate incident response without impacting operations

Oliver Rochford, VP of Security Evangelism

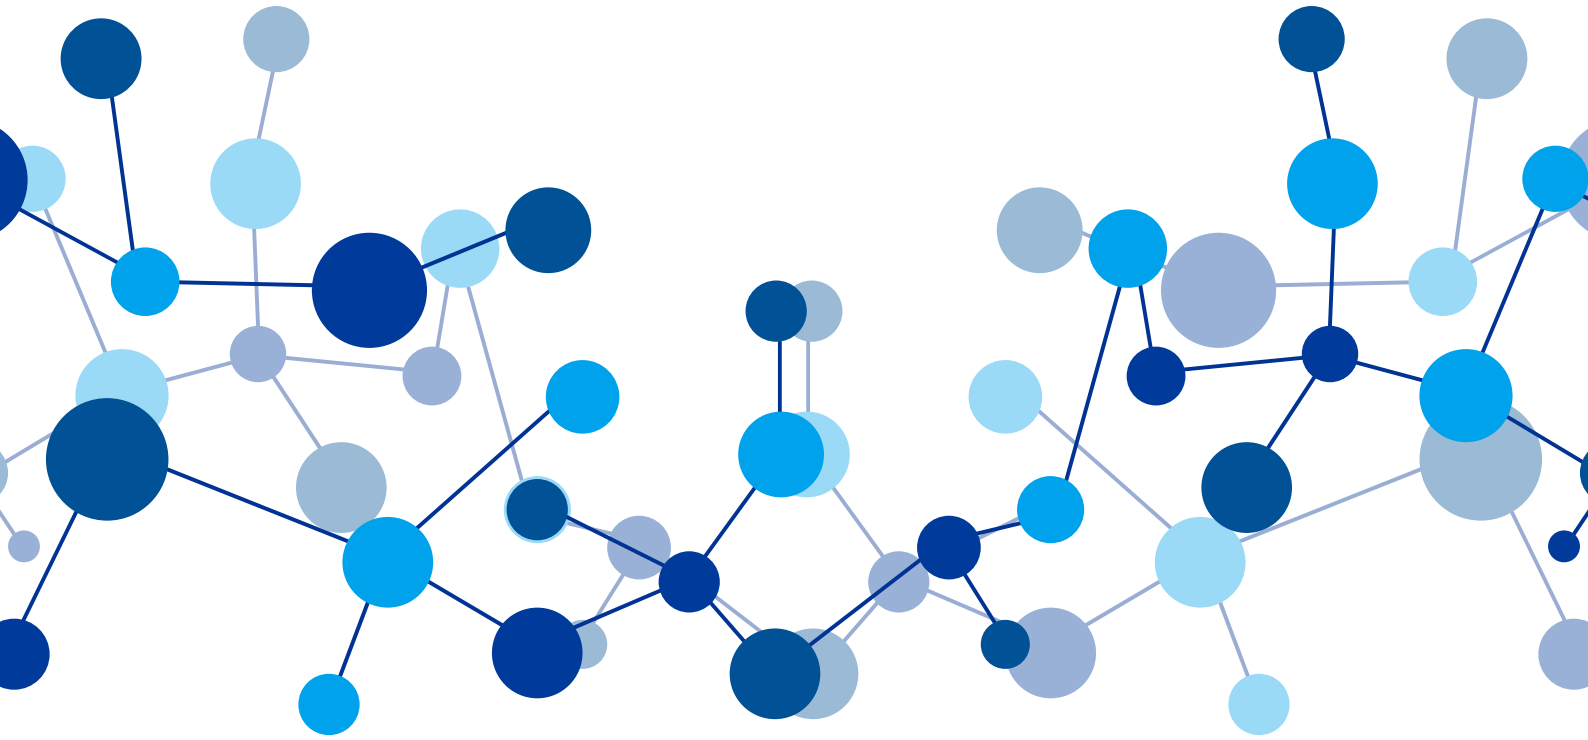

www.dflabs.com

Introduction

Automation in Security Operations and Incident Response is a topic that has recently experienced a lot of attention and is undergoing a renaissance. This is driven by the ongoing cyber security skills shortage, an escalation in the volume and sophistication of cyber threats, and the growing adoption of automation by threat actors themselves.

Manual processes cannot achieve the velocity required to effectively and rapidly respond to attacks that are often not detected until a threat actor has almost completed the cyber kill chain[1]. Attacks such as Ransomware or Phishing especially stress the effectiveness and efficiency of manual incident response, frequently resulting in Disaster Recovery rather than Threat Containment.

Gartner states that "The speed, accuracy and, therefore, effectiveness of response are also major factors to consider. Detection and mitigation must be executed swiftly and with precision to limit the damage from a breach and interrupt the kill chain before data exfiltration."[2]

The consistent feedback is that automation is highly desirable, at least by security teams. But this desire has been inhibited by doubt and fear: Doubt about the accuracy of the detection of threats, and fear of the consequences of automating the containment or mitigation responses and the prospect of detrimental impact and damage resulting from doing this wrong.

Enterprises have been promised automated attack containment capabilities before. Antispam and Intrusion Prevention Systems, and the chaos that these premature attempts at automating containment caused, based on too much faith in their ability to reliably identify anomalies and attacks, has meant that IT Operations and Executive Management are reluctant to pass such power to machines.

Many organizations own an IPS, but run it in non-blocking mode, demoting them to Intrusion Detection Systems. This trend is not isolated, with organizations that have automation capabilities built into existing technologies such as Security Information and Event Management, Endpoint Detection & Response and Security Automation & Orchestration solutions not trusting these to automate much beyond basic tasks such as sending out notifications or running a threat intelligence query.

This is despite detection capabilities having dramatically improved in recent years, especially using behavioural modelling and machine-learning driven approaches.

Solving the 3 Automation Challenges

There are 3 basic challenges that Security teams face when considering Automation. Each one of these challenges can be solved.

The SecOps team can assess the impact of the risk, but NOT the impact on production.

The SecOps team sitting in its ivory tower, focused purely on the risk and impact of the threat, will struggle to build up and maintain an awareness of what is going on in production.

- Is the affected system mission critical, is the system unstable, or is it a legacy system?
- Is the system currently being used to process the annual financial reports?
- Is a paying customer using it?

Disabling a seemingly innocuous user account may seem trivial – but that user account may be used to run critical processes. Dependencies, complexities and unknowns are the bane of automation. These are all data points that most security operations team either lack, or the information that pertains to this may be stale - but can have a huge impact on how the incident response or remediation process must be conducted. It's not that anyone is advocating that the incidents or the vulnerabilities do not have to be addressed at all – but this may require additional time or a specific way of approaching it.

"The speed, accuracy and, therefore, effectiveness of response are also major factors to consider. Detection and mitigation must be executed swiftly and with precision to limit the damage from a breach and interrupt the kill chain before data exfiltration."

Gartner

You can automate the actions, but not the decision

The 3 Security Automation Challenges

1.The SecOps team can assess the impact of the threat, but not the impact on production

2.You can automate the actions, but not the decision

3. IT Operations do not trust automation

The actual containment or remediation response is not the only thing that can be automated. We can automate a wide variety of tasks, including prioritizing an incident, fetching additionally required information and context or notifying stakeholders. Using automation, we can make our people more efficient and we can automate away some of the more menial tasks. We may even be able to use machine learning to compute an analysis that a human would take millennia to do manually, or could not do at all due to complexity. But somewhere along the line a human will be needed to do the rest. In addition, the more that we automate the easy tasks, the more complex and demanding the remaining tasks will be. But we can automate the action itself if it has been vetted. Gartner recommends "Rather than to seek full automation of all SOC activities, enterprises should seek "automatability" – the capability of being automated as higher levels of confidence are achieved". In the simplest scenario, this means sending out a notification to the IT Operations team that outlines the issue – what is the problem, what is the potential impact and what action is required to resolve it – and asking them to either confirm that this can be executed or to reject the automated action to do it manually. We can automate the action, without automating the decision.

> "Rather than to seek full automation of all SOC activities, enterprises should seek "automatability" – the capability of being automated as higher levels of confidence are achieved".
>
> Gartner

IT Operations do not trust automation

The downside to getting IT Operations to vet an action is that IT Operations teams are frequently overloaded, so that a handoff occurs from SecOps to IT Ops with a long delay in response. In the case of incidents such as ransomware, this delay can mean the difference between containment and disaster recovery, between an incident and a full-blown breach. The Security Operations team can help to alleviate this and build trust and confidence.

This can be achieved by keeping track of what actions are done manually – how many times the same action was done by a human instead of a machine – and working out the difference in time and effort between the two. The idea is that if someone receives the same notification for similar incidents requiring the same manual actions a multitude of times, we can demonstrate to them that this could have been safely automated. After all, we have the audit trail to prove it and the data to build a business case. More importantly, we will also garner data on what and where we can't safely automate. You can then expand automation as trust and confidence increase.

An automation may be safe in one business unit, but not acceptable in another. To accommodate this, process must support granularity, whether when gathering metrics or configuring the automations. Ideally, whatever automation technology you use must support this approach and provide the metrics that this requires. Technology can help to build trust, but when all is said and done, it's going to require that it is experienced by the people you expect to trust you.

Safely Automating means selectively automating

- Safe Automation
- Continuous Data Collection
- Threat

The most dangerous aspect of automating Security Operations and Incident Response is automating the containment of a threat. It is here where the greatest potential for detrimental impact on operational integrity can occur. But there are a lot of actions and tasks involved in effective incident response that can be safely automated, primarily focused on 4 core task categories listed below.

Continuous Data Collection

Ingesting, normalizing, parsing and correlating incoming security data from disparate security technologies and sources is something that can be safely and trivially automated. SIEM's are the most commonly utilized technology utilized to do this, but focus primarily on correlating log and event sources. Security Automation and Orchestration solutions, like DFLabs IncMan, provide a broader focus in terms of 3rd party security data sources that can be ingested, and also provide granular and customizable playbooks that can execute additional and specific data collection actions based on conditional workflows and triggers.

Continuous data collection means that SOC teams and incident responders can be proactively alerted of suspicious and malicious activity and events that occur in the environment that they are monitoring, and makes additional data required for incident qualification, analysis and investigation immediately available.

Threat Intelligence and context fusion and context enrichment

Sending out requests for related threat intelligence or for example submitting malware hashes to an external 3rd party provider can enable that party to infer that the submitting organization has been impacted by a specific threat or threat actor.

Automatically fusing incident data with external threat intelligence, or enriching it with additional operational context such related observables or IoC's is also generally safe and trivial. There is one caveat to this that applies especially to sensitive environments or organizations with critical security requirements – sending out requests for related threat intelligence or submitting malware hashes to an external 3rd party provider can enable that party to infer that the submitting organization has been impacted by a specific threat or threat actor. Automating threat intelligence fusion and context enrichment can be a great time saver, reducing the time and effort required to qualify and verify an incident, assess its impact and shorten the subsequent time from discovery to containment.

Triage and Notification

Triaging incidents to ensure that they are assigned to the correct security analyst or incident responder can be safely automated, as can notifying relevant stakeholders such as HR, Legal, executive management or related IT Operations Teams. Similarly, opening tickets or cases in IT Helpdesk and Service Management systems reduces the amount of menial work that the SOC or CSIRT must conduct, speed up incident resolution and also ensure that a closed-loop incident response process is followed.

Forensic Evidence Gathering

Gathering and fetching related forensic evidence such as process lists, application inventories, configuration settings, activity logs and disk images can also be safely automated, providing these do not lock out active users, shut down running processes or initiate system shutdowns or reboots.

Forensic evidence gathering must be conducted in a manner that does not tamper with or destroy relevant evidence, and must collect and store that evidence in a way that ensures that it is compliant with legal and regulatory mandates. Automation, if done correctly, can aid to assure this as well.

Automating Containment

The greatest risk and danger in automating incident response is when applying it to containment and remediation of threats, and so poses the greatest challenge. When considering automating the containment of a specific threat, three questions are relevant and should be posed:

How reliable is the detection and identification?

Three questions to pose when automating containment

- How reliable is the detection and identification?
- What is the potential detrimental impact if the automation goes wrong?
- What is the potential risk if this is not

The degree of confidence in the detection and identification of a specific threat or attack is a major factor in deciding whether to automate containment. This has been the historical inhibitor for enabling full blocking and containment, for example when deploying an IPS. Generally, there are two types of approach to this.

The first approach considers how reliable it is in general to detect a specific attack. Some types of malware or exploitation can be easily identified. This is the case for example when a malware file hash has been confirmed via multiple queries, for example via a local AV solution and an external service such as VirusTotal.

The second approach is based on an analysis, essentially a consolidated score, that factors in multiple related IoC's, exploited attack vectors and other observables to derive a weighting of certainty which will be used to determine whether containment will be automated or not. Once a high certainty threshold is exceeded, containment can be automated. Behaviorial analysis and related machine learning capabilities are frequently used for this, although simpler methods such as Correlation can also be used providing sufficient IoC's have been evaluated.

What is the potential detrimental impact if the automation goes wrong?

Whether the cure is worse than the disease must be carefully considered when deciding what to automate. Erroneously automating threat containment on crucial infrastructure or in operationally critical environments, or when related to priority customer or mission critical processes can detrimentally impact operational integrity and in a worst-case scenario cause the loss of revenue. Trading Platforms, Internet retail portals and Medical or Energy infrastructure are good examples where automated response is best avoided unless done very selectively and with caution. One thing to consider though, is that the same reasons why we are hesitant to automate – because these are critical and sensitive – also means that incident response in these cases must often be executed rapidly to contain threats to avoid the same potentially negative impacts that hap hazardous automation may cause.

What is the potential risk if this is not automated?

This question needs to consider the potential impact and associated risk of not automating the containment. Many types of incidents do not require an immediate response. For example, detecting a port scan against an external asset does not necessarily pose an immediate danger – it just indicates that a malicious actor or automated tool is probing your infrastructure. Malicious activity in a test environment, providing that it is not used for R&D, will represent a lower risk than if the same activity is detected in a mission critical environment.

On the opposite end of the spectrum, if activity is detected in a critical or sensitive environment or targeting privileged users, or an attack in the latter phases of the cyber kill chain are discovered, automating the containment response may be highly desirable or even necessary to prevent the loss of sensitive IP or a detrimental impact on operational integrity.

Generally, the further along the kill chain an attack has progressed, the quicker the response must be and the higher the need for automated containment.

White and Black Listing

The safest way to approach and implement the automated containment of threats is to work with White and Black lists. These are used to identify threats and incident types, environments and infrastructure where automation is acceptable and desirable, or where it is absolutely not safe and permitted.

These lists will be adapted and amended as confidence in detection and automation increases and improves, and as the threat landscape, regulatory drivers and organizational priorities and objectives evolve.

Most organizations will also work with is effectively a Grey List – composed of threats or incident types that have been encountered for the first time and will require careful evaluation in regards to whether they can be automated safely, and how to go about this.

[1]The Cyber Kill Chain is a registered trademark of Lockheed Martin. See the reference below for a detailed overview of the Cyber Kill Chain http://www.lockheedmartin.com/us/what-we-do/aerospace-defense/cyber/cyber-kill-chain.html

[2] See "The Five Characteristics of an Intelligence-Driven Security Operations Center", Gartner 2015, Oliver Rochford & Neil MacDonald

Acronyms

EDR Endpoint Detection and Response
IoC Indicator of Compromise
IPS Intrusion Prevention System
R&D Research and Development
SAO Security Automation and Orchestration
SIEM Security Information and Event Management
SOC Security Operations Center

About DFLabs

DFLabs IncMan is the leading Security Automation and Orchestration technology platform to manage, measure and orchestrate security operations tasks including security incident qualification, triage and escalation, threat hunting & investigation and threat containment. IncMan harnesses machine learning and automation capabilities to augment human analysts to maximize the effectiveness and efficiency of security operations teams, reducing the time from breach discovery to resolution and increasing the return on investment for existing security technologies.

Our vision is to enable full spectrum intelligence-driven command and control of your security operations, orchestrating the entire incident and investigation lifecycle for SOC and CSIRT teams, with technology that empowers security analysts, forensic investigators and incident responders to respond to, track, predict and visualize cybersecurity incidents, and for security managers and CISO's to manage and measure operational performance and cyber risk.