WHITE PAPER

Automate Your Incident Response Safely By Automating Selectively.

Overcome security automation challenges within incident response by safely and selectively automating processes and tasks without impacting operations.

JOHN MORAN, SENIOR PRODUCT MANAGER

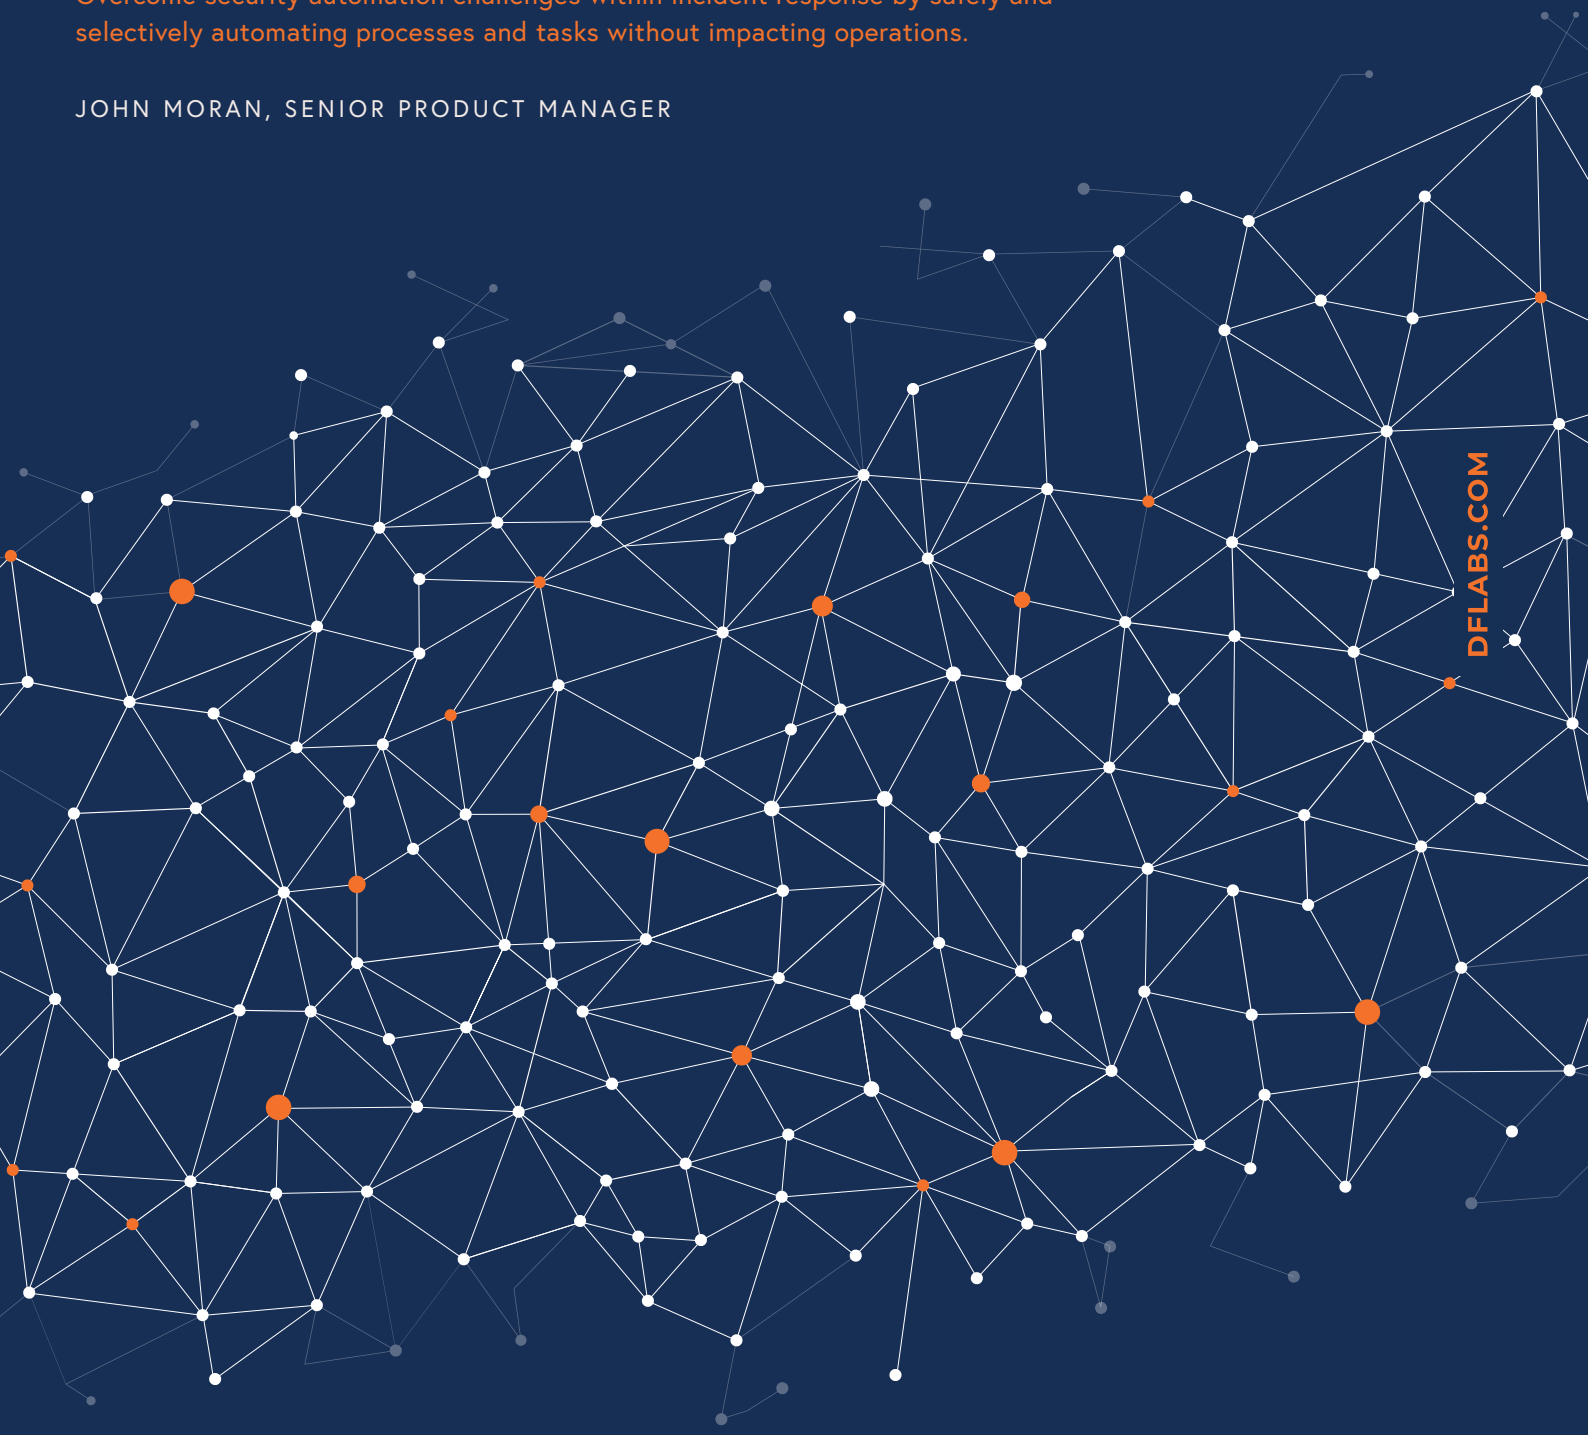

Automate.
Orchestrate.
Measure.

DFLABS
CYBER INCIDENTS UNDER CONTROL

Contents.

DFLABS.COM

Automate.
Orchestrate.
Measure.

DFLABS
CYBER INCIDENTS UNDER CONTROL

Introduction.

Automation in security operations and incident response is a topic that is now more on the radar than ever before. 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 can not achieve the velocity required to effectively and rapidly respond to attacks that are often not detected until a threat actor has almost completed the kill chain. 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.

The consistent feedback is that automation is highly desirable, at least by security teams, who are the ones struggling on a daily basis from work overload and alert fatigue, but this desire to implement effective automation is often 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 the detrimental impact and damage resulting from potentially doing this wrong.

Automation is not new and enterprises have been promised automated containment capabilities before, but previous premature attempts such as through antispam and intrusion prevention systems for example, which lacked the ability to reliably identify anomalies and attacks, has since led to IT operations and some executive management teams being more reluctant to pass such powers to machines. This is despite detection capabilities having dramatically improved in recent years, especially using behavioural modelling and machine learning driven approaches.

"Manual processes can not achieve the velocity required to effectively and rapidly respond to attacks that are often not detected until a threat actor has almost completed the kill chain."

Three Common Automation Challenges.

Now let's take a look at the three basic challenges that security teams face when considering automation and how they can be overcome so automation can be successfully implemented.

1. The SecOps team can assess the impact of the risk, but not the impact on production
2. Not all decisions can always be completely automated
3. IT operations do not trust automation

SecOps Can Assess the Impact of the Risk but Not the Impact on Production.

Security operations teams are often focused purely on the risk and impact of the threat and in their own silo struggle to build up and maintain an awareness of what is going on in production and who it may affect. E.g. Is the affected system mission critical, is the system unstable, or is it a legacy system? Is the system currently being used to process critical business internal financial reports, or is a customer using it and being affected when they are paying for a service you should be providing? Disabling a seemingly harmless user account may actually 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 relates to this may be out of date, but either way this can have a huge impact on how the incident response or remediation process must be conducted. The incidents or vulnerabilities should still be addressed but this may require additional time, tasks and a specific way of approaching it, and this is likely to vary from organization to organization. Regardless of this, it is important for departments to be interlinked as much as possible and for processes and procedures and related documents to be regularly updated to ensure critical information used and kept on file is always correct.

Not All Decisions Can Always Be Completely Automated.

The actual containment or remediation response is not the only thing that can be automated. We can automate a wide variety of tasks, including prioritization of an incident, fetching additionally required information and context or more simply notifying and creating tasks for stakeholders.

Through using automation we can make people more efficient and can use automation to take away some of the more menial and repetitive tasks. We can even use machine learning to compute an analysis that a human would take a millennia to do manually, or could not do at all due to its complexity. But somewhere along the line a human analyst may well still be required to carry out a manual decision if needed.

The more we automate the easy tasks, the more complex and demanding the remaining tasks will be, but we can still automate the next actions to be taken regardless of this if they have been manually vetted. Analysts will be able to spend more time handling and vetting these more complex manual decisions, rather than wasting their valuable time carrying our laborious, mundane and repetitive tasks

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. It would include what the problem is, the potential impact, and what action is required to resolve it. It would ask them to either confirm that this can be executed automatically or require them to reject the automated action and for them to carry it out manually.

We can therefore successfully automate the action without automating the decision as and when required, based on the levels of automation we are comfortable with in our operations processes and workflows and this is also open to change over time and as experience and knowledge grows.

"The actual containment or remediation response is not the only thing that can be automated."

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 and between an incident and a full blown breach. The security operations team can help to alleviate this and by building trust and confidence.

This can be achieved by keeping track of what actions are done manually including how many times the same action was taken 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, SecOps can demonstrate to them that this could have been safely automated. There will also be an audit trail to prove it and the data to build a business case if required. More importantly the team will be able to gather data on what tasks can be automated safely and those that couldn't be, with their potential resulting consequences. The level of automation can then be expanded as needed as trust and confidence increases.

Why Safely Automating Means Selectively Automating.

The most dangerous aspect of automating security operations and incident response processes and tasks is automating the containment of a threat. It is here where the greatest potential 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 four core categories outlined below.

"Automating threat intelligence fusion and context enrichments can overall be a great time saver."

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. Security Information Event Management (SIEM) tools are the most commonly utilized technology to do this, but its focus is primarily on correlating log and event sources. Security Orchestration, Automation and Response (SOAR) solutions, like IncMan SOAR from DFLabs, provides a broader focus in terms of third party security data sources that can be ingested, and also provides granular and customizable playbooks that can execute additional and specific data collection actions based on conditional workflows and triggers.

Continuous data collections means that security operations center (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.

Automatically fusing incident data with external threat intelligence, or enriching it with additional context such as related observables or indicators of compromise is also generally safe. There is one possible caveat to this though that applies to sensitive environments or organizations with critical security requirements. Sending out requests for related threat intelligence or submitting malware hashes to an external third 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 enrichments can overall be a great time saver, reducing the time and effort required to qualify and verify an incident, assess the 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 also be safely automated, as can notifying relevant stakeholders such as HR, legal or 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. This essentially speeds up incident resolution and also ensures that a closed-loop incident response process is followed.

"Triaging incidents to ensure that they are assigned to the correct security analyst or incident responder can also be safely automated."

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 carried out correctly, can aid to assure this as well.

"Forensic evidence gathering must be conducted in a manner that does not tamper or destroy relevant evidence."

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 asked.

1. How reliable is the detection and identification?
2. What is the potential detrimental impact if the automation goes wrong?
3. What is the potential risk if this is not contained?

Let's now briefly look at these one by one.

How reliable is the detection and identification?

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 intrusion prevention system (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.

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

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

Whether the curse is worse than the disease must be carefully considered when deciding what to automate. Erroneously automating threat containment on critical 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 could 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 or with caution. One thing to consider though, is that the same reasons why we are hesitant to automate, because they are critical and sensitive, also means that incident response in these cases must often be executed more rapidly to contain threats to avoid the same potentially negative impacts that hap hazardous automation may cause.

"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."

What is the potential risk if this is not contained?

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 research and development (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.

"One of the safest ways to approach and implement the automated containment of threats is to work with White and Black lists."

The Safest Way to Automate Containment.

One of the safest ways to approach and implement the automated containment of threats is to work with White and Black lists. These are used to identify threats and identify 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.

For cases where the potential risk is high if the threat is not contained, however the detection and identification confidence is low and/or the detrimental impact if the automation goes wrong is high, semi-automated containment actions may be a better alternative than fully manual actions. Semi-automated actions (referred to as User Choice Decisions within DFLabs' IncMan SOAR solution) pause the automated workflow, allow analysts to review the previously gathered intelligence, and make a human decision regarding the next appropriate course of action.

In contrast to a completely manual decision, a semi-automated decision contains two or more predefined paths for the analyst to choose from as appropriate, and allows the automated workflow to continue after the decision is made. In the case of containment actions, a semi-automated decision may be used to allow an analyst to view previously gathered intelligence on the source and destination of malicious network traffic, The analyst can then manually determine if blocking the external host or internal asset is necessary and most appropriate course of action to take.

Conclusion.

An automated action may be safe in one business unit, but not acceptable in another, therefore safely automating means selectively automating. To accommodate this, processes must support granularity, whether gathering metrics or gathering the automations themselves. Ideally, whenever automation technology is used, there must be the correct level of support for this approach by all teams involved, including SecOps, IT Ops and other teams potentially affected within the business to ensure its success. Technology can help to build trust, but when all is said and done, it is going to require that it is experienced by the people you expect to trust you.

About Us.

DFLabs is an award-winning and recognized global leader in Security Orchestration, Automation and Response (SOAR) technology.

Its pioneering purpose-built platform, IncMan SOAR, is designed to manage, measure and orchestrate security operations tasks, including security incident qualification, triage and escalation, threat hunting & investigation and threat containment.

IncMan SOAR 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.

As its flagship product, IncMan SOAR has been adopted by Fortune 500 and Global 2000 organizations worldwide.

The company's management team has helped shape the cyber security industry, which includes co-editing several industry standards such as ISO 27043 and ISO 30121.

DFLabs has operations in Europe, North America and EMEA.

For more information, visit our website www.dflabs.com or connect with us on Twitter @DFLabs.

DFLABS.COM

DFLABS
CYBER INCIDENTS UNDER CONTROL

CONTACT US:

BOSTON - UNITED STATES
150 State Street
Boston, 02109
T – +1 201 579 0893
E – sales@dflabs.com

LONDON - UNITED KINGDOM
1 Primrose Street
London, EC2A 2EX
T – +44 203 286 4193
E – sales@dflabs.com

MILAN - ITALY
Via Bergognone, 31
20144, Milan
T – +39 0373 82416
E – sales@dflabs.com

CUSTOMER SUPPORT:

T – +39 0373 82416
E – support@dflabs.com

Automate.
Orchestrate.
Measure.