The Essential Guide to Phishing Investigation and Response

Introduction

Phishing has been around since the mid-90s and can be legitimately considered a "classic" hacker tool. The word was coined by hackers for their activities tied to luring AOL users to part with their password and account information. And like all classics, it has stood the test of time as a lucrative technique for cybercriminals, becoming more prolific and sophisticated. The 2021 Ponemon Cost of Phishing Study has shown that the average cost of phishing has more than tripled from $3.8 million in 2015 to $14.8 million in 2021.[1]

The study also discovered that loss of employee productivity was a significant component of the cost of phishing. As a security operations incident responder or security analyst working in a SOC, this data point probably comes as no surprise. The phishing response workflow consists of multiple time-consuming tasks associated with cleaning up infected systems and conducting incident investigations, and to some extent, documentation and end user communications (and awareness education). Ponemon found that the cost of resolving malware infections has doubled the total cost of phishing.

The cost of phishing has more than tripled since 2015. The average annual cost of phishing has increased from $3.8 million in 2015 to $14.8 million in 2021.

– The Ponemon 2021 Cost of Phishing Study

Part of this cost is business disruption due to ransomware. And as we know, phishing is a popular gateway to ransomware and other major breaches. We cover the topic of ransomware response in more detail in our Essential Guide to Post-Intrusion Ransomware Response.

Your Automation Strategy: Some Considerations

Before you start on your automation journey, we recommend you consider the following questions:

- **Identify repetitive tasks**: How cumbersome is the current manual process? How important is manual intervention? Is this task often performed the same way? Are your security analysts using their time efficiently? Does your team need to touch this task, or are they just performing rudimentary data collection?
- **Respond rapidly**: How critical is it that the organization responds quickly? How much time can you save through automation? Are responses currently delayed due to a lack of resources? Does the current manual process meet the required service-level agreement?
- **Integrate across multiple components**: How many different systems are involved? How many teams work on resolution? What level of expertise do you require to master individual systems? Do security analysts have the appropriate permissions to access the required systems? Is the data properly formatted and compatible across systems?
- **Reduce errors**: Is the current manual process prone to errors? Are processes followed consistently? Is all data accessible within a single system?

The average total cost of ransomware last year was $5.66 million, and the average percentage rate of ransomware attacks from phishing was 17.6 percent.

– The Ponemon 2021 Cost of Phishing Study

Phishing Investigation and Response

Phishing incident response is the poster child for automation for several reasons. Phishing incidents can be frequent and require a quick response. The response workflow is relatively straightforward and repeatable. As mentioned above, there are high-efficiency costs associated with capturing intelligence, evaluating intelligence, and tracking and cleaning up infected systems and end user communications.

So let's take a look at the following automation scenario with Cortex® XSOAR.

Your SecOps team has set up a phishing mailbox on the email system, which they monitor using Cortex XSOAR, and they request that users forward all suspected phishing messages to that mailbox. A user has forwarded one such email. Alternatively, you might be using a phishing detection tool that would then send alerts to Cortex XSOAR to trigger a series of automation responses.

1. The 2021 Cost of Phishing Study," Ponemon Institute, June 2021.

When the new email or phishing alert arrives, Cortex XSOAR retrieves it as an event and creates an incident. As part of the incident response, Cortex XSOAR automatically executes a playbook to analyze the email and, optionally, to automatically respond if the email contains phishing or malware content. Cortex XSOAR starts the basic analysis by retrieving the original message that the end user forwarded. In addition to extracting the email headers, including domain and IP address indicators, the analysis extracts URL indicators. To determine if unknown URL indicators are malicious, Cortex XSOAR submits them to WildFire. WildFire® is a Palo Alto Networks cloud service that analyzes files and email links to detect threats. You can use any of your preferred malware analysis products with our extensive out-of-the-box integrations. When the analysis completes, before taking any further action, the assigned SecOps analyst uses Cortex XSOAR to review the attack information that Cortex XSOAR has automatically gathered.

Automatic Response Options

If you want Cortex XSOAR to take actions based on the results, this automation scenario includes some advanced options. The first option is to search the email system for all recipients of the original message and if the analysis determines that the message contained a malicious phishing URL, remove the message from all mailboxes system-wide. The second option assumes that users at your organization connect

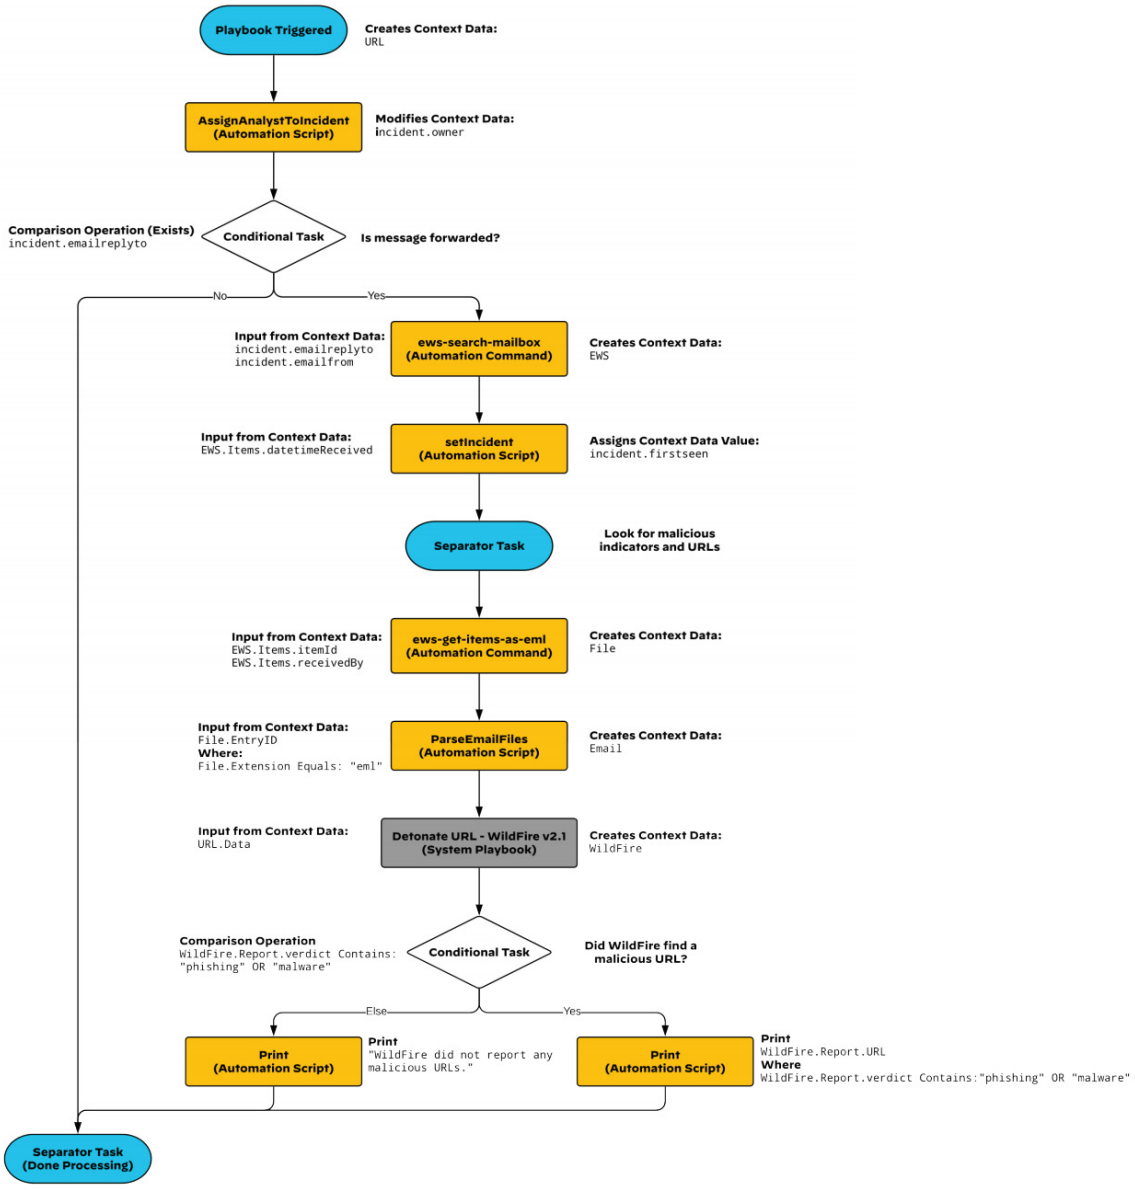

Figure 1: Automated phishing playbook (basic)

to the network through Prisma® Access. Cortex XSOAR searches the Prisma Access logs stored in Cortex Data Lake to determine if any users tried to access the malicious URL in the original message. If the logs confirm that users did try and connect, then you can notify the users. You can use this information to create new incidents to interact with users who might need further attention, such as attending a training session focused on phishing awareness or initiating a forensic examination of their computer system.

The Cortex XSOAR automated playbook flows:

After creating an incident, a Cortex XSOAR playbook performs an automated investigation and response:

- It retrieves the email from your Microsoft 365 (or other) email system
- Accesses reputation data for the content from AutoFocus™
- Submits any URLs contained in the message body to WildFire for malware dynamic analysis
- Notifies the case owner to review the results of the investigation that the playbook captures throughout the incident

These optional tasks enhance the basic playbook:

- You can modify Cortex XSOAR to add custom incident fields that you use to summarize phishing information across multiple incidents.
- You then append tasks to the basic playbook, including tasks that extract incident-specific data and populate the custom fields.
- You can create a custom report by using custom widgets that summarize information collected across all phishing incidents.
- You can modify your basic playbook to add a response action. If you have confirmed that the phishing email contains malicious content, your playbook searches the Microsoft 365 (or other) email system to identify users who have received copies of the phishing email.
- If the case owner approves, the playbook deletes the email message system-wide.
- You can further modify the playbook to determine if any users tried to access any malicious content contained in the email. This option assumes that your organization uses Prisma Access for mobile user access. The playbook queries Cortex Data Lake for security log entries that match malicious URLs in the email message. If Cortex Data Lake returns any matches, the playbook sends emails to notify affected users.

Leveraging Machine Learning to Train Cortex XSOAR

In the last five years or so, we have become closely acquainted with security operation center (SOC) teams that use Cortex XSOAR. One of the first things we learned was that reviewing potential phishing incidents consumes a significant amount of time. And many of the suspected phishing incidents turn out to be false positives.

So when we analyzed what most analysts do to investigate a potential phishing attack, we realized that this is a classic challenge for machine learning (ML). The result? A phishing email classifier aimed to help organizations detect malicious phishing emails with a high degree of accuracy.

Here are some of the considerations we took into account to make the model work for as many SOCs/SecOps teams as possible:

- When investigating phishing incidents, SecOps teams use different tools and services that provide enriched data on indicators found within the emails (IP addresses, attachments, URLs, and so on) to see if any proof of malice exists. The phishing classifier adds additional perspective as a text classifier, which means that it's trained based on the text of the email. By learning word patterns that correspond with phishing and non-phishing emails, it can predict whether a given new email is phishing or not. Therefore, it can supplement existing tools which often do not take the email text into account.
- The training of the phishing classifier is done using historical emails identified and classified as malicious by the organization. Phishing attempts might differ significantly from one organization to another. Therefore, training a classifier based on your own emails has the potential to yield a classifier that is adjusted to your typical phishing email and, therefore, a lot more accurate than a generic solution.
- The training of the phishing classifier is done within your environment. We are aware that phishing incidents contain sensitive data. Therefore, we designed the phishing classifier as an integrated part of Cortex XSOAR so that your emails do not need to be sent to a third party. The pre-processing of the emails and the training of the phishing classifier are all done in Cortex XSOAR and do not require any export of emails.

- You can train a phishing classifier with a relatively small number of incidents. The phishing classifier is a deep learning model. It achieves a model with relatively high precision, even if it's trained on a small number of incidents.
- It's possible to use the phishing classifier in multiple ways. Customers can choose to present the classifier's output to human SOC analysts as an additional parameter to consider. Another option is to prioritize the incidents which the classifier predicts as phishing within high confidence. The final option is full automation, where incidents can be closed automatically or emails searched and destroyed based on the classifier's prediction. The decision on how to use the model is taken based on the classifier's performance and based on where it can be of most benefit to your team.

More details on the phishing classifier can be found in this blog, and you can also download the playbooks for phishing from our Cortex XSOAR Marketplace.

In order to run these playbooks, you will need to install Cortex XSOAR. If you don't have a version, register for a Cortex XSOAR Community Edition free trial.

A Selection of Content Packs in the XSOAR Marketplace

Our content packs (with playbooks) are updated biweekly, so this is just a snapshot of content packs related to Phishing investigation and response. We encourage you to explore the Marketplace directly for products you currently use in your SecOps/SOC.

Table 1: Content Packs, Integrations, and Playbooks

Content Pack	Integrations Used	Playbook
AutoFocus	Palo Alto Networks AutoFocus v2	Automated Phishing Investigation
Cortex Data Lake	Palo Alto Networks Cortex Data Lake	Check CDL logs for URL
Exchange Web Services	Microsoft EWS v2	Automated Phishing Investigation
Palo Alto Networks WildFire	Palo Alto Networks WildFire v2	Automated Phishing Investigation
Phishing Campaign	Mail listeners	Detect and manage phishing campaigns
Malware Lateral Movement Assessment and Response	Splunk, Carbon Black, Microsoft Exchange	Remediate malware lateral movement due to phishing campaign
DBot Create Phishing Classifier	Cortex XSOAR	Create a phishing classifier

Other relevant content packs (integrations and/or threat feeds) available in our XSOAR Marketplace include: Abnormal Security, Agari Phishing Defense, Cisco Email Security, Cyren Threat Feeds, FraudWatch PhishPortal, Google Suite Security Alert Center, Gmail, GreatHorn, Group-IB Threat Intelligence, IPQualityScore Threat Intelligence, SecurityAdvisor, RiskIQ Phishing Feed, SlashNext, Trustwave SEG, and Twinwave.

ROI: Measuring Automation Success

Once you start your automation journey, your Cortex account manager can work with you and your team to assess the efficiency metrics gained from leveraging automation.

Ready to Get Started?

Check out our Cortex XSOAR Community Edition free trial and learn how to edit our playbooks in our Phishing Hands-on Workshops.

Want to learn how to leverage automation in case of a ransomware breach? Read our companion, Essential Guide to Post-Intrusion Ransomware Response.

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com