DOMAINTOOLS®

# STRATEGIES TO VET YOUR THREAT INTELLIGENCE AND REDUCE FALSE POSITIVES

## WHERE THERE'S SMOKE, THERE'S FIRE.

Alarm bells ring and first responders rush to the scene. Equipment and supplies are readied in expectation of a blaze. With masks on and hoses at the ready, firefighters barrel into a cloud of smoke. But soon, they find no flames. It was a false alarm. Confused and frustrated, they retreat, regretting the time, energy and resources wasted on nothing more than an unplanned drill.

While this may not be a likely scenario for firefighters, the burden of false alarms are a daily reality for cybersecurity practitioners. One **IDC report** found that respondents field 10,000 or more security alerts each month, 52 percent of which are false positives. In addition to killing morale and detracting resources away from real threats, false positives cost organizations an estimated 21,000 hours according to the **Ponemon Institute**. Another **source** found that 56 percent of IT professionals surveyed admit to ignoring alerts due to past false positive experiences.

Can security professionals reduce the deafening noise of alarm bells, so that the smoke they respond to is from an actual fire? Is it possible to overcome the daunting challenge of knowing everything that's going on inside and outside the network? Robert M. Lee, a SANS Institute Certified Instructor and Tarik Saleh, Senior Security Engineer and Malware Researcher at DomainTools believe the answer is yes—if security teams take new approaches to how they think about, vet and validate intelligence, indicators and adversary behaviors.

This paper will discuss detection strategies to reduce false positives, and models that improve threat hunting and investigations outcomes. It will also cover leading tools that help teams make the most of their limited time and resources.

# TYPES OF THREAT DETECTION

The model below illustrates common detections security teams use, and the value of each.

| ENVIRONMENTAL | THREAT |
|---|---|
| Weak as a detection strategy, but useful for layering on top of an investigation. | Information must be analyzed in context to turn pieces of data into true intelligence and knowledge that can answer key questions about an adversary. |
| **Modeling** from the existing environment and applying that to alerts shows what's happening on the network, but without any context or insight into actual threats. | **Threat Behaviors** based on user and behavior analytics, machine learning applications and models can add context to threat data. |
| **Configuration** based solely on the environment, industry and operations will lead to a lot of false positives, due to inherent field view bias. | **Indicators** from threat feeds and other sources provide a subset of useful data, when combined with other pieces of information, but not when relied upon alone. |

## FIRE DRILL VS. STRATEGY

At some point, every security professional will have a field of view bias that leads to belief that their visibility is better than it actually is. What we receive via intel feeds is limited to only what is known—therefore it provides an incomplete view of the threats out in the wild, or potentially even penetrating the network. In a recent talk, Robert M. Lee said, "You can download all the indicators in the world, and keep downloading and going through them, and you will always be behind."

This is because an indicator-led approach to threat detection and response is a reactive one. Like an unplanned fire drill, an indicator-led approach leaves everyone scrambling and reacting with limited knowledge about what's really going on. It is impossible to keep up with adversaries, threats and malicious infrastructure with this methodology. Indicators offer lots of context, but as we consume them, we find that some of them are incorrect—false positives that waste time and cause teams to fall behind.

Conversely, a strategic, behavioral detection approach, rooted in an understanding of the difference and relation between data, information and intelligence, can drastically improve response. With more accurate detection (i.e. fewer false positives), everyone can remain calm, focus on the real threats and increase coverage and response for when an incident occurs. A detection strategy led by the threat behaviors and adversary tactics, techniques and procedures (TTPs) that are of the greatest consequence to the organization allows investigators to focus in on the most relevant intelligence. It provides a place to pivot from, so investigators can enrich their information with additional data and indicators, and in turn learn even more about what happened. It begins to tell a story about the series of events, the impacts and why they matter. This is a powerful place to be as a defender.

*"You can download all the indicators in the world, and keep downloading and going through them, and you will always be behind."*
*— Robert M. Lee*

## FINE-TUNING THREAT DETECTION AND RESPONSE

Behavioral or TTP-led detections can cover one or more phases of an intrusion and can be simple or complex. They contain important context and can be transposed and scaled. MITRE's ATT&CK framework, a useful depiction of high level tactics and techniques that have been observed in the real world, echoes the importance of this type of approach and offers teams a methodology for building breadth and depth into their detection models.

A detection strategy that relies on TTPs, leverages indicators for enrichment and reduces false positives includes the following steps:

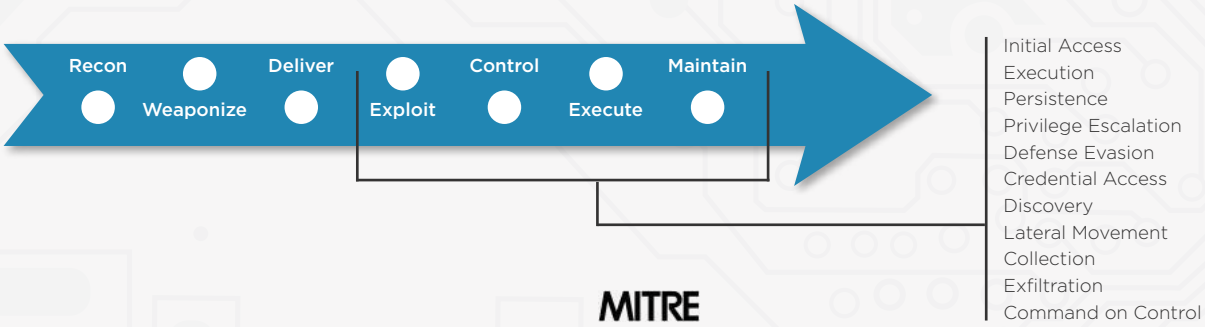


1. **Understand your threat model.** Lay the foundation of your threat model by industry and the different types of threats and activity groups that you are most worried about.
2. **Identify TTPs.** Think about the behaviors, TTPs and scenarios that you know are of the greatest consequence to your organization and for which you need to be well prepared.
3. **Map detections to MITRE ATT&CK.** Support detection engineering and identification by mapping your alert system to the MITRE framework. This ensures you are getting adequate coverage without casting too wide a net.
4. **Focus on relevant indicators:** Outline the indicators that are related to the behaviors you care about. Threat feeds and industry intelligence can help define specific pieces of data that matter to your organization.
5. **Implement and enrich:** The scope of indicators that need to be investigated will begin to narrow, to only the things that are suspicious enough to explore. From there, related indicators can be used for enrichment and forensically analyzed to mine additional intelligence.

## TTP-BASED DETECTION JOURNEY: MOCK SCENARIO

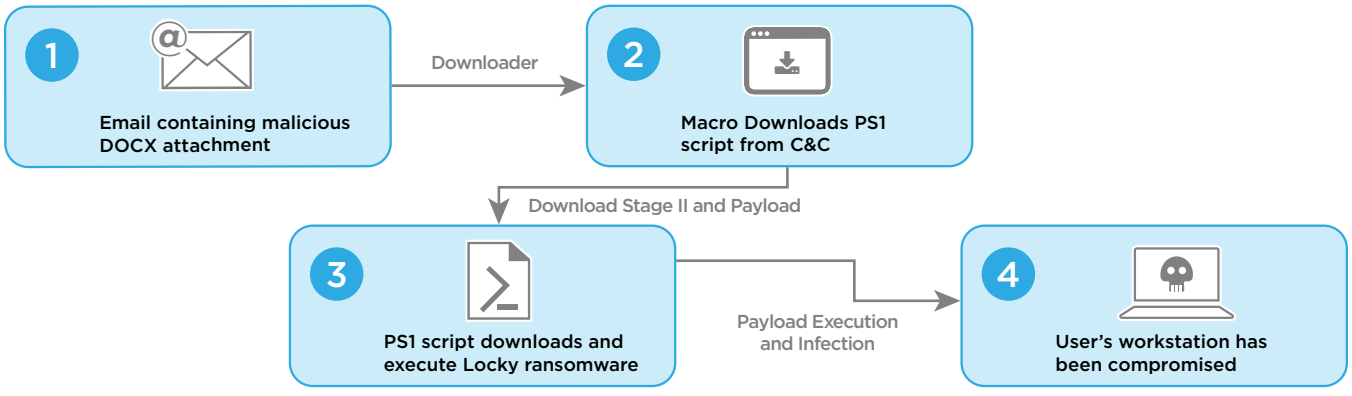

## FINE-TUNING THREAT DETECTION AND RESPONSE (CONT.)

Investigators need to focus on telling the whole story of a threat with confidence. Fewer false positives leads to faster and more accurate understanding of the root cause of an event, what happened, what was impacted, how that relates to what's happening in the wild and what the attacker's capabilities are. These insights create true security intelligence, so teams can pinpoint flaws in the defense and remediate them.

*"Investigators need to be data driven, able to reveal specific details about an attacker's activities. With every alert, we should be following strategies that help us piece the facts together and share the full story with our teams, company leadership and the broader security community."*
*— Tarik Saleh*

A "lateral movement" analytic alerts a security team to an external VPN session that moved files onto systems and executed them, followed by lateral movement. Looking at indicators alone would create a complicated map of data about an IP address that penetrated the network.

### THE BEHAVIOR-LED APPROACH NARROWS THE SCOPE:

- Right at the outset, investigators know that a malicious file was dropped. They can conduct forensics on the behavior to take a closer look at what else is related to it, and which systems are reaching out to it.
- The team can then take all of the indicators against that behavior to enrich and reveal what else is already known about the threat.
- The team builds context, connects the dots of the threat to information from other investigations and can compare it to analysis of what is happening in the wild.
- Coverage is created for the threat quickly, and the team then begins to validate the threat and the intelligence about it.

### SETTING UP AND SEGMENTING FOR THREAT ANALYSIS

To effectively investigate and validate alerts for watering hole attacks, and leverage the full capabilities of tools designed to support these efforts, the environment must be set up ahead of time.

**BE SURE TO ESTABLISH:**

- **SEGMENTATION:** An analysis machine or malware lab segmented from the network (AWS and Docker work well for this).
- **MIMICKING:** A virtual machine that can execute codes including JavaScript and ActiveX to mimic a vulnerable client (e.g. Windows Virtual machine).

### WHEN THE SMOKE CLEARS, VALIDATE

Validating alerts is an important step in further reducing the volume of false positives and strengthening the behavioral detection model. For example, watering hole attacks such as Exploit Kits, malvertising and drive-by malware present unique challenges for defenders trying to gather the whole story. By using a combination of tools to gather complete data about watering hole attempts and attacks, and taking steps to validate that data, organizations can improve the quality of their alerts for these types of threats.

## WITH THE RIGHT TOOLS AND ENVIRONMENT IN PLACE, THE METHODOLOGY FOR VALIDATING ALERTS INCLUDES:

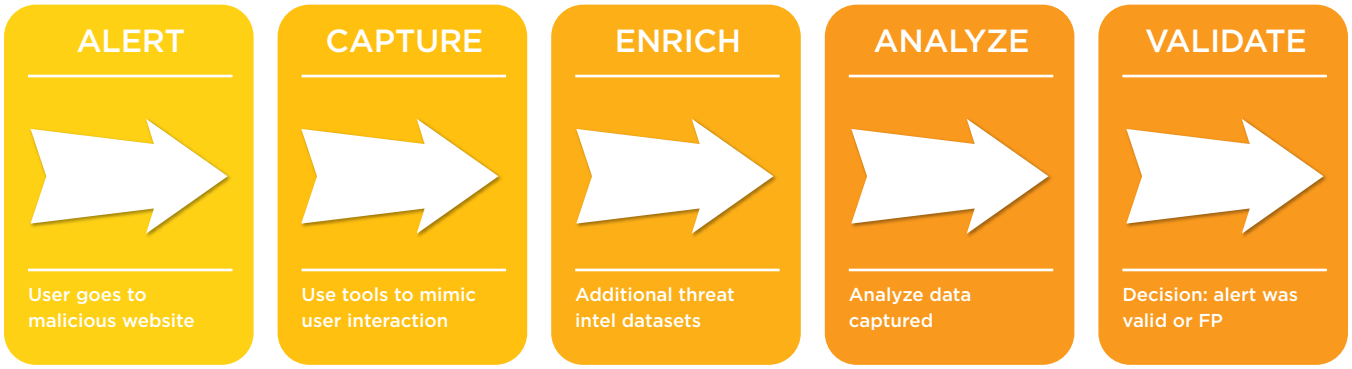


An **alert** triggers inside the security operations center.

The team uses tools to **capture** and mimic the network victim's experience.

With open source intelligence feeds and closed source curated feeds, the team begins to **enrich** the data.

Piecing data together, adding context and leveraging MITRE, investigators **analyze** key focus areas.

The team inspects whether the alert was a true positive to **validate** its intelligence—if the alert was a false positive, further examination reveals which elements of the detection must be adjusted to reduce the noise.

Illustration: Tarik Saleh

### TOOLS YOU CAN USE

There are numerous tools security teams can tap to support a behavior-led detection strategy, gather information quickly and speed up the process of validating alerts. Some of the most effective tools available include:

**CapTipper:** A Python tool used to analyze, explore and revive HTTP malicious traffic and provide PCAP analysis using a web server that mimics the server in the PCAP file. It contains interactive tools for inspection of the hosts, objects and conversations, and to slice and dice pieces of data from an alert.

**Key features include:**

- File classification that highlights when a PDF or executable is embedded in web traffic
- Visual structure and display of HTTP communications, content and traffic flow, including GET requests and response codes
- Object beautifier to make it easier to read/analyze objects and debug on the fly
- Alert validation and forensic reports support via metadata (e.g. hash sums) and plugins to open source feeds and other intelligence repositories

## TOOLS YOU CAN USE (CONT.)

**Fiddler2** + **EKFiddle**: A bundle framework based on the Fiddler web debugger to study Exploit Kits, malvertising and malicious traffic. By handling PCAP files using your analysis machine, it can execute SSL decryption of traffic and help connect the dots behind a threat.

**Key features include:**

- Allows investigators to inspect and download the actual content involved, to support alert vetting and quick, surgical response
- Provides a numerical order in which traffic proceeded on a known malicious site, alongside a breakdown of the content including CSS code and which pieces are JavaScript, binaries and other formats
- Automatically extracts lookups for IP addresses, host names, etc., and enriches captured traffic with regexes
- Can connect to DomainTools Iris to integrate additional enrichment data (domain and passive DNS)
- Collects IOCs

**YARA**: An open-source tool designed to help malware researchers identify and classify malicious code and create descriptions of malware families and malicious code families based on text or binary patterns. When intel has inconclusive results, YARA may provide answers about binaries, etc. that help tell the whole story about a threat.

**Key features include:**

- Built-in rules, or descriptions, that consist of a set of strings (regular expressions, hexadecimal and text) and metadata (hashes, dates and other valuable context) unique to a specific malware; the strings include conditional statements that identify when the rule fires on a specific threat
- Incorporates fine-grained data that may not be captured in other intelligence sources
- Filtering so that only rules that align with the investigation's goals are displayed
- Prebuilt index files and capabilities index to offer insight into what the file can do
- Workflows to proceed with next steps—such as taking the threat into a sandbox, manually reviewing or investigating further with additional tools

## CONCLUSION

Like smoke from a blazing fire, threat alerts are suffocating security teams. But when so many of the alerts are false positives, investigators can't effectively see or prioritize the fires they need to put out. A shift in process and mindset is essential. Taking a behavior-led approach, instead of an indicator-led one, can significantly reduce the false alarms. Prioritizing alerts by risk—such as focusing first on threats against a system that handles sensitive data—is another best practice. Organizations that roll up their sleeves and take the time to build a strategic approach for vetting alerts—and follow-up with a consistent threat validation process—will make it much easier for their teams to manage the growing volume of work with existing or limited resources.