SPYDERBAT

Attack Tracing and Intercept

Fast and Accurate Investigation Automation

BY SPYDERBAT

Table of Contents

Forward by Richard Stiennon

We are approaching the limits of what a team of security analysts are able to do in a SOC. The problem is that we have been addressing the symptom, not the problem. The symptom is alerts triggered by multiple layers of sensors. The problem, of course, is a rising tide of attacks as threat actors increase their use of automation, targeting, and nefarious motivations. The solution is to have context, but not the context delivered by time synchronicity, geo-location, investing threat intelligence, or correlating all these disparate factors. The solution is to map your entire computing infrastructure down to the process and system call level. In other words, if you know what all of your devices are doing all the time you can match that activity with an alert and immediately be able to diagnose if the alert is meaningful, and what to do about it.

This white paper introduces the concept of Attack Tracing and Intercept (ATI), a completely new way to address alert fatigue, threat hunting, and incident response. As explained below, it means surveying and mapping the activity of all computing devices, not to generate alerts based on their behavior, but to have complete situational awareness. Any alert generated by any security tool can be immediately mapped to the device and the process that is impacted. This is shift-left for security operations.

How did we get here? For about an eight-year period between 1995 and 2003 there were two primary network security technologies: firewalls and intrusion detection systems (IDS). The firewall, if configured properly was effective at stopping network-based attacks. IDS sensors were passive devices that would monitor network traffic and look for signature matches with known malicious activity. Thousands of signatures were developed by an open source community and installed in the sensors, which not surprisingly became very noisy. Tuning IDS systems to eliminate false positives became a valued art, yet the tendency for most security people is to “collect it all.” Most organizations either ignored all the alerts or eventually outsourced the logging of alerts to a managed security service provider (MSSP), who would ignore them for you.

As alert management became a burden, new products were created to collect, de-dupe, and normalize alerts based on time stamps. These Security Information and Event Management (SIEM) platforms became the center of the security operations center (SOC). Tier 1 analysts would be tasked with looking at every alert and performing triage, often based on intuition, or a value judgement: “This might be something, but it is targeting a low value asset, so I’ll move on to the next one.” If an alert was deemed credible it would be passed up to a Tier 2 analyst for investigation. This analyst would identify the severity of the attack and kick off an incident response process that would often call for the top Tier 3 analysts to threat hunt and investigate. Their goal is to determine what happened, how to stop it, and how to prevent it from happening again.

To assist the SOC teams there are multiple technologies being deployed. Many SIEM products are adding machine learning or Big Data for example. This new breed of SIEM technology is often called Security Analytics, but it also overlaps with XDR, the combination of Network Detection and Response, and Endpoint Detection and Response. Even threat intelligence platforms (TIPs) are getting into the game because threat intelligence is a way to extend SOC operations out towards the attackers. If you have forewarning of an attacker's intent or tools and methodologies, you can apply that to your alert data and at least identify serious attacks.

The Response part of NDR, EDR, and XDR, is an attempt to automate the more mundane responses to attacks through orchestration (sometimes called Robotic Process Control (RPC) or SOAR, Security Orchestration, Automation, and Response.) Unfortunately, configuring SOAR systems means creating playbooks for every possible response scenario, for instance: too many attempted logins, therefore lock down the account and alert the user.

ATI holds the promise of reducing the load on a SOC team by providing immediate visibility into the cause of an alert. This is valuable context; not a correlation, but a causal determination that can be used to evaluate what happened and what should be done to either stop an attack in process or remediate one that has occurred.

The expense of operating a SOC is only increasing with the level and number of attacks. The demand for skilled security analysts is high and their salaries reflect that. SOC teams also suffer from burn-out because of the pressure cooker environment. More and more sophisticated tools are being deployed to handle some of the burden but they are expensive and demand even higher skill levels to configure and maintain.

Read on to discover how Spyderbat implements Attack Tracing and Intercept to serve as a game-changer for your SOC.

About Richard Stiennon

Richard Stiennon (@stiennon and @cyberwar) is one of the foremost industry analysts in cybersecurity and founder of industry research firm IT-Harvest. Mr. Stiennon began his own career in cybersecurity in 1995 at Netrex, one of the first Managed Security Service Providers (MSSPs) and covered the network security industry as a Gartner analyst for four years. He has held positions across several security product companies including Webroot Software as VP of Threat Research, Fortinet as CMO, and Chief Strategy Officer at data erasure company, Blancco Technology Group. In addition, Stiennon is the author of several books on the cybersecurity industry including Secure Cloud Transformation: The CIO's Journey, Security Yearbook 2020: A History, and most recently, Curmudgeon: How to Succeed as an Industry Analyst.

Introduction

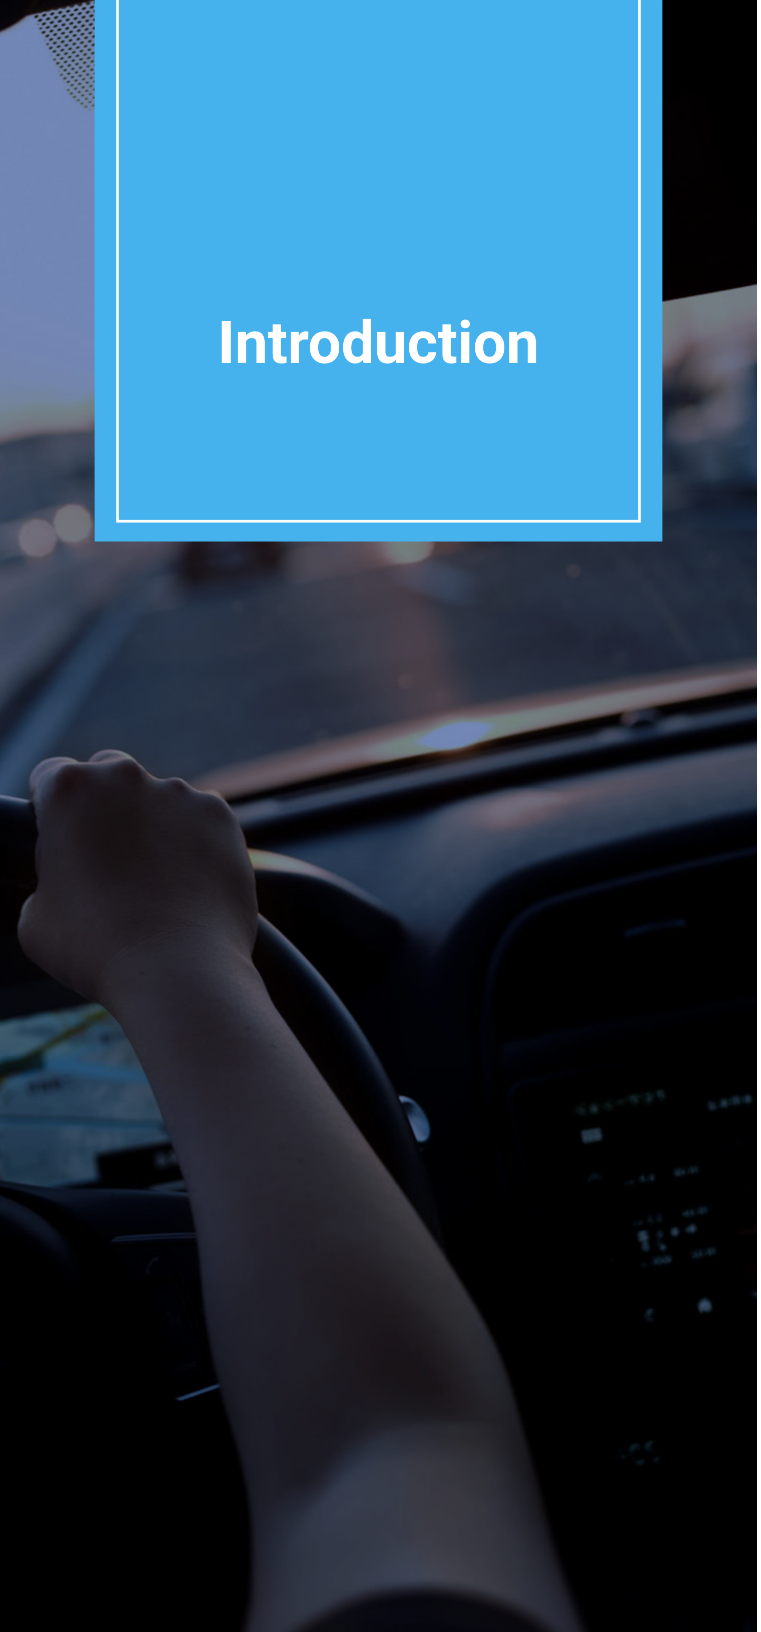

Pre-pandemic, I arrived at SFO, grabbed my rental car, and set my first meeting's directions on my phone. My phone battery was already low so I pulled over to dig out my charger only to realize I had left it at home. I was familiar enough with the city and assumed I could rely on my intuition of its grid-like streets to get to my destination. But shortly after my phone died I was sent down an unplanned detour. My memory wasn't as strong to recall which one-way streets faced which way. I became hopelessly lost and was running late. Acknowledging defeat, I turned into a gas station to ask for directions. I realized how dependent I was on using a map for navigation.

When an investigation begins, we also rely on our intuition looking through log data and security events to identify possible evidence. We try to retrace the attack's steps, attempting to recreate the attack's path through our networks, systems, and user accounts. What if instead of relying on log analysis, we were immediately presented with a clear picture of all the attack's steps? A "Google Map" that provided a clear, accurate, and complete navigation of the attack through each system it touched and user accounts it compromised since its inception.

This is the origin for Attack Tracing and Intercept.

The Chasm Between Detection and Response

The detection and response chasm characterizes the manual effort required by security analysts to investigate each alert in order to either dismiss it as a false positive or unveil the details of an attack to remediate it.

The analyst is seeking answers to these key questions:

- Is it real?
- What is the impact?
- How do I stop it and clean it up?
- How do I prevent it in the future?

The Chasm is measured in terms of depth (the number of alerts) and width (the time to resolve alerts) [Figure 1].

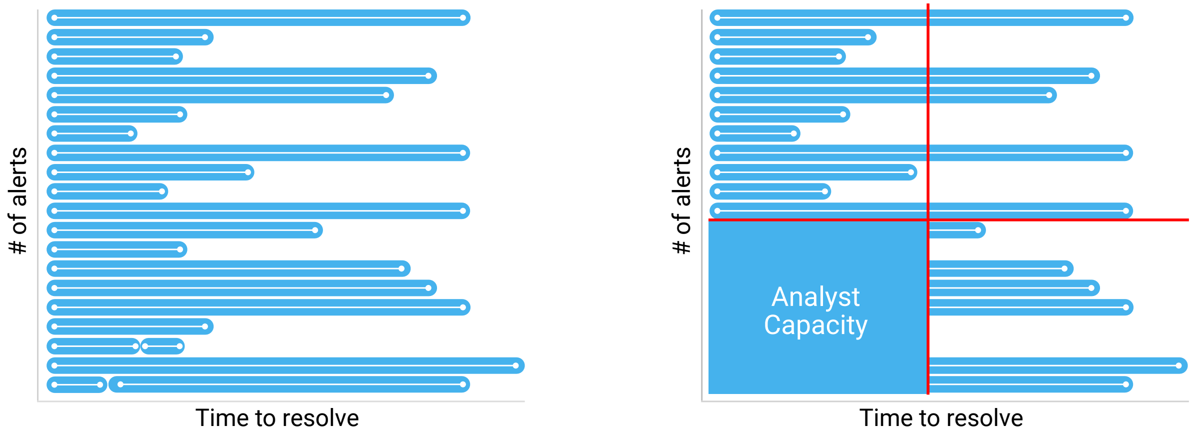

Figure 1 - Alert Workload and Analyst Capacity

The Chasm's depth and width surpass the analyst teams' capacity, forcing choices as to which alerts receive attention and creating time constraints on investigation. Without an ability to review each indicator and fully investigate potential threats, early threat activity is missed and not discovered until after the breach or late in the attack's lifecycle.

Tuning alerts out increases the risk of missing threat indicators and investigation time limits create uncertainty around the results.

The Chasm's Depth

Security teams continue to face large volumes of alerts mostly consisting of false positives. According to Fidelis' *The State of Threat Detection Report 2019*, approximately 67% surveyed felt alert overload is one of the main issues their teams face.[1] According to IDC, nearly half of alerts are false positive.[2] Due to the alert volume, analysts only triage higher priority alerts. Tuning alerts out increases the risk of missing threat indicators and constraining investigation time creates uncertainty in the results.

[1] The State of Threat Detection Report 2019, Fidelis, 2019
[2] "The Voice of the Analyst, IDC/FireEye, 2021

The Chasm's Width

When analysts do find sufficient evidence to investigate a threat, they are challenged to manually re-trace an attack to its origin due to their reliance on inference and incomplete data.

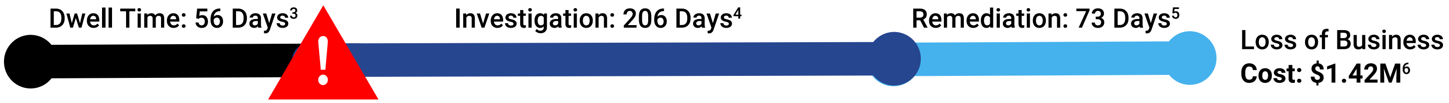

Ponemon quantifies the width of the chasm, observing that the average Dwell Time (the time between when an attacker enters your environment to the time it is first detected) is 56 days! The Dwell Time is before the investigation even starts. The long Dwell Time is a result of missed signals (ignored low priority alerts) or compressed, incomplete investigations of early signals that were dismissed as false positives.

Regarding the investigation, "on average, companies required 207 days to identify and 73 days to contain a breach in 2019, combining for an average 'lifecycle' of 280 days."[7]

The primary issue contributing to the manual investigation challenge is an excessive amount of data and the wrong data. In a survey conducted by CyberEdge, "too much data to analyze" was identified as one of the top three barriers to establishing effective defenses. As attacks cross systems, networks, users, and span across long periods of time, the more voluminous, innocuous data hides the attack's steps from analysts. Furthermore, even with time, analysts do not have access to the right data to identify causal connections. Finding the true causal connections vs correlation assumptions requires the right data to determine linkages. This makes the accurate reconstruction of events not just difficult, but impossible.

[3] Mandiant *2020 M-Trends Report*
[4] Ponemon *2019 Cost of a Breach Report*
[5] Ibid.
[6] Ibid.
[7] *2020 Cost of a Breach Report*, Ponemon and IBM Security, 2020
[8] *Cyberthreat Defense Report 2020*, CyberEdge Group, 2020

Recent trends act as multipliers to this challenge:

- **Attackers' deliberate attempts to obfuscate steps.** *Crowdstrike Front Lines 2020 Report* acknowledges "dwell time increased due in part to advanced adversaries employing countermeasures, allowing them to remain hidden longer."[9] This creates significant and often insurmountable challenges in performing historic searches from the time of an alert to retrace the attack steps.

- **New layers of abstraction from virtualization in cloud environments and cloud-native applications.** Oracle and KPMG found 92% of organizations in their 2020 survey had a "cloud security readiness gap."[10] Without a stateful representation of a system at the precise point in time it was involved in an attack, a security analyst loses the ability to retrace the attack's trajectory.

The combination of alert load and manual inference within the investigation process creates a constant state of anxiety—what did I miss? In IDC's *The Voice of the Analyst Report*, 75% of security analysts are worried about missing incidents[11].

75%

of security analysts are worried about missing incidents[11]

[9] *Crowdstrike Front Lines 2020 Report*, Crowdstrike, 2020
[10] *Cloud Threat Report*, Oracle and KPMG, 2020
[11] *The Voice of the Analyst*, IDC/FireEye, 2021

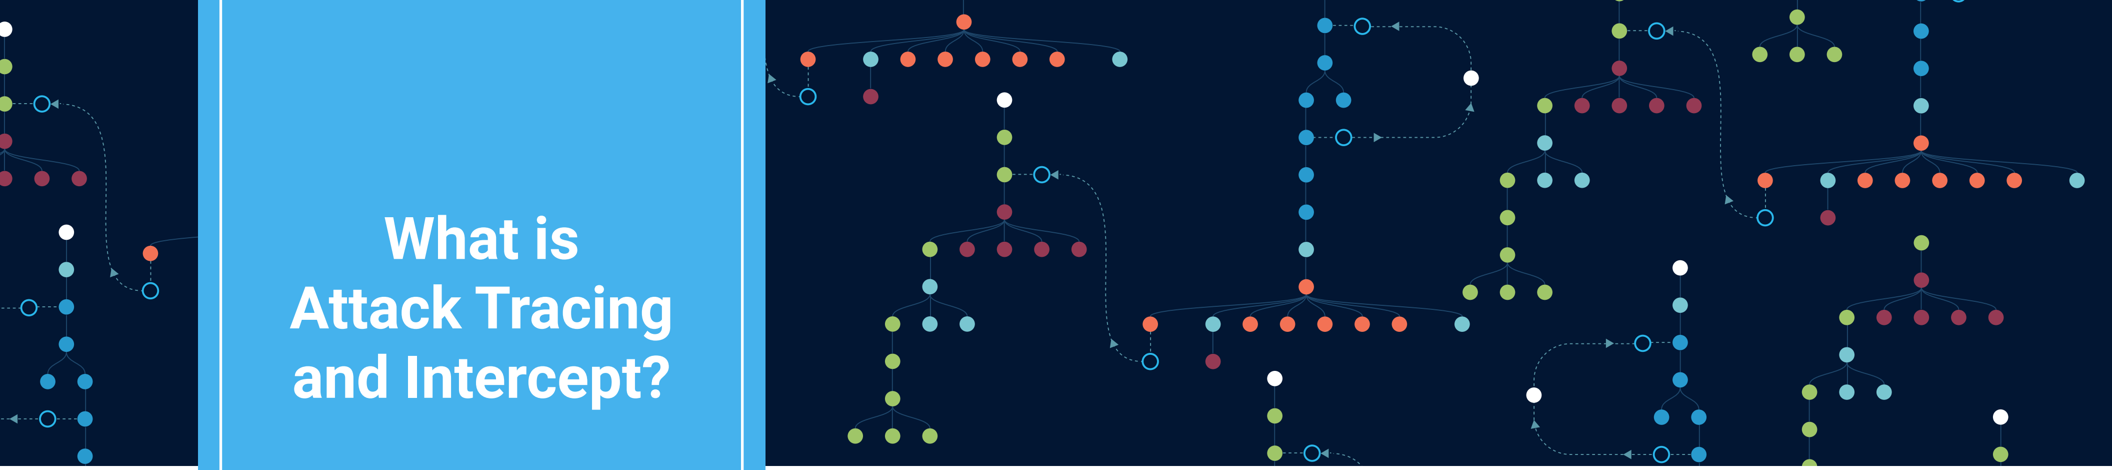

What is Attack Tracing and Intercept?

Attack Tracing and Intercept (ATI) uses ground-truth data to provide a focused view of all causal activity of an attack, regardless of whether the activity appears malicious or benign, in order to stop the full threat early in its lifecycle.

For example, if an attacker succeeds in getting remote control of a target system, all follow-on activity, including just “looking around”, lateral movement, or installation of a backdoor, is critical to uncover in the investigation to fully remediate the attack and prevent similar threats. Causal activity also includes preceding events. The objective is to track the attack to its origin which precedes the first detection.

ATI provides the linkages for all activity that precedes and follows the moment an attack is detected, regardless of the time span (e.g., attack’s may pause for weeks). To have a complete picture requires capturing all activity that is useful for the investigation, including causally connected network, file, process, and authentication activities.

Causal Connection
The link between two activities based on a causal relationship. Activity A caused Activity B. An example, an apache webserver process receives a connection, and then executes a php script. Causal connections are determined based on ground truth data.

Ground Truths
A system-level transaction between an application and the operating system.

Spyderbat's Attack Tracing and Intercept includes three key components:

The Universal Trace

is the ground-truth foundation of ATI which continuously collects and assembles all activity within and across systems. The Universal Trace proactively establishes all causal connections to produce a real-time enterprise-wide causal graph. This is independent of any security detections and serves as a complete, living historical record of all causal activity.

Security Fusion

maps security information to the Universal Trace, fusing in real-time to the ground-truth causal graph to instantly illuminate attack's paths.

Context Enrichment

Additional context is flagged on the Universal Trace to uncover additional attack details. These context flags (e.g. threat intelligence, internal events) do not warrant investigation when seen in isolation. However, they are critical to understand the attack's steps and behavior when found causally connected to security detections.

The Universal Trace differs from existing solutions by proactively establishing and maintaining causal connections across all activitiy instead of reacting to an alert. Other solutions depend on threat detection signals to begin 'recording' information. Proactively building the Universal Trace enables immediate and accurate visibility to the complete, focused attack trace at the onset of the security investigation.

Figure 2: The Universal Trace fuses security data and context with ground truths as they occur, providing a map of causal connections that trace back over time, systems, and user sessions.

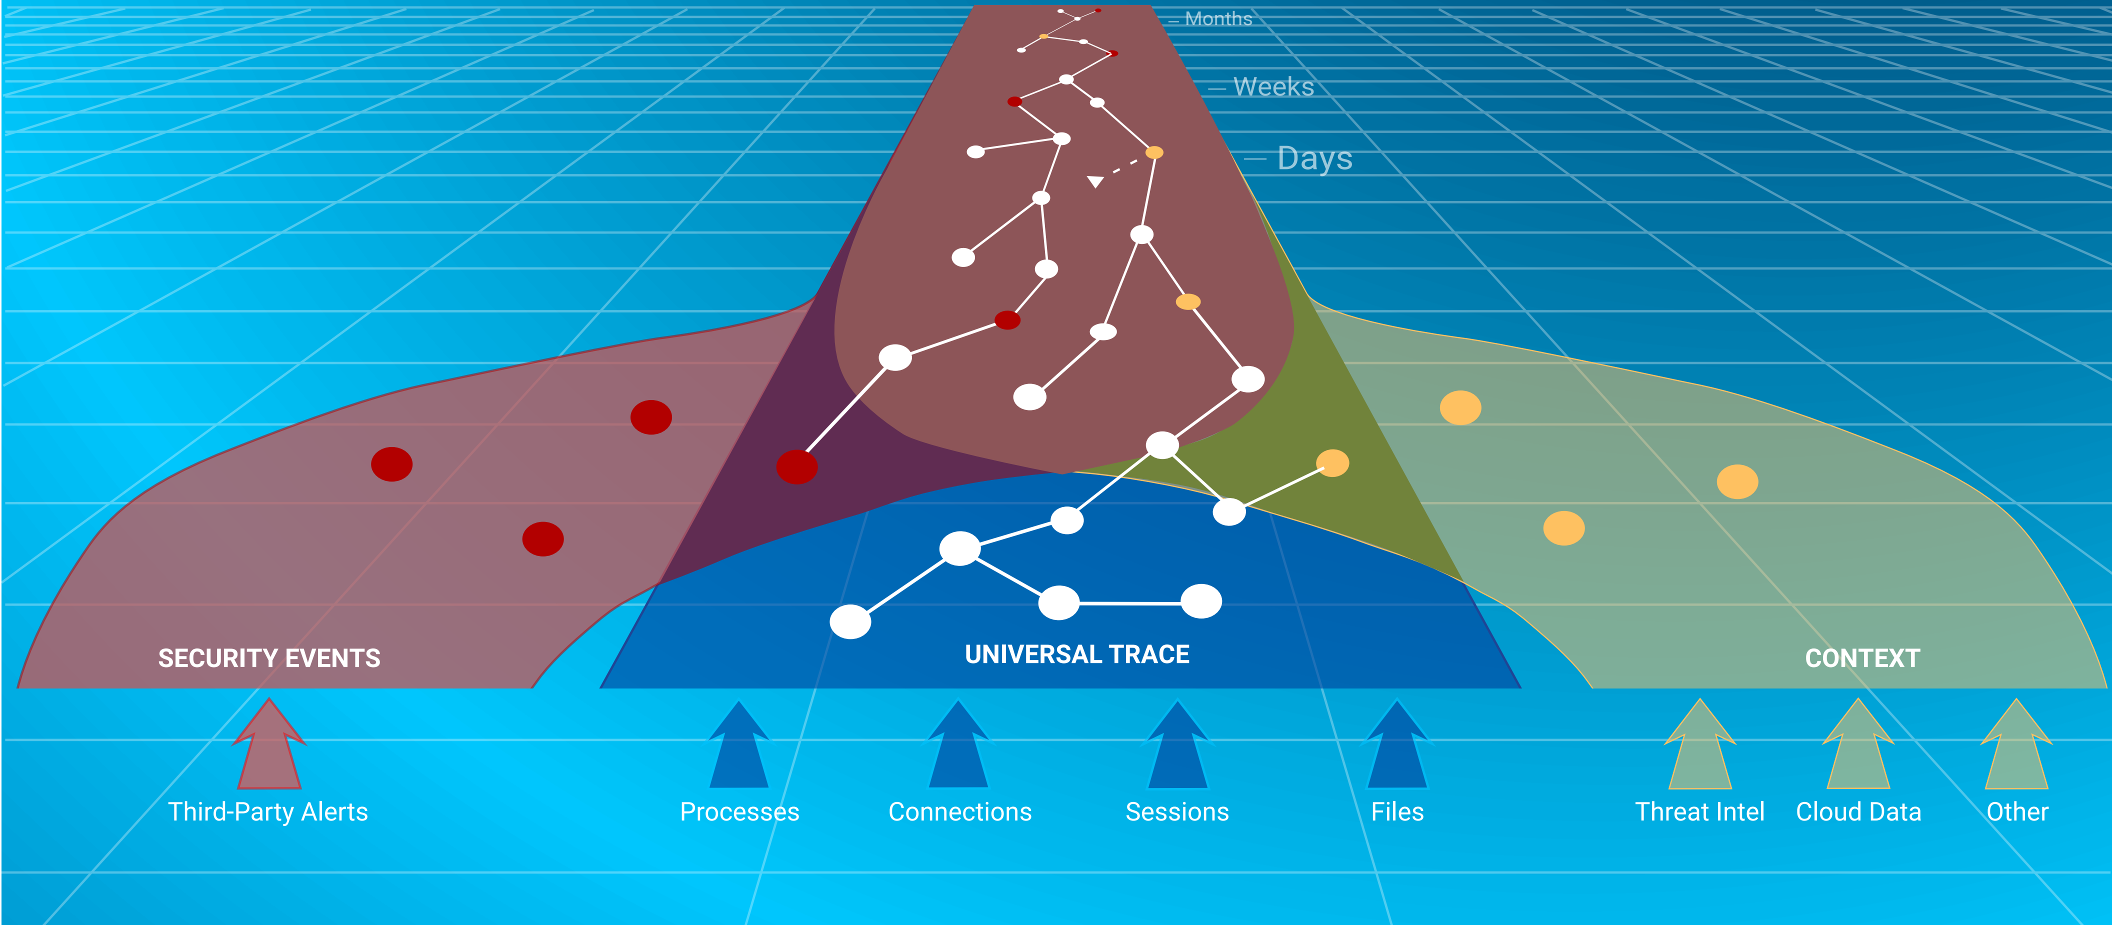

The Universal Trace

Spyderbat's Universal Trace changes the game. Alerts and security events no longer exist in isolation. They are instantly connected in context and easily viewable by the analyst. The real-time enterprise-wide causal graph eliminates manual effort and uncertainty related to log searches and endless pivots performed in the attempt to reconstruct what may have happened. The statefulness of Universal Tracing retains the deep causal history to capture potential false negatives. False positives are instantly recognized as dead-ends with no causal activity.

The Universal Trace is constantly generated and maintained proactively for accuracy, completeness, and immediate availability. It is created by capturing every interaction of every program with the operating system as they occur:

- Intra system activity e.g., user sessions, processes, and file access
- Inter system activity e.g., bi-directional internal and external network connections

Attempting to reconstruct universal tracing manually is not only extremely complex and time consuming, but highly unreliable as log and audit data are incomplete sources and simply do not contain sufficient level of detail. Instead, Spyderbat captures ground-truth system-level data from each interaction with the operating system and automatically establishes the causal links between these interactions. These include both incoming and outgoing network transmissions that span across internal and external network connections. The result is a complete, living, ever-changing graph of every transaction within and across the organizations' physical and virtual environments.

Universal Tracing Example

1. The system is named boco.us-west.internal
2. By booting, spawns a process called systemd
3. Systemd runs with the permissions of the root user
4. Systemd spawns a process called sshd
5. sshd accepts a connection from an external IP 172.16.22.51
6. The connection authenticates as ec2-user
7. By logging in, ec2-user spawns a bash process
8. In the bash session, ec2-user runs several commands
9. Ec2-user runs the command ssh to login to the system named atx.us-west.internal as bsmith
10. Bsmith runs several commands in a bash shell

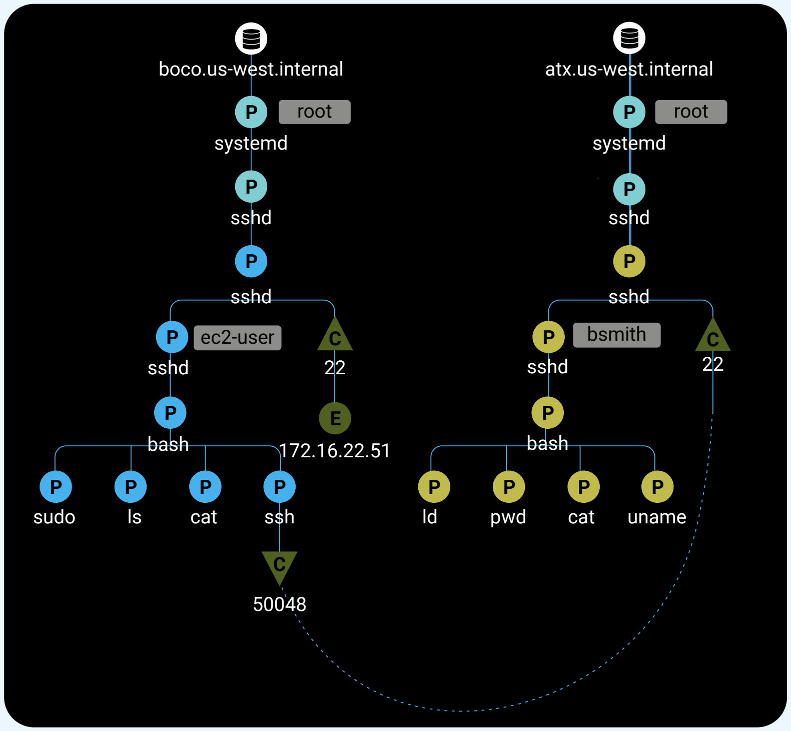

Figure 3 - Simple Causal Tree from Universal Tracing

Security Fusion and Attack Traces

The Univeral Trace is annotated with flags highlighting suspicious and malicious activities. These flags come from Spyderbat and optionally from third-party sources such as SIEM, cloud workload protection programs, NextGen Firewalls, and Endpoint Detection and Response (EDR). The Universal Trace links seemingly disparate security events and seemingly benign activity. The attack trace presents a clear and complete depiction of each causally connected step of an attack prior to and following a security alert. This causal chain may span across systems, user sessions, and even long periods of time (e.g., months).

In this example, the attack trace captures:

1. Back at the point of entry, a command injection into a vulnerable php page, which occurred several weeks prior to the alert initiating the investigation.
2. Command and control traffic from a backdoor application left on the system of the initial intrusion.
3. Lateral movement to a new system with different user credentials.
4. Any file that was created, read, or changed by the attacker.

Example of Attack Trace

Flags indicating suspicious and malicious activity are linked together by the Universal Trace. This focused view within the Universal Trace tells the 'story' of what is happening during an attack.

1. A file is written then executed from a command-insertion on the web server of atx.us-west.internal
2. The service account www-data starts a shell
3. The service account www-data sudo to user gibson
4. The user enumerates commands that can be run as root
5. The user wrote to another user's .ssh directory
6. User scanned for files with setuid bit set
7. User executed a suid file.

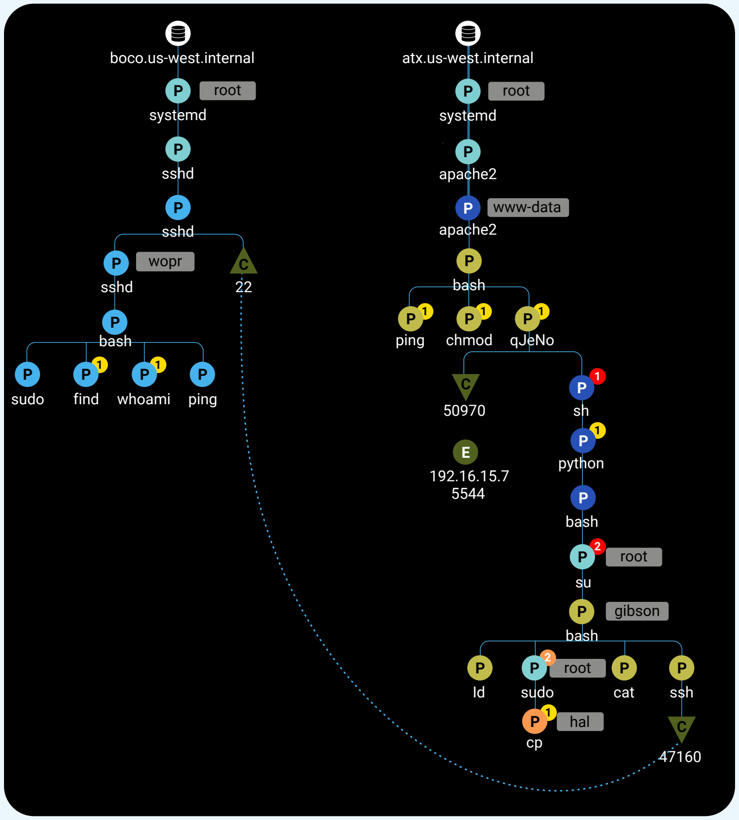

Figure 4 - Causal Tree Displays the Attack Trace

If the Universal Trace is like Google Maps, Security Fusion is like the GPS adding navigation to each stop.

Attack tracing gives an analyst:

The Story

All activities of the attack including root cause

Impact

The complete scope of the attack including compromised systems and user accounts

Intercept

The knowledge to immediately and completely stop the attack.

Cleanup

By understanding the attack's complete scope, how to fully remediate its impact.

Prevention

With the full understanding of the attack, the organization can update security controls and processes to prevent similar, future attacks

Attack tracing—the fusion of security data into the Universal Trace to identify causally connected malicious activity—results in four benefits:

1. **Immediate recognition of a false positive**
2. **Immediate recognition of full attack's steps** including the entry point, even if this occurred months previously.
3. **Capture of would-be false negatives** by capturing any new causally connected activity to resurface the attack trace.
4. **Early intercept of legitimate attacks** before significant damage.

Proactively creating the Universal trace with real-time fusing of security information and context enrichment results in accuracy, speed, and completeness:

	Without ATI	With Spyderbat ATI
Accuracy	○ Inaccuracies incurred from manually inferring correlation of activities from reviewing inscrutable and incomplete log and audit data in pivot search results. ○ Human error from investigation fatigue, skillset, and time constraints ○ Missed detail from security controls only 'recording' activity after observing known signals	Identifies attack steps using the Universal Trace's causal connections, built proactively with stateful ground-truth data. The attack attack trace starts an investigation with a focused view of the attack, removing innocuous, distracting activities that otherwise obscure evidence.
Speed	○ Investigations impeded by manual investigation and laborious attempts to reconstruct events from extraneous and incomplete data. ○ Missed early threat indicators from only investigating high-priority alerts, allowing attacks to grow in scope and complexity.	The Univeral Trace immediately suppresses false positives by exposing alerts with no causal outcomes. The attack trace compresses investigations of credible attacks by instantly providing precise and complete details of the attack's steps. Enables detection at earlier stages of an attack (reducing/eliminating dwell time) to intercept and remediate ahead of severe damage.
Completeness	○ Full threat remediation is challenged by time constraints and incomplete data to be certain the investigation has exhausted all activities by malware/threat actors.	The Universal Trace ensures full remediation by linking investigations to all causally connected activity back to its origin point, even if days, weeks or months apart. Protects against false negatives by linking previous, causally connected alerts together.

Use Cases for Attack Tracing and Intercept

	Description	Benefit
Security Investigations	Remove manual effort analyzing log data by presenting a focused trace of activities leading to and following the alert, including associated security information.	Radically compress investigations with complete and accurate representation of the attack’s steps through time and across systems, including its entry point to enable early and full intercept.
Root Cause Determination	Search systems or processes experiencing irregularities or operational challenges to recognize root issues faster by its attack trace	Immediate visibility to root causes and impact for fast and accurate remediation
Alert Triage	Identification of false positives and true positives based on causally connected activity, including other alerts	Eliminate time identifying false positives and increase efficiency by investigating causally connected alerts together
Democratizing Investigations	Automate the complex manual task of reconstructing attack steps and invite collaboration with DevOps, Developers and Application Owners familiar with application and system behavior to review an easily consumable format.	Facilitates alert triage for junior analysts and enables earlier threat detection and intercept by involving internal expertise familiar with the intended behavior of their applications.

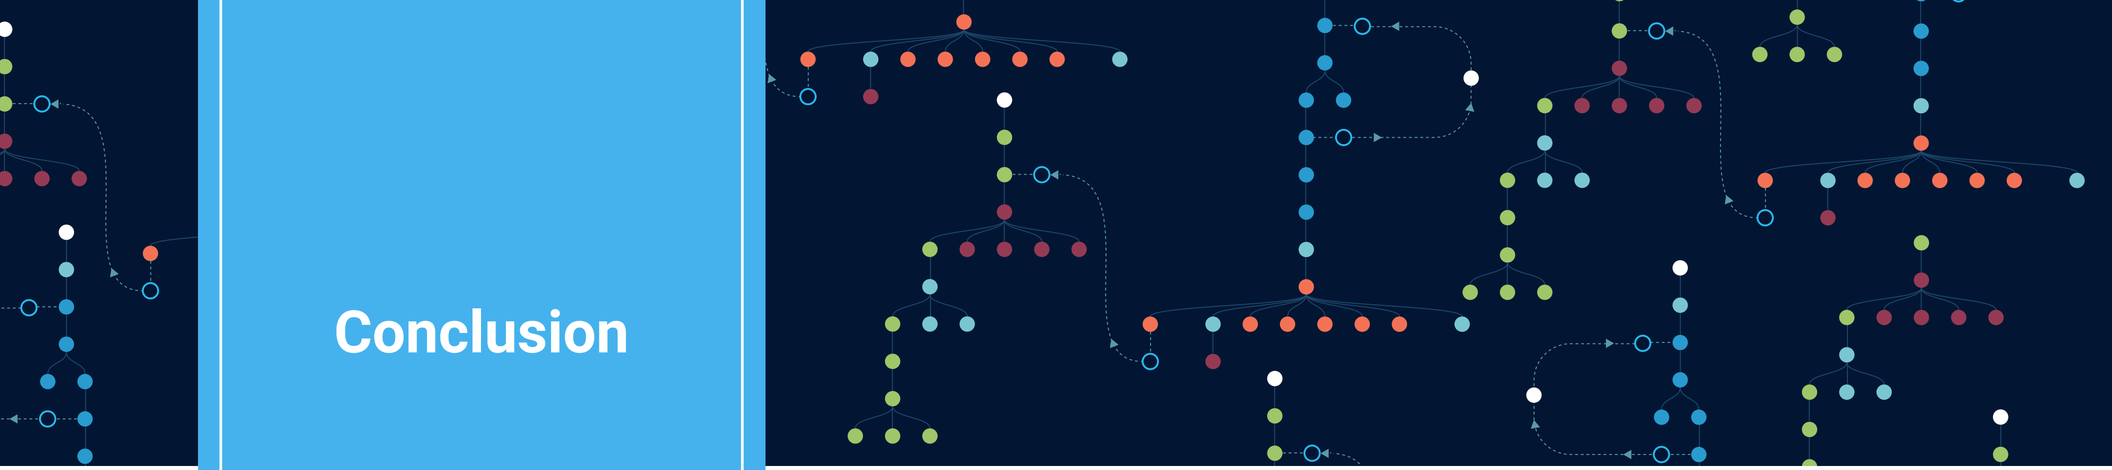

Conclusion

Attack tracing and intercept introduces revolutionary innovation to profoundly reimagine security operations by automating the most challenging manual aspects of security investigation. The unique Universal Trace acts as a springboard for continued innovation, enabling constant assessments of causal connections of ground-truths within and across systems to rapidly identify and preemptively intercept attacks.

ATI radically compresses the capacity required for performing alert triage and investigation by:

- Immediately dismissing false positives without risk of a false negative
- Combining alerts that are causally connected
- Crushing the time needed to investigate and intercept threats by immediately providing analysts with accurate and complete attack traces

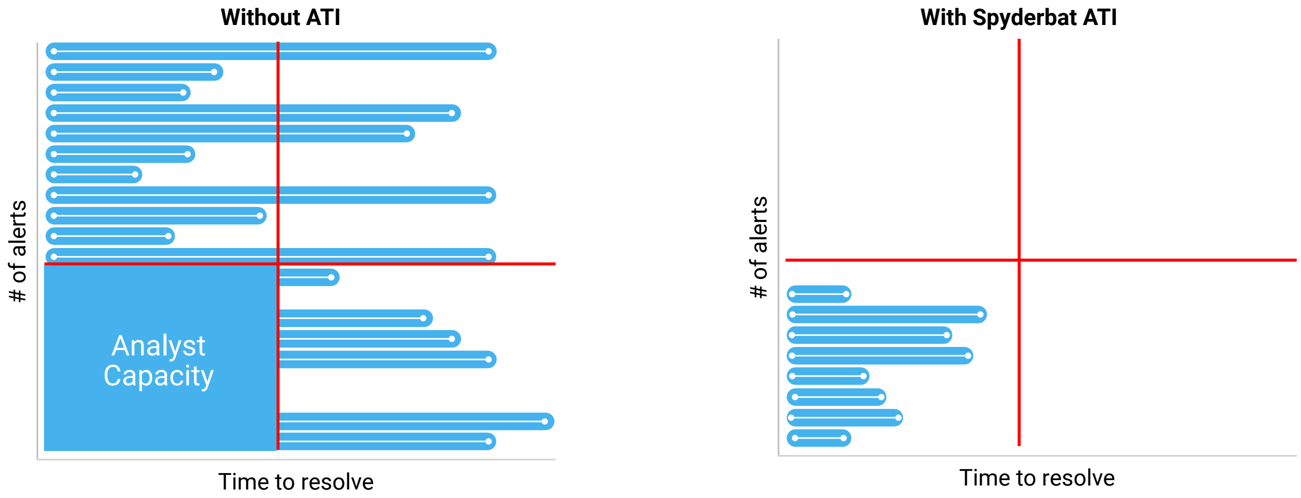

Figure 5 - Alert Workload and Analyst Capacity Updated with ATI

ATI eliminates the anxiety due to missed alerts or overlooked evidence. Spyderbat ATI allows security analysts to quickly verify threats, their root cause, and scope while avoiding time spent investigating false positives. Analysts are armed with the visual acuity to discover and intercept attacks early, before damage occurs.

About Spyderbat

Spyderbat introduces the industry’s first attack tracing and inception tool to radically change the way organizations handle early threat discovery and investigation. Spyderbat is backed by LiveOak Venture Partners, Benhamou Global Ventures, and cybersecurity veteran John McHale. To track Spyderbat’s progress, please visit **spyderbat.com**