IronNet

WHITE PAPER

Cybersecurity Market Insights:

A practical way to rule out false positives

By Dean Teffer, PhD, IronNet Vice President of Detection and Prioritization

Dean Teffer, PhD

IronNet Vice President of Detection and Prioritization

In his role at IronNet Cybersecurity, Dean oversees the Threat Research, Data Science, and Data Engineering teams. He brings more than 20 years of research and development and product engineering experience to bear on the challenges of improving relevancy and context of surfaced information to customers and increasing coverage while reducing time to value.

Prior to joining IronNet, Dean:

- Was Director of Data Science at JASK, a cloud-based NTA and SIEM, acquired by Sumo Logic
- Developed anti-submarine warfare algorithms and software for the U.S. Navy, countermeasure systems for NAVAIR, and cyberdefense / counter-intelligence systems for the U.S. Intelligence Community
- Led engineering at two other Austin-area startups, including founding one acquired by Siemens
- Obtained a PhD in Computer Engineering and a Masters in Physics from The University of Texas at Austin

There is no question that alert fatigue and staffing shortfalls continue to plague security teams. Nation-state cyber attacks have doubled over three years, and, today, highly organized cyber criminal groups are increasingly backed by nation-states. Even relatively unsophisticated attacks, but with big impact, are on the rise in large part due to the pandemic. McKinsey recently reported "a near-sevenfold increase in spear-phishing attacks" since the COVID-19 cyber chaos began.

Security operations teams simply cannot keep up. McKinsey also noted that 60% of enterprise-level SOC analysts can triage only less than 40% of their enterprises' log data. This is a function of both the complexity of modern IT systems, which include hybrid on-prem and cloud, each with targeted security coverage, and the relative rarity and high cost of trained security professionals with the skills to triage and investigate across all these ecosystems and tools. Malicious threats are going undetected and/or uninvestigated. Indeed, visibility is limited in the murky waters of the vast cyber sea.

An increasingly large ecosystem of cybersecurity products has done little to mitigate these challenges. In fact, installing more monitoring products makes "the alert cannon" even worse, flooding the SIEM with false positives. All the while the true cyber pearls for operators and threat hunters remain hidden in the dark depths. How can security operators more easily crack the code to gain full visibility of urgent, actionable threats? Do advancements in artificial intelligence (AI) and machine learning (ML) really help real threats rise to the surface?

Despite all the promises of machine learning (ML) and algorithmic threat detection, these technologies notoriously still yield so many false positives. Why? We experience high-quality algorithmic results every day using natural language processing and image processing. So why not threat detection?

As I see it, there are two primary reasons why typical threat detection is still trying to boil the ocean:

1 Actual threat events are very rare.

This might be a shocking statement to make in a post-SolarWinds, post-Log4j cyber world. But what this means in average times is that a detector that always yields a "no threat" label would be accurate almost all the time. In order to (sort of) turn up the gain on the detector to make sure we do not miss any potential threat events, we necessarily incorrectly label some events as "threat" when they should not be. We need to close the book on this chicken little approach.

2 An apple is not always an apple in security.

Even if we do tune algorithms well for both detection of all actual threat events (recall) and simultaneous minimization of false detection (precision), unlike image recognition or text, an apple is not always an apple in security. Outside the narrow scope of "known-knowns," signature-based detection, behavioral anomalies, and threat signals are manifest in real-work networks in myriad ways, which frequently vary based on specific configurations of infrastructure, IT policy, and user conventions. Accordingly, we need a way to know which apples are the ripest—the best ones to eat first.

60%

of enterprise SOC analysts triage

<40%

of log data

— MCKINSEY

Indeed, visibility is limited in the murky waters of the vast cyber sea.

A practical way to reduce false positives

Fortunately, there is an activity that is both essential to threat hunting and directly in support of reducing false positives (truly putting marketing claims to the test): identifying corroborating evidence. This is, in fact, one of the primary activities of the threat hunter and SOC operator. An alert that fires on a network log, however, indexed by IP address, is difficult to correlate to information in the Active Directory log, let alone the AWS security log. This is why it takes so long, and why the activity requires so much expertise that, as already mentioned, is difficult to recruit.

But what if all data were tagged on ingest with a device or user ID, regardless of data source, and any information across the ecosystem related to entity association (user authorization, device registration, etc.) were tracked and recorded, so that ALL events detected on a device or associated with a user could be not just searched but also automatically combined? This would, in fact, be the set of corroborating evidence a threat hunter is trying to assemble. Also, such a set of events could inform a model-driven probability of the likelihood that not just one event but a whole sequence of events, comprising, say, three distinct MITRE ATT&CK® stages, has been detected within the past two days. Such a likelihood function would yield a more robust and mathematically defensible measure of severity and confidence. That is the power of the IronNet threat engine.

Fine-tuning behavioral-based detections of network threats

The correlation of detection analytics is one thing. Behavioral analytics *enriched by human insights* is another altogether. This scenario gets us closer and closer to minimizing false positives and reducing the margin of error. Further still, correlation *across* SOC analyst teams in a Collective Defense ecosystem, as the IronNet Collective Defense platform allows, drives home the difference between crying wolf and an urgent and real need to batten the hatches against the real wolves lurking on the network.

Threat analysts and hunters spend a significant portion of their time triaging individual alerts by manually identifying corroborating evidence and related information. Given the threat volume and strained resources, they need a way to cut to the chase: Which alerts are meaningful? Which ones are the priority? IronNet's answer is a threat engine that is embedded with human threat intelligence. The goal is to send all those false positives through a cyber sieve so that only correlated, actionable alerts rise to the surface—early in the kill chain to enable threat mitigation well *before* business impact.

Detecting adversaries “left of boom” with behavioral analytics

It remains true that signatures are an important part of the detection equation, but they represent only known knowns. It is broadly accepted in the analyst community that a cyber attacker can easily change hash values, IPs, and domains (signatures)— the three lowest levels of David J. Bianco's "Pyramid of Pain" Threat Hunting Framework.

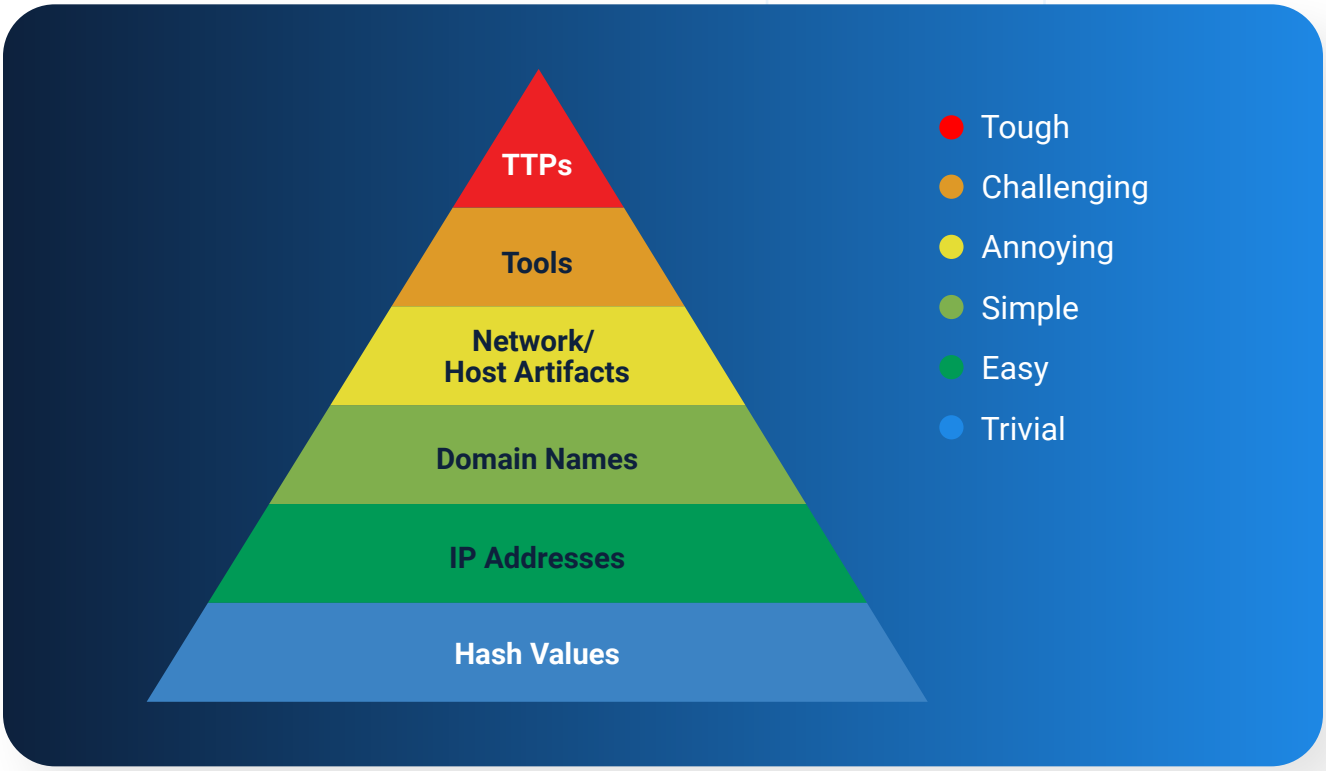

As a result, chasing after known signatures, or Indicators of Compromise (IoCs), will take threat detection efforts only so far. Tactics, techniques, and procedures (TTP), which really boil down to adversarial "behaviors" on the network, however, are hard for adversaries to change. Therefore, TTPs are the best type of indicators for defenders to focus on—vs. looking for known IoCs.

Behaviors require a data-processing infrastructure purpose-built to analyze sessions and entity lifecycle aggregations of data. Behavioral analytics like this are core to IronNet's IronDefense NDR solution. They analyze network traffic and identify likely malicious activity.

IronNet analytics cover the kill chain from reconnaissance through action. Most are focused early in the kill chain in order to detect "left of boom"—that is, before business impact. Think about it: the wrong time to detect ransomware is when your screen goes black and you see a skull and crossbones with an ominous note that says, "Hey [fill in name], if you want your computer back, pay me X Bitcoin?" No one would argue that that's the wrong time to detect

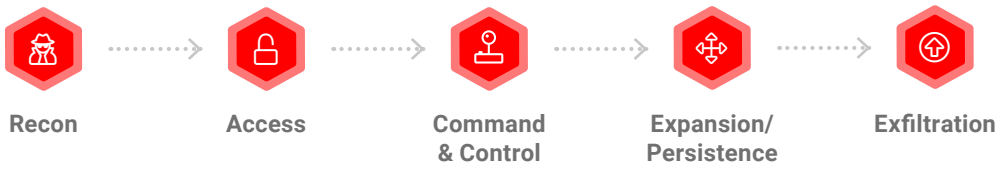

The right time to detect is in the early reconnaissance stage, whether in the initial access stage (best) or during the command and control stage (better) as they are communicating with their malware they embedded into your network, moving laterally in your network to have a bigger impact (bad). It's early in the kill chain that the detection is important, so that's where we focus most of our analytics.

We've learned over the years, though, that detecting alone isn't good enough. Just because you have detections doesn't mean you have actionable detections. You need *correlation-based* detections.

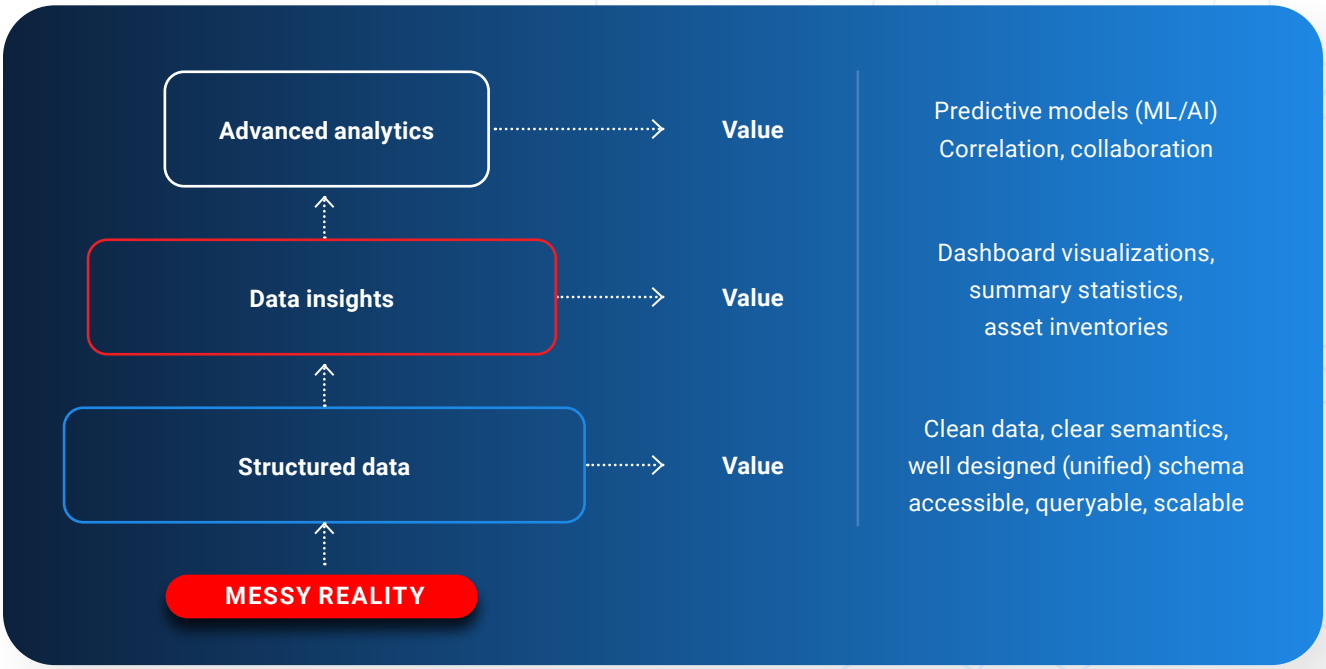

Modern data processing systems are based on a hierarchy of sequential processing steps. One-off or single-use analytics can be used for occasional needs, but for embedded, real-time analytics, one must automate a series of successively more sophisticated processing steps, each of which builds on the step that precedes it. For image recognition, this means detecting edges and colors first, then shapes, then objects, and only then classifying the image. For security, this means starting with knowns before proceeding to unknowns.

The knowns of cybersecurity are signatures, or IOCs. Signature-based alerts are table-stakes that also have a relatively low false positive rate. However, they are static and reactive, and therefore have low efficacy for new or evolving threats.

The next level is behavioral detections. These are necessarily based on more data, such as history, statistical measures, and lower-level signatures or "factual" events. Unlike signatures, behavioral detections can identify novel events and new attack types or events that vary by situation or configuration. However, they are relatively high false positive.

The next frontier in security builds upon the prior levels by automatically correlating signature and behavioral detections. These can yield dynamic, novel, and situational detections, but with relatively low false positives. In addition, since all available relevant information is automatically collated in advance, these alerts provide actionably content, unlike the simpler events produced by signature and behavioral detection.

Signature-based alerts are:
- static
- reactive
- relatively low-false positive rate
- relatively low efficacy

Behavior-based alerts are:
- dynamic / situational
- relatively high-false positive rate
- relatively high efficacy

Correlation-based alerts are:
- dynamic / situational
- relatively low-false positive rate
- relatively high efficacy
- automatically provides actionable context

“CODE-ifying” human threat intelligence

As one who lives and breathes in the world of data analytics, I will continue to shout from the rooftops that analytics without human intelligence will never cut through the cacophonous noise that every SOC analyst is forced to endure. Why not? It’s simple: their tools are missing the “So what?” that can calm even the noisiest of SOCs.

The secret ingredient here is human smarts. In the case of IronNet, our automatically correlated alerts are infused with the cumulative human intelligence of elite, Tier 3 analysts and threat hunters, allowing us first to fine-tune the detections themselves and, second, pre-package corroborating investigation information with the automated alerts. In other words, we have “CODE-ified” human intelligence to allow analysts to prioritize actionable alerts and pivot quickly to triage.

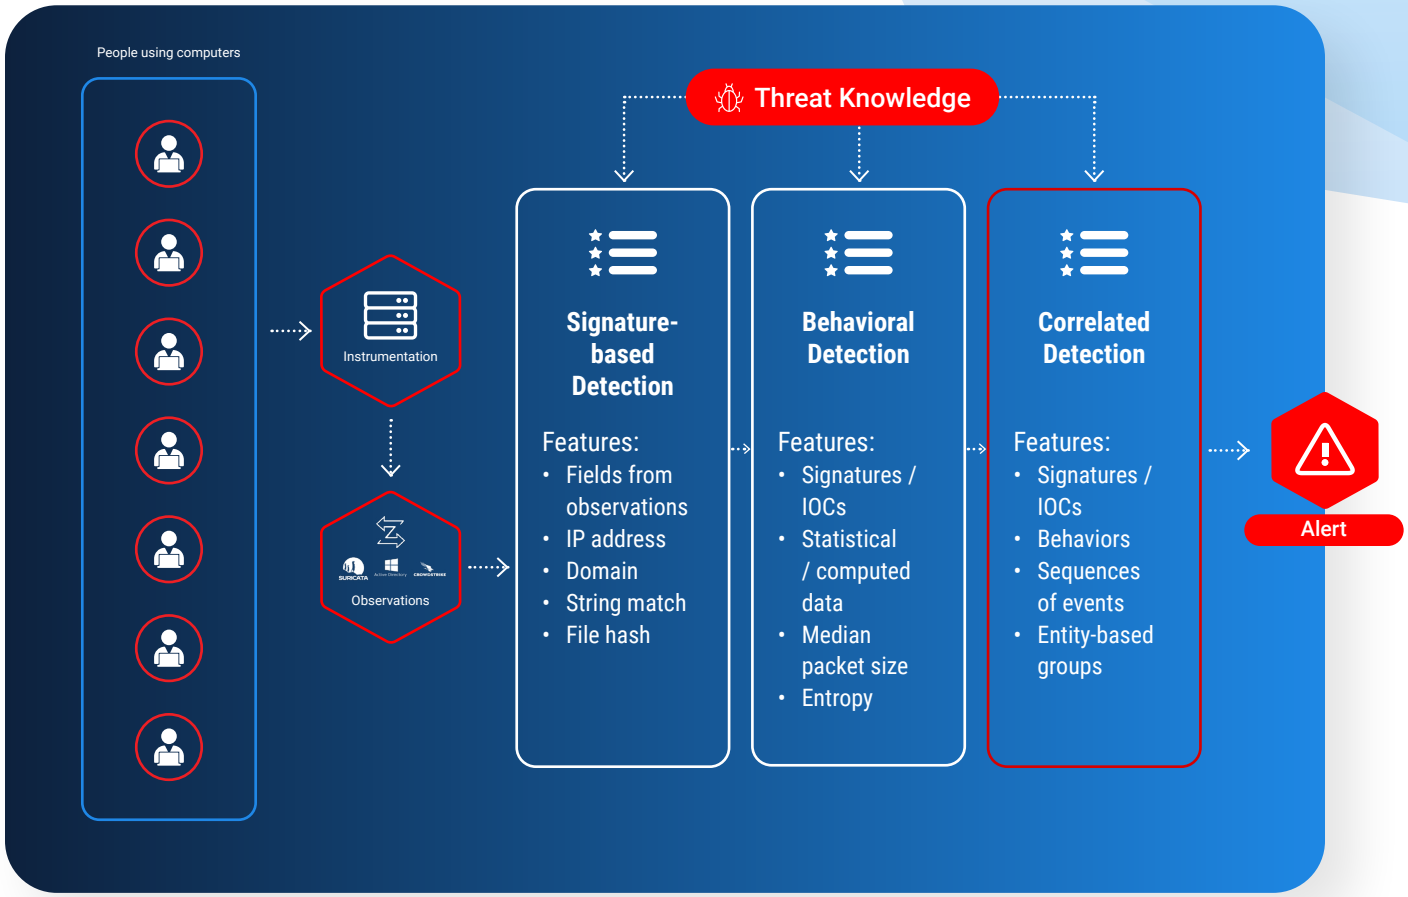

Leveraging human insights advances the power of analytics, turning the messy reality of SOC detections to more predictable, relevant, and actionable threat alerts.

Consider the way diagnostic tests work for healthcare. A cancer screening may be only 80-85% accurate, for instance. How do you fill that very relevant gap? The doctor. She or he applies human insights and intelligence to best mitigate the possibility of false results. They look at the cancer risk from all angles: patient lifestyle, family history, etc.

The fundamental approach of what the doctors do transcends medicine. This strategy is called improving the signal-to-noise ratio. What's more, if a doctor orders another diagnostic test, they are arming themselves with a way to look at the situation from all angles. In the same way, IronNet's automated correlation-based detections triangulate alerts (and sometimes bring together even more than three alerts) to package them as one incident. It's like opening a puzzle box only to discover that half of the puzzle is already put together. Think about how much tedious time you're saving by knowing immediately where to focus your attention.

It's no wonder that SOC burnout (and, now, the "Great Resignation") is a serious concern. All the data cleansing and threat analysis takes precious time. Not only does IronNet solve this problem by correlating alerts; we also corroborate information about the alerts by putting together the story from all points of view. Is there HTTP alert information? Did the endpoint alert? Are there firewall rules about the domain? And so on. Single, discrete directions now take on clear meaning. Kind of like the way foreshadowing works in a novel—all those nagging elements that don't shine the full light until you get to the end of the book. With IronNet's automated threat engine, you don't have to read the *War and Peace* version of your alert dashboard.

From single, discrete observable detections:

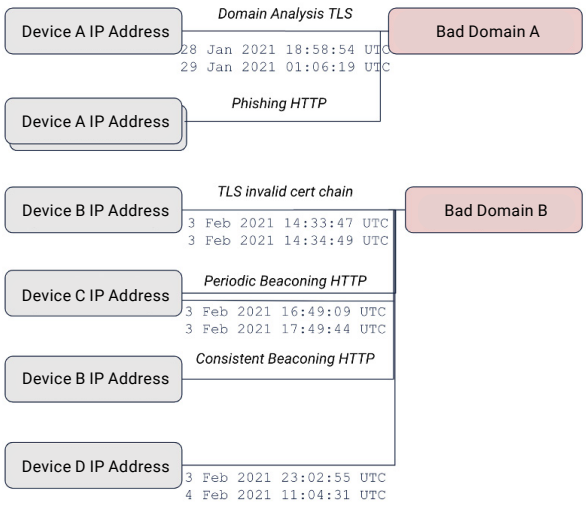

To automatically correlated observable events in a single incident:

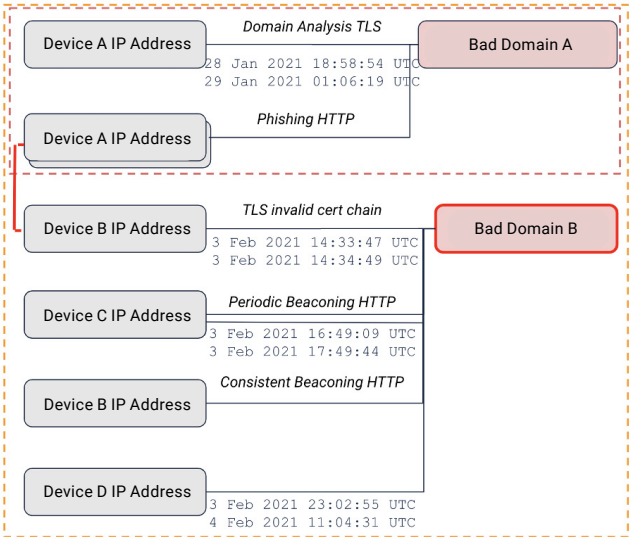

What this means is that analysts are not just getting alerts; instead, they are getting "stories" with each alert.

Every alert tells a story.

There's a mathematical concept I'd like to share: orthogonality. What this means is looking at something from different angles. Take a very high-pixelated cardboard cutout of your favorite basketball player or actor or musician. Looking straight on from afar, you may be hard pressed to know whether or not it's real. Look from the side, though, and you get a different perspective. No need to race ahead to claim your brush with fame.

In our cyber world, take an activity that looks like C2 communications. Could it really be an innocuous system update instead? Or does the alert link up with another suspicious activity like an unusual email with a strange link or attachment? And, now what: a beacon? While seeing a beacon in your network activity most certainly is often normal, if you string it along with the other two alerts, then, Houston, you have a problem. A clear story rises to the surface.

How it Works

Automated Correlated Engine

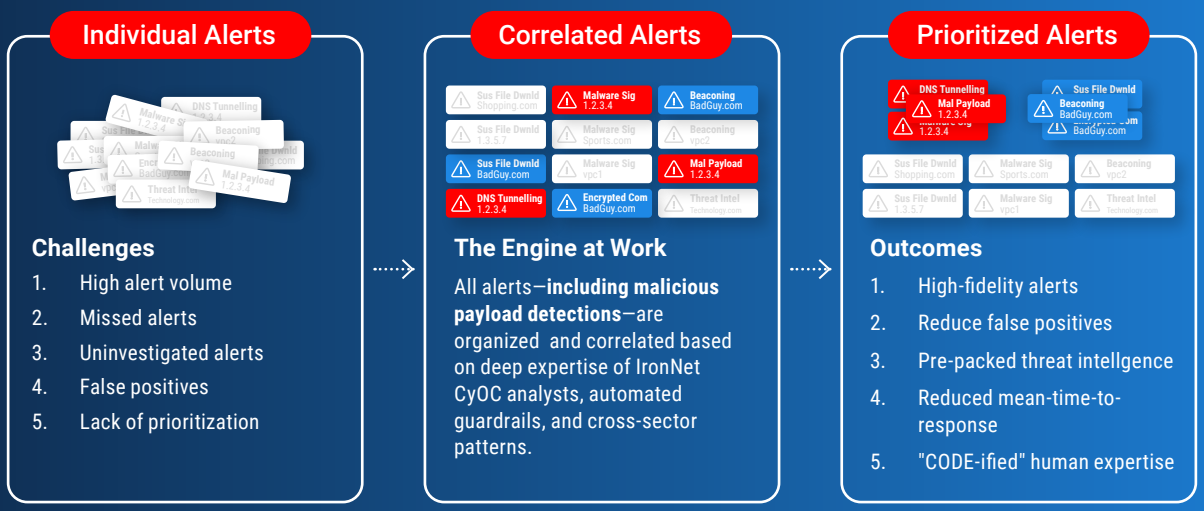

Automatically stringing together multiple alert detections across the kill chain yields the following benefits:

- Alert fidelity
- False positive reduction
- Better analyst workflow
- Customizability
- Collective defense

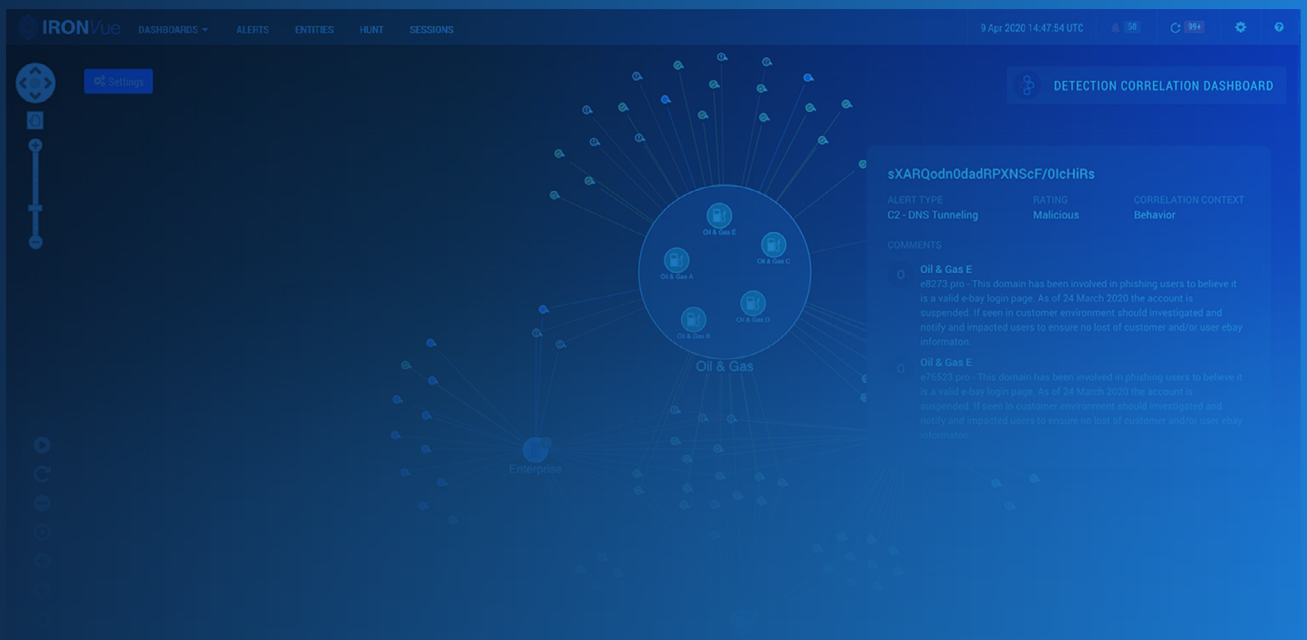

The power of alert correlation

With its unique threat engine, IronNet is bringing the power of alert correlation—the very DNA of the IronNet Collective Defense platform as it automatically corroborates threat detections across companies, sectors, supply chains, states, and governments—to individual networks. This practical way to rule out false positives advances threat detection capabilities without flooding the SOC. After all, analysts already are drowning. Now, we can help them.

IronNet

Connect with us to learn more about the IronNet threat engine.