WHITE PAPER

proofpoint.

THE ALARMING STATE OF SECOPS

TOO MANY TOOLS, ALERTS AND CUTTING CORNERS

TABLE OF CONTACTS

A cybersecurity alert is like a smoke detector sounding off in your house. If the detector is sensitive, anything will set it off, cooking bacon, burnt toast or a steamy shower. If it happens over and over with no actual emergency, you might disable it, take the battery out, or rip it off the wall and stomp on it out of sheer frustration of responding to an alert when there is nothing wrong.

Security alerts from SIEM systems and other types of detection tools share many of the same traits and annoyances of an overactive smoke detector. However, turning off SIEM and other noisy detection tools is not an option for security teams. Instead, most alerts are disregarded. There is just too many and security teams can't keep up. In a recent Ponemon survey, 74% of respondents reported that security events/alerts are simply ignored because their teams can't keep up with the deluge.[1]

The Ponemon survey further illustrates this point. More than 620 security practitioners reported receiving an average of 17,000 security alerts each week. Of those alerts, only 19% are considered reliable and only 4% are investigated.[2]

Collectively, organizations pour billions of dollars into new detection technologies every year. These tools are important. Just like a smoke detector is important to keeping your house safe. But all the smoke detectors in the world can't do anything to put a fire out once it has started. Adding more security tools works the same way. Rather than offering better protection, they generate more alerts and chaos, overwhelming already-stretched security teams. The threats that matter most can't rise above the noise, allowing infections to spread unchecked.

Given that responding to a single alert can take hours or days, security teams simply don't have the capability to respond to every one. They lack the resources to analyze and verify each alert. They don't have visibility they need to pinpoint the infection. And they don't have the expertise to handle the volume when operating with an ad-hoc incident response process.

This paper examines why today's popular detection tools are leaving security teams overworked and underprepared for today's threats. And it explains what organizations can do about it.

NEW DETECTION TOOLS MEAN MORE WORK

Organizations are always looking for ways to enhance their cyber defenses. They deploy a mishmash of security technology: next-generation firewalls (NGFW), data loss prevention (DLP), intrusion prevention systems (IPS) and security event incident management (SEIM) tools. Unfortunately, these point tools don't always offer better protection. Security teams just end up with too many tools and alerts to verify and resolve.

Detection is important—you can't respond to what you can't see. But security alerts can become too much of a good thing. The more alerts a security team gets, the less effective it becomes at prioritizing, responding and containing threats. This situation increases the chance of an attack. It gives attackers more time to do damage. And it makes cleanup costlier and more difficult.

DETECTION TOOLS NOT ONLY DETECT MALWARE, BUT THEY CAN DELIVER HUNDREDS OR THOUSANDS OF ALERTS AND INCIDENTS THAT YOU MUST ADDRESS.

DETECTION TOOLS – WHAT IS THE RETURN ON YOUR INVESTMENT?

Are detection tools really protecting companies or just adding more chaos to an already- hectic firefighting scene? Consider the leading tools.

SANDBOX ANALYSIS

Sandbox-driven detection is popular and effective at detecting zero-day threats and unknown malware. Security teams benefit from the copious forensic data that sandboxing analysis provides. This includes IP addresses, file hashes, domains, C&C infrastructure and location.

If security teams are not prepared to digest this information, the volume of alerts and follow-on investigation can be overwhelming.

All alerts are not created equal

A single threat can have multiple binaries, callback targets, and even sources for file downloads. Attacks may target multiple systems. They often drop or download hundreds of files. And they can take dozens of actions that might harm your organization. That means a multitude of alerts for a single threat. Some alerts may signal confirmed malicious behavior. Others may indicate only suspicious behavior. Any suspected malware infection, unexpected remote server connection, or potential callback warrants a further look. Matching attack data from the sandbox to forensics on an endpoint requires many manual steps. In a typical attack, security teams must:

- Obtain access to the targeted system
- Build a loaner system
- Download SysDump files
- Comb through the files to find forensic matches

That's hours of work for a single infection.

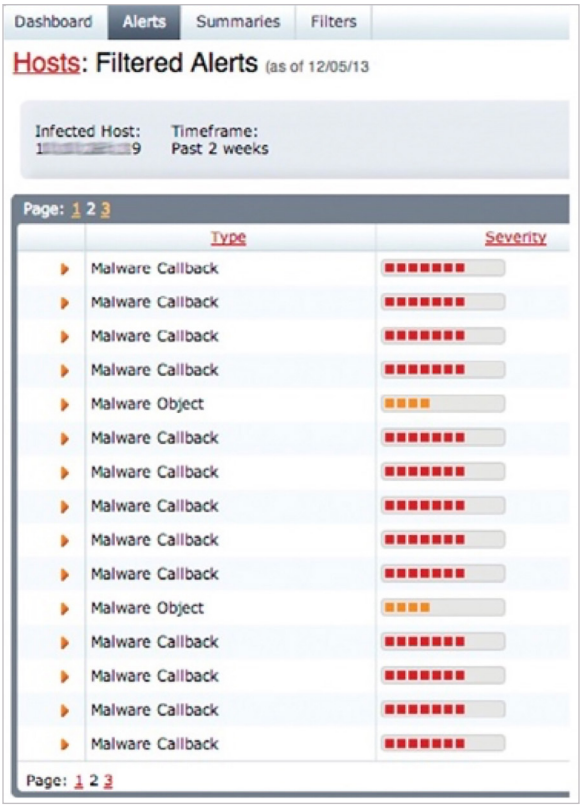

Figure 1: A single threat can represent a great deal of investigative work

Turning forensic data into value

Detecting threats and generating forensic data are not enough to aid your incident response process. Turning this data into value—and getting a return on your sandbox investment—requires an automated response. Consider incident response technology that can reduce your forensic investigation time. This technology should automatically bring together endpoint forensic data, such as processes and file-system changes, and compare it against sandbox forensics. By combining this data and connecting the dots, security teams can quickly verify and prioritize a security threat.

SIEM SYSTEMS

Security information and event management (SIEM) systems are expensive and complex. They aggregate machine and log data, then run user-defined rules against that data. The goal: uncover server and network anomalies that may point to malicious activity.

As a result, security analysts see thousands of alerts everyday. But without some type verification, alerts get easily be dismissed.

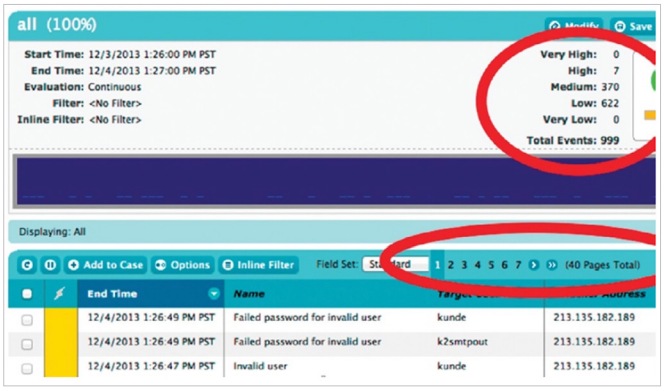

Security analysts see thousands of alerts

More rules mean more work

When a SIEM solution is aimed at uncovering security issues, it requires creating many rules to reveal security concerns. Our customers have written as many as 500 rules to filter out the "noise" that typically comes when aggregating a stream of alerts. This task can take weeks of painstaking work. First, security teams must create and tune these rules. Then they have to verify that each rule is firing properly in relation to hundreds of other rules and network traffic flows.

Even when rules are working and the system is freshly tuned, security teams still need higher-fidelity information about the security alerts they receive. SIEMs are complex and require expertise to maintain—two big reasons they are never fully deployed.

Threat Response solutions bolster SIEM value

In a recent survey by Ponemon Institute, 90% of respondents said their organization scrapped a security technology investment before or soon after deployment. And 31% said security technologies purchased by their organization over the past 24 months were never fully rolled out. SIEM-related threat intelligence was the third-most shelved technology; roughly half of all deployments were abandoned.[3]

SIEM solutions can aggregate and analyze huge volumes of information. But they still require security teams to take many manual steps to match security alerts to the machines and users being targeted. SIEM-generated security alerts can provide a better return on their detection capabilities when paired with threat response technology. For example, alerts are much more valuable when the IP and email addresses of targeted users are automatically resolved to their real identities.

DETECTION AND INTRUSION PREVENTION SYSTEMS (IDS/IPS)

Intrusion prevention systems and similar sensors are a core part of most security programs. They can filter and detect hidden attacks such as distributed denial-of-service (DDoS) attacks. And they can discover unique patterns of attack exploits.

But IDS and IPS systems have a severe downside: once they are set up, the real work begins—and never ends. Detection rules must be constantly tuned. The IPS/IDS database must be frequently updated to filter out the noise and produce high-quality alerts. It requires constant attention to adopt to changes on the network. Producing actionable alerts also requires context. If all alerts are coming through as critical, security teams lack the context needed for good decision-making or response.

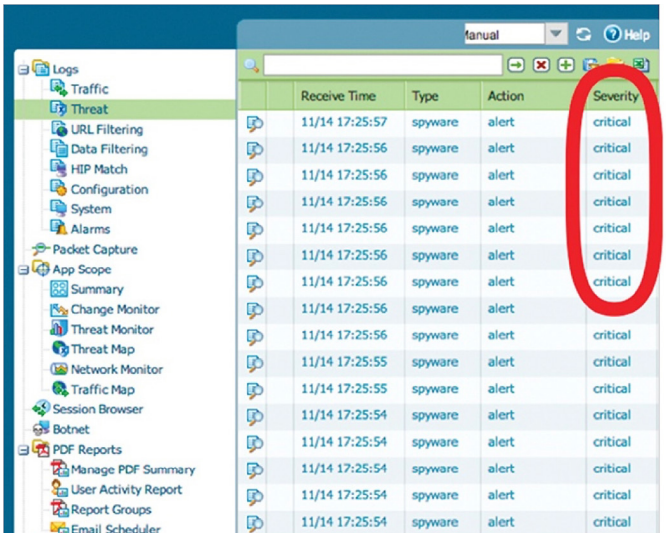

Producing high fidelity alerts required constant tuning and context to make them actionable.

When all alerts look the same

To be effective, IDS/IPS systems must generate high-fidelity alerts; security teams have to be able to objectively "trust but verify" what they're seeing. IDS/IPS systems generate alerts that include information on the potential severity of the threat. But they can't prioritize the alerts in terms of how critical targeted assets and employees are.

With IDS/IPS alone, security teams can't discern between a massive security breach and minor issues—a benign network scan, misconfiguration, change in policy, or other non-threats. Without expert tuning, IDS/IPS systems might rate ransomware attacks as medium to low severity because they often use commodity malware. Blind tuning to ignore any low- to medium-severity alerts may let in ransomware that could have been easily stopped. Similarly, the act of exfiltrating your data may mimic normal business operations. So IDS/IPS systems may rate the activity a low threat—even as your data flows to a command-and-control (C&C) server or data-staging site.

Intrusion detection systems (IDS) and other similar sensors can generate a flood of alerts for something as common as a network scan (friendly or otherwise). Lowering the priority of these alerts might seem tempting. But doing so risks filtering out a hidden DDoS or other attacks. Some teams may instead let these alerts through to gain a better picture of potential attacks on the network. But they still struggle with the workload of investigating them manually.

The obvious problem in either scenario: if all the alerts are tracked and reported as critical, how do you know which ones really matter?

The point of detecting threats is resolving them

If you can determine which alerts matter and respond to them efficiently, IPS/IDS systems may be well worth their investment. If you can't, these tools can mean more noise. And that noise may lead your security team to tune out alerts altogether.

IPS/IDS systems are much more effective with automated response technology. By automatically matching attack attributes and attackers' reputation to targeted users and their group permissions, you can prioritize alerts and act on them immediately.

DETECTION DOES NOT EQUAL PROTECTION

New detection tools may raise the awareness that you've been breached. But unless you have a plan and the ability to contain threats, you're not protected. In effect, you're watching the house burn down—along with any return you were hoping to receive from your security investment.

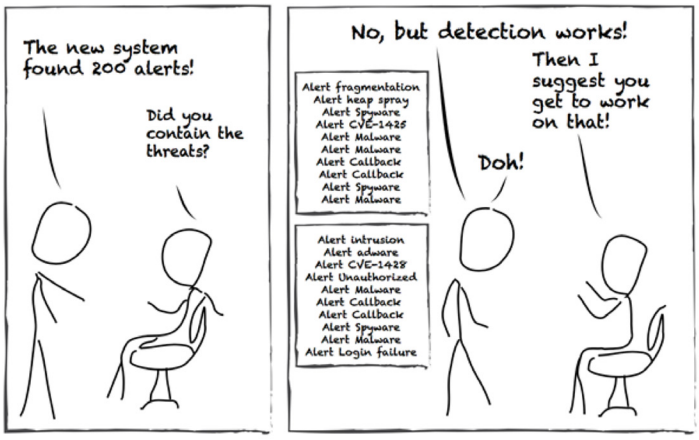

It's no wonder that a recent Ponemon survey showed that it takes enterprise security teams an average of 206 days to spot a breach. Containing it takes another 69 days.[4]

The complexity of many security systems and having the expertise available to administer them is only part of the problem. Many security teams still take an ad-hoc approach to the containment process; most times, no one person or function is accountable.

A recent Enterprise Strategy Group (ESG) survey illustrates the depth of this problem. The survey polled 184 cybersecurity professionals familiar with their company's incident response practices. Nearly 75% said that incident response tends to be based upon informal processes. And 93% of said that the effectiveness and efficiency of their incident response are limited by the burden of manual processes.[5]

RECOMMENDATIONS

To get the highest return on your investment in new detection capabilities, ask yourself these questions:

- How will a new detection capability improve the security team's ability to respond to security events detected?
- What is the ongoing work and expertise needed to reduce security noise and generate high-quality alerts?
- How will this detection capability better prioritize security alerts so that the team can respond quickly to critical threats?
- How does this solution reduce the number of manual processes the security team must perform to investigate an alert?
- What is my return of detection for this product? Does it just detect new threats or can it also block them?

CALLOUT: PROTECTING SMARTPHONES AND TABLETS WHILE NEGLECTING EMPLOYEE-OWNED LAPTOPS LEAVES A SIGNIFICANT GAP IN MOBILE DEVICE RISK MANAGEMENT

Consider technologies that complement your detection tools by matching external threat data with potential internal targets in an automated way.

PROOFPOINT THREAT RESPONSE

Proofpoint Threat Response helps you transform your detection capabilities. It automates key areas of incident response so that you can prioritize, verify and resolve threats 10 times faster than with manual processes. Threat Response helps you:

- Accelerate response. Get automatic collection of external threat information and internal target data. This gives security analysts full situational awareness. They can investigate and prioritize security alerts quickly.
- Save hours per incident. Automated, built-in infection verification dramatically reduces time spent chasing false positives and confirming infections.
- Instantly contain threats. Automated workflows trigger response actions to immediately quarantine and contain infected systems.

Don't waste your security team's time with the mind-numbing manual work of investigating security alerts. Automate these tasks instead. Threat Response is the force multiplier you need to resolve security alerts in less time with less effort.

Learn more about how Threat Response can automate your incident response process. Save hours or days per incident and engage your team in more strategic work.

[1] Ponemon Institute. “The Cost of Malware Containment.” January 2015.
[2] Ibid.
[3] Ponemon Institute. “Risk & Innovation in Cybersecurity Investments.” April 2015.
[4] Ponemon Institute. “The Cost of Malware Containment.” January 2015.
[5] Enterprise Strategy Group (ESG)

ABOUT PROOFPOINT

Proofpoint, Inc. (NASDAQ:PFPT), a next-generation cybersecurity company, enables organizations to protect the way their people work today from advanced threats and compliance risks. Proofpoint helps cybersecurity professionals protect their users from the advanced attacks that target them (via email, mobile apps, and social media), protect the critical information people create, and equip their teams with the right intelligence and tools to respond quickly when things go wrong. Leading organizations of all sizes, including over 50 percent of the Fortune 100, rely on Proofpoint solutions, which are built for today’s mobile and social-enabled IT environments and leverage both the power of the cloud and a big-data-driven analytics platform to combat modern advanced threats.

proofpoint.

www.proofpoint.com