SIEMian Monkey Business

A Commentary Article

by Mark McArdle, Chief Technology Officer
eSentire Inc.

www.esentire.com | @esentire | sales@esentire.com

If you believe deploying SIEM (Security Information and Event Management) with your perimeter security is an effective defense against the ever-growing threats facing your corporate network then read on, or you may learn the truth the hard way.

SIEM was born of SIM (Security Information Management). SIM was the result of a period of massive corporate malfeasance in the early days of the 21st century. Enron, Worldcom and others were the primary motivators of Sarbanes-Oxley. A new regulatory regime that drove compliance officers to deploy SIM as a means of providing evidence their financial control policies were in place and enforced.

As with most accounting-focused initiatives, it was rearward facing. The compliance model driving SIM delivered on weekly, monthly, quarterly and annual reporting requirements, which ultimately captured past incidents.

As SIM became commonplace in publicly traded companies (think ArcSight), some people thought that there was a security play for SIM. And just like that, SIEM was invented as a new security product category.

SIEM provided some useful behavioral capabilities if deployed and configured correctly. Being able to alert on things as basic by highly sensitive as the creation of a new administrator account. These types of accounts aren't created very often, and when they are, it should only be under strict IT control.

The need to manage security logs wasn't something new. In the early days of IDS (remember ISS RealSecure), there was quite a bit of excitement. IDS systems were rapidly deployed. By the early 2000s, they were commonplace. But the IDS systems created a new problem: they generated enormous amounts of data in the form of logs/alerts. Unfortunately, in the real world of signature-based anomaly detection, (the core brain of most IDS systems), there are a lot of false positives. IDS systems had (and still have) real limitations in their ability to produce black and white results. They produce lots of gray. Gray is a problem. Gray is noise. And noise means extra work.

The response to this noise was to outsource IDS logs to a 3rd party. Companies couldn't justify having resources sift through the massive logs in search of threats. By this time, a market called Managed Security

If you believe deploying SIEM (Security Information and Event Management) with your perimeter security is an effective defense against the ever-growing threats facing your corporate network then read on, or you may learn the truth the hard way.

Unfortunately, in the real world of signature-based anomaly detection, (the core brain of most IDS systems), there are a lot of false positives.

Service Providers (MSSP) was already in flight. This market was created because firewall management became quite difficult. Firewalls like Checkpoint's were powerful but required specific skills to manage effectively. These skills were in short supply (just as security skill remains in short supply to this day). MSSPs stepped up to concentrate the talent around a model that supported many corporate networks. It was a valuable service and so the MSSP market grew. The IDS noise problem was something MSSPs were ready and willing to help solve. However managing Firewall and IDS logs/alerts requires a different approach than a change-control firewall policy service.

Moving the noise generated by IDS systems to "experts", MSSPs solved one problem, (or at least gave the perception of solving one problem) : "We have smart people looking for threats in our IDS logs."

But the honest, often unheard truth is that ultimately relying on logs leaves you incapable of taking the appropriate action because the noise can't become a useful signal without better context.

No matter how long you stare at an IDS log event, it won't become any more informative. The same is true for the vast majority of security log events. But let's put that primary flaw in log-based security aside for a moment.

Let's think about how SIEM works in the real world. It takes feeds from firewalls, domain controllers, IDS systems, end point protection systems and other security devices. Given their capabilities, policy and threat intel, these devices do their best to make the right decision. There are a lot of very smart engineers working on security products who when given a discrete problem, can design elegant solutions. However, in the security use case, attacks aren't discrete and static. They're commonly dynamic and frustratingly innovative.

Almost every security product solution comes armed with many powerful capabilities that can, in theory, insulate a system or network against attack. Those capabilities work exceptionally well in a lab environment where known attacks are used and policies are optimized for security. But in the real world, where we all live and work, security policies are never fully optimized for security. They are a compromise between business priorities and security. This isn't news to anyone working in corporate network security today. Shutting down attack vectors often isn't possible because the technologies those vectors utilize are mission critical. So IT security managers are left with more gray and even more risk.

Today we have powerful security devices, like NGFW, IPS/IDS, endpoint and everything in-between deployed with watered-down policies that compromise their efficacy of the perimeter. And even using the word perimeter is a bit of a joke today with the mobility of endpoints and the Internet of Things.

What are the implications of this real-world scenario for the SIEM approach? Well, first off, it requires a level of trust in the perimeter defenses to identify a threat and provide an actionable event. Those of you who have looked at the kinds of things that get shipped to a SIEM via syslog or Windows Event Manager will immediately recognize a challenge: the event itself is atomic. For the majority of events, they say, "This happened. I think. Good luck."

Actually, you don't get the "Good Luck" part. It's implied.

So the first challenge we face by relying on SIEM alone is trusting that the perimeter will give us a high fidelity signal that indicates a specific threat that can be acted upon. But in the gray world of false positives, it's frankly impossible.

You need more context to turn a "This happened… I think," into a definitive threat or non-threat.

The second, and I'd argue more important challenge faced by SIEM, is the assumption that there will even be an event in the SIEM to investigate. If the NGFW, IPS/IDS system or endpoint has no reason to suspect the legitimacy of a connection, transaction or binary based on its policy (and latest Threat Intel update) then it isn't going to put anything useful in the SIEM. It will be mute. Essentially you trust the very things we know to be deficient against new and innovative attacks to provide the SIEM with any kind of actionable alert. The perimeter is useful for blocking the "background radiation" of the Internet.

These are the threats we've seen before or known vulnerabilities from actors we know to be suspicious. But it's useless when you change any one of these parameters. A NGFW never looks over its shoulder and second-guesses itself. You'll never find anything useful in your pretty SIEM reports about these kinds of attacks.

A NGFW never looks over its shoulder and second-guesses itself.

I think Amit Yoram, the CEO of RSA Security stated the problem beautifully in is 2015 RSA keynote titled "Escaping Security's Dark Ages" when he said:

"Nonetheless, many security professionals base their programs on the futile aggregation of telemetry from these virtually blind IDSes, AV platforms, and firewall logs, implementing the glorious and increasingly useless money-pit, known as the SIEM. I know it didn't surprise many of you when last year's Verizon Data Breach Investigations Report asserted that less than one percent of successful advanced threat attacks were spotted by SIEM systems. Less than 1%. The terrain has changed but we're still clinging to our old maps. It's time to realize that things are different."

Relying on a SIEM is reckless…

Relying exclusively on a SIEM to identify and manage threats is reckless; it's an accounting "rear-view mirror" perspective that can **only** inform you of known threats based only on the insights gleaned from perimeter defenses. These perimeter defenses are essentially useless when it comes to new and innovative attacks. And without additional context, you can't identify an actual threat from a mundane false positive.

If you rely exclusively on a SIEM, and haven't deployed anything in your network to allow your SOC (or your outsourced partner's SOC) to identify and forensically investigate threats then the truth is that you're only pretending to do security.

The gray signals generated by security technologies require additional context. A talented human analyst when provided with this context can quickly and effectively determine the nature of a potential threat.

At eSentire, we have always believed this requires deploying a sensor that has visibility into everything going in to and out of your network. This same sensor will identify strange behaviors using Deep Packet Inspection and a rich suite of network forensic capabilities. One of these capabilities is full packet capture archiving. This sensor, called esNETWORK^{TM}, enables our Security Operations Center (SOC) analysts to investigate any issue in "full UHDTV". That's exactly what we deliver with esNETWORK through our Managed Detection and Response^{TM} platform. We treat the logs and events from 3rd party products as useful context to help paint a richer picture for our SOC analysts. But we don't rely on them to identify threats.

It's only going to get harder to protect your networks. You have to embrace the reality that your perimeter and endpoint security products, no matter how powerful, will ultimately fail when dealing with anything other than yesterday's attacks. The security game has shifted from prevention to detection. The new game plan demands not just an effective perimeter defense to block the "background radiation", but also requires continuous monitoring that doesn't rely on a SIEM for its visibility into threats.

Security is hard. But it can be a lot easier if you focus on managing threats effectively and stop worrying about who's pretending to deliver security by staring at your logs.

Mark McArdle is a charismatic, seasoned technology executive with over 20 years experience with top technology brands, such as PGP Inc. and McAfee, where he served as Senior VP for the company's consumer product development division. Mark defines new products for eSentire to remain leading edge and competitive, and himself holds six patents in Internet security. Mark holds a Bachelor of Science degree from the University of Waterloo and is a graduate of Ivey Business School at Western University, and completed the Ivey Executive Program in 2005.