LogicHub

TRANSITIONING THE ISOC FROM ALERTS TO NEWS

INTRODUCTION

In November of 2015 Gartner released a white paper titled "The Five Characteristics of an Intelligence Driven Security Operations Center." The point of the paper was that, despite the extensive deployment of Security Operations Centers (SOCs) in the enterprise sector, the basic problem of threat visibility has not significantly improved. 82% of all enterprises consider it a problem, and, perhaps more importantly, 76% of enterprises *that already have* a SOC consider it a problem. This despite the fact that significant amounts of money and resources have been deployed to prevent, detect, and react to cyber threats. It's clear from these figures that the hacking threat is growing far faster than today's SOCs can keep up.

The solution, in Gartner's view, is for enterprises to deploy what it calls "Intelligent SOCs", or ISOCs. In essence, Gartner says you need to make your current SOC smarter. We agree. The question is how?

THEY GAVE US "GOOGLE ALERTS"

Enterprises can rightly ask: wasn't Security Intelligence Management (SIM), as promised by Arcsight, Splunk, and others supposed to fix this? As it turns out, SIM products most closely resemble Google Alerts. Strip all the promises away and you're still faced with the need to know what you're looking for, set a bunch of specific alerts for yourself, and sit back and pay tens of thousands of dollars a month to wait for intrusions to hit.

CISO's are not charged with catching specific threats: they're charged with preventing intrusions where possible and, at minimum, detecting them when they happen and mitigating the damage they cost. And that is not happening today, SOCs or not.

WE DON'T NEED "GOOGLE ALERTS", WE NEED "GOOGLE NEWS"

The Gartner report goes on to describe the main issues with the current allocation of human and technology resources in SOCs. It is a more detailed description of the Google Alerts problem.

As an industry, we know these all too well:

1. Attack patterns are seldom known ahead of time.
2. The systems in place to deal with them are not adaptable enough to accurately identify unknown attack patterns based on the known ones.
3. Current systems are not easy to automate.
4. Existing analytics packages do not do a good enough job of organizing and prioritizing potential threats for review.

Employing our Google News metaphor, if detecting threats is the same as finding the right news stories, the problems with the current batch of SOCs are:

1. We cannot identify news topics other than the ones we specify
2. Our news product does not use the topics we specify, or our behavior, to predict the news stories we wish to find.
3. We are running the news curation using mechanical turks at best.
4. There is no visual interface that lets us see which news stories we are most likely to want to read.

Gartner goes on to describe what they see as the solution, in basically a point-by-point response to the above issues:

1. Use multi-sourced threat intelligence.
2. Make the technology predict threats.
3. Automate as much as possible.
4. Create better analytics interfaces.

This all makes sense. The question is: how do you achieve these changes, given that an ISOC (whether it goes by that name or any other) has been the Holy Grail of the enterprise infosec community for years?

“GOOGLE NEWS” FOR SECURITY OPERATIONS IS NOW POSSIBLE

We are seeing the beginnings of an approach to solving the problem that incorporates and expands on Gartner’s approach.

Step 1: Round up all the data, add context

It isn’t enough to feed your known attack patterns into your SIM system and sit back and wait for similar patterns to come in. By the time those patterns belong to your dataset they’re probably not going to come in again. Your first step is to get as much data as possible into your system. This includes:

1. Events
2. Context
3. Threat intelligence
4. Entity enrichment

When Gartner suggests using multi-sourced threat intelligence, they are pointing to a subset of all the data we believe needs to be brought into the system. Which means multiple sources of data.

It is important to note that the value of this data is purely as a teaching mechanism, and what matters with this teaching mechanism is less the *events* themselves and more the *context* and supplementary information we can enrich them with.

Context data is seldom utilized when building these systems, but in fact context is far more important than event data for teaching technology to predict. To coin a term, we think that the event/context ratio used by current SIM systems is the inverse of what it should be.

We estimate that current systems use on the order of 90% of intrusion events combined with about 3% of intrusion event-related contextual data. But the deep correlation, the diseases that drive the event symptoms, are to be found in the contextual data, not the events themselves.

Our approach is to use on the order of 1% of available event data and pair it with 90% of available contextual data in order to create predictive algorithms that truly work for unknown attack patterns.

Google News doesn’t just use a list of news article headlines to evaluate what you’re interested in. In fact, for every headline they have dozens of other statistics that help them infer it – everything from what you click on to how long you view to what you share and sentiment analysis of your comments. That is the type of event/context ratio we should be targeting in our security initiatives.

Step 2: Use the Data to Feed Advanced Analytics and Find Deep Correlation

When you have as much data in the system as possible, your next step is to bring your technology stack to bear on inferring correlation wherever possible to predict Indicators of Compromise (IOCs). However, the current batch of SIMs do not take advantage of the latest advances in Machine Learning (ML) and Artificial Intelligence (AI) to infer those correlations.

Step 3: Make the Detection Logic Easy to Automate

Everything that can be automated with positive ROI should be automated. The good news is that advances in machine learning make it much less expensive to automate many repetitive processes that would have been too costly to implement in the past (an example being the detection logic used to flag IOCs). However, SIM solutions do not take full advantage of flexible automation to take often-repeating tasks off analysts' plate.

Step 4: Enable Better Interfaces

Think of the list of trending news topics that run down the left-hand side of the Google News Feed. They represent what the system thinks are the highest priority news items for you, based on how you and others have interacted with the system in the past. The main body of the news feed is broken up into categories that you've customized, within which related news stories have been prioritized.

Threat detection interfaces should similarly aid analysts in understanding the relevance of events identified using several different vectors. Those vectors should include overall learnings developed by finding deep correlation, as well as real-time customization that happens based on the specifics of the deployment.

The alternative is for the analyst to manually sift through an unending list of events and to try to apply logic to it unaided, which is totally impractical in today's security environment.

HOW INTELLIGENT IS YOUR SOC, ANYWAY?

It is possible to measure how well your current SOC setup stacks up in each of the steps we describe above:

Rounding Up the Data

1. What percentage of event data can you bring into your system?
2. What percentage of real world entities are modeled?
3. Can you enrich the entity model?
4. Do you understand the relationship between entities?
5. Can you use analytics to enrich those entities?

Advanced Analytics

1. Can you programmatically correlate your different sources of data?
2. Can you do that correlation in real time?
3. Does your system incorporate Behavior Models and Outlier Detection?
4. Can you correlate your entity model and the events you feed the system?
5. Does your system do deep level analytics
6. Can machines do the explorative analytics for you, or are you still limited to human evaluation?

Automation

1. What percentage of tasks do you do that are repetitive but not automated?
2. What percentage of automation-worthy tasks are not automated because it would take too long?
3. What percentage of automation-worthy tasks are not automated because the underlying analytics are not powerful enough to analyze what you can?
4. What percentage of automation-worthy tasks are not automated because the automation is static and hard to adapt, which lowers the ROI?

Discovery

1. Can you automate discovery reliably?
2. What is your false negative rate? Can you measure it?

CONCLUSION

If your enterprise is like the 76% of enterprises that have a SOC and still feel that threat visibility is a huge problem, then Gartner's recommendation to evolve to an ISOC is one you should take seriously. As you do so, we recommend implementing a set of technologies, processes, and analysts that address the limitations inherent in today's SIM deployments.