BIZCOMMUNITY

# For cybersecurity success, double-down on developing better detectors

By Martin Potgieter

29 Jul 2019

The basis of threat detection in security operations centres over the past decade has predominantly relied on rules, commonly defined by SIEM vendors.

Martin Potgieter, Technical Director at Nclose

These “rules” (also called alarms, alerts or use cases depending on the SIEM vendor, or as we prefer to call them “detectors”) then generate alerts that tell analysts that there is a potential threat and helps to produce data that can unearth trends in what types of attacks are common at any given time. In theory, detectors should improve cybersecurity efforts.

However, the sheer volume of alerts the average cybersecurity team have to deal with is often overwhelming. In a recent Ovum study of cybersecurity challenges in the global banking sector, 37% of banks said they receive more than 200,000 alerts per day. For 17% it was in excess of 300,000.

This barrage of alerts is causing a “signal-to-noise ratio” problem. Inevitably, security professionals miss threats simply because there are too many alerts to attend to.

It's hard to quantify how many alerts are ignored. One US study by the Cloud Security Alliance found that 31.9% of security professionals surveyed ignored alerts because there are so many false positives. Local data is unavailable, but we know that many alerts are passing by unnoticed.

Most cybersecurity vendors offer the following answers to this barrage of alerts:

- **Orchestration** – this essentially automates the response to alerts. Most basic and even some more complex responses to alerts can be automated with orchestration tools that integrate the various security tools.

  However, this automation normally takes significant effort and the ROI is not always clearly visible when looking at the price tag of the tools and effort required from a DevOps perspective.



- **Artificial Intelligence and Machine Learning** – The idea here is that the alerts are more refined and better at detecting malicious activity given the large volume of data. However, this often exacerbates the problem by actually just creating alerts on top of the alerts already generated by the traditional methods.

Maybe the answer is simpler. We can improve our detectors' performance and discard ineffective detectors. This doesn't sound as exciting as AI and Orchestration, but we have seen first-hand how effective it can be. Before we are able to improve the performance of detectors, we need the means to measure their current effectiveness.

I suggest the following four key attributes that could allow us to measure our detectors:

## #1 Simplicity

Keep detector parameters as simple as possible - complexity doesn't always improve alerts (in many cases it has the opposite effect). As soon as detectors are overly complex, they can be difficult to investigate, become expensive or just break entirely.

There are almost always multiple ways to detect a particular type of malicious behaviour, and it may be a good idea to create a range of different detectors to identify that behaviour, but measuring each detectors simplicity will allow you to prioritise the simple detectors.

For example, to detect password spray attempts, we could tap into network traffic, apply some sort of AI or ML to the captured traffic and look for outliers that would indicate this type of malicious activity.

Alternatively and more simply, we could enable a specific audit level on Active Directory, and look for a flood of authentication failures from different user names. Both detectors would work but the latter, simpler approach would be my preferred method.

## #2 Certainty

Certainty relates to how likely it is that a detector represents actual malicious behaviour. Oftentimes, detectors will pick up anomalies that are not actually malicious, but that still require further manual investigation. If it is malicious, further details about the incident have to be determined.

This manual investigation is not always a bad thing but it should be measured as a metric of the detector. If it is not producing consistently accurate alerts it has to be tuned.

## #3 Resilience

Can your detectors work in adaptive conditions? Organisations' constantly changing IT landscapes require that cyber defences are adaptable.

These changes are often referred to as “environmental drift”, where previously optimally-running processes suddenly underperform or stop working entirely. Keeping things simple certainly aids resilience, but there are other variables to consider. Attackers may be aware of common detection methods and will attempt to execute their attack without violating these rules, so how does the rule stand up to these evasive manoeuvres?

Again this is where simplicity may be on your side. Using the same example from above, the simpler detector which only relies on AD security events would be much more resilient than the detector that requires network taps, AI and complex analysis.

## #4 Relevance

This is one of the more difficult things to measure; relevance is often subjective to the organisation. Relevance also depends on multiple factors: How new or old is the attack the detector is designed to uncover?

Is the attack relevant to the organisation? For example, if a detector is designed to identify an attack specifically against the Solaris Operating system and an organisation does not have a single Solaris system, there is probably only limited value in that detector.

The reason it’s important to measure relevance is that we need to ensure our defences can detect attacks that are happening today and that may be successful in our environment. Organisational complexity and the growing sophistication of cyber attackers makes the job of the cybersecurity professional all the more difficult. Measuring detectors against these four attributes provides a useful starting point to start measuring their effectiveness. Some detectors may not score highly on all four attributes, and that’s fine.

Knowing a detector is highly resilient and its underlying rules are not overly complex gives cybersecurity professionals clarity on how that detector can be used, and how it should be measured.

And that clarity can greatly improve the signal-to-noise ratio, reduce alert fatigue, and deliver greater efficiency in an organisation’s overall cybersecurity efforts.


ABOUT THE AUTHOR

Martin Potgieter, technical director at Nclose.