FALSE POSITIVES

Chuck Everette *explains how automation can resolve the drain on cyber security teams*

One of the biggest challenges in cyber security is the sheer scale of the task. Modern IT networks are highly complex, combining hundreds of components that constantly interact in different ways. This provides threat actors with a large attack surface and myriad options for exploiting flaws to infiltrate the network. Alongside the sheer scale, cyber security is highly dynamic, with attackers continually exploiting new vulnerabilities, developing new malware and discovering ways to evade or counter existing security solutions.

The scope of the challenge means security strategies usually hinge on trying to cover as much ground as possible. Businesses will be armed with a range of security tools configured to detect known threat signatures or signs of unusual behaviour that indicate malicious activity. With a single security incident potentially leading to millions of pounds in lost business and recovery efforts, most decision makers will opt for a 'better-safe-than-sorry' approach that flags any potential risk for review.

However, casting such a broad net invariably means dredging up far more potential threats than security personnel can handle. For every genuine threat that is detected, there will inevitably be a large number of false positives – alerts that soak up security resources but represent no genuine threat to the organisation.

False positives are the result of security detection tools incorrectly flagging benign activity as a potential threat

Most security strategies revolve around security information and event management (SIEM) tools that serve as log data aggregators for solutions such as firewalls, endpoint detection and response (EDR) and antivirus. The SIEM will gather logs from detection tools across the IT infrastructure and turn them into alerts for the security team to analyse and resolve. While these solutions have improved in recent years, it's still common to find that personnel are given little in the way of context, making it hard for them to properly prioritise their response.

False positives are the result of these security scanning and detection tools incorrectly flagging benign activity as a potential threat. This is usually the result of a file or application having some similarities to the threat signature or tactics, techniques and procedures (TTPs) of known malware. Sometimes it is simply because the scanning tool lacks the fidelity to accurately tell the difference. On the other end of the scale are false negatives – genuine threats that have been overlooked by security tools.

Like the boy who cried wolf or someone repeatedly pulling a fire alarm, an influx of false positives will quickly drain the resources, attention, and patience of a security team, causing several issues that can expose the organisation to greater risk.

The most prevalent issue caused by a large volume of false positives is the sheer amount of time it takes to resolve them. Each alert needs to be properly analysed and resolved, with the amount of time required varying on the nature of the alert.

Deep Instinct's Voice of SecOps report, which surveyed over 600 security decision makers and practitioners, found that an average 10 out of every 39 hours in a working week was spent handling false positives. This means roughly a quarter of any given week is spent simply ticking off alerts that pose no danger to the organisation.

False positives usually vastly outnumber the genuine threats. In one prominent example, we worked with a large organisation that received around 75,000 alerts on a daily basis. On any given day, just two of those alerts were likely to be valid threats.

UNWANTED DISTRACTIONS

Dealing with all these false alarms distracts personnel from the real threats lurking in their to-do pile. As mentioned, SIEM tools often present alerts with little or no context, so teams will be left to slog through them in chronological order without the ability to prioritise effectively. Accordingly, it might be days or even weeks until a serious threat is properly assessed.

Threat actors love this kind of backlog because it grants them a great deal of additional dwell time. A genuine alert sitting at number 200 in the to-do list could herald a threat actor within the environment who is free to roam undisturbed for days. In many cases even being 15 minutes too late can result in losing the trail of breadcrumbs that could have been used to discover and stop an intruder, so being days behind is a total lost cause.

Aside from the increased risk exposure of such sluggish response times, having the security team spend so much of their day grinding through repetitive, low-value activities creates a poor work environment. Unsurprisingly, 90 percent of respondents in our survey stated they considered false positives to be contributing to low staff morale. Analysts feeling burnt out dealing with the endless supply of alerts will begin to perform poorly and staff turnover is likely to be elevated.

Security analysts tend to have inquisitive minds that thrive on unlocking puzzles and solving challenges. Many will rightly begin to consider the daily slog through false positives to be a waste of their talents and training and it's common to find security personnel looking for greener pastures as a result.

This is particularly problematic in an industry that already suffers from a severe shortage of qualified, experienced practitioners. Further, the best analysts are able to draw on years of previous experience. Whenever a subject matter expert leaves, the team will be left with a gap in its knowledge. This often has a knock-on effect on the value of security solutions, with tools often becoming shelfware when the resident expert moves on and the replacement lacks their expertise.

The volume of threat alerts bombarding the average SOC team is far too great for human personnel to deal with alone. Even if they could keep up with the influx of new alerts, having skilled analysts spend a quarter of their day ticking off false positives is bad for productivity, ROI and morale.

THE MOST PREVALENT ISSUE OF FALSE POSITIVES IS THE AMOUNT OF TIME IT TAKES TO RESOLVE THEM

The answer is automating through artificial intelligence. The superhuman level of speed, stamina and attention to detail required for the job can only be met by analytical solutions powered by artificial intelligence. Our research found that tools such as AI, machine learning (ML) and deep learning (DL) have made a significant impact on reducing false positives and preventing unknown threats.

AI has become a broad umbrella term covering a number of different subsets of technology, but the common trait is a solution that can ingest data in order to establish and recognise patterns. In the security field this is most commonly seen in the form of ML solutions. These tools can be trained up on threat data until they can reliably recognise different kinds of attacks, common TTPs and automatically respond as needed. Once they are trained, the solutions can crunch through huge sets of data far faster than the best human analyst, and without the risk of low morale and burnout.

Processes can be automated so that false positives and genuine, but low-level threats, are analysed and resolved in a matter of seconds without human intervention. This frees up the security personnel to concentrate on the more fulfilling and valuable aspects of their job.

However, while they are powerful tools for dealing with the constant influx of alerts, traditional ML solutions have flaws that can be manipulated by threat actors. Attackers can use their own machine learning tools to create poisoned data sets, tricking the security tool's model into mislabelling a threat as something benign. If the machine recognises that data set as safe, it will create a false negative that allows the attacker to slip into the environment and create a backdoor.

Traditional ML tools also generally rely on data feeds from tools such as AV and endpoint detection and response (EDR), which means the technology can only react to identified threats, not predict them. Threat actors have become more adept at launching attacks that will have an impact well before EDR systems can gather enough data to assess the threat.

The answer is deep learning, the highest subset of AI. This approach is marked by its extremely high processing speed, identifying potential breaches in less than 20 milliseconds. The technology is new in cybersecurity, but has seen extensive use by Tesla and YouTube for autonomous driving and image recognition.

Deep learning involves the creation of a neurological network, which is trained on raw data samples of millions of labelled files. The key difference to standard machine learning is that deep learning is not given information about which files are malicious and which are benign. Instead, it must learn to make these determinations independently.

ROUGHLY A QUARTER OF ANY WEEK IS SPENT TICKING OFF ALERTS THAT POSE NO DANGER

Eventually, the network is able to instinctively identify malicious code, moving away from reacting to EDR data to actually predicting and preventing attacks before they can truly begin. This more complex approach is also harder for criminals to crack, greatly reducing the risk posed by adversarial ML and poisoned datasets.

Alongside its exceptional ability to predict attacks and proactively identify emerging threats, deep learning will do all the heavy analytical lifting for the security team. When properly integrated into the security stack, we have found that deep learning can reduce the volume of alerts a security team is reviewing by as much as 25 percent on a weekly basis. As a result, the team will no longer be burdened by the millstone of false positives and low-level threats, releasing them for more valuable activity and ensuring they are free to react quickly when a legitimate threat does emerge.

However, it's important to bear in mind that, powerful as it is, deep learning is not a magic button that can be pressed to solve every security problem by itself. Organisations still need a solid framework of conventional security solutions in place, and deep learning works best when there is a strong, multi-layered security infrastructure for it to support.

Accordingly, organisations need to take their time when planning to move into deep learning. It can be tempting to rush into purchasing the latest shiny new technology, but it's important to carry out the due diligence first.

CISOs and other security decision makers should carefully assess their current stack and security goals, and consider how deep learning fits in. What results are they looking for? What are they going to augment or complement? Will anything not fit and need to be replaced or upgraded?

When they have solid answers to these questions, CISOs can begin the process of integrating deep learning into their security strategy. With an automated system powered by deep learning, the security team will be freed from wading through false positives and can then concentrate all of its energy on keeping the organisation safe ●

Chuck Everette is Director of cyber security advocacy at Deep Instinct.

The biggest issue caused by false positives is the amount of time it takes to resolve them